

FourierSAT: A Fourier Expansion-Based Algebraic Framework for Solving Hybrid Boolean Constraints*

Anastasios Kyrillidis, Anshumali Shrivastava, Moshe Y. Vardi, Zhiwei Zhang[†]

Rice University, Houston, TX, USA
{anastasios, anshumali, vardi, zhiwei}@rice.edu

Abstract

The Boolean SATisfiability problem (SAT) is of central importance in computer science. Although SAT is known to be NP-complete, progress on the engineering side—especially that of Conflict-Driven Clause Learning (CDCL) and Local Search SAT solvers—has been remarkable. Yet, while SAT solvers, aimed at solving industrial-scale benchmarks in Conjunctive Normal Form (CNF), have become quite mature, SAT solvers that are effective on other types of constraints (e.g., cardinality constraints and XORs) are less well-studied; a general approach to handling non-CNF constraints is still lacking. In addition, previous work indicated that for specific classes of benchmarks, the running time of extant SAT solvers depends heavily on properties of the formula and details of encoding, instead of the scale of the benchmarks, which adds uncertainty to expectations of running time.

To address the issues above, we design `FourierSAT`, an incomplete SAT solver based on Fourier analysis of Boolean functions, a technique to represent Boolean functions by multilinear polynomials. By such a reduction to continuous optimization, we propose an algebraic framework for solving systems consisting of different types of constraints. The idea is to leverage gradient information to guide the search process in the direction of local improvements. Empirical results demonstrate that `FourierSAT` is more robust than other solvers on certain classes of benchmarks.

1 Introduction

Constraint satisfaction problems (CSPs) are fundamental in mathematics, physics, and computer science. The Boolean SATisfiability problem (SAT) is a special class of CSPs, where each variable takes value from the binary set $\{\text{True}, \text{False}\}$. Solving SAT efficiently is of utmost significance in computer science, both from a theoretical and a practical perspective.

As a special case of SAT, conjunctive normal forms (CNFs) are a conjunction (and-ing) of disjunctions (or-ing)

*The author list has been sorted alphabetically by last name; this should not be used to determine the extent of authors' contributions.

[†]Corresponding author: Zhiwei Zhang.

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

of literals. Despite the NP-completeness of CNF-SAT, there has been a lot of progress on the engineering side of CNF-SAT solvers. Mainstream SAT solvers can be classified into complete and incomplete ones: A complete SAT solver will return a solution if there exists one or prove unsatisfiability if no solution exists, while an incomplete algorithm is not guaranteed to find a satisfying assignment. Clearly, an incomplete algorithm cannot prove unsatisfiability.

Most modern complete SAT solvers are based on the Conflict-Driven Clause Learning (CDCL) algorithm (Marques-Silva and Sakallah 1999), an evolution of the backtracking Davis-Putnam-Logemann-Loveland (DPLL) algorithm (Davis and Putnam 1960; Davis, Logemann, and Loveland 1962). The main techniques used in CDCL solvers include clause-based learning of conflicts, random restarts, heuristic variable selection, and effective constraint-propagation data structures (Gong and Zhou 2017). Examples of highly efficient complete SAT solvers include `MiniSat` (Eén and Sörensson 2003), `BerkMin` (Goldberg and Novikov 2007), `PicoSAT` (Biere 2008), `Lingeling` (Biere 2010), `Glucose` (Audemard and Simon 2014) and `MapleSAT` (Liang 2018). Overall, CDCL-based SAT solvers constitute a huge success for SAT problems, and have been dominating in the research of SAT solving.

Local search techniques are mostly used in incomplete SAT solvers. The number of unsatisfied clauses is often regarded as the objective function. Local search algorithms mainly include greedy local search (GSAT) (Selman, Levesque, and Mitchell 1992) and random walk GSAT (WSAT) (Selman, Kautz, and Cohen 1999). During the main loop, GSAT repeatedly checks the current assignments neighbors and selects a new assignment to maximize the number of satisfied clauses. In contrast, WSAT randomly selects a variable in an unsatisfied clause and inverts its value (Gong and Zhou 2017). On top of these basic algorithms, several heuristics for variable selection have been proposed, such as `NSAT` (McAllester, Selman, and Kautz 1997), `Sparrow` (Balint and Fröhlich 2010), and `ProbsAT` (Balint and Schöning 2012). While practical local search SAT solvers could be slower than CDCL solvers, local search techniques are still useful for solving a certain class of benchmarks, such as hard random formulas and `MaxSAT`. Local search al-

gorithms can also have nice theoretical properties (Brueggemann and Kern 2004).

Non-CNF constraints are playing important roles in theoretical computer science and other engineering areas, *e.g.*, XOR constraints in cryptography (Bogdanov, Khovratovich, and Rechberger 2011) as well as cardinality constraints (CARD) and Not-all-equal (NAE) constraints in discrete optimization (Costa et al. 2009) (Dinur, Regev, and Smyth 2005). The combination of different types of constraints enhances the expressive power of Boolean formulas; *e.g.*, CARD-XOR is necessary and sufficient for maximum likelihood decoding (MLD) (Feldman, Wainwright, and Karger 2005), one of the most crucial problems in coding theory. Nevertheless, compared to that of CNF-SAT solving, efficient SAT solvers that can handle non-CNF constraints are less well studied.

One way to deal with non-CNF constraints is to encode them in CNF. However, different encodings differ from the size, the ability to detect inconsistencies by unit propagation (arc consistency) and solution density (Prestwich 2009). It is generally observed that the running time of SAT solvers relies heavily on the detail of encodings. *E.g.*, CDCL solvers benefit from arc-consistency (Quimper and Walsh 2007), while local search solvers prefer short chains of variable dependencies (Kautz, McAllester, and Selman 1997). Finding a best encoding for a solver usually requires considerable testing and comparison (Martins, Manquinho, and Lynce 2011).

Another way is to extend the existing SAT solvers to adapt to non-CNF clauses. Work on this line includes Cryptominisat (Soos, Nohl, and Castelluccia 2009) for CNF-XOR, MiniCARD (Liffiton and Maglalog 2012) for CNF-CARD, Pueblo (Sheini and Sakallah 2006) and RoundingSAT (Elfers and Nordström 2018) for CNF-Pseudo Boolean and MonoSAT (Bayless et al. 2015) for handling CNF-graph properties formulas. Such specialized extensions, however, often require different techniques for different types of constraints. Meanwhile, general ideas for solving hybrid constraints uniformly are still lacking.

Contributions. The primary contribution of this work is the design of a novel algebraic framework as well as a versatile, robust incomplete SAT solver—`FourierSAT`—for solving hybrid Boolean constraints.

The main technique we used in our method is the Walsh-Fourier transform (Fourier transform, for short) on Boolean functions (O’Donnell 2014). By transforming Boolean functions into “nice” polynomials, numerous properties can be analyzed mathematically. Besides the success of Fourier transform in theoretical computer science (Linial, Mansour, and Nisan 1989) (Wei and Ermon 2017), recently, this tool has also found surprising uses in algorithm design (Achim, Sabharwal, and Ermon 2016) (Xue et al. 2015). In (Warners and van Maaren 1998), the authors used a polynomial representation to design a SAT solver. However, their solver was still DPLL-based and mathematical properties of polynomials were not fully exploited. More algorithmic uses of this technique are waiting to be discovered.

To our best knowledge, this paper will be the first algorithmic work to study Fourier expansions of Boolean functions in the real domain instead of only Boolean domain. After

this relaxation, we find Fourier expansions well-behaved in the real domain. Thus, we manage to reduce satisfiability of Boolean constraints to continuous optimization and apply gradient-based methods. One of the attractive properties of our method is, different types of constraints are handled uniformly—we no longer design specific methods for each type of constraints. Moreover, as long as the Fourier expansions of a new type of constraints are known, our solver can be extended trivially. In addition, we explain the intuition of why doing continuous optimization for SAT is better than doing discrete local search.

Furthermore, our study on the local landscape of Fourier expansions reveals that the existing of saddle points is an obstacle for us to design algorithms with theoretical guarantees. Previous research shows that gradient-based algorithms are in particular susceptible to saddle point problems (Ge et al. 2015). Although the study of (Jin et al. 2019), (Lee et al. 2016) indicate that stochastic gradient descent with random noise is enough to escape saddle points, strict saddle property, which is not valid in our case, is assumed in the work above. Therefore, we design specialized algorithms for optimizing Fourier expansions.

Finally, we demonstrate, by experimental results, that for certain class of hybrid formulas, `FourierSAT` is more robust compared to existing tools.

We believe that the natural reduction from SAT to continuous optimization, combined with state-of-the-art constrained-continuous optimization techniques, opens a new line of research on SAT solving.

2 Notations and Preliminaries

Boolean Formulas and Clauses

Let $x = (x_1, \dots, x_n)$ be a sequence of n Boolean variables. A Boolean function $f(x)$ is a mapping from a Boolean vector $\{\text{True}, \text{False}\}^n$ to $\{\text{True}, \text{False}\}$. A vector $a \in \{\text{True}, \text{False}\}^n$ is called an assignment and $f(a)$ denotes the value of f on a . A literal is either a variable x_i or its negation $\neg x_i$. A formula $f = c_1 \wedge c_2 \wedge \dots \wedge c_m$ is the conjunction of m Boolean functions, where each c_i is called a clause and belongs to a type from the list below:

- **CNF**: A CNF clause is a disjunction of elements in a literal set, which is satisfied when at least one literal is true. *E.g.*, $(x_1 \vee x_2 \vee x_3)$.
- **CARD**: Given a set L of variables and an integer $k \in [n]$, a cardinality constraint $D^{\geq k}(L)$ (resp. $D^{\leq k}(L)$) requires the number of `True` variables to be at least (resp. most) k . *E.g.*, $D^{\geq 2}(\{x_1, x_2, x_3\})$: $x_1 + x_2 + x_3 \geq 2$.
- **XOR**: An XOR clause indicates the parity of the number of variables assigned to be `True`, which can be computed by addition on $\text{GF}(2)$. *E.g.*, $x_1 \oplus x_2 \oplus x_3$.
- **NAE**: A Not-all-equal (NAE) constraint is satisfied when not all the variables have the same value. *E.g.*, $\text{NAE}(1, 1, 1, 1, 0) = 1$; $\text{NAE}(0, 0, 0, 0, 0) = 0$

Let the set of clauses of f be $C(f)$. Let $m = |C(f)|$ and n be the number of variables of f . A solution of f is an assignment that satisfies all the clauses in $C(f)$. We aim to design a framework to find a solution of f .

Clause	Fourier Expansion
$x_1 \vee x_2$	$-\frac{1}{2} + \frac{1}{2}x_1 + \frac{1}{2}x_2 + \frac{1}{2}x_1x_2$
$D^{\geq 2}(x_1, x_2, x_3)$	$\frac{1}{2}x_1 + \frac{1}{2}x_2 + \frac{1}{2}x_3 - \frac{1}{2}x_1x_2x_3$
$x_1 \oplus x_2 \oplus x_3$	$x_1x_2x_3$
$\text{NAE}(x_1, x_2, x_3)$	$-\frac{1}{2} + \frac{1}{2}x_1x_2 + \frac{1}{2}x_2x_3 + \frac{1}{2}x_1x_3$

Table 1: Examples of Fourier Expansion on different clauses

Fourier Expansion of a Boolean Function

We define a Boolean function by $f : \{\pm 1\}^n \rightarrow \{\pm 1\}$. One point which might be counter-intuitive is that -1 is used to stand for True and $+1$ for False. An assignment now is a vector $a \in \{\pm 1\}^n$.

Fourier expansion is a method for transforming a Boolean function into a multilinear polynomial. The following theorem shows that any function defined on a Boolean hypercube has an equivalent Fourier representation.

Theorem 1. (Walsh-Fourier Transform) *Given a function $f : \{\pm 1\}^n \rightarrow \mathbb{R}$, there is a unique way of expressing f as a multilinear polynomial, with at most 2^n terms in S , where each term corresponds to one subset of $[n]$, according to:*

$$f(X) = \sum_{S \subseteq [n]} \left(\hat{f}(S) \cdot \prod_{i \in S} x_i \right),$$

where $\hat{f}(S)$'s $\in \mathbb{R}$ are called the Fourier coefficients, given S , and computed as:

$$\hat{f}(S) = \mathbb{E}_{x \sim \{\pm 1\}^n} \left[f(x) \cdot \prod_{i \in S} x_i \right] = \frac{1}{2^n} \sum_{x \in \{\pm 1\}^n} \left(f(x) \cdot \prod_{i \in S} x_i \right)$$

where $x \sim \{\pm 1\}^n$ indicates x is chosen uniformly from $\{\pm 1\}^n$. The polynomial is referred as the **Fourier expansion** of f .

Table 1 shows some examples of Fourier expansions.

3 A Reduction from SAT to Continuous Optimization

In this section, we reduce finding a solution of a Boolean formula to finding a minimum of a continuous multivariate polynomial, based on Fourier expansion. Due to the space limit, we delay most of the proofs to the supplemental material.

A Reduction to a Minimization Problem

Our basic idea is to do continuous optimization on Fourier expansions of formulas. However, in general, computing the Fourier coefficients of a Boolean function is nontrivial and often harder than SAT itself (for an example, see Proposition 1). Even if we are able to get the Fourier expansion of a Boolean formula, the polynomial can have high degree—and as a result, exponentially many terms—making it hard to store and evaluate.

Proposition 1. *Computing the Fourier Expansion of a pseudo-Boolean constraint (a generalization of a cardinality constraint, e.g., $3x_1 + 2x_2 + 7(\neg x_3) \geq 0$, $x \in \{\pm 1\}^3$) is #P-hard.*

Instead of computing Fourier coefficients of a monolithic logic formula, we take advantage of factoring, constructing a polynomial for a formula by the Fourier expansions of its clauses. Excitingly, Fourier expansions of many types of clauses with great interest can be computed easily. Details can be found in the proof of Proposition 2.

Proposition 2. *Fourier expansions of CNF, XOR, NAE clauses and cardinality constraints have closed form representations.*

For a formula f with clause set $C(f)$, we define the *objective function* associated with f , denoted by F_f , by the sum of Fourier expansions of f 's clauses, i.e.,

$$F_f = \sum_{c \in C(f)} \text{FE}_c$$

where FE_c denotes the Fourier expansion of clause c .

The degree (maximum number of variables in all the terms) of F_f equals to the maximum number of literals among all clauses. Note that, in general, F_f is not the Fourier expansion of f . Instead, it can be understood as a substitute of f 's Fourier expansion which is relatively easy to compute.

Let us provide an example to illustrate the above ideas.

Example 1. *Suppose $f = (x_1 \vee x_2) \wedge (x_3) \wedge (x_2 \vee \neg x_4)$. Then, $C(f) = \{x_1 \vee x_2, x_3, x_2 \vee \neg x_4\}$ and*

$$\begin{aligned} F_f &= \left(-\frac{1}{2} + \frac{1}{2}x_1 + \frac{1}{2}x_2 + \frac{1}{2}x_1x_2 \right) + x_3 \\ &\quad + \left(-\frac{1}{2} + \frac{1}{2}x_2 - \frac{1}{2}x_4 - \frac{1}{2}x_2x_4 \right) \\ &= -1 + \frac{1}{2}x_1 + x_2 + x_3 - \frac{1}{2}x_4 + \frac{1}{2}x_1x_2 - \frac{1}{2}x_2x_4. \end{aligned}$$

Notice that the objective function is nothing special but a polynomial. Therefore we relax the domain from discrete to continuous. An assignment is now defined as a real vector $a \in [-1, 1]^n$. The reduction is formalized in Theorem 2.

Theorem 2. (Reduction) *f is satisfiable if and only if*

$$\min_{x \in [-1, 1]^n} F_f(x) = -m.$$

Theorem 2 reduces SAT to a multivariate minimization problem over $[-1, 1]^n$. In the rest of this subsection, we will provide proof sketch of Theorem 2.

Definition 1. (Constant). *A clause is constant if it is equivalent to either True or False. The Fourier expansion of a clause is constant if it always equals to -1 or 1 .*

Lemma 1 indicates the value of multilinear polynomials is well-behaved in the cube $[-1, 1]^n$.

Lemma 1. *Let c be a non-constant clause and a be an assignment. Then:*

1. *if $a_i \in \{-1, 1\}$ for all $i \in [n]$, then $\text{FE}_c(a) \in \{-1, 1\}$.*
2. *if $a_i \in [-1, 1]$ for all $i \in [n]$, then $\text{FE}_c(a) \in [-1, 1]$.*

3. if $a_i \in (-1, 1)$ for all $i \in [n]$, then $FE_c(a) \in (-1, 1)$.

In order to formalize the procedure of converting a fractional solution to a Boolean assignment, we define the concept “feasibility”.

Definition 2. (Partially assigned function) Suppose $[n]$ is partitioned into two sets, J and $[n] - J$. Let $z \in [-1, 1]^{|J|}$ be a real vector, we write $F_{J \leftarrow z}$ for the partially assigned function of F given by fixing the coordinates in J to be the values z in F .

Definition 3. (Feasible Solution). For an objective function F of a Boolean formula and an assignment $a \in [-1, 1]^n$, let $I = \{i \mid a_i \in \{\pm 1\}\}$. a is a **feasible solution** of F if $F_{I \leftarrow a_I}$ is constant, where a_I is the projected vector of a at coordinates in I . We say a is feasible if F is clear in context.

Example 2. Let $f = x_1 \wedge x_2$; then

$$F_f = \frac{1}{2} + \frac{1}{2}x_1 + \frac{1}{2}x_2 - \frac{1}{2}x_1x_2.$$

$a = (1, -0.3)$ is a feasible solution because $I = \{1\}$ and $F_{I \leftarrow a_I} = F_{1 \leftarrow 1} = \frac{1}{2} + \frac{1}{2} + \frac{1}{2}x_2 - \frac{1}{2}x_2 = 1$ is constant not depending on x_2 .

By Definition 3 if we find a feasible solution a^* of the objective function, we can adjust it into a Boolean assignment b^* by modifying values of a^* in $(-1, 1)$ to $\{\pm 1\}$ in b^* arbitrarily. By feasibility of a^* , $F_f(a^*) = F_f(b^*)$.

Lemma 2. Let c be a non-constant clause and a be an assignment. If $FE_c(a) = -1$, then a is a feasible solution of FE_c .

Now we are ready to prove Theorem 2.

Proof of Theorem 2. Note that by the second argument in Lemma 1, we have $F_f(a) \geq -m$ for all $a \in [-1, 1]^n$.

– “ \Rightarrow ”: Suppose f is satisfiable and $\phi \in \{\pm 1\}^n$ is one of its solution. ϕ is also a solution of every clause of f . Thus for every $c \in C(f)$, $FE_c(\phi) = -1$. Therefore $F_f(\phi) = -m$ and $\min_{x \in [-1, 1]^n} F_f(x) = -m$.

– “ \Leftarrow ”: Suppose $\min_{x \in [-1, 1]^n} F_f(x) = -m$. Thus $\exists a^*$ such that $F_f(a^*) = -m$. By the second argument in Lemma 1, we have $FE_c(a^*) = -1$ for every $c \in C(f)$. By Lemma 2 a^* is a feasible solution of FE_c for every $c \in C(f)$. Thus a^* is also a feasible solution of F_f . Therefore, by feasibility, we can get a Boolean assignment which satisfies all the clauses by arbitrarily rounding all the fractional coordinates. \square

Local Landscape of Multilinear Polynomials

In this subsection we characterize the local landscape of multilinear polynomials. We also show that, in our case, local minima are meaningful and useful.

First, we formally define local minimum for constrained problems in Definition 4.

Definition 4. (Local minimum for constrained problem) For a vector $a \in \mathbb{R}^n$, we define $\mathcal{N}_\delta(a)$, the neighborhood of a with radius δ as $\mathcal{N}_\delta(a) = \{x \mid \|a - x\|_2^2 \leq \delta\}$. Given a set Δ , we say an assignment a is a **local minimum of F in Δ** if

$$\exists \delta > 0, \forall a' \in \mathcal{N}_\delta(a) \cap \Delta, F(a) \leq F(a').$$

For a multivariate function, a saddle point is a point at which all the first derivatives are zero but there exist at least one positive and one negative direction. Therefore, a saddle point is not a local minimum in an unconstrained problem.

Lemma 3. Every critical point (where all the first derivatives are zero) of a non-constant multilinear polynomial is a saddle point.

Lemma 3 indicates for unconstrained, non-constant multilinear polynomials, no local minimum exists. On the other hand, after adding the bound $[-1, 1]^n$, intuitively, a local minimum of F in $[-1, 1]^n$ can only appear on the boundary. Lemma 4 uses feasibility again to formalize this intuition.

Lemma 4. If $a^* \in [-1, 1]^n$ is a local minimum of F in $[-1, 1]^n$, then a^* is feasible.

Lemma 4 indicates that every local minimum a^* in $[-1, 1]^n$ is meaningful in the sense that it can be easily converted into a Boolean assignment b^* . Moreover, since b^* is a Boolean assignment, it is easy to see that the number of clauses satisfied by b^* is $\frac{m - F(a^*)}{2}$, which is a special case of Theorem 3 in the next section. Thus if a global minimum is too hard to get, our method is still useful for some problems, e.g., MaxSAT.

In most multilinear polynomials generated from Boolean formulas in practice, saddle points rarely exist. However, there exist “pathological” polynomials that have unaccountably infinite saddle points, as Example 3 shows.

Example 3. Let $F(x) = x_1x_2x_3x_4$. Then every point which has form $(x_1, 0, 0, 0)$, $(0, x_2, 0, 0)$, $(0, 0, x_3, 0)$ or $(0, 0, 0, x_4)$ is a degenerate saddle point where both the gradient and Hessian are zero.

4 Why Continuous?

Before introducing our algorithm for minimizing multilinear polynomials, we would like to illustrate why doing continuous optimization might be better than discrete local search.

Our explanation is, compared to local search SAT solvers which only make progress when the number of satisfied clauses increases, our tool makes progress as long as the value of the polynomial decreases, even by a small amount.

Theorem 3 formalizes the intuition. It indicates that when we decrease the polynomial, we are in fact increasing the expectation of the number of satisfied clauses, after a randomized rounding.

Definition 5. (Randomized Rounding) The randomized rounding function, denoted by $R : [-1, 1]^n \rightarrow \{\pm 1\}^n$ is defined by:

$$\begin{cases} \mathbb{P}(R(a)_i = -1) = -\frac{1}{2}a_i + \frac{1}{2} \\ \mathbb{P}(R(a)_i = +1) = +\frac{1}{2}a_i + \frac{1}{2} \end{cases}$$

where $i \in [n]$ and $a \in [-1, 1]^n$. Note that the closer a_i is to -1 , the more likely $R(a)_i$ will be -1 and vice versa.

Theorem 3. Let f be a formula with m clauses and F be the multilinear polynomial associated with f . For $a \in [-1, 1]^n$, let $R(a) \in \{\pm 1\}^n$ be the vector given by rounding a . Let

$m_{\text{SAT}}(R(a))$ be the number of satisfied clauses by $R(a)$, then

$$\mathbb{E}_{R(a)}(m_{\text{SAT}}(R(a))) = \frac{m - F(a)}{2}.$$

In particular, if $F(a) = -m$ then $\mathbb{E}_{R(a)}(m_{\text{SAT}}(R(a))) = m$.

Since $m_{\text{SAT}}(b) \leq m$ for any $b \in \{\pm 1\}^n$, $m_{\text{SAT}}(R(a)) = m$ and $R(a)$ is a solution of f .

Research showed that local search SAT solvers, such as GSAT, spend most of the time on the so-called “sideway” moves (Selman, Levesque, and Mitchell 1992). Sideway moves are moves that do not increase or decrease the total number of unsatisfied clauses. Although heuristics and adding noise for sideway moves lead the design of efficient solvers, e.g., WalkSAT, local search SAT solvers fail to provide any guarantee when making sideway moves.

It is illustrative to think of a cardinality constraint, e.g., the majority constraint which requires at least half of all the variables to be `True`. If we start from the assignment of all `False`’s, local search solvers need to flip at least half of all bits to make progress. In other words, local search solvers will encounter a neighbourhood with exponential size where their movements will be “blind”. In contrast, guided by the gradient, our method will behave like “flipping” all the variables by a small amount towards the solution.

5 A Gradient Descend-Based Algorithm for Minimizing Multilinear Polynomials

In this section, we propose an algorithm for minimizing multilinear polynomials associated with Boolean formulas. We aim to show promising theoretical guarantees that gradient-based algorithms can enjoy. In practice, there are elaborate packages available and we used them in our experiments.

Since the objective function is constrained, continuous and differentiable, projected gradient descent (PGD) (Nesterov 2014) is a candidate for solving our optimization problem. As a non-convex problem, the initialization plays a key role on the result. We use random initialization for each iteration and return the best result within a time limit.

Efficient Evaluation of Objective Function. Although theoretically the objective function can have up to 2^n terms, we do not actually store all the coefficients and do a naive evaluation to get $F(a)$ for $a \in [-1, 1]^n$. Instead, we leverage the symmetry of clauses to compute the value of Fourier expansion of each clause c , denoted as $\text{FE}_c(a)$, separately. By this trick, we are able to evaluate the objective function in $O(\sum_{c \in \mathcal{C}(f)} k_c^2)$ time (in worst case $O(n^2 m)$), where k_c is the length of clause c . By this trick, we bypass considering data structures that store the multilinear polynomials.

First note that for every negative literal, say $\neg x_i$, we can convert it to positive by flipping the sign of a_i , since applying a negation on formulas is equivalent to adding a minus sign to Fourier expansions (assume each variable appears at most once in a clause c).

Now suppose c is a clause from $\{\text{CNF}, \text{CARD}, \text{XOR}, \text{NAE}\}$ with no negative literals. Then c is symmetric,

which means its Fourier coefficient at S only depends on $|S|$. Thus the set of Fourier coefficients can be denoted as $\text{Coef}(\text{FE}_c) = (\kappa(\emptyset), \kappa([1]), \kappa([2]), \dots, \kappa([k_c]))$. Let $s(a) = (1, \sum_{i=1}^{k_c} a_i, \sum_{i < j} a_i a_j, \dots, \prod_{i=1}^{k_c} a_i)$. One can easily verify that $\text{FE}_c(a) = \text{Coef}(\text{FE}_c) \cdot s(a)$, where “ \cdot ” represents the inner product of two vectors.

$s(a)$ can be obtained by computing the coefficients of t^0, t^1, \dots, t^{k_c} in the expansion of the following polynomial.

$$\prod_{i=1}^{k_c} (a_i + t) = t^{k_c} + \left(\sum_{i=1}^{k_c} a_i \right) \cdot t^{k_c-1} + \dots + \prod_{i=1}^{k_c} a_i \cdot t^0$$

The coefficients of the polynomial above can be computed in $O(k_c^2)$ time, given $a \in [-1, 1]^{k_c}$. Thus we can evaluate the objective function in $O(\sum_{c \in \mathcal{C}(f)} k_c^2)$.

Since the gradient and Hessian of a multilinear polynomial are still multilinear (O’Donnell 2014), we are able to calculate them analytically by the method above. In experiments, we observed feeding gradients to the optimizer significantly accelerated the minimization process.

Our PGD-Based Multilinear Optimization Algorithm.

In Algorithm 1, we propose a PGD-based algorithm for multilinear optimization problem. In gradient descent, the iteration for minimizing an objective function F is:

$$x'_{t+1} = x_t - \eta \cdot \nabla F(x_t), \quad x_{t+1} = x'_{t+1},$$

where $\eta > 0$ is a step size.

For a constrained domain other than \mathbb{R}^n , x'_{t+1} may be outside of the domain. In PGD, we choose the point nearest to x'_{t+1} in $[-1, 1]^n$ as x_{t+1} (He 2016), i.e., the Euclidean projection of x'_{t+1} onto the set $[-1, 1]^n$, denoted as $\Pi_{[-1, 1]^n}(x'_{t+1})$.

Definition 6. The Euclidean projection of a point y , onto a set Δ , denoted by $\Pi_\Delta(y)$, is defined as

$$\Pi_\Delta(y) = \arg \min_{x \in \Delta} \frac{1}{2} \|x - y\|_2^2.$$

In our case $\Delta \equiv [-1, 1]^n$; computing such a projection is almost free, as shown in Proposition 3.

Proposition 3.

$$\Pi_{[-1, 1]^n}(y)_i = \begin{cases} y_i, & \text{if } y_i \in [-1, 1], \\ \text{sgn}(y_i), & \text{otherwise.} \end{cases}$$

Combining the above, the main iteration of PGD can be rewritten as

$$x_{t+1} = \Pi_{[-1, 1]^n}(x_t - \eta \cdot \nabla F(x_t)) = x_t - \eta G(x_t),$$

where $G(x) = \frac{1}{\eta}(x_t - \Pi_{[-1, 1]^n}(x_t - \eta \nabla F(x_t)))$ is regarded as the *gradient mapping*.

In Algorithm 1, we start at a uniformly random point in $[-1, 1]^n$. When gradient mapping $G(\cdot)$ is large, we follow it to decrease the function value. Otherwise, it means the algorithm either reaches a local minimum in $[-1, 1]^n$, or falls

Algorithm 1: PGD for multilinear F

Input : Polynomial F , $\eta > 0$, $\varepsilon > 0$.**Output**: Approximately minimizer \hat{x} .

```
1 for  $j = 1, \dots, J$  do
2    $x_0 \sim \mathcal{U}[-1, 1]^n$ 
3   for  $t = 0, \dots$  do
4      $G(x_t) = \frac{1}{\eta} (x_t - \Pi_{[-1, 1]^n} (x_t - \eta \nabla F(x_t)))$ 
5     if  $\|G(x_t)\|_2 > 0$  then
6        $x_{t+1} = x_t - \eta \cdot G(x_t)$ 
7     else
8       if  $x_t$  not feasible then
9          $x_{t+1} = \text{DecInnerSaddle}(F, x_t, \eta)$ 
10      else
11         $I = \{i \mid (x_t)_i \in \{-1, 1\}\}$ 
12        if  $\nabla F(x_t)_i \neq 0, \forall i \in I$  then
13          Break // Prop. 5
14        else
15          ( $\text{LocalMinFlag}, x_{t+1}$ ) =
16            useHessian( $F, x_t$ )
17          if  $\text{LocalMinFlag} = \text{True}$  or
18             $\text{Unknown}$  then
19            Break
18 Until convergence
19 return  $x_j$  with the lowest  $F(x)$  after  $J$  iterators
```

into a saddle point where the original gradient $\nabla F(\cdot)$ is negligible. If the first case happens we are done for this iterator. Else, we still try to escape the saddle point by additional methods.

In Theorem 4, we show that Algorithm 1 is guaranteed to converge to a point where the projected mapping is small, $\|G(x_t)\|_2 = 0$,¹ and depends polynomially on n , m and tolerance ε . In the proof, we used techniques from (Jin et al. 2019) and (He 2016).

Theorem 4. (Convergence speed) *With the step size $\eta = \frac{1}{nm}$, Algorithm 1 converges to a ε -projected-critical point (where $\|G(x)\|_2 < \varepsilon$) in $O(\frac{nm^2}{\varepsilon^2})$ iterators.*

The proof of Theorem 4 relies on the fact that multilinear polynomials on $[-1, 1]^n$ are relatively smooth as Proposition 4 indicates.

Proposition 4. (Lipschitz) *Let $F : [-1, 1]^n \rightarrow [-m, m]$ be a multilinear polynomial. Then, for every $x, y \in [-1, 1]^n$,*

1. $|F(x) - F(y)| \leq m\sqrt{n} \cdot \|x - y\|_2$. I.e., F is $(m\sqrt{n})$ -Lipschitz continuous.
2. $\|\nabla F(x) - \nabla F(y)\|_2 \leq mn \cdot \|x - y\|_2$. I.e., F has (mn) -Lipschitz continuous gradients.

Proposition 5 shows that in Algorithm 1 we can identify local minima by gradient information.

Proposition 5. *When algorithm 1 reaches line 13, x is a local minimum. i.e, if x is feasible, $G(x) = 0$ and $\nabla F(x)_i \neq 0$ for all $i \in I$, then x is a local minimum in $[-1, 1]^n$ of F .*

¹In practice, a stopping criterion to use is $\|G(x_t)\|_2 < \varepsilon$ for a small accuracy level $\varepsilon > 0$.

We design `DecInnerSaddle` and `useHessian` to escape saddle points. `DecInnerSaddle` uses the idea from Lemma 3 to give a negative direction, as long as the point is not feasible. `useHessian` leverages second order derivatives to give a negative direction, or identifies a local minimum when the first order derivatives are too small. Due to space limit, we leave subprocedures `DecInnerSaddle` and `useHessian`, as well as their properties, in the supplemental material.

Weighted Case

The weighted version of the objective function is defined as

$$F_f = \sum_{c \in C(f)} w_f(c) \cdot \text{FE}_c$$

where $w_f : C(f) \rightarrow \mathbb{R}$ is the weight function. It is easy to verify that the weighted version of Theorem 2 and Lemma 4 still hold. The weighted case is useful in two cases:

- **Vanishing Gradient.** When a clause contains too many of variables (e.g., a global cardinality constraint), we observed that the gradient given by this clause on a single variable becomes negligible at most points. By assigning large weights to long clauses, gradient vanishing can be alleviated.
- **Weighted MAX-SAT.** By returning a local minimum, our tool can be used to solve weighted MAX-SAT problems.

6 Experimental Results

In this section we compare our tool, `FourierSAT` with other SAT solvers on random and realistic benchmarks. The objective of our experimental evaluation is to answer these three research questions:

RQ1. How does `FourierSAT` compare to local search SAT solvers with cardinality constraints and XORs encoded in CNF?

RQ2. How does `FourierSAT` compare to specialized CARD-CNF solvers with respect to handling cardinality constraints?

RQ3. How does `FourierSAT` compare to CDCL SAT solvers with cardinality constraints/XORs encoded in CNF?

Experimental Setup

We choose `SLSQP` (Kraft 1988), implemented in `scipy` (Jones et al. 2001) as our optimization oracle, after comparing available algorithms for non-convex, constrained optimization.

Since random restart is used, different iterators are independent. Thus, we parallelized `FourierSAT` to take advantage of multicore computation resources. Each experiment was run on an exclusive node in a homogeneous Linux cluster. These nodes contain 24-processor cores at 2.63 GHz each with 1 GB of RAM per node. The time cap for each job is 1 minute.

We compare our method with the following SAT solvers:

- `Cryptominisat` (Soos, Nohl, and Castelluccia 2009) (CMS), an advanced SAT solver supporting CNF+XOR clauses by enabling Gauss Elimination.

- WalkSAT (Selman, Kautz, and Cohen 1999), an efficient local search SAT solver. We used a fast C implementation due to (Cai, Su, and Luo 2013).
- MiniCARD (Liffiton and Maglalang 2012), a MiniSAT-based CARD-CNF solver. It has been implemented in `pysat` (Ignatiev, Morgado, and Marques-Silva 2018).
- MonoSAT (Bayless et al. 2015), a SAT Modulo Theory solver which supports a large set of graph predicates.

MiniCARD can handle cardinality constraints natively. For MonoSAT, we reduced cardinality constraints to max-flow. For CMS and WalkSAT, we applied a set of cardinality encodings, which includes Sequential Counter (Sinz 2005), Sorting Network (Batcher 1968), Totalizer (Bailleux and Boufkhad 2003) and Adder (Eén and Sörensson 2006). For solvers that do not support XORs, we used a linear encoding due to (Li 2000) to decompose them into 3-CNF. To address vanishing gradient, the weight of each clause equals to its length.

Benchmarks

We generated hybrid Boolean constraints using natural encodings of the following benchmark problems.

Benchmark 1: Approximate minimum vertex covering. Minimum vertex covering is a well-known NP-optimization problem (NPO) problem. For each $n \in \{50, 100, 150, 200, 250\}$, we generated 100 random cubic graphs. For each graph, Gurobi (Gurobi Optimization 2019) was used to compute the minimum vertex cover, denoted by opt . Then we added one cardinality constraint, $D^{\leq k}(X)$ where $k = \lceil 1.1 \cdot |\text{opt}| \rceil$ to the CNF encoding of vertex covering.

Benchmark 2: Parity learning with error. Parity Learning is to identify an unknown parity function given I/O samples. Technically, in this task one needs to find a solution of an XOR system while tolerating some constraints to be violated. Suppose there are n candidate variables and m I/O samples (XOR constraints). A solution to this problem is allowed to be incorrect on at most $(e \cdot m)$ I/O samples. For $e = 0$ the problem is in P (Gaussian Elimination); for $0 < e < \frac{1}{2}$, whether the problem is in P still remains open. Parity learning is a well-known hard for SAT solvers as well, especially for local search solvers (Crawford and Kearns 1994).

We chose $N \in \{8, 16, 32\}$, $e = \frac{1}{4}$ and $m = 2N$ to generate hard cases. For FourierSAT, this problem can be encoded into solving M XOR clauses where we allow at most eM clauses to be violated. For WalkSAT and CMS, we used the encoding due to (Hoos and Stützle 2000).

Benchmark 3: Random CNF-XOR-CARD formulas. For each $n \in \{50, 100, 150\}$, $r \in \{1, 1.5\}$, $s \in \{0.2, 0.3\}$, $l \in \{0.1, 0.2\}$ and $k \in \{0.4n, 0.5n\}$, we generated random benchmarks with rn 3-CNF clauses, sn XOR clauses with length ln and 1 global cardinality constraint $D^{\leq k}(X)$. Those parameters are chosen to guarantee that all problems




Figure 1: Results on 500 vertex covering problems




Figure 2: Results on 300 parity learning with error problems

are satisfiable (Dudek, Meel, and Vardi 2016) (Pote, Joshi, and Meel 2019).

Results

The experimental results of three benchmarks above are shown in Figure 1, 2 and 3 respectively.

RQ1. Figure 1 shows FourierSAT solved more problems than WalkSAT did with all the cardinality encodings except SeqCounter. For parity learning problem, WalkSAT with CNF encoding only solved 79 problems, while FourierSAT was able to solve all 300 problems. For the random hybrid constraints benchmarks shown in Figure 3, FourierSAT solved 2957 problems while WalkSAT with SeqCounter, SortNetwork and Totalizer solved 2272, 2444 and 2452 problems respectively. In our experiment, FourierSAT is more robust than WalkSAT, which agrees on our intuition in Section 4.

RQ2. In Figure 1 it is clear that FourierSAT solved more problems (472) than MonoSAT (420) and MiniCard (261) did, which indicates that FourierSAT is capable of handling cardinality constraints efficiently.




Figure 3: Results on 4800 random CNF-XOR-CARD problems

RQ3. For the parity learning problem shown in Figure 2, the competition of `FourierSAT` and CMS with CNF+XOR formula is roughly a tie, while `FourierSAT` outperformed CMS with pure CNF formula. For other two benchmarks shown in Figure 1 and 3, we observed that CMS has the best performance, especially for solving random hybrid constraints, despite of the choice of encodings.

Due to the maturity of CMS, it is not surprising that CMS performed better than `FourierSAT` did on some benchmarks, especially when dealing with long XOR clauses. However, the experimental result shows that as a versatile solver, `FourierSAT` is comparable with many existing, well-developed, specialized solvers.

Furthermore, we notice that although `FourierSAT` is usually slower than other solvers on easy problems, it seems to scale better. One potential reason is, due to the continuous nature of `FourierSAT`, it avoids suffering from scaling exponentially too early. This behavior of `FourierSAT` increases our confidence in that algebraic methods are worth exploring for SAT solving in the future.

7 Conclusion and Future Directions

In this paper, we propose a novel algebraic framework for solving Boolean formulas consisting of hybrid constraints. Fourier expansion is used to reduce Boolean Satisfiability to multilinear optimization. Our study on the landscape and properties of multilinear polynomials leads to the discovery of attractive features of this reduction. Furthermore, to show the algorithmic benefits of this reduction, we design algorithms with certain theoretical guarantee. Finally, we implement our method as `FourierSAT`. The experimental results indicate that in practical, `FourierSAT` is comparable with state-of-the-art solvers on certain benchmarks of hybrid constraints by leveraging parallelization.

We also list a few future directions in the following:

- **Complete algebraic SAT solvers.** Besides giving solutions to satisfiable formulas, we also aim to prove unsatisfiability algebraically. In other words, we hope to design

a complete SAT solver based on algebraic approaches.

Several theoretical tools, such as Hilbert’s Nullstellensatz and Gröebner basis (e.g., see (Cox, Little, and O’shea 1994)) can be used for giving an algebraic proof of unsatisfiability. Previous algorithmic work on this direction considered specific polynomial encodings on problems such as graph coloring (Romero Barbosa, Julian 2016) (Loera et al. 2011). We are interested to see the behavior of these algebraic tools on more problems with Fourier transform as the encoding in practice.

- **Escaping local minima and saddle points.** As a local-information-based approach, `FourierSAT` suffers from getting stuck in local minima and relies heavily on random restart as well as parallelization. We believe the idea of “sideway moves” from Local Search SAT solvers (Selman, Levesque, and Mitchell 1992) is a good direction to address these issues. For example, one can use the solution of `WSAT/FourierSAT` as the initial point of `FourierSAT/WSAT`. We will also consider purely first-order gradient descent approaches, probably with random noise (Jin et al. 2019), and how they perform on escaping saddle points both in theory and practice.
- **Alternative objective functions.** The Fourier expansions of clauses with large number of variables and low solution density can be much less smooth, which is one of the main difficulties for `FourierSAT` in practice. To refine the objective function, one possible way is to develop better weight functions, which can be static or dynamic. Due to the techniques for learning Fourier coefficients (Bruck and Smolensky 1992) (Linial, Mansour, and Nisan 1989), it is also promising to use the low-degree approximation of Fourier expansions as the objective function.

Acknowledgement

Work supported in part by NSF grants IIS-1527668, CCF-1704883, and IIS-1830549.

References

- Achim, T.; Sabharwal, A.; and Ermon, S. 2016. Beyond Parity Constraints: Fourier Analysis of Hash Functions for Inference. In *ICML*, 2254–2262.
- Audemard, G., and Simon, L. 2014. Lazy Clause Exchange Policy for Parallel SAT Solvers. In *SAT 2014*, 197–205.
- Bailleux, O., and Boufkhad, Y. 2003. Efficient CNF Encoding of Boolean Cardinality Constraints. In *CP*.
- Balint, A., and Fröhlich, A. 2010. Improving Stochastic Local Search for SAT with a New Probability Distribution. In *SAT 2010*, 10–15.
- Balint, A., and Schöning, U. 2012. Choosing Probability Distributions for Stochastic Local Search and the Role of Make versus Break. In *SAT 2012*, 16–29.
- Batcher. 1968. Sorting Networks and Their Applications. In *Spring Joint Computer Conference*, AFIPS ’68 (Spring).
- Bayless, S.; Bayless, N.; Hoos, H. H.; and Hu, A. J. 2015. SAT Modulo Monotonic Theories. In *AAAI*.
- Biere, A. 2008. PicoSAT Essentials. *JSAT* 4:75–97.

- Biere, A. 2010. Lingeling , Plingeling , PicoSAT and PrecoSAT at SAT Race 2010.
- Bogdanov, A.; Khovratovich, D.; and Rechberger, C. 2011. Biclique Cryptanalysis of the Full AES. In *ASIACRYPT 2011*.
- Bruck, J., and Smolensky, R. 1992. Polynomial Threshold Functions, AC0 Functions, and Spectral Norms. *SIAM J. Comput.* 21(1):33–42.
- Brueggemann, T., and Kern, W. 2004. An Improved Deterministic Local Search Algorithm for 3-SAT. *Theoretical Computer Science* 329(1):303 – 313.
- Cai, S.; Su, K.; and Luo, C. 2013. Improving Walksat for Random k -Satisfiability Problem with $k > 3$. In *AAAI*.
- Costa, M.-C.; de Werra, D.; Picouleau, C.; and Ries, B. 2009. Graph Coloring with Cardinality Constraints on the Neighborhoods. *Discrete Optimization* 6(4):362 – 369.
- Cox, D. A.; Little, J.; and O’shea, D. 1994. Ideals, Varieties and Algorithms.
- Crawford, J. M., and Kearns, M. J. 1994. The Minimal Disagreement Parity Problem as a Hard Satisfiability Problem.
- Davis, M., and Putnam, H. 1960. A Computing Procedure for Quantification Theory. *Journal of the ACM (JACM)* 7(3):201–215.
- Davis, M.; Logemann, G.; and Loveland, D. 1962. A Machine Program for Theorem-Proving. *Communications of the ACM* 5(7):394–397.
- Dinur, I.; Regev, O.; and Smyth, C. 2005. The Hardness of 3-Uniform Hypergraph Coloring. *Combinatorica* 25(5):519–535.
- Dudek, J. M.; Meel, K. S.; and Vardi, M. Y. 2016. Combining the k -CNF and XOR Phase-Transitions. In *IJCAI*.
- Eén, N., and Sörensson, N. 2003. An Extensible SAT-Solver. In *SAT*, 502–518.
- Eén, N., and Sörensson, N. 2006. Translating Pseudo-Boolean Constraints into SAT. *JSAT* 2:1–26.
- Elffers, J., and Nordström, J. 2018. Divide and Conquer: Towards Faster Pseudo-Boolean Solving.
- Feldman, J.; Wainwright, M. J.; and Karger, D. R. 2005. Using Linear Programming to Decode Binary Linear Codes. *IEEE Transactions on Information Theory* 51(3):954–972.
- Ge, R.; Huang, F.; Jin, C.; and Yuan, Y. 2015. Escaping From Saddle Points — Online Stochastic Gradient for Tensor Decomposition. *arXiv e-prints* arXiv:1503.02101.
- Goldberg, E., and Novikov, Y. 2007. BerkMin: A Fast and Robust Sat-solver. *Discrete Applied Mathematics* 155(12):1549 – 1561. SAT 2001.
- Gong, W., and Zhou, X. 2017. A Survey of SAT Solver. *AIP Conference Proceedings* 1836(1):020059.
- Gurobi Optimization, L. 2019. Gurobi Optimizer Reference Manual.
- He, N. 2016. Lecture notes in IE 598: Big Data Optimization.
- Hoos, H., and Stützle, T. 2000. SATLIB: An Online Resource for Research on SAT.
- Ignatiev, A.; Morgado, A.; and Marques-Silva, J. 2018. PySAT: A Python Toolkit for Prototyping with SAT Oracles. In *SAT*, 428–437.
- Jin, C.; Netrapalli, P.; Ge, R.; Kakade, S. M.; and Jordan, M. I. 2019. Stochastic Gradient Descent Escapes Saddle Points Efficiently. *arXiv e-prints* arXiv:1902.04811.
- Jones, E.; Oliphant, T.; Peterson, P.; et al. 2001. SciPy: Open source scientific tools for Python.
- Kautz, H.; McAllester, D.; and Selman, B. 1997. Exploiting Variable Dependency in Local Search. In *Abstracts of the Poster Sessions of IJCAI-97*, 9–7.
- Kraft, D. 1988. A Software Package for Sequential Quadratic Programming. *Forschungsbericht- Deutsche Forschungs- und Versuchsanstalt für Luft- und Raumfahrt*.
- Lee, J. D.; Simchowitz, M.; Jordan, M. I.; and Recht, B. 2016. Gradient Descent Converges to Minimizers. *arXiv e-prints* arXiv:1602.04915.
- Li, C. M. 2000. Integrating Equivalency Reasoning into Davis-Putnam Procedure. In *AAAI*, 291–296. AAAI Press.
- Liang, J. H. 2018. *Machine Learning for SAT Solvers*. Ph.D. Dissertation, University of Waterloo.
- Liffiton, M. H., and Maglalat, J. C. 2012. A Cardinality Solver: More Expressive Constraints for Free. In *SAT 2012*.
- Linial, N.; Mansour, Y.; and Nisan, N. 1989. Constant Depth Circuits, Fourier Transform, and Learnability. In *ASFCs*.
- Loera, J. A. D.; Lee, J.; Malkin, P. N.; and Margulies, S. 2011. Computing Infeasibility Certificates for Combinatorial Problems through Hilbert’s Nullstellensatz. *J. Symb. Comput.* 46:1260–1283.
- Marques-Silva, J. P., and Sakallah, K. A. 1999. GRASP: A Search Algorithm For Propositional Satisfiability. *IEEE Transactions on Computers* 48(5):506–521.
- Martins, R.; Manquinho, V.; and Lynce, I. 2011. Exploiting Cardinality Encodings in Parallel Maximum Satisfiability. In *ICTAI*, 313–320.
- McAllester, D.; Selman, B.; and Kautz, H. 1997. Evidence for Invariants in Local Search.
- Nesterov, Y. 2014. Introductory lectures on convex optimization: A basic course.
- O’Donnell, R. 2014. *Analysis of Boolean Functions*. New York, NY, USA: Cambridge University Press.
- Pote, Y.; Joshi, S.; and Meel, K. S. 2019. Phase Transition Behavior of Cardinality and XOR Constraints. In *IJCAI-19*.
- Prestwich, S. 2009. CNF Encodings, Handbook of Satisfiability: Volume 185 Frontiers in Artificial Intelligence and Applications.
- Quimper, C.-G., and Walsh, T. 2007. Decomposing Global Grammar Constraints. In *CP 2007*, 590–604.
- Romero Barbosa, Julian. 2016. Applied Hilbert’s Nullstellensatz for Combinatorial Problems.
- Selman, B.; Kautz, H.; and Cohen, B. 1999. Local Search Strategies for Satisfiability Testing. *Second DIMACS Implementation Challenge* 26.
- Selman, B.; Levesque, H.; and Mitchell, D. 1992. A New Method for Solving Hard Satisfiability Problems.
- Sheini, H. M., and Sakallah, K. A. 2006. Pueblo: A hybrid pseudo-boolean sat solver. *JSAT* 2:165–189.
- Sinz, C. 2005. Towards an Optimal CNF Encoding of Boolean Cardinality Constraints. In *CP 2005*, 827–831.
- Soos, M.; Nohl, K.; and Castelluccia, C. 2009. Extending SAT Solvers to Cryptographic Problems. In *SAT*, 244–257.
- Warners, J. P., and van Maaren, H. 1998. A Two-phase Algorithm for Solving a Class of Hard Satisfiability Problems. *Operations Research Letters* 23(3):81 – 88.
- Wei, C., and Ermon, S. 2017. General Bounds on Satisfiability Thresholds for Random CSPs via Fourier Analysis.
- Xue, Y.; Ermon, S.; Le Bras, R.; Gomes, C. P.; and Selman, B. 2015. Variable Elimination in the Fourier Domain. *arXiv e-prints* arXiv:1508.04032.