Distributive lattice-structured ontologies

Hans Bruun[1], Dion Coumans[2], and Mai Gehrke[2] *

[1] Technical University of Denmark, Denmark
[2] Radboud University Nijmegen, The Netherlands

Abstract. In this paper we describe a language and method for deriving ontologies and ordering databases. The ontological structures arrived at are distributive lattices with attribution operations that preserve $\vee$, $\wedge$ and $\perp$. The preservation of $\wedge$ allows the attributes to model the natural join operation in databases. We start by introducing ontological frameworks and knowledge bases and define the notion of a solution of a knowledge base. The import of this definition is that it specifies under what condition all information relevant to the domain of interest is present and it allows us to prove that a knowledge base always has a smallest, or terminal, solution. Though universal or initial solutions almost always are infinite in this setting with attributes, the terminal solution is finite in many cases.
We describe a method for computing terminal solutions and give some conditions for termination and non-termination. The approach is predominantly coalgebraic, using Priestley duality, and calculations are made in the terminal coalgebra for the category of bounded distributive lattices with attribution operations.

1 Introduction

In traditional relation database models used in commercial applications, information is stored in tuples. Research in knowledge representation however seeks models of information that are able to accommodate more complex information forms. A key issue in this respect is classification structures in the form of formal ontologies. While lattices are just what is called for to model the subsumption relation of concepts captured by ontologies, they fall short of providing tuples supported in the relational database model. In [9] Fischer Nilsson showed that lattices with additional (unary) operations preserving not only $\perp$ and $\vee$, but also $\wedge$ provide a common framework generalising ontologies and traditional databases. The tuples in relational databases may be represented as meets of basic attribution terms $a_i(c_i)$ and the database relations are then represented as lattice joins of such meets of basic attribution terms. This representation together with preservation of $\wedge$ and $\perp$ by the attribution operations provide database natural join simply as lattice meet. What makes this framework interesting from the application perspective is thus that it combines and enriches relational databases with lattice classifications. A preliminary algorithmization was implemented for the OntoQuery project [11] which was concerned with content-based querying

* The authors would like to thank the referees for their very useful input.

of large text databases. A description of an updated version of the algorithm is available in [4]. Here we describe the setting and the mathematical ideas behind the algorithm and analyse it in mathematical terms.

In algebraic terms, we want to solve a given generators and relations problem (modelling the dependencies between the basic concepts of the expert domain in question) that imposes the appropriate order and algebra structure on a given set of terms (which correspond to the current entries of the database/ontology). This approach to generative knowledge representation generalises the work of Oles which was used in IBM's medical knowledge representation tool [10]. Oles' work differs from ours in two aspects. It doesn't allow additional operations in the type and the only solution that is identified is the universal solution. Even in the plain distributive lattice setting this solution is typically much too big and involves many irrelevant concepts.

In this paper we introduce what we believe to be a key concept, namely that of classification of a term w.r.t. a solution of an ontological framework, and show that, given a set of terms, a generators and relations problem has a minimal solution that classifies all the given terms. We call this the terminal solution and it is the smallest meaningful solution - also in Oles' setting. In the setting with attributes, the universal solution of a generators and relations problem is typically infinite, and thus the generative approach to knowledge representation is impossible without the key notion of a terminal solution. While the terminal solution is typically finite, this is not always the case. We give sufficient conditions for it to be finite and infinite, respectively. We conjecture that our condition for non-termination is essentially sharp, but this remains an open problem.

The algebraic logics treated here are related to description logic [2] in the sense that both formalisms concern generator and relations problems for distributive lattices with additional operations, see e.g. [12, 8, 1]. The operations in description logic are more general as they may just preserve $\wedge$ or $\vee$ while ours preserve both. However, the main difference lies in the problems treated: the main issue in description logic is, given two compound concept descriptions (i.e., terms in the free distributive lattice with attribution operations over the set of basic concepts), to determine whether the equivalence class of one of them subsumes the equivalence class of the other in the universal solution of the generators and relations problem. This is a very local piece of information that concerns only the two terms in question. By contrast, we give the decomposition of a term into a join of join-irreducibles relative to the universal solution (that is, when such decomposition exists). This yields global information about that term, and when we do this for each term of interest, we obtain the full terminal solution. This is a global solution. In addition, finding the normal form for a term identifies the irreducible building blocks of the ontology which may not be readily identifiable from the generators and relation problem. This is the sense in which these ontologies are generative. Furthermore, our methods should readily generalise to the setting of operations preserving only $\wedge$ or $\vee$ using the finitely generated final coalgebras for modal logic as developed in [3].

In Section 2 we introduce the key concepts of this work, including ontological frameworks, knowledge bases, classification of a term and terminal solutions. In Section 3 we recall a few facts from duality theory and canonical extensions, and in Section 4 we apply these facts to obtain the existence as well as a useful description of terminal solutions. In Section 5 we describe the final coalgebra for the variety we work in and in the final section we use this to give a method for computing terminal solutions. In Section 6 we also supply some results on termination and non-termination.

2 Ontological frameworks and knowledge bases

An *ontological framework*, $\mathcal{O} = (C, A, \Pi)$, consists of three finite sets. The first, C, is a set of basic concept names. We think of these as generating elements. The second, A, is a set of attribution operation, or attribute, symbols. We assume these operations to be unary. We write $T_{L_\bot A}(C)$ for the collection of terms built up from concepts from C using attributes in A and the lower bounded lattice operation symbols in $L_\bot = \{\vee, \wedge, \bot\}$. The third set, Π, is a set of so-called terminological axioms. The elements of Π are pairs (r, s) of terms from $T_{L_\bot A}(C)$. The idea is that the terminological axioms specify pairs of expressions that for ontological reasons must be identified.

As a simple example one may consider an ontological framework on real estate. The set C then consists of concepts relevant to real estate, like 'flat', 'villa', names of geographical regions, different sizes, etc. We may express the fact that the geographical region r_2 is a subregion of r_1 by the terminological axiom $(r_1, r_1 \vee r_2)$. The set A may contain for instance an attribute L for 'located in', which maps a region like r_1 to the concept $L(r_1)$ which (extensionally) should designate all real estates located in region r_1.

There are various reasons for which we only allow the use of $\bot$ in the specification of ontological frameworks. First of all, it is quite common in ontology to include $\bot$ as the inconsistent concept. However, a universal concept is less meaningful, especially in the database setting. Secondly, for our attribution operations, $\top$ and $\bot$ play very different roles. We will elaborate on this further on.

The fundamental idea is that an ontological framework specifies a class of ontological structures that are the solutions for the framework. In this work we restrict ourselves to considering solutions that have a lower bounded distributive lattice structure with attribution operations, one for each $a \in A$. We require the attribution operations to preserve $\vee$, $\wedge$, and $\bot$. Requiring the preservation of $\wedge$ allows the attributes to model the so called natural join operation in databases. In the mathematical work below we add $\top$ to the type as well as the laws $a(\top) = \top$ for $a \in A$ and call the corresponding algebras DLAs. This facilitates our construction and does not interact with our solutions: removing the top from a solution will always yield a subalgebra of type $L_\bot A$. Thus the named element $\top$ is not a part of the ontologically meaningful solution and is just added to make the mathematics smoother. It is important that in ontologies (such as ontologically ordered databases) the top need not be preserved by the attribution operations. The top we add here is just a mathematical gadget and it will always remain an unreachable join as the terminological axioms never include top. It does not influence the action of the attributes on the elements of the actual ontology and in particular, in case it exists, the action of the attributes on the top of the ontology.

Every DLA generated by C is a homomorphic image of $F_{DLA}(C)$, the free bounded distributive lattice with attribution operations generated by C. Thus if we want solutions of the ontological framework $\mathcal{O} = (C, A, \Pi)$ to be DLAs that are generated by C we only have to look among homomorphic images of $F_{DLA}(C)$. As said, the terminological axioms, that is, the elements of Π, are pairs of $DLA_{\perp}$ terms in the basic concept names. Each such axiom may be seen as an identification of particular elements in $F_{DLA}(C)$. This leads to the following definition:

Definition 1. *Let $\mathcal{O} = (C, A, \Pi)$ be an ontological framework. A* solution of $\mathcal{O}$ *is a homomorphic image of $F_{DLA}(C)$ in which r is identified with s for each pair (r, s) in Π.*

Since homomorphic images correspond to quotients of the domain by congruences, solutions are those quotients that are given by congruences containing Π. One of the most fundamental facts from general algebra about congruences is that there is a least congruence containing all the pairs in Π. This is the join of the principal congruences $\theta_{(r,s)}$ for $(r, s) \in \Pi$. The corresponding quotient of $F_{DLA}(C)$ is the least-collapsed algebra satisfying the terminological axioms and all other quotients satisfying all the terminological axioms are quotients of it. This yields the following theorem.

Theorem 1. *Let $\mathcal{O} = (C, A, \Pi)$ be an ontological framework. There exists a quotient $h_{\mathcal{O}} : F_{DLA}(C) \to F_{\mathcal{O}}$ of $F_{DLA}(C)$ satisfying the following conditions:*

1. *$h_{\mathcal{O}} : F_{DLA}(C) \to F_{\mathcal{O}}$ is a solution of $\mathcal{O} = (C, A, \Pi)$;*
2. *If $h : F_{DLA}(C) \to D$ is any solution of $\mathcal{O} = (C, A, \Pi)$, then there is a unique homomorphism $h_D : F_{\mathcal{O}} \to D$ so that the diagram*

$$\begin{array}{ccc} F_{DLA}(C) & \longrightarrow & D \\ & \searrow \quad \nearrow & \\ & F_{\mathcal{O}} & \end{array}$$

commutes.

The solutions of $\mathcal{O}$ are exactly the quotients of $F_{\mathcal{O}}$.

Definition 2. *Let $\mathcal{O} = (C, A, \Pi)$ be an ontological framework. We call the quotient $h_{\mathcal{O}} : F_{DLA}(C) \to F_{\mathcal{O}}$ the* universal solution of $\mathcal{O}$.

This universal solution may also be described by:

$$F_{\mathcal{O}} = F_{DLA}(C)/\theta_{\Pi}$$

$$\text{with } \theta_{\Pi} = \bigvee\{\theta_{(r,s)} : (r, s) \in \Pi\} = \bigcap\{\theta : \Pi \subseteq \theta\}.$$

The universal solution is typically much too large in the sense that there will be many points that are not relevant to the underlying domain or database for which the ontology is made. Accordingly, we give a definition of knowledge base that also specifies a set of domain terms.

Definition 3. *A* Knowledge Base *(KB), $\mathcal{B} = (C, A, \Pi, I)$, is an ontological framework $\mathcal{O} = (C, A, \Pi)$ together with a specified finite set I of DLA terms. We call the elements of I* inserted *or* inhabited terms *and $\mathcal{O}$ the* associated ontological framework*.*

The idea of this definition is that the inserted terms are the ones that actually correspond to concepts of interest, or, in a more database oriented view, to terms for which data is available and has to be classified by the ontology. In the setting of our real estate example, the inserted terms could be the real estates some agent has for sale. The significance of this definition lies in the companion definition of a solution of a KB. First we define the notion of a classification of a term w.r.t. a solution of an ontological framework.

Definition 4. *Let $\mathcal{O} = (C, A, \Pi)$ be an ontological framework, $h : F_{DLA}(C) \to D$ a solution of $\mathcal{O}$, and $t \in F_{DLA}(C)$. Then $t_1, \ldots, t_n \in F_{DLA}(C)$ is a* classification of t with respect to *$h : F_{DLA}(C) \to D$ provided for each i we have $t_i \leq t$ and*

$$h(t) = h(t_1) \vee \ldots \vee h(t_n).$$

This formal definition generalises the idea of classifications as one has in taxonomies: e.g., the animal kingdom is divided into a disjunction of subclasses (mammals, ...).

Notice that if $t_1, \ldots, t_n \in F_{DLA}(C)$ is a classification of t, then $t_1 \vee \ldots \vee t_n \leq t$ in $F_{DLA}(C)$. If we actually have equality in $F_{DLA}(C)$, then it is of course a classification w.r.t. any ontological framework, but these classifications are trivial. The interesting ones are the ones forced by Π. Notice also that any classification of a term t w.r.t. the universal solution is also a classification of t w.r.t. any other solution of $\mathcal{O}$. However, the implication does not hold in the other direction. For the trivial quotient, $F_{DLA}(C) \to \mathbf{1}$, the term $\bot$ by itself is a classification of any term t. This of course tells us that $F_{DLA}(C) \to \mathbf{1}$ is not a very useful solution. A solution of a knowledge base is a solution of the associated ontological framework that is faithful to the universal solution in terms of classifying inserted terms. We are now ready to give the formal definition.

Definition 5. *Let $\mathcal{B} = (C, A, \Pi, I)$ be a KB. A* solution of $\mathcal{B}$ *is a solution $h : F_{DLA}(C) \to D$ of the associated ontological framework $\mathcal{O} = (C, A, \Pi)$ with the additional property that, for each inserted term $t \in I$, every classification $t_1, \ldots, t_n$ of t w.r.t. $h : F_{DLA}(C) \to D$ is also a classification of t w.r.t. $h_{\mathcal{O}} : F_{DLA}(C) \to F_{\mathcal{O}}$.*

Note that, for a knowledge base $\mathcal{B}$, a solution $h : F_{DLA}(C) \to D$ of the associated ontological framework is a solution of $\mathcal{B}$ if and only if, for each inserted term t, we have that $h(t) \leq h(s)$ implies $h_{\mathcal{O}}(t) \leq h_{\mathcal{O}}(s)$ for all terms s.

Clearly, the universal solution of the ontological framework associated to a knowledge base is a solution of the knowledge base. The import of the definition is that it allows us to work with smaller solutions than the universal one which still have all information relevant to the domain that has to be searched or classified.

Consider the concept human. Let $C = \{h, a, c, m, f\}$, where we think of these five concepts as human, adult, child, male and female. The free distributive lattice generated by these five unrelated concepts has more than 7000 elements. Introducing a terminological axiom to identify human with the disjuction of adult and child $(h, a \vee c)$ and one to identify human with the disjunction of male and female $(h, m \vee f)$, gives us an (attribute free) ontological framework whose universal solution $F_{\mathcal{O}}$ has 49 elements and is depicted in Figure 1(a). The boldface $\mathbf{h}, \mathbf{a}, \mathbf{c}, \mathbf{m}$ and $\mathbf{f}$ indicate the map $C \to F_{\mathcal{O}}$. Inserting just the concept h we obtain a knowledge base for which the least solution (which is what we will define to be the terminal solution) is the 16 element Boolean lattice depicted in Figure 1(b). Four new concepts: man $(m \wedge a)$, boy $(m \wedge c)$, girl $(f \wedge c)$, and woman $(f \wedge a)$ have been identified as the four concepts pertinent to the classification of human in this ontological framework and the terminal solution is the lattice generated by these four as irreducibles.

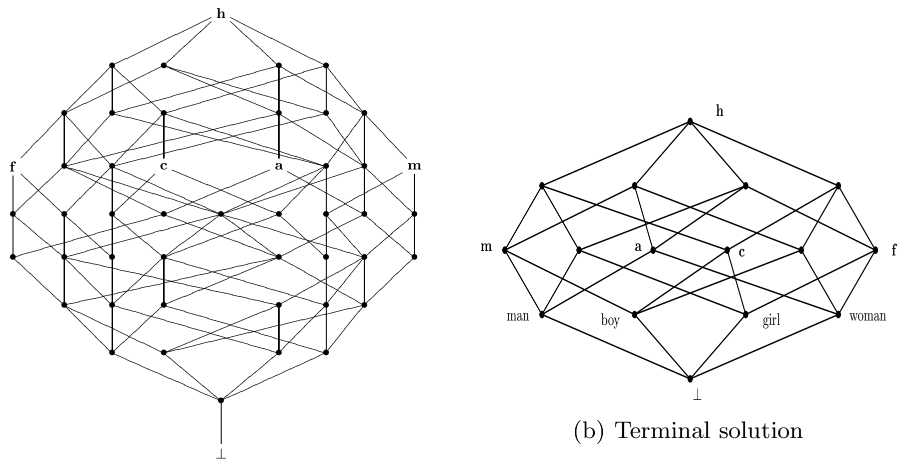

(a) Universal solution

(b) Terminal solution

Fig. 1. The human example

We will now use duality to obtain a more transparent description of the terminal solution. Most importantly, we will see that a knowledge base always has a terminal solution and, for a broad class of knowledge bases, this terminal solution is finite. Proving these facts as well as giving a description of an algorithm for finding the terminal solution of a knowledge base is best done coalgebraically or using duality theory.

3 Duality for distributive lattices

We will solve the problem of finding a solution of a knowledge base dually using Priestley duality for distributive lattices extended to incorporate the representation of attribution operations. The (Priestley) dual of a distributive lattice D is an ordered topological frame $(P_D, \tau, \leq)$. The set P_D may be taken to be the set of prime filters of D equipped with the reversed inclusion order. We define for each $d \in D$ a set $\widehat{d} = \{\wp \in P_D : d \in \wp\}$. The topology τ on P_D is the topology generated by $\{\widehat{d}, (\widehat{d})^c : d \in D\}$, where $(\widehat{d})^c$ denotes the complement of $\widehat{d}$ in P_D. The dual lattice of this topological frame is the lattice of its clopen downsets, which is isomorphic to the original lattice D.

We also want to view the additional operations dually. The notion of canonical extension is particularly well-suited for this. The canonical extension encodes the topological dual frame in algebraic terms. A thorough explanation of canonical extensions may be found in [7]. One may obtain an incarnation of the canonical extension D^σ of a bounded distributive lattice D as the lattice of all downsets of the frame $(P_D, \tau, \leq)$ which is denoted by $\mathcal{D}(P_D)$. The original lattice, D, embeds in D^σ via:

$$\begin{aligned} D &\to D^\sigma \\ d &\mapsto \widehat{d} = \{\wp \in P_D : d \in \wp\}. \end{aligned}$$

On the other hand, the points of the dual space P_D may be obtained as the *completely join irreducible* elements of the canonical extension of the lattice. These are the points satisfying the infinitary version of join irreducibility: p is completely join irreducible provided $p = \bigvee_I u_i$ implies $p = u_i$ for some $i \in I$. We denote this set by $J^\infty(D^\sigma)$.

The power of canonical extensions lies in the fact that the assignment is categorical, that is, any bounded lattice homomorphism $h : D \to E$ extends to a complete homomorphism between the canonical extensions, $h^\sigma : D^\sigma \to E^\sigma$. The extension is effected by using the fact that the canonical extension of D is generated from the image of D by means of arbitrary meets and joins, see Theorem 3.25 in [7].

For a lattice homomorphism $h : D \to E$, the dual of h is given by the lower adjoint of h^σ restricted to the dual space, $P_E = J^\infty(E^\sigma)$, of E, i.e. by the function $(h^\sigma)^\flat \restriction J^\infty(E^\sigma) : P_E \to P_D$ which is continuous and order preserving. This function is well-defined as h^σ sends completely join-irreducible elements of E^σ to completely join irreducible elements of D^σ. Every attribution operation on D may be viewed as a lattice homomorphism $D \to D$. Hence, every attribute a on D dually is captured by a continuous, order-preserving function $f_a : P_D \to P_D$.

One may show that a DLA-quotient $h : D \to E$ corresponds dually to a topologically closed subset $P_E \subseteq P_D$, closed under the actions of the maps f_a, dual to the attributes a on D. The extension $h^\sigma : D^\sigma \to E^\sigma$ may in that case be seen as the map

$$\begin{aligned} h^\sigma : \mathcal{D}(P_D) &\to \mathcal{D}(P_E) \\ U &\mapsto U \cap P_E. \end{aligned}$$

It is clear that this map is a complete homomorphism and its lower adjoint is the map

$$(h^\sigma)^\flat : \mathcal{D}(P_E) \to \mathcal{D}(P_D)$$
$$V \mapsto \downarrow V$$

where $\downarrow V$ denotes the downset of V in P_D.

4 The terminal solution of a knowledge base

In the approach to ontological frameworks presented here, as in [10], our solutions are presented dually. However, the dual point of view is also particularly useful for understanding our newly introduced notion of a solution of a knowledge base.

To see this, recall that solutions of an ontological framework correspond to quotients of its universal solution. That is, every solution of $\mathcal{O}$ is of the form $h = g \circ h_{\mathcal{O}} : F_{DLA}(C) \to F_{\mathcal{O}} \to D$. The maps $h_{\mathcal{O}}$ and g are DLA morphisms and thus in particular bounded lattice morphisms. As a consequence their canonical extensions are complete lattice homomorphisms, which have lower adjoints. We obtain the following characterization of solutions of knowledge bases.

Theorem 2. *Let $\mathcal{B} = (C, A, \Pi, I)$ be a knowledge base, $\mathcal{O} = (C, A, \Pi)$ the associated ontological framework, and $h = g \circ h_{\mathcal{O}} : F_{DLA}(C) \to F_{\mathcal{O}} \to D$ a solution of $\mathcal{O}$. The following conditions are equivalent:*

1. *h is a solution of the knowledge base $\mathcal{B}$;*
2. *The retraction $F^\sigma_{\mathcal{O}} \xrightarrow{g^\sigma} D^\sigma \xrightarrow{(g^\sigma)^\flat} F^\sigma_{\mathcal{O}}$ fixes $h_{\mathcal{O}}(t)$ for each $t \in I$;*
3. *For each $t \in I$, $max(\widehat{h_{\mathcal{O}}(t)})$ is contained in P_D;*

where $\widehat{h_{\mathcal{O}}(t)}$ denotes the clopen downset of $P_{\mathcal{O}}$ (the frame dual to $F_{\mathcal{O}}$), corresponding to $h_{\mathcal{O}}(t)$, $max(\widehat{h_{\mathcal{O}}(t)})$ is the set of its maximal elements and P_D is the closed subspace of $P_{\mathcal{O}}$ dual to the quotient $g : F_{\mathcal{O}} \to D$.

Proof. As explained in the previous section, the composition in condition *2* may be viewed as

$$\mathcal{D}(P_{\mathcal{O}}) \xrightarrow{g^\sigma} \mathcal{D}(P_D) \xrightarrow{(g^\sigma)^\flat} \mathcal{D}(P_{\mathcal{O}})$$
$$S \mapsto S \cap P_D \mapsto \downarrow(S \cap P_D).$$

It is a basic fact of duality theory that every clopen S has enough maximal points in the sense that every element of S is below a maximal element of S. Thus it is clear that a clopen S is fixed by the above composition if and only if the maximal points of S are elements of P_D. This shows that conditions *2* and *3* of the theorem are equivalent. Their equivalence with *1* is an exercise in duality theory.

For any knowledge base $\mathcal{B} = (C, A, \Pi, I)$, the universal solution $h_{\mathcal{O}} : F_{DLA}(C) \to F_{\mathcal{O}}$ of the associated ontological framework $\mathcal{O} = (C, A, \Pi)$ is the greatest solution of $\mathcal{B}$ in the sense that any other solution factors through it. Every KB also has a least solution in the sense of the following theorem.

Theorem 3. *Let $\mathcal{B} = (C, A, \Pi, I)$ be a knowledge base. Then there exists a quotient $h_{\mathcal{B}} : F_{DLA}(C) \to D_{\mathcal{B}}$ of $F_{DLA}(C)$ satisfying the following conditions:*

1. *$h_{\mathcal{B}} : F_{DLA}(C) \to D_{\mathcal{B}}$ is a solution of $\mathcal{B} = (C, A, \Pi, I)$;*
2. *If $h : F_{DLA}(C) \to D$ is any solution of $\mathcal{B} = (C, A, \Pi, I)$, then there is a unique homomorphism $h_D : D \to D_{\mathcal{B}}$ so that the diagram*

$$\begin{array}{ccc} F_{DLA}(C) & \longrightarrow & D_{\mathcal{B}} \\ & \searrow \quad \nearrow & \\ & D & \end{array}$$

commutes.

Proof. By Theorem 2 and the general duality facts stated in Section 3, a solution of $\mathcal{B} = (C, A, \Pi, I)$ corresponds dually to a subspace Q of $P_{\mathcal{O}}$ that is topologically closed, is closed under the map f_a for each $a \in A$, and for which $\max(\widehat{h_{\mathcal{O}}(t)}) \subseteq Q$ for each $t \in I$. Since each of these three requirements on Q is preserved under arbitrary intersection, it follows that $\mathcal{B}$ has a least solution.

Note that this argument goes through unscathed in the setting of arbitrary modalities: We just need Q closed under relational image w.r.t. the dual Kripke relations R_a instead of under the f_as. In fact, using the canonical extension perspective, one can see that the proof goes through even for monotone operations (not necessarily join or meet preserving).

Definition 6. *Let $\mathcal{B} = (C, A, \Pi, I)$ be a knowledge base. We call the quotient $h_{\mathcal{B}} : F_{DLA}(C) \to D_{\mathcal{B}}$ the* terminal solution of $\mathcal{B}$.

The terminal solution may also be described dually in a generative way.

Theorem 4. *Let $\mathcal{B} = (C, A, \Pi, I)$ be a knowledge base. The terminal solution of $\mathcal{B}$ is described dually by*

$$P_{\mathcal{B}} = \overline{\left(\bigcup \{ f_w(max(\widehat{h_{\mathcal{O}}(t)})) \mid w \in A^*,\ t \in I \} \right)},$$

where, for every word $w = a_0 \ldots a_n$ in A^*, $f_w = f_{a_n} \circ \ldots \circ f_{a_0}$, the map dual to $a_0 \circ \ldots \circ a_n$; and for a subset $P \subseteq P_{\mathcal{O}}$, $\overline{P}$ denotes the topological closure of P.

5 Description of $F_{DLA}(C)$ and its dual frame

Since every solution is a quotient of the free DLA, we study the structure of $F_{DLA}(C)$ and its dual frame. These are particularly simple to understand on the basis of DLs. Let $A^*(C)$ denote the free algebra over C of type A. That is, $A^*(C) = \{w(c) : w \in A^* \text{ and } c \in C\}$ where we equate $\lambda(c)$ with c for each $c \in C$ (λ being the empty word) and we have operations $a^{A^*(C)} : A^*(C) \to A^*(C), w(c) \mapsto aw(c)$ for each $a \in A$. Generate the free bounded distributive lattice over $A^*(C)$, $F_{DL}(A^*(C))$, and define, for each $a \in A$, a unary operation $a^F : F_{DL}(A^*(C)) \to F_{DL}(A^*(C))$ to be the unique bounded lattice homomorphism given by $x \mapsto a^{A^*(C)}(x)$ for $x \in A^*(C)$. Then $(F_{DL}(A^*(C)), (a^F)_{a \in A})$ is the free DLA over C.

The dual space of $F_{DL}(A^*(C))$ is the ordered Cantor space $2^{A^*(C)}$ (with the order inherited from 2). In particular, the prime filters of $F_{DL}(A^*(C))$ are in one-to-one correspondence with the subsets of $A^*(C)$. The principal filters, i.e. the ones with a minimum element, are given by $\uparrow d$, where d is a join-irreducible element of $F_{DL}(A^*(C))$. These correspond to the finite subsets of, or finite conjunctions over, $A^*(C)$. The remaining subsets of $A^*(C)$ (an uncountable number of them) correspond to the non-principal filters or infinite conjunctions over $A^*(C)$. Hence, we may view the underlying set of the dual frame of $F_{DL}(A^*(C))$ as the collection of all conjunctions over $A^*(C)$. This is in fact precisely what it is in canonical extension terms. We will denote both the dual frame and its underlying set by $P(A^*(C))$. Dual to each attribution operation, we have a continuous order preserving function $f_a : P(A^*(C)) \to P(A^*(C))$. It is not hard to see that a conjunction x is sent to $a\backslash x$ which is the conjunction of all those $w(c)$ such that $aw(c)$ is one of the conjuncts in x.

Definition 7. *Let x be a conjunction over $A^*(C)$. The* A-depth of x, *notation $d_A(x)$, is n provided n is the smallest natural number such that $x \in P(A^{\leq n}(C))$ where $A^{\leq n}(C) = \bigcup_{k=0}^{n} A^k(C)$, and, for each set X, $P(X)$ denotes the set of all conjunctions over X, otherwise the A-depth of x is infinite. We will call conjunctions of finite depth* basic conjunctions *and call this part of the dual space the* finite part of the dual space.

Organising the basic conjunctions according to their A-depth, the poset $P(A^*(C))$ may be depicted as shown in figure 2. We omit the meet operation, writing for example $ca(c)$ for $c \wedge a(c)$.

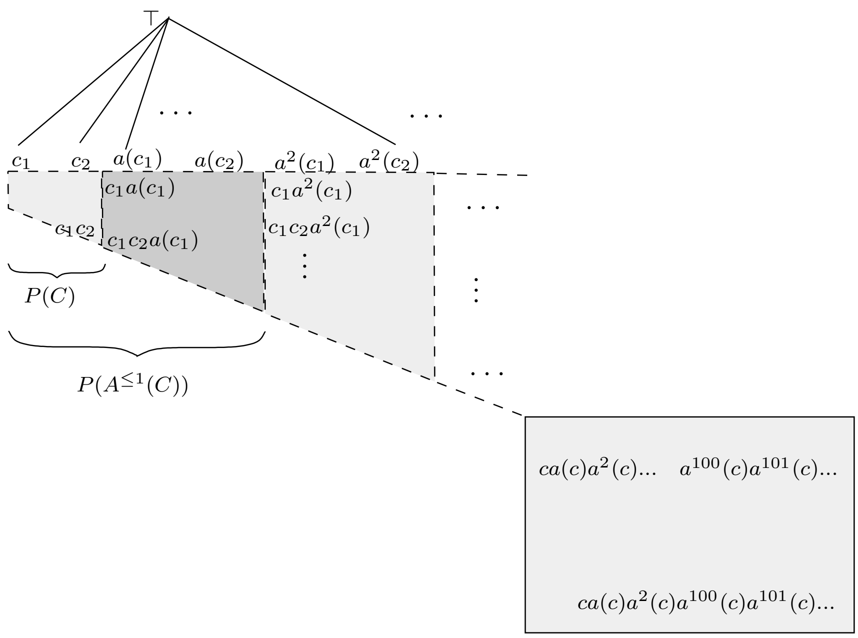

Fig. 2. $P(A^*(C))$

This graded view of the dual space is closely related to the coalgebraic construction that takes advantage of the fact that DLA is $Alg(H)$ for a functor on DL since the axioms for the attribution operations are all of rank 1, see e.g. [3]. In the coalgebraic setting $P(A^*(C))$ is obtained as the inverse limit of the posets

$$P(C) \twoheadleftarrow P(A^{\leq 1}(C)) \twoheadleftarrow P(A^{\leq 2}(C)) \dots$$

where the projection $P(A^{\leq k-1}(C)) \twoheadleftarrow P(A^{\leq k}(C))$ is the map that, from a conjunction, removes all conjuncts of depth k. Notice that we also have the embeddings going forward, making these into retractions. The inverse limit of the projections is the dual space, and the direct limit of the embeddings is what we call the finite part of the dual space. Finally, for each $a \in A$, the map $f_a : P(A^*(C)) \to P(A^*(C))$ dual to the attribute a, is given by the maps $f_a^k : P(A^{\leq k}(C)) \to P(A^{\leq k-1}(C))$ which have the same action as f_a as described above.

6 Computing the terminal solution

In this section we will discuss how to find the terminal solution of a given knowledge base. We start by considering knowledge bases without attributes. In this case the construction of solutions is considerably simpler as the free algebra and its dual frame are finite and the topology on the dual frame is the trivial topology.

6.1 The attribute-free case

In the attribute-free case we work over the finite poset $P(C)$ as described above, and we want to calculate

$$P_{\mathcal{B}} = \big(\bigcup\{max(\widehat{h_{\mathcal{O}}(t)}) \mid t \in I\}\big).$$

That is, we need to calculate $max(\widehat{h_{\mathcal{O}}(t)}))$ for each $t \in I$ where $h_{\mathcal{O}}$ is the quotient map corresponding to the congruence $\theta_{\Pi} = \bigvee\{\theta_{(r,s)} \mid (r,s) \in \Pi\}$. Dually, this last equality tells us $P_{\mathcal{O}} = \bigcap\{P_{(r,s)} \mid (r,s) \in \Pi\}$ where $P_{(r,s)}$ is the dual of $F_{DL}(C)/\theta_{(r,s)}$. This will allow us to devise a method to find $P_{\mathcal{B}}$ in which we compute with the posets $P_{(r,s)}$. Since we want to compute maximal antichains it makes more sense to work with a slightly different representation of the duality where the lattice dual to a poset is given in a way similar to the Hoare powerdomain construction.

Definition 8. *For any poset P, let $H(P)$ be the set of all anti-chains of P with the Hoare order:*

$$T \leq T' \quad \textit{if and only if} \quad \forall p \in T\ \exists p' \in T' \quad p \leq p'.$$

It is not hard to see that this makes $H(P)$ into a lattice with the operations given by, for $T, T' \in H(P)$,

$$T \vee T' = max(T \cup T'),$$
$$T \wedge T' = max\{p \in P : \exists q \in T\ \exists q' \in T' \quad (p \leq q \textit{ and } p \leq q')\}.$$

For a given finite poset P, $H(P)$ is isomorphic to the collection $\mathcal{D}(P)$ of downsets of P by $T \mapsto {\downarrow}T$ and with inverse $S \mapsto max(S)$. For a finite lattice D, writing $J(D)$ for the set of join-irreducible elements of D, $D \cong \mathcal{D}(J(D)) \cong H(J(D))$ with $d \mapsto \widehat{d} = {\downarrow}d \cap J(D) \mapsto max({\downarrow}d \cap J(D))$ the latter of which we will denote by $\tilde{d}$. Note that in the case of a term $t \in F_{DL}(C)$, the anti-chain $\tilde{t}$ consists of the pure conjunctions in the disjunctive normal form of t. That is, computing $\tilde{t}$ is just computing the disjunctive normal form of t.

Remark 1. For a poset P and a subposet $P' \subseteq P$, the corresponding quotient map $h : H(P) \to H(P')$ is given by $T \mapsto max[({\downarrow}T) \cap P']$. A nice property of this representation of the lattice dual to a finite poset is that the lower adjoint of h, $h^\flat : H(P') \to H(P)$, is just the identity map, and in particular, since adjoints satisfy $h^\flat h(T) \leq T$, we have $h(T) \leq T$ for all $T \in H(P)$. Further, if $h : H(P) \to H(P')$ factors through $h' : H(P) \to H(P'')$, i.e. if $P' \subseteq P'' \subseteq P$, then $h(T) \leq h'(T) \leq T$ for all $T \in H(P)$.

For a pair $(r, s) \in \Pi$, the quotient $F_{DL}(C)/\theta_{(r,s)}$ corresponds to the subset $P_{(r,s)}$ of all $p \in P(C)$ satisfying $p \leq r \iff p \leq s$, or equivalently $p \leq r \vee s \Rightarrow p \leq r \wedge s$. Thus we may assume our pairs (r, s) are given with $r \geq s$ and then

$$P_{(r,s)} = P(C) - Q_{(r,s)} \quad \text{where} \quad Q_{(r,s)} = \{p \in P(C) \mid p \leq r \text{ but } p \not\leq s\}.$$

The following proposition provides the basic step of the method for computing $P_{\mathcal{B}}$ for an attribute-free knowledge base $\mathcal{B}$.

Proposition 1. *Let $(r, s) \in F_{DL}(C)^2$ and $t \in F_{DL}(C)$. Then the quotient map $h_{(r,s)}$ corresponding to $\theta_{(r,s)}$ is given by*

$$\begin{array}{rcl} h_{(r,s)} : H(P(C)) & \to & H(P_{(r,s)}) \\ \tilde{t} & \mapsto & max(\bigcup_{p \in \tilde{t}} h_{(r,s)}(p)) \end{array}$$

where

$$h_{(r,s)}(p) = \begin{cases} p & \text{if } p \not\leq \tilde{r} \\ max(\{pp' : p' \in \tilde{s}\}) & \text{if } p \leq \tilde{r} \end{cases}$$

Note that $p \leq \tilde{r}$ provided there is $p' \in \tilde{r}$ with $p \leq p'$ and the latter holds if and only if $p' \subseteq p$ when these are viewed as subsets of C. Also, the conjunction of pure conjunctions $p, p' \in P(C)$ is just the union of these viewed as subsets of C.

Proof. First of all $h_{(r,s)}$ is join preserving, so $h_{(r,s)}(\tilde{t}) = \bigvee_{p \in \tilde{t}} h_{(r,s)}(p))$ which is $max(\bigcup_{p \in \tilde{t}} h_{(r,s)}(p))$. Also $P_{(r,s)} = P(C) - \{p \in P(C) \mid p \leq r \text{ but } p \not\leq s\}$. Thus it is clear that if $p \not\leq \tilde{r}$ then $h_{(r,s)}(p) = p$. If on the other hand $p \leq \tilde{r}$, then $h_{(r,s)}(p) = h^\flat_{(r,s)}(h_{(r,s)}(p)) = p \wedge \tilde{s}$. Also $p \wedge \tilde{s} = max(\{pp' : p' \in \tilde{s}\})$ since $P(C)$ is closed under $\wedge$.

Now we are able to compute $h_{\mathcal{O}}(\tilde{t})$. We start by computing $\tilde{t}$, i.e. the disjunctive normal form of t. Next we generate a sequence $T_0, T_1, T_2, \ldots$ by applying the subroutine described above repeatedly for varying choices of $(r, s) \in \Pi$.

Let $(r_0, s_0), \dots, (r_k, s_k)$ be a list of the elements in Π and define an infinite sequence $\{(u_n, v_n)\}$ by repeating this list time after time. That is, $\{(u_n, v_n)\}$ is the sequence $(r_0, s_0), \dots, (r_k, s_k), (r_0, s_0), \dots, (r_k, s_k), \dots$. We then define

$$T_0 = \tilde{t}$$
$$T_n = h_{(u_n,v_n)}(T_{n-1}).$$

By Remark 1 it follows that

$$\tilde{t} = T_0 \geq T_1 \geq T_2 \geq \dots \qquad \geq h_{\mathcal{O}}(\tilde{t}).$$

Since $H(P(C))$ is a finite set, this sequence must eventually be constant. In fact, one may show that once it has stayed constant through a whole cycle of choices of $(r, s) \in \Pi$, we have reached $h_{\mathcal{O}}(\tilde{t})$.

Having computed $h_{\mathcal{O}}(\tilde{t})$ for all inserted terms t, one finds $P_{\mathcal{B}}$ by taking the union of all their maximal elements.

6.2 The general case

Adding attribution operations clearly complicates matters. We have to deal with an infinite lattice and with its dual space, which carries a non-trivial topology. However, some things remain the same: all clopen downsets of $P(A^*(C))$ are in fact downsets of finite anti-chains of the finite part of $P(A^*(C))$, and, with the graded or coalgebraic perspective on $P(A^*(C))$, we can work by approximating in finite DLs of the same type as in the attribute-free case (the generating set is just $A^{\leq n}(C)$ for some n rather than C). In the sequel we will not make a notational difference, for DLA terms t, between the term and the corresponding anti-chain $\tilde{t}$, nor between $h_{\mathcal{O}}(t)$, $\widehat{h_{\mathcal{O}}(t)}$ and $max(\widehat{h_{\mathcal{O}}(t)})$.

We first study the structure of $P_{\mathcal{O}}$. Recall, $P_{\mathcal{O}} = \bigcap_{(r,s)\in\Pi} P_{(r,s)}$, where $P_{(r,s)}$ is the subspace of $P(A^*(C))$ corresponding dually to the quotient $F_{DLA}(C) \to F_{DLA}(C)/\theta_{(r,s)}$. In the attribute-free setting $P_{(r,s)}$ is obtained by removing all points below r that are not below s (assuming, as above, without loss of generality that $r > s$). However, with a set of attributes involved, $P_{(r,s)}$ is not given by $P(A^*(C)) - (\downarrow r - \downarrow s)$, as this set is not closed under the actions of the maps f_a. Dually, this corresponds to the fact that requiring $r = s$ induces $w(r) = w(s)$ for every composition of attributes $w \in A^*$. However, we may understand DLA-quotients of $F_{DLA}(C)$ via DL-quotients of $F_{DL}(A^*(C))$. We will use superscripts DLA and DL to indicate which setting we are working in.

As $Q^{DL}_{(r,s)} = \downarrow r - \downarrow s$ is open in $P(A^*(C))$ and f_w is continuous, it follows that $P(A^*(C)) - \bigcup_{w\in A^*} f_w^{-1}(Q^{DL}_{(r,s)})$ is a closed subspace of $P(A^*(C))$. One easily proves that it is the largest closed subspace of $P(A^*(C))$ closed under the $f_a's$ for which r and s get identified under the dual quotient map. By the adjunction property it follows that $f_w(x) \in Q^{DL}_{(r,s)}$ iff $x \in Q^{DL}_{(w(r),w(s))}$, hence $f_w^{-1}(Q^{DL}_{(r,s)}) = Q^{DL}_{(w(r),w(s))}$ and thus

$$P^{DLA}_{\mathcal{O}} = P(A^*(C)) - \bigcup_{\substack{w\in A^* \\ (r,s)\in\Pi}} Q^{DL}_{(w(r),w(s))}.$$

This means we can attempt to find $h_{\mathcal{O}}(t)$ by applying the same subroutine as in the attribute-free case: Starting from a term t (which is necessarily an anti-chain in the finite part) we repeatedly apply quotient maps $h^{DL}_{(x,y)}$ (as described in Proposition 1) for $(x,y) \in \Pi^* = \{(w(r), w(s)) : (r,s) \in \Pi, w \in A^*\}$. Note that for any term t' we reach in this way, $t \geq t' \geq h_{\mathcal{O}}(t)$. If at some point we reach a term t' which satisfies $h^{DL}_{(x,y)}(t') = t'$ for all $(x,y) \in \Pi^*$, then $t' = h_{\mathcal{O}}(t)$.

To check whether or not $h^{DL}_{(x,y)}(t') = t'$ for *all* $(x,y) \in \Pi^*$, we only have to consider finitely many pairs in Π^*. To explain why this is the case, we start with two definitions.

Definition 9. *Let t be a term and $(x,y) \in F_{DL_\perp A}(C)^2$. We say t is* rejected by (x,y) *if $h^{DL}_{(x,y)}(t) \neq t$.*

We generalize the notion of A-depth given in Definition 7 for basic conjunctions.

Definition 10. *Let $t \in F_{DLA}(C)$, $t \neq \top, t \neq \bot$. We have identified t with a non-empty finite anti-chain of non-empty basic conjunctions (namely, the ones that occur in its disjunctive normal form). We define the* A-depth *of t, $d_A(t)$, to be the maximum of $d_A(p)$ for p a basic conjunction in t.*

For two basic conjunctions p and p', $p \leq p'$ implies $d_A(p) \geq d_A(p')$. Hence, a term t' being rejected by $(x,y) \in \Pi^*$, implies that $d_A(t') \geq d_A(p)$ for some p in x. So to check whether t' is rejected, we only have to consider the finite collection of elements of the form $(w(r), w(s))$ with $(r,s) \in \Pi$ and $w \in A^{\leq d_A(t)}$.

There do however exist knowledge bases with inserted terms t for which $h_{\mathcal{O}}(t)$ is not in the finite part of the space $P(A^*(C))$. Consider for example $\mathcal{B} = (\{c\}, \{a\}, \{(c, ca(c))\}, \{c, a(c)\})$. Starting from the inserted term c we get the following chain:

$$c \xrightarrow{(c,ca(c))} ca(c) \xrightarrow{(a(c),a(c)a^2(c))} ca(c)a^2(c) \xrightarrow{(a^2(c),a^2(c)a^3(c))} ca(c)a^2(c)a^3(c)\ldots$$

Above each arrow we have indicated which element of Π^* was used. In this case, $h_{\mathcal{O}}(c)$ is the conjunction $p = \bigwedge_{n\geq 0} a^n(c)$, an infinite point of the dual frame. The dual of the terminal solution is the two element chain, $\{p, q = \bigwedge_{n\geq 1} a^n(c)\}$ for which $f_a(p) = p$ and $f_a(q) = p$. Obtaining solutions in which f_a has fixpoints or maps downwards is impossible in the finite part of $P(A^*(C))$.

In the above example, the solution is finite but requires using points of the infinite part of $P(A^*(C))$. There are also KBs for which the terminal solution is infinite, e.g. $\mathcal{B} = (\{c,d\}, \{a\}, \{(c, ca(c) \vee cd)\}, \{c\})$. We call a solution of a KB *totally finite* if it lies within the finite part of $P(A^*(C))$. This implies in particular that the solution is finite. The algorithm for finding the terminal solution described above terminates if and only if the terminal solution is totally finite. We would like to be able to determine whether or not a given KB has a totally finite terminal solution. This problem has not been solved as far as we are aware. We do give two sufficient conditions for termination and a condition that implies non-termination, but these are not exhaustive.

Definition 11. *Let $\mathcal{B} = (C, A, \Pi, I)$ be a KB, p be a basic conjunction and $c \in C$. We define the* A_c-depth *of p to be the smallest $n \in \mathbb{N} \cup \{-1\}$ such that*

$p \in P(\, A^*(C - \{c\}) \cup A^{\leq n}(\{c\})\,)$. In general, for a term $t \in F_{DLA}(C) - \{\bot, \top\}$, and a concept $c \in C$, $d_{A_c}(t)$ is defined to be the maximum of $d_{A_c}(p)$ for p a basic conjunction in t.

From now on, we will omit A and just write $d(t)$ and $d_c(t)$.

Theorem 5. *(Termination condition I) Let $\mathcal{B} = (C, A, \Pi, I)$ be a KB. If, for every terminological axiom $(r, s) \in \Pi$, there is a concept c that occurs in r, i.e. $d_c(r) \geq 0$, with, for all $(u, v) \in \Pi$, $d_c(v) \leq d_c(u)$, then the terminal solution of $\mathcal{B}$ is totally finite.*

Proof. We will show that, under the condition stated in the theorem, $h_{\mathcal{O}}(t)$ is finite, i.e. it is a finite set of basic conjunctions, for every term t. It thus follows that the terminal solution is totally finite.

Let t be a term and define $\Pi' = \{(w(r), w(s)) : (r, s) \in \Pi, w \in A^{\leq d(t)}(C)\}$. This is a finite set of terminological axioms and therefore we may apply the algorithm for the attribute-free case and find a term t' that is not rejected by any of the terminological axioms in Π'. We claim that t' is not rejected by any element of Π^*.

Let $(x, y) = (w(r), w(s))$ be an element of Π^* that is not in Π'. By assumption there exists a concept c that occurs in r such that for all $(u, v) \in \Pi$, $d_c(v) \leq d_c(u)$. Using the fact that t' is the result of finitely many applications of the subroutine of Proposition 1, one easily shows that $d_c(t') = d_c(t)$. By assumption $(x, y) \notin \Pi'$ and hence $d_c(x) > d(t) \geq d_c(t) = d_c(t')$ and thus $t' \not\leq x$ and t' is not rejected by (x, y).

Theorem 6. *(Termination condition II) Let $\mathcal{B} = (C, A, \Pi, I)$ be a KB such that for every terminological axiom $(r, s) \in \Pi$, $d(r) = d(s)$. Then the terminal solution of $\mathcal{B}$ is totally finite.*

Proof. Again it suffices, for a given term t, to consider the set $\Pi' = \{(w(r), w(s)) : (r, s) \in \Pi, w \in A^{\leq d(t)}\}$ as, under the given condition, the depth of a term does not change when applying the subroutine.

It is easy to find KBs with totally finite terminal solutions that do not satisfy either of the above conditions. Finding a sharp syntactic termination condition does not seem probable to us but the above conditions cover most actual applications as concepts and their attribution translates typically are disjoint (e.g., in an ontology for real estate, c = a geographical region and the concept $a(c)$ of being a piece of real estate located in region c, only have the inconsistent concept as common subsumer and a^2 of anything is inconsistent). Nevertheless, finding a sharp condition for total finiteness seems worthwhile and we expect that an algorithmic approach may be more fruitful. A first step in this direction is made by the non-termination condition below. For simplicity, we formulate it only for basic conjunctions, but it is readily extended to incorporate disjunctions of basic conjunctions. We start with a definition.

Definition 12. *For two words w, w', we say w' is an* extension *of w, notation $w \sqsubseteq w'$, if there exists a word w'' such that $w' = ww''$.*

Theorem 7. *(Non-termination condition) Let $\mathcal{B} = (C, A, \Pi, I)$ be a KB such that all terminological axioms are pairs of basic conjunctions and let $t \in I$ be a basic conjunction. If there exists a chain of rejections:*

$$t = t_0 \xrightarrow{w_0(r_0,s_0)} t_1 \xrightarrow{w_1(r_1,s_1)} \cdots \xrightarrow{w_{L-1}(r_{L-1},s_{L-1})} t_L$$

such that there exist $n < L$ and $w_L \in A^*$ with $f_{w_n}(t_n) = f_{w_L}(t_L)$ and for every k with $n \leq k \leq L-1$, $w_k \sqsubseteq w_{k+1}$, then $h_{\mathcal{O}}(t)$ is an infinite point of the dual frame.

Proof. It suffices to show that we may extend the chain of rejections given above ad infinitum. We may assume $n = 0$. As, by assumption, $f_{w_L}(t_L) = f_{w_0}(t_0) \leq r_0$, by adjunction $t_L \leq w_L(r_0)$. Similarly one shows $t_L \not\leq w_L(s_0)$. Hence, t_L is rejected by $w_L(r_0, s_0)$ and we define $t_{L+1} = t_L \wedge w_L(s_0)$. As $w_0 \sqsubseteq w_1$, there exists a word w_1' such that $w_1 = w_0 w_1'$. One may show that t_{L+1} is rejected by $(w_L w_1')(r_1, s_1)$ and we define $t_{L+2} = t_{L+1} \wedge (w_L w_1')(s_1)$. It is readily seen that one may continue this way, thereby forming an infinite chain of rejections.

This non-termination condition is not sharp. Consider for example the knowledge base $\mathcal{B} = (\{c, d, e\}, \{a\}, \{z_1, z_2, z_3\}, \{pa(p)\})$, where $p = cde$ and

$$z_1 = (cde,\ cdea(c)),\ \ z_2 = (a(c)a(d),\ a(c)a(d)a^2(e)),\ \ z_3 = (a^2(c),\ a^2(c)a^2(d)).$$

There exists no chain of rejections satisfying the non-termination condition given above. Nevertheless, $h_{\mathcal{O}}(pa(p))$ is an infinite conjunction as we have the following unending chain of rejections,

$$pa(p) \overset{z_2}{\to} pa(p)a^2(e) \overset{az_1}{\to} pa(p)a^2(c)a^2(e) \overset{z_3}{\to} pa(p)a^2(p) \overset{az_2}{\to} pa(p)a^2(p)a^3(e) \overset{a^2z_1}{\to} \ldots$$

However, when we add the pair az_1 to our set Π, we do recognize the chain above as one satisfying our non-termination condition. So we may broaden the class of KB's that are recognized as having an infinite terminal solution by adding some of the elements of Π^* to the terminological axioms Π. This raises the question whether there is a general way to extend, for a given knowledge base $\mathcal{B} = (C, A, \Pi, I)$, the set of terminological axioms Π to a finite set $\Pi' \subseteq \Pi^*$, such that $\mathcal{B}$ has a totally finite terminal solution if and only if $\mathcal{B}' = (C, A, \Pi', I)$ does not satisfy the above non-termination condition (extended to disjunctions). We have no counterexample to this, the question is still open.

7 Related and future work

The approach to generative knowledge representation presented in this paper is a generalisation of the work of Oles [10]. However, Oles only identifies the universal solution of a given ontological framework thus making work with additional operators impossible. Algorithmically, however, Oles also only computes (max of) $h_{\mathcal{O}}(t)$ for the terms t that are of interest. For this reason we expect that in the applications to medical expert systems that he mentions, it is in fact the terminal solution that one works with. By introducing the concept of a classification of a term w.r.t. a solution of an ontological framework we are able to identify this solution as the least one with a certain property, thereby giving its mathematical significance. This in turn makes it feasible to work with additional operations. While we consider the identification of the concept of terminal solution a significant advancement, the fact that these need not be finite raises new questions. Our termination and non-termination conditions cover most practical applications, but finding a sharp termination condition seems an interesting and worthwhile problem. We feel that a coalgebraic approach to this problem is most promising. Since the dual of a DL is up-to exponentially bigger than the generating set, this is the worst case complexity, but in practice the algorithm performs better and this should be studied experimentally. Further, we have restricted ourselves to deterministic modalities, but the concept and existence of a terminal solution does not depend on this. We expect that our results readily can be extended to non-deterministic modalities. Finally, thanks to helpful input from the referees, we have made first steps in understanding the relation of this work to description logic. It requires further study and discussion and we hope the publication of this paper will help to achieve this goal.

References

1. Baader, F., Lutz, C., Sturm, H., Wolter, F., Fusions of Description Logics and Abstract Description Systems, *J. Artif. Intell. Res. (JAIR)* **16**: 1-58, 2002.
2. Baader, F., and Nutt, W., Basic description logics, in F. Baader et al. (eds), *The description logic handbook: theory, implementation, and applications*, Cambridge University Press, 43–95, 2003.
3. Bezhanishvili, N., Kurz, A.; *Free modal algebras: a coalgebraic perspective*, in Lecture Notes in Computer Science 4624, Springer Berlin, pp. 143-157, 2007
4. Bruun, H., *The developement of a lattice structured database.* IMM Technical report, Technical University of Denmark, DTU, 2006. http://www2.imm.dtu.dk/pubdb/views/publication_details.php?id=4967
5. Bruun, H., Fischer Nilsson, J., Gehrke, M.; *Lattice ordered ontologies*, preprint.
6. Davey, B. A., Priestley, H. A.; *Introduction to Lattices and Order*, Second Edition, Cambridge University Press, 2002.
7. Gehrke, M., Jónsson, B., Bounded distributive lattice expansions. *Mathematica Scandinavica* **94**(1) (2004). 13–45.
8. Ghilardi, S., Santocanale, L., Algebraic and Model-Theoretic Techniques for Fusion Decidability in Modal Logic. Vardi M., Voronkov A. (eds.) *Logic for Programming, Artificial Intelligence and Reasoning (LPAR 03)*, Springer *LNAI* **2850**, 152-166, 2003.
9. Nilsson, J.F., Relational data base model simplified and generalized as algebraic lattice with attribution, FQAS 1994, 101-104.
10. Oles, F. J.; An application of lattice theory to knowledge representation, *Theoretical Computer Science* **249** (2000), 163–196.
11. OntoQuery project net site: http://www.ontoquery.dk.
12. Sofronie-Stokkermans, V., Automated theorem proving by resolution in non-classical logics, *Annals of Mathematics and Artificial Intelligence* **49** (1-4), 221-252, 2007.