

Solving Model-Based Diagnosis Problems with Max-SAT Solvers and Vice Versa

Alexander Feldman^{1,4}, Gregory Provan², Johan de Kleer³,
Stephan Robert⁴, and Arjan van Gemund¹

¹ Delft University of Technology, Mekelweg 4, 2628 CD, Delft, The Netherlands
{a.b.feldman,a.j.c.vangemund}@tudelft.nl

² University College Cork, College Road, Cork, Ireland
g.provan@cs.ucc.ie

³ Palo Alto Research Center, 3333 Coyote Hill Road, Palo Alto, CA 94304, USA
dekleer@parc.com

⁴ Haute Ecole d'Ingénierie et de Gestion du Canton de Vaud
Route de Cheseaux 1, CH-1401 Yverdon-les-Bains, Switzerland
stephan.robert@heig-vd.ch

ABSTRACT

In this paper we bring closer computation of consistency-based cardinality-minimal diagnosis and solving Max-SAT. We propose two algorithms for translating between those: (1) DIORAMA (DIagnOsis-based algoRithm for mAx-sat optiMizAtion) for translating cardinality-minimal consistency based diagnosis to Max-SAT and (2) MERIDIAN (Max-sat-basEd algoRithm for DIAgNosis) for the other way around. While the former approach has been studied, solving Max-SAT instances with a diagnostic solver is, to the best of our knowledge, novel. We configure MERIDIAN with the Stochastic Local Search (SLS) solvers from the UBCSAT suite, perform extensive experimentation on fault-models of the 74XXX/ISCAS85 circuits and compare the resulting optimality (in terms of approximation error in the minimal-cardinality) to the one of the MBD algorithm SAFARI. The results show that the optimality of SAFARI is up to several-orders-of-magnitude better than that of the SLS-based Max-SAT solvers. We configure DIORAMA with SAFARI and experiment on instances from the Max-SAT competitions. While the performance of DIORAMA/SAFARI on crafted Max-SAT problems is slightly worse compared to UBCSAT, DIORAMA/SAFARI outperforms at least several-orders-of-magnitude all UBCSAT algorithms on small industrial Max-SAT instances.

1 INTRODUCTION

Model-Based Diagnosis (MBD) inference algorithms for propositional diagnosis models have, in

This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

general, been created specifically for the diagnosis problem, e.g., GDE (de Kleer and Williams, 1987), and more recently SAFARI (Feldman *et al.*, 2010). These algorithms have been improved over the years by taking advantage of properties of the diagnosis problem, e.g., by focusing on the most likely diagnoses, or using a notion of continuity¹ of the diagnosis space (Feldman *et al.*, 2010).

With regard to domain-independent problems, significant progress has been made in developing powerful solvers for the *satisfiability* (SAT) problem, e.g., (Gomes *et al.*, 2007). This success has prompted the use of SAT-solvers for many other problems, such as planning (Castellini *et al.*, 2003) and circuit test-case generation (Iyer *et al.*, 2003). Using a SAT-solver for another (non-SAT) problem P entails rewriting P in SAT format; although this rewriting process can increase the size of the problem, the high efficiency of SAT solvers often makes the rewriting process worthwhile.

MBD is a version of propositional abduction, which is a more complex problem than SAT. Hence, although one can make calls to a SAT solver during the process of computing diagnoses (e.g., SAFARI uses an incomplete SAT solver), one cannot use a standard SAT solver directly for diagnostic inference. This article shows how one can use an extension of the SAT problem, called Max-SAT, to solve MBD problems.

Max-SAT is an optimization extension of SAT. Given a formula Φ in Conjunctive Normal Form (CNF), a Max-SAT (Hoos and Stützle, 2004) solution is a variable assignment that maximizes the number of satisfied clauses in Φ (in most cases of interest Φ is unsatisfiable, otherwise any variable assignment which satisfies Φ is also a Max-SAT solution). In partial Max-SAT, some of the clauses

¹A continuous subspace is loosely defined as a subset of the true-points of a Boolean function, such that for each true-point it includes another true-point, different in one bit exactly.

in Φ are designated as hard, the others are “soft”. A solution to the partial Max-SAT problem should satisfy all “hard” clauses and maximize the number of satisfied “soft” clauses. Similarly, in weighted Max-SAT a weight is assigned to each clause in Φ and a solution maximizes the sum of the weights of the satisfied clauses.

The contributions of this paper are as follows.

(1) We are the first to cast the Max-SAT problem as an MBD problem, for which we propose an algorithm called DIORAMA (DIagnOsis-based algoRithm for mAx-sat optiMizAtion). (2) We show that DIORAMA/SAFARI outperforms the traditional UBCSAT (Tompkins and Hoos, 2005) Max-SAT algorithms by at least two-orders-of-magnitude on a class of industrial Max-SAT problems, even though it performs slightly worse than Stochastic Local Search (SLS) Max-SAT algorithms on crafted Max-SAT competition problems. (3) We propose an algorithm, called MERIDIAN (Max-sat-basEd algoRithm for DIAgNosis), that translates an MBD problem to a Max-SAT problem. (4) We empirically show that MERIDIAN configured with traditional Max-SAT is less optimal than specialized MBD solvers such as SAFARI thereby revealing a large class of Max-SAT problems that expose continuous properties amenable to greedy algorithms like DIORAMA/SAFARI.

2 RELATED WORK

MBD has resemblance to Max-SAT (Hoos and Stützle, 2004) and we have conducted extensive experimentation with both complete Max-SAT (partial and weighted) and Max-SAT based on Stochastic Local Search (SLS). Empirical evidence shows that although Max-SAT can compute diagnoses in many of the cases, the performance of Max-SAT degrades when increasing the circuit size or the cardinality of the injected faults.

Fu and Malik (2006) construct a partial Max-SAT algorithm that uses UNSAT cores provided by SAT solvers. The algorithm of Fu and Malik iteratively relaxes UNSAT cores until the CNF input becomes satisfiable. The difference from their approach and DIORAMA is that they do not explicitly use a diagnostic algorithm to find a single minimal unsatisfiable core.

An interesting approach to solving Max-SAT is proposed by de Givry, *et al.* (2003) where they cast a Max-SAT problem as a weighted constraint satisfaction problem. On the other side, solving the diagnosis problem as a COP (Constraint Optimization Problem) is well-known from Williams and Ragno (2007).

On the side of solving diagnosis with Max-SAT, Kutsuna *et al.* (2009) use a partial Max-SAT algorithm to solve several diagnostic automotive problems. Similarly, Chen *et al.* (2009) use partial Max-SAT to solve problems of debugging sequential circuits. All these approaches differ from ours in that they solve specific diagnostic problems as opposed to empirically studying the general performance characteristics of Max-SAT and diagnostic algorithms.

3 TECHNICAL BACKGROUND

A *model* of an artifact is represented as a propositional formula over some set of variables. We discern subsets of these variables as *assumable* and *observable*.

Definition 1 (Diagnostic System). A diagnostic system DS is defined as the triple $DS = \langle SD, COMPS, OBS \rangle$, where SD is a propositional theory over a set of variables V , $COMPS \subseteq V$, $OBS \subseteq V$, COMPS is the set of assumables, and OBS is the set of observables.

Throughout this paper we assume that $OBS \cap COMPS = \emptyset$ and $SD \not\models \perp$.

Not all propositional theories are of interest to MBD. Traditionally MBD uses propositional theories that describe nominal and, optionally, faulty behavior of an interconnected set of components. We first define weak-fault models, i.e., models that describe nominal behavior only, a concept referred also to as ignorance of abnormal behavior.

Definition 2 (Weak-Fault Model). A diagnostic system $DS = \langle SD, COMPS, OBS \rangle$ belongs to the class **WFM** iff for $COMPS = \{h_1, h_2, \dots, h_n\}$, SD is equivalent to $(h_1 \Rightarrow F_1) \wedge (h_2 \Rightarrow F_2) \wedge \dots \wedge (h_n \Rightarrow F_n)$ and $COMPS \cap V' = \emptyset$, where V' is the set of all variables appearing in the propositional formulae F_1, F_2, \dots, F_n .

Modeling of faults makes the problem of computing diagnoses more complex (de Kleer *et al.*, 1992), but can increase the precision of a diagnostic algorithm. Models that have knowledge of faults are formalized below.

Definition 3 (Strong-Fault Model). A diagnostic system $DS = \langle SD, COMPS, OBS \rangle$ belongs to the class **SFM** iff SD is equivalent to $(h_1 \Rightarrow F_{1,1}) \wedge (\neg h_1 \Rightarrow F_{1,2}) \wedge \dots \wedge (h_n \Rightarrow F_{n,1}) \wedge (\neg h_n \Rightarrow F_{n,2})$ such that $1 \leq i, j \leq n, k \in \{1, 2\}$, $\{h_i\} \subseteq COMPS$, $F_{\{j,k\}}$ is a propositional formula, and none of h_i appears in $F_{j,k}$.

In this paper, in addition to **WFM**, we experiment with stuck-at-zero (S-A-0) and stuck-at-one (S-A-1) models. S-A-0 and S-A-1 are subclasses of **SFM** in which the output of a malfunctioning component is assumed either \perp or \top .

Definition 4 (Diagnosis). Given a diagnostic system $DS = \langle SD, COMPS, OBS \rangle$, an observation α over some variables in OBS, and a conjunction of literals ω , ω is a diagnosis iff $SD \wedge \alpha \wedge \omega \not\models \perp$.

Given a conjunction of literals ω we denote the set of negative literals in ω as $Lit^-(\omega)$.

Definition 5 (Subset-Minimal Diagnosis). A diagnosis $\omega \subseteq$ is defined as subset-minimal, if no other diagnosis $\tilde{\omega} \subseteq$, $\tilde{\omega} \subseteq \neq \omega \subseteq$, exists such that $Lit^-(\tilde{\omega} \subseteq) \subset Lit^-(\omega \subseteq)$.

In the MBD literature, a range of types of “preferred” diagnosis has been proposed. In the following definition we consider the important from the practical perspective cardinality ordering.

Definition 6 (Cardinality-Minimal Diagnosis). A diagnosis $\tilde{\omega}^{\leq}$ is defined as cardinality-minimal if no other diagnosis $\tilde{\omega}^{\leq}$, $\tilde{\omega}^{\leq} \neq \omega^{\leq}$, exists such that $Lit^-(\tilde{\omega}^{\leq}) < Lit^-(\omega^{\leq})$.

On the Max-SAT side we leave out all definitions as there is an agreement in the literature. Note that most SAT and Max-SAT solvers (as well as many diagnostic solvers) accept SD in CNF only. Any propositional formula can be converted to CNF taking into consideration a number of complexity and other issues (Feldman *et al.*, 2010).

4 MBD FRAMED AS MAX-SAT

In this section we demonstrate the use of Max-SAT for solving MBD problems.

4.1 A Max-SAT-Based MBD Algorithm

We propose an algorithm, called MERIDIAN (Max-sat-basEd algoRithm for DIAgNosis), for computing cardinality-minimal diagnoses (see Def. 6). MERIDIAN uses the approach of Sang *et al.* (2007) for encoding Most Probable Explanation (MPE) as Max-SAT. Computing MPE is identical to computing a most-probable diagnosis in a more general framework. Algorithm 1 computes diagnoses by calling a Max-SAT oracle.

Note that the diagnostic problems we solve in this paper can be translated to multiple optimization problems which can be solved with SAT-based methods (Giunchiglia and Maratea, 2006). The Maximum Satisfiable Subset (MSS) problem, for example, is dual to the Minimal Unsatisfiable Subset problem (Bailey and Stuckey, 2005) and the two can be solved with Max-SAT and Min-UNSAT solvers, respectively (Liffiton and Sakallah, 2005). From those, we have found preference in the research community towards Max-SAT, and for practical reasons we therefore compare SAFARI to Max-SAT.

Algorithm 1 adds a unit clause with weight 1 for each assumable (line 6). Note that these weight/clause pairs are added to W which is first set (lines 2 – 4) to contain all the clauses of the model SD. As a result, in line 8, W contains both the clauses of the original system description SD and the unit clauses added for each assumable in line 5 – 7. The weight of each input clause is set to a value greater than the number of all assumables (line 3). The loop in lines 8 – 11 computes a diagnosis with a call to Max-SAT and if a diagnosis exists, it is added to the result (line 10), and its negation is added to the original set of clauses (line 9) to prevent subsequent computation of the same diagnosis. Note that the negation of a term is, conveniently, a clause.

Depending on the implementation of the MAX-SAT call in line 8 of Alg. 1 we have a family of MAX-SAT algorithms for diagnosis: (1) if MAX-SAT is a partial Max-SAT solver, Alg. 1 computes diagnoses ordered by cardinality; (2) if MAX-SAT is a weighted Max-SAT solver, Alg. 1 computes diagnoses ordered by probability; and (3) if MAX-SAT is based on SLS, not every iteration of the

Algorithm 1 MERIDIAN: an algorithm for MBD based on weighted Max-SAT

```

1: function MERIDIAN(DS,  $\alpha$ ) returns a set of
   diagnoses
   inputs: DS, diagnostic system
             DS =  $\langle$ SD, COMPS, OBS $\rangle$ 
              $\alpha$ , term, observation
   local variables:  $W$ , set of weight and
                     clause pairs
                      $\Omega$ , set of diagnoses
                      $\omega$ , diagnosis term
                      $c_i$ , clause
                      $h_i$ , variable

2:   for all  $c_i \in$  CLAUSES(SD) do
3:      $W \leftarrow W \cup \langle \infty, c_i \rangle$ 
4:   end for
5:   for all  $h_i \in$  COMPS do
6:      $W \leftarrow W \cup \langle 1, h_i \rangle$ 
7:   end for
8:   while  $\omega \leftarrow$  MAX-SAT( $W$ ) do
9:      $W \leftarrow W \cup \langle \infty, \neg\omega \rangle$ 
10:     $\Omega \leftarrow \Omega \cup \omega$ 
11:   end while
12:   return  $\Omega$ 
13: end function

```

main loop yields a diagnosis. We have run extensive experiments with all three Max-SAT variants, which we describe in the following sub-sections.

4.2 Experimental Results with SLS Max-SAT

In the experiments that follow we compare the optimality of MERIDIAN (in terms of approximation error in the minimal-cardinality) configured with a number of SLS-based Max-SAT algorithms from the UBCSAT suite (Tompkins and Hoos, 2005). We compare the results to SAFARI, a state-of-the-art stochastic MBD algorithm (Feldman *et al.*, 2010).

There are two issues that complicate the use of SLS Max-SAT in diagnostic algorithms. First, there is no simple termination criterion in diagnostic algorithms based on SLS Max-SAT, i.e., we keep the local diagnosis and restart SAFARI after a number of successive “unsuccessful” flips, while there is no notion of “unsuccessful” flip (from the viewpoint of diagnosis) in Max-SAT. As we will see from our experimentation, flipping a variable which decreases the weight (or number) of currently satisfied clauses may be necessary to escape plateaus and/or local optima, hence the accumulation of such flips cannot be used as a termination criterion. Second, diagnostic Max-SAT problems have two types of constraints: hard and soft. The hard constraints are the clauses of the original (“nominal”) model, while the soft constraints are the unit clauses received from the assumable variables. An SLS Max-SAT algorithm does not distinguish between those hard and soft clauses; if such an algorithm guaranteed the satisfaction of the hard-constraints it would be classified as hy-

brid and not stochastic. These reasons make the use of algorithms based on SLS Max-SAT problematic in practical diagnosis. Despite that, we have conducted extensive experimentation with UBCSAT in order to evaluate the potential of SLS Max-SAT in MBD.

To overcome the termination problems with SLS Max-SAT, for the following experiments, we have chosen observations leading to known single faults. For each **WFM** we have chosen 50 observations. We have configured the SLS Max-SAT search to terminate after 100 000 variable flips and we have modified Alg. 1 to terminate after 10 calls to Max-SAT, this 10 runs. The resulting optimality of algorithms based on SLS Max-SAT in computing single fault diagnoses is shown in Table 1. We have run experiments with all algorithms or algorithm variants implemented by the UBCSAT suite, covering the algorithms RGSAT, Schöning, CDRW, through GSAT.

The data in Table 1 show best cases. From each of the 500 Max-SAT invocations per algorithm/circuit (50 single faults, 10 runs per experiment) we have (1) ignored all results which do not satisfy all hard constraints, (2) recorded the best diagnostic cardinality achieved in the hill climbing (recall that these are single-faults hence the best result is 1) and (3) recorded the number of steps (bit flips) in which this best diagnostic cardinality was achieved (the number of bit-flips are given in parentheses below the optimality number in Table 1).

Table 1 shows the generally poor performance of SLS Max-SAT algorithms. In most of the cases the algorithm could either never satisfy all hard-constraints or achieved increasingly worse cardinality with the growth of the circuit. Exceptions are the two variants of SAPS (Hutter *et al.*, 2002) and we attribute this relatively good optimality of SAPS to its mechanism for assigning and updating weights to clauses based on the clause length. Recall that in our diagnostic problems clauses of assumable literals have unit weights while hard-constraints have weights greater than the number of assumable literals. Despite that, in the best case for *c7552*, SAPS needed 77 264 bit flips to find the optimal single-fault diagnosis. In comparison SAFARI performed 11 bit flips, and although an LTMS/SAT consistency check of SAFARI is strictly more expensive than the consistency checking of SLS Max-SAT (the former is worst-case NP-hard while the latter is in P), SAFARI is computationally more efficient on average.

Figure 1 illustrates the progress of two SLS Max-SAT invocations. The Conflict-Directed Random Walk (CDRW) (Papadimitriou, 1991) starts with a random variable assignment and flips the most profitable (for increasing the satisfied weight) variable. This often leads to violated hard-constraints (due to flipping of non-assumable variables), and the restarts which are needed for escaping those situations lead to the relatively noisy ascent of CDRW. Other SLS Max-SAT algorithms like HSAT (Gent and Walsh, 1993) avoid down-

ward flips (flips which decrease the currently satisfied weight), quickly increasing the satisfied weight but ultimately get stuck in local optima. A close inspection of Fig. 1 reveals that HSAT oscillates forever short of satisfying all hard constraints.

5 MAX-SAT FRAMED AS MBD

In what follows we discuss the use of MBD for solving Max-SAT problems.

5.1 An MBD-Based Max-SAT Algorithm

Algorithm 2, called DIORAMA (DIagnOsis-based algoRithm for mAx-sat optiMizAtion), shows a very simple translation from a Max-SAT problem in CNF to a diagnostic problem.

Algorithm 2 DIORAMA: an algorithm for Max-SAT optimization based on MBD

```

1: function DIORAMA( $\Phi$ ) returns a term
   inputs:  $\Phi$ , set of clauses
   local variables: DS = (SD, COMPS,
                        OBS),
                    diagnostic system
                     $c_i$ , clause
                     $h_i$ , variable
2:   for all  $c_i \in \Phi$  do
3:     SD  $\leftarrow$  SD  $\wedge$   $\{h_i \Rightarrow c_i\}$ 
4:     COMPS  $\leftarrow$  COMPS  $\cup$   $h_i$ 
5:   end for
6:   return MBD(DS,  $\top$ )
7: end function

```

The loop in lines 2 - 4 of Alg. 2 modifies each clause in the input problem Φ . Note that line 3 adds exactly one literal to each input clause c_i as, given a clause $c = x_1 \vee x_2 \vee \dots \vee x_n$, we have $h \Rightarrow (x_1 \vee x_2 \vee \dots \vee x_n) \equiv \neg h \vee x_1 \vee x_2 \vee \dots \vee x_n$ and the right-hand side of the last equivalence is also a clause. Line 4 adds a total of $|\Phi|$ assumable variables to $|\text{COMPS}|$ where $|\Phi|$ is the number of clauses in Φ .

Algorithm 2 always creates a system description $\text{SD} \in \mathbf{WFM}$ (cf. Def. 2). Note as well that Alg. 2 invokes the MBD oracle in line 6 with an empty observation (for any propositional formula Φ we have $\Phi \wedge \top \equiv \Phi$).

In a stricter paper one can formally show the correctness of DIORAMA, i.e., one can prove that Alg. 2 always computes an optimal Max-SAT solution if it is configured with an MBD oracle that computes at least one cardinality-minimal diagnosis. The complexity of DIORAMA is dominated by the complexity of the Max-SAT solver. The complexity of Alg. 2 is $O(|\Phi|) + \Psi$ where Ψ is the complexity of the MBD oracle. We will, however, leave this discussion short in order to provide more extensive empirical evidence on the optimality of DIORAMA.

5.2 Experimental Results with a Stochastic MBD Oracle

In our first series of Max-SAT experiments we have configured Alg. 2 with the stochastic MBD oracle

Table 1: Optimality of SLS-based Max-SAT MERIDIAN and SAFARI on 74XXX/ISCAS85 WFMs and the number of steps (in parentheses) in which this optimality has been achieved

Name	RGSAT	Schöeningg	CDRW	URW	IROTS	ROTS	G2WSAT	SAPs ^a	SAPs ^b	Adaptive Novelty ⁺	Novelty ⁺	Novelty	WALKSAT/TABU	WALKSAT	HSAT	GWSAT	GSAT	SAFARI
74182	1 (4331)	1 (163)	1 (501)	9 (29896)	1 (1265)	1 (19)	1 (40)	1 (76)	1 (42260)	1 (28)	1 (56)	1 (21)	1 (37)	1 (26)	1 (186)	1 (38)	1 (2)	1
74L85	1 (65623)	1 (276)	1 (353)	-	1 (7236)	3 (27)	1 (39)	1 (143)	1 (80782)	1 (9815)	1 (1267)	5 (47)	1 (116)	1 (34)	1 (629)	1 (34)	1 (4)	1
74283	1 (63429)	1 (1134)	1 (573)	-	1 (3644)	1 (26)	1 (21806)	1 (112)	1 (25003)	1 (31900)	1 (435)	4 (35)	1 (1034)	1 (29)	1 (187)	1 (32)	1 (2)	1
74181	4 (48395)	1 (4525)	1 (6542)	-	1 (49967)	13 (1498)	1 (1814)	1 (453)	1 (1045)	1 (81883)	1 (2060)	10 (75)	1 (2123)	3 (67)	1 (1139)	3 (54)	1 (3)	1
c432	28 (43176)	1 (1866)	1 (7452)	-	4 (82496)	36 (101)	4 (10262)	1 (564)	7 (20211)	1 (75488)	8 (16393)	-	1 (905)	13 (133)	1 (24164)	16 (167)	1 (2)	1
c499	-	1 (24203)	1 (95242)	-	29 (99999)	-	6 (98683)	1 (42534)	7 (2112)	5 (87589)	13 (49516)	-	1 (47102)	-	1 (45465)	-	1 (1)	1
c880	-	1 (36269)	1 (8388)	-	-	-	26 (9561)	1 (4244)	35 (5680)	23 (80926)	30 (34989)	-	1 (9825)	-	1 (27026)	46 (304)	1 (5)	1
c1355	-	2 (93549)	1 (34235)	-	-	-	44 (83944)	1 (4975)	52 (84777)	51 (67454)	53 (85365)	-	1 (34786)	-	1 (93656)	-	1 (5)	1
c1908	240 (68649)	-	-	-	-	-	91 (97567)	1 (19664)	94 (79442)	98 (79442)	101 (43541)	-	1 (50660)	125 (654)	1 (94695)	139 (643)	1 (6)	1
c2670	-	-	-	-	-	-	122 (93799)	1 (6635)	126 (19000)	135 (67042)	144 (6532)	-	1 (83292)	173 (1008)	1 (85581)	177 (976)	1 (5)	1
c3540	-	-	-	-	-	-	167 (79188)	1 (17776)	170 (88052)	177 (68342)	203 (68663)	-	1 (99209)	-	16 (93639)	-	1 (9)	1
c5315	-	-	-	-	-	-	261 (78093)	1 (48388)	270 (64079)	275 (96973)	329 (48497)	-	14 (98557)	-	70 (89462)	-	1 (9)	1
c6288	-	-	-	-	-	-	329 (23191)	4 (73752)	360 (36373)	372 (91636)	-	-	58 (94673)	-	-	-	1 (3)	1
c7552	-	-	-	-	-	-	403 (12905)	1 (42284)	412 (39046)	440 (90485)	449 (32732)	-	115 (90449)	492 (2543)	181 (99340)	507 (2552)	1 (11)	1

^aClause penalties are initialized to the clause weights and smoothed back to their initial values.

^bClause penalties are initialized to the clause weights.




Figure 1: Progress of two SLS Max-SAT algorithms in a weak-fault model of c432, single fault observation

SAFARI (Feldman *et al.*, 2010). SAFARI is an approximation-based algorithm and we have configured it to compute guaranteed subset-minimal diagnoses (it cannot be configured to compute guaranteed cardinality-minimal diagnoses). These subset-minimal diagnoses are used as an approximation to cardinality-minimal diagnoses. The resulting algorithm DIORAMA/SAFARI is similar to SLS-based Max-SAT algorithms like the one discussed in Sec. 4.2.

Table 2 compares the optimality of DIORAMA/SAFARI to the algorithms from the UBCSAT suite. The experiments are on the problems from the Second Max-SAT Evaluation 2007. The majority of those problems (680 out of a total of 815) are random 2-SAT² and 3-SAT. We have configured UBCSAT to terminate after 100 000 steps and we have run it 10 times for each experiment. We can see in Table 2 that the optimality of DIORAMA/SAFARI is slightly worse but comparable to the optimality of the UBCSAT algorithms. In general, the optimality, of all UBCSAT algorithms and DIORAMA/SAFARI is similar which means that there are either (1) continuous diagnostic subspaces in the Max-SAT instances³ or (2) the Max-SAT algorithms and DIORAMA/SAFARI cannot climb after the initial variable assignment.

Table 3 shows the optimality of the UBCSAT Max-SAT algorithms and DIORAMA/SAFARI on the eight smallest instances of the Max-SAT 2009 industrial benchmark. The c3, c5315, c6288, c7552, mot_comb1, mot_comb2, mot_comb3, and s15850 columns in Table 3 correspond to the c3_DD_s3_f1_e1_v1-bug-onevec-gate-0, c5315-bug-gate-0, c6288-bug-gate-0, c7552-bug-gate-0, mot_comb1_red-gate-0, mot_comb2_red-gate-0, mot_comb3_red-gate-0, and s15850-bug-onevec-gate-0 instances in the Max-SAT benchmark. We can see that DIORAMA/SAFARI outperforms the traditional SLS-based algorithms by two to three orders-of-magnitude. This is not surprising as the c5315, c6288, c7552 instances come from the ISCAS85 benchmark and we have seen the

²Recall that although the 2-SAT decision problem is easy, the optimization Max-2-SAT problem is already NP-hard.

³See Feldman, *et al.* 2010 for defining continuity.

good performance of SAFARI on these instances in Sec. 4. What is more interesting is that these results hold for other benchmark instances from formal verification. s15850, for example, comes from ISCAS89 and has 534 D-type flip-flops. Note that all these problem instances result in solutions of very small cardinality.

6 CONCLUSION

This paper offers extensive empirical research on the use of consistency-based diagnosis for solving Max-SAT problems and vice-versa. The main contribution of this paper is solving more than 800 Max-SAT instances with DIORAMA/SAFARI and UBCSAT and more than 700 74XXX/ISCAS85 problems with MERIDIAN/UBCSAT and SAFARI. We have experimented with 20 algorithms for SLS-based Max-SAT. The good result of DIORAMA/SAFARI on small industrial instances show that many Max-SAT problems of real-world importance can be optimally and efficiently solved with greedy stochastic algorithms.

REFERENCES

- (Bailey and Stuckey, 2005) James Bailey and Peter J. Stuckey. Discovery of minimal unsatisfiable subsets of constraints using hitting set dualization. In *Proc. PADL'05*, pages 174–186, 2005.
- (Castellini *et al.*, 2003) Claudio Castellini, Enrico Giunchiglia, and Armando Tacchella. SAT-based planning in complex domains: Concurrency, constraints and nondeterminism. *Artificial Intelligence*, 147(1-2):85–117, 2003.
- (Chen *et al.*, 2009) Yibin Chen, Sean Safarpour, Andreas G. Veneris, and João P. Marques Silva. Spatial and temporal design debug using partial MaxSAT. In *ACM Great Lakes Symposium on VLSI*, pages 345–350, 2009.
- (de Givry *et al.*, 2003) Simon de Givry, Javier Larrosa, Pedro Meseguer, and Thomas Schiex. Solving Max-SAT as weighted CSP. In *CP*, pages 363–376, 2003.
- (de Kleer and Williams, 1987) Johan de Kleer and Brian Williams. Diagnosing multiple

Table 2: Optimality of UBCSAT algorithms and DIORAMA/SAFARI and the number of steps (in parentheses) in which this optimality has been achieved

Set Name	RGSAT	Schöningh	CDRW	URW	SAMD	IROTS	ROTS	G2WSAT Novelty+	G2WSAT	Adaptive Novelty+	Novelty+	Novelty	WalkSAT/TABU	WalkSAT	HWSAT	HSAT	GSAT/TABU	GWSAT	GSAT	DIORAMA/SAFARI
MAX3SAT/40	30 (801)	39 (4586)	39 (4425)	48 (4515)	30 (183)	30 (109)	30 (108)	30 (247)	30 (240)	30 (1846)	31 (2852)	31 (3141)	34 (4272)	32 (3699)	31 (1505)	33 (29)	30 (231)	30 (747)	33 (32)	58 (—)
MAX3SAT/50	26 (2192)	37 (4486)	37 (4659)	48 (4546)	26 (406)	26 (241)	26 (255)	26 (654)	26 (695)	26 (2388)	27 (3374)	27 (3484)	31 (3999)	29 (3892)	27 (1522)	29 (47)	26 (384)	26 (1513)	29 (77)	54 (—)
MAX3SAT/60	22 (2947)	34 (4667)	34 (4622)	48 (4755)	21 (328)	21 (289)	21 (271)	21 (755)	21 (867)	22 (2161)	23 (3094)	23 (3209)	26 (3868)	25 (3633)	23 (1696)	24 (82)	21 (310)	21 (1488)	24 (108)	48 (—)
MAX3SAT/70	19 (3735)	32 (4480)	31 (4591)	48 (4976)	18 (847)	18 (631)	18 (622)	18 (1527)	18 (1697)	19 (2516)	20 (3205)	20 (3093)	23 (3648)	22 (3464)	20 (1604)	21 (147)	18 (770)	18 (2316)	21 (279)	43 (—)
MAX3SAT/40	31 (812)	39 (4411)	39 (4398)	48 (4540)	31 (209)	31 (130)	31 (126)	31 (277)	31 (296)	31 (2007)	32 (2841)	32 (2819)	34 (3654)	33 (3371)	32 (1386)	34 (40)	31 (220)	31 (829)	34 (41)	58 (—)
MAX3SAT/60	23 (2768)	34 (3962)	34 (3957)	49 (4324)	22 (337)	22 (234)	22 (247)	22 (735)	22 (803)	23 (2176)	24 (3022)	24 (3178)	27 (3367)	26 (3245)	24 (1259)	25 (128)	22 (363)	22 (1442)	25 (142)	49 (—)
MAX3SAT/80	18 (3376)	31 (3711)	31 (3722)	49 (4380)	17 (534)	17 (458)	17 (458)	17 (1727)	17 (1689)	18 (1972)	19 (2581)	19 (2769)	21 (3356)	21 (3049)	18 (1427)	20 (148)	17 (593)	17 (2109)	20 (235)	40 (—)
MAX2SAT/60	93 (794)	110 (4754)	110 (4673)	132 (4765)	93 (88)	93 (63)	93 (64)	93 (2204)	93 (2136)	97 (3774)	102 (4348)	102 (4445)	104 (4442)	102 (4346)	93 (444)	93 (58)	93 (95)	93 (387)	93 (84)	143 (—)
MAX2SAT/100	81 (3320)	105 (4611)	105 (4659)	138 (4710)	80 (246)	80 (238)	80 (221)	82 (2930)	83 (3080)	87 (3651)	96 (4178)	95 (4237)	95 (4444)	95 (4184)	81 (574)	81 (233)	80 (248)	80 (1227)	81 (357)	131 (—)
MAX2SAT/140	68 (4080)	93 (4404)	93 (4384)	139 (4752)	66 (554)	66 (522)	66 (689)	70 (3056)	70 (2915)	73 (3533)	83 (3975)	83 (3966)	80 (3960)	81 (4001)	66 (764)	66 (574)	66 (542)	66 (1824)	67 (672)	113 (—)
SPINGLASS	64 (3628)	86 (4249)	87 (4484)	113 (4492)	58 (692)	58 (431)	58 (792)	63 (3488)	63 (3771)	66 (3368)	75 (4418)	74 (4345)	70 (4100)	69 (4369)	60 (914)	60 (587)	58 (964)	59 (2739)	60 (712)	105 (—)
RAMSEY	11 (1243)	16 (1721)	16 (1625)	21 (2596)	10 (189)	9 (368)	9 (261)	10 (818)	10 (724)	10 (672)	11 (1063)	11 (1132)	14 (1115)	14 (1214)	10 (657)	11 (40)	10 (198)	9 (687)	11 (45)	673 (—)
DIMACS	233 (601)	248 (3878)	248 (3975)	252 (4153)	233 (89)	233 (59)	233 (67)	233 (391)	233 (367)	238 (3365)	242 (3885)	242 (3785)	242 (3974)	242 (3785)	234 (863)	235 (22)	233 (111)	233 (510)	235 (25)	267 (—)
RANDOM	182 (2459)	212 (4710)	212 (4896)	223 (4534)	182 (303)	182 (205)	182 (204)	182 (2638)	182 (2698)	193 (4299)	201 (4544)	201 (4818)	202 (4718)	202 (4865)	185 (1298)	188 (33)	182 (344)	182 (1931)	187 (37)	239 (—)
SPINGLASS	113 (4558)	156 (4008)	155 (4370)	199 (4887)	98 (899)	96 (1981)	97 (2289)	112 (3246)	112 (3724)	116 (3673)	136 (3931)	136 (3615)	118 (3556)	126 (4543)	99 (1826)	99 (1153)	98 (1052)	99 (3553)	100 (1202)	172 (—)

Table 3: Optimality of the UBC SLS-based Max-SAT algorithms and SAFARI/DIORAMA on small industrial Max-SAT 2009 instances

	c3	c5315	c6288	c7552	mot_comb1	mot_comb2	mot_comb3	s15850
RGSAT	1 215	526	440	647	531	1362	1 758	3 483
Schöening	1 556	146	410	262	1	805	2 534	6 055
CDRW	1 532	145	423	255	2	832	2 533	6 012
URW	4 936	1 098	1 939	1 520	1 432	3 654	6 838	12 093
SAMD	342	500	132	695	9	86	598	2 175
IRoTS	261	78	86	86	24	83	436	1 972
RoTS	347	129	130	122	5	108	565	2 170
G2WSAT Novelty ⁺	237	13	99	54	1	4	544	2 124
G2WSAT	206	16	101	60	1	3	484	2 084
Adaptive Novelty ⁺	238	38	106	64	1	184	522	2 504
Novelty ⁺	339	13	111	66	1	49	800	3 367
Novelty	314	16	121	63	1	46	769	3 335
WalkSAT/TABU	368	24	129	84	1	676	810	2 819
WalkSAT	555	21	177	94	1	27	1 088	3 281
HWSAT	351	417	123	547	4	66	579	2 187
HSAT	354	498	130	798	18	116	576	2 145
GSAT/TABU	354	98	126	98	7	147	589	2 144
GWSAT	446	371	117	474	1	229	839	2 629
GSAT	372	392	130	347	18	149	596	2 191
DIORAMA/SAFARI	1	2	3	1	2	2	2	1

- faults. *Artificial Intelligence*, 32(1):97–130, 1987.
- (de Kleer *et al.*, 1992) Johan de Kleer, Alan Mackworth, and Raymond Reiter. Characterizing diagnoses and systems. *Artificial Intelligence*, 56(2-3):197–222, 1992.
- (Feldman *et al.*, 2010) Alexander Feldman, Gregory Provan, and Arjan van Gemund. Approximate model-based diagnosis using greedy stochastic search. *JAIR*, 2010.
- (Fu and Malik, 2006) Zhaohui Fu and Sharad Malik. On solving the partial max-sat problem. In *SAT*, pages 252–265, 2006.
- (Gent and Walsh, 1993) Ian P. Gent and Toby Walsh. Towards an understanding of hill-climbing procedures for SAT. In *Proc. AAAI’93*, pages 28–33, 1993.
- (Giunchiglia and Maratea, 2006) Enrico Giunchiglia and Marco Maratea. Solving optimization problems with DLL. In *Proc. ECAI’06*, pages 377–381, 2006.
- (Gomes *et al.*, 2007) Carla P. Gomes, Henry Kautz, Ashish Sabharwal, and Bart Selman. Satisfiability solvers. *Handbook of Knowledge Representation*, 2007.
- (Hoos and Stützle, 2004) Holger Hoos and Thomas Stützle. *Stochastic Local Search: Foundations and Applications*. Morgan Kaufmann Publishers Inc., 2004.
- (Hutter *et al.*, 2002) Frank Hutter, Dave A. D. Tompkins, and Holger H. Hoos. Scaling and probabilistic smoothing: Efficient dynamic local search for SAT. In *Proc. CP’02*, pages 233–248, 2002.
- (Iyer *et al.*, 2003) Madhu K. Iyer, Ganapathy Parthasarathy, and Kwang Ting Ting Cheng. SATORI - a fast sequential SAT engine for circuits. In *Proc. ICCAD’03*, pages 320–325, 2003.
- (Kutsuna *et al.*, 2009) Takuro Kutsuna, Shuichi Sato, and Naoya Chujo. Diagnosing automotive control systems using abstract model-based diagnosis. In *Proc. DX’09*, pages 99–105, 2009.
- (Liffiton and Sakallah, 2005) Mark H. Liffiton and Karem A. Sakallah. On finding all minimally unsatisfiable subformulas. In *Proc. SAT’05*, pages 173–186, 2005.
- (Papadimitriou, 1991) Christos Papadimitriou. On selecting a satisfying truth assignment. In *Proc. FOCS’91*, pages 163–169, 1991.
- (Sang *et al.*, 2007) Tian Sang, Paul Beame, and Henry A. Kautz. A dynamic approach for MPE and weighted MAX-SAT. In *Proc. IJCAI’07*, pages 173–179, 2007.
- (Tompkins and Hoos, 2005) Dave A. D. Tompkins and Holger H. Hoos. UBCSAT: An implementation and experimentation environment for SLS algorithms for SAT and MAX-SAT. In *Proc. SAT’04*, pages 306–320, 2005.
- (Williams and Ragno, 2007) Brian Williams and Robert Ragno. Conflict-directed A* and its role in model-based embedded systems. *Journal of Discrete Applied Mathematics*, 155(12):1562–1595, 2007.