

Status *QIO*: Conjunctive Query Entailment Is Decidable

Birte Glimm

Oxford University Computing Laboratory, UK
birte.glimm@comlab.ox.ac.uk

Sebastian Rudolph

AIFB, Karlsruhe Institute of Technology, DE
rudolph@kit.edu

Abstract

Description Logics (DLs) are knowledge representation formalisms that provide, for example, the logical underpinning of the W3C OWL standards. Conjunctive queries (CQs), the standard query language in databases, have recently gained significant attention for querying DL knowledge bases. Several different techniques are available for a wide range of DLs. Nevertheless, for OWL 1 DL and OWL 2 DL, decidability of CQ entailment is an open problem. So far, the combination of nominals, inverse roles, and number restrictions caused unsolvable problems. We tackle this problem and present a decidability result for entailment of unions of CQs in a DL with all three problematic constructors. For queries with only simple roles, our result also shows decidability in the logic that underpins OWL 1 DL and we believe that the presented results will pave the way for further progress towards CQ entailment decision procedures for OWL.

Introduction

Since conjunctive queries were introduced in the context of Description Logics (DLs) (Calvanese, De Giacomo, and Lenzerini 1998), the topic has gained significant attention. Existing techniques for conjunctive query entailment fail if the queried knowledge base (KB) contains inverse roles, nominals, and functionality (or number restrictions) simultaneously; in this case even decidability was an open problem. We tackle this problem and present a decidability result for entailment of conjunctive queries (CQs) in the very expressive Description Logic (DL) *ALCHOIQb* (Baader et al. 2003), which contains all three problematic constructors simultaneously.

Given the variety of recent publications which show a great interest in the problem of conjunctive query entailment over expressive DLs, it is very interesting that for the DLs *SHIF*, *SHOIN*, and *SROIQ* that underpin the widely adopted standards OWL Lite, OWL 1 DL, and OWL 2 DL, respectively, decidability of conjunctive query entailment has only been established for OWL Lite. The critical combination of inverse roles (*I*), nominals (*O*), and number restrictions/counting quantifiers (*F* stands for functionality, *N* for unqualified number restrictions, and *Q* for qualified number restrictions) caused also a major hurdle

Copyright © 2010, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

in the development of implementable algorithms for knowledge base satisfiability in *SHOIN* and extensions thereof, but in 2005, Horrocks et al. devised a tableaux-based decision procedure (Horrocks and Sattler 2005) that has since been extended to *SROIQ*. Meanwhile also alternative approaches such as resolution (Kazakov and Motik 2008), and hypertableaux based procedures (Motik, Shearer, and Horrocks 2009) are available and implemented.

The key obstacle in establishing a decision procedure is the existence of potentially infinitely many *new nominals*, i.e., elements that are uniquely identifiable in any model of a KB. For an example, consider a KB \mathcal{K} containing the axioms $\{o_1\} \sqsubseteq \exists f.\exists s.\exists f^-. \{o_2\}$, $\{o_2\} \sqsubseteq \exists f.\exists s.\exists f^-. \{o_3\}$, $\{o_3\} \sqsubseteq \exists f.\exists s.\exists f^-. \{o_1\}$ for f a functional role (see Figure 1). Due to the functionality of f , each nominal o_i can have only one f successor, which forces the existence of an s -cycle. A cyclic Boolean query such as $\{s(x, y), s(y, z), s(z, x)\}$ that checks for the existence of such a cycle cannot be answered by replacing variables with individual names (o_i) nor can we rewrite the query into an equivalent tree-shaped query. The elements in the cycle behave as if they were nominals, but we do not have names for them.




Figure 1: A representation for a model of \mathcal{K} , where the three elements in the s -cycle are so-called *new nominals*.

We tackle the problem of conjunctive query entailment in a very expressive DL that contains all the three problematic constructors simultaneously and prove decidability of (unions of) conjunctive queries. The most challenging part is to establish finite representability of countermodels in case the query given as input is not entailed by the knowledge base. Our results also hold for *SHOIQ* knowledge bases, i.e., with some roles declared as transitive, provided that the queries contain only simple roles (roles that are neither transitive nor have a transitive subrole). This is essentially the same restriction that is placed on roles that can occur in number restrictions since otherwise the standard reasoning tasks become undecidable. Under this restriction, we can use standard techniques for eliminating transitivity (Kazakov and Motik 2008). Hence, we also show decidability of

conjunctive query entailment in OWL DL, for queries with only simple roles.

We believe that our work is also valuable for understanding, in general, the structure of models in DLs that contain nominals, inverse roles, and number restrictions. Furthermore, we devise non-trivial extensions of standard techniques such as unraveling, which we believe will prove useful when working with such expressive DLs.

Full proofs and additional material can be found in the accompanying technical report (Glimm and Rudolph 2009).

Related Work

Conjunctive queries have been introduced in the context of Description Logics (DLs) by Calvanese, De Giacomo, and Lenzerini 1998. In particular in recent years, the problem of decidability of conjunctive query entailment and the complexity of the problem in different logics has gained significant attention. For the DLs *SHIQ* and *SHOQ* decidability and 2-EXPTIME-completeness of the problem is known (Glimm et al. 2008; Glimm, Horrocks, and Sattler 2008; Lutz 2008; Eiter et al. 2009). Conjunctive query entailment is already 2-EXPTIME-hard in the relatively weak DL *ALCC* (Lutz 2008), which was initially attributed to inverse roles. Recently, it was shown, however, that also transitive roles together with role hierarchies as in the DL *SH* make conjunctive query entailment 2-EXPTIME-hard (Eiter et al. 2009). The techniques by Glimm et al. for *SHIQ* and *SHOQ* (Glimm et al. 2008; Glimm, Horrocks, and Sattler 2008) reduce query entailment to the standard reasoning task of knowledge base satisfiability checking in the DL extended with role conjunctions. An alternative technique is the so-called knots technique (Ortiz, Simkus, and Eiter 2008), which is an instance of the mosaic technique originating in Modal Logic. This technique also gives worst-case optimal algorithms for *SHIQ* and several of its sub-logics. Further, there are automata-based decision procedures for positive existential path queries (Calvanese, Eiter, and Ortiz 2007; 2007). Positive existential path queries generalize unions of conjunctive queries and, therefore, decision procedures for this kind of query also provides decision procedures for unions of conjunctive queries. In particular the most recent extension (Calvanese, Eiter, and Ortiz 2007) is very close to a conjunctive query entailment decision procedure for OWL 2, which corresponds to the DL *SROIQ*, because it covers *SRIQ*, *SROQ*, and *SROI*. The use of the three problematic constructors for nominals, inverses, and number restrictions is, however, not covered.

Regarding data complexity, i.e., the complexity with respect to the ABox (the data) only, CQ entailment is usually NP-complete for expressive logics. For example, for DLs from *ALÉ* up to *SHIQ* this is the case (Glimm et al. 2008) and this holds also for CQ entailment in the two variable guarded fragment with counting (Pratt-Hartmann 2009). The latter work is quite closely related since many Description Logics can be translated into the two variable guarded fragment with counting, i.e., the results of Pratt-Hartmann also holds for *SHIQ* with only simple roles (roles that are not transitive and have no transitive subrole) in the query.

Query entailment and answering have also been studied in the context of databases with incomplete information (Rosati 2006b; van der Meyden 1998; Grahne 1991). In this setting, DLs can be used as schema languages, but the expressivity of the considered DLs is usually much lower than the expressivity of the DL *ALCHOIQb* that we consider here and reasoning in them is usually tractable. For example, the constructors provided by logics of the DL-Lite family (Calvanese et al. 2007) are chosen such that the standard reasoning tasks are in PTIME and query entailment is in LOGSPACE with respect to data complexity. Thus, TBox reasoning can be done independently of the ABox and the ABox can be stored and accessed using a standard database SQL engine. Another tractable DL is *ÉL* (Baader 2003). Conjunctive query entailment in *ÉL* is, however, not tractable as the complexity increases to coNP-complete (Rosati 2007b). Moreover for *ÉL⁺⁺* (Baader, Brandt, and Lutz 2005), a still tractable extension of *ÉL*, query entailment is even undecidable (Krötzsch, Rudolph, and Hitzler 2007). This is mainly because in *ÉL⁺⁺*, one can use unrestricted role compositions. This allows for encoding context-free languages, and conjunctive queries can then be used to check the intersection of such languages, which is known to be an undecidable problem. Since the logics used in databases with incomplete information are considerably less expressive than *ALCHOIQb*, the techniques developed in that area do not transfer to our setting.

Given that query entailment is a (computationally) harder task than, for example, knowledge base satisfiability, it is not very surprising that decidability of the latter task does not necessarily transfer to the problem of CQ entailment. Most of the undecidability results can be transferred from FOL since many DLs can directly be translated into an equivalent FOL theory. For example, it is known that conjunctive query entailment is undecidable in the two variable fragment of First-Order Logic \mathcal{L}_2 (Rosati 2007a), and Rosati identifies a relatively small set of constructors that cause the undecidability (most notably role negation axioms, i.e., axioms of the form $\forall x, y (\neg R(x, y) \rightarrow P(x, y))$ for R, P binary predicates). Pratt-Hartmann 2009 recently established decidability for CQ entailment in the two variable guarded fragment with counting (\mathcal{GC}_2). It is worth noting that Pratt-Hartmann assumes that the background theory (that is the knowledge base in our case) is constant free and formulae of the form $\exists_{=1} x (P(x))$, which can be used to simulate constants/nominals, are not considered guarded. His result covers, therefore, only the DL *ALCHIQb* and is not applicable to the case, when the input knowledge base (the background theory) contains nominals (individual constants).

Most of the implemented DL reasoners, e.g., KAON2,¹ Pellet, and RacerPro,² provide an interface for conjunctive query answering, although KAON2 and RacerPro consider only named individuals in the ABox for the assignments of variables. Under that restriction queries do no longer have the standard FOL semantics and decidability is obviously not an issue since conjunctive query answering with this re-

¹<http://kaon2.semanticweb.org>

²<http://www.racer-systems.com>

striction can be reduced to standard instance retrieval by replacing the variables with individual names from the ABox and then testing entailment of each conjunct separately. Pellet goes beyond that and also provides an interface for conjunctive queries with FOL semantics under the restriction that the queries have a kind of tree shape. Under this restriction decidability is known since CQs can then be expressed as normal concepts (possibly by adding role conjunctions).

The Big Picture

Before going into the technical details, we will describe our overall line of argumentation establishing decidability of conjunctive query entailment in *ALCHOIQb*.

Decidability via Finitely Representable Countermodels. Let \mathcal{K} be an *ALCHOIQb* knowledge base and q be the conjunctive query in question, i.e., we aim to determine whether

$$\mathcal{K} \models q.$$

Clearly, as *ALCHOIQb* is a fragment of first-order predicate logic with equality, \mathcal{K} can be translated into a FOL sentence $FOL(\mathcal{K})$. Likewise we find a FOL sentence $FOL(q)$ for q being just an existentially quantified formula. Hence, checking the above entailment is equivalent to determining whether the first-order theory $FOL(\mathcal{K})$ entails $FOL(q)$. As a result of the completeness theorem for FOL (Gödel 1929), the consequences of a finite FOL theory are recursively enumerable, which provides us with a procedure that terminates if $\mathcal{K} \models q$. Hence, we can establish decidability by providing another algorithm that terminates iff the entailment above does not hold – i.e., if there is a so-called *countermodel* being a model \mathcal{I} of \mathcal{K} for which $\mathcal{I} \not\models q$.

We will provide such an algorithm by showing that, whenever such a countermodel \mathcal{I} exists at all, there is also a countermodel $\bar{\mathcal{I}}$ that is finitely representable. More precisely, $\bar{\mathcal{I}}$ can be encoded into a word $Rep(\bar{\mathcal{I}})$ of finite length over a finite alphabet, whereby the encoding Rep has the property that for every such finite word it can be effectively checked whether it represents a countermodel for a given knowledge base and query.

As a consequence thereof, we can create the desired algorithm that enumerates all words, checks each for being a countermodel, and terminates as soon as it has found one.

Finite Representability by Bounding Nominals and Blocking. We now outline how we are going to show that there is always a finitely representable countermodel, if there is one at all. We do this by taking an arbitrary countermodel and cautiously transforming it into a countermodel that is finitely representable. Cautiously means that we have to make sure that the transformation does preserve the two properties of 1) being a model of the underlying knowledge base \mathcal{K} and 2) not entailing the considered query q .

The result of the overall transformation is going to be a *regular model*, i.e., a structure where substructures are being in a certain sense periodically repeated. It is common practice in DL theory to construct this kind of models from arbitrary ones by blocking techniques, whereby certain element configurations occurring twice in the original model are detected and the new model is generated by infinitely

stringing together the same finite substructure that is delimited by those two configurations.

In the case we consider, this technique cannot be applied directly to the original countermodel. This is due to an intricate interplay of nominals, inverse roles and cardinality constraints by which an arbitrary – even an infinite – number of domain elements can be forced to “behave” like nominals; this is why those elements are usually referred to as *new nominals* in a DL setting. In FOL, nominals are often called kings and the new nominals are called the court. In our case, the presence of infinitely many new nominals in the model may prevent the existence of repeated configurations needed for blocking.

We overcome this difficulty by first applying a transformation by means of which the original countermodel is converted into a countermodel with only finitely many new nominals. This guarantees that the subsequent blocking-based transformation is applicable and will yield the desired regular (and thus finitely representable) model.

Bounding Nominals by Transformations of Forest Quasi-Models. For our argumentation, we introduce the notion of *forest quasi-models*. These are structures not satisfying the originally considered knowledge base but a weakened form of it. In return to this concession, they exhibit a proper forest structure that is easier to handle and manipulate.

We employ two techniques to turn “proper” models into forest quasi-models and vice versa: a model can be *unraveled* yielding a forest quasi-model. A forest quasi-model can be *collapsed* to obtain a “proper” model. Both techniques preserve certain structural properties.

Our strategy to construct a countermodel with finitely many nominals consists of the following three steps:

- Take an arbitrary countermodel and unravel it.
- Transform the obtained forest quasi-model by substituting critical parts by well-behaved ones,
- Collapse the obtained structure into a (proper) model.

The mentioned “critical parts” are those giving rise to new nominals. They have to be – at least largely – avoided (we do not care about a finite set of those critical parts remaining).

The central question is: where do these mysterious well-behaved substitutes come from? Fortunately, the plethora of critical parts brings about its own remedy. We can use infinite sets of critical parts to construct well-behaved ones in an infinite approximation process. We thereby obtain parts which have not been present in our structure before, but are well compatible with it and can hence be used for its reorganization.

After having informally introduced our main line of argumentation, we now move on to the technical details.

Preliminaries

The basic elements in a DL are atomic concepts (unary predicates), atomic roles (binary predicates), and individuals (constants). In the basic DL *ALC*, complex concepts can be built from atomic ones by using negation ($\neg C$), conjunction ($C_1 \sqcap C_2$), disjunction ($C_1 \sqcup C_2$), or by quantification

over a role ($\forall r.C$ and $\exists r.C$), which have typical set-theoretic first-order logic interpretations. The DL *ALCOIFb* further allows for nominals, which are concepts defined as a singleton set containing a constant ($\{o\}$), inverse roles (r^- interpreted as $\{\langle y, x \rangle \mid \langle x, y \rangle \in r^{\mathcal{I}}\}$), functionality constraints ($\text{func}(f)$), which require that the interpretation of f is a functional relation, and safe Boolean combinations of roles. A Boolean role expression is “safe” if its disjunctive normal form contains a positive conjunct in every disjunct. The DL *ALCHOIQb* further allows for role hierarchies and qualified number restrictions (counting quantifiers), but these features can be eliminated by a polynomial reduction, while preserving query (non-)entailment (Rudolph, Krötzsch, and Hitzler 2008) and, w.l.o.g., we consider only *ALCOIFb*.

In the remainder, we use A and B for atomic concepts, o for an individual name, r for an atomic role, U for a safe Boolean role expression, and f for a role (atomic or inverse) that is declared functional.

An *ALCOIFb* knowledge base is a finite set of general concept inclusions (GCIs) $C \sqsubseteq D$ ($C \equiv D$ abbreviates $C \sqsubseteq D$ and $D \sqsubseteq C$) and functionality constraints, where C and D are *ALCOIFb* concepts. W.l.o.g., we do not consider ABoxes since with nominals, the ABox can be internalized into the TBox, and we assume that GCIs are *simplified* into the following forms:

$$\prod A_i \sqsubseteq \prod B_j \mid A \equiv \{o\} \mid A \sqsubseteq \forall U.B \mid A \sqsubseteq \exists U.B$$

If $i = 0$, we interpret $\prod A_i$ as \top and if $j = 0$, we interpret $\prod B_j$ as \perp . We use $\text{con}(\mathcal{K})$, $\text{rol}(\mathcal{K})$, and $\text{nom}(\mathcal{K})$ to denote, respectively, the set of concept, role, and individual names occurring in \mathcal{K} , and $\text{cl}(\mathcal{K})$ to denote the closure of \mathcal{K} . A role f is (inverse) functional in \mathcal{K} if \mathcal{K} contains an axiom $\text{func}(f)$ ($\text{func}(f^-)$).

Let N_V be a countably infinite set of variables containing x and y . An *atom* is an expression $A(x)$ or $r(x, y)$. A *Boolean conjunctive query* q is a non-empty set of atoms. We use $\text{var}(q)$ to denote the set of (existentially quantified) variables occurring in q and $\#(q)$ for the number of atoms in q . For $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ an interpretation, $A(x), r(x, y)$ atoms, and $\pi: \text{var}(q) \rightarrow \Delta^{\mathcal{I}}$ a total function, we write (i) $\mathcal{I} \models^{\pi} A(x)$ if $\pi(x) \in A^{\mathcal{I}}$ and (ii) $\mathcal{I} \models^{\pi} r(x, y)$ if $\langle \pi(x), \pi(y) \rangle \in r^{\mathcal{I}}$. If $\mathcal{I} \models^{\pi} \text{At}$ for all atoms $\text{At} \in q$, we write $\mathcal{I} \models^{\pi} q$ and say that \mathcal{I} *satisfies* q . We write $\mathcal{I} \models q$ if there exists a function π such that $\mathcal{I} \models^{\pi} q$ and call π a *match* for q in \mathcal{I} . If $\mathcal{I} \models \mathcal{K}$ implies $\mathcal{I} \models q$, we say that \mathcal{K} *entails* q and write $\mathcal{K} \models q$. W.l.o.g., we assume that queries are connected. Given a KB \mathcal{K} and a CQ q , the *query entailment problem* is to decide whether $\mathcal{K} \models q$.

Unless stated otherwise, we use q for a connected Boolean conjunctive query, \mathcal{K} for a simplified *ALCOIFb* knowledge base, and \mathcal{I} for an interpretation $(\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$. As a running example, we use a KB \mathcal{K} containing the axioms $\{o\} \sqsubseteq \exists r.A, A \sqsubseteq \exists r.A, A \sqsubseteq \exists s.B, B \sqsubseteq C \sqcup D, C \sqsubseteq \exists f.E, D \sqsubseteq \exists g.E, E \sqsubseteq B \sqcup \{o\}, \text{func}(f^-), \text{func}(g^-)$. Figure 2 a) (p. 6) displays a representation of a model for \mathcal{K} .

Model Construction

We first introduce interpretations and models that have a kind of forest shape. However, the notion of a forest is very

weak since we do also allow for arbitrary relations between tree elements and roots.

Definition 1. A tree T is a non-empty, prefix-closed subset of \mathbb{N}^* . A forest F is a subset of $R \times \mathbb{N}^*$, where R is a countable, possibly infinite set of elements $\{\rho_1, \dots, \rho_n\}$ such that, for each $\rho \in R$, the set $\{w \mid (\rho, w) \in F\}$ is a tree. Each pair $(\rho, \varepsilon) \in F$ is called a root of F . For $(\rho, w), (\rho', w') \in F$, we call (ρ', w') a successor (predecessor) of (ρ, w) if $\rho' = \rho$ and $w' = w \cdot c$ ($w = w' \cdot c$) for some $c \in \mathbb{N}$, where “ \cdot ” denotes concatenation; (ρ', w') is a neighbor of (ρ, w) if (ρ', w') is a successor of (ρ, w) or vice versa. A node (ρ, w) is an ancestor (descendant) of a node (ρ', w') if $\rho = \rho'$ and w is a prefix of w' (w' is a prefix of w). We use $|w|$ to denote the length of w . The branching degree $d(w)$ of a node w in a tree T is the number of successors of w .

A forest interpretation of \mathcal{K} is an interpretation \mathcal{I} that satisfies:

F11 $\Delta^{\mathcal{I}}$ is a forest with roots R ;

F12 there is a total and surjective function $\lambda: \text{nom}(\mathcal{K}) \rightarrow R \times \{\varepsilon\}$ s.t. $\lambda(o) = (\rho, \varepsilon)$ iff $o^{\mathcal{I}} = (\rho, \varepsilon)$;

F13 for each role $r \in \text{rol}(\mathcal{K})$, if $\langle (\rho, w), (\rho', w') \rangle \in r^{\mathcal{I}}$, then either (a) $w = \varepsilon$ or $w' = \varepsilon$, or (b) (ρ, w) is a neighbor of (ρ', w') .

If $\mathcal{I} \models \mathcal{K}$ we say that \mathcal{I} is a forest model for \mathcal{K} . With $\text{nomFree}(\mathcal{K})$, we denote a knowledge base obtained from \mathcal{K} by replacing each nominal concept $\{o\}$ with $o \in \text{nom}(\mathcal{K})$ with a fresh concept name N_o . A forest quasi-interpretation for \mathcal{K} is an interpretation \mathcal{J} that satisfies F11 and F13, and the adapted version F12' of F12 that there is a total and surjective function $\lambda: \text{nom}(\mathcal{K}) \rightarrow R \times \{\varepsilon\}$ s.t. $\lambda(o) = (\rho, \varepsilon)$ iff $(\rho, \varepsilon) \in N_o^{\mathcal{J}}$ (there might be other $(\rho, w) \in \Delta^{\mathcal{J}}$ with $w \neq \varepsilon$ s.t. $(\rho, w) \in N_o^{\mathcal{J}}$). If $\mathcal{J} \models \text{nomFree}(\mathcal{K})$ we say that \mathcal{J} is a forest quasi-model for \mathcal{K} . A forest (quasi) interpretation \mathcal{I} is a strict forest (quasi) interpretation if, in condition F13, only (b) is allowed; it is a tree interpretation, if it has a single root. If there is a k such that $d(w) \leq k$ for each $(\rho, w) \in \Delta^{\mathcal{I}}$, then we say that \mathcal{I} has branching degree k .

Let $\mathcal{I}, \mathcal{I}'$ be two forest interpretations of \mathcal{K} with $\delta_1, \delta_2 \in \Delta^{\mathcal{I}}, \delta'_1, \delta'_2 \in \Delta^{\mathcal{I}'}$. The pairs $\langle \delta_1, \delta_2 \rangle, \langle \delta'_1, \delta'_2 \rangle$ are isomorphic w.r.t. \mathcal{K} , written $\langle \delta_1, \delta_2 \rangle \cong_{\mathcal{K}} \langle \delta'_1, \delta'_2 \rangle$ iff

- $\langle \delta_1, \delta_2 \rangle \in r^{\mathcal{I}}$ iff $\langle \delta'_1, \delta'_2 \rangle \in r^{\mathcal{I}'}$ for each $r \in \text{rol}(\mathcal{K})$,
- $\delta_i \in A^{\mathcal{I}}$ iff $\delta'_i \in A^{\mathcal{I}'}$ for each $i \in \{1, 2\}, A \in \text{con}(\mathcal{K})$,
- $\delta_i = o^{\mathcal{I}}$ iff $\delta'_i = o^{\mathcal{I}'}$ for each $i \in \{1, 2\}, o \in \text{nom}(\mathcal{K})$.

We say that \mathcal{I} and \mathcal{I}' are isomorphic w.r.t. \mathcal{K} , written: $\mathcal{I} \cong_{\mathcal{K}} \mathcal{I}'$, if there is a bijection $\varphi: \Delta^{\mathcal{I}} \rightarrow \Delta^{\mathcal{I}'}$ such that, for each $\delta_1, \delta_2 \in \Delta^{\mathcal{I}}, \langle \delta_1, \delta_2 \rangle \cong_{\mathcal{K}} \langle \varphi(\delta_1), \varphi(\delta_2) \rangle$ and δ_1 is a successor of δ_2 iff $\varphi(\delta_1)$ is a successor of $\varphi(\delta_2)$.

If clear from the context, we omit the subscript \mathcal{K} of $\cong_{\mathcal{K}}$ and we extend the definition to forest quasi-interpretations in the obvious way. Forest quasi-models represent, intuitively, an intermediate step between arbitrary models of \mathcal{K} and forest models of \mathcal{K} .

Since KBs are assumed to be simplified, it can mostly be checked locally (by looking at an element of the domain and its direct neighbors) whether an interpretation \mathcal{I} is a model

of \mathcal{K} . Only nominals impose a global restriction on the cardinality of concepts. We call an element $\delta \in \Delta^{\mathcal{I}}$ *locally \mathcal{K} -consistent* if it satisfies each GCI in \mathcal{K} (a functionality restriction $\text{func}(f)$ is satisfied if δ has at most one f -neighbor); \mathcal{I} is a model of \mathcal{K} if each $\delta \in \Delta^{\mathcal{I}}$ is locally \mathcal{K} -consistent and, for each $o \in \text{nom}(\mathcal{K})$, there is exactly one element $\delta \in \Delta^{\mathcal{I}}$ such that $o^{\mathcal{I}} = \delta$. We now show how we can obtain a forest quasi-model from a model of \mathcal{K} by using an adapted version of unraveling (Glimm et al. 2008).

Definition 2. Let \mathcal{I} be a model for \mathcal{K} and choose a function that returns, for a concept $C = \exists U.B \in \text{cl}(\mathcal{K})$ and $\delta \in C^{\mathcal{I}}$ some $\delta_{C,\delta} \in \Delta^{\mathcal{I}}$ s.t. $\langle \delta, \delta_{C,\delta} \rangle \in U^{\mathcal{I}}$ and $\delta_{C,\delta} \in B^{\mathcal{I}}$. For each $\delta \in C_1^{\mathcal{I}} \cap C_2^{\mathcal{I}}$ with $C_i = \exists U_i.B_i \in \text{cl}(\mathcal{K})$ and choose $(C_i, \delta) = \delta_i$ for $i \in \{1, 2\}$, w.l.o.g., we assume that if $\langle \delta, \delta_1 \rangle \cong \langle \delta, \delta_2 \rangle$, then $\delta_1 = \delta_2$.

An unraveling for some $\delta \in \Delta^{\mathcal{I}}$, denoted $\downarrow(\mathcal{I}, \delta)$, is an interpretation obtained from \mathcal{I} and δ as follows: we define the set $S \subseteq (\Delta^{\mathcal{I}})^*$ of sequences to be the smallest set such that δ is a sequence and $\delta_1 \cdots \delta_n \cdot \delta_{n+1}$ is a sequence, if

- $\delta_1 \cdots \delta_n$ is a sequence,
- if $n > 2$ and $\langle \delta_n, \delta_{n-1} \rangle \in f^{\mathcal{I}}$ for some functional role f , then $\delta_{n+1} \neq \delta_{n-1}$,
- $\delta_{n+1} = \text{choose}(C, \delta_n)$ for some $C = \exists U.B \in \text{cl}(\mathcal{K})$.

Now fix a set $F \subseteq \{\delta\} \times \mathbb{N}^*$ and a bijection $\lambda: F \rightarrow S$ such that (i) F is a forest, (ii) $\lambda(\delta, \varepsilon) = \delta$, (iii) if $(\delta, w), (\delta, w \cdot c) \in F$ with $w \cdot c$ a successor of w , then $\lambda(\delta, w \cdot c) = \lambda(\delta, w) \cdot \delta_{n+1}$ for some $\delta_{n+1} \in \Delta^{\mathcal{I}}$. For each $(\delta, w) \in F$, set $\text{Tail}(\delta, w) = \delta_n$ if $f(\delta, w) = \delta_1 \cdots \delta_n$. The unraveling for δ is the interpretation \mathcal{J} with $\Delta^{\mathcal{J}} = F$ and, for each $(\delta, w) \in \Delta^{\mathcal{J}}$:

- for each $o \in \text{nom}(\mathcal{K})$, $N_o^{\mathcal{J}} = \{(\delta, w) \in \Delta^{\mathcal{J}} \mid \text{Tail}(\delta, w) \in o^{\mathcal{I}}\}$, for $N_o \in N_C$ a fresh concept name;
- for each concept name $A \in \text{con}(\mathcal{K})$, $(\delta, w) \in A^{\mathcal{J}}$ iff $\text{Tail}(\delta, w) \in A^{\mathcal{I}}$;
- for each role name $r \in \text{rol}(\mathcal{K})$, $\langle (\delta, w), (\delta, w') \rangle \in r^{\mathcal{J}}$ iff (δ, w') is a neighbor of (δ, w) , and $\langle \text{Tail}(\delta, w), \text{Tail}(\delta, w') \rangle \in r^{\mathcal{I}}$.

Let $R \subseteq \Delta^{\mathcal{I}}$ contain each δ s.t. $o^{\mathcal{I}} = \delta$ for some $o \in \text{nom}(\mathcal{K})$. The union of all $\downarrow(\mathcal{I}, \delta)$ with $\delta \in R$ is called an unraveling for \mathcal{I} , denoted $\downarrow(\mathcal{I})$, where unions of interpretations are defined in the natural way.

Please note that the function Tail can also be seen as a homomorphism (up to signature extension) from the elements in the unraveling to elements in the original model. Figure 2 b) shows the unraveling for our example KB and model. The dotted lines for the non-root elements labeled N_o indicate that a copy of the whole tree should be appended.

Unravelings are the first step in the process of transforming an arbitrary model of \mathcal{K} into a forest model since the resulting interpretation is a strict forest quasi-model of \mathcal{K} .

Lemma 3. Let \mathcal{I} be a model of \mathcal{K} , then $\mathcal{J} = \downarrow(\mathcal{I})$ is a strict forest quasi-model for \mathcal{K} with branching degree bounded in $|\text{cl}(\mathcal{K})|$.

We now show how we can collapse an unraveling back into a model. Without any modifications to the unraveling, we have not gained much other than a “tidier” model that

does not contain unnecessary elements, but we later introduce modifications for an unraveling such that we can collapse a modified unraveling into a model that contains only a finite number of nominals. In the following steps, we traverse a forest quasi-model in an order in which elements with smaller tree depth are always of smaller order than elements with greater tree depth. Elements with the same tree depth are ordered lexicographically. The bounded branching degree of unravelings then guarantees that, after a finite number of steps, we go on to the next level in the forest and process all nodes eventually. Further, we can merge nodes such that, finally, all nominal placeholders (in the extension of some N_o) can be interpreted as nominals without violating functionality restrictions. In fact, we do not only have to merge nominal placeholders, but also elements that are related to a nominal placeholder by an inverse functional role since, by definition of the semantics, these elements have to correspond to the same element in a model. In order to identify such elements, we define *backwards counting paths* as follows:

Definition 4. Let \mathcal{I} be a (quasi) forest model for \mathcal{K} . We call $p = \delta_1 \cdots \delta_n$ a path from δ_1 to δ_n if, for each i with $1 \leq i < n$, $\langle \delta_i, \delta_{i+1} \rangle \in r_i^{\mathcal{I}}$ for some role $r_i \in \text{rol}(\mathcal{K})$. The length $|p|$ of a path p is $n - 1$. We write $\delta_1 \xrightarrow{U_1} \delta_2 \cdots \xrightarrow{U_{n-1}} \delta_n$ to denote that $\langle \delta_i, \delta_{i+1} \rangle \in U_i^{\mathcal{I}}$ for each $1 \leq i < n$. The path p is a descending path if there is some $(\rho, \varepsilon) \in \Delta^{\mathcal{I}}$ s.t., for each $1 \leq i \leq n$, $\delta_i = (\rho, w_i)$ and, for each $1 \leq i < n$, $|w_i| < |w_{i+1}|$; p is a backwards counting path (BCP) in \mathcal{I} if $\delta_n \in o^{\mathcal{I}}$ ($\delta_n \in N_o^{\mathcal{I}}$) for some $o \in \text{nom}(\mathcal{K})$ and, for each $1 \leq i < n$, $\langle \delta_i, \delta_{i+1} \rangle \in f_i^{\mathcal{I}}$ for some inverse functional role f_i ; p is a descending BCP if it is descending and a BCP.

Given a BCP $p = \delta_1 \xrightarrow{f_1} \delta_2 \cdots \xrightarrow{f_n} \delta_{n+1}$ with $\delta_{n+1} \in o^{\mathcal{J}}$ ($\delta_{n+1} \in N_o^{\mathcal{J}}$), we call the sequence $f_1 \cdots f_n$ a path sketch of p .

Let $<$ be a strict total order over N_I , \mathcal{K} a consistent *ALCOIFb* KB, and \mathcal{J} a forest quasi-interpretation for \mathcal{K} . We extend the order to elements in $\Delta^{\mathcal{J}}$ as follows: let $w_1 = w_p \cdot c_1^1 \cdots c_1^n$, $w_2 = w_p \cdot c_2^1 \cdots c_2^m \in \mathbb{N}^*$ where $w_p \in \mathbb{N}^*$ is the longest common prefix of w_1 and w_2 , then $w_1 < w_2$ if either $n < m$ or both $n = m$ and $c_1^1 < c_2^1$. For $i \in \{1, 2\}$ and $(\rho_i, \varepsilon) \in \Delta^{\mathcal{J}}$, let $o_i \in \text{nom}(\mathcal{K})$ be the smallest nominal such that $(\rho_i, \varepsilon) \in N_{o_i}^{\mathcal{J}}$. Now $(\rho_1, w_1) < (\rho_2, w_2)$ if either (i) $|w_1| < |w_2|$ or (ii) $|w_1| = |w_2|$ and $o_1 < o_2$ or (ii) $|w_1| = |w_2|$, $o_1 = o_2$ and $w_1 < w_2$. When collapsing, we create new elements of the form $(\rho w, w')$ from $(\rho, w w')$. We extend, therefore, the order as follows: $(\rho_1 w_1, w'_1) < (\rho_2 w_2, w'_2)$ if $(\rho_1, w_1 w'_1) < (\rho_2, w_2 w'_2)$.

Please note that (ρ, w) is already a descending BCP if $(\rho, w) \in o^{\mathcal{J}}$ ($N_o^{\mathcal{J}}$). We now show how we can “collapse” a forest quasi-model into a forest model provided it satisfies some admissibility restrictions. During the traversal, we distinguish two situations: (i) we encounter an element (ρ, w) that starts a descending BCP and we have not seen another element before that starts a descending BCP with the same path sketch. In this case, we promote (ρ, w) to become a new root node of the form $(\rho w, \varepsilon)$ and we shift the subtree rooted in (ρ, w) with it; (ii) we encounter a node (ρ, w) that




Figure 2: a) A representation of a model for the running example (left). b) Result of unraveling this model (middle). c) Result of collapsing this unraveling with infinitely many new root elements displayed in the top line (right).

starts a descending BCP, but we have already seen a node (ρ', w') that starts a descending BCP with that path sketch and which is now a root of the form $(\rho'w', \varepsilon)$. In this case, we delete the subtree rooted in (ρ, w) and identify (ρ, w) with $(\rho'w', \varepsilon)$. If (ρ, w) is an f -successor of its predecessor for some inverse functional role f , we delete all f -successors of $(\rho'w', \varepsilon)$ and their subtrees in order to satisfy the functionality restriction. We use a notion of collapsing admissibility to characterize models s.t. the predecessor of (ρ, w) satisfies the same atomic concepts as the deleted successor of (ρ', w') , which ensures that local consistency is preserved.

Definition 5. Let $\mathcal{K}' = \text{nomFree}(\mathcal{K})$ and \mathcal{J} a forest quasi-interpretation for \mathcal{K} . We define \sim as the smallest equivalence relation on $\Delta^{\mathcal{J}}$ that satisfies $\delta_1 \sim \delta_2$ if δ_1, δ_2 start descending BCPs with identical path sketches.

If \mathcal{J} is a strict forest quasi-model for \mathcal{K} , we call $\mathcal{J}_0 = \mathcal{J}$ an initial collapsing for \mathcal{J} and the smallest element $(\rho_0, w_0) \in \Delta^{\mathcal{J}_0}$ with $w_0 \neq \varepsilon$ that starts a descending BCP the focus of \mathcal{J}_0 . Let \mathcal{J}_i be a collapsing for \mathcal{J} and $(\rho_i, w_i) \in \Delta^{\mathcal{J}_i}$ the focus of \mathcal{J}_i . We obtain a collapsing \mathcal{J}_{i+1} for \mathcal{J} from \mathcal{J}_i with focus (ρ_{i+1}, w_{i+1}) the smallest element starting a descending BCP and $(\rho_{i+1}, w_{i+1}) > (\rho_i, w_i)$ according to the following two cases:

1. There is no element $(\rho, \varepsilon) \in \Delta^{\mathcal{J}_i}$ s.t. $(\rho, \varepsilon) < (\rho_i, w_i)$ and $(\rho, \varepsilon) \sim (\rho_i, w_i)$. Then \mathcal{J}_{i+1} is obtained from \mathcal{J}_i by renaming each element $(\rho_i, w_i w'_i) \in \Delta^{\mathcal{J}_i}$ to $(\rho_i w'_i, w'_i)$.
2. There is an element $(\rho, \varepsilon) \in \Delta^{\mathcal{J}_i}$ s.t. $(\rho, \varepsilon) < (\rho_i, w_i)$ and $(\rho, \varepsilon) \sim (\rho_i, w_i)$. Let (ρ, ε) be the smallest such element.
 - (a) $\Delta^{\mathcal{J}_{i+1}} = \Delta^{\mathcal{J}_i} \setminus (\{(\rho_i, w_i w'_i) \mid w'_i \in \mathbb{N}^*\} \cup \{(\rho, w) \mid w = c \cdot w', c \in \mathbb{N}, w' \in \mathbb{N}^*, (\rho_i, w_i)$ has a predecessor (ρ_i, w'_i) such that $\langle (\rho_i, w'_i), (\rho_i, w_i) \rangle \in f^{\mathcal{J}_i}$ for an inverse functional role f in $\text{rol}(\mathcal{K})$ and $\langle (\rho, c), (\rho, \varepsilon) \rangle \in f^{\mathcal{J}_i} \})$;
 - (b) for each $A \in \text{con}(\mathcal{K})$ and $\delta \in \Delta^{\mathcal{J}_{i+1}}$, $\delta \in A^{\mathcal{J}_{i+1}}$ iff $\delta \in A^{\mathcal{J}_i}$;
 - (c) for each $r \in \text{rol}(\mathcal{K})$ and $\delta_1, \delta_2 \in \Delta^{\mathcal{J}_{i+1}}$, $\langle \delta_1, \delta_2 \rangle \in r^{\mathcal{J}_{i+1}}$ iff (a) $\langle \delta_1, \delta_2 \rangle \in r^{\mathcal{J}_i}$ or (b) δ_1 is the predecessor of (ρ_i, w_i) in \mathcal{J}_i , $\delta_2 = (\rho, \varepsilon)$, and $\langle \delta_1, (\rho_i, w_i) \rangle \in r^{\mathcal{J}_i}$.

For a collapsing \mathcal{J}_i , $\text{safe}(\mathcal{J}_i)$ is the restriction of \mathcal{J}_i to elements (ρ, w) s.t. $(\rho, w) \in \mathcal{J}_j$ for all $j \geq i$. With \mathcal{J}_ω we denote the non-disjoint union of all interpretations $\text{safe}(\mathcal{J}_i)$ obtained from subsequent collapsings \mathcal{J}_i for \mathcal{J} . The interpretation obtained from \mathcal{J}_ω by interpreting each

$o \in \text{nom}(\mathcal{K})$ as $(\rho, \varepsilon) \in N_o^{\mathcal{J}_\omega}$ is denoted by $\text{collapse}(\mathcal{J})$ and called a purified interpretation w.r.t. \mathcal{J} . If $\text{collapse}(\mathcal{J}) \models \mathcal{K}$, we call $\text{collapse}(\mathcal{J})$ a purified model of \mathcal{K} .

We use the notion of collapsing-admissibility to characterize quasi-interpretations that become forest models of the KB after we collapse them.

Definition 6. Let \mathcal{J} be a forest quasi-interpretation for \mathcal{K} , then \mathcal{J} is collapsing-admissible if there exists a function $\text{ch}: (\text{cl}(\mathcal{K}) \times \Delta^{\mathcal{J}}) \rightarrow \Delta^{\mathcal{J}}$ s.t., for each $C = \exists U.B \in \text{cl}(\mathcal{K})$ and $\delta \in C^{\mathcal{J}}$,

1. $\langle \delta, \text{ch}(C, \delta) \rangle \in U^{\mathcal{J}}$, $\text{ch}(C, \delta) \in B^{\mathcal{J}}$ and, if there is no functional role f s.t. $\langle \delta, \text{ch}(C, \delta) \rangle \in f^{\mathcal{J}}$, then $\text{ch}(C, \delta)$ is a successor of δ ,
2. if there is some $\delta' \in C^{\mathcal{J}}$ s.t. δ and δ' start descending BCPs with identical path sketches, then $\langle \delta, \text{ch}(C, \delta) \rangle \cong \langle \delta', \text{ch}(C, \delta') \rangle$.

Since, in unravelings, elements that start (descending) BCPs with identical path sketches have been generated from the same element in the unravelled model, collapsing-admissibility is immediate and the function ch can be defined using the function choose from the unraveling. We can use collapsing-admissibility to show that, whenever we delete an element and the subtree rooted in it during the collapsing, the predecessor of the focus is a suitable replacement.

Lemma 7. Let \mathcal{J} be a strict forest quasi-model for \mathcal{K} with branching degree b that is collapsing-admissible. Then $\text{collapse}(\mathcal{J})$ is a forest model for \mathcal{K} that still has branching degree b .

Since unravelings are strict forest quasi-models with bounded branching degree by Lemma 3 and collapsing-admissible, it is an immediate consequence of Lemma 7 that collapsing an unraveling yields a forest model with bounded branching degree. At this point, the number of roots might still be infinite and we could have obtained the same result by unraveling an arbitrary model, where we take all elements on BCPs as roots instead of taking just the nominals and creating new roots in the collapsing process. In the next sections, however, we show how we can transform an unraveling of a counter-model for the query such that it remains collapsing-admissible and such that it can in the end be collapsed into a forest model with a finite number of roots that is still a counter model for the query. For

this transformation it is much more convenient to work with real (strict) trees and forests, which is why we use (strict) forest quasi-interpretations.

Quasi-Entailment in Quasi-Models

In this section, we provide a characterization for query entailment in forest quasi-models that mirrors query entailment for the corresponding “proper models”. In our further argumentation, we talk about the initial part of a tree, i.e., the part that remains if one cuts branches down to a fixed length. For a forest interpretation \mathcal{I} and some $n \in \mathbb{N}$, we denote, therefore, with $\text{cut}_n(\mathcal{I})$ the interpretation obtained from \mathcal{I} by restricting $\Delta^{\mathcal{I}}$ to those pairs (ρ, w) for which $|w| \leq n$. One can show that, in the case of purified models, we find only finitely many unraveling trees of depth n that “look different” (i.e., that are non-isomorphic).

Lemma 8. *If \mathcal{K} is consistent, then there is a purified interpretation \mathcal{I} such that $\mathcal{I} \models \mathcal{K}$ and, for every $n \in \mathbb{N}$, there are only finitely many non-isomorphic trees of depth n .*

For our further considerations, we introduce the notion of “anchored n -components”. These are certain substructures of forest quasi-interpretations that we use to define a notion of quasi entailment.

Definition 9. *Let \mathcal{J} be a forest quasi-interpretation and $\delta \in \Delta^{\mathcal{J}}$. An interpretation \mathcal{C} is called anchored n -component of \mathcal{J} with witness δ if \mathcal{C} can be created by restricting \mathcal{J} to a set $W \subseteq \Delta^{\mathcal{J}}$ obtained as follows:*

- Let \mathcal{J}_δ be the subtree of \mathcal{J} that is started by δ and let $\mathcal{J}_{\delta,n} := \text{cut}_n(\mathcal{J}_\delta)$. Select a subset $W' \subseteq \Delta^{\mathcal{J}_{\delta,n}}$ that is closed under predecessors.
- For every $\delta' \in W'$, let P be a finite set (possibly empty) of descending BCPs \mathbf{p} starting from δ' and let $W_{\delta'}$ contain all nodes from all $\mathbf{p} \in P$.
- Set $W = W' \cup \bigcup_{\delta' \in W'} W_{\delta'}$.

The following definition and lemma employ the notion of anchored n -components to come up with the notion of *quasi-entailment* (short for quasi-entailment), a criterion that reflects query-entailment in the world of forest quasi-models.

Definition 10. *Let \mathcal{J} be a forest quasi-model for \mathcal{K} and q a CQ with $\sharp(q) = n$ and $V = \text{var}(q)$. We say that \mathcal{J} *quasi-entails* q , written $\mathcal{J} \approx q$, if \mathcal{J} contains anchored n -components $\mathcal{C}_1, \dots, \mathcal{C}_\ell$ and there are variable assignment functions $\mu_i: V \rightarrow 2^{\Delta^{\mathcal{C}_i}}$ such that:*

- Q1 For every $x \in V$, there is at least one \mathcal{C}_i , such that $\mu_i(x) \neq \emptyset$
- Q2 For all $A(x) \in q$, we have $\mu_i(x) \subseteq A^{\mathcal{J}}$ for some i .
- Q3 For every $r(x, y) \in q$ there is a \mathcal{C}_i such that there are $\delta_1 \in \mu_i(x)$ and $\delta_2 \in \mu_i(y)$ such that $\langle \delta_1, \delta_2 \rangle \in r^{\mathcal{J}}$.
- Q4 If, for some $x \in V$, there are anchored n -components \mathcal{C}_i and \mathcal{C}_j with $\delta \in \mu_i(x)$ and $\delta' \in \mu_j(x)$, then there is
 - a sequence $\mathcal{C}_{n_1}, \dots, \mathcal{C}_{n_k}$ with $n_1 = i$ and $n_k = j$ and
 - a sequence $\delta_1, \dots, \delta_k$ with $\delta_1 = \delta$ and $\delta_k = \delta'$ as well as $\delta_m \in \mu_{n_m}(x)$ for all $1 \leq m < k$,
such that, for every m with $1 \leq m < k$, we have that
 - \mathcal{C}_{n_m} contains a descending BCP p_1 started by δ_m ,




Figure 3: An example for the correspondence between entailment and quasi-entailment.

- $\mathcal{C}_{n_{m+1}}$ contains a descending BCP p_2 started by δ_{m+1} ,
- p_1 and p_2 have the same path sketch.

Note that an anchored component may contain none, one or several instantiations of a variable $x \in V$. Intuitively, the definition ensures, that we find matches of query parts which when fitted together by identifying BCP-equal elements yield a complete query match. In Figure 3, we illustrate the correspondence between entailment and quasi-entailment for the query $q = \{A(x_1), D(x_3), s(x_1, x_2), g(x_3, x_2), s(x_4, x_3), r(x_1, x_4)\}$. The dotted lines indicate a match into the “proper model” (left) and the dashed lines a match into the quasi-model (right) for a single anchored n -component. Note that the function witnessing the quasi-entailment maps x_2 to a set containing two elements, but these two elements will be merged during the collapsing process because they are on a BCP.

Lemma 11. *For any model \mathcal{I} of \mathcal{K} , $\downarrow(\mathcal{I}) \approx q$ implies $\mathcal{I} \models q$ and, for any collapsing-admissible strict forest quasi-model \mathcal{J} of \mathcal{K} , $\text{collapse}(\mathcal{J}) \approx q$ implies $\mathcal{J} \approx q$.*

Limits and Forest Transformations

One of the major obstacles for a decision procedure for CQ entailment is that for DLs including inverses, nominals, and cardinality restrictions (or alternatively functionality), there are potentially infinitely many new roots. If we want to eliminate new roots such that only finitely many remain, they have to be replaced by “uncritical” elements. We will construct such elements as “environment-limits” – new domain elements which can be approximated with arbitrary precision by already present domain elements – possibly without themselves being present in the domain.³

Definition 12. *Let \mathcal{I} be a model of \mathcal{K} and let $\delta \in \Delta^{\mathcal{I}}$. A tree interpretation \mathcal{J} is said to be generated by δ , written: $\mathcal{J} \triangleleft \delta$, if it is isomorphic to the restriction of $\downarrow(\mathcal{I}, \delta)$ to elements of $\{(\delta, cw) \mid (\delta, cw) \in \Delta^{\downarrow(\mathcal{I}, \delta)}, c \notin H\}$ for some $H \subseteq \mathbb{N}$. The set of limits of \mathcal{I} , written $\text{lim } \mathcal{I}$, is the set of all tree interpretations \mathcal{J} s.t., for every $k \in \mathbb{N}$, there are infinitely many $\delta \in \Delta^{\mathcal{I}}$ with $\text{cut}_k(\mathcal{L}) \cong \text{cut}_k(\mathcal{J})$ for some $\mathcal{L} \triangleleft \delta$.*

³As an analogy, consider the fact that any real number can be approximated by a sequence of rational numbers, even if it is itself irrational.




Figure 4: From left to right: a) a 2-secure replacement for δ in b) the quasi-forest model for \mathcal{K} . c) The result of replacing δ by the 2-secure replacement (the inserted component is highlighted). d) The result of collapsing the forest quasi-model from c).

Figure 4 a) displays a limit element for our example model. The following lemma gives some useful properties of limits.

Lemma 13. *Let $\mathcal{K}' = \text{nomFree}(\mathcal{K})$, \mathcal{I} a purified model of \mathcal{K} , and n some fixed natural number. Then the following hold:*

1. *Let \mathcal{L}' be a tree interpretation such that there are infinitely many $\delta \in \Delta^{\mathcal{I}}$ with $\mathcal{L}' \cong \text{cut}_n(\mathcal{L})$ for some $\mathcal{L} \triangleleft \delta$. Then, there is at least one limit $\mathcal{J} \in \lim \mathcal{I}$ such that $\text{cut}_n(\mathcal{J}) \cong \mathcal{L}'$.*
2. *Every $\mathcal{J} \in \lim \mathcal{I}$ is locally \mathcal{K}' -consistent apart from its root (ρ, ε) .*
3. *For every $\mathcal{J} \in \lim \mathcal{I}$, every root (ρ, ε) in \mathcal{J} has no BCP to any $(\rho, w) \in \Delta^{\mathcal{J}}$.*
4. *Every $\mathcal{J} \in \lim \mathcal{I}$ is collapsing-admissible.*

Having defined and justified limit elements as convenient building blocks for restructuring forest quasi-interpretations, the following definition states how this restructuring is carried out.

Definition 14. *Let \mathcal{I} be a model for \mathcal{K} and \mathcal{J} some forest quasi-model for \mathcal{K} with $\delta \in \Delta^{\mathcal{J}}$. A strict tree quasi-interpretation $\mathcal{J}' \in \lim \mathcal{I}$ is called an n -secure replacement for δ if (i) $\text{cut}_n(\downarrow(\mathcal{J}, \delta))$ is isomorphic to $\text{cut}_n(\mathcal{J}')$ and (ii) for every anchored n -component of \mathcal{J}' with witness δ' , there is an isomorphic anchored n -component of \mathcal{J} with witness δ . If $\delta \in \Delta^{\mathcal{J}}$ has an n -secure replacement in $\lim \mathcal{I}$, δ is n -replaceable w.r.t. \mathcal{I} and it is n -irreplaceable w.r.t. \mathcal{I} otherwise.*

Now, let $(\rho, w) \in \Delta^{\mathcal{J}}$ be n -replaceable w.r.t. \mathcal{I} and \mathcal{J}' an according n -replacement for (ρ, w) from $\lim \mathcal{I}$ with root (ς, ε) . The result of replacing (ρ, w) by \mathcal{J}' is an interpretation \mathcal{R} with $\Delta^{\mathcal{R}} = \Delta_{\text{red}}^{\mathcal{J}} \cup \{(\rho, ww') \mid (\varsigma, w'') \in \Delta^{\mathcal{J}'}\}$ for $\Delta_{\text{red}}^{\mathcal{J}} = (\Delta^{\mathcal{J}} \setminus \{(\rho, ww') \mid |w'| > 1\})$ s.t.

- for each $A \in \text{con}(\mathcal{K}')$, $A^{\mathcal{R}} = (A^{\mathcal{J}} \cap \Delta_{\text{red}}^{\mathcal{J}}) \cup \{(\rho, ww') \mid (\varsigma, w') \in A^{\mathcal{J}'}\}$
- for each $r \in \text{rol}(\mathcal{K}')$, $r^{\mathcal{R}} = (r^{\mathcal{J}} \cap \Delta_{\text{red}}^{\mathcal{J}} \times \Delta_{\text{red}}^{\mathcal{J}}) \cup \{(\rho, ww'), (\rho, ww'') \mid ((\varsigma, w'), (\varsigma, w'')) \in r^{\mathcal{J}'}\}$

For $\mathcal{J} = \downarrow(\mathcal{I})$, an interpretation \mathcal{J}' is called an n -secure transformation of \mathcal{J} if it is obtained by (possibly infinitely) repeating the following step:

Choose one unvisited w.r.t. tree-depth minimal node (ρ, w) that is n -replaceable w.r.t. \mathcal{I} . Replace (ρ, w) with one of its n -secure replacements from $\lim \mathcal{I}$ and mark (ρ, w) as visited.

Figure 4 a) displays a 2-secure replacement in the considered unraveling of our example model shown in Figure 4 b). Figure 4 c) displays the result of carrying out this replacement step on our example. The following lemma ensures that not too many elements (actually defined in terms of the original model) are exempt from being replaced.

Lemma 15. *Every purified model \mathcal{I} of \mathcal{K} contains only finitely many distinct elements that start a BCP and are the cause for n -irreplaceable nodes in the unraveling of \mathcal{I} .*

Proof. Assume the converse: let a purified model \mathcal{I} of \mathcal{K} contain an infinite set D of elements giving rise to n -irreplaceable nodes in $\downarrow(\mathcal{I})$. Then there must be an \mathcal{L}' such that there is an infinite set $D' \subseteq D$ such that every $d' \in D'$ generates an \mathcal{L} for which $\text{cut}_n(\mathcal{L}) \cong \mathcal{L}'$ (since by Lemma 8, there are only finitely many non-isomorphic choices for \mathcal{L}'). This set D' can be used to guide the construction of a specific limit element $\mathcal{J} \in \lim \mathcal{I}$ according to Lemma 13.1. Now, for an element (ρ, w) from \mathcal{J} starting a BCP, let $l_{(\rho, w)} \in \mathbb{N}$ be the length of the shortest such BCP starting from (ρ, w) . Then, let k be the maximum over all $l_{(\rho, w)}$ of individuals (ρ, w) from \mathcal{J} that start a BCP and for which $|w| \leq n$. By construction, D' contains one element d'' generating an \mathcal{L} with $\text{cut}_k(\mathcal{L}) \cong \text{cut}_k(\mathcal{J})$ (actually infinitely many). By the choice of k , we can conclude that \mathcal{J} is an n -secure replacement for the irreplaceable $\downarrow(\mathcal{I})$ -node caused by d'' which gives us a contradiction. \square

Next, we can show that the process of unraveling, n -secure transformation and collapsing preserves the property of being a model of a knowledge base and (with the right choice of n) also preserves the property of *not* entailing a conjunctive query. Moreover, this model conversion process ensures that the resulting model contains only finitely many new nominals (witnessed by a bound on the length of BCPs). Figure 4 d) illustrates these properties for our example model. Note that only one new nominal is left whereas collapsing the original unraveling yields infinitely many.

Lemma 16. *Let \mathcal{I} be a purified model of \mathcal{K} , $\mathcal{J} = \downarrow(\mathcal{I})$, and \mathcal{J}' an n -secure transformation of \mathcal{J} . Then the following hold:*

1. $\text{collapse}(\mathcal{J}')$ is a model of \mathcal{K} .
2. There is a natural number m such that \mathcal{J}' does not contain any node whose shortest descending BCP has a length greater than m .
3. If, for some CQ q , we have $\mathcal{J} \not\models q$ and $n > \sharp(q)$, then $\mathcal{J}' \not\models q$.
4. If, for some CQ q , we have $\mathcal{I} \not\models q$ and $n > \sharp(q)$, then $\text{collapse}(\mathcal{J}') \not\models q$.

Now we are able to establish our first milestone on the way to showing finite representability of countermodels.

Theorem 17. *For every $\mathcal{ALCOIFb}$ KB \mathcal{K} and CQ q s.t. $\mathcal{K} \not\models q$, there is a forest model \mathcal{I} of \mathcal{K} with finitely many roots and bounded branching degree s.t. $\mathcal{I} \not\models q$.*

Finite Representations of Models

We can now use standard techniques from tableau algorithms (adapted to work on models) to construct finite representations for a forest model of \mathcal{K} with finitely many roots. In particular the tableau algorithm with n -tree-blocking, $n \geq \sharp(q)$, for deciding CQ entailment in \mathcal{SHIQ} , \mathcal{SHOQ} , and \mathcal{SHOI} with only simple roles in the query (Ortiz 2008; Ortiz, Calvanese, and Eiter 2008) works exactly like that. We call an interpretation, on which we applied n -tree-blocking and discarded all blocked elements, an n -representation. Such an n -representation corresponds to a complete and clash-free completion graph in tableau algorithms.

Definition 18. *Let $n \in \mathbb{N}$ be a fixed natural number and \mathcal{I} with $(\delta, w) \in \Delta^{\mathcal{I}}, w \neq \varepsilon$ a forest interpretation for \mathcal{K} . An n -blocking-tree w.r.t. (δ, w) , denoted $\text{block}_{\mathcal{I}}^n(\delta, w)$, is the interpretation obtained from \mathcal{I} by restricting \mathcal{I} to elements in $\{(\delta, ww') \mid |w'| \leq n\} \cup \{(\rho, \varepsilon) \mid (\rho, \varepsilon) \in \Delta^{\mathcal{I}}\}$. An n -blocking-tree $\text{block}_{\mathcal{I}}^n(\delta, w)$ n -blocks an n -blocking-tree $\text{block}_{\mathcal{I}}^n(\delta, ww')$ if*

1. $\text{block}_{\mathcal{I}}^n(\delta, w)$ and $\text{block}_{\mathcal{I}}^n(\delta, ww')$ have disjoint domains except for root elements,
2. there is a bijection φ from elements in $\text{block}_{\mathcal{I}}^n(\delta, w)$ to elements in $\text{block}_{\mathcal{I}}^n(\delta, ww')$ that witnesses $\text{block}_{\mathcal{I}}^n(\delta, w) \cong \text{block}_{\mathcal{I}}^n(\delta, ww')$, and
3. for each descendant (δ, wv) of (δ, w) , there is no inverse functional role f and root $(\rho, \varepsilon) \in \Delta^{\mathcal{I}}$ such that $\langle (\delta, wv), (\rho, \varepsilon) \rangle \in f^{\mathcal{I}}$.

A node $(\delta, v) \in \Delta^{\mathcal{I}}$ is n -blocked, if (δ, v) is either directly or indirectly n -blocked; (δ, v) is indirectly n -blocked, if one of its ancestors is n -blocked; (δ, v) is directly n -blocked if none of its ancestors is n -blocked and (δ, v) is a leaf of some n -blocking-tree $\text{block}_{\mathcal{I}}^n(\delta, ww')$ in \mathcal{I} that is n -blocked. W.l.o.g., we assume that the n -blocking-trees are minimal w.r.t. the order over $\Delta^{\mathcal{I}}$ (cf. Def. 4).

A forest interpretation $\mathcal{I} = (\Delta^{\mathcal{I}}, \mathcal{I})$ for \mathcal{K} is an n -representation of \mathcal{K} if (i) $\Delta^{\mathcal{I}}$ is finite, (ii) $\Delta^{\mathcal{I}}$ contains no indirectly n -blocked nodes, (iii) each $\delta \in \Delta^{\mathcal{I}}$ that is not directly n -blocked is locally \mathcal{K} -consistent, and (iv) for each $o \in \text{nom}(\mathcal{K})$, there is one $\delta \in \Delta^{\mathcal{I}}$ s.t. $\delta = (\rho, \varepsilon)$ and $o^{\mathcal{I}} = \delta^{\mathcal{I}}$.

Since we fixed a bound on the number of roots in Theorem 17 and otherwise only consider the closure $\text{cl}(\mathcal{K})$ of \mathcal{K} , one can show that there are only finitely many non-isomorphic n -blocking-trees even though we take links back to roots into account. It is worth noting that, for DLs as expressive as $\mathcal{ALCOIFb}$, n has to be greater than 0 (at least trees of depth 1) if we want to transform n -representations into models of the knowledge base. We now show that each knowledge base has an n -representation for some fixed $n \in \mathbb{N}$ and, afterwards, that we can use an n -representation to build a model for the knowledge base.

Lemma 19. *Let $n \geq \sharp(q)$. If $\mathcal{K} \not\models q$, then there is an n -representation \mathcal{R} of \mathcal{K} s.t. $\mathcal{R} \not\models q$, and from \mathcal{R} one can build a model \mathcal{I} of \mathcal{K} such that $\mathcal{I} \not\models q$.*

Sketch. By Theorem 17, there is a forest model \mathcal{I} of \mathcal{K} with finitely many roots and branching degree bounded in $|\text{cl}(\mathcal{K})|$ s.t. $\mathcal{I} \not\models q$. Using an argumentation similar to the one in the proof of Lemma 8, one can show that there are only finitely many non-isomorphic n -blocking trees. Together with the fact that \mathcal{I} is obtained from a collapsing and relations from elements within a tree to a root in collapsings are never for inverse functional roles, this shows that there is an n -representation of \mathcal{I} because for each tree rooted in a node $(\delta, \varepsilon) \in \Delta^{\mathcal{I}}$ trees can only grow to a limited depth before two nodes (δ, w) and (δ, ww') occur s.t. $\text{block}_{\mathcal{I}}^n(\delta, w)$ n -blocks $\text{block}_{\mathcal{I}}^n(\delta, ww')$. We can then simply discard indirectly n -blocked nodes from \mathcal{I} to obtain the desired n -representation.

Since $\mathcal{I} \not\models q$ and the n -representation is a restriction of \mathcal{I} , non-entailment of q is clearly preserved. \square

Please note that we would not obtain such a bound if we had not fixed a bound on the number of new nominals (roots) beforehand and that we cannot use the standard tableau algorithms to obtain this result.

In order to obtain a model for a knowledge base \mathcal{K} from some n -representation, the techniques for tableau algorithms with tree blocking (Ortiz 2008; Ortiz, Calvanese, and Eiter 2008) for building a tableau from a complete and clash-free completion graph can straightforwardly be adapted to our case.

Lemma 20. *Let \mathcal{K} be a consistent $\mathcal{ALCOIFb}$ knowledge base, q a conjunctive query, and $n \geq \sharp(q)$. If \mathcal{R} is an n -representation of \mathcal{K} such that $\mathcal{R} \not\models q$, then there is a model \mathcal{I} of \mathcal{K} s.t. $\mathcal{I} \not\models q$.*

Now Lemma 19 guarantees that, in case $\mathcal{K} \not\models q$, there is always a finite n -representation \mathcal{R} for \mathcal{K} such that $\mathcal{R} \not\models q$ and Lemma 20 guarantees that \mathcal{R} can be transformed into a model \mathcal{I} of \mathcal{K} such that $\mathcal{I} \not\models q$. This suffices to show that we can enumerate all (finite) n -representations for \mathcal{K} and check whether they entail the query. Since role hierarchies and qualified number restrictions can be encoded in $\mathcal{ALCOIFb}$, we get, together with the semi-decidability result for FOL (Gödel 1929), the desired theorem.

Theorem 21. *It is decidable whether $\mathcal{K} \models q$ for \mathcal{K} an $\mathcal{ALCHOIQb}$ knowledge base and q a Boolean conjunctive query.*

Conclusions

This solves the long-standing open problem of deciding conjunctive query entailment in the presence of nominals, inverse roles, and qualified number restrictions. Our result generalizes to unions of conjunctive queries and to *SHOIQ* provided the query contains only simple roles, and we are confident that the technique also extends to *SR_QIQ* under the same restriction.

Entailment of unions of conjunctive queries is also closely related to the problem of adding rules to a DL knowledge base, e.g., in the form of Datalog rules. Augmenting a DL KB with an arbitrary Datalog program easily leads to undecidability (Levy and Rousset 1998). To ensure decidability, the interaction between the Datalog rules and the DL knowledge base can be restricted by imposing a safety condition. The *DL+log* framework (Rosati 2006a) provides the least restrictive integration proposed so far and Rosati presents an algorithm that decides the consistency of a *DL+log* knowledge base by reducing the problem to entailment of unions of conjunctive queries. Notably, his results (Rosati 2006a, Thm. 11) imply that the consistency of an *ALCH_QIQ_b* knowledge base extended with (weakly-safe) Datalog rules is decidable if and only if entailment of unions of conjunctive queries in *ALCH_QIQ_b* is decidable, which we have established.

Corollary 22. *The consistency of ALCH_QIQ_b+log-knowledge bases (both under FOL semantics and under non-monotonic semantics) is decidable.*

Another related reasoning problem is *query containment*. Given a schema (or TBox) \mathcal{S} and two queries q and q' , we have that q is contained in q' w.r.t. \mathcal{S} iff every interpretation \mathcal{I} that satisfies \mathcal{S} and q also satisfies q' . It is well known that query containment w.r.t. a TBox can be reduced to deciding entailment for unions of conjunctive queries w.r.t. a knowledge base (Calvanese, De Giacomo, and Lenzerini 1998). Decidability of unions of conjunctive query entailment in *ALCH_QIQ_b* implies, therefore, also decidability of query containment w.r.t. to an *ALCH_QIQ_b* TBox.

There are two obvious avenues for future work. Firstly, we will embark on extending our results in order to allow non-simple roles as query predicates. This is a non-trivial task as our current approach heavily relies on a certain locality of query matches, which has to be relinquished when considering non-simple roles. Secondly, since the approach is purely a decision procedure, the computational complexity of the problem remains open, and we are eager to determine the associated computational complexities and provide techniques that can form the basis for implementable algorithms.

Acknowledgements. Sebastian Rudolph was supported by a scholarship of the German Academic Exchange Service (DAAD) and the DFG projects ReaSem and ExpresST. Birte Glimm is funded by the EPSRC project Hermit: Reasoning with Large Ontologies.

References

Baader, F.; Calvanese, D.; McGuinness, D. L.; Nardi, D.; and Patel-Schneider, P. F. 2003. *The Description Logic*

Handbook. Cambridge University Press.

Baader, F.; Brandt, S.; and Lutz, C. 2005. Pushing the \mathcal{EL} envelope. In *Proc. of the 19th Int. Joint Conf. on Artificial Intelligence (IJCAI-05)*. Morgan Kaufmann.

Baader, F. 2003. Terminological cycles in a description logic with existential restrictions. In *Proc. of the 18th Int. Joint Conf. on Artificial Intelligence (IJCAI-03)*, 325–330.

Calvanese, D.; De Giacomo, G.; Lembo, D.; Lenzerini, M.; and Rosati, R. 2007. Tractable reasoning and efficient query answering in description logics: The DL-Lite family. *J. of Automated Reasoning* 39(3):385–429.

Calvanese, D.; De Giacomo, G.; and Lenzerini, M. 1998. On the decidability of query containment under constraints. In *Proc. of the 17th ACM SIGACT SIGMOD Sym. on Principles of Database Systems (PODS-98)*, 149–158. ACM Press and Addison Wesley.

Calvanese, D.; Eiter, T.; and Ortiz, M. 2007. Answering regular path queries in expressive description logics: An automata-theoretic approach. In *Proc. of the Nat. Conf. on Artificial Intelligence (AAAI-07)*.

Eiter, T.; Lutz, C.; Ortiz, M.; and Simkus, M. 2009. Query answering in description logics with transitive roles. In *Proc. of IJCAI-09*, 759–764. AAAI Press/The MIT Press.

Glimm, B., and Rudolph, S. 2009. Nominals, inverses, counting, and conjunctive queries or Why infinity is your friend! Technical report, University of Oxford. <http://www.comlab.ox.ac.uk/files/2175/paper.pdf>.

Glimm, B.; Horrocks, I.; Lutz, C.; and Sattler, U. 2008. Conjunctive query answering for the description logic *SHIQ*. *J. of Artificial Intelligence Research* 31:151–198.

Glimm, B.; Horrocks, I.; and Sattler, U. 2008. Unions of conjunctive queries in *SHOQ*. In *Proc. of KR-08*. AAAI Press/The MIT Press.

Gödel, K. 1929. *Über die Vollständigkeit des Logikkalküls*. Ph.D. Dissertation, Universität Wien.

Grahne, G. 1991. *Problem of Incomplete Information in Relational Databases*. LNCS. Springer.

Horrocks, I., and Sattler, U. 2005. A tableaux decision procedure for *SHOIQ*. In *Proc. of the 19th Int. Joint Conf. on Artificial Intelligence (IJCAI-05)*.

Kazakov, Y., and Motik, B. 2008. A resolution-based decision procedure for *SHOIQ*. *J. of Automated Reasoning* 40(2–3):89–116.

Krötzsch, M.; Rudolph, S.; and Hitzler, P. 2007. Conjunctive queries for a tractable fragment of OWL 1.1. In *Proc. of the Int. Semantic Web Conf. (ISWC-07)*, volume 4825 of LNCS, 310–323. Springer.

Levy, A. Y., and Rousset, M.-C. 1998. Combining horn rules and description logics in CARIN. *Artificial Intelligence* 104(1–2):165–209.

Lutz, C. 2008. The complexity of conjunctive query answering in expressive description logics. In *Proc. of the 4th Int. Joint Conf. on Automated Reasoning (IJCAR-08)*, 179–193. LNCS.

- Motik, B.; Shearer, R.; and Horrocks, I. 2009. Hypertableau reasoning for description logics. Submitted to a journal. <http://www.hermit-reasoner.com/publications/msh08hypertableau-journal.pdf>.
- Ortiz, M.; Calvanese, D.; and Eiter, T. 2008. Data complexity of query answering in expressive description logics via tableaux. *J. of Automated Reasoning* 41(1):61–98.
- Ortiz, M.; Simkus, M.; and Eiter, T. 2008. Conjunctive query answering in sh using knots. In *Proc. of the 2008 Description Logic Workshop (DL-08)*. CEUR.
- Ortiz, M. 2008. Extending CARIN to the description logics of the \mathcal{SH} family. In *Proc. of JELIA-08*, 324–337. LNAI.
- Pratt-Hartmann, I. 2009. Data-complexity of the two-variable fragment with counting quantifiers. Forthcoming in *Information and Computation*. <http://arxiv.org/abs/0806.1636>.
- Rosati, R. 2006a. DL+log: Tight integration of description logics and disjunctive datalog. In *Proc. of the 10th Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR-06)*, 68–78.
- Rosati, R. 2006b. On the decidability and finite controllability of query processing in databases with incomplete information. In *Proc. of the 25th ACM SIGACT SIGMOD Symposium on Principles of Database Systems (PODS 2006)*, 356–365. ACM Press and Addison Wesley.
- Rosati, R. 2007a. The limits of querying ontologies. In *Proc. of the Int. Conf. on Database Theory (ICDT-07)*, volume 4353 of *LNCS*, 164–178. Springer.
- Rosati, R. 2007b. On conjunctive query answering in \mathcal{EL} . In *Proc. of the 2007 Description Logic Workshop (DL-07)*. CEUR.
- Rudolph, S.; Krötzsch, M.; and Hitzler, P. 2008. Terminological reasoning in \mathcal{SHIQ} with ordered binary decision diagrams. In *Proc. of AAAI-08*, 529–534. AAAI Press/The MIT Press.
- van der Meyden, R. 1998. Logical approaches to incomplete information: A survey. In *Logics for Databases and Information Systems*, 307–356. Kluwer Academic Publisher.