

Enumerating Projective Planes of Order Nine with Proof Verification

Daniel Dallaire¹, Curtis Bright^{1,*}

¹University of Windsor, Windsor, Ontario, Canada

Abstract

In this paper we describe a method of enumerating projective planes of order nine. The enumeration was previously completed by Lam, Kolesova, and Thiel using highly optimized and customized search code. Despite the importance of this result in the classification of projective geometries, the previous search relied on unverified code and has never been independently verified. Our enumeration procedure—which is still a work in progress—uses a hybrid satisfiability (SAT) solving and symbolic computation approach. SAT solving performs an enumerative search, while symbolic computation removes symmetries from the search space. Certificates are produced which demonstrate the enumeration completed successfully.

Keywords

Satisfiability Checking, Symbolic Computation, Combinatorial Enumeration, Computer-assisted Proof

1. Introduction

Projective geometry is a form of geometry developed by Renaissance artists in order to describe how to represent a three dimensional scene onto a canvas. Projective geometry differs from the more familiar Euclidean geometry because there are no parallel lines in a projective geometry. For example, a pair of train tracks—parallel lines in three dimensions—are no longer parallel when projected onto a canvas because the tracks will meet on the horizon.

The classification of projective geometries is an important and long-standing mathematical problem which despite intense study is still incomplete. Even in the finite case—when the geometry consists of a finite number of points—a classification is missing in the two-dimensional case. Two dimensional projective geometries are known as *projective planes* and in this work we consider one of the few cases for which a classification is known (albeit one relying on custom-written and unverified search code). In this ongoing project, we develop a computer-assisted method that we plan to use to complete the classification of projective planes that have exactly ten points on each line—known as projective planes of order nine.

The method fits into the “SC-square” paradigm of relying on both *satisfiability checking* and *symbolic computation*—two fields of computer science which developed with little interaction [1] but recently have been combined and applied to a variety of problems like circuit

7th International Workshop on Satisfiability Checking and Symbolic Computation, August 12, 2022, Haifa, Israel

*Corresponding author.

✉ dallaird@uwindsor.ca (D. Dallaire); cbright@uwindsor.ca (C. Bright)

🌐 <http://www.curtisbright.com/> (C. Bright)

🆔 0000-0002-0462-625X (C. Bright)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

verification [2], finding new algorithms for matrix multiplication [3], and improving cylindrical algebraic decomposition algorithms [4]. Our method uses satisfiability (SAT) solvers to search for projective planes and symbolic computation to detect when two partial projective planes are isomorphic to each other. During the search many isomorphic subplanes are detected and pruned from the search, thereby dramatically improving the efficiency of the solver.

Crucially, our work does not rely on trusting the output of either the SAT solver or the computer algebra system (CAS). Both tools produce certificates that can be used to *verify* the output without taking their claims on faith. Although this is a work in progress, we are confident that our approach will successfully complete a classification of projective planes of order nine without requiring the trust of any search code. This is particularly important in computational classification which requires the generation of a complete list of *all* instances of a given object—it is in general quite difficult to prove that the list is complete. Moreover, it is not feasible to prove that a complicated algorithm (like a search procedure) generates correct results in all cases [5]. A similar SAT+CAS approach has also successfully been used [6] to generate proof certificates verifying Lam et al.’s experimental proof of the nonexistence of projective planes of order ten [7].

2. Background

The objects we will be interested in this paper are primarily finite projective planes, which we introduce here. These projective planes have a few different representations, the first of which is perhaps the easiest to understand. After giving this as the definition, we discuss another representation which will be more useful for the purpose of doing an exhaustive computer search for these objects. Also playing a role in this work is the notion of a latin square which we define here. Lastly, we describe the notion of a Boolean satisfiability problem or SAT problem.

To start, we define a projective plane of a given order n :

Definition 1. *A projective plane of order n is a collection of $n^2 + n + 1$ lines and $n^2 + n + 1$ points such that:*

- (1) *every line contains $n + 1$ points,*
- (2) *every point is on $n + 1$ lines,*
- (3) *any two distinct lines intersect at exactly one point, and*
- (4) *any two distinct points lie on exactly one line.*

From this definition, we can see that projective planes are objects which have a natural interpretation as an incidence structure—that is, two disjoint sets equipped with a relation between them which we call the incidence relation [8]. One of the simplest examples of this is the *Fano plane*, which is a projective plane of order 2 (i.e., it has $2^2 + 2 + 1 = 7$ points and lines) as seen in Figure 1. With such structures, we may associate a bipartite graph, whose parts in our case are the lines and points of the projective plane respectively, and whose edge set is determined by the incidence relation. Specifically, in this bipartite graph, we make an edge between a point and a line if and only if that point lies on the line. Suppose now that we order the points and lines of the projective plane as: $p_1, p_2, \dots, p_{n^2+n+1}$ and $\ell_1, \ell_2, \dots, \ell_{n^2+n+1}$.




Figure 1: The Fano Plane. The points are represented by dots and the lines of the plane are represented by straight or curved lines between the points.

Then we may represent the associated bipartite graph as an $(n^2 + n + 1) \times (n^2 + n + 1)$ binary *incidence matrix*, in which each row represents a point and each column represents a line. The (i, j) th entry of this matrix will be 1 if the point p_i lies on the line ℓ_j , and it will be 0 otherwise. This definition is more workable for the purpose of using a SAT solver as we can encode the incidence matrix as $(n^2 + n + 1)^2$ Boolean variables. To be able to utilize this representation, we must also translate the axioms of the projective plane (1)–(4) given above in terms of the incidence matrix. Two binary vectors are said to *intersect* if they both contain 1s in the same entry of the vector (or rather, their inner product is greater than 0). One can check that an $(n^2 + n + 1) \times (n^2 + n + 1)$ incidence matrix is one which arises from a projective plane of order n if the following properties are satisfied:

- (1) each column sum of the matrix is $n + 1$,
- (2) each row sum of the matrix is $n + 1$,
- (3) two distinct columns of the matrix intersect exactly once, and
- (4) two distinct rows of the matrix intersect exactly once.

Next, we also introduce latin squares.

Definition 2. A $k \times k$ latin square is a $k \times k$ array consisting of the integers $1, 2, \dots, k$ such that:

1. each row contains each of the numbers $1, 2, \dots, k$ exactly once, and
2. each column contains each of the numbers $1, 2, \dots, k$ exactly once.

For example, the following is a latin square of order 4:

$$\begin{bmatrix} 1 & 2 & 3 & 4 \\ 2 & 3 & 4 & 1 \\ 3 & 4 & 1 & 2 \\ 4 & 1 & 2 & 3 \end{bmatrix}$$

The role that latin squares play in our search for projective planes will become more clear after our discussion in section 2.2.

Lastly, we describe SAT problems—which are useful because heuristic algorithms exist that often solve SAT instances extremely efficiently. A *literal* is a Boolean variable or a negated

Boolean variable, and a *clause* is a disjunction of literals. Given a list of clauses, we say the list is *satisfiable* if we can assign true and false values to its variables such that every clause evaluates to be true. A *SAT problem* is the problem of determining if a given list of clauses is satisfiable or not. This problem is known to be NP-complete and there is no known polynomial time algorithm solving it. However, there are good heuristic SAT solvers like MapleSAT [9] which can often solve SAT instances.

2.1. Problem Overview

With this background established, we now give more details about the problem we're interested in. As we discussed, we can represent a hypothetical projective plane of order n as an $(n^2 + n + 1) \times (n^2 + n + 1)$ incidence matrix satisfying certain properties. For the order 9 case, we're interested in enumerating the 91×91 incidence matrices (note: $91 = 9^2 + 9 + 1$) satisfying certain properties up to a certain symmetry. We now comment on what this symmetry is.

Notice that given a fixed projective plane of order 9, one can obtain a distinct, but similar, projective plane by relabelling some of the points, and relabelling some of the lines. Realistically however, these "new" planes are the same as the original one, thus we wish to ignore these extra planes in our search if possible. These relabellings correspond to column and row permutations of the incidence matrix. Consequently, we can give a group theoretic interpretation of this: Let G be the group $S_{91} \times S_{91}$. We may view the first component of this group as the row permutations, and the second as the column permutations. In this way, we have a group action of G on the set of incidence matrices corresponding to projective planes. Then we can say that two planes are equivalent if they are in the same *orbit* under this action. Two planes are in the same orbit, if one can be obtained from the other via action of the group G , or rather, by applying column and row permutations.

Searching for only a representative of each orbit makes our search drastically more feasible. This computational search was first done by Lam, Kolesova, and Thiel [10]. Before this search was done, there were four distinct projective planes of order 9 known to exist. Lam et al.'s search showed that these were the only four planes up to isomorphism. We repeat this search using a SAT solver and by exploiting the symmetry group mentioned above to reduce the search space. By applying the elements of the symmetry group to a hypothetical projective plane, we can fix certain structure for the plane. This is described in Section 2.2.

2.2. Structure of Planes of Order 9

Here we detail the structure we impose on the incidence matrix of our hypothetical projective plane for the purpose of reducing the search space of our problem. The first important thing to note is that a hypothetical projective plane must contain a *triangle*—that is, a set of three non-collinear points. We may suppose that the first three points p_1, p_2 , and p_3 form this triangle, and furthermore that ℓ_1 is the unique line joining p_2 and p_3 , ℓ_2 is the unique line joining p_1 and p_2 , and lastly that ℓ_3 is the unique line joining p_1 and p_3 . We may impose this order simply by applying the needed row and column permutations to the incidence matrix. With this ordering




Figure 2: The normalized incidence matrix from Lam et al.'s 1991 paper [10]. The red entries (those in the first 19 columns or the first 19 rows) are fixed in advance. The blue entries (those in the submatrices B_1, \dots, B_8 which appear in columns 20–27) are not fixed in advance, and all possibilities for these entries must be enumerated during the search. One way of completing the submatrices B_1, \dots, B_8 is given in the figure.

of the points and lines, the upper left 3×3 sub-matrix of the incidence matrix will look like:

$$\begin{bmatrix} 0 & 1 & 1 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix}$$

Once this is done, we know by the first two axioms of a projective plane, that there should be 8 more 1's in each of the first three rows and columns. By applying further row and column permutations, we may impose a staircase like structure on these entries. For example, we can make the first three rows in columns 4 to 27 will look like:

$$\begin{bmatrix} 1 & 1 & \dots & 1 & 0 & 0 & \dots & 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & \dots & 0 & 1 & 1 & \dots & 1 & 0 & 0 & \dots & 0 \\ 0 & 0 & \dots & 0 & 0 & 0 & \dots & 0 & 1 & 1 & \dots & 1 \end{bmatrix}$$

with the entries in columns 28 and above being 0 in those rows.

Then, using further column and row permutations, we may fix many of the entries in the first 27 columns of the incidence matrix shown in Figure 2. The entries in the submatrix formed by columns 20–27 and rows 28–91 will not be fixed but they do have a nice structure.

One can check that the 8×8 blocks B_1, B_2, \dots, B_8 given in Figure 2 will be 8×8 permutation matrices. Consequently, these blocks correspond to permutations $\pi_1, \pi_2, \dots, \pi_8 \in S_8$. Furthermore, we know that we can't have $\pi_i(k) = \pi_j(k)$ where $i \neq j$ since otherwise we would have two distinct rows intersecting twice. This property ensures that we may encode this set of

8 permutations as an 8×8 latin square. A more detailed explanation of this normalized form for the partial plane can be found in Kolesova’s thesis [11, Prop. 4.1].

Thus, for a given 8×8 latin square, we may obtain the 27 columns of a partial projective plane, and in this way, all such partial planes may be obtained if all 8×8 latin squares can be enumerated. Thus, our approach for generating the projective planes of order 9 will begin by generating the latin squares of order 8. However, two distinct latin squares can give equivalent partial planes. Thus we will be interested in equivalence classes of latin squares as detailed in the next section.

3. Enumeration Method

Our enumeration method proceeds in two steps. First, all 8×8 latin squares are enumerated up to isomorphism using a SAT+CAS method. Second, for every 8×8 latin square generated in the first step the initial 27 columns of a projective plane of order nine are defined in terms of the structure revealed in Section 2.2. For each possible way of completing the first 27 columns a SAT instance is created which is satisfiable exactly when the 27 columns can be extended to 40 columns of a projective plane.

3.1. Step 1: Latin Squares

As the structure we identified in the last section suggests, a good first step to generate the projective planes of order 9 will be to first generate all 8×8 latin squares up to a certain equivalence. More specifically, we will generate a representative of each *main class* of latin squares. Two 8×8 latin squares are said to belong to the same main class if they differ only by a permutation of the rows, columns, or symbols, or also possibly an exchange of the roles of these three things. Thus we can view the main class of a latin square as its orbit under the natural action of the group $(S_8 \times S_8 \times S_8) \rtimes S_3$ [12, Sec 5.5]¹ on the set of latin squares. To generate all such representatives, we reduce the problem to SAT, and then utilize a SAT solver to find all solutions of the instance. The exhaustive search is performed by adjoining a “solution blocking clause” to the instance whenever a solution is found. To reduce the computation, we also provide a mechanism for doing isomorphism checking along the way; if two partial squares belong to the same main class, we throw away the duplicates to reduce the computation time.

The reduction of the latin square problem into SAT is well known, and can be found in a variety of sources [13, 14]. Before utilizing the SAT solver, we first normalize the latin squares by insisting that the numbers 1, 2, . . . , 8 appear in order in the first column and row. This is done by adding unit clauses specifying that the entries in the first row have such values. This is justified since we may apply column and row permutations to impose this structure. With this done, we use MapleSAT to extend one row at a time up to row 4 of the latin square, and then from row 4 to row 8. After each row extension, isomorphism removal is also done.

Isomorphism removal is done by translating our latin squares to graphs and then checking if the corresponding graphs are isomorphic to each other using pynauty [15]. We do this in

¹Semi-direct products (denoted by \rtimes) are like regular products, except elements from the two groups in the product may not commute. Rather, the elements may be commuted at the cost of applying the action of the 2nd group element on the first. For more details, see the reference cited.

such a way that the corresponding graphs will be isomorphic if and only if the original latin squares belonged to the same main class. Our translation of a latin square $(A_{i,j})$ to a graph is as follows: The vertex set is $\{V_{i,j} : 1 \leq i, j \leq n\}$ and we draw an edge between $V_{i,j}$ and $V_{i',j'}$ if $i = i', j = j'$, or $A_{i,j} = A_{i',j'}$. More details of this construction and its validity can found in Miller's paper [16].

Before performing the final isomorphism removal after generating all 8 rows of the latin squares, there were 43,791,204 solutions found by MapleSAT. Once the isomorphism removal is applied, there were 283,657 representatives of the main classes of latin squares which remained, which serve as the starting point for the next step in our computation. The main classes of latin squares of order 8 have been enumerated by others including Lam et al. [17] and the number of latin squares produced by our computation matches the number previously reported. Solving the SAT instances takes around 20 minutes on a single CPU core and performing the isomorphism removal takes around 85 hours.

3.2. Step 2: Column 40 Extension

The next major step in our enumeration of the projective planes to extend each of the 283,657 partial planes (each with 27 columns to start) to partial planes of 40 columns. We accomplish this by once again using a SAT solver. However, we take a slightly different approach than we did with the latin squares since the axioms of a projective plane don't translate as nicely into SAT. A partial plane $(A_{i,j})$ has a natural encoding as a set of Boolean variables by defining a variable $a_{i,j}$ for each entry. We encode in SAT only axioms (3) and (4)—namely, that two distinct rows or columns intersect exactly once. First, the requirement that they intersect at most once is given by the *quadfree* clauses:

$$\neg a_{i,j} \vee \neg a_{i',j} \vee \neg a_{i,j'} \vee \neg a_{i',j'}$$

for $i < i'$ and $j < j'$. If one of these clauses is false this means there is a rectangle in the incidence matrix whose corners are 1s—which is exactly what happens if two rows or columns intersect more than once.

Then, we check that a given column intersects the first 27 columns of the partial plane each at least once. Note that because the first 19 columns are the only ones which are fixed among all partial planes, we will need a different list of clauses for each starting point of the extension. If column $j \leq 27$ has 1s in the entries $(i_1, j), (i_2, j), \dots, (i_{10}, j)$, then for $j' > 27$, we include the clause

$$a_{i_1,j'} \vee a_{i_2,j'} \vee \dots \vee a_{i_{10},j'}$$

which ensures that column j' intersects column j at least once. In addition to these clauses, we also use unit clauses to impose some structure on the first 19 rows (an exact transpose of the first 19 columns in fact) in order to remove symmetry.

This is currently work in progress but in practice the extension takes about 45 seconds for the entire pipeline (MapleSAT, GRATgen, isomorphism removal, and isomorphism removal verification) to run for each 27-column partial plane. This still needs to be done for each of the 283,657 latin squares which will require Compute Canada's servers. The few 40-column partial planes that result from this can typically be extended to a full 91 columns (or shown not

to complete at all) almost instantly [10]. The final step, which still needs to be implemented, will be to perform isomorphism removal and verification on each of the complete projective planes found by the solver.

4. Verification

The SAT solver MapleSAT returns DRAT proofs which can be checked using a proof verifier such as DRAT-trim [18] or GRATgen [19]. This way only the proof verifier—which is much simpler than the SAT solver—needs to be trusted.

In step 1, we do several row extensions with MapleSAT to produce the final list of latin squares. For each of these, we generate and verify the proofs for these steps. This involves adding a blocking clause for each solution of the SAT instance and then showing a conflict can be derived from the original list of clauses in conjunction with the blocking clauses. In step 2, we have a separate list of clauses for each of the 27 column partial planes, thus we end up with a separate proof for each one, all of which need to be verified independently just as in step 1. Lastly, we must also do this for the final extension from column 40 to the full 91 columns.

In addition to verifying the proofs above, we will verify the correctness of the isomorphism removal which is done at several stages. After generating solutions, we create a new file with one representative of the isomorphism class of each of the solutions in the first file. When doing so, we store the relabelling of the corresponding graph which turns it into the graph of its representative in the new file. This can then be verified easily by checking that each relabelling turns the associated solution into the one at the specified index in the new file. Representatives are chosen by determining the canonical form of the associated graph with pynauty and then using the adjacency lists of these new graphs as keys in a Python dictionary in order to only keep one of each.

The DRAT proof from step 1 can be verified by DRAT-trim in about an hour, and the isomorphism removal can be verified in about 25 hours.

5. Conclusion

This method of enumerating of the projective planes of order nine relies on two main components: the generation of solutions with our SAT solver (MapleSAT), and the isomorphism removal of partial solutions with our CAS (pynauty). Both components are crucial—if either component were removed this work would not be feasible to complete in a reasonable amount of time. To the best of our knowledge, the resulting enumeration will be the first independent verification of the search of Lam, Kolesova, and Thiel [10]. Moreover, an enumeration using a SAT+CAS system can be trusted to a higher degree of certainty than an enumeration via custom-written search code. Our enumeration has been designed to require trusting a *certificate verifier* rather than trusting a *search procedure*.

References

- [1] E. Abraham, Building bridges between symbolic computation and satisfiability checking, in: Proceedings of the 2015 ACM on International Symposium on Symbolic and Algebraic Computation, ACM, 2015, pp. 1–6. doi:10.1145/2755996.2756636.
- [2] D. Kaufmann, A. Biere, M. Kauers, SAT, computer algebra, multipliers, in: EPiC Series in Computing, EasyChair, 2020, pp. 1–18. doi:10.29007/j8cm.
- [3] M. J. H. Heule, M. Kauers, M. Seidl, New ways to multiply 3×3 -matrices, Journal of Symbolic Computation 104 (2021) 899–916. doi:10.1016/j.jsc.2020.10.003.
- [4] R. Bradford, J. H. Davenport, M. England, A. Sadeghimanesh, A. Uncu, The DEWCAD project: pushing back the doubly exponential wall of cylindrical algebraic decomposition, ACM Communications in Computer Algebra 55 (2021) 107–111. doi:10.1145/3511528.3511538.
- [5] C. W. H. Lam, Opinion, The Mathematical Intelligencer 12 (1990) 8–12. doi:10.1007/bf03023977.
- [6] C. Bright, K. K. H. Cheung, B. Stevens, I. Kotsireas, V. Ganesh, A SAT-based resolution of Lam’s problem, in: Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence, 2021, pp. 3669–3676. URL: <https://ojs.aaai.org/index.php/AAAI/article/view/16483>.
- [7] C. W. H. Lam, L. Thiel, S. Swiercz, The non-existence of finite projective planes of order 10, Canadian Journal of Mathematics 41 (1989) 1117–1123. doi:10.4153/cjm-1989-049-4.
- [8] C. Godsil, G. F. Royle, Algebraic Graph Theory, Springer Science & Business Media, 2001.
- [9] J. H. Liang, V. Ganesh, P. Poupart, K. Czarnecki, Learning rate based branching heuristic for SAT solvers, in: Theory and Applications of Satisfiability Testing – SAT 2016, Springer International Publishing, 2016, pp. 123–140. doi:10.1007/978-3-319-40970-2_9.
- [10] C. W. H. Lam, G. Kolesova, L. Thiel, A computer search for finite projective planes of order 9, Discrete Mathematics 92 (1991) 187–195. doi:10.1016/0012-365x(91)90280-f.
- [11] G. I. Kolesova, Enumeration of the Finite Projective Planes of Order Nine, Master’s thesis, Concordia University, 1989.
- [12] D. S. Dummit, R. M. Foote, Abstract Algebra, John Wiley and Sons, Inc., 2004.
- [13] J. Jin, Y. Lv, C. Ge, F. Ma, J. Zhang, Investigating the existence of Costas Latin squares via satisfiability testing, in: Theory and Applications of Satisfiability Testing – SAT 2021, Springer International Publishing, 2021, pp. 270–279. doi:10.1007/978-3-030-80223-3_19.
- [14] C. Bright, J. Gerhard, I. Kotsireas, V. Ganesh, Effective problem solving using SAT solvers, in: Communications in Computer and Information Science, Springer International Publishing, 2020, pp. 205–219. doi:10.1007/978-3-030-41258-6_15.
- [15] B. D. McKay, A. Piperno, Practical graph isomorphism, II, Journal of Symbolic Computation 60 (2014) 94–112. doi:10.1016/j.jsc.2013.09.003.
- [16] G. L. Miller, On the $n^{\log n}$ isomorphism technique (a preliminary report), in: Proceedings of the tenth annual ACM symposium on Theory of computing - STOC ’78, ACM Press, 1978, pp. 51–58. doi:10.1145/800133.804331.
- [17] G. Kolesova, C. W. H. Lam, L. Thiel, On the number of 8×8 Latin squares, Journal of Combinatorial Theory, Series A 54 (1990) 143–148. doi:10.1016/0097-3165(90)90015-o.

- [18] N. Wetzler, M. J. H. Heule, W. A. Hunt, DRAT-trim: Efficient checking and trimming using expressive clausal proofs, in: *Lecture Notes in Computer Science*, Springer International Publishing, 2014, pp. 422–429. doi:10.1007/978-3-319-09284-3_31.
- [19] P. Lammich, Efficient verified (UN)SAT certificate checking, *Journal of Automated Reasoning* 64 (2019) 513–532. doi:10.1007/s10817-019-09525-z.