The Puzzling Role of Simplification in Propositional Satisfiability

Inês Lynce and João Marques-Silva *

Technical University of Lisbon,
IST/INESC/CEL, Lisbon, Portugal
{ines,jpms}@sat.inesc.pt

Abstract. The ability to reduce either the number of variables or clauses in instances of the Propositional Satisfiability problem (SAT) is expected to impact the computational effort of solving a given instance. Moreover, the ability to *infer* new clauses is also expected to impact the computational effort of SAT solvers. Even though new clauses are inferred and added to the formula, the task of solving a target problem instance can become significantly *simplified*. The main objectives of this paper are to review existing simplification techniques for formulas in Conjunctive Normal Form (CNF), and to evaluate the effectiveness of the different simplification techniques, based on existing propositional simplifiers and state-of-the-art SAT solvers. Experimental results, obtained on representative hard real-world problem instances, indicate that simplification can play a puzzling role in Propositional Satisfiability. The main experimental conclusion is that reducing the number of variables or inferring additional clauses is not necessarily useful in terms of reducing the amount of search performed by backtrack search SAT algorithms, but that depends on which simplification techniques are actually used.

1 Introduction

Recent years have seen the proposal of several effective approaches for solving Propositional Satisfiability (SAT). These approaches include, among others, local search and variations [18], backtrack search improved with different search pruning techniques [2, 16, 24], backtrack search with randomization and restarts [8], continuous formulations [19] and algebraic manipulation [9, 12, 20]. These different approaches have allowed efficiently solving different classes of instances of SAT. It is generally accepted that whereas most approaches can be competitive in proving satisfiability for different classes of instances, backtrack search is preferred when the objective is to prove unsatisfiability. Nevertheless, formula simplification solutions are also known to be competitive for proving unsatisfiability in specific contexts [12].

* This work is partially supported by FCT - Fundação para a Ciência e Tecnologia under research projects PRAXIS/P/EEI/11249/1998 and POSI/34504/CHS/2000.

The main goal of this paper is to briefly review existing formula simplification techniques, categorize these techniques, and illustrate the effectiveness of the simplification techniques as a preprocessing tool for SAT algorithms. Moreover, we analyze the application of the proposed simplification techniques in real-world instances of SAT. The obtained experimental results indicate that simplification can play a puzzling role in Propositional Satisfiability. The main experimental observation is that reducing the number of variables or inferring additional clauses is not necessarily useful in terms of reducing the amount of search performed by backtrack search SAT algorithms. The practical usefulness of reducing the number of variables or inferring clauses depends on which techniques are actually utilized by the simplification algorithm. As a result, one conclusion that can be drawn is that simplification techniques ought to be evaluated in the context in which they are to be used, and not solely by their intrinsic simplification ability.

The paper is organized as follows. We start with a few definitions in Section 2. Next we address formula simplification, and propose a categorization of the different techniques. Section 4 analyzes the existing propositional simplifiers. Experimental results on applying these simplifiers on real-world instances of SAT are presented in Section 5. Moreover, the obtained results are compared and discussed. Finally, Section 6 concludes the paper.

2 Definitions

This section introduces the notational framework used throughout the paper. Propositional variables are denoted $x_1, \ldots, x_n$, and can be assigned truth values 0 (or F) or 1 (or T). In addition to letter x, and whenever necessary, we will use letters y, w and z to denote variables. To denote specific variables we may also use $x_i, x_j, x_k, \ldots$. A literal l is either a variable x_i (i.e. a positive literal) or its complement $\neg x_i$ (i.e. a negative literal). A clause ω is a disjunction of literals and a CNF formula φ is a conjunction of clauses. When referring to specific clauses we will utilize subscripts $a, b, \ldots$, and when referring to sub-formulas of a CNF formula we will utilize subscripts $r, s, \ldots$. Disjunctions of literals, not necessarily representing clauses will be represented as $\alpha, \beta, \gamma, \delta, \epsilon$. A literal l_j of a clause ω_a that is assigned truth value 1 satisfies the clause, and the clause is said to be *satisfied*. If the literal is assigned truth value 0 then it is removed from the clause. A clause with a single literal is said to be *unit*. The derivation of an *empty* clause indicates that the formula is *unsatisfied*. The formula is *satisfied* if all its clauses are satisfied. The SAT problem consists of deciding whether there exists a truth assignment to the variables such that the formula becomes satisfied.

3 Simplification Techniques

In this section we review existing techniques for CNF formula simplification. We start with the most well-known and fundamental techniques, then analyze pattern-based simplification techniques, and conclude by describing advanced techniques for identification of necessary assignments.

3.1 Basic Techniques

This section covers well-known formula simplification techniques, that include the unit clause rule, the pure literal rule, restricted resolution and two-variable equivalence conditions.

Simple simplifications. We start by observing that a clause containing a literal l_j and its complement $\neg l_j$ is satisfied, i.e. $l_j \vee \neg l_j \equiv 1$. Moreover, if a clause contains two copies of a literal l_j, then one copy can be removed, i.e. $l_j \vee l_j \equiv l_j$. Another simplification consists in removing duplicated clauses, which often occurs due to simplifications on the original formula.

Unit clause rule and pure literal rule. The most widely used formula simplification procedures in backtrack search SAT algorithms are certainly the *unit clause rule* and the *pure literal rule* [6].

The unit clause rule applies whenever a unit clause $\omega_i = (l_j)$ is identified. In this case, all clauses containing literal l_j are declared satisfied, and literal $\neg l_j$ is removed from the clauses containing it. The unit clause rule can be iterated, and this process is usually referred to as *Boolean Constraint Propagation (BCP)* [23].

The pure literal rule applies whenever a variable x_j only occurs as a pure literal, either $l_j = x_j$ or $l_j = \neg x_j$. In this situation, all clauses containing literal l_j can be declared satisfied. (Notice that literal l_j can be assigned truth value 1, satisfying all of its clauses.)

Restricted resolution. Resolution is known to be a sound and complete proof procedure for Propositional Satisfiability. It is in general not computationally feasible to utilize unrestricted forms of resolution. Hence, different algorithms for simplifying propositional formulas have used different forms of constraining the application of the resolution operation [22, 9]. For example, we can consider *restricted* forms of resolution by which *only* simplified clauses can be derived. Considering clauses $\omega_1 = (x_1 \vee x_2 \vee \neg x_3)$ and $\omega_2 = (\neg x_1 \vee x_2 \vee \neg x_3)$, resolution allows deriving clause $(x_2 \vee \neg x_3)$. This can be written as $\omega_3 = \rho(\omega_1, \omega_2, x_1) = (x_2 \vee \neg x_3)$. In general, given $\omega_r = (x_i \vee \alpha)$ and $\omega_s = (\neg x_i \vee \alpha)$, where α is a disjunction of literals, resolution allows deriving the clause,

$$\omega_t = \rho(\omega_r, \omega_s, x_i) = \rho((x_i \vee \alpha), (\neg x_i \vee \alpha), x_i) = (\alpha). \tag{1}$$

Two-variable equivalence. Two-variable equivalence is a well-known formula simplification procedure. Let us consider the pair of clauses $(\neg x \vee y) \wedge (x \vee \neg y)$, that can be represented as $(x \rightarrow y) \wedge (y \rightarrow x)$ (and also as $(\neg y \rightarrow \neg x) \wedge (\neg x \rightarrow \neg y)$). Then x and y are said to be equivalent, i.e. for any assignment that satisfies the two clauses, the truth values of x and y must be the same. Consequently, we can replace y by x on *all* occurrences of y, thus eliminating variable y from the CNF formula. Moreover, additional two-variable equivalence conditions can be established, namely based on loops of binary implications [1] or obtaining the *2-closure* of the 2CNF sub-formula [7].

3.2 Pattern-Based Simplification

This section describes pattern-based simplification techniques [12, 15]. Pattern-based clause inference conditions serve two main purposes:

- To enable variable equivalence conditions, thus allowing variables to be eliminated, and some clauses to become satisfied.
- To further complete the CNF formula, thus establishing additional relations between the problem instance variables, and so reducing the amount of expected search.

Equivalency Reasoning. Recent work on pattern-based simplification was done by C.-M. Li [12] and is referred to as *equivalency reasoning.* This work motivated and was superseded by other pattern-based formula simplification techniques [15], which are described in the following paragraphs. Nevertheless, equivalency reasoning is applied during search and pattern-based formula simplification techniques are currently solely applied during preprocessing.

Inference of Binary Clauses. In what follows, we analyze subsets of clauses for inferring binary clauses [15]. We should note that all proposed conditions can be explained by resorting to resolution. Nonetheless, the objective of studying sets of clauses is to indirectly select to which sets of clauses the resolution operation should be applied to.

Let us start by considering an illustrative example. Let $\varphi_c = (y_1 \vee \neg x) \wedge (y_2 \vee \neg x) \wedge (\neg y_1 \vee \neg y_2 \vee z)$ be a sub-formula. The application of the resolution operation between the three clauses allows deriving $(\neg x \vee z)$. Similarly, for the sub-formula $\varphi_d = (y_1 \vee \neg z) \wedge (y_2 \vee \neg z) \wedge (\neg y_1 \vee \neg y_2 \vee x)$, the resolution operation allows deriving the clause $(x \vee \neg z)$. It is interesting to observe that we have just illustrated an alternative approach for proving two-variable equivalence.

In general, let us consider a sub-formula of the form $\varphi_s = (l_1 \vee \neg l_2) \wedge (l_1 \vee \neg l_3) \wedge (l_2 \vee l_3 \vee l_4)$, where $l_1, \ldots, l_4$ are any literals. The application of resolution allows deriving the binary clause $(l_1 \vee l_4)$. Since 2 binary clauses and 1 ternary clause are involved in deriving the resulting binary clause, we say that 2B/1T reasoning was applied.

Additional forms of mB/nT reasoning can be established. As an example we analyze 0B/4T reasoning. Suppose the following set of four ternary clauses:

$$(l_1 \vee l_2 \vee l_3) \wedge (\neg l_1 \vee \neg l_2 \vee l_3) \wedge (l_1 \vee \neg l_2 \vee \neg l_4) \wedge (\neg l_1 \vee l_2 \vee \neg l_4)$$

Then the binary clause $(l_3 \vee \neg l_4)$ can be inferred, again by restricted resolution.

3.3 Identification of Necessary Assignments

From [23], we know that given a set of variable assignments, BCP identifies necessary assignments due to the unit clause rule in linear time on the number of literals of the CNF formula. However, BCP does not identify all necessary assignments given a set of variable assignments and a CNF formula. Consider for example the clauses $\omega_1 = (x_1 \vee \neg x_2)$ and $\omega_2 = (x_1 \vee x_2)$. For any assignment to variable x_2, variable x_1 must be assigned value 1 for preventing a conflict. Nevertheless, BCP applied to this CNF formula would not produce this straightforward conclusion. Moreover, the backtrack search procedure can and has been augmented with different probing techniques, namely by using *Recursive Learning* [11, 14] and by applying the branch-merge rule of *Stålmark's Method* [20, 9]. Recursive Learning and Stålmark's Method can be respectively interpreted as concrete forms of *clause probing* and *variable probing*. We should note that there are other, more specific, types of probing that are briefly described in Section 4.

Clause probing [11, 14] consists of the recursive evaluation of clause satisfiability requirements for identifying common assignments to variables. Common assignments are deemed necessary for a given clause to become satisfiable and consequently for the instance of SAT to be satisfiable. Let us consider the formula $\varphi_a = (x_1 \vee x_2) \wedge (\neg x_1 \vee x_3) \wedge (\neg x_2 \vee x_3)$. Clause probing would assign $x_1 = 1$ and then $x_2 = 1$ to satisfy clause $(x_1 \vee x_2)$, and in each case it would propagate the implied assignments. For this example observe that every way of satisfying clause $(x_1 \vee x_2)$ implies $x_3 = 1$. Hence, $x_3 = 1$ is a necessary assignment.

The main goal of variable probing is to identify common assignments to variables, detecting and merging equivalent branches in the search tree. In other words, variable probing is defined by the recursive application of the branch-merge rule to each variable [9]. Let us now consider the formula $\varphi_b = (x_1 \vee x_2) \wedge (\neg x_1 \vee x_3) \wedge (\neg x_2 \vee x_4) \wedge (\neg x_3 \vee x_4)$. We will start by applying variable probing on variable x_1. We conclude that both assignments, i.e. $x_1 = 0$ and $x_1 = 1$, imply $x_4 = 1$. Hence, $x_4 = 1$ is a necessary assignment.

Note that both clause and variable probing can be applied using recursion of *arbitrary depth*, respectively for clauses and for variables, thus providing complete proof procedures for SAT. Moreover, both approaches are based on similar reasoning, despite one being based on clauses and the other on variables. More interestingly, observe that these methods can be integrated into backtrack search SAT algorithms and can also be used together.

4 Formula Simplification Algorithms

Having described most of the existing simplification techniques in the previous section, we now characterize different formula simplification algorithms in terms of these techniques. For this purpose we selected state-of-the-art simplification algorithms, namely `Compact` [5], `Compactor` [13] [1], `Simplify` [15] and `CompressLite` [3][2].

[1] Available from http://www.laria.u-picardie.fr/~cli/my_compact.c.
[2] Available from the SATLIVE! web site, http://www.satlive.org/.

Before describing the techniques specific to each simplification algorithm, we enumerate the ones common to all the algorithms considered:

- Boolean constraint propagation (BCP).
- Pure literal rule (PLR).
- Elimination of duplicated clauses and literals in clauses, and elimination of satisfied clauses.

Compact. `Compact` is a propositional simplifier developed by James Crawford [5], that implements the following techniques:

- *Singleton rule*: if a variable literal has a single occurrence, then apply resolution to eliminate that variable.
- *Failed literal rule*: if the assignment $l_1 = 1$ yields a conflict (due to BCP), then we can assign $l_1 = 0$.
- *Binary failed literal rule*: if the assignments $l_1 = 1$ and $l_2 = 1$ yield a conflict (due to BCP), then we can create the binary clause $(\neg l_1 \vee \neg l_2)$.

Compactor. `Compactor` corresponds to the SATZ algorithm [13], without the backtrack search phase. `Compactor` includes the following simplification techniques:

- Failed literal rule.
- Application of resolution for deriving resolvents such that each resolvent has less than four literals.

Simplify. `Simplify` is another simplification algorithm that incorporates the following simplification techniques:

- Two-variable equivalence.
- Restricted resolution (i.e. only applied to binary clauses and to infer unit clauses).
- Clause inference conditions, based on the 2B/1T and 0B/4T clause patterns (see Section 3.2).

CompressLite. `CompressLite` [3] is a recent simplifier that implements alternative forms of variable probing:

- Utilization of BCP for probing variable assignments. The main objective is to identify two-variable equivalences. If $x = 1$ implies $y = 1$ and $x = 0$ implies $y = 0$, then x and y can be deemed equivalent.

Table 1. Example Instances

Application Domain	Instance	Satisfiable?
Circuit Testing (DIMACS) [10]	bf1355-075	N
	bf2670-001	N
	ssa2670-130	N
	ssa2670-141	N
Bounded Model Checking (BMC) [4]	barrel6	N
	barrel7	N
	longmult5	N
	longmult6	N
	queueinvar12	N
	queueinvar16	N
Combinational Equivalence Checking (CEC) [11]	c1908	N
	c1908_bug	Y
	c2670	N
	c2670_bug	Y
Superscalar Processor Verification (SPV) [21]	dlx2_cc_bug02	Y
	dlx2_cc_bug08	Y
Data Encryption Standard (DES) [17]	cnf-r1-b4-k1.1	Y
	cnf-r3-b3-k1.2	Y
	cnf-r3-b4-k1.2	Y

5 Experimental Results

This section evaluates the practical application of the different simplification techniques described in this paper, according to two main metrics:

- The ability to effectively simplify the original formula.
- The effective reduction in the amount of search, when using formula simplification techniques as a preprocessing engine for backtrack search SAT algorithms.

With the objective of evaluating the formula simplification ability of the simplifiers described in Section 4 (`Compact`, `Compactor`, `Simplify` and `CompressLite`), several representative problem instances from different real-world applications were used. The problem instances are shown in Table 1, and can be obtained from the different authors. Moreover, for all the results presented in this section, a PIII-866MHz Linux machine with 512 MByte of RAM was used.

Formula Simplification. The first experiment consisted in applying the proposed formula simplification techniques to the different problem instances. Table 2 shows the results, namely the number of variables and clauses before applying the simplification techniques and its variations (in percentage) after using the four simplifiers. (The maximum allowed CPU time was set to 1,000 seconds.)

As can be concluded, for most of the instances considered, large reductions in the number of variables and clauses can be achieved. In some cases (e.g. *barrel6*

Table 2. Formula Simplification — Number of Variables and Clauses

Instance	Initial Formula		`Compact`		`Compactor`		`Simplify`		`CompressLite`	
	Var^s	Cl^s	Var^s	Cl^s	Var^s	Cl^s	Var^s	Cl^s	Var^s	Cl^s
barrel6	2306	8931	75%	80%	96%	112%	22%	25%	75%	79%
barrel7	3523	13765	76%	81%	97%	111%	23%	25%	78%	82%
bf1355-075	2180	6778	unsat		77%	82%	30%	44%	31%	40%
bf2670-001	1393	3434	unsat		86%	86%	28%	43%	31%	37%
c1908_bug	1919	5100	99%	101%	100%	99%	34%	44%	51%	63%
c1908	1917	5096	99%	101%	100%	99%	34%	44%	51%	63%
c2670_bug	2708	6696	90%	94%	94%	96%	43%	53%	50%	60%
c2670	2703	6756	90%	93%	94%	96%	43%	18%	53%	63%
cnf-r1-b4-k1.1	397893	7089	—		sat		0%	102%	0%	87%
cnf-r3-b3-k1.2	676896	27503	—		1%	86%	1%	101%	1%	99%
cnf-r3-b4-k1.2	939040	35963	—		1%	86%	1%	101%	1%	99%
dlx2_cc_bug02	1515	12808	93%	125%	98%	102%	98%	109%	99%	99%
dlx2_cc_bug08	1515	12808	93%	125%	98%	102%	98%	108%	99%	99%
longmult5	2397	7431	61%	62%	77%	75%	62%	70%	46%	49%
longmult6	2848	8853	64%	67%	80%	79%	64%	72%	49%	52%
queueinvar12	1112	7335	66%	67%	97%	97%	94%	135%	95%	99%
queueinvar16	1168	6496	92%	105%	97%	98%	93%	110%	94%	98%
ssa2670-130	1359	3321	unsat		94%	92%	32%	49%	33%	40%
ssa2670-141	986	2315	45%	49%	98%	98%	29%	45%	30%	39%

and *barrel7* for `Simplify`) the final number of variables is one fourth of the original number [3]. Similar reductions can be observed in the number of clauses. Nevertheless, for some instances (e.g. *queueinvar16* and *dlx2_cc_bug08*) the amount of simplification in the number of variables is negligible. It is also interesting to note that some simplifiers detected some of the instances to be sat/unsat in this preprocessing phase. Furthermore, for some instances (e.g. *dlx2_cc_bug02* and *dlx2_cc_bug08*) it is possible to observe an increase in the number of clauses, due to the application of clause inference techniques. These clauses can eventually be utilized later to reduce the search effort.

Table 3 shows the CPU Time for simplifying each instance by each simplifier. Clearly, the results depend on the instance and on the simplifier. Nonetheless, and considering in concrete the Data Encryption Standard instances, and the `Compact` and `CompressLite` simplifiers, we can conclude that some simplification techniques can sometimes introduce so much computational overhead that their utilization may not be worthwhile.

Search Results. The next experiment consisted of evaluating the actual usefulness of the simplification techniques in solving instances of SAT. Table 4 shows the results of running GRASP [16], rel_sat [2], SATZ [13] and EQSATZ [12] on

[3] For the DES instances the reductions in the number of variables are even more dramatic, but all these variables are known to be trivially removed.

Table 3. Formula Simplification — CPU Time

Instance	`Compact`	`Compactor`	`Simplify`	`CompressLite`
barrel6	0.43	0.52	2.68	0.05
barrel7	1.53	0.90	7.37	0.07
bf1355-075	0.09	0.16	0.17	1.52
bf2670-001	0.00	0.03	0.08	0.07
c1908_bug	0.09	0.06	0.19	1.25
c1908	0.07	0.06	0.19	1.23
c2670_bug	0.07	0.07	0.22	0.56
c2670	0.06	0.06	0.22	0.55
cnf-r1-b4-k1.1	>1000	0.68	1.77	9.33
cnf-r3-b3-k1.2	>1000	3.79	5.30	448.79
cnf-r3-b4-k1.2	>1000	6.74	9.01	693.37
dlx2_cc_bug02	0.54	1.14	2.18	0.05
dlx2_cc_bug08	0.52	1.14	2.16	0.05
longmult5	2.28	0.20	0.28	1.17
longmult6	5.45	0.26	0.34	2.18
queueinvar12	0.36	0.82	1.87	0.22
queueinvar16	0.21	0.65	0.78	0.13
ssa2670-130	0.02	0.03	0.09	0.11
ssa2670-141	0.03	0.04	0.05	0.07

the different problem instances *before* applying the simplification techniques [4]. The maximum allowed CPU time was set to 3,000 seconds. The results obtained *after* applying the simplification techniques are shown in Tables 5 [5] and 6 [6].

Discussion of the Results. The obtained experimental results are somewhat surprising, since formula simplification by different simplifiers may yield instances that require fairly different run times for each SAT solver. As a result, we can establish general trends, denoting which simplification techniques are indeed effective in reducing the CPU times of SAT solvers.

We start by analyzing the results for different simplifiers and solvers, and afterwards highlight which simplification techniques are in general useful for reducing the run times of SAT solvers.

For the `Compact` simplifier, we can point out two main conclusions:

1. A solver can sometimes need more time to solve an instance after applying formula simplification (e.g. *longmult* instances).
2. The reductions in the number of variables and clauses on the original formula do not necessarily imply a reduction in the search effort (e.g *dlx2_cc_bug* instances, for which `Compact` obtains the best reductions for the original formula).

[4] SATZ and EQSATZ stopped with error conditions when solving the DES instances.
[5] Instances that were solved during preprocessing are marked with *simp*.
[6] EQSATZ stopped with an error condition for instance *barrel7* simplified by `Simplify`.

Table 4. Time Results before Formula Simplification

Instance	GRASP	rel_sat	SATZ	EQSATZ
barrel6	>3000	643.0	1100.1	0.5
barrel7	>3000	>3000	24.4	0.7
bf1355-075	0.3	0.0	2.5	11.5
bf2670-001	0.1	0.0	0.1	0.6
c1908_bug	1126.0	>3000	>3000	0.7
c1908	2465.5	>3000	>3000	525.3
c2670_bug	6.4	1.0	0.4	>3000
c2670	302.9	>3000	>3000	>3000
cnf-r1-b4-k1.1	25.9	>3000	—	—
cnf-r3-b3-k1.2	1858.0	>3000	—	—
cnf-r3-b4-k1.2	1520.4	>3000	—	—
dlx2_cc_bug02	649.9	>3000	>3000	>3000
dlx2_cc_bug08	2.6	>3000	>3000	8.4
longmult5	21.1	9.0	7.1	4.2
longmult6	354.9	64.0	36.7	18.9
queueinvar12	10.8	10.0	39.0	1.9
queueinvar16	21.1	15.0	705.7	2.8
ssa2670-130	0.4	1.0	62.5	0.2
ssa2670-141	0.7	1.0	233.1	1.1

In contrast, `Compactor` *never* presents the best reductions in the number of variables and clauses for simplifying a given instance, when compared with the other simplifiers. Nevertheless, `Compactor` is more effective than `Compact` in reducing the search after formula simplification.

Formula simplifcation with `Simplify` allows solving problem instances that, otherwise, the different SAT solvers would be unable to solve in the allowed CPU time. In concrete, *all* the instances that were not solved in 3,000 seconds by GRASP and rel_sat before formula simplification are now solved in the allowed CPU time. In addition, all the simplified formulas are in general solved faster than the original formulas. There are however a few exceptions. For example, for one instance (*dlx2_cc_bug08*), and for rel_sat, the small reduction in the number of variables and the increase in the number of clauses, causes an increase in the amount of search. This result is explained by the dynamic nature of rel_sat's branching heuristic. Moreover, it is interesting to observe that the reduction in the number of variables and clauses allows problem instances *barrel6* and *barrel7* to be solved by GRASP in a few seconds, in contrast with the original formula, for which the GRASP SAT solver quits after 3,000 CPU seconds. Finally, and with respect to EQSATZ, it is also interesting to observe that formula simplification applied by `Simplify` during preprocessing cooperates effectively with the techniques for formula simplification applied by EQSATZ during the search. Moreover, one instance (*c2670_bug*) would not be solved by EQSATZ in 3,000 CPU seconds if it was not for the simplifications produced by `Simplify`. Finally,

Table 5. Time Results after Formula Simplification — `Compact` and `Compactor`

Instance	`Compact`				`Compactor`			
	GRASP	rel_sat	SATZ	EQSATZ	GRASP	rel_sat	SATZ	EQSATZ
barrel6	>3000	296.0	674.4	0.3	>3000	667.0	1099.1	0.2
barrel7	>3000	2945.0	5.1	0.6	>3000	>3000	3.1	0.3
bf1355-075	*simp*	*simp*	*simp*	*simp*	3.1	0.0	2.8	12.6
bf2670-001	*simp*	*simp*	*simp*	*simp*	0.1	0.0	0.1	0.5
c1908_bug	>3000	>3000	>3000	7.2	427.0	>3000	>3000	0.9
c1908	>3000	>3000	>3000	1358.9	>3000	>3000	>3000	393.6
c2670_bug	263.3	1.0	0.5	44.3	0.3	0.0	0.5	>3000
c2670	>3000	>3000	>3000	>3000	635.7	>3000	>3000	>3000
cnf-r1-b4-k1.1	—	—	—	—	*simp*	*simp*	*simp*	*simp*
cnf-r3-b3-k1.2	—	—	—	—	16.3	10.0	>3000	3.0
cnf-r3-b4-k1.2	—	—	—	—	12.8	13.0	7.5	5.2
dlx2_cc_bug02	760.4	>3000	>3000	>3000	1408.8	2133.0	>3000	>3000
dlx2_cc_bug08	97.7	>3000	>3000	>3000	17.6	9.0	>3000	8.3
longmult5	14.5	>3000	>3000	>3000	18.9	8.0	6.8	4.0
longmult6	263.6	>3000	>3000	>3000	241.5	67.0	35.4	18.2
queueinvar12	7.9	7.0	40.9	1.7	10.9	11.0	38.4	1.8
queueinvar16	20.1	14.0	739.6	2.4	21.2	18.0	704.4	2.8
ssa2670-130	*simp*	*simp*	*simp*	*simp*	0.4	0.0	67.0	0.2
ssa2670-141	0.0	0.0	0.1	0.0	0.4	1.0	171.0	2.2

`CompressLite` simplifications do not allow significant reductions in the amount of search.

Next, we present more detailed results for `Simplify`. Despite the simplifications observed in the number of variables and clauses, it is also interesting to analyze the contribution of each simplification technique. Table 7 shows the results for two-variable equivalence, restricted resolution, 2B/1T and 0B/4T inference of clauses, and the pure literal rule. Observe that none of the proposed techniques can be elected as the one solely responsible for simplifying the formulas. Actually, some techniques are mostly applied to specific classes of instances (e.g restricted resolution and 0B/4T for the *barrel* instances and the pure literal rule for the *bf* instances). Hence, we can conclude that the ability to effectively simplify the CNF formulas is obtained by applying the proposed techniques *as a whole*.

From the experimental results and previous observations, we can conclude that the *singleton rule* and the binary *failed literal rule* used in `Compact` do not in general contribute to reducing the run times of SAT solvers (quite the contrary). The restricted form of probing utilized by `CompressLite` is also not essential for reducing the run times of SAT solvers. In contrast, the failed-literal rule, two-variable equivalence, restricted resolution (as used in `Simplify`), and the 2B/1T and 0B/4T clause inference conditions, yield formula simplifications that effectively help SAT solvers reducing their run times.

Table 6. Time Results after Formula Simplification — `Simplify` and `CompressLite`

Instance	`Simplify`				`CompressLite`			
	GRASP	rel_sat	SATZ	EQSATZ	GRASP	rel_sat	SATZ	EQSATZ
barrel6	1.2	0.0	0.1	0.1	1999.5	276.0	>3000	485.4
barrel7	124.8	1.0	0.1	—	>3000	920.0	>3000	>3000
bf1355-075	0.1	0.0	0.9	0.2	0.1	0.0	3.2	0.1
bf2670-001	0.0	0.0	0.0	0.1	0.0	0.0	0.0	0.1
c1908_bug	6.3	2.0	>3000	1.0	1285.1	>3000	>3000	6.7
c1908	10.0	4.0	>3000	0.9	>3000	>3000	>3000	408.8
c2670_bug	0.1	0.0	0.1	0.5	0.1	0.0	0.2	1.9
c2670	80.9	1097.0	>3000	>3000	357.2	>3000	>3000	>3000
cnf-r1-b4-k1.1	0.3	1.0	0.5	0.4	0.2	1.0	0.5	0.4
cnf-r3-b3-k1.2	82.8	15.0	15.0	4.0	68.4	14.0	149.8	4.4
cnf-r3-b4-k1.2	80.7	16.0	>3000	7.2	29.9	14.0	>3000	7.9
dlx2_cc_bug02	107.4	869.0	>3000	>3000	>3000	361.0	>3000	>3000
dlx2_cc_bug08	0.5	520.0	>3000	3.8	19.5	3.0	>3000	>3000
longmult5	11.3	5.0	3.5	8.6	9.8	6.0	2.7	7.4
longmult6	249.2	43.0	19.8	48.2	150.2	53.0	15.5	26.5
queueinvar12	11.4	12.0	25.5	2.0	9.5	12.0	24.4	2.2
queueinvar16	20.4	11.0	430.4	3.1	19.1	14.0	416.9	2.8
ssa2670-130	0.1	0.0	3.5	0.3	0.1	0.0	12.7	0.1
ssa2670-141	0.1	0.0	12.8	0.2	0.1	0.0	9.8	0.2

6 Conclusions

This paper analyzes existing techniques for the simplification of propositional formulas, and conducts an experimental evaluation of these techniques, with the goal of deciding which ones are indeed useful in the context of solving hard real-world instances of SAT.

The experimental results, obtained for different existing simplifiers and on representative real-world instances of SAT, clearly demonstrate the effectiveness of these simplification techniques, when formula simplification is the main objective. However, and interestingly, the experimental results also indicate that some formula simplification techniques, despite yielding significant reductions in the number of variables and allowing the inference of a large number of clauses, may not necessarily contribute to solving specific classes of problem instances. One conclusion of this experimental data is that formula simplification techniques should be evaluated not only on their ability to reduce the number of variables or to infer new clauses, but also on their ability to yield CNF formulas that SAT solvers can more easily solve. This suggests a two-phase evaluation process, both for evaluating new formula simplication techniques and for selecting which of the existing techniques are indeed helpful for SAT solvers.

The presented experimental results depend on the problem instances considered. The next natural step is to conduct a thorough evaluation of the existing formula simplication techniques on *all* existing suites of real-world problem in-

Table 7. Formula Simplification — `Simplify` Techniques

Instance	Equivalences	Restricted Resolution	2B/1T	0B/4T	PLR	Time
barrel6	1575	126	0	9702	0	2.68
barrel7	2436	168	0	17472	0	7.37
bf1355-075	1028	34	327	0	63	0.17
bf2670-001	804	1	361	13	16	0.08
c1908_bug	1259	0	800	0	0	0.19
c1908	1258	0	803	0	0	0.19
c2670_bug	1401	66	634	0	0	0.22
c2670	1338	66	664	0	0	0.22
cnf-r1-b4-k1.1	0	0	318	0	0	1.77
cnf-r3-b3-k1.2	0	0	398	0	0	5.3
cnf-r3-b4-k1.2	0	0	521	0	0	9.01
dlx2_cc_bug02	3	18	1481	0	0	2.18
dlx2_cc_bug08	3	18	1406	0	0	2.16
longmult5	695	1	302	0	0	0.28
longmult6	791	1	375	0	0	0.34
queueinvar12	63	0	2931	28	0	1.87
queueinvar16	80	0	905	30	0	0.78
ssa2670-130	851	2	328	4	4	0.09
ssa2670-141	686	2	189	0	0	0.05

stances, of which the problem instances considered in this paper are representative examples.

References

1. B. Aspvall, M. F. Plass, and R. E. Tarjan. A linear-time algorithm for testing the truth of certain quantified boolean formulas. *Information Processing Letters*, 8(3):121–123, March 1979.
2. R. Bayardo Jr. and R. Schrag. Using CSP look-back techniques to solve real-world SAT instances. In *Proceedings of the National Conference on Artificial Intelligence*, pages 203–208, 1997.
3. D. L. Berre. Exploiting the real power of unit propagation lookahead. In *LICS Workshop on Theory and Applications of Satisfiability Testing*, June 2001.
4. A. Biere, A. Cimatti, E. Clarke, and Y. Zhu. Symbolic model checking without BDDs. In W. R. Cleaveland, editor, *Proceedings of Tools and Algorithms for the Construction and Analysis of Systems*, LNCS 1579. Springer-Verlag, March 1999.
5. J. Crawford and L. Auton. Experimental results on the cross-over point in satisfiability problems. In *Proceedings of the National Conference on Artificial Intelligence*, pages 22–28, 1993.
6. M. Davis and H. Putnam. A computing procedure for quantification theory. *Journal of the Association for Computing Machinery*, 7:201–215, 1960.
7. A. V. Gelder and Y. K. Tsuji. Satisfiability testing with more reasoning and less guessing. In D. S. Johnson and M. A. Trick, editors, *Second DIMACS Implementation Challenge*. American Mathematical Society, 1993.

8. C. P. Gomes, B. Selman, and H. Kautz. Boosting combinatorial search through randomization. In *Proceedings of the National Conference on Artificial Intelligence*, July 1998.
9. J. Groote and J. Warners. The propositional formula checker heerhugo. In I. Gent, H. van Maaren, and T. Walsh, editors, *SAT 2000*, pages 261–281. IOS Press, 2000.
10. D. S. Johnson and M. A. Trick, editors. *Second DIMACS Implementation Challenge*. American Mathematical Society, 1993. (DIMACS SAT instances available from URL ftp://Dimacs.Rutgers.EDU/pub/challenge/sat/benchmarks/cnf).
11. W. Kunz and D. Stoffel. *Reasoning in Boolean Networks*. Kluwer Academic Publishers, 1997.
12. C.-M. Li. Integrating equivalency reasoning into davis-putnam procedure. In *Proceedings of the National Conference on Artificial Intelligence*, pages 291–296, July 2000.
13. C. M. Li and Anbulagan. Look-ahead versus look-back for satisfiability problems. In *Proceedings of the International Conference on Principles and Practice of Constraint Programming*, 1997.
14. J. P. Marques-Silva. Improving satisfiability algorithms by using recursive learning. In *Proceedings of the International Workshop on Boolean Problems*, pages 47–58, September 1998.
15. J. P. Marques-Silva. Algebraic simplification techniques for propositional satisfiability. In *International Conference on Principles and Practice of Constraint Programming*, pages 537–542, September 2000.
16. J. P. Marques-Silva and K. A. Sakallah. GRASP-A search algorithm for propositional satisfiability. *IEEE Transactions on Computers*, 48(5):506–521, May 1999.
17. F. Massacci. Using Walk-SAT and Rel-SAT for cryptographic key search. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 290–295, 1999. (DES instances available from URL http://www.uni-koblenz.de/∼massacci).
18. B. Selman and H. Kautz. Domain-independent extensions to GSAT: Solving large structured satisfiability problems. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 290–295, 1993.
19. Y. Shang and B. W. Wah. A discrete lagrangian-based global-search method for solving satisfiability problems. *Journal of Global Optimization*, 12(1):61–99, January 1998.
20. G. Stålmarck. A system for determining propositional logic theorems by applying values and rules to triplets that are generated from a formula, 1989. Swedish Patent 467 076 (Approved 1992), US Patent 5 276 897 (approved 1994), European Patent 0 403 454 (approved 1995).
21. M. N. Velev and R. E. Bryant. Superscalar processor verification using efficient reductions from the logic of equality with uninterpreted functions to propositional logic. In *Proceedings of Correct Hardware Design and Verification Methods*, LNCS 1703, pages 37–53, September 1999.
22. F. Vlach. Simplification in a satisfiability checker for VLSI applications. *Journal of Automated Reasoning*, 10(1):115–136, February 1993.
23. R. Zabih and D. A. McAllester. A rearrangement search strategy for determining propositional satisfiability. In *Proceedings of the National Conference on Artificial Intelligence*, pages 155–160, 1988.
24. H. Zhang. SATO: An efficient propositional prover. In *Proceedings of the International Conference on Automated Deduction*, pages 272–275, July 1997.