ECAI 2016
G.A. Kaminka et al. (Eds.)

doi:10.3233/978-1-61499-672-9-795

A Computational Approach to Consensus-Finding

Eric Grégoire[1] and **Jean-Marie Lagniez**[2]

Abstract. Consensus-finding plays a ubiquitous role in A.I. In this paper, a consensus among agents is defined as a non-contradictory fragment of all the information conveyed by the agents such that this fragment does not logically conflict with any of the agents. This concept is investigated in modal logic S5 in order to meet representation needs that are put in light by this concept of consensus itself. Interestingly, an optimization-based approach to compute maximal consensuses is developed and shown experimentally efficient very often for both the standard Boolean and S5 frameworks.

1 INTRODUCTION

Consensus-finding plays a ubiquitous role in A.I. For example, interacting agents [4, 17] can need to target a consensual, shared, goal whereas negotiation can amount to finding and settling on a consensual agreement. More generally, a notion of consensus can prove helpful for reconciling several information[3] sources, hereafter often simply called sources or agents. For example, cautious A.I. systems might rely on uncontroversial, consensual, fragments of the global information conveyed by several mutually conflicting belief sources [10, 7].

However, consensus paradigms are rarely defined in a precise way, especially in the context of several agents who are equipped with inferential reasoning capabilities. In this paper, a consensus among several sources is defined as a non-contradictory fragment of all the information conveyed by the sources such that this fragment does not logically conflict with any source. Interestingly, such a consensus as proposed in [5], might not contain only the information shared by every source; it can also contain some additional information that is in some sense possibly acceptable from the point of view of each source since it does not contradict it. Hence, each source might endorse the information in the consensus as this information either also belongs to the source or does not conflict with it.

First, we discuss this notion of consensus that has been defined in a Boolean setting in [5]. Then, it is extended to modal logic S5 (see for example [3, 2] for an introduction to modal logics) in order to meet representation needs that are put in light by this concept of consensus itself. The focus is on the practical computational extraction of maximal consensuses. Noticeably, it is stressed that the computation of one maximal consensus diverges from the well-studied search for maximal satisfiable subsets of all the information conveyed by the sources. Interestingly, we provide an optimization schema that proves experimentally efficient even for very large sources for both the standard Boolean and S5 frameworks, extending also [5] in this latter aspect.

[1] CRIL, Univ. Artois & CNRS, 62300 Lens, France gregoire@cril.fr
[2] CRIL, Univ. Artois & CNRS, 62300 Lens, France lagniez@cril.fr
[3] In this paper, no difference is made between information, plans, desires, goals, knowledge and beliefs.

The paper is organized as follows. In the next section, we provide the main logical preliminaries and notations used throughout the paper. In section 3, we introduce several basic concepts of consensus before we push the envelope in section 4 by investigating consensuses in modal logic S5. Section 5 presents a practical computational approach for the extraction of one consensus whereas section 6 extends it to S5. Section 7 reports our experimental study. The paper ends with a discussion and some promising perspectives for further research.

2 LOGICAL PRELIMINARIES

We consider the standard (Boolean logic) language $\mathcal{L}$ of formulas, based on a denumerable set of Boolean variables $\mathcal{P}$, which are written $a, b, \dots$ and can be assigned either *true* or *false*. The conjunctive, disjunctive, negation, material implication and equivalence connectives are written $\wedge, \vee, \neg, \rightarrow, \equiv$, respectively. Formulas and sets of formulas are denoted $\alpha, \beta, \dots$ and $\Phi, \Gamma, \dots$, respectively. In the following, we assume n agents Φ_i where $i \in [1..n]$ and identify each agent with her knowledge (actually, with the part of her knowledge that is concerned by the search for one consensus). Without loss of generality, we often assume that each Φ_i is under clausal form (CNF): namely, is a conjunction of clauses, where a clause is a disjunction of literals and a literal is a possibly negated Boolean variable.

Φ_i is satisfiable iff there exists a truth value assignment of every variable such that all formulas in Φ_i are *true* according to usual compositional rules. Such an assignment is called a model of Φ_i. Any Boolean formula can be rewritten in linear time in CNF that is equivalent with respect to satisfiability. $\vdash$ denotes the deduction relation: $\Phi_i \vdash \alpha$ iff α is *true* in all models of Φ_i. A tautology is *true* under any assignment; in other words, we have that $\vdash \alpha$ for every tautology α. Logically equivalent formulas are considered indistinguishable: for example, we do not distinguish between $a \vee b$ and $b \vee a$. $Th(\Phi_i)$ represents the set of deductive consequences (also called the deductive closure) of Φ_i, namely $Th(\Phi_i) = \{\alpha \in \mathcal{L} \text{ s.t. } \Phi_i \vdash \alpha\}$. Arrays of sets of formulas are called profiles and noted $\mathcal{S}, \mathcal{V}, \dots$ The cardinality of a set Φ is noted $\#\Phi$. We will use the concept of prime implicate, defined as follows. A prime implicate of a finite set Δ of formulas is any clause δ such that $\Delta \vdash \delta$, and, at the same time, $\vdash (\delta' \equiv \delta)$ for every clause δ' s.t. ($\Delta \vdash \delta'$ and $\delta' \vdash \delta$). For readability reason, we will present elements of modal logic when needed.

3 CONSENSUS: BASIC DEFINITIONS AND PROPERTIES

Assume $\mathcal{S} = [\Phi_1, \dots, \Phi_n]$ represents n sources Φ_i where each Φ_i is such that $\Phi_i \subset \mathcal{L}$ and is satisfiable. In this paper, we consider consensuses that are included within $\bigcup_{i=1}^{n} \Phi_i$, as defined in [5]:

Definition 1. *A set $\Gamma \subset \mathcal{L}$ is a consensus for $\mathcal{S}$ iff $\Gamma \subseteq \bigcup_{i=1}^{n} \Phi_i$ and $\forall\ \Phi_i \in \mathcal{S} : \Gamma \cup \Phi_i$ is satisfiable.*

Accordingly, a consensus is always satisfiable since no set is satisfiable together with an unsatisfiable one. Since each Φ_i is satisfiable, there exists always at least one consensus, which can be the empty set. When $\bigcup_{i=1}^{n} \Phi_i$ is satisfiable, it forms the only maximal consensus. In the general case, a consensus need not be unique. No logical consequence of a consensus contradicts any source: hence, a consensus can be identified with its deductive closure. Notice that it would be possible to require a consensus to be a subset of $Th(\bigcup_{i=1}^{n} \Phi_i)$ vs. a subset of $\bigcup_{i=1}^{n} \Phi_i$, however such an extended definition would not allow for the practical computational approach that we develop in this paper. Two natural definitions for maximal consensuses are as follows [5], depending on whether maximality is considered with respect to set-theoretic inclusion or set cardinality.

Definition 2. *A consensus Γ for $\mathcal{S}$ is max$_\subseteq$ iff $\forall\ \Theta$ s.t. $\Gamma \subset \Theta \subseteq \bigcup_{i=1}^{n} \Phi_i$, $\exists\ \Phi_i \in \mathcal{S}$ s.t. $\Theta \cup \Phi_i$ is unsatisfiable.*
A consensus Γ for $\mathcal{S}$ is max$_\#$ iff $\forall\ \Theta$ s.t. $\Theta \subseteq \bigcup_{i=1}^{n} \Phi_i$ and $\#\Theta > \#\Gamma$, $\exists\ \Phi_i \in \mathcal{S}$ s.t. $\Theta \cup \Phi_i$ is unsatisfiable.

Clearly, any *max*$_\#$ consensus is a *max*$_\subseteq$ consensus whereas the converse does not hold. From now on, we consider *max*$_\subseteq$ and *max*$_\#$ consensuses, only. We write *consensus* when there is no need to differentiate between *max*$_\subseteq$ and *max*$_\#$. Notice that consensuses and maximal (w.r.t. $\subseteq$ or $\#$) satisfiable subsets of $\bigcup_{i=1}^{n} \Phi_i$ (in short MSSes) are closely related but different concepts. In the general case, a consensus is not one MSS that is satisfiable with each Φ_i: although every consensus is included in some MSSes, an MSS is not necessary a consensus. Actually, the sets of MSSes and of consensuses can be disjoint.

Example 1. *Assume that three political groups are negotiating for a possible coalition program while their individual political agendas are mutually conflicting. One maximal consensus is searched within all elements of the programs such that it does not contradict any of the agendas. Although a political group might find within the consensus some elements that do not belong to its own agenda, it might endorse the consensus since these elements do not contradict its own objectives. Assume for example that $\{it, tss, ids\}$ is a set of Boolean variables standing for Increase Taxation, Trim Social Security and Increase Defense Spendings, respectively. Let $\mathcal{S}$ be the agendas of these three political groups: $\mathcal{S} = [\Phi_1, \Phi_2, \Phi_3]$ with $\Phi_1 = \{it, \neg tss, \neg it \rightarrow \neg ids\}$, $\Phi_2 = \{tss, it \rightarrow ids\}$ and $\Phi_3 = \{\neg ids\}$. It is easy to see that there are three max$_\#$ consensuses for $\mathcal{S}$: $\Gamma_1 = \{\neg it \rightarrow \neg ids, it \rightarrow ids\}$, $\Gamma_2 = \{\neg ids, \neg it \rightarrow \neg ids\}$ and $\Gamma_3 = \{it, \neg it \rightarrow \neg ids\}$. For example, Γ_3 states "Increase taxation and if we do not increase taxation then we do not increase the defense spendings". Remember that any consensus can be identified with its deductive closure: hence Γ_3 can be identified with $Th(\{ids\})$. Notice that none of the consensuses is an MSS of $\bigcup_{i=1}^{3} \Phi_i$.*

The consensus concept is easily extended by requiring forms of integrity constraints to be obeyed. In the following, we assume that a set of integrity constraints, noted Ψ, is included in $\bigcup_{i=1}^{n} \Phi_i$ [5]. Ψ can need to be included in any consensus or, as in the next definition, be a deductive consequence of any consensus (definitions for maximal consensuses are easily adapted, too). In a negotiation setting, integrity constraints can express elements that must belong to any consensus.

Definition 3. *A set $\Gamma \subset \mathcal{L}$ is a consensus for $\mathcal{S}$ under the constraints Ψ iff $\Gamma \subset \mathcal{L}$ is a consensus for $\mathcal{S}$ and $\Gamma \vdash \Psi$.*

Example 2. *In the previous example, $\Gamma = \{it, \neg it \rightarrow \neg ids\}$ is a consensus for $\mathcal{S}$ under the constraint $\Psi = \{it\}$. For example, there is no consensus for $\mathcal{S}$ under the constraint $\Psi = \{\neg tss\}$ since tss is logically conflicting with Φ_2.*

Notice that the last definition could be easily extended in such a way that a consensus includes or entails a set of integrity constraints that is not included in $\bigcup_{i=1}^{n} \Phi_i$ or even $Th(\bigcup_{i=1}^{n} \Phi_i)$.

From a computational point of view, extracting one *max*$_\subseteq$ consensus amounts to computing one MSS under an additional constraint of satisfiability with each Φ_i. Consequently, computing such consensuses in this way is as hard as computing MSSes, which is known to be intractable in the worst case. Indeed, the computation of one MSS$_\subseteq$ belongs to the $FP^{NP}[wit, log]$ class, i.e., the set of function problems that can be computed in polynomial time by executing a logarithmic number of calls to an NP oracle that returns a witness for the positive outcome [13]. Even worse, computing one MSS$_\#$ belongs to the $Opt\text{-}P$ class of problems [16], i.e., the class of functions computable by taking the maximum of the output values over all accepting paths of an NP machine.

It is also important to notice that the number of different consensuses is exponential in the number of clauses in $[\Phi_1, \ldots, \Phi_n]$ in the worst case. However, in many difficult negotiation problems and applications where a consensus needs to be found, extracting *one* maximal consensus is often the actual and sufficient problem to be solved. It can also be a useful starting point for further successful discussions. Accordingly, the focus in the rest of the paper is on the search for one maximal consensus; more precisely, we investigate the computation of one *max*$_\#$ consensus (remember that any *max*$_\#$ consensus is also a *max*$_\subseteq$ consensus).

4 CONSENSUS IN A MODAL LOGIC OF POSSIBILITY AND NECESSITY

4.1 Motivations

We claim that the *possibility* and *necessity* modalities (noted $\Diamond$ and $\Box$, respectively) of standard modal logics can be of specific interest in situations where consensus-finding must take place. In this section, we motivate the needs for the additional expressive power provided by these modal logics; we take advantage of these motivating examples to provide some early intuitive grasp about how consensus-finding will be actually implemented by reduction to standard Boolean logic. To this end, the reader only needs to remember at this stage that a formula $\Diamond\alpha$ is intended to assert that α is "possible", i.e., is *true* in some *possible world*.

Actually, the concept of consensus already allows an agent Φ_i to express that a formula α should be "possible" with respect to any consensus in the following sense. By definition, when $\alpha \in \Phi_i$ (and more generally when $\Phi_i \vdash \alpha$), no consensus for $\mathcal{S}$ contradicts α or, equivalently, for any consensus Γ, we have that $\Gamma \not\vdash \neg\alpha$. In the example from the previous section, if one political agenda Φ_i contains $\neg$*taxes-increase* then *taxes-increase* does not belong to any consensus Γ and is not inferable from Γ. Hence, any consensus conveys the information that $\neg$*taxes-increase* remains possible. This feature will be exploited to some extent to handle some occurrences of the modality operator $\Diamond$ in the search for consensuses for $\mathcal{S}$. However, this cannot be extended in the general case. Indeed, inserting every formula of a set Θ inside Φ_i to express that each formula Θ must be possible does not always yield the intended result. The simplest example of this is as follows. Assume Φ_i requires both a formula α and its contrary $\neg\alpha$ to be *true* in some (possibly different) possible worlds. Inserting both α and its contrary $\neg\alpha$ inside Φ_i will make Φ_i become unsatisfiable and collapse: no consensus will exist. On the contrary, both modal formulas $\Diamond\alpha$ and $\Diamond\neg\alpha$ are not mutually contradictory since they express that α and $\neg\alpha$ need to be *true* in some (different) possible worlds. Interestingly, by means of additional Boolean variables, it will be possible to reduce $\mathcal{S}$ expressed using modal logic S5 into mere (clausal) Boolean logic, and reuse a computational approach to consensus-finding that has been developed in the standard Boolean framework [5].

In the political agenda example, Φ_i might require both *taxes-increase* and $\neg$*taxes-increase* to remain possible: any consensus should prevent any of those formulas from being derivable. In this way, the taxation issue will remain an open question for possible further negotiations since it is not hindered by any consensus, which can form a first useful acquired result among the agents. Note that if both formulas were inserted within Φ_i then Φ_i would become unsatisfiable and no consensus could exist, hence the use of modalities to prevent this from happening. More generally, we claim that in negotiation and consensus-finding situations, there can be a need to express and handle requirements from agents asserting that some formulas should be possible, or in a dual way, that some formulas should not be derivable. Indeed, an agent might not only express negative desires under the form of standard logic formulas $\neg\alpha$ but also weaker desires that simply require some α not to be derivable. Additionally, since the latter form of desire is logically weaker than the former one, weakening some desires into mere "possible" or "not derivable" forms can be a way to allow for consensuses to exist when no consensus for the agents' initial requirements exists.

Let us also give another motivating example for the use of the possibility modality $\Diamond$. Assume that a consensus is to be found among two engineers who have conflicting diagnoses about a same device fault. The first one claims that the reason for failure is to be found in three (possibly cumulative) device faults. Let us tentatively represent this by $\Phi_1 = \{cause_1 \vee cause_2 \vee cause_3\}$. The second one is convinced that the reasons for failure are to be found in some of the first two causes, only. Assume that this is represented by $\Phi_2 = \{cause_1 \vee cause_2\}$. The only *max* consensus is $\Phi_1 \cup \Phi_2$ since $\Phi_1 \cup \Phi_2$ is satisfiable. From this consensus we can deduce $cause_1 \vee cause_2$. It can be argued that this result is counter-intuitive with respect to the intended role of a consensus since Φ_1 requires *cause*$_3$ to be also considered as a possible cause of failure whereas the same agent Φ_1 does not agree that the only two possible reasons for failure are *cause*$_1$ and *cause*$_2$. Actually, what the engineer Φ_1 needs to express is that $\alpha = cause_1 \vee cause_2 \vee cause_3$ is a prime implicate of her own knowledge: namely, no strict sub-clause of α is derivable. She must thus express that $\Phi_1 \not\vdash \alpha$ for any α that is a strict sub-clause of $cause_1 \vee cause_2 \vee cause_3$, or equivalently, that any corresponding formula $\neg\alpha$ must be *possible*. In this way, no consensus can allow one to conclude for example $cause_1 \vee cause_2$ without contradicting her own requirements. Notice that, like in the example about *taxes-increase* and $\neg$*taxes-increase*, the set of standard logic formulas $\neg\alpha$ such that α is a strict sub-clause of $cause_1 \vee cause_2 \vee cause_3$ is unsatisfiable. Hence, introducing this set within Φ_i would make Φ_i collapse and prevent any consensus from existing: this also justifies our use of modal logic.

4.2 Consensus in modal logic S5

Modal logic S5 is a canonical logic of possibility and necessity that can also be considered as a logic of knowledge. Its language $\mathcal{L}_M$ extends the language of standard Boolean logic by means of two modality connectives: $\Diamond$ and $\Box$. Main basic definitions and concepts about S5 are recalled in Appendix: for the understanding of this paper, it is sufficient to know that (1). $\Diamond$ and $\Box$ are used as additional unary connectives: for example $\Diamond(\Box a \vee b \vee \Diamond\neg c)$ is a well-formed formula of $\mathcal{L}_M$. (2). modalities are dual in the sense that $\Diamond\alpha = \neg\Box\neg\alpha$. (3). Truth values of modal formulas can vary depending on the considered so-called possible world, these worlds being connected by an equivalence accessibility relation in S5; the satisfiability paradigm is adapted accordingly. Interestingly, the satisfiability problem in S5 is NP-complete [11], just like SAT.

We assume that $\forall i \in [1..n] : \Phi_i \subset \mathcal{L}_M$ is S5-satisfiable and we consider $[\Phi_1, \ldots, \Phi_n]$ as the profile $\mathcal{S}$ for which a consensus needs to be found. Now, all definitions from Section 3 directly apply in this modal framework, using the satisfiability paradigm of S5. In this last respect, let us just introduce the two basic cases involving formulas with modalities. We will indicate how to handle more complex modal formulas later in the paper by reduction to these formulas and non-modal ones. Assume that Φ_i contains the formula $\Diamond\alpha$ where α does not contain any modalities: any consensus for $\mathcal{S}$ cannot allow one to deduce $\neg\alpha$ since any consensus should not conflict with Φ_1, which asserts that α is possible, i.e., is *true* in some possible world. When Φ_i contains $\Box\alpha$ then any non-empty consensus must contain α since it cannot conflict with Φ_1, which asserts that α is *true* in any possible world. Consequently, even when all Φ_i are satisfiable, the existence of a non-empty consensus is not guaranteed: when one agent asserts $\Box\alpha$ and a second one can deduce $\neg\alpha$, no non-empty consensus can exist. In this respect, $\Box\alpha$ can be interpreted as requiring α to be one integrity constraint for any non-empty consensus. Similarly, if one agent can deduce $\Box\alpha$ whereas another one can deduce $\Diamond\neg\alpha$ then no non-empty consensus exists.

Proposition 1. *Let $\mathcal{S} = [\Phi_1, \ldots, \Phi_n]$ such that $\forall i \in [1..n] : \Phi_i \subset \mathcal{L}_M$ is satisfiable. There is no non-empty consensus for $\mathcal{S}$ iff $\exists\, i, j \in [1..n]$ s.t. for some some α in $\mathcal{L}_M$ we have ($\Phi_i \models_{S5} \neg\alpha$ and $\Phi_j \models_{S5} \Box\alpha$) or ($\Phi_i \models_{S5} \Diamond\neg\alpha$ and $\Phi_j \models_{S5} \Box\alpha$).*

Example 3. *Let us come back to Example 1 and assume now that agent Φ_3 strengthens her desires and does not want to leave open any possibility in the consensus of having an increase of defense spendings: Φ_3 is now $\{\Box\neg ids\}$ (or equivalently $\{\neg\Diamond ids\}$). There remains only one max$_\#$ consensus, namely $\Gamma_2 = \{\Box\neg ids, \neg it \rightarrow \neg ids\}$. Note that $\Box\neg ids$ entails $\neg ids$ in S5. Now, if any Φ_i is then augmented with $\Diamond ids$ then no non-empty consensus exists anymore.*

5 COMPUTING ONE MAX$_\#$ CONSENSUS

Let us focus first on the computation of one *max*$_\#$ consensus in the standard Boolean framework. We assume that each Φ_i in $\mathcal{S}$ is a satisfiable set of clauses of $\mathcal{L}$.

If one MSS$_\#$ of $\bigcup_{i=1}^{n} \Phi_i$ were to be extracted, we could directly use SAT-related techniques to address this issue. Specifically, the problem would amount to solving a variant of Max-SAT($\bigcup_{i=1}^{n} \Phi_i$) that would return one MSS$_\#$ of its argument (instead of yielding merely the cardinality of this MSS). Actually, in order to deliver one *max*$_\#$ consensus, the computation of one MSS$_\#$ must also take into account the additional constraint requiring the result to be satisfiable with each Φ_i taken individually. Notice that since the Φ_i can be mutually conflicting, in the general case it is not possible to simply replace this multiple constraint by a unique one stating that the result should be satisfiable with $\bigcup_{i=1}^{n} \Phi_i$. A naive direct approach to compute one *max*$_{\#}$ consensus could however consist of the following steps. First, initialize Γ with $\bigcup_{i=1}^{n} \Phi_i$. Then, consider each Φ_j, successively. At each step, trim the current contents of Γ so that it becomes satisfiable with the current Φ_j. Clearly, at the end of the process, Γ would be satisfiable with each Φ_j and would be a consensus for $\mathcal{S}$. However, there would be no guarantee that Γ is a maximal consensus. Indeed, some clauses from Γ might have been dropped to ensure satisfiability with, say, Φ_j; at some subsequent step, some other clauses could have been discarded from Γ to ensure satisfiability with another source whereas dropping only these other clauses would have been sufficient to ensure that Γ is satisfiable with Φ_j. Hence to ensure that Γ is a *max*$_{\#}$ consensus, we would need to consider every possible ordering of all Φ_i and for each of them, consider every Φ_j and record the corresponding various minimal subsets of clauses to be dropped in order to ensure satisfiability with Φ_j. Based on all this information, we might finally select the clauses to be expelled to give rise to one maximal consensus. Clearly, such an approach is doomed to face a combinatorial blow-up very often, especially since the number of possible orderings to consider is exponential.

We have followed another path in [5] and have adapted a method, called *Transformational Method*, introduced in [1] to circumvent a close combinatorial issue consisting in extracting one maximal subset Γ from a set of clauses such that, at the same time, Γ is satisfiable with several possibly mutually conflicting contexts. The method is based on the transformation of the initial problem into one single-step optimization problem. Intuitively, the satisfiability of Γ with one given Φ_i is interpreted as a sub-problem, using its own range of fresh Boolean variables. All sub-problems are then linked together with the use of additional variables, called linking variables. The use of a Partial Max-SAT solver allows then to extract one subset of clauses that is a solution to the initial problem. There has been very significant progress these last years about the design of experimentally efficient (Partial) Max-SAT solvers: see for example [15] and the related international competitions `http://www.maxsat.udl.cat/`. Interestingly, despite the increase of the problem size that is linear with respect to the number of contexts, this approach proves experimentally far more efficient and scalable than the above naive method [1]. In order to adapt this method to compute one *max*$_{\#}$ consensus, we use Partial Max-SAT, which belongs to the *Opt-P* class of problems [16].

Definition 4. *Let Σ_1 and Σ_2 be two sets of clauses. Partial Max-SAT(Σ_1, Σ_2) computes one cardinality maximal subset of Σ_1 that is satisfiable with Σ_2. Σ_1 and Σ_2 are called the sets of soft and hard constraints, respectively.*

Algorithm 1 depicts the method. The problem of having Γ being a subset of $\bigcup_{i=1}^{n} \Phi_i$ that is satisfiable with each Φ_j is first treated as n independent subproblems; these subproblems will be then linked together to form one single optimization problem through one single call to Partial Max-SAT. Each clause δ_i^j from any Φ_i is augmented with an additional disjunct $\neg\epsilon_i^j$ using a new fresh variable (line 2): this yields a set Σ. These ϵ_i^j variables will be used to *link* the various subproblems. Each subproblem is created by unioning Σ with one Φ_i and by renaming all variables except the ϵ_i^j (l. 4-7). All together, the subproblems form the set of hard clauses; these ones are all simultaneously satisfiable (just assign all ϵ_i^j to *true*). The set of soft clauses is made of all unit clauses ϵ_i^j (l. 3). The instance of the Partial Max-SAT problem with these sets of hard and soft clauses will search one truth-value assignment such that all hard clauses and one maximal number of clauses ϵ_i^j are satisfied. Accordingly, all clauses δ_i^j corresponding to the satisfied ϵ_i^j form one *max*$_{\#}$ consensus for $\mathcal{S}$. Notice that the use of additional variables and clauses is a paradigm that has long been exploited in Max-SAT computation and the extraction of minimal unsatisfiable sets of clauses by other authors, see for example [6] and [14].

Algorithm 1: Compute one *max*$_{\#}$ consensus for $\mathcal{S}$

input : $\mathcal{S} = [\Phi_1, \ldots, \Phi_n]$: a profile of n satisfiable sets of Boolean clauses;
Assume that the clauses of Φ_i are noted $\delta_i^1, \delta_i^2, \ldots$;
output: One *max*$_{\#}$ consensus for $\mathcal{S}$;

1 $\Gamma_{\text{Hard}} \leftarrow \emptyset$; $\Gamma_{\text{Soft}} \leftarrow \emptyset$;
2 $\Sigma \leftarrow \bigcup_{\Phi_i \in S} \{\neg\epsilon_i^j \vee \delta_i^j \mid \delta_i^j \in \Phi_i$ and ϵ_i^j are new variables$\}$;
3 $\Gamma_{\text{Soft}} \leftarrow \{\epsilon_i^j\}_{i,j}$;
4 **foreach** $\Phi_i \in S$ **do**
5     $\Phi_i \leftarrow \Sigma \cup \Phi_i$;
6   Rename all var. in Φ_i (except the ϵ_i^j) with new ones;
7     $\Gamma_{\text{Hard}} \leftarrow \Gamma_{\text{Hard}} \cup \Phi_i$;
8 $\Sigma \leftarrow$Partial Max-SAT$(\Gamma_{\text{Soft}}, \Gamma_{\text{Hard}})$;
9 **return** $(\{\delta_i^j \in \mathcal{S} \mid \epsilon_i^j \in \Sigma\})$;

Proposition 2. *Let m be the number of clauses in $\mathcal{S} = [\Phi_1, \ldots, \Phi_n]$. The Transformational Approach computes one max$_{\#}$ consensus for $\mathcal{S}$. It requires one call to a Partial Max-SAT solver on a set of hard constraints made of mn clauses and a set of soft constraints made of m clauses.*

6 COMPUTING ONE MAX$_{\#}$ CONSENSUS IN S5

Assume now that $\forall i \in [1..n]$: $\Phi_i \subset \mathcal{L}_M$ and Φ_i is S5-satisfiable.

Consensus-finding grounds itself in multiple occurrences of satisfiability constraints since it amounts to finding a subset of formulas that is satisfiable with each source. Interestingly, as already mentioned, the satisfiability problem for S5 is NP-complete [11]: hence, there exists a polynomial transformation allowing this latter problem to be rewritten as SAT. Accordingly, there has been already much research about SAT-based deduction and satisfiability in S5 and other usual modal logics: see for example [8] and [9]. In the same vein, one way to compute consensuses in S5 can thus amount to translating the profile $\mathcal{S}$ in $\mathcal{L}_M$ into a profile in $\mathcal{L}$, while preserving (un)satisfiability.

The challenge is to devise such a transformation so that the practical efficiency of the approach in Boolean logic to extract one *max*$_{\#}$ consensus remains experimentally efficient. Especially, the number of additional variables and clauses that need to be introduced by the translation process into standard Boolean logic must be as minimal as possible to address large instances since the consensus-finding method in the Boolean framework itself roughly multiplies the size of the instance by the number of sources. Interestingly, the subsequent optimization technique to extract consensuses does not require the so-called *nominals* and other related concepts that are often introduced in the Boolean language by existing techniques, like for example [19], when modal S5 formulas are translated into Boolean ones. On the contrary, we use a plain direct translation, as proposed in [18], to transform the S5 modal formulas into equi-satisfiable CNF.

The first step consists in rewriting the formulas of $\mathcal{S}$ into modal CNF that preserves (un)satisfiability. We used the axiom schemata of S5, as well as De Morgan laws, which are valid in S5, to simplify nested occurrences of connectives and modalities, and transform each Φ_i into its Normal Negation From (NNF), which requires (1). $\neg$ to only occur immediately before a Boolean variable and (2). the absence of occurrences the $\rightarrow$ connectives. In a second step, the NNF was then transformed into a set of Boolean clauses using the Tseitin encoding technique [20], which introduces new variables to encode sub-formulas, and using other additional variables to encode the various possible worlds as follows. It is easy to see that it is sufficient to consider a number of different possible worlds that is bounded by the number of occurrences of modal operators in the NNF. Intuitively, a formula $\Diamond a$ is rewritten as $a_{pw1} \vee ... \vee a_{pwk}$ where a_{pwi} are new fresh variables referring to a in world i, and k is the total number of possible worlds that need be considered.

Notice that this transformation procedure of an S5 formula into an equi-satisfiable CNF standard Boolean logic is linear in the number of occurrences of modal operators in the initial formula. The number of Boolean variables in the NNF Boolean form depends on the number of variables in the initial modal formula with a proportional factor that is the number of involved possible worlds.

7 EXPERIMENTAL STUDY

All the experimentations have been conducted on Intel Xeon E5-2643 (3.30GHz) processors with 8Gb RAM with Linux CentOS. We used the Weighted Partial Max-SAT solver `MaxHS` from `http://www.maxhs.org/` and have implemented all the tested algorithms in C++ on top of *Glucose* (`http://www.labri.fr/perso/lsimon/glucose/`). All software, data and results are available at `http://cril.univ-artois.fr/consensus`.

7.1 $max_\#$ consensus in the Boolean framework

The profiles $\mathcal{S}$ in the standard Boolean case were based on the 291 (mostly real-world) unsatisfiable instances from the 2011 MUS competition `http://www.satcompetition.org/2011/` about the extraction of (set-inclusion) MUSes (Minimal Unsatisfiable Subsets). MUSes and MSSes are naturally related: each MUS is a minimal hitting set on the set of Co-MSSes (a Co-MSS is the set-theoretical complement of an MSS) whereas Co-MSS are hitting sets on the set of MUSes (see [12] for more on the use of this duality to compute MUSes). The selected instances are highly challenging: they are made of up to more than 15983000 clauses and 4426000 variables (457459 clauses using 139139 different variables on average). Each instance has been randomly divided into $n \in [3, 5, 7, 10]$ same-size (modulo n) sources Φ_i to yield all the $\mathcal{S}$. Time-out for each single $max_\#$ consensus extraction was set to 900 seconds.

Table 1 summarizes the average results for the extraction of one $max_\#$ consensus per value of n. It lists the number of successful extractions, the average time in seconds to extract one $max_\#$ consensus, the average numbers of clauses and variables in the transformed instance and, finally, the average number of clauses to drop to deliver the consensuses. A drop of performance can be observed when n increases (from 235 successful extractions to 207): this is due to both the increase of size of the representation of the transformed instance and additional satisfiability tests when n increases. These results show the viability of the approach and its scalability. Let us stress again that these benchmarks were selected to test extreme computational limits of the approach. Hopefully, these benchmarks should be harder and bigger than most real-life consensus-finding applications.

	$n=3$	$n=5$	$n=7$	$n=10$
#solved	235	223	210	207
time (seconds)	96	109	119	150
#variables	303643	329599	380110	460194
#clauses	1325632	1855884	2386137	3181517
#clauses$_{removed}$	7	2	2	2

Table 1. Computing one $max_\#$ consensus in the standard Boolean framework.

7.2 $max_\#$ consensus in S5

For the S5 logic framework, we have considered all the unsatisfiable modal logic benchmarks from `http://www.ps.uni-saarland.de/theses/goetzmann/` that we were able to split into $n \in [2, 3, 5]$ S5-satisfiable Φ_i using the following procedure (there were 93 such successfully split benchmarks). For each of them, we have built the various modal logic profiles $\mathcal{S}$ made of n sources as follows. $\mathcal{S}$ is initialized with one source made of the initial benchmark, which is treated as a unique formula. We define the size of a modal formula as the number of edges within its usual NNF representation. While the number of sources in S is less than n, the largest formula Σ in S is replaced by two formulas obtained from Σ in the following way: let Σ' be a sub-formula of Σ such that the size of Σ' is as close as possible to half the size of Σ. Then, we replace Σ' in Σ by a fresh variable $s_{\Sigma'}$ and we insert Σ' within $\mathcal{S}$ as an additional source, together with the information that $s_{\Sigma'} \equiv \Sigma'$. In the whole process, we make sure that all Φ_i are S5-satisfiable. Then, each resulting modal profile was translated into a standard Boolean logic one according to the aforementioned transformation technique. Finally, one $max_\#$ consensus was then searched using our transformational Partial Max-SAT-based technique.

One $max_\#$ consensus was delivered for all the 93 benchmarks when $n = 2$. One $max_\#$ consensus was found for 88 and 74 benchmarks when $n = 3$ and $n = 5$, respectively. We explain the drop of performance when n increases by the additional clauses that are needed to express the $s_{\Sigma'} \equiv \Sigma'$ constraints in NNF format, and by the fact that the Partial Max-SAT step increases the size of the representation by a factor n. Figure 1 illustrates the number of successful extractions of $max_\#$ consensuses according to the CPU time spent (in seconds) to compute such a consensus for these instances, and according to n.

Table 2 summarizes the average parameters values of the benchmarks and of their transformations after the different steps, as well as the average time spent in the process. The first column gives n. The four next ones list the main parameters of the initial modal formulas and of their corresponding modal profiles $\mathcal{S}$: namely, $\#vars$ is the average number of variables in the initial instance and $\#\{\Diamond,\Box\}$ is the average number of occurrences of modal operators; sum($|\Phi_i|$) is the average total size of the n Φ_i (size is the number of edges in the NNF representation) and, for convenience (since this can be computed from the previous columns) avg($|\Psi_i|$) gives the average size of each Φ_i. The next columns list the average values about the CNF transformation: namely, #vars gives the average number of Boolean variables used to encode the S5 formula into one CNF; sum($|\#\text{cls}_i|$) is the average total number of clauses in the standard logic profile whereas avg($|\#\text{cls}_i|$) is the average number of clauses in each

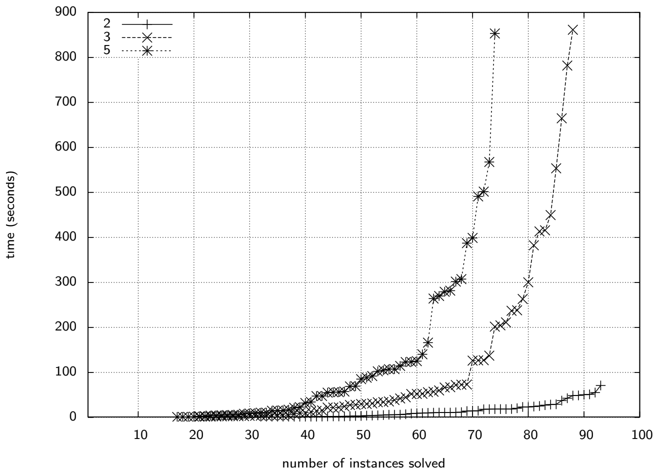

Figure 1. Number of modal profiles $\mathcal{S}$ for which one $max_{\#}$ consensus was extracted, depending on the CPU time spent and n.

Φ_i. Then, Time-Tr. gives the average CPU time in seconds spent to achieve the CNF encoding. The next group of columns provide average values about the Partial Max-SAT step: #vars and #cls are the average numbers of Boolean variables and clauses used in the encoding of the optimization problem. Finally, Time-Opt. is the average time in seconds to extract a consensus from the standard Boolean representation whereas Time-Opt. + Time-Tr. gives the average time required to compute the consensus plus the time spent for the translation process, when the whole operation was successful. Hence the values in this column slightly differ from the sum of values from the corresponding columns Time-Opt. and Time-Tr. when $n = 3$ and $n = 5$.

It is important to stress that the increase of size of the modal representation with n is simply due to the additional clauses that are needed to express equivalences during the splitting process to form the modal profiles $\mathcal{S}$. Accordingly, the actual "size" that is relevant is not the "size" of the modal benchmarks but the number of total clauses in the the modal $\mathcal{S}$, namely sum(Φ_i) in Table 2. For the same reason, the decrease of performance with the increase of n should not be interpreted as being only an increase of difficulty due to a larger number of sources, as the Partial Max-SAT step increases the size of the CNF representation by a factor n. Actually, this is also in part the expected consequence of having to consider larger modal profiles in these specific experimentations. The number of occurrences of modal operators in the benchmarks also plays a crucial role: it directly influences the size of the CNF representation. Actually, this parameter appears to the most limiting one, together with n. Accordingly, we believe that these experimentations show the viability of the approach for real-life applications, whose size in terms of formulas, numbers of sources and of occurrences of modal operators remain in the range of the ones in the tested benchmarks.

8 CONCLUSION AND PERSPECTIVES

Consensus-finding is a ubiquitous issue in real-life and in many A.I.-related applications. In this paper, a concept of consensus has been investigated, with a focus on practical computational issues. This concept of consensus does not merely capture what is shared by several agents: it provides a satisfiable fragment of all the information conveyed by the agents that is satisfiable with every agent. Hence, such a consensus might be endorsed by the agents as it does not contradict them. The paper focused on the most appealing consensuses, namely maximal ones. They have been studied in both standard Boolean logic and in modal logic S5, as the necessity and possibility representation paradigms are highly relevant in consensus-related domains, like negotiation. Interestingly, we have proposed and experimented a single-step optimization technique that delivers one maximal consensus in an efficient way very often, for both logical frameworks.

At this point, this study opens paths for various promising further research. Let us here simply mention three of them. First, a natural extension of this study would consist in allowing various maximality preferences in consensuses. For example, an agent might pre-order her desires according to her priorities and any consensus should attempt to obey these preferences. Interestingly, the use of *weighted* Partial Max-SAT vs. Partial Max-SAT could allow for a direct handling of these kinds of preferences, not only in the standard Boolean framework as this has been done in [5], but also in the S5 logic. Another natural challenging issue would be the extension of the concept of consensus for a multiple-agents modal S5. However, extending our transformational approach accordingly does not seem viable since satisfiability in such a logic is P-SPACE complete, making a polynomial translation into SAT out of reach. Finally, a promising challenge would be to build more elaborate transformation procedures for the modal logic framework, so that the structure of the initial formulas is not lost in the CNF representation and is fully exploited in the checks for satisfiability during the subsequent operations in the search for a maximal consensus.

ACKNOWLEDGEMENTS

We would like to thank the referees for their useful comments, which helped improve the presentation of this paper. All experimentations have been conducted on a cluster that has been funded in part by the *Conseil Régional du Nord/Pas-de-Calais* and an EC FEDER grant.

APPENDIX

The language $\mathcal{L}_M$ of modal logic S5 extends the language $\mathcal{L}$ of Boolean logic by allowing two modalities $\Diamond$ and $\Box$ to be used as additional unary connectives and that are dual in the sense that $\Box \stackrel{def}{=} \neg\Diamond\neg$. An axiomatic system for S5 is given by the following axiom schemata and rules:

All Boolean logic tautologies. $\mathcal{K} : \Box(A \to B) \to (\Box A \to \Box B)$
$\mathcal{T} : \Box A \to A \quad \mathcal{B} : A \to \Box\Diamond A \quad 4 : \Box A \to \Box\Box A$
(or $5 : \Diamond A \to \Box\Diamond A$ instead of $\mathcal{B}$ and 4)

$$MP : \frac{A \quad (A \to B)}{B} \qquad Nec : \frac{A}{\Box A}$$

The *possible worlds* semantics of S5 is based on Kripke frames, which are pairs (W, R) where W is a non-empty set of possible worlds and R is an accessibility relation between worlds. In S5, the relation R is reflexive, transitive and symmetric, i.e, is an equivalence relation. We note wRw' to express that w' is accessible from w. A valuation V assigns a subset of W to each atomic proposition p: namely, the worlds where p is *true*. Given a Kripke frame (W, R), a valuation V and a world w, the satisfaction relation $\models$ is defined inductively as follows.

$[(W, R), V, w] \models p$ iff $w \in V(p)$

	Initial modal formulas and modal profiles				CNF transformation step				Partial Max-SAT step			Total (secs.)
n	#vars	#{$\Diamond$,$\Box$}	sum($\lvert\Phi_i\rvert$)	avg($\lvert\Phi_i\rvert$)	#vars	sum($\lvert$#cls$_i\rvert$)	avg($\lvert$#cls$_i\rvert$)	Time-Tr.	#vars	#cls	Time-Opt.	Time-Tr. + Time-Opt.
2	80	85	464	232	22022	279983	139991	1.28	324029	1119933	10.30	11.58
3	81	101	569	189	68695	632146	210715	2.07	838233	3160731	90.64	92.49
5	83	106	640	128	114239	924416	184883	2.70	1495611	6470913	96.08	97.90

Table 2. Computing one $max_{\#}$ consensus in the S5 framework.

$[(W,R),V,w] \models \neg p$ iff $w \notin V(p)$
$[(W,R),V,w] \models A \wedge B$ iff $[(W,R),V,w] \models A$ and $[(W,R),V,w] \models B$
$[(W,R),V,w] \models A \vee B$ iff $[(W,R),V,w] \models A$ or $[(W,R),V,w] \models B$
$[(W,R),V,w] \models \Diamond A$ iff $\exists w' \in W$ s.t. wRw': $[(W,R),V,w'] \models A$
$[(W,R),V,w] \models \Box A$ iff $\forall w' \in W$ s.t. wRw': $[(W,R),V,w'] \models A$

The S5 satisfiability problem is: given a formula α of $\mathcal{L}_M$, determine whether there exists a Kripke frame (W,R), a valuation V and a world $w \in W$ s.t. $[(W,R),V,w] \models \alpha$.

The deduction relation in S5 is noted $\models_{S5}$.

REFERENCES

[1] Philippe Besnard, Éric Grégoire, and Jean-Marie Lagniez, 'On computing maximal subsets of clauses that must be satisfiable with possibly mutually-contradictory assumptive contexts', in *Proceedings of the Twenty-Ninth National Conference on Artificial Intelligence AAAI '15*, pp. 3710–3716, (2015).

[2] Patrick Blackburn, M. de Rijke, and Y. Venema, *Modal Logic*, Cambridge University Press, 2001.

[3] Brian F. Chellas, *Modal Logic*, Cambridge University Press, 1980.

[4] Eithan Ephrati and Jeffrey S. Rosenschein, 'Deriving consensus in multiagent systems', *Artificial Intelligence*, **87**(1-2), 21–74, (1996).

[5] Éric Grégoire, Sébastien Konieczny, and Jean-Marie Lagniez, 'On consensus extraction', in *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI'16)*, (2016).

[6] Zhaohui Fu and Sharad Malik, 'On solving the partial MAX-SAT problem', in *Proceedings of the 9th International Conference on Theory and Applications of Satisfiability Testing (SAT 2006)*, pp. 252–265, (2006).

[7] Olivier Gauwin, Sébastien Konieczny, and Pierre Marquis, 'Conciliation through iterated belief merging', *Journal of Logic and Computation*, **17**(5), 909–937, (2007).

[8] E. Giunchiglia, F. Giunchiglia, R. Sebastiani, and A. Tacchella, 'SAT vs. translation based decision procedures for modal logics: a comparative evaluation', *Journal of Applied Non-Classical Logics*, **10**(2), (2000).

[9] E. Giunchiglia, A. Tacchella, and F. Giunchiglia, 'SAT-based decision procedures for classical modal logics', *Journal of Automated Reasoning*, **28**(2), 143–171, (2000).

[10] Audun Jøsang, 'The consensus operator for combining beliefs', *Artificial Intelligence*, **141**(1/2), 157–170, (2002).

[11] R.E. Ladner, 'The computational complexity of provability in systems of modal propositional logic', *SIAM Journal of Computing*, **6**(3), 467–480, (1977).

[12] Mark H. Liffiton and Karem A. Sakallah, 'On finding all minimally unsatisfiable subformulas', in *Proceedings of the 8th International Conference on Theory and Applications of Satisfiability Testing (SAT'05)*, volume 3569 of *Lecture Notes in Computer Science*, pp. 173–186. Springer, (2005).

[13] Joao Marques-Silva and Mikolás Janota, 'On the query complexity of selecting few minimal sets', *Electronic Colloquium on Computational Complexity (ECCC)*, **21**, 31, (2014).

[14] Alexander Nadel, 'Boosting minimal unsatisfiable core extraction', in *Proceedings of 10th International Conference on Formal Methods in Computer-Aided Design (FMCAD 2010)*, pp. 221–229, (2010).

[15] Nina Narodytska and Fahiem Bacchus, 'Maximum satisfiability using core-guided maxsat resolution', in *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI'14)*, pp. 2717–2723, (2014).

[16] Christos H. Papadimitriou and Mihalis Yannakakis, 'Optimization, approximation, and complexity classes', *Journal of Computer and System Sciences*, **43**(3), 425 – 440, (1991).

[17] Wei Ren, R.W. Beard, and E.M. Atkins, 'A survey of consensus problems in multi-agent coordination', in *Proceedings of the 2005 American Control Conference*, volume 3, pp. 1859–1864, (2005).

[18] Yakoub Salhi, Saïd Jabbour, and Lakhdar Sais, *Information Search, Integration and Personalization: International Workshop, ISIP 2012, Sapporo, Japan, October 11-13, 2012. Revised Selected Papers*, chapter Graded Modal Logic GS5 and Itemset Support Satisfiability, 131–140, Springer, 2013.

[19] Yakoub Salhi and Michael Sioutis, 'A resolution method for modal logic S5', in *Proceedings of the Global Conference on Artificial Intelligence (CCAI 2015)*, volume 36 of *EPiC Series in Computer Science*, pp. 152–262, (2015).

[20] G.S. Tseitin, *Structures in Constructives Mathematics and Mathematical Logic, Part II*, chapter On the Complexity of Derivations in the Propositional Calculus, 115–125, Steklov Mathematical Institute, 1968.