

Efficient Reasoning with Range and Domain Constraints

Dmitry Tsarkov and Ian Horrocks

Department of Computer Science
The University of Manchester
Manchester, UK
{tsarkov|horrocks}@cs.man.ac.uk

Abstract. We show how a tableaux algorithm for $SHIQ$ can be extended to support role boxes that include range and domain axioms, prove that the extended algorithm is still a decision procedure for the satisfiability and subsumption of $SHIQ$ concepts w.r.t. such a role box, and show how support for range and domain axioms can be exploited in order to add a new form of absorption optimisation called role absorption. We illustrate the effectiveness of the optimised algorithm by analysing the performance of our FaCT++ implementation when classifying terminologies derived from realistic ontologies.

1 Introduction

Many modern ontology languages (e.g., OIL [5], DAML+OIL [12] and OWL [4]) are based on expressive description logics, and in particular on the $SHIQ$ family of description logics [14]. These ontology languages typically support domain and range constraints on roles, i.e., axioms asserting that if an individual x is related to an individual y by a role R , then x must be an instance of the concept that is the domain of R and y must be an instance of the concept that is the range of R [4, 17]. Such axioms are not directly supported by $SHIQ$, but can trivially be transformed into *general inclusion axioms* (GCIs), i.e., an axiom asserting a subsumption relationship between two arbitrary concept terms. In particular, restricting the domain of a role R to be concept C is equivalent to adding an axiom asserting that the concept whose instances are related to some other individual by role R is subsumed by C (i.e., $\exists R.\top \sqsubseteq C$), and restricting the range of a role R to be concept D is equivalent to adding an axiom asserting that the most general concept is subsumed by the concept whose instances are related by role R only to instances of D (i.e., $\top \sqsubseteq \forall R.D$) [11].

The problem with this transformation is that such GCIs are not amenable to *absorption*, an optimisation technique that tries to rewrite GCIs so that they can be efficiently dealt with using the *lazy unfolding* optimisation [9]. Absorption is one of the crucial optimisations that enable state of the art DL reasoners such as FaCT [10], Racer [7] and Pellet [18] to deal effectively with large knowledge bases (KBs), and these reasoners perform much less well with KBs containing significant numbers of unabsorbable GCIs. Unfortunately, many ontologies contain large numbers of different roles, each with a range and domain constraint, and the resulting KBs therefore contain many unabsorbable GCIs.

It has already been shown that, in order for the Racer system to be able to classify¹ large KBs containing many range and domain constraints, it is necessary to give a special treatment to the GCIs introduced by range and domain axioms [6]. The approach used by Racer is to extend the lazy unfolding optimisation so that concepts equivalent to those that would be introduced by the GCIs are introduced only as necessary. In the approach presented here, we extend the tableaux satisfiability testing algorithm so that range and domain axioms are directly supported. The advantage with this approach is that we are able to extend the formal correctness proof to demonstrate that the extended algorithm is still a decision procedure for \mathcal{SHIQ} satisfiability (i.e., it returns *satisfiable* iff the input concept is satisfiable).

As well as allowing range and domain to be dealt with very efficiently, this algorithm also allows us to implement an extended version of the absorption optimisation, called *role absorption*, that transforms GCIs into domain constraints. Role absorption can provide alternative and perhaps more effective ways to absorb certain forms of GCI, and can also be applied to some otherwise unabsorbable forms of GCI. This can lead to dramatic performance improvements for KBs that contain significant numbers of such GCIs. We demonstrate this (as well as the performance improvements resulting from support for range and domain axioms) with an empirical analysis of the performance of the extended algorithm when classifying several KBs derived from realistic ontologies.

2 Preliminaries

We first introduce the syntax and semantics of the \mathcal{SHIQ} logic, including the semantics of role boxes extended with range and domain axioms. Most details of the logic and the tableaux algorithm are little changed from those presented in [14]. We will, therefore, focus mainly on the parts that have been added in order to deal with range and domain axioms, and refer the reader to [14] for complete information on the remainder.

Definition 1. Let \mathbf{C} and \mathbf{R} be disjoint sets of concept names and role names respectively. The set of \mathcal{SHIQ} -roles is $\mathbf{R} \cup \{R^- \mid R \in \mathbf{R}\}$. To avoid considering roles such as R^{--} , we define a function Inv on roles such that $\text{Inv}(R) = R^-$ if R is a role name, and $\text{Inv}(R) = S$ if $R = S^-$. For R and S \mathcal{SHIQ} -roles and C a \mathcal{SHIQ} -concept, a role axiom is either a role inclusion of the form $R \sqsubseteq S$, a transitivity axiom of the form $\text{Trans}(R)$, or a constraint axiom of the form $\text{Domain}(R, C)$ or $\text{Range}(R, C)$. A role box \mathcal{R} is a finite set of role axioms.

A role R is called *simple* if, for \sqsubseteq^* the transitive reflexive closure of \sqsubseteq on \mathcal{R} and for each role S , $S \sqsubseteq^* R$ implies $\text{Trans}(S) \notin \mathcal{R}$ and $\text{Trans}(\text{Inv}(S)) \notin \mathcal{R}$.

The set of concepts is the smallest set such that every concept name is a concept, and, for C and D concepts, R a role, S a simple role and n a non-negative integer, then $C \sqcap D$, $C \sqcup D$, $\neg C$, $\exists R.C$, $\forall R.C$, $\geq n.S.C$ and $\leq n.S.C$ are also concepts.

The semantics is given by means of an interpretation $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ consisting of a non-empty set $\Delta^{\mathcal{I}}$, called the domain of \mathcal{I} , and a valuation $\cdot^{\mathcal{I}}$ which maps every concept to a subset of $\Delta^{\mathcal{I}}$ and every role to a subset of $\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$ such that, for all concepts C ,

¹ I.e., to compute the subsumption partial ordering of named concepts in a KB.

Concepts & Roles	Syntax	Semantics
atomic concept C	A	$A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$
atomic role R	R	$R^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$
inverse role	R^{-}	$\{\langle x, y \rangle \mid \langle y, x \rangle \in R^{\mathcal{I}}\}$
conjunction	$C \sqcap D$	$(C \sqcap D)^{\mathcal{I}} = C^{\mathcal{I}} \cap D^{\mathcal{I}}$
disjunction	$C \sqcup D$	$(C \sqcup D)^{\mathcal{I}} = C^{\mathcal{I}} \cup D^{\mathcal{I}}$
negation	$\neg C$	$(\neg C)^{\mathcal{I}} = \Delta^{\mathcal{I}} \setminus C^{\mathcal{I}}$
exists restriction	$\exists R.C$	$(\exists R.C)^{\mathcal{I}} = \{x \mid \exists y. \langle x, y \rangle \in R^{\mathcal{I}} \text{ and } y \in C^{\mathcal{I}}\}$
value restriction	$\forall R.C$	$(\forall R.C)^{\mathcal{I}} = \{x \mid \forall y. \langle x, y \rangle \in R^{\mathcal{I}} \text{ implies } y \in C^{\mathcal{I}}\}$
atleast restriction	$\geq n.S.C$	$(\geq n.S.C)^{\mathcal{I}} = \{x \mid \#\{y. \langle x, y \rangle \in S^{\mathcal{I}}\} \cap C^{\mathcal{I}} \geq n\}$
atmost restriction	$\leq n.S.C$	$(\leq n.S.C)^{\mathcal{I}} = \{x \mid \#\{y. \langle x, y \rangle \in S^{\mathcal{I}}\} \cap C^{\mathcal{I}} \leq n\}$
Role Axioms	Syntax	Semantics
role inclusion	$R \sqsubseteq S$	$R^{\mathcal{I}} \subseteq S^{\mathcal{I}}$
transitive role	$\text{Trans}(R)$	$R^{\mathcal{I}} = (R^+)^{\mathcal{I}}$
role domain	$\text{Domain}(R, C)$	$\langle x, y \rangle \in R^{\mathcal{I}} \text{ implies } x \in C^{\mathcal{I}}$
role range	$\text{Range}(R, C)$	$\langle x, y \rangle \in R^{\mathcal{I}} \text{ implies } y \in C^{\mathcal{I}}$

Fig. 1. Syntax and semantics of *SHIQ*

D , roles R , S , and non-negative integers n , the properties in Figure 1 are satisfied, where $\#M$ denotes the cardinality of a set M .

An interpretation satisfies a role axiom if it satisfies the semantic conditions given in Figure 1. An interpretation satisfies a role box \mathcal{R} if it satisfies each role axiom in \mathcal{R} .

A terminology or TBox \mathcal{T} is a finite set of general concept inclusion axioms, $\mathcal{T} = \{C_1 \sqsubseteq D_1, \dots, C_n \sqsubseteq D_n\}$, where C_i, D_i are arbitrary *SHIQ*-concepts. An interpretation \mathcal{I} satisfies \mathcal{T} iff $C_i^{\mathcal{I}} \subseteq D_i^{\mathcal{I}}$ holds for all $C_i \sqsubseteq D_i \in \mathcal{T}$.

A *SHIQ*-concept C is satisfiable w.r.t. a role box \mathcal{R} and a terminology \mathcal{T} iff there is an interpretation \mathcal{I} with $C^{\mathcal{I}} \neq \emptyset$ that satisfies both \mathcal{R} and \mathcal{T} . Such an interpretation is called a model of C w.r.t. \mathcal{R} and \mathcal{T} . A concept C is subsumed by a concept D w.r.t. \mathcal{R} and \mathcal{T} iff $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$ for each interpretation \mathcal{I} satisfying \mathcal{R} and \mathcal{T} .

Theorem 1. Satisfiability and subsumption of *SHIQ*-concepts w.r.t. terminologies and role boxes is polynomially reducible to (un)satisfiability of *SHIQ*-concepts w.r.t. role boxes.

Subsumption problems can trivially be reduced to satisfiability problems: $C \sqsubseteq D$ iff $D \sqcap \neg C$ is not satisfiable. A procedure called *internalisation* can be used to reduce (polynomially) a satisfiability problem for a *SHIQ*-concept w.r.t. a terminology and role box to a satisfiability problem for a *SHIQ*-concept w.r.t. a role box [14].

3 Tableaux Reasoning with Range and Domain

Here we present an algorithm for deciding the satisfiability of a *SHIQ*-concept C w.r.t. a role box \mathcal{R} ; it is an extension of the *SHIQ* tableaux algorithm from [14].

For ease of Tableaux construction, we assume C and all concepts in (range and domain axioms in) \mathcal{R} to be in *negation normal form* (NNF), that is, negation occurs

only in front of concept names. Any *SHIQ*-concept can easily be transformed into an equivalent one in NNF by pushing negations inwards [8]; with $\sim C$ we denote the NNF of $\neg C$. We define $\text{RD}(\mathcal{R})$ as the set of concepts s.t. $C \in \text{RD}(\mathcal{R})$ iff $\text{Domain}(R, C) \in \mathcal{R}$ or $\text{Range}(R, C) \in \mathcal{R}$ for some role R . We define $\text{c1}(C, \mathcal{R})$ as the smallest set of concepts that is a superset of $C \cup \text{RD}(\mathcal{R})$ and is closed under subconcepts and \sim .

Definition 2. Let D be a *SHIQ*-concept in NNF, \mathcal{R} a role box, and \mathbf{R}_D the set of roles occurring in D and \mathcal{R} together with their inverses. Then $T = (\mathbf{S}, \mathcal{L}, \mathcal{E})$ is a tableau for D w.r.t. \mathcal{R} iff \mathbf{S} is a set of individuals, $\mathcal{L} : \mathbf{S} \rightarrow 2^{\text{c1}(D, \mathcal{R})}$ maps each individual to a set of concepts, $\mathcal{E} : \mathbf{R}_D \rightarrow 2^{\mathbf{S} \times \mathbf{S}}$ maps each role to a set of pairs of individuals, and there is some individual $s \in \mathbf{S}$ such that $D \in \mathcal{L}(s)$. Furthermore, for all $s, t \in \mathbf{S}$, $C, C_1, C_2 \in \text{c1}(D, \mathcal{R})$, and $R, S \in \mathbf{R}_D$, it holds that:

1. if $C \in \mathcal{L}(s)$, then $\neg C \notin \mathcal{L}(s)$,
2. if $C_1 \sqcap C_2 \in \mathcal{L}(s)$, then $C_1 \in \mathcal{L}(s)$ and $C_2 \in \mathcal{L}(s)$,
3. if $C_1 \sqcup C_2 \in \mathcal{L}(s)$, then $C_1 \in \mathcal{L}(s)$ or $C_2 \in \mathcal{L}(s)$,
4. if $\forall S.C \in \mathcal{L}(s)$ and $\langle s, t \rangle \in \mathcal{E}(S)$, then $C \in \mathcal{L}(t)$,
5. if $\exists S.C \in \mathcal{L}(s)$, then there is some $t \in \mathbf{S}$ such that $\langle s, t \rangle \in \mathcal{E}(S)$ and $C \in \mathcal{L}(t)$,
6. if $\forall S.C \in \mathcal{L}(s)$ and $\langle s, t \rangle \in \mathcal{E}(R)$ for some $R \sqsubseteq S$ with $\text{Trans}(R)$, then $\forall R.C \in \mathcal{L}(t)$,
7. $\langle s, t \rangle \in \mathcal{E}(R)$ iff $\langle t, s \rangle \in \mathcal{E}(\text{Inv}(R))$,
8. if $\langle s, t \rangle \in \mathcal{E}(R)$ and $R \sqsubseteq S$, then $\langle s, t \rangle \in \mathcal{E}(S)$,
9. if $(\leq n S C) \in \mathcal{L}(s)$, then $\#S^T(s, C) \leq n$,
10. if $(\geq n S C) \in \mathcal{L}(s)$, then $\#S^T(s, C) \geq n$,
11. if $(\bowtie n S C) \in \mathcal{L}(s)$ and $\langle s, t \rangle \in \mathcal{E}(S)$ then $C \in \mathcal{L}(t)$ or $\sim C \in \mathcal{L}(t)$,
12. if $\langle s, t \rangle \in \mathcal{E}(S)$ and $\text{Domain}(S, C) \in \mathcal{R}$, then $C \in \mathcal{L}(s)$,
13. if $\langle s, t \rangle \in \mathcal{E}(S)$ and $\text{Range}(S, C) \in \mathcal{R}$, then $C \in \mathcal{L}(t)$,

where we use \bowtie as a placeholder for both \leq and \geq and we define

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

Lemma 1. A *SHIQ*-concept D is satisfiable w.r.t. a role box \mathcal{R} iff D has a tableau w.r.t. \mathcal{R} .

Proof. The proof found in [14] is easily extended to deal with range and domain axioms. When constructing a model from a tableau, conditions (12) and (13) ensure that the semantics of range and domain axioms are satisfied by the model. Note that if the tableau contains a set of pairs $\{\langle s_1, s_2 \rangle, \dots, \langle s_{n-1}, s_n \rangle\} \subseteq \mathcal{E}(S)$ for some transitive role S , then the interpretation of S will include the transitive edge $\langle s_1, s_n \rangle$. In this case, however, $\langle s_1, s_2 \rangle \in \mathcal{E}(S)$ and condition (12) mean that the semantics of domain axioms will still be satisfied w.r.t. $\langle s_1, s_n \rangle$, while $\langle s_{n-1}, s_n \rangle \in \mathcal{E}(S)$ and condition (13) mean that the semantics of range axioms will still be satisfied w.r.t. $\langle s_1, s_n \rangle$. Similarly, the semantics ensure that we can easily transform a model into a tableau. \square

3.1 An Extended Tableaux Algorithm

In order to make the following description easier, we will abuse notation by using $\text{Domain}(R)$ and $\text{Range}(R)$ to mean the sets of concepts corresponding to the domain and range axioms in \mathcal{R} that apply to a role R , i.e., $\text{Domain}(R) = \{C \mid \text{Domain}(R, C) \in \mathcal{R}\}$, and $\text{Range}(R) = \{C \mid \text{Range}(R, C) \in \mathcal{R}\}$.

Definition 3. A completion tree for a concept D is a tree where each node x of the tree is labelled with a set $\mathcal{L}(x) \subseteq \text{c1}(D, \mathcal{R})$ and each edge $\langle x, y \rangle$ is labelled with a set $\mathcal{L}(\langle x, y \rangle)$ of (possibly inverse) roles occurring in $\text{c1}(D, \mathcal{R})$; explicit inequalities between nodes of the tree are recorded in a binary relation \neq that is implicitly assumed to be symmetric.

Given a completion tree, a node y is called an R -successor of a node x iff y is a successor of x and $S \in \mathcal{L}(\langle x, y \rangle)$ for some S with $S \sqsubseteq R$. A node y is called an R -neighbour of x iff y is an R -successor of x , or if x is an $\text{Inv}(R)$ -successor of y . Predecessors and ancestors are defined as usual.

A node is blocked iff it is directly or indirectly blocked. A node x is directly blocked iff none of its ancestors are blocked, and it has ancestors x' , y and y' such that

1. x is a successor of x' and y is a successor of y' and
2. $\mathcal{L}(x) = \mathcal{L}(y)$ and $\mathcal{L}(x') = \mathcal{L}(y')$ and
3. $\mathcal{L}(\langle x', x \rangle) = \mathcal{L}(\langle y', y \rangle)$.

A node y is indirectly blocked iff one of its ancestors is blocked, or it is a successor of a node x and $\mathcal{L}(\langle x, y \rangle) = \emptyset$.²

For a node x , $\mathcal{L}(x)$ is said to contain a clash iff $\{A, \neg A\} \subseteq \mathcal{L}(x)$ or if, for some concept C , some role S , and some $n \in \mathbb{N}$: $(\leq n S C) \in \mathcal{L}(x)$ and there are $n + 1$ S -neighbours y_0, \dots, y_n of x such that $C \in \mathcal{L}(y_i)$ and $y_i \neq y_j$ for all $0 \leq i < j \leq n$. A completion tree is called clash-free iff none of its nodes contains a clash; it is called complete iff none of the expansion rules in Figure 2 is applicable.

For a \mathcal{SHIQ} -concept D , the algorithm starts with a completion tree consisting of a single node x with $\mathcal{L}(x) = \{D\}$ and $\neq = \emptyset$. It applies the expansion rules in Figure 2, stopping when a clash occurs, and answers “ D is satisfiable” iff the completion rules can be applied in such a way that they yield a complete and clash-free completion tree.

Note that the only change w.r.t. [14] is addition of the *domain* and *range*-rules that add concepts to node labels as required by domain and range axioms.

Lemma 2. Let D be an \mathcal{SHIQ} -concept.

1. The tableaux algorithm terminates when started with D .
2. If the expansion rules can be applied to D such that they yield a complete and clash-free completion tree, then D has a tableau.
3. If D has a tableau, then the expansion rules can be applied to D such that they yield a complete and clash-free completion tree.

² A more complex but more efficient form of blocking is described in [13].

\sqcap -rule:	if 1. $C_1 \sqcap C_2 \in \mathcal{L}(x)$, x is not indirectly blocked, and 2. $\{C_1, C_2\} \not\subseteq \mathcal{L}(x)$ then $\mathcal{L}(x) \longrightarrow \mathcal{L}(x) \cup \{C_1, C_2\}$
\sqcup -rule:	if 1. $C_1 \sqcup C_2 \in \mathcal{L}(x)$, x is not indirectly blocked, and 2. $\{C_1, C_2\} \cap \mathcal{L}(x) = \emptyset$ then $\mathcal{L}(x) \longrightarrow \mathcal{L}(x) \cup \{C\}$ for some $C \in \{C_1, C_2\}$
\exists -rule:	if 1. $\exists S.C \in \mathcal{L}(x)$, x is not blocked, and 2. x has no S -neighbour y with $C \in \mathcal{L}(y)$, then create a new node y with $\mathcal{L}(\langle x, y \rangle) = \{S\}$ and $\mathcal{L}(y) = \{C\}$
\forall -rule:	if 1. $\forall S.C \in \mathcal{L}(x)$, x is not indirectly blocked, and 2. there is an S -neighbour y of x with $C \notin \mathcal{L}(y)$ then $\mathcal{L}(y) \longrightarrow \mathcal{L}(y) \cup \{C\}$
\forall_+ -rule:	if 1. $\forall S.C \in \mathcal{L}(x)$, x is not indirectly blocked, and 2. there is some R with $\text{Trans}(R)$ and $R \boxplus S$, 3. there is an R -neighbour y of x with $\forall R.C \notin \mathcal{L}(y)$ then $\mathcal{L}(y) \longrightarrow \mathcal{L}(y) \cup \{\forall R.C\}$
<i>choose</i> -rule:	if 1. $(\boxtimes n S C) \in \mathcal{L}(x)$, x is not indirectly blocked, and 2. there is an S -neighbour y of x with $\{C, \sim C\} \cap \mathcal{L}(y) = \emptyset$ then $\mathcal{L}(y) \longrightarrow \mathcal{L}(y) \cup \{E\}$ for some $E \in \{C, \sim C\}$
\geq -rule:	if 1. $(\geq n S C) \in \mathcal{L}(x)$, x is not blocked, and 2. there are not n S -neighbours y_1, \dots, y_n of x with $C \in \mathcal{L}(y_i)$ and $y_i \neq y_j$ for $1 \leq i < j \leq n$ then create n new nodes y_1, \dots, y_n with $\mathcal{L}(\langle x, y_i \rangle) = \{S\}$, $\mathcal{L}(y_i) = \{C\}$, and $y_i \neq y_j$ for $1 \leq i < j \leq n$.
\leq -rule:	if 1. $(\leq n S C) \in \mathcal{L}(x)$, x is not indirectly blocked, and 2. $\#S^T(x, C) > n$ and there are two S -neighbours y, z of x with $C \in \mathcal{L}(y)$, $C \in \mathcal{L}(z)$, y is not an ancestor of x , and not $y \neq z$ then 1. $\mathcal{L}(z) \longrightarrow \mathcal{L}(z) \cup \mathcal{L}(y)$ and 2. if z is an ancestor of x then $\mathcal{L}(\langle z, x \rangle) \longrightarrow \mathcal{L}(\langle z, x \rangle) \cup \text{Inv}(\mathcal{L}(\langle x, y \rangle))$ else $\mathcal{L}(\langle x, z \rangle) \longrightarrow \mathcal{L}(\langle x, z \rangle) \cup \mathcal{L}(\langle x, y \rangle)$ 3. $\mathcal{L}(\langle x, y \rangle) \longrightarrow \emptyset$ 4. Set $u \neq z$ for all u with $u \neq y$
<i>domain</i> -rule	if 1. $C \in \text{Domain}(S)$, x is not indirectly blocked, and 2. there is an S -neighbour y of x and $C \notin \mathcal{L}(y)$ then $\mathcal{L}(x) \longrightarrow \mathcal{L}(x) \cup \{C\}$
<i>range</i> -rule	if 1. $C \in \text{Range}(S)$, x is not indirectly blocked, and 2. there is an S -neighbour y of x with $C \notin \mathcal{L}(y)$ then $\mathcal{L}(y) \longrightarrow \mathcal{L}(y) \cup \{C\}$

Fig. 2. The complete tableaux expansion rules for \mathcal{SHIQ}

Proof. Again, only a small extension of the proof in [14] is required, and we will only consider the new conditions due to domain and range axioms.

For termination (claim 1), the *domain* and *range*-rules do not cause any new nodes to be added to the expansion tree, and nodes are still labeled with subsets of $c1(D, \mathcal{R})$, so the same arguments apply.

For soundness (claim 2), we can obtain a tableau $T = (\mathbf{S}, \mathcal{L}, \mathcal{E})$ from a complete and clash-free completion tree \mathbf{T} by *unravelling* \mathbf{T} in the usual way. For an edge $\langle s, t \rangle \in \mathcal{E}(S)$, s, t correspond to nodes x, y in \mathbf{T} with $\mathcal{L}(s) = \mathcal{L}(x)$ and $\mathcal{L}(t) = \mathcal{L}(y)$, and either y is an S -neighbour of x or there is an S -neighbour z of x s.t. y blocks z (note that, in the latter case, the definition of blocking means that $\mathcal{L}(y) = \mathcal{L}(z)$ and that y is also an S -neighbour of some node w). If property (12) in Definition 2 is *not* satisfied, then there is some domain constraint $\text{Domain}(S, C) \in \mathcal{R}$ s.t. $C \notin \mathcal{L}(s)$. In this case, however, the *domain*-rule would be applicable to x and \mathbf{T} would not be complete. A similar argument applies w.r.t. property (13) in Definition 2.

For completeness (claim 3), we can again use a tableau T to guide the application of the non-deterministic rules (i.e., the \sqcup , \leq and *choose*-rules) so that we obtain a complete and clash-free completion tree \mathbf{T} . We do this by defining (inductively) a mapping π from nodes in \mathbf{T} to individuals in T s.t. $\mathcal{L}(x) \subseteq \mathcal{L}(\pi(x))$ for $\pi(x) \in \mathbf{S}$, for each pair of nodes x, y and each role R , if y is an R -neighbour of x , then $\langle \pi(x), \pi(y) \rangle \in \mathcal{E}(R)$, and $x \neq y$ implies $\pi(x) \neq \pi(y)$. Properties (12) and (13) in Definition 2 ensure that applications of the *domain* and *range*-rules do not lead to a clash. \square

The following theorem is an immediate consequence of Lemmas 1, 2 and Theorem 1.

Theorem 2. *The tableaux algorithm is a decision procedure for the satisfiability and subsumption of SHIQ-concepts with respect to role boxes.*

4 Role Absorption

Given that the new algorithm is able to deal directly with range and domain axioms, it makes sense to transform GCIs of the form $\exists R.\top \sqsubseteq C$ and $\top \sqsubseteq \forall R.C$ into range and domain axioms respectively. We call this new form of absorption *role absorption* in contrast to the usual form of absorption we will refer to as *concept absorption* (see [15]).

Role absorption is important because in ontology derived KBs range and domain constraints will often have been transformed into GCIs. This is because tools such as OilEd [1] and Protégé [16] are designed to work with range of DL reasoners, some of which (e.g., FaCT) do not support range and domain axioms. Moreover, these forms of GCI are not, in general, amenable to standard concept absorption techniques.

This simple form of role absorption, which we will refer to as *basic role absorption*, is formalised in the following theorem:

Theorem 3. *Let \mathcal{R} be a SHIQ role box.*

1. *An interpretation \mathcal{I} satisfies \mathcal{R} and $\exists R.\top \sqsubseteq C$ iff \mathcal{I} satisfies $\mathcal{R} \cup \{\text{Domain}(R, C)\}$.*
2. *An interpretation \mathcal{I} satisfies \mathcal{R} and $\top \sqsubseteq \forall R.C$ iff \mathcal{I} satisfies $\mathcal{R} \cup \{\text{Range}(R, C)\}$.*

Proof. The proof follows directly from the semantics. An interpretation \mathcal{I} satisfies a role box \mathcal{R} iff it satisfies every axiom in \mathcal{R} . For the first claim, if \mathcal{I} does not satisfy $\text{Domain}(R, C)$, then there is some $\langle x, y \rangle \in R^{\mathcal{I}}$ s.t. $x \notin C^{\mathcal{I}}$. In this case, however, $x \in (\exists R. \top)^{\mathcal{I}}$ and $(\exists R. \top)^{\mathcal{I}} \not\subseteq C^{\mathcal{I}}$, so \mathcal{I} does not satisfy $\exists R. \top \sqsubseteq C$. Similarly, if \mathcal{I} does not satisfy $\exists R. \top \sqsubseteq C$, then there is some $x \in (\exists R. \top)^{\mathcal{I}}$ s.t. $x \notin C^{\mathcal{I}}$, so \mathcal{I} does not satisfy $\text{Domain}(R, C)$.

For the second claim, if \mathcal{I} does not satisfy $\text{Range}(R, C)$, then there is some $\langle x, y \rangle \in R^{\mathcal{I}}$ s.t. $y \notin C^{\mathcal{I}}$. In this case, however, $x \notin (\forall R. C)^{\mathcal{I}}$ and $\top^{\mathcal{I}} \not\subseteq (\forall R. C)^{\mathcal{I}}$, so \mathcal{I} does not satisfy $\top \sqsubseteq \forall R. C$. Similarly, if \mathcal{I} does not satisfy $\top \sqsubseteq \forall R. C$, then there is some $x \in (\forall R. C)^{\mathcal{I}}$ and hence some $\langle x, y \rangle \in R^{\mathcal{I}}$ s.t. $y \notin C^{\mathcal{I}}$, so \mathcal{I} does not satisfy $\text{Range}(R, C)$. \square

4.1 Extended Role Absorption

Rewriting techniques similar to those used in concept absorption can be used to extend the basic role absorption technique to deal with a wider range of axioms. An axiom of the form $\exists R. C \sqsubseteq D$ can be absorbed into a domain constraint $\text{Domain}(R, D \sqcup \neg \exists R. C)$ by rewriting it as $\exists R. \top \sqsubseteq D \sqcup \neg \exists R. C$: from the semantics it is easy to see that $(\exists R. \top \sqcap \exists R. C)^{\mathcal{I}} = (\exists R. C)^{\mathcal{I}}$, and $(\exists R. \top \sqcap \exists R. C)^{\mathcal{I}} \subseteq D^{\mathcal{I}}$ iff $(\exists R. \top)^{\mathcal{I}} \subseteq D^{\mathcal{I}} \cup (\neg \exists R. C)^{\mathcal{I}}$. Similarly, an axiom of the form $D \sqsubseteq \forall R. C$ can be absorbed into a domain constraint $\text{Domain}(R, \neg D \sqcup \neg \exists R. \neg C)$: again, it is easy to see that $D \sqsubseteq \forall R. C$ iff $\neg \forall R. C \sqsubseteq \neg D$, and $\neg \forall R. C \equiv \exists R. \neg C$.

Additional rewritings and simplifications can be used to further extend the range of axioms that can be dealt with using role absorption. In some cases these techniques can be applied in different ways such that an axiom could be absorbed using either role or concept absorption. For example, if the TBox contains an axiom $A \sqcap \exists R. B \sqsubseteq C$, then it could be rewritten as $\exists R. B \sqsubseteq C \sqcup \neg A$ and dealt with using extended role absorption, or it could be rewritten as $A \sqsubseteq C \sqcup \neg \exists R. B$ and (possibly) dealt with using concept absorption.

5 Implementation

We have implemented the extended tableaux algorithm and role absorption optimisation in the FaCT++ DL reasoner. FaCT++ is a next generation of the well-known FaCT reasoner [10], being developed as part of the EU WonderWeb project (see <http://wonderweb.semanticweb.org/>); it is based on the same tableaux algorithms as the original FaCT, but has a different architecture and is written in C++ instead of Lisp.

Absorption in FaCT++ uses the same basic approach as FaCT [15, 9]. Given a TBox \mathcal{T} , the absorption algorithm constructs a triple of TBoxes $\langle \mathcal{T}_{\text{def}}, \mathcal{T}_{\text{sub}}, \mathcal{T}_{\text{g}} \rangle$ such that:

- \mathcal{T}_{def} is a set of axioms of the form $A \equiv C$ (equivalent to a pair of axioms $\{A \sqsubseteq C, C \sqsubseteq A\} \subseteq \mathcal{T}$), where $A \in \mathbf{C}$ (i.e., A is a concept name) and there is most one such axiom for each $A \in \mathbf{C}$. Such an axiom is often called a *definition* (of A).
- \mathcal{T}_{sub} consists of a set of axioms of the form $A \sqsubseteq D$, where $A \in \mathbf{C}$ and there is no axiom $A \equiv C$ in \mathcal{T}_{def} .

- \mathcal{T}_g contains all the remaining axioms from \mathcal{T} .

The lazy unfolding optimisation allows the axioms in \mathcal{T}_{def} and \mathcal{T}_{sub} to be dealt with more efficiently than those in \mathcal{T}_g . Therefore, during the absorption process, **FaCT++** processes the axioms in \mathcal{T}_g one at a time, trying to absorb them into \mathcal{T}_{sub} . Those axioms that are not absorbed remain in \mathcal{T}_g .

To simplify the formulation of the absorption algorithm, each axiom $C \sqsubseteq D$ is viewed as a clause $\mathbf{G} = \{D, \neg C\}$, corresponding to the axiom $\top \sqsubseteq C \rightarrow D$, which is equivalent to $C \sqsubseteq D$. The concepts in \mathbf{G} are also assumed to be in negation normal form. For each such axiom, **FaCT++** applies the absorption steps described in Fig. 3, with $\sqcup(\{C_1, \dots, C_n\})$ being used to denote $C_1 \sqcup \dots \sqcup C_n$.

<p>B Beginning of the absorption cycle.</p> <p>C Concept absorption. If there is a concept $\neg A \in \mathbf{G}$ such that $A \in \mathbf{C}$ and there is no axiom of the form $A \equiv C$ in \mathcal{T}_{def}, then add $A \sqsubseteq \sqcup(\mathbf{G} \setminus \{\neg A\})$ to \mathcal{T}_{sub} and exit.</p> <p>R Role absorption. If there is a concept $\neg \exists R.C \in \mathbf{G}$, then add $\text{Domain}(R, \sqcup(\mathbf{G}))$ to \mathcal{R} and exit.</p> <p>S Simplification.</p> <ol style="list-style-type: none"> 1. For every $C \in \mathbf{G}$ such that C is of the form $(C_1 \sqcup \dots \sqcup C_n)$, change \mathbf{G} as follows: $\mathbf{G} = \mathbf{G} \cup \{C_1, \dots, C_n\} \setminus \{C\}$. 2. For every $A \in \mathbf{G}$ (resp. $\neg A \in \mathbf{G}$), if there is an axiom $A \equiv C$ in \mathcal{T}_{def}, then substitute $A \in \mathbf{G}$ (resp. $\neg A \in \mathbf{G}$) with C (resp. $\sim C$). 3. If any simplification rule was applied, then return to step B. <p>E If there is some $C \in \mathbf{G}$ such that C is of the form $(C_1 \sqcap \dots \sqcap C_n)$, then for each C_i try to absorb (recursively) $\mathbf{G} \cup \{C_i\} \setminus \{C\}$, and exit. Otherwise, absorption of \mathbf{G} has failed; leave \mathbf{G} in \mathcal{T}_g, and exit.</p>

Fig. 3. **FaCT++** absorption algorithm

In contrast to the **FaCT** approach, **FaCT++** applies all possible simplifications (except recursive absorption) in a single step. This usually leads to several possible concept and role absorption options, with the intention that heuristics will be used to select the “best” absorption. The development of suitable heuristics is, however, still part of future work.

As a first step towards investigating such heuristics, **FaCT++** can vary the ordering of the concept absorption (C), role absorption (R) and simplification (S) steps within the absorption algorithm. In the following experiments, the ordering is indicated by, e.g., “S,C,R”, indicating that simplification was performed first, followed by concept absorption and then role absorption. Steps that are irrelevant or not performed may be omitted from the ordering indication. E.g., when concept absorption is not applicable to any axiom, the ordering may be indicated as “S,R” or “R,S”. Finally, basic role absorption is applied *only* to axioms of the form $\exists R.C \sqsubseteq D$, and when this form of role absorption is used the role absorption step is always performed first.

6 Empirical Evaluation

We have tested FaCT++’s performance when classifying several TBoxes derived from realistic ontologies. In each case range and domain constraints from the ontology had already been transformed into GCIs of the form $\exists R.T \sqsubseteq C$ and $T \sqsubseteq \forall R.C$ as described above. All tests used FaCT++ version 0.90 beta running under Linux on an Athlon 2000+ machine with 1Gb of memory.

For each ontology we present results using different kinds of absorption—none (–), basic (B) or extended (E)—and different orderings of the absorption and simplification steps in the absorption algorithm. In each case we give the number of concept and role absorptions performed, the time in CPU seconds taken to classify the Tbox and the total number of basic operations performed by the tableaux algorithm (ops). In addition, we give the total number of \exists -rule applications (exists), the total number of \sqcup -rule applications (or) and the total number of state saves performed by the algorithm (saves)—these are most complicated and time-consuming operations performed by the algorithm, and give a useful indication of how the optimisations affect its behavior.

6.1 The NCI Ontology

NCI is a large ontology build by the National Cancer Institute [3]. It contains more than 27,000 concepts (with a very simple structure), 70 roles with 140 range and domain constraints and no other GCIs.³

The results of classifying the Tbox derived from the NCI ontology are presented in Table 1. As there are no other GCIs, only basic role absorption was relevant. It can be seen that classification time and number of operations reduced by approximately 2 orders of magnitude after applying basic role absorption, and that the operation of the algorithm became deterministic (there were no \sqcup -rule applications or state saves).

Role abs.	Order of steps	NCI						
		C-abs	R-abs	time (s)	ops	exists	or	saves
–	–	0	0	2,447.7	344,142,434	1,956,387	150,467,100	21,295,060
B	R	0	140	63.9	1,580,206	172,514	0	0

Table 1. Classification test results for the NCI TBox

6.2 The Wine Ontology

The Wine ontology forms part of the OWL Guide and Test Suite [19, 2]. It contains 346 concepts (with medium-complex structure), 16 roles with 23 range and domain constraints and 151 other GCIs, some of which could be absorbed by either concept or role absorption.

³ See <http://www.mindswap.org/2003/CancerOntology>

The results of classifying the TBox derived from the Wine ontology are presented in Table 2. Here a difference between orders of absorption operations became apparent, with much better results being obtained when concept absorption and simplification (in either order) are performed before role absorption. The reason for this may be the relatively small number of roles and the frequent use of these roles in the TBox, so that absorbing into domain axioms causes GCIs to be more widely applied than is the case when concept absorption is used. For ontologies containing large numbers of roles, it may be better to try role absorption first. Determining this, and experimenting with the use of other heuristics to select different kinds of absorption, will be part of future work.

Role abs.	Order of steps	Wine						
		C-abs	R-abs	time (s)	ops	exists	or	saves
–	–	0	0	4406.6	234,065,470	3,021,098	151,491,192	42,831,327
	C,S	163	0	3176.8	240,520,201	5,899,924	154,204,855	43,734,548
	S,C	163	0	3181.9	240,520,201	5,899,924	154,204,855	43,734,548
B	R	0	23	1430.3	196,959,970	2,845,855	133,602,203	36,937,387
	R,C,S	163	23	47.4	10,488,142	739,066	5,284,936	1,046,458
	R,S,C	163	23	47.1	10,488,142	739,066	5,284,936	1,046,458
E	C,R,S	82	104	485.3	83,073,060	4,664,895	42,490,360	12,203,812
	C,S,R	163	23	47.2	10,488,142	739,066	5,284,936	1,046,458
	S,C,R	163	23	46.9	10,488,142	739,066	5,284,936	1,046,458
	R,C,S	18	168	185.1	36,816,877	1,752,535	20,305,522	4,848,005
	R,S,C	18	168	184.7	36,816,877	1,752,535	20,305,522	4,848,005
	S,R,C	18	168	122.9	27,310,521	1,786,860	14,529,696	2,569,501

Table 2. Classification test results for the Wine TBox

Other variations are mainly due to random factors in the non-deterministic absorption procedure—apart from the ordering of the absorption steps, the current implementation makes an arbitrary choice of possible absorptions. For example, in this test the S,R,C ordering leads to 20% more possible absorption variants than R,C,S or R,S,C, and the resulting absorption turned out to be “better”, even though the distribution of role and concept absorptions remains the same.

6.3 The RTIMS Ontology

The RTIMS ontology is taken from a publish and subscribe application where it is used by document publishers to annotate documents so that they can be routed to the appropriate subscribers [20]. The ontology contains about 250 concepts (with medium-complex structure), 76 range and domain constraints and 14 GCIs that are not absorbable by concept absorption; it was this that first inspired our investigation of extended role absorption.

The results of classifying the TBox derived from the RTIMS ontology are presented in Table 3. It can be seen that classification time and number of operations reduced

by approximately 1 order of magnitude after applying basic role absorption, and by a further 60% (approximately) after applying extended role absorption.

Role abs.	Order of calls	RTIMS						
		C-abs	R-abs	time (s)	ops	exists	or	saves
-	-	0	0	4.72	1,011,467	8,354	499,154	93,733
B	R	0	76	0.63	143,813	5,447	90,252	18,812
E	R,S	0	90	0.24	36,501	4,557	9,007	2,087

Table 3. Absorption test results for the RTIMS Tbox

6.4 Multi RTIMS

RTIMS is the most interesting TBox in terms of extended role absorption, but it is too small to show significant gains in performance. In order to give an indication of the effects of extended role absorption on larger Tboxes containing proportionately more GCIs, we duplicating the RTIMS TBox, systematically renaming concepts and roles, and generated larger TBoxes by unioning together several (from 1 to 100) copies of the the original TBox.




Fig. 4. Classification time (left) and \sqcup -rule applications (right) for multi-RTIMS TBoxes

The results of our experiments with these Tboxes are shown in Figures 4 and 5, with the problem size (number of copies of the original TBox) on the x-axis and classification time in CPU seconds, number of \sqcup -rule applications, \exists -rule applications and state saves on the y-axis (using a logarithmic scale). It can be seen that without role absorption the classification time (and other y-axis parameters) increases rapidly with problem size, and without extended (basic) role absorption a TBox consisting of 28 (8) copies of the original already takes several thousand CPU seconds to classify. Memory usage also




Fig. 5. \exists -rule applications (left) and state saves (right) for multi-RTIMS TBoxes

increases rapidly in these cases, and system memory was exhausted when trying to classify a TBox consisting of 29 (9) copies of the original without extended (basic) role absorption. In contrast, when using extended role absorption, a TBox consisting of 100 copies of the original could be classified in a little over 100 CPU seconds.

7 Discussion

We have shown how a tableaux algorithm for *SHIQ* can be extended to support role boxes that include range and domain axioms, and proved that the extended algorithm is still a decision procedure for the satisfiability and subsumption of *SHIQ* concepts w.r.t. such a role box. It should be straightforward to similarly extend tableau algorithms for related DLs such as *SHOQ*. We have also shown how support for range and domain axioms can be exploited in order to add a new form of absorption optimisation called role absorption.

We have implemented the extended algorithm and the role absorption optimisation in the *FaCT++* reasoner, and we have illustrated their effectiveness by analysing the behaviour of *FaCT++* when classifying several KBs derived from realistic ontologies. The analysis shows that, not only are the new techniques highly effective, but also that the ordering of different absorption steps can have a significant effect on performance. Future work will include a more detailed study of this effect with a view to devising heuristics that can select the most effective absorption for each GCI.

References

1. S. Bechhofer, I. Horrocks, C. Goble, and R. Stevens. OilEd: a reason-able ontology editor for the semantic web. In *Proc. of the Joint German/Austrian Conf. on Artificial Intelligence (KI 2001)*, number 2174 in Lecture Notes in Artificial Intelligence, pages 396–408. Springer-Verlag, 2001.
2. Jeremy J. Carroll and Jos De Roo. OWL web ontology language test cases. W3C Proposed Recommendation, 15 December 2003. Available at <http://www.w3.org/TR/owl-test/>.

3. Peter A. Covitz, Frank Hartel, Carl Schaefer, Sherri De Coronado, Gilberto Fragoso, Himanso Sahni, Scott Gustafson, and Kenneth H. Buetow. caCORE: A common infrastructure for cancer informatics. *Bioinformatics*, 19(18):2404–2412, 2003.
4. Mike Dean, Dan Connolly, Frank van Harmelen, James Hendler, Ian Horrocks, Deborah L. McGuinness, Peter F. Patel-Schneider, and Lynn Andrea Stein. OWL web ontology language 1.0 reference. W3C Proposed Recommendation, 15 December 2003. Available at <http://www.w3.org/TR/owl-ref/>.
5. D. Fensel, F. van Harmelen, I. Horrocks, D. McGuinness, and P. F. Patel-Schneider. OIL: An ontology infrastructure for the semantic web. *IEEE Intelligent Systems*, 16(2):38–45, 2001.
6. Volker Haarslev and Ralf Möller. High performance reasoning with very large knowledge bases: A practical case study. In *Proc. of the 17th Int. Joint Conf. on Artificial Intelligence (IJCAI 2001)*, pages 161–168, 2001.
7. Volker Haarslev and Ralf Möller. RACER system description. In *Proc. of the Int. Joint Conf. on Automated Reasoning (IJCAR 2001)*, volume 2083 of *Lecture Notes in Artificial Intelligence*, pages 701–705. Springer, 2001.
8. Bernhard Hollunder, Werner Nutt, and Manfred Schmidt-Schauß. Subsumption algorithms for concept description languages. In *Proc. of the 9th Eur. Conf. on Artificial Intelligence (ECAI'90)*, pages 348–353, London (United Kingdom), 1990. Pitman.
9. I. Horrocks. Implementation and optimisation techniques. In Franz Baader, Diego Calvanese, Deborah McGuinness, Daniele Nardi, and Peter F. Patel-Schneider, editors, *The Description Logic Handbook: Theory, Implementation, and Applications*, pages 306–346. Cambridge University Press, 2003.
10. Ian Horrocks. Using an expressive description logic: FaCT or fiction? In *Proc. of the 6th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR'98)*, pages 636–647, 1998.
11. Ian Horrocks and Peter F. Patel-Schneider. Reducing OWL entailment to description logic satisfiability. In *Proc. of the 2003 Description Logic Workshop (DL 2003)*, volume 81 of *CEUR* (<http://ceur-ws.org/>), pages 1–8, 2003.
12. Ian Horrocks, Peter F. Patel-Schneider, and Frank van Harmelen. Reviewing the design of DAML+OIL: An ontology language for the semantic web. In *Proc. of the 18th Nat. Conf. on Artificial Intelligence (AAAI 2002)*, pages 792–797. AAAI Press, 2002.
13. Ian Horrocks and Ulrike Sattler. Optimised reasoning for *SHIQ*. In *Proc. of the 15th Eur. Conf. on Artificial Intelligence (ECAI 2002)*, pages 277–281, July 2002.
14. Ian Horrocks, Ulrike Sattler, and Stephan Tobies. Practical reasoning for expressive description logics. In Harald Ganzinger, David McAllester, and Andrei Voronkov, editors, *Proc. of the 6th Int. Conf. on Logic for Programming and Automated Reasoning (LPAR'99)*, number 1705 in *Lecture Notes in Artificial Intelligence*, pages 161–180. Springer, 1999.
15. Ian Horrocks and Stephan Tobies. Reasoning with axioms: Theory and practice. In *Proc. of the 7th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR 2000)*, pages 285–296, 2000.
16. N. F. Noy, M. Sintek, S. Decker, M. Crubezy, R. W. Ferguson, and M. A. Musen. Creating semantic web contents with Protégé-2000. *IEEE Intelligent Systems*, 16(2):60–71, 2001.
17. Peter F. Patel-Schneider, Patrick Hayes, and Ian Horrocks. OWL web ontology language semantics and abstract syntax. W3C Proposed Recommendation, 15 December 2003. Available at <http://www.w3.org/TR/owl-semantic/>.
18. Pellet OWL reasoner. Maryland Information and Network Dynamics Lab, <http://www.mindswap.org/2003/pellet/>.
19. Michael K. Smith, Chris Welty, and Deborah L. McGuinness. OWL web ontology language guide. W3C Proposed Recommendation, 15 December 2003. Available at <http://www.w3.org/TR/owl-guide/>.

20. Michael Uschold, Peter Clark, Fred Dickey, Casey Fung, Sonia Smith, Stephen Uczekaj Michael Wilke, Sean Bechhofer, and Ian Horrocks. A semantic infosphere. In Dieter Fensel, Katia Sycara, and John Mylopoulos, editors, *Proc. of the 2003 International Semantic Web Conference (ISWC 2003)*, number 2870 in Lecture Notes in Computer Science, pages 882–896. Springer, 2003.