



CNF Satisfiability in a Subspace and Related Problems

V. Arvind¹ · Venkatesan Guruswami²

Received: 15 October 2021 / Accepted: 6 March 2022 / Published online: 31 March 2022

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2022

Abstract

We introduce the problem of finding a satisfying assignment to a CNF formula that must further belong to a prescribed input subspace. Equivalent formulations of the problem include finding a point outside a union of subspaces (the Union-of-Subspace Avoidance (USA) problem), and finding a common zero of a system of polynomials over \mathbb{F}_2 each of which is a product of affine forms. We focus on the case of k -CNF formulas (the k -SUB-SAT problem). Clearly, k -SUB-SAT is no easier than k -SAT, and might be harder. Indeed, via simple reductions we show that 2-SUB-SAT is NP-hard, and W[1]-hard when parameterized by the co-dimension of the subspace. We also prove that the optimization version Max-2-SUB-SAT is NP-hard to approximate better than the trivial 3/4 ratio even on satisfiable instances. On the algorithmic front, we investigate fast exponential algorithms which give non-trivial savings over brute-force algorithms. We give a simple branching algorithm with running time $O^*(1.5)^r$ for 2-SUB-SAT, where r is the subspace dimension, as well as an $O^*(1.4312)^n$ time algorithm where n is the number of variables. Turning to k -SUB-SAT for $k \geq 3$, while known algorithms for solving a system of degree k polynomial equations already imply a solution with running time $\approx 2^{r(1-1/2k)}$, we explore a more combinatorial approach. Based on an analysis of critical variables (a key notion underlying the randomized k -SAT algorithm of Paturi, Pudlak, and Zane), we give an algorithm with running time $\approx \binom{n}{\leq t} 2^{n-n/k}$ where n is the number of variables and t is the co-dimension of the subspace. This improves upon the running time of the polynomial equations approach for small co-dimension. Our combinatorial approach also achieves

A conference version of this work was presented at the 2021 International Symposium on Parameterized and Exact Computation (IPEC). Portions of this work were done during visits to the Institute of Mathematical Sciences, Chennai. Research supported in part by the US National Science Foundation Grant CCF-1908125 and a Simons Investigator Award.

✉ V. Arvind
arvind@imsc.res.in
Venkatesan Guruswami
venkatg@cs.cmu.edu

¹ The Institute of Mathematical Sciences (HBNI), Chennai, India

² Computer Science Department, Carnegie Mellon University, Pittsburgh, USA

polynomial space in contrast to the algebraic approach that uses exponential space. We also give a PPZ-style algorithm for $k - \text{SUB} - \text{SAT}$ with running time $\approx 2^{n-n/2k}$. This algorithm is in fact oblivious to the structure of the subspace, and extends when the subspace-membership constraint is replaced by any constraint for which partial satisfying assignments can be efficiently completed to a full satisfying assignment. Finally, for systems of $O(n)$ polynomial equations in n variables over \mathbb{F}_2 , we give a fast exponential algorithm when each polynomial has bounded degree irreducible factors (but can otherwise have large degree) using a degree reduction trick.

Keywords Satisfiability problem · CNF formulas · Linear equations · Fast exponential algorithms

Contents

1	Introduction	3277
1.1	Hardness Results	3278
1.2	Algorithmic Results	3279
1.2.1	Algorithms for 2-Sub-Sat	3280
1.2.2	Algorithms for k -Sub-Sat	3281
1.3	Equivalent and Related Problems to Sub-Sat	3282
2	Algorithmic Results for k -Sub-Sat	3283
2.1	An $O^*(\binom{n}{t} \cdot 2^{n(1-1/k)})$ Time Randomized Algorithm: Co-dimension t Case	3283
2.2	An $O^*(2^{n-n/2k+n/2k^2})$ Time PPZ-Based Algorithm for k -Sub-Sat	3287
2.3	An $O^*(1.5^t)$ time algorithm for 2-Sub-Sat	3289
2.4	2-Sub-Sat in a Co-dimension t Subspace	3290
3	Hardness Results	3292
3.1	NP-Hardness of 2-Sub-Sat	3292
3.2	W[1]-Hardness of 2-Sub-Sat Parameterized by Co-dimension	3293
3.3	Approximability of Max-2-Sub-Sat	3294
3.3.1	Easy Approximation Algorithms	3294
3.3.2	Tight Inapproximability via Simple Reductions	3294
3.3.3	Inapproximability for Satisfiable Instances	3295
4	System of Polynomial Equations Over Binary Field: Effect of Reducibility	3296
	References	3298

1 Introduction

Given an n -variate Boolean formula Φ along with an affine subspace $A \subseteq \mathbb{F}_2^n$ (given by a system of \mathbb{F}_2 -linear equations) as input, we explore the complexity of testing if Φ has a satisfying assignment in A . This is a natural twist on Boolean constraint satisfaction problems that studies the effects of linear algebra on Boolean logic. Our focus shall be on the case when Φ is presented in Conjunctive Normal Form (CNF). We refer to this problem as *satisfiability in a subspace* and denote it by $\text{SUB} - \text{SAT}$. This framework can capture non-Boolean problems such as Graph K -Colorability indicating the richness of combining the problem of Boolean CNF-satisfiability with a linear-algebraic constraint. The combination of linear and boolean constraints in satisfiability problems have been studied earlier. Chen and Santhanam [6] studied the satisfiability problem for mixed instances which is a more general framework

that includes the SUB – Sat problem. Also, Lokshtanov et al. [19] have studied the algorithmic problem of solving polynomial equations over finite fields which is a more general problem to which SUB – Sat is easily reducible. We also note that in the area of practical SAT solvers there is interest in CNF satisfiability conjuncted with XOR constraints [26, 27].

Further, SUB – Sat has two other equivalent interesting formulations. The first of these is *union of subspace avoidance*, USA for short: Given affine subspaces $A_1, A_2, \dots, A_m \subseteq \mathbb{F}_2^n$ is there an $x \in \mathbb{F}_2^n$ that is not in the union $\bigcup_{i=1}^m A_i$? A different formulation is a special case of finding a solution to a bunch of polynomial equations $p_i = 0$ over \mathbb{F}_2^n , namely when each p_i is a product of affine forms. We refer to this reformulation as PAF – SAT. We will describe these (easy) equivalences in Sect. 1.3.

For most of the paper, we restrict attention to the case when Φ is a k -CNF formula (a CNF formula with clauses of width at most k) for a fixed k , referred to as the k – SUB – SAT problem. Clearly, k – SUB – SAT is a generalization of the well-studied k -SAT (k -CNF satisfiability). In terms of the two reformulations above, k – SUB – SAT corresponds to the USA problem when the spaces A_i have co-dimension at most k , and for the PAF – SAT problem, each polynomial p_i is the product of up to k affine forms.

We present both hardness results and algorithms for k – SUB – SAT, described in Sects. 1.1 and 1.2 below respectively. Owing to the NP-hardness of the problems, the algorithmic focus is on exponential time algorithms that give non-trivial improvements over brute-force.

There are two possible angles from which to view the study of k – SUB – SAT. The first is as a problem intermediate between satisfiability of k -CNF formula and a system of degree k polynomial equations. The second is as a specific instance of a constraint satisfaction problem (CSP) obtained by combining two fundamental types of constraints. There have been a few works [6, 21] giving algorithms beating brute-force for some natural problems with mixed constraints, but we are still far from a general picture of how to obtain fast exponential algorithms for a combined template of constraints when each constraint type does admit such non-trivial algorithms. In this context, tackling the combination of k -CNF formulas and linear equations is a good starting point, and one that could hopefully spur a more systematic study in the future. There have been a few investigations [7, 15, 16, 18] into the fine-grained complexity of CSPs via the algebraic approach based on (partial) polymorphisms. This theory has developed the tools to compare the optimal exponents of different constraint types, identifying for instance the “easiest” NP-hard CSP within some classes. However, with the exception of [3], polymorphisms have not been leveraged to design fast exponential algorithms with competitive exponents.

1.1 Hardness Results

Since k – SUB – SAT is a generalization of k -SAT, k – SUB – SAT inherits all the intractability results of k -SAT for $k \geq 3$. This leaves the interesting case of 2 – SUB – SAT. This turns out to be much harder than the polynomial time solvable 2 – SAT. We establish the following theorem, showing not just hardness (even for FPT algorithms)

of the exact version, but also a tight inapproximability for the approximation version (even on satisfiable instances). The proofs (given in Sects. 3.1, 3.2) are based on short, simple reductions, once an appropriate problem to reduce from is chosen.¹ The W[1]-hardness answers a question posed in [2] on the fixed-parameter complexity of 2-SAT with a global modular constraint, parameterized by the modulus.

Theorem 1 *1. 2-SUB-SAT is NP-hard. It is further W[1]-hard when parameterized by the co-dimension of the affine space A in which we seek a satisfying assignment.*
2. Given a satisfiable instance of 2-SUB-SAT, it is NP-hard to find an assignment in the input space A that satisfies more than $3/4 + \epsilon$ of the 2SAT clauses, for any $\epsilon > 0$.

1.2 Algorithmic Results

Analogous to seeking k -SAT algorithms faster than brute-force, we investigate fast exponential time algorithms for k -SUB-SAT that beat the naive brute-force $2^{\dim(A)}$ time algorithm, where $A \subseteq \mathbb{F}_2^n$ is the subspace in which we seek a solution. Algorithms for k -SAT have received much attention and are central to the burgeoning field of fast exponential-time algorithms. The algorithmic theory is closely connected to fixed parameter tractability and parameterized complexity [9, 11]. The accompanying hardness theory [13, 14], based on the exponential-time hypothesis (ETH) and the strong exponential-time hypothesis (SETH), is a sanity check to the quest for faster algorithms for k -SAT and other NP-complete problems.

There are several interesting k -SAT algorithms with running time $O^*(2^{n(1-\Theta(1/k))})$.² We only mention two significant algorithms from among these: one by Paturi, Pudlak, Zane [22] and another due to Schöning [24]. Both algorithms are simple to describe with delightfully clever and elegant analyses. The PPZ algorithm considers variables in a random order, and gives each a random value unless its value is forced by a clause and previously set values. It achieves a running time of $O^*(2^{n(1-1/k)})$. Schöning’s algorithm starts with a random assignment and in each step fixes an unsatisfied clause by flipping the value of a random one of its variables. It achieves a running time of $O^*((2 - 2/k)^n)$.

Given that k -SUB-SAT generalizes k -SAT, it is natural to seek exponential algorithms with similar running times for k -SUB-SAT. For SUB-SAT with input space $A \subseteq \mathbb{F}_2^n$, the brute-force algorithm in fact runs in time $O^*(2^{\dim(A)})$. A natural question is whether we can get similar improvements in the exponent of the $O^*(2^{\dim(A)})$ running time.

An algorithm [19] with running time about $O^*(2^{r(1-1/5k)})$ is known for checking satisfiability of a collection of arbitrary degree k polynomial equations in r variables: Let $P_i \in \mathbb{F}_2[x_1, x_2, \dots, x_r]$, $1 \leq i \leq m$, be polynomials over the field \mathbb{F}_2 . Following [19], the POLY-EQS problem is solving the system of polynomial equations $P_i = 0$, $1 \leq i \leq m$ over \mathbb{F}_2^r : to check if there exists a solution in \mathbb{F}_2^r and compute one

¹ The NP-hardness would also follow from Schaefer’s dichotomy theorem for Boolean CSP [23], though that is an overkill hammer for this result.

² The notation $O^*(f(n))$ for running time bounds suppresses polynomial factors.

if it exists. When P_i are all of degree bounded by k we denote this special case by k -POLY-Eqs. The k -POLY-Eqs problem generalizes k -SUB-SAT by the following easy transformation: Suppose the subspace A where we seek a satisfying assignment is r dimensional. Then we can express the i th clause in the k -SUB-SAT instance as a disjunction of k affine linear forms in r variables: $C_i = (\ell_{i,1} \vee \ell_{i,2} \vee \dots \vee \ell_{i,k})$. We define the corresponding polynomial $P_i = \prod_{j=1}^k (\ell_{i,j} + 1)$. Now, the k -SUB-SAT instance is satisfiable iff the k -POLY-Eqs instance $P_i = 0, 1 \leq i \leq m$ has a solution in \mathbb{F}_2^r .

The algorithm [19] is a novel application of the Razborov-Smolensky “polynomial method,” originally developed as a lower bound technique, used to define low-degree probabilistic polynomials for approximating the OR gate. The same idea allows for replacing a system of polynomial equations by a single probabilistic polynomial (without significant increase in degree), followed by a partial table lookup search. The article [19] presents more general results applicable to all finite fields \mathbb{F}_q . Recently, in [8], the running time for the case of \mathbb{F}_2 has been improved to $O^*(2^{r(1-1/2^k)})$ by a refinement of the search method in [19].

Since k -SUB-SAT is a special case of solving a system of polynomial equations over \mathbb{F}_2 , it raises the natural question of improving the running time further to match the $O^*(2^{r(1-1/k)})$ running time of the PPZ randomized algorithm for k -SAT. We are only able to achieve this speed-up in some special cases. However, on the positive side, our algorithms turn out to be *polynomial space bounded*, unlike the polynomial equations based method which requires exponential space [8, 19].

1.2.1 Algorithms for 2-Sub-Sat

For 2 -SUB-SAT a simple deterministic branch-and-bound algorithm achieves a running time of $O^*(3^{r/2})$ where r is the dimension of the subspace A . We can improve on this with a randomized branching strategy to a running time of $O^*(1.5^r)$. This improves over the randomized $O^*(1.6181^r)$ algorithm given by the polynomial method [8] for solving a system of quadratic equations over \mathbb{F}_2 . There is also a simple deterministic branching algorithm with $O^*((1+\sqrt{5})/2)^r$ running time for 2 -SUB-SAT. This is based on the same branching strategy for k -SAT [20, Theorem, pp. 295] with its running time governed by the generalized Fibonacci numbers.

When $\dim(A) = n - t$, we can adapt the algorithm from [2, Algorithm 4.1] (for solving 2-SAT with a single abelian group constraint) to obtain an $O^*(\binom{n}{\leq t})$ time algorithm.³

The result of Theorem 1 shows that this problem is not in FPT parameterized by the co-dimension t , answering a question posed in [2] on whether 2-SAT with a global abelian group constraint might be fixed-parameter tractable, parameterized by the group size. More generally, the work [2] systematically studied the effect of a global modular constraint on the complexity of Boolean constraint satisfaction problems, exposing many interesting phenomena and connections.

Balancing the two running times of $O^*(1.5^r)$ and $O^*(\binom{n}{n-r})$ algorithm when $r \geq n/2$ (the exponents of the two bounds become equal at $r = (1-\eta)n$ for $\eta \approx 0.115816$)

³ For nonnegative integers n, t , the notation $\binom{n}{\leq t}$ stands for $\sum_{i=0}^t \binom{n}{i}$.

yields a $O^*(1.4312^n)$ time randomized algorithm for 2 – SUB – SAT on n variables. The following records these results (proof is in Sect. 2.3).

Theorem 2 *There is a randomized $O^*(1.5^r)$ algorithm for 2 – SUB – SAT where r is the dimension of the input space, as well as a deterministic $O^*\left(\binom{n}{\leq t}\right)$ time algorithm where t is the co-dimension. Together, these imply a randomized $O^*(1.4312^n)$ time algorithm as a function of the number n of variables.*

1.2.2 Algorithms for k -Sub-Sat

We explore combinatorial algorithms for k – SUB – SAT based on the notion of *critical variables* (which was introduced in [22] and plays an important role in their satisfiability algorithm). Let Φ be a satisfiable CNF formula in n variables $x_i, i \in [n]$, and let $\bar{a} \in \mathbb{F}_2^n$ be a satisfying assignment.

Definition 3 [22] We say x_i is a *critical variable* for \bar{a} with respect to Φ if the assignment $\bar{a} + e_i$ falsifies Φ , where e_i is the i th elementary vector with 1 in the i th coordinate and zero elsewhere (so $\bar{a} + e_i$ is just \bar{a} with x_i flipped). If the formula Φ is clear from context, we simply say that x_i is a critical variable for assignment \bar{a} .

The key idea in our combinatorial algorithms is *plucking* of non-critical variables based on the following simple observation: if Φ is an n -variate CNF formula and \bar{a} is a satisfying assignment such that variable x_i is non-critical for it, then the formula Φ' obtained by plucking x_i (i.e., dropping all occurrences of x_i and its complement from Φ) remains satisfiable with $\bar{a}' \in \mathbb{F}_2^{n-1}$ as a satisfying assignment, where \bar{a}' is obtained from \bar{a} by dropping the i th coordinate.

The important property of Φ' is that given any satisfying assignment for Φ' we can set x_i to either 0 or 1 to recover a satisfying assignment for Φ . This facilitates searching for a satisfying assignment in an affine space A : if the plucked variable x_i occurs in a linear constraint defining A then we can drop that linear constraint while seeking a satisfying assignment for Φ' , because that linear constraint can always be satisfied by choosing the right value of x_i which still remains overall a satisfying assignment for Φ . Based on this idea we obtain the following algorithms for k – SUB – SAT:

- The first result here is a randomized $O^*\left(\binom{n}{t}2^{n-n/k}\right)$ time algorithm for k – SUB – SAT where $t = \text{codim}(A)$. This algorithm is essentially governed by the running time of the PPZ satisfiability algorithm [22] combined with an iterative “search and pluck” operation to remove t non-critical variables from the t linear equations defining A . This running time is superior to the $O^*(2^{r-r/2k})$ time randomized algorithm based on solving polynomial equations for small values of $t = o(n)$. This result is presented in Sect. 2.1.
- The second result is a general randomized $O^*(2^{n-n/2k+n/2k^2})$ time algorithm for k – SUB – SAT, nearly matching the $\approx 2^{r-r/2k}$ running time of the polynomial equations algorithm [8, 19] for r close to n . It again uses the PPZ satisfiability algorithm as a subroutine combined with simple applications of the plucking step: if the number of critical variables is fewer than $n/2$, it randomly guesses and plucks non-critical variables. This algorithm does not need to look at the linear

equations defining A . In fact, it works for any Boolean constraint $C(x_1, x_2, \dots, x_n)$ (replacing membership in the affine space A) with a polynomial-time algorithm that takes a partial assignment and extends it to an assignment that satisfies C . For example, C can be a Horn or dual Horn formula.

- It is pleasing to note that we can apply the idea of plucking non-critical variables to 2-SUB-SAT and obtain an $O^*\binom{n}{\leq t}$ deterministic algorithm (cf. [2]), where $t = \text{codim}(A)$. Exploiting the structure of 2-CNF formulas, we can find the non-critical variables efficiently.

The proof of the following result is presented in Sect. 2.2.

Theorem 4 *The k -SUB-SAT problem admits two randomized algorithms, one running in time $O^*(2^{n-n/2k+n/2k^2})$, and another running in $O^*\binom{n}{t}2^{n-n/k}$ when the input subspace has co-dimension $t \leq n/2$.⁴ Both algorithms use space bounded by a polynomial in n .*

Remark 5 Satisfiability algorithms based on the switching lemma (which converts k -CNF to decision trees of moderate term size and number of terms) are known in the literature (e.g., see [12]). We can easily adapt this algorithm to solve k -SUB-SAT, because once we have a decision tree for the underlying k -CNF formula, for the k -SUB-SAT instance each path of the decision tree will give rise to a system of linear equations over \mathbb{F}_2 . For each path, therefore, we can even count the number of satisfying assignments. Counting over all the paths of the decision tree gives the total number of satisfying assignments for the k -SUB-SAT instance in randomized time $O^*(2^{n(1-1/c \cdot k)})$ for some suitable large constant $c > 0$. Furthermore, the algorithm is also polynomial space-bounded. In terms of running time, however, it is a much weaker bound in comparison to [19] or even the algorithms of Theorem 4. In this context, we note that for $\#k$ -SAT there is a *deterministic* $O^*(2^{n(1-1/c \cdot k)})$ time algorithm based on the polynomial method (albeit using exponential space) [5]. We do not know of any such deterministic algorithm for counting satisfying assignments to k -SUB-SAT.

Finally, motivated by the (unbounded CNF) SUB-SAT problem, we revisit the general problem solving a system of polynomial equations $p_i = 0$, $1 \leq i \leq m$ over \mathbb{F}_2 , where $m = O(n)$, where each p_i is given by an arithmetic circuit of $\text{poly}(n)$ degree. In the case when each p_i has small degree irreducible factors, we get a $2^{r(1-\alpha)}$ time randomized algorithm, where α depends on the number of equations m and the degree bound on the irreducible factors (Theorem 28).

1.3 Equivalent and Related Problems to Sub-Sat

Recall the USA problem: Given a collection of affine subspaces $A_1, A_2, \dots, A_m \subseteq \mathbb{F}_2^n$ (where each A_i is given by a bunch of affine linear equations over \mathbb{F}_2) the problem is to determine if there is a point $x \in \mathbb{F}_2^n \setminus \bigcup_{i=1}^m A_i$.

Clearly, the complement $\mathbb{F}_2^n \setminus \bigcup_{i=1}^m A_i$ is expressible as an AND of ORs of affine linear forms $\bigoplus_{i \in S} x_i + b$, $b \in \{0, 1\}$. Thus, USA is clearly reducible to SUB-SAT. The

⁴ Of course, there is also a trivial $O^*(2^{n-t})$ time brute force algorithm.

converse reduction is also easy: given a CNF formula Φ and an affine subspace $A \subseteq \mathbb{F}_2^n$ we first convert it to an AND of ORs of affine linear forms. An assignment $x \in A$ satisfies Φ if and only if it satisfies $C_1 \wedge C_2 \wedge \dots \wedge C_m$, where each clause C_i is an OR of affine linear forms. The set A_i of satisfying assignments of the complement $\overline{C_i}$ is an affine subspace of \mathbb{F}_2^n , and Φ is satisfiable by $x \in A$ if and only if $x \in \mathbb{F}_2^n \setminus \bigcup_{i=1}^m A_i$.

For the equivalence to PAF – SAT, suppose $\Phi = C_1 \wedge C_2 \wedge \dots \wedge C_m$, where each clause C_i is an OR of affine linear forms $C_i = \vee_{j=1}^t L_{ij}$. As already discussed in Sect. 1.2, the assignment $x \in \mathbb{F}_2^n$ satisfies C_i if and only if it satisfies the polynomial equation $\prod_{j=1}^m (L_{ij} + 1) = 0$. Thus, the satisfiability of Φ is reducible to a system of m polynomial equations $p_i = 0$, where each p_i is a product of affine linear forms. The converse reduction is also easy which we omit.

Organization of the paper We present the results in a different order than in the introduction. In Sect. 2 we first present the algorithms for k – SUB – SAT and then for 2 – SUB – SAT. In Sect. 3 we present our hardness results for 2 – SUB – SAT. Finally, in Sect. 4 we present the algorithm for POLY – EQS for $O(n)$ equations $p_i = 0$, where each p_i has unrestricted degree but constant-degree irreducible factors.

2 Algorithmic Results for k -Sub-Sat

As mentioned in the introduction, the k – SUB – SAT problem seems intermediate in difficulty, between k – SAT and the problem k – POLY – Eqs of solving a system of degree- k polynomial equations over \mathbb{F}_2 . The latter problem has an $O^*(2^{r(1-1/2k)})$ time algorithm [1, 8, 19], which yields an $O^*(2^{r(1-1/2k)})$ time algorithm for k – SUB – SAT, where $r = \dim(A)$.

Ideally, we would like an algorithm for k – SUB – SAT with running time $O^*(2^{r(1-1/k)})$, with savings in the exponent similar to that of the PPZ algorithm [22] for k -SAT.

We present some algorithms in this direction: For 2 – SUB – SAT there is a simple $O^*(1.5^r)$ time randomized algorithm which improves on the $O^*(2^{r(1-1/2k)})$ bound for $k = 2$. For a special case of k – SUB – SAT, when $r = \dim(A)$ is close to the number of variables n , we are able to adapt the PPZ algorithm to essentially get an $O^*(2^{r(1-1/2k)})$ time algorithm. Writing $t = n - r = \text{codim}(A)$, we can even obtain an $O^*(\binom{n}{\leq t} \cdot 2^{n(1-1/k)})$ time algorithm for the problem, also based on the PPZ satisfiability algorithm, which yields the desired $1/k$ savings in the exponent for small t .

2.1 An $O^*(\binom{n}{t} \cdot 2^{n(1-1/k)})$ Time Randomized Algorithm: Co-dimension t Case

As outlined in Sect. 1.2, the algorithm will use the PPZ satisfiability algorithm [22] as a subroutine, combined with variable plucking steps to solve k – SUB – SAT in randomized time $O^*(\binom{n}{t} \cdot 2^{n(1-1/k)})$, when $\text{codim}(A) = t$. In particular, for $\text{codim}(A) = o(n)$ the algorithm has running time $O^*(2^{n(1-1/k+o(1))})$.

The variable plucking is based on analyzing the critical variables for a solution $\bar{a} \in \mathbb{F}_2^n$ of a given k – SUB – SAT instance (Φ, A) , depending on whether or not they occur in the linear equations defining A .

For an instance (Φ, A) we partition the variables into two sets

$$\{x_i \mid i \in [n]\} = V_{\text{in}} \sqcup V_{\text{out}},$$

where V_{in} is the subset of variables that have nonzero coefficient in at least one of the t linear equations defining A , and V_{out} is the remaining set of variables. By abuse of notation, we will also treat $V_{\text{in}} \sqcup V_{\text{out}}$ as a partition of the index set $[n]$. We consider the following two cases.

Case 1 Suppose (Φ, A) has the property that for every solution $\bar{a} \in \mathbb{F}_2^n$ each variable in V_{in} is critical for \bar{a} w.r.t Φ . There is no variable plucking required in this case. It only involves the application of the PPZ satisfiability algorithm on Φ and checking that the assignment found belongs to A . We need the following lemma which is analogous to [22, Lemma 4]. The proof of the lemma is by an induction argument like in [22].

Lemma 6 *Let S be a nonempty subset of \mathbb{F}_2^n . For each $\bar{a} \in S$, let $I_{\text{out}}(\bar{a}) = \{i \in V_{\text{out}} \mid \bar{a} + e_i \notin S\}$, where e_i is the i th elementary vector. Then we have*

$$\sum_{\bar{a} \in S} 2^{|I_{\text{out}}(\bar{a})| - |V_{\text{out}}|} \geq 1. \quad (1)$$

Proof If $|V_{\text{out}}| = 0$, then $I_{\text{out}}(\bar{a}) = \emptyset$ for every $\bar{a} \in S$, and the left hand side of (1) equals $|S|$ which is at least 1.

So assume $|V_{\text{out}}| \geq 1$ and without loss of generality that $1 \in V_{\text{out}}$. Let $S_0 = \{\bar{a} \in S \mid a_1 = 0\}$ and $S_1 = \{\bar{a} \in S \mid a_1 = 1\}$, and also denote $V'_{\text{out}} = V_{\text{out}} \setminus \{1\}$.

First consider the case when both S_0 and S_1 are nonempty. For $\bar{a} \in S_0$, define $I_{\text{out}}^{(0)}(\bar{a}) = \{i \in V'_{\text{out}} \mid \bar{a} + e_j \notin S_0\}$ and likewise for $\bar{a} \in S_1$, define $I_{\text{out}}^{(1)}(\bar{a}) = \{i \in V'_{\text{out}} \mid \bar{a} + e_i \notin S_1\}$. By induction hypothesis, applied w.r.t V'_{out} , and pairs S_0 and $I_{\text{out}}^{(0)}(\bar{a})$, as well as S_1 and $I_{\text{out}}^{(1)}(\bar{a})$, we know that

$$\sum_{\bar{a} \in S_0} 2^{|I_{\text{out}}^{(0)}(\bar{a})| - |V'_{\text{out}}|} \geq 1 \quad \text{and} \quad \sum_{\bar{a} \in S_1} 2^{|I_{\text{out}}^{(1)}(\bar{a})| - |V'_{\text{out}}|} \geq 1. \quad (2)$$

Now if index $j \in I_{\text{out}}^{(0)}(\bar{a})$ for some $\bar{a} \in S_0 \subset S$, then $\bar{a} + e_j \notin S_0$ and as the first coordinate of $\bar{a} + e_j$ is also 0, we have $\bar{a} + e_j \notin S$, and thus $j \in I_{\text{out}}(\bar{a})$. Thus $|I_{\text{out}}(\bar{a})| \geq |I_{\text{out}}^{(0)}(\bar{a})|$ for all $\bar{a} \in S_0$. Likewise, $|I_{\text{out}}(\bar{a})| \geq |I_{\text{out}}^{(1)}(\bar{a})|$ for all $\bar{a} \in S_1$. Since $|V'_{\text{out}}| = |V_{\text{out}}| - 1$, using these in (2), we conclude (1) in this case, as desired.

Next, suppose $S = S_0$ and $S_1 = \emptyset$ (the case when $S_0 = \emptyset$ is handled the same way). In this case, for every $\bar{a} \in S$, $1 \in I_{\text{out}}(\bar{a})$, as $S_1 = \emptyset$ and thus flipping the first bit will always lead to a vector outside S . Thus $|I_{\text{out}}(\bar{a})| = |I_{\text{out}}^{(0)}(\bar{a})| + 1$. Using this together with $|V'_{\text{out}}| = |V_{\text{out}}| - 1$ in the first inequality of (2), we conclude (1) in this case as well. \square

Now, let $\bar{a} \in \mathbb{F}_2^n$ be some solution of the k - SUB - SAT instance (Φ, A) . Then, by the assumption of Case 1 and the preceding discussion \bar{a} has $|V_{\text{in}}| + |I_{\text{out}}(\bar{a})|$ critical variables w.r.t Φ .

Following the analysis in [22], if we now run one iteration of the PPZ algorithm on the instance Φ , the probability that \bar{a} is output is at least

$$\frac{1}{n^2} \cdot 2^{-n+(|V_{in}|+|I_{out}(\bar{a})|)/k}.$$

Let $S \subset \mathbb{F}_2^n$ denote the subset of solutions to the instance (Φ, A) . Summing up over all $\bar{a} \in S$, the probability that some solution \bar{a} is output is given by

$$\begin{aligned} \sum_{\bar{a} \in S} \frac{1}{n^2} \cdot 2^{-n+(|V_{in}|/k+|I_{out}(\bar{a})|/k)} &= \frac{1}{n^2} 2^{-n+n/k} \cdot \sum_{\bar{a} \in S} 2^{(-|V_{out}|/k+|I_{out}(\bar{a})|/k)} \\ &\geq \frac{1}{n^2} 2^{-n+n/k} \cdot \sum_{\bar{a} \in S} 2^{(-|V_{out}|+|I_{out}(\bar{a})|)} \geq \frac{1}{n^2} 2^{-n+n/k}, \end{aligned}$$

where the last step uses Lemma 6. This finishes the analysis of Case 1.

Remark 7 Notice in the probability analysis that S is the set of solutions to (Φ, A) and not all solutions to Φ . The crucial property that for every $\bar{a} \in S$, each variable in V_{in} is critical w.r.t Φ yields that there are $|V_{in}| + |I_{out}(\bar{a})|$ critical variables for \bar{a} w.r.t Φ . Intuitively, as the variables in V_{out} do not occur in the linear equations, the PPZ algorithm when run on Φ will be able to deterministically set, on average, $|I_{out}(\bar{a})|/k$ many of the critical variables in V_{out} without any interaction with the linear equations defining A .

Case 2 We now consider the case when not all variables in V_{in} are critical to all solutions to (Φ, A) . We will show that there is a subset of at most t variables in V_{in} that can be plucked from Φ and reduce the transformed instance to Case 1. We will argue that the algorithm can do an exhaustive search for this subset of V_{in} of size at most t .

Lemma 8 *In the $k - \text{SUB} - \text{SAT}$ instance (Φ, A) , let $Bx = b$ be the system of t linear equations defining A . Suppose variable x_1 occurs in the first equation $\sum_{j=1}^n B_{1j}x_j = b_1$ (i.e., $B_{11} \neq 0$). Further, suppose x_1 is not critical for some solution to (Φ, A) . Let Φ' be the formula obtained by plucking x_1 from Φ . Let A' be the affine space of co-dimension $t - 1$ defined by dropping the first linear equation $\sum_{j=1}^n B_{1j}x_j = b_1$ after eliminating x_1 from the other linear equations by row operations. Then (Φ', A') is satisfiable and any solution \bar{a}' to (Φ', A') can be extended to a solution \bar{a} of (Φ, A) .*

Proof By assumption, there is a solution \hat{a} to (Φ, A) for which x_1 is non-critical. Let $\hat{a}' \in \mathbb{F}_2^{n-1}$ be the assignment to x_2, x_3, \dots, x_n obtained from \hat{a} by dropping the x_1 -coordinate. Clearly, \hat{a}' is a solution to (Φ', A') . Hence, (Φ', A') is satisfiable. Furthermore, suppose \bar{a}' is some solution to (Φ', A') . Then the assignment \bar{a}' to the $n - 1$ variables x_2, x_3, \dots, x_n can be extended by choosing x_1 such that the constraint $\sum_{j=1}^n B_{1j}x_j = b_1$ is satisfied. The resulting assignment \bar{a} satisfies Φ and all t constraints defining A . □

Lemma 8 describes a pluck/eliminate step applied to the non-critical variable x_1 : namely, pluck x_1 from Φ and eliminate it from the equations describing A .

Clearly, for some sequence of $s \leq t$ pluck/eliminate steps applied successively transforms (Φ, A) to (Φ_s, A_s) for which Case 1 holds. Since we do not have an efficient test for checking non-criticality, the algorithm has to do an exhaustive search for the sequence of s variables to pluck/eliminate. The number of variable sequences to consider is bounded by n^t . However, as we argue in the next claim, it suffices to consider each unordered subset U of size $s \leq t$ variables and apply pluck/eliminate steps to its variables in the natural order x_1, \dots, x_n . Thus, we can bound the exhaustive search to $\binom{n}{\leq t}$ subsets of variables. Let (Φ_U, A_U) be the resulting instance after pluck/eliminate applied to variables in U in the natural order.

Lemma 9 *Let (Φ, A) be a satisfiable instance of k – SUB – SAT with $\text{codim}(A) = t$. There is a subset U of variables of size at most t , such that (Φ_U, A_U) is a satisfiable Case 1 instance of k – SUB – SAT.*

Proof Suppose $x_{i_1}, x_{i_2}, \dots, x_{i_s}$ is a sequence of $s \leq t$ variables to which the pluck/eliminate steps applied results in a satisfiable Case 1 instance (Φ_s, A_s) . Let the t equations $Bx = b$ define the affine space A . The row operations applied with the pluck/eliminate steps transforms this system into the following equations (also defining A):

$$\ell_j = x_{i_j}, 1 \leq j \leq s \text{ and } \ell_j = 0, s + 1 \leq j \leq t, \quad (3)$$

for affine linear forms $\ell_j, j \in [t]$ in which none of the variables $x_{i_1}, x_{i_2}, \dots, x_{i_s}$ occur. Moreover, the $t - s$ equations $\ell_j, j > s$ define A_s , and for every solution \bar{a} to (Φ_s, A_s) all variables occurring in these $t - s$ equations are critical for \bar{a} w.r.t Φ_s .

Now, suppose we apply the pluck/eliminate steps in the natural order to the variable subset $U = \{x_{i_1}, x_{i_2}, \dots, x_{i_s}\}$ resulting in (Φ_U, A_U) . Formulas Φ_U and Φ_s are identical (as both are obtained by plucking variables from U). The accompanying row operations for the eliminate steps could result in a different set of equations (defining A): $\ell'_j = x_{i_j}, 1 \leq j \leq s$ and $\ell'_j = 0, s + 1 \leq j \leq t$. The variables in U do not occur in $\ell'_j, j \in [t]$, and the affine space A_U is defined by the $t - s$ equations $\ell'_j = 0, j > s$. Since any solution to these equations uniquely determines the values to the variables in U , and all equations together define A , we can conclude that $A_U = A_s$. \square

The $O^*\left(\binom{n}{\leq t} \cdot 2^{n-n/k}\right)$ time Algorithm.

On input (Φ, A) , the algorithm proceeds as follows:

For each subset $U \subset V_{\text{in}}$ of size at most t do the following:

1. Pluck the variables in U from Φ to obtain Φ_U .
2. For each variable $x_i \in U$ (in any order): pick some equation in which x_i occurs; remove x_i from other equations by adding the picked equation to it; drop the picked equation from the system.
3. Run the PPZ algorithm on the resulting instance (Φ_U, A_U) as if Case 1 were applicable. More precisely, run PPZ on Φ_U for $O^*(2^{n-n/k})$ steps; for each solution

obtained, if it satisfies A_U then output an extension of it to a solution to (Φ, A) and exit,⁵ else continue the for-loop for the next choice of subset U .

To see the correctness, suppose (Φ, A) is satisfiable. By Lemma 9, for some choice of U with $|U| \leq t$, (Φ_U, A_U) is a Case 1 instance. Hence, the PPZ satisfiability algorithm will output a solution to (Φ_U, A_U) in time $O^*(2^{n-n/k})$ with high probability. This solution can be uniquely extended to a solution to (Φ, A) using the linear equations.

We have thus shown the following.

Theorem 10 *There is a randomized $O^*(\binom{n}{t} \cdot 2^{n-n/k})$ time algorithm for k -SUB-SAT for subspaces of co-dimension t . In particular, for $t = o(n)$ we have a randomized $O^*(2^{n(1-1/k+o(1))})$ time algorithm.*

2.2 An $O^*(2^{n-n/2k+n/2k^2})$ Time PPZ-Based Algorithm for k -Sub-Sat

Let (Φ, A) be a k -SUB-SAT instance. Our algorithmic strategy is essentially based on the PPZ algorithm for k -SAT, and our objective is a randomized algorithm with running time $2^{n-(1-\nu)n/k}$ for as small a ν as possible (ideally, $\nu = 0$ which would match the running time of the PPZ algorithm for k -SAT).

To this end, we can first apply Valiant-Vazirani Lemma [28] to increase the number of constraints (thereby reducing the rank of A) and getting an instance (Φ, A') such that Φ has a unique solution in A' with high probability (i.e., inverse polynomial probability as guaranteed by Valiant-Vazirani).

If $\dim(A') \leq n - (1 - \nu)n/k$ we can brute force search in A' in deterministic time $2^{\dim(A')} \leq 2^{n-(1-\nu)n/k}$. Thus, we can assume that $\dim(A') = n - t$ and A' is the solution space of $t < (1 - \nu)n/k$ independent affine linear equations.

Let now $\bar{a} \in \mathbb{F}_2^n$ be the unique solution to the k -SUB-SAT instance (Φ, A') . We partition the variable set into $V_{\text{in}} \sqcup V_{\text{out}}$ as before.

Claim 11 *Every variable in V_{out} is critical for the satisfying assignment \bar{a} of Φ .*

Proof of Claim Suppose $x_i \in V_{\text{out}}$ is not critical for \bar{a} . Then $\bar{a} + e_i$ is also a satisfying assignment for Φ . Moreover, since x_i does not occur in V_{in} , $\bar{a} + e_i$ satisfies the linear equations defining A' . Hence $\bar{a} + e_i$ is a solution to (Φ, A') contradicting the uniqueness of \bar{a} .

The variable plucking algorithm If \bar{a} has more than $(1 - \nu)n$ many critical variables (ν to be fixed in the analysis) then by running the PPZ satisfiability algorithm [22] for $O^*(2^{n-(1-\nu)n/k})$ iterations we will find it with high probability.

Otherwise, there are more than νn many variables in V_{in} that are *not critical* for Φ at \bar{a} .

1. Repeat the following two steps at most t times.
2. (The plucking step) Randomly pluck a variable x_i from V_{in} and drop it from the formula Φ to obtain its *shrinking* Φ_1 . Take a linear equation $\ell = b$ in which x_i

⁵ From a solution to (Φ_U, A_U) we can reconstruct the solution to (Φ, A) as the values to variables in U are uniquely determined via the linear equations from the values to the other variables.

occurs. By row operations eliminate x_i from all other linear equations in which x_i occurs and then drop the equation $\ell = b$. Let the affine space described by the new set of at most $t - 1$ linear equations be A_1 . We claim that (Φ_1, A_1) also has a unique solution \bar{a}_1 (obtained from \bar{a} by dropping the i th coordinate).

- Let $n_1 = n - 1$. Run the PPZ algorithm for $2^{n_1 - (1-\nu)n_1/k}$ time on Φ_1 . If we do not find the unique solution \bar{a}_1 then repeat the plucking step.

At the end of t successful plucking steps we are left with a $k - \text{SAT}$ instance Φ_t with a unique solution (the subspace A_t is \mathbb{F}_2^n) and PPZ will find that solution from which we can compute \bar{a} by recovering the unique values of the plucked variables using the linear equations.

Analysis At the j th iteration of the plucking step, the probability that all j steps pluck off non-critical variables is at least ν^j . Thus, the running time of the search for unique solutions for the (Φ_j, A_j) over all t steps is bounded by $\sum_{j=0}^t O^*(\frac{1}{\nu^j} \cdot 2^{n_j - (1-\nu)n_j/k})$.

Letting $\alpha = 2^{1-(1-\nu)/k}$ and noting that $n_j = n - j$ we can rewrite and bound the above sum as

$$\begin{aligned} O^*(2^{n-(1-\nu)n/k}) \cdot \sum_{j=0}^t \frac{1}{\nu^j \cdot \alpha^j} &\leq O^*(2^{n-(1-\nu)n/k}) \cdot t \cdot \frac{1}{\nu^t \cdot \alpha^t} \\ &\leq O^*(2^{n-(1-\nu)n/k}) \cdot t \cdot \left(\frac{1}{2\nu}\right)^{(1-\nu)n/k} \cdot 2^{(1-\nu)n/k^2}, \end{aligned}$$

as the sum $\sum_{j=0}^t \frac{1}{\nu^j \cdot \alpha^j}$ is bounded by $t \frac{1}{\nu^t \cdot \alpha^t}$ for $\nu\alpha < 1$ and $t \leq (1 - \nu)n/k$.

The overall running time of the algorithm is, therefore, $O^*(2^{n-n/k}) \cdot 2^{\nu n/k} \cdot \left(\frac{1}{2\nu}\right)^{(1-\nu)n/k} \cdot 2^{(1-\nu)n/k^2}$, which is minimized at $\nu = 1/2$ as we argue below, and is given by $O^*(2^{n-n/2k+n/2k^2})$.

Ignoring the last factor, we need to minimize $2^{\nu n/k} \cdot \left(\frac{1}{2\nu}\right)^{(1-\nu)n/k}$. In other words, we need to minimize

$$2^\nu \cdot \left(\frac{1}{2\nu}\right)^{1-\nu},$$

Or, equivalently, minimize

$$\nu \log(4\nu) - \log(2\nu) \text{ over } \nu \in [0, 0.5].$$

This is minimized at $\nu = 0.5$ and the minimum value is also 0.5.

Remark 12 (Extension beyond linear-algebraic constraints) We note some aspects about the algorithm and explain its adaptation to the more general setting of k -CNF satisfiability in the presence of a global boolean constraint $C(x_1, x_2, \dots, x_n)$ with the property that given a partial assignment to the variables x_i we can extend the assignment to the remaining variables that satisfies the constraint C , if such an extension exists. We set $\nu = 1/2$ and $t = n/2k$. Note that the algorithm need not partition the variables into V_{in} and V_{out} . If there are over $n/2$ non-critical variables, the algorithm

can "obliviously" pluck one with probability $1/2$. Oblivious in the sense that it does not need to see the constraint C . After $t = n/2k$ plucking steps, there are at most $n - n/2k$ remaining variables. We add a final step to the algorithm which is a brute-force search over all $2^{n-n/2k}$ assignments to the remaining variables. For each assignment to these that satisfies Φ_t we can check, in polynomial time, if there is an extension to it that satisfies C . This search will succeed for the unique solution \bar{a} . An interesting example for constraint C would be Horn formulas. As clause size is unrestricted in Horn formulas, notice that neither a direct application of the PPZ satisfiability algorithm, nor an application of the polynomial equations algorithms would give constant savings in the exponent for the running time bound.

More generally, call a Boolean constraint $C(x_1, x_2, \dots, x_n)$ $T(n)$ -easy if there is a $T(n)$ time-bounded algorithm that searches for a satisfying extension of a given partial assignment to the variables x_i .

Theorem 13 *There is a randomized $O^*(2^{n-n/2k+n/2k^2} \cdot T(n))$ time algorithm that takes any k -CNF formula and a $T(n)$ -easy boolean constraint $C(x_1, x_2, \dots, x_n)$ as input and computes a satisfying assignment for the formula and C .*

Corollary 14 *There is a randomized $O^*(2^{n-n/2k+n/2k^2})$ time algorithm for k – SUB – SAT.*

2.3 An $O^*(1.5^r)$ time algorithm for 2-Sub-Sat

Theorem 15 *Given a 2 – SUB – SAT instance (Φ, A) , where Φ is a 2-CNF formula and $A \subset \mathbb{F}_2^n$ is an r -dimensional affine subspace given by linear equations, there is a randomized $O^*(1.5^r)$ time algorithm to check if Φ has a satisfying assignment in A and if so to compute it.*

Proof Let $X = \{x_1, x_2, \dots, x_r, \dots, x_n\}$ be the variable set. Without loss of generality, we can assume that x_1, x_2, \dots, x_r are independent variables and for $j > r$ we have $x_j = \ell_j$, where ℓ_j is a linear form in x_1, x_2, \dots, x_r . The literal \bar{x}_j is the affine linear form $\ell_j + 1$.

Thus, we can treat the instance (Φ, A) as a conjunction Ψ of disjunctions $(\ell \vee \ell')$, where ℓ and ℓ' are affine linear forms in x_1, x_2, \dots, x_r . We can think of this satisfiability problem as picking one affine form from each such 2-disjunction $(\ell \vee \ell')$ and setting it to true such that the resulting equations are all consistent (i.e. the equations have a solution in \mathbb{F}_2^r).

We describe below a randomized algorithm that builds a system of independent linear equations over x_1, x_2, \dots, x_r such that any satisfying assignment \bar{a} is a solution to this system of linear equations with probability at least $(2/3)^r$, and, moreover, any solution to this system satisfies Ψ . Clearly repeating this algorithm $O^*(1.5^r)$ times will find a satisfying assignment to the 2 – SUB – SAT instance Ψ if one exists.

Here is a description of the algorithm to convert Ψ to a system of linear equations:

1. The algorithm runs in stages $i = 0, 1, \dots$ where in the i th Stage, it has a system of linear equations $\ell'_j = 1, 1 \leq j \leq i$ for a collection of linearly independent affine forms ℓ'_j . We start off with the empty system at stage 0.

2. (Stage $i + 1$): Take a clause $(\ell \vee \ell')$. If either $\ell = 1$ or $\ell' = 1$ is implied by the equations from stage i (which can be checked by solving linear equations) then we can discard that clause as satisfied and examine the next clause. If both $\ell = 0$ and $\ell' = 0$ are implied by the equations then this is a rejecting computation and algorithm outputs “fail”. If $\ell = 0$ is implied by the equations and ℓ' is independent of the ℓ'_j then we include the equation $\ell' = 1$ and go to Stage $i + 2$ (if there are any clauses left). Finally, if both ℓ and ℓ' are independent of the ℓ'_j then we randomly pick one of *three* linear forms ℓ , ℓ' and $\ell + \ell'$, include the equation setting it to 1 and go to Stage $i + 2$ (if there are any clauses left).
3. Let the final stage be r' . Note that $r' \leq r$ since the equations $\ell'_j = 1$ are all independent. At this stage we have no clauses left and any solution to the linear equations $\ell'_j = 1, 1 \leq j \leq r'$ satisfies Ψ . Output an arbitrary such solution.

We now analyze the success probability of the algorithm. Suppose $\bar{a} \in \mathbb{F}_2^r$ is a satisfying assignment for Ψ . We claim that the probability that \bar{a} satisfies the final system of equations $\ell'_j = 1, 1 \leq j \leq r'$ is at least $(2/3)^{r'}$. We will prove this by an induction on the stage number i : the induction hypothesis is that \bar{a} satisfies the set of equations at stage i with probability at least $(2/3)^i$. Clearly, it holds at $i = 0$.

For the induction step, suppose after Stage i , the assignment \bar{a} satisfies $\ell'_j = 1, 1 \leq j \leq i$. Then notice that in Stage $i + 1$ we either deterministically add the equation $\ell' = 1$ which \bar{a} must satisfy since it does not satisfy $\ell = 1$ (indeed ℓ must evaluate to 0 at \bar{a}), or we randomly pick one of ℓ, ℓ' and $\ell + \ell'$. Clearly, \bar{a} must satisfy exactly two of these three linear forms. Hence at the end of Stage $i + 1$ the assignment \bar{a} satisfies the system $\ell'_j = 1, 1 \leq j \leq i + 1$ with probability at least $(2/3)^{i+1}$. It follows that at the end of stage $r' \leq r$, \bar{a} satisfies the equations with probability at least $(2/3)^{r'}$. \square

Remark 16 The running time of $O^*(1.5^r)$ that we obtain improves on the polynomial equations based algorithms, where for $k = 2$ the best run time so far is $O^*(1.618^r)$ [8]. For $k = 3$ a similar randomized branching strategy gives an algorithm with running time $O^*((7/4)^r)$. For larger k the running time degrades to $O^*((2 - 1/2^{k-1})^r)$. This running time bound is obtained similarly as for Theorem 15: fix a satisfying assignment \bar{a} of the k -SUB-SAT instance. For a clause $(\ell_1 \vee \ell_2 \vee \dots \vee \ell_k)$ of k linearly independent linear forms a random (nonzero) linear combination $\sum_{i=1}^k \alpha_i \ell_i$ evaluates to 1 at \bar{a} with probability exactly $\frac{2^{k-1}}{2^k - 1}$.

2.4 2-Sub-Sat in a Co-dimension t Subspace

In this section we consider 2-SUB-SAT where we are seeking a solution in an affine space A such that $\text{codim}(A) = t$.

Given a formula Φ we will identify a canonical satisfying assignment \bar{a} for Φ based on which we will define critical variables. Since 2-SAT is in polynomial-time, we can detect non-critical variables in Φ w.r.t. \bar{a} in polynomial time. Now the plucking step will try all the possible $\binom{t}{1}$ choices of plucking non-critical variables, recalling that a non-critical variable plucked from a linear constraint defining A allows us to drop that constraint.

Theorem 17 *There is an $O^*\binom{n}{t}$ time deterministic algorithm for checking if a 2 – SUB – SAT instance (Φ, A) is satisfiable where the affine space A has co-dimension t .*

Proof Let Φ be a 2-CNF formula in variables $x_i, i \in [n]$.

We first do a standard preprocessing of Φ by considering its implication graph on the $2n$ literals $x_i, \bar{x}_i, i \in [n]$, where for each clause $u \vee u'$, for literals u and u' , we have two directed edges (\bar{u}, u') and (\bar{u}', u) . The literals that form strongly connected components must all take the same value in any satisfying assignment and, therefore, can be replaced by a single variable. This shrinks the implication graph to a DAG and also reduces the number of variables. Thus, without loss of generality, we can assume the implication graph of Φ is a DAG, and we refer to Φ as a *reduced 2 – SAT* formula. *Computing a canonical satisfying assignment* A standard linear-time 2 – SAT algorithm computes a canonical satisfying assignment \bar{a} for Φ (if satisfiable) by the following algorithm:

- (a) All literals of outdegree 0 in the implication DAG are assigned true.
- (b) The formula Φ is simplified after this substitution and the new implication DAG computed. If the DAG is non-empty we repeat Step(a).

The following claim uses the above algorithm to identify non-critical variables for some satisfying assignment for Φ .

Claim 18 *Let Φ be a 2 – SAT formula with implication DAG G . Let $u \in \{x_i, \bar{x}_i\}$ be an outdegree 0 literal in G . If Φ is not satisfiable with $u = 0$ then x_i is critical for every satisfying assignment of Φ , and if Φ is satisfiable with $u = 0$ then x_i is non-critical for every satisfying assignment for Φ that sets $u = 0$.*

Proof of Claim. If there is no satisfying assignment for Φ with $u = 0$ then clearly x_i is critical for every satisfying assignment. Conversely, suppose \bar{a} is a satisfying assignment with $u = 0$. Then we note that x_i is not critical for \bar{a} because $\bar{a} + e_i$ is also a satisfying assignment for Φ . More precisely, because u has outdegree 0 in the implication graph we can set $u = 1$, while retaining the other values in \bar{a} , and it remains a satisfying assignment.

More generally, given Φ we can partition the literals occurring in its implication DAG G as $S_0 \sqcup S_1 \sqcup \dots \sqcup S_w$, where S_0 is the set of outdegree 0 literals in G , S_1 is the set of outdegree 0 literals in DAG $G_1 = G \setminus S_0$, and in general S_i is the set of outdegree 0 literals in the DAG $G_{i+1} = G_i \setminus S_i$. For a variable x_i let $\text{depth}(x_i)$ be the least index j such that x_i or its complement is in S_j .

We observe the following claim which is an easy consequence of the previous one.

Claim 19 *Let Φ' be the 2 – SAT formula obtained by setting all literals in $S_0 \sqcup S_i \dots \sqcup S_{i-1}$ to true. For $u \in S_i$, if Φ' has no satisfying assignment with $u = 0$ then u is critical for every satisfying assignment for Φ that sets all literals in $S_0 \sqcup S_i \dots \sqcup S_{i-1}$ to true. If Φ' has a satisfying assignment with $u = 0$ then u is non-critical for every satisfying assignment of Φ that sets all literals in $S_0 \sqcup S_i \dots \sqcup S_{i-1}$ to true.*

We can immediately conclude the following.

Claim 20 *If there is a satisfying assignment for Φ in which all variables are critical that has to be the canonical satisfying assignment.*

We describe the basic search procedure used by the algorithm.

1. Let $\Phi_0 = \Phi$ and $A_0 = A$.
2. Repeat the following for steps $s = 0$ to $t - 1$.
3. Find the canonical satisfying assignment for Φ_s .
4. If it satisfies the linear equations $\ell_i = 0, i \in [t - s]$ defining A_s then output and stop (we can extend it uniquely to the s plucked non-critical variables using the linear equations).
5. Else a variable occurring in some ℓ_i is non-critical for Φ_s in the solution assignment.
6. Pick a non-critical variable x_j with minimum depth(x_j) and pluck it from Φ_s to get Φ_{s+1} . We take a linear equation $\ell_i = 0$ where x_j occurs in ℓ_i , eliminate x_j from all other equations by row operations using ℓ_i , and finally drop the constraint $\ell_i = 0$ to obtain a new affine space A_{s+1} . Continue with the repeat step.

Clearly, as long as the canonical satisfying assignment for Φ_s does not satisfy the system of equations $\ell_i = 0$ we can remove a non-critical variable occurring in one of the ℓ_i from Φ_s .

Correctness of the algorithm follows from noting that (Φ_s, A_s) is satisfiable if and only if (Φ_{s+1}, A_{s+1}) is satisfiable, and if t non-critical variables are plucked then the problem reduces to a 2 – SAT instances (without any linear constraints).

To complete the overall algorithm, in the basic iteration procedure we need to cycle through all possible choices of non-critical x_j at minimum depth depth(x_j). Since we are going to pluck at most t non-critical variables, this can be done by a brute-force search over all $\binom{n}{t}$ subsets of the variables. The running time bound also follows. \square

3 Hardness Results

In this section we prove our hardness results for subspace satisfiability. Since k – SAT itself is NP-hard for $k \geq 3$, so is k – SUB – SAT for $k \geq 3$. So we focus on the case $k = 2$.

3.1 NP-Hardness of 2-Sub-Sat

While 2 – SAT is polynomial time solvable, the following theorem shows that 2 – SUB – SAT is NP-hard. Note that this follows from Schaefer’s dichotomy theorem for Boolean CSP as the combination of 2 – SAT constraints and linear equations (even with 3 variables per equation) is not one of the six tractable cases, and thus NP-hard. Below we give a direct proof based on a simple reduction.

Theorem 21 *2 – SUB – SAT is NP-hard.*

Proof We show that we can express the NP-hard problem Graph 4-Colorability as an instance of 2 – SUB – SAT, or equivalently 2 – PAF – SAT. Indeed, given a graph $G = (V, E)$, the instance of 2 – PAF – SAT consists of two Boolean variables $x_{u,1}, x_{u,2}$

for each $u \in V$, which will encode the 2-bit representation of the 4 possible colors we can assign to u . For each edge $e = (u, v) \in E$, we include the polynomial equation

$$(x_{u,1} + x_{v,1} + 1) \cdot (x_{u,2} + x_{v,2} + 1) = 0. \tag{4}$$

Note that this equation is satisfied iff $x_{u,1} \neq x_{v,1}$ or $x_{u,2} \neq x_{v,2}$, i.e., when $(x_{u,1}, x_{u,2}) \neq (x_{v,1}, x_{v,2})$, which captures the fact the vertices u and v get different colors. The simultaneous satisfiability of the equations (4) for all $e \in E$ is thus equivalent to G being 4-colorable. \square

3.2 W[1]-Hardness of 2-Sub-Sat Parameterized by Co-dimension

We now strengthen the hardness result of Theorem 21 and show that $2 - \text{SUB} - \text{SAT}$ is unlikely to even be fixed-parameter tractable when parameterized by the co-dimension t of the subspace in which we seek a satisfying assignment to the 2CNF formula. On the other hand, recall that (as shown in [2] and also Sect. 2.4), for fixed co-dimension t , $2 - \text{SUB} - \text{SAT}$ can be solved in polynomial time. Our W[1]-hardness answers (in the negative) a question posed in [2] on whether $2 - \text{SAT}$ with a single modular constraint modulo M is fixed-parameter tractable when parameterized by M (they gave an algorithm with complexity $n^{O(M)}$).

Theorem 22 *Consider the $2 - \text{SUB} - \text{SAT}$ where the input subspace within which one has to satisfy the $2 - \text{SAT}$ formula has co-dimension t . Parameterized by t , $2 - \text{SUB} - \text{SAT}$ is W[1]-hard.*

Proof We give a reduction from the problem MULTICOLORED- CLIQUE. The input to MULTICOLORED- CLIQUE consists of a graph G , an integer t , and a partition (V_1, V_2, \dots, V_t) of the vertices of G , and the task is to decide if there is a t -clique in G containing exactly one vertex from each part V_i . The parameter associated with the problem is t . The problem MULTICOLORED- CLIQUE parameterized by t is known to be W[1]-hard [10, Lemma 1].

The variables in the $2 - \text{SUB} - \text{SAT}$ instance correspond to the vertices of the graph. Let us denote these variables by x_v for $v \in V := V_1 \cup V_2 \cup \dots \cup V_k$. The 2CNF clauses in the instance will be the following:

- For all $i \in \{1, 2, \dots, t\}$ and $v \neq v' \in V_i$, the clause $(\neg x_v \vee \neg x_{v'})$. These clauses ensure that at most one x_v can be set to 1 in each part.
- If (u, v) is not an edge in the graph with G , the clause $(\neg x_u \vee \neg x_v)$. These clauses ensure that the set $\{u \mid x_u = 1\}$ must induce a clique in G .

Note that this instance of $2 - \text{SAT}$ is trivial to satisfy by setting all variables to 0. The affine space A we will use to make this an instance of $2 - \text{SUB} - \text{SAT}$ is defined by the following equations:

$$\sum_{u \in V_i} x_u = 1 \quad \text{for } i = 1, 2, \dots, t. \tag{5}$$

We stress that the above equations are over \mathbb{F}_2 , and thus stipulate that there are an odd number of variables set to 1 in each part. But together with the 2CNF clauses which ensure that at most one variable in each part can be set to 1, it follows that satisfying assignments of this 2 – SUB – SAT instance are in one-one correspondence with t -cliques of G that include exactly one vertex from each V_i . The proof is now complete by noting that the co-dimension of the affine space A defined by (5) equals t . Parameterizing MULTICOLORED- CLIQUE by the clique size is thus equivalent to parameterizing the constructed 2 – SUB – SAT instance by the co-dimension. \square

3.3 Approximability of Max-2-Sub-Sat

Given the hardness of deciding exact satisfiability of 2 – SUB – SAT instance, we now turn to approximate satisfiability. In the MAX – 2 – SUB – SAT problem, the goal is to satisfy the maximum number of 2SAT clauses with an assignment that belongs to the input affine space A . Thus, the affine constraints are treated as hard constraints. We allow clauses of width 1. If unary clauses are disallowed in the 2CNF formula, and each clause involves exactly two distinct variables, we call the problem MAX – E2 – SUB – SAT.

3.3.1 Easy Approximation Algorithms

We can assume that no variable is forced to 0 or 1 by the affine space A , since if that happens we can just set and remove that variable and work on the reduced instance. If we pick a random assignment from A , it will satisfy at least $1/2$ of the clauses of the 2CNF formula in expectation, and in fact at least an expected fraction $3/4$ of the clauses when each clause involves two distinct variables. The algorithms are easily derandomized. For satisfiable instances of MAX-2-SUB-SAT, one can find a $3/4$ approximate solution, as one can eliminate all the unary clauses, and add those conditions to the subspace inside which we want to find an assignment to the 2CNF formula. So we get the following trivial algorithmic guarantees.

Observation 23 *In polynomial time, one can get a factor $1/2$ approximate solution to instances of MAX-2-SUB-SAT, a factor $3/4$ approximate solution to instances of MAX-E2-SUB-SAT, and a factor $3/4$ approximate solution to satisfiable instances of MAX-2-SUB-SAT.*

We will now show that all the above guarantees are best possible, with matching NP-hardness results.

3.3.2 Tight Inapproximability via Simple Reductions

For the hardness results and rest of the section, it is convenient to work with the PAF – SAT formulation of SUB – Sat. The Max-LIN2 problem, of maximizing the number of satisfied equations in a system of affine equations mod 2, trivially reduces to MAX-2-PAF-SAT (with each equation being degree 1 instead of degree 2). By Håstad's seminal tight inapproximability for Max-LIN2, we have the following.

Observation 24 For any $\epsilon > 0$, MAX-2-PAF-SAT (and thus MAX-2-SUB-SAT) is NP-hard to approximate within a factor of $(1/2 + \epsilon)$, and this holds for almost satisfiable instances that admit an assignment satisfying a fraction $(1 - \epsilon)$ of equations.

We also get a tight hardness (matching Observation 23) for the MAX-E2-SUB-SAT or equivalently when each polynomial equation is the product of exactly two (linearly independent) affine forms.

Lemma 25 For any $\epsilon > 0$, MAX-E2-PAF-SAT is NP-hard to approximate within a factor of $(3/4 + \epsilon)$, and this holds for almost satisfiable instances that admit an assignment satisfying a fraction $(1 - \epsilon)$ of equations.

Proof This follows from a simple reduction from Max-LIN2. Suppose we are given a system of affine equations $A_1 = 0, A_2 = 0, \dots, A_m = 0$, where the A_i 's are distinct affine forms in Boolean variables x_1, x_2, \dots, x_n . We produce a system of $\binom{m}{2}$ quadratic equations $A_i \cdot A_j = 0$ for $1 \leq i < j \leq m$ in the same variables x_1, x_2, \dots, x_n . If an assignment to the x_i 's violates r affine constraints $A_j = 0$, then the same assignment violates $\binom{r}{2}$ of the quadratic constraints. When $r = \epsilon m$, the fraction of violated quadratic constraints is $\approx \epsilon^2$, and when $r = 1/2 - \epsilon$, the fraction of violated quadratic constraints is $\approx 3/4 - O(\epsilon)$. The claimed hardness now follows from Håstad's inapproximability result for Max-LIN2. \square

3.3.3 Inapproximability for Satisfiable Instances

The above inapproximability results do not apply to *satisfiable* instances of 2-SUB-SAT. They are obtained by reductions from linear equations whose exact satisfiability can be easily checked. We now prove that approximating MAX-2-SUB-SAT doesn't get easier on satisfiable instances.

Theorem 26 For every $\epsilon > 0$, it is NP-hard to approximately solve satisfiable instance of MAX-E2-SUB-SAT within a factor of $3/4 + \epsilon$. That is, it is NP-hard to find, given as input a satisfiable instance of 2-SUB-SAT, an assignment satisfying a fraction $3/4 + \epsilon$ of the 2SAT constraints.

Proof Consider the arity 3 Boolean CSP which is defined by the predicate $\text{OXR} : \{0, 1\}^3 \rightarrow \{0, 1\}$, defined by

$$\text{OXR}(x_1, x_2, x_3) = x_1 \vee (x_2 \oplus x_3)$$

applied to literals. En route his celebrated tight inapproximability for satisfiable Max-3SAT, Håstad proved that the CSP defined by OXR (and with negations allowed on variables) is NP-hard to approximate within a factor of $(3/4 + \epsilon)$ even on satisfiable instances, for arbitrary $\epsilon > 0$. (Note that independent random choices of the bits x_1, x_2, x_3 makes $\text{OXR}(x_1, x_2, x_3) = 1$ with probability $3/4$, so the hardness factor of $3/4$ is tight.) Now the constraint $\text{OXR}(x_1, x_2, x_3) = 1$ is equivalent to the equation

$$(x_1 + 1)(x_2 + x_3 + 1) = 0$$

stipulating that a product of two affine forms vanishes. Thus the CSP defined by OXR can be equivalently expressed as a 2 – SUB – SAT instance, and the claimed inapproximability of MAX – E2 – SUB – SAT on satisfiable instances follows. \square

4 System of Polynomial Equations Over Binary Field: Effect of Reducibility

We now examine a special case of the problem of solving a system of polynomial equations over \mathbb{F}_2 studied in [1, 8, 19]. For motivating background, we recall according to the strong exponential time hypothesis (SETH) that SAT, that is n -variable CNF satisfiability of unrestricted clause width, cannot be essentially solved faster than 2^n time. However, Schuler [25] and Calabro et al [4] have shown the special case that sparse instances of SAT (with $c \cdot n$ clauses) can be solved in $O^*(2^{n(1-\alpha)})$ time, where α is a constant depending on the clause density c . It is natural to ask if there is an analogous result for SUB – Sat (satisfiability of conjunctions of unbounded disjunctions of affine linear forms). In this section we show a more general algorithmic result in the setting of systems of polynomial equations over \mathbb{F}_2 .

Let $P_i \in \mathbb{F}_2[x_1, x_2, \dots, x_n]$, $1 \leq i \leq m$ be polynomials over the field \mathbb{F}_2 as input instance to the POLY – EQS problem. The problem is denoted k – POLY – Eqs when the degrees are bounded by k which generalizes k – SUB – SAT as already explained in the introduction.

The unrestricted degree case is significantly different, because we can easily combine the m equations into a single equation as follows. Define

$$P = 1 + \prod_{i=1}^m (1 + P_i).$$

Clearly, the system $P_i = 0$, $1 \leq i \leq m$ has a solution iff $P = 0$ has a solution.

Thus, assuming SETH, there is no algorithm essentially faster than 2^n for solving $P = 0$.

Remark 27 There is also the question of how the polynomials P_i are given as part of the input. If $\deg P_i \leq k$ for all P_i then we can in polynomial-time compute their sparse representation as a linear combination of the n^k many monomials of degree at most k . However, in the above reduction of combining the P_i into a single polynomial, P is a small arithmetic formula. In fact, for the case of POLY – EQS we consider, where the instance is a system of equations $P_i = 0$, $1 \leq i \leq m$ such that $m = O(n)$ and each P_i has constant degree irreducible factors, we can assume that the P_i are given as arithmetic circuits.

We now show that POLY – EQS instances $P_i = 0$, $1 \leq i \leq m$ can be solved faster than 2^n if m is linear in n and the irreducible factors of each P_i are of constant degree. This can be seen as a “polynomial equations” analogue of Schuler’s SAT algorithm for sparse instances with unrestricted clause width [4, 25]. We note that a different degree reduction method, based on a rank argument, is used in [19, Section 4] to solve

systems of polynomial equations $p_i = 0$, where each p_i is given by a sum of product of affine linear forms.

Theorem 28 *Let $P_i = 0, 1 \leq i \leq c \cdot n$, for a constant $c > 0$, be an instance of POLY – EQS, such that the degree of each irreducible factor of each P_i is bounded by a constant b . There is a randomized algorithm for POLY – EQS that runs in time $2^{n(1-\alpha)}$ for such instances, where $\alpha > 0$ is a constant that depends on c and b .*

Proof We can factorize each polynomial P_i into its irreducible factors in randomized polynomial time using Kaltofen’s algorithm [17]. Let

$$P_i = \prod_{j=1}^{r_i} Q_{ij}$$

be this factorization for each i . Define polynomials $R_{ij} = 1 + Q_{ij}$ for each i and j , and note that $\deg R_{ij} \leq b$. For $a_{ijs} \in \mathbb{F}_2$ picked independently and uniformly at random define polynomials

$$\tilde{R}_{is} = \sum_{j=1}^{r_i} a_{ijs} R_{ij}, 1 \leq s \leq \log m + 2.$$

Finally, we define the polynomials

$$\tilde{R}_i = \prod_{s=1}^{(\beta+1) \log c} (1 + \tilde{R}_{is}), 1 \leq i \leq m,$$

where $\beta > 0$ is a constant to be fixed later in the analysis.

Notice that $\deg \tilde{R}_i \leq b \cdot (\beta + 1) \log c$ for each i .

Claim 29 *If $P_i = 0, 1 \leq i \leq m$ is unsatisfiable then $\tilde{R}_i = 0, 1 \leq i \leq m$ is also unsatisfiable.*

To see this, suppose $P_i(\bar{a}) = 1$ at assignment $\bar{a} \in \mathbb{F}_2^n$. Then $Q_{ij}(\bar{a}) = 1$ for each j which implies each $R_{ij}(\bar{a}) = 0$ for each j . It follows that $\tilde{R}_{is} = 0$ for all s and hence $\tilde{R}_i = 1$.

On the other hand, we have:

Claim 30 *If $\bar{a} \in \mathbb{F}_2^n$ is a solution to the system of equations $P_i = 0, 1 \leq i \leq m$ then with probability at least e^{-n/c^β} \bar{a} is a solution to the system of equations $\tilde{R}_i = 0, 1 \leq i \leq m$.*

The probability that \bar{a} is a solution to the single equation $\tilde{R}_i = 0$ is given by $1 - \frac{1}{c^{\beta+1}}$. Since the events are independent, the probability that \bar{a} is a solution to the system $\tilde{R}_i, 1 \leq i \leq m$ is given by

$$\left(1 - \frac{1}{c^{\beta+1}}\right)^m = \left(1 - \frac{1}{c^{\beta+1}}\right)^{cn}$$

$$\approx e^{-n/c^\beta}.$$

Now the system of equations \tilde{R}_i , $1 \leq i \leq m$ is an instance of $k - \text{POLY} - \text{Eqs}$, where $k = b(\beta + 1) \log c$ is a constant. Applying one of the algorithms [1, 8, 19] yields an $O^*(2^{n(1-1/2k)})$ algorithm with success probability e^{-n/c^β} . We can boost the success probability to a constant with an overall run time of $O^*(2^{n(1-1/2k)} \cdot e^{n/c^\beta})$, which can be optimized by choosing β appropriately. \square

Acknowledgements We thank anonymous reviewers for their useful comments and pointers to the literature.

References

1. Björklund, A., Kaski, P., Williams, R.: Solving systems of polynomial equations over $\text{GF}(2)$ by a parity-counting self-reduction. In: Baier, C., Chatzigiannakis, I., Flocchini, P., Leonardi, S. (eds.) 46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9–12, 2019, Patras, Greece. LIPIcs, vol. 132, pp. 26:1–26:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2019)
2. Brakensiek, J., Gopi, S., Guruswami, V.: CSPs with global modular constraints: algorithms and hardness via polynomial representations. In: Proceedings of the 51st Annual ACM Symposium on Theory of Computing (STOC), pp. 590–601 (2019)
3. Brakensiek, J., Guruswami, V.: Bridging between 0/1 and linear programming via random walks. In: Proceedings of the 51st Annual ACM Symposium on Theory of Computing, pp. 568–577 (2019)
4. Calabro, C., Impagliazzo, R., Paturi, R.: A duality between clause width and clause density for SAT. In: 21st Annual IEEE Conference on Computational Complexity (CCC 2006), 16–20 July 2006, Prague, Czech Republic, pp. 252–260. IEEE Computer Society (2006)
5. Chan, T.M., Williams, R.R.: Deterministic APSP, orthogonal vectors, and more: quickly derandomizing Razborov–Smolensky. *ACM Trans. Algorithms* **17**(1), 2:1–2:14 (2021)
6. Chen, R., Santhanam, R.: Satisfiability on mixed instances. In: Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, pp. 393–402 (2016)
7. Couceiro, M., Haddad, L., Lagerkvist, V.: Fine-grained complexity of constraint satisfaction problems through partial polymorphisms: a survey. In: 2019 IEEE 49th International Symposium on Multiple-Valued Logic (ISMVL), pp. 170–175 (2019)
8. Dinur, I.: Improved algorithms for solving polynomial systems over $\text{GF}(2)$ by multiple parity-counting. In: Marx, D. (ed.) Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10–13, 2021, pp. 2550–2564. SIAM (2021)
9. Downey, R.G., Fellows, M.R.: Parameterized Complexity. Springer, Berlin (2012)
10. Fellows, M.R., Hermelin, D., Rosamond, F.A., Vialette, S.: On the parameterized complexity of multiple-interval graph problems. *Theor. Comput. Sci.* **410**(1), 53–61 (2009)
11. Flum, J., Grohe, M.: Parameterized Complexity Theory (Texts in Theoretical Computer Science. An EATCS Series). Springer, Berlin (2006)
12. Impagliazzo, R., Matthews, W., Paturi, R.: A satisfiability algorithm for AC^0 . In: Rabani, Y. (ed.) Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17–19, 2012, pp. 961–972. SIAM (2012)
13. Impagliazzo, R., Paturi, R.: On the complexity of k -sat. *J. Comput. Syst. Sci.* **62**(2), 367–375 (2001)
14. Impagliazzo, R., Paturi, R., Zane, F.: Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.* **63**(4), 512–530 (2001)
15. Jonsson, P., Lagerkvist, V., Nordh, G., Zanuttini, B.: Strong partial clones and the time complexity of SAT problems. *J. Comput. Syst. Sci.* **84**, 52–78 (2017)
16. Jonsson, P., Lagerkvist, V., Roy, B.: Fine-grained time complexity of constraint satisfaction problems. *ACM Trans. Comput. Theory* **13**(1), 2:1–2:32 (2021)
17. Kaltofen, E.: Factorization of polynomials given by straight-line programs. *Adv. Comput. Res.* **5**, 375–412 (1989)

18. Lagerkvist, V., Wahlström, M.: Which NP-hard SAT and CSP problems admit exponentially improved algorithms? CoRR, abs/1801.09488 (2018). [arXiv:1801.09488](https://arxiv.org/abs/1801.09488)
19. Lokshtanov, D., Paturi, R., Tamaki, S., Williams, R.R., Yu, H.: Beating brute force for systems of polynomial equations over finite fields. In: Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 2190–2202 (2017)
20. Monien, B., Speckenmeyer, E.: Solving satisfiability in less than 2^n steps. Discrete Appl. Math. **10**(3), 287–295 (1985)
21. Patrascu, M., Williams, R.: On the possibility of faster SAT algorithms. In: Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1065–1075 (2010)
22. Paturi, R., Pudlák, P., Zane, F.: Satisfiability coding lemma. In: 38th Annual Symposium on Foundations of Computer Science (FOCS), pp. 566–574 (1997)
23. Schaefer, T.J.: The complexity of satisfiability problems. In: Lipton, R.J., Burkhard, W.A., Savitch, W.J., Friedman, E.P., Aho, A.V. (eds.) Proceedings of the 10th Annual ACM Symposium on Theory of Computing, May 1–3, 1978, San Diego, California, USA, pp. 216–226. ACM (1978)
24. Schönning, U.: A probabilistic algorithm for k-sat and constraint satisfaction problems. In: 40th Annual Symposium on Foundations of Computer Science, pp. 410–414 (1999)
25. Schuler, R.: An algorithm for the satisfiability problem of formulas in conjunctive normal form. J. Algorithms **54**(1), 40–44 (2005)
26. Soos, M., Gocht, S., Meel, K.S.: Tinted, detached, and lazy CNF-XOR solving and its applications to counting and sampling. In: Lahiri, S.K., Wang, C. (eds.) Computer Aided Verification—32nd International Conference, CAV 2020, Los Angeles, CA, USA, July 21–24, 2020, Proceedings, Part I. Lecture Notes in Computer Science, vol. 12224, pp. 463–484. Springer (2020)
27. Soos, M., Meel, K.S.: BIRD: engineering an efficient CNF-XOR SAT solver and its applications to approximate model counting. In: The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27–February 1, 2019, pp. 1592–1599. AAAI Press (2019)
28. Valiant, L.G., Vazirani, V.V.: NP is as easy as detecting unique solutions. Theor. Comput. Sci. **47**(3), 85–93 (1986)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.