



The enumeration of k -sets of mutually orthogonal Latin squares

JG Benadé* AP Burger† JH van Vuuren‡

Abstract

Latin squares and sets of k mutually orthogonal Latin squares (k -MOLS) have application in various scheduling problems, from providing effective ways to access parallel memory structures to scheduling transmissions from sensor arrays. MOLS also play an important role in sports tournament scheduling where every structurally different MOLS provides the scheduler with an additional degree of scheduling freedom. The existence of 3-MOLS have been resolved for all orders of Latin squares, except for order 10. We consider a backtracking algorithm for the enumeration of structurally different MOLS which partitions the search space in such a way that it is possible to estimate bounds for the enumeration of higher-order MOLS. A contribution towards the celebrated question of the existence of a 3-MOLS of order 10 is made by investigating the feasibility of using this algorithm in conjunction with specific computing paradigms in search of such a design.

Key words: Enumeration, mutually orthogonal Latin squares (MOLS), volunteer computing.

1 Introduction

A Latin square of order n is an $n \times n$ array in which every cell contains a single symbol with the property that each symbol occurs exactly once in each row and column of the array [6, Definition 1.1], and two Latin squares are *orthogonal* if each of the n^2 superimposed ordered pairs of symbols, one pair for every (row, column)-position in the arrays, is distinct. Four examples of Latin squares of order 4 may be seen in Figure 1. Note that all three pairs of Latin squares from the set $\{\mathbf{L}_1^*, \mathbf{L}_2^*, \mathbf{L}_3^*\}$ are orthogonal.

$$\mathbf{L}_1^* = \begin{bmatrix} 0 & 1 & 2 & 3 \\ 3 & 2 & 1 & 0 \\ 1 & 0 & 3 & 2 \\ 2 & 3 & 0 & 1 \end{bmatrix} \quad \mathbf{L}_2^* = \begin{bmatrix} 0 & \mathbf{1} & 2 & 3 \\ 2 & 3 & 0 & \mathbf{1} \\ 3 & 2 & \mathbf{1} & 0 \\ \mathbf{1} & 0 & 3 & 2 \end{bmatrix} \quad \mathbf{L}_3^* = \begin{bmatrix} 0 & 1 & 2 & 3 \\ 1 & 0 & 3 & 2 \\ 2 & 3 & 0 & 1 \\ 3 & 2 & 1 & 0 \end{bmatrix} \quad \mathbf{L}_4^* = \begin{bmatrix} 0 & 3 & 1 & 2 \\ 1 & 2 & 0 & 3 \\ 2 & 1 & 3 & 0 \\ 3 & 0 & 2 & 1 \end{bmatrix}$$

Figure 1: Some Latin squares of order 4.

Latin squares were first formally studied by Leonard Euler when he considered the so-called “36-Officers problem” asking whether it is possible to arrange thirty-six soldiers of six different ranks and from six different regiments in a square platoon with the properties that every row and

*Department of Logistics and MIH Media Lab, Stellenbosch University, Private Bag X1, Matieland, 7602, South Africa, email: jgbenade@m1.sun.ac.za

†Department of Logistics, Stellenbosch University, Private Bag X1, Matieland, 7602, South Africa, email: apburger@sun.ac.za

‡(Fellow of the Operations Research Society of South Africa), Department of Industrial Engineering, Stellenbosch University, Private Bag X1, Matieland, 7602, South Africa, email: vuuren@sun.ac.za

column contains exactly one soldier of every rank and one soldier from every regiment¹ [8]. Euler was unable to find such an arrangement of soldiers (corresponding to a pair of orthogonal Latin squares of order 6) and continued to propose what has become known as *Euler's Conjecture*, namely that no pair of orthogonal Latin squares of order n exists for $n = 4m + 2$, where m is an integer [8]. In 1900, the French mathematician Gaston Tarry proved Euler's Conjecture correct for $n = 6$, but sixty years later Bose, Shrikhande and Parker [4] showed that it is possible to construct such orthogonal pairs for all cases of Euler's Conjecture other than $n = 6$, thereby disproving the conjecture in general.

The notion of orthogonality may be generalised to sets of k *mutually orthogonal Latin squares*, abbreviated to k -MOLS, which have the property that any two distinct Latin squares in the k -MOLS $\mathcal{M} = \{\mathbf{L}_0, \mathbf{L}_1, \dots, \mathbf{L}_{k-1}\}$ are pairwise orthogonal. The set $\mathcal{M}_1^* = \{\mathbf{L}_1^*, \mathbf{L}_2^*\}$ is therefore an example of a 2-MOLS of order 4, while the set $\mathcal{M}_2^* = \{\mathbf{L}_1^*, \mathbf{L}_2^*, \mathbf{L}_3^*\}$ is a 3-MOLS of order 4.

It has been shown that k -MOLS have important applications in coding theory [16], various subfields of statistics (including experimental design) [9, 10], distributed database systems [1] and numerous scheduling problems, including the scheduling of sports tournaments [11, 13, 20]. Moreover, every structurally different set of orthogonal Latin squares provides the scheduler with an alternative schedule, and some of these schedules may be more desirable than others due to *ad hoc* constraints or preferences in the scheduling problem. The known numbers of structurally different k -MOLS for order $n \in \{3, 4, \dots, 8\}$ appear in Table 1. Additionally, it is known that there are 19 structurally different 8-MOLS of order 9 [19] and that no k -MOLS of order 10 exists for $k \in \{7, 8, 9\}$ [7, 15].

n	k					
	2	3	4	5	6	7
3	1					
4	1	1				
5	1	1	1			
6	0	0	0	0		
7	7	1	1	1	1	
8	2 165	39	1	1	1	1

Table 1: The numbers of structurally different k -MOLS of order n for $n \in \{3, 4, \dots, 8\}$.

The objective in this paper is to consider an algorithm for the enumeration of structurally different k -MOLS of order n , demonstrating the correctness of this algorithm by replicating the known results in Table 1, and to produce estimates of the sizes of the search spaces for 3-MOLS of orders 9 and 10, which are yet to be enumerated. These estimates should shed light on the current and short-term future feasibility of any further enumeration attempts using this approach.

2 Mathematical preliminaries

Let $S(\mathbf{L})$ denote the symbol set of a Latin square \mathbf{L} and let $R(\mathbf{L})$ and $C(\mathbf{L})$ denote its row and column indexing sets, respectively. For any $i \in R(\mathbf{L})$ and $j \in C(\mathbf{L})$, let $\mathbf{L}(i, j) \in S(\mathbf{L})$ denote the element in the i -th row and the j -th column of \mathbf{L} . In the remainder of this paper it is assumed, without any subsequent loss of generality, that $R(\mathbf{L}) = C(\mathbf{L}) = S(\mathbf{L}) = \mathbb{Z}_n$, the set of residues of the integers after division by the natural number n . The *transpose* of \mathbf{L} , denoted \mathbf{L}^T , is the Latin square for which $\mathbf{L}^T(j, i) = \mathbf{L}(i, j)$ for all $i \in R(\mathbf{L})$ and $j \in C(\mathbf{L})$. Note, for example, that example, $\mathbf{L}_1^{*T} = \mathbf{L}_4^*$.

The notion of a universal was introduced by Burger *et al.* [5] in 2010 to facilitate the enumeration of specific classes of Latin squares. A *universal* of a Latin square \mathbf{L} is a set of n distinct, ordered pairs (i, j) , one from each row and column, all containing the same symbol. Universals may be expressed in permutation form such that the *universal permutation* u_ℓ of the symbol ℓ maps i to j if $\mathbf{L}(i, j) = \ell$ and, as such, it is possible to find the cycle structure

¹A similar and even older puzzle posed by Claude Gaspard Bachet in 1624 concerns the number of fundamentally different ways in which the 16 court cards from a standard deck of cards may be arranged in a 4×4 square such that every row and column contains a card of every rank and every suit [3].

and inverse of any universal. The *relative cycle structure* of any pair of universals u_1 and u_2 of the same order is defined to be the cycle structure of $u_2 \circ u_1^{-1}$, where \circ is the traditional composition operator for permutations. In \mathbf{L}_2^* , for example, the entries in boldface correspond to the universal $\{(0, 1), (1, 3), (2, 2), (3, 0)\}$ of the symbol 1, which may be written in permutation notation as $\begin{pmatrix} 0 & 1 & 2 & 3 \\ 1 & 3 & 2 & 0 \end{pmatrix}$, abbreviated here as $\langle 1320 \rangle$. Let $U(\mathcal{M})$ denote the set of universal permutations of some k -MOLS \mathcal{M} , and let $u_\ell(m) \in U(\mathcal{M})$ denote the universal permutation of the symbol ℓ in the m -th square $\mathbf{L}_m \in \mathcal{M}$. The set of all universals of \mathcal{M}_1^* is therefore $U(\mathcal{M}_1^*) = \{\langle 0312 \rangle, \langle 1203 \rangle, \langle 2130 \rangle, \langle 3021 \rangle, \langle 0231 \rangle, \langle 1320 \rangle, \langle 2013 \rangle, \langle 3102 \rangle\}$, while the universal permutation of the symbol 2 in the third Latin square of \mathcal{M}_2^* is $u_2(2) = \langle 2301 \rangle$. The relative cycle structure of $u_0(1) \in U(\mathcal{M}_1^*)$ and $u_3(0) \in U(\mathcal{M}_1^*)$ is the cycle structure of the permutation $\langle 3021 \rangle \circ \langle 0231 \rangle^{-1} = \langle 2013 \rangle$, which may be denoted as $z_1^1 z_3^1$ as it consists of one cycle of length 1 and one cycle of length 3.

Latin squares which can be generated from one another by changing the order of their rows and/or columns, and/or by renaming their symbols, are said to be *isotopic*, while Latin squares formed by uniformly applying a permutation to all n^2 3-tuples $(i, j, \mathbf{L}(i, j))$ are called *conjugates*. For example, applying the permutation $\begin{pmatrix} 0 & 1 & 2 \\ 1 & 0 & 2 \end{pmatrix}$ to the 3-tuple $(i, j, \mathbf{L}(i, j))$ yields the transpose $(j, i, \mathbf{L}(i, j))$ of \mathbf{L} . A maximal set of isotopic Latin squares, together with all their conjugates, form a *main class* of Latin squares. It is possible to show that \mathbf{L}_1^* , \mathbf{L}_2^* , \mathbf{L}_3^* and \mathbf{L}_4^* are all in the same main class by reordering the rows of \mathbf{L}_1^* to find \mathbf{L}_2^* , transposing \mathbf{L}_1^* to form \mathbf{L}_4^* and, finally, reordering the columns of \mathbf{L}_4^* to form \mathbf{L}_3^* .

The notions of isotopic and conjugate Latin squares, as well as that of main classes, may be extended to k -MOLS. All k -MOLS which may be generated by row, column and symbol permutations from a given k -MOLS are isotopic, with the additional constraint that the same row or column permutation must be applied to all k Latin squares in the k -MOLS in order to maintain orthogonality (the symbol sets, however, may be renamed independently). Conjugates, in this case, are k -MOLS formed by uniformly applying permutations to the $(k + 2)$ -tuples $(i, j, \mathbf{L}_0(i, j), \dots, \mathbf{L}_{k-1}(i, j))$ and a main class consists of a given k -MOLS, together with its $(k + 2)!$ conjugates as well as their respective isotopic k -MOLS.

It is possible to define a *lexicographical ordering*, denoted by the symbol \prec , on a main class of k -MOLS by comparing the universals lexicographically in such a way that every main class has a unique smallest element, called the *class representative*. Two k -MOLS, \mathcal{M} and \mathcal{M}' are ordered in this way by comparing corresponding universals, starting with $u_0(0) \in U(\mathcal{M})$ and $u'_0(0) \in U(\mathcal{M}')$, followed by $u_0(1) \in U(\mathcal{M})$ and $u'_0(1) \in U(\mathcal{M}')$, etc. until it is either found that the one k -MOLS is lexicographically smaller than the other, or until all universals have been compared, in which case \mathcal{M} and \mathcal{M}' are lexicographically equal and therefore the same k -MOLS. For example, when comparing the two 2-MOLS $\mathcal{M}^* = \{\mathbf{L}_1^*, \mathbf{L}_2^*\}$ and $\mathcal{M}' = \{\mathbf{L}_1^*, \mathbf{L}_4^*\}$ of order 4, it is seen that $u_0(0) = u'_0(0)$, $u_0(1) = u'_0(1)$ and $u_1(0) = u'_1(0)$ but $u_1(1) = \langle 1320 \rangle \prec u'_1(1) = \langle 2013 \rangle$, implying that $\mathcal{M}^* \prec \mathcal{M}'$.

3 Exhaustive enumeration of k -MOLS

An exhaustive enumeration of k -MOLS of order n may be carried out by the orderly generation of the class representatives of every main class. The pseudo-code of such an enumeration procedure is given as Algorithm 1. A backtracking tree-search is implemented in Algorithm 1 for constructing k -MOLS of order n , one universal at a time, in such a way that, for $i \in \mathbb{Z}_n$ and $m \in \mathbb{Z}_k$, the active nodes on level $i.m$ of the search tree correspond to the lexicographically smallest partial k -MOLS whose Latin squares $\mathbf{L}_0, \dots, \mathbf{L}_m$ each contains $i + 1$ universals and whose Latin squares $\mathbf{L}_{m+1}, \dots, \mathbf{L}_k$ each contains i universals. The inactive nodes in the search tree represent those partial k -MOLS which cannot be completed to a class representative or in

which the partial Latin squares are no longer pairwise orthogonal. On level $i.(k-1)$ of the search tree the universal for the symbol i has been inserted in all the Latin squares $\mathbf{L}_0, \dots, \mathbf{L}_{k-1}$ of the partial k -MOLS and the next universal to insert is $u_{i+1}(0)$; as this level marks the completion of the partial k -MOLS up to the symbol i , it is also referred to simply as level i .

Suppose that the partial k -MOLS \mathcal{P} has been constructed on level $i.\ell$ of the search tree, in other words, the next universal to insert into \mathcal{P} is $u_i(\ell+1)$, or $u_{i+1}(0)$ if $\ell = k-1$. Let $U(\mathcal{P})$ be the set of all universals in the partial k -MOLS \mathcal{P} , $U(\mathcal{P}_{\ell+1})$ the set of all universals in \mathcal{P} , excluding the universals of $\mathbf{L}_{\ell+1}$ (the Latin square into which a universal is currently being added) and denote the set of feasible candidate universals by $\mathcal{C}(\mathcal{P})$. The node in the search tree representing \mathcal{P} thus has $|\mathcal{C}(\mathcal{P})|$ children, any number of which may be inactive.

To verify orthogonality in a child $\mathcal{P} \cup c$ of \mathcal{P} , for some candidate universal $c \in \mathcal{C}(\mathcal{P})$, it is necessary to confirm that the relative cycle structure of c and every permutation $p \in U(\mathcal{P}_{\ell+1})$ has exactly one fixed point. The following result by Kidd *et al.* [12, Theorem 4.3.2] provides an easy way of determining whether a partial k -MOLS \mathcal{M} is the lexicographically smallest partial k -MOLS in its main class.

Theorem 1 [12, Theorem 4.3.2] *If $\mathcal{M} = (\mathbf{L}_0, \dots, \mathbf{L}_{k-1})$ is the lexicographically smallest k -MOLS of order n in its main class, then (a) $u_0(0)$ is the identity permutation, (b) $u_0(1)$ is a cycle structure representative, and (c) the relative cycle structure of two universal permutations $u_i(j), u_\ell(m)$ is not lexicographically smaller than the cycle structure of $u_0(1) \in U(\mathcal{M})$ for all $i, j \in \mathbb{Z}_n$ and $j, m \in \mathbb{Z}_n$.*

According to Theorem 1 (a) and (b) there is a very limited number of feasible zero universals in \mathbf{L}_0 and \mathbf{L}_1 , and by Theorem 1 (c) no relative cycle structure calculated while verifying orthogonality may be smaller than the cycle structure of $u_0(1)$ if $\mathcal{P} \cup c$ is to be the lexicographically smallest partial k -MOLS in its main class.

If $\mathcal{P} \cup c$ passes this test, then all possible pairs of universals $u_a(j), u_b(m)$ in $\mathcal{P} \cup c$ or its transpose $(\mathcal{P} \cup c)^T$ with a relative cycle structure equal to the cycle structure of $u_0(1)$ are mapped to the pair of universals $u_0(0), u_0(1)$ to form a new partial k -MOLS $(\mathcal{P} \cup c)'$ in the same main class, which is then subjected to a restricted number of row, column and symbol permutations in an attempt to find a lexicographically smaller partial MOLS. More specifically, in order to ensure that the universal $u_0(0)$ in $(\mathcal{P} \cup c)'$ remains unchanged, it is necessary to apply any potential permutation

Algorithm 1: enumerateMOLS(\mathcal{P})

input : A partial k -MOLS \mathcal{P}

output: All completed class representatives in the subtree rooted at \mathcal{P}

```

1 begin
2   if  $\mathcal{P}$  is complete then
3     if none of the conjugates of  $\mathcal{P}$  has smaller isotopics then
4       output  $\mathcal{P}$  as class representative
5       return
6     else
7       return
8   for every candidate universal  $c$  do
9     if  $c$  preserves orthogonality and is valid by Theorem 1 (c) then
10      if  $\mathcal{P} \cup c$  has no smaller isotopic  $k$ -MOLS then
11        enumerateMOLS( $\mathcal{P} \cup c$ )

```

to both the row and column indices. However, for $u_0(1)$ with cycle structure $z_1 z_2^{n_2} \dots z_p^{n_p}$ to be unaffected by the permutations, the set of potential permutations is restricted to only the $\prod_{i=1}^{i \leq p} i^{n_i} n_i!$ permutations that are found by rotating and reordering the cycles of $u_0(1)$. This step of the enumeration process, referred to in line 10 of Algorithm 1, is called the `isSmallest` test. If such a smaller partial MOLS is found, the node representing $\mathcal{P} \cup c$ becomes inactive and the next candidate universal is inspected for insertion into \mathcal{P} . Otherwise, a new list of candidate universals are generated for insertion into $\mathcal{P} \cup c$ and the search restarts one level lower down the tree. Whenever there are no more candidate universals to inspect, the search returns to the previous level. For a completed k -MOLS \mathcal{P} on level $n - 1$, the mappings and transformations described above are performed on all of the conjugates of \mathcal{P} to confirm that none of these conjugates have a lexicographically smaller isotopic k -MOLS than \mathcal{P} .

This enumeration process for 2-MOLS of order 5 is represented in Figure 2 (the same example may be found in [12]). According to Theorem 1, $u_0(0)$ must be the identity permutation and $u_0(1)$ a cycle structure representative, of which there are two possibilities for order 5, namely $z_1 z_2^2$ and $z_1 z_4$ (note that there must be exactly one 1-cycle to ensure orthogonality with the identity permutation). Two partial k -MOLS are said to be in the same *section* of the search tree if the respective $u_0(1)$ universals are the same cycle structure representative; the enumeration of 2-MOLS of order 5 therefore consists of two sections. Where branches become inactive it is indicated that either (a) no candidate universals preserve orthogonality, (b) a lexicographically smaller partial MOLS has been found in the same main class, or (c) a class representative has been found. One 2-MOLS is found in the section of $z_1 z_1^2$ and no structurally different 2-MOLS is found in the section of $z_1 z_4$, although a completed candidate 2-MOLS is uncovered which, upon inspection, is shown to be in the same main class as the first one, but lexicographically larger.

The known results in Table 1 were replicated in a validation attempt and details on the enumeration results for 3-MOLS of order 8 are given in Table 2. The number of active nodes found on every level is identical to that found by Kidd [12], while the serialized runtime has been improved from approximately 36 days to just under 10 days, although this improvement may be partially due to the use of different computing platforms. There are 45 active nodes on level 0 (after all of the zero universals have been inserted and an `isSmallest` test has been performed) and these nodes were given as the starting positions from which all of the subtrees were enumerated. It was found that there are 259 and 1700 active nodes on level 0 of the search trees for orders 9 and 10, respectively, which may be partitioned into 7 and 8 sections, respectively. Interestingly, the runtime increased from 6 seconds for the enumeration of 3-MOLS of order 7 to just under 10 days for the 3-MOLS of order 8, raising serious concerns over the feasibility of the enumeration of 3-MOLS of order 9 and higher.

Section	Level							Time (s)	
	0	1	2	3	4	5	6		7
$z_1 z_2^2 z_3$	17	12 501 028	1 484 518 094	18 814 494	55	23	22	20	775 321
$z_1 z_2^1 z_5$	14	3 358 273	61 708 802	63 157	97	92	84	17	60 011
$z_1 z_3 z_4$	5	52 059	5 283	1	0	0	0	0	93
$z_1 z_7$	9	37 403	9 079	82	64	53	53	2	111
Total	45	15 948 763	1 546 241 258	318 877 734	216	168	159	39	835 537

Table 2: The number of active nodes in every section and on every level of the search tree for enumerating 3-MOLS of order 8, together with the time in seconds that the enumeration of every section took on a 3.2 GHz processor with 8 Gb of RAM.




Figure 2: The backtracking enumeration search tree for 2-MOLS of order 5. At every leaf it is either indicated that (a) no candidate universals preserve orthogonality, or that (b) a lexicographically smaller partial MOLS has been found in the same main class, or that (c) a class representative has been found.

4 On the enumerability of larger order search spaces

In order to determine the feasibility of enumerating 3-MOLS of orders 9 and 10, the algorithm was modified so that it only examines MOLS that are isotopic to a partial MOLS \mathcal{P} after universals of the i -th symbol have been inserted into every Latin square in \mathcal{P} . Although this increases the total number of branches of the search tree that survive to level i , it decreases the total number of `isSmallest` tests performed during the enumeration, as all branches that would otherwise have been pruned earlier must necessarily have been subjected to at least one `isSmallest` test. Furthermore, the effect on the search tree as a whole is minimised, as the exact same number of branches will pass the `isSmallest` and proceed to the next symbol. The sizes of the subsequent search trees for orders 9 and 10 were approximated by estimating the total number of nodes in the absence of the `isSmallest` test before applying the expected pruning effect of the `isSmallest` test to determine the number of active nodes on every level of the tree. Finally, a small number of nodes from one of these levels were used as starting points for the enumeration algorithm so that the total time it would take to traverse the entire trees could be estimated. The enumeration tree for 3-MOLS of order 8 was also traversed to determine the average number of universals that preserve orthogonality and are valid by Theorem 1 (c), *i.e.* the universals that pass the test on line 9 of Algorithm 1, for partial 3-MOLS on different levels of the search tree.

It was found that that this average number of feasible candidate universals, which corresponds to the number of children of a node representing any partial 3-MOLS on level $i.l$ for $l \in \mathbb{Z}_{k-1}$, depends sensitively on the cycle structure of $u_0(1)$, but remains largely constant within a given section of the tree. Evidence of this may be seen for the 45 active nodes on level 0 of the enumeration tree for 3-MOLS of order 8 in Figure 3 for the two sets of universals $u_1(j)$ and $u_2(j)$ with $j \in \mathbb{Z}_k$. Notice in the figure, that the average number of feasible candidate solutions decreases with every additional universal in \mathcal{P} as it becomes harder to preserve orthogonality. This regularity in the number of children of a node of the search tree, as well as its sensitive dependence on the cycle structure of $u_1(0)$ was also observed in the search trees for 3-MOLS of orders 7, 9 and 10.

These properties make it possible to estimate the average number of children of any partial 3-MOLS by only examining a very small random selection of partial 3-MOLS that are on the

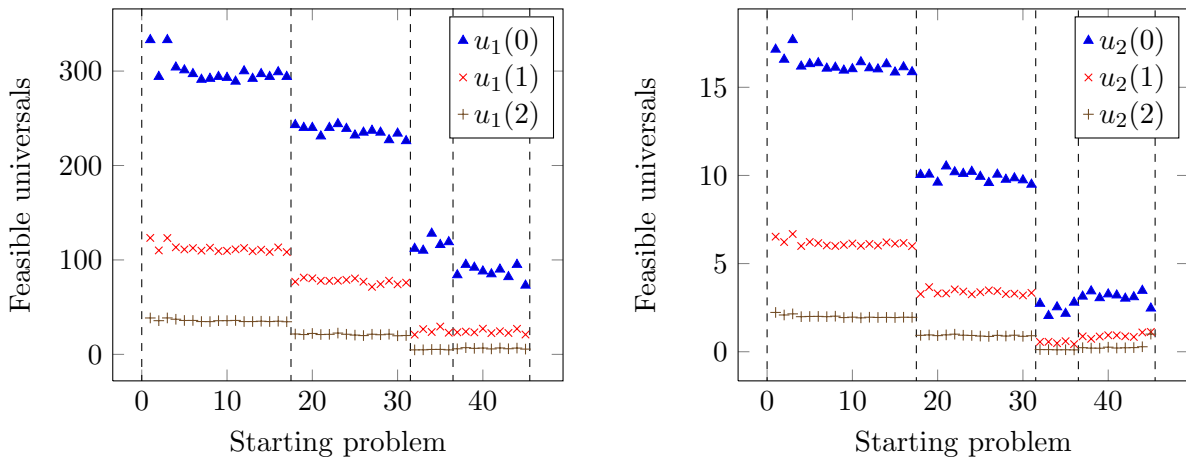


Figure 3: The average number of feasible candidate universals $u_i(j)$ found for $i = 1, 2$ and $j \in \mathbb{Z}_k$ in the enumeration of 3-MOLS of order 8 for each of the 45 partial 3-MOLS which pass the `isSmallest` test on level 0 of the search tree. The dashed lines indicate in which section the starting position resides, *i.e.* whether the permutation $u_0(1)$ in the initial partial 3-MOLS has the cycle structure $z_1 z_2^2 z_3$, $z_1 z_2 z_5$, $z_1 z_3 z_4$ or $z_1 z_7$, in that order.

	Order 8		Order 9	Order 10	n	6	7	8
	Actual	Estimated	Estimated	Estimated				
Level 1	2.61×10^7	2.60×10^7	5.79×10^{10}	2.41×10^{14}	Level 0	0.15	0.07	0.03
Level 2	4.34×10^9	3.74×10^9	3.39×10^{15}	9.67×10^{21}	Level 1	0.55	0.48	0.57
Level 3	9.96×10^8	9.31×10^8	2.15×10^{16}	—	Level 2	0	0.54	0.51

Table 3: A comparison of the actual and estimated total number of nodes on levels 0, 1, 2 and 3 of the search tree for 3-MOLS of order 8, together with similar estimates for orders 9 and 10.

Table 4: The proportions of nodes which pass the `isSmallest` test on levels 0, 1 and 2 for 3-MOLS of orders 6, 7 and 8.

same level and in the same section of the tree. This process was repeated on every level of the tree in order to estimate the total number of nodes in the search tree for 3-MOLS of orders 8, 9 and 10. This estimate proved to be fairly accurate for order 8, as may be seen in Table 3.

In order to estimate the number of active nodes on levels 1 and 2 of the search tree, the pruning effect of the `isSmallest` test must be applied to these estimated total numbers of nodes on every level of the tree. Let p_i denote the percentage of partial 3-MOLS which pass the `isSmallest` test on level i . The values of p_0 , p_1 and p_2 for orders 6, 7 and 8 may be seen in Table 4. Notice that less than 10% of the nodes on level 0 are active, and that this value is approximately 50% for levels 1 and 2. Based on this evidence, the numbers of active nodes on levels 1 and 2 of the search trees for orders 9 and 10 were estimated for three values of $p = p_1 = p_2$, specifically $p = 0.5$ together with expected over and under estimate values, $p = 0.4$ and $p = 0.6$. Note that the pruning effect is carried forward through the tree, *i.e.* if $p = 0.5$, then 50% of the nodes on level 1 are considered inactive, which implies that half the nodes on level 2 would not have been reached at all so that only 25% of the total number of nodes on level 2 are considered active. For order 9 the number of active nodes of level 1 (*i.e.* the number of partial 3-MOLS with all 0 and 1 universals filled in which pass the `isSmallest` test) is estimated to be between 2.32×10^{10} and 3.47×10^{10} , depending on the value of p , and for order 10 this number grows to approximately 1.21×10^{14} . The remainder of the estimated numbers of active nodes may be found in Table 5.

To gather insight into the potential total runtime of the enumeration algorithm for 3-MOLS of orders 9 and 10, a representative sample of active nodes on level 1 of the respective search trees was used as starting points for Algorithm 1, after which the number of active nodes was multiplied by the weighted average time to completion. To enable comparison between computing systems of different speeds the estimated time to completion is expressed in GHz-days, the number of days that a single 1GHz processor would take to complete the computation. It is expected that a complete enumeration of 3-MOLS of order 9 would take approximately 5.64×10^8 GHz-days, while for order 10 this is expected to take approximately 1.42×10^{18} GHz-days (these estimates may also be found in Table 5).

5 Conclusion

The serialized estimated enumeration time on a single 3.2 GHz core of 465 219 years for 3-MOLS of order 9, and 1.17×10^{14} years for order 10 is currently beyond the capabilities of most research computing clusters. For example, the high performance cluster, *Rhasatsha*, at Stellenbosch University currently consists of a hundred and thirty six 2.83 Ghz cores, thirty two 2.4 Ghz cores and three hundred and seventy six 2.1 GHz cores for a daily maximum throughput of approximately 1250 GHz-days. The performance of this cluster is dwarfed by, for example, the *Great Internet Mersenne Prime Search* (GIMPS), a distributed computing project which makes use of volunteers' computing power to find extremely large prime numbers

p	Order 9			Order 10		
	0.4	0.5	0.6	0.4	0.5	0.6
Level 1	2.32×10^{10}	2.89×10^{10}	3.47×10^{10}	9.65×10^{13}	1.21×10^{14}	1.44×10^{14}
Level 2	5.43×10^{14}	8.48×10^{14}	1.22×10^{15}	1.55×10^{21}	2.42×10^{21}	3.48×10^{21}
Level 3	1.37×10^{15}	2.68×10^{15}	4.64×10^{15}	—	—	—
Time	4.51×10^8	5.64×10^8	6.77×10^8	9.11×10^{17}	1.42×10^{18}	2.05×10^{18}

Table 5: The estimated total number of active nodes on different levels of the search tree for the enumeration of 3-MOLS of orders 9 and 10, as well as the estimated time (measured in GHZ-days) that the enumeration would take.

[17] and Seti@Home, a distributed project examining large datasets for signs of extraterrestrial intelligence [2]. GIMPS has an average daily throughput of approximately 100 000 GHZ-days [17], while SETI@Home averages 362 000 GHZ-days daily [2]. If the enumeration of k -MOLS were to take place with the computing power that is available to these distributed projects, the enumeration of 3-MOLS of order 9 would take approximately 15.5 years (at 100 000 GHZ-days daily) and it would be possible to answer the celebrated question of the existence of 3-MOLS of order 10 in approximately 3.9×10^{10} years. These estimates ignore the fact that the computing power of desktop computers, which are indispensable to distributed computing projects, have been estimated to double every 1.5 years over the last five decades [14, 18]. Assuming that this rate of growth continues, a distributed computing project of the scale of GIMPS may enumerate the 3-MOLS of order 9 in as little as 4.5 years, and the 3-MOLS of order 10 in approximately 51 years.

The enumeration of 3-MOLS of orders 9 and 10 therefore seems to be feasible as part of a long-term distributed, volunteer computing project. Moreover, the enumeration attempt would benefit greatly from a significant technical breakthrough in computing power or an important theoretical breakthrough (such as the design of a very effective pruning rule for the search tree or a speed-up of the `isSmallest` test). To resolve the question of the existence of 3-MOLS of order 10, however, it is only necessary to find a single 3-MOLS, making the estimated enumeration time a worst-case scenario that will only be reached if no 3-MOLS of order 10 exists.

References

- [1] ABDEL-GHAFFAR KAS, 1996. *On the number of mutually orthogonal partial latin squares*, *Ars Combinatorica* **42**, pp. 259–286.
- [2] ANDERSON DP, 2013. *SETI@Home*, [Online], [Cited 2013, June 4th], Available from: <http://setiathome.berkeley.edu/>.
- [3] BACHET CG, 1624. *Problèmes plaisants & délectables: Qui se font par les nombres*, Pierre Rigaud & Associates, Lyons.
- [4] BOSE RC, SHRIKHANDE SS & PARKER ET, 1960. *Further results on the construction of mutually orthogonal Latin squares and the falsity of Euler’s conjecture*, *Canadian Journal of Mathematics* **12**, pp. 189–203.
- [5] BURGER AP, KIDD MP & VAN VUUREN JH, 2010. *Enumerasie van self-ortogonale Latynse vierkante van orde 10*, *Litnet Akademies (Natuurwetenskappe)* **1(1)**, pp. 1–22.
- [6] COLBOURN CJ & DINITZ JH, 2006. *Handbook of combinatorial designs*, Chapman & Hall/CRC, Boca Raton (FL).
- [7] DUKES P & HOWARD L, 2012. *Group divisible designs in MOLS of order ten*, *Designs, Codes and Cryptography*, pp. 1–9, [Online], [Cited 2013, June 4th], Available from: <http://link.springer.com/article/10.1007/s10623-012-9729-8>
- [8] EULER, L, 1782. *Recherches sur une nouvelle espèce de quarrés magiques*, *Verhandelingen uitgegeven door het zeeuwsh Genootschap der Wetenschappen te Vlissingen* **9**, pp. 85–239.

- [9] FISHER RA, 1925. *Statistical methods for research workers*, Oliver & Boyd, Edinburgh.
- [10] FISHER RA, 1935. *The design of experiments*, Oliver & Boyd, Edinburgh.
- [11] KEEDWELL AD, 2000. *Designing tournaments with the aid of Latin squares: A presentation of old and new results*, *Utilitas Mathematica* **58**, pp. 65–86.
- [12] KIDD MP, 2012. *On the existence and enumeration of sets of two or three mutually orthogonal Latin squares with application to sports tournament scheduling*, PhD Dissertation, Stellenbosch University, Stellenbosch.
- [13] KIDD MP, 2010. *A tabu-search for minimising the carry-over effects value of a round-robin tournament*, *ORiON* **26(2)**, pp. 125–141.
- [14] KOOMEY JG, BERARD S, SANCHEZ M, WONG H, 2011. *Implications of historical trends in the electrical efficiency of computing*, *IEEE Annals of the History of Computing* **33(3)**, pp. 46–54.
- [15] LAM CWH, THIEL L & SWIERCZ S, 1989. *The non-existence of finite projective planes of order 10*, *Canadian Journal of Mathematics* **41(6)**, pp. 1117–1123.
- [16] LAYWINE CF & MULLEN GL, 1998. *Discrete mathematics using Latin squares*, Wiley, New York (NY).
- [17] MERSENNE RESEARCH, INC., 2013. *Great Internet Mersenne Prime Search*, [Online], [Cited 2013, June 4th], Available from: <http://www.mersenne.org/>.
- [18] NORDHAUS WD, 2007. *Two centuries of productivity growth in computing*, *Journal of Economic History* **67(1)**, pp. 128–159.
- [19] OWENS PJ & PREECE DA, 1995. *Complete sets of pairwise orthogonal Latin squares of order 9*, *Journal of Combinatorial Mathematics and Combinatorial Computing* **18**, pp. 83–96.
- [20] ROBINSON DF, 1981. *Constructing an annual round-robin tournament played on neutral grounds*, *Mathematics Chronicle* **10**, pp. 73–82.