# Consequence-Based and Fixed-Parameter Tractable Reasoning in Description Logics 

František Simančík and Boris Motik and Ian Horrocks<br>Department of Computer Science, University of Oxford, UK


#### Abstract

In this paper we investigate the consequence-based algorithms that are nowadays commonly used for subsumption reasoning with description logic ontologies, presenting the following novel results. First, we present a very general consequence-based reasoning algorithm that can be instantiated so as to capture the essential features of the known algorithms. Second, we use this algorithm to develop a novel framework for a quantitative and parametric analysis of the complexity of subsumption reasoning in description logic ontologies. Our approach is based on the notion of a decomposition-a graph-like structure that, roughly speaking, summarizes the models relevant during ontology reasoning. We identify width and length as decomposition parameters that determine the "level" of combinatorial reasoning. We prove that, when parameterized by a decomposition, our consequencebased algorithm runs in time that is fixed-parameter tractable in width and length. Third, we briefly discuss how length and width characterize the behavior of tableau algorithms. Fourth, we show that the width and length of existing ontologies are reasonably small, which, we believe, explains the good performance of consequence-based algorithms in practice. We thus obtain a sound foundation for a practical implementation of ontology reasoners, as well as a way to analyze the reasoners' performance.


Keywords: description logics, parameterized complexity, treewidth

## 1. Introduction

An ontology describes an application domain in terms of concepts (i.e., the kinds of objects in the domain), roles (i.e., the relationships between objects), and axioms (i.e., statements that formally describe concepts and roles). Axioms are logical statements and can thus be automatically processed via reasoning. For example, an ontology describing the structure of a university might use concepts Employee, Professor, and AssistantProfessor to represent employees, professors, and assistant professors, respectively, and the university's organizational hierarchy might be captured by axioms stating that Professor is a subconcept of Employee, and that AssistantProfessor is a subconcept of Professor. An ontology reasoner can use such statements to conclude that AssistantProfessor is also a subconcept of Employee. This is an example of subsumption reasoning (i.e., reasoning about the hierarchy of ontology concepts), which is a central computational problem in ontology-based applications in areas as diverse as biology [1], medicine [2], geography [3], astronomy [4], agriculture [5], and defense [6].

The Web Ontology Language (OWL) [7] and its successor OWL $2[8,9]$ are standard ontology languages developed by the World Wide Web Consortium (W3C). Their logical underpinning is provided by description logics (DLs) [10]-a well-known family of knowledge representation languages. Reasoning with expressive DLs can be quite complex, so for simplicity in this paper we focus on the basic description logic $\mathcal{A L C I}$ that provides the fundamental constructs such as inverse roles, full Boolean connectives on concepts, and existential and universal quantification; however, quantification is guarded [11], which ensures that subsumption reasoning is decidable, albeit of high complexity (ExpTime-complete). The reason for this high complexity is the size and number of relevant models [10, Ch. 3]: due to an interaction between existential and universal quantifiers, an ontology may be satisfied only in models containing an exponential number of objects (an effect known as and-branching), and due to disjunctions, an ontology may admit an exponential number of "relevant" models (an effect known as or-branching). State of the art OWL reasoners, including Pellet [12], FaCT++ [13], and HermiT [14], are all based on variants of the tableau algorithm, which try to construct a finite representation of a model of an ontology. The basic algorithm has been heavily optimized so as to curb and- and or-branching in practice [10, Ch. 9], which allows DL reasoners to successfully process many large and nontrivial ontologies. The performance of such reasoners is, however, relatively brittle, and there exist ontologies (for example the GALEN ontology ${ }^{1}$ ) that none of these reasoners is able to process. Although there has been some work on identifying the qualitative features of an ontology that may degrade the performance of a given reasoner [15], no existing method can provide a quantitative measure of the "difficulty" of reasoning with a particular ontology.

In order to provide more robust performance of reasoning, one can reduce language expressivity so as to make reasoning easier. By removing disjunctions one obtains the Horn family of DLs [16]. Horn ontologies can always be translated into

[^0]Horn clauses, which eliminates all or-branching; however, while this has been empirically shown to make reasoning easier [14], reasoning with Horn- $\mathcal{A L C \mathcal { I }}$ is still ExpTime-hard due to and-branching [17]. By further removing universal quantification and inverse roles from Horn- $\mathcal{A L C \mathcal { I }}$ one obtains the $\mathcal{E} \mathcal{L}$ family of DLs [18], for which subsumption reasoning is tractable; similarly, by removing universal quantification, qualified existential quantification, and conjunction from Horn- $\mathcal{A L C I}$ one obtains the DLLite family of DLs [19, 20], for which subsumption reasoning is also tractable. Although Horn DLs have been successfully applied in practice, many applications require inverse roles (which precludes the use of $\mathcal{E} \mathcal{L}$ ), qualified existential quantification (which precludes the use of DL-Lite), or concept covering constraints (which requires disjunction and hence precludes the use of Horn DLs altogether). Thus, improving reasoning performance by reducing expressivity is not always feasible.

An alternative is to explore new reasoning techniques that may be better adapted to practical ontologies. Consequence-based algorithms [21, 22] have recently proved to be very effective in practice; for example, the CB reasoner is currently the only reasoner that can process the GALEN ontology in its entirety [21], and the ELK reasoner has revolutionized the processing of large $\mathcal{E} \mathcal{L}$ ontologies such as SNOMED $\mathrm{CT}^{2}$ [23]. Intuitively, such algorithms can be seen as variants of resolution optimized to reduce the number of clauses produced on typical ontologies. However, although consequence-based reasoners have been very successful, it is still not possible to quantify the difficulty of reasoning with a specific ontology.

We believe that methods for the quantitative analysis of the difficulty of reasoning with a specific ontology $\mathcal{O}$ would have major benefits in practice. On the one hand, users would be able to predict the effectiveness of reasoners on $\mathcal{O}$, and analyze performance problems if they did arise; and on the other hand, algorithm designers would gain a deeper understanding as to why reasoning with practical ontologies is often easy notwithstanding the high worst case complexity of the problem, and could exploit this understanding in order to design more robustly scalable reasoning algorithms. Crucially, this analysis must go beyond worstcase complexity and consider the salient features of a specific ontology. One can frame this problem in terms of parameterized complexity [24], which measures the difficulty of a computational problem not only w.r.t. the size $n$ of a problem instance, but also a parameter $k$ that quantifies specific aspects of the instance. Of particular interest are fixed-parameter tractable (FPT) problems, which can be solved in time $f(k) \cdot n^{c}$ for a fixed constant $c$ and a fixed computable function $f$; such problems are interesting because one can hope to solve large instances provided that the parameter remains "small." For example, many hard graphtheoretic problems (including graph homomorphism and 3-colorability) are FPT when the parameter is the graph's treewidth [25]-a measure of the graph's similarity to a tree. Furthermore, the notion of treewidth has also been extended to propositional formulae and has been used to obtain FPT results for propositional satisfiability [26]. A key question in parameterized complexity is to identify natural parameters that accurately characterize the problem's difficulty.

In this paper, we present what we believe to be the first framework for a quantitative, parameterized analysis of the complexity of subsumption reasoning with $\mathcal{A L C I}$. Via well-known reductions of role inclusions and role transitivity axioms, our approach can also be applied to $\mathcal{S H} \mathcal{I}$ ontologies, and so it covers a significant subset of OWL 2 DL (the latter further allows constructs such as nominals and number restrictions). We characterize the difficulty of subsumption reasoning in DLs using a graph-like structure called a decomposition. Furthermore, we present a family of reasoning algorithms based on decompositions that facilitate FPT reasoning. Finally, we show how the good performance of consequence-based reasoners on several widely used ontologies can be explained by an analysis of their decompositions. We next summarize the main contributions of our paper.

In Section 3 we introduce a general consequence-based framework for subsumption reasoning in $\mathcal{A L C I}$ and $\mathcal{S H}$, and then we discuss possibilities for instantiating the framework using different initialization and expansion strategies. Our algorithms are not only refutationally complete: they can actively derive queries (i.e., axioms of a specific form) that logically follow from a given ontology. A specific instantiation of our framework captures the algorithm used in CB [21] and is closely related to the algorithm used in ConDOR [22]. As well as capturing existing consequence-based approaches, our framework incorporates redundancy deletion and optimized inference rules that have the potential to further improve the practical performance of consequence-based algorithms. Thus, we see our framework as an important contribution in its own right.

In Section 4 we introduce the central notion of a decomposition $\mathcal{D}$ of an ontology $\mathcal{O}$ and a set of queries $\mathcal{Q}$. Roughly speaking, $\mathcal{D}$ is a graph-like structure that "summarizes" the models of $\mathcal{O}$ relevant for answering the queries in $\mathcal{Q}$; each vertex of $\mathcal{D}$ identifies a propositional subproblem, and each edge of $\mathcal{D}$ identifies a pathway for the exchange of information between such subproblems. We then show how to instantiate our consequence-based framework from Section 3 so that it can be applied to $\mathcal{D}$, and we identify the length and width of $\mathcal{D}$ as parameters that characterize the and- and or-branching encountered in this process. Finally, we show that the algorithm can be implemented so that it is fixed-parameter tractable w.r.t. decomposition length and width.

In Section 4.3 we show that decompositions also explain to an extent the difficulty of tableau reasoning. In particular, decomposition width bounds the size of sets of concepts annotating tableau vertices, so length and width bound the size of the constructed model representations. We also point out, however, that tableau algorithms are not FPT because the way in which they construct model representations may result in redundant computations.

In Section 4.4 we present an algorithm for computing a decomposition of $\mathcal{O}$ and $\mathcal{Q}$. One would naturally want to compute a decomposition of smallest length and width; however, in Section 7 we show that there is a tension between these two parameters

[^1]by exhibiting a family of ontologies for which each decomposition of minimum width has exponential length. To address this problem, our decomposition construction algorithm is parameterized by a collection of parameters called a control that imposes an upper bound on decomposition length so that a unique decomposition of $\mathcal{O}$ and $\mathcal{Q}$ can be computed in polynomial time.

In Section 5 we further extend our results using ideas from reasoning with propositional problems of bounded treewidth. More specifically, we present the notion of $\epsilon$-refinement of a decomposition $\mathcal{D}$, which, roughly speaking, is obtained by replacing each vertex of $\mathcal{D}$ with a tree decomposition of the propositional problem corresponding to the vertex. This can reduce the width of a decomposition while increasing the length only by a linear factor, so it can reduce the complexity of reasoning. As for tree decompositions, computing an $\epsilon$-refinement of bounded width (if one exists) is FPT.

In Section 6 we present several soundness- and admissibility-preserving transformations for decompositions that can be used to eliminate redundant information. In Section 7 we then establish bounds on the sizes of the decompositions of $\mathcal{O}$ and $\mathcal{Q}$. For the lower bound, we show that ontologies exist for which all decompositions of minimal width have exponential length; for the upper bound, we use the transformations from Section 6 to show that $\mathcal{O}$ and $\mathcal{Q}$ always admit a decomposition of minimal width with at most exponential length.

In Section 8 we investigate the width and length of several ontologies commonly used in practice. Our results show that almost all ontologies admit a decomposition of length about twice the number of concepts in the ontology, and of width below 30. These results, we believe, explain the good performance of consequence-based reasoning algorithms in practice.

To the best of our knowledge, these are the first results on quantitative analysis and fixed-parameter complexity of ontology reasoning. The only related work we are aware of is the analysis of conditions under which query answering for a given ontology (independently of the DL language) becomes tractable w.r.t. data complexity [27]; however, these results are not quantitative and they do not address the problem of subsumption reasoning. Thus, we believe that this paper provides a completely new perspective on the problem of identifying the sources of complexity in ontology reasoning. Moreover, our results are amenable to practical implementation, and we believe they can significantly improve the scalability and robustness of reasoners.

## 2. Preliminaries

In this section we recapitulate some well-known definitions and introduce the notation that we use in the rest of our paper.

### 2.1. Description Logics

Description logics $\mathcal{A L C I}$ and $\mathcal{S H \mathcal { I }}$ are defined w.r.t. a fixed signature $\Sigma=\left\langle\Sigma_{A}, \Sigma_{T}\right\rangle$, where $\Sigma_{A}$ and $\Sigma_{T}$ are disjoint sets of atomic concepts and atomic roles, respectively.

An inverse role is an expression $T^{-}$for $T$ an atomic role. A role is an atomic role or an inverse role, and $\Sigma_{R}$ is the set of all roles. For an atomic role $T$, let $\operatorname{inv}(T):=T^{-}$and $\operatorname{inv}\left(T^{-}\right):=T$.

The set of concepts is the smallest set that contains $\top$ (the top concept), $\perp$ (the bottom concept), each atomic concept $A$, and $\neg D$ (negation), $D_{1} \sqcap D_{2}$ (conjunction), $D_{1} \sqcup D_{2}$ (disjunction), $\exists R . D$ (existential restriction), and $\forall R . D$ (universal restriction) for each role $R$ and all concepts $D, D_{1}$, and $D_{2}$. We assume that $\{\top, \perp\} \cap \Sigma_{A}=\emptyset$-that is, $T$ and $\perp$ are not atomic concepts. A literal is a concept of the form $A, \exists R . A$, or $\forall R . A$, for $A$ an atomic concept and $R$ a role, and $\Sigma_{L}$ is the set of all literals. For $\mathbf{L}$ a set of literals, $\mathbf{L}^{\exists}$ is the subset of $\mathbf{L}$ containing all literals of the form $\exists R . A$.

Unless otherwise stated, certain letters of the alphabet (possibly with sub- and/or superscripts) are used consistently throughout this paper as specified below, so we do not qualify these letters whenever no ambiguity can arise:

- letters $A, B$, and $C$ denote atomic concepts,
- letter $D$ denotes a concept,
- letter $L$ denotes a literal,
- letter $T$ denotes an atomic role,
- letters $R$ and $S$ denote roles,
- letter $K$ denotes a conjunction of literals, and
- letter $M$ denotes a disjunction of literals.

We identify conjunctions and disjunctions of literals with sets of literals (i.e., conjunctions and disjunctions of literals are unordered and without repeated literals) and we use them in standard set operations; furthermore, we identify the empty conjunction and the empty disjunction with $\top$ and $\perp$, respectively. Conversely, for $U$ a set of literals, $\sqcap U$ and $\bigsqcup U$ are the conjunction and the disjunction, respectively, of the literals in $U$; furthermore, to simplify the notation, in some cases we will explicitly note that we abbreviate $\sqcap U$ as just $U$. The cardinality of a set $N$ is written $|N|$.

An axiom is an expression of the form $D_{1} \sqsubseteq D_{2}$ (general concept inclusion), $R_{1} \sqsubseteq R_{2}$ (role inclusion), or $\operatorname{Tra}(R)$ (role transitivity axiom). A $\mathcal{S H \mathcal { I }}$ ontology $\mathcal{O}$ is a set of axioms. In this paper we consider only terminological reasoning, so we do not allow for facts (aka ABoxes) in ontologies. The size of $\mathcal{O}$, written $\|\mathcal{O}\|$, is the number of symbols in $\mathcal{O}$.

```
            Interpretation of Roles and Concepts
            \(\left(T^{-}\right)^{\mathcal{I}}=\left\{\left\langle\delta_{2}, \delta_{1}\right\rangle \in \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}} \mid\left\langle\delta_{1}, \delta_{2}\right\rangle \in T^{\mathcal{I}}\right\}\)
            \(\mathrm{T}^{\mathcal{I}}=\Delta^{\mathcal{I}}\)
            \(\perp^{\mathcal{I}}=\emptyset\)
        \((\neg D)^{\mathcal{I}}=\Delta^{\mathcal{I}} \backslash D^{\mathcal{I}}\)
\(\left(D_{1} \sqcap D_{2}\right)^{\mathcal{I}}=D_{1}^{\mathcal{I}} \cap D_{2}^{\mathcal{I}}\)
\(\left(D_{1} \sqcup D_{2}\right)^{\mathcal{I}}=D_{1}^{\mathcal{I}} \cup D_{2}^{\mathcal{I}}\)
    \((\exists R . D)^{\mathcal{I}}=\left\{\delta_{1} \in \Delta^{\mathcal{I}} \mid \exists \delta_{2} \cdot\left\langle\delta_{1}, \delta_{2}\right\rangle \in R^{\mathcal{I}} \wedge \delta_{2} \in D^{\mathcal{I}}\right\}\)
    \((\forall R . D)^{\mathcal{I}}=\left\{\delta_{1} \in \Delta^{\mathcal{I}} \mid \forall \delta_{2} .\left\langle\delta_{1}, \delta_{2}\right\rangle \in R^{\mathcal{I}} \rightarrow \delta_{2} \in D^{\mathcal{I}}\right\}\)
```

|  | Satisfaction of Axioms in an Interpretation |
| :--- | :--- |
| $\mathcal{I} \vDash D_{1} \sqsubseteq D_{2}$ | if $\left(D_{1}\right)^{\mathcal{I}} \subseteq\left(D_{2}\right)^{\mathcal{I}}$ |
| $\mathcal{I} \vDash R_{1} \sqsubseteq R_{2}$ | if $\left(R_{1}\right)^{\mathcal{I}} \subseteq\left(R_{2}\right)^{\mathcal{I}}$ |
| $\mathcal{I} \vDash \operatorname{Tra}(R)$ | if $\forall \delta_{1} \forall \delta_{2} \forall \delta_{3} .\left\langle\delta_{1}, \delta_{2}\right\rangle \in R^{\mathcal{I}} \wedge\left\langle\delta_{2}, \delta_{3}\right\rangle \in R^{\mathcal{I}} \rightarrow\left\langle\delta_{1}, \delta_{3}\right\rangle \in R^{\mathcal{I}}$ |

Ontologies are interpreted using Tarski-style semantics. An interpretation $\mathcal{I}=\left\langle\Delta^{\mathcal{I}}, \mathcal{I}^{\mathcal{I}}\right\rangle$ is a pair where $\Delta^{\mathcal{I}}$ is a nonempty set called the domain, and ${ }^{{ }^{\mathcal{I}}}$ is a function that assigns a set $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$ to each atomic concept $A \in \Sigma_{A}$, and a set $T^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$ to each atomic role $T \in \Sigma_{T}$. Function ${ }^{\mathcal{I}}$ is extended to roles and concepts as shown in the upper part of Table 1. The satisfaction of an axiom $\alpha$ in $\mathcal{I}$, written $\mathcal{I} \vDash \alpha$, is defined as shown in the lower part of Table 1. Interpretation $\mathcal{I}$ is a model of an ontology $\mathcal{O}$, written $\mathcal{I} \vDash \mathcal{O}$, if $\mathcal{I}$ satisfies all axioms in $\mathcal{O}$. An ontology is $\mathcal{O}$ satisfiable if a model of $\mathcal{O}$ exists; furthermore, $\mathcal{O}$ entails an axiom $\alpha$, written $\mathcal{O} \vDash \alpha$, if $\mathcal{I} \vDash \alpha$ for each model $\mathcal{I}$ of $\mathcal{O}$. Note that $\mathcal{O}$ is unsatisfiable if and only if $\mathcal{O} \vDash \top \sqsubseteq \perp$.

A clause is a general concept inclusion of the form $\prod_{i=1}^{m} L_{i} \sqsubseteq \bigsqcup_{i=m+1}^{n} L_{i}$ where $0 \leq m \leq n$ and each $L_{i}$ is a literal; the clause is normal if each $L_{i}$ with $1 \leq i \leq m$ is an atomic concept; and the clause is a query if each $L_{i}$ with $m+1 \leq i \leq n$ is an atomic concept. Conjunction $K$ in a clause $K \sqsubseteq M$ is the antecedent, and disjunction $M$ is the consequent. An ontology $\mathcal{O}$ is normalized if each general concept inclusion in $\mathcal{O}$ is a normal clause. Each $\mathcal{S H}$ I ontology $\mathcal{O}$ can be transformed in linear time to a normalized $\mathcal{S H} \mathcal{I}$ ontology $\mathcal{O}^{\prime}$ such that $\mathcal{O}^{\prime}$ is a conservative extension of $\mathcal{O}$. Several such transformations are well known (see, e.g., [21, 14]), so we omit the details for brevity. Given a set of literals $\mathbf{L}$, a clause $K \sqsubseteq M$ is over $\mathbf{L}$ if $K \cup M \subseteq \mathbf{L}$. A clause $K^{\prime} \sqsubseteq M^{\prime}$ is a strengthening of a clause $K \sqsubseteq M$ if $K^{\prime} \subseteq K$ and $M^{\prime} \subseteq M ;{ }^{3}$ furthermore, $K \sqsubseteq M \hat{\in} \mathcal{N}$ means that a set of clauses $\mathcal{N}$ contains at least one strengthening of $K \sqsubseteq M$.

We consider several fragments of $\mathcal{S H}$, which, for simplicity, we define only for normalized ontologies. Let $\mathcal{O}$ be a normalized $\mathcal{S H I}$ ontology. Due to normalization, $\mathcal{O}$ cannot contain axioms of the form $\exists R . A \sqcap K \sqsubseteq M$, so the following definitions of various fragments are not standard; however, it is straightforward to see that they correspond to the conventional definitions used in the literature. Furthermore, note that $\exists R$. $T$ is not a literal since we do not consider $T$ to be an atomic concept; however, one can always replace $T$ with a fresh atomic concept $A_{\top}$ and axiomatize the latter using a clause $\top \sqsubseteq A_{\top}$. Now $\mathcal{O}$ is Horn if the consequent of each clause in $\mathcal{O}$ contains at most one literal; $\mathcal{O}$ is an $\mathcal{A L C I}$ ontology if each axiom in $\mathcal{O}$ is a general concept inclusion; $\mathcal{O}$ is an $\mathcal{E L}$ ontology if $\mathcal{O}$ is a Horn- $\mathcal{A L C I}$ ontology, each role occurring in $\mathcal{O}$ in an existential restriction is an atomic role, and each role occurring in $\mathcal{O}$ in a universal restriction is an inverse role; and $\mathcal{O}$ is a DL-Lite horn ontology if $\mathcal{O}$ is a Horn- $\mathcal{A L C I}$ ontology and universal restrictions occur in $\mathcal{O}$ only in axioms of the form $T \sqsubseteq \forall R . B$.

A normalized $\mathcal{S H I}$ ontology $\mathcal{O}$ can be transformed into a normalized $\mathcal{A L C I}$ ontology by eliminating all role inclusions and role transitivity axioms. Towards this goal, let $\sqsubseteq_{\mathcal{O}}$ be the smallest reflexive and transitive binary relation on roles such that $R \sqsubseteq_{\mathcal{O}} S$ and $\operatorname{inv}(R) \sqsubseteq_{\mathcal{O}} \operatorname{inv}(S)$ for each role inclusion $R \sqsubseteq S \in \mathcal{O}$; furthermore, for each role $R$ and each atomic concept $C$ occurring in $\mathcal{O}$, let $A_{R, C}$ and $B_{R, C}$ be fresh atomic concepts unique for $R$ and $C$. The elimination proceeds in three steps: step 1 ensures that universal restrictions occur only in clauses of the form $A \sqsubseteq \forall R . C$, step 2 encodes transitivity axioms using the wellknown "box pushing" method, and step 3 eliminates role inclusions by expanding the role hierarchy into universal restrictions. We capture these steps by defining three ontologies obtained by transforming $\mathcal{O}$ as follows.

1. For each clause $K \sqsubseteq M \in \mathcal{O}$, ontology $\mathcal{O}_{1}$ contains the clause obtained from $K \sqsubseteq M$ by replacing each literal $\forall R . C \in M$ with $A_{R, C}$; furthermore, for each literal $\forall R . C$ occurring in a clause in $\mathcal{O}$, ontology $\mathcal{O}_{1}$ contains the axiom $A_{R, C} \sqsubseteq \forall R . C$.
[^2]| Hyp-rule | If $\quad A_{1} \sqcap \ldots \sqcap A_{m} \sqsubseteq L_{1} \sqcup \ldots \sqcup L_{n} \in \mathcal{O}$, and an individual $s \in \operatorname{ind}\left(\mathcal{A}_{0}\right)$ exists that is not indirectly blocked such that $A_{i}(s) \in \mathcal{A}_{0}$ for each $1 \leq i \leq m$ and $L_{j}(s) \notin \mathcal{A}_{0}$ for each $1 \leq j \leq n$, then $\mathcal{A}_{1}:=\mathcal{A}_{0} \cup\{\perp\}$ if $n=0$; and $\mathcal{A}_{j}:=\mathcal{A}_{0} \cup\left\{L_{j}(s)\right\}$ for $1 \leq j \leq n$ if $n>0$. |
| :---: | :---: |
| ヨ-rule | If an individual $s$ exists that is not blocked such that $\exists R . A(s) \in \mathcal{A}_{0}$ and no individual $t$ exists such that $\{R(s, t), A(t)\} \subseteq \mathcal{A}_{0}$, then $\mathcal{A}_{1}:=\mathcal{A}_{0} \cup\{R(s, u), A(u)\}$ for $u$ a fresh successor of $s$. |
| $\forall^{+}$-rule | If individuals $s$ and $t$ exist that are not indirectly blocked such that $\{\forall R . B(s), R(s, t)\} \subseteq \mathcal{A}_{0}$ and $B(t) \notin \mathcal{A}_{0}$, then $\mathcal{A}_{1}:=\mathcal{A}_{0} \cup\{B(t)\}$. |
| $\forall^{-}$-rule | If individuals $s$ and $t$ exist that are not indirectly blocked such that $\{[\forall \operatorname{inv}(R) . B](t), R(s, t)\} \subseteq \mathcal{A}_{0}$ and $B(s) \notin \mathcal{A}_{0}$, then $\mathcal{A}_{1}:=\mathcal{A}_{0} \cup\{B(s)\}$. |

2. Ontology $\mathcal{O}_{2}$ contains each clause in $\mathcal{O}_{1}$; furthermore, for each clause $A \sqsubseteq \forall R . C \in \mathcal{O}_{1}$ and each role $S$ such that $S \sqsubseteq_{\mathcal{O}} R$ and either $\operatorname{Tra}(S) \in \mathcal{O}$ or $\operatorname{Tra}(\operatorname{inv}(S)) \in \mathcal{O}$, ontology $\mathcal{O}_{2}$ contains clauses

$$
A \sqsubseteq \forall S . B_{S, C} \quad \text { and } \quad B_{S, C} \sqsubseteq \forall S . B_{S, C} \quad \text { and } \quad B_{S, C} \sqsubseteq C .
$$

3. Ontology $\mathcal{O}_{3}$ contains each clause in $\mathcal{O}_{2}$; furthermore, for each clause $A \sqsubseteq \forall R . C \in \mathcal{O}_{2}$ and each role $S$ such that $S \sqsubseteq_{\mathcal{O}} R$, ontology $\mathcal{O}_{3}$ contains clause $A \sqsubseteq \forall S . C$.

For each clause $K \sqsubseteq M$ where $K$ and $M$ contain only atomic concepts, $\mathcal{O} \vDash K \sqsubseteq M$ if and only if $\mathcal{O}_{3} \vDash K \sqsubseteq M$; this can be proved using a straightforward adaptation of the proofs by Demri and de Nivelle [28].

The algorithms presented in this paper take a normalized $\mathcal{A L C I}$ ontology $\mathcal{O}$ and a set of queries $\mathcal{Q}$, and they compute the status of $\mathcal{O} \vDash K \sqsubseteq M$ for each query $K \sqsubseteq M \in \mathcal{Q}$; hence, our algorithms compute functions rather than solving decision problems. For our complexity results, however, we consider a related decision problem of determining whether $\mathcal{O} \vDash K \sqsubseteq M$ holds for each query $K \sqsubseteq M \in \mathcal{Q}$-that is, the decision problem checks the conjunction of all such entailments.

### 2.2. The Hypertableau Algorithm

Most OWL 2 DL reasoners are currently based on variants of tableau algorithms. To relate our techniques to the state of the art in ontology reasoning, we next present a variant of the hypertableau algorithm by Motik et al. [14].

We assume the existence of a countably infinite set of named individuals. An individual is a finite string of the form $a . i_{1} \ldots i_{n}$ where $a$ is a named individual and $i_{1} \ldots i_{n}$ is a possibly empty sequence of integers. An individual of the form s. $i$ is unnamed, and it is a successor of individual $s$; predecessor is the inverse of successor; and ancestor and descendant are transitive closures of predecessor and successor, respectively. An ABox $\mathcal{A}$ is a finite set of facts of the form $\perp, L(s)$, or $R(s, t)$, for $L$ a literal, $R$ a role, and $s$ and $t$ individuals; furthermore, $\operatorname{ind}(\mathcal{A})$ is the set of individuals occurring in $\mathcal{A}$.

Termination of the algorithm is ensured via anywhere blocking; for $\mathcal{A L C \mathcal { L }}$, the single anywhere variant suffices. Let $\triangleleft$ be an arbitrary strict order (i.e., an irreflexive and transitive relation) on the set of all individuals compatible with the ancestor relation (i.e., $s \triangleleft t$ holds for all individuals $s$ and $t$ such that $s$ is an ancestor of $t$ ). The label of an individual $s$ in an ABox $\mathcal{A}$ is defined as $\mathcal{L}_{\mathcal{A}}(s)=\{L \mid L(s) \in \mathcal{A}\}$. By induction on $\triangleleft$, each individual $s$ occurring in $\mathcal{A}$ is assigned a status as follows:

- $s$ is directly blocked by an individual $s^{\prime}$ in $\mathcal{A}$ if both $s$ and $s^{\prime}$ are unnamed, no ancestor of $s^{\prime}$ is blocked, $\mathcal{L}_{\mathcal{A}}(s)=\mathcal{L}_{\mathcal{A}}\left(s^{\prime}\right)$, and $s^{\prime} \triangleleft s$;
- $s$ is indirectly blocked if $s$ is unnamed and some ancestor of $s$ is directly blocked; and
- $s$ is blocked if it is directly or indirectly blocked.

Given a normalized $\mathcal{A L C I}$ ontology $\mathcal{O}$ and an $\operatorname{ABox} \mathcal{A}$ containing only named individuals, the hypertableau algorithm constructs a derivation $D=(V, E, \rho)$ for $\mathcal{O}$ and $\mathcal{A}$, where $V$ and $E$ are the vertices and edges of a finite tree, and $\rho$ labels each vertex with an ABox. A vertex $v \in V$ is closed if $\perp \in \rho(v)$; otherwise, $v$ is open. Labeling $\rho$ must satisfy the following conditions:

- $\rho\left(v_{0}\right)=\mathcal{A}$ for $v_{0}$ the root of the tree;
- vertex $v_{0}$ is a leaf if $v_{0}$ is closed or no derivation rule from Table 2 is applicable to $\mathcal{A}_{0}=\rho\left(v_{0}\right)$;
- in all other cases, the children $v_{1}, \ldots, v_{n}$ of vertex $v_{0}$ are labeled by ABoxes $\mathcal{A}_{1}, \ldots, \mathcal{A}_{n}$ obtained by applying one (arbitrarily chosen) derivation rule from Table 2 to $\mathcal{A}_{0}=\rho\left(v_{0}\right)$.

The algorithm is sound and complete: for each derivation $D$, we have that $\mathcal{O} \cup \mathcal{A}$ is satisfiable if and only if $D$ contains an open leaf. Thus, checking whether $\mathcal{O} \vDash A \sqsubseteq B$ holds can be solved using the hypertableau algorithm by constructing a derivation for $\mathcal{O} \cup\{C \sqsubseteq A, B \sqcap C \sqsubseteq \perp\}$ and $\{C(a)\}$ where $C$ is a fresh atomic concept, and then checking whether each derivation leaf is closed.

Due to inverse roles, blocking is dynamic: an individual can change its blocking status more than once as the inference rules are applied, which is why individuals can be indirectly blocked. This, in turn, ensures that the algorithm runs in N2ExpTime: the algorithm nondeterministically constructs a tree of individuals that can have exponential depth and a linear branching factor. Hence, although the number of nonblocked and directly blocked individuals is at most exponential, the number of indirectly blocked individuals in the tree can be doubly exponential; Motik et al. [14] discuss these issues in more detail.

### 2.3. Fixed-Parameter Tractable Problems

We now recapitulate several definitions from parameterized complexity [24]. A parameterized problem is a set $P \subseteq \Sigma^{*} \times \mathbb{N}$, where $\Sigma$ is a finite alphabet and $\mathbb{N}$ is the set of nonnegative integers. Each pair $\langle x, k\rangle \in \Sigma^{*} \times \mathbb{N}$ is called a problem instance; $x$ is called the input; $k$ is called the parameter; and $\|x\|$ is the length of $x$. Problem $P$ is fixed-parameter tractable (FPT) if a computable function $f: \mathbb{N} \rightarrow \mathbb{N}$, a constant $c$, and an algorithm exist such that, given an arbitrary pair $\langle x, k\rangle \in \Sigma^{*} \times \mathbb{N}$, the algorithm decides $\langle x, k\rangle \in P$ in at most $f(k) \cdot\|x\|^{c}$ steps. FPT is the class of all fixed-parameter tractable problems.

Some authors use an alternative definition of parameterized complexity. Just like a regular problem, a parameterized problem $P$ is defined as a subset of $\Sigma^{*}$. Furthermore, a parameterization is a function $\kappa: \Sigma^{*} \rightarrow \mathbb{N}$. A problem $P$ is fixed-parameter tractable w.r.t. a parameterization $\kappa$ if a computable function $f: \mathbb{N} \rightarrow \mathbb{N}$, a constant $c$, and an algorithm exist such that, for an arbitrary $x \in \Sigma^{*}$, the algorithm decides $x \in P$ in at most $f(\kappa(x)) \cdot\|x\|^{c}$ steps. The two definitions are clearly equivalent if, for each $x \in \Sigma^{*}$, value $\kappa(x)$ is computable in polynomial time. Furthermore, even if the latter is not the case, the two definitions coincide if a computable function $g$ and a constant $d$ exist such that, for each fixed $k$, one can check whether $\kappa(x) \leq k$ holds using at most $g(k) \cdot\|x\|^{d}$ steps; in other words, checking $\kappa(x) \leq k$ must be fixed-parameter tractable as well. The results we present in this paper involve the computation of a tree decomposition of a fixed width, which satisfies the latter condition.

### 2.4. Tree Decompositions and Treewidth of Hypergraphs and Propositional Problems

We recapitulate the notions of tree decompositions and treewidth [25], which were extensively used to obtain FPT results for a number of intractable graph-theoretic problems. For convenience, we extend these notions to hypergraphs by treating each hyperedge as the clique of all of its vertices, but this does not alter the definitions in any significant way.

Given a finite set $N$ of nodes, an $N$-hypergraph $\mathcal{H}$ is a subset of $2^{N}$; the elements of $\mathcal{H}$ are called hyperedges. A tree decomposition of $\mathcal{H}$ is a tuple $\mathcal{T}=\langle\mathcal{V}, \mathcal{E}, \mathcal{L}\rangle$ where $\mathcal{V}$ and $\mathcal{E}$ are sets of vertices and edges, respectively, of an undirected tree, and $\mathcal{L}: \mathcal{V} \rightarrow 2^{N}$ is a labeling of vertices with subsets of $N$ such that the following conditions hold:
(T1) for each hyperedge $h \in \mathcal{H}$, a vertex $v \in \mathcal{V}$ exists such that $h \subseteq \mathcal{L}(v)$; and
(T2) for each element $n \in N$, the set $\{v \in \mathcal{V} \mid n \in \mathcal{L}(v)\}$ is connected in $\mathcal{E}$.
The width of $\mathcal{T}$ is $\operatorname{wd}(\mathcal{T}):=\max _{v \in \mathcal{V}}|\mathcal{L}(v)|$, and the treewidth of $\mathcal{H}$ is the minimum width over all tree decompositions of $\mathcal{H}$. In related literature width is actually defined as $\operatorname{wd}(\mathcal{T}):=\max _{v \in \mathcal{V}}|\mathcal{L}(v)|-1$ in order to ensure that tree-like hypergraphs have width one, but we drop the term -1 from our definition for uniformity with Section 4 . For a fixed integer mwd, a tree decomposition of width at most mwd can be computed in linear time [29]. This result was originally formulated for tree decompositions of graphs; however, by treating hyperedges as cliques, this result can also be straightforwardly applied to our definitions.

Lemma 1 ([29]). Let $\mathcal{H}$ be an $N$-hypergraph, and let $m w d$ be an integer. One can compute a tree decomposition of $\mathcal{H}$ of width at most mwd , or determine that such a tree decomposition does not exist, in time $O(f(\mathrm{mwd}) \cdot|N|)$ where $f$ is a computable function. The resulting tree decomposition has at most $|\mathcal{H}|$ vertices.

We next recapitulate how these notions can be used to obtain an FPT procedure for propositional satisfiability [26]. To unify the notation throughout this paper, we write propositional clauses as implications $K \sqsubseteq M$ where $K$ is a conjunction of atoms, and $M$ is a disjunction of atoms; furthermore, we consider a propositional interpretation $J$ to be a set of atoms; finally, we write $J \vDash K \sqsubseteq M$ if $K \subseteq J$ implies $M \cap J \neq \emptyset$.

Let $\mathcal{N}$ be a set of propositional clauses. Then, we define $\mathcal{H}$ as the hypergraph that contains a hyperedge $K \cup M \in \mathcal{H}$ for each clause $K \sqsubseteq M \in \mathcal{N}$. Furthermore, a tree decomposition and the treewidth of $\mathcal{N}$ are defined as a tree decomposition and the treewidth, respectively, of $\mathcal{H}$. Now assume that the treewidth of $\mathcal{N}$ is bounded by some integer mwd. The following procedure checks the satisfiability of $\mathcal{N}$ and is fixed-parameter tractable w.r.t. mwd.

1. Compute a tree decomposition $\mathcal{T}=\langle\mathcal{V}, \mathcal{E}, \mathcal{L}\rangle$ of $\mathcal{H}$ of width at most mwd. By Lemma 1, this step is FPT w.r.t. mwd.
2. Associate with each vertex $v \in \mathcal{V}$ a set $\mathcal{P}_{v}$ of propositional interpretations as follows: set $\mathcal{P}_{v}:=2^{\mathcal{L}(v)}$, and then eliminate each $J \in \mathcal{P}_{v}$ for which a clause $K \sqsubseteq M \in \mathcal{N}$ exists such that $K \cup M \subseteq \mathcal{L}(v)$ and $J \not \equiv K \sqsubseteq M$. Since $|\mathcal{L}(v)| \leq$ mwd, this step requires time polynomial in $2^{\text {mwd }}$ and $|\mathcal{N}|$, and is thus FPT w.r.t. mwd.
3. Using a bottom-up approach starting from the leaves of $\mathcal{T}$, for each non-root vertex $v$ and its parent $v^{\prime}$, eliminate from $\mathcal{P}_{v^{\prime}}$ each $J^{\prime} \in \mathcal{P}_{v^{\prime}}$ for which no $J \in \mathcal{P}_{v}$ exists such that $J \cap \mathcal{L}(v) \cap \mathcal{L}\left(v^{\prime}\right)=J^{\prime} \cap \mathcal{L}(v) \cap \mathcal{L}\left(v^{\prime}\right)$. We can achieve this by iterating over all pairs of $J^{\prime} \in \mathcal{P}_{v^{\prime}}$ and $J \in \mathcal{P}_{v}$; since $\mathcal{T}$ has at most $|\mathcal{N}|$ vertices by Lemma 1, this step is FPT w.r.t. mwd.
4. Set $\mathcal{N}$ is unsatisfiable if and only if $\mathcal{P}_{r}=\emptyset$, where $r$ is the root of $\mathcal{T}$.

To see why this algorithm is correct, assume that $\mathcal{P}_{r} \neq \emptyset$; we can then construct an interpretation $J$ for $\mathcal{N}$ as follows. Select an arbitrary propositional interpretation $J_{r} \in \mathcal{P}_{r}$; since $J_{r}$ has not been deleted in step 3 , for $v_{1}, \ldots, v_{k}$ the children of $r$, we can select interpretations $J_{v_{1}}, \ldots, J_{v_{k}}$ such that $J_{r} \cap \mathcal{L}(r) \cap \mathcal{L}\left(v_{i}\right)=J_{v_{i}} \cap \mathcal{L}(r) \cap \mathcal{L}\left(v_{i}\right)$ holds for each $1 \leq i \leq k$. By inductively repeating this argument from the root to the leaves, we associate with each vertex $v \in \mathcal{V}$ an interpretation $J_{v}$, and then we define $J=\bigcup_{v \in \mathcal{V}} J_{v}$. Now since $\mathcal{T}$ satisfies properties (T1) and (T2), it is straightforward to see that $J \vDash \mathcal{N}$. The converse direction can be shown analogously, and we omit it for the sake of brevity. Thus, the treewidth of $\mathcal{N}$ can be taken as an indicator of the "hardness" of $\mathcal{N}$, and it shows the size of the sets of atomic concepts that we need to consider to decide the satisfiability of $\mathcal{N}$.

## 3. Consequence-Based Framework for $\mathcal{A L C I}$

In this section we present our consequence-based framework for solving the following reasoning problem: given a normalized $\mathcal{A L C I}$ ontology $\mathcal{O}$ and a finite set of queries $\mathcal{Q}$, determine whether $\mathcal{O} \vDash K \sqsubseteq M$ holds for each query $K \sqsubseteq M \in \mathcal{Q}$. Note that this problem captures ontology classification, which involves only queries of the form $A \sqsubseteq B$ for $A$ and $B$ atomic concepts from $\mathcal{O}$; furthermore, by using the transformations from Section 2, the framework can also be applied to $\mathcal{S H \mathcal { I }}$ ontologies. In Section 3.1 we discuss the intuitions behind our framework, in Section 3.2 we introduce the framework formally, in Section 3.3 we discuss certain redundancy elimination techniques that can be used to optimize reasoning, and finally in Section 3.4 we discuss several concrete instantiations of our framework.

### 3.1. Intuitions

In order to capture the essence of the known similar algorithms, our framework can be instantiated in many different ways by varying several parameters. In this section we describe the framework at a high level, and we discuss possible parameterizations in Sections 3.3 and 3.4. In the rest of this section, we fix the ontology $\mathcal{O}$ and the query $q$ as specified in Example 2, and then we show how one can prove that $\mathcal{O} \vDash q$ holds using our framework.

Example 2. Let $\mathcal{O}$ be the ontology consisting of clauses (1)-(10), and let $q=A \sqsubseteq G$; one can readily verify that $\mathcal{O} \vDash q$.

$$
\begin{array}{rlrr}
A \sqsubseteq A_{i} & \text { for } 1 \leq i \leq n & \text { (1) } & C \sqsubseteq \exists R . B \\
A_{i} \sqsubseteq C_{i} & \text { for } 1 \leq i \leq n & \text { (2) } & C \sqsubseteq \forall R . D \\
B \sqsubseteq B_{i} & \text { for } 1 \leq i \leq n & \text { (3) } & B \sqcap D \sqsubseteq E \sqcup \forall R^{-} . G \\
B_{i} \sqsubseteq C_{i} & \text { for } 1 \leq i \leq n & \text { (4) } & E \sqsubseteq \forall R^{-} . F \\
C_{1} \sqcap \ldots \sqcap C_{n} & \sqsubseteq C & & \text { (5) }
\end{array}
$$

Resolution decision procedures [30] can decide whether $\mathcal{O} \vDash A \sqsubseteq G$ holds by transforming $\mathcal{O}$ into a set of clauses, adding clauses $A(a)$ and $\neg G(a)$ obtained from the negation of the theorem $A \sqsubseteq G$ that is to be proved, and then saturating the result using a suitable first-order resolution variant [31]. Such algorithms are typically worst-case optimal, and they can solve many practically-relevant problems; however, on complex ontologies they can easily run into a combinatorial explosion [32]. In our example, from clauses (2)-(5) resolution can derive $3^{n}$ clauses of the form $L_{1} \sqcap \ldots \sqcap L_{n} \sqsubseteq C$ with $L_{i} \in\left\{A_{i}, B_{i}, C_{i}\right\}$, thus covering all possible combinations of atomic concepts that might be relevant. In contrast, when applied to $\mathcal{O}$, the hypertableau algorithm from Section 2.2 does not run into this problem: the algorithm is initialized using the fact $A(a)$, and then it derives $A_{i}(a)$ and $C_{i}(a)$ for each $1 \leq i \leq n$ and $C(a)$; thus, the algorithm does not consider the "irrelevant" combinations of $A_{i}, B_{i}$, and $C_{i}$.

The hypertableau algorithm can thus be seen as being more "goal directed" than resolution. However, as we discussed in Section 2.2, the hypertableau algorithm can construct a very large tree of individuals most of which are indirectly blocked, and can thus perform a lot of redundant computation since each indirectly-blocked individual is a "copy" of a nonblocked or a directly blocked individual. Resolution is not susceptible to such problems: clauses introduced by resolution are universally quantified and are (unlike individuals in the tableau algorithm) not localized to a specific part of a model. Resolution can thus describe all relevant combinations of atomic concepts using exponentially many clauses, thereby avoiding redundant computation.

Our consequence-based algorithm can be understood as a hybrid between resolution and the hypertableau algorithm. It does not explicitly construct a model; instead, it uses resolution to compute clauses that describe a model $\mathcal{I}$ of $\mathcal{O}$ refuting the relevant

|  | A | B, D |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  | Succ[15+16]: | $\exists R . B \quad$ (17) |  | ucc[23+24]: | $\exists$ R.B |
| Initialization: | T¢ ${ }^{\text {a }}$ |  | (11) | Succ[15+16]: | T¢ $B$ |  | (18) |
| Hyper[1+11]: | T¢ $A_{i}$ | for $1 \leq i \leq n$ | (12) | Succ[15+16]: | T¢ $D$ |  | (19) |
| Hyper[2+12]: | T $\subseteq C_{i}$ | for $1 \leq i \leq n$ | (13) | Hyper[3+18]: | $\top \sqsubseteq B_{i}$ | for $1 \leq i \leq n$ | (20) |
| Hyper[5+13]: | T¢ $C$ |  | (14) | Hyper[4+20]: | $\top \sqsubseteq C_{i}$ | for $1 \leq i \leq n$ | (21) |
| Hyper[6+14]: | $\top \sqsubseteq \exists R . B$ |  | (15) | Hyper[5+21]: | T¢ $C$ |  | (22) |
| Hyper[7+14]: | $\mathrm{T} \subseteq \forall R . D$ |  | (16) | Hyper[6+22]: | T¢ $\preceq$ R. $B$ |  | (23) |
| Pred[15+17+26]: | T巨 $F \sqcup G$ |  | (28) | Hyper[7+22]: | $T \sqsubseteq \forall R . D$ |  | (24) |
| Hyper[10+11+28]: | T¢ $G$ |  | (30) | Hyper[8+18+19]: | $\top \sqsubseteq E \sqcup \forall R^{-} . G$ |  | (25) |
|  |  |  |  | Hyper[9+25]: | $\top \sqsubseteq \forall R^{-} . F \sqcup \forall R^{-} . G$ |  | (26) |
|  |  |  |  | Pred[23+27+26]: | T¢ $F \sqcup G$ |  | (29) |

Figure 1: Example Inferences of the Consequence-Based Algorithm
queries in $\mathcal{Q}$. The inferences of the algorithm are not localized to a part of $\mathcal{I}$, which allows our algorithm to avoid redundant computation: although there are many technical differences, our algorithm is somewhat related to the tableau algorithms with caching [33, 34]. To be goal oriented, our algorithm constructs a context structure-a graph-like structure whose vertices are called contexts. Intuitively, each context describes one or more elements in the model $\mathcal{I}$, and the edges between contexts capture the relations between the corresponding elements of $\mathcal{I}$. Each context $v$ is associated with a set core $(v)$ of core atomic concepts. Intuitively, the concepts in core $(v)$ hold for each element of $\mathcal{I}$ that corresponds to $v$; hence, core $(v)$ determines the "type" of $v$ and the corresponding model elements. Furthermore, each context $v$ is associated with a set of clauses $\mathcal{S}(v)$ that, as in propositional resolution, describe the domain elements in $\mathcal{I}$ that correspond to $v$. Each clause $K \sqsubseteq M \in \mathcal{S}(v)$ is "relative" to core $(v)$ and should be interpreted as core $(v) \sqcap K \sqsubseteq M$. In other words, the concepts in core $(v)$ hold in all model elements corresponding to $v$, so we drop core $(v)$ from the clauses in $\mathcal{S}(v)$ for the sake of clarity. We next demonstrate how to use the derivation rules in Table 3 on page 13 to prove $\mathcal{O} \vDash A \sqsubseteq G$. We will construct the context structure shown in Figure 1; each context is shown as a circle, the core of each context is shown above the circle, the clauses belonging to the context are shown below the circle, and the numbers next to the clauses correspond to the order of inference rule applications.

To avoid the drawbacks outlined earlier, our algorithm does not perform inferences between the clauses in $\mathcal{O}$; instead, each inference involves either a single set of clauses $\mathcal{S}(v)$ and possibly the ontology $\mathcal{O}$, or a pair of sets of clauses $\mathcal{S}(v)$ and $\mathcal{S}(u)$. To initiate the inference process, we initialize the algorithm according to the target query. Since our goal is to prove $\mathcal{O} \vDash A \sqsubseteq G$, we introduce a single context $v_{1}$ with $\operatorname{core}\left(v_{1}\right)=\{A\}$, and we add clause (11) to $\mathcal{S}\left(v_{1}\right)$. Intuitively, this says that the model $\mathcal{I}$ must contain at least one element in which $A$ holds, which is similar to initializing the hypertableau algorithm by $A(a)$. We then use the Hyper rule to derive clauses (12)-(16). Since only $A$ is assumed to hold in context $v_{1}$, we derive only a linear number of clauses: no $B_{i}$ holds in $v_{1}$, so we do not derive clauses with irrelevant combinations of atomic concepts. These inferences are analogous to the inferences of the Hyp-rule in the hypertableau algorithm, with the difference that the conclusions are not localized: they hold for each element of $\mathcal{I}$ that corresponds to $v_{1}$.

Clause (15) says that the elements in $\mathcal{I}$ corresponding to $v_{1}$ must have a successor in which $B$ holds. To satisfy this requirement, we use the Succ rule to introduce context $v_{2}$ and add the edge (17) from $v_{1}$ to $v_{2}$; the edge is labeled with the concept $\exists R . B$ that it satisfies. The Succ rule combines the $\exists$-rule and the $\forall^{+}$-rule from the hypertableau algorithm by means of sets of atomic concepts $\mathbf{B}_{k}$ and $\mathbf{B}_{p}$. Set $\mathbf{B}_{k}$ contains atomic concepts known to hold in $v_{2}$ due to universal restrictions-that is, those concepts $L$ for which $T \sqsubseteq \forall R$. $L$ has been derived in $v_{1}$. In our example, clause (16) ensures that $D$ holds in $v_{2}$, and so we have $\mathbf{B}_{k}=\{D\}$. Thus, set $\{B\} \cup \mathbf{B}_{k}=\{B, D\}$ provides us with an "upper bound" on the core of the new context: both $B$ and $D$ necessarily hold in context $v_{2}$, but we can use an arbitrary subset of $\{B, D\}$ as the core of $v_{2}$. Choosing smaller cores might reduce the number of contexts because it allows us to reuse the same context to satisfy different applications of the Succ rule, but it might also increase the number of clauses per context as the latter is then "less precise." The decision which subset to use is determined by a strategy, which is supplied as parameter to our algorithm. In this example we use the eager strategy that always uses the maximal subset and thus sets core $\left(v_{2}\right)=\{B, D\}$; however, we discuss other reasonable strategies and the tradeoffs involved in the choice of a core in more detail in Section 3.4. The Succ rule also initializes $\mathcal{S}\left(v_{2}\right)$ with clauses (18) and (19) specifying that $B$ and $D$ hold in each element of $\mathcal{I}$ that is represented by context $v_{2}$. Set $\mathbf{B}_{p}$ contains atomic concepts that can possibly hold in $v_{2}$ due to universal restrictions-that is, those concepts $L$ for which $K \sqsubseteq M \sqcup \forall R . L$ has been derived in $v_{1}$ for some $K$ and $M$. Thus, set $\{B\} \cup \mathbf{B}_{p}$ provides us with an "upper bound" on the concepts that might need to be considered in $v_{2}$. Since (16) is the only clause in $\mathcal{S}\left(v_{1}\right)$ containing a universal restriction, in our example we have $\mathbf{B}_{p}=\mathbf{B}_{k}=\{D\}$. Sets $\mathbf{B}_{k}$ and $\mathbf{B}_{p}$ need not coincide in general, and the Succ rule adds a clause of the form $L \sqsubseteq L$ for each concept $L$ from $\mathbf{B}_{p}$ that is not in the core of the new context.

We next use the Hyper rule to derive clauses (20)-(26) in $\mathcal{S}\left(v_{2}\right)$. As in the hypertableau algorithm, since only $B$ holds in context $v_{2}$, we do not derive a clause involving $A_{i}$. Now clause (23) requires an edge labeled with $\exists R . B$; due to clause (24), the core of the target context can be any subset of $\{B, D\}$; and due to the eager strategy, the core of the target context will be $\{B, D\}$. But then, there is no need to introduce a fresh context: our eager strategy can instruct the Succ rule to reuse the existing context $v_{2}$ since we already have core $\left(v_{2}\right)=\{B, D\}$. Consequently, we introduce edge (27) using the Succ rule. This "reuse" of $v_{2}$ is similar to blocking in the hypertableau algorithm, but is actually much more effective in eliminating redundant computations. First, we never need more than exponentially many contexts, whereas the hypertableau algorithm can construct trees of doubly exponential size. Second, the clauses belonging to each context are not localized to a specific place in $\mathcal{I}$, and so our algorithm draws the inferences for a particular core only once. In contrast, to ensure that the labels of $s$ and $s^{\prime}$ coincide and thus ensure the blocking condition, the hypertableau algorithm must draw the same conclusions for $s$ and $s^{\prime}$; furthermore, an individual can directly block exponentially many other individuals, so the potential for redundant work is even higher.

Clause (26) says that each element in $\mathcal{I}$ corresponding to a predecessor of an element represented by $v_{2}$ must satisfy $F$ or $G$; hence, we apply the Pred rule to edge (17) and clauses (15) and (26) to derive clause (28). Furthermore, we also apply the Pred rule to edge (27) and thus derive clause (29). The Pred rule essentially "pulls" the information from the successor to the predecessor; however, unlike the $\forall^{-}$-rule in the hypertableau algorithm, it simultaneously deals with several universal restrictions.

Finally, we use the Hyper rule to derive clause (30), at which point no further inferences are possible. Since all clauses are "relative" to the core of the corresponding context, clause (30) actually corresponds to $A \sqsubseteq G$, so we have proved $\mathcal{O} \vDash A \sqsubseteq G$. In fact, due to (30), we know that $\mathcal{O} \vDash K \sqsubseteq M$ for each query $K \sqsubseteq M$ such that $A \sqsubseteq G$ is a strengthening of $K \sqsubseteq M$. Our algorithm is thus not just refutationally complete: for each query $K \sqsubseteq M$ such that $\mathcal{O} \vDash K \sqsubseteq M$, it derives at least one strengthening of $K \sqsubseteq M$ in each context that covers (cf. Definition 4) the query.

We finish this section with a note that unrestricted application of the Hyper rule can be very prolific, so we use the ordered resolution variant [31]: we parameterize our algorithm with an ordering on literals, and, for each clause $K \sqsubseteq M \in \mathcal{S}(v)$, we apply the inference rules only to literals that are maximal in $M$ w.r.t. the ordering. The Pred rule, however, introduces a complication: clause (26), which is needed for the Pred rule, can be derived from (25) only if $\forall R^{-} . G$ is smaller than $E$ in the ordering. Therefore, we use an ordering $<_{v}$ per context $v$; furthermore, if $v$ has an incoming edge labeled with $\exists R . A$, then $<_{v}$ must be $R$-admissible (cf. Definition 3)-that is, each literal of the form $\forall \operatorname{inv}(R) . B$ must be smallest in $<_{\nu}$. These restrictions ensure that all clauses that can participate in the Pred rule, such as (26), are derived in $\mathcal{S}(v)$ and so the rule can be applied. Please note that orderings are optional: the trivial empty ordering (i.e., the ordering under which all literals are incomparable) is $R$-admissible for each role $R$, and it can therefore be used in each context. Moreover, orderings are not necessary for the fixed-parameter tractability results we present in Sections 4 and 5: these hold even if each context uses the trivial empty ordering.

### 3.2. Formalizing the Consequence-Based Framework

In this section we formalize the intuitions that we discussed in Section 3.1. Recall that in Section 2.1 we introduced $\Sigma_{L}$ as the set of all literals, and $\Sigma_{L}^{\exists}$ as the set of all existential restrictions over the signature $\Sigma=\left\langle\Sigma_{A}, \Sigma_{T}\right\rangle$. Moreover, recall that we often treat conjunctions and disjunctions of literals as sets of their respective conjuncts and disjuncts. We first introduce the notion of a literal ordering, which we will use later to restrict applicability of our inference rules.

Definition 3. A literal ordering < is a strict partial order (i.e., an irreflexive and transitive relation) on the set $\Sigma_{L}$ of all literals. A literal $L \in \Sigma_{L}$ is <-minimal if no literal $L^{\prime} \in \Sigma_{L}$ exists such that $L^{\prime}<L$; moreover, a set of literals $N$ is $\prec$-minimal if each literal in $N$ is <-minimal. A literal $L \in \Sigma_{L}$ is <-maximal w.r.t. a set of literals $N$, written $L \nprec N$, if no literal $L^{\prime} \in N$ exists such that $L<L^{\prime}$. Given a role $R$, ordering $<$ is $R$-admissible if each literal of the form $\forall \operatorname{inv}(R)$. $A$ is <-minimal.

By Definition 3, a literal ordering < is a partial order so two literals can be incomparable w.r.t. <. Thus, a literal $L$ being maximal in some set of literals $N$ does not imply that $L^{\prime}<L$ holds for each literal $L^{\prime} \in N$; instead, this only means that no literal in $N$ is larger than $L$ w.r.t. $<$, which we reflect by using notation $L \nprec N$ for the notion of maximality.

Throughout the rest of this paper we fix a "global" countably infinite set of contexts $\mathbf{X}$. Sets $\mathbf{X}, \Sigma_{L}$, and $\Sigma_{L}^{\exists}$ provide us with building blocks for the following definition of a context structure.
Definition 4. A context structure is a tuple $\mathcal{D}=\left\langle\mathcal{V}, \mathcal{E}\right.$, core, $\langle \rangle$, where $\mathcal{V} \subseteq \mathbf{X}$ is a finite set of contexts, $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V} \times \Sigma_{L}^{\exists}$ is a finite set of edges labeled by existential restrictions, function core : $\mathcal{V} \rightarrow 2^{\Sigma_{L}}$ labels each context with a finite set of literals, and function $<$ assigns to each context $v \in \mathcal{V}$ a literal ordering $\prec_{v}$. Such $\mathcal{D}$ is admissible if, for each edge $\langle v, u, \exists R . A\rangle \in \mathcal{E}$, the literal ordering $<_{u}$ is $R$-admissible. Furthermore, $\mathcal{D}$ is over a set of literals $\mathbf{L}$ if all literals in $\mathcal{D}$ are contained in $\mathbf{L}$. Finally, given a context $v$, set core $(v)$ is often treated as the conjunction of its atomic concepts, thus allowing core $(v)$ to occur in concepts and axioms.

Let $v \in \mathcal{V}$ be an arbitrary context of $\mathcal{D}$; then, $v$ is trivial if $\operatorname{core}(v)=\emptyset$ and $<_{v}=\emptyset$. In addition, let $K \sqsubseteq M$ be an arbitrary query; then, $v$ is sound for $K \sqsubseteq M$ if $\operatorname{core}(v) \subseteq K$; $v$ is complete for $K \sqsubseteq M$ if $M$ is $<_{v}$-minimal; and $v$ covers $K \sqsubseteq M$ if $v$ is both sound and complete for $K \sqsubseteq M$.

A clause system, which we formalize next, assigns a set of clauses $\mathcal{S}(v)$ to each context $v$. As we discussed in Section 3.1, our consequence-based algorithm produces a context structure $\mathcal{D}$ and a clause system $\mathcal{S}$. The clauses in each set $\mathcal{S}(v)$ will be
"relative" to core( $v$ )—that is, for each clause $K \sqsubseteq M \in \mathcal{S}(v)$, we will have $\mathcal{O} \vDash \operatorname{core}(v) \sqcap K \sqsubseteq M$. Moreover, we will be able to use $\mathcal{S}$ to decide query entailment: for each query $K \sqsubseteq M$ and each context $v$ covering $K \sqsubseteq M$ initialized so that $K \sqsubseteq L \hat{\epsilon} \mathcal{S}(v)$ holds for each literal $L \in K$, we will have $\mathcal{O} \vDash K \sqsubseteq M$ if and only if $K \sqsubseteq M \hat{\in} \mathcal{S}(v)$. The notion of context soundness in Definition 4 ensures that query $K \sqsubseteq M$ is compatible with the intended reading of the clauses in $\mathcal{S}(v)$, and the notion of context completeness ensures that ordering $<_{v}$ does not prevent the derivation of a strengthening of $K \sqsubseteq M$ in such $\mathcal{S}(v)$.
Definition 5. A clause system for a context structure $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, $\langle \rangle$ is a function $\mathcal{S}$ that assigns to each context $v \in \mathcal{V}$ a finite set of clauses $\mathcal{S}(v)$.

To simplify the presentation, in the rest of this section we fix a normalized $\mathcal{A L C I}$ ontology $\mathcal{O}$ and a finite set of queries $\mathcal{Q}$-that is, we assume that all subsequent definitions and theorems in this section are implicitly parameterized with $\mathcal{O}$ and $\mathcal{Q}$. Furthermore, we fix $\mathbf{L}$ to be the set of literals occurring in $\mathcal{O} \cup \mathcal{Q}$; since $\mathcal{O}$ and $\mathcal{Q}$ are finite sets, set $\mathbf{L}$ is finite as well.

As we explained in Section 3.1, our algorithm is based on the inference rules shown in Table 3. The Hyper, Succ, and Pred rules are responsible for completeness, and we discussed the intuitions behind these rules in Section 3.1. The Elim rule supports redundancy elimination and is not needed for completeness or the complexity results in Sections 4 and 5. We discuss the intuitions underlying the Elim rule in more detail in Section 3.3; however, in order to present the soundness, completeness, and termination proofs in one place, we introduce the rule here. The completeness of our algorithm is guaranteed by Theorem 6, the proof of which is given in Appendix A. The theorem essentially says that, given a context structure $\mathcal{D}$ and a clause system $\mathcal{S}$ saturated under the the Hyper, Pred, and Succ rules, we can read off the consequences of the form $K \sqsubseteq M$ from the sets $\mathcal{S}(v)$ provided that $K \sqsubseteq M$ is a query, the context $v$ is complete (cf. Definition 4) for $K \sqsubseteq M$, and set $\mathcal{S}(v)$ is appropriately initialized. The initialization is similar to asserting $L(a)$ for each literal $L \in K$ at the beginning of a hypertableau test of $\mathcal{O} \vDash K \sqsubseteq M$. Please note that Theorem 6 depends on the Succ rule being not applicable, and so the exact definition of strategy used in the postcondition of the Succ rule is not relevant for completeness. In the rest of this section we discuss how to initialize the sets $\mathcal{S}(v)$ and how to satisfy the precondition of the Succ rule to obtain a sound and complete algorithm.

Theorem 6 (Completeness). Let $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, $\langle \rangle$ be an admissible context structure, and let $\mathcal{S}$ be a clause system for $\mathcal{D}$ such that the Hyper, Pred, and Succ rules from Table 3 are not applicable to $\mathcal{D}$ and $\mathcal{S}$. Then, $K \sqsubseteq M \hat{\in} \mathcal{S}(v)$ holds for each query $K \sqsubseteq M$ and each context $v \in \mathcal{V}$ that satisfy all of the following three conditions:

- $\mathcal{O} \vDash K \sqsubseteq M$,
- context $v$ is complete for query $K \sqsubseteq M$, and
- $K \sqsubseteq L \hat{\epsilon} \mathcal{S}(v)$ for each literal $L \in K$.

As we discussed in Section 3.1, our framework is parameterized by a function strategy, which we formalize in Definition 7. Intuitively, when the Succ rule identifies a context $v$ and an existential restriction $\exists R . A$ for which the rule's precondition is not satisfied, strategy $\left(\exists R . A, \mathbf{B}_{k}, \mathcal{D}\right)$ returns a triple $\left\langle u\right.$, core', $\left.<^{\prime}\right\rangle$ that specifies how to modify $\mathcal{D}$ and $\mathcal{S}$ so that the precondition becomes satisfied. The strategy has two options. First, the strategy can decide to extend $\mathcal{D}$ by returning a fresh context $u$; in such a case, $u$ is added to $\mathcal{D}$, and core' and $<^{\prime}$ specify how to initialize core $(u)$ and $<_{u}$. Second, the strategy can decide to reuse a context $u$ that is already a part of $\mathcal{D}$; if so, then ordering $<_{u}$ is intersected with $<^{\prime}$ to ensure $R$-admissibility, but also preserve admissibility for all other roles. Definition 7 requires each strategy to be bounded: if the signature is finite, then the strategy should return only finitely many different results on all possible inputs. Please note that a finite signature does not necessarily bound the size of $\mathcal{D}$, so the number of different arguments (and return values) for strategy is not bounded. Bounded strategies can thus introduce only finitely many fresh contexts, which ensures termination.

Definition 7. An expansion strategy is a function strategy computable in polynomial time that takes an existential restriction $\exists R . A$, a set of atomic concepts $\mathbf{B}_{k}$, and a context structure $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, $\prec\rangle$. The result of strategy $\left(\exists R . A, \mathbf{B}_{k}, \mathcal{D}\right)$ is a triple $\left\langle u\right.$, core' $\left.^{\prime},<^{\prime}\right\rangle$ where

- core' is a subset of $\{A\} \cup \mathbf{B}_{k}$,
- $\prec^{\prime}$ is an $R$-admissible literal ordering, and
- either $u \in \mathbf{X} \backslash \mathcal{V}$ is a fresh context, or $u \in \mathcal{V}$ is a context in $\mathcal{D}$ such that $\operatorname{core}(u)=$ core' $^{\prime}$.

Each expansion strategy must be bounded, which is the case if, for each finite set of literals $\mathbf{L}$, the number $\mid$ strategy $\left.\right|_{\mathbf{L}}$ of different values that strategy can return for all possible arguments over $\mathbf{L}$ is finite.

We next define when a context structure $\mathcal{D}$ and a clause system $\mathcal{S}$ are sound for $\mathcal{O}$. Recall that, for a context $v$, clause $K \sqsubseteq M$ in $\mathcal{S}(v)$ is "relative" to core $(v)$-that is, the clause should be interpreted as core $(v) \sqcap K \sqsubseteq M$. Soundness for a clause simply means that $\mathcal{O} \vDash \operatorname{core}(v) \sqcap K \sqsubseteq M$ should hold for each clause $K \sqsubseteq M \in \mathcal{S}(v)$. To understand the notion of soundness for a context structure, consider an arbitrary edge $\langle v, u, \exists R . A\rangle$ in $\mathcal{D}$. As we explained in Section 3.1, the Pred rule "pulls" information from $\mathcal{S}(u)$ into $\mathcal{S}(v)$; however, the clauses in $\mathcal{S}(v)$ and $\mathcal{S}(u)$ are "relative" to core $(v)$ and core $(u)$, respectively, so core $(u)$ and
core $(v)$ must be related as specified in Definition 8 to make the Pred rule sound. Please remember that core $(u)$ and core $(v)$ should be understood as conjunctions of the respective sets of atomic concepts.

Definition 8. A context structure $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, $\langle \rangle$ is sound for $\mathcal{O}$ if $\mathcal{O} \vDash \operatorname{core}(v) \sqcap \exists R . A \sqsubseteq \exists R$.[core $(u) \sqcap A]$ holds for each edge $\langle v, u, \exists R . A\rangle \in \mathcal{E}$. Furthermore, a clause system $\mathcal{S}$ is sound for $\mathcal{O}$ if $\mathcal{O} \vDash \operatorname{core}(v) \sqcap K \sqsubseteq M$ holds for each context $v \in \mathcal{V}$ and each clause $K \sqsubseteq M \in \mathcal{S}(v)$.

Proposition 9 shows that the inference rules in Table 3 preserve admissibility and soundness of a context structure and a clause system. The former follows immediately from Definition 7, and the proof of the latter is analogous to the soundness proof for first-order resolution [31]: we show that each clause introduced by an inference rule is a logical consequence of $\mathcal{O}$.

Proposition 9 (Soundness). Let $\mathcal{D}_{1}=\left\langle\mathcal{V}, \mathcal{E}\right.$, core, $\langle \rangle$ be a context structure, let $\mathcal{S}_{1}$ be a clause system for $\mathcal{D}_{1}$, and let $\mathcal{D}_{2}$ and $\mathcal{S}_{2}$ be the context structure and the clause system, respectively, obtained by applying an inference rule from Table 3 to $\mathcal{D}_{1}$ and $\mathcal{S}_{1}$. If $\mathcal{D}_{1}$ is admissible, then $\mathcal{D}_{2}$ is admissible. Furthermore, if both $\mathcal{D}_{1}$ and $\mathcal{S}_{1}$ are sound for $\mathcal{O}$, then both $\mathcal{D}_{2}$ and $\mathcal{S}_{2}$ are sound for $\mathcal{O}$.

Proof. Assume that $\mathcal{D}_{1}$ is admissible and that the Succ rule is applied to $\mathcal{D}_{1}$ as shown in Table 3. By Definition 7, literal ordering $<^{\prime}$ is $R$-admissible; furthermore, for arbitrary literal orderings $<_{1}$ and $<_{2}$ that are $R$ - and $S$-admissible, respectively, set $<_{1} \cap<_{2}$ is a literal ordering that is both $R$ - and $S$-admissible. Therefore, regardless of whether context $u$ is fresh or not, $\mathcal{D}_{2}$ is admissible. Now assume that both $\mathcal{D}_{1}$ and $\mathcal{S}_{1}$ are sound; we next show that both $\mathcal{D}_{2}$ and $\mathcal{S}_{2}$ are sound as well. To this end, let $\mathcal{I}=\left\langle\Delta^{\mathcal{I}}, .^{\mathcal{I}}\right\rangle$ be an arbitrary model of $\mathcal{O}$, and consider all possible inference rules that derive a clause in $\mathcal{S}_{2}$ or modify $\mathcal{D}_{2}$.
(Hyper rule) Assume that the rule is applied as in Table 3; we will show that $\mathcal{O} \vDash \operatorname{core}(v) \sqcap\rceil_{i=1}^{n} K_{i} \sqsubseteq M \sqcup \bigsqcup_{i=1}^{n} M_{i}$. To this end, consider an arbitrary element $\delta \in \Delta^{\mathcal{I}}$ such that $\delta \in\left(\operatorname{core}(v) \sqcap \prod_{i=1}^{n} K_{i}\right)^{\mathcal{I}}$. Now $\mathcal{S}_{1}$ is sound for $\mathcal{O}$ so, for each $1 \leq i \leq n$, we have $\mathcal{O} \vDash \operatorname{core}(v) \sqcap K_{i} \sqsubseteq M_{i} \sqcup A_{i}$, which implies $\delta \in\left(M_{i} \sqcup A_{i}\right)^{\mathcal{I}}$. If $\delta \in M_{i}^{\mathcal{I}}$ for some $1 \leq i \leq n$, then clearly $\delta \in\left(\bigsqcup_{i=1}^{n} M_{i}\right)^{\mathcal{I}}$. Otherwise, we have $\delta \in\left(\prod_{i=1}^{n} A_{i}\right)^{\mathcal{I}}$, but then $\prod_{i=1}^{n} A_{i} \sqsubseteq M \in \mathcal{O}$ implies $\delta \in M^{\mathcal{I}}$. In either case, we have $\delta \in\left(M \sqcup \bigsqcup_{i=1}^{n} M_{i}\right)^{\mathcal{I}}$. Since $\delta$ was chosen arbitrarily, $\mathcal{I}$ satisfies the conclusion of the rule, as required.
(Pred rule) Assume that the rule is applied as in Table 3; we will show that $\mathcal{O} \vDash \operatorname{core}(v) \sqcap\rceil_{i=0}^{n} K_{i} \sqsubseteq \bigsqcup_{i=0}^{n} M_{i} \sqcup \bigsqcup_{j=1}^{m} C_{j}$. To this end, consider an arbitrary element $\delta \in \Delta^{\mathcal{I}}$ such that $\delta \in\left(\operatorname{core}(v) \sqcap \prod_{i=0}^{n} K_{i}\right)^{\mathcal{I}}$. As in the previous case, if $\delta \in M_{i}^{\mathcal{I}}$ for some $0 \leq i \leq n$, then clearly $\delta \in\left(\bigsqcup_{i=0}^{n} M_{i}\right)^{\mathcal{I}}$. Otherwise, as $\mathcal{S}_{1}$ is sound for $\mathcal{O}$, we have $\delta \in(\exists R . A)^{\mathcal{I}}$ and $\delta \in\left(\forall R . B_{i}\right)^{\mathcal{I}}$ for each $1 \leq i \leq n$. Since $\langle v, u, \exists R . A\rangle \in \mathcal{E}$ and $\mathcal{D}_{1}$ is sound, we have $\mathcal{O} \vDash \operatorname{core}(v) \sqcap \exists R . A \sqsubseteq \exists R$.[core $\left.(u) \sqcap A\right]$; therefore, since $\delta \in[\operatorname{core}(v) \sqcap \exists R . A]^{\mathcal{I}}$, we have $\delta \in(\exists R$.[core $(u) \sqcap A])^{\mathcal{I}}$. Hence, an element $\gamma \in \Delta^{\mathcal{I}}$ exists such that $\langle\delta, \gamma\rangle \in R^{\mathcal{I}}$ and $\gamma \in\left[\operatorname{core}(u) \sqcap A \sqcap \prod_{i=1}^{n} B_{i}\right]^{\mathcal{I}}$. But then, due to either of the two alternative premises of the rule, we also have $\gamma \in\left[\bigsqcup_{j=1}^{m} \forall \operatorname{inv}(R) . C_{j}\right]^{\mathcal{I}}$; thus, $\delta \in\left(\bigsqcup_{j=1}^{m} C_{j}\right)^{\mathcal{I}}$ holds. Since $\delta$ was chosen arbitrarily, $\mathcal{I}$ satisfies the conclusion of the rule, as required.
(Succ rule) Assume that the rule is applied as in Table 3; in the rest of this proof, we consider set $\mathbf{B}_{k}$ to be equivalent to the conjunction of the atomic concepts contained in it. Each clause introduced by the rule is of the form $L \sqsubseteq L$, so clearly $\mathcal{I} \vDash \operatorname{core}(u) \sqcap L \sqsubseteq L$. We next show that the new edge $\langle v, u, \exists R . A\rangle$ introduced by the rule satisfies the axiom from Definition 8 . For each atomic concept $B \in \mathbf{B}_{k}$, we have $\top \sqsubseteq \forall R . B \in \mathcal{S}(v)$ by the rule preconditions, so $\mathcal{O} \vDash \operatorname{core}(v) \sqsubseteq \forall R . B$ since $\mathcal{S}_{1}$ is sound for $\mathcal{O}$; but then, we have $\mathcal{O} \vDash \operatorname{core}(v) \sqcap \exists R . A \sqsubseteq \exists R .\left[A \sqcap \mathbf{B}_{k}\right]$. Furthermore, since core $(u)=$ core $\subseteq A \sqcap \mathbf{B}_{k}$ holds, we also have $\mathcal{O} \vDash \operatorname{core}(v) \sqcap \exists R . A \sqsubseteq \exists R$. [core $(u) \sqcap A]$, as required.

The following notion of a covering mapping maps each query $q \in \mathcal{Q}$ to a context in a context structure. Intuitively, our algorithm will use such a mapping to find for each query $q \in \mathcal{Q}$ a context that can be used to verify $\mathcal{O} \vDash q$.

Definition 10. Let $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, $\langle \rangle$ be a context structure. A covering mapping from $\mathcal{Q}$ to $\mathcal{D}$ is a function $\vartheta: \mathcal{Q} \rightarrow \mathcal{V}$ such that each query $q \in \mathcal{Q}$ is covered in the context $\vartheta(q)$.

We are finally ready to formally define our consequence-based reasoning algorithm. Apart from $\mathcal{O}$ and $\mathcal{Q}$, the algorithm is parameterized by an expansion strategy strategy, a context structure $\mathcal{D}$ with no edges, and a covering mapping $\vartheta$. By passing $\mathcal{D}$ and $\vartheta$ as arguments, the algorithm's users can decide how to initialize the contexts necessary for checking the queries in $\mathcal{Q}$; we discuss reasonable possibilities in Section 3.4.1. Since the algorithm is given a context structure without edges, $\mathcal{D}$ is trivially sound for $\mathcal{O}$; by inductively applying Proposition 9, the resulting clause system is sound for $\mathcal{O}$. Furthermore, the algorithm is complete because steps 1 and 2 satisfy the preconditions of Theorem 6 .

Algorithm 11. The consequence-based reasoning algorithm for $\mathcal{A L C \mathcal { I }}$ takes $\mathcal{O}$, $\mathcal{Q}$, a context structure $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, $\langle \rangle$ over $\mathbf{L}$ with $\mathcal{E}=\emptyset$, an expansion strategy strategy, and a covering mapping $\vartheta$ from $\mathcal{Q}$ to $\mathcal{D}$. The algorithm (nondeterministically) extends $\mathcal{D}$ and computes a clause system $\mathcal{S}$ for the extended context structure as follows.

1. Set $\mathcal{S}(v):=\{\top \sqsubseteq L \mid L \in \operatorname{core}(v)\}$ for each context $v \in \mathcal{V}$.
2. For each query $K \sqsubseteq M \in \mathcal{Q}$, let $v:=\vartheta(K \sqsubseteq M)$; then, for each literal $L \in K \backslash \operatorname{core}(v)$, add $L \sqsubseteq L$ to $\mathcal{S}(v)$.
3. Exhaustively apply the inference rules from Table 3.

Please note that Algorithm 11 is nondeterministic: the conclusion of an inference rule is added to $\mathcal{S}(v)$ only if $\mathcal{S}(v)$ does not contain a strengthening of the conclusion, and the order of inference rule application is not predetermined; therefore, the resulting $\mathcal{D}$ and $\mathcal{S}$ are not unique. This, however, is don't-care nondeterminism: any order of inference rule applications suffices for soundness and completeness.

In Proposition 13 we formalize the soundness and completeness argument outlined earlier. Towards this goal, in Lemma 12 we show that the Elim rule never deletes the "relevant" consequences derived by our algorithm: whenever a clause $K \sqsubseteq M$ is deleted from some set $\mathcal{S}(v)$, after deletion the set still contains a strengthening of the deleted clause.

Lemma 12. If, at some point in the execution of Algorithm 11, $K \sqsubseteq M \hat{\in} \mathcal{S}(v)$ holds for some clause $K \sqsubseteq M$ and context v, then $K \sqsubseteq M \hat{\in} \mathcal{S}(v)$ holds at all future points as well.

Proof. The proof is by a straightforward induction on the application of the rules. In particular, rules Hyper, Pred, and Succ just add clauses, so their application clearly preserves this property. Furthermore, the Elim rule removes a clause $K_{2} \sqsubseteq M_{2}$ from $\mathcal{S}(v)$ only if $K_{1} \sqsubseteq M_{1} \in \mathcal{S}(v)$ where $K_{1} \sqsubseteq M_{1}$ is a strengthening of $K_{2} \sqsubseteq M_{2}$; since the strengthening relation on clauses is transitive, an application of the Elim rule clearly preserves this property as well.

Proposition 13. Let $\mathcal{S}$ be a clause system obtained by applying Algorithm 11 to $\mathcal{O}, \mathcal{Q}$, a context structure $\mathcal{D}$, an expansion strategy strategy, and a covering mapping $\vartheta$. Then, for each query $q \in \mathcal{Q}$, we have $\mathcal{O} \vDash q$ if and only if $q \hat{\in} \mathcal{S}(\vartheta(q))$.
Proof. Consider an arbitrary query $K \sqsubseteq M \in \mathcal{Q}$ and let $v:=\vartheta(K \sqsubseteq M)$. Context $v$ covers $K \sqsubseteq M$, so $v$ is sound for $K \sqsubseteq M$ and thus core $(v) \subseteq K$ holds by Definition 4, and $v$ is also complete for $K \sqsubseteq M$.

Assume that $K \sqsubseteq M \hat{\in} \mathcal{S}(v)$, so a clause $K^{\prime} \sqsubseteq M^{\prime} \in \mathcal{S}(v)$ exists such that $K^{\prime} \subseteq K$ and $M^{\prime} \subseteq M$. For each clause $T \sqsubseteq L$ added to some $\mathcal{S}(u)$ in step 1 , we have $\mathcal{O} \vDash \operatorname{core}(u) \sqsubseteq L$ since $L \in \operatorname{core}(u)$; furthermore, for each clause $L \sqsubseteq L$ added to $\mathcal{S}(u)$ in step 2 , we clearly have $\mathcal{O} \vDash \operatorname{core}(u) \sqcap L \sqsubseteq L$; finally, Algorithm 11 is applied to a sound context structure, so $\mathcal{D}$ and $\mathcal{S}$ are sound for $\mathcal{O}$ by Proposition 9 , and $\mathcal{O} \vDash \operatorname{core}(v) \sqcap K^{\prime} \sqsubseteq M^{\prime}$ holds. Finally, $K^{\prime} \cup \operatorname{core}(v) \subseteq K$ and $M^{\prime} \subseteq M$ imply $\mathcal{O} \vDash K \sqsubseteq M$, as required.

Conversely, assume that $\mathcal{O} \vDash K \sqsubseteq M$. Consider an arbitrary literal $L \in K$; if $L \in \operatorname{core}(v)$, then $T \sqsubseteq L$ is added to $\mathcal{S}(v)$ in step 1 , and if $L \in K \backslash \operatorname{core}(v)$, then $L \sqsubseteq L$ is added to $\mathcal{S}(v)$ in step 2 ; in either case, we have that $K \sqsubseteq L \hat{\in} \mathcal{S}(v)$ holds after step 2, and Lemma 12 ensures that this property is preserved during the algorithm's execution. But then, since $v$ is complete for $K \sqsubseteq M$, we have $K \sqsubseteq M \hat{\epsilon} \mathcal{S}(v)$ by Theorem 6, as required.

Finally, we determine the complexity of our algorithm.
Proposition 14 (Termination). When applied to $\mathcal{O}, \mathcal{Q}$, a context structure $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, $\langle \rangle$, an expansion strategy strategy, and a covering mapping $\vartheta$, Algorithm 11 terminates in time polynomial in $4^{|\mathbf{L}|^{2}}$, |strategy $\left.\right|_{\mathbf{L}}+|\mathcal{V}|$, and $\|\mathcal{O}\|+\|\mathcal{Q}\|$.

Proof. When applied to arguments over $\mathbf{L}$, the result of the expansion strategy is clearly over $\mathbf{L}$ as well; therefore, the context structure and the clause system computed using Algorithm 11 are over $\mathbf{L}$. The number of different clauses over $\mathbf{L}$ is bounded by $c=2^{|\mathbf{L}|} \cdot 2^{|\mathbf{L}|}=4^{|\mathbf{L}|}$ since each literal from $\mathbf{L}$ can independently occur in the clause antecedent and/or the clause consequent. Algorithm 11 is applied to a context structure with $|\mathcal{V}|$ contexts, and the expansion strategy can introduce at most $|s t r a t e g y|_{\mathbf{L}}$ additional contexts; hence, the total number of contexts introduced in the algorithm is bounded by $m=\mid$ strategy $\left.\right|_{\mathbf{L}}+|\mathcal{V}|$.

We call all objects needed to apply an inference rule from Table 3 premises; for the Succ rule, this includes set $\mathbf{B}_{p}$. We next show that, for the Hyper, Pred, and Succ rules, the number of different premises is polynomial in $c^{|\mathbf{L}|}, m$, and $\|\mathcal{O}\|+\|\mathcal{Q}\|$; recall that $|\mathbf{L}| \leq\|\mathcal{O}\|+\|\mathcal{Q}\|$. For the Hyper rule, context $v \in \mathcal{V}$ can be chosen in at most $m$ ways, clause $\prod_{i} A_{i} \sqsubseteq M \in \mathcal{O}$ can be chosen in at most $|\mathcal{O}|$ ways, and each of the $n \leq|\mathbf{L}|$ clauses $K_{i} \sqsubseteq M_{i} \sqcup A_{i} \in \mathcal{S}(v)$ can be chosen in at most $c$ ways, so there are at most $c^{|\mathbf{L}|}$ ways of choosing such a set of clauses. For the Pred rule, edge $\langle v, u, \exists R . A\rangle \in \mathcal{E}$ can be chosen in at most $m^{2} \cdot|\mathbf{L}|$ ways, clause $A \sqcap \prod_{i} B_{i} \sqsubseteq \bigsqcup_{j} \forall \operatorname{inv}(R) . C_{j} \in \mathcal{S}(u)$ or $\prod_{i} B_{i} \sqsubseteq \bigsqcup_{j} \forall \operatorname{inv}(R) . C_{j} \in \mathcal{S}(u)$ can be chosen in at most $c$ ways, and each of the $n+1 \leq|\mathbf{L}|$ clauses $K_{0} \sqsubseteq M_{0} \sqcup \exists R . A \in \mathcal{S}(v)$ and $K_{i} \sqsubseteq M_{i} \sqcup \forall R . B_{i} \in \mathcal{S}(v)$ can be chosen in at most $c$ ways, so there are at most $c^{|\mathbf{L}|}$ ways of choosing such a set of clauses. For the Succ rule, context $v \in \mathcal{V}$ can be chosen in at most $m$ ways, literal $\exists R . A$ can be chosen in at most $|\mathbf{L}|$ ways, clause $K \sqsubseteq M \sqcup \exists R . A \in \mathcal{S}(v)$ can be chosen in at most $c$ ways, and there are at most $2^{|\mathbf{L}|}$ different sets $\mathbf{B}_{p}$.

Each rule in Table 3 adds a clause $K \sqsubseteq M$ to some $\mathcal{S}(v)$ only if $K \sqsubseteq M \ddot{\mathcal{G}}(v)$; thus, due to Lemma 12, each such $K \sqsubseteq M$ can be added to some $\mathcal{S}(v)$ at most once; but then, the Elim rule can eliminate such $K \sqsubseteq M$ from $\mathcal{S}(v)$ at most once as well. Thus, no inference rule is applied twice to the same premises, so the number of inference rule applications is bounded by a number that is polynomial in $c^{|\mathbf{L}|}, m$, and $\|\mathcal{O}\|+\|\mathcal{Q}\|$. Moreover, strategy is computable in polynomial time, so each inference rule can be applied to fixed premises in polynomial time. Consequently, the algorithm can be implemented so that it runs in time polynomial in $c^{|\mathbf{L}|}, m$, and $\|\mathcal{O}\|+\|\mathcal{Q}\|$.

### 3.3. Redundancy Elimination

Resolution can often derive redundant clauses-that is, clauses that are not necessary for a proof. Such clauses can give rise to a large number of other redundant clauses, so it is beneficial to detect and eliminate redundant clauses whenever possible. To

| $\stackrel{\text { ¢ }}{\text { ¢ }}$ | $\begin{array}{ll} \text { If } \quad & \prod_{i=1}^{n} A_{i} \sqsubseteq M \in \mathcal{O}, \\ & K_{i} \sqsubseteq M_{i} \sqcup A_{i} \in \mathcal{S}(v) \text { with } A_{i} \nprec_{v} M_{i} \text { for } 1 \leq i \leq n, \\ \quad & \text { and } \prod_{i=1}^{n} K_{i} \sqsubseteq M \sqcup \bigsqcup_{i=1}^{n} M_{i} \ddot{\notin \mathcal{S}(v),} \\ \text { then } \operatorname{add} \prod_{i=1}^{n} K_{i} \sqsubseteq M \sqcup \bigsqcup_{i=1}^{n} M_{i} \text { to } \mathcal{S}(v) . \end{array}$ |
| :---: | :---: |
| - | $\text { If } \begin{array}{ll}  & \langle v, u, \exists R . A\rangle \in \mathcal{E}, \\ & A \sqcap \prod_{i=1}^{n} B_{i} \sqsubseteq \bigsqcup_{j=1}^{m} \forall \operatorname{inv}(R) . C_{j} \in \mathcal{S}(u) \quad \text { or } \quad \prod_{i=1}^{n} B_{i} \sqsubseteq \bigsqcup_{j=1}^{m} \forall \operatorname{inv}(R) . C_{j} \in \mathcal{S}(u), \\ & K_{0} \sqsubseteq M_{0} \sqcup \exists R . A \in \mathcal{S}(v) \text { with } \exists R . A \nprec_{v} M_{0}, \\ & K_{i} \sqsubseteq M_{i} \sqcup \forall R . B_{i} \in \mathcal{S}(v) \text { with } \forall R . B_{i} \nprec_{v} M_{i} \text { for } 1 \leq i \leq n, \\ & \text { and } \prod_{i=0}^{n} K_{i} \sqsubseteq \bigsqcup_{i=0}^{n} M_{i} \sqcup \bigsqcup_{j=1}^{m} C_{j} \notin \mathcal{S}(v), \\ \text { then add } \prod_{i=0}^{n} K_{i} \sqsubseteq \bigsqcup_{i=0}^{n} M_{i} \sqcup \bigsqcup_{j=1}^{m} C_{j} \text { to } \mathcal{S}(v) . \end{array}$ |
| U | If $\quad K \sqsubseteq M \sqcup \exists R . A \in \mathcal{S}(v)$ with $\exists R . A \nprec_{v} M$, and <br> for $\mathbf{B}_{p}:=\left\{B \mid K^{\prime} \sqsubseteq M^{\prime} \sqcup \forall R . B \in \mathcal{S}(v)\right.$ and $\left.\forall R . B \nprec_{v} M^{\prime}\right\}$, <br> no edge $\langle v, u, \exists R . A\rangle \in \mathcal{E}$ exists such that $L \sqsubseteq L \hat{\in} \mathcal{S}(u)$ for each $L \in\{A\} \cup \mathbf{B}_{p}$, <br> then let $\left\langle u\right.$, core $\left.^{\prime},<^{\prime}\right\rangle:=\operatorname{strategy}\left(\exists R . A, \mathbf{B}_{k}, \mathcal{D}\right)$ where $\mathbf{B}_{k}:=\{B \mid \top \sqsubseteq \forall R . B \in \mathcal{S}(v)\}$; <br> if $u \in \mathcal{V}$, then let $<_{u}:=<_{u} \cap<^{\prime}$, and <br> otherwise let $\mathcal{V}:=\mathcal{V} \cup\{u\}, \quad$ core $(u):=$ core' $^{\prime}, \quad<_{u}:=<^{\prime}, \quad$ and $\mathcal{S}(u):=\{T \sqsubseteq L \mid L \in \operatorname{core}(u)\} ;$ <br> $\operatorname{add}\langle v, u, \exists R . A\rangle$ to $\mathcal{E}$; and <br> for each $L \in\left[\{A\} \cup \mathbf{B}_{p}\right] \backslash$ core $(v)$ such that $L \sqsubseteq L \notin \mathcal{S}(u)$, add $L \sqsubseteq L$ to $\mathcal{S}(u)$. |
| $\underline{\underline{E}}$ | If $\quad K_{1} \sqsubseteq M_{1} \in \mathcal{S}(v)$, <br> $K_{2} \sqsubseteq M_{2} \in \mathcal{S}(v)$, and <br> $K_{1} \sqsubseteq M_{1}$ is a strengthening of $K_{2} \sqsubseteq M_{2}$ and the two clauses are distinct, then remove $K_{2} \sqsubseteq M_{2}$ from $\mathcal{S}(v)$. |

this end, modern first-order theorem provers employ a number of redundancy elimination rules that eliminate or simplify certain clauses; Weidenbach [35] presents an overview of these techniques. Inspired by redundancy elimination in first-order theorem provers, we equipped our consequence-based framework with analogous possibilities, which we discuss in this section. Please note that redundancy elimination is optional: it may optimize the theorem proving process but is not needed for completeness. Similarly, the complexity of the algorithms that we present in Sections 4 and 5 is independent from any redundancy elimination techniques, but redundancy elimination can still improve the performance of these algorithms in practice.

First, the completeness Theorem 6 only requires each set of clauses $\mathcal{S}(v)$ to contain a strengthening of each conclusion obtained by applying the inference rules from Table 3 ; therefore, our inference rules derive a conclusion in $\mathcal{S}(v)$ only if $\mathcal{S}(v)$ does not already contain a strengthening of the conclusion. This is analogous to forward subsumption in first-order theorem proving: a newly derived clause is kept only if the clause is not redundant given the clauses derived thus far.

Second, the Elim rule deletes a clause $K_{2} \sqsubseteq M_{2}$ from $\mathcal{S}(v)$ if $\mathcal{S}(v)$ contains a clause $K_{1} \sqsubseteq M_{1}$ such that $K_{1} \sqsubseteq M_{1}$ is a strengthening of $K_{2} \sqsubseteq M_{2}$ and the two clauses are distinct. Note that $K_{2} \sqsubseteq M_{2}$ can be derived before $K_{1} \sqsubseteq M_{1}$, in which case the technique described in the previous paragraph will not eliminate $K_{2} \sqsubseteq M_{2}$, and so $K_{2} \sqsubseteq M_{2}$ can potentially participate in further redundant inferences. The Elim rule is thus analogous to backward subsumption in first-order theorem proving.

In the rest of this section we discuss certain aspects of our redundancy elimination techniques. We first consider syntactic tautologies, which are clauses of the form $K \sqcap L \sqsubseteq M \sqcup L$ with either $K$ or $M$ (or both) nonempty; note that this definition excludes clauses of the form $L \sqsubseteq L$, which are needed for completeness in Theorem 6 and are thus not redundant. The following proposition straightforwardly implies that our algorithm never derives syntactic tautologies; in practice this means that, whenever a possible consequence of the inference rule is a syntactic tautology, one can eagerly eliminate the clause (without checking whether the target clause set contains a strengthening of the consequence clause).

Proposition 15. If, at some point in the execution of Algorithm 11, $K \sqsubseteq M \hat{\in} \mathcal{S}(v)$ holds for some clause $K \sqsubseteq M$ and context $v$, then $L \sqsubseteq L \hat{\in} \mathcal{S}(v)$ also holds at this point for each literal $L \in K$.

Proof. The proof is by induction on the application of inference rules. Each clause introduced in steps 1 and 2 of Algorithm 11 or by the Succ rule obviously satisfies this property. Furthermore, an application of the Elim rule preserves this property by Lemma 12. Now consider an application of the Hyper or the Pred rule as shown in Table 3; then, for each literal $L$ occurring in
the antecedent of the conclusion, integer $i$ exists such that $L \in K_{i}$; but then, the premise corresponding to $K_{i}$ satisfies this property by the induction assumption, so $L \sqsubseteq L \hat{\epsilon} \mathcal{S}(v)$ holds.

The Elim rule can also be used to define more powerful redundancy elimination rules. Assume that, for some context $v$, set $\mathcal{S}(v)$ contains clauses $K_{1} \sqsubseteq M_{1} \sqcup L$ and $K_{2} \sqcap L \sqsubseteq M_{2}$ with $K_{1} \subseteq K_{2}$ and $M_{1} \subseteq M_{2}$. Clause $K_{1} \sqcap K_{2} \sqsubseteq M_{1} \sqcup M_{2}$ is equal to $K_{2} \sqsubseteq M_{2}$ and it is a logical consequence of $\mathcal{S}(v)$; now even though clause $K_{2} \sqsubseteq M_{2}$ is not derived by an inference rule from Table 3 (in particular, note that the Hyper rule does not derive this clause), extending $\mathcal{S}(v)$ with this clause is sound. Thus, if we add $K_{2} \sqsubseteq M_{2}$ to $\mathcal{S}(v)$, we can delete clause $K_{2} \sqcap L \sqsubseteq M_{2}$ from $\mathcal{S}(v)$ using the Elim rule; essentially, we thus obtain a simplification rule that eliminates $L$ from clause $K_{2} \sqcap L \sqsubseteq M_{2}$. Analogously, if the clauses satisfy $K_{2} \subseteq K_{1}$ and $M_{2} \subseteq M_{1}$, then we can eliminate $L$ from clause $K_{1} \sqsubseteq M_{1} \sqcup L$. This is closely related to the contextual literal cutting rule used in the first-order prover E [36].

Redundancy elimination may be difficult to implement fully in practice: even with complex index structures, checking whether some $\mathcal{S}(v)$ contains a strengthening of some clause may be inefficient. Please note, however, that both forms of redundancy elimination are optional-that is, they are not needed for completeness. Therefore, instead of $\hat{\epsilon}$, one can use $\in$ to check rule applicability. Furthermore, one does not need to apply $\hat{\epsilon}$ and the Elim rule exhaustively; for example, one can restrict redundancy checking to clauses of the form $T \sqsubseteq A$ and $L \sqsubseteq L$ and index the latter using a hash table. Finally, an advanced implementation can "switch on" redundancy elimination only if set $\mathcal{S}(v)$ grows beyond a certain size. In most cases, however, eliminating syntactic tautologies should not cause any noticeable overhead.

### 3.4. Instantiating the Consequence-Based Framework

We next discuss ways of instantiating our consequence-based framework presented in Section 3.2. In particular, in Section 3.4.1 we discuss possibilities for initializing the context structure and the covering mapping that are given to Algorithm 11. Then, in Section 3.4.2 we discuss possible expansion strategies.

### 3.4.1. Initializing the Context Structure and the Covering Mapping

Algorithm 11 takes as input a context structure $\mathcal{D}$ without any edges and a covering mapping $\vartheta$, and so the initialization of $\mathcal{D}$ and $\vartheta$ is deferred to the algorithm's users. We next discuss several reasonable approaches to initialization.

The first possibility is to introduce just one trivial (cf. Definition 4) context $v_{0}$. In such a case, step 2 of Algorithm 11 initializes $\mathcal{S}\left(v_{0}\right)$ with many clauses $L \sqsubseteq L$; this may give rise to a large number of inferences at context $v_{0}$, and may be further exacerbated by the fact that $<_{v_{0}}$ is empty. Hence, such an approach is unlikely to be efficient apart from in very simple cases.

The second possibility is to introduce a separate context $v_{q}=\vartheta(q)$ for each query $q=K \sqsubseteq M \in \mathcal{Q}$, define core $\left(v_{q}\right):=K$, and define $<_{v_{q}}$ such that $M$ is ${<_{v_{q}}}$-minimal and the remaining literals are ordered arbitrarily. This has the opposite effect from the first approach: step 2 of Algorithm 11 initializes each set $\mathcal{S}\left(v_{q}\right)$ with the least possible number of clauses $L \sqsubseteq L$, and ordering $<_{v_{q}}$ is as fine-grained as possible, which maximally reduces the number of inferences at context $v_{q}$. The main drawback, however, is that the number of contexts needed is linear in the size of $\mathcal{Q}$.

The third possibility is to introduce a context $v_{K}$ for each conjunction $K$ occurring in the antecedent of some query in $\mathcal{Q}$, define core $\left(v_{K}\right):=K$, and define $<_{v_{K}}$ so that each literal occurring in the consequent of some query in $\mathcal{Q}$ is $<_{v_{K}}$-minimal while all remaining literals are ordered arbitrarily. We thus strike a balance between the number of contexts and the inferences performed.

Description logic reasoners are often used to classify an ontology $\mathcal{O}$, in which case $\mathcal{Q}$ contains $A \sqsubseteq B$ for all pairs of atomic concepts $A$ and $B$. The second possibility then requires a quadratic number of contexts, which can be prohibitively large even on ontologies of moderate size. In contrast, the third possibility requires all atomic concepts in $\mathcal{O}$ to be minimal, which essentially "turns off" most ordering constraints. To address the former problem in hypertableau-based reasoners, Glimm et al. [37] developed a classification algorithm that gathers information about known and possible subsumptions from the ABoxes encountered during the algorithm's execution in order to reduce the number of pairs of atomic concepts for which $\mathcal{O} \vDash A \sqsubseteq B$ needs to be checked. This algorithm can be adapted to the consequence-based framework as well. In particular, let $\mathcal{S}$ be a clause system
 such that core $(v) \subseteq\{A\}$ and $A \sqsubseteq B \hat{\in} \mathcal{S}(v)$, by the soundness of our algorithm we have that $\mathcal{O} \vDash A \sqsubseteq B$-in other words, subsumption $A \sqsubseteq B$ is known. Furthermore, Proposition 16, shown below, tells us how to identify subsumptions that are not possible. By tightly integrating the consequence-based algorithm with the appropriately modified classification algorithm, we can extend $\mathcal{Q}$ as needed, thus considerably reducing the size of this set. Furthermore, whenever the classification algorithm adds a query $A \sqsubseteq B$ to $\mathcal{Q}$, we have two options: we can introduce a fresh context $v_{A \sqsubseteq B}$; or we can introduce or reuse context $v_{A}$, keeping in mind that, whenever we reuse an existing context, we must refine the literal ordering $<_{v_{A}}$ so that $B$ becomes $<_{v_{A}}$-minimal.

Proposition 16. Let $\mathcal{S}$ be a clause system obtained by applying Algorithm 11 to a normalized $\mathcal{A L C I}$ ontology $\mathcal{O}$ and a finite set of queries $\mathcal{Q}$. Then, $\mathcal{O} \not \vDash A \sqsubseteq B$ holds for all atomic concepts $A$ and $B$ for which there exists a context $v$ such that $A \sqsubseteq A \hat{\in} \mathcal{S}(v)$, $A \sqsubseteq B \notin \mathcal{S}(v)$, and no clause of the form $K \sqsubseteq M \sqcup B \in \mathcal{S}(v)$ exists that satisfies $K \subseteq\{A\}$ and $B \nprec_{v} M$.

Proof. The proof of this proposition relies on the proof of Theorem 6 presented in Appendix A, the notion of pre-models introduced in Appendix A.1, and the symbol $\vdash$ introduced in Appendix A.4.
 empty and therefore $<_{v}$-minimal; together with $A \sqsubseteq A \hat{\in} \mathcal{S}(v)$, these observations imply $v \nvdash A \sqsubseteq \perp$. Thus, the construction of a pre-model $I=\langle\Delta, E, J\rangle$ in Appendix A. 4 introduces an element $\delta_{A \sqsubseteq \perp}^{v} \in \Delta$ whose literal interpretation $J\left(\delta_{A \sqsubseteq \perp}^{v}\right)$ is constructed as specified in Appendix A.2. Now consider an arbitrary clause $K \sqsubseteq M \sqcup B \in \mathcal{S}(v)$ : by our assumption, at least one of $K \subseteq\{A\}$ or $B \nprec_{v} M$ does not hold; thus, the conditions of Lemma 64 are not satisfied, and so clause $K \sqsubseteq M \sqcup B$ is not productive. Since this holds for all clauses in $\mathcal{S}(v)$ with $B$ in the consequent, we have $B \notin J\left(\delta_{A \sqsubseteq \perp}^{v}\right)$. Furthermore, due to $A \sqsubseteq A \hat{\in} \mathcal{S}(v)$ and Lemma 65, we have $A \in J\left(\delta_{A \sqsubseteq \perp}^{v}\right)$. Thus, we have $I \not \vDash A \sqsubseteq B$, and so $\mathcal{O} \not \vDash A \sqsubseteq B$ holds by Corollary 61 and Claim 72 .

### 3.4.2. Expansion Strategies

In this section we discuss reasonable expansion strategies. We use the ontology $\mathcal{O}$ and query $q$ defined in the following example to demonstrate how the choice of strategy affects the inferences of our algorithm.

Example 17. Let $\mathcal{O}$ be the ontology containing axioms (31)-(38), and let $q=A \sqsubseteq E$. Clearly, $\mathcal{O} \vDash q$ holds due to (31)-(33).

$$
\begin{align*}
A & \sqsubseteq \exists R . B  \tag{31}\\
A & \sqsubseteq \forall R . C_{1}  \tag{35}\\
B \sqcap C_{1} & \sqsubseteq \forall R^{-} . E
\end{align*}
$$

$$
\begin{align*}
& A \sqsubseteq \exists S . B  \tag{36}\\
& A \sqsubseteq D \sqcup \forall S . C_{2} \tag{37}
\end{align*}
$$

$$
\begin{align*}
A \sqcap B & \sqsubseteq \exists T . F_{1} \\
C_{1} \sqcap C_{2} & \sqsubseteq C \\
B \sqcap C & \sqsubseteq \exists T . F_{2} \tag{38}
\end{align*}
$$

The simplest possibility is to always reuse the trivial context $v_{0}$; thus, we define $\operatorname{trivial}\left(\exists R . A, \mathbf{B}_{k}, \mathcal{D}\right)=\left\langle v_{0}, \emptyset, \emptyset\right\rangle$. This strategy makes most sense if the trivial context $v_{0}$ is also used for initialization. The inferences of our algorithm on $\mathcal{O}$ with the trivial strategy are shown in Figure 2. As one can see, the algorithm is prolific: inferences (47)-(57) are not needed in order to derive (46). Thus, such an algorithm resembles various resolution-based decision procedures for DLs [31].

To reduce the number of irrelevant inferences, we can use a cautious approach and introduce a context $v_{A}$ for each atomic concept $A$; thus, we define cautious $\left(\exists R . A, \mathbf{B}_{k}, \mathcal{D}\right)=\left\langle v_{A},\{A\},<\right\rangle$, where $<$ is an arbitrary $R$-admissible ordering. The number of contexts is then limited to the number of atomic concepts occurring in $\mathcal{O}$ and $\mathcal{Q}$. Please remember that, whenever context $v_{A}$ is reused, its literal ordering must be refined so that it stays admissible for each relevant role. As one can see in Figure 3, the cautious approach considerably reduces the inferences of our algorithm; for example, it does not derive clauses containing $\exists T . F_{1}$. Nevertheless, since we reuse context $v_{B}$ to satisfy existential restrictions (59) and (66), the algorithm derives clause (70), which then leads to the derivation of (71) and the introduction of context $v_{F_{2}}$.

We can further reduce the number of irrelevant inferences by adopting an eager approach and introducing a context $v_{\mathbf{A}}$ for each set of atomic concepts $\mathbf{A}$; thus, we define eager $\left(\exists R . A, \mathbf{B}_{k}, \mathcal{D}\right)=\left\langle v_{\{A\} \cup \mathbf{B}_{k}},\{A\} \cup \mathbf{B}_{k},<\right\rangle$, where $\langle$ is an arbitrary $R$-admissible ordering. As one can see in Figure 4, the algorithm now does not even derive clauses containing $\exists T . F_{2}$ and is considerably more efficient than in the previous two cases. In general, however, the eager strategy can be problematic because the number of introduced contexts can be exponential in the number of atomic concepts occurring in $\mathcal{O}$ and $\mathcal{Q}$.

One can refine all of the strategies discussed above by using a different context per role $R$. For example, the cautious strategy can be refined to use a context $v_{A}^{R}$ for each atomic concept $A$ and role $R$; thus, we define cautious ${ }_{R}\left(\exists R . A, \mathbf{B}_{k}, \mathcal{D}\right)=\left\langle v_{A}^{R},\{A\},\langle \rangle\right.$, where $<$ is an arbitrary $R$-admissible ordering. A benefit of the refined cautious strategy is that context $v_{A}^{R}$ is reused only to satisfy existential restrictions of the form $\exists R . A$, so the literal ordering used in the context never needs to be refined.

The cautious and the eager strategies appear to be most natural, striking a different balance between the number of contexts and the number of inferences in each context. In Section 8 we investigate the effectiveness of these strategies in practice. Our results suggest that, in most practical cases, the eager strategy may be most appropriate because it does not introduce too many contexts; however, we also identified cases in which the cautious strategy is more appropriate. This suggests that a hybrid approach may be useful: one uses the eager strategy until some predetermined number of contexts is introduced, after which one switches to the cautious strategy.

Although the eager strategy can, in general, introduce an exponential number of contexts, as we argue next, on $\mathcal{E L}$ and DL-Lite ${ }_{\text {horn }}$ ontologies the eager strategy introduces only linearly many contexts. First, if $\mathcal{O}$ is an $\mathcal{E} \mathcal{L}$ ontology, then all existential restrictions in $\mathcal{O}$ contain only atomic roles, and all universal restrictions in $\mathcal{O}$ contain only inverse roles; hence, the Succ rule is applicable only with $\mathbf{B}_{k}=\emptyset$, and so the eager strategy coincides with the cautious strategy. Thus, similarly to existing $\mathcal{E} \mathcal{L}$ reasoning algorithms [23, 18], the eager strategy introduces only one context for each atomic concept. Second, if $\mathcal{O}$ is a DL-Lite $_{\text {horn }}$ ontology, then all universal restrictions occur in $\mathcal{O}$ only in axioms of the form $\top \sqsubseteq \forall R . B$; hence, if we apply the Hyper rule before the Succ rule, then all clauses $T \sqsubseteq \forall R . A$ will be eagerly derived in each context, and so the Succ rule will be applicable to an existential restriction $\exists R . A$ only with $\mathbf{B}_{k}=\{B \mid \top \sqsubseteq \forall R . B \in \mathcal{O}\}$; since this set is uniquely identified by $\exists R . A$, the eager strategy will introduce at most one context for each existential restriction. Furthermore, if all existential restrictions occurring in $\mathcal{O}$ are of the form $\exists R . A_{\top}$ as discussed in Section 2.1, then, similarly to existing DL-Lite ${ }_{\text {horn }}$ reasoning algorithms [19, 20], the eager strategy will introduce at most one context for each role occurring in an existential restriction.


| Initialization: | $A \sqsubseteq A$ | (39) |
| :---: | :---: | :---: |
| Hyper[31+39]: | $A \sqsubseteq \exists R . B$ | (40) |
| Hyper[32+39]: | $A \sqsubseteq \forall R . C_{1}$ | (41) |
| Succ[40+41]: | $B \sqsubseteq B$ | (43) |
| Succ[40+41]: | $C_{1} \sqsubseteq C_{1}$ | (44) |
| Hyper[33+43+44]: | $B \sqcap C_{1} \sqsubseteq \forall R^{-} . E$ | (45) |
| Pred[40+41+45]: | $A \sqsubseteq E$ | (46) |
| Hyper[34+39]: | $A \sqsubseteq \exists S . B$ | (47) |
| Hyper[35+39]: | $A \sqsubseteq D \sqcup \forall S . C_{2}$ | (48) |
| Succ[47+48]: | $C_{2} \sqsubseteq C_{2}$ | (50) |
| Hyper[36+39+43]: | $A \sqcap B \sqsubseteq \exists T . F_{1}$ | (51) |
| Succ[51]: | $F_{1} \sqsubseteq F_{1}$ | (53) |
| Hyper[37+44+50]: | $C_{1} \sqcap C_{2} \sqsubseteq C$ | (54) |
| Hyper[38+43+44+50]: | $B \sqcap C_{1} \sqcap C_{2} \sqsubseteq \exists T_{\text {c }} F_{2}$ | (55) |
| Succ[55]: | $F_{2} \sqsubseteq F_{2}$ | (57) |

Figure 2: Inferences with the Trivial Strategy


Figure 3: Inferences with the Cautious Strategy

| Succ[75+76]: | $\top \sqsubseteq B$ | $(78)$ | Initialization: | $\top \sqsubseteq A$ | $(74)$ | Succ[75+83]: | $\top \sqsubseteq B$ |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | ---: |
| Succ[75+76]: | $\top \sqsubseteq C_{1}$ | $(79)$ | Hyper[31+74]: | $\top \sqsubseteq \exists R \cdot B$ | $(75)$ | Succ[82+83]: | $C_{2} \sqsubseteq C_{2}$ |
| Hyper[33+78+79]: | $\top \sqsubseteq \forall R^{-} . E$ | $(80)$ | Hyper[32+74]: | $\top \sqsubseteq \forall R \cdot C_{1}$ | $(76)$ |  |  |
|  |  |  | Pred[75+80]: | $\top \sqsubseteq E$ | $(81)$ |  |  |
|  |  |  | Hyper[34+74]: | $\top \sqsubseteq \exists S . B$ | $(82)$ |  |  |
|  |  |  | Hyper[35+74]: | $\top \sqsubseteq D \sqcup \forall S . C_{2}$ | $(83)$ |  |  |

Figure 4: Inferences with the Eager Strategy

Finally, we relate the above strategies to existing consequence-based algorithms. Kazakov [21] considers only the eager strategy. Simančík et al. [22] present their algorithm with the eager strategy, but their "context partitioning" optimizations can be seen as providing more cautious strategies; moreover, their ConDOR reasoner implements only the refined cautious strategy.

## 4. Analyzing And-Branching Using $\boldsymbol{\epsilon}$-Free Decompositions

In the rest of this paper we develop our framework for a parametric analysis of the complexity of reasoning. As we mentioned in the introduction, and-branching and or-branching are two main sources of complexity [10]. In this section we focus on quantifying the effects of and-branching, whereas in Section 5 we turn our attention to or-branching.

### 4.1. Intuitions

Let $\mathcal{O}$ and $q$ be as specified in Example 17, and let $\mathbf{L}$ be the set of all literals occurring in $\mathcal{O}$ and $q$; note that $|\mathbf{L}|=13$. Proposition 14 provides us with an estimate of the complexity of deciding $\mathcal{O} \vDash q$ : we can introduce at most $2^{\mid \mathbf{L |}}=2^{13}$ contexts using the eager strategy, and each set $\mathcal{S}(v)$ can contain at most $4^{|\mathbf{L}|}=4^{13}$ clauses, so the algorithm runs in time exponential in $|\mathbf{L}|=13$. This, however, is a crude estimate because both the number of contexts and the sizes of sets $\mathcal{S}(v)$ are vastly overestimated.

To obtain a better estimate of these values, in this section we extend the notion of a context structure to a decomposition of an ontology and a set of queries; since decompositions are proper extensions of context structures, one can apply our consequencebased algorithm to a decomposition by only slightly adapting the algorithm's initialization phase. The high-level principle behind our decompositions is similar to tree decompositions: both aim to identify sets of literals whose subsets can be relevant during reasoning. For example, given a vertex $v$ in a tree decomposition, set $\mathcal{L}(v)$ determines the propositional variables that should be considered in "local" propositional interpretation associated with $v$, and so the number of such interpretations is bounded by $2^{|\mathcal{L}(v)|}$. Our notion of decomposition aims for a similar goal: each context in a decomposition will describe the clauses that the consequence-based algorithm can derive, which will then determine the worst-case complexity of the algorithm. Towards this goal, please remember that there are two sources of complexity in description logic reasoning [10]: existential and universal quantification cause and-branching, and disjunction causes or-branching. The latter corresponds to or-branching in propositional logic and can be characterized using tree decompositions; however, the former is unique to description logics and requires completely new techniques. Our decompositions therefore consist of two kinds of edges: $\exists R . A$-edges are defined as in context structures and capture information propagation due to and-branching, and $\epsilon$-edges capture information propagation due to orbranching. In the rest of this section we focus on and-branching, so the decompositions we shall consider will be $\epsilon$-free (i.e., they will not contain $\epsilon$-edges). Please note that $\epsilon$-free decompositions are completely different from and do not generalize tree decompositions. In Section 5 we then show how $\epsilon$-edges can be used to capture or-branching.

As we discussed in Section 3.1, each context in a context structure describes one or more domain elements in a model of $\mathcal{O}$. A context structure can thus be seen as decomposing a description logic reasoning problem into several of interconnected propositional problems; given such a representation, the Hyper rule then locally solves each propositional subproblem, and the Pred and Succ rules handle the connections between subproblems. An $\epsilon$-free decomposition $\mathcal{D}$ for $\mathcal{O}$ and $q$ can in similar vein be understood as breaking up the problem of checking $\mathcal{O} \vDash q$ into several interconnected propositional problems, but with several differences. First, a context structure is constructed during reasoning, and so it cannot provide us with an a priori complexity estimate. In contrast, $\mathcal{D}$ is independent from any reasoning and contains "sufficiently many" contexts (i.e., all contexts needed by the Succ rule). Second, each context $v$ in $\mathcal{D}$ is associated with three sets of literals core $(v)$, $\operatorname{knw}(v)$, and $\operatorname{poss}(v)$ that describe the clauses one can derive in $\mathcal{S}(v)$. Set core $(v)$ serves the same purpose as in context structures. Furthermore, set knw( $v$ ) contains the known literals-that is, for each literal $L \in \operatorname{knw}(v)$, our algorithm will derive $T \sqsubseteq L$ in $\mathcal{S}(v)$. Finally, set poss $(v)$ contains the possible literals-that is, each clause $K \sqsubseteq M \in \mathcal{S}(v)$ will satisfy $K \cup M \subseteq \operatorname{poss}(v)$. Clearly, all core literals are known, and all known literals are possible, so $\mathcal{D}$ should satisfy core $(v) \subseteq \operatorname{knw}(v) \subseteq \operatorname{poss}(v)$ for each context $v$.

To ensure that each clause obtained by applying our consequence-based algorithm to $\mathcal{D}$ is a logical consequence of $\mathcal{O}$, we require $\mathcal{D}$ to be sound: in addition to the condition on $\exists R . A$-edges from Definition 8 , the known literals in each context of $\mathcal{D}$ must logically follow from the context's core-that is, $\mathcal{O} \vDash \operatorname{core}(v) \sqsubseteq L$ should hold for each context $v$ in $\mathcal{D}$ and each literal $L \in \operatorname{knw}(v)$. Furthermore, to ensure completeness of our algorithm, we require $\mathcal{D}$ to be admissible: roughly speaking, this notion requires $\mathcal{D}$ to contain all contexts that might be needed during our algorithm's execution, and sets core $(v)$, $\operatorname{knw}(v)$, and poss $(v)$ to correctly describe the types of clauses that our algorithm can derive in $\mathcal{S}(v)$. This will ensure that (i) the Succ rule is never applicable, and (ii) each clause in $\mathcal{S}(v)$ is of the form $T \sqsubseteq L$ with $L \in \operatorname{knw}(v)$, or of the form $K \sqsubseteq M$ with $K \cup M \subseteq \operatorname{poss}(v) \backslash \operatorname{knw}(v)$. These observations allow us to capture the complexity of our algorithm using the following two parameters: decomposition length $\ln (\mathcal{D})$, defined as the number of contexts in $\mathcal{D}$, captures the combinatorial difficulty due to (i); and decomposition width $\mathrm{wd}(\mathcal{D})$, defined as the maximum cardinality of $\operatorname{poss}(v) \backslash k n w(v)$, captures the combinatorial difficulty due to (ii). We show that, when applied to $\mathcal{D}$, our consequence-based algorithm runs in time polynomial in $\ln (\mathcal{D})$ and $4^{\mathrm{wd}(\mathcal{D})^{2}}$; the latter factor is due to the fact that each set $\mathcal{S}(v)$ can contain at most $|k n w(v)|+4^{|\operatorname{poss}(v)| k n w(v) \mid}$ clauses.

An ontology and a set of queries can admit infinitely many sound, admissible, and $\epsilon$-free decompositions, so a decomposition $\mathcal{D}$ with least $\ln (\mathcal{D})$ and $\operatorname{wd}(\mathcal{D})$ would provide us with a "general" difficulty of reasoning with a particular ontology and a set


Figure 5: Examples of $\epsilon$-Free Decompositions
of queries. This would be analogous to propositional satisfiability, where the treewidth of a propositional problem is defined as the smallest width over all tree decompositions of the problem. However, identifying such $\mathcal{D}$ is difficult for two reasons. First, checking semantic conditions in the definition of soundness requires reasoning, so verifying decomposition soundness can be as hard as the reasoning problem whose complexity we aim to analyze. Second, ontologies and queries exist for which all decompositions of minimal width have exponential length; we prove this claim for general (i.e., not only $\epsilon$-free) decompositions, so we postpone the formal treatment of this property until Section 7. Hence, minimizing width can increase length and vice versa, which essentially prevents us from obtaining a "general" characterization of reasoning difficulty.

As a possible solution to these challenges, in Section 4.4 we present a practical algorithm that can construct in polynomial time an $\epsilon$-free decomposition of $\mathcal{O}$ and $\mathcal{Q}$. To address the first issue, our algorithm approximates the entailment relation used in the definition of soundness. Thus, for some context $v$, set $k n w(v)$ approximates the set of literals entailed by core $(v)$, and for each literal in $\operatorname{knw}(v)$ the relevant entailment can be established in polynomial time. To address the second issue, the algorithm is given a set of parameters $\mathcal{C}$ called a control that consists of an expansion strategy, an integer mln that imposes a bound on decomposition length, and two other parameters that determinize the decomposition construction process. For each $\mathcal{C}$, the $\epsilon$ free decomposition of $\mathcal{O}$ and $\mathcal{Q}$ that is deterministically constructed by our algorithm is called the $\mathcal{C}$-decomposition of $\mathcal{O}$ and $\mathcal{Q}$. Parameter min bounds the length of the $\mathcal{C}$-decomposition, which determines the decomposition width. Thus, the width and length of a $\mathcal{C}$-decomposition provide us with an account of the reasoning difficulty relative to $\mathcal{C}$.

To illustrate these ideas, Figure 5 shows two $\epsilon$-free decompositions of $\mathcal{O}$ and $q$ from Example 17 on page 15 obtained using the cautious and the eager expansion strategy. We use the same notation as for context structures, but show sets core $(v)$, $\operatorname{knw}(v) \backslash \operatorname{core}(v)$, and $\operatorname{poss}(v) \backslash \mathrm{knw}(v)$ in a table below each context.

The $\epsilon$-free decomposition $\mathcal{D}_{1}$ constructed from $\mathcal{O}$ and $q$ using the cautions strategy is shown in Figure 5a. Since our goal is to check the entailment of $q=A \sqsubseteq E$, we introduce context $v_{A}$ and add $A$ to core $\left(v_{A}\right)$, thus capturing the fact that the consequencebased algorithm initializes $\mathcal{S}\left(v_{A}\right)$ with (58). But then, from clauses (31), (32), and (34), the consequence-based algorithm can derive (59), (60), and (66); to reflect this, we add literals $\exists R . B, \forall R . C_{1}$, and $\exists S . B$ to $\mathrm{knw}\left(v_{A}\right)$. Furthermore, using clause (35), the consequence-based algorithm will derive (67); however, the consequent of the clause contains a disjunction, so neither of the literals is derived deterministically, and therefore we add $D$ and $\forall S . C_{2}$ to poss $\left(v_{A}\right)$. Next, we must satisfy existential restrictions $\exists R . B$ and $\exists S . B$; since we use the cautious strategy, we introduce a single context $v_{B}$, add the two edges from $v_{A}$ to $v_{B}$, and initialize core $\left(v_{B}\right)$ to $B$; furthermore, due to universal restrictions $\forall R . C_{1} \in \operatorname{knw}\left(v_{A}\right) \subseteq \operatorname{poss}\left(v_{A}\right)$ and $\forall S . C_{2} \in \operatorname{poss}\left(v_{A}\right)$, we add $C_{1}$ and $C_{2}$ to $\operatorname{poss}\left(v_{B}\right)$. Note that, due to the reuse of $v_{B}$, we cannot add $C_{1}$ to $\mathrm{knw}\left(v_{B}\right)$ : not all domain elements represented by $v_{B}$ necessarily satisfy $C_{1}$. We now continue our estimation for $v_{B}$ : literal $\forall R^{-} . E$ is possible due to (33), concept $C$ is possible due to (37), and literal $\exists T . F_{2}$ is possible due to (38). Moreover, since $\forall R^{-} . E \in \operatorname{poss}\left(v_{B}\right)$, literal $E$ is possible for $v_{A}$. Note that the consequence-based algorithm derives (65); however, $\forall R^{-} . E$ is not in $\operatorname{knw}\left(v_{B}\right)$ due to context reuse, so we do not have sufficient information to add $E$ to $\mathrm{knw}\left(v_{A}\right)$; in other words, $\mathrm{knw}\left(v_{A}\right)$ provides us only with a lower bound on "truly" known literals. Finally, we introduce context $v_{F_{2}}$ to satisfy $\exists T . F_{2} \in \operatorname{poss}\left(v_{B}\right)$ and thus obtain an admissible $\epsilon$-free decomposition. Clearly, $\operatorname{wd}\left(\mathcal{D}_{1}\right)=5$ and $\ln \left(\mathcal{D}_{1}\right)=3$ so, when applied to $\mathcal{D}_{1}$, our inference rules can derive $3 \cdot 4^{5}$ clauses, which is a much more accurate estimate.

The $\epsilon$-free decomposition $\mathcal{D}_{2}$ constructed from $\mathcal{O}$ and $q$ using the eager strategy is shown in Figure 5 b. The main difference compared to the previous case is that, due to the eager strategy, we now use distinct contexts, $v_{\left\{B, C_{1}\right\}}$ and $v_{\{B\}}$, to satisfy existential restrictions $\exists R . B$ and $\exists S . B$ at context $v_{\{A\}}$. Since literal $\forall R . C_{1}$ is known at $v_{\{A\}}$, we can now add $C_{1}$ to the core of $v_{\left\{B, C_{1}\right\}}$, which in turn allows us to make $E$ known at $v_{\{A\}}$. Furthermore, literal $\exists T . F_{2}$ is now not possible at $v_{\{B\}}$. In this case, we have $w d\left(\mathcal{D}_{2}\right)=2$ and $\ln \left(\mathcal{D}_{1}\right)=3$ so, when applied to $\mathcal{D}_{2}$, our inference rules can derive at most $3 \cdot 4^{2}$ clauses; thus, we have obtained an even better estimate of the running time of our consequence-based algorithm.

Reasoning with Horn- $\mathcal{A L C I}$ ontologies is generally easier in practice: such ontologies admit a single canonical model, which eliminates all or-branching; however, due to and-branching, reasoning is still ExpTime-hard [17]. This is reflected in our notions of decompositions. In particular, if control $\mathcal{C}$ uses the eager expansion strategy and allows for "sufficiently many" contexts, the $\mathcal{C}$-decomposition $\mathcal{D}$ of a Horn ontology $\mathcal{O}$ has zero width. Then, all possible literals in $\mathcal{D}$ are also known, no or-branching is required, and so $\mathcal{D}$ contains a complete solution to the reasoning problem. By reducing the maximum number of contexts
in $\mathcal{C}$, however, we can "trade off" and-branching for or-branching. Then, $\mathcal{D}$ does not necessarily have zero width, so further reasoning is needed; moreover, when applied to $\mathcal{D}$, our consequence-based algorithm derives only Horn clauses, and the size of the antecedent of each clause is bounded by $\operatorname{wd}(\mathcal{D})$. Furthermore, as we discussed in Section 3.4.2, if $\mathcal{O}$ is an $\mathcal{E} \mathcal{L}$ or a DL-Lite ${ }_{\text {horn }}$ ontology, the eager strategy introduces only linearly many contexts. Therefore, if $\mathcal{C}$ uses the eager strategy, the $\mathcal{C}$-decomposition of $\mathcal{O}$ has zero width and linear length, thus explaining the tractability of subsumption reasoning in $\mathcal{E} \mathcal{L}$ and DL-Lite ${ }_{\text {horn }}$.

It is instructive to consider the case when $\mathcal{O}$ and $\mathcal{Q}$ contain only propositional clauses-that is, if all literals in $\mathcal{O}$ and $\mathcal{Q}$ are atomic concepts. Then, for each admissible $\epsilon$-decomposition $\mathcal{D}$ of $\mathcal{O}$ and $\mathcal{Q}$, the $\exists R$. $A$-edges are not needed and can be removed, but the decomposition may still contain multiple contexts that are used to (dis)prove entailment of different queries in $\mathcal{Q}$. Thus, only the Hyper rule is used when the consequence-based calculus is applied to $\mathcal{D}$, which essentially amounts to solving the propositional problem(s) using hyperresolution. However, since $k n w(v)$ contains literals that deterministically follow from $\mathcal{O}$ for each context $v$, hyperresolution is restricted to the unknown literals.

Please note that neither redundancy elimination not literal orderings are needed to obtain our complexity results: the results hold even if each context is associated with the empty literal ordering, and if no redundant clauses are eliminated. We leave a study of whether redundancy elimination and literal orderings can provide us with a tighter complexity bound for future work.

### 4.2. Formalization

Throughout this section we fix a normalized $\mathcal{A L C I}$ ontology $\mathcal{O}$ and a finite set of queries $\mathcal{Q}$; furthermore, we let $\mathbf{L}$ be the set of literals that occur in $\mathcal{O} \cup \mathcal{Q}$, and we let $\epsilon$ be a special symbol not occurring in $\mathbf{L}$. We next formalize the notion of decompositions and then generalize the notion of context structure soundness to decompositions. In order to simplify the presentation, Definitions 18 and 19 both consider $\epsilon$-edges that are used only in Section 5 for the analysis of or-branching; however, in this section we subsequently consider only decompositions without $\epsilon$-edges.

Definition 18. A decomposition of $\mathcal{O}$ and $\mathcal{Q}$ is a tuple $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, knw, poss, $\langle, \vartheta\rangle$ where $\mathcal{V} \subseteq \mathbf{X}$ is a finite set of contexts, set $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V} \times\left(\mathbf{L}^{\exists} \cup\{\epsilon\}\right)$ contains edges labeled by existential restrictions ( $\exists R$. $A$-edges) or by $\epsilon$ ( $\epsilon$-edges), functions core, knw, poss : $\mathcal{V} \rightarrow 2^{\mathbf{L}}$ label contexts with sets of literals, function $<$ assigns to each context $v \in \mathcal{V}$ a literal ordering $<_{v}$, and function $\vartheta: \mathcal{Q} \rightarrow \mathcal{V}$ maps queries to contexts. For each vertex $v \in \mathcal{V}$, let $u k n w(v):=\operatorname{poss}(v) \backslash \operatorname{knw}(v)$. The length of $\mathcal{D}$ is $\ln (\mathcal{D}):=|\mathcal{V}|$, and the width of $\mathcal{D}$ is $\operatorname{wd}(\mathcal{D}):=\max _{v \in \mathcal{V}} \mid$ uknw $(v) \mid$. Decomposition $\mathcal{D}$ is $\epsilon$-free if it contains no $\epsilon$-edges.

Definition 19. A decomposition $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, knw, poss, $\langle, \vartheta\rangle$ of $\mathcal{O}$ and $\mathcal{Q}$ is sound if it satisfies all of the following conditions.
(N1) $\mathcal{O} \vDash \operatorname{core}(v) \sqcap \exists R . A \sqsubseteq \exists R$. $[\operatorname{core}(u) \sqcap A]$ for each $\exists R$.A-edge $\langle v, u, \exists R . A\rangle \in \mathcal{E}$.
$(N 2) \operatorname{core}(v)=\operatorname{core}(u)$ for each $\epsilon$-edge $\langle v, u, \epsilon\rangle \in \mathcal{E}$.
(N3) $\mathcal{O} \vDash \operatorname{core}(v) \sqsubseteq L$ for each context $v \in \mathcal{V}$ and each $L \in \operatorname{knw}(v)$.
Condition (N1) of Definition 19 is the same as in Definition 8; condition (N2) is relevant only for decompositions that are not $\epsilon$-free and will be explained in Section 5; and condition (N3) says that known literals should hold in all domain elements represented by a specific context. Please note that condition (N3) imposes only an upper bound on the known literals; thus, $\mathrm{knw}(v)$ can be obtained using a sound, but incomplete technique such as the one we present in Section 4.4.

Next, we extend the notion of context structure admissibility to decompositions. The ordering condition is the same as in context structures. Furthermore, to encapsulate all relevant parameters for the consequence-based algorithm, a decomposition contains a covering mapping, and the covering condition restates Definition 10. The ontology condition says that, if $\mathcal{O}$ contains a clause $K \sqsubseteq M$ and all of the literals in $K$ are possible at context $v$, then the Hyper rule can derive $M$ and so all literals in $M$ should be possible at $v$ too. Finally, the structural condition (S1) captures the intuitive relationship between the core, known, and possible literals; condition (S2) ensures that a decomposition correctly estimates all consequences of the Pred rule; and condition (S3) ensures that a decomposition contains sufficiently many contexts so that the Succ rule is never applicable.

Definition 20. Let $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, knw, poss, $\langle, \vartheta\rangle$ be an $\epsilon$-free decomposition of $\mathcal{O}$ and $\mathcal{Q}$. Let $v \in \mathcal{V}$ be an arbitrary context of $\mathcal{D}$, and let $K \sqsubseteq M$ be an arbitrary query; then, $v$ is sound for $K \sqsubseteq M$ if core $(v) \subseteq K$; $v$ is complete for $K \sqsubseteq M$ if $K \subseteq \operatorname{poss}(v)$ and $M$ is $<_{v}$-minimal; and $v$ covers $K \sqsubseteq M$ if $v$ is both sound and complete for $K \sqsubseteq M$. Finally, $\mathcal{D}$ is admissible if all of the following conditions are satisfied.

- Structural conditions:
(S1) For each context $v \in \mathcal{V}$, we have $\operatorname{core}(v) \subseteq \operatorname{knw}(v) \subseteq \operatorname{poss}(v)$.
(S2) For each edge $\langle v, u, \exists R . A\rangle \in \mathcal{E}$ and each literal $\forall \operatorname{inv}(R) . C \in \operatorname{poss}(u)$, we have $C \in \operatorname{poss}(v)$.
(S3) For each literal $\exists R . A \in \operatorname{poss}(v)$, an edge $\langle v, u, \exists R . A\rangle \in \mathcal{E}$ exists such that $\operatorname{poss}(u) \supseteq\{A\} \cup\{B \mid \forall R . B \in \operatorname{poss}(v)\}$.
- Ordering condition: For each edge $\langle v, u, \exists R . A\rangle \in \mathcal{E}$, the literal ordering $\prec_{u}$ is $R$-admissible.
- Ontology condition: For each context $v \in \mathcal{V}$ and each clause $K \sqsubseteq M \in \mathcal{O}$ with $K \subseteq \operatorname{poss}(v)$, we have $M \subseteq \operatorname{poss}(v)$.
- Covering condition: Each query $q \in \mathcal{Q}$ is covered in the context $\vartheta(q)$.

We can apply the consequence-based algorithm to a sound, admissible, and $\epsilon$-free decomposition $\mathcal{D}$ as specified in Algorithm 21. Step 1 initializes each $\mathcal{S}(v)$ with a clause $T \sqsubseteq L$ for each known literal $L \in \operatorname{knw}(v)$. This allows us to prove in Lemma 22 that the algorithm derives only clauses of the form $K \sqsubseteq M$ with $K \cup M \subseteq u k n w(v)$, which is central to our complexity argument in Proposition 24. Please note that, if step 1 were to consider only the literals in core( $v$ ), our algorithm would eventually derive a clause $T \sqsubseteq L$ for each known literal $L \in \operatorname{knw}(v)$, but this might involve clauses containing combinations of literals in $\operatorname{knw}(v)$; thus, step 1 introduces derivation "shortcuts" for the known literals, which can improve the algorithm's performance in practice. Step 2 is analogous to the way the Succ rule initializes contexts, and it ensures that each atomic concept $B$ that can be propagated to a context $u$ through existential or universal quantification is relevant for $u$. Step 3 ensures that each query $q \in \mathcal{Q}$ is covered in context $\vartheta(q)$. Finally, step 4 applies exhaustively the consequence-based inference rules, but without the Succ rule. The latter is possible because $\mathcal{D}$ is admissible and thus contains all relevant contexts, which allows us to prove in Lemma 23 that the Succ rule is never applicable during the algorithm's execution. By combining this observation with the fact that $\mathcal{D}$ is sound, in Proposition 24 we show that our algorithm is sound and complete.

Algorithm 21. The $\epsilon$-free decomposition-based reasoning algorithm for $\mathcal{A L C \mathcal { I }}$ takes a sound, admissible, and $\epsilon$-free decomposition $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, knw, poss, $\langle, \vartheta\rangle$ of $\mathcal{O}$ and $\mathcal{Q}$. The algorithm (nondeterministically) computes a clause system $\mathcal{S}$ for $\mathcal{D}$ as follows.

1. Set $\mathcal{S}(v):=\{\top \sqsubseteq L \mid L \in \operatorname{knw}(v)\}$ for each context $v \in \mathcal{V}$.
2. For each edge $\langle v, u, \exists R . A\rangle \in \mathcal{E}$ and each concept $B \in \operatorname{uknw}(u)$ such that $B=A$ or $\forall R . B \in \operatorname{poss}(v)$, add $B \sqsubseteq B$ to $\mathcal{S}(u)$.
3. For each query $K \sqsubseteq M \in \mathcal{Q}$, let $v:=\vartheta(K \sqsubseteq M)$; then, for each literal $L \in K \backslash \operatorname{knw}(u)$, add $L \sqsubseteq L$ to $\mathcal{S}(v)$.
4. Exhaustively apply rules Hyper, Pred, and Elim from Table 3.

Lemma 22. Let $\mathcal{S}$ be an arbitrary clause system encountered during step 4 of Algorithm 21, and let $v \in \mathcal{V}$ be an arbitrary context in $\mathcal{D}$. Then, each clause in $\mathcal{S}(v)$ is of the form $T \sqsubseteq L$ with $L \in \operatorname{knw}(v)$, or $K \sqsubseteq M$ with $K \cup M \subseteq u k n w(v)$.

Proof. Let $\mathcal{S}$ be a clause system at the beginning of step 4 , and let $v \in \mathcal{V}$ be an arbitrary context. By Definition 20, if context $v$ covers some query $K \sqsubseteq M$, then $K \subseteq \operatorname{poss}(v)$ holds; thus, each clause added to $\mathcal{S}(v)$ in steps $1-3$ is clearly of the required form. We next show that this property is preserved in step 4.

Assume that the Hyper rule is applicable to some clause $\prod_{i=1}^{n} A_{i} \sqsubseteq M \in \mathcal{O}$. By the induction assumption, each premise from $\mathcal{S}(v)$ is of the form $K_{i} \subseteq M_{i} \sqcup A_{i}$ with $K_{i} \cup M_{i} \subseteq u k n w(v)$ and $A_{i} \in \operatorname{poss}(v)$. Thus, $\square_{i=1}^{n} A_{i} \subseteq \operatorname{poss}(v)$, so $M \subseteq \operatorname{poss}(v)$ by the ontology condition of Definition 20. The Hyper rule produces the clause $\prod_{i=1}^{n} K_{i} \sqsubseteq M \sqcup \bigsqcup_{i=1}^{n} M_{i}$. Now assume that there exists a literal $L \in M \cap \operatorname{knw}(v)$; then, $T \sqsubseteq L \hat{\in} \mathcal{S}(v)$ due to step 1 of Algorithm 21 and Lemma 12, which contradicts the assumption that the Hyper rule is applicable. Hence, $M \subseteq u k n w(v)$, and the produced clause is of the required form.

Assume that the Pred rule is applicable to an edge $\langle v, u, \exists R . A\rangle, \in \mathcal{E}$. By the induction assumption, the premises of the rule are of the form $K_{0} \sqsubseteq M_{0} \sqcup \exists R . A$ and $K_{i} \sqsubseteq M_{i} \sqcup \forall R . B_{i}$, where $K_{i} \cup M_{i} \subseteq u k n w(v)$ for each $0 \leq i \leq n$. The Pred rule produces the clause $\prod_{i=0}^{n} K_{i} \sqsubseteq \bigsqcup_{i=0}^{n} M_{i} \sqcup \bigsqcup_{j=1}^{m} C_{j}$. Now for each $1 \leq j \leq m$, by the induction assumption we have $\forall \operatorname{inv}(R) . C_{j} \in \operatorname{poss}(u)$, so $C_{j} \in \operatorname{poss}(v)$ by condition (S2) of Definition 20. Furthermore, if $C_{j} \in \operatorname{knw}(v)$, then we have T $\sqsubseteq C_{j} \hat{\epsilon} \mathcal{S}(v)$ due to step 1 of Algorithm 21 and Lemma 12, which contradicts the assumption that the Pred rule is applicable; consequently, we have $C_{j} \in \operatorname{uknw}(v)$. Hence, the clause produced by the rule is of the required form.

Lemma 23. The Succ rule is never applicable during step 4 of Algorithm 21.
Proof. Let $\mathcal{S}$ be a clause system encountered in step 4 of Algorithm 21; let $v \in \mathcal{V}$ be a context; let $K \sqsubseteq M \sqcup \exists R . A \in \mathcal{S}(v)$ be a clause; and let $\mathbf{B}_{p}$ be as specified in Table 3. By Lemma 22, we have $\{B \mid \forall R . B \in \operatorname{poss}(v)\} \supseteq \mathbf{B}_{p}$ and $\exists R . A \in \operatorname{poss}(v)$. Thus, by condition (S3) of Definition 20, an edge $\langle v, u, \exists R . A\rangle \in \mathcal{E}$ exists with $\operatorname{poss}(u) \supseteq\{A\} \cup\{B \mid \forall R . B \in \operatorname{poss}(v)\} \supseteq\{A\} \cup \mathbf{B}_{p}$; furthermore, due to steps 1 and 2 of Algorithm 21 and Lemma 12, we have that $L \sqsubseteq L \hat{\in} \mathcal{S}(u)$ holds for each $L \in\{A\} \cup \mathbf{B}_{p}$. Thus, the Succ rule is not applicable to context $v$ and clause $K \sqsubseteq M \sqcup \exists R . A$ in $\mathcal{S}(v)$.

Proposition 24. Let $\mathcal{S}$ be a clause system obtained by applying Algorithm 21 to $\mathcal{O}, \mathcal{Q}$, and $\mathcal{D}$. Then, for each query $q \in \mathcal{Q}$, we have $\mathcal{O} \vDash q$ if and only if $q \hat{\in} \mathcal{S}(\vartheta(q))$.

Proof. Consider an arbitrary query $q=K \sqsubseteq M \hat{\in} \mathcal{Q}$ and let $v=\vartheta(q)$; context $v$ covers $q$ by the covering condition of Definition 20, Decomposition $\mathcal{D}$ is sound, so, by condition (N3) of Definition 19, we have $\mathcal{O} \vDash \operatorname{core}(u) \sqsubseteq L$ for each clause $\top \sqsubseteq L$ added to some $\mathcal{S}(u)$ in step 1 of Algorithm 21; furthermore, clauses added in steps 2 and 3 are tautologies; consequently, $\mathcal{S}$ is sound for $\mathcal{O}$. But then, since context $v$ is sound for $q$, we have that $K \sqsubseteq M \hat{\in} \mathcal{S}(v)$ implies $\mathcal{O} \vDash K \sqsubseteq M$ in the same way as in the
proof of Proposition 13. Furthermore, for an arbitrary literal $L \in K$, either $\mathrm{T} \sqsubseteq L$ is added to $\mathcal{S}(v)$ in step 1 , or $L \sqsubseteq L$ is added to $\mathcal{S}(v)$ in step 3 ; either way, we have $K \sqsubseteq L \hat{\epsilon} \mathcal{S}(v)$, and this property is preserved during the algorithm's execution by Lemma 12. Finally, Lemma 23 ensures that $\mathcal{S}$ is closed under the Hyper, Pred, and Succ rules; but then, since context $v$ is complete for $q$, we have that $\mathcal{O} \vDash K \sqsubseteq M$ implies $K \sqsubseteq M \hat{\in} \mathcal{S}(v)$ by Theorem 6 .

Proposition 25 provides us with our ultimate goal in this section: it shows that decomposition width and length provide us with a more accurate estimate of the complexity of subsumption reasoning.
Proposition 25 (Termination). When applied to some $\mathcal{O}, \mathcal{Q}$, and a sound, admissible, and $\epsilon$-free decomposition $\mathcal{D}$ of $\mathcal{O}$ and $\mathcal{Q}$, Algorithm 21 terminates in time polynomial in $4^{\operatorname{wd}(\mathcal{D})^{2}}, \ln (\mathcal{D})$, and $\|\mathcal{O}\|+\|\mathcal{Q}\|$.

Proof. Let $d=\operatorname{wd}(\mathcal{D})$ and $m=\ln (\mathcal{D})$, and recall that $|\mathbf{L}| \leq\|\mathcal{O}\|+\|\mathcal{Q}\|$. We next show that the number of inferences is polynomial in $4^{d^{2}}, m$, and $\|\mathcal{O}\|+\|\mathcal{Q}\|$; this observation implies the claim of this proposition in the same way as in the proof of Proposition 14. By Lemma 22, for each context $v$, each clause in $\mathcal{S}(v)$ is of the form

- T $\sqsubseteq L$ with $L \in \operatorname{knw}(v)$, and set $\mathcal{S}(v)$ can contain at most $|\mathbf{L}|$ such clauses, or
- $K \sqsubseteq M$ with $K \cup M \subseteq u k n w(v)$, and set $\mathcal{S}(v)$ can contain at most $c=4^{d}$ such clauses.

For the Hyper rule, a context $v \in \mathcal{V}$ can be selected in at most $m$ ways, and a clause $\prod_{i} A_{i} \sqsubseteq M \in \mathcal{O}$ can be selected in at most $|\mathcal{O}|$ ways. For each $A_{i} \in \operatorname{uknw}(v)$ there are at most $c$ ways of choosing a matching clause $K_{i} \sqsubseteq M_{i} \sqcup A_{i} \in \mathcal{S}(v)$, and for each $A_{i} \in \operatorname{knw}(v)$ there is exactly one way of choosing a matching clause $T \sqsubseteq A_{i} \in \mathcal{S}(v)$. Consequently, there are at most $c^{d}$ ways of choosing the premises $K_{i} \sqsubseteq M_{i} \sqcup A_{i} \in \mathcal{S}(v)$.

For the Pred rule, there are at most $m^{2} \cdot|\mathbf{L}|$ ways of choosing the edge $\langle v, u, \exists R . A\rangle \in \mathcal{E}$, at most $|\mathbf{L}|+4^{d}$ ways of choosing the clause $A \sqcap\rceil_{i} B_{i} \sqsubseteq \bigsqcup_{j} \forall \operatorname{inv}(R) . C_{j} \in \mathcal{S}(u)$ or $\prod_{i} B_{i} \sqsubseteq \bigsqcup_{j} \forall \operatorname{inv}(R) . C_{j} \in \mathcal{S}(u)$, and, analogously to the Hyper rule, $c^{d}$ ways of choosing the premises $K_{0} \sqsubseteq M_{0} \sqcup \exists R . A \in \mathcal{S}(v)$ and $K_{i} \sqsubseteq M_{i} \sqcup \forall R . A \in \mathcal{S}(v)$.

We would like to point out that the factor $4^{\mathrm{wd}(\mathcal{D})^{2}}$ in Proposition 25 is due to the fact that the Hyper and the Pred rules are based on hyperresolution: they fix one of the $4^{\mathrm{wd}(\mathcal{D})}$ clauses, and then try to find matches for each of the wd $(\mathcal{D})$ literals in the antecedent, which gives rise to $4^{\mathrm{wd}(\mathcal{D})^{2}}$ combinations. It is, however, possible to develop a binary version of these inference rules, in which case the algorithm's complexity would be polynomial in $4^{\mathrm{wd}(\mathcal{D})}$. In our experience, hyperresolution tends to be more effective than binary resolution in practice, which is why we opted for this style of presentation.

## 4.3. $\epsilon$-Free Decompositions and the Hypertableau Algorithm

In this section we show that $\epsilon$-free decompositions explain to an extent the performance of the hypertableau algorithm and are thus not specific to consequence-based algorithms. Since the hypertableau algorithm cannot answer several queries at once, we consider only the case when set $\mathcal{Q}$ contains a single query of the form $A \sqsubseteq \perp$ for $A$ an atomic concept-that is, when our goal is to determine the satisfiability of an atomic concept. The following theorem provides us with a key insight.

Theorem 26. Let $\mathcal{O}$ be a normalized $\mathcal{A L C I}$ ontology, let $q=A \sqsubseteq \perp$, let $D$ be a derivation of the hypertableau algorithm for $\mathcal{O}$ and the ABox $\{A(a)\}$, and let $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, knw, poss, $\langle, \vartheta\rangle$ be an admissible and $\epsilon$-free decomposition of $\mathcal{O}$ and $\{q\}$. Then, for each ABox $\mathcal{A}_{0}$ labeling a vertex of $D$, there exists a mapping $\mu: \operatorname{ind}\left(\mathcal{A}_{0}\right) \rightarrow \mathcal{V}$ such that
(a) if $L(s) \in \mathcal{A}_{0}$, then $L \in \operatorname{poss}(\mu(s))$, and
(b) if $R(s, t) \in \mathcal{A}_{0}$, then an atomic concept $A$ exists such that $\langle\mu(s), \mu(t), \exists R . A\rangle \in \mathcal{E}$.

Proof. We prove the claim by induction of the depth of $D$, but we strengthen property $(b)$ to the following property $\left(b^{\prime}\right)$ :
$\left(b^{\prime}\right)$ if $R(s, t) \in \mathcal{A}_{0}$, then some $A$ exists such that $\langle\mu(s), \mu(t), \exists R . A\rangle \in \mathcal{E}$ and $\operatorname{poss}(\mu(t)) \supseteq\{A\} \cup\{B \mid \forall R . B \in \operatorname{poss}(\mu(s))\}$.
For the induction base, let $\mathcal{A}_{0}=\{A(a)\}$ be the ABox labeling the root of $D$, and let $\mu(a)=\vartheta(q)$. Then $\mathcal{L}_{\mathcal{A}_{0}}(a)=\{A\} \subseteq \operatorname{poss}(\mu(a))$ holds by the covering condition from Definition 20, as required for property ( $a$ ); moreover, property ( $b^{\prime}$ ) holds vacuously. For the inductive step, let $\mathcal{A}_{0}$ be an ABox satisfying the property for some $\mu$, and consider an ABox obtained from $\mathcal{A}_{0}$ as follows.

Assume that the Hyp-rule derives $L_{j}(s) \in \mathcal{A}_{j}$ from $\left\{A_{1}(s), \ldots, A_{m}(s)\right\} \subseteq \mathcal{A}_{0}$ and $A_{1} \sqcap \ldots \sqcap A_{m} \sqsubseteq L_{1} \sqcup \ldots \sqcup L_{n} \in \mathcal{O}$. By the induction assumption, we have $\left\{A_{1}, \ldots, A_{m}\right\} \subseteq \mathcal{L}_{\mathcal{A}_{0}}(s) \subseteq \operatorname{poss}(\mu(s))$; but then, since $\mathcal{D}$ satisfies the ontology condition from Definition 20, we have $L_{j} \in \operatorname{poss}(\mu(s))$, as required for property (a).

Assume that the $\exists$-rule derives $\{R(s, t), A(t)\} \subseteq A_{1}$ from $\exists R . A(s) \in \mathcal{A}_{0}$. Then $\exists R . A \in \operatorname{poss}(\mu(s))$ holds by the induction assumption, so a context $u \in \mathcal{V}$ exists that satisfies condition (S3) from Definition 20. We extend $\mu$ by defining $\mu(t):=u$. ABox $\mathcal{A}_{1}$ and the extended mapping $\mu$ then clearly satisfy properties $(a)$ and $\left(b^{\prime}\right)$.

Assume that the $\forall^{+}$-rule derives $B(t) \in \mathcal{A}_{1}$ from $\{\forall R . B(s), R(s, t)\} \subseteq \mathcal{A}_{0}$. By the induction assumption, property (a) then implies $\forall R . B \in \operatorname{poss}(\mu(s))$, and property $\left(b^{\prime}\right)$ is satisfied for some atomic concept $A$, which immediately implies $B \in \operatorname{poss}(\mu(t))$, as required for property $(a)$.

Assume that the $\forall^{-}$-rule derives $B(s) \in \mathcal{A}_{1}$ from $\{[\forall \operatorname{inv}(R) . B](t), R(s, t)\} \subseteq \mathcal{A}_{0}$. By the induction assumption, property (a) then implies $\forall \operatorname{inv}(R) \cdot B \in \operatorname{poss}(\mu(t))$, and property $\left(b^{\prime}\right)$ is satisfied for some atomic concept $A$. But then, condition (S2) from Definition 20 implies $B \in \operatorname{poss}(\mu(s))$, as required for property $(a)$.

Intuitively, Theorem 26 says that each ABox $\mathcal{A}_{0}$ labeling a derivation vertex can be "embedded" via some mapping $\mu$ into each admissible and $\epsilon$-free decomposition $\mathcal{D}$ of $\mathcal{O}$ and $\mathcal{Q}$. Thus, for each individual $s \in \operatorname{ind}\left(\mathcal{A}_{0}\right)$, set $\operatorname{poss}(\mu(s))$ provides us with an upper bound on $\mathcal{L}_{\mathcal{A}_{0}}(s)$ —a set used in the definition of blocking. Now let $\ell=\max _{v \in V}|\operatorname{poss}(v)|$; then, for each context in $\mathcal{D}$ we can have at most $2^{\ell}$ different labels, so the total number of different labels is bounded by $\wp=\ln (\mathcal{D}) \cdot 2^{\ell}$; in other words, $\wp$ provides us with an upper bound on the number of nonblocked individuals in $\mathcal{A}_{0}$, which is analogous to Proposition 25. As we explained in Section 2.2, however, the number of indirectly blocked individuals can be exponentially larger due to dynamic blocking, but our decompositions do not analyze the computations caused by this effect.

Furthermore, assume that we modify the algorithm in two ways: (i) we explicitly maintain the mapping $\mu$ from Theorem 26, and (ii) whenever the $\exists$-rule introduces a new individual $t$, we immediately derive $L(t)$ for each literal $L \in \operatorname{knw}(\mu(t))$. Provided that $\mathcal{D}$ is sound, these changes do not affect the soundness of the algorithm, but they ensure that also $L(s) \in \mathcal{A}_{0}$ holds in Theorem 26 for each individual $s \in \operatorname{ind}\left(\mathcal{A}_{0}\right)$ and each literal $L \in \operatorname{knw}(\mu(s))$. Then, the number of different labels for each context is bounded by $2^{\mathrm{wd}(\mathcal{D})}$, so the total number of labels and the number of nonblocked individuals is bounded by $\ln (\mathcal{D}) \cdot 2^{\mathrm{wd}(\mathcal{D})}$. Please note that the standard hypertableau algorithm would eventually derive all literals introduced in (ii), but since the order in which derivation rules are applied is don't-care nondeterministic, this may occur only after a substantial amount of work. In contrast, modification (ii) eagerly introduces all known facts, which, through "earlier" blocking, can improve the algorithm's performance.

One can analogously use decompositions to obtain an automata-based algorithm running in time polynomial in $\ln (\mathcal{D})$ and $2^{\mathrm{wd}(\mathcal{D})}$, but we do not discuss the technical details any further for the sake of brevity.

### 4.4. Constructing $\epsilon$-Free Decompositions

We now present our algorithm for computing $\epsilon$-free decompositions that we outlined in Section 4.1. We first formalize the notion of a control, which encapsulates the parameters that guide our algorithm in the construction process. Please recall that $\mathbf{X}$ is the countably infinite set of all contexts that was fixed in Section 3.

Definition 27. A control is a tuple $\mathcal{C}=\langle\mathrm{mln}$, initialization, strategy, precedence $\rangle$ with the following components.

- mln is a positive integer.
- initialization is a polynomial-time computable function that takes a normalized $\mathcal{A L C I}$ ontology $\mathcal{O}$, a finite set of queries $\mathcal{Q}$, and a positive integer mln . The result of initialization $(\mathcal{O}, \mathcal{Q}, \mathrm{mln})$ is a pair $\langle\mathcal{D}, \vartheta\rangle$ where
- $\mathcal{D}$ is a context structure over the literals occurring in $\mathcal{O} \cup \mathcal{Q}$ such that $\mathcal{D}$ contains no edges, $\ln (\mathcal{D}) \leq m \ln$, and $\mathcal{D}$ contains a trivial context, and
$-\vartheta$ is a covering mapping from $\mathcal{Q}$ to $\mathcal{D}$.
- strategy is a polynomial-time computable function that takes an existential restriction $\exists R . A$, a set of atomic concepts $\mathbf{B}_{k}$, a decomposition $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, knw, poss, $\langle, \vartheta\rangle$ that contains a trivial context, and a positive integer min. The result of strategy $\left(\exists R . A, \mathbf{B}_{k}, \mathcal{D}, \mathrm{mln}\right)$ is a triple $\left\langle u\right.$, core $\left.^{\prime},<^{\prime}\right\rangle$ where
- core' is a subset of $\{A\} \cup \mathbf{B}_{k}$,
- $\prec^{\prime}$ is an $R$-admissible literal ordering, and
- either $u \in \mathbf{X} \backslash \mathcal{V}$ is a fresh context, or $u \in \mathcal{V}$ is a context in $\mathcal{D}$ such that $\operatorname{core}(u)=$ core'.

If $|\mathcal{V}| \geq \mathrm{mln}$, then strategy must return a context $u \in \mathcal{V}$ from $\mathcal{D}$; since $\mathcal{D}$ contains a trivial context, a context satisfying the required conditions is guaranteed to exist.

- precedence is a strict total order on the set of all pairs of the form $\langle\exists R . A, v\rangle$ with $\exists R . A \in \Sigma_{L}^{\exists}$ and $v \in \mathbf{X}$.

Intuitively, a control $\mathcal{C}$ specifies which contexts to introduce for a given ontology and a set of queries. In particular, function initialization specifies how to assign the queries to contexts; the function returns a context structure and a covering mapping that our algorithm will use as a starting point for the construction of an $\epsilon$-free decomposition: the algorithm will immediately extend the given context structure into an $\epsilon$-free decomposition by setting $\operatorname{poss}(v):=\operatorname{knw}(v):=\operatorname{core}(v)$ for each context $v$. Furthermore, function strategy is analogous to an expansion strategy and it specifies which contexts to introduce in order to satisfy existential and universal restrictions. Moreover, as we discussed in Section 4.1, the width and the length of a decomposition are not independent, and an exponential length may be needed to minimize the width. To overcome this problem, mln imposes an upper bound on decomposition length, and initialization and strategy are both required to honor this restriction. In particular, strategy is required to reuse a context whenever this bound has been exceeded; the decomposition being constructed will always contain a trivial context, so context reuse will always be possible. Finally, precedence determines a unique order in which fresh contexts are introduced, so that the resulting $\epsilon$-free decomposition is uniquely determined by $\mathcal{C}$, the ontology, and the set of queries.


Throughout the rest of section we fix a normalized $\mathcal{A L C I}$ ontology $\mathcal{O}$, a finite set of queries $\mathcal{Q}$, and a control $\mathcal{C}$; furthermore, we let $\mathbf{L}$ be the set of literals that occur in $\mathcal{O} \cup \mathcal{Q}$. We are now ready to present our decomposition construction algorithm.

Algorithm 28. The $\epsilon$-free decomposition construction algorithm takes a control $\mathcal{C}=\langle$ min, initialization, strategy, precedence $\rangle, \mathcal{O}$, and $\mathcal{Q}$, and it constructs an $\epsilon$-free decomposition $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, knw, poss, $\langle, \vartheta\rangle$ of $\mathcal{O}$ and $\mathcal{Q}$ as follows.

1. Let $\langle\mathcal{D}, \vartheta\rangle:=$ initialization $(\mathcal{O}, \mathcal{Q}, \mathrm{mln})$ and then extend $\mathcal{D}$ to the decomposition $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, knw, poss, $\langle, \vartheta\rangle$ by setting $\operatorname{poss}(v):=\operatorname{knw}(v):=\operatorname{core}(v)$ for each context $v \in \mathcal{V}$.
2. For each query $K \sqsubseteq M \in \mathcal{Q}$, let $v:=\vartheta(K \sqsubseteq M)$; then, set $\operatorname{poss}(v):=\operatorname{poss}(v) \cup K$.
3. Exhaustively apply all rules from Table 4 apart from rule $\mathbf{S}$.
4. Stop and return $\mathcal{D}$ if rule $\mathbf{S}$ is not applicable.
5. Let $\langle\exists R . A, v\rangle$ be the minimal pair w.r.t. precedence such that rule $\mathbf{S}$ is applicable to $\exists R . A$ and $v$; then, apply rule $\mathbf{S}$ to $\exists R . A$ and $v$, and proceed to step 3 .

A decomposition produced by this algorithm is called $a \mathcal{C}$-decomposition of $\mathcal{O}$ and $\mathcal{Q}$.
Algorithm 28 can be intuitively understood as follows. Step 1 uses function initialization to introduce the contexts needed to cover the queries in $\mathcal{Q}$, and then it initializes the sets of known and possible literals for each context to be equal to the context's core. Next, for each query $q=K \sqsubseteq M \in \mathcal{Q}$ and each literal $L \in K$, step 2 of Algorithm 28 adds $L$ to $\operatorname{poss}(\vartheta(q))$ in order to reflect the fact that step 3 of Algorithm 21 introduces a clause of the form $L \sqsubseteq L$ for each such $L$. Finally, in steps 3-5, the algorithm applies the rules shown in Table 4. Rules $\mathbf{H}_{k}$ and $\mathbf{H}_{p}$ correspond to the Hyper rule; rules $\mathbf{P}_{k}$ and $\mathbf{P}_{p}$ correspond to the Pred rule; and rule $\mathbf{S}$ corresponds to the Succ rule. Rules $\mathbf{H}_{k}$ and $\mathbf{P}_{k}$, however, are applied to the sets of known literals, and so they determine a lower bound on the known consequences of rules Hyper and Pred, respectively. In contrast, rules $\mathbf{H}_{p}$ and $\mathbf{P}_{p}$ are applied to the sets of possible literals, and so they determine an upper bound on the possible consequences of rules Hyper and Pred, respectively. Finally, rule $\mathbf{S}$ introduces all possible contexts that may be introduced by the Succ rule; the rule is applied with the lowest priority and in accordance with the precedence precedence from $\mathcal{C}$ in order to make the algorithm deterministic. In the rest of this section we prove the following theorem, which summarizes the formal properties of Algorithm 28.

Theorem 29. When applied to $\mathcal{O}, \mathcal{Q}$, and $\mathcal{C}$, Algorithm 28 terminates in time polynomial in mln and $\|\mathcal{O}\|+\|\mathcal{Q}\|$. The result is a unique, sound, admissible, and $\epsilon$-free decomposition $\mathcal{D}$ of $\mathcal{O}$ and $\mathcal{Q}$ satisfying $\ln (\mathcal{D}) \leq m \ln$.

Let $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, knw, poss, $\langle, \vartheta\rangle$ be a decomposition obtained by the decomposition construction algorithm. Clearly, $\mathcal{D}$ refers only to the literals occurring in $\mathcal{O} \cup \mathcal{Q}$, and it does not contain $\epsilon$-edges; hence, $\mathcal{D}$ is an $\epsilon$-free decomposition of $\mathcal{O}$ and $\mathcal{Q}$. We next prove the remaining claims of Theorem 29.

Claim 30. The decomposition construction algorithm terminates in time polynomial in $\mathrm{m} \ln$ and $\|\mathcal{O}\|+\|\mathcal{Q}\|$, and the resulting decomposition $\mathcal{D}$ satisfies $\ln (\mathcal{D}) \leq m \ln$.

Proof. Since initialization is computable in polynomial time, steps 1 and 2 can be performed in time polynomial in mln and $\|\mathcal{O}\|+\|\mathcal{Q}\|$. By Definition 27, step 1 introduces at most mln contexts due to the requirement on initialization, and rule $\mathbf{S}$ can introduce a new context in the call to strategy only if the total number of contexts is smaller than min; thus, $|\mathcal{V}| \leq$ mln clearly holds. Since $\mathcal{D}$ refers only to literals from $\mathbf{L}$, we have $|\mathcal{E}| \leq \operatorname{mln}^{2} \cdot|\mathbf{L}|$ and, for each context $v \in \mathcal{V}$, we have

$$
|\operatorname{core}(v)| \leq|\mathbf{L}|, \quad|\operatorname{knw}(v)| \leq|\mathbf{L}|, \quad \text { and } \quad|\operatorname{poss}(v)| \leq|\mathbf{L}| .
$$

The decomposition construction algorithm is monotonic-that is, it never removes elements from the decomposition being constructed. Also, the precondition of each rule contains a negation of the rule's postcondition; thus, whenever a rule is applicable, its application actually modifies the decomposition.

Now let $m_{\mathcal{D}}=\mathrm{m} \mathrm{m}^{2} \cdot|\mathbf{L}|+\mathrm{mln} \cdot 3 \cdot|\mathbf{L}|$. The decomposition construction algorithm can modify the decomposition at most $m_{\mathcal{D}}$ times: it can add each of the $\mathrm{mln}^{2} \cdot|\mathbf{L}|$ possible edges, and it can add $|\mathbf{L}|$ concepts to core $(v)$, $\operatorname{knw}(v)$, and poss $(v)$ for each context $v \in \mathcal{V}$. Thus, after at most $m_{\mathcal{D}}$ modifications, no rule can further modify the decomposition. Therefore, since each rule application modifies the decomposition, the decomposition construction algorithm terminates after at most $m_{\mathcal{D}}$ rule applications. Furthermore, due to the fact that each rule precondition refers to a bounded number of vertices, the preconditions of each rule instance can be checked in polynomial time. Since strategy is computable in polynomial time, each rule can be applied in polynomial time as well.

Thus, the algorithm requires a polynomial number of modifications of the decomposition, each of which requires polynomial time. Since $|\mathbf{L}|<\|\mathcal{O}\|+\|\mathcal{Q}\|$, the algorithm can be implemented so that it runs in time polynomial in min and $\|\mathcal{O}\|+\|\mathcal{Q}\|$.

## Claim 31. Decomposition $\mathcal{D}$ is unique.

Proof. In all rules apart from $\mathbf{S}$, the negative precondition merely checks whether the rule's postcondition is satisfied. Therefore, the rules are monotonic and they have a unique least fixpoint-that is, the order of rule applications is irrelevant. Furthermore, rule $\mathbf{S}$ is applied only if no other rule is applicable, and precedence ensures that the rule is applied in a deterministic way. Thus, the result of the decomposition construction algorithm is uniquely determined by $\mathcal{O}, \mathcal{Q}$, and $\mathcal{C}$.

## Claim 32. Decomposition $\mathcal{D}$ is sound.

Proof. The proof is by induction on rule application. After the initial step 1 , we have $\mathcal{E}=\emptyset$ and $\operatorname{core}(v)=k n w(v)$ for each context $v$, so $\mathcal{D}$ is sound. We next consider all rules from Table 4.
$\left(\mathbf{H}_{k}\right)$ If $K \subseteq \operatorname{knw}(v)$ and $K \sqsubseteq L \in \mathcal{O}$, then, by the induction hypothesis, we have $\mathcal{O} \vDash \operatorname{core}(v) \sqsubseteq K$; but then, we clearly have $\mathcal{O} \vDash \operatorname{core}(v) \sqsubseteq L$, so adding $L$ to $\mathrm{knw}(v)$ preserves soundness.
$\left(\mathbf{P}_{k}\right)$ If $\langle v, u, \exists R . A\rangle \in \mathcal{E}, \exists R . A \in \operatorname{knw}(v)$, and $\forall \operatorname{inv}(R) . C \in \operatorname{knw}(u)$, then, by the induction hypothesis, all of the following entailments hold:

$$
\mathcal{O} \vDash \operatorname{core}(v) \sqcap \exists R . A \sqsubseteq \exists R .[\operatorname{core}(u) \sqcap A] \quad \mathcal{O} \vDash \operatorname{core}(v) \sqsubseteq \exists R . A \quad \mathcal{O} \vDash \operatorname{core}(u) \sqsubseteq \forall \operatorname{inv}(R) . C
$$

These entailments imply $\mathcal{O} \vDash \operatorname{core}(v) \sqsubseteq \exists R$. $\forall \operatorname{inv}(R) . C$, which then implies $\mathcal{O} \vDash \operatorname{core}(v) \sqsubseteq C$. Consequently, adding $C$ to $\operatorname{knw}(v)$ preserves soundness.
(S) By the induction hypothesis, we have $\mathcal{O} \vDash \operatorname{core}(v) \sqsubseteq \forall R . B$ for each $B \in \mathbf{B}_{k}$, so $\mathcal{O} \vDash$ core $(v) \sqcap \exists R . A \sqsubseteq \exists R .\left[A \sqcap \mathbf{B}_{k}\right]$ holds; but then, $\mathcal{O} \vDash \operatorname{core}(v) \sqcap \exists R . A \sqsubseteq \exists R$. [core $(u) \sqcap A]$ holds since we have core $(u)=$ core $^{\prime} \subseteq\{A\} \cup \mathbf{B}_{k}$ by the requirement on strategy in Definition 27. Consequently, adding the edge $\langle v, u, \exists R . A\rangle$ to $\mathcal{E}$ preserves soundness.

Since no other rule affects $\mathcal{E}$ or knw, this concludes the proof of this claim.

## Claim 33. Decomposition $\mathcal{D}$ is admissible.

Proof. We check that $\mathcal{D}$ satisfies all conditions of Definition 20. The ontology condition and the structural conditions (S2) and (S3) hold trivially since the construction is closed under rules $\mathbf{H}_{p}, \mathbf{P}_{p}$, and $\mathbf{S}$, respectively.

For condition (S1), we show that core $(v) \subseteq \operatorname{knw}(v)$ and $k n w(v) \subseteq \operatorname{poss}(v)$ hold for each context $v \in \mathcal{V}$. The first inclusion holds because each set $\operatorname{knw}(v)$ is initialized in step 1 and in rule $\mathbf{S}$ with $\operatorname{knw}(v):=$ core $(v)$, and core $(v)$ remains constant during the construction while $\operatorname{knw}(v)$ only increases. The second inclusion holds because each set poss $(v)$ is initialized with $\operatorname{poss}(v):=\operatorname{knw}(v)$, and the rules $\mathbf{H}_{p}$ and $\mathbf{P}_{p}$ extending poss $(v)$ are applicable whenever the rules $\mathbf{H}_{k}$ and $\mathbf{P}_{k}$ extending $\operatorname{knw}(v)$ are applicable.

The initial context structure produced by initialization contains no edges, and whenever an edge $\langle v, u, \exists R . A\rangle$ is added to $\mathcal{E}$ in rule $\mathbf{S}$, strategy ensures that ${<_{u}}_{u}$ is $R$-admissible. Thus, the ordering condition holds throughout the construction.

Finally, for the covering condition, consider an arbitrary query $K \sqsubseteq M \in \mathcal{Q}$. After step 1 , context $v=\vartheta(K \sqsubseteq M)$ covers $K \sqsubseteq M$ in the sense of Definition 7-that is, core $(v) \subseteq K$ and $M$ is $<_{v}$-minimal. Then, step 2 ensures also that $K \subseteq$ poss $(v)$; thus, $v$ covers $K \sqsubseteq M$, as required by Definition 20 .


Figure 6: An $\epsilon$-Free Decomposition of $\mathcal{O}$ and $q$ and the Related Inferences

We are now ready to combine all of the results presented thus far and formulate what we believe to be the first result on fixed-parameter tractability of subsumption reasoning in description logics.

Theorem 34. For each control $\mathcal{C}$, the following problem is fixed-parameter tractable:

- Inputs: a normalized $\mathcal{A L C I}$ ontology $\mathcal{O}$ and a set of queries $\mathcal{Q}$
- Parameter: an integer mwd
- Problem: return "yes" if the $\mathcal{C}$-decomposition of $\mathcal{O}$ and $\mathcal{Q}$ has width at most mwd, and if also $\mathcal{O} \vDash K \sqsubseteq M$ holds for each query $K \sqsubseteq M \in \mathcal{Q}$

Proof. Immediate by Theorem 29, and Propositions 24 and 25.
Please note that Theorem 34 requires the input ontology to be normalized. As we mentioned in Section 2.1, many different normalizing transformations have been proposed in the literature, and they can produce syntactically very different ontologies. Thus, for an ontology that is not normalized, the width, the length, and our fixed-parameter tractability results implicitly depend on the normalization transformation. Investigating how normalization affects width and length, as well as what kind of transformation can minimize these two parameters is interesting, but we leave it for our future work.

## 5. Analyzing Or-Branching Using General Decompositions

We now generalize our framework to also allow for a quantitative analysis of or-branching. We first discuss the intuitions behind our ideas, and then we proceed with a formal presentation of our results.

### 5.1. Intuitions

Example 35. Let $\mathcal{O}$ be the ontology consisting of axioms (87)-(93), and let $q=A \sqsubseteq E$. One can readily check that $\mathcal{O} \vDash q$ holds.

$$
\begin{equation*}
A \sqsubseteq \exists R . B_{1} \tag{87}
\end{equation*}
$$

$$
\begin{align*}
B_{1} \sqcap B_{2} & \sqsubseteq C_{0} \sqcup D_{0}  \tag{89}\\
C_{i-1} & \sqsubseteq C_{i} \sqcup D_{i} \quad \text { for } 1 \leq i \leq n  \tag{88}\\
D_{i-1} & \sqsubseteq C_{i} \sqcup D_{i} \quad \text { for } 1 \leq i \leq n
\end{align*}
$$

$$
\begin{align*}
& C_{n} \sqsubseteq \forall R^{-} . E \\
& D_{n} \sqsubseteq \forall R^{-} . E \tag{93}
\end{align*}
$$

The $\mathcal{C}$-decomposition $\mathcal{D}_{1}$ of $\mathcal{O}$ and $q$ obtained using the eager strategy is shown in Figure 6a, and the inferences of Algorithm 21 on $\mathcal{D}_{1}$ are shown in Figure 6b. Clause (100) stands for clauses with $M \subseteq\left\{C_{0}, \ldots, C_{n}, D_{0}, \ldots, D_{n}, \forall R^{-} . E\right\}$; exactly which clauses of this form are derived depends on the literal ordering $<_{v_{2}}$. For example, with $\forall R^{-} . E<_{v_{2}} C_{n}<_{v_{2}} D_{n}<_{v_{2}} \ldots<_{v_{2}} C_{0}<_{v_{2}} D_{0}$ we derive a linear number of clauses, and with $\forall R^{-} . E<_{v_{2}} C_{0}<_{v_{2}} D_{0}<_{v_{2}} \ldots<_{v_{2}} C_{n}<_{v_{2}} D_{n}$ we derive an exponential number of clauses. In practice, it is difficult to identify a literal ordering that minimizes the number of derived clauses, so the only guarantee on the algorithm's performance is exponential in $n$; this is in agreement with the fact that $\operatorname{wd}\left(\mathcal{D}_{1}\right)=2 n+3$.

Clauses in $\mathcal{S}\left(v_{2}\right)$ essentially solve a propositional problem consisting of clauses (89)-(93) and clause $\alpha=B_{1} \sqcap B_{2} \sqsubseteq \forall R^{-} . E$ : the antecedent of $\alpha$ captures the information that can be propagated from $v_{1}$ to $v_{2}$, and the consequent of $\alpha$ captures the information that can be propagated from $v_{2}$ to $v_{1}$. As we explained in Section 2.4, the treewidth of propositional problems serves as a measure of the difficulty of such problems, and the bottom-up algorithm provides an FPT decision procedure for propositional satisfiability. Thus, to explain the difficulty of or-branching in description logics, we next incorporate tree decompositions into our framework. To this end, we replace context $v_{2}$ with a tree decomposition of the mentioned propositional problem, where we label the newly introduced edges by a special symbol $\epsilon$; similarly to edges in a tree decomposition, our $\epsilon$-edges will always be
symmetric. Decomposition $\mathcal{D}_{2}$ obtained in this way is shown in Figure 7 a. The set of contexts $\left\{v_{2}^{0}, \ldots, v_{2}^{n}\right\}$ obtained by replacing $\nu_{2}$ is called an $\epsilon$-component of $\mathcal{D}_{2}$. Intuitively, each $\epsilon$-component $\mathcal{W}$ of $\mathcal{D}_{2}$ corresponds to one or more domain elements in a model of $\mathcal{O}$, and each context $v \in \mathcal{W}$ provides a partial interpretation for the literals in poss(v). Condition ( N 2 ) of Definition 19 requires all contexts in an $\epsilon$-component $\mathcal{W}$ of a sound decomposition to have the same core; thus, all partial descriptions of the elements represented by $\mathcal{W}$ must coincide on the "assumed" literals. However, the contexts in $\mathcal{W}$ can differ on known and possible literals, which will allow us to reduce decomposition width.

We will extend the notion of decomposition admissibility so that it guarantees completeness of a modified consequence-based algorithm shown in Table 5 on page 28. The Hyper and Pred rules are as in Table 3, but modified so that a clause is added to some $\mathcal{S}(v)$ only if all of its literals are possible at context $v$. In addition, we introduce a new Epsilon rule, which essentially performs hyperresolution across $\epsilon$-edges: given an edge $\langle v, u, \epsilon\rangle$, a clause in $\prod_{i=1}^{n} L_{i} \sqsubseteq M \in \mathcal{S}(u)$ is chosen such that all literals in $M$ are possible at context $v$, each literal $L_{i}$ is resolved away using a clause in $\mathcal{S}(v)$, and the result is added to $\mathcal{S}(v)$. The Epsilon rule can be seen as a practical variant of the bottom-up algorithm for propositional satisfiability described in Section 2.4.

The inferences of our new algorithm on $\mathcal{O}, q$, and $\mathcal{D}_{2}$ are shown in Figure 7b. As before, we first initialize each $\mathcal{S}(v)$ with $\top \sqsubseteq L$ for each literal $L$ known at context $v$, thus obtaining clauses (103)-(113). Moreover, for each edge $\langle v, u, \epsilon\rangle$ and each literal $L$ that is possible at both $u$ and $v$, we initialize $\mathcal{S}(u)$ and $\mathcal{S}(v)$ with $L \sqsubseteq L$, thus obtaining clauses (114)-(125). Finally, using rules from Table 5 on page 28, we derive clauses (126)-(138). At each context $v_{2}^{i}$ with $0 \leq i<n$, inferences follow the same pattern: using clauses (90) and (91) from $\mathcal{O}$, we derive a clause of the form $\top \sqsubseteq C_{i+1} \sqcup D_{i+1}$, which the Epsilon rule then "transfers" over an $\epsilon$-edge into context $v_{2}^{i+1}$. Finally, in context $v_{2}^{n}$ we derive clause (137), and then by the Pred rule we derive the goal clause (138). In order to ensure that the Epsilon rule derives all relevant consequences, our notion of decomposition admissibility imposes an additional restriction on the literal orderings of contexts. For example, clause (128) is relevant for the Epsilon rule and is derived using the Hyper rule from clause (127); however, to facilitate the latter derivation, concepts $D_{1}$ and $C_{1}$ must not be larger in $<_{\nu_{2}^{0}}$ than concept $D_{0}$. Thus, the new restriction on literal orderings is analogous to $R$-admissibility for $\exists R$. $A$-edges, but adapted to $\epsilon$-edges. To estimate the complexity of our algorithm, we note that $\operatorname{wd}\left(\mathcal{D}_{2}\right)=4$; hence, the number of clauses derived at each context is bounded by $4^{\mathrm{wd}\left(\mathcal{D}_{2}\right)}=4^{4}$, and this bound does not depend on the literal ordering or the number $n$. In this way, $\operatorname{wd}\left(\mathcal{D}_{2}\right)$ provides us with a better estimate of the or-branching involved in deciding $\mathcal{O} \vDash q$.

As in Section 4, the definitions of soundness and admissibility for general decompositions do not directly provide us with an algorithm for computing decompositions. As a possible remedy, in Section 5.3 we introduce the notion of an $\epsilon$-refinement, which is obtained from an $\epsilon$-free decomposition by replacing each context $v$ with a tree decomposition of the propositional problem associated with $v$. In our example, $\mathcal{D}_{2}$ is an $\epsilon$-refinement of $\mathcal{D}_{1}$. We will show that computing an $\epsilon$-refinement involves determining a linear number of tree decompositions and is thus fixed-parameter tractable w.r.t. decomposition width.

It is again instructive to consider the case when $\mathcal{O}$ and $\mathcal{Q}$ contain only propositional clauses. Then, for each admissible decomposition $\mathcal{D}$ of $\mathcal{O}$ and $\mathcal{Q}$, the $\exists R$. $A$-edges are not needed and can be removed, but $\mathcal{D}$ can contain several $\epsilon$-components, each of which corresponds to a tree decomposition of a subset of $\mathcal{O}$ but with the known literals removed. Since the known literals are removed, the width of such $\mathcal{D}$ can be lower than the treewidth of $\mathcal{O}$, and so our decompositions generalize tree decompositions of propositional problems. Furthermore, only the Hyper and Epsilon are used when our consequence-based calculus is applied to such $\mathcal{D}$, and the calculus can be seen as a practical variant of the algorithm for propositional logic presented in Section 2.4.

### 5.2. Formalization

Decompositions and decomposition soundness (with or without $\epsilon$-edges) were already introduced in Section 4.2, so we next focus on admissibility of decompositions with $\epsilon$-edges. In the rest of this section we fix a normalized $\mathcal{A L C \mathcal { I }}$ ontology $\mathcal{O}$ and a finite set of queries $\mathcal{Q}$; furthermore, we let $\mathbf{L}$ be the set of literals that occur in $\mathcal{O} \cup \mathcal{Q}$.

Definition 36. Let $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, knw, poss, $\langle, \vartheta\rangle$ be a decomposition of $\mathcal{O}$ and $\mathcal{Q}$, and let $\mathcal{W}$ be an arbitrary set such that $\mathcal{W} \subseteq \mathcal{V}$. Then, let $\operatorname{poss}(\mathcal{W}):=\bigcup_{w \in \mathcal{W}} \operatorname{poss}(w)$, and let $\operatorname{knw}(\mathcal{W})$ and $\operatorname{core}(\mathcal{W})$ be defined analogously. The $\epsilon$-projection of $\mathcal{D}$ w.r.t. $\mathcal{W}$, written $\mathcal{D}_{\mathcal{W}}$, is the graph whose set of vertices is $\mathcal{W}$ and that contains the undirected edge $\{u, v\}$ for each $\langle u, v, \epsilon\rangle \in \mathcal{E}$ with $u, v \in \mathcal{W}$. Set $\mathcal{W}$ is $\epsilon$-connected if $\mathcal{D}_{\mathcal{W}}$ is connected; furthermore, $\mathcal{W}$ is an $\epsilon$-component of $\mathcal{D}$ if $\mathcal{W}$ is $\epsilon$-connected, and no $\epsilon$-connected set of vertices $\mathcal{W}^{\prime}$ exists such that $\mathcal{W} \subsetneq \mathcal{W}^{\prime} \subseteq \mathcal{V}$.

Let $v \in \mathcal{V}$ be an arbitrary context of $\mathcal{D}$, let $\mathcal{W}$ be the $\epsilon$-component of $\mathcal{D}$ such that $v \in \mathcal{W}$, and let $K \sqsubseteq M$ be an arbitrary query; then, $v$ is sound for $K \sqsubseteq M$ if core $(v) \subseteq K$; $v$ is complete for $K \sqsubseteq M$ if $K \subseteq \operatorname{poss}(v), M \cap \operatorname{poss}(\mathcal{W}) \subseteq \operatorname{poss}(v)$, and $M$ is $<_{v}$-minimal; and $v$ covers $K \sqsubseteq M$ if $v$ is both sound and complete for $K \sqsubseteq M$. Finally, $\mathcal{D}$ is admissible if all of the following conditions are satisfied.

- Epsilon conditions:
(E1) For each $\langle w, u, \epsilon\rangle \in \mathcal{E}$, we have $\langle u, w, \epsilon\rangle \in \mathcal{E}$.
(E2) For each $\epsilon$-component $\mathcal{W}$ of $\mathcal{D}$, graph $\mathcal{D}_{\mathcal{W}}$ is an undirected tree.
(E3) For each $\epsilon$-component $\mathcal{W}$ of $\mathcal{D}$ and each literal $L \in \operatorname{poss}(\mathcal{W})$, set $\{w \in \mathcal{W} \mid L \in \operatorname{poss}(w)\}$ is $\epsilon$-connected.
- Structural conditions:

| $\underline{B_{1}, B_{2}}$ |  |
| :--- | :--- |
|  | $\overline{C_{2}, D_{2}}$ |
| $C_{1}, D_{1}$ | $\underline{v_{1}, B_{2}}$ |
|  |  |
| $C_{1}, D_{1}$, |  |
| $C_{0}, D_{0}$ |  |




| $\stackrel{\text { ® }}{\text { D }}$ | $\begin{array}{ll} \text { If } \quad & \prod_{i=1}^{n} A_{i} \sqsubseteq M \in \mathcal{O} \text { with } M \subseteq \operatorname{poss}(v) \\ & K_{i} \sqsubseteq M_{i} \sqcup A_{i} \in \mathcal{S}(v) \text { with } A_{i} \nprec_{v} M_{i} \\ & \text { for } 1 \leq i \leq n, \\ & \text { and } \prod_{i=1}^{n} K_{i} \sqsubseteq M \sqcup \bigsqcup_{i=1}^{n} M_{i} \ddot{\notin \mathcal{S}(v),} \\ \text { then } \text { add } \prod_{i=1}^{n} K_{i} \sqsubseteq M \sqcup \bigsqcup_{i=1}^{n} M_{i} \text { to } \mathcal{S}(v) . \end{array}$ |
| :---: | :---: |
| 흔 | $\begin{array}{ll} \text { If } & \langle v, u, \exists R . A\rangle \in \mathcal{E}, \\ & A \sqcap \prod_{i=1}^{n} B_{i} \sqsubseteq \bigsqcup_{j=1}^{m} \forall \operatorname{inv}(R) \cdot C_{j} \in \mathcal{S}(u) \quad \text { or } \quad \prod_{i=1}^{n} B_{i} \sqsubseteq \bigsqcup_{j=1}^{m} \forall \operatorname{inv}(R) \cdot C_{j} \in \mathcal{S}(u) \\ & \text { with } C_{j} \in \operatorname{poss}(v) \text { for each } 1 \leq j \leq m, \\ & K_{0} \sqsubseteq M_{0} \sqcup \exists R \cdot A \in \mathcal{S}(v) \text { with } \exists R . A \nprec_{v} M_{0}, \\ & K_{i} \sqsubseteq M_{i} \sqcup \forall R . B_{i} \in \mathcal{S}(v) \text { with } \forall R . B_{i} \nprec_{v} M_{i} \text { for } 1 \leq i \leq n, \\ & \text { and } \prod_{i=0}^{n} K_{i} \sqsubseteq \bigsqcup_{i=0}^{n} M_{i} \sqcup \bigsqcup_{j=1}^{m} C_{j} \nexists \mathcal{S}(v), \\ \text { then } \text { add } \prod_{i=0}^{n} K_{i} \sqsubseteq \bigsqcup_{i=0}^{n} M_{i} \sqcup \bigsqcup_{j=1}^{m} C_{j} \text { to } \mathcal{S}(v) . \end{array}$ |
|  | $\begin{array}{ll} \text { If } & \langle v, u, \epsilon\rangle \in \mathcal{E}, \\ & \prod_{i=1}^{n} L_{i} \sqsubseteq M \in \mathcal{S}(u) \text { with } M \subseteq \operatorname{poss}(v), \\ \quad K_{i} \sqsubseteq M_{i} \sqcup L_{i} \in \mathcal{S}(v) \text { with } \\ \quad & \text { and } \prod_{i=1}^{n} K_{i} \sqsubseteq M M_{i} \text { for } 1 \leq i \leq n, \\ \text { then } \quad \text { add } \prod_{i=1}^{n} K_{i} \sqsubseteq M \sqcup M_{i=1}^{n} M_{i} \nsubseteq \mathcal{S}(v), \\ M_{i} \text { to } \mathcal{S}(v) . \end{array}$ |

(S1) For each context $v \in \mathcal{V}$, we have $\operatorname{core}(v) \subseteq \operatorname{knw}(v) \subseteq \operatorname{poss}(v)$.
(S2) For each $\epsilon$-component $\mathcal{U}$ of $\mathcal{D}$, each edge $\langle w, u, \exists R . A\rangle \in \mathcal{E}$ with $u \in \mathcal{U}$, and each concept $\forall \operatorname{inv}(R) . C \in \operatorname{poss}(\mathcal{U})$, we have $C \in \operatorname{poss}(w)$ and $\forall \operatorname{inv}(R) . C \in \operatorname{poss}(u)$.
(S3) For each $\epsilon$-component $\mathcal{W}$ of $\mathcal{D}$ and each $\exists R . A \in \operatorname{poss}(\mathcal{W})$, some edge $\langle w, u, \exists R . A\rangle \in \mathcal{E}$ with $w \in \mathcal{W}$ exists such that
$-\operatorname{poss}(w) \supseteq\{\exists R . A\} \cup\{\forall R . B \mid \forall R . B \in \operatorname{poss}(\mathcal{W})\}$, and
$-\operatorname{poss}(u) \supseteq\{A\} \cup\{B \mid \forall R . B \in \operatorname{poss}(\mathcal{W})\}$.

- Ordering conditions:
(P1) For each $\exists R$.A-edge $\langle w, u, \exists R . A\rangle \in \mathcal{E}$, literal ordering $<_{u}$ is $R$-admissible.
(P2) For each $\epsilon$-edge $\langle w, u, \epsilon\rangle \in \mathcal{E}$, set $\operatorname{poss}(w) \cap \operatorname{poss}(u)$ is $<_{u}$-minimal.
- Ontology condition: For each $\epsilon$-component $\mathcal{W}$ of $\mathcal{D}$ and each clause $K \sqsubseteq M \in \mathcal{O}$ with $K \subseteq \operatorname{poss}(\mathcal{W})$, a context $w \in \mathcal{W}$ exists such that $K \cup M \subseteq \operatorname{poss}(w)$.
- Covering condition: Each query $q \in \mathcal{Q}$ is covered in the context $\vartheta(q)$.

Definition 36 can be intuitively understood as follows. The set of contexts in $\mathcal{D}$ can be partitioned into $\epsilon$-componentsmaximal sets of contexts connected by $\epsilon$-edges. Each $\epsilon$-component can be understood as a tree decomposition of a propositional problem, so conditions (E1)-(E3) require the contexts in the $\epsilon$-component to form an undirected tree satisfying the connectedness condition (T2) of tree decompositions. Structural conditions (S1)-(S3) extend the respective conditions from Definition 20. To understand condition (S2), let $\langle w, u, \exists R . A\rangle$ be an arbitrary edge of $\mathcal{D}$, let $\mathcal{U}$ be the $\epsilon$-component of $\mathcal{D}$ that $u$ belongs to, and assume that $\forall \operatorname{inv}(R) . C$ is possible at some context $u^{\prime} \in \mathcal{U}$ different from $u$. Since $u^{\prime}$ and $u$ describe the same elements in a model of $\mathcal{O}$, the elements represented by context $v$ might need to satisfy $C$, which is taken care of by the Pred rule. To ensure that the rule can be applied to edge $\langle w, u, \exists R . A\rangle$, condition (S2) requires $\forall \operatorname{inv}(R) . C$ to be possible at $u$, and $C$ to be possible at $v$ (just like in Definition 20). To understand condition (S3), let $\mathcal{W}$ be an arbitrary $\epsilon$-component of $\mathcal{D}$, assume that $\exists R . A$ is possible at some context $w^{\prime} \in \mathcal{W}$, and assume that $\mathcal{D}$ contains an edge $\langle w, u, \exists R . A\rangle$ with $w$ different from $w^{\prime}$. Contexts $w$ and $w^{\prime}$ jointly describe elements of a model of $\mathcal{O}$, so the existential restriction $\exists R . A$ also "applies" to $w$ and the mentioned edge as well; thus, all literals $\forall R . B$ possible at $\mathcal{W}$ must be possible at $w$, and all corresponding concepts $B$ must be possible at $u$ in order to satisfy the semantics of $\exists R$. $A$-edges. The ontology condition is modified in similar vein; in particular, assume that $\mathcal{D}$ contains an $\epsilon$-component $\mathcal{W}$ such that each antecedent literal of a clause $K \sqsubseteq M \in \mathcal{O}$ is possible at some context in $\mathcal{W}$; then, the domain elements represented by $\mathcal{W}$ may satisfy $K$; but then, we must find a context $w \in \mathcal{W}$ in which we can apply the Hyper rule to $K \sqsubseteq M$. The notion of covering is extended analogously. Finally, the ordering condition restricts the literal ordering as we discussed in Section 5.1. Please note that the notions of admissibility introduced in Definitions 20 and 36 coincide on $\epsilon$-free decompositions.

To generalize our consequence-based algorithm as described in Section 5.1, we modify the Hyper and the Pred rules as shown in Table 5 (the modifications are underlined), and we add a new Epsilon rule. Theorem 37 provides us with the completeness claim for the new algorithm; the proof is technically involved, so we defer it to Appendix B. In addition, Proposition 38 provides us with a matching soundness claim.

Theorem 37 (Completeness). Let $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, knw, poss, $\langle, \vartheta\rangle$ be an admissible decomposition of $\mathcal{O}$ and $\mathcal{Q}$, and let $\mathcal{S}$ be a clause system for $\mathcal{D}$ such that
(II) no inference rule in Table 5 is applicable to $\mathcal{S}$,
(I2) $\mathrm{T} \sqsubseteq L \hat{\in} \mathcal{S}(v)$ for each context $v \in \mathcal{V}$ and each literal $L \in \operatorname{knw}(v)$,
(I3) $B \sqsubseteq B \hat{\in} \mathcal{S}(u)$ for each edge $\langle v, u, \exists R . A\rangle \in \mathcal{E}$ and each concept $B \in \operatorname{poss}(u)$ such that either $B=A$ or $\forall R . B \in \operatorname{poss}(v)$, and
(I4) $L \sqsubseteq L \hat{\in} \mathcal{S}(u)$ for each edge $\langle v, u, \epsilon\rangle \in \mathcal{E}$ and each literal $L \in \operatorname{poss}(v) \cap \operatorname{poss}(u)$.
Then, $K \sqsubseteq M \hat{\in} \mathcal{S}(v)$ holds for each query $K \sqsubseteq M$ and each context $v \in \mathcal{V}$ that satisfy all of the following three conditions:

- $\mathcal{O} \vDash K \sqsubseteq M$,
- context $v$ is complete for query $K \sqsubseteq M$, and
- $K \sqsubseteq L \hat{\epsilon} \mathcal{S}(v)$ for each literal $L \in K$.

Proposition 38 (Soundness). Let $\mathcal{D}=\left\langle\mathcal{V}, \mathcal{E}\right.$, core, $\langle \rangle$ be a sound decomposition of $\mathcal{O}$ and $\mathcal{Q}$, and let $\mathcal{S}_{1}$ be a clause system for $\mathcal{D}$ sound for $\mathcal{O}$. Then, each clause system $\mathcal{S}_{2}$ obtained by applying an inference rule from Table 5 to $\mathcal{D}$ and $\mathcal{S}_{1}$ is sound for $\mathcal{O}$.
Proof. Rules Hyper and Pred in Table 5 are more restrictive than the corresponding rules in Table 3, so the proof of Proposition 9 applies. Assume that the Epsilon rule is applied as shown in Table 5, so we show that $\mathcal{O} \vDash \operatorname{core}(v) \sqcap\rceil_{i=1}^{n} K_{i} \sqsubseteq M \sqcup \bigsqcup_{i=1}^{n} M_{i}$. To this end, let $\mathcal{I}=\left\langle\Delta^{\mathcal{I}},,^{\mathcal{I}}\right\rangle$ be an arbitrary model of $\mathcal{O}$ and consider an arbitrary element $\delta \in \Delta^{\mathcal{I}}$ such that $\delta \in\left[\operatorname{core}(v) \sqcap \prod_{i=1}^{n} K_{i}\right]^{\mathcal{I}}$. Now $\mathcal{S}_{1}$ is sound for $\mathcal{O}$, so, for each $1 \leq i \leq n$, we have $\mathcal{I} \vDash \operatorname{core}(v) \sqcap K_{i} \sqsubseteq M_{i} \sqcup L_{i}$, which together with our assumption implies $\delta \in\left(M_{i} \sqcup L_{i}\right)^{\mathcal{I}}$. If $\delta \in M_{i}^{\mathcal{I}}$ for some $1 \leq i \leq n$, then $\delta \in\left(\bigsqcup_{i=1}^{n} M_{i}\right)^{\mathcal{I}}$. Otherwise, we have $\delta \in L_{i}^{\mathcal{I}}$ for each $1 \leq i \leq n$; but then, $\mathcal{I} \vDash \operatorname{core}(u) \sqcap \prod_{i=1}^{n} L_{i} \sqsubseteq M$ holds by the induction assumption and the fact that $\mathcal{S}_{1}$ is sound for $\mathcal{O}$, and core $(v)=\operatorname{core}(u)$ holds by the soundness condition (N2); together, these observations imply $\delta \in M^{\mathcal{I}}$. Either way, $\delta \in\left(M \sqcup \bigsqcup_{i=1}^{n} M_{i}\right)^{\mathcal{I}}$ holds. Since $\delta$ was chosen arbitrarily, we have $\mathcal{I} \vDash \operatorname{core}(v) \sqcap \prod_{i=1}^{n} K_{i} \sqsubseteq M \sqcup \bigsqcup_{i=1}^{n} M_{i}$, as required.

We next extend the consequence-based algorithm for $\epsilon$-free decompositions (i.e., Algorithm 21) so that it can be applied to decompositions with $\epsilon$-edges. Apart from applying the modified rules from Table 5, the only other difference is in step 3, which is needed for the Epsilon rule. The Elim rule and redundancy elimination are the same as in Section 3.3.

Algorithm 39. The decomposition-based reasoning algorithm for $\mathcal{A L C I}$ takes $\mathcal{O}, \mathcal{Q}$, and a sound and admissible decomposition $\mathcal{D}$ of $\mathcal{O}$ and $\mathcal{Q}$. The algorithm (nondeterministically) computes a clause system $\mathcal{S}$ for $\mathcal{D}$ as follows.

1. Set $\mathcal{S}(v):=\{\top \sqsubseteq L \mid L \in \operatorname{knw}(v)\}$ for each context $v \in \mathcal{V}$.
2. For each edge $\langle v, u, \exists R . A\rangle \in \mathcal{E}$ and each concept $B \in \operatorname{uknw}(u)$ such that $B=A$ or $\forall R . B \in \operatorname{poss}(v)$, add $B \sqsubseteq B$ to $\mathcal{S}(u)$.
3. For each $\epsilon$-edge $\langle v, u, \epsilon\rangle \in \mathcal{E}$ and each literal $L \in \operatorname{poss}(v) \cap u k n w(u)$, add $L \sqsubseteq L$ to $\mathcal{S}(u)$.
4. For each query $K \sqsubseteq M \in \mathcal{Q}$, let $v:=\vartheta(K \sqsubseteq M)$; then, for each literal $L \in K \backslash \operatorname{knw}(v)$, add $L \sqsubseteq L$ to $\mathcal{S}(v)$.
5. Exhaustively apply rules Hyper, Pred, and Epsilon from Table 5 and rule Elim from Table 3.

Lemma 12 (showing that if $K \sqsubseteq M \hat{\in} \mathcal{S}(v)$ holds at some point during algorithm's execution, then this also holds at all future points) clearly applies to Algorithm 39 as well: the Epsilon rule only adds clauses, and the Elim rule is the same as in Section 3.2. Furthermore, the proof of Lemma 22 (showing that each clause in $\mathcal{S}(v)$ is of the form $\top \sqsubseteq L$ with $L \in \operatorname{knw}(v)$ or of the form $K \sqsubseteq M$ with $K \cup M \subseteq u k n w(v)$ ) can be straightforwardly extended to Algorithm 39. This allows us to establish in Proposition 40 the soundness and completeness of our algorithm, and then determine in Proposition 41 the algorithm's complexity.
Proposition 40. Let $\mathcal{S}$ be a clause system obtained by applying Algorithm 39 to $\mathcal{O}, \mathcal{Q}$, and $\mathcal{D}$. Then, for each query $q \in \mathcal{Q}$, we have $\mathcal{O} \vDash q$ if and only if $q \hat{\in} \mathcal{S}(\vartheta(q))$.
Proof. By Lemma 12, steps 2-4 ensure conditions (I2)-(I4) and step 5 ensures the required condition on the queries in Theorem 37; then, the claim of this proposition follows from Theorem 37 and Proposition 38 as in the proof of Proposition 24.
Proposition 41 (Termination). Algorithm 39 terminates in time polynomial in $4^{\mathrm{wd}(\mathcal{D})^{2}}, \ln (\mathcal{D})$, and $\|\mathcal{O}\|+\|\mathcal{Q}\|$.
Proof. Analogous to the proof of Proposition 25.

### 5.3. Constructing Decompositions via $\epsilon$-Refinement

We next formalize the notion of $\epsilon$-refinement of an $\epsilon$-free decomposition $\mathcal{D}$ that we described in Section 5.1.
Definition 42. Let $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, knw, poss, $\langle, \vartheta\rangle$ be an $\epsilon$-free decomposition of $\mathcal{O}$ and $\mathcal{Q}$. For each edge $\langle v, u, \exists R . A\rangle \in \mathcal{E}$, sets $\delta_{v, u}^{\exists R . A}$ and $\varrho_{v, u}^{\exists R . A}$ are defined as follows:

$$
\begin{aligned}
& \delta_{v, u}^{\exists R . A}:=\{\exists R . A\} \cup\{\forall R . B \mid \forall R . B \in \operatorname{poss}(v)\} \cup \quad\{C \mid \forall \operatorname{inv}(R) . C \in \operatorname{poss}(u)\} \\
& \varrho_{v, u}^{\exists R . A}:=\quad\{A\} \cup \quad\{B \mid \forall R . B \in \operatorname{poss}(v)\} \cup\{\forall \operatorname{inv}(R) . C \mid \forall \operatorname{inv}(R) . C \in \operatorname{poss}(u)\}
\end{aligned}
$$

Function $\mathcal{H}$ for $\mathcal{D}$ maps each context $v \in \mathcal{V}$ to an $\mathbf{L}$-hypergraph $\mathcal{H}_{v}$ such that each $\mathcal{H}_{v}$ is the smallest set satisfying all of the following conditions:
(R1) $(K \cup M) \cap \operatorname{uknw}(v) \in \mathcal{H}_{v}$ for each context $v \in \mathcal{V}$ and each clause $K \sqsubseteq M \in \mathcal{O}$ such that $K \cup M \subseteq \operatorname{poss}(v)$,
(R2) $\delta_{v, u}^{\exists R . A} \cap \mathrm{uknw}(v) \in \mathcal{H}_{v}$ and $\varrho_{v, u}^{\exists R . A} \cap \mathrm{uknw}(u) \in \mathcal{H}_{u}$ for each edge $\langle v, u, \exists R . A\rangle \in \mathcal{E}$, and
(R3) $(K \cup M) \cap \operatorname{uknw}(v) \in \mathcal{H}_{v}$ for each query $q=K \sqsubseteq M \in \mathcal{Q}$ and context $v=\vartheta(q)$.
A decomposition $\mathcal{D}^{\prime}=\left\langle\mathcal{V}^{\prime}, \mathcal{E}^{\prime}\right.$, core' ${ }^{\prime} \mathrm{knw}^{\prime}$, poss $^{\prime},\left\langle^{\prime}, \vartheta^{\prime}\right\rangle$ of $\mathcal{O}$ and $\mathcal{Q}$ is an $\epsilon$-refinement of $\mathcal{D}$ if there exists a function $\mathcal{T}$ mapping each context $v \in \mathcal{V}$ to a tree decomposition $\mathcal{T}_{v}=\left\langle\mathcal{V}_{v}, \mathcal{E}_{v}, \mathcal{L}_{v}\right\rangle$ of $\mathcal{H}_{v}$ such that $\mathcal{V}_{v} \cap \mathcal{V}_{w}=\emptyset$ and $\mathcal{V}_{v} \subseteq \mathbf{X}$ for all $v, w \in \mathcal{V}$ with $v \neq w$, and all of the following conditions are satisfied.
(R4) $\mathcal{V}^{\prime}=\bigcup_{v \in \mathcal{V}} \mathcal{V}_{v}$.
(R5) For each context $v \in \mathcal{V}$ and each vertex $w \in \mathcal{V}_{v}$, we have

$$
\operatorname{core}^{\prime}(w)=\operatorname{core}(v), \quad \operatorname{knw}^{\prime}(w)=\operatorname{knw}(v), \quad \text { and } \quad \operatorname{poss}^{\prime}(w)=\operatorname{knw}(v) \cup \mathcal{L}_{v}(w)
$$

(R6) For all $v, u \in \mathcal{V}$, each $w \in \mathcal{V}_{v}$, and each $z \in \mathcal{V}_{u}$, we have $\langle w, z, \epsilon\rangle \in \mathcal{E}^{\prime}$ if and only if $v=u$ and $\{w, z\} \in \mathcal{E}_{v}$.
(R7) For all $v, u \in \mathcal{V}$, each $w \in \mathcal{V}_{v}$, each $z \in \mathcal{V}_{u}$, and each literal $\exists R . A \in \mathbf{L}$, we have $\langle w, z, \exists R . A\rangle \in \mathcal{E}^{\prime}$ if and only if

$$
\langle v, u, \exists R . A\rangle \in \mathcal{E}, \quad \delta_{v, u}^{\exists R . A} \cap \mathrm{uknw}(v) \subseteq \mathcal{L}_{v}(w), \quad \text { and } \quad \varrho_{v, u}^{\exists R . A} \cap \mathrm{uknw}(u) \subseteq \mathcal{L}_{u}(z) .
$$

(R8) For each context $v \in \mathcal{V}$, each vertex $w \in \mathcal{V}_{v}$, and all literals $L_{1}, L_{2} \in \mathbf{L}$, we have $L_{1} \prec_{w}^{\prime} L_{2}$ if and only if $L_{1}<_{v} L_{2}$ and no $\epsilon$-edge $\langle u, w, \epsilon\rangle \in \mathcal{E}^{\prime}$ exists such that $L_{2} \in \operatorname{poss}^{\prime}(u) \cap \operatorname{poss}^{\prime}(w)$.
(R9) For each query $q=K \sqsubseteq M \in \mathcal{Q}$, context $v=\vartheta(q)$, and context $w=\vartheta^{\prime}(q)$, we have $w \in \mathcal{V}_{v}$ and $(K \cup M) \cap \operatorname{uknw}(v) \subseteq \mathcal{L}_{v}(w)$.
Definition 42 can be intuitively understood as follows. Each hypergraph $\mathcal{H}_{v}$ captures the structure of a propositional problem that must be solved at context $v$ : condition (R1) ensures that, for each context $v$ of $\mathcal{D}$, hypergraph $\mathcal{H}_{v}$ contains the "unknown part" of each clause $K \sqsubseteq M \in \mathcal{O}$ for which all literals are possible at $v$; furthermore, condition (R2) ensures that, for each edge $\langle v, u, \exists R . A\rangle$ of $\mathcal{D}$, hypergraphs $\mathcal{H}_{v}$ and $\mathcal{H}_{u}$ contain all unknown literals that might be needed to apply the Pred rule; and finally, condition (R3) ensures that, for each query $q \in \mathcal{Q}$ and $v=\vartheta(q)$, hypergraph $\mathcal{H}_{v}$ contains the "unknown part" of $q$. Conditions (R4)-(R9) then essentially capture the intuition that each context $v$ of $\mathcal{D}$ is replaced by a tree decomposition $\mathcal{T}_{v}$ of $\mathcal{H}_{v}$. Please note that, for each edge $\langle v, u, \exists R . A\rangle$ of $\mathcal{D}$, by condition (R2) and property (T1) of tree decompositions, there exists at least one pair of $w$ and $z$ satisfying condition (R7). Furthermore, for each query $q \in \mathcal{Q}$ and $v=\vartheta(q)$, context $\vartheta^{\prime}(q)$ can be an arbitrary vertex $w \in \mathcal{V}_{v}$ satisfying $(K \cup M) \cap u k n w(v) \subseteq \mathcal{L}_{v}(w)$; due to condition (R3) and property (T1) of tree decompositions, such $w$ is guaranteed to exist. Finally, please note that decomposition $\mathcal{D}$ can admit many $\epsilon$-refinements; however, by the following theorem, if $\mathcal{D}$ is sound and admissible, then each $\epsilon$-refinement of $\mathcal{D}$ is sound and admissible too.

Theorem 43. For each sound, admissible, and $\epsilon$-free decomposition $\mathcal{D}$ of $\mathcal{O}$ and $\mathcal{Q}$, each $\epsilon$-refinement $\mathcal{D}^{\prime}$ of $\mathcal{D}$ is also a sound and admissible decomposition of $\mathcal{O}$ and $\mathcal{Q}$.

Proof. We first show that $\mathcal{D}^{\prime}$ is sound. Consider arbitrary contexts $w, z \in \mathcal{V}$, and let $v, u \in \mathcal{V}$ be such that $w \in \mathcal{V}_{v}$ and $z \in \mathcal{V}_{u}$. By condition (R5), we have core' $(w)=\operatorname{core}(v), \operatorname{knw}^{\prime}(w)=\operatorname{knw}(v)$, $\operatorname{core}^{\prime}(z)=\operatorname{core}(u)$, and $\operatorname{knw}^{\prime}(z)=\operatorname{knw}(u)$.

Consider an arbitrary $\exists R . A$-edge $\langle w, z, \exists R . A\rangle \in \mathcal{E}^{\prime}$. By condition (R7), then $\mathcal{D}$ contains the edge $\langle v, u, \exists R . A\rangle \in \mathcal{E}$. Since $\mathcal{D}$ is sound, $\mathcal{O} \vDash \operatorname{core}(v) \sqcap \exists R . A \sqsubseteq \exists R$.[core $(u) \sqcap A]$ holds; but then, $\mathcal{O} \vDash \operatorname{core}^{\prime}(w) \sqcap \exists R . A \sqsubseteq \exists R$.[core' $\left.(z) \sqcap A\right]$ holds as well.

Consider an arbitrary $\epsilon$-edge $\langle w, z, \epsilon\rangle \in \mathcal{E}^{\prime}$. By condition (R6), we have $v=u$, which clearly implies $\operatorname{core}^{\prime}(w)=\operatorname{core}^{\prime}(z)$, as required.

Finally, since $\mathcal{D}$ is sound, we have $\mathcal{O} \vDash \operatorname{core}(v) \sqsubseteq A$ for each $A \in \operatorname{knw}(v)$; but then, $\mathcal{O} \vDash \operatorname{core}^{\prime}(w) \sqsubseteq A$ for each $A \in \operatorname{knw}^{\prime}(w)$ clearly holds as well.

We next show that $\mathcal{D}^{\prime}$ is admissible. Clearly, $\mathcal{D}^{\prime}$ refers only to the literals occurring in $\mathcal{O} \cup \mathcal{Q}$, so $\mathcal{D}^{\prime}$ is a decomposition of $\mathcal{O}$ and $\mathcal{Q}$. Furthermore, for each context $v \in \mathcal{V}$ the following conditions are satisfied.

$$
\begin{align*}
& \operatorname{knw}^{\prime}\left(\mathcal{V}_{v}\right)=\bigcup_{w \in \mathcal{V}_{v}} \operatorname{knw}^{\prime}(w)=\operatorname{knw}(v)  \tag{139}\\
& \operatorname{poss}^{\prime}\left(\mathcal{V}_{v}\right)=\bigcup_{w \in \mathcal{V}_{v}} \operatorname{poss}^{\prime}(w) \subseteq \operatorname{poss}(v) \tag{140}
\end{align*}
$$

Finally, for each $\epsilon$-component $\mathcal{W}$ of $\mathcal{D}^{\prime}$, a context $v \in \mathcal{V}$ exists such that $\mathcal{W}=\mathcal{V}_{v}$; conversely, for each context $v \in \mathcal{V}$, set $\mathcal{V}_{v}$ is an $\epsilon$-component of $\mathcal{D}$. We next check that $\mathcal{D}^{\prime}$ satisfies all the admissibility conditions listed in Definition 36.

## Epsilon conditions:

(Condition E1) This property follows immediately from (R6) and the fact that the edges in $\mathcal{E}_{v}$ are undirected.
(Condition E2) By (R6), for each $\epsilon$-component $\mathcal{V}_{v}$ of $\mathcal{D}^{\prime}$, we have that $\mathcal{D}_{\mathcal{V}_{v}}^{\prime}$ is equal to the graph $\left\langle\mathcal{V}_{v}, \mathcal{E}_{v}\right\rangle$, and the latter, being a tree decomposition, is an undirected tree.
(Condition E3) Consider an arbitrary $\epsilon$-component $\mathcal{V}_{v}$ of $\mathcal{D}^{\prime}$ and literal $L \in \operatorname{poss}^{\prime}\left(\mathcal{V}_{v}\right)$, and let $\Gamma:=\left\{w \in \mathcal{V}_{v} \mid L \in \operatorname{poss}^{\prime}(w)\right\}$. We have $L \in \operatorname{poss}(v)$ due to (140), so either $L \in \operatorname{knw}(v)$ or $L \in \operatorname{uknw}(v)$. In the former case, we have $L \in \operatorname{posss}^{\prime}(w)$ for each $w \in \mathcal{V}_{v}$, so $\Gamma=\mathcal{V}_{v}$ and $\Gamma$ is clearly $\epsilon$-connected. In the latter case, we have $\Gamma=\left\{w \in \mathcal{V}_{v} \mid L \in \mathcal{L}_{v}(w)\right\}$, so $\Gamma$ is $\epsilon$-connected by property (T2) of tree decompositions.

## Structural conditions:

(Condition S1) By (R5), for each $v \in \mathcal{V}$ and each $w \in \mathcal{V}$, we have $\operatorname{core}^{\prime}(w)=\operatorname{core}(v)$ and $\operatorname{knw}^{\prime}(w)=\operatorname{knw}(v) \subseteq \operatorname{poss}^{\prime}(w)$. Furthermore, core $(v) \subseteq \operatorname{knw}(v)$ holds since $\mathcal{D}$ satisfies (S1), so we have $\operatorname{core}^{\prime}(w) \subseteq \operatorname{knw}^{\prime}(w) \subseteq \operatorname{poss}^{\prime}(w)$.
(Condition S2) Consider an arbitrary $\epsilon$-component $\mathcal{V}_{u}$ of $\mathcal{D}^{\prime}$, an arbitrary edge $\langle w, z, \exists R . A\rangle \in \mathcal{E}^{\prime}$ such that $z \in \mathcal{V}_{u}$, and an arbitrary literal $\forall \operatorname{inv}(R) . C \in \operatorname{poss}^{\prime}\left(\mathcal{V}_{u}\right)$; furthermore, let $v \in \mathcal{V}$ be the context of $\mathcal{D}$ such that $w \in \mathcal{V}_{v}$. By condition (R7) and $\langle w, z, \exists R . A\rangle \in \mathcal{E}^{\prime}$, we have

$$
\begin{equation*}
\langle v, u, \exists R . A\rangle \in \mathcal{E}, \quad \delta_{v, u}^{\exists R . A} \cap \operatorname{uknw}(v) \subseteq \mathcal{L}_{v}(w), \quad \text { and } \quad \varrho_{v, u}^{\exists R . A} \cap \mathrm{uknw}(v) \subseteq \mathcal{L}_{u}(z) . \tag{141}
\end{equation*}
$$

Now, by property (140) and $\forall \operatorname{inv}(R) . C \in \operatorname{poss}^{\prime}\left(\mathcal{V}_{u}\right)$, we have $\forall \operatorname{inv}(R) . C \in \operatorname{poss}(u)$; hence, by admissibility condition (S2) for $\mathcal{D}$, we have $C \in \operatorname{poss}(v)$. But then, we have $C \in \delta_{v, u}^{\exists R . A}$ and $\forall \operatorname{inv}(R) . C \in \varrho_{v, u}^{\exists R . A}$; together with (141), these conditions imply (142), so $\mathcal{D}^{\prime}$ clearly satisfies admissibility condition (S2).

$$
\begin{equation*}
C \in \mathcal{L}_{v}(w) \cup \operatorname{knw}(v)=\operatorname{poss}^{\prime}(w) \quad \forall i n v(R) . C \in \mathcal{L}_{u}(z) \cup \operatorname{knw}(u)=\operatorname{poss}^{\prime}(z) \tag{142}
\end{equation*}
$$

(Condition S3) Consider an arbitrary $\epsilon$-component $\mathcal{V}_{v}$ of $\mathcal{D}^{\prime}$ and an arbitrary literal $\exists R . A \in \operatorname{poss}^{\prime}\left(\mathcal{V}_{v}\right)$. By property (140), we have $\exists R . A \in \operatorname{poss}(v)$, so, by admissibility condition (S3) for $\mathcal{D}$, an $\exists R . A$-edge $\langle v, u, \exists R . A\rangle \in \mathcal{E}$ exists satisfying (143).

$$
\begin{equation*}
\operatorname{poss}(u) \supseteq\{A\} \cup\{B \mid \forall R \cdot B \in \operatorname{poss}(v)\} \tag{143}
\end{equation*}
$$

Since $\mathcal{T}_{v}$ and $\mathcal{T}_{u}$ are tree decompositions of $\mathcal{H}_{v}$ and $\mathcal{H}_{u}$, respectively, by property (R2) of Definition 42 and property (T1) of tree decompositions, contexts $w \in \mathcal{V}_{v}$ and $z \in \mathcal{V}_{u}$ exist such that the following two properties hold.

$$
\begin{align*}
& \delta_{v, u}^{\exists R, A} \cap \mathrm{uknw}(v) \subseteq \mathcal{L}_{v}(w)  \tag{144}\\
& \varrho_{v, u}^{\exists R . A} \cap \operatorname{uknw}(u) \subseteq \mathcal{L}_{u}(z) \tag{145}
\end{align*}
$$

But then, by condition (R7), we have $\langle w, z, \exists R . A\rangle \in \mathcal{E}^{\prime}$. We next show that this edge satisfies condition (S3)-that is, that
(i) $\operatorname{poss}^{\prime}(z) \supseteq\{A\} \cup\left\{B \mid \forall R . B \in \operatorname{poss}^{\prime}\left(\mathcal{V}_{v}\right)\right\}$, and
(ii) $\operatorname{poss}^{\prime}(w) \supseteq\{\exists R . A\} \cup\left\{\forall R . B \mid \forall R . B \in \operatorname{poss}^{\prime}\left(\mathcal{V}_{v}\right)\right\}$.

By Definition 42 we have $\exists R . A \in \delta_{v, u}^{\exists R . A}$ and $A \in \varrho_{v, u}^{\exists R . A}$; furthermore, for each $\forall R . B \in \operatorname{poss}^{\prime}\left(\mathcal{V}_{v}\right)$, we have $\forall R . B \in \operatorname{poss}(v)$ by (140), and $B \in \operatorname{poss}(u)$ by (143), which imply $\forall R . B \in \delta_{v, u}^{\exists \mathcal{P . A}}$ and $B \in \varrho_{v, u}^{\exists R . A}$. But then, by (144) and (145), we have

$$
\begin{aligned}
\{A\} \cup\left\{B \mid \forall R . B \in \operatorname{poss}^{\prime}\left(\mathcal{V}_{v}\right)\right\} & \subseteq \mathcal{L}_{u}(z) \cup \operatorname{knw}(u)=\operatorname{poss}^{\prime}(z) & \text { which proves (i), and } \\
\{\exists R . A\} \cup\left\{\forall R . B \mid \forall R . B \in \operatorname{posss}^{\prime}\left(\mathcal{V}_{v}\right)\right\} & \subseteq \mathcal{L}_{v}(w) \cup \operatorname{knw}(w)=\operatorname{poss}^{\prime}(w) & \text { which proves (ii). }
\end{aligned}
$$

## Ontology condition:

Consider an arbitrary $\epsilon$-component $\mathcal{V}_{v}$ of $\mathcal{D}^{\prime}$ and an arbitrary clause $K \sqsubseteq M \in \mathcal{O}$ with $K \subseteq \operatorname{poss}^{\prime}\left(\mathcal{V}_{v}\right)$. By (140), $K \subseteq \operatorname{poss}^{\prime}\left(\mathcal{V}_{v}\right)$ implies $K \subseteq \operatorname{poss}(v)$, so by the ontology condition for $\mathcal{D}$ we have $M \subseteq \operatorname{poss}(v)$. By condition (R1) and property (T1) of tree
decompositions, a context $w \in \mathcal{V}_{v}$ exists such that $(K \cup M) \cap u k n w(v) \subseteq \mathcal{L}_{v}(w)$. But then, $K \cup M \subseteq \mathcal{L}_{v}(w) \cup \operatorname{knw}(v)=\operatorname{poss}^{\prime}(w)$ holds, so $\mathcal{D}^{\prime}$ satisfies the ontology condition.

## Ordering conditions:

(Condition P1) Consider an arbitrary $\exists R . A$-edge $\langle w, z, \exists R . A\rangle \in \mathcal{E}^{\prime}$. By condition (R7), then $v, u \in \mathcal{V}$ exist such that $w \in \mathcal{V}_{v}$, $z \in \mathcal{V}_{u}$, and $\langle v, u, \exists R . A\rangle \in \mathcal{E}$. Since $\mathcal{D}$ satisfies admissibility condition (P1), each literal $\forall \operatorname{inv}(R) . C \in \mathbf{L}$ is $<_{u}$-minimal. But then, by condition (R8), each such $\forall \operatorname{inv}(R) . C$ is also $<_{z}^{\prime}$-minimal; hence, $<_{z}^{\prime}$ is $R$-admissible.
(Condition P2) By condition (R8), for each $\epsilon$-edge $\langle u, w, \epsilon\rangle \in \mathcal{E}^{\prime}$, all literals in $\operatorname{poss}^{\prime}(u) \cap \operatorname{poss}^{\prime}(w)$ are $<_{w}^{\prime}$-minimal.

## Covering condition:

Consider an arbitrary query $q=K \sqsubseteq M \in \mathcal{Q}$, and let $v=\vartheta(q)$ and $w=\vartheta^{\prime}(q)$. By the covering condition for $\mathcal{D}$, context $v$ cov-
 This and condition (R5) imply that

$$
\operatorname{core}^{\prime}(w)=\operatorname{core}(v) \subseteq K \subseteq \operatorname{knw}(v) \cup \mathcal{L}_{v}(w)=\operatorname{poss}^{\prime}(w)
$$

Furthermore, by (140) and condition (R5), we have

$$
M \cap \operatorname{poss}^{\prime}\left(\mathcal{V}_{v}\right) \subseteq M \cap \operatorname{poss}(v)=M \cap[\operatorname{knw}(v) \cup u k n w(v)] \subseteq \operatorname{knw}(v) \cup[M \cap u k n w(v)] \subseteq \operatorname{knw}(v) \cup \mathcal{L}_{v}(w)=\operatorname{poss}^{\prime}(w)
$$

Finally, by condition (R8), disjunction $M$ is $\prec_{w}^{\prime}$-minimal because $M$ is $<_{\nu}$-minimal. Thus, context $w$ covers $K \sqsubseteq M$ in $\mathcal{D}^{\prime}$.
Definition 42 can be straightforwardly turned into an algorithm for computing an $\epsilon$-refinement $\mathcal{D}^{\prime}$ of $\mathcal{D}$ : first, we compute all hypergraphs satisfying conditions (R1)-(R3); second, we compute a tree decomposition $\mathcal{T}_{v}$ for each hypergraph $\mathcal{H}_{v}$; third, we construct decomposition $\mathcal{D}^{\prime}$ so that conditions (R4)-(R9) are satisfied; and fourth, for each context $v$ of $\mathcal{D}$ and each $w \in \mathcal{V}_{v}$, we define $<_{w}^{\prime}$ by modifying $<_{v}$ to ensure condition (R8). As explained earlier, in condition (R9) we can don't-care nondeterministically choose $\vartheta(q)$, and at least one suitable context is guaranteed to exist. Such an algorithm is fixed-parameter tractable, as shown by the following theorem.

Theorem 44. For each integer mwd, one can compute an $\epsilon$-refinement $\mathcal{D}^{\prime}$ of $\mathcal{D}$ with $\mathrm{wd}\left(\mathcal{D}^{\prime}\right) \leq \mathrm{mwd}$, or determine that such $\mathcal{D}^{\prime}$ does not exist, in time polynomial in $f(\mathrm{mwd}), \ln (\mathcal{D})$, and $\|\mathcal{O}\|+\|\mathcal{Q}\|$, for $f$ a computable function.
Proof. The L-hypergraphs $\mathcal{H}_{v}$ corresponding to conditions (R1)-(R3) of Definition (42) can clearly be constructed in time polynomial in $\ln (\mathcal{D})$ and $\|\mathcal{O}\|+\|\mathcal{Q}\|$.

An $\epsilon$-refinement $\mathcal{D}^{\prime}$ with $\operatorname{wd}\left(\mathcal{D}^{\prime}\right) \leq m w d$ exists if and only if, for each $v \in \mathcal{V}$, hypertraph $\mathcal{H}_{v}$ admits a tree decomposition $\mathcal{T}_{v}$ with $\operatorname{wd}\left(\mathcal{T}_{v}\right) \leq$ mwd. For each $v \in \mathcal{V}$, all literals occurring in $\mathcal{N}_{v}$ are contained in $\mathcal{O} \cup \mathcal{Q}$, so by Lemma 1, computing one $\mathcal{T}_{v}$ with $\mathrm{wd}\left(\mathcal{T}_{v}\right) \leq \mathrm{mwd}$, or determining that one does not exist can be done in time $O(f(\mathrm{mwd}) \cdot(\|\mathcal{O}\|+\|\mathcal{Q}\|))$.

Decomposition $\mathcal{D}^{\prime}$ can be obtained from $\mathcal{D}$ by simply following conditions (R4)-(R8). Furthermore, for each query $q \in \mathcal{Q}$, one can define $\vartheta^{\prime}(q)$ as an arbitrary element of $\mathcal{V}_{\vartheta(q)}$ satisfying condition (R9); due to condition (R3) and property (T1) of tree decomposition, such an element is guaranteed to exist. This construction is polynomial in $\ln \left(\mathcal{D}^{\prime}\right)$ and $\|\mathcal{O}\|+\|\mathcal{Q}\|$. Since the length of each $\mathcal{D}_{\nu}$ is bounded by the time it takes to compute it, the length of $\mathcal{D}^{\prime}$ is bounded by $O(\ln (\mathcal{D}) \cdot f(\mathrm{mwd}) \cdot(\|\mathcal{O}\|+\|\mathcal{Q}\|)$ ), so the entire construction runs in time polynomial in $f(\mathrm{mwd}), \ln (\mathcal{D})$, and $\|\mathcal{O}\|+\|\mathcal{Q}\|$.

This allows us to generalize our result on fixed-parameter tractability of subsumption reasoning from Section 4.
Theorem 45. For each control $\mathcal{C}$, the following problem is fixed-parameter tractable:

- Inputs: a normalized $\mathcal{A L C I}$ ontology $\mathcal{O}$ and a set of queries $\mathcal{Q}$
- Parameter: an integer mwd
- Problem: return "yes" if an $\epsilon$-refinement $\mathcal{D}$ of the $\mathcal{C}$-decomposition of $\mathcal{O}$ and $\mathcal{Q}$ exists with $\operatorname{wd}(\mathcal{D}) \leq m w d$, and if also $\mathcal{O} \vDash K \sqsubseteq M$ holds for each query $K \sqsubseteq M \in \mathcal{Q}$

Proof. Immediate by Theorems 29, 43, and 44, and Propositions 40 and 41.

## 6. Soundness- and Admissibility-Preserving Decomposition Transformations

In this section we present several soundness- and admissibility-preserving decomposition transformations that, under certain conditions, can delete edges and/or contexts from decompositions. These transformations can be useful in practice as they can reduce the number of inferences of the consequence-based algorithm; moreover, we use these transformations in Section 7 to establish an exponential upper bound on the decomposition size. Throughout this section, we fix a normalized $\mathcal{A} \mathcal{L C I}$ ontology $\mathcal{O}$, a finite set of queries $\mathcal{Q}$, and a sound and admissible decomposition $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, knw, poss, $\langle, \vartheta\rangle$ of $\mathcal{O}$ and $\mathcal{Q}$.

Definition 46 introduces a notion of redundancy for $\epsilon$-edges, as well as a notion of $\epsilon$-edge contraction that extends edge contraction from graph theory: given an edge $\left\langle v_{1}, v_{2}, \epsilon\right\rangle$ in $\mathcal{D}$, we merge contexts $v_{1}$ and $v_{2}$ into a new context $v$, eliminate the self-loop from $v$ to $v$ in order to preserve the tree shape of the enclosing $\epsilon$-component, and define the core, known, and possible literals and the literal ordering of $v$ to reflect both $v_{1}$ and $v_{2}$. Lemma 47 captures the properties of $\epsilon$-edge contraction.
Definition 46. An $\epsilon$-edge $\left\langle v_{1}, v_{2}, \epsilon\right\rangle \in \mathcal{E}$ is redundant in $\mathcal{D}$ if $\operatorname{uknw}\left(v_{1}\right) \subseteq \operatorname{poss}\left(v_{2}\right)$. The result of contracting $a$ (not necessarily redundant) $\epsilon$-edge $\left\langle v_{1}, v_{2}, \epsilon\right\rangle \in \mathcal{E}$ is obtained as follows.
(1) Remove $v_{1}$ and $v_{2}$ from $\mathcal{V}$, and add a fresh context $v$ to $\mathcal{V}$.
(2) Replace each occurrence of $v_{1}$ and $v_{2}$ in $\vartheta$ and in $\mathcal{E}$ by $v$, and then remove the edge $\langle v, v, \epsilon\rangle$ from $\mathcal{E}$.
(3) Set core $(v):=\operatorname{core}\left(v_{1}\right), \operatorname{knw}(v):=\operatorname{knw}\left(v_{1}\right) \cup \operatorname{knw}\left(v_{2}\right), \operatorname{poss}(v):=\operatorname{poss}\left(v_{1}\right) \cup \operatorname{poss}\left(v_{2}\right)$, and $<_{v}:=<_{\nu_{1}} \cap<_{v_{2}}$.

Lemma 47. Contracting an $\epsilon$-edge preserves soundness and admissibility of a decomposition. Contracting a redundant $\epsilon$-edge does not increase the width of a decomposition.
Proof. It is straightforward to check that contracting an $\epsilon$-edge preserves all conditions of Definitions 19 and 36. Assume now that a redundant $\epsilon$-edge $\left\langle v_{1}, v_{2}, \epsilon\right\rangle \in \mathcal{E}$ is contracted. Then, $\operatorname{uknw}\left(v_{1}\right) \subseteq \operatorname{poss}\left(v_{2}\right)$ implies

$$
\begin{aligned}
\operatorname{uknw}(v)= & \operatorname{poss}(v) \backslash \operatorname{knw}(v)=\left[\operatorname{poss}\left(v_{1}\right) \cup \operatorname{poss}\left(v_{2}\right)\right] \backslash\left[\operatorname{knw}\left(v_{1}\right) \cup \operatorname{knw}\left(v_{2}\right)\right] \\
& \left.\subseteq\left[\operatorname{poss}\left(v_{1}\right) \backslash \operatorname{knw}\left(v_{1}\right)\right] \cup \operatorname{poss}\left(v_{2}\right)\right] \backslash \operatorname{knw}\left(v_{2}\right)=\left[u k n w\left(v_{1}\right) \cup \operatorname{poss}\left(v_{2}\right)\right] \backslash \operatorname{knw}\left(v_{2}\right)=\operatorname{poss}\left(v_{2}\right) \backslash \operatorname{knw}\left(v_{2}\right)=u k n w\left(v_{2}\right) .
\end{aligned}
$$

Consequently, the width of a decomposition can only decrease as a result of contraction.
Definition 48 specifies when a context $w$ is broader than a context $u$ in $\mathcal{D}$; roughly speaking, $w$ is then "less specific" regarding core literals, but allows "more" possible literals. Given such $w$ and $u$, we can then redirect each $\exists R . A$-edge ending at $u$ so that the edge ends at $w$. That this operation preserves decomposition soundness and admissibility follows immediately from Lemma 49; moreover, this transformation manipulates only edges, so it clearly preserves decomposition width and length.
Definition 48. A context $w \in \mathcal{V}$ of $\mathcal{D}$ is broader than a context $u \in \mathcal{V}$ of $\mathcal{D}$ if $\operatorname{core}(w) \subseteq \operatorname{core}(u)$, $\operatorname{poss}(w) \supseteq \operatorname{poss}(u),<_{w} \subseteq<_{u}$, and $\operatorname{poss}(\mathcal{W}) \subseteq \operatorname{poss}(\mathcal{U})$, where $\mathcal{W}$ and $\mathcal{U}$ are the $\epsilon$-components of $\mathcal{D}$ such that $w \in \mathcal{W}$ and $u \in \mathcal{U}$. For such $w$ and $u$, the result of redirecting $u$ to $w$ in $\mathcal{D}$ is obtained by replacing each edge $\langle v, u, \exists R . A\rangle$ in $\mathcal{E}$ with $\langle v, w, \exists R . A\rangle$, and by replacing $u$ with $w$ in $\vartheta$.
Lemma 49. Let $w, u \in \mathcal{V}$ be contexts of $\mathcal{D}$ such that $w$ is broader than $u$. Then,

1. replacing an arbitrary $\exists R . A$-edge $\langle v, u, \exists R . A\rangle$ in $\mathcal{E}$ with $\langle v, w, \exists R . A\rangle$ preserves soundness and admissibility, and
2. each query $q$ that is covered in $u$ is also covered in $w$.

Proof. (Claim 1) Assume that some $\langle v, u, \exists R . A\rangle$ in $\mathcal{E}$ is replaced with $\langle v, w, \exists R . A\rangle$. Only conditions (S2), (S3), and (P1) of Definition 36 concern $\exists R . A$-edges. To show that the "new" edge satisfies (S2), consider an arbitrary literal $\forall \operatorname{inv}(R) . C \in \operatorname{poss}(\mathcal{W})$; then $\operatorname{poss}(\mathcal{W}) \subseteq \operatorname{poss}(\mathcal{U})$ implies $\forall \operatorname{inv}(R) . C \in \operatorname{poss}(\mathcal{U})$; hence, by applying condition (S2) to the "old" edge $\langle v, u, \exists R . A\rangle$, we have $C \in \operatorname{poss}(v)$ and $\forall \operatorname{inv}(R) . C \in \operatorname{poss}(u) \subseteq \operatorname{poss}(w)$, as required. Condition (S3) is clearly preserved due to $\operatorname{poss}(u) \subseteq \operatorname{poss}(w)$. Condition ( P 1 ) is preserved since $<_{w} \subseteq<_{u}$ and a subset of an $R$-admissible literal ordering is $R$-admissible.
(Claim 2) Consider an arbitrary query $K \sqsubseteq M$ that is covered in $u$; hence, core $(u) \subseteq K \subseteq \operatorname{poss}(u), M \cap \operatorname{poss}(\mathcal{U}) \subseteq \operatorname{poss}(u)$, and $M$ is $<_{u}$-minimal. Since context $w$ is broader than context $u$, we have core $(w) \subseteq K \subseteq \operatorname{poss}(w), M \cap \operatorname{poss}(\mathcal{W}) \subseteq \operatorname{poss}(w)$, and $M$ is $<_{w}$-minimal; but then, $K \sqsubseteq M$ is also covered in $w$.

Definition 50 introduces a notion of redundancy for $\exists R . A$-edges. Intuitively, an $\exists R . A$-edge $\langle v, u, \exists R . A\rangle$ is redundant in $\mathcal{D}$ if either $\exists R . A$ is not possible in the $\epsilon$-component of $v$ so condition (S3) of Definition 36 is satisfied vacuously, or $\mathcal{D}$ contains another $\exists R . A$-edge that satisfies condition (S3) for $\exists R . A$. It is obvious that deleting a redundant $\exists R . A$-edge in $\mathcal{D}$ preserves soundness and admissibility; moreover, doing so can be particularly useful if $\mathcal{D}$ is an $\epsilon$-refinement of an $\epsilon$-free decomposition. In particular, for each $v, u \in \mathcal{V}$ in condition (R7) of Definition 42, decomposition $\mathcal{D}$ then contains an $\exists R$. $A$-edge from each $w \in \mathcal{V}_{v}$ to each $z \in \mathcal{V}_{u}$ that satisfy $\delta_{v, u}^{\exists R . A} \cap u k n w(v) \subseteq \mathcal{L}_{v}(w)$ and $\varrho_{v, u}^{\exists R . A} \cap u k n w(u) \subseteq \mathcal{L}_{u}(z)$; however, only one such edge may be needed to satisfy admissibility condition ( S 3 ), so deleting the redundant $\exists R$. $A$-edges might be beneficial in practice.
Definition 50. An $\exists R . A$-edge $\langle v, u, \exists R . A\rangle \in \mathcal{E}$ is redundant in $\mathcal{D}$ if either $\exists R . A \notin \operatorname{poss}(\mathcal{W})$, or another $\exists R . A$-edge $\langle w, z, \exists R . A\rangle \in \mathcal{E}$ with $w \in \mathcal{W}$ exists that satisfies condition $(\mathrm{S} 3)$ of Definition 36 for $\exists R . A \in \operatorname{poss}(\mathcal{W})$, where $\mathcal{W}$ is the $\epsilon$-component containing $v$.

Definition 51 says that an entire $\epsilon$-component $\mathcal{U}$ of $\mathcal{D}$ may be redundant if, for each context $u \in \mathcal{U}$, either (i) there is no $\exists R$. $A$-edge from another $\epsilon$-component to $u$, and $u$ is not used to cover a query in $\mathcal{Q}$, or (ii) $u$ can be redirected to a context in another $\epsilon$-component. In such a case, we can redirect appropriately all contexts satisfying condition (ii) to obtain a decomposition in which no context in $\mathcal{U}$ has an incoming $\exists R . A$-edge from another $\epsilon$-component; but then, we can clearly delete all contexts in $\mathcal{U}$ and all the related edges without affecting decomposition admissibility.

Definition 51. An $\epsilon$-component $\mathcal{U}$ of $\mathcal{D}$ is redundant in $\mathcal{D}$ if at least one of the following holds for each context $u \in \mathcal{U}$ :

- $v \in \mathcal{U}$ for each $\exists R$.A-edge $\langle v, u, \exists R . A\rangle \in \mathcal{E}$, and $\vartheta(q) \neq u$ for each query $q \in \mathcal{Q}$; or
- a context $w \in \mathcal{V} \backslash \mathcal{U}$ exists that is broader than $u$.

A result of eliminating such $\mathcal{U}$ from $\mathcal{D}$ is obtained from $\mathcal{D}$ as follows.

1. For each context $u \in \mathcal{U}$ such that either there exists an edge $\exists R$.A-edge $\langle v, u, \exists R . A\rangle \in \mathcal{E}$ with $v \notin \mathcal{U}$, or there exists a query $q \in \mathcal{Q}$ with $\vartheta(q)=u$, redirect $u$ to some context $w \in \mathcal{V} \backslash \mathcal{U}$ that is broader than $u$.
2. Remove from the result of the previous step all contexts in $\mathcal{U}$ and all edges that are incident to a context in $\mathcal{U}$.

Lemma 52. Eliminating a redundant $\epsilon$-component preserves soundness and admissibility, and it does not increase the width.
Proof. By Lemma 49, step 1 preserves soundness and admissibility. Furthermore, by the definition of when $\mathcal{U}$ is redundant in $\mathcal{D}$, after step 1 each context $u \in \mathcal{U}$ neither occurs in an $\exists R . A$-edge $\langle v, u, \exists R . A\rangle$ with $v \notin \mathcal{U}$, nor is it used to cover a query from $\mathcal{Q}$; hence, step 2 preserves decomposition admissibility. Furthermore, removal of edges and contexts clearly preserves decomposition soundness. Finally, steps 1 and 2 clearly do not increase the width of a decomposition.

Please note that context redirection, elimination of $\exists R . A$-edges, and elimination of redundant $\epsilon$-components also apply to context structures (with obvious modifications), and they can be applied during execution of Algorithm 11. In this way, we obtain "structural" redundancy elimination rules that can further improve the performance of the consequence-based algorithm.

## 7. Bounds on Decomposition Length

In this section we discuss the bounds on the length of decompositions of minimum width. In particular, in Section 7.1 we consider the lower bound, and in Section 7.2 we turn our attention to the upper bound.

### 7.1. Lower Bound

We now exhibit a family of ontologies $\left\{\mathcal{O}_{n}\right\}$ and a fixed set of queries $\mathcal{Q}$ such that each sound and admissible decomposition of $\mathcal{O}_{n}$ and $\mathcal{Q}$ of minimum width necessarily has exponential length. Intuitively, this is because each ontology $\mathcal{O}_{n}$ has a "canonical" model $\mathcal{I}$ containing exponentially many domain elements, each of which satisfies a distinct combination of atomic concepts. Therefore, to obtain a decomposition of $\mathcal{O}_{n}$ and $\mathcal{Q}$ of minimal width, we must introduce a context for each domain element of $\mathcal{I}$, and thus such a decomposition will actually reflect the structure of $\mathcal{I}$. This is interesting because it shows that, in general, one cannot minimize decomposition width and still expect to obtain a practically useful decomposition (i.e., a decomposition of polynomial size). We use this observation as a justification for our practical approach to decomposition construction that we presented in Sections 4.4 and 5.3.

Theorem 53. Let $\mathcal{Q}=\left\{C_{0} \sqsubseteq \perp\right\}$. A family of $\mathcal{A L C I}$ ontologies $\left\{\mathcal{O}_{n}\right\}$ exists such that, for each $\mathcal{O}_{n}$,

1. a sound, admissible, and $\epsilon$-free decomposition of $\mathcal{O}_{n}$ and $\mathcal{Q}$ of width 0 exists, and
2. each sound and admissible decomposition of $\mathcal{O}_{n}$ and $\mathcal{Q}$ of width 0 has length at least exponential in $\left\|\mathcal{O}_{n}\right\|$.

For readability, we split the proof of Theorem 53 into several claims. Let $\mathcal{Q}=\left\{C_{0} \sqsubseteq \perp\right\}$, let $n$ be a positive integer, and let $\mathcal{O}_{n}$ be the ontology ( of size linear in $n$ ) containing the following axioms for each $1 \leq i \leq n$.

$$
\begin{equation*}
C_{i-1} \sqsubseteq \exists R . C_{i} \tag{146}
\end{equation*}
$$

$$
\begin{align*}
C_{i-1} & \sqsubseteq \exists S . C_{i}  \tag{150}\\
C_{i-1} & \sqsubseteq \forall S . B_{i}  \tag{147}\\
A_{i} & \sqsubseteq \forall S . A_{i}  \tag{148}\\
B_{i} & \sqsubseteq \forall S . B_{i} \tag{149}
\end{align*}
$$

Let $\mathbf{L}=\left\{C_{0}\right\} \cup\left\{C_{i}, A_{i}, B_{i}, \exists R . C_{i}, \exists S . C_{i}, \forall R . A_{i}, \forall S . A_{i}, \forall R . B_{i}, \forall S . B_{i} \mid 1 \leq i \leq n\right\}$ be the set of literals that occur in $\mathcal{O}_{n}$ and $\mathcal{Q}$. An ABC-number is a set of atomic concepts of the form $X=\left\{X_{1}, \ldots, X_{k}, C_{k}\right\}$, where $0 \leq k \leq n$ and each $X_{i}$ is either $A_{i}$ or $B_{i}$. By a slight abuse of notation, we often treat $X$ as the conjunction of the atomic concepts contained in $X$; furthermore, the rank of $X$ is $k$; finally, $X^{R}$ and $X^{S}$ are ABC-numbers of rank $k+1$ defined as

$$
X^{R}:=\left\{X_{1}, \ldots, X_{k}, A_{k+1}, C_{k+1}\right\} \quad \text { and } \quad X^{S}:=\left\{X_{1}, \ldots, X_{k}, B_{k+1}, C_{k+1}\right\} .
$$

Note that $\mathcal{O}_{n} \vDash X \sqsubseteq \exists R . X^{R}$ and $\mathcal{O}_{n} \vDash X \sqsubseteq \exists S . X^{S}$. Finally, for each ABC-number $X$, we define the set $\Gamma(X)$ of told subsumees of $X$ w.r.t. $\mathcal{O}_{n}$ as follows:

$$
\Gamma(X):=X \cup\left\{L \in \mathbf{L} \mid \text { an atomic concept } D \in X \text { exists such that } D \sqsubseteq L \in \mathcal{O}_{n}\right\}
$$

One can readily check that $\mathcal{O}_{n}$ entails only told subsumptions-that is, for each ABC-number $X$ and each $L \in \mathbf{L}$, we have

$$
\begin{array}{ll}
\mathcal{O}_{n} \vDash X \sqsubseteq L & \text { if and only if } L \in \Gamma(X) ; \\
\mathcal{O}_{n} \vDash X \sqsubseteq \exists R . L & \text { if and only if } L \in \Gamma\left(X^{R}\right) ; \\
\mathcal{O}_{n} \models X \sqsubseteq \exists S . L & \text { if and only if } L \in \Gamma\left(X^{S}\right) .
\end{array}
$$

The following claim proves the first item of Theorem 53.
Claim 54. There exists a sound, admissible, and $\epsilon$-free decomposition $\mathcal{D}$ of $\mathcal{O}_{n}$ and $\mathcal{Q}$ such that $\operatorname{wd}(\mathcal{D})=0$.
Proof. Let $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, knw, poss, $\langle, \vartheta\rangle$ be defined as follows:

$$
\begin{aligned}
\mathcal{V} & :=\left\{v_{X} \mid X \text { is an ABC-number }\right\} \\
\operatorname{core}\left(v_{X}\right) & :=X \\
\operatorname{knw}\left(v_{X}\right) & :=\operatorname{poss}\left(v_{X}\right):=\Gamma(X) \\
\prec_{v_{X}} & :=\emptyset \\
\vartheta\left(C_{0} \sqsubseteq \perp\right) & :=v_{\left\{C_{0}\right\}}^{\mathcal{E}}
\end{aligned}:=\left\{\left\langle v_{X}, v_{X^{R}}, \exists R . C_{k+1}\right\rangle,\left\langle v_{X}, v_{X^{s}}, \exists S . C_{k+1}\right\rangle \mid \text { for each ABC-number } X \text { of rank } k \text { with } 0 \leq k<n\right\}
$$

Clearly, $\mathcal{D}$ is an $\epsilon$-free decomposition of $\mathcal{O}_{n}$ and $\mathcal{Q}$. Furthermore, it is straightforward to check that $\mathcal{D}$ is sound. Finally, $\mathcal{D}$ satisfies all the admissibility conditions of Definition 20: conditions (S1) and (S2) hold trivially, the definition of $\mathcal{E}$ ensures that condition (S3) is satisfied, the definition of poss ensures that the ontology condition is satisfied, the trivial literal ordering $<_{v_{X}}$ always satisfies the ordering condition, and the query $C_{0} \sqsubseteq \perp$ is covered in context $v_{\left\{C_{0}\right\}}$. Finally, since $\operatorname{knw}\left(v_{X}\right)=\operatorname{poss}\left(v_{X}\right)$ for each context $v_{X} \in \mathcal{V}$, we have $\operatorname{wd}(\mathcal{D})=0$, as required.

To prove the second item of Theorem 53 , let $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, knw, poss, $\langle, \vartheta\rangle$ be an arbitrary sound and admissible decomposition of $\mathcal{O}_{n}$ and $\mathcal{Q}$ such that $\operatorname{wd}(\mathcal{D})=0$. Every $\epsilon$-edge in a decomposition of width 0 is redundant in the sense of Definition 46, so we can assume without loss of generality that $\mathcal{D}$ is $\epsilon$-free. We prove in Claim 56 that, for each ABC-number $X$, there exists a context $v_{X} \in \mathcal{V}$ such that $\operatorname{poss}\left(v_{X}\right)=\Gamma(X)$. Now $\Gamma(X)$ is distinct for each ABC-number $X$, so $\mathcal{V}$ clearly contains at least $2^{n}$ contexts, which concludes the proof of Theorem 53. Towards proving Claim 56, we first prove the following auxiliary claim.
Claim 55. Let $v \in \mathcal{V}$ be a context and let $X$ be an $A B C$-number. If $\operatorname{core}(v) \subseteq \Gamma(X)$ and $\operatorname{poss}(v) \supseteq X$, then $\operatorname{poss}(v)=\Gamma(X)$.
Proof. Let $v \in \mathcal{V}$ be a context, and let $X$ be an ABC-number such that core $(v) \subseteq \Gamma(X)$ and $\operatorname{poss}(v) \supseteq X$ hold. By the ontology condition of Definition 20, set $\operatorname{poss}(v)$ is closed under $\Gamma$; thus, poss $(v) \supseteq X$ implies poss $(v) \supseteq \Gamma(X)$. To show that $\operatorname{poss}(v) \subseteq \Gamma(X)$ holds as well, consider an arbitrary literal $L \in \operatorname{poss}(v)$. Since $\operatorname{wd}(\mathcal{D})=0$, we have $L \in \operatorname{knw}(v)$. But then, by soundness condition (N3) of Definition 19, we have $\mathcal{O}_{n} \vDash \operatorname{core}(v) \sqsubseteq L$. Furthermore, we have $\mathcal{O}_{n} \vDash X \sqsubseteq \Gamma(X)$ by property (154), and core $(v) \subseteq \Gamma(X)$ clearly implies $\mathcal{O}_{n} \vDash \Gamma(X) \sqsubseteq$ core $(v)$. Together, all these observations imply that $\mathcal{O}_{n} \vDash X \sqsubseteq L$ holds, which, by property (154), implies $L \in \Gamma(X)$. Consequently, we have $\operatorname{poss}(v)=\Gamma(X)$, as required.

Claim 56. For each $A B C$-number $X$, a context $v_{X} \in \mathcal{V}$ exists such that $\operatorname{poss}\left(v_{X}\right)=\Gamma(X)$.
Proof. The proof is by induction on the rank of $X$. For the base case, we have $X=\left\{C_{0}\right\}$. Since $X \sqsubseteq \perp \in \mathcal{Q}$, query $X \sqsubseteq \perp$ is covered in the context $v_{X}=\vartheta(X \subseteq \perp)$, and so we have core $\left(v_{X}\right) \subseteq X \subseteq \operatorname{poss}\left(v_{X}\right)$ by the definition of covering. But then, Claim 55 implies that $\operatorname{poss}\left(v_{X}\right)=\Gamma(X)$ holds, as required.

For the induction step, assume that the claim holds for each ABC-number of rank $k-1$, and consider an arbitrary ABCnumber of the form $Y=\left\{X_{1}, \ldots, X_{k-1}, X_{k}, C_{k}\right\}$; furthermore, let $X=\left\{X_{1}, \ldots, X_{k-1}, C_{k-1}\right\}$. By the induction hypothesis, a context $v_{X} \in \mathcal{V}$ exists such that $\operatorname{poss}\left(v_{X}\right)=\Gamma(X)$. We now consider the two possible forms of $X_{k}$.

Assume that $X_{k}=A_{k}$, which implies $Y=X^{R}$. By applying admissibility condition (S3) of Definition 20 to an existential restriction $\exists R . C_{k} \in \Gamma(X)=\operatorname{poss}\left(v_{X}\right)$, an edge $\left\langle v_{X}, v_{Y}, \exists R . C_{k}\right\rangle \in \mathcal{E}$ exists such that

$$
\operatorname{poss}\left(v_{Y}\right) \supseteq\left\{C_{k}\right\} \cup\left\{D \mid \forall R . D \in \operatorname{poss}\left(v_{X}\right)\right\}=X^{R} .
$$

Soundness condition (N1) of Definition 19 implies that $\mathcal{O}_{n} \vDash \operatorname{core}\left(v_{X}\right) \sqcap \exists R . C_{k} \sqsubseteq \exists R$.[core $\left.\left(v_{Y}\right) \sqcap C_{k}\right]$. Now $\mathcal{O}_{n} \vDash X \sqsubseteq \Gamma(X)$ holds by property (154); furthermore, [core $\left.\left(v_{X}\right) \sqcap \exists R . C_{k}\right] \subseteq \operatorname{poss}\left(v_{X}\right)=\Gamma(X)$ holds by the definition of poss, which clearly implies $\mathcal{O}_{n} \vDash \Gamma(X) \sqsubseteq\left[\operatorname{core}\left(v_{X}\right) \sqcap \exists R . C_{k}\right]$; thus, we have $\mathcal{O}_{n} \vDash X \sqsubseteq \exists R$. $\left[\operatorname{core}\left(v_{Y}\right) \sqcap C_{k}\right]$; but then, core $\left(v_{Y}\right) \subseteq \Gamma\left(X^{R}\right)$ holds by (155). Thus, for $Y=X^{R}$, we have core $\left(v_{Y}\right) \subseteq \Gamma\left(X^{R}\right)$ and $\operatorname{poss}\left(v_{Y}\right) \supseteq X^{R}$, so Claim 55 implies $\operatorname{poss}\left(v_{Y}\right)=\Gamma\left(X^{R}\right)=\Gamma(Y)$, as required.

The analysis for the case $X_{k}=B_{k}$ is analogous to the one above, with the difference that role $S$ is used instead of role $R$.

### 7.2. Upper Bound

We now show that, for each ontology $\mathcal{O}$ and each set of queries $\mathcal{Q}$, a sound and admissible decomposition of $\mathcal{O}$ and $\mathcal{Q}$ exists that has minimum width and at most exponential length. Intuitively, this is because, given an arbitrary decomposition of minimal width, we can obtain the required decomposition as follows. First, we repeatedly eliminate redundant $\epsilon$-edges; this ensures that each $\epsilon$-component contains at most a polynomial number of contexts (cf. Lemma 57), so the number of distinct $\epsilon$-components is exponential. Second, we eliminate "duplicates" by repeatedly eliminating redundant $\epsilon$-components. We thus obtain a decomposition of $\mathcal{O}$ and $\mathcal{Q}$ of at most exponential length; since none of these transformations increases decomposition width, the resulting decomposition is of minimal width as well (cf. Theorem 58).

Lemma 57. Let $\mathcal{O}$ be a normalized $\mathcal{A L C I}$ ontology, let $\mathcal{Q}$ be a finite set of queries, and let $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, knw, poss, $\langle, \vartheta\rangle$ be an admissible decomposition of $\mathcal{O}$ and $\mathcal{Q}$ that has no redundant $\epsilon$-edges. Then, for each $\epsilon$-component $\mathcal{W}$ of $\mathcal{D}$, we have $|\mathcal{W}| \leq \max (|U|, 1)$ for $U=\bigcup_{w \in \mathcal{W}} \operatorname{uknw}(w)$.

Proof. Let $\mathcal{W}$ be an arbitrary $\epsilon$-component of $\mathcal{D}$ and let $n=|\mathcal{W}|$. If there exists a context $v \in \mathcal{W}$ with $u k n w(v)=\emptyset$, since $\mathcal{D}$ has no redundant $\epsilon$-edges, then $v$ is the only context in $\mathcal{W}$, so $|\mathcal{W}|=1$ and the lemma holds. Thus, in the rest of the proof, we assume that uknw $(v) \neq \emptyset$ for each context $v \in \mathcal{W}$.

Let $v_{1}, \ldots, v_{n}$ be an arbitrary ordering of $\mathcal{W}$ obtained by an arbitrary depth-first traversal of the tree $\mathcal{D}_{\mathcal{W}}$. We define a function $\lambda: U \rightarrow\{1, \ldots, n\}$ by setting $\lambda(L):=\min \left\{i \mid L \in \operatorname{uknw}\left(v_{i}\right)\right\}$. We claim that $\lambda$ is surjective-that is, for each $1 \leq i \leq n$, a literal $L$ exists such that $\lambda(L)=i$. For $i=1$, since uknw $\left(v_{1}\right) \neq \emptyset$, a literal $L \in u k n w\left(v_{1}\right)$ exists such that $\lambda(L)=1$. Consider now an arbitrary $i$ with $1<i \leq n$, and let $v_{j}$ be the parent of $v_{i}$ in the depth-first traversal. Since $\mathcal{D}$ has no redundant $\epsilon$-edges, we have uknw $\left(v_{i}\right) \nsubseteq \operatorname{poss}\left(v_{j}\right)$, so a literal $L$ exists such that $L \in \operatorname{uknw}\left(v_{i}\right) \backslash \operatorname{poss}\left(v_{j}\right)$. Now if $\lambda(L)<i$, then $v_{\lambda(L)}$ is visited in the depth-first traversal before $v_{i}$; thus, the unique path between $v_{\lambda(L)}$ and $v_{i}$ goes through $v_{j}$, so, by the connectedness condition (E3) of admissibility, we have $L \in \operatorname{poss}\left(v_{j}\right)$; however, this contradicts our choice of $L$. Consequently, we have $\lambda(L)=i$, and so $\lambda$ is surjective from $U$ onto $\{1, \ldots, n\}$. But then, $|\mathcal{W}|=n \leq|U|$.

Theorem 58. Let $\mathcal{O}$ be a normalized $\mathcal{A L C I}$ ontology, let $\mathcal{Q}$ be a finite set of queries, and let $\mathbf{L}$ be the set of literals occurring in $\mathcal{O} \cup \mathcal{Q}$. For each sound and admissible decomposition $\mathcal{D}^{\prime}$ of $\mathcal{O}$ and $\mathcal{Q}$, there exists a sound and admissible decomposition $\mathcal{D}$ of $\mathcal{O}$ and $\mathcal{Q}$ such that $\mathrm{wd}(\mathcal{D}) \leq w d\left(\mathcal{D}^{\prime}\right)$ and $\ln (\mathcal{D}) \leq 2^{O\left(\left.\mathbf{L}\right|^{2}\right)}$.

Proof. Let $\mathcal{D}^{\prime}$ be an arbitrary sound and admissible decomposition of $\mathcal{O}$ and $\mathcal{Q}$, and let $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, knw, poss, $\langle, \vartheta\rangle$ be an arbitrary decomposition obtained from $\mathcal{D}^{\prime}$ by applying the following (nondeterministic) transformations.

1. Set $<_{v}:=\emptyset$ for each context $v \in \mathcal{V}$.
2. While there are redundant $\epsilon$-edges, pick one such $\epsilon$-edge and contract it.
3. While there are redundant $\epsilon$-components, pick one such $\epsilon$-component and eliminate it.

By Lemmas 47 and 52, steps 2 and 3 preserve soundness and admissibility, and they do not increase the width; furthermore, step 1 trivially satisfies the two properties as well. To conclude our proof, we next show that $\ln (\mathcal{D}) \leq 2^{O\left(\left.\mathbf{L}\right|^{2}\right)}$.

Due to step 2, decomposition $\mathcal{D}$ contains no redundant $\epsilon$-edges; hence, by Lemma 57, the cardinality of each $\epsilon$-component of $\mathcal{D}$ is bounded by $|\mathbf{L}|$. We define the type of an $\epsilon$-component $\mathcal{W}$ of $\mathcal{D}$ to be the pair $\langle\operatorname{core}(\mathcal{W}),\{\operatorname{poss}(v) \mid v \in \mathcal{W}\}\rangle$. Now assume that $\mathcal{U}$ and $\mathcal{W}$ are two different $\epsilon$-components of $\mathcal{D}$ with the same type; then, for each context $u \in \mathcal{U}$, we can find a context $w \in \mathcal{W}$ with $\operatorname{poss}(u)=\operatorname{poss}(w)$; such $w$ is broader than $u$ because

- $<_{w}=\emptyset=<_{u}$ by step 1 ,
- $\operatorname{core}(w)=\operatorname{core}(\mathcal{W})=\operatorname{core}(\mathcal{U})=\operatorname{core}(u)$ by soundness condition (N2), and
- $\operatorname{poss}(\mathcal{W})=\bigcup_{v \in \mathcal{W}} \operatorname{poss}(v)=\bigcup_{v \in \mathcal{U}} \operatorname{poss}(v)=\operatorname{poss}(\mathcal{U})$.

Consequently, such $\mathcal{U}$ is redundant in $\mathcal{D}$; however, due to step 3 , decomposition $\mathcal{D}$ contains no redundant $\epsilon$-components, so we obtain a contradiction. Therefore, all distinct $\epsilon$-components of $\mathcal{D}$ are of a different type.

Thus, we can bound the number of $\epsilon$-components of $\mathcal{D}$ by computing the number of possible types. For each $\epsilon$-component $\mathcal{W}$ of $\mathcal{D}$, set $\operatorname{core}(\mathcal{W})$ is a subset of $\mathbf{L}$, so there are at most $2^{|\mathbf{L}|}$ such subsets; furthermore, $\operatorname{set}\{\operatorname{poss}(v) \mid v \in \mathcal{W}\}$ is a subset of $2^{\mathbf{L}}$ of cardinality at most $|\mathbf{L}|$, so the number of such subsets is bounded by

$$
\sum_{k=0}^{|\mathbf{L}|}\binom{2^{|\mathbf{L}|}}{k} \leq \sum_{k=0}^{|\mathbf{L}|}\left(2^{|\mathbf{L}|}\right)^{k} \leq\left(2^{|\mathbf{L}|}\right)^{(\mathbf{L} \mid+1)}
$$

Consequently, the number of $\epsilon$-components in $\mathcal{D}$ is bounded by $2^{|\mathbf{L}|} \cdot\left(2^{|\mathbf{L}|}\right)^{(\mathbf{L} \mid+1)}$. Since each $\epsilon$-component contains at most $|\mathbf{L}|$ contexts, we have $\ln (\mathcal{D}) \leq|\mathbf{L}| \cdot 2^{|\mathbf{L}|(\mathbf{L} \mid+2)}$, as required.

## 8. Decomposition Width and Length in Practice

In this section we discuss the results of our experiments, whose goal was to measure decomposition width and length of realistic ontologies using several different expansion strategies.

Our experiments are based on our ontology repository that contains a diverse set of ontologies, including standard benchmarks and numerous ontologies from the life sciences domain. Each ontology in the repository can be accessed using a unique fivedigit identifier. ${ }^{4}$ To obtain our test corpus, we focused mostly on those ontologies in the repository that contain more than 1000 axioms and are not expressed in $\mathcal{E L}$ or DL-Lite $_{\text {horn }}$ : ontologies expressed in one of these two languages have a decomposition of polynomial length and zero width, so we did not consider such ontologies interesting. However, we also included several smaller "toy" ontologies, such as Pizza and Wine, which were developed to demonstrate various ontology modeling constructs. In this way we obtained a corpus of 44 ontologies shown in Table 6 . Since our framework is applicable only to normalized $\mathcal{A L C I}$ ontologies, we further transformed each ontology as follows. First, we eliminated from the ontology all axioms that are not supported in $\mathcal{S H} \mathcal{I}$. Second, we transformed the result into a normalized $\mathcal{S H I}$ ontology using a normalization procedure similar to the one presented by Motik et al. [14]. Third, we eliminated role inclusion and transitivity axioms as described in Section 2 to obtain a normalized $\mathcal{A} \mathcal{L C} \mathcal{I}$ ontology. For each test ontology, Table 6 shows (i) the number of terminological axioms in the original ontology, (ii) the number of atomic concepts, atomic roles, and axioms in the $\mathcal{S H} \mathcal{I}$ ontology after elimination of unsupported axioms, and (iii) the number of literals, roles, and clauses in the final $\mathcal{A L C I}$ ontology after normalization; ontologies are grouped to match Table 7, and the rationale behind the grouping will be explained shortly. For each normalized ontology $\mathcal{O}$, we defined the set $\mathcal{Q}_{\mathcal{O}}$ to contain a query $A \sqsubseteq B$ for all pairs of atomic concepts $A$ and $B$ in $\mathcal{O}$ that were not introduced during normalization; thus, $\mathcal{Q}_{\mathcal{O}}$ contains all queries relevant for ontology classification.

We developed three controls, $\mathcal{C}_{1}, \mathcal{C}_{2}$, and $\mathcal{C}_{3}$, that use the cautious, refined cautious, and eager expansion strategies, respectively, described in Section 3.4.2. Each control $\mathcal{C}_{i}$ used $\mathrm{mln}=\infty$-that is, we did not restrict the number of contexts introduced by the strategies. Although this is not formally allowed in our presentation of the algorithm, it allowed us to see how the strategies perform when run freely, without being forced to reuse contexts. The maximum number of contexts that could be introduced by $\mathcal{C}_{1}$ and $\mathcal{C}_{2}$ was thus polynomial, but for $\mathcal{C}_{3}$ it was exponential in the ontology size. Finally, each $\mathcal{C}_{i}$ used an initialization function that, for each atomic concept $A$ occurring, introduces a single context $v_{A}$ with core $\left(v_{A}\right)=\{A\}$ to cover all queries of the form $A \sqsubseteq B$. Please note that the expansion strategies of $\mathcal{C}_{1}$ and $\mathcal{C}_{3}$ can reuse these "initial" contexts, but the refined cautious expansion strategy of $\mathcal{C}_{2}$ cannot since a context of the form $v_{A}$ is distinct from a context of the form $v_{A}^{R}$ for some role $R$.

For each normalized ontology $\mathcal{O}$, we conducted the following experiments. First, for each $1 \leq i \leq 3$, we computed the $\mathcal{C}_{i^{-}}$ decomposition $\mathcal{D}_{i}$ of $\mathcal{O}$ and $\mathcal{Q}_{\mathcal{O}}$ using the decomposition construction algorithm from Section 4.4. Second, for each $1 \leq i \leq 3$, we computed the width of an $\epsilon$-refinement $\mathcal{D}_{i}^{\prime}$ of $\mathcal{D}_{i}$ using the approach from Section 5.3. We computed tree decompositions of the relevant hypergraphs using the TreeD library; ${ }^{5}$ however, since our hypergraphs were large, we used a heuristic mode in which the library returns only an approximation of the treewidth, without the actual tree decomposition. Thus, we were able to establish only an upper bound on the width, but not the length of $\mathcal{D}_{i}^{\prime}$; however, by Lemma 57, we know that the size of each $\epsilon$-component of $\mathcal{D}_{i}^{\prime}$ is at $\operatorname{most} \mathrm{wd}\left(\mathcal{D}_{i}\right)$, so an upper bound on $\ln \left(\mathcal{D}_{i}^{\prime}\right)$ is given by $\ln \left(\mathcal{D}_{i}\right) \cdot \mathrm{wd}\left(\mathcal{D}_{i}\right)$.

Table 7 summarizes the results of our experiments. For readability, the ontologies were grouped into three subtables: Table 7a contains those ontologies for which all three strategies produce decompositions of different widths; Table 7 b contains those ontologies for which the cautious and the refined cautious strategies produce decompositions of equal widths; and Table 7c contains those ontologies for which the choice of a strategy does not affect decomposition width. Moreover, each subtable is split into two parts, where the bottom part contains ontologies for which $\operatorname{wd}\left(\mathcal{D}_{3}\right)=0$ : these are all Horn ontologies on which the decomposition construction algorithm with the eager strategy fully solves the reasoning problem. On the "go-anatomy-importer" ontology, the cautious and the refined cautious strategies produce decompositions of very large width, which the TreeD library was unable to process. On the "nif-gross" ontology, the eager strategy seems to produce a decomposition of very large length, and our decomposition construction algorithm ran out of memory after introducing about five million contexts.

Our experiments show that decomposition length, even with the eager strategy, is in most cases commensurate with the number of atomic concepts. This is the main reason why consequence-based algorithms perform well in practice: our test ontologies do not seem to incur a substantial amount of and-branching; moreover, unlike the hypertableau algorithms, consequence-based algorithms do not suffer from unnecessary and-branching. The only outlier is the "nif-gross" ontology, on which the eager strategy produces very large decompositions; the cautious strategy, however, produces a decomposition of reasonably small width. This suggests that the hybrid strategy from Section 3.4 .2 might actually be very useful in practice.

The refined cautious strategy does not seem to provide much practical benefit over the cautious strategy: there are only ten ontologies (all shown in Table 7a) on which $\operatorname{wd}\left(\mathcal{D}_{1}\right)>\operatorname{wd}\left(\mathcal{D}_{2}\right)$ holds, and in all cases the difference between $\operatorname{wd}\left(\mathcal{D}_{1}\right)$ and $\operatorname{wd}\left(\mathcal{D}_{2}\right)$ is negligible. Furthermore, only four decompositions satisfy $\operatorname{wd}\left(\mathcal{D}_{1}^{\prime}\right)>\operatorname{wd}\left(\mathcal{D}_{2}^{\prime}\right)$-that is, in all other cases, computing an

[^3]Table 6: Statistics for test ontologies

| Ontology | ID | Original logic | TBox | Reduced to $\mathcal{S H I}$ |  |  | Normalized |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  |  | A | $T$ | axioms | $L$ | $R$ | clauses |
| acgt | 00001 | $\mathcal{S R O I Q}$ (D) | 5,457 | 1,750 | 247 | 4,981 | 3,851 | 417 | 9,039 |
| fma-cons | 00285 | $\mathcal{A L C O I F}(\mathcal{D})$ | 123,090 | 41,646 | 139 | 116,270 | 85,576 | 243 | 121,833 |
| go-anatomy-importer | 00040 | $\mathcal{S R I Q}$ | 130,480 | 58,193 | 209 | 130,411 | 177,934 | 356 | 279,928 |
| nif-gross | 00354 | $\mathcal{S R O I F}(\mathcal{D})$ | 6,630 | 4,042 | 62 | 6,580 | 6,126 | 79 | 9,493 |
| uberon | 00658 | $\mathcal{S R I \mathcal { L }}$ | 23,644 | 7,838 | 82 | 23,575 | 42,960 | 118 | 77,373 |
| envo-xp | 00448 | $\mathcal{S R I F}$ | 74,103 | 35,702 | 95 | 74,065 | 50,836 | 51 | 75,154 |
| galen7 | 00792 | $\mathcal{A L E R I F}(\mathcal{D})$ | 45,260 | 28,447 | 964 | 44,531 | 341,544 | 1,610 | 370,519 |
| go-xp-all | 00483 | $\mathcal{S R I}$ | 115,252 | 89,926 | 191 | 115,214 | 238,923 | 319 | 295,029 |
| pato | 00762 | SHIF | 7,616 | 2,307 | 24 | 2,820 | 1,737 | 34 | 2,641 |
| plant-trait-xp | 00573 | $\mathcal{S R I F}$ | 127,603 | 60,968 | 96 | 127,559 | 85,831 | 46 | 130,235 |
| cdao | 00410 | $\mathcal{S R O I Q}(\mathcal{D})$ | 3,175 | 890 | 106 | 3,056 | 2,934 | 192 | 6,152 |
| dolce-all | 00024 | $\mathcal{S H O I N}(\mathcal{D})$ | 1,544 | 204 | 313 | 1,502 | 10,679 | 626 | 16,442 |
| ero | 00450 | $\mathcal{S H O I F}(\mathcal{D})$ | 3,277 | 2,397 | 103 | 3,146 | 3,965 | 189 | 5,075 |
| gardiner-wafa | 00055 | $\mathcal{S H I N}$ | 246 | 152 | 20 | 242 | 277 | 27 | 317 |
| influenza-ontology | 00507 | $\mathcal{S R O I N}(\mathcal{D})$ | 1,552 | 745 | 66 | 1,504 | 3,739 | 116 | 5,963 |
| lipid | 00512 | $\mathcal{A L C H I N}$ | 2,375 | 710 | 46 | 2,192 | 3,055 | 92 | 6,470 |
| obi | 00350 | $\mathcal{S H O I N}(\mathcal{D})$ | 9,926 | 2,634 | 69 | 9,866 | 6,009 | 125 | 19,626 |
| pizza | 00793 | SHOIN | 701 | 97 | 7 | 686 | 391 | 14 | 1,355 |
| propreo | 00772 | $\mathcal{S H I N}$ | 565 | 474 | 30 | 548 | 994 | 48 | 1,086 |
| sweet-space | 00788 | $\mathcal{S H O I N}(\mathcal{D})$ | 2,430 | 1,514 | 123 | 2,132 | 2,774 | 210 | 3,658 |
| sweet-numerics | 00789 | $\mathcal{S H O I N}(\mathcal{D})$ | 2,499 | 1,499 | 137 | 2,185 | 2,768 | 239 | 3,599 |
| sweet-phenomena | 00790 | $\mathcal{S H O I N}(\mathcal{D})$ | 2,658 | 1,721 | 112 | 2,367 | 2,871 | 191 | 3,734 |
| fly-anatomy | 00460 | $\mathcal{S R I}$ | 19,477 | 7,798 | 21 | 19,466 | 10,391 | 27 | 19,534 |
| galen-doctored | 00029 | $\mathcal{A L E H I F}{ }^{+}$ | 4,762 | 2,748 | 413 | 4,586 | 11,447 | 534 | 13,769 |
| galen-undoctored | 00032 | $\mathcal{A L E H I F}{ }^{+}$ | 5,005 | 2,748 | 413 | 4,828 | 11,767 | 535 | 14,158 |
| go | 00764 | $\mathcal{S R I F}$ | 76,286 | 36,495 | 13 | 76,279 | 82,191 | 22 | 123,506 |
| worm-phenotype-xp | 00675 | SHIF | 90,522 | 51,413 | 113 | 90,497 | 74,594 | 95 | 100,966 |
| aeo | 00002 | $\mathcal{S H I F}(\mathcal{D})$ | 20,317 | 759 | 47 | 3,372 | 1,487 | 92 | 5,598 |
| mammalian-phenotype-xp | 00518 | $\mathcal{S R}$ | 6,684 | 9,809 | 24 | 6,680 | 32,111 | 48 | 38,994 |
| nci-thesaurus | 00554 | $\mathcal{S H}(\mathcal{D})$ | 127,739 | 91,226 | 123 | 127,729 | 157,186 | 244 | 226,658 |
| nif-cell | 00352 | $\mathcal{S R O I F}(\mathcal{D})$ | 3,482 | 2,770 | 61 | 3,431 | 3,212 | 62 | 4,366 |
| pharmacogenomics | 00566 | $\mathcal{S H O I N}(\mathcal{D})$ | 52,813 | 45,000 | 201 | 52,616 | 48,056 | 279 | 53,797 |
| protein | 00791 | $\mathcal{S}$ | 37,266 | 35,196 | 7 | 37,266 | 46,240 | 10 | 61,424 |
| sao | 00624 | $\mathcal{S H I N}(\mathcal{D})$ | 2,802 | 764 | 36 | 2,608 | 1,231 | 58 | 4,595 |
| sct-sep | 00778 | SH | 54,978 | 54,974 | 9 | 54,977 | 333,067 | 13 | 503,379 |
| software | 00636 | $\mathcal{S H O I Q}(\mathcal{D})$ | 2,661 | 903 | 27 | 2,501 | 1,721 | 31 | 3,064 |
| vaccine | 00668 | $\mathcal{S R I Q}$ | 12,516 | 5,410 | 137 | 12,470 | 12,412 | 201 | 20,730 |
| wine | 00783 | $\mathcal{S H O I N}(\mathcal{D})$ | 395 | 98 | 17 | 159 | 302 | 30 | 388 |
| bams-simplified | 00004 | $\mathcal{S H I F}$ | 18,822 | 1,110 | 12 | 18,818 | 4,413 | 14 | 19,575 |
| iedm | 00293 | $\mathcal{A L U N}(\mathcal{D})$ | 4,310 | 395 | 888 | 1,877 | 4,821 | 1,771 | 4,532 |
| molecular-function-xp | 00533 | $\mathcal{S R I F}$ | 84,408 | 42,165 | 82 | 84,368 | 73,981 | 20 | 106,124 |
| neuron-ontology | 00006 | $\mathcal{A L E H I}{ }^{+}$ | 1,260 | 1,119 | 25 | 1,245 | 1,419 | 1 | 1,208 |
| sequence | 00627 | $\mathcal{S H I}$ | 2,812 | 2,136 | 24 | 3,029 | 3,179 | 26 | 4,074 |
| world-fact-book | 00104 | $\mathcal{A L C H O I}(\mathcal{D})$ | 1,158 | 1,080 | 4 | 1,122 | 1,090 | 1 | 1,120 |

Table 7: Decomposition width and length for test ontologies

|  | cautious |  |  | cautious |  |  |  | eager |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :---: |
| Ontology | $\ln \left(\mathcal{D}_{1}\right)$ | $\operatorname{wd}\left(\mathcal{D}_{1}\right)$ | $\operatorname{wd}\left(\mathcal{D}_{1}^{\prime}\right)$ | $\ln \left(\mathcal{D}_{2}\right)$ | $\operatorname{wd}\left(\mathcal{D}_{2}\right)$ | $\operatorname{wd}\left(\mathcal{D}_{2}^{\prime}\right)$ | $\ln \left(\mathcal{D}_{3}\right)$ | $\operatorname{wd}\left(\mathcal{D}_{3}\right)$ | $\operatorname{wd}\left(\mathcal{D}_{3}^{\prime}\right)$ |  |

(a) $\operatorname{wd}\left(\mathcal{D}_{1}\right), \operatorname{wd}\left(\mathcal{D}_{2}\right)$, and $\operatorname{wd}\left(\mathcal{D}_{3}\right)$ are all different:

| acgt | 1,754 | 235 | 11 | 1,883 | 151 | 9 | 1,851 | 162 | 9 |
| :--- | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: |
| fma-cons | 41,646 | 418 | 22 | 80,606 | 389 | 22 | 84,647 | 387 | 22 |
| go-anatomy-importer | 58,192 | 12,909 | $?$ | 79,533 | 12,582 | $?$ | 94,915 | 491 | 28 |
| nif-gross | 4,069 | 686 | 178 | 4,653 | 589 | 152 | $>5 \cdot 10^{6}$ | $?$ | $?$ |
| uberon | 7,383 | 4,931 | 369 | 10,966 | 4,600 | 115 | 10,241 | 199 | 18 |
| envo-xp | 35,702 | 2 | 1 | 50,247 | 1 | 1 | 35,704 | 0 | 0 |
| galen7 | 30,472 | 3,850 | 47 | 43,521 | 3,833 | 47 | 49,344 | 0 | 0 |
| go-xp-all | 89,926 | 1,456 | 76 | 107,797 | 1,404 | 49 | 92,801 | 0 | 0 |
| pato | 2,306 | 2 | 1 | 2,331 | 1 | 1 | 2,308 | 0 | 0 |
| plant-trait-xp | 60,968 | 2 | 1 | 84,271 | 1 | 1 | 60,973 | 0 | 0 |

(b) $\operatorname{wd}\left(\mathcal{D}_{1}\right)=\operatorname{wd}\left(\mathcal{D}_{2}\right)>\operatorname{wd}\left(\mathcal{D}_{3}\right)$ and $\operatorname{wd}\left(\mathcal{D}_{1}^{\prime}\right)=\operatorname{wd}\left(\mathcal{D}_{2}^{\prime}\right) \geq \operatorname{wd}\left(\mathcal{D}_{3}^{\prime}\right)$ :

| cdao | 890 | 369 | 90 | 1,203 | 369 | 90 | 1,449 | 82 | 25 |
| :--- | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: |
| dolce-all | 236 | 2,180 | 72 | 336 | 2,180 | 72 | 453 | 1,819 | 62 |
| influenza-ontology | 842 | 554 | 32 | 1,032 | 554 | 32 | 1,060 | 237 | 12 |
| lipid | 860 | 1,010 | 186 | 983 | 1,010 | 186 | 3,561 | 1,008 | 186 |
| obi | 2,826 | 293 | 39 | 3,249 | 293 | 39 | 3,287 | 233 | 17 |
| pizza | 101 | 101 | 34 | 147 | 101 | 34 | 392 | 41 | 10 |
| propreo | 478 | 147 | 22 | 483 | 147 | 22 | 561 | 75 | 13 |
| sweet-space | 1,517 | 60 | 6 | 1,522 | 60 | 6 | 1,540 | 38 | 5 |
| ero | 2,399 | 100 | 11 | 2,479 | 100 | 11 | 2,456 | 99 | 11 |
| gardiner-wafa | 158 | 34 | 6 | 169 | 34 | 6 | 1167 | 33 | 6 |
| sweet-numerics | 1,503 | 44 | 5 | 1,508 | 44 | 5 | 1,518 | 34 | 5 |
| sweet-phenomena | 1,723 | 44 | 5 | 1,728 | 44 | 5 | 1,739 | 43 | 5 |
| fly-anatomy | 7,798 | 5 | 2 | 10,356 | 5 | 2 | 7,873 | 0 | 0 |
| galen-doctored | 3,020 | 139 | 14 | 3,824 | 139 | 14 | 3,364 | 0 | 0 |
| galen-undoctored | 3,042 | 139 | 14 | 4,091 | 139 | 14 | 3,460 | 0 | 0 |
| go | 36,494 | 140 | 16 | 46,813 | 140 | 16 | 36,859 | 0 | 0 |
| worm-phenotype-xp | 51,413 | 6 | 2 | 66,320 | 6 | 2 | 51,421 | 0 | 0 |

(c) $\operatorname{wd}\left(\mathcal{D}_{1}\right)=\operatorname{wd}\left(\mathcal{D}_{2}\right)=\operatorname{wd}\left(\mathcal{D}_{3}\right)$ and $\operatorname{wd}\left(\mathcal{D}_{1}^{\prime}\right)=\operatorname{wd}\left(\mathcal{D}_{2}^{\prime}\right)=\operatorname{wd}\left(\mathcal{D}_{3}^{\prime}\right):$

| aeo | 851 | 32 | 6 | 1,124 | 32 | 6 | 963 | 32 | 6 |
| :--- | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: |
| mammalian-phenotype-xp | 9,809 | 16 | 5 | 13,668 | 16 | 5 | 9,809 | 16 | 5 |
| nci-thesaurus | 91,226 | 48 | 11 | 107,788 | 48 | 11 | 103,977 | 48 | 11 |
| nif-cell | 2,769 | 48 | 7 | 2,796 | 48 | 7 | 2,771 | 48 | 7 |
| pharmacogenomics | 45,015 | 321 | 26 | 46,844 | 321 | 26 | 45,348 | 321 | 26 |
| protein | 35,196 | 22 | 5 | 36,636 | 22 | 5 | 35,196 | 22 | 5 |
| sao | 769 | 108 | 9 | 843 | 108 | 9 | 856 | 108 | 9 |
| sct-sep | 54,973 | 2,602 | 278 | 76,545 | 2,602 | 278 | 54,973 | 2,602 | 278 |
| software | 1,120 | 9 | 5 | 1,422 | 9 | 5 | 1,234 | 9 | 5 |
| vaccine | 5,691 | 1,137 | 246 | 7,087 | 1,137 | 246 | 6,584 | 1,137 | 246 |
| wine | 121 | 122 | 27 | 122 | 122 | 27 | 698 | 122 | 27 |
| bams-simplified | 1,187 | 0 | 0 | 3,424 | 0 | 0 | 3,154 | 0 | 0 |
| iedm | 395 | 0 | 0 | 395 | 0 | 0 | 395 | 0 | 0 |
| molecular-function-xp | 42,165 | 0 | 0 | 60,140 | 0 | 0 | 42,165 | 0 | 0 |
| neuron-ontology | 1,119 | 0 | 0 | 1,419 | 0 | 0 | 1,119 | 0 | 0 |
| sequence | 2,136 | 0 | 0 | 2,446 | 0 | 0 | 2,136 | 0 | 0 |
| world-fact-book | 1,080 | 0 | 0 | 1,080 | 0 | 0 | 1,080 | 0 | 0 |

$\epsilon$-refinement of $\mathcal{D}_{1}$ and $\mathcal{D}_{2}$ produces decompositions of equal width. The refined cautious strategy, however, can still be useful in practice: as we discussed in more detail in Section 3.4.2, contexts introduced by this strategy can have more refined literal orderings, but our framework cannot estimate the effects of a literal ordering on the performance of reasoning.

Particularly interesting are the results shown in the bottom part of Table 7c: all ontologies shown there contain universal restrictions and/or inverse roles, but they "behave" like $\mathcal{E} \mathcal{L}$ ontologies in that there is no interaction between existential and universal restrictions and/or inverse roles; this is reflected by the fact that the cautious and the eager strategy introduce in most cases the same numbers of contexts. The only exception is the "bams-simplified" ontology, on which there is some interaction between existential and universal restrictions; however, even in this case, the cautious strategy can produce a decomposition of zero width. To understand how this can happen, consider the ontology $\{A \sqsubseteq \exists R . B, A \sqsubseteq \forall R . C, B \sqsubseteq C\}$ : with the cautious strategy we get a context $v_{B}$ with core $\left(v_{B}\right)=\{B\}$, whereas with the eager strategy we get a context $v_{\{B, C\}}$ with core $\left(v_{\{B, C\}}\right)=\{B, C\}$; however, in the former case $C$ is a known consequence of $B$, so introducing a context with a "suboptimal" core does not increase decomposition width.

Finally, our experiments show that computing an $\epsilon$-refinement can substantially reduce decomposition width: $\operatorname{wd}\left(\mathcal{D}_{i}^{\prime}\right)$ is substantially lower than $\operatorname{wd}\left(\mathcal{D}_{i}\right)$ in all cases except when $\operatorname{wd}\left(\mathcal{D}_{i}\right)$ is already very small. Moreover, the eager strategy seems to produce decompositions of least width: with the exception of four ontologies ("dolce-all", "lipid", "sct-sep", and "vaccine"), decomposition $\mathcal{D}_{3}^{\prime}$ has width below 30, and often much less than 30 . This suggests that our approach to reducing or-branching from Section 5 could improve the performance of consequence-based reasoners in practice.

We finish this section with the observation that, in all cases, decomposition width is substantially smaller than the number of literals in the ontology, so our results provide a tighter estimate of the algorithm's complexity than that obtained from the usual complexity arguments. This is particularly true for Horn ontologies: decomposition length on such ontologies is manageable and decomposition width is zero, which explains why such ontologies are generally easier to reason with. On non-Horn ontologies, however, the upper complexity bound obtained using our results may still be very large; for example, for a decomposition of width 30 , the number of clauses derived in each context is bounded by $4^{30} \approx 1.15 \cdot 10^{18}$-clearly, no algorithm deriving that many clauses can be deemed practical. This bound, however, only analyzes what could happen in the worst case and, similarly to propositional resolution, such worst cases are unlikely to occur in practice. We observed that decomposition widths of our test ontologies are mainly determined by universal restrictions involving inverse roles: in order to guarantee that all clauses needed to apply the Pred rule are derived, sets $\varrho_{v, u}^{\exists R . A}$ are needed in condition (R7) of Definition 42, and these sets can be large. In our experience, however, clauses with many literals of the form $\operatorname{\forall inv}(R) . C$ are rarely derived during reasoning, which is why most ontologies can still be handled in practice. Therefore, a small decomposition width provides an indication, but not a definite guarantee, that reasoning with a particular ontology is easy. In contrast, decomposition length measures directly the number of contexts constructed in the consequence-based algorithm, so length is generally a good measure of ontology hardness.

## 9. Conclusion

In this paper we presented a very general framework for consequence-based reasoning in description logics. Then, we introduced a notion of decomposition that allows us to quantify the effects of and- and or-branching during reasoning. Finally, we presented what we believe to be the first result on fixed-parameter tractability of description logic reasoning. We see several theoretical and practical challenges for our future work.

On the theoretical side, we will try to extend our results to logics with number restrictions and nominals and thus obtain a consequence-based algorithm that can handle all of OWL 2 DL. Number restrictions have already been successfully integrated into resolution-based procedures so, while not trivial, incorporating them into consequence-based algorithms should be possible. We see two main challenges towards that goal. First, we believe that the notion of a context core will have to include two sets of literals (representing the types of an individual and its predecessor in a model) and a set of roles (describing the connections between the two individuals); the reason for this is analogous to why the hypertableau algorithm must use pairwise instead of single blocking [14] when an ontology contains both number restrictions and inverse roles. Second, we believe that our clauses will have to represent disjunctive properties of an individual's successors, which one can likely achieve using a skolemized clausal form as in first-order theorem proving [31]. In contrast, the combination of nominals, number restrictions, and inverse roles is a significant source of complexity, and it is currently unclear whether and how this combination of constructs can be handled in a consequence-based framework.

On the practical side, our results suggest several significant improvements to existing consequence-based reasoners. First, context reuse may allow reasoners to handle ontologies on which the eager strategy introduces a large number of contexts. Second, redundancy elimination techniques may turn out to further reduce the search space. Third, although computing tree decompositions in advance may incur significant overhead, we believe that computing $\epsilon$-refinements "on demand" for contexts with large widths has the potential to further reduce the number of clauses derived. We will incorporate these ideas into a consequence-based reasoner and try to determine the extent to which they optimize the performance of reasoning in practice.

## Appendix A. Proof of Theorem 6

Throughout Appendix A, we fix a normalized $\mathcal{A L C \mathcal { L }}$ ontology $\mathcal{O}$, an admissible context structure $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, $\langle \rangle$, and a clause system $\mathcal{S}$ for $\mathcal{D}$ such that no inference rule from Table 3 is applicable to $\mathcal{S}$; furthermore, we fix $\mathbf{L}$ to be the set of all literals that occur in $\mathcal{O}, \mathcal{D}$, or $\mathcal{S}$. We prove Theorem 6 by showing the following contrapositive claim: $\mathcal{O} \not \vDash K \sqsubseteq M$ holds for each query $K \sqsubseteq M$ for which a context $v \in \mathcal{V}$ exists such that

- $v$ is complete for $K \sqsubseteq M$,
- $K \sqsubseteq L \hat{\epsilon} \mathcal{S}(v)$ for each $L \in K$, and
- $K \sqsubseteq M \ddot{\notin} \mathcal{S}(v)$.

To this end, we construct a model $\mathcal{I}$ of $\mathcal{O}$ that refutes each query satisfying the above conditions. The construction is organized as follows. Instead of directly constructing a model, for simplicity we construct a pre-model-a graph-like structure with nodes labeled by sets of literals. We introduce this notion in Appendix A.1, and we explain how to convert a pre-model into a model. In Appendix A. 2 we show how to construct a propositional interpretation satisfying a set of clauses closed under hyperresolution. In Appendix A. 3 we prove certain properties of our inference rules that will allow us to use the construction from Appendix A.2. Finally, in Appendix A. 4 we use these propositional interpretations to construct the desired pre-model of $\mathcal{O}$.

## Appendix A.1. Pre-Interpretations and Pre-Models

In this section we introduce the notions of a pre-interpretation and a pre-model of an ontology, which we use to simplify the presentation of the subsequent proofs.

Definition 59. A set of literals $J$ satisfies a clause $K \sqsubseteq M$, written $J \vDash K \sqsubseteq M$, if $K \subseteq J$ implies $M \cap J \neq \emptyset$; otherwise, $J$ refutes $K \sqsubseteq M$, written $J \not \vDash K \sqsubseteq M$.

A pre-interpretation is a labeled graph $I=\langle\Delta, E, J\rangle$ where $\Delta$ is a nonempty set of nodes, $E \subseteq \Delta \times \Delta \times \Sigma_{R}$ is a set of rolelabeled edges, and $J: \Delta \rightarrow 2^{\Sigma_{L}}$ is a labeling of nodes by sets of literals satisfying the following two conditions.
( $\left.I_{\exists}\right)$ For each node $x \in \Delta$ and each literal $\exists R . A \in J(x)$, an edge $\langle x, y, R\rangle \in E$ exists such that $A \in J(y)$.
( $\left.I_{\forall}\right)$ For each edge $\langle x, y, R\rangle \in E$ and all literals $\forall R . B$ and $\forall \operatorname{inv}(R) . C$, we have that

- $\forall R . B \in J(x)$ implies $B \in J(y)$, and
- $\forall \operatorname{inv}(R) . C \in J(y)$ implies $C \in J(x)$.

Pre-interpretation $I=\langle\Delta, E, J\rangle$ satisfies a clause $K \sqsubseteq M$, written $I \models K \sqsubseteq M$, if $J(x) \vDash K \sqsubseteq M$ for each node $x \in \Delta$; otherwise, $I$ refutes $K \sqsubseteq M$, written $I \not \vDash K \sqsubseteq M$. Finally, I is a pre-model of $\mathcal{O}$ if I satisfies each (normal) clause in $\mathcal{O}$.

Given a pre-interpretation $I=\langle\Delta, E, J\rangle$, we can construct an interpretation $\mathcal{I}$ for all atomic concepts $A \in \Sigma_{A}$ and all atomic roles $T \in \Sigma_{T}$ as follows.

$$
\begin{align*}
& \Delta^{\mathcal{I}}:=\Delta  \tag{A.1}\\
& A^{\mathcal{I}}:=\{x \mid A \in J(x)\}  \tag{A.2}\\
& T^{\mathcal{I}}:=\{\langle x, y\rangle \mid\langle x, y, T\rangle \in E\} \cup\left\{\langle y, x\rangle \mid\left\langle x, y, T^{-}\right\rangle \in E\right\} \tag{A.3}
\end{align*}
$$

Interpretation $\mathcal{I}$ satisfies the following property for each literal $L \in \Sigma_{L}$ and each node $x \in \Delta$ :

$$
\begin{equation*}
L \in J(x) \quad \text { implies } \quad x \in L^{\mathcal{I}} . \tag{A.4}
\end{equation*}
$$

For atomic concepts, property (A.4) holds by the definition of $A^{\mathcal{I}}$; furthermore, for literals of the form $\exists R . A$ and $\forall R . A$, property (A.4) holds due to conditions ( $I_{\exists}$ ) and ( $I_{y}$ ), respectively. The converse of property (A.4), however, holds only for atomic concepts, which is why all clauses in $\mathcal{O}$ must have only atomic concepts in their antecedents, and all queries must have only atomic concepts in their consequents. We next show that the interpretation $\mathcal{I}$ obtained as specified above is indeed a model of $\mathcal{O}$.
Lemma 60. Let $I=\langle\Delta, E, J\rangle$ be a pre-interpretation, and let $\mathcal{I}$ be the interpretation obtained from $I$ as specified in (A.1)-(A.3). If I satisfies a normal clause $\alpha$, then $\mathcal{I} \vDash \alpha$. Furthermore, if I refutes a query $q$, then $\mathcal{I} \not \vDash q$.
Proof. Let $\prod_{i=1}^{m} A_{i} \sqsubseteq \bigsqcup_{j=1}^{n} L_{j}$ be an arbitrary normal clause that is satisfied in $I$, and consider an arbitrary node $x \in \Delta$ such that $x \in A_{i}^{\mathcal{I}}$ for each $1 \leq i \leq m$. By (A.2), we have $A_{i} \in J(x)$ for each $1 \leq i \leq m$; but then, since $I$ satisfies the clause by the assumption, we have $L_{j} \in J(x)$ for some $1 \leq j \leq n$; hence, we have $x \in L_{j}^{\mathcal{I}}$ for some $1 \leq j \leq n$ by (A.4). Since this holds for arbitrary $x \in \Delta$, we have $\mathcal{I} \vDash \prod_{i=1}^{m} A_{i} \sqsubseteq \bigsqcup_{j=1}^{n} L_{j}$.

Let $\prod_{i=1}^{m} L_{i} \sqsubseteq \bigsqcup_{j=1}^{n} A_{j}$ be an arbitrary query that is refuted in $I$. Thus, a node $x \in \Delta$ exists such that $L_{i} \in J(x)$ for each $1 \leq i \leq m$, and $A_{j} \notin J(x)$ for each $1 \leq j \leq n$. But then, $x \in L_{i}^{\mathcal{I}}$ for each $1 \leq i \leq m$ by (A.4), and $x \notin A_{j}^{\mathcal{I}}$ for each $1 \leq j \leq n$ by (A.2); consequently, we have $\mathcal{I} \not \vDash \prod_{i=1}^{m} L_{i} \sqsubseteq \bigsqcup_{j=1}^{n} A_{j}$.

Corollary 61. If a pre-model I of $\mathcal{O}$ exists that refutes a query $K \sqsubseteq M$, then $\mathcal{O} \not \models K \sqsubseteq M$.

## Appendix A.2. Constructing Literal Interpretations

In this section, we show how the sets of clauses $\mathcal{S}(v)$ and the corresponding literal orderings $<_{v}$ can be used to construct the sets $J(x)$ of a pre-model $I=\langle\Delta, E, J\rangle$ that satisfies $\mathcal{O}$ but refutes a query $K \sqsubseteq M$. Our construction is independent from the selected context $v$; therefore, we consider here just a single set of clauses $\mathcal{N}$.

Throughout Appendix A. 2 we fix a literal ordering $<$, a set of clauses $\mathcal{N}$ over $\mathbf{L}$, and a clause $K \sqsubseteq M$ over $\mathbf{L}$ satisfying the following conditions:

$$
\begin{align*}
& M \text { is <-minimal, }  \tag{A.5}\\
& K \sqsubseteq L \hat{\in} \mathcal{N} \text { for each literal } L \in K \text {, and }  \tag{A.6}\\
& K \sqsubseteq M \nsubseteq \mathcal{N} . \tag{A.7}
\end{align*}
$$

Next, in Definition 62 we introduce a notion that determines when set $\mathcal{N}$ contains "sufficiently many" hyperresolution consequences w.r.t. literal ordering $<$.

Definition 62. Set $\mathcal{N}$ is closed under <-hyperresolution with a clause $\prod_{i=1}^{n} L_{i} \sqsubseteq M^{\prime}$ (not necessarily contained in $\mathcal{N}$ ) if, for each choice of $n$ clauses $K_{i} \sqsubseteq M_{i} \sqcup L_{i} \in \mathcal{N}, 1 \leq i \leq n$, each satisfying $L_{i} \nprec M_{i}$, we have $\prod_{i=1}^{n} K_{i} \sqsubseteq M^{\prime} \sqcup \bigsqcup_{i=1}^{n} M_{i} \hat{\in} \mathcal{N}$.

The following lemma follows trivially from Definition 62 and the notion of clause strengthening.
Lemma 63. Let $\alpha_{1}$ and $\alpha_{2}$ be arbitrary clauses such that $\alpha_{1}$ is a strengthening of $\alpha_{2}$. If $\mathcal{N}$ is closed under <-hyperresolution with $\alpha_{1}$, then $\mathcal{N}$ is also closed under <-hyperresolution with $\alpha_{2}$.

For $<, \mathcal{N}$, and $K \sqsubseteq M$ fixed as specified above, we next construct a set of literals $J$ that refutes $K \sqsubseteq M$, but that satisfies each clause $\prod_{i=1}^{n} L_{i} \sqsubseteq M^{\prime}$ for which $\mathcal{N}$ is closed under <-hyperresolution. We achieve this by essentially adapting the standard techniques from resolution theorem proving [31] to our setting. Towards this goal, we proceed as follows: we first present the construction of $J$; in Lemma 64 we prove certain properties of the productive clauses in $\mathcal{N}$; in Lemma 65 we show that $J$ satisfies each clause $K^{\prime} \sqsubseteq M^{\prime}$ for which a strengthening is contained in $\mathcal{N}$ and that satisfies $K^{\prime} \subseteq K$; in Lemma 66 we show that $J$ refutes $K \sqsubseteq M$; and finally in Lemma 67 we prove the desired result.

Let $L_{1}, L_{2}, \ldots, L_{\ell}$ be an arbitrary total ordering of $\mathbf{L}$ that extends $<$ such that the elements of $M$ precede all the remaining literals from $\mathbf{L}$-that is, for all $i$ and $j$, we have that

$$
\begin{align*}
& L_{i}<L_{j} \text { implies } i<j, \text { and }  \tag{A.8}\\
& i<j \text { and } L_{j} \in M \text { imply } L_{i} \in M . \tag{A.9}
\end{align*}
$$

Since $M$ is <-minimal by (A.5) and set $\mathbf{L}$ is finite, at least one such total ordering exists. For each $0 \leq k \leq \ell$, let $\mathbf{L}_{k}:=\left\{L_{1}, \ldots, L_{k}\right\}$, and let $J_{k}$ be an inductively defined set of literals such that $J_{0}:=\emptyset$ and, for each $0<k \leq \ell$, we have

$$
J_{k}:= \begin{cases}J_{k-1} \cup\left\{L_{k}\right\} & \text { if a clause } K^{\prime} \sqsubseteq M^{\prime} \sqcup L_{k} \in \mathcal{N} \text { exists such that } K^{\prime} \subseteq K, M^{\prime} \cap J_{k-1}=\emptyset, \text { and } M^{\prime} \subseteq \mathbf{L}_{k-1}  \tag{A.10}\\ J_{k-1} & \text { otherwise. }\end{cases}
$$

Let $J:=J_{\ell}$. Each set $J$ that can be obtained by such a construction (which is nondeterministic due to the choice of the total ordering) is called a literal interpretation of $\mathbf{L}$ w.r.t. $\mathcal{N}$ and $<$ refuting $K \sqsubseteq M$. Each clause $K^{\prime} \sqsubseteq M^{\prime} \sqcup L_{k} \in \mathcal{N}$ that satisfies the first condition in (A.10) for some $k$ is said to be a productive clause, and the clause is said to produce the literal $L_{k}$ in $J$. Clearly, each literal in $J$ is produced by at least one productive clause in $\mathcal{N}$.

Lemma 64. Each productive clause $K^{\prime} \sqsubseteq M^{\prime} \sqcup L_{k} \in \mathcal{N}$ satisfies $K^{\prime} \subseteq K, M^{\prime} \cap J=\emptyset$, and $L_{k} \nprec M^{\prime}$.
Proof. Let $K^{\prime} \sqsubseteq M^{\prime} \sqcup L_{k} \in \mathcal{N}$ be an arbitrary clause producing literal $L_{k}$ in $J$ for some integer $k$. By (A.10), we have $K^{\prime} \subseteq K$, $M^{\prime} \cap J_{k-1}=\emptyset$, and $M^{\prime} \subseteq \mathbf{L}_{k-1}$. Clearly $J_{k-1}=J \cap \mathbf{L}_{k-1}$, so we have $M^{\prime} \cap J=\emptyset$. Finally, by (A.8), we have $L_{k} \nprec L$ for each $L \in \mathbf{L}_{k-1}$; hence, due to $M^{\prime} \subseteq \mathbf{L}_{k-1}$, we have $L_{k} \nprec M^{\prime}$.
Lemma 65. For each clause $K^{\prime} \sqsubseteq M^{\prime}$ such that $K^{\prime} \sqsubseteq M^{\prime} \hat{\in} \mathcal{N}$ and $K^{\prime} \subseteq K$, we have $M^{\prime} \cap J \neq \emptyset$.
Proof. Let $K^{\prime} \sqsubseteq M^{\prime}$ be an arbitrary clause satisfying $K^{\prime} \sqsubseteq M^{\prime} \hat{\in} \mathcal{N}$ and $K^{\prime} \subseteq K$. Then, a clause $K_{1} \sqsubseteq M_{1} \in \mathcal{N}$ exists such that $K_{1} \subseteq K^{\prime} \subseteq K$ and $M_{1} \subseteq M^{\prime}$ hold. Clause $K_{1} \sqsubseteq \perp$ is thus a strengthening of $K \sqsubseteq M$, but $K \sqsubseteq M \ddot{\notin \mathcal{N}}$ by assumption (A.7); thus, $M_{1} \neq \emptyset$. Let $k$ be the largest integer for which $L_{k} \in M^{\prime}$ holds; hence, clause $K_{1} \sqsubseteq M_{1}$ is of the form $K_{1} \sqsubseteq M_{2} \sqcup L_{k}$ where $M_{2} \subseteq \mathbf{L}_{k-1}$. The claim of this lemma holds trivially if $M_{2} \cap J \neq \emptyset$. Furthermore, if $M_{2} \cap J=\emptyset$, then by (A.10) clause $K_{1} \sqsubseteq M_{2} \sqcup L_{k}$ produces $L_{k} \in J$, so we again have $M^{\prime} \cap J \neq \emptyset$, as required.

Lemma 66. We have $J \not \vDash K \sqsubseteq M$-that is, $K \subseteq J$ and $M \cap J=\emptyset$.

Proof. ( $K \subseteq J$ ) Consider an arbitrary literal $L \in K$. Then $K \subseteq L \hat{\in} \mathcal{N}$ by assumption (A.6), so a conjunction $K^{\prime} \subseteq K$ exists such that $K^{\prime} \sqsubseteq \perp \in \mathcal{N}$ or $K^{\prime} \sqsubseteq L \in \mathcal{N}$. The former contradicts $K \sqsubseteq M \ddot{\not} \mathcal{N}$, which is assumption (A.7); hence, we have $K^{\prime} \sqsubseteq L \in \mathcal{N}$, so by (A.10) this clause produces $L \in J$.
$(M \cap J=\emptyset)$ Assume that a literal $L_{k} \in M \cap J$ exists, and let $K^{\prime} \sqsubseteq M^{\prime} \sqcup L_{k} \in \mathcal{N}$ be a clause that produces $L_{k} \in J$. By (A.10), we have $M^{\prime} \subseteq \mathbf{L}_{k-1}$ and $K^{\prime} \subseteq K$; furthermore, by (A.9) and $L_{k} \in M$, we have $M^{\prime} \cup\left\{L_{k}\right\} \subseteq M$. But then, $K^{\prime} \sqsubseteq M^{\prime} \sqcup L_{k} \in \mathcal{N}$ contradicts $K \sqsubseteq M \notin \mathcal{N}$, which is assumption (A.7).
Lemma 67. For each clause $\alpha$ such that $\mathcal{N}$ is closed under <-hyperresolution with $\alpha$, we have $J \vDash \alpha$.
Proof. Assume that $\alpha=\prod_{i=1}^{n} L_{i} \sqsubseteq M^{\prime}$ and that $L_{i} \in J$ for each $1 \leq i \leq n$. For each $1 \leq i \leq n$, let $K_{i} \sqsubseteq M_{i} \sqcup L_{i}$ be some clause that produces $L_{i} \in J$; then $K_{i} \subseteq K, M_{i} \cap J=\emptyset$, and $L_{i} \nprec M_{i}$ by Lemma 64. Since $\mathcal{N}$ is closed under <-hyperresolution with $\alpha$, we have $\prod_{i=1}^{n} K_{i} \sqsubseteq M^{\prime} \sqcup \bigsqcup_{i=1}^{n} M_{i} \hat{\in} \mathcal{N}$. By $\left(\prod_{i=1}^{n} K_{i}\right) \subseteq K$ and Lemma 65, we have $\left(M^{\prime} \sqcup \bigsqcup_{i=1}^{n} M_{i}\right) \cap J \neq \emptyset$. Finally, since $M_{i} \cap J=\emptyset$ holds for each $1 \leq i \leq n$, we have $M^{\prime} \cap J \neq \emptyset$. Consequently, we have $J \vDash \alpha$, as required.

## Appendix A.3. Properties of the Consequence-Based Inference Rules

To apply the construction from Appendix A. 2 to our clause system $\mathcal{S}$ and the context structure $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, < $\rangle$ (which were fixed at the beginning of Appendix A), we next prove the following two auxiliary lemmas about the consequence-based inference rules. Lemma 68 allows us to show that, for each context $v \in \mathcal{V}$, the set of literals obtained by applying the above construction to $\mathcal{S}(v)$ satisfies each clause in $\mathcal{O}$; the lemma follows trivially from Definition 62 and the fact that the Hyper rule is not applicable to $\mathcal{S}$. Furthermore, Lemma 69 shows a similar property for the clauses that participate in the Pred rule.

Lemma 68. For each context $v \in \mathcal{V}$ and each clause $\alpha \in \mathcal{O}$, set $\mathcal{S}(v)$ is closed under $<_{\nu}$-hyperresolution with $\alpha$.
Lemma 69. For each edge $\langle v, u, \exists R . A\rangle \in \mathcal{E}$ and each clause $\alpha=A \sqcap \prod_{i=1}^{n} B_{i} \sqsubseteq \bigsqcup_{j=1}^{m} \forall \operatorname{inv}(R) . C_{j}$ such that $\alpha \hat{\in} \mathcal{S}(u)$, set $\mathcal{S}(v)$ is closed under $<_{\nu}$-hyperresolution with $\exists R . A \sqcap \prod_{i=1}^{n} \forall R . B_{i} \sqsubseteq \bigsqcup_{j=1}^{m} C_{j}$.
Proof. Consider an arbitrary edge $\langle v, u, \exists R . A\rangle \in \mathcal{E}$ and an arbitrary clause $\alpha=A \sqcap \prod_{i=1}^{n} B_{i} \sqsubseteq \bigsqcup_{j=1}^{m} \forall \operatorname{inv}(R) . C_{j}$ such that $\alpha \hat{\in} \mathcal{S}(u)$. By the definition of $\hat{\epsilon}$, a subset $\left\{B_{i}^{\prime}\right\}_{i=1}^{n^{\prime}}$ of $\left\{B_{i}\right\}_{i=1}^{n}$ and a subset $\left\{C_{j}^{\prime}\right\}_{j=1}^{m^{\prime}}$ of $\left\{C_{j}\right\}_{j=1}^{m}$ exist such that

$$
A \sqcap \prod_{i=1}^{n^{\prime}} B_{i}^{\prime} \sqsubseteq \bigsqcup_{j=1}^{m^{\prime}} \forall \operatorname{inv}(R) \cdot C_{j}^{\prime} \in \mathcal{S}(u) \quad \text { or } \quad \prod_{i=1}^{n^{\prime}} B_{i}^{\prime} \sqsubseteq \bigsqcup_{j=1}^{m^{\prime}} \forall \operatorname{inv}(R) . C_{j}^{\prime} \in \mathcal{S}(u) .
$$

The Pred rule is not applicable to clause system $\mathcal{S}$; thus, for each clause $K_{0} \sqsubseteq M_{0} \sqcup \exists R . A \in \mathcal{S}(v)$ satisfying $\exists R . A \nprec_{v} M_{0}$, and for each choice of $n^{\prime}$ clauses $K_{i} \sqsubseteq M_{i} \sqcup \forall R . B_{i}^{\prime}, 1 \leq i \leq n^{\prime}$, satisfying $\forall R . B_{i}^{\prime} \not \kappa_{v} M_{i}$, we have

$$
K_{0} \sqcap \prod_{i=1}^{n^{\prime}} K_{i} \sqsubseteq M_{0} \sqcup \bigsqcup_{i=1}^{n^{\prime}} M_{i} \sqcup \bigsqcup_{j=1}^{m^{\prime}} C_{j}^{\prime} \hat{\in} \mathcal{S}(v) .
$$

Thus, $\mathcal{S}(v)$ is closed under $<_{\nu}$-hyperresolution with clause $\exists R . A \sqcap \prod_{i=1}^{n^{\prime}} \forall R . B_{i}^{\prime} \sqsubseteq \bigsqcup_{j=1}^{m^{\prime}} C_{j}^{\prime}$. By Lemma 63, then $\mathcal{S}(v)$ is also closed under $<_{\nu}$-hyperresolution with clause $\left.\exists R . A \sqcap\right\rceil_{i=1}^{n} \forall R . B_{i} \sqsubseteq \bigsqcup_{j=1}^{m} C_{j}$ since the former is a strengthening of the latter clause.

## Appendix A.4. The Completeness Claim

We now complete the proof of Theorem 6 by constructing a pre-interpretation that satisfies $\mathcal{O}$ but refutes the relevant queries. To simplify the presentation, given a context $v \in \mathcal{V}$ and a clause $K \sqsubseteq M$, we write $v \nvdash K \sqsubseteq M$ if and only if

- $M$ is $<_{v}$-minimal (i.e., $v$ is complete for $K \sqsubseteq M$ ),
- $K \sqsubseteq L \hat{\in} \mathcal{S}(v)$ for each literal $L \in K$, and
- $K \sqsubseteq M \notin \mathcal{S}(v)$.

Claim 70. For each context $v \in \mathcal{V}$ and each clause $K \sqsubseteq M$, if $K \sqsubseteq L \hat{\in} \mathcal{S}(v)$ for each $L \in K$ and $K \sqsubseteq M \notin \mathcal{S}(v)$, then $K \subseteq \mathbf{L}$.
Proof. Assume that there exists a literal $L \in K \backslash \mathbf{L}$; then, $L$ does not occur in $\mathcal{S}(v)$, so $K \sqsubseteq L \hat{\in} \mathcal{S}(v)$ implies $K \sqsubseteq \perp \hat{\epsilon} \mathcal{S}(v)$, which contradicts the assumption that $K \sqsubseteq M \notin \mathcal{S}(v)$. Thus, $K \subseteq \mathbf{L}$ holds.

Theorem 6 holds vacuously if no context $v \in \mathcal{V}$ and no query $K \sqsubseteq M$ exist such that $v \nvdash K \sqsubseteq M$. Thus, we assume that $v \nvdash K \sqsubseteq M$ holds for at least one context $v \in \mathcal{V}$ and at least one query $K \sqsubseteq M$; by the definition of clause strengthening, we then also have $v \nvdash K \sqsubseteq M^{\prime}$ for $M^{\prime}=M \cap \mathbf{L}$; the latter clause is over $\mathbf{L}$ by Claim 70. Let $I=\langle\Delta, E, J\rangle$ be defined as follows.

- Set $\Delta$ is the smallest set that contains a distinguished element $\delta_{K \sqsubseteq M}^{v}$ for each context $v \in \mathcal{V}$ and each clause $K \sqsubseteq M$ over $\mathbf{L}$ such that $v \nvdash K \sqsubseteq M$. Set $\Delta$ is not empty due to the above assumption.
- For each $\delta_{K \sqsubseteq M}^{v} \in \Delta$, let $J\left(\delta_{K \sqsubseteq M}^{v}\right)$ be an arbitrary literal interpretation of $\mathbf{L}$ w.r.t. $\mathcal{S}(v)$ and $<_{\nu}$ refuting $K \sqsubseteq M$.
- Set $E$ is the smallest set containing $\left\langle\delta_{K_{1} \sqsubseteq M_{1}}^{v}, \delta_{A \sqcap K_{2} \sqsubseteq M_{2}}^{u}, R\right\rangle$ for each element $\delta_{K_{1} \sqsubseteq M_{1}}^{v} \in \Delta$, each literal $\exists R . A \in J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right)$, and each edge $\langle v, u, \exists R . A\rangle \in \mathcal{E}$ such that $u \nvdash A \sqcap K_{2} \sqsubseteq M_{2}$, where $K_{2}$ and $M_{2}$ are as specified below.

$$
\begin{align*}
& \left.K_{2}=\right\rceil\left\{B \mid \forall R . B \in J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right)\right\}  \tag{A.11}\\
& M_{2}=\bigsqcup\left\{\forall \operatorname{inv}(R) . C \in \mathbf{L} \mid C \notin J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right)\right\} \tag{A.12}
\end{align*}
$$

In the last item, we have $K_{2} \cup M_{2} \subseteq \mathbf{L}$ and $A \in \mathbf{L}$; thus, $u \nvdash A \sqcap K_{2} \sqsubseteq M_{2}$ ensures $\delta_{A \sqcap K_{2} \sqsubseteq M_{2}}^{u} \in \Delta$, and so $E$ is correctly defined. Next, in Claim 71 we show that $I$ is a pre-interpretation, in Claim 72 we show that $I$ is a pre-model of $\mathcal{O}$, and in Claim 73 we show that $I$ refutes the relevant queries. The claim of Theorem 6 then follows from Claims 72 and 73, and Corollary 61.

Claim 71. Structure I satisfies conditions $\left(I_{\exists}\right)$ and $\left(I_{\Downarrow}\right)$ of pre-interpretations from Definition 59.
Proof. (Property $I_{\exists}$ ) Consider an arbitrary element $\delta_{K_{1} \sqsubseteq M_{1}}^{v} \in \Delta$ and an arbitrary literal $\exists R . A \in J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right)$; we next show that an edge $\left\langle\delta_{K_{1} \sqsubseteq M_{1}}^{v}, \gamma, R\right\rangle \in E$ exists such that $A \in J(\gamma)$. The Succ rule is not applicable to $\mathcal{S}$, so an edge $\langle v, u, \exists R . A\rangle \in \mathcal{E}$ exists satisfying

$$
\begin{equation*}
L \sqsubseteq L \hat{\in} \mathcal{S}(u) \text { for each } L \in\{A\} \cup \mathbf{B}_{p} \text {, where } \mathbf{B}_{p}=\left\{B \mid K^{\prime} \sqsubseteq M^{\prime} \sqcup \forall R . B \in S(v) \text { and } \forall R . B \nprec_{v} M^{\prime}\right\} \text {. } \tag{A.13}
\end{equation*}
$$

Let $K_{2}$ and $M_{2}$ be as specified in (A.11) and (A.12). The following observations prove that $u \nvdash A \sqcap K_{2} \sqsubseteq M_{2}$.

- Admissibility of $\mathcal{D}$ and $\langle v, u, \exists R . A\rangle \in \mathcal{E}$ imply, by Definition 4, that $<_{u}$ is $R$-admissible, so $M_{2}$ is $<_{u}$-minimal.
- Consider an arbitrary literal $L \in A \sqcap K_{2}$; we show that $A \sqcap K_{2} \sqsubseteq L \hat{\in} \mathcal{S}(u)$. Assume that $L \in K_{2}$, so $L$ is an atomic concept $L=B$. Then, by (A.11), we have $\forall R . B \in J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right)$; thus, there exists at least one clause $K^{\prime} \sqsubseteq M^{\prime} \sqcup \forall R . B \in S(v)$ with $\forall R . B \nprec_{v} M^{\prime}$ that produces $\forall R . B$ in $J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right)$; hence $B \in \mathbf{B}_{p}$. But then $L \sqsubseteq L \hat{\in} \mathcal{S}(u)$ holds by (A.13), which is stronger than the required $A \sqcap K_{2} \sqsubseteq L \hat{\in} \mathcal{S}(u)$. The case when $L=A$ follows directly from (A.13).
- Assume for contradiction that $A \sqcap K_{2} \sqsubseteq M_{2} \hat{\epsilon} \mathcal{S}(u)$. By Lemma 69 , set $\mathcal{S}(v)$ is closed under $<_{\nu}$-hyperresolution with clause

$$
\exists R . A \sqcap\rceil\left\{\forall R . B \mid \forall R . B \in J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right)\right\} \sqsubseteq\left\lfloor\left\{C \mid \forall \operatorname{inv}(R) . C \in \mathbf{L} \text { and } C \notin J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right)\right\} .\right.
$$

Recall that $\exists R . A \in J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right)$, so the above clause is clearly refuted in $J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right)$. However, by Lemma 67 , the clause is satisfied in $J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right)$, which is a contradiction. Consequently, we have $A \sqcap K_{2} \sqsubseteq M_{2} \notin \mathcal{S}(u)$, as required.
Thus, we have $u \nvdash A \sqcap K_{2} \sqsubseteq M_{2}$, which implies $\left\langle\delta_{K_{1} \sqsubseteq M_{1}}^{v}, \delta_{A \sqcap K_{2} \sqsubseteq M_{2}}^{u}, R\right\rangle \in E$ by the definition of $E$. Finally, $A \in J\left(\delta_{A \sqcap K_{2} \sqsubseteq M_{2}}^{u}\right)$ holds by Lemma 66. All of these observations prove that $I$ satisfies property ( $I_{\exists}$ ) of Definition 59.
(Property $I_{Y}$ ) Consider an arbitrary edge $\left\langle\delta_{K_{1} \sqsubseteq M_{1}}^{v}, \delta_{A \sqcap K_{2} \sqsubseteq M_{2}}^{u}, R\right\rangle \in E$ as in the definition of $E$, as well as arbitrary literals $\forall R . B$ and $\forall \operatorname{inv}(R) . C$.

- Assume that $\forall R . B \in J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right)$. Then $B \in K_{2}$ holds by (A.11); but then, Lemma 66 implies $J\left(\delta_{A \sqcap K_{2} \sqsubseteq M_{2}}^{u}\right) \not \vDash A \sqcap K_{2} \sqsubseteq M_{2}$. Consequently, we have $B \in J\left(\delta_{A \sqcap K_{2} \sqsubseteq M_{2}}^{u}\right)$, as required.
- Assume that $\forall \operatorname{inv}(R) . C \in J\left(\delta_{A \sqcap K_{2} \sqsubseteq M_{2}}^{u}\right)$. Then $\forall \operatorname{inv}(R) . C \in \mathbf{L}$ holds since $J\left(\delta_{A \sqcap K_{2} \sqsubseteq M_{2}}^{u}\right)$ is a literal interpretation of $\mathbf{L}$. Now, if $C \notin J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right)$, then $\forall \operatorname{inv}(R) . C \in M_{2}$ by (A.12), so $\forall \operatorname{inv}(R) . C \notin J\left(\delta_{A \sqcap K_{2} \sqsubseteq M_{2}}^{u}\right)$ by Lemma 66, which contradicts our assumption. Consequently, we have $C \in J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right)$, as required.
Claim 72. Pre-interpretation I is a pre-model of $\mathcal{O}$.
Proof. Consider an arbitrary element $\delta_{K \sqsubseteq M}^{v} \in \Delta$ and an arbitrary normal clause $\alpha \in \mathcal{O}$. By Lemma 68, set $\mathcal{S}(v)$ is closed under $<_{\nu}$-hyperresolution with $\alpha$, so $J\left(\delta_{K \subseteq M}^{v}\right) \vDash \alpha$ by Lemma 67. This holds for arbitrary $\delta_{K \subseteq M}^{v} \in \Delta$, so we have $I \vDash \alpha$. Finally, the latter holds for arbitrary $\alpha \in \mathcal{O}$, so $I$ is a pre-model of $\mathcal{O}$.

Claim 73. Pre-interpretation I refutes each query $K \sqsubseteq M$ for which there exists a context $v \in \mathcal{V}$ such that $v$ is complete for $K \sqsubseteq M, K \sqsubseteq L \hat{\in} \mathcal{S}(v)$ for each $L \in K$, and $K \sqsubseteq M \nexists \mathcal{S}(v)$.

Proof. Let $K \sqsubseteq M$ be an arbitrary query satisfying the preconditions of the claim, and let $M^{\prime}=M \cap \mathbf{L}$. Then $v \nvdash K \sqsubseteq M^{\prime}$ holds as well, and this clause is over $\mathbf{L}$ by Claim 70. But then, $J\left(\delta_{K \sqsubseteq M^{\prime}}^{v}\right) \not \vDash K \sqsubseteq M^{\prime}$ by Lemma 66. Finally, each literal from $M \backslash M^{\prime}$ is not from $\mathbf{L}$ and so it cannot occur in $J\left(\delta_{K \sqsubseteq M^{\prime}}^{v}\right)$; thus, $J\left(\delta_{K \sqsubseteq M^{\prime}}^{v}\right) \not \models K \sqsubseteq M$ holds as well, so $I$ refutes $K \sqsubseteq M$, as required.

## Appendix B. Proof of Theorem 37

We proceed similarly as in the proof of Theorem 6 in Appendix A. In the rest of Appendix B, we fix a normalized $\mathcal{A L C L}$ ontology $\mathcal{O}$, an admissible decomposition $\mathcal{D}=\langle\mathcal{V}, \mathcal{E}$, core, knw, poss, $\langle, \vartheta\rangle$ of $\mathcal{O}$ and an arbitrary finite set of queries, and a clause system $\mathcal{S}$ for $\mathcal{D}$ satisfying the preconditions of Theorem 37; furthermore, we fix $\mathbf{L}$ to be the set of all literals that occur in $\mathcal{O}, \mathcal{D}$, or $\mathcal{S}$. We prove Theorem 37 by constructing a pre-interpretation that satisfies $\mathcal{O}$ but refutes the relevant queries. To simplify the presentation, given an $\epsilon$-component $\mathcal{W}$ of $\mathcal{D}$, a context $v \in \mathcal{W}$, and a clause $K \sqsubseteq M$, we write $v \nvdash K \sqsubseteq M$ if and only if

- $K \subseteq \operatorname{poss}(v), \quad M \cap \operatorname{poss}(\mathcal{W}) \subseteq \operatorname{poss}(v)$, and $M$ is $<_{\nu}$-minimal (i.e., $v$ is complete for $K \sqsubseteq M$ ),
- $K \sqsubseteq L \hat{\in} \mathcal{S}(v)$ for each literal $L \in K$, and

Theorem 37 holds vacuously if no context $v \in \mathcal{V}$ and no query $K \sqsubseteq M$ exist such that $v \nvdash K \sqsubseteq M$. Thus, we assume that $v \nvdash K \sqsubseteq M$ holds for at least one context $v \in \mathcal{V}$ and query $K \sqsubseteq M$; by the definition of clause strengthening, we then also have $v \nvdash K \sqsubseteq M^{\prime}$ for $M^{\prime}=M \cap \mathbf{L}$; the latter clause is over $\mathbf{L}$ by Claim 70. Let $I=\langle\Delta, E, J\rangle$ be defined as follows.
- Set $\Delta$ is the smallest set that contains a distinguished element $\delta_{K \sqsubseteq M}^{v}$ for each context $v \in \mathcal{V}$ and each clause $K \sqsubseteq M$ over $\mathbf{L}$ such that $v \nvdash K \sqsubseteq M$. Set $\Delta$ is not empty due to the above assumption.
- The value of function $J$ is defined on each element $\delta_{K \sqsubseteq M}^{v} \in \Delta$ as follows. Let $\mathcal{W}$ be the $\epsilon$-component of $\mathcal{D}$ that contains $v$, and let $w_{0}, w_{1}, \ldots, w_{n}$ be an ordering of the elements of $\mathcal{W}$ obtained by an arbitrary depth-first traversal of $\mathcal{D}_{\mathcal{W}}$ starting from $w_{0}=v$. Conjunctions $K^{0}, K^{1}, \ldots, K^{n}$, disjunctions $M^{0}, M^{1}, \ldots, M^{n}$, and literal interpretations $J^{0}, J^{1}, \ldots, J^{n}$ are defined inductively as follows.
- Case $i=0$. Let $K^{0}:=K$, let $M^{0}:=M$, and let $J^{0}$ be an arbitrary literal interpretation of $\mathbf{L}$ w.r.t. $\mathcal{S}\left(w_{0}\right)$ and $\prec_{w_{0}}$ refuting $K^{0} \sqsubseteq M^{0}$.
- Case $i>0$. Let $w_{j}$ be the parent of $w_{i}$ in the depth-first traversal of $\mathcal{D}_{\mathcal{W}}$; since $w_{j}$ is considered before $w_{i}$ in the depth-first traversal, $K^{j}, M^{j}$, and $J^{j}$ have been defined. Let $K^{i}$ and $M^{i}$ be as follows:

$$
\begin{align*}
K^{i} & :=\bigvee\left\{C \in \operatorname{poss}\left(w_{j}\right) \cap \operatorname{poss}\left(w_{i}\right) \mid C \in J^{j}\right\}  \tag{B.1}\\
M^{i} & :=\bigsqcup\left\{C \in \operatorname{poss}\left(w_{j}\right) \cap \operatorname{poss}\left(w_{i}\right) \mid C \notin J^{j}\right\} \tag{B.2}
\end{align*}
$$


Based on these definitions for $J^{0}, J^{1}, \ldots, J^{n}$, the value of function $J$ on element $\delta_{K \sqsubseteq M}^{v}$ is defined by

$$
\begin{equation*}
J\left(\delta_{K \sqsubseteq M}^{v}\right):=\operatorname{poss}(\mathcal{W}) \cap\left(\bigcup_{i=0}^{n} J^{i}\right) . \tag{B.3}
\end{equation*}
$$

- Set $E$ is the smallest set containing $\left\langle\delta_{K_{1} \sqsubseteq M_{1}}^{v}, \delta_{A \sqcap K_{2} \sqsubseteq M_{2}}^{u}, R\right\rangle$ for all $\epsilon$-components $\mathcal{W}$ and $\mathcal{U}$ of $\mathcal{D}$, each element $\delta_{K_{1} \sqsubseteq M_{1}}^{v} \in \Delta$ with $v \in \mathcal{W}$, each literal $\exists R . A \in J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right)$, and each edge $\langle v, u, \exists R . A\rangle \in \mathcal{E}$ with $u \in \mathcal{U}$ such that $u \nvdash A \sqcap K_{2} \sqsubseteq M_{2}$, where $K_{2}$ and $M_{2}$ are as specified below.

$$
\begin{align*}
K_{2} & =\rceil\left\{B \mid \forall R . B \in J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right)\right\}  \tag{B.4}\\
M_{2} & =\bigsqcup\left\{\forall \operatorname{inv}(R) \cdot C \in \operatorname{poss}(\mathcal{U}) \mid C \notin J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right)\right\} \tag{B.5}
\end{align*}
$$

In the last item, we have $K_{2} \cup M_{2} \subseteq \mathbf{L}$ and $A \in \mathbf{L}$; thus, $u \nvdash A \sqcap K_{2} \sqsubseteq M_{2}$ ensures $\delta_{A \sqcap K_{2} \sqsubseteq M_{2}}^{u} \in \Delta$, and so $E$ is correctly defined. In contrast, the definition of $J\left(\delta_{K \sqsubseteq M}^{v}\right)$ relies on the existence of literal interpretations that refute certain clauses, and it is not obvious that these literal interpretations exist; hence, the following lemma proves that this is the case.

Claim 74. The value of $J\left(\delta_{K \sqsubseteq M}^{v}\right)$ is correctly defined for each element $\delta_{K \sqsubseteq M}^{v} \in \Delta$.
Proof. Consider an arbitrary element $\delta_{K \sqsubseteq M}^{v}$, and let $\mathcal{W}, w_{0}, w_{1}, \ldots, w_{n}, K^{0}, K^{1}, \ldots, K^{n}, M^{0}, M^{1}, \ldots, M^{n}$, and $J^{0}, J^{1}, \ldots, J^{n}$ be as specified above. We next inductively prove that $w_{i} \nvdash K^{i} \sqsubseteq M^{i}$ for each $0 \leq i \leq n$; this is sufficient to guarantee existence of the corresponding literal interpretation $J^{i}$. For $i=0$, we have $w_{0} \nvdash K^{0} \sqsubseteq M^{0}$ immediately from the facts that $K^{0}=K, M^{0}=M$, and $\delta_{K \sqsubseteq M}^{v} \in \Delta$. Consider now an arbitrary integer $i>0$, and let $w_{j}$ be the parent of $w_{i}$ in the depth-first traversal of $\mathcal{D}_{\mathcal{W}}$. Since the $\epsilon$-edges of $\mathcal{D}$ are symmetric by admissibility condition (E1), the presence of the undirected edge $\left\{w_{j}, w_{i}\right\}$ in $\mathcal{D}_{\mathcal{W}}$ implies that $\left\langle w_{j}, w_{i}, \epsilon\right\rangle \in \mathcal{E}$. Now $w_{i} \nvdash K^{i} \sqsubseteq M^{i}$ follows from the following observations.

- Condition $K^{i} \subseteq \operatorname{poss}\left(w_{i}\right)$ holds trivially by the definition of $K^{i}$.
- Condition $M^{i} \cap \operatorname{poss}(\mathcal{W}) \subseteq \operatorname{poss}\left(w_{i}\right)$ holds trivially since $M^{i} \subseteq \operatorname{poss}\left(w_{i}\right)$ by the definition of $M^{i}$.
- Disjunction $M^{i}$ is $<_{w_{i}}$-minimal by admissibility condition (P2) since $\left\langle w_{j}, w_{i}, \epsilon\right\rangle \in \mathcal{E}$ and $M^{i} \subseteq \operatorname{poss}\left(w_{j}\right) \cap \operatorname{poss}\left(w_{i}\right)$.
- Consider an arbitrary literal $L \in K^{i}$; we show $K^{i} \sqsubseteq L \hat{\in} \mathcal{S}\left(w_{i}\right)$. By the definition of $K^{i}$, we have $L \in \operatorname{poss}\left(w_{j}\right) \cap \operatorname{poss}\left(w_{i}\right)$; hence, $\left\langle w_{j}, w_{i}, \epsilon\right\rangle \in \mathcal{E}$ implies $L \sqsubseteq L \hat{\in} \mathcal{S}\left(w_{i}\right)$ by condition (I4) of Theorem 37, which is stronger than $K^{i} \sqsubseteq L \hat{\in} \mathcal{S}\left(w_{i}\right)$.
- Assume for contradiction that $K^{i} \sqsubseteq M^{i} \hat{\in} \mathcal{S}\left(w_{i}\right)$. Since $\left\langle w_{j}, w_{i}, \epsilon\right\rangle \in \mathcal{E}$ and $\mathcal{S}$ is closed under the Epsilon rule, $\mathcal{S}\left(w_{j}\right)$ is closed under $<_{w_{j}}$-hyperresolution with some strengthening of $K^{i} \sqsubseteq M^{i}$. By Lemmas 63 and 67 , we then have $J^{j} \models K^{i} \sqsubseteq M^{i}$; however, by the definition of $K^{i}$ and $M^{i}$ we clearly have $J^{j} \not \vDash K^{i} \sqsubseteq M^{i}$. Thus, $K^{i} \sqsubseteq M^{i} \notin \mathcal{S}\left(w_{i}\right)$ holds, as required.

Next, we prove several important properties of $J$.
Claim 75. Let $\delta_{K \sqsubseteq M}^{v} \in \Delta$ be an arbitrary element, and let $\mathcal{W}$ be the $\epsilon$-component of $\mathcal{D}$ such that $v \in \mathcal{W}$. Then,

1. $\operatorname{knw}(\mathcal{W}) \subseteq J\left(\delta_{K \subseteq M}^{v}\right) \subseteq \operatorname{poss}(\mathcal{W})$,
2. $J\left(\delta_{K \sqsubseteq M}^{v}\right) \not \vDash K \sqsubseteq M$, and
3. $J\left(\delta_{K \sqsubseteq M}^{v}\right) \vDash K^{\prime} \sqsubseteq M^{\prime}$ for each clause $K^{\prime} \sqsubseteq M^{\prime}$ for which a context $w \in \mathcal{W}$ exists such that $K^{\prime} \cup M^{\prime} \subseteq \operatorname{poss}(w)$ and $\mathcal{S}(w)$ is closed under $<_{w}$-hyperresolution with $K^{\prime} \sqsubseteq M^{\prime}$.

Proof. Consider an arbitrary element $\delta_{K \sqsubset M}^{v} \in \Delta$; clearly, we have $v \nvdash K \sqsubseteq M$. Furthermore, let $\mathcal{W}$ be the $\epsilon$-component of $\mathcal{D}$ such that $v \in \mathcal{W}$, and let $w_{0}, w_{1}, \ldots, w_{n}, K^{0}, K^{1}, \ldots, K^{n}, M^{0}, M^{1}, \ldots, M^{n}$, and $J^{0}, J^{1}, \ldots, J^{n}$ be as specified in the definition of the pre-interpretation I. Before proceeding, we first prove the following auxiliary property ( $*$ ).

For all integers $0 \leq i, j \leq n$ and each literal $L \in \operatorname{poss}\left(w_{i}\right) \cap \operatorname{poss}\left(w_{j}\right)$, we have $L \in J^{i}$ if and only if $L \in J^{j}$.
Since $\mathcal{D}_{\mathcal{W}}$ is a tree by admissibility condition (E2), the shortest path between $w_{i}$ and $w_{j}$ is unique; we now prove (*) by induction on the length of this path. Towards this goal, consider an arbitrary literal $L \in \operatorname{poss}\left(w_{i}\right) \cap \operatorname{poss}\left(w_{j}\right)$; we have the following cases.

- If $i=j$, then $J^{i}=J^{j}$, so clearly $L \in J^{i}$ if and only if $L \in J^{j}$.
- Assume that $w_{j}$ is the parent of $w_{i}$ in the traversal of $\mathcal{D}_{\mathcal{W}}$. Then $K^{i} \cup M^{i}=\operatorname{poss}\left(w_{i}\right) \cap \operatorname{poss}\left(w_{j}\right)$ by the definition of $K^{i}$ and $M^{i}$, which implies $L \in K^{i} \cup M^{i}$. Furthermore, by Lemma 66 we have $J^{i} \not \models K^{i} \sqsubseteq M^{i}$. But then, if $L \in K^{i}$, then by the definition of $K^{i}$ we have $L \in J^{j}$, and $J^{i} \not \vDash K^{i} \sqsubseteq M^{i}$ implies $L \in J^{i}$. Analogously, if $L \in M^{i}$, then by the definition of $M^{i}$ we have $L \notin J^{j}$, and $J^{i} \not \neq K^{i} \sqsubseteq M^{i}$ implies $L \notin J^{i}$.
- The case when $w_{i}$ is the parent of $w_{j}$ is symmetric to the previous one.
- The remaining possibility is that the length of the shortest path between $w_{i}$ and $w_{j}$ is greater than one, so let $w_{k}$ be an arbitrary vertex on this path different from $w_{i}$ and $w_{j}$. By admissibility condition (E3), we have $L \in \operatorname{poss}\left(w_{k}\right)$. Furthermore, the paths between $w_{i}$ and $w_{k}$, and between $w_{k}$ and $w_{j}$ are both shorter than the path between $w_{i}$ and $w_{j}$, so property (*) holds for $i$ and $k$, and for $k$ and $j$. But then, property ( $*$ ) clearly holds for $i$ and $j$ as well.

Property $(*)$ clearly implies that, for each integer $0 \leq i \leq n$ and each literal $L \in \operatorname{poss}\left(w_{i}\right)$, we have $L \in J\left(\delta_{K \subseteq M}^{v}\right)$ if and only if $L \in J^{i}$; we call this property $(\dagger)$. We are now ready to prove Properties $1-3$ of this claim.
(Property 1) The definition of $J\left(\delta_{K \sqsubseteq M}^{v}\right)$ in (B.3) clearly implies $J\left(\delta_{K \sqsubseteq M}^{v}\right) \subseteq \operatorname{poss}(\mathcal{W})$. In order to prove $\operatorname{knw}(\mathcal{W}) \subseteq J\left(\delta_{K \sqsubseteq M}^{v}\right)$, consider an arbitrary literal $L \in \operatorname{knw}(\mathcal{W})$; clearly, a context $w_{i} \in \mathcal{W}$ exists such that $L \in \operatorname{knw}\left(w_{i}\right)$. By condition (I2) of Theorem 37, we have $T \sqsubseteq L \hat{\in} \mathcal{S}\left(w_{i}\right)$; hence $L \in J^{i}$ by Lemma 65 . Consequently, $L \in J\left(\delta_{K \sqsubseteq M}^{v}\right)$ holds by property ( $\dagger$ ), as required.
(Property 2) By Lemma 66, $K^{0}=K$, and $M^{0}=M$, we have $J^{0} \not \models K \sqsubseteq M$-that is, $K \subseteq J^{0}$ and $M \cap J^{0}=\emptyset$. Now $v \nvdash K \sqsubseteq M$ implies $K \subseteq \operatorname{poss}(v)$ and $M \cap \operatorname{poss}(\mathcal{W}) \subseteq \operatorname{poss}(v)$, so property $(\dagger)$ implies $K \subseteq J\left(\delta_{K \subseteq M}^{v}\right)$ and $[M \cap \operatorname{poss}(\mathcal{W})] \cap J\left(\delta_{K \sqsubseteq M}^{v}\right)=\emptyset$. Furthermore, since $\operatorname{poss}(\mathcal{W}) \cap J\left(\delta_{K \sqsubseteq M}^{v}\right)=J\left(\delta_{K \sqsubseteq M}^{v}\right)$, we have $M \cap J\left(\delta_{K \sqsubseteq M}^{v}\right)=\emptyset$; thus, $J\left(\delta_{K \sqsubseteq M}^{v}\right) \not \vDash K \sqsubseteq M$ holds, as required.
(Property 3) Let $K^{\prime} \sqsubseteq M^{\prime}$ be a clause and let $w_{i} \in \mathcal{W}$ be a context such that $K^{\prime} \cup M^{\prime} \subseteq \operatorname{poss}\left(w_{i}\right)$ and set $\mathcal{S}\left(w_{i}\right)$ is closed under $<_{w_{i}}$-hyperresolution with $K^{\prime} \sqsubseteq M^{\prime}$. By Lemma 67 , we have $J^{i} \vDash K^{\prime} \sqsubseteq M^{\prime}$. But then, since $K^{\prime} \cup M^{\prime} \subseteq \operatorname{poss}\left(w_{i}\right)$, property ( $\dagger$ ) implies $J\left(\delta_{K \sqsubseteq M}^{v}\right) \vDash K^{\prime} \sqsubseteq M^{\prime}$, as required.

To complete the proof of Theorem 37, the following claims show that structure $I$ is a pre-interpretation, and that it satisfies $\mathcal{O}$ but refutes the relevant queries.

Claim 76. Structure I satisfies conditions $\left(I_{\exists}\right)$ and $\left(I_{\Downarrow}\right)$ of pre-interpretations from Definition 59.

Proof. (Condition $I_{\exists}$ ) Consider an arbitrary element $\delta_{K_{1} \sqsubseteq M_{1}}^{v} \in \Delta$ and an arbitrary literal $\exists R . A \in J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right)$; we show that an edge $\left\langle\delta_{K_{1} \sqsubseteq M_{1}}^{v}, \gamma, R\right\rangle \in E$ exists such that $A \in J(\gamma)$. Let $\mathcal{W}$ be the $\epsilon$-component of $\mathcal{D}$ such that $v \in \mathcal{W}$. By admissibility condition (S3), an edge $\langle w, u, \exists R . A\rangle \in \mathcal{E}$ exists such that $w \in \mathcal{W}$ and conditions (B.6) and (B.7) are both satisfied.

$$
\begin{align*}
& \operatorname{poss}(u) \supseteq\{A\} \cup\{B \mid \forall R . B \in \operatorname{poss}(\mathcal{W})\}  \tag{B.6}\\
& \operatorname{poss}(w) \supseteq\{\exists R . A\} \cup\{\forall R . B \mid \forall R . B \in \operatorname{poss}(\mathcal{W})\} \tag{B.7}
\end{align*}
$$

Let $\mathcal{U}$ be the $\epsilon$-component of $\mathcal{D}$ such that $u \in \mathcal{U}$, and let $K_{2}$ and $M_{2}$ be as specified in (B.4) and (B.5), respectively. The following observations prove that $u \nvdash A \sqcap K_{2} \sqsubseteq M_{2}$ holds.

- Condition $A \sqcap K_{2} \subseteq \operatorname{poss}(u)$ follows from (B.4), Property 1 of Claim 75, and (B.6).
- Admissibility condition (S2) implies $M_{2} \cap \operatorname{poss}(\mathcal{U}) \subseteq \operatorname{poss}(u)$.
- Disjunction $M_{2}$ is $<_{u}$-minimal since, by admissibility condition ( P 1 ), $\langle w, u, \exists R . A\rangle \in \mathcal{E}$ implies that $<_{u}$ is $R$-admissible.
- Consider an arbitrary literal $L \in A \sqcap K_{2}$; we show that $A \sqcap K_{2} \sqsubseteq L \hat{\in} \mathcal{S}(u)$. Assume that $L \in K_{2}$, so $L$ is an atomic concept $L=B$. Then, we have $\forall R . B \in J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right) \subseteq \operatorname{poss}(\mathcal{W})$ by (B.4) and Property 1 of Claim 75. Thus, we have $B \in \operatorname{poss}(u)$ and $\forall R . B \in \operatorname{poss}(w)$ by (B.6) and (B.7), and then $B \sqsubseteq B \hat{\in} \mathcal{S}(u)$ by condition (I3) of Theorem 37, which is stronger than the required $A \sqcap K_{2} \sqsubseteq B \hat{\in} \mathcal{S}(u)$. The proof for the case $L=A$ is analogous.
- Assume for contradiction that $A \sqcap K_{2} \sqsubseteq M_{2} \hat{\in} \mathcal{S}(u)$. By admissibility condition (S2), we have $C \in \operatorname{poss}(w)$ for each $\forall \operatorname{inv}(R) . C \in M_{2}$. Thus, since the Pred rule from Table 5 is not applicable to $\mathcal{S}(v)$, the Pred from Table 3 is not applicable to $\mathcal{S}(v)$ either. Hence, by Lemma 69 , set $\mathcal{S}(w)$ is closed under $<_{w}$-hyperresolution with the clause

$$
\exists R . A \sqcap\rceil\left\{\forall R . B \mid \forall R . B \in J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right)\right\} \sqsubseteq \bigsqcup\left\{C \mid \forall \operatorname{inv}(R) . C \in \operatorname{poss}(\mathcal{U}) \text { and } C \notin J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right)\right\} .
$$

Recall that $\exists R . A \in J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right)$, so the above clause is clearly refuted in $J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right)$. However, all literals on the left-hand side of this clause are contained in $\operatorname{poss}(w)$ by (B.7), and each concept $C$ on its right-hand side is contained in $\operatorname{poss}(w)$; but then, Property 3 of Claim 75 implies that this clause is satisfied in $J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right)$, which is a contradiction. Consequently, we have $A \sqcap K_{2} \sqsubseteq M_{2} \notin \mathcal{S}(u)$, as required.

Thus, we have $u \nvdash A \sqcap K_{2} \sqsubseteq M_{2}$, which implies $\left\langle\delta_{K_{1} \subseteq M_{1}}^{v}, \delta_{A \sqcap K_{2} \sqsubseteq M_{2}}^{u}, R\right\rangle \in E$ by the definition of $E$. Finally, $A \in J\left(\delta_{A \sqcap K_{2} \sqsubseteq M_{2}}^{u}\right)$ holds by Property 2 of Claim 75. This concludes the proof of condition ( $I_{\exists}$ ).
(Condition $I_{y}$ ). Consider an arbitrary edge $\left\langle\delta_{K_{1} \sqsubseteq M_{1}}^{v}, \delta_{A \sqcap K_{2} \sqsubseteq M_{2}}^{u}, R\right\rangle \in E$ and arbitrary literals $\forall R . B$ and $\forall \operatorname{inv}(R) . C$, and let $\mathcal{U}$ be the $\epsilon$-component of $\mathcal{D}$ such that $u \in \mathcal{U}$.

- Assume that $\forall R . B \in J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right)$. Then $B \in K_{2}$ by (B.4); hence Property 2 of Claim 75 implies $J\left(\delta_{A \sqcap K_{2} \sqsubseteq M_{2}}^{u}\right) \not \vDash A \sqcap K_{2} \sqsubseteq M_{2}$; therefore $B \in J\left(\delta_{A \sqcap K_{2} \subseteq M_{2}}^{u}\right)$, as required.
- Assume that $\forall \operatorname{inv}(R) . C \in J\left(\delta_{A \sqcap K_{2} \sqsubseteq M_{2}}^{u}\right)$. Then $\forall \operatorname{inv}(R) . C \in \operatorname{poss}(\mathcal{U})$ by Property 1 of Claim 75. Now, if $C \notin J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right)$, then (B.5) implies $\forall \operatorname{inv}(R) . C \in M_{2}$, so $\forall \operatorname{inv}(R) . C \notin J\left(\delta_{A \sqcap K_{2} \subseteq M_{2}}^{u}\right)$ follows from Property 2 of Claim 75, which contradicts the assumption. Thus, we have $C \in J\left(\delta_{K_{1} \sqsubseteq M_{1}}^{v}\right)$, as required.


## Claim 77. Pre-interpretation I is a pre-model of $\mathcal{O}$.

Proof. Consider an arbitrary element $\delta_{K \sqsubseteq M}^{v} \in \Delta$ and an arbitrary normal clause $K^{\prime} \sqsubseteq M^{\prime} \in \mathcal{O}$; we show that $J\left(\delta_{K \sqsubseteq M}^{v}\right)$ satisfies $K^{\prime} \sqsubseteq M^{\prime}$. This is obvious if $K^{\prime} \nsubseteq J\left(\delta_{K \subseteq M}^{v}\right)$, so assume that $K^{\prime} \subseteq J\left(\delta_{K \subseteq M}^{v}\right)$. Let $\mathcal{W}$ be the $\epsilon$-component of $\mathcal{D}$ such that $v \in \mathcal{W}$. Then, Property 1 of Claim 75 implies $K^{\prime} \subseteq \operatorname{poss}(\mathcal{W})$. Furthermore, by the ontology condition of Definition 36, a context $w \in \mathcal{W}$ exists such that $K^{\prime} \cup M^{\prime} \subseteq \operatorname{poss}(w)$. Since $\mathcal{S}$ is closed under the Hyper rule from Table 5 , set $\mathcal{S}(w)$ is closed under ${ }{ }_{w}$-hyperresolution with $K^{\prime} \sqsubseteq M^{\prime}$. But then, Property 3 of Claim 75 implies $J\left(\delta_{K \sqsubseteq M}^{v}\right) \vDash K^{\prime} \sqsubseteq M^{\prime}$, as required.

Claim 78. Pre-interpretation I refutes each query $K \sqsubseteq M$ for which there exists a context $v \in \mathcal{V}$ such that $v$ is complete for $K \sqsubseteq M, K \sqsubseteq L \hat{\epsilon} \mathcal{S}(v)$ for each $L \in K$, and $K \sqsubseteq M \notin \mathcal{S}(v)$.

Proof. Let $K \sqsubseteq M$ be an arbitrary query satisfying the preconditions of the claim, and let $M^{\prime}=M \cap \mathbf{L}$. Then $v \nvdash K \sqsubseteq M^{\prime}$ holds as well, and this clause is over $\mathbf{L}$ by Claim 70. But then, $J\left(\delta_{K \sqsubseteq M^{\prime}}^{v}\right) \not \vDash K \sqsubseteq M^{\prime}$ by Property 2 of Claim 75. Finally, each literal from $M \backslash M^{\prime}$ is not from $\mathbf{L}$ so it cannot occur in $J\left(\delta_{K \sqsubseteq M^{\prime}}^{v}\right)$; thus, $J\left(\bar{\delta}_{K \sqsubseteq M^{\prime}}^{v}\right) \not \models K \sqsubseteq M$ holds as well, so $I$ refutes $K \sqsubseteq M$, as required.

## References

[1] A. Sidhu, T. Dillon, E. Chang, B. S. Sidhu, Protein Ontology Development using OWL, in: Proc. of the OWL: Experiences and Directions Workshop (OWLED 2005), volume 188 of CEUR WS Proceedings, Galway, Ireland.
[2] C. Golbreich, S. Zhang, O. Bodenreider, The Foundational Model of Anatomy in OWL: Experience and Perspectives, Journal of Web Semantics 4 (2006) 181-195.
[3] J. Goodwin, Experiences of using OWL at the Ordnance Survey, in: Proc. of the OWL: Experiences and Directions Workshop (OWLED 2005), volume 188 of CEUR WS Proceedings, Galway, Ireland.
[4] S. Derriere, A. Richard, A. Preite-Martinez, An Ontology of Astronomical Object Types for the Virtual Observatory, in: Proc. of the 26th meeting of the IAU: Virtual Observatory in Action: New Science, New Technology, and Next Generation Facilities, Prague, Czech Republic, pp. 17-18.
[5] D. Soergel, B. Lauser, A. Liang, F. Fisseha, J. Keizer, S. Katz, Reengineering Thesauri for New Applications: The AGROVOC Example, Journal of Digital Information 4 (2004).
[6] L. Lacy, G. Aviles, K. Fraser, W. Gerber, A. Mulvehill, R. Gaskill, Experiences Using OWL in Military Applications, in: Proc. of the OWL: Experiences and Directions Workshop (OWLED 2005), volume 188 of CEUR WS Proceedings, Galway, Ireland.
[7] P. F. Patel-Schneider, P. Hayes, I. Horrocks, OWL Web Ontology Language: Semantics and Abstract Syntax, W3C Recommendation, http://www.w3.org/TR/owl-semantics/, 2004.
[8] B. Motik, P. F. Patel-Schneider, B. Parsia, OWL 2 Web Ontology Language: Structural Specification and Functional-Style Syntax, W3C Recommendation, http://www.w3.org/TR/owl2-syntax/, 2009.
[9] B. Cuenca Grau, I. Horrocks, B. Motik, B. Parsia, P. Patel-Schneider, U. Sattler, OWL 2: The next step for OWL, Journal of Web Semantics: Science, Services and Agents on the World Wide Web 6 (2008) 309-322.
[10] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, P. F. Patel-Schneider (Eds.), The Description Logic Handbook: Theory, Implementation and Applications, Cambridge University Press, 2nd edition, 2007.
[11] H. Andréka, J. van Benthem, I. Németi, Modal Languages and Bounded Fragments of Predicate Logic, Journal of Philosophical Logic 27 (1998) $217-274$.
[12] B. Parsia, E. Sirin, Pellet: An OWL-DL Reasoner, Poster at the 3rd Int. Semantic Web Conference (ISWC 2004), 2004.
[13] D. Tsarkov, I. Horrocks, FaCT++ Description Logic Reasoner: System Description, in: Proc. of the 3rd Int. Joint Conf. on Automated Reasoning (IJCAR 2006), volume 4130 of LNAI, Springer, Seattle, WA, USA, 2006, pp. 292-297.
[14] B. Motik, R. Shearer, I. Horrocks, Hypertableau Reasoning for Description Logics, Journal of Artificial Intelligence Research 36 (2009) 165-228.
[15] R. S. Gonçalves, B. Parsia, U. Sattler, Performance Heterogeneity and Approximate Reasoning in Description Logic Ontologies, in: P. Cudré-Mauroux, J. Heflin, E. Sirin, T. Tudorache, J. Euzenat, M. Hauswirth, J. X. Parreira, J. Hendler, G. Schreiber, A. Bernstein, E. Blomqvist (Eds.), Proc. of the 11th Int. Semantic Web Conference (ISWC 2012), volume 7649 of LNCS, Springer, Boston, MA, USA, 2012, pp. 82-98.
[16] U. Hustadt, B. Motik, U. Sattler, Data Complexity of Reasoning in Very Expressive Description Logics, in: Proc. of the 19th Int. Joint Conf. on Artificial Intelligence (IJCAI 2005), Morgan Kaufmann Publishers, Edinburgh, UK, 2005, pp. 466-471.
[17] M. Krötzsch, S. Rudolph, P. Hitzler, Complexity Boundaries for Horn Description Logics, in: Proc. of the 22nd National Conference on Artificial Intelligence (AAAI 2007), AAAI Press, Vancouver, BC, Canada, 2007, pp. 452-457.
[18] F. Baader, S. Brandt, C. Lutz, Pushing the $\mathcal{E} \mathcal{L}$ Envelope, in: L. P. Kaelbling, A. Saffiotti (Eds.), Proc. of the 19th Int. Joint Conference on Artificial Intelligence (IJCAI 2005), Morgan Kaufmann Publishers, Edinburgh, UK, 2005, pp. 364-369.
[19] D. Calvanese, G. De Giacomo, D. Lembo, M. Lenzerini, R. Rosati, Tractable Reasoning and Efficient Query Answering in Description Logics: The DL-Lite Family, Journal of Automated Reasoning 9 (2007) 385-429.
[20] A. Artale, D. Calvanese, R. Kontchakov, M. Zakharyaschev, The DL-Lite Family and Relations, Journal of Artificial Intelligence Research 36 (2009) 1-69.
[21] Y. Kazakov, Consequence-Driven Reasoning for Horn SHIQ Ontologies, in: C. Boutilier (Ed.), Proc. of the 21st Int. Joint Conf. on Artificial Intelligence (IJCAI 2009), Pasadena, CA, USA, pp. 2040-2045.
[22] F. Simančík, Y. Kazakov, I. Horrocks, Consequence-Based Reasoning beyond Horn Ontologies, in: T. Walsh (Ed.), Proc. of the 22nd Int. Joint Conf. on Artificial Intelligence (IJCAI 2011), Barcelona, Spain, pp. 1093-1098.
[23] Y. Kazakov, M. Krötzsch, F. Simančík, Concurrent Classification of $\mathcal{E L}$ Ontologies, in: L. Aroyo, C. Welty, H. Alani, J. Taylor, A. Bernstein, L. Kagal, N. F. Noy, E. Blomqvist (Eds.), Proc. of the 10th Int. Semantic Web Conference (ISWC 2011), volume 7031 of LNCS, Springer, Boston, MA, USA, 2011, pp. 305-320.
[24] R. G. Downey, M. R. Fellows, Parameterized Complexity, Springer, 1999.
[25] N. Robertson, P. Seymour, Graph minors. III. Planar tree-width, Journal of Combinatorial Theory, Series B 36 (1984) 49-64.
[26] S. Szeider, On Fixed-Parameter Tractable Parameterizations of SAT, in: E. Giunchiglia, A. Tacchella (Eds.), Proc. of the 6th Int. Conf. on Theory and Applications of Satisfiability Testing (SAT 2003), volume 2919 of LNCS, Springer, Santa Margherita Ligure, Italy, 2003, pp. 188-202.
[27] C. Lutz, F. Wolter, Non-Uniform Data Complexity of Query Answering in Description Logics, in: G. Brewka, T. Eiter, S. A. McIlraith (Eds.), Proc. of the 13th Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR 2012), AAAI Press, Rome, Italy, 2012, pp. 297-307.
[28] S. Demri, H. de Nivelle, Deciding Regular Grammar Logics with Converse Through First-Order Logic, Journal of Logic, Language and Information 14 (2005) 289-329.
[29] H. L. Bodlaender, A Linear-Time Algorithm for Finding Tree-Decompositions of Small Treewidth, SIAM Journal on Computing 25 (1996) $1305-1317$.
[30] C. G. Fermüller, A. Leitsch, U. Hustadt, T. Tammet, Resolution Decision Procedures, in: A. Robinson, A. Voronkov (Eds.), Handbook of Automated Reasoning, volume II, Elsevier Science, 2001, pp. 1791-1849.
[31] L. Bachmair, H. Ganzinger, Resolution Theorem Proving, in: A. Robinson, A. Voronkov (Eds.), Handbook of Automated Reasoning, volume I, Elsevier Science, 2001, pp. 19-99.
[32] B. Motik, U. Sattler, A Comparison of Reasoning Techniques for Querying Large Description Logic ABoxes, in: M. Hermann, A. Voronkov (Eds.), Proc. of the 13th Int. Conf. on Logic for Programming Artificial Intelligence and Reasoning (LPAR 2006), volume 4246 of LNCS, Springer, Phnom Penh, Cambodia, 2006, pp. 227-241.
[33] R. Goré, L. A. Nguyen, EXPTIME Tableaux with Global Caching for Description Logics with Transitive Roles, Inverse Roles and Role Hierarchies, in: N. Olivetti (Ed.), Proc. of the 16th Int. Conf. on Automated Reasoning with Tableaux and Related Methods (TABLEAUX 2007), volume 4548 of $L N C S$, Springer, Aix en Provence, France, 2007, pp. 133-148.
[34] R. Goré, L. A. Nguyen, ExpTime Tableaux for $\mathcal{A L C}$ Using Sound Global Caching, Journal of Automated Reasoning 50 (2013) 355-381.
[35] C. Weidenbach, Combining superposition, sorts and splitting, in: A. Robinson, A. Voronkov (Eds.), Handbook of Automated Reasoning, volume II, Elsevier Science, 2001, pp. 1965-2013.
[36] S. Schulz, System Description: E 0.81 , in: D. A. Basin, M. Rusinowitch (Eds.), Proc. of the 2nd Int. Conf. on Automated Reasoning (IJCAR 2004), volume 3097 of $L N C S$, Springer, Cork, Ireland, 2004, pp. 223-228.
[37] B. Glimm, I. Horrocks, B. Motik, G. Stoilos, A Novel Approach to Ontology Classification, Journal of Web Semantics: Science, Services and Agents on the World Wide Web 14 (2012) 84-101.


[^0]:    ${ }^{1}$ http://opengalen.org/

[^1]:    ${ }^{2}$ http://ihtsdo.org/

[^2]:    ${ }^{3}$ Note that, by this definition, each clause is a strengthening of itself.

[^3]:    ${ }^{4}$ An ontology with ID XXXXX can be readily downloaded from http://www.cs.ox.ac.uk/isg/ontologies/UID/XXXXX.owl, and a general description of our repository is available at http://www.cs.ox.ac.uk/isg/ontologies/.
    ${ }^{5}$ http://www.itu.dk/people/sathi/treed/

