

# Satisfiability Checking + Symbolic Computation: A New Approach to Combinatorial Mathematics

CURTIS BRIGHT, University of Windsor, Canada

ILIAS KOTSIREAS, Wilfrid Laurier University, Canada

VIJAY GANESH, University of Waterloo, Canada

Advances in the fields of satisfiability solving (SAT) and computer algebra have recently led to the development of search tools that can solve mathematical problems significantly larger than ever before—and in a faster, more verifiable way. In this article we provide an overview of this “satisfiability solving & computer algebra” paradigm by describing some of its recent successes. In particular, we describe how it produces significantly faster and more verifiable solutions of Lam’s problem from geometry (uncovering bugs in previous solutions in the process) and has discovered over 100,000 new Williamson matrices from combinatorics. The paradigm is well-suited for solving problems that require both efficient search and sophisticated mathematics.

## ACM Reference Format:

Curtis Bright, Ilias Kotsireas, and Vijay Ganesh. 2020. Satisfiability Checking + Symbolic Computation: A New Approach to Combinatorial Mathematics. 1, 1 (September 2020), 11 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## INTRODUCTION

Mathematicians have long been fascinated by objects that exhibit exceptionally nice combinatorial properties. However, it is often difficult to determine whether or not objects satisfying a given combinatorial property exist. Sometimes, the only feasible method of definitively answering the question of existence is simply to perform a mathematical search. A famous example of this is the proof of the *four colour theorem*—that four colours suffice to colour the regions of a planar map so that adjacent regions are coloured differently [3]. The theorem has been known to be true since 1977, but every known proof relies on computer calculations in an essential way. Mathematical arguments are used to reduce the search for counterexamples to a finite number of cases, and the cases are then exhaustively checked using a custom-written computer program in order to rule out any counterexamples.

Independently, over the last fifty years computer scientists have made much progress on developing general-purpose programs that can automatically solve many kinds of mathematical problems. *Satisfiability solving* and *symbolic computation* are two important branches of computer science that each specialize in solving mathematical problems. Both of these fields have long histories and have produced impressive tools—satisfiability (SAT) solvers in the former case and computer algebra systems (CASs) in the latter case. Originally, SAT solvers were designed to solve problems in logic, while CASs were tools to manipulate and simplify algebraic expressions. As we will see, these tools have since found an abundance of new applications outside these original domains.

---

Authors’ addresses: Curtis Bright, University of Windsor, Canada; Ilias Kotsireas, Wilfrid Laurier University, Canada; Vijay Ganesh, University of Waterloo, Canada.

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2020 Association for Computing Machinery.

Manuscript submitted to ACM

Manuscript submitted to ACM

Despite their common specialization in solving mathematical problems, the SAT and CAS communities have developed independently of each other [1]. Broadly speaking, the SAT community has focused on effective search methods, while the CAS community has focused on effective mathematical algorithms. Recently, these two communities have started to collaborate in crossover initiatives like the SC-square project<sup>1</sup> [2]. Since the insights of these communities are largely complementary, bringing them together has resulted in new solutions to problems that were out-of-reach of either community separately and has produced a wealth of new applications [14]. In this overview, we focus on our own contribution to this ongoing project—a hybrid SAT solver and CAS system called MathCheck<sup>2</sup> that we have applied to mathematical problems in graph theory [37], finite geometry [4], combinatorics [8], and number theory [10].

### Satisfiability Solving

A satisfiability (SAT) solver is a program that solves the satisfiability problem from Boolean logic—given a formula in conjunctive normal form, is there an assignment to its variables that makes the expression true? At first glance, SAT solvers seem disconnected from the kinds of problems that most mathematicians and engineers care about. However, stunning progress in applied SAT solving over the last several decades [35] has lead to a surprising diversity of applications for SAT solvers—from generating or solving puzzles like Sudoku [7] to software verification [25] and hardware design [31]. This “SAT revolution” has even lead to resolution of decades-old mathematical problems such as the Boolean Pythagorean triples problem [22] and the determination of the fifth Schur number [21]. These problems are solved via a reduction to a set of constraints in Boolean logic, which are then given as input to SAT solvers that are used to search for solutions to these constraints or to provide nonexistence proofs when no solutions exist.

The search spaces in the aforementioned problems are enormous and could never be exhaustively searched without very clever and powerful search methods. The fact that SAT solvers are currently the only tools that can solve these problems speak to their exceptional search ability. However, SAT solvers do struggle with some kinds of problems—particularly those with an underlying mathematical structure that is unknown to the SAT solver.

As a simple example of this phenomenon, consider the following problem: Find a way to put  $n$  pigeons into  $n - 1$  holes given that each hole is only large enough to contain a single pigeon. A moment’s thought reveals that the problem is not solvable due to a simple counting argument. This problem is also straightforwardly expressed in Boolean logic—but embarrassingly, SAT solvers are known to take an infeasible amount of time to solve it using the most straightforward encoding [19]. The issue is that during the reduction to Boolean logic the mathematical context of the problem is lost; SAT solvers simply don’t realize that a counting argument suffices to rule out the existence of a solution. In this case, there are additional mathematical facts that can be encoded into Boolean logic and that allow SAT solvers to effectively show the problem is unsatisfiable even for large  $n$ . However, in many problems—such as the ones we consider in this article—the mathematical facts that greatly reduce the search space are not easily expressible in Boolean logic. It would seem as if we either have to ignore the mathematical facts or avoid using a SAT solver entirely.

### Symbolic Computation

A computer algebra system (CAS) is a program that can manipulate and simplify mathematical expressions and objects. They have a long history and have been developed for wide-ranging branches of mathematics such as graph theory, optimization, and group theory. A stunning amount of research—including the 1999 Nobel prize in physics<sup>3</sup>—relies on

<sup>1</sup>[www.sc-square.org](http://www.sc-square.org)

<sup>2</sup>[www.uwaterloo.ca/mathcheck](http://www.uwaterloo.ca/mathcheck)

<sup>3</sup>[www.nobelprize.org/prizes/physics/1999/press-release/](http://www.nobelprize.org/prizes/physics/1999/press-release/)

the precise calculations made possible by computer algebra systems. For example, they were essential in the resolution of Kepler’s 1611 conjecture that the most efficient way of packing spheres is in a pyramid shape [20]. This proof made use of numerous CAS functionalities including linear programming, global nonlinear optimization, and interval arithmetic.

Although CASs can solve an impressive number of mathematical problems there are certain kinds of problems for which they have not been tuned. For example, CASs have not typically been optimized to perform the kind of general-purpose search with learning as done by modern SAT solvers [1]. Note that the search spaces for mathematical objects are typically of an exponential size and will quickly overwhelm any system that has not been designed to cope with this immenseness. Simply put, finding a needle in a exponentially-large haystack requires tools specifically designed to find clever ways of pruning the search space down to a manageable size and this has not traditionally been in the domain of symbolic computation research.

### The Best of Both Worlds

Although the SAT and CAS worlds have both been applied to solve very hard mathematical problems, they have traditionally been applied in isolation—with SAT solvers applied to Boolean constraint problems involving huge search spaces and CASs applied to complicated mathematical problems involving minimal search. While each tool is very effective at solving problems in their respective worlds, a driving motivation of our work is that *there are many problems whose solutions demand effectiveness in both worlds*.

In 2015, Ábrahám proposed to join the techniques of the SAT and CAS worlds together in order to solve problems that are intractable using the techniques of either world in isolation [1]. Independently, the MathCheck project was started and demonstrated the effectiveness of this combination by improving the best known bounds in some graph theoretic conjectures [38]. However, MathCheck can be applied much more generally and over the past five years MathCheck has solved problems from a variety of branches of mathematics. In the remainder of this article we describe how MathCheck has been applied to long-standing problems in finite geometry and combinatorial matrix theory.

### LAM’S PROBLEM

The roots of Lam’s Problem date back to 300 BC, when Euclid defined five postulates that he believed characterized geometry. While all five postulates were considered obvious, the fifth or “parallel” postulate was significantly more complicated than his other postulates. For over 2000 years this bothered some mathematicians and many attempts were made to prove the parallel postulate from Euclid’s other postulates.

Surprisingly, it was not until the 1800s when it was finally realized that this task was impossible—because there are alternate geometries that satisfy Euclid’s first four postulates but in which the parallel postulate fails. For example, this happens in a *projective plane* which is a collection of points and lines satisfying the following two axioms:

- (1) There is exactly one line through any two points.
- (2) Any two lines intersect at a unique point.

These axioms are visually demonstrated in Figure 1. The first axiom holds in the usual Euclidean plane and is visualized in Figure 1a. The second axiom does not hold in the usual Euclidean plane—if two lines are skew (shown in Figure 1b) they *will* intersect in a unique point, but if the lines are parallel they will not meet at all, let alone in a unique point. To remedy this situation, the usual Euclidean plane may be augmented with a “line at infinity” on which any two formerly parallel lines will intersect (shown in Figure 1c). The resulting structure satisfies both axioms (1) and (2).




Fig. 1. Visual demonstrations of the axioms of a projective plane.

An interesting question can now be raised: are there any other structures that also satisfy these axioms? For example, are there examples of projective planes with a finite number of points? There are some trivial ways of satisfying the axioms—for example, if the plane consists of a single line containing every point. Disqualifying these trivial examples, counting arguments can be used to show that if a finite projective plane exists then it contains the same number of lines and points and every point lies on the same number of lines. If each point lies on  $n + 1$  lines then the plane is said to be of *order*  $n$  and it will contain exactly  $n^2 + n + 1$  points. Determining the set of possible orders for projective planes has been of significant mathematical interest for over 200 years and today it remains an major open problem.

Projective planes in the orders up to three may be explicitly visualized through the structures shown in Figure 2. By inspection, one can verify the axioms that any two points lie on a unique line and any two lines intersect in a unique point. Some might protest that the “lines” in these planes are not necessarily represented by straight lines—but so long as the lines satisfy the axioms we can draw them however we like.




Fig. 2. Visual depictions of projective planes in orders 1–3. Each point and line are drawn in a separate colour.

Figure 3 demonstrates the small orders for which projective planes are known to exist. Immediately what jumps out is that projective planes *do* generally exist in small orders—except in order six, when a projective plane would contain 43 points. What is special about order six? In 1949, Bruck and Ryser proved that if the order of a projective plane is  $n \equiv 1, 2 \pmod{4}$  then  $n$  must be the sum of two integer squares [11]. Since six satisfies the congruence condition but is not the sum of two squares it follows that no projective plane of order six can exist. However, the Bruck–Ryser theorem cannot be used to rule out

1	2	3	4	5	6	7	8	9	10
✓	✓	✓	✓	✓	✗	✓	✓	✓	?

Fig. 3. The orders for which projective planes are explicitly known to exist or known to not exist through a theoretical obstruction.

order ten, as ten is the sum of  $3^2$  and  $1^2$ . This case has attracted a huge amount of interest since no projective plane of order ten is known—yet no one knows any theoretical reason why one shouldn't exist. *Lam's problem* is to resolve this dilemma and determine how to correctly complete the missing entry in Figure 3.

### A Computational Solution

Lam's problem was resolved in the 1970s and 1980s by using a sophisticated case breakdown and an enormous amount of computing to separately search each case. No cases were found to lead to a solution, thereby disproving the existence of a projective plane of order ten.

The case breakdown was based on properties that a binary error-correcting code associated with a hypothetical projective plane of order ten must satisfy. This code is derived from the projective plane's *incidence matrix*—the  $\{0, 1\}$ -matrix that contains a 1 in its  $(i, j)$ th entry exactly when the  $i$ th line is incident with the  $j$ th point (see Figure 4). A *codeword* is a  $\{0, 1\}$ -vector in the row space (mod 2) of this incidence matrix and the *weight* of a codeword is the number of nonzero entries it contains. Analyzing the properties of these codewords through some powerful theorems of coding theory shows that a projective plane of order ten must produce at least one codeword of weight 15, 16, or 19—thus splitting the search into three cases.

The weight 15 case was resolved in 1973 using several hours on a mainframe computer [30], the weight 16 case was resolved in 1986 using 2,000 hours on a supermini computer [28] and 100 hours on a supercomputer [12], and the weight 19 case was resolved in 1989 using about 20,000 hours on a supermini computer and 2,000 hours on a supercomputer [29]. These searches were performed using code specifically written for each case and the authors acknowledged that mistakes were a real possibility. The source codes are not publicly available, and even if they were, they were designed to be run on specialized computers that are no longer produced. In order to believe that Lam's problem has in fact been resolved, we now have to resolve a different and serious problem, namely, how can these searches be verified?

### A New Hope

Relatively recent developments in the fields of satisfiability checking and symbolic computation offer the promise that not only can Lam's Problem be resolved more efficiently, but also be resolved to a higher standard of rigor.

This is because SAT solvers—though complicated pieces

of software—are well-tested and produce certificates or proofs that allow their results to be validated by external certificate verifiers. This means that it is no longer necessary to trust the implementation of a complicated search algorithm. Instead, one need only trust the certificate verifier.



Fig. 5. The red area can be ignored due to vertical symmetry.

1	1	0	1	0	0	0
0	1	1	0	1	0	0
0	0	1	1	0	1	0
0	0	0	1	1	0	1
1	0	0	0	1	1	0
0	1	0	0	0	1	1
1	0	1	0	0	0	1

Fig. 4. The incidence matrix of a projective plane of order 2.

*“I want to emphasize that this is only an experimental result and it desperately needs an independent verification.”*

C. W. H. Lam [27]

Naturally, the SAT solver will be used as the “combinatorial workhorse” that searches for a projective plane in each possible case. But of what use is the CAS in the search? To answer this, suppose you have a haystack for which you know the left and right sides are perfectly symmetrical, as in Figure 5. If you want to find a needle in this haystack a brilliant insight is to split the haystack down the middle and only search through one side—since anything in the left side appears on the right side and vice versa.

The same reasoning also applies more generally. Whenever a symmetry is detected in the search space, the space should be split up and reduced to a single nonsymmetric component. Fortunately, computing symmetries is one of the things that CASs excel at—including the complicated non-obvious symmetries that arise in the searches of Lam’s problem. Thus, a SAT+CAS system is perfectly suited to take advantage of the strengths of both SAT solvers and CASs.

### A SAT+CAS Solution

How should the SAT solver and CAS communicate? The basic connection is outlined in Figure 6. As the SAT solver is searching it finds “partial” projective planes—structures that satisfy the projective plane constraints up to some point. These are provided to the CAS and the CAS will send back additional “symmetry blocking” constraints that remove other partial projective planes (symmetric to the one that was found) from the search space.

This kind of connection takes inspiration from the Davis–Putnam–Logemann–Loveland DPLL( $T$ ) algorithm [32], with the first-order logic  $T$ -solver replaced by a CAS. We stress that this connection is very general and can be applied in many other circumstances than symmetry breaking (as we will see in the next section).

Indeed, one of the strengths of CASs is their huge amount of mathematical functionality—something that can be exploited to learn a variety of mathematical facts that the SAT solver would otherwise have no way of knowing.

We are now able to compare the SAT+CAS approach with previous approaches. The weight 15 search that was first solved in 1973 [30] has since been confirmed by at least four independent implementations, some requiring up to 80 minutes on a modern desktop [5]. The fastest previous implementation that we are aware of uses highly optimized C code and requires about 30 seconds to complete [13]. Using a straightforward SAT reduction this case can be solved in about 6 minutes with a modern SAT solver—already faster than some searches *specifically written* to perform this search. This speaks to the efficiency of modern SAT solvers, as this is not taking advantage of the symmetry that a mathematical analysis of the search space reveals. Using a SAT+CAS solver that does this (as described in Figure 6) the search completes in about 7 seconds, the current fastest known approach [5].

The weight 16 search is significantly more involved and has a much larger search space compared with the weight 15 search. Following the exhaustive searches of the space completing in 1986 [12, 28] we are aware of only a single previous independent confirmation: an optimized C implementation that required 16,000 hours on a cluster of desktops in 2011 [34]. In comparison, the SAT+CAS approach resolves this case in 30 hours [4]. Certain partial projective planes in this case have over a thousand symmetries—leading to dramatic reductions in the search space when these symmetries can be detected and removed.

The weight 19 search is even more involved than the weight 16 search, and following the exhaustive search in 1989 [29] the only independent confirmation we are aware of required 19,000 hours in 2011 [34]. Unfortunately, in this case the main search is usually not very symmetric, meaning we cannot expect to achieve a big speedup through symmetry breaking. Even still, the SAT+CAS method completed the search 25% faster than the special-purpose search code used in the 2011 confirmation of Lam’s problem—and without using a special-purpose search algorithm.




Fig. 6. An outline of the SAT+CAS method as applied to searching for projective planes. The SAT solver provides a partial projective plane to the CAS which computes the symmetries present. This example shows blocking the symmetry of rotating the partial plane by 180 degrees.



A final advantage of the SAT+CAS method is that the certificates produced by the solver may be verified by an independent party. They only need to trust the SAT encoding and the CAS-generated constraints—not the actual procedure used to generate the certificates. This does not prove our searches are error-free but does significantly reduce the amount of trust necessary. This is a particularly important consideration for “experimental” results lacking formal proof because we can be almost certain that custom-written special-purpose programs contain bugs. This is not a sleight on the authors but a simple reality of software development: without extensive formal verification even the most well-used and well-tested software cannot be made bug-free.

Indeed, our searches uncovered errors in several previous searches, including both the original 1989 search and its 2011 confirmation. Errors can be detected when the intermediate results of these searches (e.g., the claimed nonexistence of a partial solution) contradict the partial solutions that are found by the SAT solver. For example, the 2011 confirmation of the weight 15 case was based upon proving the nonexistence of a certain 51-column partial projective plane—which in fact actually exists and is explicitly shown in Figure 7. Note that it is easy to check that such a matrix is in fact a partial projective plane, as every column contains the same number of 1s and any two distinct columns share exactly a single 1 in the same location.

In summary, the SAT+CAS method nicely applies to Lam’s problem because it allows searching for projective planes using the powerful search routines of a modern SAT solver while simultaneously allowing the mathematical properties revealed by a computer algebra system to greatly constrain the search space. Furthermore, it removes the need to write a custom special-purpose search algorithm and therefore ultimately makes executing the search a more straightforward and trustworthy process.

## WILLIAMSON CONJECTURE

Next, we discuss the impact the SAT+CAS method has had on combinatorial matrix theory. We begin with a motivating example: say you need to communicate with someone over a noisy channel, i.e., the data that you send may not be the data that is received. How can you increase the likelihood the recipient can perfectly recover your original message? One way to do this is to first *encode* your message using a set of codewords that you and your recipient agree on in advance. Because the codewords may be corrupted while being sent it makes sense to choose the codewords to be as different as possible. For example, if you have  $n$  binary codewords of length  $n$  you could try to maximize the minimum pairwise Hamming distance of the codewords. Plotkin’s bound from coding theory [33] says that this minimum distance cannot be strictly larger than  $n/2$ . In certain cases it is possible to reach this bound and find a set of  $n$  binary codewords of length  $n$  such that the Hamming distance of any two distinct codewords is exactly  $n/2$ .

In this context it is convenient to represent a binary codeword as a  $\{\pm 1\}$ -vector. Then two codewords of length  $n$  have a Hamming distance of  $n/2$  exactly when they are orthogonal vectors and a set of  $n$  such codewords with a pairwise Hamming distance of  $n/2$  form the rows of a matrix  $A$  such that the off-diagonal entries of  $AA^T$  are zero. Such a matrix is called a *Hadamard matrix* after the French mathematician Jacques Hadamard who studied them in 1893 [18].




Fig. 7. A 51-column partial projective plane previously claimed to not exist. Red entries denote 1s known in advance and blue entries denote 1s determined by MathCheck.

Even though they have been extensively studied for over 125 years there remain many open problems concerning them—we still don’t even know all the orders in which they exist. Hadamard was able to prove that if a Hadamard matrix has order larger than two then its order must necessarily be a multiple of four. The *Hadamard conjecture* is that this necessary condition is also sufficient.

$A$	$B$	$C$	$D$
$-B$	$A$	$-D$	$C$
$-C$	$D$	$A$	$-B$
$-D$	$-C$	$B$	$A$

Fig. 8. The general form of Hadamard matrices constructed via the Williamson construction.

Since a proof of the Hadamard conjecture has not been forthcoming, mathematicians and computer scientists have done the next best thing and constructed Hadamard matrices in as many orders as possible; currently  $n = 167$  is the smallest case for which a Hadamard matrix of order  $4n$  is not known. Because the search space for general Hadamard matrices is so enormous it makes sense to focus on special constructions that have more structure to exploit—one such construction was presented by the mathematician John Williamson in 1944 [36]. In this construction we have four symmetric  $\{\pm 1\}$ -matrices  $A, B, C, D$  of order  $n$  which for simplicity we assume are circulant (each row is a cyclic shift of the previous row). If  $A^2 + B^2 + C^2 + D^2$  is the identity matrix scaled by a factor of  $4n$  then a Hadamard matrix of order  $4n$  exists and is explicitly presented in Figure 8. In this case the matrices  $A, B, C, D$  are known as a *set of Williamson matrices*.

In the 1960s, while developing codes for communicating with spacecraft, scientists at NASA’s Jet Propulsion Laboratory discovered a set of Williamson matrices of order twenty-three and conjectured that Williamson matrices exist for all values of  $n$  [17]. This conjecture was disproven in 1993 when the counterexample  $n = 35$  was discovered by exhaustive computer search [15]. It was noted that this was the smallest *odd* counterexample but not much was known about the even orders—though Williamson himself found some examples in even orders including all powers of two up to 32.

MathCheck was able to greatly improve our knowledge of Williamson matrices in even orders by exhaustively searching for Williamson matrices in all even orders  $n \leq 70$  (see Figure 9 for one example). Two immediate surprises were uncovered: first, Williamson matrices exist in all such orders, and second, Williamson matrices are much more plentiful when their order is divisible by a large power of two. For example, there are over 70,000 sets of Williamson matrices in order 64 alone and fewer than 100 sets of Williamson matrices exist in all the odd orders up to 64. Moreover, analyzing the structure of the Williamson matrices in order 64 revealed a structure that can be generalized to all powers of two. Thus, as a result of the searches of MathCheck we have improved Williamson’s result that Williamson matrices of order  $2^k$  exist with  $k \leq 5$ ; we now know that Williamson matrices of this form actually exist for all  $k$  [9].

### A SAT+CAS Solution

SAT solvers can be used to search for Williamson matrices because the arithmetical defining relationship of Williamson matrices may be directly encoded using straightforward arithmetic circuits. However, this encoding is not at all competitive with special-purpose solvers and tops out around order thirty [6].




Fig. 9. A Hadamard matrix of order 280 generated from a set of Williamson matrices of order 70. Coloured entries represent 1s and white entries represent  $-1$ s.



The reason special-purpose solvers are much more effective is because they can exploit mathematical properties that Williamson matrices are known to satisfy, such as the power spectral density or *PSD condition*.

Briefly, the PSD of a vector  $X = [x_0, \dots, x_{n-1}]$  is the vector of squared magnitudes of the discrete Fourier transform of  $X$ . In other words, the  $k$ th entry of  $\text{PSD}(X)$  is  $|\sum_{j=0}^{n-1} x_j \omega^{jk}|^2$  where  $\omega$  is a primitive  $n$ th root of unity. Amazingly, if  $A$  is the first row of a Williamson matrix then all entries of  $\text{PSD}(A)$  are at most  $4n$ . This is a very useful condition since the vast majority of  $\{\pm 1\}$ -vectors do not satisfy it—thus exploiting it dramatically reduces the size of the search space. The problem, of course, is that SAT solvers have no conception of what a PSD value is and the condition is not easily encoded into Boolean logic.

However, CASs are adept at computing PSD values using the fast Fourier transform. Thus, when a SAT solver finds a matrix that it thinks might appear in a set of Williamson matrices it can provide this matrix to a CAS. The CAS checks if can be ruled out using the PSD condition, and if so—because the first row of the matrix has a PSD value larger than  $4n$ —then the matrix is removed from the search space. The SAT solver will continue to search for matrices that satisfy the constraints until either a set of Williamson matrices is found or it is able to certify that no Williamson matrices exist in the given order. This general method is outlined in Figure 10.

The SAT+CAS method using PSD filtering easily outperforms approaches using either SAT solvers alone or CASs alone—particularly in large orders where the SAT+CAS method can perform several orders of magnitude faster [8]. The naive size of the search space for Williamson matrices of order 70 is about  $2^{4(70/2)} \approx 10^{42}$  making it rather surprising that such a space can be searched *exhaustively*. But using the powerful search routines of SAT solvers coupled with the sophisticated mathematical capabilities available in CASs (and some additional powerful filtering criteria [16]) make this search tractable to complete on a modern CPU in under 1,000 hours.

## CONCLUSION

We’ve seen that the coupling of SAT solvers with computer algebra systems can effectively search large spaces specified via mathematical constraints—combining the power of SAT solvers with the expressiveness of CASs. Although we’ve focused on applications to finite geometry and combinatorics the paradigm is very general: a huge number of mathematical problems stand to benefit from this unique combination of SAT and CAS tools. The type of problems particularly attractive for pursuing with SAT+CAS methods have two primary characteristics. First, the problem is “Booleanizable” to an extent—in other words, a large part of the search space can be specified in Boolean logic, allowing the search power of SAT solvers to be exploited. Second, there are mathematical properties, theorems, or invariants that *cannot* easily be specified in Boolean logic but significantly decrease the search space size. These kind of constraints lie perfectly in the wheelhouse of a CAS and as we’ve outlined can dramatically improve the running time of the SAT solver. Best of all, they are not limited to any particular branch of mathematics, and indeed there have been applications of SAT+CAS methods to a surprising variety of problems in areas like cryptanalysis [23], program synthesis [24], and circuit verification [26]. The increasing visibility of initiatives like the SC-square project [14] makes it likely that




Fig. 10. A outline of the SAT+CAS method as applied to searching for Williamson matrices. The matrix provided to the CAS is blocked when it has a PSD value that is too large.

SAT+CAS methods will be with us for a long time to come—solving problems larger and larger than previously thought possible.

## REFERENCES

- [1] E. Ábrahám. Building bridges between symbolic computation and satisfiability checking. In S. Linton, editor, *Proceedings of the 2015 ACM on International Symposium on Symbolic and Algebraic Computation*, pages 1–6. ACM, 2015.
- [2] E. Ábrahám, J. Abbott, B. Becker, A. M. Bigatti, M. Brain, A. Cimatti, J. H. Davenport, M. England, P. Fontaine, S. Forrest, V. Ganesh, A. Griggio, D. Kroening, and W. M. Seiler.  $SC^2$ : when satisfiability checking and symbolic computation join forces. In G. Reger and D. Traytel, editors, *ARCADE 2017, 1st International Workshop on Automated Reasoning: Challenges, Applications, Directions, Exemplary Achievements*, volume 51 of *EPiC Series in Computing*, pages 6–10. EasyChair, 2017.
- [3] K. Appel and W. Haken. Every planar map is four colorable. Part I: Discharging. *Illinois Journal of Mathematics*, 21(3):429–490, 1977.
- [4] C. Bright, K. K. H. Cheung, B. Stevens, I. Kotsireas, and V. Ganesh. Unsatisfiability proofs for weight 16 codewords in Lam’s problem. *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, pages 1460–1466, 2020.
- [5] C. Bright, K. K. H. Cheung, B. Stevens, D. Roy, I. Kotsireas, and V. Ganesh. A nonexistence certificate for projective planes of order ten with weight 15 codewords. *Applicable Algebra in Engineering, Communication and Computing*, 31(3):195–213, 2020.
- [6] C. Bright, V. Ganesh, A. Heinle, I. Kotsireas, S. Nejati, and K. Czarnecki. MATHCHECK2: A SAT+CAS verifier for combinatorial conjectures. In V. P. Gerdt, W. Koepf, W. M. Seiler, and E. V. Vorozhtsov, editors, *Proceedings of the 18th International Workshop on Computer Algebra in Scientific Computing*, pages 117–133. Springer, 2016.
- [7] C. Bright, J. Gerhard, I. Kotsireas, and V. Ganesh. Effective problem solving using SAT solvers. In *Maple in Mathematics Education and Research*, volume 1125 of *Communications in Computer and Information Science*, pages 205–219. Springer, 2020.
- [8] C. Bright, I. Kotsireas, and V. Ganesh. Applying computer algebra systems with SAT solvers to the Williamson conjecture. *Journal of Symbolic Computation*, 100:187–209, 2020.
- [9] C. Bright, I. Kotsireas, and V. Ganesh. New infinite families of perfect quaternion sequences and Williamson sequences. *IEEE Transactions on Information Theory*, to appear.
- [10] C. Bright, I. Kotsireas, A. Heinle, and V. Ganesh. Complex Golay pairs up to length 28: A search via computer algebra and programmatic SAT. *Journal of Symbolic Computation*, 102:153–172, 2021.
- [11] R. H. Bruck and H. J. Ryser. The nonexistence of certain finite projective planes. *Canadian Journal of Mathematics*, 1(1):88–93, 1949.
- [12] J. L. Carter. *On the existence of a projective plane of order ten*. PhD thesis, University of California, Berkeley, 1974.
- [13] K. Clarkson and S. Whitesides. On the non-existence of maximal 6-arcs in projective planes of order 10. In *Poster session at IWOCOA 2014, the 25th International Workshop on Combinatorial Algorithms*, 2014.
- [14] J. H. Davenport, M. England, A. Griggio, T. Sturm, and C. Tinelli. Symbolic computation and satisfiability checking. *Journal of Symbolic Computation*, 100:1–10, 2020.
- [15] D. Ž. Đoković. Williamson matrices of order  $4n$  for  $n = 33, 35, 39$ . *Discrete Mathematics*, 115(1-3):267–271, 1993.
- [16] D. Ž. Đoković and I. S. Kotsireas. Compression of periodic complementary sequences and applications. *Designs, Codes and Cryptography*, 74(2):365–377, 2015.
- [17] S. W. Golomb and L. D. Baumert. The search for Hadamard matrices. *The American Mathematical Monthly*, 70(1):12–17, 1963.
- [18] J. Hadamard. Résolution d’une question relative aux déterminants. *Bulletin des Sciences Mathématiques*, 17(1):240–246, 1893.
- [19] A. Haken. The intractability of resolution. *Theoretical Computer Science*, 39:297–308, 1985.
- [20] T. C. Hales and S. P. Ferguson. *The Kepler Conjecture: The Hales-Ferguson Proof*. Springer Science & Business Media, 2011.
- [21] M. J. H. Heule. Schur number five. In S. A. McIlraith and K. Q. Weinberger, editors, *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18)*, pages 6598–6606. AAAI Press, 2018.
- [22] M. J. H. Heule and O. Kullmann. The science of brute force. *Communications of the ACM*, 60(8):70–79, 2017.
- [23] J. Horáček. *Algebraic and Logic Solving Methods for Cryptanalysis*. PhD thesis, University of Passau, 2020.
- [24] J. P. Inala, S. Gao, S. Kong, and A. Solar-Lezama. REAS: Combining numerical optimization with SAT solving. *arXiv preprint arXiv:1802.04408*, 2018.
- [25] R. Jhala and R. Majumdar. Software model checking. *ACM Computing Surveys (CSUR)*, 41(4):1–54, 2009.
- [26] D. Kaufmann, A. Biere, and M. Kauers. SAT, computer algebra, multipliers. In L. Kovács and A. Voronkov, editors, *Vampire 2018 and Vampire 2019. The 5th and 6th Vampire Workshops*, volume 71 of *EPiC Series in Computing*, pages 1–18. EasyChair, 2020.
- [27] C. W. H. Lam. The search for a finite projective plane of order 10. *The American Mathematical Monthly*, 98(4):305–318, 1991.
- [28] C. W. H. Lam, L. Thiel, and S. Swiercz. The nonexistence of code words of weight 16 in a projective plane of order 10. *Journal of Combinatorial Theory, Series A*, 42(2):207–214, 1986.
- [29] C. W. H. Lam, L. Thiel, and S. Swiercz. The non-existence of finite projective planes of order 10. *Canadian Journal of Mathematics*, 41(6):1117–1123, 1989.
- [30] F. J. MacWilliams, N. J. A. Sloane, and J. G. Thompson. On the existence of a projective plane of order 10. *Journal of Combinatorial Theory, Series A*, 14(1):66–78, 1973.

- [31] J. P. Marques-Silva and K. A. Sakallah. Boolean satisfiability in electronic design automation. In *Proceedings of the 37th Annual Design Automation Conference*, pages 675–680, 2000.
- [32] R. Nieuwenhuis, A. Oliveras, and C. Tinelli. Solving SAT and SAT modulo theories: From an abstract Davis–Putnam–Logemann–Loveland procedure to DPLL( $T$ ). *Journal of the ACM (JACM)*, 53(6):937–977, 2006.
- [33] M. Plotkin. Binary codes with specified minimum distance. *IRE Transactions on Information Theory*, 6(4):445–450, 1960.
- [34] D. J. Roy. Confirmation of the non-existence of a projective plane of order 10. Master’s thesis, Carleton University, 2011.
- [35] M. Y. Vardi. Boolean satisfiability: Theory and engineering. *Communications of the ACM*, 57(3):5, 2014.
- [36] J. Williamson. Hadamard’s determinant theorem and the sum of four squares. *Duke Mathematical Journal*, 11(1):65–81, 1944.
- [37] E. Zulkoski, C. Bright, A. Heinle, I. Kotsireas, K. Czarnecki, and V. Ganesh. Combining SAT solvers with computer algebra systems to verify combinatorial conjectures. *Journal of Automated Reasoning*, 58(3):313–339, 2017.
- [38] E. Zulkoski, V. Ganesh, and K. Czarnecki. MathCheck: A math assistant via a combination of computer algebra systems and SAT solvers. In A. P. Felty and A. Middeldorp, editors, *Proceedings of the 25th International Conference on Automated Deduction*, volume 9195 of *Lecture Notes in Computer Science*, pages 607–622. Springer, 2015.