In: Proc. 7th International Conference on Logic for Programming and Automated Reasoning (LPAR 2000), Reunion, France, pp. 1-19. LNCS 1955, © 2000 Springer.

On the Complexity of Theory Curbing

Thomas Eiter and Georg Gottlob

Institut für Informationssysteme, TU Wien
Favoritenstraße 9–11, A-1040 Wien, Austria
eiter@kr.tuwien.ac.at, gottlob@dbai.tuwien.ac.at

Abstract. In this paper, we determine the complexity of propositional theory curbing. Theory Curbing is a nonmonotonic technique of common sense reasoning that is based on model minimality but unlike circumscription treats disjunction inclusively. In an earlier paper, theory curbing was shown to be feasible in PSPACE, but the precise complexity was left open. In the present paper we prove it to be PSPACE-complete. In particular, we show that both the model checking and the inferencing problem under curbed theories are PSPACE complete. We also study relevant cases where the complexity of theory curbing is located – just as for plain propositional circumscription – at the second level of the polynomial hierarchy and is thus presumably easier than PSPACE.

1 Introduction

Circumscription [15] is a well-known technique of nonmonotonic reasoning based on model-minimality. The (total) circumscription $Circ(T)$ of a theory T, which is a finite set of sentences, consists of a formula whose set of models is equal to the set of all *minimal* models of T. For various variants of circumscription, see [14].

As noted by various authors [5, 6, 17–20], reasoning under minimal models runs into problems in connection with disjunctive information. The minimality principle of circumscription often enforces the *exclusive* interpretation of a disjunction $a \vee b$ by adopting the models in which either a or b is true but not both. There are many situations in which an *inclusive* interpretation is desired and seems more natural (for examples, see Section 2).

To redress this problem, and to be able to handle inclusive disjunctions of positive information properly, the method of *theory curbing* was introduced in [8]. This method is based on the notion of a *good model* of a theory. Roughly, a good model of a theory T is either a minimal model, or a model of T that constitutes a minimal upper bound of a set of good models of T. The sentence $Curb(T)$ has as its model precisely the good models of T. When T is a first-order theory, $Curb(T)$ is most naturally expressed as a third-order formula. However, in [8], it was shown that $Curb(T)$ is expressible in second-order logic.

Circumscription is usually not applied to *all* predicates of a theory, but only to the members of a list $\mathbf{p}$ of predicates, where the predicates from a list $\mathbf{z}$ disjoint with $\mathbf{p}$, called the *floating* predicates, may be selected such that the predicates in $\mathbf{p}$ become as small as possible; the remaining predicates not occurring in $\mathbf{p}$ and $\mathbf{z}$ (called *fixed* predicates) are treated classically. In analogy to this, in [8], formulas of the form

$Curb(T; \mathbf{p}, \mathbf{z})$ are defined, where curbing is applied to the predicates in list $\mathbf{p}$ only, while those from list $\mathbf{z}$ (the floating predicates) are interpreted in the standard way. In the propositional case, the lists $\mathbf{p}$ and $\mathbf{q}$ of predicate symbols are lists of propositional variables (corresponding to zero-ary predicates).

Since its introduction in [8], the curbing technique has been used and studied in a number of other papers. For instance, Scarcello, Leone, and Palopoli [21], provide a fixpoint semantics for propositional curbing and derive complexity results for curbing Krom theories, i.e., clausal theories where each clause contains at most two literals. Liberatore [11, 12] bases a belief update operator on a restricted version of curbing. Note that curbing is a purely model-theoretic and thus syntax-independent method. In particular, for two logically equivalent theories T and T', it holds that $Curb(T)$ is logically equivalent to $Curb(T')$. Curbing can be applied to arbitrary logical theories and not just to logic programs. In the context of disjunctive logic programming, various syntax-dependent methods of reasoning that do not treat disjunction exclusively were defined in [5, 18, 17, 19, 20, 6].

In [8], the following two major reasoning problems under curbing where shown to be in PSPACE:

Curb Model Checking: Given a propositional theory T, an interpretation M of T, and disjoint lists $\mathbf{p}$ and $\mathbf{z}$ of propositional variables, decide whether M is a good model of T w.r.t. $\mathbf{p}$ and $\mathbf{z}$ (i.e., decide whether M is a model of $Curb(T; \mathbf{p}, \mathbf{z})$).

Curb Inference : Given a propositional theory T, disjoint lists $\mathbf{p}$ and $\mathbf{z}$ of propositional variables, and a propositional formula G, decide whether $Curb(T; \mathbf{p}, \mathbf{z}) \models G$.

The precise complexity of curbing, for both model checking and inferencing, was left open in [8]. Note that model checking for propositional circumscription is coNP complete [3] and inferencing under propositional circumscription is Π_2^P complete [7]. It was conjectured in [21, 11] that curbing is of higher complexity than circumscription. This is intuitively supported by a result of Bodenstorfer [2] stating that in an explicitly given set of models, witnessing that some particular model is good may involve an exponential number of smaller good models (for a formal statement of this result, see Section 3).

The main result of this paper answers the above questions. We prove that Curb Model Checking and Curb Inference are PSPACE-complete. Both problems remain PSPACE-hard even in case of *total* curbing, i.e., when curbing is applied to *all* propositional variables, and thus the list $\mathbf{z}$ of floating propositional variables is empty and no propositional variables are fixed. The proof takes Bodenstorfer's construction as a starting point and shows how to reduce the evaluation of quantified Boolean formulas to theory curbing.

The PSPACE-completeness result strongly indicates that curbing is a much more powerful reasoning method than circumscription, and that it can not be reduced in polynomial time to circumscription. Thus, circumscriptive theorem provers can not be efficiently used for curb reasoning. On the other hand, a curb theorem prover could be based on a QBF solver (see [10, 4, 16, 1, 9]).

After proving our main result, we identify classes of theories for which the complexity of curbing is located at a lower complexity level. Specifically, we show that if a theory T has the *lub property*, that is, every set of good models of T has a *least* (unique minimal) upper bound, then propositional Curb Model Checking is in Σ_2^P, while Curb Inference is feasible in Π_2^p. Note that relevant classes of theories have this property. For example, as shown by Scarcello, Leone, and Palopoli [21], Krom theories enjoy the lub property. More specifically, in [21] it is shown that the *union* of any pair of good models of a Krom theory is a good model, too. This is clearly a special case of the lub property; in in [21], this special property is used to show that Curb Model Checking for propositional Krom theories is in Σ_2^P. The lub property can be further generalized. We show that following less restrictive *weak least upper bound property (weak lub property)* also leads to complexity results at the second level of the polynomial hierarchy: T has the weak least upper bound (weak lub) property, if every non-minimal good model of φ is the lub of *some* collection $\mathcal{M}$ of good models of T. The lub and the weak lub property are of interest not only in the case of propositional circumscription, but also in case of predicate logic. We therefore discuss these properties in the general setting.

The rest of this paper is organized as follows. In the next Section 2, we review some examples from [8] and give a formal definition of curbing. We then prove in Section 3 the main result stating that propositional Curb Model Checking and Curb Inference are both PSPACE-complete. In Section 4 we discuss the lub property, and the final Section 5 the weak lub property.

2 Review of Curbing

In this section, we review the concept of "good model" and give a formal definition of curbing. The presentation follows very closely the exposition in [8]; the reader familiar with [8] may skip the rest of this section.

2.1 Good Models

Let us first describe two scenarios in which an inclusive interpretation of disjunction is desirable. Models are represented by their positive atoms.

Example 1: Suppose there is a man in a room with a painting, which he hangs on the wall if he has a hammer and a nail. It is known that the man has a hammer or a nail or both. This scenario is represented by the theory T_1 in Figure 1. The desired models are h, n, and hnp, which are encircled. Circumscribing T_1 by minimizing all variables yields the two minimal models h and n (see Figure 1). Since p is false in the minimal models, circumscription tells us that the man does not hang the painting up. One might argue that the variable p should not be minimized but fixed when applying circumscription. However, starting with the model of T_1 where h, n and p are all true and then circumscribing with respect to h and p while keeping p true, we obtain the models hp and np, which are not very intuitive. If we allow p to vary in minimizing h and n, the outcome is the same as for minimizing all variables. On the other hand, the model hnp seems plausible. This model corresponds to the inclusive interpretation of the disjunction $h \vee n$. □

Example 2: Suppose you have invited some friends to a party. You know for certain that one of Alice, Bob, and Chris will come, but you don't know whether Doug will

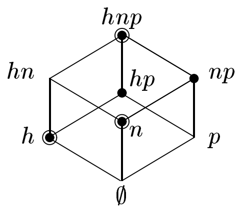

$$T_1 = \{\, h \vee n,\ (h \wedge n) \rightarrow p \,\}$$

Fig. 1. The hammer-nail-painting example

come. You know in addition the following habits of your friends. If Alice and Bob go to a party, then Chris or Doug will also come; if Bob and Chris go, then Alice or Doug will go. Furthermore, if Alice and Chris go, then Bob will also go. This is represented by theory T_2 in Figure 2. Now what can you say about who will come to the party? Look

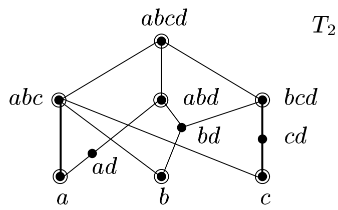

$$\begin{aligned} T_2 = \{\, & a \vee b \vee c, \\ & (a \wedge b) \rightarrow (c \vee d), \\ & (b \wedge c) \rightarrow (a \vee d), \\ & (a \wedge c) \rightarrow b \,\} \end{aligned}$$

Fig. 2. The party example

at the models of T_2 in Figure 2. Circumscription yields the minimal models a, b, and c, which interpret the clause $a \vee b \vee c$ exclusively in the sense that it is minimally satisfied. However, there are other plausible models. For example, abc. This model embodies an inclusive interpretation of a and b within $a \vee b \vee c$; it is also minimal in this respect. abd is another model of this property. Similarly, bcd is a minimal model for an inclusive interpretation of b and c. The models ad, bd, and cd are not plausible, however, since a scenario in which Doug and only one of Alice, Bob or Chris are present does not seem well-supported. □

In the light of these examples, the question arises how circumscription can be extended to work satisfactory. An important insight is that such an extension must take disjunctions of positive events seriously and allow inclusive (hence non-minimal) models, even if such models contain positive information that is not contained in any minimal model. On the other hand, the fruitful principle of minimality should not be abandoned by adopting models that are intuitively not concise. The idea of curbing is based on the synthesis of both: adopt the minimal inclusive models. That is, adopt for minimal models M_1, M_2 any model M which includes both M_1 and M_2 and is a minimal such model; in other words, M is a *minimal upper bound* (*mub*) for M_1 and M_2.

To illustrate, in Example 1 hnp is a mub for h and n (notice that hn is not a model), and in Example 2 abc is a mub for a and c; abd is another one, so several mub's can exist. In order to capture general inclusive interpretations, mub's of arbitrary collections $M_1, M_2, M_3, \ldots$ of minimal models are adopted.

It appears that in general not all "good" models are obtainable as mub's of collections of minimal models. The good model $abcd$ in Example 2 shows this. It is, however, a mub of the good models a and bcd (as well as of abc and abd). This suggests that not only mub's of collections of minimal models, but mub's of any collection of good models should also be good models.

The curbing approach to extend circumscription for inclusive interpretation of disjunctions is thus the following: adopt as good models the least set of models which contains all circumscriptive (i.e. minimal) models and which is closed under including mub's. Notice that this approach yields in Examples 1 and 2 the sets of intuitively good models, which are encircled in Figs. 1 and 2.

2.2 Formal Definition of Curbing

In this section we state the formal semantical definition of good models of a first-order sentence as defined in [8].

As for circumscription, we need a language of higher-order logic (cf. [22]) over a set of predicate and function symbols, i.e. variables and constants of finite arity $n \geq 0$ of suitable type. Recall that 0-ary predicate symbols are identified with propositional symbols.

A sentence is a formula φ in which no variable occurs free; it is of order $n+1$ if the order of any quantified symbol occurring in it is $\leq n$ [22]. We use set notation for predicate membership and inclusion. A *theory* T is a finite set of sentences. As usual, we identify a theory T with the sentence φ_T which is the conjunction $\bigwedge_{\varphi \in T} \varphi$ of all sentences in T.

A structure M consists of a nonempty set $|M|$ and an assignment $\mathcal{I}(M)$ of predicates, i.e. relations (resp. functions), of suitable type over $|M|$ to the predicate (resp. function) constants. The object assigned to constant C, i.e. the extension of C in M, is denoted by $[\![C]\!]_M$ or simply C if this is clear from the context. Equality is interpreted as identity. A model for a sentence φ is any structure M such that φ is true in M (in symbols, $M \models \varphi$). $\mathcal{M}[\varphi]$ denotes all models of φ.

Let $\mathbf{p} = p_1, \ldots, p_n$ be a list of first-order predicate constants and $\mathbf{z} = z_1, \ldots, z_m$ a list of first-order predicate or function constants disjoint with $\mathbf{p}$. For any structure M, let $\mathcal{M}^M_{\mathbf{p};\mathbf{z}}$ be the class of structures M' such that $|M| = |M'|$, and $[\![C]\!]_M = [\![C]\!]_{M'}$ for every constant C not occurring in $\mathbf{p}$ or $\mathbf{z}$. The pre-order $\leq^M_{\mathbf{p};\mathbf{z}}$ on $\mathcal{M}^M_{\mathbf{p};\mathbf{z}}$ is defined by $M_1 \leq^M_{\mathbf{p};\mathbf{z}} M_2$ iff $[\![p_i]\!]_{M_1} \subseteq [\![p_i]\!]_{M_2}$ for all $1 \leq i \leq n$. The pre-order $\leq_{\mathbf{p};\mathbf{z}}$ is the union of all $\leq^M_{\mathbf{p};\mathbf{z}}$ over all structures. We write $\mathcal{M}^M_{\mathbf{p}}$ etc. if $\mathbf{z}$ is empty; $\leq^M_{\mathbf{p}}$ and $\leq_{\mathbf{p}}$ are partial orders on $\mathcal{M}^M_p$ resp. all structures.

The circumscription of $\mathbf{p}$ in a first-order sentence $\varphi(\mathbf{p}, \mathbf{z})$ with $\mathbf{z}$ floating is the second-order sentence [13]

$$\varphi(\mathbf{p}, \mathbf{z}) \wedge \neg\exists \mathbf{p}', \mathbf{z}'(\varphi(\mathbf{p}', \mathbf{z}') \wedge (\mathbf{p}' \subset \mathbf{p}))$$

which will be denoted by $Circ(\varphi(\mathbf{p}, \mathbf{z}))$ ($\mathbf{p}$ and $\mathbf{z}$ will be always presupposed). Here $\mathbf{p}'$, $\mathbf{z}'$ are lists of predicate and function variables matching $\mathbf{p}$ and $\mathbf{z}$ and $\mathbf{p} \subset \mathbf{p}'$ stands for

$(\mathbf{p}' \subseteq \mathbf{p}) \wedge (\mathbf{p}' \neq \mathbf{p})$, where $(\mathbf{p}' \subseteq \mathbf{p})$ is the conjunction of all $(p_i' \subseteq p_i)$, $1 \leq i \leq n$. The following is a straightforward consequence of the definitions.

Proposition 2.1. [13] *$M \models \mathit{Circ}(\varphi(\mathbf{p}, \mathbf{z}))$ iff M is $\leq_{\mathbf{p};\mathbf{z}}$-minimal among the models of $\varphi(\mathbf{p}, \mathbf{z})$.*

We formally define the concept of a "good" model as follows. First define the property that a set of models is closed under minimal upper bounds.

Definition 2.1. *Let $\varphi(\mathbf{p}, \mathbf{z})$ be a first-order sentence. A set $\mathcal{M}$ of models of $\varphi(\mathbf{p}, \mathbf{z})$ is $\leq_{\mathbf{p};\mathbf{z}}$-closed iff, for every $\mathcal{M}' \subseteq \mathcal{M}$ and any model M of $\varphi(\mathbf{p}, \mathbf{z})$, if M is $\leq_{\mathbf{p};\mathbf{z}}$-minimal among the models of $\varphi(\mathbf{p}, \mathbf{z})$ which satisfy $M' \leq_{\mathbf{p};\mathbf{z}} M$ for all $M' \in \mathcal{M}'$ then $M \in \mathcal{M}$.*

Clearly the set of all models is closed. Further, every closed set must contain all $\leq_{\mathbf{p};\mathbf{z}}$-minimal models of $\varphi(\mathbf{p}, \mathbf{z})$ (let $\mathcal{M}' = \emptyset$); the empty set is closed iff $\varphi(\mathbf{p}, \mathbf{z})$ has no minimal model. We define goodness as follows.

Definition 2.2. *A model M of $\varphi(\mathbf{p}, \mathbf{z})$ is good with respect to $\mathbf{p}; \mathbf{z}$ iff M belongs to the least $\mathbf{p}; \mathbf{z}$-closed set of models of $\varphi(\mathbf{p}, \mathbf{z})$.*

Notice that good models only exist if a unique smallest closed set exists. The latter is immediately evident from the following characterization of goodness.

Proposition 2.2 ([8]). *A model M of $\varphi(\mathbf{p}, \mathbf{z})$ is good with respect to $\mathbf{p}; \mathbf{z}$ iff M belongs to the intersection of all $\mathbf{p}; \mathbf{z}$-closed sets.*

In [8], it was shown how to capture goodness by a sentence $\mathit{Curb}(\varphi(\mathbf{p}, \mathbf{z}); \mathbf{p}, \mathbf{z})$ whose models are precisely the good models of $\varphi(\mathbf{p}, \mathbf{z})$. Similar to circumscription, $\mathbf{p}$ are the minimized predicates (here under the *inclusive* interpretation of disjunction), $\mathbf{z}$ are the floating predicates, and all other predicates are fixed. Curbing is naturally formalized as a sentence of third-order logic, given that the definition of the set of good models of a theory involves sets of sets of models. However, in [8] it was also shown that curbing can be formalized in second-order logic.

In the present paper we do not need the formal definitions of $\mathit{Curb}(\varphi(\mathbf{p}, \mathbf{z}); \mathbf{p}, \mathbf{z})$ in third or second order logic, but we are interested in the problems Curb Inference and Curb Model Checking as defined in the introduction.

2.3 Previous Complexity Results on Propositional Curbing

Recall that in the propositional case, a structure M is a truth-value assignment to the propositional variables. The problems *Curb Model Checking* and *Curb Inference* were described in the introduction. In [8] it was shown that both problems are in PSPACE, and in fact can be solved in quadratic space.

Two possibilities to approximate the full set of good models by a subset are discussed in [8]. The first approximation is to limit iterated inclusion of minimal upper bounds. Let us define the notion of α-goodness for ordinals α.

Definition 2.3. *A model M of $\varphi(\mathbf{p}, \mathbf{z})$ is 0-good with respect to $\mathbf{p}$ and $\mathbf{z}$, if M is $\leq_{\mathbf{p};\mathbf{z}}$-minimal among the models of φ.*

A model M of $\varphi(\mathbf{p}, \mathbf{z})$ is α-good with respect to $\mathbf{p}$ and $\mathbf{z}$, if M is a $\leq_{\mathbf{p};\mathbf{z}}$ minimal upper bound of a set of models $\mathcal{M}$ of φ, such that for each model $M' \in \mathcal{M}$ there exists an ordinal $\beta < \alpha$ such that M' is β-good w.r.t. $\mathbf{p}$ and $\mathbf{z}$.

Informally, in the approximation, one chooses only the models that are α-good for some α such that $\|\alpha\| \leq \|\delta\|$, where the ordinal δ is a limit on the depth in building minimal upper bounds. The operator corresponding to such a restricted version of curbing is denoted by $Curb^{\delta}$. Notice that circumscription appears as the case $\delta = 0$, i.e. $Curb^0(\varphi(\mathbf{p}, \mathbf{z}); \mathbf{p}, \mathbf{z})$ is equivalent to $Circ(\varphi(\mathbf{p}, \mathbf{z}); \mathbf{p}, \mathbf{z})$.

Concerning the computational complexity, the following was shown in [8]:

Theorem 2.1. *For $Curb^{\delta}$ (with fixed constant δ) the model checking problem is Σ_2^P complete, while inferencing is Π_2^P complete.*

Thus, the inference problem is in the propositional case for finite constant δ as easy (and as hard) as circumscription.

Another potential approximation to curbing studied in [8] is to limit the cardinality of model sets from which minimal upper bounds are formed. Intuitively, this corresponds to limiting the number of inclusively interpreted disjuncts by a cardinal $\kappa > 0$. The concept of closed$_\kappa$ set is defined by adding in the definition of closed set the condition "$\|\mathcal{M}'\| \leq \kappa$"; goodness$_\kappa$ is the relative notion of goodness.

Clearly, goodness$_1$ is equivalent to circumscription. For $\kappa \geq 2$, (i.e. $|M|$ is finite) the following result was proven:

Theorem 2.2 ([8]). *Over finite structures, for every $\kappa \geq 2$ a model of $\varphi(\mathbf{p}, \mathbf{z})$ is good$_\kappa$ with respect to $\mathbf{p}; \mathbf{z}$ iff it is good with respect to $\mathbf{p}; \mathbf{z}$.*

This result, which fails for arbitrary structures, implies a dichotomy result on the expressivity of κ-bounded disjuncts: Either we get only the minimal models, or all models obtainable by unbounded disjuncts. Thus the method of bounded disjunction is not a really useful approximation.

3 Main Result: PSPACE Completeness of Theory Curbing

In this section, we shall prove that inference as well as model checking under curbing is PSPACE-complete. Intuitively, the problems have this high complexity since checking whether a model is good requests a "proof", given by a proper collection of models, which may have non-polynomial size in general.

That such large proofs are necessary has been shown by Bodenstorfer [2]. A *support* of a model M in a collection $\mathcal{F}$ of models is a subset $\mathcal{F}' \subseteq \mathcal{F}$ containing M such that every $M' \in \mathcal{F}'$ is in $\mathcal{F}$ a mub of some models $\mathcal{M} \subseteq \mathcal{F}' \setminus \{M'\}$. Note that every minimal model $M \in \mathcal{F}$ has a support $\{M\}$ and that all models in a support are good models. Furthermore, every good model of $\mathcal{F}$ has some support.

Bodenstorfer has defined a family $\mathcal{F}_n$, $n \geq 0$, of sets of models on an alphabet of $O(n)$ propositional atoms, such that $\mathcal{F}_n$ contains exponentially many models (in n), and

$\mathcal{F}_n$ itself is the only support of the unique maximal model M_n of $\mathcal{F}_n$. Informally, $\mathcal{F}_0 = \{\{a_0\}\}$, and the family $\mathcal{F}_n$ is constructed inductively by cloning $\mathcal{F}_{n-1}$ and adding some sets which ensure that the ¡maximal model needs all models for a proof of goodness (see Figure 3).

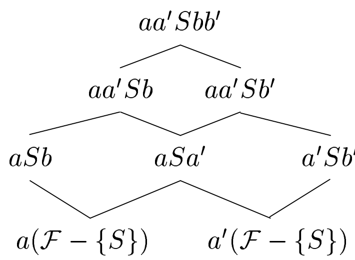

Fig. 3. Cloning a family $\mathcal{F}$ with unique maximal model S

3.1 Describing the exponential support family $\mathcal{F}_n$

We describe Bodenstorfer's family $\mathcal{F}_n$ by a formula Φ_n, such that $\mathcal{F}_n = mod(\Phi_n)$. The letters we use are $At_n = \{a_i, a'_i, b_i, b'_i \mid 1 \leq i \leq n\} \cup \{a_0\}$. We define the formula Φ_n inductively, where we set $\Phi_0 = a_0$ and $M_0 = \{a_0\}$, and for $n > 1$:

$$\Phi_n = (M_{n-1} \wedge \gamma_n) \vee (\neg M_{n-1} \wedge \Phi_{n-1} \wedge (a_n \leftrightarrow \neg a'_n) \wedge \neg b_n \wedge \neg b'_n),$$

where

$$\begin{aligned}\gamma_n = {} & (a_n \wedge b_n \wedge \neg a'_n \wedge \neg b'_n) \vee (a_n \wedge a'_n \wedge \neg b_n \wedge \neg b'_n) \vee \\ & (a'_n \wedge b'_n \wedge \neg a_n \wedge \neg b_n) \vee (a_n \wedge b_n \wedge a'_n \wedge \neg b'_n) \vee \\ & (a'_n \wedge b'_n \wedge a_n \wedge \neg b_n) \vee (a_n \wedge b_n \wedge a'_n \wedge b'_n);\end{aligned}$$

$$M_n = M_{n-1} \cup \{a_n, a'_n, b_n, b'_n\}.$$

Note that the left disjunct of Φ_n gives rise to six models, which extend M_{n-1} by the following sets of atoms:

$A_{n,1} = \{a_n, b_n\}$, $A_{n,0} = \{a'_n, b'_n\}$, $B_n = \{a_n, a'_n\}$, $C_{n,1} = \{a_n, a'_n, b_n\}$, $C_{n,0} = \{a_n, a'_n, b'_n\}$, and $D_n = \{a_n, a'_n, b_n, b'_n\}$.

Informally, $A_{n,1}$ (resp., $A_{n,0}$) represents the assignment of true (resp., false) to the atom a_n. The right disjunct of Φ_n generates recursively assignments to the other atoms a_{n-1}, …, a_1, such that certain minimal models of Φ_n represent truth assignments to the atoms $a_1, \ldots, a_n$ (see Figure 4).

Note that $M_n = M_{n-1} \cup D_n$ (i.e., all atoms are true) is, as easily seen, the unique maximal model of the formula Φ_n. The set of models of Φ_n over At_n, $mod(\Phi_n)$, defines the family $\mathcal{F}_n$ as described in [2]. Thus, each model $M \in mod(\Phi_n)$ is good, and M_n requires an exponential size support.

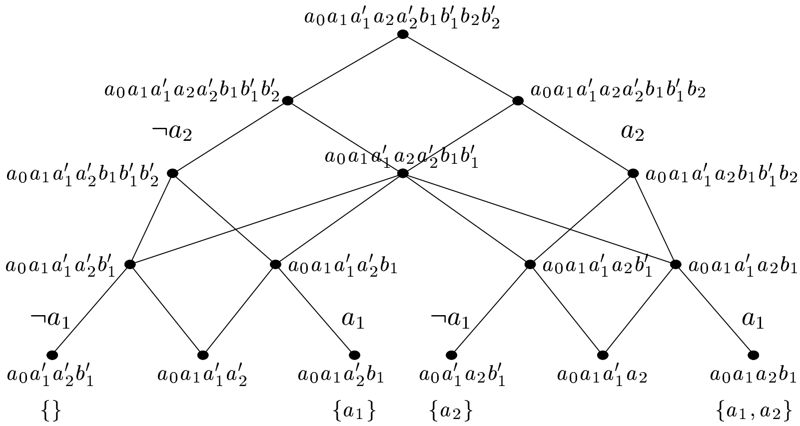

Fig. 4. The set of models $mod(\Phi_2)$

3.2 Evaluating a quantified Boolean formula on $mod(\Phi_n)$

We now show that a quantified Boolean formula (QBF)

$$F = Q_n a_n Q_{n-1} a_{n-1} \cdots Q_1 a_1 \varphi,$$

where each $Q_i \in \{\forall, \exists\}$ and φ is a Boolean formula over atoms $a_1, \ldots, a_n$, can be "evaluated" on the collection $mod(\Phi_n)$ of good models exploiting the curbing principle.

Roughly, the idea is as follows: $mod(\Phi_n)$ can be layered into n overlapping layers of models, where each layer i contains the models which are recursively generated by the left disjunct of the formula Φ_i. In each layer we have three levels of models. Neighbored layers i and $i-1$ overlap such that the bottom level of i is the top level of $i-1$ (see Figure 5). The minimal models in $mod(\Phi_n)$ are the bottom models of layer 1, and might be considered as the top model of an artificial layer 0. Similarly, the maximal model M_n in $mod(\Phi_n)$ might be viewed as a bottom model of an artificial layer $n+1$.

In order to "evaluate" the QBF F, we will obtain a formula $\Psi(F)$ from F by adding conjunctively a set of formulas $\Gamma(F)$ to Φ_n. Thus $\Psi(F) = \Phi_n \wedge \Gamma(F)$. The formulas in Γ will be chosen such that the overall structure of the set of good models of $\Psi(F)$ does not differ from the one of the set of models of Φ_n. In particular, each model M of Φ_n will correspond to some good model $f(M)$ of $\Psi(F)$ which augments M by certain atoms that describe the truth status of subformulas of F.

By adjoining $\Gamma(F)$ to Φ_n, we "adorn" the models in $mod(\Phi_n)$ with additional atoms which help us in evaluating the formula F along the layers. At a layer i in $mod(\Phi_n)$, we have fixed an assignment to the variables $a_{i+1}, \ldots, a_n$ already, where a_j is true if a_j occurs in the model, and a_j is false if a_j' occurs in the model, for all $j \geq i+1$ (there are some ill-defined assignments in top elements of layer i, in which both a_{i+1} and a_{i+1}' occur; these assignment will be ignored). Then, at two sets at the bottom of the layer i which correspond to the possible extensions of the assignment to $a_{i+1}, \ldots, a_n$ by setting a_i either true (effected by the set $A_{i,1}$) or to false (by $A_{i,0}$),

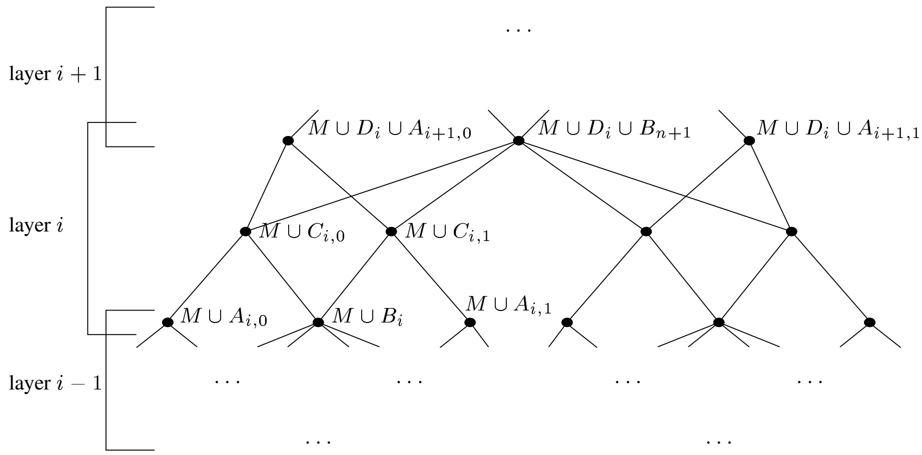

Fig. 5. Layers in $mod(\Phi_n)$

we "evaluate" the formula $Q_{i-1}a_{i-1} \cdots Q_1 a_1 \varphi(a_i, a_{i+1}, \ldots, a_n)$ where the variables $a_i, \ldots, a_n$ are fixed to the assignment. If that formula evaluates to true, then if a_i is true an atom v_i is included (resp., if a_i is false an atom v_i') at this bottom element. The quantifier Q_i is then evaluated by including in the top element "above" the two bottom sets an atom t_i if, in case of $Q_i = \exists$, either v_i or v_i' occurs in one of the two bottom elements, and in case of $Q_i = \forall$, v_i resp. v_i' occur in the bottom elements. The top element is itself a bottom element at the next layer $i+1$, and the atom t_i is used there to see whether the formula $Q_i a_i \cdots Q_1 a_1 \varphi(a_{i+1}, \ldots, a_n)$ evaluates to true.

In what follows, we formalize this intuition. We introduce a set of new atoms $At_n' = \{v_i, v_i', t_i \mid 1 \leq i \leq n\} \cup \{t_0\}$.

The following formulas are convenient for our purpose:

$$
\begin{aligned}
ass_i &= a_i \leftrightarrow \neg a_i', \quad 1 \leq i \leq n; \\
\lambda_i &= (\neg b_{i+1} \vee \neg b_{i+1}') \wedge (a_{i+1} \wedge a_{i+1}' \rightarrow \neg b_{i+1} \wedge \neg b_{i+1}'), \quad 1 \leq i \leq n; \\
\Lambda_i &= \lambda_i \wedge \neg \lambda_{i-1}, \quad 2 \leq i \leq n; \\
\Lambda_1 &= \lambda_1.
\end{aligned}
$$

Informally, ass_i tells whether the model considered assigns the atom a_i legally a truth value. The formula λ_i says that the model is at layer i or below. The formula Λ_i says that the model is at layer i. The models at the bottom of layer i which are of interest to us are those in which ass_i is true; all other models of the entire layer violate ass_i.

At layer $i \geq 1$, we evaluate the formula using the following formulas:

$$\Lambda_i \wedge ass_i \wedge t_{i-1} \wedge a_i \rightarrow v_i$$

$$\Lambda_i \wedge ass_i \wedge t_{i-1} \wedge a_i' \rightarrow v_i'$$

For $i = 1$, we add

$$\varphi \rightarrow t_0,$$

which under curbing evaluates the quantifier-free part after assigning all variables. Depending on the quantifier Q_i, we add a clause as follows. If $Q_i = \exists$, then we add

$$\Lambda_i \wedge (v_i \vee v_i') \rightarrow t_i;$$

otherwise, if $Q_i = \forall$, then we add

$$\Lambda_i \wedge v_i \wedge v_i' \rightarrow t_i.$$

For "garbage collection" of the new atoms used at lower layers, we use a formulas $trap_i$ which adds all values v_j, v_j', t_j' of lower layers to all elements of layer i which correspond to an illegal assignment to a_i:

$$trap_i = \Lambda_i \wedge \neg ass_i \rightarrow t_0 \wedge \bigwedge_{j=1}^{i-1} v_j \wedge v_j' \wedge t_j.$$

Informally, models corresponding to different extensions of an assignment will always have a mub which is upper bounded by the bottom model at layer i which is an illegal assignment.

Let the conjunction of all formulas introduced for layer i, where $1 \leq i \leq n$, be Γ_i, and let $\Gamma(F) = \bigwedge_{i=1}^{n} \Gamma_i$. Then we define

$$\Psi(F) = \Phi_n \wedge \Gamma(F).$$

Note that $\Phi(F)$ has a unique maximal model M_F, which is given by $M_F = M_n \cup \{v_i, v_i', t_i \mid 1 \leq i \leq n\}$ (i.e., all atoms are true).

Let us call a model $M \in mod(\Psi(F))$ an *assignment model*, if either $M \cap At_n = M_n$, or (b) $M \models \Lambda_i \wedge ass_i$, i.e., either M extends the maximal model of Φ_n or M is at the bottom of layer i and assigns a_i a unique truth value. In case (a), we view M at the bottom of an artificial layer $n+1$. M represents a (partial) assignment σ_M to a_i,…,a_n defined by $\sigma_M(a_j) =$ true if $a_j \in M$ and $\sigma_M(a_j) =$ false if $a_j' \in M$, for all $j = i, \ldots, n$.

We show the following

Lemma 3.1. *For each model $M \in mod(\Phi_n)$, there exists a good model $f(M)$ of $mod(\Psi(F))$, such that:*

1. *$f(M) \cap At_n = M$ (i.e., $f(M)$ coincides with M on the atoms of Φ_n);*
2. *if M is an assignment model at layer $i \in \{1, \ldots, n+1\}$, then $f(M)$ contains t_{i-1} iff the formula*

$$F_i = Q_{i-1}a_{i-1}Q_{i-2}a_2 \cdots Q_1 a_1 \varphi(a_1, \ldots, a_{i-1}, \sigma_M(a_i), \ldots, \sigma_M(a_n))$$

is true

3. *If M is at layer $i \in \{1, \ldots, n\}$ but not an assignment model, then*

$$f(M) = \begin{cases} M \cup At'_{i-1}, & \text{if } M = M_{n-1} \cup B_n; \\ f(M_{n-1} \cup A_{n,k}) \cup f(M_{n-1} \cup B_n), & \text{if } M = M_{n-1} \cup C_{n,k},\ k \in \{0,1\}. \end{cases}$$

4. *$f(M_n)$ is the unique maximal good model of $\Psi(F)$, and if $Q_n = \forall$, then $t_n \in f(M_n)$ iff $f(M_n) = At_n \cup At'_n$.*

An example of the construction of $f(\cdot)$ for the formula $F = \forall a_2 \exists a_1 (a_2 \to a_1)$ is shown in Figure 6.

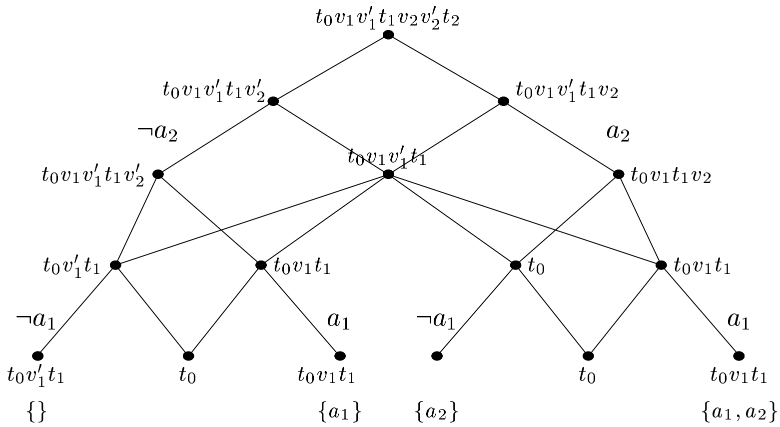

Fig. 6. Evaluating $F = \forall a_2 \exists a_1 (a_2 \to a_1)$: Extending M to $f(M) = M \cup X$ (X shown)

Proof. We first note that each model M' of $\Psi(F)$ is of the form $M \cup S$, where $M \in mod(\Phi_n)$ and $S \subseteq At'_n$, and each $M \in mod(\Phi_n)$ gives rise to at least one such M' (just add At'_n to M).

We prove the lemma showing by induction on $n \geq 0$ how to construct such a correspondence $f(M)$.

The base case $n = 0$ (in which F contains no variables and is either truth or falsity) is easy: $mod(\Phi_0) = \{\{a_0\}\}$ and, if F is truth, then $mod(\Psi(F)) = \{\{a_0, t_0\}\}$ and $f(\{a_0\}) = \{a_0, t_0\}$, and if F is falsity, then $mod(\Psi(F)) = \{\{a_0\}, \{a_0, t_0\}\}$ and $f(\{a_0\}) = \{a_0\}$.

Consider the case $n > 1$ and suppose the statement holds for $n - 1$. Let $M \in mod(\Phi_n)$. We consider two cases.

(1) $M \models \lambda_{n-1}$ and $M \not\models a_n a'_n$. Then, $M \models a_n \leftrightarrow \neg a'_n$, and either M is an assignment model at the bottom of layer n (in which case, M satisfies the left disjunct of Φ_n) or some model not at layer n (in which case M satisfies the right disjunct of M). In any case, $N = M \setminus \{a_n, a'_n, b_n, b'_n\}$ is a model of Φ_{n-1}. By the induction hypothesis, it follows that for N we have a good model $\hat{f}(N)$ of $\Psi(F')$, where $F' =$

$Q_{n-1}a_{n-1}\cdots Q_1a_1\varphi'$ and $\varphi' = \varphi[a_n/\top]$ (where $\top$ is truth) if $a_n \in M$ and $\varphi' = \varphi[a_n/\bot]$ (where $\bot$ is falsity) if $a'_n \in M$ (i.e., $a_n \notin M$), such that $\hat{f}(N)$ fulfills the items in the lemma. We define $f(M)$ as follows. If $N \subset M_{n-1}$, then $f(M) := M \cup \hat{f}(N)$; otherwise, if $N = M_{n-1}$, then $f(M) = M \cup f(N) \cup S_M$, where

$$S_M = \begin{cases} \emptyset, & \text{if } t_{i-1} \notin \hat{f}(N); \\ \{v_n, t_n\}, & \text{if } t_{i-1} \in \hat{f}(N), Q_n = \exists, \text{ and } a_i \in M; \\ \{v'_n, t_n\}, & \text{if } t_{i-1} \in \hat{f}(N), Q_n = \exists, \text{ and } a'_i \in M; \\ \{v_n\}, & \text{if } t_{i-1} \in \hat{f}(N), Q_n = \forall, \text{ and } a_i \in M; \\ \{v'_n\}, & \text{if } t_{i-1} \in \hat{f}(N), Q_n = \forall, \text{ and } a'_i \in M. \end{cases}$$

As easily checked, $f(M)$ is a model of $\Psi(F)$. Furthermore, $f(M)$ is either a minimal model of $\Psi(F)$ (if $n = 1$), or the mub of good models $f(M_1)$ and $f(M_2)$ such that $M_1, M_2 \in mod(\Phi_{n-1})$, $M_1, M_2 \subset M$, and M is a mub of M_1, M_2 in $mod(\Phi_{n-1})$. (If not, then $\hat{f}(N)$ were not a mub of $\hat{f}(N_1)$, $\hat{f}(N_2)$ in $mod(\Psi(F'))$, which is a contradiction.) We can see that $f(M)$ fulfills the items 1-3 in the lemma.

(2) $M \not\models \lambda_{n-1}$ or $M \models a_n a'_n$, i.e., M is at layer n but not an assignment model at its bottom. We consider the following possible cases for M:

(2.1) $M = M_{n-1} \cup B_n$: If $n = 1$, then M is a minimal model of Φ_n, and $f(M) = M \cup \{t_0\}$ is a minimal model of $\Psi(F)$, thus $f(M)$ is a good model of $\Psi(F)$; otherwise (i.e., $n > 2$), M is a mub of any arbitrary models M_1, $M_2 \in mod(\Phi_n)$ such that M_1 contains a_n and M_2 contains a'_n, respectively, and $M_i \setminus \{a_n, a'_n, b_n, b'_n\} \subset M_{n-1}$, for $i \in \{1, 2\}$. Since, by construction, $\hat{f}(M_i) \subseteq M_{n-1} \cup At'_{n-1} =: f(M)$, this set is an upper bound of $f(M_1)$ and $f(M_2)$ in $mod(\Psi(F))$; from formula $trap_{n-1}$ it follows that $f(M)$ is a mub of $f(M_1)$,$f(M_2)$. Thus, $f(M)$ is a good model of $\Psi(F)$.

(2.2) $M = M_{n-1} \cup C_{n,k}$, $k \in \{0, 1\}$: As easily checked, $f(M) = f(M_{n-1} \cup A_{n,k}) \cup f(M_{n-1} \cup B_n)$ (=$M_{n-1} \cup B_n \cup S_{M_{n-1} \cup A_{n,k}}$) is a model of $\Psi(F)$. Since, as already shown, both $f(M_{n-1} \cup A_{n,k})$ and $f(M_{n-1} \cup B_n)$ are good models of $\Psi(F)$, clearly $f(M)$ is a mub of them and thus a good model of $\Psi(F)$.

(2.3) $M = M_n$: We define

$$f(M) = f(M_{n-1} \cup C_{n,0}) \cup f(M_{n-1} \cup C_{n,1}) \cup \begin{cases} \{t_n\}, & \text{if } Q_n = \forall \text{ and } v_n, v'_n \in X; \\ \emptyset, & \text{otherwise.} \end{cases}$$

Observe that $f(M) = M_n \cup At'_{n-1} \cup X$, where $X \subseteq \{v_n, v'_n, t_n\}$. Then, as easily checked, $f(M)$ is a model of $\Psi(F)$. Clearly, $f(M)$ is a mub of $f(M_{n-1} \cup C_{n,0})$ and $f(M_{n-1} \cup C_{n,1})$, and thus, $f(M)$ is a good model of $\Psi(F)$.

We now show that $f(M)$ in (2.1)–(2.3) satisfies items 1-3 in the lemma. Obviously, this is true for (2.1) and (2.2). For the case (2.3), from the definitions of $f(\cdot)$ in (1) and (2.1)–(2.2) it follows that $t_n \in f(M)$ if and only if $t_{n-1} \in f(M_{n-1} \cup A_{n,k})$ holds for for some $k \in \{0, 1\}$ if $Q_n = \exists$ and for both $k \in \{0, 1\}$ if $Q_n = \forall$. By the induction hypothesis, $t_{n-1} \in f(M_{n-1} \cup A_{n,k})$ is true iff the QBF $Q_{n-1}a_{n-1}\cdots Q_1a_1\varphi'$, where $\varphi' = \varphi[a_n/\top]$ if $k = 1$ and $\varphi' = \varphi[a_n/\bot]$ if $k = 0$, is true. Thus, $t_n \in f(M)$ iff the QBF F is true. Hence, $f(M)$ satisfies items 1-3 of the lemma.

As for property 4, Furthermore, in the case where $Q_n = \forall$, we have by definition of $f(M)$ that $t_n \in f(M)$ iff $f(M) = M_n \cup At'_{n-1} \cup \{v_n, v'_n, t_n\} = At_n \cup At'_n$.

Finally, it remains to show that $f(M_n)$ is the unique maximal good model of $\Psi(F)$. As easily seen, every finite propositional theory which has a unique maximal model has a unique maximal good model, thus $\Psi(F)$ has a unique maximal good model M'. From the induction hypothesis, it follows that $M_k = f(M_{n-1} \cup A_{n,k})$ is the unique maximal good model M'_k of $\Psi(F)$ such that $M' \cap At_n \subseteq M_{n-1} \cup A_{n,k}$, for $k \in \{0, 1\}$. Since $M_2 = f(M_{n-1} B_n)$ is the unique maximal good model N of $\Psi(F)$ such that $N \cap At_n \subseteq M_{n-1} \cup B_n$, we conclude from the structure of layer n, which has the lub property (see Section 4), that M' is a mub of M_0, M_1, M_2. Since, by construction, $f(M)$ is an upper bound of M_1, M_2, M_3, it follows $M' = f(M)$.

This proves that the claimed statement holds for n, and completes the induction. □

We thus obtain the following result.

Theorem 3.1. *1. Given a propositional formula G and a model M of G, deciding whether M is a good model of G is* PSPACE-*hard.*
2. *Given a propositional formula G and an atom p, deciding whether $Curb(G) \models p$ is* PSPACE-*hard.*

Proof. By items 2 and 4 in Lemma 3.1, $M = At_n \cup At'_n$ is a good model of $\Psi(F)$ for a QBF $F = \forall a_n Q_{n-1} a_{n-1} \cdots Q_1 a_1 \varphi$ iff F is true. Furthermore, F is false if and only if no good model of $\Psi(F)$ contains t_n. Deciding whether any given QBF of this form is true (resp. false) is clearly PSPACE-hard, and the formula $\Psi(F)$ is easily constructed in polynomial time from F. This proves the result. □

Combined with the previous results [8] that Curb Inference and Curb Model Checking are in PSPACE, we obtain the main result of this section.

Theorem 3.2. *1.* Curb Model Checking*, i.e., given a propositional theory T and sets $\mathbf{p}$, $\mathbf{z}$ of propositional letters, deciding whether M is a $\mathbf{p};\mathbf{z}$-good model of T is* PSPACE-*complete.*
2. Curb-Inference*, i.e., given a propositional theory T, sets $\mathbf{p};\mathbf{z}$ of propositional letters, and a propositional formula G, deciding whether $Curb(T; \mathbf{p}, \mathbf{z}) \models G$ is* PSPACE-*complete.*

4 The Lub Property

While curbing of general theories is PSPACE-complete, it is possible to identify specific classes of theories on which curbing has lower complexity. In this section, we identify a relevant fragment of propositional logic for which curb-inference is in Π_2^P.

Definition 4.1. *A theory T has the lub property iff every nonempty set $\mathcal{S}$ of good models has a least upper bound (lub) M.*

Lemma 4.1. *Let $\mathcal{S}_1, \mathcal{S}_2$ be nonempty sets of good models of theory T such that $\mathcal{S}_1 \subseteq \mathcal{S}_2$, and let M_1, M_2 be mubs of $\mathcal{S}_1$ and $\mathcal{S}_2$, respectively. If M_1 is the lub of $\mathcal{S}_1$, then $M_1 \leq M_2$.*

Theorem 4.1. *If theory T has the lub property, then a model is good iff it is 1-good.*

Proof. Prove by induction on α that if model M is α-good, then it is 1-good. Obvious for $\alpha \leq 1$. Assume $\alpha > 1$. Then, M is a mub of $\mathcal{S} = \{M' : (< \alpha)\text{-good}(M'), M' \leq M\}$. Now, by the hypothesis, each $M' \in \mathcal{S}$ is the mub of some $\mathcal{S}' \subseteq \mathcal{S}$ which contains only minimal models. Let $\mathcal{S}_m$ be the minimal models from $\mathcal{S}$. If $\mathcal{S}_m = \emptyset$, then M is a minimal model and the statement holds. Else $\mathcal{S}_m$ has a lub M_m. From the unique mub property and Lemma 4.1, it follows that $M' \leq M_m$ for each $M' \in \mathcal{S}$. Thus M_m is an upper bound of $\mathcal{S}$, hence $M \leq M_m$. On the other hand, since $\mathcal{S}_m \subseteq \mathcal{S}$, it follows from Lemma 4.1 that $M_m \leq M$. Since $\leq$ is a partial order, it follows $M_m = M$. Thus M is 1-good and the statement holds. $\square$

Corollary 4.1. *For propositional theories T having the lub property,* Curb Inference *is in Π_2^P, and* Curb Model Checking *is in Σ_2^P.*

Proof. To show $Curb(T) \not\models F$, guess a model M of $Curb(T)$ such that $M \not\models F$. To verify M, guess k from $\{0, \ldots, |V|\}$, where V is the variable set, and minimal models $M_1, \ldots, M_k$ of T such that M is a mub of them. Use an NP oracle for testing whether M_i is minimal (is in coNP) and for testing if M is a mub of the M_i (is in coNP). $\square$

Notice the following characterization of lub theories.

Definition 4.2. *A theory T is mub-compact over a domain iff every good model is a mub of a finite set of good models.*

Theorem 4.2. *Let T be a mub-compact theory over some domain. Then T has the lub property iff every pair of good models has a lub.*

Proof. (Sketch) To show the *if* direction, demonstrate by induction on finite cardinality κ that every set $\mathcal{S}$ such that $\|\mathcal{S}\| \leq \kappa$ has a lub. For $\kappa \leq 2$, this is obvious. For $\kappa > 2$, let $M \in \mathcal{S}$ be a maximal element in $\mathcal{S}$. By the hypothesis, $\mathcal{S} - \{M\}$ has a lub M'. M and M' have a lub M'', which must (Lemma 4.1) be the lub of $\mathcal{S}$. $\square$

Corollary 4.2. *If the domain is finite and the models of T form an upper semi-lattice, then T has the lub property and a model is* good *iff it is* 1-good*.*

As already mentioned in the introduction, Scarcello, Leone, and Palopoli [21] derived complexity results for curbing Krom theories, i.e., clausal theories where each clause contains at most two literals. They showed that Curb Model Checking for propositional Krom theories is in Σ_2^P. To establish this result, they showed that the *union* of any pair of good models of a propositional Krom theory is also a good model. From this it clearly follows that propositional Krom theories enjoy the (more general) lub property. Hence their Σ_2^P upper bound and, in addition, a Π_2^P upper bound for curb inferencing can also be derived via our more general results.

5 Good Models and Least Upper Bounds

The lub property defined in Section 4 requires that *all* nonempty collections of good models of a theory have a lub. Let us weaken this property by requiring merely that for every non-minimal good model M there exists a collection of models whose lub is M.

Definition 5.1. *A theor T has the weak least upper bound (weak lub) property, if every non-minimal good model of T is the lub of some collection $\mathcal{M}$ of good models of T.*

Notice that the lub property implies the weak lub property, but not vice versa. This is shown by the following example.

Example 5.1. Suppose the models of a propositional theory T are the ones shown in Figure 7. All models are good, and $M_1 = \{a, b, c\}$, $M_2 = \{b, c, d\}$ are the lubs of

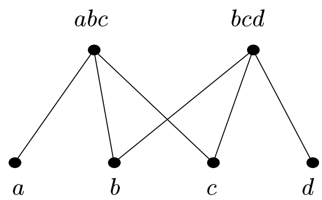

Fig. 7. The weak lub property does not imply the lub property

the collections $\{\{a\}, \{b\}, \{c\}\}$ and $\{\{b\}, \{c\}, \{d\}\}$, respectively. However, the good models $\{b\}$ and $\{c\}$ do not have a lub; thus, the theory satisfies the weak lub property but not the lub property.

Intuitively, if a theory satisfies the weak lub property, then any good model M in a collection $\mathcal{M}$ of good models can be replaced by a collection $\mathcal{M}'$ of good models whose lub is M, without affecting the mubs of the collection, i.e., $\mathcal{M}$ has the same mubs as $\mathcal{M} \setminus \{M\} \cup \mathcal{M}'$. By repeating this replacement, $\mathcal{M}$ can be replaced by a collection $\mathcal{M}^*$ of minimal models that has the same mubs as $\mathcal{M}$. This is actually the case, provided that the collection of good models has the following property.

Definition 5.2. *The collection of good models of a sentence φ is well-founded if every decreasing chain $M_0 \supseteq M_1 \supseteq \cdots$ of good models has a smallest element.*

Notice that in the context of circumscription, theories were sometimes called well-founded if every model M of a sentence φ includes a minimal model of φ [14]. That notion of well-foundedness is different from the one employed here.

The collection of good models of a theory is not necessarily well-founded, as shown by the following example.

Example 5.2. Consider the theory T on the domain $\mathbb{Z}$ of all integers:

$$\varphi = (\forall x)(p(x) \longleftrightarrow x < 0)) \;\vee\; (\exists x \geq 0)(\forall y)(p(y) \longleftrightarrow \neg(1 \leq y \leq x)) \;\vee$$
$$(\exists x \geq 0)(\forall y)(p(y) \longleftrightarrow (y > x) \vee (-x \leq y \leq 0))$$

Informally, T says that the numbers having property p are either all the negative numbers ($\mathbb{Z}^- = \{-1, -2, \ldots\}$), all numbers except some interval $[1, 2, \ldots, k]$, $k \geq 0$, or all nonnegative numbers where the interval $[0, k]$, $k \geq 0$, is replaced by the interval $[-k, -0]$. All models of T are good. The minimal models are $\mathbb{Z}^-$ and $N_k = (N_0 \setminus [0, k]) \cup [-k, -0]$, $k \geq 0$; every model $M_k = \mathbb{Z} \setminus [1, k]$, $k \geq 0$, is a mub of the models $\mathbb{Z}^-$ and N_k (see Figure 8). Clearly, $M_0 \supseteq M_1 \supseteq \cdots \supseteq M_i \supseteq \cdots$, $i \in \omega$, forms

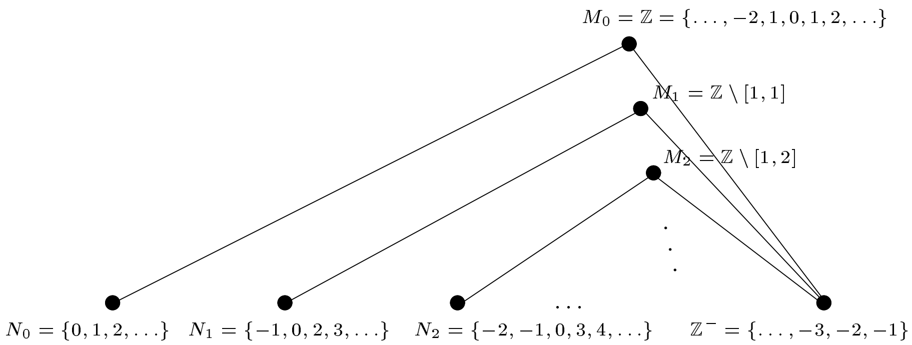

Fig. 8. A collection of good models that is not well-founded.

a decreasing chain of good models. This chain has no smallest element, and hence the collection of good models of T is not well-founded. □

Theorem 5.1. *Let φ be a first-order sentence such that the collection of good models of φ is well-founded. If φ hast the weak lub property, then every good model is either minimal or the lub of some collection of minimal models.*

Proof. We show this by contradiction. Assume the contrary holds. Let $\mathcal{B}$ be the set of good models which are not the lub of some collection of minimal models; note that $\mathcal{B}$ is not empty. Since the collection of good models is well-founded, $\mathcal{B}$ must have a minimal element M. (To obtain such an M, construct a maximal chain in $\mathcal{B}$, and take the unique minimal element of this chain, which must exist). Since φ has the weak lub property, M is the mub of some collection $\mathcal{S}$ of good models. The definition of $\mathcal{B}$ and the weak lub property of φ imply that every $M' \in S$ is the lub of a collection $\mathcal{S}_{M'}$ of minimal models. Let $\mathcal{S}'$ be the union of all these $\mathcal{S}_{M'}$. We show that M is the lub of $\mathcal{S}'$. Clearly, M is an upper bound of $\mathcal{S}'$. Assume then that M is not a minimal. Then there exists a good model $M' < M$ which is an upper bound of $\mathcal{S}'$. But this M' is also an upper bound of $\mathcal{S}$. This means that M is not a mub of $\mathcal{S}$, which is a contradiction. It follows that M is a mub of $\mathcal{S}'$. On the other hand, every upper bound M' of $\mathcal{S}'$ must satisfy $M \leq M'$.

Therefore, M is the unique mub of $\mathcal{S}'$. Consequently, M is the unique minimal upper bound of a collection of minimal models. By definition, this means $M \notin \mathcal{B}$. This is a (global) contradiction. □

The converse of this theorem (which is equivalent to the statement that a theory, if every good model is either minimal or the lub of some collection of minimal models, is well-founded) is not true. This is shown by Example 5.2. Furthermore, this theorem does not hold if the collection of good models is arbitrary. This is shown by the following example.

Example 5.3. Replace in Example 5.2 every model M_i, $i \in \omega$, by the two models $M_i^a = M_i \cup \{a\}$ and $M_i^b = M_i \cup \{b\}$ and extend the domain with the new elements a and b.

In the resulting collection of models, which is clearly axiomatizable by a first-order sentence φ, every model is good and the lub of some collection of good models (M_i^a is the lub of $\{N_i, M_{i+1}^a\}$, and M_i^b of $\{N_i, M_{i+1}^b\}$; all other models are minimal). However, no M_i^a is the lub of a collection of minimal models. Notice that each good model is the lub of two good models and 1-good. □

From Theorems 5.1 and 2.1, we immediately get the following complexity results for propositional theories.

Theorem 5.2. *For propositional theories which enjoy the weak lub property, the problem* Curb Model Checking *is in Σ_2^P, while the problem* Curb Inference *is in Π_2^P.*

A possible attempt to strengthen the weak-lub property is to use ordinals. Say that the collection of good models of a theory has the *inductive weak-lub property*, if every non-minimal α-good model is the lub of a collection of $(< \alpha)$-good models. Notice that collection of good models in Example 5.2 has the inductive weak-lub property (which, as a consequence, does not imply well-foundedness). However, the following result is an easy consequence of our results from above.

Theorem 5.3. *Let φ be a first-order sentence whose collection of good models is well-founded. Then, it has the inductive weak-lub property if and only if it has the weak-lub property.*

Proof. The *only if* direction is trivial. The *if* direction follows from Theorem 5.1. □

Acknowledgments. This work was supported by the Austrian Science Fund (FWF) Project N. Z29-INF.

References

1. bddlib. http://www.cs.cmu.edu/ modelcheck/bdd.html.
2. B. Bodenstorfer. How Many Minimal Upper Bounds of Minimal Upper Bounds. *Computing*, 56:171–178, 1996.

3. M. Cadoli. The Complexity of Model Checking for Circumscriptive Formulae. *Information Processing Letters*, 44:113–118, 1992.
4. M. Cadoli, A. Giovanardi, and M. Schaerf. An Algorithm to Evaluate Quantified Boolean Formulae. In *Proc. AAAI/IAAI-98*, pp. 262–267, 1998.
5. E. Chan. A Possible Worlds Semantics for Disjunctive Databases. *IEEE Trans. Knowledge and Data Engineering*, 5(2):282–292, 1993.
6. H. Decker and J. C. Casamayor. Sustained Models and Sustained Answers in First-Order Databases. In A. Olivé, editor, *Proc. 4th International Workshop on the Deductive Approach to Information Systems and Databases (DAISD 1993), 1993, Lloret de Mar, Catalonia*, pp. 267–286, Report de recerca, LSI/93-25-R, Departament de Llenguatges i Sistemes Informatics, Universitat Politecnica de Catalunya, 1990.
7. T. Eiter and G. Gottlob. Propositional Circumscription and Extended Closed World Reasoning are $\mathbf{\Pi}_2^P$-complete. *Theoretical Computer Science*, 114(2):231–245, 1993. Addendum 118:315.
8. T. Eiter, G. Gottlob, and Y. Gurevich. Curb Your Theory! A circumscriptive approach for inclusive interpretation of disjunctive information. In R. Bajcsy, editor, *Proc. 13th Intl. Joint Conference on Artificial Intelligence (IJCAI-93)*, pp. 634–639. Morgan Kaufmann, 1993.
9. R. Feldmann, B. Monien, and S. Schamberger. A Distributed Algorithm to Evaluate Quantified Boolean Formulae. In *Proc. National Conference on AI (AAAI'00)*, Austin, Texas, 2000. AAAI Press.
10. A. Flögel, M. Karpinski, and H. Kleine Büning. Resolution for Quantified Boolean Formulas. *Information and Computation*, 117:12–18, 1995.
11. P. Liberatore. The Complexity of Iterated Belief Revision. In F. Afrati and P. Kolaitis, editors, *Proc. 6th Intl. Conference on Database Theory (ICDT-97)*, LNCS 1186, pp. 276–290, Springer, 1997.
12. P. Liberatore. The Complexity of Belief Update. *Artificial Intelligence*, 119(1-2):141–190, 2000.
13. V. Lifschitz. Computing Circumscription. In *Proc. International Joint Conference on Artificial Intelligence (IJCAI-85)*, pp. 121–127, 1985.
14. V. Lifschitz. Circumscription. In D. Gabbay, C. Hogger, and J. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume III, pages 297–352. Clarendon Press, Oxford, 1994.
15. J. McCarthy. Circumscription – A Form of Non-Monotonic Reasoning. *Artificial Intelligence*, 13:27–39, 1980.
16. J. Rintanen. Improvements to the Evaluation of Quantified Boolean Formulae. In *Proc. IJCAI '99*, pp. 1192–1197. AAAI Press, 1999.
17. K. Ross. The Well-Founded Semantics for Disjunctive Logic Programs. In W. Kim, J.-M. Nicholas, and S. Nishio, editors, *Proc. First Intl. Conf. on Deductive and Object-Oriented Databases (DOOD-89)*, pp. 352–369. Elsevier Science Pub., 1990.
18. K. Ross and R. Topor. Inferring Negative Information From Disjunctive Databases. *Journal of Automated Reasoning*, 4(2):397–424, 1988.
19. C. Sakama. Possible Model Semantics for Disjunctive Databases. In W. Kim, J.-M. Nicholas, and S. Nishio, editors, *Proc. First Intl. Conf. on Deductive and Object-Oriented Databases (DOOD-89)*, pp. 337–351. Elsevier Science Pub., 1990.
20. C. Sakama and K. Inoue. Negation in Disjunctive Logic Programs. In *Proc. ICLP-93*, Budapest, Hungary, June 1993. MIT-Press.
21. F. Scarcello, N. Leone, and L. Palopoli. Curbing Theories: Fixpoint Semantics and Complexity Issues. In M. Alpuente and M. I. Sessa, editors, *Proc. 1995 Joint Conference on Declarative Programming (GULP-PRODE'95)*, pp. 545–554. Palladio Press, 1995.
22. J. van Benthem and K. Doets. Higher Order Logic. In D. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic, Vol.I*, chapter I.4, pp. 275–329. D. Reidel Pub., 1983.