Symmetries That Latin Squares Inherit from 1-Factorizations

Ian M. Wanless,[1] Edwin C. Ihrig[2]
[1]*Department of Computer Science, Australian National University, ACT 0200 Australia, E-mail: imw@cs.anu.edu.au*

[2]*Department of Mathematics, Arizona State University, Tempe, Arizona, 85287-1804, E-mail: edwin.ihrig@asu.edu*

Received October 21, 2003; revised October 26, 2004

Published online in Wiley InterScience (www.interscience.wiley.com).
DOI 10.1002/jcd.20045

Abstract: A 1-factorization of a graph is a decomposition of the graph into edge disjoint perfect matchings. There is a well-known method, which we call the $\mathbb{K}$-construction, for building a 1-factorization of $K_{n,n}$ from a 1-factorization of K_{n+1}. The 1-factorization of $K_{n,n}$ can be written as a latin square of order n. The $\mathbb{K}$-construction has been used, among other things, to make perfect 1-factorizations, subsquare-free latin squares, and atomic latin squares. This paper studies the relationship between the factorizations involved in the $\mathbb{K}$-construction. In particular, we show how symmetries (automorphisms) of the starting factorization are inherited as symmetries by the end product, either as automorphisms of the factorization or as autotopies of the latin square. Suppose that the $\mathbb{K}$-construction produces a latin square L from a 1-factorization F of K_{n+1}. We show that the main class of L determines the isomorphism class of F, although the converse is false. We also prove a number of restrictions on the symmetries (autotopies and paratopies) which L may possess, many of which are simple consequences of the fact that L must be symmetric (in the usual matrix sense) and idempotent. In some circumstances, these restrictions are tight enough to ensure that L has trivial autotopy group. Finally, we give a cubic time algorithm for deciding whether a main class of latin squares contains any square derived from the $\mathbb{K}$-construction. The algorithm also detects symmetric squares and totally symmetric squares (latin squares that equal their six conjugates). © 2005 Wiley Periodicals, Inc. J Combin Designs 13: 157–172, 2005.

Keywords: *1-factorization; perfect* 1*-factorization; latin square; autotopy, paratopy; main class; totally symmetric; idempotent; atomic latin square*

1. INTRODUCTION

A *k-factor* in a graph is a k-regular spanning subgraph. In particular, a 1-*factor* (also called a *perfect matching*) is a set of edges in a graph which cover every vertex

exactly once. A 1-*factorization* is a collection of 1-factors which partitions the edges of the graph. A *rooted* 1-*factorization*, (F, v), of a graph G is a pair consisting of a 1-factorization F of G together with a vertex v from G, called the *root*. As part of its definition we also require that (F, v) includes an implicit total order on the vertices of G.

This paper studies a well-known method for obtaining a 1-factorization of the complete bipartite graph $K_{n,n}$ from a rooted 1-factorization of the complete graph K_{n+1}. We call this method the $\mathbb{K}$-construction. It seems to have become part of the folklore of the subject and is commonly used [5, 10, 12] without accreditation to an earlier source in the literature. Dénes and Keedwell [7, p. 116] attribute it to "Rosa and others." An early source containing the basic ideas behind the $\mathbb{K}$-construction is the paper of Kotzig [11].

Before we begin, we give notation for some representations of graphs. Let U and V be two disjoint sets with n elements each. By $K_{U,V}$, we mean the complete bipartite graph with U and V as the *parts* (maximal independent sets). This gives a representation of $K_{n,n}$. Similarly, K_U denotes the complete graph with vertex set U.

We now show how to produce a 1-factorization of a complete bipartite graph from a 1-factorization of a complete graph. Suppose that we are given a rooted 1-factorization (F, v) of K_W. Assume W has cardinality $n + 1$, where n is odd, so K_W is isomorphic to K_{n+1}. We will show how (F, v) can be used to define a 1-factorization on $K_{U,V}$ where $U = (W - \{v\}) \times \{1\}$ and $V = (W - \{v\}) \times \{2\}$. Given any 1-factor $f \in F$, we define f^*, a 1-factor of $K_{U,V}$, as follows. If x is connected to v in f, then connect $(x, 1)$ to $(x, 2)$ in f^*. If x is connected to y in f where $x \neq v$ and $y \neq v$, then $(x, 1)$ is connected to $(y, 2)$ in f^*. The new 1-factorization of $K_{U,V}$, denoted by $\mathcal{F}(F, v)$, is $\{f^* : f \in F\}$. Throughout this paper, the asterisk $(^*)$ will be used exclusively in the above sense. That is, x^* will always denote a 1-factor related to the 1-factor x in the same way that f^* is related to f above.

We will next show how to build a latin square from a 1-factorization of a complete bipartite graph. Any reader who is not familiar with the definitions of *latin square*, *isotopic*, and *main class* may find it useful to read §2 before proceeding. For an earlier paper which spells out the relationship between latin squares and 1-factorizations, consult Dénes and Keedwell [5].

Let U and V be two disjoint sets with n elements each. We will produce a latin square $\mathcal{L}(F)$ of order n from an ordered 1-factorization F of $K_{U,V}$. By an *ordered* 1-*factorization* of $K_{U,V}$, we mean a 1-factorization of $K_{U,V}$ together with a total order on each of the parts, a designation of one part as the first part (by default, the other part becomes known as the second part), and a total order on the 1-factors. Given such an ordered 1-factorization, construct a latin square by making k the entry in row i and column j if the kth 1-factor connects the ith vertex in the first part with the jth vertex in the second part.

The above construction shows the combinatorial equivalence between ordered 1-factorizations of $K_{U,V}$ and latin squares of order n. Note that an ordering of the 1-factorization is needed to define this construction, and different orderings of the same 1-factorization will give rise to different latin squares. However, the latin squares constructed from different orderings will all be in the same main class. In fact, any two isomorphic 1-factorizations will produce latin squares in the same main class, irrespective of what orderings are used in the construction above. Thus, if only

the main class of the latin square is of interest, only the isomorphism class of the 1-factorization is needed for the construction.

Now we are ready to give the $\mathbb{K}$-construction. In this construction, we start with (F, v), a rooted 1-factorization of a complete graph, K_W, and use it to produce a latin square $\mathcal{L}(F, v)$. While we have no trouble first producing a 1-factorization of a complete bipartite graph, we cannot proceed further without knowing how to put an order on this 1-factorization. This is why we insisted that a rooted 1-factorization of K_W includes a total order on its vertex set. This allows us to give the bipartite 1-factorization an order as follows. First, give $U = (W - \{v\}) \times \{1\}$ and $V = (W - \{v\}) \times \{2\}$ the natural order inherited from the order on W. Second, call U the first part. Finally, order the 1-factors in F by saying f precedes f' if the vertex connected to the root by f precedes the vertex connected to the root by f'. This ordering of the 1-factors in F gives rise to an ordering of the 1-factors in $\mathcal{F}(F, v)$. In this way, $\mathcal{F}(F, v)$ becomes an ordered 1-factorization, and $\mathcal{L}\big(\mathcal{F}(F, v)\big)$ becomes well defined. We will use $\mathcal{L}(F, v)$ as a shorthand for $\mathcal{L}(\mathcal{F}(F, v))$.

While different orderings of the vertices in the rooted 1-factorization give rise to different latin squares, these latin squares will all be isotopic. Thus, if two latin squares that are isotopic are treated as being the same, then the vertex order may be omitted from the definition of rooted factorizations and $\mathcal{L}(F, v)$ will still be well defined. In fact, $\mathcal{L}(F, v)$ will only depend on the isomorphism class of the rooted 1-factorization. (Here, two rooted 1-factorizations are considered to be isomorphic if there is an isomorphism between the 1-factorizations which takes the root of one to the root of the other.) However, the main class of the resulting latin square can change if the root is changed. We will, in fact, show (Theorem 16) that $\mathcal{L}(F, v)$ and $\mathcal{L}(F, w)$ belong to the same main class if and only if there is an automorphism of F which takes v to w. Thus, any 1-factorization which does not have a vertex transitive automorphism group will give rise to several distinct main classes of latin squares depending on how the root is chosen.

The present paper was motivated by the observation that the main class of the latin square resulting from the $\mathbb{K}$-construction is root dependent. This was observed independently by Keedwell [10] and Wanless [19] (see also [13]). Here we develop parts of a theory which enable us to identify and classify the results of the $\mathbb{K}$-construction.

The structure of the paper is as follows. In the next section, we cover the basic terminology of latin squares, in which many of our results will be formulated. The basic graph theoretic definitions have already been made above. In §3, we review what is perhaps the original and most important application of the $\mathbb{K}$-construction, which is the construction of perfect 1-factorizations. In §4, we then study in some depth the symmetries which factorizations involved in the $\mathbb{K}$-construction can or must possess. Finally, in §5, we look briefly at some algorithmic questions related to the $\mathbb{K}$-construction.

2. LATIN SQUARES TERMINOLOGY

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 will always assume that the symbol set coincides with the set of indices of the rows

and columns, which will usually be $[n] = \{1, 2, \ldots, n\}$. 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. A latin square that includes the triple (x, x, x) for each $x \in [n]$, is said to be *idempotent*. It will turn out that $\mathcal{L}(F, v)$ is always idempotent.

For each latin square L, there are $3! = 6$ *conjugate* squares obtained by uniformly permuting the coordinates in each triple. We denote these six conjugates by L^{123}, L^{213}, L^{132}, L^{231}, L^{312}, and L^{321}. The superscript in each case designates the new order of the co-ordinates—it is the permutation, in image notation, which was applied to the triples. For example, $L^{123} = L$ and L^{213} is the transpose of L. If $L = L^{213}$ then L is said to be *symmetric*. If it is equal to all six of its conjugates then L is said to be *totally symmetric*. We will show that $\mathcal{L}(F, v)$ is always symmetric, but need not be totally symmetric. This observation will be developed further in §5.

An *isotopy* of a latin square L is a permutation of its rows, permutation of its columns, and permutation 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*. In the special case when the same permutation is applied to the rows, columns, and symbols, we say that the isotopy is an *isomorphism*. An isotopy that maps L to itself is called an *autotopy* of L and an autotopy that happens to be an isomorphism is called an *automorphism*. The *main class* of L is the set of squares which are isotopic to some conjugate of L. Latin squares belonging to the same main class are said to be *paratopic* and a map that combines an isotopy with conjugation is called a *paratopy*. A paratopy that maps a latin square to itself is an *autoparatopy* of the square.

Our notation for isotopies will make use of permutations, which will always be denoted by lower case Greek letters. In particular, ε will always denote the identity permutation. Also, we use the convention that permutations act on the left. So, for example, $x\pi$ denotes the image of x under the permutation π. Also, $L\langle\alpha, \beta, \gamma\rangle$ will denote the result on L of the isotopy $\langle\alpha, \beta, \gamma\rangle$, which applies permutations α to the rows, β to the columns, and γ to the symbols. If $L = L\langle\pi, \pi, \pi\rangle$ then we say that π is an automorphism of L. An isotopy $\langle\alpha, \beta, \gamma\rangle$ is called a *principal isotopy* if $\gamma = \varepsilon$. Much of our terminology, including paratopy and principal isotopy, comes from the seminal writings of Albert Sade (e.g., [16], [17]).

If A and B are two latin squares of the same order then we write the following.

1. $A \sim B$ to indicate that A is paratopic to B.
2. $A \simeq B$ to indicate that A is isotopic to B.
3. $A \cong B$ to indicate that A is isomorphic to B.

Each successive statement is stronger, so 3 implies 2, which implies 1.

3. PERFECTION

Suppose that we have a 1-factorization F of some graph. The union of any two distinct 1-factors in F is a 2-factor, that is, a collection of cycles. If, regardless of which two 1-factors we choose in F, their union is a single cycle then we say that F is *perfect*.

The original motivation for the $\mathbb{K}$-construction is the following result. The "if" part of this theorem was proved by Laufer [12], and this is the result that is usually

quoted. However, the "only if" direction is important as well (for example, it was recently used in [2] and [20]).

Theorem 1. *Let (F, v) be a rooted 1-factorization of K_{n+1}. Then $\mathcal{F}(F, v)$ is perfect if and only if F is perfect.*

Proof. Let f and g denote two 1-factors in F with f^* and g^* the corresponding 1-factors in $\mathcal{F}(F, v)$.

We examine how the cycles in $f \cup g$ correspond to the cycles in $f^* \cup g^*$. Let C be a cycle in $f \cup g$ which does not contain the vertex v. Note that C has even length, since it is known to be 2-colorable. Suppose that the vertices in order round C are $w_1, w_2, w_3, \ldots, w_k$. There are two corresponding cycles in $f^* \cup g^*$. One has vertices $(w_1, 1), (w_2, 2), (w_3, 1), (w_4, 2), \ldots, (w_k, 2)$ in cyclic order, and the other has vertices $(w_1, 2), (w_2, 1), (w_3, 2), (w_4, 1), \ldots, (w_k, 1)$. Thus any k-cycle in $f \cup g$ that does not contain v corresponds to two k-cycles in $f^* \cup g^*$.

Next assume $v, w_1, w_2, w_3, \ldots, w_k$ is the $(k+1)-$cycle in $f \cup g$ which does contain v. The corresponding cycle in $f^* \cup g^*$ is

$$(w_1, 1), (w_2, 2), (w_3, 1), \ldots, (w_k, 1), (w_k, 2), (w_{k-1}, 1), (w_{k-2}, 2), \ldots, (w_1, 2),$$

which is a $2k$-cycle.

This shows that $f \cup g$ has j distinct cycles iff $f^* \cup g^*$ has $2j - 1$ distinct cycles. In particular, $f \cup g$ is a cycle iff $f^* \cup g^*$ is a cycle. □

Not all perfect 1-factorizations of $K_{n,n}$ can be constructed by using the $\mathbb{K}$-construction. In fact, Wanless [20] showed that there are 37 non-isomorphic perfect 1-factorizations of $K_{9,9}$ and only one comes from a perfect 1-factorization of K_{10}. In §5, we will discuss the issue of how to recognize those 1-factorizations which can be produced by the $\mathbb{K}$-construction.

The spectrum of values of n for which a perfect 1-factorization of K_{n+1} exists is not known, although it is thought to be the set of odd natural numbers. Three infinite families [2] and some constructions for sporadic values have been found but there is still much work to be done. See Wallis's book [21] for a survey of what was known in 1997, and [2], [20] for some recent constructions. Also, see Xu et al. [22] for an application of perfect 1-factorizations of complete graphs to coding theory.

For complete bipartite graphs, the situation is little better. All the results for complete graphs can be employed via Theorem 1 and there is one additional infinite family known [1]. However, we are still a long way from proving that, as is likely, $K_{n,n}$ has a perfect 1-factorization iff n is odd or $n = 2$.

Next we define a cycle structure of a latin square and relate it to the cycle structure of the 2-factors obtained as pairwise unions of 1-factors.

Suppose that a latin square L encodes a particular 1-factorization F of $K_{n,n}$; that is, $L = \mathcal{L}(F)$. For each symbol s of L, let f_s denote the 1-factor in F which corresponds to s and define the permutation $\tau_s : [n] \mapsto [n]$ by $c\tau_s = d$ if $L_{cd} = s$. Next, for each ordered pair (s, t) of distinct symbols of L, define $\sigma_{s,t} = \tau_s \tau_t^{-1}$. The permutation $\sigma_{s,t}$ can be written as a product of disjoint cycles in the standard way. These cycles correspond exactly to cycles in $f_s \cup f_t$ in the sense that every cycle in $\sigma_{s,t}$ of length j corresponds to a cycle of length $2j$ in $f_s \cup f_t$. Notice that the cycle lengths of the permutation τ_s are not an isotopy invariant. Since τ_s is a map from the row indices to

the column indices, and an isotopy can permute the column indices while leaving the row indices alone, then any τ_s can be changed to any other permutation by performing a suitable isotopy. However, the same is not true for $\sigma_{s,t}$. This permutation takes the row indices to row indices, so it will be just conjugated after an isotopy is applied. Hence the cycle structure of $\sigma_{s,t}$ is an invariant under isotopy.

If $\sigma_{s,t}$ consists of a single n-cycle then we say it is *Hamiltonian*. If $\sigma_{s,t}$ is Hamiltonian for every possible choice of (s,t) then we say that L is *symbol-Hamiltonian*. From these remarks, the following result is easily deduced.

Theorem 2. *Let $L = \mathcal{L}(F)$ for a 1-factorization F of a complete bipartite graph. Then L is symbol-Hamiltonian if and only if F is perfect. Moreover, if two latin squares are isotopic, one is symbol-Hamiltonian if and only if the other is symbol-Hamiltonian.*

The original interest in symbol-Hamiltonian latin squares arose from the fact that they contain no non-trivial latin subsquares [7, p. 116]. They have also been studied by Bryant et al. [1] and Wanless [19] (although both these papers deal with squares which are called ''pan-Hamiltonian'' and are defined in terms of cycles between two rows rather than two symbols, the concept is related to the present topic by conjugacy). A related idea is that of an *atomic latin square*, which is a latin square for which all of its conjugates are symbol-Hamiltonian. Atomic latin squares have been studied in [2], [13], [14], [19], and [20].

4. INHERITED SYMMETRIES

There are three combinatorial objects related by the $\mathbb{K}$-construction. We have a rooted 1-factorization (F, v) of a complete graph, a 1-factorization $\mathcal{F}(F, v)$ of a complete bipartite graph, and a latin square $L = \mathcal{L}(F, v)$. Each of these objects may have symmetries. In the case of the 1-factorizations, the symmetries are automorphisms, which we define below. For the latin square, the appropriate notion of symmetry is autotopy, or more generally, autoparatopy. In this section, we investigate which symmetries our three objects might have. In particular, we are interested in inherited properties. What does a symmetry of one of the objects tell us about symmetries of the others?

Two 1-factorizations $F = [f_1, f_2, \ldots, f_n]$ and $F' = [f'_1, f'_2, \ldots, f'_n]$ of the graphs G, and G', respectively, are said to be *isomorphic* if there is a bijection Φ from the vertices of G to the vertices of G' and a permutation $\pi : [n] \mapsto [n]$ such that $f_i\Phi = f'_{i\pi}$ for each $i \in [n]$. Here $f_i\Phi$ denotes the set of all edges $\{x\Phi, y\Phi\}$ of G' such that $\{x, y\} \in f_i$. An *automorphism* of F is an isomorphism from F to itself. For rooted 1-factorizations, isomorphisms (including automorphisms) have the additional requirement that Φ must map the root of the initial 1-factorization to the root in the image.

In a connected bipartite graph, there are exactly two parts. Since the image of a part under a graph automorphism must be a part, any graph automorphism of a connected bipartite graph must permute the two parts. An automorphism is called *part-preserving* if it maps each part to itself and *part-reversing* if it maps each part to the other part. In connected bipartite graphs, these are the only two possibilities.

Theorem 3. *Let F be an ordered 1-factorization of $K_{n,n}$ and let $L = \mathcal{L}(F)$. The group of part-preserving automorphisms of F is isomorphic to the autotopy group of $\mathcal{L}(F)$.*

Proof. Let the two parts of $K_{n,n}$ be U and V, where U corresponds to the rows in L and V corresponds to the columns of L. The symbols of L correspond to the 1-factors in F. As a result, from any autotopy $\langle \alpha, \beta, \gamma \rangle$ of L, we find an automorphism of F by applying α to U and β to V (the role of γ is to provide the permutation π required by the definition of isomorphism given above).

This gives a homomorphism Θ from the autotopy group of L to the automorphism group of F. We first show that Θ is one-to-one by showing it has trivial kernel. If $\langle \alpha, \beta, \gamma \rangle$ is in the kernel of Θ then clearly $\alpha = \beta = \varepsilon$. But any two components of an autotopy determine the third uniquely (see Theorem 4.1.4 in [6]), and $\langle \varepsilon, \varepsilon, \varepsilon \rangle$ is certainly an autotopy of L, so $\gamma = \varepsilon$ as desired.

We finish by showing that Θ maps onto the set of part-preserving automorphisms of F. Any such automorphism is a permutation of the vertex set of $K_{n,n}$ which acts as α applied to U and β applied to V. Moreover, this automorphism also permutes the 1-factors of $K_{n,n}$. Since the 1-factors of $K_{n,n}$ are labeled by the symbols of L, this induces a permutation γ of the symbol set of L. It is easy to check that the $\langle \alpha, \beta, \gamma \rangle$ defined in this way will be an autotopy of F which is mapped by Θ to the given automorphism. Hence Θ is the isomorphism required to complete the proof. □

Of course, in some cases, F will also have part-reversing automorphisms, in which case (since the underlying graph is connected) the automorphisms identified in Theorem 3 will form a subgroup of index 2 in the full automorphism group of F. In general, the automorphism group of F is a subgroup of the wreath product $S_n \wr S_2$, where S_i denotes the full symmetric group on $[i]$. The part-reversing automorphisms of F correspond to autoparatopies of $\mathcal{L}(F)$, which are composed of an isotopy together with transposition. Autoparatopies of $\mathcal{L}(F)$ that involve taking conjugates other than the identity and transpose are not so natural when considered as symmetries of F. The autoparatopy group of any latin square of order n is a subgroup of $S_n \wr S_3$.

One case of present interest where F has a part-reversing automorphism is when F is produced by applying the $\mathbb{K}$-construction to a rooted 1-factorization of K_W. In that case, our next result tells us that $\mathcal{L}(F)$ is symmetric, which means that F has a part-reversing automorphism which simply interchanges $(w, 1)$ with $(w, 2)$ for each $w \in W$ other than the root.

Lemma 4. *Let (F, v) be a rooted 1-factorization of K_W where W has $n+1$ elements. Then $\mathcal{L}(F, v)$ is a symmetric idempotent latin square of order n. Conversely, any symmetric idempotent latin square of order n is $\mathcal{L}(F, v)$ for some rooted 1-factorization (F, v) of K_{n+1}.*

Proof. By the construction of $\mathcal{F}(F, v)$, if f is a 1-factor with an edge between w and v then $(w, 1)$ is connected to $(w, 2)$ by an edge of f^*. If w is the ith vertex in $W \backslash \{v\}$, then this means (i, i, i) will be in $\mathcal{L}(F, v)$. As this is true for each i, $\mathcal{L}(F, v)$ is idempotent. Furthermore, if neither x nor y is connected to v in f, then both $(x, 1)$ is connected to $(y, 2)$ and $(y, 1)$ is connected to $(x, 2)$ in f^*. Let x be the ith element,

y the jth element, and z the k-th element of $W\backslash\{v\}$ (where z is connected to v in f). Then this means both (i,j,k) and (j,i,k) are in $\mathcal{L}(F,v)$, so $\mathcal{L}(F,v)$ is symmetric.

To prove the second assertion of the lemma, suppose that L is a symmetric idempotent latin square of order n. Let the vertices of K_{n+1} be $\{1,2,\ldots,n,v\}$ with the usual order. Let the root be v. For each symbol i in L, we define a 1-factor of K_{n+1} that contains the edge $\{i,v\}$ as well as edges $\{a,b\}$ whenever $(a,b,i)\in L$ and $a\neq b$. This is well defined because L is symmetric and idempotent. (Idempotence implies that if $a\neq b$ then neither a nor b is i.) The set of such 1-factors forms a 1-factorization F, because of the latin property of L. Finally, it is routine to check from the definition that $L=\mathcal{L}(F,v)$. □

Since the $\mathbb{K}$-construction produces symmetric idempotent latin squares, we now embark on a sequence of lemmas that explore the properties of such squares.

Lemma 5. *Let A and B be idempotent latin squares related by $A=B\langle\rho,\rho,\tau\rangle$. Then $\tau=\rho$, so that $A\cong B$.*

Proof. If x is any symbol of B then B contains the triple (x,x,x) because B is idempotent. But this means that A must contain the triple $(x\rho,x\rho,x\tau)$, which implies that $x\rho=x\tau$ because A is idempotent. Since x was arbitrary, $\rho=\tau$ as required. □

Lemma 6. *If A and B are idempotent symmetric latin squares and $A\simeq B$ then $A\cong B$. In particular, if $A=B\langle\rho,\sigma,\tau\rangle$ then $A=B\langle\tau,\tau,\tau\rangle$.*

Proof. Suppose that A and B are idempotent symmetric latin squares and $A=B\langle\rho,\sigma,\tau\rangle$. Let $C=B\langle\rho,\rho,\rho\rangle$ so that $A=C\langle\varepsilon,\pi,\rho^{-1}\tau\rangle$ where $\pi=\rho^{-1}\sigma$. Note that C inherits the property of being idempotent and symmetric from B.

By the symmetry of A and C, it follows that $A=C\langle\pi,\varepsilon,\rho^{-1}\tau\rangle=A\langle\pi,\pi^{-1},\varepsilon\rangle$. Suppose that π has even order, say $2m$. Then

$$A=A\langle\pi,\pi^{-1},\varepsilon\rangle^m=A\langle\pi^m,\pi^{-m},\varepsilon\rangle=A\langle\pi^m,\pi^m,\varepsilon\rangle.$$

Now by Lemma 5, we see that $\pi^m=\varepsilon$, which contradicts the definition of m unless $2m=m=0$. However, if $m=0$ then $\pi=\varepsilon$ so $\rho=\sigma$. As this case is covered by Lemma 5, we may assume that π has odd order, say $2m+1$. Then $A=A\langle\pi^m,\pi^{-m},\varepsilon\rangle=C\langle\pi^{m+1},\pi^{-m},\rho^{-1}\tau\rangle=C\langle\pi^{m+1},\pi^{m+1},\rho^{-1}\tau\rangle$.

Hence $\pi^{m+1}=\rho^{-1}\tau$ by Lemma 5 and $A=C\langle\rho^{-1}\tau,\rho^{-1}\tau,\rho^{-1}\tau\rangle=B\langle\tau,\tau,\tau\rangle$ as claimed. □

Corollary 7. *If A is an idempotent symmetric latin square with trivial automorphism group then every autotopy of A is a principal autotopy.*

The corollary above shows that principal autotopies are of some interest, so next we investigate the subgroup of the autotopy group consisting of the principal autotopies.

Lemma 8. *Let P be the group of principal autotopies of a symmetric latin square A of odd order. Then P is Abelian and has odd order. Also if $\langle\rho,\sigma,\varepsilon\rangle\in P$ then $\rho=\sigma^{-1}$.*

Proof. Consider the involution f that maps a typical principal autotopy $\langle \rho, \sigma, \varepsilon \rangle$ to $\langle \sigma, \rho, \varepsilon \rangle$. Note that $f : P \mapsto P$ since A is symmetric and hence f is an automorphism of P. If we can show that $\langle \varepsilon, \varepsilon, \varepsilon \rangle$ is the only fixed point of f then our result will follow from a lemma due to Burnside (see [18, p. 131]). So suppose that $\langle \sigma, \sigma, \varepsilon \rangle \in P$ is a fixed point of f. For any i, there is a j with $(i, i\sigma, j) \in A$. Hence $(i\sigma, i\sigma^2, j) \in A$, which means $\sigma^2 = \varepsilon$ since A is symmetric. Because the order of A is odd, σ must have a fixed point, k say. Now for any i, there is a j with $(i, k, j) \in A$. Applying $\langle \sigma, \sigma, \varepsilon \rangle$ to this triple gives us $(i\sigma, k, j)$, which means that $i = i\sigma$. As i was arbitrary, $\sigma = \varepsilon$ as required. □

For the next result, we remind the reader that a *regular* permutation is one in which all cycles have the same length.

Lemma 9. *If $\langle \rho, \sigma, \tau \rangle$ is an autotopy of a symmetric latin square then $\rho\sigma^{-1}$ is a regular permutation. If $\langle \rho, \sigma, \varepsilon \rangle$ is a principal autotopy of a symmetric latin square of odd order, then $\rho = \sigma^{-1}$ is a regular permutation.*

Proof. If A is symmetric then $A = A\langle \rho, \sigma, \tau \rangle = A\langle \sigma, \rho, \tau \rangle = A\langle \sigma^{-1}, \rho^{-1}, \tau^{-1} \rangle$ so $A = A\langle \rho, \sigma, \tau \rangle \langle \sigma^{-1}, \rho^{-1}, \tau^{-1} \rangle = A\langle \pi, \pi^{-1}, \varepsilon \rangle$ where $\pi = \rho\sigma^{-1}$. Suppose, by way of contradiction that π is not regular. Let l be the smallest integer such that π has cycles of length l. Then π^l has a fixed point f but is not the identity. However, $A = A\langle \pi, \pi^{-1}, \varepsilon \rangle^l = A\langle \pi^l, \pi^{-l}, \varepsilon \rangle$. The autotopy $\langle \pi^l, \pi^{-l}, \varepsilon \rangle$ fixes row f and all symbols in it, from which we conclude that it must fix all columns. But this implies the contradiction $\pi^l = \varepsilon$, thereby proving the first statement.

The second statement in the lemma follows from the same argument with ρ in place of π, since $\sigma = \rho^{-1}$ by Lemma 8. □

Let $\mathcal{P}_1$ be the projection of P onto its first co-ordinate (that is, $\mathcal{P}_1 = \{\rho : \langle \rho, \sigma, \varepsilon \rangle \in P \text{ for some } \sigma\}$). By Lemma 8, we know that $|P_1| = |P|$ and hence P_1 is isomorphic to P. The argument in Lemma 9 shows that ε is the only permutation in P_1 with a fixed point. In other words, P_1 acts 'semi-regularly' on $[n]$ and hence $|P_1|$ divides n.

Lemma 10. *Let P be the group of principal autotopies of an idempotent symmetric latin square A. If P is trivial then every autotopy of A is an automorphism of A.*

Proof. Suppose that $A = A\langle \rho, \sigma, \tau \rangle$. Then by the argument in the proof of Lemma 9, we know that $\langle \rho\sigma^{-1}, \sigma\rho^{-1}, \varepsilon \rangle \in P$. So if P is trivial then $\rho\sigma^{-1} = \varepsilon$ and $\rho = \sigma$. But now Lemma 5 implies that $\rho = \sigma = \tau$ as required. □

Our next result involves the well-known *patterned factorizations*. See [8] and [21] for background on this concept. For a general (possibly non-Abelian) group G of odd order, the latin square derived from the patterned factorization on G can be defined by

$$\mathcal{L}(GK(G), \infty) = \{(xy, x^{-1}y, y) : x, y \in G\}. \tag{1}$$

Alternatively, replacing x by xy^{-1} gives the formulation

$$\mathcal{L}(GK(G), \infty) = \{(x, yx^{-1}y, y) : x, y \in G\}.$$

If G is Abelian, $\mathcal{L}(GK(G), \infty)$ is isotopic to the Cayley table of G; that is, to the latin square $\{(s, t, st) : s, t \in G\}$. The isotopy is given by $\langle \rho, \rho, \varepsilon \rangle$ where ρ is defined by $x\rho = x^{(|G|+1)/2}$ for each $x \in G$.

Lemma 11. *Let A be an idempotent symmetric latin square for which the group of principal autotopies, P, acts transitively on the rows. Define*

$$\mathcal{P}_1 = \{\rho : \langle \rho, \sigma, \varepsilon \rangle \in P \text{ for some } \sigma\}$$

and let i be any fixed symbol. Then

$$A = \left\{(i\sigma\tau, i\sigma\tau^{-1}, i\sigma) : \sigma, \tau \in \mathcal{P}_1\right\}.$$

Moreover, A is isomorphic to $\mathcal{L}(GK(P), \infty)$ and hence isotopic to the Cayley table of P.

Proof. Note that A must have odd order since there are no idempotent symmetric latin squares of even order. Hence, P (which is isomorphic to $\mathcal{P}_1$) is Abelian of odd order and acts simply transitively on the rows by Lemma 8 and Lemma 9.

For any $\sigma, \tau \in \mathcal{P}_1$, observe that $(i\sigma\tau, i\sigma\tau^{-1}, i\sigma) \in A$ since $\langle \tau, \tau^{-1}, \varepsilon \rangle \in P$ and $(i\sigma, i\sigma, i\sigma) \in A$, given that A is idempotent.

Assume $(j, k, l) \in A$. Since $\mathcal{P}_1$ acts transitively there is a $\sigma \in \mathcal{P}_1$ such that $l = i\sigma$. Furthermore, there is a $\tau \in \mathcal{P}_1$ such that $i\tau = j\sigma^{-1}$. We have $(j, k, l) = (i\sigma\tau, k, i\sigma) \in A$. But we also have $(i\sigma\tau, i\sigma\tau^{-1}, i\sigma) \in A$. Since these two triples agree in two components, they must be equal. We conclude that every triple of A can be written in the form $(i\sigma\tau, i\sigma\tau^{-1}, i\sigma)$ for some $\sigma, \tau \in \mathcal{P}_1$.

The last statement follows from (1) using the isomorphism $\langle \rho, \rho, \rho \rangle$ where ρ maps σ to $i\sigma$ for each $\sigma \in \mathcal{P}_1$. The fact that P is Abelian is used to show that $\langle \rho, \rho, \rho \rangle$ has the correct image and also to justify that $\mathcal{L}(GK(P), \infty)$ is isotopic to the Cayley table of P. □

Note that when G is Abelian, $\mathcal{L}(GK(G), \infty)$ does have a principal autotopy group P, isomorphic to G, which acts transitively on the rows. Thus, Lemma 11 characterizes symmetric idempotent latin squares with transitive P.

Lemma 12. *Let A be an idempotent symmetric latin square of prime order and suppose that A is not isotopic to the Cayley table of a cyclic group. Then every autotopy of A is an automorphism of A.*

Proof. Suppose that $A = A\langle \rho, \sigma, \tau \rangle$. By Lemma 6, $A = A\langle \tau, \tau, \tau \rangle$, from which it follows that $A = A\langle \rho\tau^{-1}, \sigma\tau^{-1}, \varepsilon \rangle$. Now, by Lemma 9, $\rho\tau^{-1} = \tau\sigma^{-1}$ is a regular permutation, which is therefore either trivial or transitive since A has prime order. The transitive case is ruled out by Lemma 11 and our assumptions on A, so we conclude that $\rho\tau^{-1} = \tau\sigma^{-1} = \varepsilon$. Thus $\rho = \sigma = \tau$ as required. □

The following example shows the restriction to prime orders in Lemma 12 cannot be abandoned. It is a symmetric idempotent latin square of order 15 with a trivial automorphism group but with a non-trivial autotopy $\langle \alpha, \alpha^{-1}, \varepsilon \rangle$ where $\alpha = (1, 2, 3)$ $(4, 5, 6)(7, 8, 9)(10, 11, 12)(13, 14, 15)$. Note that this autotopy is principal, as

predicted by Corollary 7, and is of the form predicted by Lemma 8. Also, $\alpha = (\sigma_{1,6})^5$, where $\sigma_{i,j}$ is defined in §3. Indeed, $\sigma_{1,6}$ is Hamiltonian and any principal autotopy $\langle \beta, \beta^{-1}, \varepsilon \rangle$ must be such that β commutes with each $\sigma_{i,j}$. This information is sufficient to deduce that α will be a power of $\sigma_{1,6}$.

$$\begin{pmatrix}
1 & 3 & 2 & 10 & 12 & 11 & 4 & 5 & 6 & 13 & 15 & 14 & 9 & 7 & 8 \\
3 & 2 & 1 & 12 & 11 & 10 & 5 & 6 & 4 & 15 & 14 & 13 & 7 & 8 & 9 \\
2 & 1 & 3 & 11 & 10 & 12 & 6 & 4 & 5 & 14 & 13 & 15 & 8 & 9 & 7 \\
10 & 12 & 11 & 4 & 6 & 5 & 13 & 15 & 14 & 7 & 8 & 9 & 1 & 3 & 2 \\
12 & 11 & 10 & 6 & 5 & 4 & 15 & 14 & 13 & 8 & 9 & 7 & 3 & 2 & 1 \\
11 & 10 & 12 & 5 & 4 & 6 & 14 & 13 & 15 & 9 & 7 & 8 & 2 & 1 & 3 \\
4 & 5 & 6 & 13 & 15 & 14 & 7 & 9 & 8 & 1 & 3 & 2 & 10 & 12 & 11 \\
5 & 6 & 4 & 15 & 14 & 13 & 9 & 8 & 7 & 3 & 2 & 1 & 12 & 11 & 10 \\
6 & 4 & 5 & 14 & 13 & 15 & 8 & 7 & 9 & 2 & 1 & 3 & 11 & 10 & 12 \\
13 & 15 & 14 & 7 & 8 & 9 & 1 & 3 & 2 & 10 & 12 & 11 & 4 & 6 & 5 \\
15 & 14 & 13 & 8 & 9 & 7 & 3 & 2 & 1 & 12 & 11 & 10 & 6 & 5 & 4 \\
14 & 13 & 15 & 9 & 7 & 8 & 2 & 1 & 3 & 11 & 10 & 12 & 5 & 4 & 6 \\
9 & 7 & 8 & 1 & 3 & 2 & 10 & 12 & 11 & 4 & 6 & 5 & 13 & 15 & 14 \\
7 & 8 & 9 & 3 & 2 & 1 & 12 & 11 & 10 & 6 & 5 & 4 & 15 & 14 & 13 \\
8 & 9 & 7 & 2 & 1 & 3 & 11 & 10 & 12 & 5 & 4 & 6 & 14 & 13 & 15
\end{pmatrix}$$

This example appears to be highly composite, and one is left to wonder if there exist such examples without proper subsquares (a *subsquare* is a submatrix that is itself a latin square and it is *proper* if it has order at least 2 but is not the whole square). The following result shows this is not the case.

Lemma 13. *Let A be an idempotent symmetric latin square. Then the order of any principal autotopy divides the order of A. Also if m divides the order of a principal autotopy, then each triple (i,i,i) is part of a subsquare of order m which is isomorphic to $\mathcal{L}(GK(\mathbb{Z}_m), \infty)$. In particular, any idempotent symmetric latin square with no proper subsquares but which has a non-trivial principal autotopy must be isomorphic to $\mathcal{L}(GK(\mathbb{Z}_p), \infty)$ for some prime p.*

Proof. Assume (i,i,i) and m are given as above. By Lemma 9, m divides the order of A. There is a subgroup P' of the cyclic group generated by the principal autotopy, which has order m. Define $\mathcal{P}'_1 = \{\rho : (\rho, \sigma, \varepsilon) \in P'\}$ and

$$A' = \left\{(i\sigma\tau, i\sigma\tau^{-1}, i\sigma) : \sigma, \tau \in \mathcal{P}'_1\right\}.$$

It is straightforward to verify that A' is a latin square given that m is odd (Lemma 8) and P' consists of regular permutations (Lemma 9). Also, A' is a subsquare of A because A is idempotent and $A\langle \tau, \tau^{-1}, \varepsilon \rangle = A$. Lemma 11 says A' is isomorphic to $\mathcal{L}(GK(P'), \infty)$ since P' acts transitively on the rows of A'. □

The following Corollary is immediate, given Lemma 10.

Corollary 14. *Let A be an idempotent symmetric latin square. Assume A has no proper subsquare isomorphic to $\mathcal{L}(GK(\mathbb{Z}_p), \infty)$, for any prime p. If A has trivial automorphism group then A has trivial autotopy group.*

We need one more lemma before giving our main result.

Lemma 15. *Let M be a main class containing a symmetric latin square S. Then every symmetric square in M is isotopic to S.*

Proof. Let M be a main class containing symmetric latin squares S and T. Since S is symmetric we know that $S^{123} = S^{213}$, $S^{132} = S^{231}$, and $S^{312} = S^{321}$. As T is in the main class of S, it must be isotopic to some conjugate of S. There are three cases to consider:

Case 1. T is isotopic to $S^{123} = S^{213}$.
Then T is isotopic to S and there is nothing to prove.

Case 2. T is isotopic to $S^{132} = S^{231}$.
Since T is symmetric, S^{132} must be isotopic to its transpose, namely S^{312}. But then $S^{132} = S^{231} \simeq S^{312} = S^{321}$, and four conjugates cannot be isotopic unless all six are, since the number of isotopic conjugates must divide 6. So M consists of a single isotopy class and S is isotopic to T.

Case 3. T is isotopic to $S^{312} = S^{321}$.
Since T is symmetric, S^{312} must be isotopic to its transpose, namely S^{132}. But then $S^{132} = S^{231} \simeq S^{312} = S^{321}$, and we reach the same conclusion as in the previous case. □

Having developed some of the theoretical framework, we are now in a position to prove the main result of the paper. Among other things, it says that two latin squares produced by the $\mathbb{K}$-construction are in the same main class only when they are isomorphic and so are the factorizations they came from.

Theorem 16. *Suppose that (F, u) and (G, v) are two rooted 1-factorizations of K_n. Then the following three statements are equivalent:*

(A) $\mathcal{L}(F, u) \sim \mathcal{L}(G, v)$.
(B) $\mathcal{L}(F, u) \cong \mathcal{L}(G, v)$.
(C) (F, u) is isomorphic to (G, v).

Proof. We first show that (A) is equivalent to (B). The implication (B) $\Rightarrow$ (A) is trivial since isomorphism is a special case of paratopy. So suppose (A). Then Lemma 4 together with Lemma 15 shows that $\mathcal{L}(F, u) \simeq \mathcal{L}(G, v)$ from which (B) then follows from Lemma 6. Hence (A) and (B) are equivalent.

Suppose (C), so that there is some isomorphism $\tilde{\rho}$ from the vertex set U of F to the vertex set V of G, which takes F to G and u to v. Define ρ a function from $[n]$ to $[n]$ by $i\rho = j$ iff $x\tilde{\rho} = y$ where x is the ith element of $U - \{u\}$ and y is the jth element of $V - \{v\}$. It is routine to verify that $\langle \rho, \rho, \rho \rangle$ maps $\mathcal{L}(F, u)$ to $\mathcal{L}(G, v)$, so that (C) $\Rightarrow$ (B).

The argument for the reverse implication is similar. If $\langle \rho, \rho, \rho \rangle$ is an isomorphism from $\mathcal{L}(F, u)$ to $\mathcal{L}(G, v)$ then we define a map $\tilde{\rho}$ from U to V by $u\tilde{\rho} = v$ and $x\tilde{\rho} = y$ where x is the ith element of $U - \{u\}$, y is the jth element of $V - \{v\}$ and $i\rho = j$. A routine check shows that this provides the required isomorphism from (F, u) to (G, v). □

An important family of 1-factorizations are those which are starter-induced (for the definition see [9] or [21], for example). By Theorem 16, there is either a unique main class or precisely two main classes of latin squares which can be built by applying the $\mathbb{K}$-construction to a given starter-induced 1-factorization of K_{n+1}. The first possibility arises when the factorization has a doubly transitive automorphism group. Cameron and Korchmáros [3] found a complete catalog of such 1-factorizations, showing that they only exist when $n \in \{5, 11, 27\}$ or $n + 1$ is a power of 2. Ihrig [9, p. 135] has recently shown that the only other possibility for starter-induced factorizations is that the automorphism group has exactly two distinct orbits on the vertices.

With regard to condition (C) in Theorem 16, we note that it is crucial that the isomorphism maps u to v. The examples in [10], [13], and [19] show that latin squares which are not paratopic to each other can be found in the case $F = G$ simply by varying the choice of the root. Thinking about the case $F = G$ also leads to the following.

Theorem 17. *The automorphism group of $\mathcal{L}(F, v)$ is isomorphic to the stabilizer of v in the group of automorphisms of F.*

Proof. Apply the proof that (B) and (C) are equivalent in Theorem 16 to the special case when $F = G$ and $u = v$. □

Combining this last result with Lemma 12, we immediately get the following.

Theorem 18. *Let (F, v) be a rooted 1-factorization of K_{p+1} for some prime p. Then either (F, v) is isomorphic to the patterned factorization $\big(GK(\mathbb{Z}_p), \infty\big)$, or else the autotopy group of $\mathcal{L}(F, v)$ is isomorphic to the stabilizer of v in the automorphism group of the (unrooted) factorization F.*

The above theorem explains why one of the atomic squares in [13] has trivial autotopy group. The square in question is $\mathcal{L}(F, v)$ where F is the 1-factorization labeled 'C' in the catalog of perfect 1-factorizations of K_{12} found by Petrenyuk and Petrenyuk [15], and v is any of the ten vertices other than 1 and 11 (using the labels from [15]). Petrenyuk and Petrenyuk noted that the automorphism group of factorization C is generated by a single cycle on these ten vertices, each of which therefore has a trivial stabilizer.

5. ALGORITHMIC ISSUES

In this final section, we investigate some algorithmic questions related to the work in the previous sections. Our primary goal is to bound the complexity of checking whether a given latin square can be constructed, up to paratopy, by the $\mathbb{K}$-construction.

By definition, the $\mathbb{K}$-construction always produces a symmetric latin square. However, it sometimes produces totally symmetric latin squares. A secondary aim of this section is to develop tools which might help to diagnose when this happens. One simple answer is that $\mathcal{L}(F, v)$ is totally symmetric if and only if F is a so-called *Steiner 1-factorization* and v is the infinity point (see, for example, [4, p. 24]). However, we may wish to interpret the question more broadly, to include latin

squares *which are paratopic* to a latin square produced by the $\mathbb{K}$-construction. We know from [13] that there are totally symmetric latin squares of this type of order 11, an order for which there are no Steiner triple systems.

In the results which follow we say that a latin square has a *symmetric form* if it is isotopic to a symmetric square and it has a *totally symmetric form* if it is isotopic to a totally symmetric square. Our next result will show that our algorithm can be used to diagnose how many isotopy classes there are in the main class of a given symmetric square.

Lemma 19. *If S is a symmetric latin square then either (a) S has a totally symmetric form or (b) S is not isotopic to any of its conjugates except itself and its transpose.*

Proof. Let M be the main class of a symmetric latin square S. Then M consists of either 3 or 1 isotopy classes. The first case gives case (b), so we may as well suppose that M consists of a single isotopy class. In that case, there must be an isotopy $\langle \rho, \sigma, \tau \rangle$ for which $S\langle \rho, \sigma, \tau \rangle = R$, where $R = S^{321}$. Define $T = S\langle \sigma, \sigma, \tau \rangle = R\langle \rho^{-1}\sigma, \varepsilon, \varepsilon \rangle$. Then T is symmetric since S is. Also $T = T^{132}$ because

$$R = S^{321} = (S^{213})^{321} = S^{312} = (S^{321})^{132} = R^{132}.$$

These two inherited properties ($T = T^{213}$ and $T = T^{132}$) are enough to guarantee that T is totally symmetric, so that we have case (a). □

Lemma 20. *A latin square L has a symmetric form if and only if there is some permutation which, when applied to the rows of L, produces a symmetric square.*

Proof. The "if" direction is trivial, so suppose there is some isotopy $\langle \rho, \sigma, \tau \rangle$ such that $L\langle \rho, \sigma, \tau \rangle$ is a symmetric square. Then $L\langle \rho, \sigma, \tau \rangle \langle \sigma^{-1}, \sigma^{-1}, \tau^{-1} \rangle = L\langle \rho\sigma^{-1}, \varepsilon, \varepsilon \rangle$ must also be symmetric, which proves the lemma. □

The following result was obtained by Sade [17] and can be proved by techniques similar to those used above.

Lemma 21. *A symmetric latin square S has a totally symmetric form if and only if there is some permutation which, when applied to the symbols of L, produces a totally symmetric square.*

Theorem 22. *Given a latin square L of order n, it can be established in $O(n^3)$ time whether L*

(a) *has a symmetric form*,
(b) *has a totally symmetric form*,
(c) *is paratopic to $\mathcal{L}(F, v)$ for some rooted 1-factorization (F, v) of K_{n+1}.*

Proof. To test part (a) we use Lemma 20. Choose each row of L in turn to be the first row of S, a supposed symmetric form of L. Once we know the first row, symmetry dictates the first column, which determines the order in which we must permute the rows of L if we hope to create S. Building the square S and testing if it is indeed symmetric takes $O(n^2)$ time, and there are n rows of L that we must consider as the possible first row of S.

To test whether L has a totally symmetric form T, we first test whether it has a symmetric form S using the above algorithm. Assuming that we find S, we then test if S is isotopic to some W satisfying $W = W^{132}$. This second test is conjugate to the first, and hence can also be done in cubic time (it can be viewed as applying the criterion in Lemma 21). For the existence of T, it is necessary and sufficient by Lemma 19 that both tests return a positive result.

To test (c), we check first the necessary condition that L has odd order. Assuming that is the case, we simply test (a) for each of the conjugates of L. If the result is positive for any conjugate then the answer to (c) is yes, otherwise the answer is no. This is because the squares which are $\mathcal{L}(F, v)$ for some F and v are exactly the symmetric idempotent squares, by Lemma 4. Any symmetric latin square of odd order has a transversal on its main diagonal and can be made idempotent and symmetric by a simple relabeling of its symbols. Since we have to run test (a) at most 6 times, the time complexity is still $O(n^3)$. (As an aside, it actually suffices to test just three conjugates of L, but this is not important to our result). □

For test cases for our algorithm, we refer the reader to the seven main classes of atomic squares given in [13]. Precisely six of those can be constructed by the $\mathbb{K}$-construction, and exactly two of those have a totally symmetric form.

Finally, we remark that although we phrased Theorem 22 in terms of yes/no decision problems, the method of proof shows that the algorithms can easily be written to be constructive. That is, in the case of a "yes" answer, we can return a structure which proves the answer, rather than just claiming that one exists. A question we have not considered, but which would be of interest for further research is this: Is it possible to identify structural properties of those latin squares which have a symmetric form? Ideally, such invariants would be powerful enough so that, at least in some cases, the presence or not of a symmetric form could be decided without use of the above algorithm.

An important special case is when we know our latin square is the Cayley table of a group. We finish with an observation about this case.

Lemma 23. *Suppose L is the Cayley table of a finite group G. If G is Abelian then L has a totally symmetric form, otherwise it does not even have a symmetric form.*

Proof. Let $g_1, \ldots, g_n$ be the elements of G, where g_1 is the identity. Define a latin square M, with rows and columns indexed by the elements of G, by the rule that $g_i^{-1} g_j^{-1}$ occurs in row g_i, column g_j. It is a simple matter to verify that M is isotopic to the Cayley table of G and that it is totally symmetric if G is Abelian.

It remains to show that L has no symmetric form unless G is Abelian. Suppose that the algorithm in Theorem 22 finds a symmetric form S of L when row h is chosen to be the first row. The first row (and hence also the first column) of S is therefore $[h \;\; hg_2 \;\; hg_3 \;\; \cdots \;\; hg_n]$. Thus the entry of S in a general row g_i and column g_j is hg_ig_j. For S to be symmetric, we must have $hg_ig_j = hg_jg_i$, which implies $g_ig_j = g_jg_i$, for arbitrary group elements g_i and g_j. That is, G is Abelian. □

ACKNOWLEDGMENT

The authors record their gratitude for their useful discussions with B. M. Maenhaut, D. Bryant and A. Rosa.

REFERENCES

[1] D. Bryant, B. M. Maenhaut, and I. M. Wanless, A family of perfect factorisations of complete bipartite graphs, J Combin Theory Ser A 98 (2002), 328–342.

[2] D. Bryant, B. M. Maenhaut, and I. M. Wanless, New families of atomic Latin squares and perfect one-factorisations, submitted for publication.

[3] P. J. Cameron and G. Korchmáros, One-factorizations of complete graphs with a doubly transitive automorphism group, Bull London Math Soc 25 (1993), 1–6.

[4] C. J. Colbourn and A. Rosa, Triple systems, Clarendon, Oxford, 1999.

[5] J. Dénes and A. D. Keedwell, Latin squares and 1-factorizations of complete graphs, I, Connections between the enumeration of latin squares and rectangles and r-factorizations of labelled graphs, Ars Combin 25A (1988), 109–126.

[6] J. Dénes and A. D. Keedwell, Latin squares and their applications, Akadémiai Kiadó, Budapest, 1974.

[7] J. Dénes and A. D. Keedwell, Latin squares: New developments in the theory and applications, Annals Discrete Math, Vol. 46, North-Holland, Amsterdam, 1991.

[8] D. Duncan and E. Ihrig, The structure of symmetry groups of $GK(2n, G)$, J Combin Des 2 (1994), 341–349.

[9] E. Ihrig, Conditions ensuring starter induced 1-factorizations are not isomorphic, J Combin Des 11 (2003), 124–143.

[10] A. D. Keedwell, Uniform P-circuit designs, quasigroups, and Room squares, Utilitas Math 14 (1978), 141–159.

[11] A. Kotzig, Groupoids and partitions of complete graphs, Combinatorial Structures and their Applications, Gordon and Breach, New York, 1970, pp. 215–221.

[12] P. J. Laufer, On strongly Hamiltonian complete bipartite graphs, Ars Combin 9 (1980), 43–46.

[13] B. M. Maenhaut and I. M. Wanless, Atomic latin squares of order eleven, J Combin Des 12 (2004), 12–34.

[14] P. J. Owens and D. A. Preece, Some new non-cyclic latin squares that have cyclic and Youden properties, Ars Combin 44 (1996), 137–148.

[15] L. Petrenyuk and A. Petrenyuk, Intersection of perfect one-factorizations of complete graphs, Cybernetics 16 (1980), 6–9.

[16] A. Sade, Quasigroupes parastrophiques: Expressions et identités, Math Nachr 20 (1959), 73–106.

[17] A. Sade, Critères d'isotopie d'un quasigroupe avec un quasigroupe demi-symétrique, Univ Lisboa Revista Fac Ci A 11(2) (1964/1965), 121–136.

[18] M. Suzuki, Group theory II, Grundlehren der Mathematischen Wissenschaften, 248, Springer-Verlag, New York, 1986.

[19] I. M. Wanless, Perfect factorisations of bipartite graphs and latin squares without proper subrectangles, Electron J Combin 6 (1999), R9, 16 pp.

[20] I. M. Wanless, Atomic latin squares based on cyclotomic orthomorphisms, (submitted for publication).

[21] W. D. Wallis, One-factorizations, Math Appl, Vol. 390, Kluwer Academic, Dordrecht, 1997.

[22] L. Xu, V. Bohossian, J. Bruck, and D. G. Wagner, Low-density MDS codes and factors of complete graphs, IEEE Trans Inform Theory 45 (1999), 1817–1826.