

# Improving SAT Solvers on Orthogonal Latin Square Problems

Aaron Barnoff<sup>1</sup>, Curtis Bright<sup>2</sup>

<sup>1</sup>*School of Computer Science, University of Windsor, Canada*

<sup>2</sup>*School of Computer Science, University of Waterloo, Canada*

## Abstract

Latin squares are  $n \times n$  matrices containing  $n$  symbols, where each symbol appears exactly once in each row and column. They were studied by Euler, later popularized through Sudoku, and remain a rich source of difficult combinatorial search problems. Two Latin squares are orthogonal mates if, when overlaid, no ordered pair of symbols repeats. Pairs of orthogonal Latin squares exist for every order except 2 and 6, but finding orthogonal Latin squares computationally can be challenging.

Satisfiability (SAT) solvers are strong at combinatorial search and have been used to resolve a number of various kinds of orthogonal Latin square problems. On the other hand, SAT solvers lack domain knowledge about Latin squares, such as the Euler–Parker algorithm for orthogonal mate construction. In this paper, we propose a hybrid method combining a SAT solver with the Euler–Parker algorithm (implemented using a Diophantine system solver) and show that the resulting solver is effective at finding certain kinds of orthogonal Latin squares. For example, certain pairs of  $10 \times 10$  orthogonal Latin squares whose existence was unknown for over 25 years were recently found by Bright, Keita, and Stevens using a SAT solver. The hardest cases could not be solved by the SAT solver CADICAL within seven days, but CADICAL augmented with an external Euler–Parker algorithm solves these cases in a median of around 5,100 seconds.

## Keywords

Satisfiability (SAT), Orthogonal Latin Squares, Euler–Parker Algorithm, Exact Cover

## 1. Introduction

A Latin square of order  $n$  is an  $n \times n$  array on  $n$  symbols in which each symbol occurs exactly once in each row and exactly once in each column (see Figure 1a). Two Latin squares of the same order are *orthogonal mates* if, when overlaid, every ordered pair of symbols occurs exactly once (see Figures 1b–1c). The study of orthogonal Latin squares has a long history. Euler conjectured in 1782 [1] that pairs of orthogonal Latin squares do not exist for orders  $n$  of the form  $4k + 2$ . The conjecture is trivial for  $n = 2$ , but it took over a hundred years before Tarry [2] proved the next case of the conjecture for  $n = 6$  in 1900. In the early 1900s, a number of mathematicians attempted to settle Euler’s conjecture, resulting in three independent publications laying claim to proving Euler’s conjecture [3, 4, 5]. It was thus quite surprising when Bose, Shrikhande, and Parker [6, 7, 8] in 1959–1960 showed by construction that in fact orthogonal Latin squares **do** exist for every order  $n$  with the exception of 2 and 6 and thus the “proofs” of Euler’s conjecture in the literature were flawed.

Bose, Shrikhande, and Parker’s resolution of Euler’s conjecture was completed without computer assistance. Though electronic computers were in their infancy, computers were used to search for orthogonal Latin squares around the same time Euler’s conjecture was disproven. For example, in 1960, Paige and Tompkins [9] published the results of a computer search for a pair of  $10 \times 10$  orthogonal Latin squares. They were unsuccessful in finding such a pair, and estimated their program would need around  $5 \cdot 10^{11}$  hours to determine if a given Latin square of order 10 has an orthogonal mate. On the other hand, in 1959, Parker [10] had also written a computer program

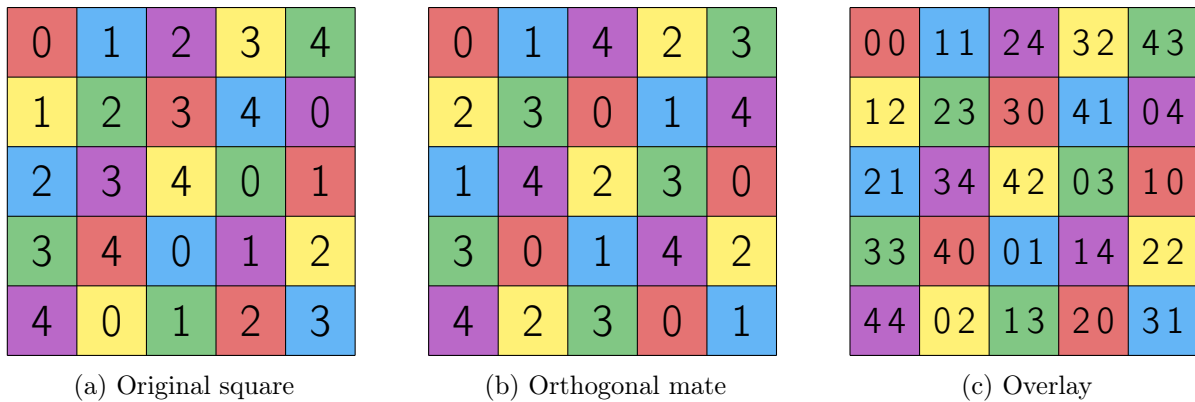
---

*11th International Workshop on Satisfiability Checking and Symbolic Computation, July 13, 2026, Oldenburg, Germany*

✉ barnoffa@uwindsor.ca (A. Barnoff); cbright@uwaterloo.ca (C. Bright)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



**Figure 1:** (a) A Latin square of order 5 decomposed into 5 non-overlapping transversals, with each transversal highlighted in a separate colour. (b) The orthogonal mate of the square in (a) produced by labelling the cells of each colour with a distinct symbol. (c) The original square and orthogonal mate overlaid, showing that every pair of symbols appears exactly once.

that only needed about an hour to find all orthogonal mates of a given Latin square of order 10. The primary reason why Parker’s program was so much faster than Paige and Tompkins’ program was because Parker’s program exploited a clever observation of Euler—that a Latin square of order  $n$  has an orthogonal mate if and only if it can be decomposed into  $n$  disjoint (i.e., non-overlapping) transversals. A *transversal* in a Latin square of order  $n$  is a collection of  $n$  cells in the square with exactly one cell from each row, exactly one cell from each column, and containing exactly one occurrence of each symbol. For example, in Figure 1a there are five transversals highlighted (the cells of each distinct colour form a distinct transversal).

Say  $L$  is the Latin square for which we want to find all orthogonal mates. Instead of searching for orthogonal mates of  $L$  cell-by-cell, Parker used an approach that decomposes the problem into two stages. The first stage finds all transversals of  $L$ , and the second stage finds all ways of combining those transversals together in non-overlapping ways to form a complete square. Whenever  $n$  non-overlapping transversals are found, an orthogonal mate is constructed by Euler’s observation (see Figure 1), and consequently this approach is known as the Euler–Parker algorithm for orthogonal mate construction. Knuth estimates that the Euler–Parker approach of searching for orthogonal mates of a  $10 \times 10$  square reduces the search cost by an astonishing factor of  $10^{12}$  [11].

Another method of searching for orthogonal pairs of Latin squares uses automated reasoning tools such as satisfiability (SAT) solvers [12], constraint or integer programming solvers [13, 14], or finite model searching tools [15]. Such tools tend not to be competitive with the Euler–Parker approach if one only wants a pair of Latin squares and does not require them to be of any particular special form. However, in many problems involving orthogonal Latin squares it is not enough to construct an arbitrary pair—one wants to find a pair of orthogonal Latin squares with a special structure, and in such problems automated reasoning tools have proven to be extremely useful. In this setting, an issue with the Euler–Parker approach is that it requires an initial Latin square to start from. For small orders, a significant proportion of Latin squares have orthogonal mates. For example, McKay et al. [16] estimate that around 61% of Latin squares of order 10 have orthogonal mates. In practice, one can find Latin squares with a mate through random generation or by a theoretical construction. However, when one seeks a pair of a special form, in order to use Euler–Parker one must find a square with a mate but also with a mate *of the special form*. Such squares may be extremely rare and with no obvious approach for finding them. In such cases, automated reasoning tools like SAT solvers can be effective when the extra constraints defining the special form can be provided as additional constraints. That way the solver is able to direct its search towards orthogonal Latin squares of the specific form in question.

Despite the success of automated reasoning on many orthogonal Latin square problems (see

Section 2.4), automated reasoning tools have no knowledge of the Euler–Parker approach. A workaround is to encode into the problem instance that the symbols of one Latin square partition the other Latin square into disjoint transversals (see Section 3.1). This is essentially a “declarative” form of the Euler–Parker algorithm, and in fact this encoding performs decently well in practice [17]. The goal of our work is to show that it can be significantly beneficial to call an external implementation of the Euler–Parker algorithm during solving, rather than relying on the declarative Euler–Parker orthogonality constraints.

In our approach, when the solver finds a partial solution that represents a complete Latin square  $L$ , we call an implementation of the Euler–Parker algorithm to determine if  $L$  has an orthogonal mate. If not, we tell the solver to backtrack immediately. In contrast, a pure SAT approach would not distinguish between the square-finding and transversal-finding stages. When a Latin square  $L$  is found, the solver will continue to try to find  $n$  disjoint transversals of  $L$ , but there is no guarantee this step will run to exhaustion: if the solver encounters a conflict when looking for transversals, it will backtrack and in doing so may choose to explore a new Latin square instead rather than run the Euler–Parker approach to completion. Even worse, in general not all transversals will be explored, because the solver will typically have transversals already appearing in its partial solution. If the transversals in the partial solution do not extend to  $n$  disjoint transversals, then an orthogonal mate of  $L$  will not be found without backtracking, even when  $L$  has an orthogonal mate.

## 1.1. Benchmarks

As a case study for benchmarking our hybrid implementation combining SAT with a symbolic Euler–Parker implementation, we evaluate the performance of our solver on a collection of difficult orthogonal Latin square problems originating in the work of Myrvold [18] in 1999 (see Section 2.3 for more details). Myrvold was studying the notoriously difficult problem of finding a collection of three mutually orthogonal Latin squares of order 10, also known as a 3 MOLS(10). Despite attracting an immense amount of interest since at least the 1950s, the question of existence of a 3 MOLS(10) is still open and is considered to be an extremely difficult problem—Knuth lists it as a research problem of difficulty 50, his highest possible ranking [11, Ex. 15].

Myrvold considered the case of a 3 MOLS(10) in which one of the squares has a  $4 \times 4$  Latin subsquare and then considered the form that the other two squares (a 2 MOLS(10)) must take on. She showed that the 2 MOLS(10) must be one of twenty-eight possible forms, and she was able to eliminate twenty of the twenty-eight possibilities by theoretical arguments. However, the remaining eight cases proved to be more difficult and resisted any approach of ruling them out. For over twenty-five years it was unknown if some more complicated argumentation could eliminate them, or if 2 MOLS(10)s existed in some or all of the eight cases. This was the state of affairs until recently, when Bright, Keita, and Stevens [17] used a SAT solver to explicitly construct 2 MOLS(10)s in all eight of Myrvold’s remaining cases.

Bright, Keita, and Stevens [17] found examples in Myrvold’s eight cases using a pure SAT encoding and the off-the-shelf SAT solver Kissat [19]. Because SAT solvers inherently have an element of randomness in their running times, for the purposes of decreasing the amount of variance, each of the eight cases had 49 instances that were run independently, each with a different random seed. They report that the SAT solver took a median of around 580,000 seconds to solve the hardest of the eight cases, and 24 of the 49 instances in the hardest case timed out after one week without a solution being found.

In contrast, in this paper we implement a new hybrid approach that combines the SAT solver CADICAL [20] with an external call to the Euler–Parker algorithm (see Section 3.2). The hybrid approach significantly improves the running time of a pure SAT solver: our hybrid version of CADICAL solves the hardest of Myrvold’s eight cases in a median of about 5,100 seconds, over 100 times faster than Kissat. Note that the SAT solver Kissat (used by Bright, Keita, and Stevens) is more effective than CADICAL at solving these instances. Even still, our hybrid solver

(based on CADICAL) dramatically outperformed both Kissat and pure CADICAL, solving all the instances in the hardest case in around 12 hours, whereas Kissat only solved half of the hardest instances within a week and pure CADICAL solved none of the hardest instances within a week. See Section 4 for a summary of our results.

## 2. Background

In this section we give the background required for the rest of the paper. In particular, we outline satisfiability solving in Section 2.1, linear Diophantine system solving as used by the Euler–Parker approach in Section 2.2, and a description of Myrvold’s 2 MOLS(10) instances that we use for benchmarking in Section 2.3. We close the section with a description of related work using automated reasoning tools on orthogonal Latin square problems in Section 2.4.

### 2.1. Satisfiability solving

The Boolean satisfiability problem (SAT) asks whether a Boolean formula admits a truth assignment that makes the formula evaluate to true. It was the first problem shown to be NP-complete [21], and it remains a central problem in computational complexity. Perhaps surprisingly, SAT solving has also become a useful tool in computational discrete mathematics—the webpage [sat4math.com](http://sat4math.com) lists over 100 papers in which SAT solvers were used to solve problems in mathematics. This success stems from the fact that over the last several decades advances in SAT solving have produced highly effective general-purpose search tools [22]. Their search capabilities have proved useful in a wide range of areas, including discrete geometry [23, 24, 25], finite geometry [26], infinite graph theory [27], and puzzles in discrete mathematics [28].

Modern SAT solvers typically require the input formula to be given in *conjunctive normal form* (CNF), defined in terms of literals and clauses. A *literal* is a Boolean variable  $p$  or its negation  $\neg p$ . A *clause* is a disjunction of literals, and a formula is in *CNF* if it is a conjunction of clauses. For example,  $(p \vee q) \wedge p \wedge (\neg p \vee q)$  is a CNF formula; its second clause is a *unit clause* because it consists of one literal. We also use implication notation as shorthand, with  $p \rightarrow q$  denoting  $\neg p \vee q$ . A *satisfying assignment* is a truth assignment to the variables such that every clause evaluates to true.

To apply SAT solving to a search problem, one encodes the problem as a CNF formula so that the original problem has a solution if and only if the formula is satisfiable. A satisfying assignment can then be decoded into a solution of the original search problem. Conversely, if the formula is unsatisfiable, then the original search problem has no solution.

#### 2.1.1. Programmatic SAT with IPASIR-UP

CADICAL supports a programmatic interface based on IPASIR-UP [29], which allows user-defined callback functions to interact with the solver during search. In particular, the interface makes it possible to monitor assignments to selected variables, detect backtracking events, and add clauses dynamically based on information not encoded directly in CNF.

In our setting, the IPASIR-UP interface is used to detect when all variables encoding a Latin square have been assigned. At that point, we invoke an Euler–Parker implementation to determine whether the completed square admits an orthogonal mate. If the square does not have an orthogonal mate, the solver is then given a clause that blocks the square from being a solution and forces the solver to backtrack.

The principal IPASIR-UP callbacks used in our method are

- `notify_assignment()`, which reports when a variable has been assigned;
- `notify_backtrack()`, which reports when a backtrack occurs, causing some assigned variables to become unassigned; and

- `has_external_clause()`, which indicates that an externally generated clause is ready to be added to the solver. In our method, whenever the Euler–Parker algorithm fails to construct a mate of a Latin square  $L$ , we add a clause blocking  $L$  from being considered again.

### 2.1.2. Combining SAT solving and symbolic computation

Interfaces like IPASIR-UP are useful when part of the instance is more naturally handled outside the base SAT encoding. For example, one may invoke a computer algebra system (CAS) on a structured subproblem determined by the current partial assignment. This “SAT+CAS” paradigm has shown to be beneficial in a number of mathematical problems [30].

More generally, combinations of satisfiability checking (SC) and symbolic computation (SC) have led to a number of successes, resulting in what has become known as the “SC<sup>2</sup>” paradigm [31]. The paradigm had its inception around 2015, when Ábrahám [32] and Zulkoski et al. [33] proposed combining the fields of satisfiability checking and symbolic computation to solve mathematical problems. Both fields were mature with well-established methods for solving mathematical problems, but there had been limited interaction between the fields up to that point, despite the fields’ complementary strengths. Satisfiability checking is effective at general search and learning, while symbolic computation is effective at solving problems that admit clever solutions by exploiting mathematical structure. The SC<sup>2</sup> approach therefore excels at solving problems that can benefit from both effective search and exploiting mathematical structure.

The SC<sup>2</sup> workshop, now in its eleventh year, hosts work that combines satisfiability checking and symbolic computation. What is particularly notable is the sheer variety of problems that benefit from the SC<sup>2</sup> framework. For example, SC<sup>2</sup> approaches have been used on graph enumeration up to isomorphism [34, 35], factorizing integers with known bits of the factors [36, 37], searching for hash function collisions [38], and proving theorems in extremal combinatorics [39].

## 2.2. Linear Diophantine system solving

For a Latin square  $L$  of order  $n$ , an orthogonal mate exists if and only if the cells of  $L$  can be partitioned into  $n$  disjoint transversals. This leads naturally to an exact cover viewpoint: one seeks a collection of transversals covering each cell exactly once. In general, an *exact cover* of a set  $X$  is a family of sets  $S_1, \dots, S_N$  where  $\bigcup_{i=1}^N S_i = X$ , and the sets  $S_i$  are pairwise disjoint ( $S_i \cap S_j = \emptyset$  for all  $1 \leq i < j \leq N$ ). The corresponding decision problem is: given  $X$  and a family of sets  $S_1, \dots, S_N$ , does there exist a subfamily of the family of sets forming an exact cover of  $X$ ?

The Euler–Parker algorithm uses a two-stage approach in order to construct the orthogonal mates of a Latin square  $L$  of order  $n$ . First, all transversals of  $L$  are found. Each transversal is represented as a set of  $n$  cells. Second,  $n$  pairwise disjoint transversals of  $L$  are found. These  $n$  transversals form an exact cover of the set of all  $n^2$  cells  $\{(0, 0), (0, 1), \dots, (n - 1, n - 1)\}$ .

Our implementation of the Euler–Parker algorithm uses the library LIBEXACT [40] in order to solve the exact cover problems. LIBEXACT is a library that finds all solutions of the linear Diophantine system

$$Ax = \mathbf{b}, \quad \mathbf{0} \leq \mathbf{x} \leq \mathbf{u}$$

where  $A = [a_{ij}]$  is a specified  $M \times N$  matrix with entries in  $\{0, 1\}$ ,  $\mathbf{b}$  and  $\mathbf{u}$  are specified column vectors of positive integers, and the inequalities hold entrywise. A solution of the system is represented by the column vector  $\mathbf{x} = (x_1; \dots; x_N)$  of nonnegative integers. For the purposes of solving the exact cover problem described above, we take  $\mathbf{b}$  and  $\mathbf{u}$  to be  $\mathbf{1}$ , the vector of all 1s of appropriate dimension.

Combinatorially, each row of  $A$  represents a constraint (e.g., that a specific cell must appear in a cover) and each column of  $A$  represents a choice (e.g., if a transversal of  $L$  will be used in a cover or not). The column  $j$  covers row  $i$  precisely when  $a_{ij} = 1$ . The equation  $A\mathbf{x} = \mathbf{1}$  requires

that each row  $i$  be covered exactly once, while the bounds  $0 \leq x_j \leq 1$  require that each column  $j$  be used at most once. Each variable  $x_j$  records if the corresponding choice is to be used in a cover, with  $x_j$  corresponding to  $S_j$ , where  $S_1, \dots, S_N$  is the family of all possible sets that could appear in a cover. The set  $S_j$  will be used in the cover when  $x_j = 1$ , and the set  $S_j$  will not be used in the cover when  $x_j = 0$ .

### 2.3. Myrvold’s 2 MOLS(10) results

In 1999, Myrvold studied the problem of extending an orthogonal pair of  $10 \times 10$  Latin squares to an orthogonal triple under the assumption that one of the squares contains a  $4 \times 4$  Latin subsquare [18]. She derived a highly restrictive classification of the possible orthogonal pair types in this setting, ruling out twenty of twenty-eight cases. The remaining eight cases were left unresolved. Subsequent work by Bright et al. [17] showed that orthogonal pairs do exist in each of these eight cases, although the existence of an orthogonal triple remains unknown. Consequently, any argument excluding the remaining cases must necessarily involve the third square, rather than only the orthogonal pair.

The full details of Myrvold’s framework can be found in [17, 18], and we review only the aspects needed here. Instead of working directly with a 2 MOLS(10), Myrvold considers a *transversal representation pair* (TRP) of Latin squares  $(P, Q)$ . In such a pair, each row of  $Q$  represents a transversal of  $P$ , and the  $n$  rows of  $Q$  decompose the cells of  $P$  into  $n$  disjoint transversals. Thus  $P$  and  $Q$  are not themselves orthogonal mates, but they imply the existence of an orthogonal pair and it is straightforward to construct a 2 MOLS(10) from  $(P, Q)$  and vice versa.

Myrvold assumes that there exists a Latin square  $L$  whose transversals are represented by both  $P$  and  $Q$ . The square  $L$  is taken to contain a  $4 \times 4$  Latin subsquare  $\Omega$  in its lower-right corner, using the symbols  $\{0, 1, 2, 3\}$ . This assumption induces a colouring on the entries of  $P$  and  $Q$ . Entries containing symbols in  $\{0, 1, 2, 3\}$  are called *white*. Among the entries in the first six columns, exactly two nonwhite entries from each column are coloured *dark*, and the remaining nonwhite entries are called *light*.

There are four transversal types, denoted  $p_1, \dots, p_4$ , classified by the number of white entries in the last four columns. A transversal of type  $p_i$  contains exactly  $i$  white entries in its last four columns and exactly  $2i - 2$  dark entries. Myrvold studied the possible ways in which  $L$  could be decomposed into 10 transversals and showed that there were exactly seven general types (determined by the number of transversals of each type  $p_1, \dots, p_4$ ) which she denoted R, S, T, U, V, W, and X. For example, a decomposition of  $L$  using eight transversals of type  $p_1$  and two transversals of type  $p_4$  was said to be of type R.

Under the assumption that  $L$  belongs to an orthogonal triple and therefore is decomposable into 10 disjoint transversals in two ways, Myrvold showed that the only pairs of potentially feasible decomposition types for  $L$  were UU, SX, UW, WW, VX, UX, WX, and XX. She was able to theoretically rule out each of the other twenty possible pairings (e.g.,  $L$  cannot be decomposed into R and S simultaneously), although she was unable to confirm or rule out any of the pairings UU, SX, UW, WW, VX, UX, WX, and XX. It turns out this is because transversal representation pairs in these eight cases actually exist. An example of a type-UW transversal representation pair is shown in Figure 2. Note that  $(P, Q)$  does *not* extend to a triple  $(P, Q, L)$ , and it appears to be a very difficult problem to find such a triple. The most we can say is that there exists a  $4 \times 4$  Latin subsquare  $\Omega$  of  $L$  compatible with  $(P, Q)$ , because the SAT encoding used to generate the pair included “ $\Omega$  compatibility” constraints (see [17, Sec. 4.4]).

### 2.4. Related work

Orthogonal Latin square problems have been approached with several optimization and search methods. Appa, Mourtos, and Magos [41] investigated the orthogonal Latin squares problem using a combination of integer and constraint programming. Appa, Magos, and Mourtos [42, 43]

|   |   |   |   |   |   |   |   |   |   |
|---|---|---|---|---|---|---|---|---|---|
| 0 | 2 | 3 | 4 | 5 | 6 | 1 | 7 | 8 | 9 |
| 1 | 0 | 4 | 9 | 7 | 2 | 6 | 5 | 3 | 8 |
| 3 | 9 | 8 | 2 | 4 | 1 | 0 | 6 | 7 | 5 |
| 4 | 1 | 7 | 0 | 3 | 5 | 9 | 8 | 6 | 2 |
| 5 | 8 | 0 | 1 | 2 | 7 | 3 | 9 | 4 | 6 |
| 8 | 3 | 1 | 6 | 9 | 0 | 4 | 2 | 5 | 7 |
| 2 | 7 | 6 | 8 | 0 | 9 | 5 | 3 | 1 | 4 |
| 9 | 6 | 2 | 5 | 1 | 8 | 7 | 4 | 0 | 3 |
| 6 | 4 | 5 | 7 | 8 | 3 | 2 | 1 | 9 | 0 |
| 7 | 5 | 9 | 3 | 6 | 4 | 8 | 0 | 2 | 1 |

|   |   |   |   |   |   |   |   |   |   |
|---|---|---|---|---|---|---|---|---|---|
| 0 | 1 | 8 | 3 | 7 | 9 | 2 | 4 | 5 | 6 |
| 1 | 6 | 7 | 4 | 2 | 0 | 8 | 3 | 9 | 5 |
| 2 | 5 | 4 | 6 | 3 | 1 | 7 | 9 | 8 | 0 |
| 3 | 0 | 2 | 7 | 9 | 6 | 5 | 8 | 4 | 1 |
| 5 | 9 | 6 | 0 | 1 | 3 | 4 | 7 | 2 | 8 |
| 6 | 8 | 3 | 5 | 4 | 2 | 9 | 0 | 1 | 7 |
| 7 | 2 | 1 | 9 | 8 | 5 | 3 | 6 | 0 | 4 |
| 8 | 4 | 9 | 2 | 0 | 7 | 1 | 5 | 6 | 3 |
| 9 | 3 | 0 | 8 | 5 | 4 | 6 | 1 | 7 | 2 |
| 4 | 7 | 5 | 1 | 6 | 8 | 0 | 2 | 3 | 9 |

**Figure 2:** A type-UW transversal representation pair, with white, light, and dark colouring. Each row of each square represents a transversal of the other square. For example, consider the first row of the first square  $[0, 2, 3, 4, \dots]$ . Going through the second square column-by-column and extracting the cell from each column containing these symbols in sequence yields the cells  $(0, 0)$ ,  $(6, 1)$ ,  $(5, 2)$ ,  $(1, 3)$ ,  $\dots$  which form a transversal of the second square (the cells highlighted with a red border).

later used integer programming both to search for  $2\text{MOLS}(n)$  and to prove the nonexistence of  $2\text{MOLS}(6)$ . In later work, they studied  $2\text{MOLS}(n)$  and  $3\text{MOLS}(n)$  more generally [13]. Rubin, Bright, Cheung, and Stevens [14] compared integer and constraint programming encodings for  $2\text{MOLS}(n)$ , and found that an indexing-based constraint programming formulation performed well on orders  $5 \leq n \leq 12$  and  $3\text{MOLS}(n)$  for  $n \leq 9$ . Ma and Zhang [44] applied finite model generation to the search for  $2\text{MOLS}(n)$ .

SAT-based methods have also been used extensively for orthogonal Latin square variants. Zaikin, Kochemazov, and Semenov [45] compared several SAT encodings for systems of orthogonal Latin squares and diagonal orthogonal Latin squares. Zaikin and Kochemazov [46] (and later in conjunction with Zhuravlev and Vatutin [47]) studied diagonal Latin square systems and searched for pairs and triples of mutually orthogonal diagonal Latin squares of order 10 using SAT-based methods and volunteer computing. Vatutin, Zaikin, Manzyuk, and Nikitina [48] later applied a cube-and-conquer approach to the search for orthogonal Latin squares.

Lu, Liu, and Zhang [49] searched for doubly self-orthogonal Latin squares using SAT, constraint programming, and also developed a custom exhaustive search algorithm based on exact cover. Falc3n, Falc3n, and N3n3ez [50] computed totally symmetric and totally conjugate orthogonal partial Latin squares using a SAT solver. Huang, Liu, Ge, Ma, and Zhang [51] investigated orthogonal golf designs via satisfiability testing. Jin, Lv, Ge, Ma, and Zhang [52] used SAT to study Costas Latin squares and orthogonal Costas Latin squares. More broadly, Zhang [53] surveys Latin square problems as a natural application area for SAT encodings.

Zaikin, Vatutin, and Bright used a SAT solver with the same orthogonality encoding that we use in this paper to enumerate all extended self-orthogonal diagonal Latin squares of orders up to 10 [54]. Bright, Keita, and Stevens also used this orthogonality encoding to re-enumerate all orthogonal pairs of Latin squares of order 10 whose associated nets have at least two nontrivial relations [55].

### 3. A Hybrid SAT + Euler–Parker Approach

In this section, we describe our hybrid SAT + Euler–Parker approach. First, we describe the basic SAT encoding in Section 3.1. Next, we describe how we implement the Euler–Parker algorithm using the LIBEXACT library in Section 3.2, and also describe how the Euler–Parker algorithm is programmatically called while the SAT solver is solving. Lastly, we describe how to specialize

our hybrid approach specifically to the 2MOLS(10) cases considered by Myrvold in Section 3.3.

### 3.1. SAT encoding

First, we describe the SAT encoding which is used to ensure an  $n \times n$  matrix  $P$  is a Latin square. To represent the entries of  $P$ , we introduce variables  $P_{i,j,k}$  for  $0 \leq i, j, k < n$  denoting that cell  $(i, j)$  of  $P$  contains symbol  $k$ . To ensure that  $P$  is a Latin square, we impose *exactly-one* constraints requiring that each cell contains exactly one symbol and that each symbol appears exactly once in each row and column. A simple way to encode such constraints is the pairwise encoding, which combines one *at-least-one* clause of length  $n$  with  $n(n-1)/2$  *at-most-one* clauses of length 2. For example, the condition that cell  $(i, j)$  of  $P$  contains exactly one symbol can be written as

$$\bigvee_{0 \leq k < n} P_{i,j,k} \quad (\text{at least one symbol}) \quad \text{and} \quad \bigwedge_{0 \leq k < l < n} (\neg P_{i,j,k} \vee \neg P_{i,j,l}) \quad (\text{at most one symbol}).$$

However, Bright et al.’s [17] experiments found that replacing the pairwise exactly-one encoding with a cardinality encoding based on the “totalizer” encoding [56] performed better in practice. We denote a cardinality constraint which says that exactly  $k$  of the variables in a set  $S$  are true by the notation  $\sum_{s \in S} s = k$ . Enforcing that cell  $(i, j)$  of  $P$  has exactly one symbol can then be done with  $\sum_{0 \leq k < n} P_{i,j,k} = 1$ .

Orthogonality of two squares  $P$  and  $R$  can be encoded directly into the SAT instance, but a naive encoding uses  $\Theta(n^6)$  clauses of length 4 [55, Eq. 4]. In practice, a better way of enforcing the orthogonality of  $P$  and  $R$  is to instead assert that for all symbols  $0 \leq k < n$  the cells containing the symbol  $k$  in  $R$  form a transversal in  $P$ . In order to do this, we form a new square  $Q$  whose  $k$ th row stores the symbols of  $P$  in the transversal formed by the cells with symbol  $k$  in  $R$ . This is done with the constraints

$$(R_{i,j,k} \wedge P_{i,j,l}) \rightarrow Q_{k,j,l} \quad \text{for all } 0 \leq i, j, k, l < n.$$

Once  $Q$  has been introduced, asserting that the cells with symbol  $k$  in  $R$  form a transversal of  $P$  is done by asserting that row  $k$  of  $Q$  contains each symbol exactly once. By construction, the  $n$  transversals formed in this way will necessarily be disjoint because  $R$  is a Latin square. Thus, the constraints on  $Q$  imply that there exist  $n$  disjoint transversals in  $P$ , and the Euler–Parker method gives that  $R$  is an orthogonal mate of  $P$ . Additionally, Bright et al. derive additional structure on the squares  $P$ ,  $Q$ , and  $R$ , including defining a notion of a “product” square for which  $P = QR$ . They prove this product equation is equivalent to  $(P, R)$  being an orthogonal pair of Latin squares and  $(P, Q)$  being a transversal representation pair of Latin squares (see [17]).

### 3.2. Programmatic Euler–Parker

The orthogonality encoding described in Section 3.1 can be considered a “static” implementation of the Euler–Parker approach. As described in Section 1, this has limitations when used with a standard backtracking SAT solver. This motivated us to augment a SAT solver with a programmatic implementation of the Euler–Parker algorithm that can be called whenever the SAT solver finds a complete Latin square during the search. If the complete square is determined to not have an orthogonal mate by Euler–Parker, the solver will immediately backtrack (which is earlier than it would have otherwise). We now describe how we implemented the two stages of the Euler–Parker method using two calls to the library LIBEXACT. In the following, say  $L$  is a Latin square of order  $n$  and we want to compute all orthogonal mates of  $L$ .

#### 3.2.1. Stage 1: Construct all transversals

A transversal of  $L$  is a set of  $n$  cells containing exactly one cell from each row, exactly one cell from each column, and with  $L$  containing each symbol exactly once amongst the transversal’s

cells. Constructing a transversal naturally reduces to solving a linear Diophantine system where a solution of the system provides a transversal  $t$  of  $L$ . Let  $x_{i,j} \in \{0, 1\}$  be a variable that will be 1 when cell  $(i, j)$  appears in  $t$ . Since  $t$  contains exactly one cell from row  $i$ , we have  $\sum_{j=0}^{n-1} x_{i,j} = 1$ , since  $t$  contains exactly one cell from column  $j$ , we have  $\sum_{i=0}^{n-1} x_{i,j} = 1$ , and since  $t$  contains an occurrence of symbol  $k$  in  $L$  exactly once we have  $\sum_{L[i,j]=k} x_{i,j} = 1$ .

Since there are  $n$  rows,  $n$  columns, and  $n$  symbols, this provides a linear Diophantine system with  $n^2$  variables and  $3n$  equations. Given  $L$ , our implementation symbolically forms the above system and uses LIBEXACT to find all integer 0–1 solutions of the system. Each solution is transformed into a transversal and the input to the second stage is the set of all transversals of  $L$ .

### 3.2.2. Stage 2: Construct $n$ disjoint transversals

Given a set  $T$  of transversals of a Latin square  $L$ , constructing a set of  $n$  disjoint transversals also naturally reduces to solving a linear Diophantine system, where a solution of the system provides  $n$  disjoint transversals  $T^* \subseteq T$ . Let  $x_t \in \{0, 1\}$  be a variable that will be 1 when transversal  $t \in T$  is one of the  $n$  transversals in  $T^*$ . Since every cell  $(i, j)$  appears in exactly one transversal in  $T^*$  (because the transversals are disjoint and there are  $n$  of them), we have the equations

$$\sum_{t \in T, (i,j) \in t} x_t = 1 \quad \text{for every cell } (i, j).$$

This is a linear system with  $|T|$  variables and  $n^2$  equations. Our implementation symbolically forms this system and uses LIBEXACT to solve it exhaustively. Each solution has exactly  $n$  variables with  $x_t = 1$ , providing  $n$  disjoint transversals. Denoting these  $n$  transversals as  $t_0, \dots, t_{n-1}$ , we create an orthogonal mate to  $L$  by assigning symbol  $i$  to the cells of  $t_i$  for all  $0 \leq i < n$ .

### 3.2.3. Injecting Euler–Parker into the SAT solver

The hybrid aspect of the method comes from using IPASIR-UP to call LIBEXACT to perform the above two-stage process during SAT search. Using the callback routines described in Section 2.1.1, the interface tracks the assignment status of the variables encoding a Latin square. Once all variables encoding a square have been assigned, the interface takes the complete square (say  $L$ ) and begins stage 1. If fewer than  $n$  transversals are found, then  $L$  cannot admit an orthogonal mate. In this case, the interface constructs a blocking clause from the literals encoding  $L$  and returns it to the SAT solver. This forces the solver to backtrack and prevents  $L$  from being considered again.

If at least  $n$  transversals are found, they are passed to stage 2, which attempts to select a disjoint set of  $n$  transversals. If this second stage fails, the complete square  $L$  is again blocked. If it succeeds, then an orthogonal mate has been found. If there happen to be additional constraints on the orthogonal mate, they can be checked at this point. Alternatively, it may be possible to run the Euler–Parker algorithm in such a way that it will *always* construct orthogonal mates that satisfy the additional constraints. For example, in Section 3.3 we show how to do this for Myrvold’s 2MOLS(10) instances.

When the values in square  $P$  have been completely assigned to form a square  $L$  which has been determined to have no orthogonal mates, the blocking clause passed to the solver is

$$\bigvee_{L[i,j]=k} \neg P_{i,j,k}$$

which naively contains  $n^2$  literals, since  $i$  and  $j$  both range over  $\{0, \dots, n-1\}$ . In fact, it is sufficient to take  $i$  and  $j$  only over  $\{0, \dots, n-2\}$ , since the upper-left  $(n-1) \times (n-1)$  subsquare of  $L$  uniquely determines the final row and column. This shrinks the length of the blocking clause to  $(n-1)^2$  literals.

### 3.3. Specialization to Myrvold’s 2 MOLS(10) instances

The encoding described so far is for the general 2 MOLS( $n$ ) problem. Our experimentation of this general encoding as  $n$  grows reveals the augmented SAT solver dramatically outperforms a pure SAT solver (see Section 4). While this is a promising sign of the usefulness of the hybrid approach, the general 2 MOLS( $n$ ) problem can be solved without a SAT solver and so this is a somewhat artificial benchmark.

To highlight the utility of our hybrid approach, we use it to solve 2 MOLS(10) instances with structure that seems intractable to handle purely theoretically. The eight 2 MOLS(10) cases left unsolved by Myrvold (see Section 2.3) are suitable for this purpose, since she was unable to solve the eight cases by theoretical arguments. Currently, the cases have only been solved with the assistance of a SAT solver [17]. In order to enforce the additional structure considered by Myrvold, we add more constraints on the pair  $(P, Q)$  from Section 3.1. For example, there are colour constraints, subsquare  $\Omega$  compatibility constraints, and symmetry breaking constraints. There are also constraints stating that the pair  $(P, Q)$  is one of Myrvold’s eight unsolved types. For example,  $(P, Q)$  being of type XX means that both  $P$  and  $Q$  contain 4 transversals of type  $p_1$  and six transversals of type  $p_2$ . A transversal is of type  $p_i$  when it contains exactly  $i$  white entries in the last four columns, so the type constraints can be implemented by adding appropriate cardinality constraints into the SAT instance. Because these constraints are not an essential component of our new hybrid approach, we do not go into them in more detail and instead refer the reader to [17, Sec. 4].

The first stage of the Euler–Parker implementation does not need modification to work on the instances encoding Myrvold’s problem. However, after the first stage, some transversals can be discarded because they cannot be part of a valid solution. For example, in the case of XX, we are looking to find 10 disjoint transversals of the first square: 4 of type  $p_1$  and 6 of type  $p_2$ . So, any transversals of type  $p_3$  found by stage 1 will be discarded before starting stage 2. Similarly, transversals that violate the  $\Omega$  compatibility constraints will also be discarded. Without loss of generality, there are two possibilities for the  $\Omega$  subsquare ( $\Omega_1$  and  $\Omega_2$ ), and the SAT instance permits solutions from either possibility (so long as solutions only contain transversals all compatible with  $\Omega_1$  or all compatible with  $\Omega_2$ ). In order to enforce this in our hybrid approach, we run the second Euler–Parker stage twice, once only with transversals compatible with  $\Omega_1$ , and once only with transversals compatible with  $\Omega_2$ .

In the second stage of the Euler–Parker implementation we make use of Myrvold’s constraints explicitly. For example, in case XX we can specify that a valid solution requires exactly 4 transversals of type  $p_1$  and 6 transversals of type  $p_2$ . Letting  $T_i \subseteq T$  denote the set of transversals of type  $p_i$ , we form the equations  $\sum_{t \in T_1} x_t = 4$  and  $\sum_{t \in T_2} x_t = 6$  and pass them to LIBEXACT. There are also colour constraints that can be directly incorporated: a valid solution must possess exactly two dark entries in each of the first six columns. Let  $D_i \subseteq T$  denote the transversals whose entry in the  $i$ th column is coloured dark. We form the six equations  $\sum_{t \in D_i} x_t = 2$  for  $0 \leq i < 6$  and pass them all to the linear system provided to LIBEXACT.

## 4. Results

In this section, we evaluate the performance of our hybrid SAT + Euler–Parker method. As a first experiment, we test its performance on the unrestricted 2 MOLS(10) problem (see Section 4.1). Our primary experiment is solving each of the eight 2 MOLS(10) cases that Myrvold left unsolved (see Section 4.2). All instances were solved using the SAT solver CADICAL version 3.0, and the Euler–Parker implementation used LIBEXACT version 1.0. Experiments were run on the Digital Research Alliance of Canada’s Fir cluster, a high-performance computing cluster consisting of AMD EPYC processors, most of which run at 2.7 GHz. Our source code is available at [https://github.com/aaronbarnoff/MOLS\\_EP](https://github.com/aaronbarnoff/MOLS_EP).

## 4.1. Results for 2 MOLS( $n$ ) instances

As pointed out in Section 1, the unrestricted 2 MOLS( $n$ ) problem has been solved by purely theoretical constructions, and therefore one does not need to use a SAT solver for this problem. Even still, it forms a useful experiment to demonstrate that providing SAT solvers with static Euler–Parker constraints (i.e., constraints asserting the symbols of one square decompose the other square into  $n$  disjoint transversals) may not be as effective as running an external Euler–Parker implementation.

Table 1 provides a detailed timing breakdown of our unrestricted 2 MOLS(10) results for  $8 \leq n \leq 12$ . The table shows that as  $n$  grows the hybrid SAT + Euler–Parker approach is orders of magnitude faster than a standard SAT solver provided with static Euler–Parker constraints. For example, for  $n = 12$ , the median time across 15 tests for the hybrid approach is at least 150,000 times faster than the median time for the pure SAT approach (for which the solver did not finish within four days).

**Table 1**

Runtime breakdown for the 2 MOLS( $n$ ) experiments, showing median times in seconds for the pure SAT and the hybrid SAT + Euler–Parker approach across 15 random instances. The median number of Euler–Parker calls before an orthogonal mate was found is also provided. The timeout was four days.

| Order $n$  | 8    | 9      | 10      | 11       | 12      |
|------------|------|--------|---------|----------|---------|
| Pure SAT   | 0.09 | 168.19 | 3486.03 | 39383.54 | timeout |
| Hybrid     | 0.26 | 0.42   | 0.78    | 1.20     | 2.23    |
| # EP calls | 160  | 31     | 2       | 1        | 1       |

These results conclusively demonstrate that off-the-shelf SAT solvers do not exploit the Euler–Parker approach to its full potential. As  $n$  increased, our experimentation showed that the static Euler–Parker constraints made the SAT solver find a square with an orthogonal mate quickly. For  $n \geq 11$ , in the median case the *very first square* the solver found was one having an orthogonal mate. Despite this, the pure SAT solver was unable to find its mate, evidently backtracking before finding  $n$  disjoint transversals in the square. Of course, the pure SAT solver is only provided static constraints telling it to find  $n$  disjoint transversals. The solver does not realize that it makes sense to first find *all* transversals and then from within those try to find  $n$  disjoint transversals. This two-stage methodology is exploited by the hybrid approach and explains why it can dramatically outperform the pure SAT approach.

These instances had CADICAL’s `--shufflerandom` option enabled, which randomizes the initial variable scores and causes the solver to branch on different variables in each experiment. We noticed that otherwise the hybrid solver would always make the same decisions and find the same solutions every time, even when run using different random seeds. This seems to be an artifact of the fact that the hybrid solver is so fast it solves the instances before the impact of the random seed is felt. The hybrid solver also had CADICAL’s `--preprocesslight` disabled because otherwise in some orders the solver would find a solution instantly as a result of the preprocessing, leading to the same decisions being made every time, even with variable shuffling enabled.

## 4.2. Results for Myrvold’s 2 MOLS(10) instances


We now describe the results of our hybrid SAT + Euler–Parker approach on Myrvold’s eight unresolved cases, and compare them with the pure SAT approach. For both approaches, we ran 15 instances of each of the eight cases for seven days, using a different random seed for each instance so that CADICAL would make different decisions each time. These results are summarized in Table 2, with a visual comparison shown in Figure 3.

We found that the best performance was obtained by limiting how often a call to Euler–Parker could be invoked. This was due to how quickly CADICAL produces completed squares, each

**Table 2**

Comparison of the hybrid and pure-SAT approaches for each of the eight Myrvold pair types. Times are in seconds. Each pair type had 15 randomly seeded instances that ran for one week.

| pair type | method   | # solved | median     | min        | max        |
|-----------|----------|----------|------------|------------|------------|
| UU        | Hybrid   | 15       | 5077.36    | 509.92     | 38 323.47  |
|           | Pure SAT | 15       | 31 721.01  | 2706.00    | 366 291.48 |
| SX        | Hybrid   | 15       | 2882.27    | 499.79     | 22 652.56  |
|           | Pure SAT | 14       | 59 447.47  | 255.08     | timeout    |
| UW        | Hybrid   | 15       | 11 304.32  | 903.01     | 147 568.43 |
|           | Pure SAT | 12       | 116 856.78 | 751.59     | timeout    |
| WW        | Hybrid   | 15       | 35 638.79  | 1360.37    | 412 264.03 |
|           | Pure SAT | 6        | timeout    | 66 699.78  | timeout    |
| VX        | Hybrid   | 15       | 20 789.28  | 484.36     | 62 917.17  |
|           | Pure SAT | 13       | 174 215.01 | 46 195.26  | timeout    |
| UX        | Hybrid   | 15       | 27 033.18  | 205.98     | 68 228.84  |
|           | Pure SAT | 13       | 243 781.81 | 9314.35    | timeout    |
| WX        | Hybrid   | 15       | 20 612.80  | 491.44     | 53 664.08  |
|           | Pure SAT | 3        | timeout    | 279 253.68 | timeout    |
| XX        | Hybrid   | 15       | 5105.64    | 491.25     | 44 225.28  |
|           | Pure SAT | 0        | timeout    | timeout    | timeout    |



**Figure 3:** Scatterplot comparing the hybrid approach (left) to the pure SAT approach (right), for the eight Myrvold pair types. The median time is indicated with a black bar.

of which makes a relatively expensive call to `LIBEXACT` in the first exact cover stage. When the call to Euler–Parker was unrestricted, it was called so often that the SAT solver had little opportunity to make progress of its own through conflict-driven search, as reflected in a low number of internal (non-programmatic) conflicts and restarts. The blocking clauses generated by Euler–Parker were added to the solver as forgettable clauses (meaning that the solver was allowed to discard them to keep memory usage low), but `CADICAL` only performs reductions after a certain number of non-programmatic conflicts. Since programmatic blocking clauses were being created so rapidly, memory usage increased steadily and typically exceeded 8 GiB after running for a few days.

To address this, we configured the solver so that the solver generates at least one non-programmatic conflict clause before calling Euler–Parker again. This threshold gave the SAT

solver more time to progress through the search space between Euler–Parker invocations. This resulted in far fewer squares being passed to Euler–Parker between clause database reductions, and this substantially reduced memory usage. The highest recorded memory usage by the hybrid solver was 0.99 GiB (after running for about 5 days). A breakdown of the median time spent in the SAT solver and in each stage of the Euler–Parker algorithm is shown in Table 3.

**Table 3**

Runtime breakdown for the eight Myrvold cases, showing median times in seconds before a 2 MOLS(10) was found. The total runtime is divided into the SAT solving time and the time for the two stages of the Euler–Parker (EP) method. The median number of calls to Euler–Parker is also provided.

| pair type  | UU        | SX      | UW        | WW        | VX        | UX        | WX        | XX        |
|------------|-----------|---------|-----------|-----------|-----------|-----------|-----------|-----------|
| SAT solver | 2036.67   | 960.89  | 6788.22   | 12 107.96 | 6847.76   | 8957.87   | 5366.88   | 1095.49   |
| EP stage 1 | 2698.06   | 1753.94 | 3929.43   | 20 634.07 | 12 179.40 | 15 456.45 | 12 957.77 | 3274.35   |
| EP stage 2 | 372.91    | 173.16  | 586.67    | 3116.00   | 1762.12   | 2618.85   | 2288.15   | 625.46    |
| # EP calls | 1 253 736 | 677 695 | 1 565 150 | 8 148 047 | 4 688 434 | 5 974 377 | 4 928 770 | 1 254 684 |

For these instances, we call Euler–Parker whenever the SAT solver finds a complete assignment of either the variables defining the square  $P$  or the square  $Q$ . When  $P$  has been determined, we run Euler–Parker to construct  $Q$ , and when  $Q$  has been determined, we run Euler–Parker to construct  $P$ . Without loss of generality, the first row of  $P$  can be taken to be one of three possibilities [17, Thm. 13], and the SAT encoding enforces the first row of  $P$  is one of these three possibilities. In cases when  $Q$  has been determined, it is also possible to add constraints into the second LIBEXACT instance enforcing that one of the three possibilities appears in the transversal decomposition of  $Q$ . However, it was more efficient to *not* provide this symmetry breaking constraint to LIBEXACT, thereby making the stage 2 instance less constrained. The downside is that it makes the LIBEXACT call less efficient. However, in general it results in fewer calls to Euler–Parker, as it allows the solver to find solutions that it otherwise would discard.

## 5. Conclusion

In this paper, we proposed and implemented a novel strategy for constructing orthogonal Latin squares by incorporating both a SAT solver and the Euler–Parker strategy for orthogonal mate construction. Our approach exploits both the powerful search-with-learning capabilities of a modern SAT solver, as well as the rich mathematical structure present in orthogonal Latin squares. We use our new strategy to solve a number of challenging orthogonal Latin square problems arising in the work of Myrvold [18]. We stress that our method crucially exploits *both* a SAT solver and the mathematical structure used by Euler [1] and Parker [10]. The SAT solver searches for an initial starting square with an orthogonal mate having structure concisely expressed in conjunctive normal form. Conversely, the Euler–Parker algorithm is used to efficiently construct an orthogonal mate from a starting square. Note that the starting square needs to be found before Euler–Parker can be used, and in many cases, including in the work of Myrvold, it is not clear how to find a square having an orthogonal mate with the proper structure.

As a case study, we build on the work of Bright, Keita, and Stevens [17], who recently reduced to SAT eight 2 MOLS(10) cases left open by Myrvold. The reduction to SAT was key to resolving these cases, as they had been left unsolved for over 25 years. Despite the success in using an off-the-shelf SAT solver to resolve these cases, we show our hybrid SAT + Euler–Parker approach dramatically outperforms the pure SAT approach. On the hardest of the eight cases, the hybrid solver typically takes 1.4 hours to solve what the pure SAT approach does not solve in a week. We believe our approach will also be useful in searching for other kinds of orthogonal Latin squares, and this will be the subject of future work.

## References

- [1] L. Euler, Recherches sur une nouvelle espèce de quarrés magiques, Verhandelingen uitgegeven door het zeeuwsch Genootschap der Wetenschappen te Vlissingen (1782) 85–239.
- [2] G. Tarry, Le problème des 36 officiers, Association Française pour l’Avancement des Sciences: Compte Rendu de la 29<sup>me</sup> session, Paris 1900 2 (1901) 170–203.
- [3] J. Peterson, Les 36 officiers, Annuaire des Mathématiciens (1902) 413–427.
- [4] P. Wernicke, Das problem der 36 offiziere, Jahresbericht der Deutschen Mathematiker-Vereinigung 19 (1910) 264–267.
- [5] H. F. MacNeish, Euler squares, Annals of Mathematics 23 (1922) 221–227. doi:10.2307/1967920.
- [6] E. T. Parker, Orthogonal Latin squares, Proceedings of the National Academy of Sciences 45 (1959) 859–862. doi:10.1073/pnas.45.6.859.
- [7] R. C. Bose, S. S. Shrikhande, On the construction of sets of mutually orthogonal Latin squares and the falsity of a conjecture of Euler, Transactions of the American Mathematical Society 95 (1960) 191–209. doi:10.2307/1993286.
- [8] R. C. Bose, S. S. Shrikhande, E. T. Parker, Further results on the construction of mutually orthogonal Latin squares and the falsity of Euler’s conjecture, Canadian Journal of Mathematics 12 (1960) 189–203. doi:10.4153/CJM-1960-016-5.
- [9] L. J. Paige, C. B. Tompkins, The size of the  $10 \times 10$  orthogonal Latin square problem, Combinatorial Analysis (1960) 71–83. doi:10.1090/psapm/010/0115927.
- [10] E. T. Parker, A computer search for Latin squares orthogonal to Latin squares of order ten, abstract 564-71, American Mathematical Society Notices 6 (1959) 798.
- [11] D. E. Knuth, The Art of Computer Programming, Volume 4A: Combinatorial Algorithms, Part 1, Addison-Wesley, Upper Saddle River, NJ, 2011.
- [12] S. Kochemazov, O. Zaikin, A. Semenov, The comparison of different SAT encodings for the problem of search for systems of orthogonal Latin squares, in: International Conference Mathematical and Information Technologies-MIT, volume 1839 of *CEUR Workshop Proceedings*, 2016, pp. 155–165. URL: <https://ceur-ws.org/Vol-1839/MIT2016-p14.pdf>.
- [13] G. Appa, D. Magos, I. Mourtos, Searching for mutually orthogonal Latin squares via integer and constraint programming, European Journal of Operational Research 173 (2006) 519–530. doi:10.1016/j.ejor.2005.01.048.
- [14] N. Rubin, C. Bright, K. Cheung, B. Stevens, Improving integer and constraint programming for Graeco-Latin squares, in: 2021 IEEE 33rd International Conference on Tools with Artificial Intelligence (ICTAI), IEEE, 2021, p. 604–608. doi:10.1109/ictai52525.2021.00096.
- [15] F. Ma, J. Zhang, Finding orthogonal Latin squares using finite model searching tools, Science China Information Sciences 56 (2011) 1–9. doi:10.1007/s11432-011-4343-3.
- [16] B. D. McKay, A. Meynert, W. Myrvold, Small Latin squares, quasigroups, and loops, Journal of Combinatorial Designs 15 (2006) 98–119. doi:10.1002/jcd.20105.
- [17] C. Bright, A. Keita, B. Stevens, Myrvold’s results on orthogonal triples of  $10 \times 10$  Latin squares: A SAT investigation, The Electronic Journal of Combinatorics 33 (2026) P1.30. doi:10.37236/13960.
- [18] W. Myrvold, Negative results for orthogonal triples of Latin squares of order 10, Journal of Combinatorial Mathematics and Combinatorial Computing 29 (1999) 95–105. URL: <https://combinatorialpress.com/jcmcc-articles/volume-029/negative-results-for-orthogonal-triples-of-Latin-squares-of-order-10/>.
- [19] A. Biere, T. Faller, K. Fazekas, M. Fleury, N. Froleyks, F. Pollitt, CaDiCaL, Gimsatul, IsaSAT and Kissat entering the SAT Competition 2024, in: M. Heule, M. Iser, M. Järvisalo, M. Suda (Eds.), Proc. of SAT Competition 2024 – Solver, Benchmark and Proof Checker Descriptions, volume B-2024-1 of *Department of Computer Science Report Series B*, University of Helsinki, 2024, pp. 8–10.

- [20] A. Biere, T. Faller, K. Fazekas, M. Fleury, N. Froleyks, F. Pollitt, CaDiCaL 2.0, in: A. Gurfinkel, V. Ganesh (Eds.), *Computer Aided Verification - 36th International Conference, CAV 2024, Montreal, QC, Canada, July 24-27, 2024, Proceedings, Part I*, volume 14681 of *Lecture Notes in Computer Science*, Springer, 2024, pp. 133–152. doi:10.1007/978-3-031-65627-9\_7.
- [21] S. A. Cook, The complexity of theorem-proving procedures, in: *Proceedings of the third annual ACM symposium on Theory of computing - STOC '71, STOC '71*, ACM Press, 1971, p. 151–158. doi:10.1145/800157.805047.
- [22] V. Ganesh, M. Y. Vardi, *On the Unreasonable Effectiveness of SAT Solvers*, Cambridge University Press, 2020, p. 547–566. doi:10.1017/9781108637435.032.
- [23] M. J. H. Heule, M. Scheucher, Happy ending: An empty hexagon in every set of 30 points, in: B. Finkbeiner, L. Kovács (Eds.), *30th International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, volume 14570 of *Lecture Notes in Computer Science*, Springer Nature Switzerland, Cham, 2024, p. 61–80. doi:10.1007/978-3-031-57246-3\_5.
- [24] B. Subercaseaux, E. Mackey, L. Qian, M. Heule, Automated symmetric constructions in discrete geometry, in: V. de Paiva, P. Koepke (Eds.), *Intelligent Computer Mathematics*, volume 16136 of *Lecture Notes in Computer Science*, Springer Nature Switzerland, 2025, p. 29–47. doi:10.1007/978-3-032-07021-0\_3.
- [25] B. Subercaseaux, J. Mackey, M. J. H. Heule, R. Martins, Automated mathematical discovery and verification: Minimizing pentagons in the plane, in: A. Kohlhase, L. Kovács (Eds.), *Intelligent Computer Mathematics*, volume 14960 of *Lecture Notes in Computer Science*, Springer Nature Switzerland, Cham, 2024, p. 21–41. doi:10.1007/978-3-031-66997-2\_2.
- [26] C. Bright, K. K. H. Cheung, B. Stevens, I. Kotsireas, V. Ganesh, A SAT-based resolution of Lam’s problem, *Proceedings of the AAAI Conference on Artificial Intelligence 35 (2021)* 3669–3676. doi:10.1609/aaai.v35i5.16483.
- [27] B. Subercaseaux, M. J. H. Heule, The packing chromatic number of the infinite square grid is 15, in: S. Sankaranarayanan, N. Sharygina (Eds.), *Tools and Algorithms for the Construction and Analysis of Systems*, volume 13993 of *Lecture Notes in Computer Science*, Springer Nature Switzerland, Cham, 2023, p. 389–406. doi:10.1007/978-3-031-30823-9\_20.
- [28] C. Bright, J. Gerhard, I. Kotsireas, V. Ganesh, Effective problem solving using SAT solvers, in: J. Gerhard, I. Kotsireas (Eds.), *Maple in Mathematics Education and Research*, volume 1125 of *Communications in Computer and Information Science*, Springer International Publishing, Cham, 2020, p. 205–219. doi:10.1007/978-3-030-41258-6\_15.
- [29] K. Fazekas, A. Niemetz, M. Preiner, M. Kirchweger, S. Szeider, A. Biere, IPASIR-UP: User Propagators for CDCL, in: M. Mahajan, F. Slivovsky (Eds.), *26th International Conference on Theory and Applications of Satisfiability Testing (SAT 2023)*, volume 271 of *Leibniz International Proceedings in Informatics (LIPIcs)*, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2023, pp. 8:1–8:13. doi:10.4230/LIPIcs.SAT.2023.8.
- [30] C. Bright, I. Kotsireas, V. Ganesh, When satisfiability solving meets symbolic computation, *Communications of the ACM 65 (2022)* 64–72. doi:10.1145/3500921.
- [31] M. England, SC-square: Overview to 2021, in: *SC<sup>2</sup> 2021: 6th International Workshop on Satisfiability Checking and Symbolic Computation*, volume 3273 of *CEUR Workshop Proceedings*, 2021, pp. 1–6. URL: <https://ceur-ws.org/Vol-3273/invited1.pdf>.
- [32] E. Abraham, Building bridges between symbolic computation and satisfiability checking, in: *Proceedings of the 2015 ACM International Symposium on Symbolic and Algebraic Computation, ISSAC’15*, ACM, 2015, p. 1–6. doi:10.1145/2755996.2756636.
- [33] E. Zulkoski, V. Ganesh, K. Czarnecki, MathCheck: A math assistant via a combination of computer algebra systems and SAT solvers, in: *Automated Deduction - CADE-25*, volume 9195 of *Lecture Notes in Computer Science*, Springer International Publishing, 2015, p. 607–622. doi:10.1007/978-3-319-21401-6\_41.
- [34] Z. Li, C. Bright, V. Ganesh, An SC-square approach to the minimum Koehen–Specker

- problem, in: SC<sup>2</sup> 2022: 7th International Workshop on Satisfiability Checking and Symbolic Computation, volume 3458 of *CEUR Workshop Proceedings*, 2022, p. 55–66. URL: <https://ceur-ws.org/Vol-3458/paper6.pdf>.
- [35] Z. Li, C. Bright, V. Ganesh, A SAT solver + computer algebra attack on the minimum Kochen–Specker problem, in: K. Larson (Ed.), *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*, International Joint Conferences on Artificial Intelligence Organization, 2024, pp. 1898–1906. doi:10.24963/ijcai.2024/210, main Track.
- [36] Y. Ajani, C. Bright, A hybrid SAT and lattice reduction approach for integer factorization, in: SC<sup>2</sup> 2023: 8th International Workshop on Satisfiability Checking and Symbolic Computation, volume 3455 of *CEUR Workshop Proceedings*, 2023, pp. 39–43. URL: <https://cs.uwaterloo.ca/~cbright/reports/sc2-fact-update.pdf>.
- [37] Y. Ajani, C. Bright, SAT and lattice reduction for integer factorization, in: *Proceedings of the 2024 International Symposium on Symbolic and Algebraic Computation, ISSAC '24*, ACM, 2024, p. 391–399. doi:10.1145/3666000.3669712.
- [38] N. Alamgir, S. Nejati, C. Bright, SHA-256 collision attack with programmatic SAT, in: SC<sup>2</sup> 2024: 9th International Workshop on Satisfiability Checking and Symbolic Computation, volume 3717 of *CEUR Workshop Proceedings*, 2024, pp. 91–110. URL: <https://ceur-ws.org/Vol-3717/paper5.pdf>.
- [39] T. Ahmed, L. Zaman, C. Bright, Symbolic sets for proving bounds on Rado numbers, in: SC<sup>2</sup> 2025: 10th International Workshop on Satisfiability Checking and Symbolic Computation, volume 4116 of *CEUR Workshop Proceedings*, 2025, pp. 55–75. URL: <https://ceur-ws.org/Vol-4116/paper139.pdf>.
- [40] P. Kaski, O. Pottonen, *libexact user’s guide*, version 1.0 (2008). URL: <https://pottonen.kapsi.fi/hiit-tr-2008-1.pdf>, HIIT Technical Reports 2008–1, Helsinki Institute for Information Technology HIIT.
- [41] G. Appa, I. Mourtos, D. Magos, Integrating Constraint and Integer Programming for the Orthogonal Latin Squares Problem, volume 2470 of *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, 2002, p. 17–32. doi:10.1007/3-540-46135-3\_2.
- [42] G. Appa, D. Magos, I. Mourtos, A branch & cut algorithm for a four-index assignment problem, *Journal of the Operational Research Society* 55 (2004) 298–307. doi:10.1057/palgrave.jors.2601655.
- [43] G. Appa, D. Magos, I. Mourtos, An LP-based proof for the non-existence of a pair of orthogonal Latin squares of order 6, *Operations Research Letters* 32 (2004) 336–344. doi:10.1016/j.orl.2003.10.010.
- [44] F. Ma, J. Zhang, Finding orthogonal Latin squares using finite model searching tools, *Science China Information Sciences* 56 (2013) 1–9. doi:10.1007/s11432-011-4343-3.
- [45] O. Zaikin, S. Kochemazov, A. Semenov, SAT-based search for systems of diagonal Latin squares in volunteer computing project SAT@home, in: *39th International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*, 2016, pp. 277–281. doi:10.1109/MIPRO.2016.7522152.
- [46] O. Zaikin, S. Kochemazov, The search for systems of diagonal Latin squares using the SAT@home project, volume 3 of *CEUR Workshop Proceedings*, 2015, pp. 4–9. URL: <https://ceur-ws.org/Vol-1502/paper6.pdf>.
- [47] O. Zaikin, A. Zhuravlev, S. Kochemazov, E. Vatutin, On the construction of triples of diagonal Latin squares of order 10, *Electronic Notes in Discrete Mathematics* 54 (2016) 307–312. doi:10.1016/j.endm.2016.09.053.
- [48] E. Vatutin, O. Zaikin, M. Manzyuk, N. Nikitina, Searching for Orthogonal Latin Squares via Cells Mapping and BOINC-Based Cube-and-Conquer, volume 1510 of *Communications in Computer and Information Science*, 2021, pp. 498–512. doi:10.1007/978-3-030-92864-3\_38.
- [49] R. Lu, S. Liu, J. Zhang, Searching for Doubly Self-orthogonal Latin Squares, volume

6876 of *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, 2011, p. 538–545. doi:10.1007/978-3-642-23786-7\_41.

- [50] R. M. Falcón, Ó. J. Falcón, J. Núñez, Computing the sets of totally symmetric and totally conjugate orthogonal partial Latin squares by means of a SAT solver, in: Proceedings of the 17th International Conference on Computational and Mathematical Methods in Science and Engineering, 2017, pp. 841–852. URL: <http://hdl.handle.net/11441/62076>.
- [51] P. Huang, M. Liu, C. Ge, F. Ma, J. Zhang, Investigating the existence of orthogonal golf designs via satisfiability testing, in: ISSAC '19: Proceedings of the 2019 International Symposium on Symbolic and Algebraic Computation, 2019, pp. 203–210. doi:10.1145/3326229.3326232.
- [52] J. Jin, Y. Lv, C. Ge, F. Ma, J. Zhang, Investigating the Existence of Costas Latin Squares via Satisfiability Testing, volume 12831 of *Lecture Notes in Computer Science*, 2021, pp. 270–279. doi:10.1007/978-3-030-80223-3\_19.
- [53] H. Zhang, Specifying Latin square problems in propositional logic, in: Automated Reasoning and Its Applications: Essays in Honor of Larry Wos, MIT Press, Cambridge, MA, USA, 1997, p. 115–146. URL: <https://dl.acm.org/doi/10.5555/271101.271124>.
- [54] O. Zaikin, E. Vatutin, C. Bright, Enumerating extended self-orthogonal diagonal Latin squares of order up to 10, *Journal of Integer Sequences* 28 (2025). URL: <https://cs.uwaterloo.ca/journals/JIS/VOL28/Zaikin/zaikin4.html>, article 28.7.4.
- [55] C. Bright, A. Keita, B. Stevens, Orthogonal Latin squares of order 10 with two relations: A SAT investigation, *Discrete Mathematics, Algorithms and Applications* (2025). doi:10.1142/s1793830925501563, to appear.
- [56] O. Bailleux, Y. Boufkhad, Efficient CNF Encoding of Boolean Cardinality Constraints, in: F. Rossi (Ed.), Principles and Practice of Constraint Programming – CP 2003, volume 2833 of *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2003, pp. 108–122. doi:10.1007/978-3-540-45193-8\_8.