Graphs and Combinatorics (2004) 20:545–570
Digital Object Identifier (DOI) 10.1007/s00373-004-0567-7

Graphs and Combinatorics

Cycle Switches in Latin Squares

Ian M. Wanless*

Department of Computer Science, Australian National University, ACT 0200, Australia
e-mail: imw@cs.anu.edu.au

Abstract. Cycle switches are the simplest changes which can be used to alter latin squares, and as such have found many applications in the generation of latin squares. They also provide the simplest examples of latin interchanges or trades in latin square designs.

In this paper we construct graphs in which the vertices are classes of latin squares. Edges arise from switching cycles to move from one class to another. Such graphs are constructed on sets of isotopy or main classes of latin squares for orders up to and including eight. Variants considered are when (i) only intercalates may be switched, (ii) any row cycle may be switched and (iii) all cycles may be switched.

The structure of these graphs reveals special roles played by N_2, pan-Hamiltonian, atomic, semi-symmetric and totally symmetric latin squares. In some of the graphs parity is important because, for example, the odd latin squares may be disconnected from the even latin squares.

An application of our results to the compact storage of large catalogues of latin squares is discussed. We also prove lower bounds on the number of cycles in latin squares of both even and odd orders and show these bounds are sharp for infinitely many orders.

1. Introduction

The idea behind this paper is to construct and study graphs which show how the latin squares of a certain order are connected by cycle switches. In this section we give a number of background definitions and references. Readers who are familiar with terminology of latin squares and graphs may prefer to skip to the next section, which will explain the specifics of the problem to be investigated.

A *latin square of order n* is a matrix of order n in which each one of n symbols appears exactly once in each row and exactly once in each column. In this paper we usually assume that the symbol set is $[n] = \{1, 2, \ldots, n\}$, so that it coincides with the set of indices of the rows and columns. It is sometimes convenient to think of a latin square of order n as a set of n^2 triples of the form (row, column, symbol). The latin property means that distinct triples never agree in more than one coordinate. For each latin square there are six conjugate squares obtained by

* This work was undertaken while the author was employed by Christ Church, Oxford, UK

uniformly permuting the coordinates of each triple. The square itself is one of these conjugates, obtained by applying the identity permutation to the triples. A square is said to be *totally symmetric* if it is equal to all six of its conjugates. If a square equals three (or six) of its conjugates then it is said to be *semi-symmetric*. The three conjugates which equal a semi-symmetric (but not totally symmetric) square are necessarily the square itself and the two squares obtained by cyclically permuting its triples.

An *isotopy* of a latin square L is a permutation of its rows, permutation of its columns and relabelling of its symbols. The resulting square is said to be *isotopic* to L and the set of all squares isotopic to L is called an *isotopy class*. An isotopy which maps L to itself is called an *autotopy* of L. The *main class* of L is the set of squares which are isotopic to some conjugate of L. For orders up to seven, representatives of the main classes of latin squares can be found in [4]. Throughout this paper we shall refer to that catalogue using the notation $M_{o.n}$ to denote the n^{th} main class of order o. For order 8 there are 283657 main classes and these can be downloaded from, for example, [14]. The number of isotopy classes of each order is also given in [4], although the number for order seven should be 564 not the quoted 563, which repeats an earlier misprint in [5].

A *latin rectangle* with symbol set S is a matrix in which each symbol in S occurs exactly once in each row and at most once in each column. A *latin subrectangle* is a submatrix which is a latin rectangle. If R is a $2 \times n$ latin subrectangle of some latin square and R is minimal in that it contains no $2 \times n'$ latin subrectangle for $2 \leq n' < n$, then we say that R is a *row cycle* of length n.

Another way to think of row cycles is in terms of the permutation which maps one row to another row. Suppose that r and s are two rows of a latin square. We define a permutation $\rho : [n] \mapsto [n]$ by $\rho(L_{rj}) = L_{sj}$ for each $j \in [n]$. Each row cycle between r and s corresponds to a cycle of the permutation ρ and vice versa. If γ is a cycle of ρ then we find the corresponding row cycle by taking all occurrences in r and s of symbols which occur in γ.

Column cycles and symbol cycles can be defined similarly to row cycles, and the operations of conjugacy interchange these objects. A column cycle is a set of entries which forms a row cycle when the square is transposed. A symbol cycle is a set of entries which forms a row cycle when the square is conjugated to exchange rows and symbols. Row cycles, column cycles and symbol cycles will collectively be known as cycles.

A cycle which has length equal to the order of the square is said to be *Hamiltonian*. A latin square is *pan-Hamiltonian* if every row cycle is Hamiltonian. A latin square is *atomic* if all of its conjugates are pan-Hamiltonian. In other words, a square is atomic if all of its cycles are Hamiltonian. This terminology comes from [21], in which both pan-Hamiltonian and atomic squares are investigated in some detail. In that paper it is reported that for orders up to 10 the only main classes of atomic squares are the cyclic group tables of prime orders. There is one additional main class ($M_{7.6}$) containing a pan-Hamiltonian latin square of order 7 and 37 main classes containing pan-Hamiltonian latin squares of order 9. It has since been shown in [16] that there are exactly 7 main classes of atomic square of order 11 and that these contain 17 isotopy classes. Note that pan-Hamiltonian latin squares (and hence also atomic squares) of order n can exist only if n is odd or $n = 2$. An infinite family of pan-Hamiltonian latin squares is constructed in [3].

Pan-Hamiltonicity is an isotopy invariant and the atomic property is a main class invariant. The original interest in pan-Hamiltonian latin squares arose from the fact that they contain no proper subsquares. In fact, they contain no non-trivial latin subrectangles. Latin squares without proper subsquares are called N_∞ squares. A weaker property, labelled N_2, is the absence of intercalates (subsquares of order two). For more information on these concepts consult [6]. Note that, as well as being subsquares, intercalates are row cycles, column cycles and symbol cycles of length two.

Throughout this work, we will use basic terminology of graph theory which can be found in any good book on the subject. In particular, *components* are maximal connected subgraphs. The *eccentricity* of a vertex v is the maximum over the distances from other vertices to v, the *diameter* is the maximum eccentricity and the *radius* is the minimum eccentricity among the vertices. Vertices achieving the radius are called *central*. A *clique* is a complete subgraph, and the *clique number* is the order of the largest clique. The *girth* is the length of the shortest cycle. A *pendant vertex* is a vertex of degree 1.

2. Cycle Switching

Any cycle in a latin square L can be switched to create a slightly different Latin square L'. To switch a row cycle we move each entry in the cycle from the row it inhabits in L to the same column of the other row of the cycle. Thus, if r and s are the two rows involved and C is the set of columns involved in the cycle then the result of the switching is defined by

$$L'_{ij} = \begin{cases} L_{sj} & \text{if } i = r \text{ and } j \in C, \\ L_{rj} & \text{if } i = s \text{ and } j \in C, \\ L_{ij} & \text{otherwise.} \end{cases}$$

Switching a column cycle involving two columns c and d and a set of rows R is similar, with the result being

$$L'_{ij} = \begin{cases} L_{id} & \text{if } j = c \text{ and } i \in R, \\ L_{ic} & \text{if } j = d \text{ and } i \in R, \\ L_{ij} & \text{otherwise.} \end{cases}$$

Switching a symbol cycle on two symbols σ_1 and σ_2 is achieved by replacing every occurrence of σ_1 in the cycle by σ_2 and vice versa. Of course these three operations are essentially the same thing, related by conjugacy of the square. Examples of cycle switching are given in Fig. 1.

Other local perturbations for producing new latin squares from old are discussed in [1], [8], [12] and [19]. These methods are all strong enough to convert any

$$\begin{array}{ccccc} 1 & 2 & 3 & 4 & 5 \\ 2 & 1 & 5 & 3 & 4 \\ 3 & 5 & 4 & 2 & 1 \\ 4 & 3 & 1 & 5 & 2 \\ 5 & 4 & 2 & 1 & 3 \end{array} \quad \xrightarrow{i} \quad \begin{array}{ccccc} 1 & 2 & 3 & 4 & 5 \\ 5^* & 1 & 2^* & 3 & 4 \\ 3 & 5 & 4 & 2 & 1 \\ 4 & 3 & 1 & 5 & 2 \\ 2^* & 4 & 5^* & 1 & 3 \end{array}$$

$$\downarrow r \qquad\qquad \searrow^{s}$$

$$\begin{array}{ccccc} 1 & 2 & 3 & 4 & 5 \\ 2 & 5^* & 4^* & 3 & 1^* \\ 3 & 1^* & 5^* & 2 & 4^* \\ 4 & 3 & 1 & 5 & 2 \\ 5 & 4 & 2 & 1 & 3 \end{array} \qquad\qquad \begin{array}{ccccc} 1 & 3^* & 2^* & 4 & 5 \\ 2 & 1 & 5 & 3 & 4 \\ 3 & 5 & 4 & 2 & 1 \\ 4 & 2^* & 1 & 5 & 3^* \\ 5 & 4 & 3^* & 1 & 2^* \end{array}$$

Fig. 1. Examples of cycle switching.The other three squares are obtained from the top left hand latin square by an intercalate switch i, a row cycle switch r and a symbol cycle switch s. In each case the entries altered by the switch are marked with an asterisk

latin square into any other latin square after a finite number of applications, although in both [8] and [12] some of the intermediate stages are not latin squares. To see that cycle switches do not have the capacity to perform arbitrary changes, it is sufficient to consider that atomic squares can never be converted to non-atomic squares by cycle switching. However, cycle switches do have the advantage of conceptual and algorithmic simplicity. This paper is, in part, an effort to investigate the price of this simplicity.

Since we will require them later (in Proposition 2), we now describe the perturbations defined by Pittenger [19]. Suppose that in a latin square L we have $L_{ij} = L_{kl} = a$ and $x = L_{il} \neq L_{kj} = y$ for distinct symbols a, x, y, rows i, k and columns j, l. We then attempt to form a new latin square L' by the following algorithm.

> Interchange the symbols a and x in row i.
>
> Interchange the symbols y and a in row k.
>
> Let $c = j$.
>
> Repeat
>
>   Let r be the row of L in which symbol x occurs in column c.
>
>   Let c be the column of L in which symbol y occurs in row r.
>
>   Interchange the symbols x and y in row r.
>
> Until $c \in \{j, l\}$.

If $c = j$ on termination then the above procedure fails to achieve our objective. On the other hand, if $c = l$ when the procedure terminates then the result is a new latin square L'. We say that L' is reached from L by a *Pittenger move*. A particular choice of i, j, k, l, a, x, y will be successful in creating a Pittenger move if and only if the triples (i, l, x) and (k, j, y) of L belong to the same symbol cycle. Pittenger showed that by a sequence of his moves and/or symbol cycle switches it is possible to convert any latin square into any other latin square of the same order.

In design theory terms, Pittenger moves and cycle switches are based on *trades* (see e.g. [4] for a definition of this concept). The phrase *latin interchange* [8] is sometimes used for trades in latin square designs. Cycles in some sense provide the minimal examples of latin interchanges. If two entries must be swapped in a latin square then completing a cycle switch represents the simplest way to then restore the latin property. The entries which end up being changed form a trade, as do the entries which replace them.

As might be expected, cycle switches have been used many times to build new latin squares from old, and the terminology employed has varied considerably. Norton [17] in a paper we will discuss in §3, dealt primarily with intercalates and referred to "intercalate reversals". Parker [18] "turned" intercalates in an order ten square without transversals to create a square with 5504 transversals and 12265168 orthogonal mates [16]. Elliott and Gibbons [9] used cycle switches, which they referred to as "rotations", in a simulated annealing approach to constructing N_∞ latin squares. Jacobson and Matthews [12] call the operation a "cycle swap", whereas Pittenger [19] calls a symbol cycle switch a "name-change". It is obvious then, that cycle switching is useful when trying to enumerate latin squares, or when trying to build, by local improvement, latin squares with special properties.

A related application is to the storage of vast numbers of latin squares. Even when restricted to special classes, the number of latin squares typically grows very quickly with the order of the square, so catalogues can require vast storage facilities. Instead of storing the entirety of each square, it is better to use an idea we call *chains*. In a chain, we store only the first square and for subsequent squares, we store instructions on how to make the square from its predecessor. In the ideal case these instructions would just specify a single cycle switch, which can be designated in significantly less space than it takes to store a new square. Performing the cycle switch may also be faster than loading a new square. Of course, there are new problems created by this approach, such as deciding on an order for the catalogue which makes each square similar to its predecessor. However, many construction methods will achieve this more or less automatically. A second problem is that on average half the catalogue must be processed in order to get to a random element. This can be alleviated by breaking the catalogue into a number of smaller chains. The storage saving is still substantial, yet fewer squares need to be generated in order to reach any given square. It would also be possible to avoid both the problems just mentioned by using a more general data structure. Instead of the path which underlies a chain, we could use a tree, with the edges still corresponding to cycle switches. It should be easy to build a tree with small radius to store a large number of latin squares.

With this motivation, we now embark on a study of graphs which model how latin squares are related by cycle switches. We call such graphs *switching graphs*. Our switching graphs use sets of latin squares as their vertices and have an edge between two sets if a member of the first set can be turned into some member of the second set by a cycle switch. Since cycle switches can be combined to produce any isotopy, but in general cannot be used to take conjugates of a square, the most natural classes of squares to use as vertices are the isotopy classes, although we do consider what happens when we use main classes as vertices. We also vary the exact operations which are allowed to produce an edge, considering three different possibilities. In §5 we study the problem where only intercalates may be switched. In §6 we allow any row cycle to be switched and in §7 we allow any cycle at all to be switched. For each of the above options we construct switching graphs for all orders of latin squares up to and including 8.

The graphs that we build in this paper are simple, meaning that we do not allow loops or multiple edges, and undirected because cycle switches can always be "undone". Using simple graphs was a matter of personal preference, but can be justified by practical considerations given the size of graphs encountered. Having said that, both loops and multiple edges would have a natural interpretation in switching graphs, since it is quite common for cycle switching to preserve the isotopy class of a latin square, or for two classes to be connected by several distinct cycle switches. The former case arises fairly trivially whenever the cycle being switched is Hamiltonian, but also arises in less trivial cases. Because of its trivial nature, the switching of Hamiltonian cycles will not be considered in the subsequent discussion. Sometimes, to emphasise this point, we will refer to the switching of a non-Hamiltonian cycle as a *non-trivial switching*. Although we do not put loops on our graphs we record in each instance the number of vertices which would have loops if we did (considering only those loops resulting from non-trivial switchings).

It should be clear that atomic latin squares will always result in isolated vertices in our switching graphs, since there are no operations allowed on them which might result in an edge. For the same reason, when switching just row cycles or intercalates the pan-Hamiltonian and N_2 squares respectively will give rise to isolated vertices.

There is a more interesting way to get isolated vertices, which is to have legal operations available but for none of them to reach a square outside the class being considered. We say that such a class of squares (or any member thereof) is *self-switching*; a label which depends on the context in terms of which operations are allowed. The smallest example of self-switching squares are the non-group based squares of order 5. There are only two main classes of order 5 and each contains a single isotopy class. One of these classes contains the cyclic group table. Since 5 is prime this class is atomic, from which it follows that any switching graph on order 5 squares must consist of just two isolated vertices. However, the non-group based squares have four intercalates and must therefore be self-switching whatever cycle switches are being allowed. So, for example, all four squares in Fig. 1 are isotopic to each other. This is the only self-switching main class known to the author, when the allowable operations include more than just the switching of intercalates. Hence for the remainder of the paper the phrase self-switching should be understood in the context of intercalate switches. There are no self-switching squares of order 6 or of orders lower than 5, but self-switching squares of orders 7 and 8 will be found in §5. The following four squares are examples of self-switching squares of order 9,

$$\begin{pmatrix} \mathbf{1} & \mathbf{2} & 7 & 8 & 6 & 5 & 4 & 9 & 3 \\ \mathbf{2} & \mathbf{1} & 9 & 5 & 3 & 7 & 6 & 4 & 8 \\ 9 & 5 & \mathbf{4} & \mathbf{3} & 7 & 8 & 1 & 6 & 2 \\ 7 & 8 & \mathbf{3} & \mathbf{4} & 2 & 9 & 5 & 1 & 6 \\ 6 & 4 & 2 & 7 & 8 & 1 & 3 & 5 & 9 \\ 5 & 7 & 8 & 9 & 1 & 6 & 2 & 3 & 4 \\ 3 & 6 & 5 & 1 & 4 & 2 & 9 & 8 & 7 \\ 4 & 9 & 6 & 2 & 5 & 3 & 8 & 7 & 1 \\ 8 & 3 & 1 & 6 & 9 & 4 & 7 & 2 & 5 \end{pmatrix} \qquad \begin{pmatrix} \mathbf{1} & \mathbf{2} & 7 & 9 & 6 & 8 & 3 & 4 & 5 \\ \mathbf{2} & \mathbf{1} & 8 & 7 & 9 & 5 & 4 & 6 & 3 \\ 7 & 9 & 3 & 5 & 8 & 4 & 1 & 2 & 6 \\ 8 & 7 & 6 & 4 & 3 & 9 & 2 & 5 & 1 \\ 9 & 6 & 4 & 8 & 7 & 1 & 5 & 3 & 2 \\ 5 & 8 & 9 & 3 & 2 & 7 & 6 & 1 & 4 \\ 3 & 4 & 1 & 2 & 5 & 6 & 8 & 7 & 9 \\ 6 & 3 & 5 & 1 & 4 & 2 & 7 & 9 & 8 \\ 4 & 5 & 2 & 6 & 1 & 3 & 9 & 8 & 7 \end{pmatrix}$$

$$\begin{pmatrix} 8 & 9 & 1 & 5 & 7 & 2 & 6 & 4 & 3 \\ 7 & 8 & 9 & 2 & 6 & 1 & 3 & 5 & 4 \\ 4 & 1 & 8 & 9 & 3 & 7 & 2 & 6 & 5 \\ 3 & 5 & 2 & 8 & 9 & 4 & 1 & 7 & 6 \\ 2 & 4 & 6 & 3 & 8 & 9 & 5 & 1 & 7 \\ 6 & 3 & 5 & 7 & 4 & 8 & 9 & 2 & 1 \\ 9 & 7 & 4 & 6 & 1 & 5 & 8 & 3 & 2 \\ 5 & 6 & 7 & 1 & 2 & 3 & 4 & \mathbf{9} & \mathbf{8} \\ 1 & 2 & 3 & 4 & 5 & 6 & 7 & \mathbf{8} & \mathbf{9} \end{pmatrix} \qquad \begin{pmatrix} \mathbf{1} & \mathbf{2} & 3 & 4 & 5 & 6 & 7 & 8 & 9 \\ \mathbf{2} & \mathbf{1} & 4 & 5 & 6 & 7 & 8 & 9 & 3 \\ 3 & 4 & 5 & 6 & 7 & 1 & 9 & 2 & 8 \\ 4 & 5 & 7 & 9 & 8 & 3 & 1 & 6 & 2 \\ 5 & 6 & 1 & 3 & 9 & 8 & 2 & 7 & 4 \\ 6 & 7 & 2 & 8 & 3 & 9 & 4 & 1 & 5 \\ 7 & 8 & 9 & 2 & 4 & 5 & 6 & 3 & 1 \\ 8 & 9 & 6 & 1 & 2 & 4 & 3 & 5 & 7 \\ 9 & 3 & 8 & 7 & 1 & 2 & 5 & 4 & 6 \end{pmatrix}$$

The first two squares are semi-symmetric and have trivial autotopy groups. Their intercalates are shown in **bold**. The third square has a unique intercalate and an autotopy of order 7 although it is not isotopic to any conjugate except itself. The fourth square represents one of the ten main classes of self-switching squares of order 9 which have a single intercalate and no non-trivial symmetry.

We now explain the relevance of various graphical characteristics to our problem. Clearly the components of the graph are important when we want to convert one square into another by a sequence of switches. The diameter and radius of a component tells us the number of switches which might be required for such a transformation. This is of interest for the storage application cited earlier in this section, as is the presence of a Hamiltonian path. More generally, we would be interested in a *path cover*, which is a set of disjoint paths which between them include every vertex. The minimum possible number of paths in a path cover tells us the fewest chains which can be used in a catalogue of the vertices. The degree of a vertex tells us how many essentially different squares can be produced by a single switching. The relevance of 4-cycles will be explained in §3.

The smallest order for which there are distinct isotopy classes is order 4. Here there are two main classes, each of which is also an isotopy class. The two classes are connected by intercalate switches, so every switching graph using order 4 squares consists of two connected vertices. As already mentioned, for order 5 the situation is similar except that the two vertices are not connected.

Thus order 6 is the smallest order for which switching graphs are interesting. As order 8 involves some significant computation and higher orders are well beyond reach, we will focus subsequent discussion on the orders 6, 7 and 8.

Our switching graphs have three parameters; namely, the order of the latin squares, whether we use isotopy or main classes as vertices and what cycle switches are allowed. We use the notation $G_{xy}(z)$ to denote a switching graph. In this notation $x \in \{i, r, a\}$ records whether intercalates (i), row cycles (r) or all cycles (a) are allowable switches, $y \in \{i, m\}$ says whether the vertices are isotopy classes (i) or main classes (m) and z is the order of the squares involved. It is worth making the following observations for general x, y and z. In each case $G_{xm}(z)$ can be formed from $G_{xi}(z)$ by identifying the vertices of the 1, 2, 3 or 6 isotopy classes which form each main class. Also, because the set of operations is being broadened at each step, $G_{iy}(z)$ is a subgraph of $G_{ry}(z)$ which in turn is a subgraph of $G_{ay}(z)$. From these observations a number of obvious inequalities follow, which we illustrate using diameter as an example. The diameter of $G_{xm}(z)$ cannot exceed that of $G_{xi}(z)$, while the diameter of $G_{ay}(z)$ cannot exceed that of $G_{ry}(z)$ which cannot exceed the diameter of $G_{iy}(z)$.

Each $G_{xy}(z)$ was constructed by a program we shall call program A. The input for program A was a set of representatives of the classes of latin squares which form the vertices of the graph. For each of these squares, the program identified all possible switches of the designated kind. The result of each switch was tested using McKay's *nauty* program [14] to see which vertex it belonged to. In this way, a list of adjacencies was built up, and this formed the output of program A. Once the graph was constructed it was analysed by program B to count its cliques and cycles, find the number and diameter of components *et cetera*. In most cases this analysis was straightforward. For example, we found the eccentricity of each vertex by finding the depth of a breadth-first search tree rooted at that vertex, and then calculated the diameter and radius of each component from these eccentricities. The one graph which was too large for this näive approach will be treated separately when it is encountered in §7.

One part of the analysis was carried out by a third program. Program C was written to find a minimal path cover for the connected graph which it took as its input. Essentially, it looked for a Hamiltonian cycle using a simple algorithm based on that advocated in [2] for the travelling salesman problem. However, before looking for such a cycle, various obstacles needed to be treated. The simplest obstacles were the pendant vertices, which were tackled as follows. Suppose there were $p > 0$ pendant vertices.

(1) If p was odd then an arbitrary pendant vertex v_0 was removed.
(2) A matching was added to the remaining pendant vertices (if any).
(3) A Hamiltonian cycle C was located in the resulting graph.
(4) Any edges which were not in the original graph were removed from C.
(5) If $p = 1$ then v_0 was reattached to its original neighbour, v_1, and an edge of C involving v_1 was dropped. For odd $p > 1$ we simply put v_0 back in the graph as a path of length zero.

There was no guarantee that step (3) would succeed, but if it did the result was a path cover achieving the trivial lower bound of $\lceil p/2 \rceil$ paths when $p > 0$. The case when $p = 0$ was handled simply by looking for a Hamiltonian cycle, from which any edge could be dropped to yield a path cover of size 1. Applying this process to each of the components of a graph produced a minimal path cover for the whole graph, assuming that each search for a Hamiltonian cycle was successful. Occasionally though, there were obstacles more complicated than pendant vertices. Another obstacle encountered was a cycle in which, within any pair of consecutive vertices, at least one vertex had degree two. We call such a subgraph a *cycle obstacle* (see Fig. 2). Each cycle obstacle must contain an end-point of a path in any path cover. Program C dealt with them by deleting an edge in the cycle between a vertex of degree two and a vertex of higher degree, then treating the resulting pendant vertex in the same way as other pendant vertices. Other obstacles will be described when the graphs in which they were encountered are discussed in sections §5 to 7.

3. Norton's Work

Norton was among the first to use cycle switching. We devote this section to a discussion of his seminal paper [17]. In it he coined the word "intercalate" to refer to subsquares of order 2, and the phrase "generalized intercalate" for what would nowadays be called latin subrectangles. He was aware that his intercalates and generalized intercalates could be switched to change the structure of latin squares and he exploited this in a highly dedicated attempt to enumerate the main classes of order seven. Starting with any latin square of order 7 the idea was to switch its intercalates, one at a time, and find all of the squares which could be generated in this way. The process could then be iterated from the new squares. Eventually this process must stop creating new main classes, at which point you have what Norton called a *family of species* (species being another name for main classes). Norton also defined a *domain* to be, in our terminology, the set of all main classes reachable by cycle switching from a given main class. He points out that cyclic group tables of prime order will be a domain on their own because, as we would say, they are atomic. He also knew of the other pan-Hamiltonian isotopy class of order 7, having found it while investigating squares orthogonal to the cyclic group

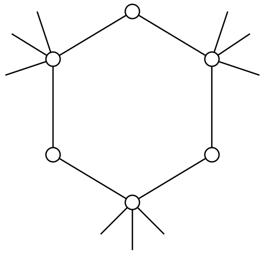

Fig. 2. A cycle obstacle of length six

table. He therefore knew of the only N_2 latin squares of order 7, which clearly must be families on their own.

What he did next was to manually compute an entire family of 144 main classes of latin squares containing intercalates. Together with the two N_2 classes, this massive effort came tantalisingly close to a complete catalogue, given that it included 146 of the 147 main classes. A representative of the missing class, which is $M_{7.147}$, was first pointed out by Sade [20]. We shall refer to the following semi-symmetric representative of $M_{7.147}$ as *Sade's square*:

$$\begin{pmatrix} \mathbf{1} & \mathbf{2} & 4 & 5 & 6 & 7 & 3 \\ \mathbf{2} & \mathbf{1} & 7 & 3 & 4 & 5 & 6 \\ 7 & 4 & 5 & 1 & 3 & 6 & 2 \\ 3 & 5 & 2 & 6 & 1 & 4 & 7 \\ 4 & 6 & 3 & 2 & 7 & 1 & 5 \\ 5 & 7 & 6 & 4 & 2 & 3 & 1 \\ 6 & 3 & 1 & 7 & 5 & 2 & 4 \end{pmatrix} \tag{1}$$

The unique intercalate of this square is shown in **bold**. The reason why Norton missed Sade's square is that it is a self-switching square belonging to a family of its own, as we will see in §5. Norton recognised that $M_{7.6}$ belongs to the same domain as his large family, and acknowledged that switching cycles of length greater than two may extend the domain further. With a few extra patient years of work he would have discovered that the Sade square belongs to the same domain as the other non-atomic squares. Of course the method cannot ever prove that a catalogue is complete, since it cannot rule out a separate undiscovered domain.

Despite narrowly failing to find a complete set of main classes, Norton's work does provide significant insight into the structure of latin squares. He made some basic observations which proved useful in writing our programs as well as providing a lot of data with which to corroborate the results. He noted that switching Hamiltonian cycles never changes the isotopy class, as we have already seen. He also argued that if a pair of rows, say, splits into just two cycles then we only need to try switching one of those cycles since switching the other will give an isotopic result.

Norton also pointed out a natural way for 4-cycles to arise in switching graphs. Suppose that a square L_1 contains two disjoint cycles γ and γ' which can be switched. Suppose further, that switching γ produces a square L_2, switching γ' produces a square L_3 and switching γ then γ' produces a square L_4. Since γ and γ' are disjoint we must also get L_4 if we start with L_1 and switch γ' then γ. If L_1, L_2, L_3 and L_4 represent distinct vertices then the result will be a 4-cycle on those four vertices. The point is that disjoint cycles can be switched independently. The final result will not depend on the order in which switching took place, but the intermediate stages may.

4. Parity

In this section we introduce ideas which will allow us to prove that many switching graphs are disconnected for a fundamental reason.

The notion of the parity for permutations is well understood. If we write any permutation as a product of transpositions then the number, modulo 2, of factors in this product is defined to be the *parity* of the permutation. Since latin squares are a 2-dimensional analogue of permutations, it is natural to define their parity.

A given row, say row r, of a latin square L defines a permutation $\sigma_r : [n] \mapsto [n]$ by $\sigma_r(j) = L_{rj}$. We say that σ_r is a *row permutation* and that it is the *permutation corresponding to row r*. Similarly, there is a permutation $\sigma_c : [n] \mapsto [n]$ corresponding to each column c defined by $\sigma_c(i) = L_{ic}$. There is also a permutation $\sigma_s : [n] \mapsto [n]$ corresponding to each symbol s defined by $\sigma_s(i) = j$ where $L_{ij} = s$.

The *row parity* of a latin square L is the sum, modulo 2, of the parities of the permutations corresponding to the rows of L. Similarly, the *column parity* of L is the sum, modulo 2, of the parities of the permutations corresponding to the columns of L and the *symbol parity* of L is the sum, modulo 2, of the parities of the permutations corresponding to the symbols of L. As conjugation permutes rows, columns and symbols it naturally permutes the row, column and symbol parities as well. For example, the column parity of L is the row parity of the transpose of L.

The row and column parities have been studied in pursuit of what is known as the Alon-Tarsi conjecture. One statement of this conjecture (see [4, p.108]) asserts that for each even order the number of latin squares with even row parity differs from the number of latin squares with odd row parity. It is known that for odd orders these two numbers are equal. Fundamental to the Alon-Tarsi conjecture is the observation that row, column and symbol parities are isotopy invariants for squares of even order, but not for squares of odd order. More generally we have:

Proposition 1. *Switching a row cycle of length l reverses the column and symbol parities if and only if l is odd, but never changes the row parity. Switching a column cycle of length l leaves the column parity unchanged and reverses the row and symbol parities if and only if l is odd. Switching a symbol cycle of length l leaves the symbol parity unchanged and reverses the row and column parities if and only if l is odd.*

Proof. We prove only the first statement, as the other two will then follow by conjugacy. Suppose that R is a row cycle of length l between the rows r_1 and r_2. To switch R we multiply the row permutation corresponding to r_1 by some permutation σ and multiply the row permutation corresponding to r_2 by the inverse permutation σ^{-1}. Since σ and σ^{-1} have the same parity we either change the parity of both rows or leave them both the same. Either way, the row parity is unchanged. Now, each of the l columns and l symbols involved in R has its corresponding permutation multiplied by a single transposition when R is switched. Hence the column and symbol parities both change by l modulo 2. □

There is another relationship between these three parities, other than the fact that conjugacy interchanges them. The following result was first proved by Janssen [13], and could also be deduced from results of Huang and Rota [11]. The proof that we outline is original and has a switching flavour appropriate to this work.

Proposition 2. *Let π_r, π_c and π_s be respectively the row, column and symbol parity of a latin square of order n. Then*

$$\pi_r + \pi_c + \pi_s \equiv \binom{n}{2} \bmod 2. \tag{2}$$

Proof. We first argue that (2) holds for one latin square of order n, namely the square T defined by $T_{ij} \equiv -i - j \bmod n$. Note that T is a totally symmetric form of the cyclic group table. Since T is totally symmetric $\pi_r = \pi_c = \pi_s$ and every row permutation of T is an involution. Any latin square of order n has to have exactly n fixed points among its row permutations, corresponding to the appearance in each column i of the symbol i. Hence there are $n^2 - n$ entries of T involved in transpositions in row permutations, and we have

$$\pi_r + \pi_c + \pi_s \equiv 3\frac{n^2 - n}{2} \equiv \binom{n}{2} \bmod 2.$$

This agrees with (2). Notice also that Proposition 1 guarantees that (2) is preserved by cycle switches. If (2) is shown to be preserved by Pittenger moves it will follow that (2) must hold for all latin squares of order n. Suppose that we have a Pittenger move P which involves t iterations of the repeat/until loop in the algorithm given in §2. A similar argument to Proposition 1 shows that P always reverses the symbol parity; reverses the row parity if and only if t is odd and reverses the column parity if and only if t is even. Hence (2) is unchanged by P. □

A latin square with respective row, column and symbol parities π_r, π_c and π_s will be referred to as an $(\pi_r\pi_c\pi_s)$-parity square. Proposition 2 tells us that for most orders there are four classes of latin squares according to their parity.

Proposition 3. *Let P_n be the set of parities of squares of order n. Then, provided $n = 3$ or $n \geq 5$,*

$$P_n = \begin{cases} \{(000), (011), (101), (110)\} & \text{if } n \equiv 0{,}1 \mod 4, \\ \{(111), (100), (010), (001)\} & \text{if } n \equiv 2{,}3 \mod 4. \end{cases} \tag{3}$$

Proof. Proposition 2 tells us that P_n is a subset of the set given in (3) for all n. For odd $n \geq 3$ it follows from Proposition 1 that all four possibilities will be achieved within each isotopy class of order n. For orders $n \geq 6$, we can then use [5, Thm 1.5.1], which says that there exists an order n square with a subsquare of order 3. Since there are four different possibilities for the parity of the subsquare the same must be true for the whole square. □

For order 2 and order 4 it is easily established that the only possible parities are (111) and (000) respectively. Since all the squares concerned are group based, [5, Thm 4.2.2] tells us that this in fact an example of a more general statement. We call L an *equal parity square* if the row, column and symbol parities of L are all equal.

Proposition 4. *If L is a latin square of even order and L is isotopic to at least 2 of its conjugates other than itself, then L is an equal parity square.*

Proof. Since L is of even order, the parities are isotopy invariant. Also, the condition on conjugates means there must be an isotopy which, in effect, cyclically permutes the roles of rows, columns and symbols. The only parities which are invariant under cyclic permutation are (000) and (111). □

In fact the idea of conjugacy yields the following directly:

Proposition 5. *For any order of latin squares the numbers of squares with parities (001), (010) and (100) are equal. Likewise the numbers of squares with parities (011), (101) and (110) are equal.*

At this stage we are not in a position to say how the number of (100)-parity squares relates to the number of (111)-parity squares or how the number of (011)-parity squares relates to the number of (000)-parity squares. To do so would require resolution of the Alon-Tarsi conjecture. However, loosely speaking we would expect them to be roughly equal since in each case switching an odd length row cycle in a square of one type creates a square of the other. Hence our next result predicts that many switching graphs will fragment into large disconnected pieces with roughly the same number of vertices. Formally, we define a *type* to be a set of vertices in a graph such that no edge joins a vertex in the set to a vertex not in the set.

Proposition 6. *Let $n \geq 6$ be an even integer. Then in $G_{ii}(n)$ the vertices are of four types corresponding to the parities given in Proposition 3. In $G_{im}(n)$ the vertices are of two types with the discrimination based on whether or not the squares have equal parity. In $G_{ri}(n)$ the squares form two types of vertices according to their row parity.*

Proof. This result follows directly from Proposition 1, by noting that an intercalate is a cycle of length 2, so switching it does not effect any parities. The reason $G_{im}(n)$ has only two types rather than four is that the non equal parity squares coalesce into one type when main classes are formed by merging isotopy classes. □

Of course there is nothing to say that vertices of a particular type must be connected. Indeed we have already seen classes of squares, such as atomic squares or self-switching squares, which give rise to isolated vertices. In fact, since N_2 squares are known [4, p.106] to exist for all orders except 2 and 4, we have:

Proposition 7. *Suppose that $G = G_{ii}(n)$ or $G = G_{im}(n)$. Then G is connected if and only if $n \leq 4$.*

Conjecture 1. *Proposition 7 holds for $G = G_{ri}(n)$ as well.*

Note that Proposition 6 proves Conjecture 1 for even orders. For odd orders it would be sufficient, but not necessary, to prove the conjecture in [21] that pan-Hamiltonian latin squares exist for all odd orders. For $G = G_{ai}(n)$ and

$G = G_{am}(n)$ we only know that G is disconnected for all orders ≥ 5 for which an atomic square exists. The set of orders for which atomic squares exist is known [21] to properly include the primes and be properly included (with the exception of $n = 2$) in the odd numbers. However it seems plausible that $G_{ai}(n)$ and $G_{am}(n)$ are connected for all even n. Certainly they are for even $n \leq 8$ as we shall see in §7. A relevant piece of evidence is this:

Proposition 8. *Let L be a latin square of even order n. Then L has at least* $\frac{1}{4}n(3n-4)$ *row cycles,* $\frac{1}{4}n(3n-4)$ *column cycles and* $\frac{1}{4}n(3n-4)$ *symbol cycles.*

Proof. By conjugacy,it is sufficient to prove the statement for row cycles. Suppose that e of the permutations corresponding to rows are even and the remaining $n-e$ are odd. Since n is even, a full cycle on n points is odd, so that Hamiltonian row cycles can only occur between rows of different parities. It follows that there are at least

$$\binom{n}{2} + \binom{e}{2} + \binom{n-e}{2}$$

row cycles. This quantity is minimised at $e = n/2$, where it is equal to the required bound. □

We next give a construction which shows that the bounds in Proposition 8 are achieved when $n = 2p$ where p is an odd prime. Let $h = \{1, 2, \ldots, p\}$ and $H = \{p+1, p+2, \ldots, 2p\}$. We use $(x)_S$ to denote that symbol in a set S which is congruent to x modulo p.

Let L be the latin square defined in four blocks as follows. The entry L_{ij} in row i and column j satisfies

$$L_{ij} = \begin{cases} (1-i-j)_h & \text{if } i, j \in h, \\ (i+j-1)_H & \text{if } i \in h \text{ and } j \in H, \\ (i-j+1)_H & \text{if } i \in H \text{ and } j \in h, \\ (1-i+j)_h & \text{if } i, j \in H. \end{cases}$$

It is routine to check from this definition that L is semi-symmetric and hence has the same number of row, column and symbol cycles. Also, since each of the four blocks of L is isotopic to the cyclic group of order p we know that any pair of rows chosen both from h or both from H will decompose into two cycles of length p. It suffices then, to show that for $a \in h$ and $b \in H$ the row cycle between rows a and b is Hamiltonian.

Suppose we start in column $c \in h$ of row a at $(1-a-c)_h$ and trace out the cycle from there. In row b we find $(b-c+1)_H$, which takes us to column $(2-a+b-c)_H$ of row a. In this column of row b the symbol is $(3-a-c)_h$. So when we return to the first block we are in column $c-2$, two places to the left of where we started. Iterating this process and noting that p is odd, we see that in following the row cycle we will visit every other column before returning to column c. This shows the validity of our construction.

Note that the bounds in Proposition 8 are achieved for orders other that those demonstrated by this construction. For example, two of the three main classes of N_2 squares of order 8, as given by Denniston [7], achieve all three bounds. The third N_2 main class achieves any two of the bounds but not the third.

The corresponding bound for odd orders is of course $\binom{n}{2}$, which is achieved by the atomic squares. Even bearing in mind that we ignore one cycle if a pair of rows does not split into at least three cycles, it would appear that a square of even order will have at least $\frac{1}{4}n(n-2)$ useful row switches available, where an odd order square may have none. Similar statements hold for the column and symbol cycles. So, a first reaction might be to expect $G_{ai}(n)$ and $G_{am}(n)$ to have high minimum degree and to be connected. However, Proposition 6 shows that the abundance of row switches is not in itself enough to ensure connectedness. Indeed, an abundance of switches need not even ensure high degrees. Our next result shows this clearly, by considering the squares with the most switches of all. Heinrich and Wallis [10] proved that the only squares achieving the trivial upper bound of $\frac{1}{4}n^2(n-1)$ on the number of intercalates belong to the main class of E_d, the Cayley table of the elementary abelian 2-group of order $n = 2^d$, for some positive integer d. Every cycle in E_d is an intercalate.

Proposition 9. *Let $n = 2^d$ for some integer $d \geq 2$. Then the vertex corresponding to E_d has degree 1 in $G_{xy}(n)$ for each $xy \in \{ii, im, ri, ai, am\}$.*

Proof. The case $d = 2$ was covered in §2. If $d > 2$ then E_d contains two disjoint copies of E_2. Switching an intercalate in one of these copies will produce some square B which does not satisfy the quadrangle criterion [5] and hence does not belong to the main class of E_d. It follows that the degree of E_d in any of the switching graphs is at least one. However, the autotopy group of E_d acts transitively on the intercalates in E_d, so switching any intercalate produces a square isotopic to B.

To see this, note that the autotopy group of E_d is the semi-direct product of $T \times T$ by the general linear group $\mathrm{GL}(d, 2)$. Here T is the translation group, so that $T \times T$ acts transitively on the cells of E_d. The stabiliser of a cell c is $\mathrm{GL}(d, 2)$ which acts transitively on the rows not containing c. It follows that the autotopy group of E_d acts transitively on the intercalates and hence the degree of E_d is exactly one. □

5. Intercalate Switches

We begin our case studies with the graphs formed when intercalate switches are the only allowable moves, which are summarised in Table 1. This is the first of three tables in which we give the following graphical parameters: Number of vertices, number of edges, number of loops, number of components, number of isolated vertices, number of pendant vertices, minimum cardinality of a path cover, minimum degree, maximum degree, diameter, radius, clique number, number of maximum cliques, number of 3-cycles and number of 4-cycles. Note

Table 1. Summary of intercalate switching graphs

	Order 6		Order 7		Order 8	
	main	isotopy	main	isotopy	main	isotopy
Vertices	12	22	147	564	283657	1676267
Edges	8	18	455	2487	1953563	11679069
Loops	6	12	30	60	5365	20008
Components	4	6	4	6	11	32
Isolated vertices	1	1	3	4	7	15
Pendant vertices	6	10	7	11	46	167
Path cover	4	8	7	10	31	109
Min degree	0	0	0	0	0	0
Max degree	2	4	14	19	45	49
Radius	2*	2*	5*	5*	9*	10*
Diameter	4*	4*	9*	9*	15*	15*
Clique num	2	2	4	4	5	4
Max cliques	8	18	4	8	27	2089
Girth	Forest	4	3	3	3	3
3-cycles	0	0	81	262	28277	109654
4-cycles	0	3	928	6375	6007654	35409032

that we adopt a few conventions for our tables. Firstly, although we cite the number of loops, these were never actually added to the graph and hence are not counted in any other field. In particular, the loops are not counted in the number of edges, nor in the degrees of vertices. Similarly, multiple copies of an edge were never included in any graph. Secondly, for a graph that is disconnected the diameter and radius are fairly meaningless. Rather than quoting a value of ∞ we provide instead the largest values which these parameters attain for a component of the graph. Values of radius and diameter which have been determined by this rule are marked with stars (*) in the tables.

The only acyclic graph which we encountered in this work is $G_{im}(6)$, which is a linear forest. As Norton [17] reported, it has four components (we knew from Proposition 6, that it would have at least two). These turn out to be paths of lengths 0, 1, 3 and 4, as can be seen from the solid lines in Fig. 3. The path of length 3 contains all the non equal-parity squares of order 6, while the other three paths contain equal parity squares.

The graph $G_{ii}(6)$ will be shown by the solid lines in Fig. 5. It has six components (two more than the lower bound derived from Proposition 6), and contains three 4-cycles. Like $G_{im}(6)$ its largest component has diameter 4.

The graph $G_{im}(7)$ is of historical interest, given the work of Norton discussed in §3. It has four components, three of which are isolated vertices. The three isolated vertices correspond to the cyclic square, the other pan-Hamiltonian square of order 7 and the Sade square, which are respectively $M_{7.7}$, $M_{7.6}$ and $M_{7.147}$. Note that the Sade square is self-switching, while the

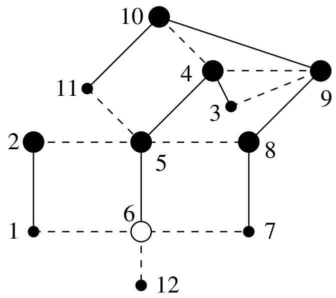

Fig. 3. The vertices of this graph are the main classes of order 6, numbered according to the catalogue in [4]. Solid edges represent intercalate switches and hence belong to the graph $G_{im}(6)$. The dotted lines indicate longer cycle switches between vertices that are not connected by an intercalate switch. Hence the dotted lines together with the solid lines show the edges of $G_{am}(6)$. The larger vertices are those for which a non-trivial switching takes the main class to itself. If there is such a switching which involves an intercalate switch the vertex is filled in, otherwise it is hollow

other two squares are N_2. The family with 144 vertices which Norton enumerated, has diameter 9. Norton's paper [17] contains a wealth of data with which to verify the accuracy of our program. The number of loops, and minimum and maximum degree of $G_{im}(7)$ were checked in this way and found to concur once multiple edges had been removed from Norton's data.

Moving to $G_{ii}(7)$ we find that there are 6 components. Four of these are isolated vertices corresponding to the N_2 squares. A fifth component is an isolated edge which involves the two isotopy classes which correspond to Sade's square (1). Again there is a giant component of diameter 9, this time with 558 vertices.

Proposition 6 predicts that $G_{im}(8)$ will have two types of vertices. The non equal-parity squares form a single component of order 207323. The equal-parity squares form ten components, with almost all the vertices belonging to a component of order 76322. Seven of the remaining components are isolated vertices, which correspond to the three N_2 main classes identified by Denniston [7] and to the following four self-switching squares:

$$S_1 = \begin{pmatrix} 1 & 2 & 4 & 3 & 6 & 5 & 8 & 7 \\ 2 & 1 & 5 & 7 & 3 & 8 & 4 & 6 \\ 4 & 5 & 3 & 1 & 2 & 7 & 6 & 8 \\ 3 & 7 & 1 & 4 & 8 & 6 & 2 & 5 \\ 6 & 3 & 2 & 8 & 5 & 1 & 7 & 4 \\ 5 & 8 & 7 & 6 & 1 & 4 & 3 & 2 \\ 8 & 4 & 6 & 2 & 7 & 3 & 5 & 1 \\ 7 & 6 & 8 & 5 & 4 & 2 & 1 & 3 \end{pmatrix}, \qquad S_2 = \begin{pmatrix} 1 & 2 & 4 & 3 & 6 & 5 & 8 & 7 \\ 2 & 1 & 5 & 6 & 7 & 8 & 4 & 3 \\ 4 & 8 & 3 & 1 & 2 & 7 & 5 & 6 \\ 3 & 7 & 1 & 5 & 4 & 2 & 6 & 8 \\ 6 & 3 & 7 & 4 & 8 & 1 & 2 & 5 \\ 5 & 4 & 8 & 7 & 1 & 6 & 3 & 2 \\ 8 & 5 & 6 & 2 & 3 & 4 & 7 & 1 \\ 7 & 6 & 2 & 8 & 5 & 3 & 1 & 4 \end{pmatrix}, \tag{4}$$

$$S_3 = \begin{pmatrix} 1 & 2 & 4 & 5 & 6 & 7 & 8 & 3 \\ 2 & 1 & 5 & 3 & 7 & 8 & 6 & 4 \\ 8 & 4 & 3 & 1 & 2 & 5 & 7 & 6 \\ 3 & 8 & 2 & 4 & 1 & 6 & 5 & 7 \\ 4 & 3 & 6 & 7 & 5 & 1 & 2 & 8 \\ 5 & 7 & 8 & 6 & 3 & 4 & 1 & 2 \\ 6 & 5 & 7 & 8 & 4 & 2 & 3 & 1 \\ 7 & 6 & 1 & 2 & 8 & 3 & 4 & 5 \end{pmatrix}, \qquad S_4 = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 \\ 2 & 1 & 4 & 5 & 6 & 7 & 8 & 3 \\ 3 & 8 & 5 & 1 & 7 & 2 & 4 & 6 \\ 4 & 6 & 7 & 8 & 3 & 5 & 2 & 1 \\ 5 & 7 & 1 & 6 & 4 & 8 & 3 & 2 \\ 6 & 5 & 2 & 7 & 8 & 3 & 1 & 4 \\ 7 & 3 & 8 & 2 & 1 & 4 & 6 & 5 \\ 8 & 4 & 6 & 3 & 2 & 1 & 5 & 7 \end{pmatrix}.$$

Squares S_1, S_2 and S_3 are semi-symmetric (in fact S_1 is totally symmetric). S_4 has no non-trivial symmetry, but flipping its unique intercalate produces a square isotopic to the transpose of S_4. The numbers of intercalates in S_1 to S_4 are 4, 1, 4 and 1 respectively. The mean number of intercalates for order 8 squares is approximately 14.0697 (see [15]).

The remaining two components are small. There is an isolated edge based on two more semi-symmetric squares. One of these is

$$\begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 \\ 2 & 1 & 4 & 5 & 6 & 7 & 8 & 3 \\ 3 & 8 & 1 & 2 & 7 & 4 & 6 & 5 \\ 4 & 3 & 6 & 1 & 2 & 8 & 5 & 7 \\ 5 & 4 & 8 & 7 & 1 & 2 & 3 & 6 \\ 6 & 5 & 7 & 3 & 8 & 1 & 2 & 4 \\ 7 & 6 & 5 & 8 & 4 & 3 & 1 & 2 \\ 8 & 7 & 2 & 6 & 3 & 5 & 4 & 1 \end{pmatrix}.$$

This square has seven intercalates all of which involve the entry in the top left corner. If any of these intercalates is switched the resulting square is also semi-symmetric, and has a unique intercalate.

Finally, there is a component which is a path of length 2, based on three latin squares, each of which has just two intercalates. The middle square of the path is as follows,

$$\begin{pmatrix} \mathbf{1} & \mathbf{2} & 7 & 5 & 4 & 3 & 6 & 8 \\ \mathbf{2} & \mathbf{1} & 6 & 8 & 5 & 7 & 4 & 3 \\ 7 & 5 & \mathbf{3} & \mathbf{4} & 6 & 8 & 1 & 2 \\ 6 & 8 & \mathbf{4} & \mathbf{3} & 1 & 2 & 5 & 7 \\ 3 & 7 & 5 & 1 & 2 & 6 & 8 & 4 \\ 4 & 6 & 8 & 2 & 3 & 5 & 7 & 1 \\ 8 & 3 & 1 & 6 & 7 & 4 & 2 & 5 \\ 5 & 4 & 2 & 7 & 8 & 1 & 3 & 6 \end{pmatrix} \tag{5}$$

and the two ends of the path can be found by flipping either one of the two intercalates, which are shown in **bold**. Flipping both intercalates produces a square isotopic to the one given in (5).

A minimal path cover for $G_{im}(8)$ has 31 paths. The nine small components require one path each. The component of order 76322 requires 15 paths, since it contains 29 pendant vertices and one more obstacle comprising a vertex whose four neighbours have degrees 1, 2, 2 and 10. The component of order 207323 requires 7 paths, since it contains 13 pendant vertices and a cycle obstacle of length 4.

Next we consider $G_{ii}(8)$, a graph on the 1676267 isotopy classes of order 8. Proposition 6 predicts four different types of vertices. The (011)-parity squares form a single component with 411313 vertices. This component has radius 10, diameter 15 and maximum degree 36. The (101)-parity squares and the (110)-parity squares each form components which are isomorphic to the (011)-parity component. The isomorphisms are provided by mapping each square to a particular conjugate of itself, and for that reason we call these components *conjugate components*.

There are 32 components in all. The equal-parity squares account for 29 of them with 15 of those being isolated vertices. The Denniston N_2 squares contribute 14 isolated vertices, and the remaining one comes from the totally symmetric self-switching square S_1 given in (4). There are seven isolated edges; two of which are conjugates based on the isolated edge in $G_{im}(8)$ and the remaining five are based on S_2, S_3 and S_4. Each of S_2 and S_3 contributes one isolated edge between the square itself and its distinct conjugate. Meanwhile, since flipping the unique intercalate in S_4 yields a square isotopic to the transpose of S_4, the six conjugates of S_4 contribute three isolated edges in $G_{ii}(8)$. There are also six conjugate components of order three derived from the component of $G_{im}(8)$ given in (5). This makes a total of 28 small components, using 47 vertices between them.

The remaining 442281 vertices form a single component with radius 10 and diameter 15. Its maximum degree is 49, which is achieved uniquely by the following square:

$$\begin{pmatrix} 1 & 3 & 4 & 2 & 6 & 5 & 7 & 8 \\ 4 & 2 & 1 & 3 & 7 & 8 & 6 & 5 \\ 2 & 4 & 3 & 1 & 5 & 7 & 8 & 6 \\ 3 & 1 & 2 & 4 & 8 & 6 & 5 & 7 \\ 6 & 8 & 5 & 7 & 3 & 1 & 2 & 4 \\ 5 & 7 & 8 & 6 & 1 & 4 & 3 & 2 \\ 7 & 5 & 6 & 8 & 4 & 2 & 1 & 3 \\ 8 & 6 & 7 & 5 & 2 & 3 & 4 & 1 \end{pmatrix} \tag{6}$$

This square has 52 intercalates and the obvious 4 subsquares of order 4. It has the most intercalates of any square of order 8 with a trivial autotopy group. However, it is semi-symmetric and isotopic to its transpose as well, so that it is isotopic to all its conjugates.

The component of order 442281 has 72 pendant vertices, so we know that any path cover contains at least 36 paths, but in fact the minimal path cover has 42 paths. There are 7 obstacles which account for this fact, the most complicated of which is shown in Figure 4. The other six are comprised of 3 cycle obstacles of length 6 (of the type pictured in Fig. 2) and 3 cycle obstacles of length 4. The subgraph shown in Fig. 4 cannot contain fewer than 7 endpoints of paths in any path cover (where an isolated vertex is interpreted as a path of length zero, with two endpoints). Since it contains two pendant vertices, this shows that a path

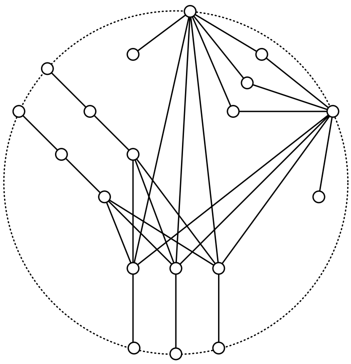

Fig. 4. Subgraph which creates obstacles to a path cover of the component of $G_{ii}(8)$ of order 442281. Vertices on the dotted circle have some neighbours not shown, but vertices inside the circle have all their adjacencies given explicitly

cover must contain at least $\lceil\frac{1}{2}(72 - 2 + 7 + 6)\rceil = 42$ paths. A path cover with this many paths was found by program C.

Each of the components of order 411313 contains two cycle obstacles of length 4 and a cycle obstacle of length 3 which is connected to the rest of the component by a single edge. There are also 23 pendant vertices. This means that any path cover must contain at least 13 paths, and a cover of this size was found by program C. Since there are 28 components of $G_{ii}(8)$ which are themselves paths, plus the four large components, we see that the minimal path cover for $G_{ii}(8)$ has $42 + 3 \times 13 + 28 = 109$ paths.

6. Row Cycle Switches

Next we consider graphs formed when the allowable operations are the switching of any row cycle. Naturally, these graphs are isomorphic to the graphs obtained by allowing just column cycle switches, or just symbol cycle switches. Since row cycles are a conjugacy dependent concept it does not make sense to use main classes as vertices in this section. If we did use main classes as vertices then we would in effect be allowing all cycle switches, which is the case treated in the next section. Also note that all intercalates are row cycles, so the allowed operations here are a proper superset of the operations allowed in the previous section. Therefore the graphs in this section have the corresponding graphs in the previous section as subgraphs.

Table 2 summarises the properties of the graphs formed by taking the isotopic classes of orders 6, 7 and 8 as the vertices.

For order six there are two components as illustrated in Fig. 5. This is our first graph achieving the lower bound prescribed by Proposition 6. The 14 vertices with odd row parity form a component with diameter 4 and radius 2. The odd row parity representative of $M_{6.5}$ is the unique central vertex. The 8 vertices with even row parity form a component with diameter 5 and radius 3. The smaller component contains a Hamiltonian path, but the larger component requires 3 paths to cover it.

Table 2. Summary of row cycle switching graphs

	Order 6 isotopy	Order 7 isotopy	Order 8 isotopy
Vertices	22	564	1676267
Edges	29	4121	23349579
Loops	12	131	27404
Components	2	3	2
Isolated vertices	0	2	0
Pendant vertices	3	2	1
Path cover	4	3	2
Min degree	1	0	1
Max degree	5	29	65
Radius	3*	4*	7*
Diameter	5*	6*	11*
Clique num	3	6	7
Max cliques	2	1	1
Girth	3	3	3
3-cycles	2	2240	7786904
4-cycles	9	20689	117996116

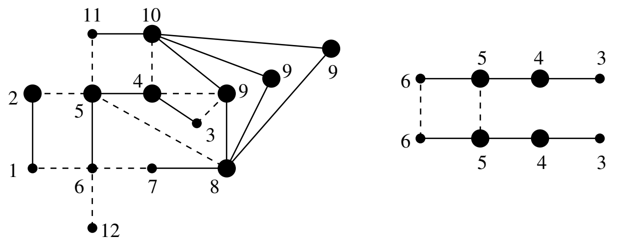

Fig. 5. The vertices of this graph are the isotopy classes of order 6, but numbered according to the catalogue of main classes in [4]. Solid edges represent intercalate switches, that is, the edges of $G_{ii}(6)$. The dotted lines indicate the edges gained by allowing row cycle switches. The dotted edges together with the solid edges form $G_{ri}(6)$. The larger vertices are those for which a non-trivial switching takes the isotopy class to itself (in each case where such a switching exists, it can be chosen to be an intercalate switching)

For order seven the graph has three components. Two of these are isolated vertices corresponding to the two pan-Hamiltonian squares of this order. The remaining 562 vertices form a single component with diameter 6, radius 4 and maximum degree 29. This contains a unique clique of order 6, which is based around two squares from each of $M_{7.9}$ and $M_{7.14}$ and one each from $M_{7.15}$ and $M_{7.60}$. Interestingly each of these squares possesses a subsquare of order 3 which is invariant under the row cycle switches which give rise to the clique. In $G_{ri}(7)$ these six isotopy classes clearly induce K_6, but in $G_{ii}(7)$ they induce $K_{2,4}$.

For order eight the graph again achieves the lower bound of two components given by Proposition 6. The smaller of the two has 822626 vertices of odd row-parity which include the vertices from two of the three components of order

411313 in $G_{ii}(8)$. This component has minimum degree 3, maximum degree 55, diameter 9 and contains a Hamiltonian cycle. The 853641 vertices of even row-parity form the second component. It has diameter 11, minimum degree 1, maximum degree 65 and contains a Hamiltonian path. The maximum degree is achieved uniquely by a close relative, call it R, of the square which maximised the degree in $G_{ii}(8)$. R can be formed by switching the intercalate which forms the four corners of the square given in (6). Like the original square, this R is semi-symmetric and isotopic to all its conjugates, despite having a trivial autotopy group. It has fewer subsquares than the square in (6), with only 45 intercalates.

Three of the seven vertices in the unique maximum clique of $G_{ri}(8)$ contain squares which, like R, are semi-symmetric, isotopic to their transpose and have trivial autotopy group. One of these squares is

$$X = \begin{pmatrix} 2_a & 1 & 3_b & 4 & 5_c & 7_d & 8 & 6 \\ 1_a & 2 & 4_b & 3 & 6_c & 5 & 7 & 8 \\ 3 & 4 & 1 & 2 & 7 & 8 & 6 & 5 \\ 4 & 3 & 2 & 1 & 8 & 6_d & 5 & 7 \\ 5 & 6 & 8 & 7 & 1 & 2 & 3 & 4 \\ 8_e & 5 & 7 & 6 & 2_f & 4 & 1 & 3 \\ 6 & 7 & 5 & 8 & 4_f & 3 & 2 & 1 \\ 7_e & 8 & 6 & 5 & 3 & 1 & 4 & 2 \end{pmatrix} \qquad (7).$$

Note that X has four subsquares of order 4. Representatives of the other six vertices in the clique can be reached by switching the row cycles a, b, c, d, e, f. Two entries from each of these row cycles have be labelled with the appropriate subscript in (7). In particular, these row cycles all lie inside one of the subsquares of X, which means that every square involved in the clique has four subsquares of order four. Just as for $G_{ri}(7)$, the vertices involved in the maximum clique induce a complete bipartite graph in the intercalate switching graph. Here, the seven vertices induce $K_{3,4}$ in $G_{ii}(8)$.

7. All Cycle Switches

Finally we consider graphs formed when the allowable operations are switches of any row, column or symbol cycle. The allowed operations here are a proper superset of the operations allowed in the previous two sections. Therefore the graphs in the previous sections are subgraphs of the corresponding graph in this section.

Taking the main classes of order 6 gives the graph as shown by taking all the edges (whether dotted or unbroken) in Fig. 3. The graph is connected with diameter 4 and radius 2. Vertex $M_{6.5}$ is the unique central vertex and also the unique vertex of maximum degree, which is 5. Note also that this graph has a Hamiltonian path. For example, take the vertices labelled 12,6,1,2,5,11,10,4,3,9,8,7 in that order.

If instead we take the isotopy classes of order 6 as vertices then both the minimum degree and the radius increase by one and the graph stays connected.

Table 3. Summary of graphs formed by switching all cycles

	Order 6		Order 7		Order 8	
	main	isotopy	main	isotopy	main	isotopy
Vertices	12	22	147	564	283657	1676267
Edges	17	51	1197	7260	7781572	46673268
Loops	7	12	97	227	17962	41241
Components	1	1	2	2	1	1
Isolated vertices	0	0	1	1	0	0
Pendant vertices	1	0	2	0	1	1
Path cover	1	1	2	2	1	1
Min degree	1	2	0	0	1	1
Max degree	5	7	32	47	95	108
Radius	2	3	3*	3*	5	6
Diameter	4	4	5*	5*	10	10
Clique num	3	4	7	7	8	9
Max cliques	2	6	1	3	97	1
Girth	3	3	3	3	3	3
3-cycles	2	26	1878	9583	5207986	29874415
4-cycles	5	57	20751	126186	90242587	518456141

Looking next at G_{am} (7) and G_{ai} (7) we find that in both graphs there are two components, but one of them is the isolated vertex corresponding to the cyclic group. Hence Norton would have found an exhaustive list of order 7 main classes if he had considered cycle switches more general than his intercalate reversals. He was aware of this possibility, but understandably daunted by the amount of work it would have involved.

Starting with Sade's square as given in (1), we can produce 145 more main classes of order 7 by the following sequence of switching operations:

*r*23 *c*47 *r*253 *r*17 *r*67 *c*47 *r*12 *r*27 *c*46 *r*26 *r*37 *r*17 *s*452 *r*15
*r*12 *s*36 *s*17 *r*23 *s*46 *s*17 *r*13 *r*26 *r*15 *c*15 *c*37 *r*56 *s*37 *r*14
*r*16 *r*27 *r*46 *s*12 *r*34 *c*45 *r*13 *r*12 *r*56 *c*35 *c*37 *r*46 *s*23 *c*67
*c*57 *r*566 *r*45 *c*15 *r*34 *c*172 *r*15 *c*13 *r*27 *c*45 *r*26 *r*24 *r*23 *c*253
*r*26 *r*57 *r*12 *s*35 *r*27 *c*35 *r*352 *r*15 *r*25 *s*23 *r*26 *r*12 *r*35 *r*45
*r*16 *s*12 *s*352 *s*12 *r*12 *s*34 *s*13 *r*143 *c*37 *s*34 *r*565 *r*16 *r*14 *r*17
*c*13 *c*15 *r*462 *r*472 *r*12 *r*352 *r*474 *r*16 *r*12 *r*16 *c*15 *r*234 *r*23 *s*23
*r*17 *s*34 *r*15 *r*27 *r*26 *c*123 *c*26 *r*23 *r*343 *r*37 *c*35 *r*45 *r*46 *s*47
*s*27 *r*37 *r*46 *s*37 *r*47 *r*16 *s*13 *s*16 *c*23 *r*23 *r*57 *r*37 *r*47 *c*24
*s*36 *c*16 *r*67 *s*47 *r*36 *c*56 *r*45 *r*67 *s*47 *c*35 *s*27 *s*16 *r*12 *r*13
*r*15 *r*14 *r*45 *c*37 *s*27

Here *rxyz* denotes the row cycle between rows x and y starting in column z. Likewise, *cxyz* is the column cycle between columns x and y starting in row z, and *sxyz* is the symbol cycle between symbols x and y starting in column z. For the majority of cycles the value of z is 1, in which case it has been omitted.

Finally we come the order 8 squares. In both $G_{am}(8)$ and $G_{ai}(8)$ the unique vertex of minimum degree is E_3, the elementary abelian 2-group. In both graphs it has degree 1 as predicted by Proposition 9, and the maximum eccentricity, which is 10.

$G_{ai}(8)$ is connected, and large enough that calculating the eccentricity of every vertex was not feasible using the straightforward algorithm which worked on all other graphs in this paper. However, it was still feasible to calculate the eccentricities of some vertices, so the radius and diameter were established as follows. We first computed that the eccentricity of E_3 is 10, and the only two vertices at distance 10 from it correspond to the square A and its transpose, where

$$A = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 \\ 2 & 1 & 4 & 5 & 6 & 7 & 8 & 3 \\ 3 & 8 & 1 & 2 & 7 & 4 & 6 & 5 \\ 4 & 3 & 6 & 1 & 2 & 8 & 5 & 7 \\ 5 & 4 & 8 & 7 & 1 & 2 & 3 & 6 \\ 6 & 5 & 7 & 3 & 8 & 1 & 2 & 4 \\ 7 & 6 & 5 & 8 & 4 & 3 & 1 & 2 \\ 8 & 7 & 2 & 6 & 3 & 5 & 4 & 1 \end{pmatrix}.$$

A is semi-symmetric, and has an autotopy group of order 42. We next found the 1951 vertices V which are at distance 5 from both A and E_3. It turns out that the minimum eccentricity of a vertex in V is 6. Hence the radius of $G_{ai}(8)$ is 6, since every vertex outside V is at distance at least 6 from either A or E_3. Next we selected $v \in V$, as represented by the semi-symmetric square:

$$\begin{pmatrix} 1 & 2 & 4 & 5 & 3 & 7 & 8 & 6 \\ 2 & 1 & 3 & 6 & 7 & 5 & 4 & 8 \\ 5 & 3 & 2 & 1 & 6 & 8 & 7 & 4 \\ 3 & 7 & 8 & 4 & 1 & 2 & 6 & 5 \\ 4 & 6 & 1 & 8 & 5 & 3 & 2 & 7 \\ 8 & 4 & 5 & 7 & 2 & 6 & 1 & 3 \\ 6 & 5 & 7 & 2 & 8 & 4 & 3 & 1 \\ 7 & 8 & 6 & 3 & 4 & 1 & 5 & 2 \end{pmatrix}$$

Our v was one of the 1729 vertices in V which are central in $G_{ai}(8)$. Let U be the set of vertices at distance 6 from v. The maximum eccentricity of the 131 vertices in U turned out to be 10. Since any two vertices outside of U could be joined to each other by a path of length at most 10 going via v, we conclude that the diameter of $G_{ai}(8)$ is 10.

The clique number of $G_{ai}(8)$ is 9 and there is a unique clique of that order. It turns out to include 3 of the vertices of the maximum clique in $G_{ii}(8)$, including the vertex represented by the square X given in (7). As they did for $G_{ri}(7)$ and $G_{ri}(8)$, the vertices involved in the maximum clique induce a complete bipartite graph in the corresponding intercalate switching graph. Here, the nine vertices induce $K_{1,8}$ in $G_{ii}(8)$, with X as the central vertex in the star. Another similarity with $G_{ri}(8)$ is that all the squares involved in the clique possess four subsquares of order 4.

Finally we note that program C found Hamiltonian paths in both $G_{am}(8)$ and $G_{ai}(8)$ (for the larger graph, we assisted the program by using the path cover of $G_{ri}(8)$ as the starting point). Hence it is possible to generate all the main or isotopy classes of order 8 in a single chain starting from E_3.

8. Concluding Remarks

In each of the even order graphs formed by intercalate switching the non equal-parity squares achieved the lower bound given in Proposition 6 but the equal parity squares were much more disconnected. Proposition 4 goes part way to explaining this behaviour, since the more symmetric squares are all equal-parity squares. However, the effect may also be partly an artifact of the small orders considered. A sample of 100000 randomly generated N_2 squares of order 10 showed that the four possible parities were roughly equally represented. This means that all four types of vertices in $G_{ii}(10)$ contain many disconnected components.

The cycle switches considered in this paper allow a substantial reduction in storage requirements for catalogues of latin squares because they allow a latin square transformation to be specified in very few bits of information. The precise details will depend on the application. However, we have found that it is possible to generate representatives of all the order 8 main or isotopy classes in a chain of switches from a single starting square. The same is true for order 7 squares except for the cyclic main class, which is atomic and hence cannot be switched to any other class. If rapid random access is required for, say, the 1676267 isotopy classes of order 8, then switching from a central vertex will never require more than 6 switches.

References

1. Asratyan, A.S., Mirumyan, A.N.: Transformations of latin squares (Russian). Diskret. Mat. **2**, 21–28 (1990)
2. Bondy J.A., Murty, U.S.R.: Graph theory with applications. American Elsevier, New York, 1976
3. Bryant, D., Maenhaut, B.M., Wanless I.M.: A family of perfect factorisations of complete bipartite graphs. J. Combin. Theory Ser. A **98**, 328–342 (2002)
4. Colbourn, C.J., Dinitz, J.H. (eds).: The CRC handbook of combinatorial designs. Boca Raton: CRC Press, 1996
5. Dénes, J., Keedwell, A.D.: Latin squares and their applications. Akadémiai Kiadó Budapest: (1974)
6. Dénes, J., Keedwell, A.D.: Latin squares: New developments in the theory and applications. Ann. Discrete Math. **46**, North-Holland, Amsterdam: (1991)
7. Denniston, R.H.F.: Remarks on latin squares with no subsquares of order two. Util. Math. **13**, 299–302 (1978)
8. Donovan, D., Mahmoodian, E.S.: An algorithm for writing any latin interchange as a sum of intercalates, Bull. Inst. Comb. Appl. **34**, 90–98 (2002)

9. Elliott, J.R., Gibbons, P.B.; The construction of subsquare free latin squares by simulated annealing. Australas. J. Comb **5**, 209–228 (1992)
10. Heinrich, K., Wallis, W.D.: The maximum number of intercalates in a latin square, Lect. Notes Math. **884**, 221–233 (1981)
11. Huang, R., Rota, G.-C.: On the relations of various conjectures on latin squares and straightening coefficients. Discrete Math **128**, 225–236 (1994)
12. Jacobson, M.T., Matthews, P.: Generating uniformly distributed random latin squares. J. Comb. Des. **4**, 405–437 (1996)
13. Janssen, J.C.M.: On even and odd latin squares. J. Comb. Theory Ser. A, **69**, 173–181 (1995)
14. McKay, B.D.: http://www.cs.anu.edu.au/∼bdm
15. McKay, B.D., Wanless, I.M.: Most latin squares have many subsquares. J. Comb. Theory Ser. A **86**, 323–347 (1999)
16. Maenhaut, B.M., Wanless, I.M.: Atomic latin squares of order eleven. J. Combin. Des. **12**, 12–34 (2004)
17. Norton, H.W.; The 7×7 squares, Ann. Eugenics **9**, 269–307 (1939)
18. Parker, E.T.: Computer investigation of orthogonal latin squares of order ten. Proc. Symp. Appl. Math. **15**, 73–81 (1963)
19. Pittenger, A.O.: Mappings of latin squares. Linear Algebra Appl. **261**, 251–268 (1997)
20. Sade, A.: An omission in Norton's list of 7×7 squares. Ann. Math. Statist. **22**, 306–307 (1951)
21. Wanless, I.M.: Perfect factorisations of bipartite graphs and latin squares without proper subrectangles. Electron. J. Comb. **6**, R9 (1999)

Received: October 10, 2002
Final version received: November 10, 2003