Complexity of Contextual Reasoning

Floris Roelofsen[1,2]
[1]University of Twente, P.O. box 217
7500 AE Enschede, Netherlands
Phone: +31 53 4893680
Fax: +31 53 4893503
roelofsen@itc.it

Luciano Serafini[2]
[2]ITC-IRST, Via Sommarive 18
38050 Povo, Trento, Italy
Phone: +39 0461 314 319
Fax: +39 0461 314 591
serafini@itc.it

Abstract

This paper delineates the computational complexity of propositional multi-context systems. We establish NP-membership by translating multi-context systems into bounded modal K_n, and obtain more refined complexity results by achieving the so-called bounded model property: the number of local models needed to satisfy a set of formulas Φ in a multi-context system MS is bounded by the number of contexts addressed by Φ plus the number of bridge rules in MS.

Exploiting this property of multi-context systems, we are able to encode contextual satisfiability into purely propositional satisfiability, providing for the implementation of contextual reasoners based on already existing specialized SAT solvers.

Finally, we apply our results to improve complexity bounds for McCarthy's propositional logic of context – we show that satisfiability in this framework can be settled in non-deterministic polynomial time $O(|\varphi|^2)$.

Keywords:
Contextual reasoning, computational complexity.

Introduction

The establishment of a solid paradigm for contextual knowledge representation and contextual reasoning is of paramount importance for the development of sophisticated theory and applications in Artificial Intelligence.

McCarthy (1987) pleaded for a formalization of context as a possible solution to the problem of *generality*, whereas Giunchiglia (1993) emphasized the principle of *locality*: reasoning based on large (common sense) knowledge bases can only be effectively pursued if confined to a manageable subset (context) of that knowledge base.

Contextual knowledge representation has been formalized in several ways. Most notable are the propositional logic of context (PLC) developed by McCarthy, Buvač and Mason (1993; 1998), and the multi-context systems (MCS) devised by Giunchiglia and Serafini (1994), which later became associated with the local model semantics (LMS) introduced by Giunchiglia and Ghidini (2001). MCS has been proven strictly more general than PLC (Bouquet & Serafini 2004).

Contexts were first implemented as microtheories into the CYC common sense knowledge base (Lenat & Guha 1990). However, while in CYC the notion of local microtheories was a choice, in contemporary settings like the semantic web the notion of local, distributed knowledge is a must. Modern architectures impose highly scattered, heterogeneous knowledge fragments, which a central reasoner cannot deal with. This engenders a high demand for distributed, contextual reasoning procedures.

More recently, the idea of grid computing (Foster 1998) has received ample attention and fostered the development of various distributed reasoning systems (G. Behrmann & Vaandrager 2000; Chrabakh & Wolski 2003) which show, from the practical point of view, that implementing logical reasoners as cooperative systems of autonomous local reasoners can indeed improve performance.

The *complexity* of contextual reasoning, however, has so far received little attention. Massacci (1996) accomplishes a non-deterministic tableaux-based decision procedure, which establishes NP-membership for PLC, but leaves MCS/LMS out of consideration. Serafini and Roelofsen (2004) provide a deterministic SAT-based decision procedure that applies to both MCS/LMS and PLC, but they do not consider the effect that introducing non-determinism may have on the inherent complexity of the problem they are facing.

The goal of this paper is exactly this: to characterize the inherent computational complexity of contextual reasoning. The lion's share of our analysis regards reasoning based on MCS/LMS. Towards the end of the paper, however, our results are shown to be applicable to PLC as well.

We proceed as follows. After defining MCS/LMS and explicating the contextual satisfiability problem we establish an equivalence result with bounded modal K_n, which directly entails NP-membership. In pursuit of more specific upper bounds, we subsequently embark upon a more direct analysis of contextual satisfiability, which leads to the so-called bounded model property for multi-context systems. Next, we encode the contextual satisfiability problem into a purely propositional one. This encoding paves the way for the implementation of contextual reasoning systems based on already existing SAT solvers. At last, we show how our results can be applied to obtain improved complexity results for PLC. We conclude with a concise recapitulation of our achievements, and some pointers to future research avenues.

Multi-Context Systems

A simple illustration of the intuitions underlying MCS/LMS is provided by the so-called "magic box" example (Ghidini & Giunchiglia 2001), depicted below.

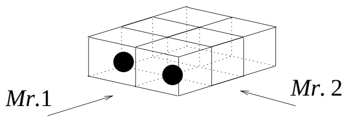

Figure 1: The magic box

Example 1 *Mr.1 and Mr.2 look at a box, which is called "magic" because the observers cannot make out its depth. Both Mr.1 and Mr.2 maintain a local representation of what they see. These representations must be coherent – if Mr.1's sees a ball, for instance, then Mr.2's must see some ball too.*

We will now show how such interrelated local representations can be captured formally. Our point of departure is a set of indices I. Each index $i \in I$ denotes a *context*, which is described by a formal (in this case propositional) language L_i. To state that a propositional formula φ in the language L_i holds in context i we utilize so-called *labeled formulas* of the form $i : \varphi$ (when no ambiguity arises we will simply refer to *labeled formulas* as *formulas*). Formulas that apply to different contexts may be related by so-called *bridge rules*. These are expressions of the form:

$$i_1 : \phi_1, \ldots, i_n : \phi_n \to i : \varphi \tag{1}$$

where $i_1, \ldots, i_n, i \in I$ and $\phi_1, \ldots, \phi_n, \varphi$ are formulas. Note that "$\to$" does not denote implication (we'll use "$\supset$" for this purpose). Also note that our language does not include expressions like $\neg(i : \varphi)$ and $(i : \varphi \wedge j : \psi)$. $i : \varphi$ is called the *consequence* and $i_1 : \phi_1, i_n : \phi_n$ are called *premises* of bridge rule (1). We write $cons(br)$ and $prem(br)$ for the consequence and the set of all premises of a bridge rule br, respectively.

Definition 1 (Propositional Multi-Context System) *A propositional multi-context system $\langle \{L_i\}_{i \in I}, \mathbb{BR} \rangle$ over a set of indices I consists of a set of propositional languages $\{L_i\}_{i \in I}$ and a set of bridge rules $\mathbb{BR}$.*

In this paper, we assume I to be (at most) countable and $\mathbb{BR}$ to be finite. Note that the latter assumption does not apply to MCSs with *schematic* bridge rules, such as provability - and multi-agent belief systems (Giunchiglia & Serafini 1994). The question whether our results may be generalized to capture these cases as well is subject to further investigation.

Example 2 *The situation described in example 1 may be formalized by an MCS with two contexts 1 and 2, described by $L_1 = L(\{l, r\})$ and $L_2 = L(\{l, c, r\})$, respectively. The constraint that Mr.2 must see a ball if Mr.1 sees one, can be captured by the following bridge rule:*

$$1 : l \vee r \quad \to \quad 2 : l \vee c \vee r$$

Let M_i denote the class of classical interpretations of L_i. An interpretation $m \in M_i$ is called a *local model* of L_i. Interpretations of entire MCSs are called *chains*. They are constructed from sets of local models.

Definition 2 (Chain) *A chain c over a set of indices I is a sequence $\{c_i\}_{i \in I}$, where each $c_i \subseteq M_i$ is a set of local models of L_i. A chain c is i-consistent if c_i is nonempty. It is point-wise if $|c_i| \leq 1$ for all $i \in I$, and set-wise otherwise.*

A chain can be thought of as a set of "epistemic states", each corresponding to a certain context (or agent). The fact that c_i contains more than one local model means that L_i can be interpreted in more than one unique way. So, set-wise chains correspond to partial knowledge, whereas point-wise chains indicate complete knowledge.

Example 3 *Consider the situation depicted in Figure 1. Both agents have complete knowledge, corresponding to a point-wise chain $\{\{\{l, r\}\}, \{\{l, \neg c, \neg r\}\}\}$. We can imagine a scenario however, in which Mr.1 and Mr.2's views are restricted to the right half and the left-most section of the box, as depicted in Figure 2.*

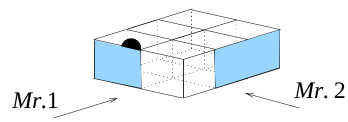

Figure 2: The partially hidden magic box

Now, both Mr.1 and Mr.2 have only partial knowledge; their observations may be interpreted in several different ways. This is reflected by the set-wise chain:

$$\left\{ \begin{array}{c} \{\{l, \neg r\}, \{\neg l, \neg r\}\}, \\ \{\{l, \neg c, \neg r\}, \{l, \neg c, r\}, \{l, c, \neg r\}, \{l, c, r\}\} \end{array} \right\}$$

The epistemic states that a chain consists of concern *one and the same* situation. Therefore, arbitrary sets of local models may not always constitute a "sensible" chain. The somewhat vague conception of "sensibility" is captured by the more formal notion of "bridge rule compliance" specified below.

Definition 3 (Compliance and Satisfiability) *Let c be a chain, φ a formula over L_i, and br an element of the set of bridge rules $\mathbb{BR}$ of a multi-context system MS.*

1. *$c \models i : \varphi$ if $m \models \varphi$ in a classical sense for all local models $m \in c_i$. We say that c satisfies $i : \varphi$.*
2. *c complies with br if either $c \models cons(br)$ or $c \not\models i : \xi$ for some $i : \xi \in prem(br)$. c complies with $\mathbb{BR}$ if it complies with every $br \in \mathbb{BR}$.*
3. *If there exists an i-consistent chain c that satisfies $i : \varphi$ and complies with $\mathbb{BR}$, we say that $i : \varphi$ is consistently satisfiable in MS.*

The contextual satisfiability problem, then, is to determine whether or not a set of labeled formulas Φ is consistently satisfiable in a multi-context system MS.

Example 4 *Consider an MCS with two contexts 1 and 2, described by $L(\{p\})$ and $L(\{q\})$, respectively, and subject to the following bridge rules:*

$$\begin{array}{rcl} 1 : p & \to & 2 : q \\ 1 : \neg p & \to & 2 : q \end{array}$$

The formula $2 : \neg q$ is satisfied in this system by the chain:

$$\left\{ \begin{array}{c} \{\{p\}, \{\neg p\}\}, \\ \{\{\neg q\}\} \end{array} \right\}$$

This example reflects that an MCS cannot be encoded into propositional logic by simply indexing propositions – such an encoding of the above system would be inconsistent.

Hereafter we refer to the set of bridge rules of MS as $\mathbb{BR}$, and to the set of contexts involved by formulas in Φ as J.

Encoding Into Bounded Modal K_n

A first insight regarding the complexity of contextual SAT may be obtained by investigating its encoding into modal K_n satisfiability. In this section we show that any contextual satisfiability problem may be reduced to that of satisfying some formula in K_n, whose depth is at most equal to one. This problem is known to be NP-complete (Ladner 1977; Halpern & Moses 1992).

Let us define a translation $(.)^*$ of labeled formulas into modal formulas:

$$(i : \phi)^* = \Box_i \phi$$

For bridge rules we have:

$$(i_1 : \phi_1, \ldots, i_n : \phi_n \rightarrow i : \phi)^* = \\ (i_1 : \phi_1)^* \wedge \ldots \wedge (i_n : \phi_n)^* \supset (i : \phi)^*$$

And a j-consistency constraint is captured by:

$$(j\text{-cons})^* = \neg \Box_j \bot$$

Theorem 1 *There is a kripke model $K = \langle W, \pi, \mathbf{R} \rangle$ such that $K, w_0 \models \psi$ for some $w_0 \in W$ and:*

$$\psi = \bigwedge_{i:\phi \in \Phi} (i : \phi)^* \wedge \bigwedge_{j \in J} (j\text{-}cons)^* \wedge \bigwedge_{br \in \mathbb{BR}} (br)^*$$

if and only if there is a J-consistent chain c^K that satisfies Φ and complies with $\mathbb{BR}$.

Proof. ($\Rightarrow$) We demonstrate how to construct c^K from K. Let m_w be the interpretation of $\bigcup_{i \in I} L_i$ associated to a world $w \in W$; for any $i \in I$, let $m_w|_i$ be the restriction of m_w to L_i and let $c_i^K = \{m_w|_i \mid w_0 R_i w\}$.

As $K, w_0 \models \Box_i \phi$, we have that $w \models \phi$ for any w with $w_0 R w$. Moreover, as $\phi \in L_i$, we have that $m_w|_i \models \phi$. This implies that $c^K \models i : \phi$. Bridge rule compliance and J-consistency are established likewise.

($\Leftarrow$) From c^K we may obtain a suitable kripke model K. Let W consist of a world w_0 plus one world w_{m_i} for each local model m_i of every component c_i^K of c^K. Let every $w_{m_i} \in W/\{w_0\}$ evaluate L_i according to m_i, and assign $True$ to the rest of $\bigcup_{i \in I} L_i$. Let w_0 evaluate every atomic proposition to $True$. For all $i \in I$, let:

$$R_i = \{\langle w_0, w_{m_i} \rangle \mid w_{m_i} \text{ corresponds to } m_i \in c_i^K\}$$

The resulting model is schematically depicted in Figure 3. One can easily verify that $K, w_0 \models \psi$. $\Box$

Contextual satisfiability clearly subsumes classical SAT, and is therefore NP-hard (Cook 1971). The above result, and the fact that satisfiability for bounded modal K_n is in NP (Ladner 1977), imply that contextual satisfiability is also in NP, and therefore NP-complete.

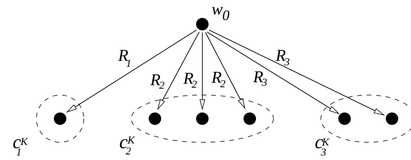

Figure 3: A schematic kripke model for ψ.

Moreover, the syntax of the formula that results from our translation is highly constrained: we obtain a conjunction of disjunctions of (negated) boxed formulas. Each disjunction comprises at most one boxed formula that is not negated, and furthermore, each boxed formula is purely propositional. This form strongly alludes to the existence of relatively efficient ways to solve the contextual satisfiability problem. Therefore, to obtain a more nuanced understanding of its complexity, we proceed with a more direct analysis.

Firsthand Analysis

We first introduce some notation and terminology. The size of a labeled formula $i : \varphi$ is denoted by $|i : \varphi|$. Let $P(i : \varphi)$ and $P(\Phi)$ be the set of propositional atoms appearing in a formula $i : \varphi$ or a set of formulas Φ. Let G_i be the number of local models contained by the i^{th} component of a chain, and let G be the total number of local models comprising that chain. Let $\Xi(br)$ and $\Xi(\mathbb{BR})$ consist of the premises and the consequence(s) of a bridge rule br or a set of bridge rules $\mathbb{BR}$. Finally, let N be the total size of the formulas in Φ and $\Xi(\mathbb{BR})$:

$$N = \sum_{i:\varphi \in \Phi} |i : \varphi| + \sum_{i:\xi \in \Xi(\mathbb{BR})} |i : \xi|$$

We first consider the *model checking problem*, that is, the problem of determining whether a given chain c consistently satisfies a set of formulas Φ in a multi-context system MS. This task can be split into three sub-tasks:

1. Check whether c satisfies Φ;
2. Check whether c complies with $\mathbb{BR}$;
3. Check whether c is J-consistent.

Theorem 2 *Model checking can be performed deterministically in time:*

$$O(\sum_{i:\varphi \in \Phi \cup \Xi(\mathbb{BR})} G_i \times |\varphi|)$$

Proof. First consider sub-task 1. Checking whether a particular formula $i : \varphi \in \Phi$ is satisfied by c can be done as follows. Let $\varphi_1, \ldots, \varphi_k$ be an ordering of the subformulas of φ, such that $\varphi_k = \varphi$ and if φ_i is a subformula of φ_j, then $i < j$. Since φ has at most $|\varphi|$ subformulas, we have $k < |\varphi|$. By induction on k' we can label each local model m in c_i with either φ_j or $\neg\varphi_j$, for $j = 1, \ldots, k'$, depending on whether or not $m \models \varphi_j$, in time $O(G_i \times k')$. As a result, checking whether c satisfies Φ can be carried out in time $O(\sum_{i:\varphi \in \Phi} G_i \times |\varphi|)$.

Sub-task 2 takes time $O(\sum_{i:\xi\in\Xi(\mathbb{BR})} G_i \times |\xi|)$, as in the worst case it involves checking whether all the consequences and premises of every bridge rule in $\mathbb{BR}$ are satisfied or not. Sub-task 3 merely consists in checking whether c_j is nonempty, for $j \in J$. This can be done in $O(|J|)$ timesteps. The result follows directly. □

Next, we consider satisfiability. We first show that MCSs enjoy the so-called *bounded model property*. More specifically, we establish that if a chain consistently satisfies Φ in MS, then it can be reduced to a chain that contains at most $|J| + |\mathbb{BR}|$ local models and still consistently satisfies Φ. Using this result, we reprove contextual satisfiability to be NP-complete, and establish an upper bound for the amount of time it requires.

Theorem 3 (Bounded Model Property) *A set of formulas Φ is consistently satisfiable in a multi-context system MS iff there is a J-consistent chain that contains at most $|J|+|\mathbb{BR}|$ local models and satisfies Φ in compliance with $\mathbb{BR}$.*

Proof. Take any J-consistent chain c that satisfies Φ in compliance with $\mathbb{BR}$. Let $\mathbb{BR}^* \subseteq \mathbb{BR}$ be the set of bridge rules whose consequences are not satisfied by c. Every $br \in \mathbb{BR}^*$ must have a premise which is not satisfied in some local model m_{br} contained by c. On the other hand, for every $j \in J$, there must be at least one local model $m_j \in c_j$ that satisfies all those formulas in Φ that apply to context j. The chain c^* obtained from c by eliminating all local models except for:

$$\bigcup_{j\in J} m_j \cup \bigcup_{br\in\mathbb{BR}^*} m_{br}$$

is J-consistent, satisfies Φ in compliance with $\mathbb{BR}$ and contains at most $|J| + |\mathbb{BR}^*| \leq |J| + |\mathbb{BR}|$ local models. □

Theorem 4 *Contextual satisfiability is NP-complete and can be settled in non-deterministic time:*

$$O((|J| + |\mathbb{BR}|) \times N)$$

Proof. We already observed that contextual satisfiability is NP-hard. Now, to determine satisfiability we may proceed as follows. First, we non-deterministically appoint a set $Cons$ of bridge rule consequences, and a set $Prem$ of bridge rule premises, such that for every $br \in \mathbb{BR}$, either br's consequence is in $Cons$, or one of br's premises is in $Prem$. Let J, I_{Cons}, and I_{Prem} be the set of contexts involved by Φ, $Cons$, and $Prem$, respectively. Furthermore, let Φ_i, $Cons_i$, and $Prem_i$ be the set of i-formulas contained by Φ, $Cons$, and $Prem$, respectively. We construct a chain c, such that:

- For all $i \in I_{Prem}$, c_i contains $|Prem_i|$ local models;
- For all $i \in J/I_{Prem}$, c_i contains exactly one local model;
- For all $i \notin J \cup I_{Prem}$, c_i is empty;
- For all $i \in I$, each $m \in c_i$ evaluates the propositional atoms *not* appearing in $\Phi_i \cup Cons_i \cup Prem_i$ to $True$.

The only "guessing" involved in constructing c, apart from the choice of $Cons$ and $Prem$, are the truth values to which each local model in c_i should evaluate the propositional atoms in $P(\Phi_i \cup Cons_i \cup Prem_i)$. Notice that c contains at most $|J|+|Prem| \leq |J|+|\mathbb{BR}|$ local models, which are distributed over those components c_i of c with $i \in J \cup I_{Prem}$; all the other components of c are empty. Consider a local model m contained in c_i for some $i \in J \cup I_{Prem}$. The number of atomic propositions $|P(\Phi_i \cup Cons_i \cup Prem_i)|$ that m should "explicitly" evaluate is clearly bounded by N. We must appoint at most $|J| + |\mathbb{BR}|$ such explicit valuations (one for each local model in c), so c can be constructed in non-deterministic time $O((|J| + |\mathbb{BR}|) \times N)$.

It remains to check whether c is J-consistent, satisfies Φ, and complies with $\mathbb{BR}$. By theorem 2 this can be done in deterministic time $O((|J| + |\mathbb{BR}|) \times N)$.

Theorem 3 assures that, if Φ is consistently satisfiable in MS, then guessing a chain as described above is bound to result in a suitable one. Thus, consistent satisfiability of Φ in MS can be determined in non-deterministic polynomial time $O((|J| + |\mathbb{BR}|) \times N)$. □

Encoding Into Propositional SAT

As contextual satisfiability is in NP, it is tractably reducible to purely propositional SAT. In providing such a reduction, we may loose the particular structure of our problem, but do lay the groundwork for the implementation of purely SAT-based contextual reasoners, which could benefit from the well-advanced techniques developed by the SAT community.

To obtain a purely propositional representation of multi-contextual satisfiability problems, we exploit the understanding we obtained while establishing the bounded model property in the previous section. The key insight there was that a set of formulas Φ is satisfied by a chain c if and only if it is satisfied by chain c^b such that:

- For each $j \in J$, c^b_j contains at least one local model m_j that satisfies all the formulas in Φ that apply to context j.
- For every bridge rule $br \in \mathbb{BR}$ whose consequence is not satisfied by c, there is at least one premise $i : \xi$ of br, such that c^b_i contains a local model m_{br} that satisfies $\neg\xi$.

Notice that to meet these requirements, the number of local models in each component of c^b can be kept down to $|\mathbb{BR}|$ (we assume that $|\mathbb{BR}| \geq 1$). Also, if a non-empty component of c^b contains less than $|\mathbb{BR}|$ local models it can be extended to comprise exactly $|\mathbb{BR}|$ models, simply by adding duplicates of already existing models. So we may say that Φ is consistently satisfiable in MS if and only if it is satisfied by a J-consistent chain c^* all of whose components are either empty or contain exactly $|\mathbb{BR}|$ local models.

Now, we construct a propositional formula ψ, which is satisfiable if and only if such a chain c^* exists. We express this formula in a language which contains a propositional atom p^k_i for every $p \in L_i$, and for every $k = 1, \ldots, |\mathbb{BR}|$. Intuitively, the truth value assigned to p^k_i by a propositional model of ψ corresponds to the truth value assigned to p by the k^{th} local model in c^*_i. The language also contains an atomic proposition e_i for each index $i \in I$. Intuitively, e_i being assigned $True$ corresponds to c^*_i being empty.

We write $K = \{1, \ldots, |\mathbb{BR}|\}$. For any formula φ, and each $i \in I$ and $k \in K$, let φ_i^k denote the formula that results from substituting every atomic proposition p in φ with p_i^k. Furthermore, let $\varphi_i^K = \bigwedge_{k \in K} \varphi_i^k$. Then, the translation of a labeled formula reads:

$$(i : \varphi)^* = e_i \vee \varphi_i^K$$

For bridge rules we have:

$$(i_1 : \varphi_1, \ldots, i_n : \varphi_n \rightarrow i : \phi)^* =$$
$$(i_1 : \phi_1)^* \wedge \ldots \wedge (i_n : \phi_n)^* \supset (i : \phi)^*$$

A j-consistency constraint is captured by:

$$(j\text{-cons})^* = \neg e_j$$

Theorem 5 *There is an assignment V to the propositions $\{p_i^k \mid i \in I$ and $k = 1, \ldots, |\mathbb{BR}|\} \cup \{e_i \mid i \in I\}$ that satisfies:*

$$\psi = \bigwedge_{i:\phi \in \Phi} (i : \phi)^* \wedge \bigwedge_{j \in J} (j\text{-}cons)^* \wedge \bigwedge_{br \in \mathbb{BR}} (br)^*$$

if and only if there is a J-consistent chain c^V that satisfies Φ in compliance with $\mathbb{BR}$.

Proof ($\Rightarrow$) From V we construct a chain c^V, such that each component c_i^V is empty if $V(e_i) = True$ and consists of exactly $|\mathbb{BR}|$ local models otherwise. In the latter case, the k^{th} local model of c_i^V is made to evaluate each propositional atom $p \in L_i$ to $True$ iff $V(p_i^k) = True$. It's easy to see that c^V is J-consistent and satisfies Φ in compliance with $\mathbb{BR}$.

($\Leftarrow$) By the above observations, if there is a J-consistent chain c that satisfies Φ in compliance with $\mathbb{BR}$, there must also be a J-consistent chain c^* each of whose components is either empty or contains exactly $|\mathbb{BR}|$ local models, and which still satisfies Φ in compliance with $\mathbb{BR}$.

From c^* we may obtain a truth assignment V as follows. To an atomic proposition e_i, let V assign $True$ iff $c_i^* = \emptyset$. To an atomic proposition p_i^k, let V assign $True$ if the k^{th} local model of c_i^* satisfies p, $False$ if the k^{th} local model of c_i^* satisfies $\neg p$, and any truth value if c_i^* is empty. It should be straightforward to see that V satisfies ψ. □

Application to PLC

We apply the results presented above to improve current complexity bounds for the propositional logic of context as described in (McCarthy 1993; McCarthy & Buvač 1998). The best result so far has been established by Massacci (1996). He proposes a tableaux-based decision procedure, which determines satisfiability of a PLC formula φ in non-deterministic polynomial time $O(|\varphi|^4)$.

We translate a PLC formula φ into a labeled formula $\epsilon : \varphi$ and a multi-context system MCS(φ), so that φ is satisfiable in PLC iff $\epsilon : \varphi$ is consistently satisfiable in MCS(φ). The translation is only sketched here – details can be found in (Bouquet & Serafini 2004). Subsequently, we demonstrate that determining whether or not $\epsilon : \varphi$ is consistently satisfiable in MCS(φ) takes non-deterministic time $O(|\varphi|^2)$.

The translation works as follows. For each nesting pattern *ist*$(k_1, \ldots$ *ist*$(k_n, \psi) \ldots)$ in φ, let MCS(φ) contain a context labeled with the sequence $k_1 \ldots k_n$. Let the language of context $k_1 \ldots k_n$ contain all the atomic propositions in ψ, plus a new atomic proposition for each formula of the form *ist*(k, χ) occurring in ψ. Finally, equip MCS(φ) with the following bridge rules[1]:

$$\bar{k}k : \psi \rightarrow \bar{k} : \mathit{ist}(k, \psi)$$
$$\bar{k} : \mathit{ist}(k, \psi) \rightarrow \bar{k}k : \psi$$
$$\bar{k} : \neg\mathit{ist}(k, \mathit{ist}(h, \psi)) \rightarrow \bar{k}k : \neg\mathit{ist}(h, \psi)$$
$$\bar{k} : \neg\mathit{ist}(k, \neg\mathit{ist}(h, \psi)) \rightarrow \bar{k}k : \mathit{ist}(h, \psi)$$

where $\bar{k} = k_1 \ldots k_n$ refers to any context of MCS(φ), whose language contains *ist*(k, ψ) or *ist*$(k,$ *ist*$(h, \chi))$, respectively.

Example 5 *Consider the following PLC formula:*

$$\varphi = p \vee \mathit{ist}(k, q \supset (\mathit{ist}(h, r \wedge s) \supset \mathit{ist}(j, q)))$$

MCS(φ) consists of four contexts which are labeled ϵ (the empty sequence), k, kh, and kj. The language of ϵ, L_ϵ, contains two propositions, p and ist$(k, q \supset ($ist$(h, r \vee s) \supset$ ist$(j, q)))$; L_k contains two propositions, q and ist$(h, r \wedge s)$; $L_{kh} = L(\{r, s\})$ and $L_{kj} = L(\{q\})$. The bridge rules of MCS(φ) are as stated above.

Theorem 6 (Bouquet & Serafini, 2004) *φ is satisfiable in PLC if and only if $\epsilon : \varphi$ is consistently satisfiable in MCS(φ).*

Theorem 7 *Satisfiability of φ in PLC can be computed in non-deterministic polynomial time $O(|\varphi|^2)$.*

Proof. By theorem 6 any satisfiability problem in PLC can be transformed into an equivalent satisfiability problem in MCS. This transformation can be established in linear time.

Every bridge rule in MCS(φ) involves at least one proposition of the form *ist*(k, ψ). Every such proposition occurs in at most four bridge rules. Every subformula of φ of the form *ist*(k, ψ) (and nothing else) results in a proposition of the form *ist*(k, ψ) in the language of exactly one context in MCS(φ). The number of subformulas of φ of the form *ist*(k, ψ) is bounded by $|\varphi|$. From these observations, we may conclude that the number of bridge rules $|\mathbb{BR}|$ in MCS(φ) is bounded by $4 \times |\varphi|$. Moreover, by construction, the sum of the lengths of the formulas involved in any bridge rule of MCS(φ) is at most four.

By theorem 4, contextual satisfiability of $\epsilon : \varphi$ in MCS(φ) can be settled in non-deterministic time:

$$O((|\Phi| + |\mathbb{BR}|) \times (\sum_{i:\varphi \in \Phi} |i : \varphi| + \sum_{br \in \mathbb{BR}} \sum_{i:\xi \in \Xi(br)} |i : \xi|))$$

In the light of the above observations, and keeping in mind that Φ merely consists of $\epsilon : \varphi$, we may rewrite this in terms of φ as:

$$O(|\varphi|^2)$$

□

[1] The first two bridge rules correspond to the notions of *entering* and *exiting* contexts (McCarthy & Buvač 1998), while the last two bridge rules correspond to the Δ axiom introduced by Buvač and Mason (1993).

Conclusion

We have analysed the complexity of contextual reasoning based on propositional multi-context systems with finite sets of bridge rules.

A first insight was obtained by establishing an encoding of contextual satisfiability into satisfiability in bounded multimodal K_n, which is well-known to be NP-complete. Next, we accomplished a more fine-grained upper bound for the complexity of contextual satisfiability by a direct investigation of its semantical properties. Herein we observed that multi-context systems enjoy the bounded model property.

We also provided a tractable encoding of contextual satisfiability problems into purely propositional ones. In doing so, we laid the groundwork for SAT-based implementations of contextual reasoning systems.

Finally, we obtained improved complexity results for the satisfiability problem in McCarthy's propositional logic of context, by translating it into the satisfiability problem that we have considered in this paper.

Future work will encompass experimentation with both native and SAT-based contextual reasoning systems. Also, we are interested whether, or to what extent, our results may be generalized so as to apply to multi-context systems with schematic bridge rules as well.

Acknowledgments

We wish to thank Johan van Benthem, Roberto Sebastiani, and the anonymous referees for their very useful comments.

References

Bouquet, P., and Serafini, L. 2004. Comparing formal theories of context in AI. *Artificial Intelligence*. Accepted for publication.

Buvač, S., and Mason, I. A. 1993. Propositional logic of context. In *National Conference on Artificial Intelligence (AAAI 93)*.

Chrabakh, W., and Wolski, R. 2003. Gridsat: a chaff-based distributed SAT solver for the grid. In *Supercomputing Conference (SC03)*.

Cook, S. 1971. The complexity of theorem proving procedures. In *ACM Symposium on the Theory of Computing*, 151–158.

Foster, I. 1998. *The Grid : Blueprint for a New Computing Infrastructure*. Morgan-Kaufmann.

G. Behrmann, T. H., and Vaandrager, F. 2000. Distributed timed model checking - how the search order matters. In *Computer Aided Design (CAV 00)*, 216–231.

Ghidini, C., and Giunchiglia, F. 2001. Local models semantics, or contextual reasoning = locality + compatibility. *Artificial Intelligence* 127(2):221–259.

Giunchiglia, F., and Serafini, L. 1994. Multilanguage hierarchical logics, or: how we can do without modal logics. *Artificial Intelligence* 65(1):29–70.

Giunchiglia, F. 1993. Contextual reasoning. *Epistemologia* XVI:345–364.

Halpern, J., and Moses, Y. 1992. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence* 54(3):311–379.

Ladner, R. 1977. The computational complexity of provability in systems of modal propositional logic. *SIAM Journal of Computing* 6:467–480.

Lenat, D. B., and Guha, R. V. 1990. *Building Large Knowledge Based Systems*. Addison Wesley.

Massacci, F. 1996. Contextual reasoning is NP complete. In *National Conference on Artificial Intelligence (AAAI 96)*, 621–626.

McCarthy, J., and Buvač, S. 1998. Formalizing context (expanded notes). In *Computing Natural Language*, volume 81 of *CSLI Lecture Notes*. 13–50.

McCarthy, J. 1987. Generality in artificial intelligence. *Communications of ACM* 30(12):1030–1035.

McCarthy, J. 1993. Notes on formalizing context. In *International Joint Conference on Artificial Intelligence (IJCAI 93)*, 555–560.

Serafini, L., and Roelofsen, F. 2004. Satisfiability for propositional contexts. In *Principles of Knowledge Representation and Reasoning (KR 04)*. Accepted for publication.