Reasoning with Axioms: Theory and Practice

Ian Horrocks
Department of Computer Science
University of Manchester, UK
horrocks@cs.man.ac.uk

Stephan Tobies
LuFg Theoretical Computer Science
RWTH Aachen, Germany
tobies@informatik.rwth-aachen.de

Abstract

When reasoning in description, modal or temporal logics it is often useful to consider axioms representing universal truths in the domain of discourse. Reasoning with respect to an arbitrary set of axioms is hard, even for relatively inexpressive logics, and it is essential to deal with such axioms in an efficient manner if implemented systems are to be effective in real applications. This is particularly relevant to Description Logics, where subsumption reasoning with respect to a terminology is a fundamental problem. Two optimisation techniques that have proved to be particularly effective in dealing with terminologies are lazy unfolding and absorption. In this paper we seek to improve our theoretical understanding of these important techniques. We define a formal framework that allows the techniques to be precisely described, establish conditions under which they can be safely applied, and prove that, provided these conditions are respected, subsumption testing algorithms will still function correctly. These results are used to show that the procedures used in the FaCT system are correct and, moreover, to show how efficiency can be significantly improved, while still retaining the guarantee of correctness, by relaxing the safety conditions for absorption.

1 MOTIVATION

Description Logics (DLs) form a family of formalisms which have grown out of knowledge representation techniques using frames and semantic networks. DLs use a class based paradigm, describing the domain of interest in terms of concepts (classes) and roles (binary relations) which can be combined using a range of operators to form more complex structured concepts [BHH+91]. A DL *terminology* typically consists of a set of asserted facts, in particular asserted subsumption (is-a-kind-of) relationships between (possibly complex) concepts.[1].

One of the distinguishing characteristics of DLs is a formally defined semantics which allows the structured objects they describe to be reasoned with. Of particular interest is the computation of implied subsumption relationships between concepts, based on the assertions in the terminology, and the maintenance of a concept hierarchy (partial ordering) based on the subsumption relationship [WS92].

The problem of computing concept subsumption relationships has been the subject of much research, and sound and complete algorithms are now known for a wide range of DLs (for example [HN90, BH91, Baa91, DM98, HST99]). However, in spite of the fundamental importance of terminologies in DLs, most of these algorithms deal only with the problem of deciding subsumption between two concepts (or, equivalently, concept satisfiability), without reference to a terminology (but see [BDS93, Cal96, DDM96, HST99]). By restricting the kinds of assertion that can appear in a terminology, concepts can be syntactically expanded so as to explicitly include all relevant terminological information. This procedure, called *unfolding*, has mostly been applied to less expressive DLs. With more expressive DLs, in particular those supporting universal roles, it is often possible to encapsulate an arbitrary terminology in a single concept. This technique can be used with satisfiability testing to ensure that the result is valid with respect to the assertions in the ter-

[1] DLs can also deal with assertions about individuals, but in this paper we will only be concerned with *terminological* (concept based) reasoning

minology, a procedure called *internalisation*.

Although the above mentioned techniques suffice to demonstrate the theoretical adequacy of satisfiability decision procedures for terminological reasoning, experiments with implementations have shown that, for reasons of (lack of) efficiency, they are highly unsatisfactory as a practical methodology for reasoning with DL terminologies. Firstly, experiments with the KRIS system have shown that integrating unfolding with the (tableaux) satisfiability algorithm (*lazy unfolding*) leads to a significant improvement in performance [BFH+94]. More recently, experiments with the FaCT system have shown that reasoning becomes hopelessly intractable when internalisation is used to deal with larger terminologies [Hor98]. However, the FaCT system has also demonstrated that this problem can be dealt with (at least for realistic terminologies) by using a combination of lazy unfolding and internalisation, having first manipulated the terminology in order to minimise the number of assertions that must be dealt with by internalisation (a technique called *absorption*).

It should be noted that, although these techniques were discovered while developing DL systems, they are applicable to a whole range of reasoning systems, independent of the concrete logic and type of algorithm. As well as tableaux based decision procedures, this includes resolution based algorithms, where the importance of minimising the number of terminological sentences has already been noted [HS99], and sequent calculus algorithms, where there is a direct correspondence with tableaux algorithms [BFH+99].

In this paper we seek to improve our theoretical understanding of these important techniques which has, until now, been very limited. In particular we would like to know exactly when and how they can be applied, and be sure that the answers we get from the algorithm are still correct. This is achieved by defining a formal framework that allows the techniques to be precisely described, establishing conditions under which they can be safely applied, and proving that, provided these conditions are respected, satisfiability algorithms will still function correctly. These results are then used to show that the procedures used in the FaCT system are correct[2] and, moreover, to show how efficiency can be significantly improved, while still retaining the guarantee of correctness, by relaxing the safety conditions for absorption. Finally, we identify several interesting directions for future research, in particular the problem of finding the "best" absorption possible.

[2]Previously, the correctness of these procedures had only been demonstrated by a relatively ad-hoc argument [Hor97].

2 PRELIMINARIES

Firstly, we will establish some basic definitions that clarify what we mean by a DL, a terminology (subsequently called a TBox), and subsumption and satisfiability with respect to a terminology, . The results in this paper are uniformly applicable to a whole range of DLs, as long as some basic criteria are met:

Definition 2.1 (Description Logic) *Let* L *be a DL based on infinite sets of atomic concepts* NC *and atomic roles* NR. *We will identify* L *with the sets of its well-formed concepts and require* L *to be closed under boolean operations and sub-concepts.*

An interpretation is a pair $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, *where* $\Delta^{\mathcal{I}}$ *is a non-empty set, called the* domain *of* $\mathcal{I}$, *and* $\cdot^{\mathcal{I}}$ *is a function mapping* NC *to* $2^{\Delta^{\mathcal{I}}}$ *and* NR *to* $2^{\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}}$. *With each DL* L *we associate a set* $\mathsf{Int}(\mathsf{L})$ *of* admissible *interpretations for* L. $\mathsf{Int}(\mathsf{L})$ *must be closed under isomorphisms, and, for any two interpretations* $\mathcal{I}$ *and* $\mathcal{I}'$ *that agree on* NR, *it must satisfy* $\mathcal{I} \in \mathsf{Int}(\mathsf{L}) \Leftrightarrow \mathcal{I}' \in \mathsf{Int}(\mathsf{L})$. *Additionally, we assume that each DL* L *comes with a semantics that allows any interpretation* $\mathcal{I} \in \mathsf{Int}(\mathsf{L})$ *to be extended to each concept* $C \in \mathsf{L}$ *such that it satisfies the following conditions:*

(I1) it maps the boolean combination of concepts to the corresponding boolean combination of their interpretations, and

(I2) the interpretation $C^{\mathcal{I}}$ *of a compound concept* $C \in \mathsf{L}$ *depends only on the interpretation of those atomic concepts and roles that appear syntactically in* C.

This definition captures a whole range of DLs, namely, the important DL $\mathcal{ALC}$ [SS91] and its many extensions. $\mathsf{Int}(\mathsf{L})$ hides restrictions on the interpretation of certain roles like transitivity, functionality, or role hierarchies, which are imposed by more expressive DLs (e.g., [HST99]), as these are irrelevant for our purposes. In these cases, $\mathsf{Int}(\mathsf{L})$ will only contain those interpretations which interpret the roles as required by the semantics of the logic, e.g., features by partial functions or transitively closed roles by transitive relations. Please note that various modal logics [Sch91], propositional dynamic logics [DL94] and temporal logics [EH85] also fit into this framework. We will use $C \rightarrow D$ as an abbreviation for $\neg C \sqcup D$, $C \leftrightarrow D$ as an abbreviation for $(C \rightarrow D) \sqcap (D \rightarrow C)$, and $\top$ as a tautological concept, e.g., $A \sqcup \neg A$ for an arbitrary $A \in \mathsf{NC}$.

A TBox consists of a set of axioms asserting subsumption or equality relations between (possibly complex) concepts.

Definition 2.2 (TBox, Satisfiability) *A* TBox $\mathcal{T}$ *for* L *is a finite set of axioms of the form* $C_1 \sqsubseteq C_2$ *or* $C_1 \doteq C_2$, *where* $C_i \in \mathsf{L}$. *If, for some* $A \in \mathsf{NC}$, $\mathcal{T}$ *contains one or more axioms of the form* $A \sqsubseteq C$ *or* $A \doteq C$, *then we say that* A *is* defined *in* $\mathcal{T}$.

Let L *be a DL and* $\mathcal{T}$ *a TBox. An interpretation* $\mathcal{I} \in \mathsf{Int}(\mathsf{L})$ *is a* model *of* $\mathcal{T}$ *iff, for each* $C_1 \sqsubseteq C_2 \in \mathcal{T}$, $C_1^{\mathcal{I}} \subseteq C_2^{\mathcal{I}}$ *holds, and, for each* $C_1 \doteq C_2 \in \mathcal{T}$, $C_1^{\mathcal{I}} = C_2^{\mathcal{I}}$ *holds. In this case we write* $\mathcal{I} \models \mathcal{T}$. *A concept* $C \in \mathsf{L}$ *is* satisfiable *with respect to a TBox* $\mathcal{T}$ *iff there is an* $\mathcal{I} \in \mathsf{Int}(\mathsf{L})$ *with* $\mathcal{I} \models \mathcal{T}$ *and* $C^{\mathcal{I}} \neq \emptyset$. *A concept* $C \in \mathsf{L}$ subsumes *a concept* $D \in \mathsf{L}$ *w.r.t.* $\mathcal{T}$ *iff, for all* $\mathcal{I} \in \mathsf{Int}(\mathsf{L})$ *with* $\mathcal{I} \models \mathcal{T}$, $C^{\mathcal{I}} \supseteq D^{\mathcal{I}}$ *holds.*

Two TBoxes $\mathcal{T}, \mathcal{T}'$ *are called* equivalent ($\mathcal{T} \equiv \mathcal{T}'$), *iff, for all* $\mathcal{I} \in \mathsf{Int}(\mathsf{L})$, $\mathcal{I} \models \mathcal{T}$ iff $\mathcal{I} \models \mathcal{T}'$.

We will only deal with concept satisfiability as concept subsumption can be reduced to it for DLs that are closed under boolean operations: C subsumes D w.r.t. $\mathcal{T}$ iff $(D \sqcap \neg C)$ is not satisfiable w.r.t. $\mathcal{T}$.

For temporal or modal logics, satisfiability with respect to a set of formulae $\{C_1, \dots, C_k\}$ asserted to be universally true corresponds to satisfiability w.r.t. the TBox $\{\top \doteq C_1, \dots, \top \doteq C_n\}$.

Many decision procedures for DLs base their judgement on the existence of models or pseudo-models for concepts. A central rôle in these algorithms is played by a structure that we will call a *witness* in this paper. It generalises the notions of *tableaux* that appear in DL tableau-algorithms [HNS90, BBH96, HST99] as well as the *Hintikka-structures* that are used in tableau and automata-based decision procedures for temporal logic [EH85] and propositional dynamic logic [VW86].

Definition 2.3 (Witness) *Let* L *be a DL and* $C \in \mathsf{L}$ *a concept. A* witness $\mathcal{W} = (\Delta^{\mathcal{W}}, \cdot^{\mathcal{W}}, \mathcal{L}^{\mathcal{W}})$ *for* C *consists of a non-empty set* $\Delta^{\mathcal{W}}$, *a function* $\cdot^{\mathcal{W}}$ *that maps* NR *to* $2^{\Delta^{\mathcal{W}} \times \Delta^{\mathcal{W}}}$, *and a function* $\mathcal{L}^{\mathcal{W}}$ *that maps* $\Delta^{\mathcal{W}}$ *to* 2^{L} *such that the following properties are satisfied:*

(W1) *there is some* $x \in \Delta^{\mathcal{W}}$ *with* $C \in \mathcal{L}^{\mathcal{W}}(x)$,

(W2) *there is an interpretation* $\mathcal{I} \in \mathsf{Int}(\mathsf{L})$ *that* stems *from* $\mathcal{W}$, *and*

(W3) *for each interpretation* $\mathcal{I} \in \mathsf{Int}(\mathsf{L})$ *that* stems *from* $\mathcal{W}$, *it holds that* $D \in \mathcal{L}^{\mathcal{W}}(x)$ *implies* $x \in D^{\mathcal{I}}$.

An interpretation $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ *is said to stem from* $\mathcal{W}$ *if it satisfies:*

1. $\Delta^{\mathcal{I}} = \Delta^{\mathcal{W}}$,
2. $\cdot^{\mathcal{I}}|_{\mathsf{NR}} = \cdot^{\mathcal{W}}$, *and*
3. *for each* $A \in \mathsf{NC}$, $A \in \mathcal{L}^{\mathcal{W}}(x) \Rightarrow x \in A^{\mathcal{I}}$ *and* $\neg A \in \mathcal{L}^{\mathcal{W}}(x) \Rightarrow x \notin A^{\mathcal{I}}$.

A witness $\mathcal{W}$ *is called* admissible *with respect to a TBox* $\mathcal{T}$ *if there is an interpretation* $\mathcal{I} \in \mathsf{Int}(\mathsf{L})$ *that stems from* $\mathcal{W}$ *with* $\mathcal{I} \models \mathcal{T}$.

Please note that, for any witness $\mathcal{W}$, (W2) together with Condition 3 of "stemming" implies that, there exists no $x \in \Delta^{\mathcal{W}}$ and $A \in \mathsf{NC}$, such that $\{A, \neg A\} \subseteq \mathcal{L}^{\mathcal{W}}(x)$. Also note that, in general, more than one interpretation may stem from a witness. This is the case if, for an atomic concept $A \in \mathsf{NC}$ and an element $x \in \Delta^{\mathcal{W}}$, $\mathcal{L}^{\mathcal{W}}(x) \cap \{A, \neg A\} = \emptyset$ holds (because two interpretations $\mathcal{I}$ and $\mathcal{I}'$, with $x \in A^{\mathcal{I}}$ and $x \in \neg A^{\mathcal{I}'}$, could both stem from $\mathcal{W}$).

Obviously, each interpretation $\mathcal{I}$ gives rise to a special witness, called the *canonical witness*:

Definition 2.4 (Canonical Witness) *Let* L *be a DL. For any interpretation* $\mathcal{I} \in \mathsf{Int}(\mathsf{L})$ *we define the* canonical witness $\mathcal{W}_{\mathcal{I}} = (\Delta^{\mathcal{W}_{\mathcal{I}}}, \cdot^{\mathcal{W}_{\mathcal{I}}}, \mathcal{L}^{\mathcal{W}_{\mathcal{I}}})$ *as follows:*

$$\begin{aligned} \Delta^{\mathcal{W}_{\mathcal{I}}} &= \Delta^{\mathcal{I}} \\ \cdot^{\mathcal{W}_{\mathcal{I}}} &= \cdot^{\mathcal{I}}|_{\mathsf{NR}} \\ \mathcal{L}^{\mathcal{W}_{\mathcal{I}}} &= \lambda x.\{D \in \mathsf{L} \mid x \in D^{\mathcal{I}}\} \end{aligned}$$

The following elementary properties of a canonical witness will be useful in our considerations.

Lemma 2.5 *Let* L *be a DL,* $C \in \mathsf{L}$, *and* $\mathcal{T}$ *a TBox. For each* $\mathcal{I} \in \mathsf{Int}(\mathsf{L})$ *with* $C^{\mathcal{I}} \neq \emptyset$,

1. *each interpretation* $\mathcal{I}'$ *stemming from* $\mathcal{W}_{\mathcal{I}}$ *is isomorphic to* $\mathcal{I}$
2. $\mathcal{W}_{\mathcal{I}}$ *is a witness for* C,
3. $\mathcal{W}_{\mathcal{I}}$ *is admissible w.r.t.* $\mathcal{T}$ *iff* $\mathcal{I} \models \mathcal{T}$

Proof.

1. Let $\mathcal{I}'$ stem from $\mathcal{W}_{\mathcal{I}}$. This implies $\Delta^{\mathcal{I}'} = \Delta^{\mathcal{I}}$ and $\cdot^{\mathcal{I}'}|_{\mathsf{NR}} = \cdot^{\mathcal{I}}|_{\mathsf{NR}}$. For each $x \in \Delta^{\mathcal{I}}$ and $A \in \mathsf{NC}$, $\{A, \neg A\} \cap \mathcal{L}^{\mathcal{W}_{\mathcal{I}}}(x) \neq \emptyset$, this implies $\cdot^{\mathcal{I}'}|_{\mathsf{NC}} = \cdot^{\mathcal{I}}|_{\mathsf{NC}}$ and hence $\mathcal{I}$ and $\mathcal{I}'$ are isomorphic.

2. Properties (W1) and (W2) hold by construction. Obviously, $\mathcal{I}$ stems from $\mathcal{W}_\mathcal{I}$ and from (1) it follows that each interpretation $\mathcal{I}'$ stemming from $\mathcal{W}_\mathcal{I}$ is isomorphic to $\mathcal{I}$, hence (W3) holds.

3. Since $\mathcal{I}$ stems from $\mathcal{W}_\mathcal{I}$, $\mathcal{I} \models \mathcal{T}$ implies that $\mathcal{W}_\mathcal{I}$ is admissible w.r.t. $\mathcal{T}$. If $\mathcal{W}_\mathcal{I}$ is admissible w.r.t. $\mathcal{T}$, then there is an interpretation $\mathcal{I}'$ stemming from $\mathcal{W}_\mathcal{I}$ with $\mathcal{I}' \models \mathcal{T}$. Since $\mathcal{I}$ is isomorphic to $\mathcal{I}'$, this implies $\mathcal{I} \models \mathcal{T}$. ■

As a corollary we get that the existence of admissible witnesses is closely related to the satisfiability of concepts w.r.t. TBoxes:

Lemma 2.6 *Let* L *be a DL. A concept* $C \in \mathsf{L}$ *is satisfiable w.r.t. a TBox* $\mathcal{T}$ *iff it has a witness that is admissible w.r.t.* $\mathcal{T}$.

Proof. For the *only if*-direction let $\mathcal{I} \in \mathsf{Int}(\mathsf{L})$ be an interpretation with $\mathcal{I} \models \mathcal{T}$ and $C^\mathcal{I} \neq \emptyset$. From Lemma 2.5 it follows that the canonical witness $\mathcal{W}_\mathcal{I}$ is a witness for C that is admissible w.r.t. $\mathcal{T}$.

For the *if*-direction let $\mathcal{W}$ be an witness for C that is admissible w.r.t. $\mathcal{T}$. This implies that there is an interpretation $\mathcal{I} \in \mathsf{Int}(\mathsf{L})$ stemming from $\mathcal{W}$ with $\mathcal{I} \models \mathcal{T}$. For each interpretation $\mathcal{I}$ that stems from $\mathcal{W}$, it holds that $C^\mathcal{I} \neq \emptyset$ due to (W1) and (W3). ■

From this it follows that one can test the satisfiability of a concept w.r.t. to a TBox by checking for the existence of an admissible witness. We call algorithms that utilise this approach *model-building algorithms.*

This notion captures tableau-based decision procedures, [HNS90, BBH96, HST99], those using automata-theoretic approaches [VW86, CDL99] and, due to their direct correspondence with tableaux algorithms [HS99, BFH+99], even resolution based and sequent calculus algorithms.

The way many decision procedures for DLs deal with TBoxes exploits the following simple lemma.

Lemma 2.7 *Let* L *be a DL,* $C \in \mathsf{L}$ *a concept, and* $\mathcal{T}$ *a TBox. Let* $\mathcal{W}$ *be a witness for* C*. If*

$$\begin{array}{lll} C_1 \sqsubseteq C_2 \in \mathcal{T} & \Rightarrow & \forall x \in \Delta^\mathcal{W}.(C_1 \rightarrow C_2 \in \mathcal{L}^\mathcal{W}(x)) \\ C_1 \doteq C_2 \in \mathcal{T} & \Rightarrow & \forall x \in \Delta^\mathcal{W}.(C_1 \leftrightarrow C_2 \in \mathcal{L}^\mathcal{W}(x)) \end{array}$$

then $\mathcal{W}$ *is admissible w.r.t.* $\mathcal{T}$.

Proof. $\mathcal{W}$ is a witness, hence there is an interpretation $\mathcal{I} \in \mathsf{Int}(\mathsf{L})$ stemming from $\mathcal{W}$. From (W3) and the fact that $\mathcal{W}$ satisfies the properties stated in 2.7 it follows that, for each $x \in \Delta^\mathcal{I}$,

$$\begin{array}{rll} C_1 \sqsubseteq C_2 \in \mathcal{T} & \Rightarrow & C_1 \rightarrow C_2 \in \mathcal{L}^\mathcal{W}(x) \\ & \Rightarrow & x \in (C_1 \rightarrow C_2)^\mathcal{I} \\ C_1 \doteq C_2 \in \mathcal{T} & \Rightarrow & C_1 \leftrightarrow C_2 \in \mathcal{L}^\mathcal{W}(x) \\ & \Rightarrow & x \in (C_1 \leftrightarrow C_2)^\mathcal{I} \end{array}$$

Hence, $\mathcal{I} \models \mathcal{T}$ and $\mathcal{W}$ is admissible w.r.t. $\mathcal{T}$. ■

Examples of algorithms that exploit this lemma to deal with axioms can be found in [DDM96, DL96, HST99], where, for each axiom $C_1 \sqsubseteq C_2$ ($C_1 \doteq C_2$) the concept $C_1 \rightarrow C_2$ ($C_1 \leftrightarrow C_2$) is added to every node of the generated tableau.

Dealing with general axioms in this manner is costly due to the high degree of nondeterminism introduced. This can best be understood by looking at tableaux algorithms, which try to build witnesses in an incremental fashion. For a concept C to be tested for satisfiability, they start with $\Delta^\mathcal{W} = \{x_0\}$, $\mathcal{L}^\mathcal{W}(x_0) = \{C\}$ and $\cdot^\mathcal{W}(R) = \emptyset$ for each $R \in \mathsf{NR}$. Subsequently, the concepts in $\mathcal{L}^\mathcal{W}$ are decomposed and, if necessary, new nodes are added to $\Delta^\mathcal{W}$, until either $\mathcal{W}$ is a witness for C, or an obvious contradiction of the form $\{A, \neg A\} \subseteq \mathcal{L}^\mathcal{W}(x)$, which violates (W2), is generated. In the latter case, backtracking search is used to explore alternative non-deterministic decompositions (e.g., of disjunctions), one of which could lead to the discovery of a witness.

When applying Lemma 2.7, disjunctions are added to the label of each node of the tableau for each general axiom in the TBox (one disjunction for axioms of the form $C_1 \sqsubseteq C_2$, two for axioms of the form $C_1 \doteq C_2$). This leads to an exponential increase in the search space as the number of nodes and axioms increases. For example, with 10 nodes and a TBox containing 10 general axioms (of the form $C_1 \sqsubseteq C_2$) there are already 100 disjunctions, and they can be non-deterministically decomposed in 2^{100} different ways. For a TBox containing large numbers of general axioms (there are 1,214 in the GALEN medical terminology KB [RNG93]), this can degrade performance to the extent that subsumption testing is effectively non-terminating. To reason with this kind of TBox we must find a more efficient way to deal with axioms.

3 ABSORPTIONS

We start our considerations with an analysis of a technique that can be used to deal more efficiently with so-called primitive or acyclic TBoxes.

Definition 3.1 (Absorption) *Let* L *be a DL and* $\mathcal{T}$ *a TBox. An* absorption *of* $\mathcal{T}$ *is a pair of TBoxes*

$(\mathcal{T}_u, \mathcal{T}_g)$ *such that* $\mathcal{T} \equiv \mathcal{T}_u \cup \mathcal{T}_g$ *and* $\mathcal{T}_u$ *contains only axioms of the form* $A \sqsubseteq D$ *and* $\neg A \sqsubseteq D$ *where* $A \in \mathsf{NC}$.

An absorption $(\mathcal{T}_u, \mathcal{T}_g)$ *of* $\mathcal{T}$ *is called* correct *if it satisfies the following condition. For each witness* $\mathcal{W}$, *if, for each* $x \in \Delta^{\mathcal{W}}$,

$$\begin{aligned}
A \sqsubseteq D \in \mathcal{T}_u \wedge A \in \mathcal{L}^{\mathcal{W}}(x) &\Rightarrow D \in \mathcal{L}^{\mathcal{W}}(x)\\
\neg A \sqsubseteq D \in \mathcal{T}_u \wedge \neg A \in \mathcal{L}^{\mathcal{W}}(x) &\Rightarrow D \in \mathcal{L}^{\mathcal{W}}(x)\\
C_1 \sqsubseteq C_2 \in \mathcal{T}_g &\Rightarrow C_1 \to C_2 \in \mathcal{L}^{\mathcal{W}}(x)\\
C_1 \doteq C_2 \in \mathcal{T}_g &\Rightarrow C_1 \leftrightarrow C_2 \in \mathcal{L}^{\mathcal{W}}(x)
\end{aligned}$$

then $\mathcal{W}$ *is admissible w.r.t.* $\mathcal{T}$. *We refer to this properties by* $(*)$. *A witness that satisfies* $(*)$ *will be called* unfolded w.r.t. $\mathcal{T}$.

If the reference to a specific TBox is clear from the context, we will often leave the TBox implicit and say that a witness is unfolded.

How does a correct absorption enable an algorithm to deal with axioms more efficiently? This is best described by returning to tableaux algorithms. Instead of dealing with axioms as previously described, which may lead to an exponential increase in the search space, axioms in $\mathcal{T}_u$ can now be dealt with in a deterministic manner. Assume, for example, that we have to handle the axiom $A \doteq C$. If the label of a node already contains A (resp. $\neg A$), then C (resp. $\neg C$) is added to the label; if the label contains neither A nor $\neg A$, then *nothing* has to be done. Dealing with the axioms in $\mathcal{T}_u$ this way avoids the necessity for additional non-deterministic choices and leads to a gain in efficiency. A witness produced in this manner will be unfolded and is a certificate for satisfiability w.r.t. $\mathcal{T}$. This technique is generally known as *lazy unfolding* of primitive TBoxes [Hor98]; formally, it is justified by the following lemma:

Lemma 3.2 *Let* $(\mathcal{T}_u, \mathcal{T}_g)$ *be a correct absorption of* $\mathcal{T}$. *For any* $C \in \mathsf{L}$, C *has a witness that is admissible w.r.t.* $\mathcal{T}$ *iff* C *has an unfolded witness.*

Proof. The *if*-direction follows from the definition of "correct absorption". For the *only if*-direction, let $C \in \mathsf{L}$ be a concept and $\mathcal{W}$ a witness for C that is admissible w.r.t. $\mathcal{T}$. This implies the existence of an interpretation $\mathcal{I} \in \mathsf{Int}(\mathsf{L})$ stemming from $\mathcal{W}$ such that $\mathcal{I} \models \mathcal{T}$ and $C^{\mathcal{I}} \neq \emptyset$. Since $\mathcal{T} \equiv \mathcal{T}_u \cup \mathcal{T}_g$ we have $\mathcal{I} \models \mathcal{T}_u \cup \mathcal{T}_g$ and hence the canonical witness $\mathcal{W}_{\mathcal{I}}$ is an unfolded witness for C. ∎

A family of TBoxes where absorption can successfully be applied are *primitive* TBoxes, the most simple form of TBox usually studied in the literature.

Definition 3.3 (Primitive TBox) *A TBox* $\mathcal{T}$ *is called* primitive *iff it consists entirely of axioms of the form* $A \doteq D$ *with* $A \in \mathsf{NC}$, *each* $A \in \mathsf{NC}$ *appears as at most one left-hand side of an axiom, and* $\mathcal{T}$ *is acyclic. Acyclicity is defined as follows:* $A \in \mathsf{NC}$ *is said to* directly use $B \in \mathsf{NC}$ *if* $A \doteq D \in \mathcal{T}$ *and* B *occurs in* D; uses *is the transitive closure of "directly uses". We say that* $\mathcal{T}$ *is* acyclic *if there is no* $A \in \mathsf{NC}$ *that uses itself.*

For primitive TBoxes a correct absorption can easily be given.

Theorem 3.4 *Let* $\mathcal{T}$ *be a primitive TBox,* $\mathcal{T}_g = \emptyset$, *and* $\mathcal{T}_u$ *defined by*

$$\mathcal{T}_u = \{A \sqsubseteq D, \neg A \sqsubseteq \neg D \mid A \doteq D \in \mathcal{T}\}.$$

Then $(\mathcal{T}_u, \mathcal{T}_g)$ *is a correct absorption of* $\mathcal{T}$.

Proof. Trivially, $\mathcal{T} \equiv \mathcal{T}_u \cup \mathcal{T}_g$ holds. Given an unfolded witness $\mathcal{W}$, we have to show that there is an interpretation $\mathcal{I}$ stemming from $\mathcal{W}$ with $\mathcal{I} \models \mathcal{T}$.

We fix an arbitrary linearisation $A_1, \ldots, A_k$ of the "uses" partial order on the atomic concept names appearing on the left-hand sides of axioms in $\mathcal{T}$ such that, if A_i uses A_j, then $j < i$ and the defining concept for A_i is D_i.

For some interpretation $\mathcal{I}$, atomic concept A, and set $X \subseteq \Delta^{\mathcal{I}}$, we denote the interpretation that maps A to X and agrees with $\mathcal{I}$ on all other atomic concepts and roles by $\mathcal{I}[A \mapsto X]$. For $0 \leq i \leq k$, we define $\mathcal{I}_i$ in an iterative process starting from an arbitrary interpretation $\mathcal{I}_0$ stemming from $\mathcal{W}$ and setting

$$\mathcal{I}_i := \mathcal{I}_{i-1}[A_i \mapsto \{x \in \Delta^{\mathcal{W}} \mid x \in D_i^{\mathcal{I}_{i-1}}\}]$$

Since, for each A_i there is exactly one axiom in $\mathcal{T}$, each step in this process is well-defined. Also, since $\mathsf{Int}(\mathsf{L})$ may only restrict the interpretation of atomic roles, $\mathcal{I}_i \in \mathsf{Int}(\mathsf{L})$ for each $0 \leq i \leq k$. For $\mathcal{I} = \mathcal{I}_k$ it can be shown that $\mathcal{I}$ is an interpretation stemming from $\mathcal{W}$ with $\mathcal{I} \models \mathcal{T}$.

First we prove inductively that, for $0 \leq i \leq k$, $\mathcal{I}_i$ stems from $\mathcal{W}$. We have already required $\mathcal{I}_0$ to stem from $\mathcal{W}$.

Assume the claim was proved for $\mathcal{I}_{i-1}$ and $\mathcal{I}_i$ does not stem from $\mathcal{W}$. Then there must be some $x \in \Delta^{\mathcal{W}}$ such that either (i) $A_i \in \mathcal{L}^{\mathcal{W}}(x)$ but $x \notin A_i^{\mathcal{I}_i}$ or (ii) $\neg A_i \in \mathcal{L}^{\mathcal{W}}(x)$ but $x \in A_i^{\mathcal{I}_i}$ (since we assume $\mathcal{I}_{i-1}$ to stem from $\mathcal{W}$ and A_i is the only atomic concept whose interpretation changes from $\mathcal{I}_{i-1}$ to $\mathcal{I}_i$). The two cases can be handled dually:

(i) From $A_i \in \mathcal{L}^{\mathcal{W}}(x)$ it follows that $D_i \in \mathcal{L}^{\mathcal{W}}(x)$, because $\mathcal{W}$ is unfolded. Since $\mathcal{I}_{i-1}$ stems from $\mathcal{W}$ and $\mathcal{W}$ is a witness, Property (W3) implies $x \in D_i^{\mathcal{I}_{i-1}}$. But this implies $x \in A_i^{\mathcal{I}_i}$, which is a contradiction.

(ii) From $\neg A_i \in \mathcal{L}^{\mathcal{W}}(x)$ it follows that $\neg D_i \in \mathcal{L}^{\mathcal{W}}(x)$ because $\mathcal{W}$ is unfolded. Since $\mathcal{I}_{i-1}$ stems from $\mathcal{W}$ and $\mathcal{W}$ is an witness, Property (W3) implies $x \in (\neg D_i)^{\mathcal{I}_{i-1}}$. Since $(\neg D_i)^{\mathcal{I}_{i-1}} = \Delta^{\mathcal{W}} \setminus D_i^{\mathcal{I}_{i-1}}$ this implies $x \notin A_i^{\mathcal{I}_i}$, which is a contradiction.

Together this implies that $\mathcal{I}_i$ also stems from $\mathcal{W}$.

To show that $\mathcal{I} \models \mathcal{T}$ we show inductively that $\mathcal{I}_i \models A_j \doteq D_j$ for each $1 \leq j \leq i$. This is obviously true for $i = 0$.

The interpretation of D_i may not depend on the interpretation of A_i because otherwise (I2) would imply that A_i uses itself. Hence $D_i^{\mathcal{I}_i} = D_i^{\mathcal{I}_{i-1}}$ and, by construction, $\mathcal{I}_i \models A_i \doteq D_i$. Assume there is some $j < i$ such that $\mathcal{I}_i \not\models A_j \doteq D_j$. Since $\mathcal{I}_{i-1} \models A_j \doteq D_j$ and only the interpretation of A_i has changed from $\mathcal{I}_{i-1}$ to $\mathcal{I}_i$, $D_j^{\mathcal{I}_i} \neq D_j^{\mathcal{I}_{i-1}}$ must hold because of (I2). But this implies that A_i occurs in D_j and hence A_j uses A_i which contradicts $j < i$. Thus, we have $\mathcal{I} \models A_j = D_j$ for each $1 \leq j \leq k$ and hence $\mathcal{I} \models \mathcal{T}$. ∎

Lazy unfolding is a well-known and widely used technique for optimising reasoning w.r.t. primitive TBoxes [BFH+94]. So far, we have only given a correctness proof for this relatively simple approach, although one that is independent of a specific DL or reasoning algorithm. With the next lemma we show how we can extend correct absorptions and hence how lazy unfolding can be applied to a broader class of TBoxes. A further enhancement of the technique is presented in Section 5.

Lemma 3.5 *Let $(\mathcal{T}_u, \mathcal{T}_g)$ be a correct absorption of a TBox $\mathcal{T}$.*

1. *If $\mathcal{T}'$ is an arbitrary TBox, then $(\mathcal{T}_u, \mathcal{T}_g \cup \mathcal{T}')$ is a correct absorption of $\mathcal{T} \cup \mathcal{T}'$.*
2. *If $\mathcal{T}'$ is a TBox that consists entirely of axioms of the form $A \sqsubseteq D$, where $A \in \mathsf{NC}$ and A is not defined in $\mathcal{T}_u$, then $(\mathcal{T}_u \cup \mathcal{T}', \mathcal{T}_g)$ is a correct absorption of $\mathcal{T} \cup \mathcal{T}'$.*

Proof. In both cases, $\mathcal{T}_u \cup \mathcal{T}_g \cup \mathcal{T}' \equiv \mathcal{T} \cup \mathcal{T}'$ holds trivially.

1. Let $C \in \mathsf{L}$ be a concept and $\mathcal{W}$ be an unfolded witness for C w.r.t. the absorption $(\mathcal{T}_u, \mathcal{T}_g \cup \mathcal{T}')$. This implies that $\mathcal{W}$ is unfolded w.r.t. the (smaller) absorption $(\mathcal{T}_u, \mathcal{T}_g)$. Since $(\mathcal{T}_u, \mathcal{T}_g)$ is a correct absorption, there is an interpretation $\mathcal{I}$ stemming from $\mathcal{W}$ with $\mathcal{I} \models \mathcal{T}$. Assume $\mathcal{I} \not\models \mathcal{T}'$. Then, without loss of generality,[3] there is an axiom $D \sqsubseteq E \in \mathcal{T}'$ such that there exists an $x \in D^{\mathcal{I}} \setminus E^{\mathcal{I}}$. Since $\mathcal{W}$ is unfolded, we have $D \rightarrow E \in \mathcal{L}^{\mathcal{W}}(x)$ and hence (W3) implies $x \in (\neg D \sqcup E)^{\mathcal{I}} = \Delta^{\mathcal{I}} \setminus (D^{\mathcal{I}} \setminus E^{\mathcal{I}})$, a contradiction. Hence $\mathcal{I} \models \mathcal{T} \cup \mathcal{T}'$ and $\mathcal{W}$ is admissible w.r.t. $\mathcal{T} \cup \mathcal{T}'$.
2. Let $C \in \mathsf{L}$ be a concept and $\mathcal{W}$ be an unfolded witness for C w.r.t. the absorption $(\mathcal{T}_u \cup \mathcal{T}', \mathcal{T}_g)$. From $\mathcal{W}$ we define a new witness $\mathcal{W}'$ for C by setting $\Delta^{\mathcal{W}'} := \Delta^{\mathcal{W}}$, $\cdot^{\mathcal{W}'} := \cdot^{\mathcal{W}}$, and definig $\mathcal{L}^{\mathcal{W}'}$ to be the function that, for every $x \in \Delta^{\mathcal{W}'}$, maps x to the set
$$\mathcal{L}^{\mathcal{W}}(x) \cup \{\neg A \mid A \sqsubseteq D \in \mathcal{T}', A \notin \mathcal{L}^{\mathcal{W}}(x)\}$$
It is easy to see that $\mathcal{W}'$ is indeed a witness for C and that $\mathcal{W}'$ is also unfolded w.r.t. the absorption $(\mathcal{T}_u \cup \mathcal{T}', \mathcal{T}_g)$. This implies that $\mathcal{W}'$ is also unfolded w.r.t. the (smaller) absorption $(\mathcal{T}_u, \mathcal{T}_g)$. Since $(\mathcal{T}_u, \mathcal{T}_g)$ is a correct absorption of $\mathcal{T}$, there exists an interpretation $\mathcal{I}$ stemming from $\mathcal{W}'$ such that $\mathcal{I} \models \mathcal{T}$. We will show that $\mathcal{I} \models \mathcal{T}'$ also holds. Assume $\mathcal{I} \not\models \mathcal{T}'$, then there is an axiom $A \sqsubseteq D \in \mathcal{T}'$ and an $x \in \Delta^{\mathcal{I}}$ such that $x \in A^{\mathcal{I}}$ but $x \notin D^{\mathcal{I}}$. By construction of $\mathcal{W}'$, $x \in A^{\mathcal{I}}$ implies $A \in \mathcal{L}^{\mathcal{W}'}(x)$ because otherwise $\neg A \in \mathcal{L}^{\mathcal{W}'}(x)$ would hold in contradiction to (W3). Then, since $\mathcal{W}'$ is unfolded, $D \in \mathcal{L}^{\mathcal{W}'}(x)$, which, again by (W3), implies $x \in D^{\mathcal{I}}$, a contradiction.

   Hence, we have shown that there exists an interpretation $\mathcal{I}$ stemming from $\mathcal{W}'$ such that $\mathcal{I} \models \mathcal{T}_u \cup \mathcal{T}' \cup \mathcal{T}_g$. By construction of $\mathcal{W}'$, any interpretation stemming from $\mathcal{W}'$ also stems from $\mathcal{W}$, hence $\mathcal{W}$ is admissible w.r.t. $\mathcal{T} \cup \mathcal{T}'$. ∎

4 APPLICATION TO FaCT

In the preceeding section we have defined correct absorptions and discussed how they can be exploited in order to optimise satisfiability procedures. However, we have said nothing about the problem of how to find an absorption given an arbitrary terminology. In this section we will describe the absorption algorithm used by FaCT and prove that it generates correct absorptions.

[3] Arbitrary TBoxes can be expressed using only axioms of the form $C \sqsubseteq D$.

Given a TBox $\mathcal{T}$ containing arbitrary axioms, the absorption algorithm used by FaCT constructs a triple of TBoxes $(\mathcal{T}_g, \mathcal{T}_{\text{prim}}, \mathcal{T}_{\text{inc}})$ such that

- $\mathcal{T} \equiv \mathcal{T}_g \cup \mathcal{T}_{\text{prim}} \cup \mathcal{T}_{\text{inc}}$,
- $\mathcal{T}_{\text{prim}}$ is primitive, and
- $\mathcal{T}_{\text{inc}}$ consists only of axioms of the form $A \sqsubseteq D$ where $A \in \mathsf{NC}$ and A is not defined in $\mathcal{T}_{\text{prim}}$.

We refer to these properties by $(*)$. From Theorem 3.4 together with Lemma 3.5 it follows that, for

$$\mathcal{T}_u := \{A \sqsubseteq D, \neg A \sqsubseteq \neg D \mid A \doteq D \in \mathcal{T}_{\text{prim}}\} \cup \mathcal{T}_{\text{inc}}$$

$(\mathcal{T}_u, \mathcal{T}_g)$ is a correct absorption of $\mathcal{T}$; hence satisfiability for a concept C w.r.t. $\mathcal{T}$ can be decided by checking for an unfolded witness for C.

In a first step, FaCT distributes axioms from $\mathcal{T}$ amongst $\mathcal{T}_{\text{inc}}$, $\mathcal{T}_{\text{prim}}$, and $\mathcal{T}_g$, trying to minimise the number of axioms in $\mathcal{T}_g$ while still maintaining $(*)$. To do this, it initialises $\mathcal{T}_{\text{prim}}, \mathcal{T}_{\text{inc}}$, and $\mathcal{T}_g$ with $\emptyset$, and then processes each axiom $X \in \mathcal{T}$ as follows.

1. If X is of the form $A \sqsubseteq C$, then
   (a) if $A \in \mathsf{NC}$ and A is not defined in $\mathcal{T}_{\text{prim}}$ then X is added to $\mathcal{T}_{\text{inc}}$,
   (b) otherwise X is added to $\mathcal{T}_g$
2. If X is of the form $A \doteq C$, then
   (a) if $A \in \mathsf{NC}$, A is not defined in $\mathcal{T}_{\text{prim}}$ or $\mathcal{T}_{\text{inc}}$ and $\mathcal{T}_{\text{prim}} \cup \{X\}$ is primitive, then X is added to $\mathcal{T}_{\text{prim}}$,
   (b) otherwise, the axioms $A \sqsubseteq C$ and $C \sqsubseteq A$ are added to $\mathcal{T}_g$

It is easy to see that the resulting TBoxes $\mathcal{T}_g, \mathcal{T}_{\text{prim}}, \mathcal{T}_{\text{inc}}$ satisfy $(*)$. In a second step, FaCT processes the axioms in $\mathcal{T}_g$ one at a time, trying to absorb them into axioms in $\mathcal{T}_{\text{inc}}$. Those axioms that are not absorbed remain in $\mathcal{T}_g$. To give a simpler formulation of the algorithm, each axiom $(C \sqsubseteq D) \in \mathcal{T}_g$ is viewed as a clause $\mathbf{G} = \{D, \neg C\}$, corresponding to the axiom $\top \sqsubseteq C \rightarrow D$, which is equivalent to $C \sqsubseteq D$. For each such axiom FaCT applies the following absorption procedure.

1. Try to absorb $\mathbf{G}$. If there is a concept $\neg A \in \mathbf{G}$ such that $A \in \mathsf{NC}$ and A is not defined in $\mathcal{T}_{\text{prim}}$, then add $A \sqsubseteq B$ to $\mathcal{T}_{\text{inc}}$, where B is the disjunction of all the concepts in $\mathbf{G} \setminus \{\neg A\}$, remove $\mathbf{G}$ from $\mathcal{T}_g$, and exit.
2. Try to simplify $\mathbf{G}$.
   (a) If there is some $\neg C \in \mathbf{G}$ such that C is of the form $C_1 \sqcap \ldots \sqcap C_n$, then substitute $\neg C$ with $\neg C_1 \sqcup \ldots \sqcup \neg C_n$, and continue with step 2b.
   (b) If there is some $C \in \mathbf{G}$ such that C is of the form $(C_1 \sqcup \ldots \sqcup C_n)$, then apply associativity by setting $\mathbf{G} = \mathbf{G} \cup \{C_1, \ldots, C_n\} \setminus \{(C_1 \sqcup \ldots \sqcup C_n)\}$, and return to step 1.
3. Try to unfold $\mathbf{G}$. If, for some $A \in \mathbf{G}$ (resp. $\neg A \in \mathbf{G}$), there is an axiom $A \doteq C$ in $\mathcal{T}_{\text{prim}}$, then substitute $A \in \mathbf{G}$ (resp. $\neg A \in \mathbf{G}$) with C (resp. $\neg C$) and return to step 1.
4. If none of the above were possible, then absorption of $\mathbf{G}$ has failed. Leave $\mathbf{G}$ in $\mathcal{T}_g$, and exit.

For each step, we have to show that $(*)$ is maintained. Dealing with clauses instead of axioms causes no problems. In the first step, axioms are moved from $\mathcal{T}_g$ to $\mathcal{T}_{\text{inc}}$ as long as this does not violate $(*)$. The second and the third step replace a clause by an equivalent one and hence do not violate $(*)$.

Termination of the procedure is obvious. Each axiom is considered only once and, for a given axiom, simplification and unfolding can only be applied finitely often before the procedure is exited, either by absorbing the axiom into $\mathcal{T}_{\text{inc}}$ or leaving it in $\mathcal{T}_g$. For simplification, this is obvious; for unfolding, this holds because $\mathcal{T}_{\text{prim}}$ is acyclic. Hence, we get the following:

Theorem 4.1 *For any TBox $\mathcal{T}$, FaCT computes a correct absorption of $\mathcal{T}$.*

5 IMPROVING PERFORMANCE

The absorption algorithm employed by FaCT already leads to a dramatic improvement in performance. This is illustrated by Figure 1, which shows the times taken by FaCT to classify versions of the GALEN KB with some or all of the general axioms removed. Without absorption, classification time increased rapidly with the number of general axioms, and exceeded 10,000s with only 25 general axioms in the KB; with absorption, only 160s was taken to classify the KB with all 1,214 general axioms.

However, there is still considerable scope for further gains. In particular, the following definition for a *stratified* TBox allows lazy unfolding to be more generally applied, while still allowing for correct absorptions.

Definition 5.1 (Stratified TBox) *A TBox $\mathcal{T}$ is called* stratified *iff it consists entirely of axioms of the*

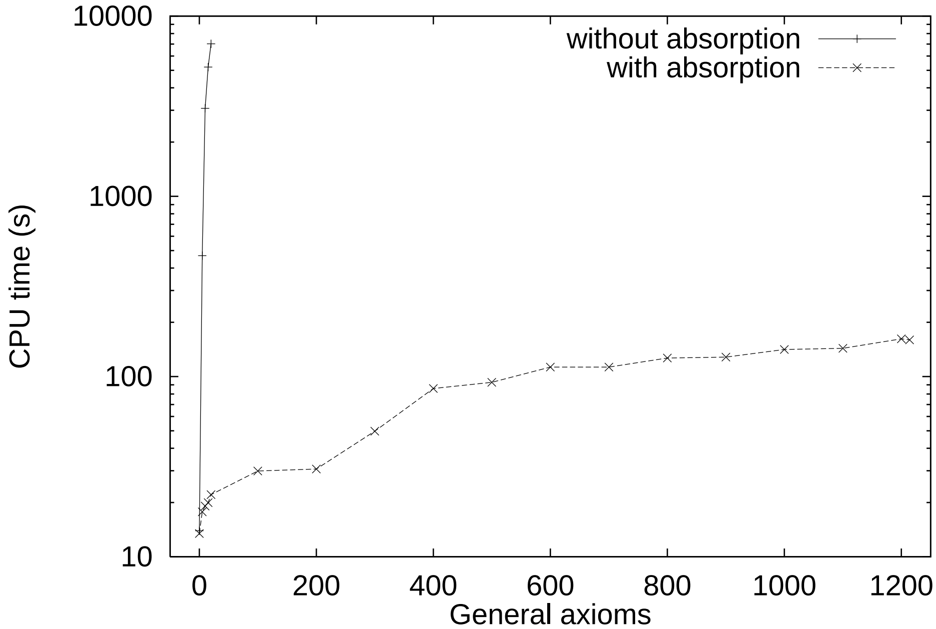

Figure 1: Classification times with and without absorption

form $A \doteq D$ with $A \in \mathsf{NC}$, each $A \in \mathsf{NC}$ appears at most once on the left-hand side of an axiom, and $\mathcal{T}$ can be arranged monotonously, i.e., there is a disjoint partition $\mathcal{T}_1 \dot\cup \mathcal{T}_2 \dot\cup \ldots \dot\cup \mathcal{T}_k$ of $\mathcal{T}$, such that

- *for all $1 \leq j < i \leq k$, if $A \in \mathsf{NC}$ is defined in $\mathcal{T}_i$, then it does not occur in $\mathcal{T}_j$, and*
- *for all $1 \leq i \leq k$, all concepts which appear on the right-hand side of axioms in $\mathcal{T}_i$ are monotone in all atomic concepts defined in $\mathcal{T}_i$.*

A concept C is monotone in an atomic concept A if, for any interpretation $\mathcal{I} \in \mathsf{Int(L)}$ and any two sets $X_1, X_2 \subseteq \Delta^{\mathcal{I}}$,

$$X_1 \subseteq X_2 \;\Rightarrow\; C^{\mathcal{I}[A \mapsto X_1]} \subseteq C^{\mathcal{I}[A \mapsto X_2]}.$$

For many DLs, a sufficient condition for monotonicity is *syntactic* monotonicity, i.e., a concept C is syntactically monotone in some atomic concept A if A does no appear in C in the scope of an odd number of negations.

Obviously, due to its acyclicity, every primitive TBox is also stratified and hence the following theorem is a strict generalisation of Theorem 3.4.

Theorem 5.2 *Let $\mathcal{T}$ be a stratified TBox, $\mathcal{T}_g = \emptyset$ and $\mathcal{T}_u$ defined by*

$$\mathcal{T}_u = \{A \sqsubseteq D, \neg A \sqsubseteq \neg D \mid A \doteq D \in \mathcal{T}\}.$$

Then $(\mathcal{T}_u, \mathcal{T}_g)$ is a correct absorption of $\mathcal{T}$.

The proof of this theorem follows the same line as the proof of Theorem 3.4. Starting from an arbitrary interpretation $\mathcal{I}_0$ stemming from the unfolded witness, we incrementally construct interpretations $\mathcal{I}_1, \ldots, \mathcal{I}_k$, using a fixed point construction in each step. We show that each $\mathcal{I}_i$ stems from $\mathcal{W}$ and that, for $1 \leq j < i \leq k$, $\mathcal{I}_i \models \mathcal{T}_j$, hence $\mathcal{I}_k \models \mathcal{T}$ and stems from $\mathcal{W}$.

Before we prove this theorem, we recall some basics of lattice theory. For any set $\mathcal{S}$, the powerset of $\mathcal{S}$, denoted by $2^{\mathcal{S}}$ forms a complete lattice, where the ordering, join and meet operations are set-inclusion $\subseteq$, union $\cup$, and intersection $\cap$, respectively. For any complete lattice $\mathcal{L}$, its n-fold cartesian product $\mathcal{L}^n$ is also a complete lattice, with ordering, join, and meet defined in a pointwise manner.

For a lattice $\mathcal{L}$, a function $\Phi : \mathcal{L} \to \mathcal{L}$ is called monotone, iff, for $x_1, x_2 \in \mathcal{L}$, $x_1 \sqsubseteq x_2$ implies $\Phi(x_1) \sqsubseteq \Phi(x_2)$.

By Tarski's fixed point theorem [Tar55], every monotone function Φ on a complete lattice, has uniquely defined least and greatest fixed points, i.e., there are elements $\overline{x}, \underline{x} \in \mathcal{L}$ such that

$$\overline{x} = \Phi(\overline{x}) \text{ and } \underline{x} = \Phi(\underline{x})$$

and, for all $x \in \mathcal{L}$ with $x = \Phi(x)$,

$$\underline{x} \sqsubseteq x \text{ and } x \sqsubseteq \overline{x}.$$

Proof of Theorem 5.2. $\mathcal{T}_u \cup \mathcal{T}_g \equiv \mathcal{T}$ is obvious. Let $\mathcal{W} = (\Delta^{\mathcal{W}}, \cdot^{\mathcal{W}}, \mathcal{L}^{\mathcal{W}})$ be an unfolded witness. We have to show that there is an interpretation $\mathcal{I}$ stemming from $\mathcal{W}$ with $\mathcal{I} \models \mathcal{T}$. Let $\mathcal{T}_1, \ldots, \mathcal{T}_k$ be the required partition of $\mathcal{T}$. We will define $\mathcal{I}$ inductively, starting with an arbitrary interpretation $\mathcal{I}_0$ stemming from $\mathcal{W}$.

Assume $\mathcal{I}_{i-1}$ was already defined. We define $\mathcal{I}_i$ from $\mathcal{I}_{i-1}$ as follows: let $\{A_1^i \doteq D_1^i, \ldots, A_m^i \doteq D_m^i\}$ be an enumeration of $\mathcal{T}_i$. First we need some auxiliary notation: for any concept $C \in \mathsf{L}$ we define

$$C^{\mathcal{W}} := \{x \in \Delta^{\mathcal{W}} \mid C \in \mathcal{L}^{\mathcal{W}}(x)\}.$$

Using this notation we define the function Φ mapping subsets $X_1, \ldots, X_m$ of $\Delta^{\mathcal{W}}$ to

$$\begin{array}{l}(\; ((A_1^i)^{\mathcal{W}} \cup (D_1^i)^{\mathcal{I}_{i-1}(X_1,\ldots,X_m)}) \setminus (\neg A_1^i)^{\mathcal{W}}, \\ \quad \ldots, \\ \quad ((A_m^i)^{\mathcal{W}} \cup (D_m^i)^{\mathcal{I}_{i-1}(X_1,\ldots,X_m)}) \setminus (\neg A_m^i)^{\mathcal{W}} \;)\end{array}$$

where

$$\mathcal{I}_{i-1}(X_1, \ldots, X_m) := \mathcal{I}_{i-1}[A_1^i \mapsto X_1, \ldots, A_m^i \mapsto X_m]$$

Since all of the D_j^i are monotone in all of the A_m^i, Φ is a monontone function. This implies that Φ has a least fixed point, which we denote by $(\underline{X}_1, \ldots, \underline{X}_m)$. We use this fixed point to define $\mathcal{I}_i$ by

$$\mathcal{I}_i := \mathcal{I}_{i-1}[A_1^i \mapsto \underline{X}_1, \ldots, A_m^i \mapsto \underline{X}_m]$$

Claim 1: For each $0 \leq i \leq k$, $\mathcal{I}_i$ stems from $\mathcal{W}$.

We show this claim by induction on i. We have already required $\mathcal{I}_0$ to stem from $\mathcal{W}$. Assume $\mathcal{I}_{i-1}$ stems from $\mathcal{W}$. Since the only thing that changes from $\mathcal{I}_{i-1}$ to $\mathcal{I}_i$ is the interpretation of the atomic concepts $A_1^i, \ldots, A_m^i$, we only have to check that $A_j^i \in \mathcal{L}^{\mathcal{W}}(x)$ implies $x \in (A_j^i)^{\mathcal{I}_i}$ and $\neg A_j^i \in \mathcal{L}^{\mathcal{W}}(x)$ implies $x \notin (A_j^i)^{\mathcal{I}_i}$.

By definition of Φ, and because $\{x \mid A_j^i \in \mathcal{L}^{\mathcal{W}}(x)\} \cap \{x \mid \neg A_j^i \in \mathcal{L}^{\mathcal{W}}(x)\} = \emptyset$, $A_j^i \in \mathcal{L}^{\mathcal{W}}(x)$ implies $x \in (A_j^i)^{\mathcal{I}_i}$. Also by the definition of Φ, $\neg A_j^i \in \mathcal{L}^{\mathcal{W}}(x)$ implies $x \notin (A_j^i)^{\mathcal{I}_i}$. Hence, $\mathcal{I}_i$ stems from $\mathcal{W}$.

Claim 2: For each $1 \leq j \leq i \leq k$, $\mathcal{I}_i \models \mathcal{T}_j$.

We prove this claim by induction over i starting from 0. For $i = 0$, there is nothing to prove. Assume the claim would hold for $\mathcal{I}_{i-1}$. The only thing that changes from $\mathcal{I}_{i-1}$ to $\mathcal{I}_i$ is the interpretation of the atomic concepts $A_1^i, \ldots A_m^i$ defined in $\mathcal{T}_i$. Since these concepts may not occur in $\mathcal{T}_j$ for $j < i$, the interpretation of the concepts in these TBoxes does not change, and from $\mathcal{I}_{i-1} \models \mathcal{T}_j$ follows $\mathcal{I}_i \models \mathcal{T}_j$ for $1 \leq j \leq i - 1$.

It remains to show that $\mathcal{I}_i \models \mathcal{T}_i$. Let $A_j^i \doteq D_j^i$ be an axiom from $\mathcal{T}_i$. From the definition of $\mathcal{I}_i$ we have

$$(A_j^i)^{\mathcal{I}_i} = ((A_j^i)^{\mathcal{W}} \cup (D_j^i)^{\mathcal{I}_i}) \setminus (\neg A_j^i)^{\mathcal{W}}. \tag{1}$$

$\mathcal{W}$ is unfolded, hence $A_j^i \in \mathcal{L}^{\mathcal{W}}(x)$ implies $D_j^i \in \mathcal{L}^{\mathcal{W}}(x)$ and, since $\mathcal{I}_i$ stems from $\mathcal{W}$, this implies $x \in (D_j^i)^{\mathcal{I}_i}$, thus

$$(A_j^i)^{\mathcal{W}} \cup (D_j^i)^{\mathcal{I}_i} = (D_j^i)^{\mathcal{I}_i} \tag{2}$$

Furthermore, $\neg A_j^i \in \mathcal{L}^{\mathcal{W}}(x)$ implies $\neg D_j^i \in \mathcal{L}^{\mathcal{W}}(x)$ implies $x \in (\neg D_j^i)^{\mathcal{I}_i}$, thus

$$(D_j^i)^{\mathcal{I}_i} \setminus (\neg A_j^i)^{\mathcal{W}} = (D_j^i)^{\mathcal{I}_i} \tag{3}$$

Taking together (1), (2), and (3) we get

$$(A_j^i)^{\mathcal{I}_i} = (D_j^i)^{\mathcal{I}_i},$$

and hence $\mathcal{I}_i \models A_j^i \doteq D_j^i$.

Together, Claim 1 and Claim 2 prove the theorem, since $\mathcal{I}_k$ is an interpretation that stems from $\mathcal{W}$ and satisfies $\mathcal{T}$. ■

This theorem makes it possible to apply the same lazy unfolding strategy as before to cyclical definitions. Such definitions are quite natural in a logic that supports inverse roles. For example, an orthopaedic procedure might be defined as a procedure performed by an orthopaedic surgeon, while an orthopaedic surgeon might be defined as a surgeon who performs only orthopaedic procedures:[4]

$$\begin{aligned}\text{o-procedure} &\doteq \text{procedure} \sqcap (\exists \textit{performs}^-.\text{o-surgeon}) \\ \text{o-surgeon} &\doteq \text{surgeon} \sqcap (\forall \textit{performs}.\text{o-procedure})\end{aligned}$$

The absorption algorithm described in Section 4 would force the second of these definitions to be added to $\mathcal{T}_g$ as two general axioms and, although both axioms would subsequently be absorbed into $\mathcal{T}_u$, the procedure would result in a disjunctive term being added to one of the definitions in $\mathcal{T}_u$. Using Theorem 5.2 to enhance the absorption algorithm so that these kinds of definition are directly added to $\mathcal{T}_u$ reduces the number of disjunctive terms in $\mathcal{T}_u$ and can lead to significant improvements in performance.

This can be demonstrated by a simple experiment with the new FaCT system, which implements the $\mathcal{SHIQ}$

[4]This example is only intended for didactic purposes.

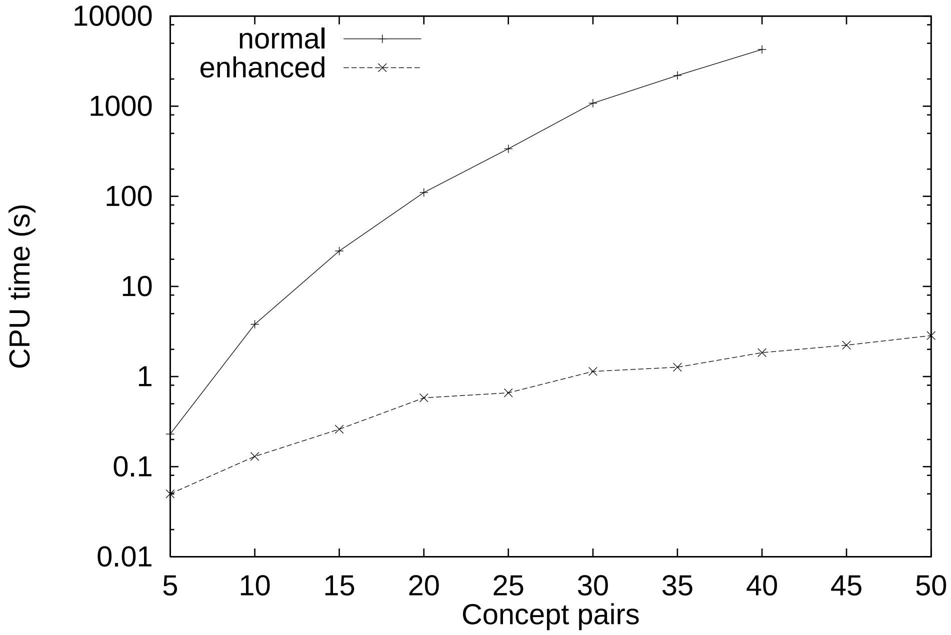

Figure 2: Classification times with and without enhanced absorption

logic [HST99] and is thus able to deal with inverse roles. Figure 2 shows the classification time in seconds using the normal and enhanced absorption algorithms for terminologies consisting of between 5 and 50 pairs of cyclical definitions like those described above for o-surgeon and o-procedure. With only 10 pairs the gain in performance is already a factor of 30, while for 45 and 50 pairs it has reached several orders of magnitude: with the enhanced absorption the terminology is classified in 2–3 seconds whereas with the original algorithm the time required exceeded the 10,000 second limit imposed in the experiment.

It is worth pointing out that it is by no means trivially true that cyclical definitions can be dealt with by lazy unfolding. Even without inverse roles it is clear that definitions such as $A \doteq \neg A$ (or more subtle variants) force the domain to be empty and would lead to an incorrect absorption if dealt with by lazy unfolding. With converse roles it is, for example, possible to force the interpretation of a role R to be empty with a definition such as $A \doteq \forall R.(\forall R^{-}.\neg A)$, again leading to an incorrect absorption if dealt with by lazy unfolding.

6 OPTIMAL ABSORPTIONS

We have demonstrated that absorption is a highly effective and widely applicable technique, and by formally defining correctness criteria for absorptions we have proved that the procedure used by FaCT finds correct absorptions. Moreover, by establishing more precise correctness criteria we have demonstrated how the effectiveness of this procedure could be further enhanced.

However, the absorption algorithm used by FaCT is clearly sub-optimal, in the sense that changes could be made that would, in general, allow more axioms to be absorbed (e.g., by also giving special consideration to axioms of the form $\neg A \sqsubseteq C$ with $A \in \mathsf{NC}$). Moreover, the procedure is non-deterministic, and, while it is guaranteed to produce a correct absorption, its specific result depends on the order of the axioms in the original TBox $\mathcal{T}$. Since the semantics of a TBox $\mathcal{T}$ does not depend on the order of its axioms, there is no reason to suppose that they will be arranged in a way that yields a "good" absorption. Given the effectiveness of absorption, it would be desirable to have an algorithm that was guaranteed to find the "best" absorption possible for any set of axioms, irrespective of their ordering in the TBox.

Unfortunately, it is not even clear how to define a sensible optimality criterion for absorptions. It is obvious that simplistic approaches based on the number or size of axioms remaining in $\mathcal{T}_g$ will not lead to a useful solution for this problem. Consider, for example, the cyclical TBox experiment from the previous section. Both the original FaCT absorption algorithm and the enhanced algorithm, which exploits Theorem 5.2, are able to compute a complete absorption of the axioms (i.e., a correct absorption with $\mathcal{T}_g = \emptyset$), but the enhanced algorithm leads to much better performance, as shown in Figure 2.

An important issue for future work is, therefore, the identification of a suitable optimality criterion for absorptions, and the development of an algorithm that is able to compute absorptions that are optimal with respect to this criterion.

Acknowledgements

This work was partially supported by the DFG, Project No. GR 1324/3-1.

References

[Baa91] F. Baader. Augmenting concept languages by transitive closure of roles: An alternative to terminological cycles. In *Proc. of the 12th International Joint Conference on Artificial Intelligence (IJCAI-91)*, pages 446–451, 1991.

[BBH96] F. Baader, M. Buchheit, and B. Hollunder. Cardinality restrictions on concepts. *Artificial Intelligence*, 88(1–2):195–213, 1996.

[BDS93] M. Buchheit, F. M. Donini, and A. Schaerf. Decidable reasoning in terminological knowledge representation systems. *J. of Artificial Intelligence Research*, 1:109–138, 1993.

[BFH+94] F. Baader, E. Franconi, B. Hollunder, B. Nebel, and H.-J. Profitlich. An empirical analysis of optimization techniques for terminological representation systems, or: Making KRIS get a move on. *Applied Artificial Intelligence*, 4:109–132, 1994.

[BFH+99] A. Borgida, E. Franconi, I. Horrocks, D. McGuinness, and P. F. Patel-Schneider. Explaining $\mathcal{ALC}$ subsumption. In P. Lambrix, A. Borgida, M. Lenzerini, R. Möller, and P. Patel-Schneider, editors, *Proc. of the International Workshop on Description Logics (DL'99)*, pages 37–40, 1999.

[BH91] F. Baader and B. Hollunder. A terminological knowledge representation system with complete inference algorithms. In *Processing declarative knowledge: International workshop PDK'91*, number 567 in Lecture Notes in Artificial Intelligence, pages 67–86, Berlin, 1991. Springer-Verlag.

[BHH+91] F. Baader, H.-J. Heinsohn, B. Hollunder, J. Muller, B. Nebel, W. Nutt, and H.-J. Profitlich. Terminological knowledge representation: A proposal for a terminological logic. Technical Memo TM-90-04, Deutsches Forschungszentrum für Künstliche Intelligenz GmbH (DFKI), 1991.

[Cal96] D. Calvanese. Reasoning with inclusion axioms in description logics: Algorithms and complexity. In Wolfgang Wahlster, editor, *Proc. of the 12th European Conference on Artificial Intelligence (ECAI'96)*, pages 303–307. John Wiley & Sons Ltd., 1996.

[CDL99] D. Calvanese, G. De Giacomo, and M. Lenzerini. Reasoning in expressive description logics with fixpoints based on automata on infinite trees. In *Proc. of the 16th Int. Joint Conf. on Artificial Intelligence (IJCAI'99)*, 1999.

[DDM96] F. Donini, G. De Giacomo, and F. Massacci. EXPTIME tableaux for $\mathcal{ALC}$. In L. Padgham, E. Franconi, M. Gehrke, D. L. McGuinness, and P. F. Patel-Schneider, editors, *Collected Papers from the International Description Logics Workshop (DL'96)*, number WS-96-05 in AAAI Technical Report, pages 107–110. AAAI Press, Menlo Park, California, 1996.

[DL94] G. De Giacomo and M. Lenzerini. Boosting the correspondence between description logics and propositional dynamic logics. In *Proc. of the 12th Nat. Conf. on Artificial Intelligence (AAAI'94)*, pages 205–212. AAAI Press/The MIT Press, 1994.

[DL96] G. De Giacomo and M. Lenzerini. TBox and ABox reasoning in expressive description logics. In Luigia C. Aiello, John Doyle, and Stuart C. Shapiro, editors, *Proc. of the 5th Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR'96)*, pages 316–327. Morgan Kaufmann, Los Altos, 1996.

[DM98] G. De Giacomo and F. Massacci. Combining deduction and model checking into tableaux and algorithms for converse-PDL.

Information and Computation: special issue on the Federated Logic Conferences. Academic Press, 1998.

[EH85] E. A. Emerson and J. Y. Halpern. Decision procedures and expressiveness in the temporal logic of branching time. *J. of Computer and System Sciences*, 30:1–24, 1985.

[HN90] B. Hollunder and W. Nutt. Subsumption algorithms for concept languages. In *Proc. of the 9th European Conference on Artificial Intelligence (ECAI'90)*, pages 348–353. John Wiley & Sons Ltd., 1990.

[HNS90] B. Hollunder, W. Nutt, and M. Schmidt-Schauss. Subsumption algorithms for concept description languages. In *ECAI-90*, Pitman Publishing, London, 1990.

[Hor97] I. Horrocks. *Optimising Tableaux Decision Procedures for Description Logics.* PhD thesis, University of Manchester, 1997.

[Hor98] I. Horrocks. Using an expressive description logic: FaCT or fiction? In A. G. Cohn, L. Schubert, and S. C. Shapiro, editors, *Principles of Knowledge Representation and Reasoning: Proc. of the Sixth International Conference (KR'98)*, pages 636–647. Morgan Kaufmann Publishers, San Francisco, California, June 1998.

[HS99] U. Hustadt and R. A. Schmidt. On the relation of resolution and tableaux proof systems for description logics. In *Proc. of the 16th International Joint Conference on Artificial Intelligence (IJCAI-99)*, pages 110–115, 1999.

[HST99] I. Horrocks, U. Sattler, and S. Tobies. Practical reasoning for expressive description logics. In *Proc. of the 6th International Conference on Logic for Programming and Automated Reasoning (LPAR'99)*, pages 161–180, 1999.

[RNG93] A. L. Rector, W A Nowlan, and A Glowinski. Goals for concept representation in the GALEN project. In *Proc. of the 17th Annual Symposium on Computer Applications in Medical Care (SCAMC'93)*, pages 414–418, Washington DC, USA, 1993.

[Sch91] K. Schild. A correspondence theory for terminological logics: Preliminary report. In *Proc. of the 12th Int. Joint Conf. on Artificial Intelligence (IJCAI'91)*, pages 466–471, Sydney, 1991.

[SS91] M. Schmidt-Schauß and G. Smolka. Attributive concept descriptions with complements. *Acta Informatica*, 48(1):1–26, 1991.

[Tar55] A. Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific J. of Mathematics*, 5:285–309, 1955.

[VW86] M. Y. Vardi and P. Wolper. Automata-theoretic techniques for modal logics of programs. *J. of Computer and System Sciences*, 32:183–221, 1986. A preliminary version appeared in *Proc. of the 16th ACM SIGACT Symp. on Theory of Computing (STOC'84)*.

[WS92] W. A. Woods and J. G. Schmolze. The KL-ONE family. *Computers and Mathematics with Applications – Special Issue on Artificial Intelligence*, 23(2–5):133–177, 1992.