

# Effective problem solving using SAT solvers

Curtis Bright<sup>1,2</sup>, Jürgen Gerhard<sup>2</sup>, Ilias Kotsireas<sup>3</sup>, and Vijay Ganesh<sup>1</sup>

<sup>1</sup> University of Waterloo, Waterloo, Canada

<sup>2</sup> Maplesoft, Waterloo, Canada

<sup>3</sup> Wilfrid Laurier University, Waterloo, Canada

**Abstract.** In this article we demonstrate how to solve a variety of problems and puzzles using the built-in SAT solver of the computer algebra system Maple. Once the problems have been encoded into Boolean logic, solutions can be found (or shown to not exist) automatically, without the need to implement any search algorithm. In particular, we describe how to solve the  $n$ -queens problem, how to generate and solve Sudoku puzzles, how to solve logic puzzles like the Einstein riddle, how to solve the 15-puzzle, how to solve the maximum clique problem, and finding Graeco-Latin squares.

**Keywords:** SAT solving · Maple ·  $n$ -queens problem · Sudoku · Logic puzzles · 15-puzzle · Maximum clique problem · Graeco-Latin squares

## 1 Introduction

“... it is a constant source of annoyance when you come up with a clever special algorithm which then gets beaten by translation to SAT.”

—Chris Jefferson

The satisfiability (SAT) problem is to determine if a given Boolean expression can be satisfied—is there some way of assigning true and false to its variables that makes the whole formula true? Despite at first seeming disconnected from most of the kinds of problems that mathematicians care about we argue in this paper that it is in the interests of mathematicians to have a familiarity with SAT solving and encoding problems in SAT. An immense amount of effort over the past several decades has produced SAT solvers that are not only practical for many problems but are actually the fastest known way of solving an impressive variety of problems such as software and hardware verification problems [3]. They have also recently been used to resolve long-standing mathematical conjectures [12] and construct large combinatorial designs [7].

Since 2018, the computer algebra system Maple has included the award-winning SAT solver MapleSAT [15] as its built-in SAT solver. This solver can be used through the `Satisfy` command of the `Logic` package. `Satisfy` returns a satisfying assignment of a given Boolean expression (if one exists) or `NULL` if no satisfying assignment exists. In this paper we demonstrate through a number of detailed examples how `Satisfy` can be an effective and efficient way of solving a variety of problems and puzzles.

Very little prerequisites are necessary to understand this paper; the main necessary background is a familiarity with Boolean logic which we outline in Section 2. We then present effective solutions to the  $n$ -queens problem (Section 3), logic puzzles like the Einstein riddle (Section 4), Sudoku puzzles (Section 5), Euler’s Graeco-Latin square problem (Section 6), the maximum clique problem (Section 7), and the 15-puzzle (Section 8). In each case we require no knowledge of any of the special-purpose search algorithms that have been proposed to solve these problems; once the problems have been encoded into Boolean logic they are automatically solved using Maple’s `Satisfy`.

All of the examples discussed in this paper were implemented and run in Maple 2018 and Maple 2019. Due to space constraints we have not included our code in this paper, but Maple worksheets containing complete implementations have been made available online through the Maple Application Center [6].

## 2 Background

A basic understanding of Boolean logic is the only prerequisite necessary to understand the solutions described in this paper. One of the main advantages of Boolean logic (but also one of its main disadvantages) is its simplicity: each variable can assume only one of two values denoted by true and false. Boolean expressions consist of variables joined by Boolean operators. The most common Boolean operators (and the ones available in the `Logic` package of Maple) are summarized in Table 1.

The  $\vee$  (or),  $\wedge$  (and), and  $\neg$  (not) operators have meanings based on their everyday English meanings:  $x_1 \vee \cdots \vee x_n$  is true exactly when at least one  $x_i$  is true,  $x_1 \wedge \cdots \wedge x_n$  is true exactly when all  $x_i$  are true, and  $\neg x$  is true exactly when  $x$  is false. More generally,  $x \Leftrightarrow y$  is true exactly when  $x$  and  $y$  have the same truth values,  $x \Rightarrow y$  is false exactly when  $y$  is true and  $x$  is false,  $x_1 \underline{\vee} \cdots \underline{\vee} x_n$  is true exactly when an odd number of  $x_i$  are true,  $x_1 \uparrow \cdots \uparrow x_n$  is true exactly when at least one  $x_i$  is false, and  $x_1 \downarrow \cdots \downarrow x_n$  is true exactly when all  $x_i$  are false.

A *literal* is an expression of the form  $x$  or  $\neg x$  where  $x$  is a Boolean variable. A *clause* is an expression of the form  $l_1 \vee \cdots \vee l_n$  where all  $l_i$  are literals. A *conjunctive*

Name	Symbol	Arity	Maple Syntax
Negation	$\neg$	1	<code>&amp;not</code>
Conjunction	$\wedge$	$n$ -ary	<code>&amp;and</code>
Disjunction	$\vee$	$n$ -ary	<code>&amp;or</code>
Implication	$\Rightarrow$	2	<code>&amp;implies</code>
Biconditional	$\Leftrightarrow$	2	<code>&amp;iff</code>
Alternative denial	$\uparrow$	$n$ -ary	<code>&amp;nand</code>
Joint denial	$\downarrow$	$n$ -ary	<code>&amp;nor</code>
Exclusive disjunction	$\underline{\vee}$	$n$ -ary	<code>&amp;xor</code>

**Table 1.** The Boolean logical operators available in Maple.

*normal form* (CNF) expression is of the form  $c_1 \wedge \cdots \wedge c_n$  where all  $c_i$  are clauses. A standard theorem of Boolean logic is that any expression can be converted into an equivalent expression in conjunctive normal form where two expressions are said to be *equivalent* if they assume the same truth values under all variable assignments.

The current algorithms used in state-of-the-art SAT solvers require that the input formula be given in conjunctive normal form. While this is convenient for the purposes of designing efficient solvers it is not convenient for the mathematician who wants to express their problem in Boolean logic—not all expressions are *naturally* expressed in conjunctive normal form. An advantage that Maple has over most current state-of-the-art SAT solvers is that Maple does not require the input to be given in conjunctive normal form. Since the algorithms used by MapleSAT require CNF to work properly, Maple internally converts expressions into CNF automatically. This is done by using a number of equivalence transformations, e.g., the expression  $x \Rightarrow y$  is rewritten as the clause  $\neg x \vee y$ .


Care has been taken to make the necessary conversion to CNF efficient. This is important because conversions that use the most straightforward equivalence rules generally require exponential time to complete. For example, the Maple command `Normalize` from the `Logic` package can be used to convert an expression into CNF. But watch out—many expressions explode in size following this conversion. For example, the expression  $x_1 \underline{\vee} \cdots \underline{\vee} x_n$  when converted into CNF contains  $2^{n-1}$  clauses. The main trick used to make the conversion into CNF efficient is the *Tseitin transformation* [23]. This transformation avoids the exponential blowup of the straightforward transformations by using additional variables to derive a new formula that is satisfiable if and only if the original formula is satisfiable. For example, the expression  $x_1 \underline{\vee} \cdots \underline{\vee} x_n$  is rewritten as  $(t \underline{\vee} x_3 \underline{\vee} \cdots \underline{\vee} x_n) \wedge C$  where  $t$  is a new variable and  $C$  is a CNF encoding of the formula  $t \Leftrightarrow (x_1 \underline{\vee} x_2)$ , namely,

$$(\neg x_1 \vee x_2 \vee t) \wedge (x_1 \vee \neg x_2 \vee t) \wedge (x_1 \vee x_2 \vee \neg t) \wedge (\neg x_1 \vee \neg x_2 \vee \neg t).$$

The transformation is then recursively applied to  $t \underline{\vee} x_3 \underline{\vee} \cdots \underline{\vee} x_n$  (the part of the formula not in CNF) until the entire formula is in CNF. The Maple command `Tseitin` of the `Logic` package can be applied to convert an arbitrary formula into CNF using this translation. Thus, Maple offers us the convenience of not requiring encodings to be in CNF while avoiding the inefficiencies associated with a totally unrestricted encoding.

### 3 The $n$ -queens problem

The  $n$ -queens problem is to place  $n$  chess queens on an  $n \times n$  chessboard such that no two queens are mutually attacking (i.e., in the same row, column, or diagonal). The problem was first proposed for  $n = 8$  by Bezzel in 1848 and the first solution for general  $n$  was given by Pauls in 1874 [2]. The problem is solvable for all  $n \geq 4$ ; a solution for  $n = 8$  (found in 0.015 seconds using `Satisfy`) is shown in Figure 1.



**Fig. 1.** A visual representation of a solution for the 8-queens problem (left) and the variables assigned to true in this solution using our SAT encoding. Following traditional chess convention, columns are indexed left to right and rows are indexed bottom to top.

The  $n$ -queens problem is a standard example of a constraint satisfaction problem [18]. The encoding that we use for this problem uses the  $n^2$  Boolean variables  $Q_{x,y}$  with  $1 \leq x, y \leq n$  to denote if there is a queen on square  $(x, y)$ . There are two kinds of constraints necessary for this problem: *positive* constraints that say that there are  $n$  queens on the board and *negative* constraints that say that queens do not attack each other. A satisfying assignment of these constraints exists exactly when the  $n$ -queens problem is solvable.

Since there are  $n$  rows and each row must contain a queen the positive constraints are of the form  $Q_{1,j} \vee \dots \vee Q_{n,j}$  for  $1 \leq j \leq n$ . Similarly, each column must contain a queen; these constraints are of the form  $Q_{i,1} \vee \dots \vee Q_{i,n}$  for  $1 \leq i \leq n$ . The negative constraints say that if  $(x, y)$  contains a queen then all squares attacked by  $(x, y)$  do not contain a queen. These constraints are represented in Boolean logic by  $Q_{x,y} \Rightarrow \neg(A_1 \vee \dots \vee A_k)$  where  $\{A_1, \dots, A_k\}$  are the variables “attacked” by a queen on  $(x, y)$ . In general this encoding uses  $\Theta(n^2)$  constraints in order  $n$ . Typically `Satisfy` is able to solve each order using slightly more time than the previous order and the last order it can solve in under a second is  $n = 32$ .

## 4 The Einstein riddle

The Einstein riddle is a logic puzzle apocryphally attributed to Albert Einstein and is often stated with the remark that it is only solvable by 2% of the world’s population. The true source of the puzzle is unknown, but a version of it appeared in the magazine *Life International* in 1962. In the puzzle there are five houses in a row with each house a different colour and each house owned by a man of a different nationality. Additionally, each of the owners have a different pet, prefer a different kind of drink, and smoke a different brand of cigarette. Furthermore, the following information is given:

1. The Brit lives in the red house.
2. The Swede keeps dogs as pets.
3. The Dane drinks tea.
4. The green house is next to the white house, on the left.
5. The owner of the green house drinks coffee.
6. The person who smokes Pall Mall rears birds.
7. The owner of the yellow house smokes Dunhill.
8. The man living in the centre house drinks milk.
9. The Norwegian lives in the first house.
10. The man who smokes Blends lives next to the one who keeps cats.
11. The man who keeps horses lives next to the man who smokes Dunhill.
12. The man who smokes Blue Master drinks beer.
13. The German smokes Prince.
14. The Norwegian lives next to the blue house.
15. The man who smokes Blends has a neighbour who drinks water.

The puzzle is: Who owns the fish?

To solve this riddle using Maple, we label the houses 1 to 5 and use the variables  $S_{i,a}$  where  $1 \leq i \leq 5$  and  $a$  is an attribute (one of the colours, nationalities, pets, drinks, or cigarette brands). For example, if  $a$  is a colour then  $a$  is in the set  $C := \{\text{red, green, white, yellow, blue}\}$  and similarly for the other attribute types; there are five distinct possible attributes for each type of attribute. In total there are  $5^2 = 25$  possible values for  $a$  and  $5^3 = 125$  variables  $S_{i,a}$ .

We know that each attribute is not shared among the five houses or their owners. Since there are exactly five houses, each attribute must appear exactly once among the five houses. The knowledge that each attribute appears at least once can be encoded as the clauses  $\bigvee_{i=1}^5 S_{i,a}$  for each attribute  $a$  and the knowledge that each attribute is not shared can be encoded as  $S_{i,a} \Rightarrow \neg S_{j,a}$  where  $j$  is a house index not equal to  $i$  and  $a$  is an attribute. Additionally, the fact that each house has some colour is encoded as  $\bigvee_{a \in C} S_{i,a}$  for each house index  $i$  and the knowledge that each house cannot have two colours can be encoded as  $S_{i,c} \Rightarrow \neg S_{i,d}$  where  $i$  is a house index and  $c$  and  $d$  are two distinct colours (and similarly for the other kinds of attributes).

The known facts can be encoded into logic fairly straightforwardly; for example, the first fact can be encoded into logic as  $S_{i,\text{Brit}} \Rightarrow S_{i,\text{red}}$  for house indices  $i$  and the last fact can be encoded as  $S_{i,\text{Blends}} \Rightarrow (S_{i-1,\text{water}} \vee S_{i+1,\text{water}})$  for  $1 < i < 5$  and  $(S_{1,\text{Blends}} \Rightarrow S_{2,\text{water}}) \wedge (S_{5,\text{Blends}} \Rightarrow S_{4,\text{water}})$ . Using `Satisfy` on these constraints produces the unique satisfying solution (that includes the equations  $S_{4,\text{German}} = \text{true}$  and  $S_{4,\text{fish}} = \text{true}$ , thereby solving the puzzle) in under 0.01 seconds.

## 5 Sudoku puzzles

Sudoku is a popular puzzle that appears in many puzzle books and newspapers. Given a 9 by 9 grid whose squares are either blank or contain a number between 1 and 9, the objective is to fill in the blank squares in such a way that each row and column contains exactly one digit between 1 and 9. Additionally, each of the

8								
		3	6					
	7			9		2		
	5				7			
				4	5	7		
			1				3	
		1					6	8
		8	5				1	
	9					4		

$S_{1,1,8}$	$S_{2,3,3}$	$S_{2,4,6}$
$S_{3,2,7}$	$S_{3,5,9}$	$S_{3,7,2}$
$S_{4,2,5}$	$S_{4,6,7}$	$S_{5,5,4}$
$S_{5,6,5}$	$S_{5,7,7}$	$S_{6,4,1}$
$S_{6,8,3}$	$S_{7,3,1}$	$S_{7,8,6}$
$S_{7,9,8}$	$S_{8,3,8}$	$S_{8,4,5}$
$S_{8,8,1}$	$S_{9,2,9}$	$S_{9,7,4}$

**Fig. 2.** A Sudoku puzzle estimated to have a difficulty rating of 11 stars, whereas the most challenging Sudoku puzzles that are usually published are given 5 stars. On the right are the starting constraints (as unit clauses) of our encoding for this puzzle.

nine 3 by 3 subgrids which compose the grid (called blocks) must also contain exactly one digit between 1 and 9. Figure 2 contains a Sudoku puzzle designed by mathematician Arto Inkala and claimed to be the world’s hardest Sudoku [9].

It is known that Sudoku can be modelled as a SAT problem [16] or a constraint satisfaction problem [20]. A straightforward encoding uses  $9^3 = 729$  variables  $S_{i,j,k}$  with  $1 \leq i, j, k \leq 9$  where  $S_{i,j,k}$  is true exactly when the square  $(i, j)$  contains the digit  $k$ . The rules of Sudoku state that each square must be filled with a digit between 1 and 9 and that the same digit cannot appear twice in the same row, column, or block. The first constraint has the form  $S_{i,j,1} \vee \dots \vee S_{i,j,9}$  for all  $1 \leq i, j \leq 9$  and the second constraint has the form  $S_{i,j,k} \Rightarrow \neg S_{i',j',k}$  for all  $1 \leq i, j, i', j', k \leq 9$  where  $(i, j)$  does not equal  $(i', j')$  but is in the same row, column, or block as  $(i, j)$ .

One can also include the constraints  $S_{i,j,k} \Rightarrow \neg S_{i,j,k'}$  for all  $1 \leq i, j, k, k' \leq 9$  with  $k \neq k'$  that say that each square can contain at most one digit. However, these constraints are unnecessary since they are logically implied by the first two constraints. In our tests, including these additional constraints slightly decreased the performance of `Satisfy`. Without the additional constraints the “world’s hardest Sudoku” was solved in 0.25 seconds and with them it was solved in 0.33 seconds.

Additionally, we developed a method of generating Sudoku puzzles with a unique solution using `Satisfy`. This allowed us to write an interactive Sudoku game where random puzzles can automatically be generated on command. To begin, `Satisfy` is used to find a solution to the above Sudoku constraints with an empty grid (no starting clues) and the produced satisfying solution generates a completed Sudoku grid  $R$ . A random seed is passed to the SAT solver so that a different solution  $R$  is generated each time; this is done by passing the following `solvoptions` in the call to `Satisfy`:

```
[rnd_init_act=true, random_seed=floor(1000*time[real]())]
```

Let  $R_{i,j}$  denote the  $(i, j)$  the entry in the solution  $R$  where  $1 \leq i, j \leq 9$ . The  $R_{i,j}$  are randomly ordered and the first 50 entries  $R_{i,j}$  are selected as the potential starting configuration of a Sudoku puzzle. This puzzle has the solution  $R$  by construction, though other solutions may also exist.

To verify that the generated solution is unique, we re-run `Satisfy` with the additional 50 unit clauses corresponding to the starting configuration along with the constraint  $\bigvee_{R_{i,j}=k} \neg S_{i,j,k}$  which blocks the solution  $R$ . If `Satisfy` returns another solution then we start over and find a new  $R$  to try. Otherwise the starting configuration forms a legal Sudoku puzzle.


Additionally, it may be the case that we can use fewer than 50 entries and still obtain a Sudoku puzzle with a unique solution. To estimate how many entries need to be assigned using only a few extra calls to the SAT solver we use a variant of binary search, letting  $l := 20$  and  $h := 50$  be lower and upper bounds on how many entries we will define in the puzzle. Next, we let  $m := \text{round}((l+h)/2)$  and repeat the first step except using only the first  $m$  entries  $R_{i,j}$ . If the resulting SAT instance is satisfiable then we need to use strictly more than  $m$  entries to ensure that a unique solution exists and if the resulting SAT instance is unsatisfiable then we can perhaps use strictly fewer than  $m$  entries. Either way, we improve the bounds on how many entries to assign (in the former case we can update  $l$  to  $m$  and in the latter case we can update  $h$  to  $m$ ) and this step can be repeated a few times to find more precise bounds on how many entries need to be assigned to ensure a unique solution exists.

## 6 Euler’s Graeco-Latin square problem

A *Latin square* of order  $n$  is an  $n \times n$  matrix containing integer entries between 1 and  $n$  such that every row and every column contains each entry exactly once. Two Latin squares are *orthogonal* if the superposition of one over the other produces all  $n^2$  distinct pairs of integers between 1 and  $n$ . A pair of orthogonal Latin squares was called a *Graeco-Latin square* by the mathematician Leonhard Euler who in 1782 used Latin characters to represent the entries of the first square and Greek characters to represent the entries of the second square [11]. Figure 3 contains a visual representation of a Graeco-Latin square.

Euler studied the orders  $n$  for which Graeco-Latin squares exist and found methods for constructing them when  $n$  was odd or a multiple of 4. Since such squares do not exist for  $n = 2$  and he was unable to find a solution for  $n = 6$  he conjectured that Graeco-Latin squares do not exist when  $n \equiv 2 \pmod{4}$ . Euler’s conjecture became famous as he was not able to resolve it in his lifetime.

The first progress on the conjecture did not come until over a hundred years later when in 1900 Tarry showed that Graeco-Latin squares of order 6 do not exist [22]. This gave credence to Euler’s conjecture and many mathematicians thought the conjecture was true—in fact, three independent proofs of the conjecture were published in the early 20th century [17,19,24]. In 1959–1960 Bose, Shrikhande, and Parker [4,5] made explosive news (even appearing on the front page of the New York Times) by showing that these proofs were invalid by giving



**Fig. 3.** On the left is a visual representation of a Graeco-Latin square of order 10 with each colour denoting a separate integer. The entries of the second Latin square (represented by the small squares in the image) are superimposed onto the entries of the first Latin square. On the right are the variables corresponding to the first column of this Graeco-Latin square that are assigned to true using our encoding.

explicit constructions for Graeco-Latin squares in all orders except two and six. As it turns out, a lot of time could have been saved if Euler had a copy of Maple—we now show that Euler’s conjecture can be automatically disproven in Maple. With `Satisfy` we are able to construct small Graeco-Latin squares without any knowledge of search algorithms or construction methods.

Our encoding for the Graeco-Latin square problem of order  $n$  uses the  $2n^3$  variables  $A_{i,j,k}$  and  $B_{i,j,k}$  with  $1 \leq i, j, k \leq n$ . The variables  $A_{i,j,k}$  will be true exactly when the  $(i, j)$ th entry of the Latin square  $A$  is  $k$  and  $B_{i,j,k}$  will be true exactly when the  $(i, j)$ th entry of the Graeco square  $B$  is  $k$ .

There are three kinds of constraints that specify that  $(A, B)$  is a Graeco-Latin square: Those that specify that every entry of  $A$  and  $B$  is an integer between 1 and  $n$ , those that specify that the rows and columns of  $A$  and  $B$  contain no duplicate entries, and those that specify that  $A$  and  $B$  are orthogonal. Additionally, there are constraints that are not logically necessary but help cut down the search space. Some work has previously been done using SAT solvers to search for special kinds of Graeco-Latin squares [25]. The encoding we use is similar but takes advantage of the fact that Maple does not require constraints to be specified in conjunctive normal form.

First, we specify that the entries of  $A$  are well-defined, i.e., consist of a single integer between 1 and  $n$ . The constraints that say that each entry of  $A$  contains at least one integer are of the form  $A_{i,j,1} \vee \dots \vee A_{i,j,n}$  for each index pair  $(i, j)$  and the constraints that say that each entry of  $A$  contains at most one integer are of the form  $A_{i,j,k} \Rightarrow \neg A_{i,j,l}$  for each index pair  $(i, j)$  and integer  $k \neq l$ . Similar constraints are also used to specify that the entries of  $B$  are well-defined.



Second, we specify that  $A$  is a Latin square, i.e., all columns and rows contain distinct entries. These have the form  $A_{i,j,k} \Rightarrow \neg A_{i',j',k}$  where  $1 \leq k \leq n$  and  $(i,j) \neq (i',j')$  but  $(i,j)$  is in the same column or row as  $(i',j')$ . Similarly, we also specify that  $B$  is a Latin square.

Third, we specify that  $A$  and  $B$  are orthogonal, i.e., for every pair  $(k,l)$  there exists some pair  $(i,j)$  such that  $A_{i,j,k} \wedge B_{i,j,l}$  holds. These constraints are of the form  $\bigvee_{i,j=1}^n (A_{i,j,k} \wedge B_{i,j,l})$  for each pair  $(k,l)$ .

Lastly, we include some “symmetry breaking” constraints. These constraints are not strictly necessary but they shrink the search space and thereby make the search more efficient. In general, when a search space splits into symmetric subspaces it is beneficial to add constraints that remove or “break” the symmetry. Graeco-Latin squares  $(A,B)$  have a number of symmetries, in particular, a row or column permutation simultaneously applied to  $A$  and  $B$  produces another Graeco-Latin square. Also, any permutation of  $\{1, \dots, n\}$  may be applied to the entries of either  $A$  or  $B$ .


The result of these symmetries is that any Graeco-Latin square can be transformed into one where the first row and column of  $A$  has entries in ascending order (by permuting rows/columns) and the first row of  $B$  has entries in ascending order (by renaming the entries of  $B$ ). Thus, we can assume the constraint  $\bigwedge_{i=1}^n (A_{1,i,i} \wedge B_{1,i,i} \wedge A_{i,1,i})$ . Altogether this encoding uses  $\Theta(n^4)$  constraints.

Using this encoding the orders up to eight can be solved in 25 total seconds (including 14 seconds to show that no Graeco-Latin squares exist in order six), a Graeco-Latin square of order nine can be found in about 45 minutes, and a Graeco-Latin square of order ten can be found in about 23 hours, thereby disproving Euler’s Graeco-Latin square conjecture.

## 7 The maximum clique problem

The maximum clique problem is to find a clique of maximum size in a given graph. A *clique* of a graph is a subset of its vertices that are all mutually connected (see Figure 4). The decision version of this problem (does a graph contain a clique of size  $k$ ?) is in NP, meaning that it is easy to verify the correctness of a solution if one can be found. By the Cook–Levin theorem [10] the problem can be encoded into a SAT instance in polynomial time. However, the reduction involves simulating the computation of a machine that solves the maximum clique problem and is therefore not very convenient to use in practice. Thus, we provide a simpler encoding into Boolean logic.

Suppose the given graph  $G$  has vertices labelled  $1, \dots, n$  and we want to find a clique of size  $k$  in  $G$ . Our encoding uses the variables  $x_1, \dots, x_n$  where  $x_i$  represents that the vertex  $i$  appears in the clique we are attempting to find. We need to enforce a constraint that says that if  $x_i$  and  $x_j$  are true (for any distinct vertices  $1 \leq i, j \leq n$ ) then the edge  $\{i, j\}$  exists in the graph  $G$ . Equivalently, if the edge  $\{i, j\}$  does not exist in the graph  $G$  then the variables  $x_i$  and  $x_j$  cannot both be true (for any vertices  $i, j$ ). In other words, for every edge  $\{i, j\}$  in the complement of  $G$  we use the clause  $\neg x_i \vee \neg x_j$ .



**Fig. 4.** On the left is a visual representation of a graph with 20 vertices and a highlighted clique of size 4 that was found in 0.025 seconds. On the right are the important variables assigned to true in the assignment returned by **Satisfy** for our encoding of the maximum clique problem for this graph.

Additionally, we need a way to enforce that the found clique is of size  $k$ . The most naive way to encode this is as a disjunction over all  $\binom{n}{k}$  conjunctions of length  $k$  on the variables  $x_1, \dots, x_n$ . However, this encoding is very inefficient in practice. A cleverer encoding uses Boolean counter variables  $s_{i,j}$  (where  $0 \leq i \leq n$  and  $0 \leq j \leq k$ ) that represent that at least  $j$  of the variables  $x_1, \dots, x_i$  are assigned to true. We know that  $s_{0,j}$  will be false for  $1 \leq j \leq k$  and that  $s_{i,0}$  will be true for  $0 \leq i \leq n$ . Additionally, we know that  $s_{i,j}$  is true exactly when  $s_{i-1,j}$  is true or  $x_i$  is true and  $s_{i-1,j-1}$  is true. This is represented by the formulas

$$s_{i,j} \Leftrightarrow (s_{i-1,j} \vee (x_i \wedge s_{i-1,j-1})) \quad \text{for } 1 \leq i \leq n \text{ and } 1 \leq j \leq k$$

or in conjunctive normal form by the clauses  $\neg s_{i-1,j} \vee s_{i,j}$ ,  $\neg x_i \vee \neg s_{i-1,j-1} \vee s_{i,j}$ ,  $\neg s_{i,j} \vee s_{i-1,j} \vee x_i$ , and  $\neg s_{i,j} \vee s_{i-1,j} \vee s_{i-1,j-1}$ . To enforce that the found clique contains at least  $k$  vertices we also assign  $s_{n,k}$  to true.

To solve the maximum clique problem for a given graph  $G$  we initialize  $k$  to 3 (assuming the graph has at least one edge, otherwise the problem is trivial) and use the above encoding to search for a clique of size  $k$ . If such a clique exists we increase  $k$  by 1 and repeat the search in this manner until a clique of size  $k$  does not exist. The last satisfying assignment found then provides a maximum clique of  $G$  using an encoding with  $\Theta(nk)$  variables and  $\Theta(n^2)$  clauses.

This implementation was tested on the maximum clique problems from the second DIMACS implementation challenge [13]. Additionally, it was compared with Maple's **MaximumClique** function from the **GraphTheory** package that uses a branch-and-bound backtracking algorithm [14]. Of the 80 benchmarks, the SAT method solved 18 in under 3 minutes and the branch-and-bound method solved 15. The SAT approach was faster in over half of the solved benchmarks and in one case solved a benchmark in 8 seconds that **MaximumClique** required 57 minutes to solve (see Table 2).

Benchmark	SAT time (sec)	Maple 2018 (sec)	Vertices	Edges	Clique size
brock200_2	23.52	<b>22.19</b>	200	9876	12
c-fat200-1	0.71	<b>0.03</b>	200	1534	12
c-fat200-2	2.32	<b>0.13</b>	200	3235	24
c-fat200-5	<b>9.07</b>	20.63	200	8473	58
c-fat500-1	4.78	<b>0.13</b>	500	4459	14
c-fat500-2	11.12	<b>0.82</b>	500	9139	26
c-fat500-5	<b>42.46</b>	132.84	500	23191	64
c-fat500-10	<b>134.42</b>	Timeout	500	46627	126
hamming6-2	<b>0.64</b>	51.10	64	1824	32
hamming6-4	0.04	<b>0.02</b>	64	704	4
hamming8-2	<b>58.51</b>	Timeout	256	31616	128
hamming8-4	<b>7.96</b>	3393.26	256	20864	16
johnson8-2-4	<b>0.01</b>	<b>0.01</b>	28	210	4
johnson8-4-4	<b>0.23</b>	7.80	70	1855	14
johnson16-2-4	<b>5.62</b>	642.24	120	5460	8
keller4	<b>7.76</b>	414.80	171	9435	11
MANN_a9	<b>0.13</b>	226.18	45	918	16
p_hat300-1	11.81	<b>3.30</b>	300	10933	8
p_hat500-1	308.87	<b>34.60</b>	500	31569	9
p_hat700-1	1281.62	<b>169.68</b>	700	60999	11

**Table 2.** A comparison of the SAT method and the `MaximumClique` function of Maple 2018 on a collection of maximum clique benchmarks with a timeout of an hour.

The SAT method has been made available in Maple 2019 by using the `method=sat` option of `MaximumClique`. By default the `MaximumClique` function in Maple 2019 will run the previous method used by Maple and the SAT method in parallel and return the answer of whichever method finishes first. This hybrid approach is the method of choice especially when more than a single core is available.

## 8 The 15-puzzle

The 15-puzzle is a classic “sliding tile” puzzle that was first designed in 1880 and became very popular in the 1880s [21]. It consists of a  $4 \times 4$  grid containing tiles numbered 1 through 15 along with one missing tile (see Figure 5). The objective of the puzzle is to arrange the tiles so that they are in ascending order when read from left to right and top to bottom and to end with the blank tile in the lower right. The only moves allowed are those that slide a tile adjacent to the blank space into the blank space. Half of the possible starting positions are solvable [1] and the hardest legal starting positions require eighty moves to complete [8].

Our encoding of the 15-puzzle is more complicated because unlike the other problems we’ve considered, a solution to the 15-puzzle is not static. In other words, our encoding must be able to deal with the state of the puzzle changing over time. To do this we use the variables  $S_{i,j,n,t}$  to denote that the entry at  $(i, j)$

5	1	7	3
9	2	11	4
13	6	15	8
	10	14	12

$S_{1,1,5,0}$	$S_{1,2,1,0}$	$S_{1,3,7,0}$	$S_{1,4,3,0}$
$S_{2,1,9,0}$	$S_{2,2,2,0}$	$S_{2,3,11,0}$	$S_{2,4,4,0}$
$S_{3,1,13,0}$	$S_{3,2,6,0}$	$S_{3,3,15,0}$	$S_{3,4,8,0}$
$S_{4,1,16,0}$	$S_{4,2,10,0}$	$S_{4,3,14,0}$	$S_{4,4,12,0}$

**Fig. 5.** On the left is a visual representation of one starting configuration of the 15-puzzle and on the right are the starting constraints (as unit clauses) for this starting configuration in our encoding.

contains tile  $n$  at timestep  $t$ . Here  $1 \leq i, j \leq 4$ ,  $1 \leq n \leq 16$  (we let 16 denote the blank tile), and  $0 \leq t \leq 80$  since each instance requires at most 80 moves to complete.

The board does not permit two tiles to occupy the same location at the same time; these constraints are of the form  $S_{i,j,n,t} \Rightarrow \neg S_{i,j,m,t}$  for each tile numbers  $n \neq m$ , valid indices  $i$  and  $j$ , and valid timesteps  $t$ . Next we need to generate constraints that tell the SAT solver how the state of the board can change from time  $t$  to time  $t + 1$ . There are two cases to consider, depending on if a square (or an adjacent square) contains the blank tile.

The easier case is when a square  $(i, j)$  and none of the squares adjacent to that square are blank. In that case, the rules of the puzzle imply that the tile in square  $(i, j)$  does not change. We define the function  $\text{doesNotChange}(i, j, t)$  to be the constraint that says that the tile in square  $(i, j)$  does not change at time  $t$ ; these constraints are of the form  $\bigwedge_{n=1}^{16} (S_{i,j,n,t} \Leftrightarrow S_{i,j,n,t+1})$ . We also use the function  $\text{adj}(i, j)$  to denote the squares adjacent to  $(i, j)$  and define  $\text{notEqualOrAdjacentToBlank}(i, j, t)$  to be  $\neg S_{i,j,16,t} \wedge \bigwedge_{(k,l) \in \text{adj}(i,j)} \neg S_{k,l,16,t}$ . The static transition constraints are of the form

$$\text{notEqualOrAdjacentToBlank}(i, j, t) \Rightarrow \text{doesNotChange}(i, j, t)$$

for all valid squares  $(i, j)$  and timesteps  $t$ .

The harder transition case is when a square  $(i, j)$  contains the blank tile. In this case we need to encode the fact that the blank tile will switch positions with exactly one of the squares adjacent to square  $(i, j)$ . If the tile on square  $(i, j)$  switches positions with the square  $(k, l)$  at time  $t$  this can be encoded as the constraint  $\bigwedge_{n=1}^{16} (S_{i,j,n,t} \Leftrightarrow S_{k,l,n,t+1})$ . We also need to enforce that all squares adjacent to  $(i, j)$  other than  $(k, l)$  do not change; these constraints are of the form  $\bigwedge_{(x,y) \text{ doesNotChange}(x,y,t)}$  where  $(x, y)$  is adjacent to  $(i, j)$  but not equal to  $(k, l)$ . Let  $\text{oneTileMoved}(i, j, k, l, t)$  denote the conjunction of the above two

constraints. Then the slide transition constraints are of the form

$$S_{i,j,16,t} \Rightarrow \bigvee_{(k,l) \in \text{adj}(i,j)} \text{oneTileMoved}(i, j, k, l, t)$$

for all valid squares  $(i, j)$  and timesteps  $t$ .

The constraint  $\text{boardSolved}(t)$  that says the board is solved at timestep  $t$  can be encoded as  $\bigwedge_{i,j=1}^4 S_{i,j,4i+j-4,t}$ . For efficiency reasons we only start looking for solutions with at most 5 moves; if no solution is found then we look for solutions using at most 10 moves and continue in this manner until a solution is found. In other words, we call `Satisfy` with the given starting constraints, the constraints of the puzzle as described above, and the constraint  $\bigvee_{t=m-4}^m \text{boardSolved}(t)$  where  $m$  is initialized to 5 and then increased by 5 every time no solution is found.

This method was applied to the starting configuration from Figure 5. It found that no solutions with at most 5 moves exist in 1.5 seconds, no solutions with at most 10 moves exist in 2.8 seconds, and found a solution with 15 moves in 6.3 seconds. It was also able to solve puzzles requiring up to 40 moves in 20 minutes. While this is not competitive with dedicated solvers for the 15-puzzle, it requires no knowledge beyond the rules of the game and makes an interesting example of how to push SAT solvers to their limits.

## 9 Conclusion

In this paper we've demonstrated how to solve a variety of problems and puzzles using the computer algebra system Maple and its SAT solver MapleSAT [15]. We discussed a number of encodings and ways for improving those encodings, e.g., by using symmetry breaking (as in Section 6) or by using auxiliary variables (as in Section 7). We also took advantage of Maple's ability to solve SAT problems not encoded in conjunctive normal form in Sections 3, 6, and 8. Maple code for all the examples covered in this paper (including code to read the output of `Satisfy` and generate the figures included in this paper) are available for download from the Maple Application Center [6].

The implementations presented in this paper can be considered examples of *declarative programming* where the programmer focuses on describing the problem but not the solution—the computer automatically decides the best way to solve the problem. This is in contrast to *imperative programming* where a programmer needs to describe precisely *how* the computation is to take place. An advantage of declarative programming is that the programmer does not need to worry about specifying a potentially complicated search algorithm. However, a disadvantage of declarative programming is that it lacks the kind of detailed control over the method of solution that can be required for optimally efficient solutions. Furthermore, not all problems are naturally expressed in a declarative way.

As we saw in the maximum clique problem, sometimes declarative solutions can outperform imperative solutions. This also occurs in the *graph colouring* (or

*chromatic number*) problem of colouring the vertices of a graph using the fewest number of colours subject to the constraint that adjacent vertices are coloured differently. For example, prior to Maple 2018 the `ChromaticNumber` function required several hours to find a minimal colouring of the  $8 \times 8$  queens graph but a SAT encoding can solve this problem in under 10 seconds [6]. The SAT approach is available in Maple 2019 using the `method=sat` option of `ChromaticNumber`.

This somewhat unconventional manner of using Maple is not applicable to all problems but we hope the examples in this paper have convinced the reader that SAT solvers are more useful and powerful than they might at first appear. With its extensive logic functionality and convenient method of expressing logical constraints, Maple is an ideal tool for experimenting with SAT solvers and logical programming.

## References

1. Archer, A.F.: A modern treatment of the 15 puzzle. *The American Mathematical Monthly* **106**(9), 793–799 (1999)
2. Bell, J., Stevens, B.: A survey of known results and research areas for  $n$ -queens. *Discrete Mathematics* **309**(1), 1–31 (2009)
3. Biere, A., Heule, M.J.H., van Maaren, H., Walsh, T.: *Handbook of satisfiability*, *Frontiers in Artificial Intelligence and Applications*, vol. 185. IOS press (2009)
4. Bose, R.C., Shrikhande, S.S., Parker, E.T.: Further results on the construction of mutually orthogonal Latin squares and the falsity of Euler’s conjecture. *Canadian Journal of Mathematics* **12**, 189–203 (1960)
5. Bose, R.C., Shrikhande, S.S.: On the falsity of Euler’s conjecture about the non-existence of two orthogonal Latin squares of order  $4t+2$ . *Proceedings of the National Academy of Sciences of the United States of America* **45**(5), 734–737 (1959)
6. Bright, C.: Maple applications by Curtis Bright. <https://www.maplesoft.com/applications/Author.aspx?mid=345070>
7. Bright, C., Kotsireas, I., Ganesh, V.: A SAT+CAS method for enumerating Williamson matrices of even order. In: McIlraith, S., Weinberger, K. (eds.) *Thirty-Second AAAI Conference on Artificial Intelligence*. pp. 6573–6580. AAAI Press (2018)
8. Brünger, A., Marzetta, A., Fukuda, K., Nievergelt, J.: The parallel search bench ZRAM and its applications. *Annals of Operations Research* **90**, 45–63 (1999)
9. Collins, N.: World’s hardest sudoku: can you crack it? <https://www.telegraph.co.uk/news/science/science-news/9359579/Worlds-hardest-sudoku-can-you-crack-it.html> (2012)
10. Cook, S.A.: The complexity of theorem-proving procedures. In: *Proceedings of the third annual ACM symposium on Theory of computing*. pp. 151–158. ACM (1971)
11. Euler, L.: *Recherches sur un nouvelle espèce de quarrés magiques*. *Verhandelingen uitgegeven door het zeeuwsch Genootschap der Wetenschappen te Vlissingen* pp. 85–239 (1782)
12. Heule, M.J.H., Kullmann, O., Marek, V.W.: Solving very hard problems: Cube-and-conquer, a hybrid SAT solving method. In: Sierra, C. (ed.) *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pp. 4864–4868. IJCAI (2017)

13. Johnson, D.S., Trick, M.A.: Cliques, coloring, and satisfiability: second DIMACS implementation challenge, October 11–13, 1993, vol. 26. American Mathematical Soc. (1996)
14. Kreher, D., Stinson, D.: Combinatorial Algorithms: Generation, Enumeration, and Search. Discrete Mathematics and Its Applications, Taylor & Francis (1998)
15. Liang, J.H., Govind V. K., H., Poupart, P., Czarnecki, K., Ganesh, V.: An empirical study of branching heuristics through the lens of global learning rate. In: International Conference on Theory and Applications of Satisfiability Testing. pp. 119–135. Springer (2017), <https://ece.uwaterloo.ca/maplesat/>
16. Lynce, I., Ouaknine, J.: Sudoku as a SAT problem. In: 9th International Symposium on Artificial Intelligence and Mathematics (2006)
17. MacNeish, H.F.: Euler squares. *Annals of Mathematics* **23**(2), 221–227 (1922)
18. Nadel, B.A.: Representation selection for constraint satisfaction: A case study using  $n$ -queens. *IEEE Intelligent Systems* **5**(3), 16–23 (1990)
19. Peterson, J.: Les 36 officiers. *Annuaire des Mathématiciens* pp. 413–427 (1902)
20. Russell, S.J., Norvig, P.: Artificial Intelligence: A Modern Approach. Pearson Education (2010)
21. Slocum, J., Sonneveld, D.: The 15 Puzzle: How It Drove the World Crazy. Slocum Puzzle Foundation (2006)
22. Tarry, G.: Le problème des 36 officiers. *Association Française pour l’Avancement des Sciences: Compte Rendu de la 29<sup>me</sup> session en Paris 1900* pp. 170–203 (1901)
23. Tseitin, G.S.: On the complexity of derivation in propositional calculus. In: Slisenko, A.O. (ed.) *Studies in Constructive Mathematics and Mathematical Logic*. pp. 115–125 (1970)
24. Wernicke, P.: Das Problem der 36 Offiziere. *Jahresbericht der Deutschen Mathematiker-Vereinigung* **19**, 264–267 (1910)
25. Zaikin, O., Kochemazov, S.: The search for systems of diagonal Latin squares using the SAT@home project. *International Journal of Open Information Technologies* **3**(11), 4–9 (2015)