Extending Consequence-Based Reasoning to $\mathcal{SHIQ}$

Andrew Bate, Boris Motik, Bernardo Cuenca Grau,
František Simančík, and Ian Horrocks

University of Oxford
Oxford, United Kingdom
`first.last@cs.ox.ac.uk`

1   Introduction

Description logics (DLs) [3] are a family of knowledge representation formalisms with numerous practical applications. $\mathcal{SHIQ}$ is a particularly important DL as it provides the formal underpinning for the Web Ontology Language (OWL). DLs model a domain of interest using *concepts* (i.e., unary predicate symbols) and *roles* (i.e., binary predicate symbols). DL applications often rely on *subsumption*—the problem of checking logical entailment between concepts—and so the development of practical subsumption procedures for DLs such as $\mathcal{SHIQ}$ has received a lot of attention.

Most DLs are fragments of the *guarded fragment* [1] of first-order logic; however, $\mathcal{SHIQ}$ provides a restricted form of counting that does not fall within the guarded fragment. Moreover, most DLs, including $\mathcal{SHIQ}$, can be captured using the *two-variable fragment of first-order logic with counting* ($\mathcal{C}^2$) [11], but this provides us with neither a practical nor a worst-case optimal reasoning procedure ($\mathcal{C}^2$ and $\mathcal{SHIQ}$ are NEXP-TIME- and EXPTIME-complete, respectively). Algorithms for more general logics thus do not satisfy the requirements of DL applications, and so numerous alternatives specific to DLs have been explored. Many DLs can be decided in the framework of resolution [18, 13], including $\mathcal{SHIQ}$ [14]. These procedures are usually worst-case optimal and can be practical, but, as we discuss in Section 3, in even very simple cases they can draw unnecessary inferences. Practically successful $\mathcal{SHIQ}$ reasoners, such as FaCT++ [26], HermiT [9], Pellet [25], and Racer [12], use variants of highly-optimised (hyper)tableau algorithms [6]—model-building algorithms that ensure termination by a variant of *blocking* [7]. Although worst-case optimal tableau algorithms are known [10], practical implementations are typically not worst-case optimal. While generally very effective, tableau algorithms still cannot process certain ontologies; for example, the GALEN ontology[1] has proved particularly challenging, mainly because tableau calculi tend to construct very large models.

A breakthrough in practical ontology reasoning came in the form of *consequence-based calculi*. Although not originally presented in the consequence-based framework, the algorithm for the DL $\mathcal{EL}$ [2] can be seen as the first such calculus. This algorithm was later reformulated and extended to Horn-$\mathcal{SHIQ}$ [15] and Horn-$\mathcal{SROIQ}$ [19]— DLs that support functional roles, but not disjunctive reasoning. Recently, consequence-based calculi were also developed for the DLs $\mathcal{ALCH}$ [24] and $\mathcal{ALCI}$ [23], which support disjunctive reasoning, but not counting. Consequence-based calculi can be seen as

---

[1] http://www.opengalen.org

a combination of resolution and hypertableau (see Section 3 for details). As in resolution, they describe ontology models by systematically deriving certain ontology consequences; and as in hypertableau, the ontology axioms can be used to guide the derivation process, and to avoid drawing unnecessary inferences. Moreover, consequence-based calculi are not just refutationally complete, but can classify an ontology in a single pass, which greatly reduces the overall work. These advantages allowed the CB system to be the first to classify all of GALEN [15].

Existing consequence-based algorithms can handle either disjunctions or counting, but not both. As we discuss in detail in Section 4, it is challenging to extend these algorithms to DLs such as $\mathcal{SHIQ}$ that combine both kinds of construct: counting quantifiers require equality reasoning, which together with disjunctions can impose complex constraints on ontology models. Unlike in existing consequence-based calculi, these constraints cannot be captured using DLs themselves; instead, a more expressive first-order fragment is needed, which makes the reasoning process much more involved.

In Section 5 we present a consequence-based calculus for $\mathcal{SHIQ}$. Borrowing ideas from resolution theorem proving, we encode the required consequences using a special kind of first-order clause; and to handle equality effectively, we base our calculus on *ordered paramodulation* [17]—a state of the art calculus for equational theorem proving used in modern systems such as E [22] and Vampire [20]. To make the calculus efficient on $\mathcal{EL}$, we have carefully constrained the rules so that, on $\mathcal{EL}$ ontologies, it mimics existing $\mathcal{EL}$ calculi. Thus, although a practical evaluation of our calculus is still pending, we believe that it is likely to perform well in practice on 'mostly-$\mathcal{EL}$' ontologies due to is close relationship with existing and well-proven calculi.

2 Preliminaries

First-Order Logic. To simplify matters technically, it is common practice in equational theorem proving to encode atoms as terms. An atomic formula $P(\vec{s})$ can be encoded as $P(\vec{s}) \approx \mathsf{t}$, where t is a new special constant, and P is considered as a function symbol rather than as a predicate symbol. Note however that, in order to avoid meaningless expressions in which predicate symbols occur at proper subterms, a multi-sorted type discipline on terms in the encoding is adopted. Thus, the set of symbols in the signature is partitioned into a set $\mathcal{P}$ of *predicate symbols* (which includes t), and a set $\mathcal{F}$ of *function symbols*.

A *term* is constructed as usual using variables and the signature symbols. Terms containing predicate symbols as their outermost symbol are called $\mathcal{P}$-*terms*, while all other terms are $\mathcal{F}$-*terms*. For example, for P a predicate and f a function symbol, both $f(P(x))$ and $P(P(x))$ are malformed; $P(f(x))$ is a well-formed $\mathcal{P}$-term; and $f(x)$ is a well-formed $\mathcal{F}$-term. An *(in)equality* is an expression of the form $s \approx t$ ($s \not\approx t$) where s and t are both either $\mathcal{F}$- or $\mathcal{P}$-terms. We assume that $\approx$ and $\not\approx$ are implicitly symmetric—that is, $s \bowtie t$ and $t \bowtie s$ are one and the same expression, for $\bowtie \in \{\approx, \not\approx\}$. A *literal* is an equality or an inequality. An *atom* is an equality of the form $P(\vec{s}) \approx \mathsf{t}$, and we write it simply as $P(\vec{s})$ whenever it is clear from the context whether $P(\vec{s})$ is intended to be a $\mathcal{P}$-term or an atom. A *clause* is an expression of the form $\Gamma \to \Delta$ where Γ is a conjunction of atoms called the *body*, and Δ is a disjunction of literals called the

Table 1. Translating Normalised $\mathcal{SHIQ}$ Ontologies into DL-Clauses

$B_1 \sqsubseteq \geqslant n\, S.B_2 \;\rightsquigarrow$	$B_1(x) \rightarrow S(x, f_i(x))$	for $1 \leq i \leq n$
	$B_1(x) \rightarrow B_2(f_i(x))$	for $1 \leq i \leq n$
	$B_1(x) \rightarrow f_i(x) \not\approx f_j(x)$	for $1 \leq i < j \leq n$

$$B_1 \sqsubseteq \leqslant n\, S.B_2 \;\rightsquigarrow\; \begin{array}{l} S(y,x) \wedge B_2(x) \rightarrow S_{B_2}(y,x) \qquad \text{for fresh } S_{B_2} \\[2pt] B_1(x) \wedge \displaystyle\bigwedge_{1 \leq i \leq n+1} S_{B_2}(x, z_i) \rightarrow \bigvee_{1 \leq i < j \leq n+1} z_i \approx z_j \end{array}$$

$$B_1 \sqsubseteq \forall S.B_2 \qquad \rightsquigarrow \qquad B_1(x) \wedge S(x, z_1) \rightarrow B_2(z_1)$$

$$\prod_{1 \leq i \leq n} B_i \sqsubseteq \bigsqcup_{1 \leq j \leq m} B_j \;\rightsquigarrow\; \bigwedge_{1 \leq i \leq n} B_i(x) \rightarrow \bigvee_{1 \leq i \leq m} B_j(x)$$

$$S_1 \sqsubseteq S_2 \qquad \rightsquigarrow \qquad S_1(x, z_1) \rightarrow S_2(x, z_1)$$

$$S_1 \sqsubseteq S_2^- \qquad \rightsquigarrow \qquad S_1(x, z_1) \rightarrow S_2(z_1, x)$$

head. We often treat conjunctions and disjunctions as sets; and we write the empty conjunction (disjunction) as $\top$ ($\bot$). We use the standard notion of subterm positions; then, $s|_p$ is the subterm of s at position p; moreover, $s[t]_p$ is the term obtained from s by replacing the subterm at position p with t; finally, position p is *proper* in t if $t|_p \neq t$.

Orders. A *term order* $\succ$ is a *strict* order on the set of all terms. The *multiset extension* $\succ_{mul}$ of $\succ$ compares multisets M and N on a universe U such that $M \succ_{mul} N$ if and only if $M \neq N$ and, for each $n \in N \setminus M$, some $m \in M \setminus N$ exists such that $m \succ n$, where $\setminus$ is the multiset difference. We extend $\succ$ to literals by identifying each $s \not\approx t$ with the multiset $\{s, s, t, t\}$ and each $s \approx t$ with the multiset $\{s, t\}$, and by comparing the result using $\succ_{mul}$. Given an order $\succ$, element $b \in U$, and subset $S \subset U$, the notation $S \succ b$ abbreviates $\exists a \in S : a \succ b$.

Description Logic $\mathcal{SHIQ}$. In this paper, a $\mathcal{SHIQ}$ ontology is represented as a set of *DL-clauses*, which we define next. Let $\mathcal{P}_1$ and $\mathcal{P}_2$ be countable sets of unary and binary predicate symbols, and let $\mathcal{F}$ be a countable set of unary function symbols. DL-clauses are constructed using the *central variable* x and variables z_i. A *DL-$\mathcal{F}$-term* has the form x, z_i, or $f(x)$ with $f \in \mathcal{F}$; a *DL-$\mathcal{P}$-term* has the form $B(z_i)$, $B(x)$, $B(f(x))$, $S(x, z_i)$, $S(z_i, x)$, $S(x, f(x))$, $S(f(x), x)$ with $B \in \mathcal{P}_1$ and $S \in \mathcal{P}_2$; and a *DL-term* is a DL-$\mathcal{F}$- or a DL-$\mathcal{P}$-term. A *DL-literal* has the form $A \approx \mathsf{t}$ with A a DL-$\mathcal{P}$-term (called a *DL-atom*), or $f(x) \bowtie g(x)$, $f(x) \bowtie z_i$, or $z_i \bowtie z_j$ with $\bowtie \in \{\approx, \not\approx\}$. A *DL-clause* contains only DL-atoms of the form $B(x)$, $S(x, z_i)$, and $S(z_i, x)$ in the body and only DL-literals in the head, and where each variable z_i occurring in the head also occurs in the body. An *ontology* $\mathcal{O}$ is a finite set of DL-clauses. A *query clause* is a DL-clause in which all atoms are of the form $B(x)$. Given an ontology $\mathcal{O}$ and a query clause $\Gamma \rightarrow \Delta$, the *query clause entailment* problem is to decide whether $\mathcal{O} \models \forall x.(\Gamma \rightarrow \Delta)$ holds; we often leave out $\forall x$ and write the latter as $\mathcal{O} \models \Gamma \rightarrow \Delta$.

$\mathcal{SHIQ}$ ontologies are commonly written using a DL-style syntax, but we can always transform such ontologies into DL-clauses without affecting the entailment of query clauses. Transitivity is encoded away as described in [21, 8], and the resulting axioms are *normalised* to the forms shown on the left-hand side of Table 1 as described in [15, 23]. The normalised axioms are translated to DL-clauses as shown in Table 1.
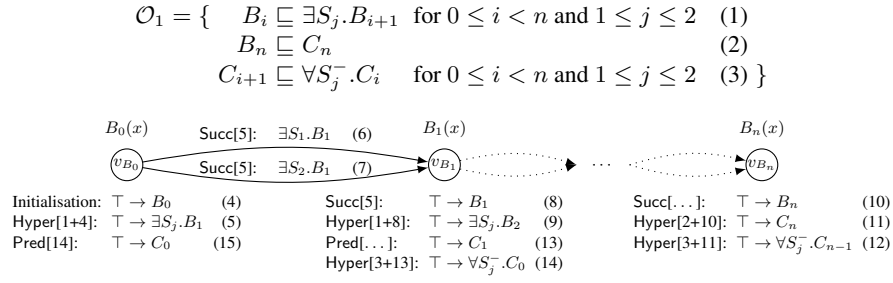
$$\mathcal{O}_1 = \{ \quad B_i \sqsubseteq \exists S_j.B_{i+1} \quad \text{for } 0 \le i < n \text{ and } 1 \le j \le 2 \quad (1)$$
$$B_n \sqsubseteq C_n \quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad (2)$$
$$C_{i+1} \sqsubseteq \forall S_j^-.C_i \quad \text{for } 0 \le i < n \text{ and } 1 \le j \le 2 \quad (3) \ \}$$

$B_0(x)$	Succ[5]:	$\exists S_1.B_1$ (6)	$B_1(x)$			$B_n(x)$	
(v_{B_0})	Succ[5]:	$\exists S_2.B_1$ (7)	(v_{B_1}) $\cdots \blacktriangleright \cdots$			(v_{B_n})	

Initialisation:	$\top \to B_0$	(4)	Succ[5]:	$\top \to B_1$	(8)	Succ[...]: $\quad \top \to B_n$	(10)
Hyper[1+4]:	$\top \to \exists S_j.B_1$	(5)	Hyper[1+8]:	$\top \to \exists S_j.B_2$	(9)	Hyper[2+10]: $\quad \top \to C_n$	(11)
Pred[14]:	$\top \to C_0$	(15)	Pred[...]:	$\top \to C_1$	(13)	Hyper[3+11]: $\quad \top \to \forall S_j^-.C_{n-1}$	(12)
			Hyper[3+13]:	$\top \to \forall S_j^-.C_0$	(14)		

Fig. 1. Example Motivating Consequence-Based Calculi

3 Why Consequence-Based Calculi?

Consider the ontology $\mathcal{O}_1$ (written using DL notation) shown in Figure 1. Axiom (3) can be reformulated as $\exists S_j.C_{i+1} \sqsubseteq C_i$, and so $\mathcal{O}_1$ is in $\mathcal{EL}$. One can readily verify that $\mathcal{O} \models B_i \sqsubseteq C_i$ holds for each $1 \le i \le n$.

To prove, say, $\mathcal{O} \models B_0 \sqsubseteq C_0$, a (hyper)tableau calculus constructs in a forward-chaining manner a tree-shaped model of depth n and of fanout two, where nodes at depth i are labelled by B_i and C_i. Forward chaining ensures that reasoning is goal-oriented and thus amenable to practical implementation. However, all nodes labelled with B_i are of the same type and behave in the same way, which reveals a weakness of (hyper)tableau calculi: the constructed model can be large (exponential in our example) and highly redundant. Techniques such as *caching* [10] or *anywhere blocking* [16] can be used to constrain model construction, but their effectiveness often depends on the order of rule applications. Thus, model size has proved to be a key limiting factor for (hyper)tableau-based reasoners in practice [16].

In contrast, resolution describes models using universally quantified clauses that 'summarise' the model. This prevents redundancy and ensures worst-case optimality of many resolution decision procedures. Nevertheless, resolution can still derive unnecessary clauses. In our example, axioms (1) and (3) are translated into clauses (16) and (17), respectively, which can be used to derive all $2n^2$ clauses of the form (18).

$$B_i(x) \to S_j(x, f_{i,j}(x)) \quad \text{for } i \in \{1, \ldots, n\} \text{ and } j \in \{1, 2\} \quad (16)$$
$$S_j(z_1, x) \wedge C_{k+1}(x) \to C_k(z_1) \quad \text{for } k \in \{1, \ldots, n\} \text{ and } j \in \{1, 2\} \quad (17)$$
$$B_i(x) \wedge C_{k+1}(f_{i,j}(x)) \to C_k(x) \quad \text{for } i,k \in \{1, \ldots, n\} \text{ and } j \in \{1, 2\} \quad (18)$$

Of these $2n^2$ clauses, only those where $i = k$ are relevant to proving $\mathcal{O} \models B_0 \sqsubseteq C_0$. Moreover, if we extend $\mathcal{O}$ with additional clauses that contain B_i and C_i, each of the $2n^2$ clauses from (18) can participate in further inferences, which can give rise to many more irrelevant conclusions. This problem is exacerbated in satisfiable cases since all resolution consequences must then be computed in full.

Consequence-based calculi combine the goal-directed reasoning of (hyper)tableau calculi with the 'summarisation' of resolution. In [23], we presented a very general framework for $\mathcal{ALCI}$ ontologies that captures the key elements of consequence-based calculi such as [2, 15, 19, 24]. We use this framework as basis for our extension to

$\mathcal{SHIQ}$ so, before presenting our extension, we explain the main concepts on $\mathcal{O}_1$. Due to space restrictions we cannot reproduce in full the inference rules from [23]; however, these are similar in spirit to our inference rules for $\mathcal{SHIQ}$ presented in Table 2.

Our calculus constructs a *context structure* $\mathcal{D} = \langle \mathcal{V}, \mathcal{E}, \mathcal{S}, \text{core}, \succ \rangle$—a graph whose vertices $\mathcal{V}$ are called *contexts* and whose directed edges are labelled with concepts of the form $\exists S.B$. Let I be a model of $\mathcal{O}$. Instead of representing each element of I individually as in (hyper)tableau calculi, we 'summarise' all elements of a certain kind using a single context v. Each context $v \in \mathcal{V}$ is associated with a (possibly empty) set core_v of *core* concepts that hold in all domain elements that v represents; thus, core_v determines the kind of context v. We use a set $\mathcal{S}_v$ of clauses to capture additional constraints that the elements represented by v must satisfy; in $\mathcal{ALCI}$, we can do so using clauses over DL concepts of the form $\prod B_i \sqsubseteq \bigsqcup B_j \sqcup \bigsqcup \exists S_k.B_k \sqcup \bigsqcup \forall S_\ell.B_\ell$. Thus, unlike in resolution where all consequences belong to a single set, we assign a consequence a particular set in order to reduce the number of inferences. Clauses in $\mathcal{S}_v$ are 'relative' to core_v: for each $\Gamma \sqsubseteq \Delta \in \mathcal{S}_v$, we have $\mathcal{O} \models \text{core}_v \sqcap \Gamma \sqsubseteq \Delta$—that is, we choose not to include core_v in clause bodies since core_v always holds. Finally, $\succ$ provides each context $v \in \mathcal{V}$ with a concept order $\succ_v$ that restricts resolution inferences in the presence of disjunctions.

Consequence-based calculi are not just refutation-complete: they actually derive the required consequences. Figure 1 shows how this is achieved for $\mathcal{O}_1 \models B_0 \sqsubseteq C_0$; the core and the clauses are shown, respectively, above and below a context. To prove $B_0 \sqsubseteq C_0$, we introduce context v_{B_0} with $\text{core}_{v_{B_0}} = \{B_0\}$ and clause (4) stating that B_0 holds in this context. Next, using the Hyper rule, we derive (5) from (1) and (4); this rule performs hyperresolution, but restricted to one context at a time.

Next, the Succ rule satisfies the existential quantifiers in (5). To this end, the rule uses a parameter called an *expansion strategy*. A strategy is given two sets of constraints that a successor of v_{B_0} must satisfy due to universal restrictions: K_1 contains constraints that must hold, and K_2 contains constraints that might hold. Given such K_1 and K_2, the strategy then decides whether to reuse an existing context or create a fresh one, and in the latter case it also determines how to initialise the new context's core. In our example, there are no universal restrictions and all information in v_{B_0} is deterministic, so $K_1 = K_2 = \{B_1\}$. For $\mathcal{EL}$, a reasonable strategy is to associate with each concept B_i a context v_{B_i} with $\text{core}_{v_{B_i}} = \{B_i\}$, and to always to satisfy existential quantifiers of the form $\exists S.B_i$ using v_{B_i}; thus, in our example we introduce v_{B_1} and initialise it with (8). Note that (5) represents two existential quantifiers, both of which we satisfy (in separate applications of the Succ rule) using v_{B_1}. Different strategies may be used with more expressive DLs; please refer to [23, Section 3.4] for an in-depth discussion.

We construct contexts $v_{B_2}, \ldots, v_{B_n}$ in a similar way, finally deriving (11) by hyperresolving (2) and (10), and then (12) by hyperresolving (3) and (11). Clause (12) imposes a constraint on the predecessor context, which we propagate backwards using the Pred rule, obtaining (13) and (15). Since, however, clauses in $\mathcal{S}_{v_{B_0}}$ are 'relative' to $\text{core}_{v_{B_0}}$, clause (15) actually represents our query clause $B_0 \sqsubseteq C_0$.

Thus, like resolution, consequence-based calculi 'summarise' models to prevent redundant computation, and, like (hyper)tableau calculi, they differentiate elements in a model of $\mathcal{O}$ to prevent the derivation of consequences such as (18).

$$\mathcal{O}_2 = \{ \qquad\quad B_0(x) \to S(f_1(x), x) \quad (19) \qquad B_1(x) \to S(x, f_i(x)) \quad \text{for } 2 \le i \le 3 \quad (22)$$
$$B_0(x) \to B_1(f_1(x)) \quad (20) \qquad B_1(x) \to B_i(f_i(x)) \quad \text{for } 2 \le i \le 3 \quad (23)$$
$$B_2(x) \wedge B_3(x) \to \bot \qquad (21) \qquad B_i(x) \to B_4(x) \qquad \text{for } 2 \le i \le 3 \quad (24)$$
$$B_1(x) \wedge \bigwedge_{1 \le j \le 3} S(x, z_i) \to \bigvee_{1 \le j < k \le 3} z_j \approx z_k \qquad (25) \ \}$$

$B_0(x)$

v_1 Succ[27+28]: f_1 (29)   Succ[32+33+37]: f_2 (38) v_3 $S(y,x), B_2(x)$

Initialisation:	$\top \to B_0(x)$	(26)
Hyper[19+26]:	$\top \to S(f_1(x), x)$	(27)
Hyper[20+26]:	$\top \to B_1(f_1(x))$	(28)
Pred[48]:	$\top \to B_2(x) \vee B_3(x)$	(49)
Hyper[24+49]:	$\top \to B_4(x) \vee B_2(x)$	(50)
Hyper[24+50]:	$\top \to B_4(x)$	(51)

Succ[32+33+37]:	$\top \to S(y,x)$	(39)
Succ[32+33+37]:	$\top \to B_2(x)$	(40)
Succ[32+32+37]:	$B_3(x) \to B_3(x)$	(41)
Hyper[21+40+41]:	$B_3(x) \to \bot$	(42)

v_2 $S(x,y), B_1(x)$ Succ[34+35]: f_3 (43)

$S(y,x), B_3(x)$ v_4

Succ[27+28]:	$\top \to S(x,y)$	(30)
Succ[27+28]:	$\top \to B_1(x)$	(31)
Hyper[22+31]:	$\top \to S(x, f_2(x))$	(32)
Hyper[23+31]:	$\top \to B_2(f_2(x))$	(33)
Hyper[22+31]:	$\top \to S(x, f_3(x))$	(34)
Hyper[23+31]:	$\top \to B_3(f_3(x))$	(35)
Hyper[25+30+31+32+34]:	$\top \to f_2(x) \approx y \vee f_3(x) \approx y \vee f_3(x) \approx f_2(x)$	(36)
Eq[35+36]:	$\top \to f_2(x) \approx y \vee f_3(x) \approx y \vee B_3(f_2(x))$	(37)
Pred[37+42]:	$\top \to f_2(x) \approx y \vee f_3(x) \approx y$	(46)
Eq[35+46]:	$\top \to B_3(y) \vee f_2(x) \approx y$	(47)
Eq[33+47]:	$\top \to B_2(y) \vee B_3(y)$	(48)

Succ[34+35]:	$\top \to S(y,x)$ (44)
Succ[34+35]:	$\top \to B_3(x)$ (45)

Fig. 2. Challenges in Extending the Consequence-Based Framework to $\mathcal{SHIQ}$

4 Extending the Framework to $\mathcal{SHIQ}$

We now present an example before formalising the calculus. Due to an interaction between counting quantifiers and inverse roles, a $\mathcal{SHIQ}$ ontology can impose more complex constraints on model elements than $\mathcal{ALCI}$. Let $\mathcal{O}_2$ be the $\mathcal{SHIQ}$ ontology shown in Figure 2; we argue that $\mathcal{O}_2 \models B_0(x) \to B_4(x)$ holds. To see why, consider an equality Herbrand interpretation I constructed from $B_0(a)$. Then, (19) and (20) derive $S(f_1(a), a)$ and $B_1(f_1(a))$; moreover, (22) and (23) derive $S(f_1(a), f_2(f_1(a)))$ and $B_2(f_2(f_1(a)))$, and $S(f_1(a), f_3(f_1(a)))$ and $B_3(f_3(f_1(a)))$. Due to (24) we derive $B_4(f_2(f_1(a)))$ and $B_4(f_3(f_1(a)))$. Finally, from (25) we derive (52).

$$f_2(f_1(a)) \approx a \vee f_3(f_1(a)) \approx a \vee f_3(f_1(a)) \approx f_2(f_1(a)) \qquad (52)$$

We must satisfy at least one disjunct in (52). Disjunct $f_3(f_1(a)) \approx f_2(f_1(a))$ cannot be satisfied due to (21); but then, regardless of whether we satisfy $f_3(f_1(a)) \approx a$ or $f_2(f_1(a)) \approx a$, we derive $B_4(a)$; hence, the inference holds.

To prove this in our consequence-based framework, we must capture constraint (52) and its consequences. However, this cannot be done using standard description logic notation because DL concepts cannot identify specific successors and predecessors of $f_1(a)$—that is, they cannot say 'either the first or the second successor is equal to the predecessor'. Thus, our main challenges are to devise a method for representing all the relevant constraints that can be induced by $\mathcal{SHIQ}$ ontologies, and to ensure that such constraints are correctly propagated between adjacent contexts.

To address these challenges, we Skolemise existential quantifiers and transform axioms into DL-clauses. Skolemisation introduces function symbols that act as names for successors. Our clauses thus contain terms of the form x, $f_i(x)$, and y which have a special meaning in our setting: variable x represents the elements that a context stands for; $f_i(x)$ represents a successor of x; and y represents the predecessor of x. This allows us to represent constraint (52) as

$$f_2(x) \approx y \vee f_3(x) \approx y \vee f_3(x) \approx f_2(x). \tag{53}$$

Table 2 shows the inference rules of our calculus that are applicable to such a clausal representation. In each clause, literals are ordered from the smallest to the largest, and so the maximal literal is always on the right; moreover, clause numbers correspond to the order of clause derivation. In the rest of this section, we discuss the rules on our running example and show how they verify $\mathcal{O}_2 \models B_0(x) \rightarrow B_4(x)$; for brevity, we present only the inferences needed to produce the desired conclusion.

We first create context v_1 and initialise it with (26); this ensures that each interpretation represented by the context structure contains an element for which B_0 holds. Next, we derive (27) and (28) using hyperresolution. At this point, we could hyperresolve (22) and (28) to obtain $\top \rightarrow S(f_1(x), f_2(f_1(x)))$; however, such inferences could easily lead to nontermination of the calculus due to increased term nesting. Therefore, we require hyperresolution to map variable x in the DL-clauses to variable x in the context clauses; thus, in each context, hyperresolution derives only consequences about x.

Function symbol f_1 in clauses (27) and (28) is akin to an existential quantifier; consequently, the Succ rule introduces a fresh context v_2. Due to Skolemisation, edges in our context structure are labelled with function symbols, rather than concepts of the form $\exists S.B$ as in [23]. The rule uses an expansion strategy analogous to the $\mathcal{EL}$ strategy from Section 3. To determine which information to propagate to a successor, Definition 2 given below introduces a set $\mathsf{Su}(\mathcal{O})$ of *successor triggers*. In our example, DL-clause (25) contains atoms $B_1(x)$ and $S(x, z_i)$ in its body, where z_i can be mapped to a predecessor or a successor of x, so a context in which hyperresolution is applied to (25) will be interested in information about its predecessors; we reflect this by adding $B_1(x)$ and $S(x, y)$ to $\mathsf{Su}(\mathcal{O})$. Thus, the Succ rule introduces context v_2, sets its core to $B_1(x)$ and $S(x, y)$, and initialises the context with (30) and (31).

We next introduce (32)–(35) using hyperresolution, at which point we have sufficient information to apply hyperresolution to (25) to derive (36). Please note how the presence of (30) is crucial for this inference. We use paramodulation to deal with equality in clause (36). As is common in resolution-based theorem proving, we order the literals in a clause and apply inferences only to maximal literals; thus, we derive (37).

Clauses (32), (33), and (37) contain function symbol f_2, so we again apply the Succ rule and introduce context v_3. Due to clause (33), we know that $B_2(x)$ must always hold in v_2, so we add $B_2(x)$ to core_{v_2}. However, $B_3(f_2(x))$ occurs in clause (37) in a disjunction, so it holds only conditionally in v_2; we reflect this by including $B_3(x)$ in the body of clause (41). This allows us derive (42) using hyperresolution.

Clause (42) essentially says '$B_3(f_2(x))$ should not hold in the predecessor', so the Pred rule propagates this information to v_2. This produces clause (46); one can intuitively understand this inference as hyperresolution of (37) and (42), where we take into account that term $f_2(x)$ in context v_2 is represented as variable x in context v_3.

7

After two paramodulation steps, we derive clause (48), which essentially says 'the predecessor must satisfy $B_2(x)$ or $B_3(x)$'. The set $\mathsf{Pr}(\mathcal{O})$ of *predecessor triggers* from Definition 2 identifies this information as relevant to v_1: the DL-clauses in (24) contain $B_2(x)$ and $B_3(x)$ in their bodies, which are represented in v_2 as $B_2(y)$ and $B_3(y)$. Hence $\mathsf{Pr}(\mathcal{O})$ contains $B_2(y)$ and $B_3(y)$, which allows the Pred rule to derive (49).

After two more hyperresolution steps, we finally derive our target clause (51). Please note, however, that we cannot derive this if $B_4(x)$ were maximal in (50); thus, for completeness we require all atoms in the head of a query clause to be smallest. A similar observation applies to $\mathsf{Pr}(\mathcal{O})$: if $B_3(y)$ were maximal in (47), we would not derive (48) and propagate it to v_1; thus, we require all atoms in $\mathsf{Pr}(\mathcal{O})$ to be smallest too.

5 Formalising the Consequence-Based Algorithm for $\mathcal{SHIQ}$

Our calculus manipulates *context clauses*, which are constructed from *context terms* and *context literals* as described in Definition 1. Unlike in general resolution, we restrict context clauses to contain only variables x and y, which have a special meaning in our setting: variable x represents a point (i.e., a term) in the model, and y represents the predecessor of x; this naming convention is important for the rules of our calculus. This is in contrast to the DL-clauses of an ontology: these can contain variables x and z_i, and the latter can refer to either the predecessor or a successor of x.

Definition 1. *A* context $\mathcal{F}$-term *is a term of the form x, y, or $f(x)$ for $f \in \mathcal{F}$; a* context $\mathcal{P}$-term *is a term of the form $B(y)$, $B(x)$, $B(f(x))$, $S(x, y)$, $S(y, x)$, $S(x, f(x))$, or $S(f(x), x)$ for $B, R \in \mathcal{P}$ and $f \in \mathcal{F}$; and a* context term *is an $\mathcal{F}$-term or a $\mathcal{P}$-term. A* context literal *is a literal of the form $A \approx \mathsf{t}$ (called a* context atom*), $f(x) \bowtie g(x)$, or $f(x) \bowtie y$, for A a context $\mathcal{P}$-term and $\bowtie \in \{\approx, \not\approx\}$. A* context clause *is a clause with only function-free context atoms in the body, and only context literals in the head.*

Definition 2 introduces sets $\mathsf{Su}(\mathcal{O})$ and $\mathsf{Pr}(\mathcal{O})$, that identify the information that must be exchanged between adjacent contexts. Intuitively, $\mathsf{Su}(\mathcal{O})$ contains atoms that are of interest to a context's successor, and it guides the Succ rule whereas $\mathsf{Pr}(\mathcal{O})$ contains atoms that are of interest to a context's predecessor and it guides the Pred rule.

Definition 2. *Let $\mathcal{O}$ be an ontology. The set $\mathsf{Su}(\mathcal{O})$ of successor triggers of $\mathcal{O}$ is the smallest set of atoms such that, for each clause $\Gamma \to \Delta \in \mathcal{O}$, we have (i) $B(x) \in \Gamma$ implies $B(x) \in \mathsf{Su}(\mathcal{O})$, (ii) $S(x, z_i) \in \Gamma$ implies $S(x, y) \in \mathsf{Su}(\mathcal{O})$, and (iii) $S(z_i, x) \in \Gamma$ implies $S(y, x) \in \mathsf{Su}(\mathcal{O})$. The set $\mathsf{Pr}(\mathcal{O})$ of predecessor triggers is defined as*

$$\mathsf{Pr}(\mathcal{O}) = \{\, A\{x \mapsto y,\ y \mapsto x\} \mid A \in \mathsf{Su}(\mathcal{O}) \,\} \cup \{\, B(y) \mid B \text{ occurs in } \mathcal{O} \,\}.$$

Similarly to resolution, our calculus uses a term order $\succ$ to restrict its inferences. Definition 3 specifies the conditions that the order must satisfy. Conditions 1 and 2 ensure that $\mathcal{F}$-terms are compared uniformly across contexts; however, $\mathcal{P}$-terms can be compared in different ways in different contexts. Conditions 1 through 4 ensure that, when grounded with x and y mapping to a term its predecessor, respectively, the order is a *simplification order* [4]—a kind of term order commonly used in equational theorem proving. Finally, condition 5 ensures that atoms that might be propagated to a context's predecessor via the Pred rule are smallest, which is important for completeness.

Definition 3. *Let $\gg$ be a total, well-founded order on function symbols. A* context term order $\succ$ *is an order on context terms satisfying the following conditions:*

1. *for each $f \in \mathcal{F}$, we have $f(x) \succ x \succ y$;*
2. *for all $f, g \in \mathcal{F}$ with $f \gg g$, we have $f(x) \succ g(x)$;*
3. *for all terms s_1, s_2, and t and each position p in t, if $s_1 \succ s_2$, then $t[s_1]_p \succ t[s_2]_p$;*
4. *for each term s and each proper position p in s, we have $s \succ s|_p$; and*
5. *for each atom $A \approx \mathsf{t} \in \mathsf{Pr}(\mathcal{O})$ and each context term $s \notin \{x, y\}$, we have $A \not\succ s$.*

Order $\succ$ is extended to literals, also written $\succ$, as described in Section 2.

A lexicographic path order (LPO) [4] over context $\mathcal{F}$-terms and context $\mathcal{P}$-terms, in which x and y are treated as constants such that $x \succ y$, satisfies conditions 1 through 4. Furthermore, $\mathsf{Pr}(\mathcal{O})$ contains only atoms of the form $B(y)$, $S(x, y)$, and $S(y, x)$, which can always be smallest in the order; thus, condition 5 does not contradict the other conditions. Hence, an LPO that is relaxed for condition 5 satisfies Definition 3. In addition to orders, redundancy elimination techniques have proven crucial to the practical effectiveness of resolution calculi. We now define a criterion compatible with our setting.

Definition 4. *A set of clauses U* contains a clause $\Gamma \to \Delta$ up to redundancy, *written $\Gamma \to \Delta \,\hat{\in}\, U$, if*

1. *terms s and s' exist such that $s \approx s \in \Delta$ or $\{s \approx s',\ s \not\approx s'\} \subseteq \Delta$, or*
2. *a clause $\Gamma' \to \Delta' \in U$ exist such that $\Gamma' \subseteq \Gamma$ and $\Delta' \subseteq \Delta$.*

Proposition 1. *For U a set of clauses and $C \in U$ a clause such that $C \,\hat{\in}\, U \setminus \{C\}$, for each clause C' with $C' \,\hat{\in}\, U$, we have $C' \,\hat{\in}\, U \setminus \{C\}$.*

Proposition 1 shows that our calculus is compatible with backward subsumption (which is captured in the Elim rule). Note that tautologies of the form $A \to A$ are *not necessarily* redundant since they can be used to initialise contexts. However, if our calculus were to derive both $A \to A$ and $A \to A \vee A'$ then the latter is always redundant.

We now formalise the notion of a context structure, and define soundness for a context structure. The latter captures the fact that core_v is not contained in the body of any context clause in $\mathcal{S}_v$.

Definition 5. *A* context structure *for an ontology $\mathcal{O}$ is a tuple $\mathcal{D} = \langle \mathcal{V}, \mathcal{E}, \mathcal{S}, \mathsf{core}, \succ \rangle$, where $\mathcal{V}$ is a finite set of* contexts, *$\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V} \times \mathcal{F}$ is a finite set of edges labelled with function symbols, function core assigns to each context v a conjunction core_v of atoms over the $\mathcal{P}$-terms from $\mathsf{Su}(\mathcal{O})$, function $\mathcal{S}$ assigns to each context v a finite set $\mathcal{S}_v$ of context clauses, and function $\succ$ assigns to each context v a context term order $\succ_v$. A context structure $\mathcal{D}$ is* sound *for $\mathcal{O}$ if the following conditions both hold:*

S1. *$\mathcal{O} \models \mathsf{core}_v \wedge \Gamma \to \Delta$ for each context $v \in \mathcal{V}$ and each clause $\Gamma \to \Delta \in \mathcal{S}_v$, and*
S2. *$\mathcal{O} \models \mathsf{core}_u \to \mathsf{core}_v\{x \mapsto f(x), y \mapsto x\}$ for each edge $\langle u, v, f \rangle \in \mathcal{E}$.*

Definition 6 introduces an expansion strategy—a parameter of our calculus that determines when and how to reuse contexts in order to satisfy existential restrictions. We have discussed the roles of expansion strategies in Section 3; moreover, in [23] we presented several expansion strategies for the DLs contained in $\mathcal{ALCI}$, and adapting these to $\mathcal{SHIQ}$ is straightforward.

Table 2. The rules of the consequence-based calculus for $\mathcal{SHIQ}$

Core	If	$A \in \mathsf{core}_v$,
		and $\top \to A \notin \mathcal{S}_v$,
	then	add $\top \to A$ to $\mathcal{S}_v$.
Hyper	If	$\bigwedge_{i=1}^{n} A_i \to \Delta \in \mathcal{O}$,
		σ is a substitution such that $\sigma(x) = x$,
		$\Gamma_i \to \Delta_i \vee A_i\sigma \in \mathcal{S}_v$ with $\Delta_i \not\succeq_v A_i\sigma$ for $i \in \{1, \dots, n\}$,
		and $\bigwedge_{i=1}^{n} \Gamma_i \to \Delta\sigma \vee \bigvee_{i=1}^{n} \Delta_i \;\hat{\notin}\; \mathcal{S}_v$,
	then	add $\bigwedge_{i=1}^{n} \Gamma_i \to \Delta\sigma \vee \bigvee_{i=1}^{n} \Delta_i$ to $\mathcal{S}_v$.
Eq	If	$\Gamma_1 \to \Delta_1 \vee s_1 \approx t_1 \in \mathcal{S}_v$ with $s_1 \succ_v t_1$ and $\Delta_1 \not\succeq_v s_1 \approx t_1$,
		$\Gamma_2 \to \Delta_2 \vee s_2 \bowtie t_2 \in \mathcal{S}_v$ with $\bowtie \in \{\approx, \not\approx\}$ and $s_2 \succ_v t_2$ and $\Delta_2 \not\succeq_v s_2 \bowtie t_2$,
		$s_2\vert_p = s_1$,
		and $\Gamma_1 \wedge \Gamma_2 \to \Delta_1 \vee \Delta_2 \vee s_2[t_1]_p \bowtie t_2 \;\hat{\notin}\; \mathcal{S}_v$,
	then	add $\Gamma_1 \wedge \Gamma_2 \to \Delta_1 \vee \Delta_2 \vee s_2[t_1]_p \bowtie t_2$ to $\mathcal{S}_v$.
Ineq	If	$\Gamma \to \Delta \vee t \not\approx t \in \mathcal{S}_v$ with $\Delta \not\succeq_v t \not\approx t$,
		and $\Gamma \to \Delta \;\hat{\notin}\; \mathcal{S}_v$,
	then	add $\Gamma \to \Delta$ to $\mathcal{S}_v$.
Factor	If	$\Gamma \to \Delta \vee s \approx t \vee s \approx t' \in \mathcal{S}_v$ with $\Delta \cup \{s \approx t\} \not\succeq_v s \approx t'$ and $s \succ_v t'$,
		and $\Gamma \to \Delta \vee t \not\approx t' \vee s \approx t' \;\hat{\notin}\; \mathcal{S}_v$,
	then	add $\Gamma \to \Delta \vee t \not\approx t' \vee s \approx t'$ to $\mathcal{S}_v$.
Elim	If	$\Gamma \to \Delta \in \mathcal{S}_v$ and
		$\Gamma \to \Delta \;\hat{\in}\; \mathcal{S}_v \setminus \{\Gamma \to \Delta\}$
	then	remove $\Gamma \to \Delta$ from $\mathcal{S}_v$.
Pred	If	$\langle u, v, f \rangle \in \mathcal{E}$,
		$\bigwedge_{i=1}^{m} A_i \to \bigvee_{i=m+1}^{m+n} A_i \in \mathcal{S}_v$,
		$\Gamma_i \to \Delta_i \vee A_i\sigma \in \mathcal{S}_u$ with $\Delta_i \not\succeq_u A_i\sigma$ for $1 \leq i \leq m$,
		$A_i \in \mathsf{Pr}(\mathcal{O})$ for each $m + 1 \leq i \leq m + n$,
		and $\bigwedge_{i=1}^{m} \Gamma_i \to \bigvee_{i=1}^{m} \Delta_i \vee \bigvee_{i=m+1}^{m+n} A_i\sigma \;\hat{\notin}\; \mathcal{S}_u$,
	then	add $\bigwedge_{i=1}^{m} \Gamma_i \to \bigvee_{i=1}^{m} \Delta_i \vee \bigvee_{i=m+1}^{m+n} A_i\sigma$ to $\mathcal{S}_u$;
	where	$\sigma = \{x \mapsto f(x), y \mapsto x\}$.
Succ	If	$\Gamma \to \Delta \vee A \in \mathcal{S}_u$ where $\Delta \not\succeq_u A$ and A contains $f(x)$, and
		no edge $\langle u, v, f \rangle \in \mathcal{E}$ exists such that $A \to A \;\hat{\in}\; \mathcal{S}_v$ for each $A \in K_2 \setminus \mathsf{core}_v$,
	then	let $\langle v, \mathsf{core}', \succ' \rangle := \mathsf{strategy}(K_1, \mathcal{D})$;
		if $v \in \mathcal{V}$, then let $\succ_v := \succ_v \cap \succ'$, and
		otherwise let $\mathcal{V} := \mathcal{V} \cup \{v\}$, $\mathsf{core}_v := \mathsf{core}'$, $\succ_v := \succ'$, and $\mathcal{S}_v := \emptyset$;
		add the edge $\langle u, v, f \rangle$ to $\mathcal{E}$; and
		add $A \to A$ to $\mathcal{S}_v$ for each $A \in K_2 \setminus \mathsf{core}_v$;
	where	$\sigma = \{x \mapsto f(x), y \mapsto x\}$,
		$K_1 = \{A \in \mathsf{Su}(\mathcal{O}) \mid \top \to A\sigma \in \mathcal{S}_u\}$, and
		$K_2 = \{A \in \mathsf{Su}(\mathcal{O}) \mid \Gamma' \to \Delta' \vee A\sigma \in \mathcal{S}_u \text{ and } \Delta' \not\succeq_u A\sigma\}$.

Definition 6. *An* expansion strategy *is a polynomial-time computable function* strategy *that takes a set of atoms K and a context structure $\mathcal{D} = \langle \mathcal{V}, \mathcal{E}, \mathcal{S}, \mathsf{core}, \succ \rangle$. The result of* $\mathsf{strategy}(K, \mathcal{D})$ *is a triple $\langle v, \mathsf{core}', \succ' \rangle$ where core' is a subset of K; either $v \notin \mathcal{V}$ is a fresh context, or $v \in \mathcal{V}$ is an existing context in $\mathcal{D}$ such that $\mathsf{core}_v = \mathsf{core}'$; and $\succ'$ is a context term order.*

We now present the main theorems. Full proofs of all technical results can be found in [5]. Theorem 1 proves that all clauses derived by our calculus are indeed conclusions of the input ontology, and Theorem 2 is our statement of completeness.

Theorem 1. *For any strategy, applying a rule from Table 2 to an ontology $\mathcal{O}$ and a context structure $\mathcal{D}$ that is sound for $\mathcal{O}$ produces a context structure that is sound for $\mathcal{O}$.*

Theorem 2. *Let $\mathcal{O}$ be an ontology, and let $\mathcal{D} = \langle \mathcal{V}, \mathcal{E}, \mathcal{S}, \mathsf{core}, \succ \rangle$ be a context structure such that no inference rule from Table 2 is applicable to $\mathcal{O}$ and $\mathcal{D}$. Then, for each query clause $\Gamma_Q \rightarrow \Delta_Q$ and each context $q \in \mathcal{V}$ that satisfy all of the following conditions, we have $\Gamma_Q \rightarrow \Delta_Q \mathrel{\hat{\in}} \mathcal{S}_q$.*

- *C1. $\mathcal{O} \models \Gamma_Q \rightarrow \Delta_Q$.*
- *C2. For each atom $A \approx \mathsf{t} \in \Delta_Q$ and each context term $s \notin \{x, y\}$ such that $A \succ_q s$, we have $s \approx \mathsf{t} \in \Delta_Q \cup \mathsf{Pr}(\mathcal{O})$.*
- *C3. For each $A \in \Gamma_Q$, we have $\Gamma_Q \rightarrow A \mathrel{\hat{\in}} \mathcal{S}_q$.*

Theorems 1 and 2 result in the following algorithm for deciding $\mathcal{O} \models \Gamma_Q \rightarrow \Delta_Q$.

1. Create an empty context structure $\mathcal{D}$, introduce a context q, and, for each $A \in \Gamma_Q$, add $\Gamma_Q \rightarrow A$ to $\mathcal{S}_q$ in order to satisfy condition C3.
2. Initialise $\succ_q$ in a way that satisfies condition C2, and select an expansion strategy.
3. Saturate $\mathcal{D}$ and $\mathcal{O}$ using the inference rules from Table 2.
4. $\Gamma_Q \rightarrow \Delta_Q$ holds if and only if $\Gamma_Q \rightarrow \Delta_Q \mathrel{\hat{\in}} \mathcal{S}_v$.

Proposition 2. *For each expansion strategy that introduces at most exponentially many contexts, the consequence-based calculus for $\mathcal{SHIQ}$ is worst-case optimal.*

Proof. The number $\wp$ of different context clauses that can be generated using the symbols in $\mathcal{O}$ is clearly at most exponential in the size of $\mathcal{O}$, and the number m of clauses participating in each inference is linear in the size of $\mathcal{O}$. Hence, with k contexts, the number of inferences is bounded by $(k \cdot \wp)^m$; if k is at most exponential in the size of $\mathcal{O}$, the number of inferences is exponential as well. Thus, if at most exponentially many contexts are introduced, our algorithm runs in exponential time, which is worst-case optimal for $\mathcal{SHIQ}$ [3]. □

6 Conclusion

We have presented the first consequence based calculus for $\mathcal{SHIQ}$, a DL that includes both disjunction and counting quantifiers. Our calculus combines ideas from state of the art resolution and (hyper)tableau calculi, including the use of ordered paramodulation for efficient equality reasoning. In spite of its increased complexity, the calculus mimics existing and well proven $\mathcal{EL}$ calculi on $\mathcal{EL}$ ontologies. Thus, although implementation and evaluation remain as future work, we believe that the calculus is likely to work well on 'mostly-$\mathcal{EL}$' ontologies—a type of ontology that occurs commonly in practice.

References

1. H. Andréka, J. van Benthem, and I. Németi. Modal Languages and Bounded Fragments of Predicate Logic. *Journal of Philosophical Logic*, 27(3):217–274, 1998.

2. F. Baader, S. Brandt, and C. Lutz. Pushing the $\mathcal{EL}$ Envelope. In L. Pack Kaelbling and A. Saffiotti, editors, *Proc. of the 19th Int. Joint Conference on Artificial Intelligence (IJCAI 2005)*, pages 364–369, Edinburgh, UK, July 30–August 5 2005. Morgan Kaufmann Publishers.

3. F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation and Applications*. Cambridge University Press, January 2003.

4. F. Baader and T. Nipkow. *Term Rewriting and All That*. Cambridge University Press, 1998.

5. A. Bate, B. Motik, B. Cuenca Grau, F. Simančík, and I. Horrocks. Extending Consequence-Based Reasoning to $\mathcal{SHIQ}$. Technical report, Department of Computer Science, University of Oxford, May 2015.

6. P. Baumgartner, U. Furbach, and I. Niemelä. Hyper Tableaux. In *Proc. of the European Workshop on Logics in Artificial Intelligence (JELIA '96)*, number 1126 in LNAI, pages 1–17, Évora, Portugal, September 30–October 3 1996. Springer.

7. P. Baumgartner and R. A. Schmidt. Blocking and Other Enhancements for Bottom-Up Model Generation Methods. In U. Furbach and N. Shankar, editors, *Proc. of the 3rd Int. Joint Conf. on Automated Reasoning (IJCAR 2006)*, volume 4130 of *LNCS*, pages 125–139, Seattle, WA, USA, August 17–20 2006. Springer.

8. S. Demri and H. de Nivelle. Deciding Regular Grammar Logics with Converse Through First-Order Logic. *Journal of Logic, Language and Information*, 14(3):289–329, 2005.

9. Birte Glimm, Ian Horrocks, Boris Motik, Giorgos Stoilos, and Zhe Wang. HermiT: An OWL 2 Reasoner. *Journal of Automated Reasoning*, 53(3):245–269, 2014.

10. Rajeev Goré and Linh Anh Nguyen. EXPTIME Tableaux with Global Caching for Description Logics with Transitive Roles, Inverse Roles and Role Hierarchies. In Nicola Olivetti, editor, *Proc. of the 16th Int. Conf. on Automated Reasoning with Tableaux and Related Methods (TABLEAUX 2007)*, volume 4548 of *LNCS*, pages 133–148, Aix en Provence, France, July 3–6 2007. Springer.

11. E. Grädel, M. Otto, and E. Rosen. Two-Variable Logic with Counting is Decidable. In *Proc. of the 12th IEEE Symposium on Logic in Computer Science (LICS '97)*, pages 306–317, Warsaw, Poland, June 29–July 2 1997. IEEE Computer Society.

12. V. Haarslev and R. Möller. RACER System Description. In R. Goré, A. Leitsch, and T. Nipkow, editors, *Proc. of the 1st Int. Joint Conf. on Automated Reasoning (IJCAR 2001)*, volume 2083 of *LNAI*, pages 701–706, Siena, Italy, June 18–23 2001. Springer.

13. U. Hustadt and R.A. Schmidt. On the Relation of Resolution and Tableaux Proof Systems for Description Logics. In Thomas Dean, editor, *Proc. of the 16th Int. Joint Conf. on Artificial Intelligence (IJCAI 1999)*, pages 110–117, Stockholm, Sweden, July 31 – August 6 1999. Morgan Kaufmann.

14. Ullrich Hustadt, Boris Motik, and Ulrike Sattler. Deciding Expressive Description Logics in the Framework of Resolution. *Information & Computation*, 206(5):579–601, 2008.

15. Y. Kazakov. Consequence-Driven Reasoning for Horn SHIQ Ontologies. In Craig Boutilier, editor, *Proc. of the 21st Int. Joint Conf. on Artificial Intelligence (IJCAI 2009)*, pages 2040–2045, Pasadena, CA, USA, July 11–17 2009.

16. B. Motik, R. Shearer, and I. Horrocks. Hypertableau Reasoning for Description Logics. *Journal of Artificial Intelligence Research*, 36:165–228, 2009.

17. R. Nieuwenhuis and A. Rubio. Theorem Proving with Ordering and Equality Constrained Clauses. *Journal of Symbolic Computation*, 19(4):312–351, 1995.

18. H. De Nivelle, R. A. Schmidt, and U. Hustadt. Resolution-Based Methods for Modal Logics. *Logic Journal of the IGPL*, 8(3):265–292, 2000.

19. M. Ortiz, S. Rudolph, and M. Simkus. Worst-Case Optimal Reasoning for the Horn-DL Fragments of OWL 1 and 2. In F. Lin, U. Sattler, and M. Truszczynski, editors, *Proc. of the 12th Int. Conf. on Knowledge Representation and Reasoning (KR 2010)*, pages 269–279, Toronto, ON, Canada, May 9–13 2010. AAAI Press.

20. A. Riazanov and A. Voronkov. The design and implementation of VAMPIRE. *AI Communications*, 15(2–3):91–110, 2002.

21. R. A. Schmidt and U. Hustadt. A Principle for Incorporating Axioms into the First-Order Translation of Modal Formulae. In F. Baader, editor, *Proc. of the 19th Int. Conf. on Automated Deduction (CADE-19)*, volume 2741 of *LNAI*, pages 412–426, Miami Beach, FL, USA, July 28–August 2 2003. Springer.

22. S. Schulz. E—A Brainiac Theorem Prover. *AI Communications*, 15(2–3):111–126, 2002.

23. František Simančík, Boris Motik, and Ian Horrocks. Consequence-Based and Fixed-Parameter Tractable Reasoning in Description Logics. *Artificial Intelligence*, 209:29–77, 2014.

24. F. Simančík, Y. Kazakov, and I. Horrocks. Consequence-Based Reasoning beyond Horn Ontologies. In Toby Walsh, editor, *Proc. of the 22nd Int. Joint Conf. on Artificial Intelligence (IJCAI 2011)*, pages 1093–1098, Barcelona, Spain, July 16–22 2011.

25. E. Sirin, B. Parsia, B. Cuenca Grau, A. Kalyanpur, and Y. Katz. Pellet: A practical OWL-DL reasoner. *Journal of Web Semantics*, 5(2):51–53, 2007.

26. D. Tsarkov and I. Horrocks. FaCT++ Description Logic Reasoner: System Description. In *Proc. of the 3rd Int. Joint Conf. on Automated Reasoning (IJCAR 2006)*, volume 4130 of *LNAI*, pages 292–297, Seattle, WA, USA, August 17–20 2006. Springer.