Using Local Search to Find and MUSes

Éric Grégoire, Bertrand Mazure, Cédric Piette

► To cite this version:

Éric Grégoire, Bertrand Mazure, Cédric Piette. Using Local Search to Find and MUSes. European Journal of Operational Research, 2009, 199 (3), pp.640-648. hal-00868114

HAL Id: hal-00868114

<https://hal.science/hal-00868114>

Submitted on 1 Oct 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Using Local Search to find MSSes and MUSes

Éric Grégoire, Bertrand Mazure and Cédric Piette

Université Lille-Nord de France, Artois, F-62307 Lens

CRIL, F-62307 Lens

CNRS UMR 8188, F-62307 Lens

rue Jean Souvraz SP18 F-62307 Lens France

Abstract

In this paper, a new complete technique to compute Maximal Satisfiable Subsets (MSSes) and Minimally Unsatisfiable Subformulas (MUSes) of sets of Boolean clauses is introduced. The approach improves the currently most efficient complete technique in several ways. It makes use of the powerful concept of critical clause and of a computationally inexpensive local search oracle to boost an exhaustive algorithm proposed by Liffiton and Sakallah. These features can allow exponential efficiency gains to be obtained. Accordingly, experimental studies show that this new approach outperforms the best current existing exhaustive ones.

Key words: SAT, MSSes, MUSes, satisfiability, hybrid algorithm

1 Introduction

This last decade, the SAT problem, namely the issue of checking whether a set of Boolean clauses is satisfiable or not, has received much attention from the AI research community. Indeed, SAT appears to be a cornerstone in many domains, like e.g. nonmonotonic reasoning, automated reasoning, model-based diagnosis, planning and knowledge bases verification and validation. However, only knowing that a SAT instance is unsatisfiable is often not satisfactory since we might prefer knowing what goes *wrong* with the instance when this latter one is expected to be satisfiable.

In this respect, the MUS (*Minimally Unsatisfiable Subformula*) concept can be crucial since a MUS can be seen as an irreducible cause for infeasibility. Indeed, a

Email addresses: gregoire@cril.fr (Éric Grégoire), mazure@cril.fr (Bertrand Mazure), piette@cril.fr (Cédric Piette).

MUS is an unsatisfiable set of clauses such that any of its subsets is satisfiable. It thus provides one explanation for unsatisfiability that cannot be made shorter in terms of the number of involved clauses. Restoring the satisfiability of an instance cannot be done without fixing all its MUSes.

Unfortunately, a same instance can exhibit several MUSes. Actually, the number of these MUSes can be exponential since a n -clauses SAT instance can exhibit $C_n^{n/2}$ MUSes in the worst case. Moreover, computing MUSes is intractable in the general case. Indeed, checking whether a set of clauses is a MUS or not is DP-complete [1] and checking whether a formula belongs to the set (clutter) of MUSes of an unsatisfiable SAT instance or not, is in Σ_2^P [2]. Fortunately, the number of MUSes remains often tractable in real-life applications. For example, in model-based diagnosis [3], it is often assumed that single faults occur most often, which can entail small numbers of MUSes.

A dual concept is the notion of *Maximal Satisfiable Subset* (MSS) of a SAT instance, and the complement of a MSS in a SAT instance is called a CoMSS. The complete set of MUSes or MSSes is an implicit encoding of the other [4]. Specifically, a CoMSS is a hitting set of the set of MUSes and represent minimal sets of clauses that should be dropped in order to restore consistency. In this paper, we are interested in exhaustive approaches to compute these three correlated concepts in the full Boolean clausal framework.

Recently, several approaches have been proposed to approximate or compute MUSes and MSSes, both in the Boolean framework and for other types of constraints. Some of them concern specific classes of clauses or remain tractable for small instances, only. Among them, let us mention the approach in [5], where it is shown how a MUS can be extracted in polynomial time through linear programming techniques for clauses exhibiting a so-called integral property. However, only restrictive classes of clauses obey such a property (mainly Horn, renamable Horn, extended Horn, balanced and matched ones). Let us also mention [6][7][8], which contain studies of the complexity and the algorithmic aspects of extracting MUSes for specific classes of clauses. In [9], an approach is proposed that approximates MUSes by means of an adaptative search guided by clauses hardness. In [10] a technique is described, that extracts MUSes by learning nogoods involved in the derivation of the empty clause by resolution. In [11], a complete and exhaustive technique to extract smallest MUSes is introduced. In [12], a DPLL-oriented approach has been presented that is based on a marked clauses concept to allow one to approximate MUSes. In [13], Grégoire, Mazure and Piette have proposed a heuristic-based incomplete approach to compute MUSes, which outperforms competing ones from a computational point of view.

Interestingly, in [14] the same authors have introduced a concept of inconsistent covers to circumvent the possible intractable number of MUSes, and presented a technique to compute them. Roughly, an inconsistent cover of an unsatisfiable

SAT instance represents a set of MUSes that covers enough independent causes of inconsistency that would allow the instance to regain consistency if they were repaired. Although an inconsistent cover does not provide us with the set of all MUSes that may be present in a formula, it does however provide us with a series of minimal explanations of inconsistency that are sufficient to explain and potentially "fix" enough causes of inconsistency in order for the whole instance to regain consistency.

These latter techniques are incomplete ones in the sense that they do not necessarily deliver all MUSes. However, in some application domains, it can be necessary to find the set of *all* MUSes, because diagnosing infeasibility is hard, if not impossible, without a complete view of its causes [4]. Obviously enough, such techniques can only remain tractable provided that the number of MUSes remains itself tractable. Likewise, the number of MSSes and CoMSSes can be exponential in the worst case. It should be noted that many domains in Artificial Intelligence like belief revision (see e.g. [15]) involve conceptual approaches to handle unsatisfiability that can require the complete sets of MUSes, MSSes, and CoMSSes to be computed in the worst case, even when additional epistemological ingredients like e.g. stratification are introduced in the logical framework.

In this paper, the focus is on complete techniques. We introduce a new complete technique to compute all MUSes, MSSes and CoMSSes of a SAT instance, provided obvious tractability limitations. It improves the currently most efficient complete technique, namely Liffiton and Sakallah's one [4] (in short L&S), which in turn was shown more competitive than previous approaches by Bailey and Stuckey [16], and by de la Banda, Stuckey and Wazny [17], which were introduced in somewhat different contexts.

Our approach exhibits two main features. First, it is a hybridization of the L&S complete approach with a local search pretreatment. A local search technique is indeed used as an oracle to find potential CoMSSes of the SAT instance, which are themselves hitting sets of MUSes. We show that such a hybridization can yield exponential efficiency gains. Second, the efficiency of the approach relies on the crucial concept of critical clause, which appears to be a powerful ingredient of our technique to locate MUSes.

The rest of the paper is organized as follows. First, the reader is provided with the necessary background about SAT, MUSes and the dual concepts of MSSes and CoMSSes. Then, Liffiton and Sakallah's exhaustive approach is briefly presented. In Section 4, we show how this technique can be valuably hybridized with a local search pretreatment, making use of the critical clause concept. It is shown how this pretreatment can be theoretically valuable from a computational point of view. In Section 5, we compare this new approach with Liffiton and Sakallah's one on various benchmarks.

2 Background

In this section, the reader is provided with basic notions about SAT, MUSes, MSSes and CoMSSes.

Let $\mathcal{L}$ be a standard Boolean logical language built on a finite set of Boolean variables, noted a, b, c , etc. The $\wedge, \vee, \neg$ and $\Rightarrow$ symbols represent the standard conjunctive, disjunctive, negation and material implication connectives, respectively. Formulas and clauses will be noted using upper-case letters such as C . Sets of formulas will be represented using Greek letters like Γ or Σ . An interpretation is a truth assignment function that assigns values from $\{true, false\}$ to every Boolean variable. A formula is satisfiable when there is at least one interpretation (called model) that satisfies it, i.e. that makes it become *true*. An interpretation will be noted by upper-case letters like I and will be represented by the set of literals that it satisfies. Actually, any formula in $\mathcal{L}$ can be represented (while preserving satisfiability) using a set (interpreted as a conjunction) of clauses, where a clause is a finite disjunction of literals, where a literal is a Boolean variable that is possibly negated. SAT is the NP-complete problem that consists in checking whether a set of Boolean clauses is satisfiable or not, i.e. whether there exists an interpretation that satisfies all clauses in the set or not.

When a SAT instance is unsatisfiable, it exhibits at least one *Minimally Unsatisfiable Subformula*, in short one *MUS*.

Definition 1 A MUS Γ of a SAT instance Σ is a set of clauses s.t.

- (1) $\Gamma \subseteq \Sigma$
- (2) Γ is unsatisfiable
- (3) $\forall \Delta \subset \Gamma, \Delta$ is satisfiable

Example 2 Let $\Sigma = \{a, \neg c, \neg b \vee \neg a, b, \neg b \vee c\}$. Σ exhibits two MUSes, namely $\{a, b, \neg b \vee \neg a\}$ and $\{\neg c, b, \neg b \vee c\}$.

A dual concept is the notion of *Maximal Satisfiable Subset* (MSS) of a SAT instance.

Definition 3 A MSS Γ of a SAT instance Σ is a set of clauses s.t.

- (1) $\Gamma \subseteq \Sigma$
- (2) Γ is satisfiable
- (3) $\forall \Delta \subseteq (\Sigma \setminus \Gamma)$ s.t. $\Delta \neq \emptyset, \Gamma \cup \Delta$ is unsatisfiable.

The set-theoretical complement of a MSS w.r.t. a SAT instance is called a *CoMSS*.

Definition 4 The CoMSS of a MSS Γ of a SAT instance Σ is given by $\Sigma \setminus \Gamma$

Example 5 *Let us consider the formula Σ from the previous example. Σ exhibits five MSSes: $\{b\}$, $\{\neg c, a\}$, $\{\neg c, \neg b \vee c\}$, $\{\neg b \vee c, \neg b \vee \neg a\}$, $\{\neg c, \neg b \vee \neg a\}$*

As shown by several authors [4], these concepts are correlated. Mainly, a CoMSS contains at least one clause from each MUS. Actually, a CoMSS is an irreducible hitting set of the set of MUSes. In a dual way, every MUS of a SAT instance is an irreducible hitting set of the CoMSSes. Accordingly, as emphasized by [4], although MINIMAL-HITTING-SET is a NP-hard problem, irreducibility is a less strict requirement than minimal cardinality. Actually, a MUS can be generated in polynomial time from the set of CoMSSes.

3 Liffiton and Sakallah's Exhaustive Approach

Liffiton and Sakallah's approach [4] to compute all MUSes (in short L&S) is based on the strong duality between MUSes and MSSes. To the best of our knowledge, it is currently the most efficient one. First it computes all MSSes before it extracts the corresponding set of MUSes. Here, the focus is on L&S first step since we shall improve it and adopt the second step as such.

L&S is integrated with a modern SAT solver and takes advantage of it. Roughly, every i th clause $C_i = x_1 \vee \dots \vee x_m$ of the SAT instance is augmented with a negated clause selector variable y_i to yield $C'_i = x_1 \vee \dots \vee x_m \vee \neg y_i$. While solving these new clauses, assigning y_i to *false* has the effect of disabling or removing C_i from the instance. Accordingly, a MSS can be obtained by finding a satisfying assignment with a minimal number of y_i variables assigned *false*. The algorithm makes use of a sliding objective approach allowing for an incremental search. A bound on the number of y_i that may be assigned to *false* is set. For each value of the bound, starting at 0 and incrementing by 1, an exhaustive search is performed for all satisfiable assignments to the augmented formula C'_i , which will find all CoMSSes having their size equal to the bound. Whenever one solution is found, it is recorded, and a corresponding clause forcing out that solution (and any supersets of it) is inserted. This blocking clause is a disjunction of the y_i variables for the clauses in that CoMSS.

Before beginning the search with the next bound, the algorithm checks that the new instance augmented with all the blocking clauses is still satisfiable without any bound on the y_i variables. If there is no such satisfying assignment, the entire set of CoMSSes has been found and the algorithm terminates.

The second part of the algorithm computes the complete set of MUSes from the set of CoMSSes in a direct way. The approach that we shall introduce will include this second step as such.

4 Local Search and Critical Clauses

In this section, it is shown how the aforementioned exhaustive search algorithm can be improved in a dramatic manner by hybridizing it with an initial local search step, which provides valuable oracles for the subsequent exhaustive search process. We shall call the new approach HYCAM (HYbridization for Computing All Muses).

First, let us motivate our approach in an intuitive manner. Clearly, a (fast) initial local search run for satisfiability on the initial instance might encounter some actual MSSes. Whenever this phenomenon happens, it can prove valuable to record the corresponding CoMSSes in order to avoid computing them during the subsequent exhaustive search. Moreover, rather than checking whether we are faced with an actual MSS or not, it can prove useful to record the corresponding candidate CoMSS that will be checked later during the exhaustive search. Obviously enough, we must study which interpretations encountered during the local search process yield candidate MSSes and criteria must be defined in order to record a limited number of potentially candidate CoMSSes only. In this respect, a concept of critical clause will prove extremely valuable in the sense that it allows us to state necessary conditions for being a CoMSS that can be checked quickly. When all the remaining candidate CoMSSes are recorded, the incremental approach by Liffiton and Sakallah allows us to exploit this information in a very valuable and efficient way. Let us describe this in a more detailed manner.

A local search algorithm is thus run on the initial SAT instance. The goal is to record as many candidate CoMSSes as possible, based on the intuitive heuristics that local search often converges towards local minima, which could translate possibly good approximations of MSSes. A straightforward approach would consist in recording for each visited interpretation the set of unsatisfied clauses. Obviously enough, we do not need to record supersets of already recorded candidate CoMSSes since they cannot be actual CoMSSes as they are not minimal with respect to set-theoretic inclusion. More generally, we have adapted the technique proposed by Zhang in [18] to sets of clauses in order to record the currently smaller candidates CoMSSes among the already encountered series of sets of unsatisfied clauses. Now, crucial ingredients in our approach are the concepts of once-satisfied and critical clauses. Moreover, we have also exploited the following concept of critical clause, which has already proved valuable for locating MUSes and inconsistent covers using an incomplete technique based on local search [13][14].

Definition 6 *A clause C is once-satisfied by an interpretation I iff exactly one literal of C is satisfied by I . A clause C that is falsified by the interpretation I is critical w.r.t. I iff the opposite of each literal of C belongs to at least one once-satisfied clause by I .*

Intuitively, a critical clause is thus a falsified clause that requires at least another

Algorithm 1: Local Search approximation

Input: a CNF formula Σ **Output:** Set of candidate CoMSS**begin**

```
  candidates  $\leftarrow \emptyset$  ;  
  #fail  $\leftarrow 0$  ;  
   $I \leftarrow \text{generate\_random\_interpretation}()$  ;  
  while (#fail < PRESET\_#FAILURES.AUTHORIZED) do  
    newcandidates  $\leftarrow$  FALSE ;  
    for  $j = 1$  to #FLIPS do  
      Let  $\Delta$  be the set of falsified clauses by  $I$  ;  
      if  $\forall C \in \Delta, C$  is critical and  $\Delta$  is not implied in candidates then  
        removeAllSetImplied( $\Delta$ , candidates) ;  
        candidates  $\leftarrow \Delta \cup$  candidates ;  
        newcandidates  $\leftarrow$  TRUE ;  
      flip( $I$ ) ;  
    if not(newcandidates) then #fail  $\leftarrow$  #fail + 1 ;  
  return candidates ;
```

end

clause to be falsified in order to become satisfied, performing a flip. The following proposition shows how this concept allows us to eliminate wrong candidates CoMSSes.

Proposition 7 *Let Σ be a SAT instance and let I be an interpretation. Let Γ be a non-empty subset of Σ s.t. all clauses of Γ are all falsified by I . When at least one clause of Γ is not critical w.r.t. I , then Γ is not a CoMSS of Σ .*

PROOF. By definition, when a clause C_f of Γ is not critical w.r.t. I , there exists at least one literal $c \in C_f$ whose truth-value can be inversed (i.e. flipped) without falsifying any other clause of Σ . Accordingly, Γ is not minimal and cannot be a CoMSS of Σ . $\square$

In practice, testing whether all falsified clauses are critical or not can be performed quickly and prevents many sets of clauses to be recorded as candidate CoMSSes.

Using these features, the local search run on the initial SAT instance, as described in Algorithm 1, yields a series of candidate CoMSSes. This information proves valuable and allows us to boost L&S complete search.

L&S is incremental in the sense that it computes CoMSSes of increasing sizes, progressively. After n iterations have been performed, all CoMSSes of cardinality lower or equal than n have been obtained. Accordingly, if we have recorded candidate CoMSSes containing $n + 1$ clauses, and if they are not supersets of already

obtained CoMSSes, we are sure that they are actual CoMSSes. In this respect, we do not need to search them, and their corresponding blocking clauses can be inserted directly. Moreover, we do not need to perform the SAT test at the end of the n -th iteration, since we are then aware of the existence of larger CoMSSes.

It is also easy to show that the insertion of these blocking clauses can allow both NP-complete and CoNP-complete tests to be avoided. Let us illustrate this on an example.

Example 8 *Let Σ be the following SAT instance.*

$$\Sigma = \left\{ \begin{array}{l} c_0 : (d) \\ c_1 : (b \vee c) \\ c_2 : (a \vee b) \\ c_3 : (a \vee \neg c) \\ c_4 : (\neg b \vee \neg e) \\ c_5 : (\neg a \vee \neg b) \\ c_6 : (a \vee e) \\ c_7 : (\neg a \vee \neg e) \\ c_8 : (b \vee e) \\ c_9 : (\neg a \vee b \vee \neg c) \\ c_{10} : (\neg a \vee b \vee \neg d) \\ c_{11} : (a \vee \neg b \vee c) \\ c_{12} : (a \vee \neg b \vee \neg d) \end{array} \right.$$

Σ is an unsatisfiable SAT instance made of 13 clauses and making use of 5 variables. It exhibits 3 MUSes, which are illustrated in Figure 1, and admits 19 MSSes. Assume that both L&S and HYCAM are run on this instance. Its clauses are augmented by the $\neg y_i$ negated clause selector variables. Assume also that the local search performed by HYCAM provides 4 candidate CoMSSes: $\{c_5\}$, $\{c_3, c_2\}$, $\{c_0, c_1, c_2\}$ and $\{c_3, c_8, c_{10}\}$.

If the branching variables are chosen based on the lexical order, then a and b are assigned to *true* and c_5 is falsified. Thus, L&S tries to prove that this clause forms a CoMSS, which requires a NP-complete test (because it has to find a model of $\Sigma \setminus \{\neg a \vee \neg b\} \cup \{a, b\}$). On the contrary, when HYCAM is run, the blocking clause y_5 is added before the first iteration of the complete algorithm is performed, since y_5 is the clause selector variable of c_5 and since the local search has already delivered

Fig. 1. MUSes of Example 8

this CoMSS. In consequence, when a and b are assigned *true*, the DPLL-algorithm backtracks immediately as the $\{y_5, \neg y_5\}$ unsatisfiable set has been obtained, without requiring any further NP-complete test.

Similarly, the introduction of additional clause selector variables by HYCAM can reduce the number of CoNP-complete tests. For example, let us assume that e is the first branching variable, that e is assigned *false* and that the next variables are selected according to the lexical order. When a and b are assigned *true*, L&S tries to prove that $\{c_5\}$ is a CoMSS. Since $\neg e$ is tautological consequence of $\Sigma \setminus \{\neg a \vee \neg b\} \cup \{a, b\}$, no model exists for $\Sigma \setminus \{\neg a \vee \neg b\} \cup \{a, b, \neg e\}$. Clearly, such a test is in CoNP. Thanks to the previously delivered candidate CoMSSes, HYCAM avoids this part of the search space to be explored. Indeed, since we know that $\{c_5\}$ is a CoMSS, when a and b are assigned *true*, no further CoNP tests are performed with respect to this partial assignment.

In fact, from a computational point of view, the preliminary non-expensive local search eliminates nodes in the search tree, avoiding both NP and CoNP tests.

The HYCAM algorithm is depicted in Algorithm 2.

5 Experimental Evaluation

HYCAM has been implemented and compared to L&S from a practical point of view. For both algorithms, the complete search step is based on the use of MiniSat [19], which is currently one of the best modern SAT solvers. As a case study, we used Walksat [20] for the local search pretreatment. The number of flips and tries of Walksat is related to the number of candidate CoMSS already found. For each try, a small number of flips is performed. If no new candidate is found during a try then a counter is incremented. When this counter exceeds a threshold (experimentally

Algorithm 2: The HYCAM algorithm

Input: a CNF formula Σ **Output:** All MSSes of Σ **begin** $cand \leftarrow \text{LS_approximation}(\Sigma); \quad /* \quad \text{algorithm 1} \quad */$ $\Sigma_y \leftarrow \text{addSelectorClauses}(\Sigma);$ $k \leftarrow 0;$ $MSS \leftarrow \emptyset;$ **while** $\text{SAT}(\Sigma_y)$ **do** $\text{removeAllSetImplied}(\{\Sigma \setminus C \mid C \in MSS\}, cand);$ $\Sigma_y \leftarrow \text{addBlockingClausesOfSize}(k, cand);$ $MSS \leftarrow MSS \cup \{\Sigma \setminus C \mid C \in cand \text{ and } |C| = k\};$ $MSS \leftarrow MSS \cup \text{SAT_with_bound}(k, \Sigma_y);$ $k \leftarrow k + 1;$ **return** $MSS;$ **end**

set to 30), we consider that no new candidate could be found by the local search. This way to end the local search pretreatment offers a good trade-off between the number of candidates found and the time spent. Besides, for all experiments, the time consumed by the local search step was less than 5% of the global time. All our experimental studies have been conducted on Intel Xeon 3GHz under Linux CentOS 4.1. (kernel 2.6.9) with a RAM memory size of 2Go. In the following, a time-out limit has been set to 3 CPU hours.

First, in Table 1, we report experimental results about the computation of MSSes on pigeon-hole and xor-chains benchmarks [21], which are globally unsatisfiable in the sense that removing any one of their clauses makes the instance become satisfiable. Obviously enough, such instances exhibit a number of CoMSSes equals to their number of clauses, and the size of any of their CoMSS is one. A significant time gap can be observed in favor of HYCAM. The efficiency gain ratio is even more significant when the size of the instance increased. For these instances, the local search run often succeeds in finding all CoMSSes, and the complete step often reduces to an unsatisfiability test. On the contrary, L&S explores many more nodes in the search space to deliver the CoMSSes.

In Table 2, experimental results on more difficult benchmarks from the annual SAT competition [21] are described. Their number of MSSes is often exponential, and computing them often remains intractable. Accordingly, we have limited the search to CoMSSes of restricted sizes, namely we have set a size limit to 5 clauses. As our experimental results illustrate, HYCAM outperforms L&S. For example, let us consider `rand_net40-30-10`. This instance contains 5831 MSSes (with the size of their corresponding CoMSSes less than 5). L&S and HYCAM deliver this result in 1748 and 386 seconds, respectively. For the `ca256` instance, HYCAM has extracted 9882 MSSes in less than 5 minutes whereas L&S did not manage to

Instance	(#var,#cla)	#MSSes	L&S (sec.)	HYCAM (sec.)
hole6	(42,133)	133	0.040	0.051
hole7	(56,204)	204	0.75	0.33
hole8	(72,297)	297	33	1.60
hole9	(90,415)	415	866	30
hole10	(110,561)	561	7159	255
x1_16	(46,122)	122	0.042	0.041
x1_24	(70,186)	186	7.7	0.82
x1_32	(94,250)	250	195	28
x1_40	(118,314)	314	2722	486

Table 1
L&S vs. HYCAM on globally unsatisfiable instances

produce this result within 3 hours. Let us note that HYCAM also delivers CoMSSes made of 5 clauses after its computation is ended since we know that all sets of 5 falsified clauses recorded by the local search run and that are not supersets of the obtained smaller CoMSSs are actually also CoMSSes.

In Table 3, experimental results on hard instances to compute the complete set of MSSes and MUSes are reported. As explained above, both L&S and HYCAM approaches require all MSSes to be obtained before MUSes are computed. By allowing complete sets of MSSes to be delivered in a shorter time, HYCAM allows the complete set of MUSes to be computed for more instances and in a faster manner than L&S does. Obviously enough, when the number of MSSes or MUSes are exponential, computing and enumerating all of them remain intractable.

For instance, L&S was unable to compute all MSSes of the php-012-011 instance within 3 hours CPU time, and could thus not discover its single MUS. HYCAM extracted it in 2597 seconds. On all instances exhibiting unique or a non-exponential number of MUSes, HYCAM was clearly more efficient than L&S. For example, on the dlx2_aa instance, L&S and HYCAM discovered the 32 MUSes within 3.12 and 0.94 seconds, respectively. Let us note that the additional time spent to compute all MUSes from the set of MSSes is often very small unless of course the number of MUSes is exponential.

Instance	(#var,#cla)	#MSSes	L&S (sec.)	HYCAM (sec.)
rand_net40-25-10	(2000,5921)	5123	893	197
rand_net40-25-5	(2000,5921)	4841	650	174
rand_net40-30-10	(2400,7121)	5831	1748	386
rand_net40-30-1	(2400,7121)	7291	1590	1325
rand_net40-30-5	(2400,7121)	5673	2145	402
ca032	(558,1606)	1173	4	1
ca064	(1132,3264)	2412	59	3
ca128	(2282,6586)	4899	691	18
ca256	(4584,13236)	9882	<i>time out</i>	277
2pipe	(892,6695)	3571	130	36
2pipe_1_ooo	(834,7026)	3679	52	30
2pipe_2_ooo	(925,8213)	5073	148	61
3pipe_1_ooo	(2223,26561)	17359	5153	2487
am_5_5	(1076,3677)	1959	68	57
c432	(389,1115)	1023	4	1
c880	(957,2590)	2408	28	3
bf0432-007	(1040,3668)	10958	233	98
velev-sss-1.0-cl	(1453,12531)	4398	1205	513

Table 2
L&S vs. HYCAM on various difficult SAT instances

6 Conclusions and Future Research

Computing all MSSes, CoMSSes and MUSes are highly intractable issues in the worst case. However, it can make sense to attempt to compute them for some real-life applications. In this paper, we have improved the currently most efficient exhaustive technique, namely Liffiton and Sakallah’s method, in several ways. Our experimental results show dramatic efficiency gains for MSSes, CoMSSes and MUSes extracting. One interesting feature of the approach lies in its anytime character for computing MSSes. MSSes of increasing sizes are computed gradually. Accordingly, we can put a bound on the maximum size of the CoMSSes to be extracted, limiting the computing resources needed to extract them. To some extent, both L&S and HYCAM prove more adapted to extract complete sets of MSSes and CoMSSes than complete sets of MUSes. Indeed, the procedure involves computing MSSes (and thus CoMSSes) first. In this respect, we agree with Liffiton

Instance	(#var,#cla)	#MSSes	L&S (sec.)	HYCAM (sec.)	#MUSes	MSSes→MUSes (sec.)
mod2-3cage-unsat-9-8	(87, 232)	232	3745	969	1	0.006
mod2-rand3bip-unsat-105-3	(105, 280)	280	2113	454	1	0.008
2pipe	(892, 6695)	10221	298	226	> 211 000	<i>time out</i>
php-012-011	(132, 738)	738	<i>time out</i>	2597	1	0.024
hcb3	(45, 288)	288	10645	6059	1	0.006
ldlx_c_mc_ex_bp_f	(776, 3725)	1763	10.4	6.8	> 350 000	<i>time out</i>
hwb-n20-02	(134, 630)	622	951	462	1	0.01
hwb-n22-02	(144, 688)	680	2183	811	1	0.025
ssa2670-141	(986, 2315)	1413	2.83	1.08	16	0.15
clqcolor-08-05-06	(116, 1114)	1114	107	62	1	0.007
dlx2_aa	(490, 2804)	1124	3.12	0.94	32	0.023
addsub.boehm	(492, 1065)	1324	35	29	> 657 000	<i>time out</i>

Table 3

L&S vs. HYCAM on computing all MUSes

and Sakallah that an interesting path for future research concerns the study of how MUSes could be computed progressively from the growing set of extracted MSSes.

Many artificial intelligence research areas have studied various problems involving the manipulation of MUSes, MSSes and CoMSSes, like model-based diagnosis, belief revision, inconsistency handling in knowledge and belief bases, etc. These studies are often conducted from a conceptual point of view, or from a worst-case complexity point of view, only. We believe that the practical computational progresses as such as the ones obtained in this paper can prove valuable in handling these problems practically. In this respect, future research could concentrate on deriving specific algorithms for these AI issues, exploiting results like the ones described in this paper.

Acknowledgments

We thank Mark Liffiton for making his system available to us. This work has been supported in part by the *Région Nord/Pas-de-Calais*. A preliminary version of this paper was published in the proceedings of IJCAI'07.

References

- [1] C. H. Papadimitriou, D. Wolfe, The complexity of facets resolved, *Journal of Computer and System Sciences* 37 (1) (1988) 2–13.

- [2] T. Eiter, G. Gottlob, On the complexity of propositional knowledge base revision, updates and counterfactual, *Artificial Intelligence* 57 (1992) 227–270.
- [3] W. Hamscher, L. Console, J. de Kleer (Eds.), *Readings in Model-Based Diagnosis*, Morgan Kaufmann, 1992.
- [4] M. Liffiton, K. Sakallah, On finding all minimally unsatisfiable subformulas, in: *Proceedings of SAT’05*, 2005, pp. 173–186.
- [5] R. Bruni, On exact selection of minimally unsatisfiable subformulae., *Annals of Mathematics and Artificial Intelligence* 43 (1) (2005) 35–50.
- [6] H. Büning, On subclasses of minimal unsatisfiable formulas, *Discrete Applied Mathematics* 107 (1–3) (2000) 83–98.
- [7] G. Davydov, I. Davydova, H. Büning, An efficient algorithm for the minimal unsatisfiability problem for a subclass of cnf, *Annals of Mathematics and Artificial Intelligence* 23 (3–4) (1998) 229–245.
- [8] H. Fleischner, O. Kullman, S. Szeider, Polynomial-time recognition of minimal unsatisfiable formulas with fixed clause-variable difference, *Theoretical Computer Science* 289 (1) (2002) 503–516.
- [9] R. Bruni, Approximating minimal unsatisfiable subformulae by means of adaptive core search, *Discrete Applied Mathematics* 130 (2) (2003) 85–100.
- [10] L. Zhang, S. Malik, Extracting small unsatisfiable cores from unsatisfiable boolean formula, in: *Sixth international conference on theory and applications of satisfiability testing (SAT’03)*, Portofino (Italy), 2003, pp. 239–249.
- [11] I. Lynce, J. Marques-Silva, On computing minimum unsatisfiable cores, in: *Seventh international conference on theory and applications of satisfiability testing (SAT’04)*, Vancouver, 2004, pp. 305–310.
- [12] Y. Oh, M. Mneimneh, Z. Andraus, K. Sakallah, I. Markov, AMUSE: a minimally-unsatisfiable subformula extractor, in: *Proceedings of the 41th Design Automation Conference (DAC 2004)*, 2004, pp. 518–523.
- [13] É. Grégoire, B. Mazure, C. Piette, Extracting MUS, in: *Proceedings of the 17th European Conference on Artificial Intelligence (ECAI’06)*, Riva del Garda, Italy, 2006, pp. 387–391.
- [14] É. Grégoire, B. Mazure, C. Piette, Tracking mus and strict inconsistent cover, in: *Proceedings of International Conference on Formal Methods in Computer-Aided Design (FMCAD’06)*, San Jose, CA, USA, 2006, pp. 39–46.
- [15] B. Bessant, É. Grégoire, P. Marquis, L. Saïs, Iterated Syntax-Based Revision in a Nonmonotonic Setting, Vol. 22 of *Applied Logic*, Kluwer Academic Publishers, Applied Logic Series, 2001, Ch. of *Frontiers in Belief Revision*, pp. 369–391.
- [16] J. Bailey, P. Stuckey, Discovery of minimal unsatisfiable subsets of constraints using hitting set dualization, in: *PADL*, 2005, pp. 174–186.

- [17] M. de la Banda, P. Stuckey, J. Wazny, Finding all minimal unsatisfiable subsets, in: Proceedings of the Fifth ACM-SIGPLAN International Conference on Principles and Practice of Declarative Programming (PPDL 2003), 2003, pp. 32–43.
- [18] L. Zhang, On subsumption removal and on-the-fly cnf simplification, in: Proceedings of the International Conference on Theory and Applications of Satisfiability Testing (SAT'05), 2005, pp. 482–489.
- [19] N. Eén, N. Sörensson, Minisat home page
<http://www.cs.chalmers.se/cs/research/formalmethods/minisat> (2004).
- [20] H. Kautz, B. Selman, D. McAllester, Walksat in the SAT'2004 competition, SAT Competition 2004 - solver description.
- [21] SATLIB, Benchmarks on SAT
<http://www.intellektik.informatik.tu-darmstadt.de/satlib/benchm.html> (2006).