Checking the Size of Circumscribed Formulae

Paolo Liberatore
DIAG
Sapienza University of Rome
Rome, Italy

***Abstract*—The circumscription of a propositional formula T may not be representable in polynomial space, unless the polynomial hierarchy collapses. This depends on the specific formula T, as some can be circumscribed in little space and others cannot. The problem considered in this article is whether this happens for a given formula or not. In particular, the complexity of deciding whether $CIRC(T)$ is equivalent to a formula of size bounded by k is studied. This theoretical question is relevant as circumscription has applications in temporal logics, diagnosis, default logic and belief revision.**

***Keywords*—Circumscription; computational complexity; belief revision.**

I. INTRODUCTION

The circumscriptive reasoning mechanism requires a set of variables to be minimized [1], [2], that is, set to the logical value false whenever possible. Similarly to the closed world assumption [3], it formalizes the assumption that lack of information on certain conditions can be considered evidence that they do not hold. Applications include temporal domains [4], [5], diagnosis [6], induction [7] and belief revision [8]. Contrary to the basic closed world assumption, circumscription takes into account all possible ways variables can be set to false; for example, $x \vee y$ is consistent with either $\neg x$ and $\neg y$ but not both, leading to the two possible cases $(x \vee y) \wedge \neg x$ and $(x \vee y) \wedge \neg y$. These may be up to 2^n, if the number of variables is n: a trivial representation of the circumscribed formula may be exponential. However, it may be equivalent to a smaller formula.

Expressing propositional circumscription as a formula of size bounded by a polynomial has been proved not possible in general [9], unless the polynomial hierarchy collapses [10], a condition generally deemed unlikely. As a result, the problem of whether propositional circumscription can be represented in space bounded by some number k has not an obvious answer: it is possible in some cases but not in others. The problem considered in this article is whether this is possible; in particular, the complexity of this problem is studied. This is similar to the problem of minimizing propositional formulae: given a formula F, is there an equivalent formula of size bounded by k [11]? For circumscription, the question is whether the circumscription of a formula is equivalent to some formula of size bounded by k. For example, the circumscription of $x \vee y$ accounts for both $(x \vee y) \wedge \neg x$ and $(x \vee y) \wedge \neg y$ to be possible; therefore, the result is the formula $((x \vee y) \wedge \neg x) \vee ((x \vee y) \wedge \neg y)$. However, this formula is equivalent to $(x \wedge \neg y) \vee (\neg x \wedge y)$. By the standard metric of formulae where size is defined as the number of variable occurrences, this formula has size 4. Therefore, the circumscription of $x \vee y$ is equivalent to a formula of size bounded by $k = 4$, but not for example $k = 1$ as no formula of a single variable is equivalent to $(x \wedge \neg y) \vee (\neg x \wedge y)$. The answer is not this easy when the formula is more complex than $x \vee y$. Indeed, it will be proved that the problem is hard for the complexity class Π^p_2, that is, harder than problems such as propositional satisfiability, vertex cover and Hamiltonian cycle [10].

The question of the size of the representation has an implementation impact. Indeed, verifying which conditions hold under the circumscription assumption amounts to $CIRC(T) \models C$, where T represents the current information and C the condition to check, and this is an hard problem [12], [13], [14]. However, if $CIRC(T)$ can be represented by a formula F of bounded size, the problem can be solved by first finding F and then solving the easier (coNP) problem $F \models C$. Once F is determined, any number of other conditions $C_1, C_2, \ldots$ can then be checked against F at the same cost.

Since circumscription is also used as the target of translation of several belief revision operators, the question concerns the dynamic of logic. Indeed, changing a formula to accommodate for new information is generally expected to produce a result of bounded size.

The article is organized as follows: the next section contains the formal definition of circumscription and the notations used in this article, plus two preliminary lemmas; in the section afterwards, the complexity of the problem of whether the circumscription of a formula can be represented in size bounded by some number is studied; the final section comments the practical implications of this analysis and its open problems.

II. PRELIMINARY RESULTS

Propositional formulae are denoted by the capital letters T and F, and are always assumed to be in Negation Normal Form (NNF). Sets of variables are denoted by X, Y and Z. Notation $X^{\neg}$ indicates the set $\{\neg x \mid x \in X\}$. The shorthand $x \not\equiv y$ indicates $(x \wedge \neg y) \vee (\neg x \wedge y)$.

Models are denoted by ω_X, where the suffix X indicates the set of variables: ω_X is a truth evaluation of the variables X, ω_Y is a truth evaluation of the variables Y, etc. Models are identified by the sets of variables they assign to true; this allows to write $\omega_X \subseteq \omega'_X$ to mean that ω'_X assigns true to all variables ω_X assigns true, but not necessarily the converse. The model assigning true to all variables X is denoted ω^+_X, the one assigning false to all ω^-_X.

The following notation is used to denote a formula that represents a single model: $Form(\omega_X) = \bigwedge\{x \mid \omega_X \models x\} \cup$

$\{\neg x \mid \omega_X \models \neg x\}$. If F is a formula over variables $X \cup Y$ and ω_X a truth evaluation over X, the notation $F|_{\omega_X}$ indicates the formula obtained by replacing each variable X in F with its truth value according to ω_X.

In this article, circumscription is defined over propositional logic, and restricted to the case where all variables are minimized. This gives rise to the following definition.

Definition 1: Given a formula T over variables X, its circumscription $CIRC(T)$ is defined as follows, where $X^{\neg} = \{\neg x \mid x \in X\}$.

$$CIRC(T) = \bigvee \left\{ T \wedge S \;\middle|\; \begin{array}{l} S \subseteq X^{\neg} \\ T \wedge S \not\models \bot \\ \forall S' \subseteq X^{\neg} \\ S \subseteq S' \Rightarrow T \wedge S' \models \bot \end{array} \right\}$$

Some formulae T have small circumscription. For example, $T = \bigwedge X$ has a circumscription equal to itself, since $S = \emptyset$ is the only subset of $X^{\neg}$ satisfying the definition. Some other formulae have larger circumscription, such as $T = \bigvee X$; indeed, for this formula $S = X^{\neg} \backslash \{x\}$ satisfies the definition for every $x \in X$. Some formulae do not even have polynomial-size equivalent representations of their circumscription [9].

Circumscription is simple to compute on formulae that imply either x, $\neg x$, or $x \not\equiv x'$ for some variables x and x':

Property 1: The following equivalences hold:

$$\begin{array}{rcl} CIRC(T \wedge x) & = & x \wedge CIRC(T|_{\omega^+_{\{x\}}}) \\ CIRC(T \wedge \neg x) & = & \neg x \wedge CIRC(T|_{\omega^-_{\{x\}}}) \\ CIRC(T \wedge (x \not\equiv x')) & = & \begin{array}{l} (x \not\equiv x') \wedge \\ (CIRC(T|_{\omega^-_{\{x\}} \omega^+_{\{x'\}}}) \vee \\ \;CIRC(T|_{\omega^+_{\{x\}} \omega^-_{\{x'\}}})) \end{array} \end{array}$$

These are well-known properties. The third equivalence allows evaluating $CIRC(T)$ separately for x true and x false, if T does not contain x'.

The size of formulae is defined by the following metrics.

Definition 2: The size of a formula F, denoted $||F||$, is the number of variable occurrences in F.

For example, the size of $(a \wedge \neg b) \vee c \vee \neg(\neg a)$ is four, since the variable a occurs twice in it and b and c once each. According to this definition, the size of a formula and of its NNF form obtained by applying the De Morgan rules coincide. A bound on the size of a formula derives from its models.

Lemma 1: If a NNF formula F has a model that satisfies a literal l but not the modified model where the value of l is inverted, then F contains l.

Proof: Let F be a formula and ω_X its model satisfying l. Let us assume, on the converse, that F does not mention the literal l. Since F is in NNF, no part of it is turned to false by changing the value of l from true to false. As a result, the model ω'_X obtained by changing the value of l in ω_X satisfies F, contradicting the assumption of the lemma. □

As a consequence, if a formula is satisfied by a model where x is true but not by the same model where x is false, and vice versa, then any formula equivalent to it contains both x and $\neg x$. Therefore, if a formula contains $x \not\equiv y$, either conjoined with a satisfiable formula not containing x and y or disjoined with a non-valid formula not containing x and y, then it must contain at least two literal occurrences for x and two for y. The following lemma shows a sufficient condition for the presence of a literal in a formula.

Lemma 2: Let F be a formula over $X \cup Y$. For any truth evaluation ω_X, no formula equivalent to F is smaller than the smallest formula equivalent to $F|_{\omega_X}$.

Proof: Let T be a formula equivalent to F. Equivalence is preserved when replacing a variable with a truth value in both formulae. As a result, $F|_{\omega_X} \equiv T|_{\omega_X}$. Furthermore, such a replacement does not increase the number of literal occurrences in T, since it only replace some variables with either true or false. As a result, the size of $T|_{\omega_X}$ is less than or equal to the size of T. Since $T|_{\omega_X}$ is a formula equivalent to $F|_{\omega_X}$, it is at least as large as the smallest formula equivalent to $F|_{\omega_X}$. Since T is larger or has the same size, the claim is proved. □

This lemma is useful when formulae contain parts that are satisfiable only for a specific truth evaluation of some variables X. Such formulae are built to the aim of generating a (relatively) large subformula whenever a condition is met.

III. The Size of Circumscriptive Formulae

In this section, we analyze the problem of deciding whether the circumscription of a formula can be represented by a formula of size bounded by an integer k, in unary notation. The unary notation is used to avoid exponentially-sized formulae to be taken into account. Equivalently, the problem could be reformulated as: is there any formula that is equivalent to $CIRC(T)$ and has size less or equal than another formula G?

Theorem 1: The problem of deciding whether $CIRC(T)$ is equivalent to a formula F with $||F|| \leq k$, where k is a number in unary notation, is in Σ^p_3.

Proof: The problem can be reformulated as follows: check whether there exists a formula F that is equivalent to $CIRC(T)$ and $||F|| \leq k$. The problem $F \models CIRC(T)$ is in coNP, since it amounts to check whether $\omega \not\subset \omega'$ for every $\omega \models T$ and $\omega' \models F$. Since coNP is a subclass of Π^p_2, this problem is also in Π^p_2. The problem $CIRC(T) \models F$ is instead Π^p_2-complete [12], [13], [14]; therefore, it is in Π^p_2. The problem under consideration can be therefore solved by guessing a formula F of size bounded by k and then checking whether $F \models CIRC(T)$ and $CIRC(T) \models F$. Since both problems are in Π^p_2, they can be checked by reversing the result of a Σ^p_2 oracle. The problem can therefore be solved by a first nondeterministic step generating all formulae F with $||F|| \leq k$ and then by calling the oracle. It is therefore in Σ^p_3. □

The problem can be proved hard for the class Π^p_2.

Theorem 2: The problem of deciding whether $CIRC(T)$ is equivalent to a formula T' with $||T'|| \leq k$ is Π^p_2-hard.

Proof: Let F be a formula over variables $X \cup Y$. The proof shows how to build in polynomial time a formula T and a number k in unary notation such that $\forall X \exists Y.F$ is valid if and only if $CIRC(T)$ is equivalent to a formula of size $\leq k$.

Let us assume, without loss of generality, that $|X| = |Y| = n$. The reduction introduces a set of new variables X' in in one-on-one correspondence with X. It also introduces a set of new variables Y' in correspondence with Y and a set of new variables Z of cardinality $m = 3n + ||F|| + 1$.

In this proof the following notations are used, where X and X' are sets of variables in one-to-one correspondence and each x corresponds to $x' = c(x)$:

$$\begin{aligned} X^{\neg} &= \{\neg x \mid x \in X\} \\ X \equiv X' &= \bigwedge\{x \equiv x' \mid x \in X,\ x' = c(x)\} \\ X \not\equiv X' &= \bigwedge\{x \not\equiv x' \mid x \in X\ , x' = c(x)\} \end{aligned}$$

Formula T and number k are as follows.

$$\begin{aligned} T &= (X \not\equiv X') \wedge \\ &\quad \Big(((Z \not\equiv Z') \wedge \bigwedge Y \wedge \bigwedge Y') \vee \\ &\quad (F \wedge (Y \not\equiv Y') \wedge \bigwedge Z^{\neg} \wedge \bigwedge Z'^{\neg})\Big) \\ k &= 14n + 3||F|| + 2 \end{aligned}$$

The reduction works as follows: $X \not\equiv X'$ allows expressing $CIRC(T)$ in terms of the disjunction of $CIRC(T|_{\omega_X})$ for all possible ω_X; if $\forall X \exists Y.F$ is true, all these formulae $CIRC(T|_{\omega_X})$ can be expressed in the same way, so that a single formula equivalent to $CIRC(T)$ exists with size bounded by k; otherwise, for the evaluation ω_X that makes F false $CIRC(T|_{\omega_X})$ alone has size greater than k.

The first step employs the third equivalence of Property 1, when applied to every $x \in X$ and its respective $x' \in X'$, since T contains $X \not\equiv X'$:

$$CIRC(T) \equiv \bigvee_{\omega_X} Form(\omega_X) CIRC(T|_{\omega_X})$$

The second step of the proof is to analyze $CIRC(T|_{\omega_X})$ for an evaluation ω_X. Formula $T|_{\omega_X}$ can be rewritten as follows.

$$\begin{aligned} T|_{\omega_X} &\equiv \Big((X \not\equiv X') \wedge \\ &\quad (((Z \not\equiv Z') \wedge \bigwedge Y \wedge \bigwedge Y') \vee \\ &\quad (F \wedge (Y \not\equiv Y') \wedge \bigwedge Z^{\neg} \wedge \bigwedge Z'^{\neg}))\Big)|_{\omega_X} \\ &\equiv (X \not\equiv X')|_{\omega_X} \wedge \\ &\quad (((Z \not\equiv Z') \wedge \bigwedge Y \wedge \bigwedge Y')|_{\omega_X} \vee \\ &\quad (F \wedge (Y \not\equiv Y') \wedge \bigwedge Z^{\neg} \wedge \bigwedge Z'^{\neg})|_{\omega_X})) \\ &\equiv Form(\omega_{X'}) \wedge \\ &\quad (((Z \not\equiv Z') \wedge \bigwedge Y \wedge \bigwedge Y') \vee \\ &\quad (F|_{\omega_X} \wedge (Y \not\equiv Y') \wedge \bigwedge Z^{\neg} \wedge \bigwedge Z'^{\neg}))) \end{aligned}$$

In this last formula, $\omega_{X'}$ is the evaluation of X' setting each variable in X' to the opposite value of the corresponding variable in X. This formula does not contain any variable in X. Therefore, $CIRC(T|_{\omega_X})$ is defined by taking into account only the other variables: X', Y, Y' and Z. Since X' has a fixed value, it holds:

$$\begin{aligned} &CIRC(T|_{\omega_X}) \equiv \\ &\quad Form(\omega_{X'}) \wedge \\ &\quad CIRC(((Z \not\equiv Z') \wedge \bigwedge Y \wedge \bigwedge Y') \vee \\ &\quad (F|_{\omega_X} \wedge (Y \not\equiv Y') \wedge \bigwedge Z^{\neg} \wedge \bigwedge Z'^{\neg})) \end{aligned}$$

The first subformula of circumscription $((Z \not\equiv Z') \wedge \bigwedge Y \wedge \bigwedge Y')$ has only models $\omega_Y^+ \cup \omega_{Y'}^+ \cup \omega_Z \cup \omega_{Z'}$ in which ω_Y^+ and $\omega_{Y'}^+$ set all variables in Y and Y' to true. This model contains a model of the second subformula if $F|_{\omega_X}$ is satisfiable. Indeed, let ω_Y be the model that satisfies $F|_{\omega_X}$. This model is contained in ω_Y^+. The model $\omega_{Y'}$ that assigns $y' \in Y'$ to true if and only if the corresponding $y \in Y$ is false in ω_Y also satisfies $(F|_{\omega_X} \wedge (Y \not\equiv Y') \wedge \bigwedge Z^{\neg} \wedge \bigwedge Z'^{\neg})$, and is contained in $\omega_{Y'}^+$. A model of the second subformula is therefore $\omega_Y \cup \omega_{Y'} \cup \omega_Z^- \cup \omega_{Z'}^-$, where $\omega_Z^- \cup \omega_{Z'}^-$ set all variables to false and are therefore contained in $\omega_Z \cup \omega_{Z'}$.

This proves that every model of the first subformula contains a model of the second, if $F|_{\omega_X}$ is satisfiable. If this is the case, the first subformula is irrelevant to circumscription. Otherwise, the second subformula is unsatisfiable.

$$\begin{aligned} &CIRC\begin{pmatrix} ((Z \not\equiv Z') \wedge \bigwedge Y \wedge \bigwedge Y') \vee \\ (F|_{\omega_X} \wedge (Y \not\equiv Y') \wedge \bigwedge Z^{\neg} \wedge \bigwedge Z'^{\neg}) \end{pmatrix} \\ &\equiv CIRC(F|_{\omega_X} \wedge (Y \not\equiv Y') \wedge \bigwedge Z^{\neg} \wedge \bigwedge Z'^{\neg}) \\ &\quad \text{if } F|_{\omega_X} \text{ is satisfiable} \\ &\equiv CIRC((Z \not\equiv Z') \wedge \bigwedge Y \wedge \bigwedge Y') \text{ otherwise} \end{aligned}$$

The rest of the proof depends on whether F is satisfiable for every ω_X. If it is, then $CIRC(T|_{\omega_X})$ is equivalent to $\omega_{X'} \wedge CIRC(F|_{\omega_X} \wedge (Y \not\equiv Y') \wedge \bigwedge Z^{\neg} \wedge \bigwedge Z'^{\neg})$ for every ω_X. As a result, $CIRC(T)$ is equivalent to $CIRC((X \not\equiv X') \wedge F \wedge (Y \not\equiv Y') \wedge \bigwedge Z^{\neg} \wedge \bigwedge Z'^{\neg})$, which is equivalent to $(X \not\equiv X') \wedge F \wedge (Y \not\equiv Y') \wedge \bigwedge Z^{\neg} \wedge \bigwedge Z'^{\neg}$ by Property 1. This formula has size $4n + ||F|| + 4n + 2m = 8n + ||F|| + 6n + 2||F|| + 2 = 14n + 3||F|| + 2 = k$.

If F is false for some ω_X, then $CIRC(T|_{\omega_X})$ is equivalent to $Form(\omega_{X'}) \wedge CIRC((Z \not\equiv Z') \wedge \bigwedge Y \wedge \bigwedge Y')$, which is also equivalent to $(Z \not\equiv Z') \wedge \bigwedge Y \wedge \bigwedge Y'$ by applying the second and third equivalence of Property 1. For every $z \in Z$, this formula has a model that makes z true, but changing only the evaluation of z results in a model not satisfying this formula. The same applies to all variables in Z and Z' and their negation, and to all variables in Y and Y'. By Lemma 1, every formula equivalent to this one has size greater than or equal to $4|Z| + 2|Y| = 4m + 2n = 4(3n + ||F|| + 1) + 2n = 12n + 4||F|| + 4 + 2n = 16n + 4||F|| + 4 > k$. By Lemma 2, every formula equivalent to $CIRC(T)$ has size greater than or equal to this amount. $\square$

IV. Conclusions

The problem of checking whether the circumscription of a formula can be represented by a formula of size bounded by k turned out to be Π_2^p-hard in Σ_3^p. These two classes are at the second and third level of the polynomial hierarchy, respectively. As a result, the problem cannot be solved by a propositional satisfiability solver. It can, however, be translated into a QBF and then passed as input to one of the existing QBF solvers [15].

An open question is how much complexity decreases if the formulae are in Horn form, and in particular if some additional restriction makes the problem tractable. If k is in binary representation rather than unary, the question is whether $CIRC(T)$ can be represented by a formula that may be exponential, but still bounded by k. The necessity of considering such large formulae is likely to make this problem harder than with k in unary notation: polynomial space may not be sufficient to solve it.

Indeed, assuming k in unary notation amounts to requiring the equivalent formula to have size comparable to that of the input data. This is equivalent to ask whether $CIRC(T)$ is equivalent to a formula of the same size of another formula G, for example. Allowing k to be stored in binary form with n bit allows the bound be as large as $2^n - 1$. As a result, even formulae of exponential size are allowed as representations of $CIRC(T)$. What complicates the analysis is that the usual guess-and-check algorithm for finding such a formula does not work in polynomial space, as this may not be enough for even storing the formula. A cycle of the minimal models of T is still feasible, but this may not allow determining the size of a formula satisfied exactly by all of them, unless such a formula is explicitly produced.

Acknowledgements

The author expresses his profound gratitude to Roberto Baldoni and Marco Schaerf for their unvaluable suggestions during the writing of this article.

References

[1] G. Brewka, I. Niemelä, and M. Truszczynski, "Nonmonotonic reasoning," in *Handbook of Knowledge Representation*, V. Lifschitz, B. Porter, and F. van Harmelen, Eds. Elsevier, pp. 239–284, 2007.

[2] K. Frankish, "Non-monotonic inference," in *The Encyclopedia of Language and Linguistics*, A. Barber, Ed. Elsevier, 2005.

[3] K. Sengupta, A. Krisnadhi, and P. Hitzler, "Local closed world semantics: Grounded circumscription for owl," in *The Semantic Web–ISWC 2011*. Springer, pp. 617–632, 2011.

[4] M. Shanahan, "The frame problem," in *The Stanford Encyclopedia of Philosophy*, E. N. Zalta, Ed., 2009.

[5] J. van Benthem, "McCarthy variations in a modal key," *Artificial Intelligence*, vol. 175, no. 1, pp. 428–439, 2011.

[6] M. Magnusson, J. Kvarnström, and P. Doherty, "Abductive reasoning with filtered circumscription," *Nonmonotonic Reasoning, Action and Change*, p. 55, 2009.

[7] K. Inoue and H. Saito, "Circumscription policies for induction," in *Inductive Logic Programming*, ser. Lecture Notes in Computer Science, R. Camacho, R. King, and A. Srinivasan, Eds., 2004, vol. 3194, pp. 164–179.

[8] P. Liberatore and M. Schaerf, "Reducing belief revision to circumscription (and viceversa)," *Artificial Intelligence*, vol. 93, no. 1–2, pp. 261–296, 1997.

[9] M. Cadoli, F. M. Donini, M. Schaerf, and R. Silvestri, "On compact representations of propositional circumscription," *Theoretical Computer Science*, vol. 182, pp. 183–202, 1997.

[10] S. Arora and B. Barak, *Computational Complexity: A Modern Approach*. Cambridge University Press, 2009.

[11] D. Buchfuhrer and C. Umans, "The complexity of boolean formula minimization," *Journal of Computer and System Sciences*, vol. 77, no. 1, pp. 142–153, 2011.

[12] T. Eiter and G. Gottlob, "Propositional circumscription and extended closed world reasoning are Π_2^p-complete," *Theoretical Computer Science*, vol. 114, pp. 231–245, 1993.

[13] M. Cadoli and M. Lenzerini, "The complexity of propositional closed world reasoning and circumscription," *Journal of Computer and System Sciences*, vol. 48, pp. 255–310, 1994.

[14] M. Thomas, "The complexity of circumscriptive inference in posts lattice," *Theory of Computing Systems*, vol. 50, no. 3, pp. 401–419, 2012.

[15] F. Donini, P. Liberatore, F. Massacci, and M. Schaerf, "Solving QBF with SMV," in *Proc. of the 8th International Conference on Principles of Knowledge Representation and Reasoning*, pp. 578–589, 2002.