

Parallelised ABox Reasoning and Query Answering with Expressive Description Logics

Andreas Steigmiller* and Birte Glimm

Ulm University, Ulm, Germany, <first name>.<last name>@uni-ulm.de

Abstract. Automated reasoning support is an important aspect of logic-based knowledge representation. The development of specialised procedures and sophisticated optimisation techniques significantly improved the performance even for complex reasoning tasks such as conjunctive query answering. Reasoning and query answering over knowledge bases with a large number of facts and expressive schemata remains, however, challenging.

We propose a novel approach where the reasoning over assertional knowledge is split into small, similarly sized work packages to enable a parallelised processing with tableau algorithms, which are dominantly used for reasoning with more expressive Description Logics. To retain completeness in the presence of expressive schemata, we propose a specifically designed cache that allows for controlling and synchronising the interaction between the constructed partial models. We further report on encouraging performance improvements for the implementation of the techniques in the tableau-based reasoning system Konclude.

1 Introduction

Description Logics (DLs) are a family of logic-based representation formalisms that provide the logical underpinning of the well-known Web Ontology Language (OWL). The knowledge expressed with DLs is typically separated into terminological (aka TBox or schema) and assertional knowledge (aka ABox or facts), where the former describes the relationships between concepts (representing sets of individuals with common characteristics) as well as roles (specifying the relationships between pairs of individuals) and the latter asserts these concepts and roles to concrete individuals of the application domain. Automated reasoning systems derive implicit consequences of the explicitly stated information, which, for example, allows for detecting modelling errors and for enriching queries by additional answers that are implied by the knowledge. Expressive DLs, such as *SROIQ* [11], allow for describing the application domain in more detail, but require sophisticated reasoning algorithms and are typically more costly in terms of computational resources. Nevertheless state-of-the-art reasoning systems are usually able to handle real-world ontologies, which often also use expressive language features, due to a large range of developed optimisation techniques.

The increasing volume of data in many application domains leads, however, also to larger amounts of assertional knowledge. For less expressive schemata (where reasoning is usually deterministic), the interest in ontology-based data access (OBDA) led

* Funded by the German Research Foundation (Deutsche Forschungsgemeinschaft, DFG) in project number 330492673

to several advancements, e.g., query rewriting, materialization techniques, or combined approaches. To cope with the reasoning challenges in the presence of an expressive schema several techniques have been developed, which often complement each other. There are, for example, summarisation [3,6] and abstraction techniques [8], which derive consequences for representative individuals and transfer the results to many other individuals with the same or a similar (syntactical) structure. These techniques do not necessarily work well for all ontologies, may be limited to certain queries or (fragments of) DLs, or require expensive computations, e.g., justifications. Several techniques also reduce reasoning to datalog [1,5,21] since datalog engines are targeted towards data intensive applications. This reduction, however, often leads to some additional overhead and, in some cases, it can be necessary to fall back to a fully-fledged DL reasoner, e.g., for handling non-deterministic features. Other approaches partition the ABox or extract modules out of it [20] such that each part can be processed independently [19]. Moreover, approaches based on big data principles such as map and reduce have been proposed [18]. However, they are typically also limited to specific language features and/or queries and do not work for arbitrary ontologies. Particularly challenging is the support of conjunctive queries with complex concept atoms or with existential variables that may bind to anonymous individuals, which are only implied by the knowledge base, since these features typically make it difficult to appropriately split the ABox upfront in such a way that queries can correctly be answered without too much data exchange.

Many state-of-the-art reasoners directly integrate techniques that improve ABox reasoning, e.g., (pseudo) model checking [10] or bulk processing with binary retrieval [9]. Most reasoners for expressive DLs are based on (variants of) tableau algorithms, which construct abstractions of models called completion graphs. By caching (parts of) the completion graph from the initial consistency check, subsequent reasoning tasks and queries can be answered more efficiently [12,16]. However, constructing and caching entire completion graphs for knowledge bases with large ABoxes requires significant amounts of (main) memory, which may be more than what is typically available.

In this paper, we propose to dynamically split the model construction process for tableau algorithms. This allows for (i) handling larger ABoxes since not everything has to be handled at once and for (ii) exploiting parallelisation. The proposed splits lead to similarly sized work packages that can be processed concurrently without direct synchronisation. To ensure that the partial models constructed in parallel are “compatible” with each other, we employ a cache where selected consequences for individuals are stored. For processing new or reprocessing incompatible parts of the knowledge base, we retrieve cached consequences and ensure with appropriate reuse and expansion strategies that the constructed partial models are eventually compatible with the cache, such that it can (asynchronously) be updated. Conjunctive query answering is supported by adapting the expansion criteria and by appropriately splitting the propagation work through the (partial) models for determining query answers.

The paper is organised as follows: Section 2 introduces some preliminaries about DLs and tableau algorithms; Section 3 describes the cache; Section 4 discusses the adaptations for query answering and Section 5 presents implementation details and results of experiments before we conclude in Section 6. An accompanying technical report [14] provides further details, examples, proofs, and evaluation results.

Table 1: Core features of *SRROIQ* ($\#M$ denotes the cardinality of the set M)

	Syntax	Semantics
<i>Individuals:</i> individual	a	$a^I \in \Delta^I$
<i>Roles:</i> atomic role	r	$r^I \subseteq \Delta^I \times \Delta^I$
inverse role	r^-	$\{\langle \gamma, \delta \rangle \mid \langle \delta, \gamma \rangle \in r^I\}$
<i>Concepts:</i> atomic concept	A	$A^I \subseteq \Delta^I$
nominal	$\{a\}$	$\{a^I\}$
top	\top	Δ^I
bottom	\perp	\emptyset
negation	$\neg C$	$\Delta^I \setminus C^I$
conjunction	$C \sqcap D$	$C^I \cap D^I$
disjunction	$C \sqcup D$	$C^I \cup D^I$
existential restriction	$\exists r.C$	$\{\delta \mid \exists \gamma \in C^I : \langle \delta, \gamma \rangle \in r^I\}$
universal restriction	$\forall r.C$	$\{\delta \mid \langle \delta, \gamma \rangle \in r^I \rightarrow \gamma \in C^I\}$
number restriction, $\bowtie \in \{\leq, \geq\}$	$\bowtie n r.C$	$\{\delta \mid \#\{\langle \delta, \gamma \rangle \in r^I \text{ and } \gamma \in C^I\} \bowtie n\}$
<i>Axioms:</i> general concept inclusion	$C \sqsubseteq D$	$C^I \subseteq D^I$
role inclusion	$r \sqsubseteq s$	$r^I \subseteq s^I$
role chains	$r_1 \circ \dots \circ r_n \sqsubseteq S$	$r_1^I \circ \dots \circ r_n^I \subseteq S^I$
<i>Assertions:</i> concept assertion	$C(a)$	$a^I \in C^I$
role assertion	$r(a, b)$	$\langle a^I, b^I \rangle \in r^I$
equality assertion	$a \approx b$	$a^I = b^I$

2 Description Logics and Reasoning Preliminaries

The syntax of DLs is defined using a vocabulary consisting of countably infinite pairwise disjoint sets N_C of *atomic concepts*, N_R of *atomic roles*, and N_I of *individuals*. A role is either atomic or an *inverse role* r^- , $r \in N_R$. The syntax and semantics of complex *concepts* and *axioms* are defined in Table 1. Note that we omit the presentation of some features (e.g., datatypes) and restrictions (e.g., number restrictions may not use “complex roles”, i.e., roles that occur on the right-hand side of role chains or are implied by such roles) for brevity. A knowledge base/ontology \mathcal{K} is a finite set of axioms. One typically distinguishes terminological axioms in the TBox \mathcal{T} (e.g., $C \sqsubseteq D$) and assertions in the ABox \mathcal{A} (e.g., $C(a)$) of \mathcal{K} , i.e., $\mathcal{K} = (\mathcal{T}, \mathcal{A})$. We use $\text{inds}(\mathcal{K})$ to refer to the individuals of \mathcal{K} . An *interpretation* $\mathcal{I} = (\Delta^I, \cdot^I)$ consists of a non-empty domain Δ^I and an *interpretation function* \cdot^I . We say that \mathcal{I} *satisfies* a general concept inclusion (GCI) $C \sqsubseteq D$, written $\mathcal{I} \models C \sqsubseteq D$, if $C^I \subseteq D^I$ (analogously for other axioms and assertions as shown in Table 1). If \mathcal{I} satisfies all axioms of \mathcal{K} , \mathcal{I} is a *model* of \mathcal{K} and \mathcal{K} is *consistent/satisfiable* if it has a model.

A tableau algorithm decides the consistency of a knowledge base \mathcal{K} by trying to construct an abstraction of a model for \mathcal{K} , a so-called *completion graph*. A completion graph G is a tuple $(V, E, \mathcal{L}, \neq)$, where each node $v \in V$ (edge $\langle v, w \rangle \in E$) represents one or more (pairs of) individuals. Each node v (edge $\langle v, w \rangle$) is labelled with a set of concepts (roles), $\mathcal{L}(v)$ ($\mathcal{L}(\langle v, w \rangle)$), which the individuals represented by v ($\langle v, w \rangle$) are instances of. The relation \neq records inequalities between nodes. We call $C \in \mathcal{L}(v)$ ($r \in \mathcal{L}(\langle v, w \rangle)$) a *concept (role) fact*, which we write as $C(v)$ ($r(v, w)$). A node v is a *nominal node* if $\{a\} \in \mathcal{L}(v)$ for some individual a and a *blockable node* otherwise. For

$r \in \mathcal{L}(\langle v, w \rangle)$ and $r \sqsubseteq^* s$, with \sqsubseteq^* the reflexive, transitive closure over role inclusions of \mathcal{K} (including their inverses), we call w an *s-successor* of v and v an *s-predecessor* of w . A node w is called an *s-neighbour* of v if w is an *s-successor* of v or v an *s⁻-successor* of w . We use *ancestor* and *descendant* as the transitive closure of the predecessor and successor relation, respectively. We say that v_n is an *implied descendant* of v_0 if there is a path v_0, v_1, \dots, v_n such that v_{i+1} is a successor of v_i for $0 \leq i < n$ and each v_j with $j > 0$ does not represent an individual of $\text{inds}(\mathcal{K})$.

A completion graph is initialised with one node for each individual in the input knowledge base. Concepts and roles are added to the node and edge labels as specified by concept and role assertions. For simplicity, we assume that, for each individual $a \in \text{inds}(\mathcal{K})$, a nominal $\{a\}$ is added to $\mathcal{L}(v_a)$. This allows for easily handling equality assertions by adding $\{b\}$ to $\mathcal{L}(v_a)$ for $a \approx b \in \mathcal{A}$. As a convention, we write v_a to refer to the node representing $a \in \text{inds}(\mathcal{K})$, i.e., $\{a\} \in \mathcal{L}(v_a)$. Note that v_a and v_b can refer to the same node if $\{a\}$ and $\{b\}$ are in its label. Complex concepts are then decomposed using a set of expansion rules, where each rule application can add new concepts to node labels and/or new nodes and edges, thereby explicating the structure of a model. The rules are applied until either the graph is *fully expanded* (no more rules are applicable), in which case the graph can be used to construct a model that is a *witness* to the consistency of \mathcal{K} , or an obvious contradiction (called a *clash*) is discovered (e.g., a node v with $C, \neg C \in \mathcal{L}(v)$), proving that the completion graph does not correspond to a model. \mathcal{K} is consistent if the rules (some of which are non-deterministic) can be applied such that they build a fully expanded, clash-free completion graph. Cycle detection techniques such as *pairwise blocking* [11] prevent the infinite generation of new nodes.

3 Caching Individual Derivations

Since tableau-based reasoning algorithms reduce (most) reasoning tasks to consistency checking, parallelising the completion graph construction has general benefits on the now ubiquitous multi-core processor systems. Partitioning the ABox upfront such that no or little interaction is required between the partitions [19] no longer works for expressive DLs, such as *SROIQ*, or complex reasoning tasks, such as conjunctive query answering (with complex concepts and/or existential variables). This is, for example, due to implied connections between individuals (e.g., due to nominals) or due to the consideration of new concept expressions at query time. The effect of parallelisation is further hindered by the multitude of optimisations, required to properly deal with real-world ontologies, which often introduce dependencies between rules and (parts of) completion graphs, resulting in the need of data synchronisation. For example, the anywhere blocking optimisation (cycle detection) investigates all previously constructed nodes in the completion graph in order to determine whether a node is blocked. Hence, a parallelisation approach where a completion graph is modified in parallel can be difficult to realise since it could require a lot of synchronisation.

For ontologies with large ABoxes, it seems more suitable to build completion graphs for parts of the ABox separately (by independent threads) and, since independence of the parts cannot be assumed, to align the results afterwards. Such an alignment can, however, be non-trivial on several levels: For example, if different non-deterministic

decisions have been made for individuals in overlapping parts or due to technical details of the often complex data structures for completion graphs, e.g., efficient processing queues, caching status of node labels, etc.

Our parallelisation approach focuses on aligning completion graphs for ABox parts and we address the challenges by employing a cache for certain derivations for individuals, which facilitates the alignment process. For this, consistency checking roughly proceeds as follows: We randomly split the ABox into equally sized parts that are distributed to worker threads. When a thread begins to process one of these ABox parts, it retrieves stored derivations from the cache for (possibly) affected individuals in that part. The thread then tries to construct a fully expanded and clash-free *local* completion graph for the ABox part by reusing cached derivations and/or by expanding the processing to individuals until they are “compatible” with the cache. Compatibility requires that the local completion graph is fully expanded as well as clash-free and that it can be expanded such that it matches the derivations for the remaining individuals in the cache. If it is required to extend the processing to some “neighbouring” individuals for achieving compatibility (e.g., if different non-deterministic decisions are required for the already processed individuals), then also the cached derivations for these individuals are retrieved and considered. If this process succeeds, the cache is updated with the new or changed derivations for the processed individuals.

If compatibility cannot be obtained (e.g., due to expansion limitations that ensure similarly sized work packages), then the cache entries of incompletely handled individuals are marked such that they are considered later separately. For this, a thread loads the data for (some) marked individuals and tries to construct a fully expanded and clash-free completion graph for them until full compatibility is obtained. If clashes occur that depend on reused (non-deterministic) derivations from the cache, then the corresponding individuals can be identified such that their expansion can be prioritized and/or the reuse of their derivations can be avoided. As a result, (in)consistency of the knowledge base can eventually be detected, as soon as all problematic individuals are directly expanded and all relevant non-deterministic decisions are investigated together.

Before describing the different aspects of the approach and the work-flow in more detail, we define a basic version of the cache and how derivations are stored and used.

Definition 1 (Individual Derivations Cache). *Let \mathcal{K} be a knowledge base. We use $\text{fclos}(\mathcal{K})$, $\text{rols}(\mathcal{K})$, and $\text{inds}(\mathcal{K})$ for the sets of concepts, roles, and individuals that can occur in \mathcal{K} or in a completion graph for \mathcal{K} . An individual derivations cache C is a (partial) mapping of individuals from $\text{inds}(\mathcal{K})$ to cache entries, where the cache entry for an individual $a \in \text{inds}(\mathcal{K})$ consists of:*

- $K^C \subseteq 2^{\text{fclos}(\mathcal{K})}$ and $P^C \subseteq 2^{\text{fclos}(\mathcal{K})}$: the sets of known and possibly instantiated concepts of a , respectively,
- $I \subseteq 2^{\text{inds}(\mathcal{K})}$: the individuals that are (indirectly) connected via nominals to a ,
- $\exists: \text{rols}(\mathcal{K}) \rightarrow \mathbf{N}_0$: mapping a role r to the number of existentially derived successors for a and r , and
- $K^R: \text{rols}(\mathcal{K}) \rightarrow 2^{\text{inds}(\mathcal{K})}$ and $P^R: \text{rols}(\mathcal{K}) \rightarrow 2^{\text{inds}(\mathcal{K})}$: mapping a role r to the sets of known and possible neighbours of a and r , respectively.

We write $K^C(a, C)$, $P^C(a, C)$, $I(a, C)$, $\exists(a, C)$, $K^R(a, C)$, and $P^R(a, C)$ to refer to the individual parts of the cache entry $C(a)$. We write $a \in C$ if C is defined for a .

The conditions for determining influenced individuals have a strong correspondence with the tableau expansion rules (cf. [11]). In fact, Condition D1, D2, and D3 basically check whether the \forall -, the choose-, or the \leq -rule could potentially be applicable, i.e., whether the handling of $\forall r.C$ or $\leq rn.C$ concepts requires the consideration of neighbours from the cache. The first part of Condition D4 further checks whether some individual is indirectly connected to a via a nominal for an implied descendant for an individual b and whether new consequences could be propagated to that descendant. This could be the case if the label for a differs to the concepts for a in the cache. The second part handles potential cases where new nominals may have to be introduced and may influence b or descendants of b . Finally, Condition D5 ensures that neighbours are integrated if individuals are newly merged such that their neighbour relations in the cache can be updated.

Instead of mirroring Conditions D1–D4 for determining influencing individuals, we use the relatively simple Conditions G1–G4 since they allow for simpler data structures and efficient cache updates. Condition G1 and G2 simply check whether some possible instances are missing in the local completion graph, which may stem from a neighbouring individual from the cache. Condition G3 analogously checks for a potentially influencing individual b that is indirectly connected via the nominal $\{a\}$ in the label of an implied descendant of b . Condition G4 checks for merges and inequality information caused by non-deterministically derived nominal expressions for other individuals.

The following example, inspired by the well-known UOBM ontology, illustrates consistency checking with the cache.

Example 1. Suppose an ABox consisting of the two parts:

$$\begin{aligned} \mathcal{A}_1 = \{ & \forall enr^-.(\forall takes.GC \sqcup \forall takes.UGC)(uni), \quad likes(stud, soccer), \quad enr(stud, uni), \\ & \forall likes^-.SoccerFan(soccer), \quad takes(stud, course)\}, \\ \mathcal{A}_2 = \{ & \exists hc.\exists likes.\{soccer\}(prof), \quad teaches(prof, course), \\ & \forall teaches.\forall takes^-. \neg TennisFan(prof), \quad likes(prof, soccer)\}. \end{aligned}$$

We abbreviate *Undergraduate Course* as *UGC*, *Graduate Course* as *GC*, *enrolled in* as *enr*, *has child* as *hc*, and *student* as *stud*. For checking \mathcal{A}_1 , we initialise a completion graph with nodes and edges that reflect the individuals and assertions in \mathcal{A}_1 (cf. upper part of Figure 1). To satisfy the universal restriction for v_{soccer} , we apply the \forall -rule, which propagates *SoccerFan* to v_{stud} . Analogously, the universal restriction for v_{uni} propagates $\forall takes.GC \sqcup \forall takes.UGC$ to v_{stud} . We assume that the disjunct $\forall takes.GC$ is checked first, i.e., it is non-deterministically added to $\mathcal{L}(v_{stud})$. Then the concept *GC* is propagated to v_{course} . The completion graph for \mathcal{A}_1 is now fully expanded and clash-free. We next extract the data for the cache (as shown in the lower part of Figure 1).

The completion graph for \mathcal{A}_2 is analogously initialised (cf. upper part of Figure 2). For the concept $\exists hc.\exists likes.\{soccer\} \in \mathcal{L}(v_{prof})$, the \exists -rule of the tableau algorithm builds a blockable *hc*-successor for v_{prof} with $\exists likes.\{soccer\}$ in its label, for which another successor is created that is merged with v_{soccer} (due to the nominal) leading to the depicted edge to v_{soccer} . Due to the universal restriction $\forall likes^-.SoccerFan$ in $\mathcal{L}(v_{soccer})$, *SoccerFan* is propagated to v_1 and to v_{prof} . For the universal restriction $\forall teaches.\forall takes^-. \neg TennisFan \in \mathcal{L}(v_{prof})$, we propagate $\forall takes^-. \neg TennisFan$ to v_{course} . Now, there are no more tableau expansion rules applicable to the constructed com-

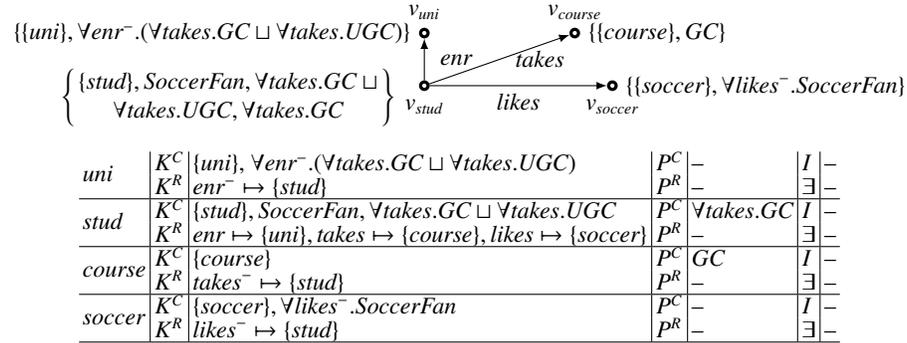


Fig. 1: Local completion graph (upper part) and entries of the individual derivations cache (lower part) for handling ABox \mathcal{A}_1 of Example 1

pletion graph, but it is not yet compatible with the cache and we have to integrate (potentially) influenced or influencing individuals. In fact, *course* causes two incompatibilities: On the one hand, Condition D1 identifies *stud* as (potentially) influenced due to $\forall takes^-. \neg TennisFan \in \mathcal{L}(v_{course})$ and because *stud* is a $takes^-$ -neighbour of *course* according to the cache (cf. Figure 1). On the other hand, Condition G2 is satisfied (since $GC \notin \mathcal{L}(v_{course})$ but $GC \in P^C(course)$) and, therefore, the neighbour *stud* listed in $K^R(course)(takes^-)$ is identified as potentially influencing. We integrate *stud* by creating the node v_{stud} , by adding the concepts $\{stud\}$ and $\forall takes.GC \sqcup \forall takes.UGC$ from the cache to $\mathcal{L}(v_{stud})$, and by creating an edge to v_{course} labelled with *takes* as well as an edge to v_{soccer} labelled with *likes*. Now, the rule application for $\forall takes^-. \neg TennisFan \in \mathcal{L}(course)$ propagates $\neg TennisFan$ to v_{stud} . In addition, by reprocessing the disjunction $\forall takes.GC \sqcup \forall takes.UGC$ for v_{stud} , we obtain $GC \in \mathcal{L}(v_{course})$ if the same disjunct is chosen. As a result, the completion graph is fully expanded and clash-free w.r.t. \mathcal{A}_2 and it is compatible with the cache. Hence, the cache can be updated resulting in the entries depicted in the lower part of Figure 2. Note that only v_{uni} has not been integrated in the completion graph for \mathcal{A}_2 , but there is usually a bigger gain for larger ABoxes.

For parallelising the work-flow with the cache, one has to update and use the entries in a consistent/atomic way such that the state is clear. This can efficiently (and asynchronously) be realised by associating an update id with each cached individual and allow them to have an “inconsistent state” in the cache (see technical report for details). If an update extracted from a constructed completion graph refers to a cache entry of an individual with a changed update id (i.e., the cache entry has been changed by another thread), then the non-deterministically derived consequences are simply added and the state of the individual is set inconsistent (e.g., by adding \perp to $P^C(a)$). If all parts of the ABox are processed, then we repeatedly retrieve individuals with inconsistent states from the cache and reprocess them until compatibility is achieved. Repeatedly deriving different consequences for individuals and correspondingly updating the cache in parallel can, however, threaten termination of the procedure. To retain termination, we simply increase the number of individuals that are processed by one thread such that, in the worst-case, all individuals with inconsistent states are eventually processed together.

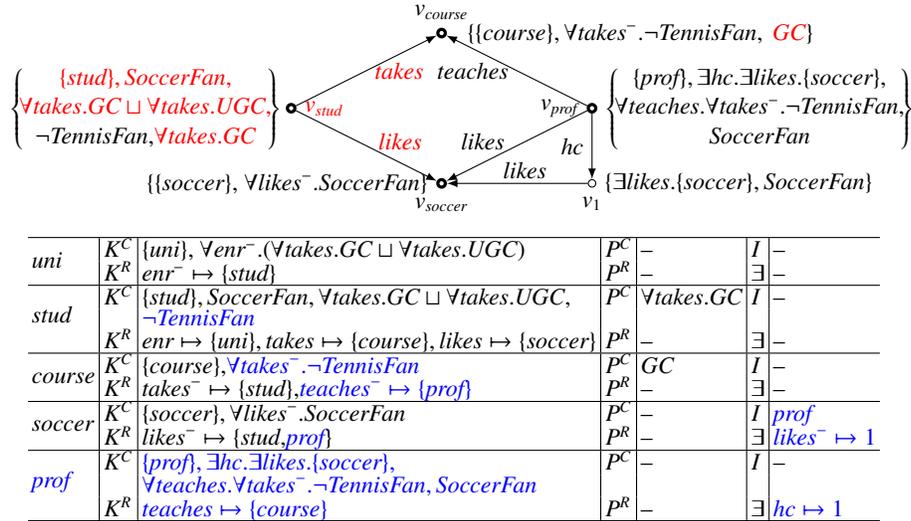


Fig. 2: Local completion graph (upper part, **expansions from cache due to incompatibilities in red**) and entries of the individual derivations cache (lower part, **changes in blue**) for handling ABox \mathcal{A}_2 of Example 1

A naive expansion to all influenced individuals can cause a significant work imbalance for knowledge bases that use complex roles or have intensively connected individuals. In fact, if many neighbour individuals are influenced by universal restrictions or by the merging of nodes, then the expansion to all of them could result in an enormous completion graph processed by one thread, which could limit the effectiveness of the parallelisation. This can be addressed by “cutting the propagation” in the completion graph, for which we then notify the cache that the states of the remaining individuals have to be considered inconsistent such that their processing can be continued later. This cannot avoid the (theoretical) worst-case of processing the whole knowledge base in the end, but it seems to work well for many real-world ontologies, as indicated by our evaluation.

4 Query Answering Support

Compared to other more sophisticated reasoning tasks, conjunctive query answering is typically more challenging since an efficient reduction to consistency checking is not easy. However, a new approach for answering (conjunctive) queries has recently been introduced, where the query atoms are “absorbed” into several simple DL-axioms [13]. These “query axioms” are of the form $C \sqsubseteq \downarrow x.S^x$, $S^x \sqsubseteq \forall r.S_r^x$, $S^x \sqcap A \sqsubseteq S_A^x$, and $S^x \sqcap S^y \sqsubseteq S^{xy}$, where $\downarrow x.S^x$ is a binder concept that triggers the creation of variable mappings in the extended tableau algorithm and S (possibly with sub- and/or superscripts) are so-called query state concepts that are associated with variable mappings, as defined in the following, in order to keep track of partial matches of the query in a completion graph.

Definition 3 (Variable Mappings). A variable mapping μ is a (partial) function from variable names to nodes. Let $G = (V, E, \mathcal{L}, \neq, \mathcal{M})$ be an (extended) completion graph, where $\mathcal{M}(C, v)$ denotes the sets of variable mappings that are associated with a concept C in $\mathcal{L}(v)$. A variable mapping $\mu_1 \cup \mu_2$ is defined by setting $(\mu_1 \cup \mu_2)(x) = \mu_1(x)$ if x is in the domain of μ_1 , and $(\mu_1 \cup \mu_2)(x) = \mu_2(x)$ otherwise. Two variable mappings μ_1 and μ_2 are compatible if $\mu_1(x) = \mu_2(x)$ for all x in the domain of μ_1 as well as μ_2 . The join $\mathcal{M}_1 \bowtie \mathcal{M}_2$ between the sets of variable mappings \mathcal{M}_1 and \mathcal{M}_2 is defined as:

$$\mathcal{M}_1 \bowtie \mathcal{M}_2 = \{\mu_1 \cup \mu_2 \mid \mu_1 \in \mathcal{M}_1, \mu_2 \in \mathcal{M}_2 \text{ and } \mu_1 \text{ is compatible with } \mu_2\}.$$

Rules of the extended tableau algorithm are shown in Table 2 (without considering blocking), which handle the new axioms and concepts by correspondingly creating and propagating variable mappings. For example, a binder concept $\downarrow x.S^x \in \mathcal{L}(v)$ is handled by adding S^x to $\mathcal{L}(v)$ and by creating a mapping $\{x \mapsto v\}$ that is associated with S^x for v , i.e., $\{x \mapsto v\} \in \mathcal{M}(S^x, v)$. Although conjunctive query answering with arbitrary existential variables is still open for *SRDIQ*, the approach works for knowledge bases where only a limited number of new nominal nodes is enforced (by using an extended analogous propagation blocking technique) [13], which is generally the case in practice.

As an example, a simple query with only the atoms $r(x, y)$ and $s(y, x)$ (with x, y both answer variables) can systematically be absorbed into the axioms $\top \sqsubseteq \downarrow x.S^x$, $S^x \sqsubseteq \forall r.S_r^x$, $S_r^x \sqsubseteq \downarrow y.S^y$, $S_r^x \sqcap S^y \sqsubseteq S^{xy}$, $S^{xy} \sqsubseteq \forall s.S_s^{xy}$, and $S_s^{xy} \sqcap S^x \sqsubseteq S^{xyx}$. The query state concept S_r^x , for example, represents the state where bindings for x are propagated to r -successors, i.e., $r(x, y)$ is satisfied. For bindings that are propagated back over s -edges via $\forall s.S_s^{xy}$, the final binary inclusion axiom checks whether the cycle is closed. If it is, the joined variable mappings are associated with S^{xyx} from which answer candidates can be extracted once a fully expanded and clash-free completion graph is found.

Note that with sophisticated absorption techniques, variable mappings can often be derived deterministically, i.e., they directly constitute query answers. Non-deterministically obtained variable mappings do, however, require a separate entailment check to verify that there exist no counter example with the query variables equally bound as in the non-deterministically derived variable mapping. This can be realised by restricting the generated binder concepts of the absorption process to only create corresponding bindings and by triggering a clash with the additional axiom $S^{xyx} \sqsubseteq \perp$.

While query answering by absorption is able to process queries for many (expressive) real-world ontologies [13], especially queries with existential variables can require a substantial amount of computation. A significant bottleneck is often the (variable mappings) *propagation task*, i.e., the creation and propagation of the variable mappings to get all potential answers from a completion graph. Building and using completion graphs for partial ABoxes (possibly in parallel) is difficult since it is unclear which joins of bindings can occur in answers and, hence, how the ABox can suitably be partitioned. The individual derivations cache can, however, also help to split the propagation work for variable mappings such that each thread can completely determine a few answer candidates over a part of the ABox. This enables (a uniform) parallelisation of the propagation task. Note that a partitioning of individuals for the first variable can be used for a naive parallelisation: We create several propagation tasks and restrict the \downarrow -rule to bind the first variable only to the individuals of the handled partition. This can, however, lead to a work imbalance if many answers are based on the same individual for the first

Table 2: Tableau rule extensions

↓-rule:
if $\downarrow x.C \in \mathcal{L}(v)$, and $C \notin \mathcal{L}(v)$ or $\{x \mapsto v\} \notin \mathcal{M}(C, v)$ then $\mathcal{L}(v) = \mathcal{L}(v) \cup \{C\}$, $\mathcal{M}(C, v) = \mathcal{M}(C, v) \cup \{\{x \mapsto v\}\}$
∀-rule:
if $\forall r.C \in \mathcal{L}(v)$, there is an r -neighbour w of v with $C \notin \mathcal{L}(w)$ or $\mathcal{M}(\forall r.C, v) \not\subseteq \mathcal{M}(C, w)$ then $\mathcal{L}(w) = \mathcal{L}(w) \cup \{C\}$, $\mathcal{M}(C, w) = \mathcal{M}(C, w) \cup \mathcal{M}(\forall r.C, v)$
⊆ ₁ -rule:
if $S^{x_1 \dots x_n} \sqsubseteq C \in \mathcal{K}$, $S^{x_1 \dots x_n} \in \mathcal{L}(v)$, and $C \notin \mathcal{L}(v)$ or $\mathcal{M}(S^{x_1 \dots x_n}, v) \not\subseteq \mathcal{M}(C, v)$ then $\mathcal{L}(v) = \mathcal{L}(v) \cup \{C\}$, $\mathcal{M}(C, v) = \mathcal{M}(C, v) \cup \mathcal{M}(S^{x_1 \dots x_n}, v)$
⊆ ₂ -rule:
if $S^{x_1 \dots x_n} \sqcap A \sqsubseteq C \in \mathcal{K}$, $\{S^{x_1 \dots x_n}, A\} \subseteq \mathcal{L}(v)$, and $C \notin \mathcal{L}(v)$ or $\mathcal{M}(S^{x_1 \dots x_n}, v) \not\subseteq \mathcal{M}(C, v)$ then $\mathcal{L}(v) = \mathcal{L}(v) \cup \{C\}$, $\mathcal{M}(C, v) = \mathcal{M}(C, v) \cup \mathcal{M}(S^{x_1 \dots x_n}, v)$
⊆ ₃ -rule:
if $S_1^{x_1 \dots x_n} \sqcap S_2^{y_1 \dots y_m} \sqsubseteq C \in \mathcal{K}$, $\{S_1^{x_1 \dots x_n}, S_2^{y_1 \dots y_m}\} \subseteq \mathcal{L}(v)$, and $(\mathcal{M}(S_1^{x_1 \dots x_n}, v) \bowtie \mathcal{M}(S_2^{y_1 \dots y_m}, v)) \not\subseteq \mathcal{M}(C, v)$ then $\mathcal{L}(v) = \mathcal{L}(v) \cup \{C\}$, $\mathcal{M}(C, v) = \mathcal{M}(C, v) \cup (\mathcal{M}(S_1^{x_1 \dots x_n}, v) \bowtie \mathcal{M}(S_2^{y_1 \dots y_m}, v))$

Algorithm 1 $\text{recPropTask}(R, i)$

Input: Variable binding restrictions R and the index of the next to be handled variable

```

1: if  $i \leq n$  then
2:    $B \leftarrow \text{recPropTask}(R, i + 1)$ 
3:   for each  $x_j$  with  $1 \leq j < i$  do
4:      $R(x_j) \leftarrow B(x_j)$ 
5:   end for
6:   while  $|B(x_i)| \geq l$  do
7:      $R(x_i) \leftarrow R(x_i) \setminus B(x_i)$ 
8:      $B_i \leftarrow \text{recPropTask}(R, i + 1)$ 
9:      $B(x_i) \leftarrow B_i(x_i)$ 
10:  end while
11:   $B(x_i) \leftarrow \emptyset$ 
12: else
13:   $G \leftarrow \text{buildComplGraph}(R, l)$ 
14:   $C \leftarrow C \cup \text{answerCands}(G)$ 
15:   $B \leftarrow \text{extractBoundIndis}(G)$ 
16: end if
17: return  $B$  ▷ Returning
    the bound individuals from the last
    constructed completion graph
    
```

variable and we cannot easily impose restrictions for the other variables since we do not know which combinations of individuals can occur in answers.

We can, however, use a dynamic approach, where we limit the number of individuals to which a variable can be bound. Each individual bound to such a “binding-limited” variable is recorded and in the next propagation task we exclude bindings to already tested individuals. This can, for example, be realised with the recursive function recPropTask sketched in Algorithm 1, which takes as input a mapping R from variables to (still) allowed bindings for individuals and the index i of the current variable (assuming that the variables are sorted in the order in which they are absorbed). The function accesses and modifies some variables via side effects, namely l , denoting the limit for the number of allowed bindings for each variable, n , standing for the number of variables in the query, and C , denoting the set of answer candidates. The function is initially called with $R(x) = \text{inds}(\mathcal{K})$ for each variable x in the query and with $i = 1$ such that the restrictions for the first variable are managed first. As long as there are more variables to handle, the function calls itself recursively for the next variable (cf. Line 2) and checks for the returned sets of bound individuals, denoted with B , from the last generated completion graph whether the limit l has been reached for the current variable. If this is the case, then the bindings for previous variables are “frozen”, i.e., they are interpreted as the only allowed bindings (cf. Line 3–5), and the used bindings for the current variable are excluded for the next propagation task (cf. Line 7). This ensures that all combina-

tions are tested step-by-step and that each propagation task only creates and propagates a limited amount of variable mappings. In fact, if the restrictions for all variables are set, then they are used for constructing the next completion graph (Line 13), where R and l are checked by an adapted \downarrow -rule. Subsequently, we can extract the found answer candidates (Line 14) and the individuals that have been used for bindings (Line 15).

Note that the concepts from the query absorption typically cause an expansion of the local completion graph due to influence criteria, e.g., if $\forall r.S_r^x$ is in the label of some node, then Condition G1 identifies all r -neighbours from the cache as (potentially) influenced and the corresponding nodes need to be integrated into the completion graph to propagate the associated variable mappings to them. This can result in significant propagation work, in particular, for complex roles and individuals with many neighbours. Moreover, exhausted binding restrictions for the next variable might prevent us from actually using the mappings. To address this, one can impose propagation restrictions for universal restrictions of the form $\forall r.S_r^x$ such that the local completion graph is only expanded to nodes for which bindings are possible for the next variable. This can easily be implemented by adapting the query absorption to annotate universal restrictions with the variable of the role atom to which the propagation occurs. For example, a concept of the form $\forall r.S_r^x$ resulting from the query atom $r(x, y)$ is annotated with y , denoted as $\forall r_{\rightarrow y}.S_r^x$. Condition D1 is then adapted to only identify individuals as influenced that are allowed as bindings for the labelled variable of the universal restriction.

Definition 4 (Query Propagation Influence). *Let $G = (V, E, \mathcal{L}, \neq, \mathcal{M})$ be an (extended) completion graph for a knowledge base \mathcal{K} , C an individual derivations cache (c.f. Def. 1), $v_a \in V$ a node representing the individual a , and y a query variable. A restriction set $R(y) \subseteq \text{inds}(\mathcal{K})$ for y restricts the individuals to which y can be bound, i.e., only to a node v_a if $a \in R(y)$. An individual $b \in C$ such that no node in V contains $\{b\}$ in its label is (query propagation) influenced if*

$$\text{Q1 } \forall r_{\rightarrow y}.S_r^{x_1, \dots, x_n} \in \mathcal{L}(v_a), b \in K^R(a)(r) \cup P^R(a)(r), \text{ and } b \in R(y).$$

As mentioned, if the restrictions for a variable are not known upfront, then one can collect them dynamically by only imposing a limit for the number of individuals for the restriction set. While we check whether an individual b is query propagation influenced w.r.t. a variable y and the amount of individuals in the restriction set $R(y)$ is less than the limit, we simply add b to $R(y)$ such that Condition Q1 is satisfied. Analogously, we add b to $R(y)$ when we test whether we can bind y to b for \downarrow -concepts and the limit is not yet reached. When the limit is reached, no more individuals are added to the restriction set and, therefore, no other (combinations of) variable mappings are created and the completion graph is not further expanded to other individuals. The collected restrictions are then used in the next propagation task to enforce the exploration of other (combinations of) variable mappings. Note, however, that steering the expansion with the query propagation influence condition cannot straightforwardly be used for roles with recursive role inclusion axioms (e.g., transitive roles) due to the unfolding process and since it would be too restrictive. One could possibly improve the handling for complex roles with non-trivial adaptations to the tableau algorithm, but it is unclear whether this is worth the effort. In particular, even if the individuals are expanded, the binding restrictions are already adhered to by the \downarrow -rule and one can simply prioritise the absorption of other roles first.

5 Implementation and Experiments

We implemented the individual derivations cache with the sketched extensions in the tableau-based reasoning system Konclude with minor adaptations to fit the architecture and the optimisations of Konclude. In particular, we use Konclude’s efficient, but incomplete saturation procedure [15] to initialise the cache entries for all individuals. If completeness of the saturation cannot be guaranteed for an individual, we mark the corresponding cache entry as inconsistent such that it is reprocessed with the tableau algorithm. Parallel processing (via small batches) is straightforward for the saturation as individuals with their assertions are handled separately. This automatically leads to a very efficient handling of the “simple parts” of an ABox and it only remains to implement the (repeated) reprocessing of individuals with inconsistent cache entries.

Since tableau algorithms are usually quite memory intensive, scalability of the parallelisation not only depends on the CPUs but also on the memory bandwidth and access. Hence, the memory allocator must scale well and the data must be organised in a way that allows for effectively using the CPU caches (e.g., by writing the data of entries with one thread in cohesive memory areas). We investigated different memory allocators (hoard, tcmalloc, jemalloc) and integrated jemalloc [7] since it seems to work best in our scenario. The worker threads for constructing completion graphs only extract the data for cache updates. A designated thread then integrates the cache updates, based on the update ids introduced on page 8, which reduces blocking, improves memory management, and allows for more sophisticated update mechanisms.

To further improve the utilisation of multi-processor systems and to avoid bottlenecks, we also parallelised some other processing steps, e.g., parsing of large RDF triple files, some preprocessing aspects (i.e., extracting internal representations from RDF triples), and indexing of the cache entries for retrieving candidates for query answering. Also note that some higher-level reasoning tasks of Konclude are already (naively) parallelised by creating and processing several consistency checking problems in parallel.

For evaluating the approach,¹ we used the large ontologies and the appertaining queries from the PAGOdA evaluation [21], which includes the well-known LUBM and UOBM benchmarks as well as the real-world ontologies ChEMBL, Reactome, and Uniprot² from the European Bioinformatics Institute. To improve the evaluation w.r.t. the computation of large amounts of answers, we further include the queries from tests for the datalog engine VLog [4], but we use them w.r.t. the original TBoxes. We run the evaluations on a Dell PowerEdge R730 server with two Intel Xeon E5-2660V3 CPUs at 2.4 GHz and 512 GB RAM under a 64bit Ubuntu 18.04.3 LTS. For security reasons and due to multi user restrictions, we could, however, only utilise 480 GB RAM and 8 CPU cores of the server in a containerised environment (via LXD).

Metrics of the evaluated ontologies are depicted on the left-hand side of Table 3, whereas the right-hand side shows the (concurrent) parsing times for the ontologies in seconds, where K-1, K-2, K-4, and K-8 stand for the versions of Konclude, where 1,

¹ Source code, evaluation data, all results, and a Docker image (koncludeeval/parqa) are available at online, e.g., at <https://zenodo.org/record/4606566>.

² We evaluated query answering on a sample (denoted with Uniprot₄₀) since the full Uniprot ontology (Uniprot₁₀₀) is inconsistent and, hence, not interesting for evaluating query answering.

Table 3: Evaluated ontologies, number of queries (from the PAGOdA+VLog evaluation), and parsing times with different number of threads in seconds (speedup factor in parentheses)

Ontology	DL	#Axioms	#Assertions	#Q	K-1	K-2	K-4	K-8
ChEMBL	<i>SRIQ(D)</i>	3,171	$255.8 \cdot 10^6$	6 + 3	1830	935 (2.0)	497 (3.7)	268 (6.8)
LUBM ₈₀₀	<i>ALSHI⁺(D)</i>	93	$110.5 \cdot 10^6$	35 + 3	363	184 (2.0)	102 (3.6)	56 (6.5)
Reactome	<i>SHIN(D)</i>	600	$87.6 \cdot 10^6$	7 + 3	66	34 (1.9)	19 (3.6)	11 (6.2)
Uniprot ₁₀₀	<i>ALCHOIQ(D)</i>	608	$109.5 \cdot 10^6$	- ²	409	229 (1.8)	119 (3.5)	63 (6.7)
Uniprot ₄₀	<i>ALCHOIQ(D)</i>	608	$42.8 \cdot 10^6$	13 + 3	215	113 (1.9)	59 (3.6)	33 (6.5)
UOBM ₅₀₀	<i>SHIN(D)</i>	246	$127.4 \cdot 10^6$	20 + 0	431	227 (1.9)	121 (3.5)	66 (6.5)

2, 4, and 8 threads are used, respectively. Since the parallel parsing hardly requires any synchronisation and only accesses the memory in a very restricted way, it can be seen as a baseline for the achievable scalability (there are minor differences based on how often the different types of assertions occur).

The left-hand side of Table 4 shows the (concurrent) pre-computation times, i.e., the time that is required to get ready for query answering after parsing the ontology, which includes the creation of the internal representation, preprocessing the ontology (e.g., absorption), saturating the concepts and individuals, repeatedly reprocessing the individuals with inconsistent cache entries, classifying the ontology, and preparing data structures for an on-demand/lazy realization. Consistency checking clearly dominates the (pre-)computation time such that the other steps can mostly be neglected for the evaluation (e.g., classification takes only a few milliseconds for these ontologies). As shown in Table 4, our parallelisation approach with the individual derivations cache is able to significantly reduce the time required for consistency checking, but the scalability w.r.t. the number of threads depends on the ontology. For LUBM and ChEMBL, the approach scales almost as well as the parsing process, whereas the scalability w.r.t. Reactome seems limited. The Reactome ontology intensively relies on (inverse) functional roles such that many and large clusters of same individuals are derived in the reasoning process. With a naive implementation of the cache, we would store, for each individual in a cluster, all derived neighbour relations, which easily becomes infeasible if large clusters of same individuals are linked. For our implementation of the cache, we identify and utilise representative individuals to store the neighbour relations more effectively, but we require consistent cache entries for this. If the clusters of same individuals are updated in parallel, which often leads to inconsistent cache entries, more neighbour relations must be managed and, thus, the parallelisation of ontologies such as Reactome only works to a limited extent with the current implementation. Also note that the enormous amounts of individuals in these ontologies make it impossible for the previous version of Konclude to build full completion graphs covering the entire ABox, i.e., the version of Konclude without the cache quickly runs out of memory for these ontologies. Also note that the individuals from the cache are mostly picked in the order in which they are indexed, i.e., more or less randomly due to hashing of pointers. Nominals, however, are indexed first and, hence, are prioritised in the (re-)processing. Clearly, the processing order can have a significant influence on how much (re-)processing is re-

Table 4: (Pre-)computation and accumulated query answering times for the evaluated ontologies with different numbers of threads in seconds (speedup factor in parentheses)

Ontology	(Pre-)computing				Query answering			
	K-1	K-2	K-4	K-8	K-1	K-2	K-4	K-8
ChEMBL	2421	1244 (1.9)	663 (3.7)	397 (6.1)	12767	8927 (1.4)	4507 (2.8)	3231 (4.0)
LUBM ₈₀₀	2793	1658 (1.7)	831 (3.4)	437 (6.4)	2777	1829 (1.5)	1026 (2.7)	569 (4.8)
Reactome	1408	687 (2.0)	427 (3.3)	361 (3.9)	935	524 (1.8)	333 (2.8)	232 (4.0)
Uniprot ₁₀₀	1343	742 (1.8)	429 (3.1)	302 (4.4)	N/A ²	N/A ²	N/A ²	N/A ²
Uniprot ₄₀	1090	532 (2.0)	289 (3.7)	198 (5.5)	28	21 (1.3)	16 (1.8)	14 (2.0)
UOBM ₅₀₀	1317	735 (1.8)	394 (3.3)	245 (5.4)	3774	1799 (2.1)	947 (4.0)	554 (6.8)

quired, but the runs for the evaluated real-world ontologies showed hardly any variance since most consequences could be derived locally.

The right-hand side of Table 4 reveals the query answering times (and scalability), accumulated for each ontology. Since not all steps are parallelised and the version of Konclude with only one thread uses specialised and more efficient implementations in some cases (e.g., an optimised join algorithm for results from several sub-queries, whereas the parallelised version is based on several in-memory map-reduce steps), query answering scalability leaves still room for improvement. Nevertheless, without splitting the propagation tasks, several queries cannot be computed, i.e., the version of Konclude without the presented (query answering) splitting techniques cannot answer all of the queries within the given memory and time limits. Moreover, the parallelisation significantly improves the query answering times and the improvements are larger, the more computation is required. As a comparison, PAGOdA requires 19,666 s for loading and preprocessing all ontologies and more than 101,817 s for query answering, where it reached the memory limit for one query and for two the time limit of 10 hours.

6 Conclusions

We show how the now ubiquitous multi-core processors can be used for parallelising reasoning tasks such as consistency checking and (conjunctive) query answering for expressive Description Logics. For this, we split the assertional knowledge of an ontology to similarly sized work packages for tableau-based reasoning. The technical foundation is a cache that stores chosen consequences derived for individuals and appropriate expansion as well as cache maintenance strategies to ensure correctness and termination. Our experiments with the reasoning system Konclude show promising performance improvements. The approach may even be a suitable basis for distributed reasoning in a compute cluster, where cache entries are distributed over different machines, and for incremental/stream reasoning, where a few assertions are (frequently) added or removed.

References

1. Allocca, C., Calimeri, F., Civili, C., Costabile, R., Cuteri, B., Fiorentino, A., Fuscà, D., Germano, S., Laboccetta, G., Manna, M., Perri, S., Reale, K., Ricca, F., Veltri, P., Zangari, J.:

- Large-scale reasoning on expressive Horn ontologies. In: Proc. 3rd Int. Workshop on the Resurgence of Datalog in Academia and Industry. vol. 2368. CEUR (2019)
2. Baader, F., Calvanese, D., McGuinness, D., Nardi, D., Patel-Schneider, P. (eds.): The Description Logic Handbook: Theory, Implementation, and Applications. Cambridge University Press, second edn. (2007)
 3. Bajraktari, L., Ortiz, M., Simkus, M.: Compiling model representations for querying large ABoxes in expressive DLs. In: Proc. 27th Int. Joint Conf. on Arti. Intelligence (2018)
 4. Carral, D., Dragoste, I., González, L., Jacobs, C.J.H., Krötzsch, M., Urbani, J.: VLog: A rule engine for knowledge graphs. In: Proc. 18th Int. Semantic Web Conf. (ISWC'19) (2019)
 5. Carral, D., González, L., Koopmann, P.: From Horn-SRIQ to datalog: A data-independent transformation that preserves assertion entailment. In: Proc. 33rd AAAI Conf. on Artificial Intelligence (AAAI'19). AAAI Press (2019)
 6. Dolby, J., Fokoue, A., Kalyanpur, A., Schonberg, E., Srinivas, K.: Scalable highly expressive reasoner (SHER). *J. Web Semantics* **7**(4) (2009)
 7. Evans, J.: A scalable concurrent malloc(3) implementation for FreeBSD. In: Proc. 3rd Technical BSD Conf. (BSDCan'06) (2006)
 8. Glimm, B., Kazakov, Y., Tran, T.: Ontology materialization by abstraction refinement in Horn *SHOIQ*. In: Proc. 31st AAAI Conf. on Artif. Intelligence (AAAI'17). pp. 1114–1120. AAAI Press (2017)
 9. Haarslev, V., Möller, R.: On the scalability of description logic instance retrieval. *J. Automated Reasoning* **41**(2) (2008)
 10. Haarslev, V., Möller, R., Turhan, A.Y.: Exploiting pseudo models for TBox and ABox reasoning in expressive description logics. In: Proc. 1st Int. Joint Conf. on Automated Reasoning (IJCAR'01). LNCS, vol. 2083. Springer (2001)
 11. Horrocks, I., Kutz, O., Sattler, U.: The even more irresistible *SROIQ*. In: Proc. 10th Int. Conf. on Principles of Knowledge Representation and Reasoning. AAAI Press (2006)
 12. Sirin, E., Cuenca Grau, B., Parsia, B.: From wine to water: Optimizing description logic reasoning for nominals. In: Proc. 10th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR'06). AAAI Press (2006)
 13. Steigmiller, A., Glimm, B.: Absorption-based query answering for expressive description logics. In: Proc. 18th Int. Semantic Web Conf. (ISWC'19). LNCS, Springer (2019)
 14. Steigmiller, A., Glimm, B.: Parallelised ABox reasoning and query answering with expressive description logics – technical report. Tech. rep., Ulm University, Ulm, Germany (2021), available online at https://www.uni-ulm.de/fileadmin/website_uni_ulm/iui.inst.090/Publikationen/2021/StGl2021-PARQA-TR-ESWC.pdf
 15. Steigmiller, A., Glimm, B., Liebig, T.: Coupling tableau algorithms for expressive description logics with completion-based saturation procedures. In: Proc. 7th Int. Joint Conf. on Automated Reasoning (IJCAR'14). LNCS, Springer (2014), accepted
 16. Steigmiller, A., Glimm, B., Liebig, T.: Completion graph caching for expressive description logics. In: Proc. 28th Int. Workshop on Description Logics (DL'15) (2015)
 17. Tsarkov, D., Horrocks, I., Patel-Schneider, P.F.: Optimizing terminological reasoning for expressive description logics. *J. Automated Reasoning* **39** (2007)
 18. Urbani, J., Kotoulas, S., Maassen, J., van Harmelen, F., Bal, H.E.: Webpie: A web-scale parallel inference engine using mapreduce. *J. Web Semantics* **10** (2012)
 19. Wandelt, S., Möller, R.: Distributed island-based query answering for expressive ontologies. In: Proc. 5th Int. Conf. on Advances in Grid and Pervasive Computing (GPC'10) (2010)
 20. Wandelt, S., Möller, R.: Towards ABox modularization of semi-expressive description logics. *J. Applied Ontology* **7**(2), 133–167 (2012)
 21. Zhou, Y., Cuenca Grau, B., Nenov, Y., Kaminski, M., Horrocks, I.: PAGoDA: Pay-as-you-go ontology query answering using a datalog reasoner. *J. Artif. Intelligence Research* **54** (2015)