ALGEBRAIC COMBINATORICS

Peter J. Dukes & Kate I. Nimegeers

The linear system for Sudoku and a fractional completion threshold

Volume 7, issue 5 (2024), p. 1283-1305.
https://doi.org/10.5802/alco.375

Algebraic Combinatorics is published by The Combinatorics Consortium
and is a member of the Centre Mersenne for Open Scientific Publishing
www.tccpublishing.org www.centre-mersenne.org
e-ISSN: 2589-5486

Algebraic Combinatorics
Volume 7, issue 5 (2024), p. 1283–1305
https://doi.org/10.5802/alco.375

The linear system for Sudoku and a fractional completion threshold

Peter J. Dukes & Kate I. Nimegeers

Abstract We study a system of linear equations associated with Sudoku latin squares. The coefficient matrix M of the normal system has various symmetries arising from Sudoku. From this, we find the eigenvalues and eigenvectors of M, and compute a generalized inverse. Then, using linear perturbation methods, we obtain a fractional completion guarantee for sufficiently large and sparse rectangular-box Sudoku puzzles.

1. Introduction

A *latin square* of order n is an $n \times n$ array with entries from a set of n symbols (often taken to be $[n] := \{1, 2, \dots, n\}$) having the property that each symbol appears exactly once in every row and every column. A *partial latin square* of order n is an $n \times n$ array whose cells are either empty or filled with one of n symbols in such a way that each symbol appears at most once in every row and every column. A partial latin square can be identified with a set of ordered triples in a natural way: if symbol k appears in row i and column j, we include the ordered triple (i, j, k). A *completion* of a partial latin square P is a latin square L which contains P in the sense of ordered triples; that is, every symbol occurring in P also occurs in the corresponding cell of L.

It is natural to ask how dense a partial latin square can be while still having a completion. Daykin and Häggkvist conjectured [7] that a partial latin square in which any row, column and symbol is used at most $n/4$ times should have a completion. They proved a weaker version of this claim, with $n/4$ replaced by $2^{-9}n^{1/2}$. Chetwynd and Häggkvist [6] and later Gustavsson [8] obtained the first such completion guarantee which was linear in n. Let us say that a partial latin square is ϵ-dense if no row, column, or symbol is used more than ϵn times. Bartlett [3] built on the preceding work to show that all ϵ-dense partial latin squares have a completion for $\epsilon = 9.8 \times 10^{-5}$. Then, over two papers, this was improved to roughly $\epsilon = 0.04$, provided n is large. One paper [4] of Bowditch and Dukes obtained this threshold for a fractional relaxation of the problem, and the other paper [2] by Barber, Kühn, Lo, Osthus and Taylor showed using absorbers and balancing graphs that the fractional threshold suffices for very large instances of the (exact) completion problem.

Let h and w be integers with $h, w \geqslant 2$, and put $n = hw$. A *Sudoku latin square* (or briefly *Sudoku*) of type (h, w) is an $n \times n$ latin square whose cells are partitioned

Manuscript received 6th November 2023, revised 16th April 2024, accepted 20th April 2024.

Keywords. Sudoku, latin square, coherent configuration, perturbation.

Acknowledgements. Research of Peter Dukes is supported by NSERC Discovery Grant RGPIN–2017–03891.

into a $w \times h$ pattern of $h \times w$ subarrays where every symbol appears exactly once in each subarray. The subarrays are called *boxes*, or sometimes also called *cages*. A partial Sudoku and completion of such is defined analogously as above for latin squares. The completion problem for partial Sudoku in the case $h = w = 3$ is a famous recreational puzzle. A mathematical discussion of Sudoku solving strategies can be found in [5, 14]. By contrast, we are interested here in the fractional relaxation of partial Sudoku completion, essentially following the approach used in [4] for latin squares.

Let us explain the fractional relaxation in more detail. Working from a partial latin square P, an empty cell can be assigned a convex combination of symbols instead of a single symbol. More formally, P can be represented as a function $f_P : [n]^3 \to \{0,1\}$ in which $f_P(i,j,k)$ is the number of times symbol k appears in cell (i,j). A *fractional completion* of P is a function $f : [n]^3 \to [0,1]$ such that, for any $i,j,k \in [n]$,

- $f_P(i,j,k) = 1$ implies $f(i,j,k) = 1$; and
- $\sum_{i=1}^n f(i,j,k) = \sum_{j=1}^n f(i,j,k) = \sum_{k=1}^n f(i,j,k) = 1$.

Viewing this as an array, cell (i,j) is assigned a fractional occurrence of symbol k with value $f(i,j,k)$. The first condition ensures that filled cells of P are left unchanged. The second condition ensures that every symbol appears with a total occurrence of one in each column and each row, and that every cell is used with a total value of one. For the Sudoku setting, we can add an extra family of constraints, namely that for all boxes b and symbols k, $\sum_{(i,j)\in b} f(i,j,k) = 1$, where the sum is over all ordered pairs (i,j) belonging to box b. We remark that when f is $\{0,1\}$-valued, it corresponds to an exact completion of P, whether for general latin squares or Sudoku.

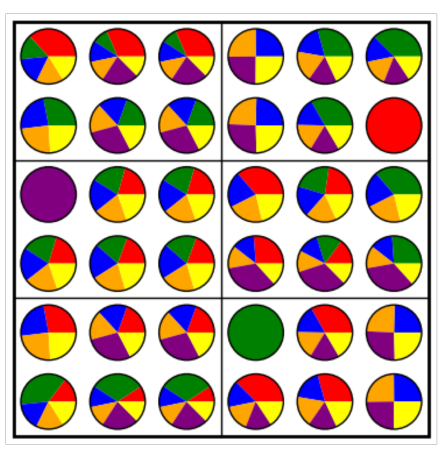

Figure 1. Illustration of a fractional completion of a partial Sudoku of type $(2,3)$

Figure 1 depicts a fractional Sudoku of type $(2,3)$, where the solid disks correspond to pre-filled cells of a partial Sudoku, and the multi-colored disks correspond to a fractional completion.

The notion of ϵ-dense needs to be strengthened in the Sudoku setting. First, it is natural to impose a constraint on number of filled cells in any box. Otherwise, completion can be blocked by placing symbols $1,\dots,n-1$ in the top-left box and symbol n in line with the remaining empty cell of that box. This uses each symbol only once and each row and column at most $\max\{h,w\}$ times. For $h \approx w$, this is sub-linear in n. Separately from this, it is also natural to prevent symbol occurrences from being too unbalanced relative to the box partition. In more detail, for a Sudoku of type (h, w), it is possible to force a given symbol in, say, the (1,1)-entry by placing it outside of the top-left box in rows $2, \ldots, h$ and in columns $2, \ldots, w$. We can arrange for this to occur for different symbols, say 1 and 2, making completion impossible. But this uses no row or column more than twice and uses each symbol only $h+w-2$ times, which could again be sub-linear in n. An illustration of the obstructions for $h = w = 3$ are given in Figure 2. To address this, we can strengthen the ϵ-dense condition for a partial Sudoku to:

- each row, column, and box has at most ϵn filled cells;
- each symbol occurs at most ϵh times in any bundle of h rows corresponding to the box partition, and likewise at most ϵw times in any bundle of w columns.

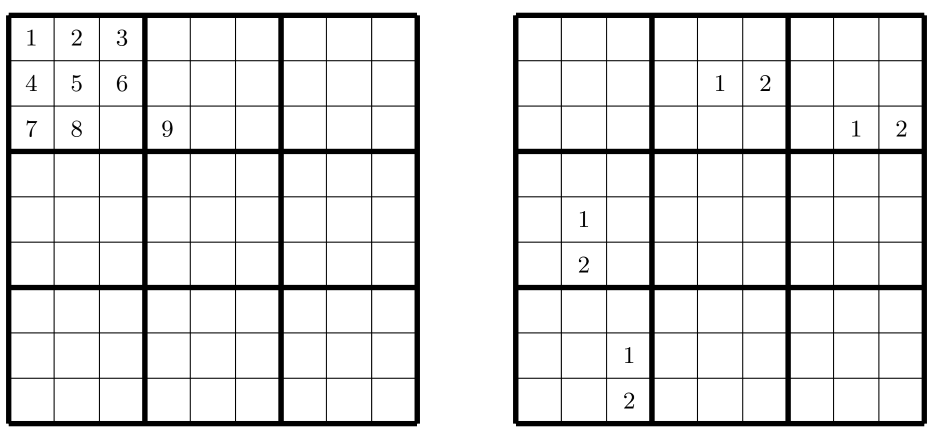

FIGURE 2. Sparse partial latin squares with no Sudoku completion

Our main result gives a guarantee on fractional completion of ϵ-dense Sudoku latin squares.

THEOREM 1.1. *Let $\epsilon < 1/101$. For sufficiently large integers h and w, every ϵ-dense partial Sudoku of type (h, w) has a fractional completion.*

It turns out that our methods to prove Theorem 1.1 really only require a weaker notion of density. Roughly speaking, we need each empty cell (i, j) to have a large proportion of all symbols k available to be placed there, and analogous availability when the roles of rows, columns and symbols are permuted, and boxes are introduced. This is made more precise later.

The outline of the paper is as follows. In the next section, we reformulate Sudoku completion as a certain graph decomposition problem, and then give a linear system of equations governing the fractional relaxation. Starting with the empty Sudoku, the rank and a basis for the nullspace are computed and interpreted combinatorially. The ϵ-dense case can be viewed as a perturbed version of the empty case, closely following the approach in [4] for latin squares. This motivates a study of the linear algebra for empty Sudoku in more detail. In Section 3, we observe that all relevant computations take place in an adjacency algebra of fixed dimension, independent of h and w. Eigenvalues and eigenvectors relevant for the problem are described in Section 4.1. Then, using computer-assisted symbolic algebra, a certain generalized inverse is computed and upper-bounded in ∞-norm. This bound ensures solutions to our linear system remain nonnegative. Section 5 discusses in more detail the perturbation as we pass to the ϵ-dense setting. By the end of this section, all the ingredients are in place to prove Theorem 1.1. The last section contains some concluding remarks and a brief discussion of possible extensions and generalizations.

2. Preliminaries and set-up

2.1. A graph decomposition model. In a Sudoku of type (h,w), let $\mathrm{box}(i,j)$ denote the box containing cell (i,j). If boxes are numbered left to right then top to bottom, then we have the formula $\mathrm{box}(i,j) = h\lfloor \frac{i-1}{h} \rfloor + \lfloor \frac{j-1}{w} \rfloor + 1$.

We define the graph G_{hw} as follows. Its vertex set is

$$V(G_{hw}) = \{r_1,\dots,r_n\} \cup \{c_1,\dots,c_n\} \cup \{b_1,\dots,b_n\} \cup \{s_1,\dots,s_n\},$$

with the four sets corresponding to rows, columns, boxes, and symbols, respectively. Its edge set is

$$E(G_{hw}) = \bigcup_{i,j=1}^{n} \{\{r_i,c_j\},\{r_i,s_j\},\{c_i,s_j\},\{b_i,s_j\}\}. \tag{1}$$

In other words, exactly one edge is present for every combination of row-column, row-symbol, column-symbol, and box-symbol. As a point of notation, G_{hw} depends only on $n = hw$; indeed, if we omit indices in (1) it is seen to be simply a blow-up of the graph $K_3 + e$ by independent sets of size n. With this in mind, the subscript in G_{hw} can reasonably be interpreted as the product of h and w, though it will be useful below to keep these parameters separate. Note that the subgraph of G_{hw} induced by rows, columns, and symbols is just the complete 3-partite graph $K_{n,n,n}$.

A *tile* in G_{hw} is a copy of $K_3 + e$ induced by four vertices r_i, c_j, b_ℓ, s_k for which $\mathrm{box}(i,j) = \ell$. This tile represents the act of placing symbol k in cell (i,j), and also keeps track of the box used. Let $T(G_{hw})$ denote the set of all n^3 tiles in G_{hw}.

Given a partial Sudoku S of type (h,w), let G_S denote the subgraph of G_{hw} obtained by removing the edge sets of tiles corresponding to filled cells of S. In other words, $V(G_S) = V(G_{hw})$ and $E(G_S)$ contains:

- $\{r_i, c_j\}$ if and only if cell (i,j) is empty;
- $\{r_i, s_k\}$ if and only if symbol k is missing in row i;
- $\{c_j, s_k\}$ if and only if symbol k is missing in column j;
- $\{b_\ell, s_k\}$ if and only if symbol k is missing in box ℓ.

Let $T(G_S)$ be the set of tiles in $T(G_{hw})$ all of whose edges are in G_S.

An equivalent but slightly different model can be obtained by including row-box and column-box edges. That is, we could change tiles into cliques K_4, and change the host graph G_{hw} into a multigraph G^*_{hw} with the same vertex set and all the edges of G_{hw}, and additionally including the edges

- $\{r_i, b_\ell\}$ with multiplicity w if and only if row i is incident with box ℓ; and
- $\{c_j, b_\ell\}$ with multiplicity h if and only if column j is incident with box ℓ.

Likewise, given a partial Sudoku S, we could define G^*_S to be the subgraph of G^*_{hw} obtained by removing the edges of all 4-cliques on $\{r_i, c_j, b_\ell, s_k\}$ whenever symbol k occurs in cell (i,j) (and box ℓ) of S.

For graphs F and G, we say that G has an *F-decomposition* if its edge set $E(G)$ can be partitioned into subgraphs isomorphic to F. This extends naturally to multigraphs G, where now repeated edges are distinguished. That is, the number of copies of F containing two vertices $u \neq v$ equals the multiplicity of edge $\{u,v\}$ in G. Many problems in combinatorics can be formulated in terms of graph decompositions. For example, K_3-decompositions of $K_{n,n,n}$ are equivalent to latin squares of order n; see for instance [2, 4, 12]. The following is an analog for Sudoku using our graphs above.

Proposition 2.1. *The partial Sudoku S has a completion if and only if the graph G_S has an edge-decomposition into tiles, or equivalently if and only if G_S^* has a K_4-decomposition.*

Proof. Suppose S has a completion S'. If cell (i,j) was blank in S but filled in S', say with symbol k, we use the tile defined by $\{r_i, c_j, s_k, b_\ell\}$, where $\ell = \mathrm{box}(i,j)$. Each such tile belongs to $T(G_S)$ because (i,j) was blank in S and because k occurs only once in S' in row i, column j and box ℓ. Consider the set $\mathcal{T}$ of these tiles induced by cells that were blank in S and filled in S'. These tiles are edge-disjoint, again because S' has no repeated symbols in any any row, column, or box. We check that $\mathcal{T}$ gives an edge-decomposition of G_S into tiles. Any row-column edge of G_S, say $\{r_i, c_j\}$, is in the tile corresponding to the symbol placed at entry (i,j) of S'. Consider a row-symbol edge, say $\{r_i, s_k\} \in E(G_S)$. The presence of this edge means k was missing from row i in S. It occurs somewhere in row i of S', say at entry (i,j). This entry was blank in S, so r_i, c_j, s_k define a tile in $\mathcal{T}$, along with the box containing (i,j). A similar verification holds for edges of type column-symbol and box-symbol in G_S.

For the converse, the argument is reversible. Given a set $\mathcal{T}$ of tiles that form an edge-decomposition of G_S, we complete S by placing symbol k in entry (i,j) whenever r_i, c_j, s_k belong to a tile of $\mathcal{T}$. Since the tiles of $\mathcal{T}$ are edge-disjoint, every entry (i,j) is filled at most once and no row, column, or box contains repeats. Since the edges within $\mathcal{T}$ partition those in G_S, it follows that every blank entry of S gets filled, and every symbol occurs in every row, column, and box.

The claim about G_S^* having a K_4-decomposition is nearly identical. In the forward implication, we note that a row-box edge $\{r_i, b_\ell\}$ in G_{hw}^* occurs in total w times counting $E(G_S^*)$ and the decomposition. This is because the completion S' has w entries in row i and box ℓ. Similarly, column-box edges of G_{hw}^* occur a total of h times. □

The model using row-box and column-box edges has the advantage that all 4-cliques in G_S^* correspond to valid tiles. However, since no new information is carried by those extra edges, we henceforth work with tiles in G_S, omitting the implied edges of type row-box and column-box.

This paper is concerned with partial Sudoku latin squares which are nearly empty. Recall the definition of ϵ-dense discussed in Section 1 and strengthened for the case of Sudoku. The definition leads easily to various degree bounds in G_S, which we summarize here.

Lemma 2.2. *Suppose S is ϵ-dense. Then in the graph G_S, the number of edges from vertex*

- *c_j to the row partite set is at least $(1-\epsilon)n$;*
- *s_k to any row bundle is at least $(1-\epsilon)h$;*
- *r_i to the column partite set is at least $(1-\epsilon)n$;*
- *s_k to any column bundle is at least $(1-\epsilon)w$;*
- *r_i, c_j or b_ℓ to the symbol partite set is at least $(1-\epsilon)n$;*
- *s_k to the box partite set is at least $(1-\epsilon)n$.*

2.2. The linear system for Sudoku. Consider an empty $n \times n$ Sudoku, to be filled with entries from $[n]$. Let x_{ijk} denote the number/fraction of symbols s_k placed in cell (i,j), where $(i,j,k) \in [n]^3$. Latin square and Sudoku constraints naturally correspond to linear equations on these variables. The condition that every cell have exactly one entry becomes $\sum_k x_{ijk} = 1$ for each $(i,j) \in [n]^2$. The condition that every row contains every symbol exactly once becomes $\sum_j x_{ijk} = 1$ for each $(i,k) \in [n]^2$. Similarly, that every column contains every symbol exactly once becomes $\sum_i x_{ijk} = 1$

for each $(j,k) \in [n]^2$. Together, these $3n^2$ equations yield a linear system for (fractional) latin squares. The additional condition relevant for Sudoku is that every box contains every symbol exactly once, or

$$\sum_{(i,j)\in \mathrm{box}(\ell)} x_{ijk} = 1$$

for each $(k,\ell) \in [n]^2$.

This results in a $4n^2 \times n^3$ linear system

$$W\mathbf{x} = \mathbb{1}, \tag{2}$$

where $\mathbb{1}$ denotes the all-ones vector and W is the $\{0,1\}$ inclusion matrix of $E(G_{hw})$ versus $T(G_{hw})$; that is, $W(e,t) = 1$ if $e \in t$ and is 0 otherwise. In this paper, we will mainly consider the (square) normal system, with coefficient matrix $M = WW^\top$. An entrywise nonnegative solution $\mathbf{y}$ to $M\mathbf{y} = \mathbb{1}$ implies the existence of a solution $\mathbf{x} = W^\top\mathbf{y} \geqslant \mathbf{0}$ to (2). This, in turn, produces a fractional edge-decomposition of G_{hw} into tiles and a fractional Sudoku of type (h,w).

The rank (over the reals) of both W and M can be found by exhibiting a basis for their range consisting of tiles in G_{hw}. For convenience, set punctuation will be omitted from edges and tiles; we abbreviate these by juxtaposing vertices in different partite sets of $V(G_{hw})$.

Proposition 2.3. *We have* $\mathrm{rank}(M) = \mathrm{rank}(W) = n^3-(n-1)^3+(n-1)(h-1)(w-1)$.

Proof. Let $\mathcal{T}_1$ be the set of $n^3-(n-1)^3$ tiles in G_{hw} which intersect at least one of r_1, c_1, s_1. It was shown in [4, Proposition 2.3] that $\mathcal{T}_1$ is linearly independent in the vector space of functions from $E(K_{n,n,n})$ to $\mathbb{R}$. Thus, $\mathcal{T}_1$ is also linearly independent in $\mathbb{R}^{E(G_{hw})}$. Let $\mathcal{T}_2$ be any set of $(n-1)(h-1)(w-1)$ tiles of the form $r_ic_js_kb_\ell$, where $k = 2,\dots,n-1$ and the b_ℓ range over all boxes disjoint from both row 1 and column 1. Since the tiles in $\mathcal{T}_2$ use distinct box-symbol edges which are not present in $\mathcal{T}_1$, it is clear that $\mathcal{T}_1 \cup \mathcal{T}_2$ is linearly independent.

We next show that $\mathcal{T}_1 \cup \mathcal{T}_2$ generates any given column of W, say the one corresponding to a tile $\{r_i, c_j, s_k, b_\ell\}$. Suppose $i \leqslant h$ and $j \leqslant w$. Then $\ell = \mathrm{box}(i,j) = 1$. We can form the target tile as a linear combination in $\mathcal{T}_1$, namely as

$$r_ic_js_kb_1 = r_1c_js_kb_1 + r_ic_1s_kb_1 + r_ic_js_1b_1 - r_1c_1s_kb_1 - r_1c_js_1b_1 - r_ic_1s_1b_1 + r_1c_1s_1b_1.$$

Suppose next that $i > h$ and $j \leqslant w$. As above, we have the linear combination

$$r_ic_js_kb_\ell = r_1c_js_kb_1 + r_ic_1s_kb_\ell + r_ic_js_1b_\ell - r_1c_1s_kb_1 - r_1c_js_1b_\ell - r_ic_1s_1b_1 + r_1c_1s_1b_1.$$

Similarly, $\mathcal{T}_1$ generates any tile with $i \leqslant h$ and $j > w$. Suppose, then, that $i > h$ and $j > w$. If $k = 1$, the corresponding tile belongs to $\mathcal{T}_1$, so assume $k > 1$. Put $p = \mathrm{box}(1,j)$ and $q = \mathrm{box}(i,1)$, and note that all tiles meeting these boxes are in the span of $\mathcal{T}_1$, as shown above. Since $i > h$, $j > w$, and $k > 1$, we know that s_k and b_ℓ occur together in some tile $r_{i'}c_{j'}s_kb_\ell \in \mathcal{T}_2$. Using this and other tiles generated so far, we compute

$$\begin{aligned} &r_ic_js_kb_\ell \\ &= r_ic_js_1b_\ell + r_1c_js_kb_p + r_ic_1s_kb_q - r_1c_1s_kb_1 - r_1c_js_1b_p - r_ic_1s_kb_q + r_{i'}c_{j'}s_kb_\ell \\ &- r_{i'}c_{j'}s_1b_\ell - r_1c_{j'}s_kb_p - r_{i'}c_1s_kb_q + r_1c_1s_kb_1 + r_1c_{j'}s_1b_p + r_{i'}c_1s_1b_q - r_1c_1s_1b_1. \end{aligned}$$

We have shown that $\mathcal{T}_1 \cup \mathcal{T}_2$ spans each column of W, and hence is a basis for $\mathrm{range}(W)$. □

Suppose now that some cells of our Sudoku have been pre-filled, resulting in the graph G_S as described earlier. Let W_S denote the $\{0,1\}$ inclusion matrix of edges versus tiles in G_S. Then the system

$$W_S \mathbf{x} = \mathbb{1} \tag{3}$$

has a solution $\mathbf{x} \geqslant \mathbf{0}$ if and only if G_S admits a fractional edge-decomposition into tiles. Note that for non-empty S, the dimensions in the system (3) are smaller than in (2). The tile weights are given by entries x_{ijk} of $\mathbf{x}$. By Proposition 2.1, the existence of such a solution is equivalent to our partial Sudoku S having a completion.

We may again consider the normal system with coefficient matrix $M_S = W_S W_S^\top$. Even though many possible solutions of (3) are lost in doing so, the normal system has the advantage of allowing eigenvalue and perturbation methods, as was done in [4]. We extend these methods to the Sudoku setting in Sections 4 and 5 to follow.

2.3. The kernel. For our analysis of the linear systems (2) and (3) above, it is important to study the nullspace/kernel of M, or equivalently the left nullspace of W. This can be viewed as the set of all edge-weightings in G_{hw} in which each tile has a vanishing total weight (over its four edges).

After some simplification, Proposition 2.3 gives

$$\dim \ker(M) = 4n^2 - \operatorname{rank}(M) = 3n + (h+w)(n-1). \tag{4}$$

We would like to find a basis for $\ker(M)$. First, a spanning set is described in three categories of vectors below.

(A) Choose a row r_i and define the vector $\mathbf{v}$, coordinates indexed by $E(G_{hw})$, where $\mathbf{v}(r_i c_j) = 1$ for all $j \in [n]$, $\mathbf{v}(r_i s_k) = -1$ for all $k \in [n]$, and otherwise $\mathbf{v}(e) = 0$. Similar classes of kernel vectors exist with the roles of row, column and symbol permuted.

Consider the characteristic vector $\mathbf{t}$ of some tile t. If $r_i \in t$, then since t contains exactly one column and exactly one symbol we have $\mathbf{t} \cdot \mathbf{v} = 1 - 1 = 0$. On the other hand, if $r_i \notin t$, the support of t is disjoint from the support of $\mathbf{v}$, hence we again have $\mathbf{t} \cdot \mathbf{v} = 0$. In plain language, this kernel vector encodes the condition that the number of times a column is used with row i equals the number of times a symbol is used in row i. Verification is similar for the permuted varieties.

(B) Choose a box b_ℓ and define the vector $\mathbf{v}$ in which $\mathbf{v}(r_i c_j) = 1$ for all $(i,j) \in b_\ell$, $\mathbf{v}(b_\ell s_k) = -1$ for all $k \in [n]$, and otherwise $\mathbf{v}(e) = 0$.

As before, let $\mathbf{t}$ be the characteristic vector of a tile t. If $b_\ell \in t$, then since t contains exactly one row, column and symbol, we have $\mathbf{t} \cdot \mathbf{v} = 1 - 1 = 0$. On the other hand, if $b_\ell \notin t$, the supports of $\mathbf{t}$ and $\mathbf{v}$ are disjoint. This kernel vector encodes the condition that the number of entries filled in box ℓ equals the number of symbols used in box ℓ.

(C) Choose a bundle of rows $\{r_{hp+1}, \dots, r_{h(p+1)}\}$ and a symbol s_k. Define the vector $\mathbf{v}$ with $\mathbf{v}(r_i s_k) = 1$ and $\mathbf{v}(b_\ell s_k) = -1$ for all $hp < i, \ell \leqslant h(p+1)$, and otherwise $\mathbf{v}(e) = 0$. A similar class of vectors exists for column bundles.

Once again, let $\mathbf{t}$ be the characteristic vector of a tile t. Suppose $s_k \in t$ and t intersects the row bundle defining $\mathbf{v}$. Since t contains exactly one row and exactly one box meeting this bundle, we have $\mathbf{t} \cdot \mathbf{v} = 1 - 1 = 0$. On the other hand, if either $s_k \notin t$ or t intersects a different row bundle, the supports of $\mathbf{t}$ and $\mathbf{v}$ are disjoint. The encoded condition states that a given symbol s_k appears the same number of times in a row bundle as in the corresponding box bundle. The column bundle case has a similar verification.

To see that the vectors above span $\ker(M)$, it is convenient to examine the subspace U spanned by the row and column varieties of type (A), along with all vectors of type (B). Note that all vectors in U are invariant under symbol permutation. It is straightforward to see that $\dim(U) = 3n$, since basis vectors have disjoint supports.

Next, let V be the subspace of $\ker(M)$ spanned by the type (C) vectors. It is easy to see that $\dim(V) = (h+w)n$ and $\dim(U \cap V) \leqslant h + w$, the number of row or column bundles. Therefore,

$$\dim(U+V) = \dim(U) + \dim(V) - \dim(U \cap V) \geqslant 3n + (h+w)(n-1) = \dim \ker(M).$$

For a basis of $\ker(M)$, it suffices to take the row and column varieties of type (A) vectors, all type (B) vectors, and those of type (C) which avoid a particular symbol, say s_n. More details on this can be found in the second author's thesis [13].

Let K be the $4n^2 \times 4n^2$ matrix which projects onto $\ker(M)$. We have $K^2 = K = K^\top$. Let $K[S]$ denote the principal submatrix of K whose rows and columns correspond to the edges of G_S. The following orthogonality relations are similar to [4, Proposition 2.5].

Proposition 2.4. *The range of $K[S]$ is orthogonal to the all-ones vector and to the range of M_S. That is,* (a) $K[S]\mathbb{1} = \mathbf{0}$*; and* (b) $K[S]M_S = O$.

Proof. For matrices and vectors indexed by $E(G_{hw})$, sort the indices so that those corresponding to $E(G_S)$ come first. Let L be the inclusion map from edges of G_S to edges of G_{hw} and let Q be the inclusion map from tiles of G_S to tiles of G_{hw}. As matrices, L and Q have the structure $[I \mid O]^\top$.

Let $\mathbb{1}_S = (\mathbb{1} \mid \mathbf{0})$ be the $4n^2 \times 1$ zero-one indicator vector of $E(G_S)$ in $E(G_{hw})$. Alternatively, $\mathbb{1}_S$ is obtained from the $4n^2 \times 1$ all-ones vector by subtracting indicator vectors of tiles in S. It follows that $\mathbb{1}_S$ is contained in the range of $W^\top$, and hence is orthogonal to $\ker(W^\top)$. We now compute

$$K[S]\mathbb{1} = L^\top K L \mathbb{1} = L^\top K \mathbb{1}_S = \mathbf{0}.$$

This proves (a). With our matrix partition, we have

$$W = \left[\begin{array}{c|c} W_S & * \\ \hline O & * \end{array}\right]$$

and $LW_S = WQ$. Working from these,

$$K[S]M_S = L^\top K L W_S W_S^\top = L^\top K W Q W_S^\top = O,$$

since $KW = O$. This proves (b). ☐

Next, we recall [4, Lemma 2.6]. The idea lets us solve an under-determined system $A\mathbf{x} = \mathbf{b}$ by inverting an additive shift of A. We use this later in Section 5 with A taking the role of M_S, B a multiple of $K[S]$, and $\mathbf{b} = \mathbb{1}$.

Lemma 2.5 (see [4]). *Let A and B be symmetric $N \times N$ real matrices with $AB = O$, $A + B$ nonsingular, and $B\mathbf{b} = \mathbf{0}$. Then $A(A+B)^{-1}\mathbf{b} = \mathbf{b}$.*

3. The adjacency algebra

3.1. A coherent configuration for Sudoku. Let X be a finite set. A *coherent configuration* on X is a partition of $X \times X$ into a set of relations $\mathcal{R} = \{R_1, \dots, R_d\}$ satisfying the following properties:

(1) the union of some relations in $\mathcal{R}$ equals the diagonal $\{(x,x) : x \in X\}$;
(2) for each R in $\mathcal{R}$, the transpose relation $R^\top = \{(y,x) : (x,y) \in R\}$ is also in $\mathcal{R}$;
(3) for each i, j, k, there exists a constant p^k_{ij} such that for any x, z with $(x,z) \in R_k$, there are exactly p^k_{ij} elements y such that $(x,y) \in R_i$ and $(y,z) \in R_j$.

Given a group G acting on a set X, the set of orbits of the induced action on $X \times X$ is a coherent configuration, [9]. Here, we set up a coherent configuration on the ground set $X = E(G_{hw})$ using a group of Sudoku symmetries. The wreath product $\mathcal{S}_h \wr \mathcal{S}_w$ acts on rows and preserves the row bundle partition. Similarly, $\mathcal{S}_w \wr \mathcal{S}_h$ acts on columns. The direct product of these acts on $[n]^2$ and in particular the second factors $\mathcal{S}_w \times \mathcal{S}_h$ act on the n boxes. Finally, if we take a product with $\mathcal{S}_n$ for symbol symmetries, we have the group

$$\Gamma_{hw} = (\mathcal{S}_h \wr \mathcal{S}_w) \times (\mathcal{S}_w \wr \mathcal{S}_h) \times \mathcal{S}_n \leqslant \mathrm{Aut}(G_{hw})$$

acting on $E(G_{hw}) = \{r_i c_j, r_i s_k, c_j s_k, b_\ell s_k : i, j, k, \ell \in [n]\}$.

We describe the relations on ordered pairs of edges in more detail. Given two rows r_i and $r_{i'}$, we write $r_i \sim r_{i'}$ if and only if $\lfloor (i-1)/h \rfloor = \lfloor (i'-1)/h \rfloor$. Similarly, given two columns c_j and $c_{j'}$, we write $c_j \sim c_{j'}$ iff $\lfloor (j-1)/w \rfloor = \lfloor (j'-1)/w \rfloor$. In other words, $\sim$ tracks whether two rows or two columns belong to the same bundle. From the definition, it is clear that $\sim$ is an equivalence relation on both rows and columns. Write $r_i \ncong r_{i'}$ if $r_i \sim r_{i'}$ but $r_i \neq r_{i'}$. Define $\ncong$ similarly for columns.

Given two boxes b_ℓ and $b_{\ell'}$, write $b_\ell \smile b_{\ell'}$ if $\lfloor (\ell-1)/h \rfloor = \lfloor (\ell'-1)/h \rfloor$. Informally, this keeps track of whether the two boxes occur in the same row bundle. Write $b_\ell \frown b_{\ell'}$ iff $\ell \equiv \ell' \pmod{h}$; this is the analog for boxes in the same column bundle. By abuse of notation, we write $r_i \smile b_\ell$ to mean that the corresponding row and box intersect, and similarly for $c_j \frown b_\ell$.

Ordered pairs of vertices in G_{hw} are partitioned into relations according to Table 1. A blank indicates the trivial relation. Moving from vertices to edges, there are 69 relations induced on ordered pairs of edges. These are labelled and displayed in Figure 3.

	rows	cols	symbols	boxes
rows	$=, \ncong, \nsim$			$\smile, \not\smile$
cols		$=, \ncong, \nsim$		$\frown, \not\frown$
symbols			$=, \neq$	
boxes	$\smile, \not\smile$	$\frown, \not\frown$		$=, \smile, \frown, \not\asymp$

Table 1. Relations on vertices of G_{hw}

Proposition 3.1. *The relations $R_1, \ldots, R_{69}$ given in Figure 3 define a coherent configuration on on $E(G_{hw})$.*

Proof. The relations $=, \ncong, \nsim$ are orbits of $\mathcal{S}_h \wr \mathcal{S}_w$ (respectively $\mathcal{S}_w \wr \mathcal{S}_r$) acting on pairs of rows (columns). The direct product of second factors $\mathcal{S}_w \times \mathcal{S}_h$ induces relations $=, \smile$ $, \frown, \asymp$ on ordered pairs of boxes. The relations $\smile, \not\smile$ (respectively $\frown, \not\frown$) also define orbits on ordered pairs of rows (columns) with boxes. Finally, the relations $=, \neq$ define orbits of $\mathcal{S}_n$ acting on pairs of symbols. It follows that the set of relations $\{R_1, \ldots, R_{69}\}$ matches the coherent configuration from the action of Γ_{hw} on $E(G_{hw})$. ☐

The diagonal relation $\{(x, x) : x \in X\}$ is a union of $R_1, R_{16}, R_{32}, R_{62}$. For any relation R_i, its transpose $R_i^\top$ can be identified directly from Figure 3. With extensive case analysis, it would be technically possible to demonstrate formulas for the structure constants p_{ij}^k. However, to avoid presenting such details and to reduce errors, we implemented the following computer-assisted procedure:

(1) argue that p_{ij}^k belongs to $\mathbb{Z}[h, w]$, and is at most quadratic in each of h, w;
(2) compute all structure constants explicitly for the nine cases $2 \leqslant h, w \leqslant 4$;
(3) interpolate this data to arrive at symbolic expressions for p_{ij}^k.

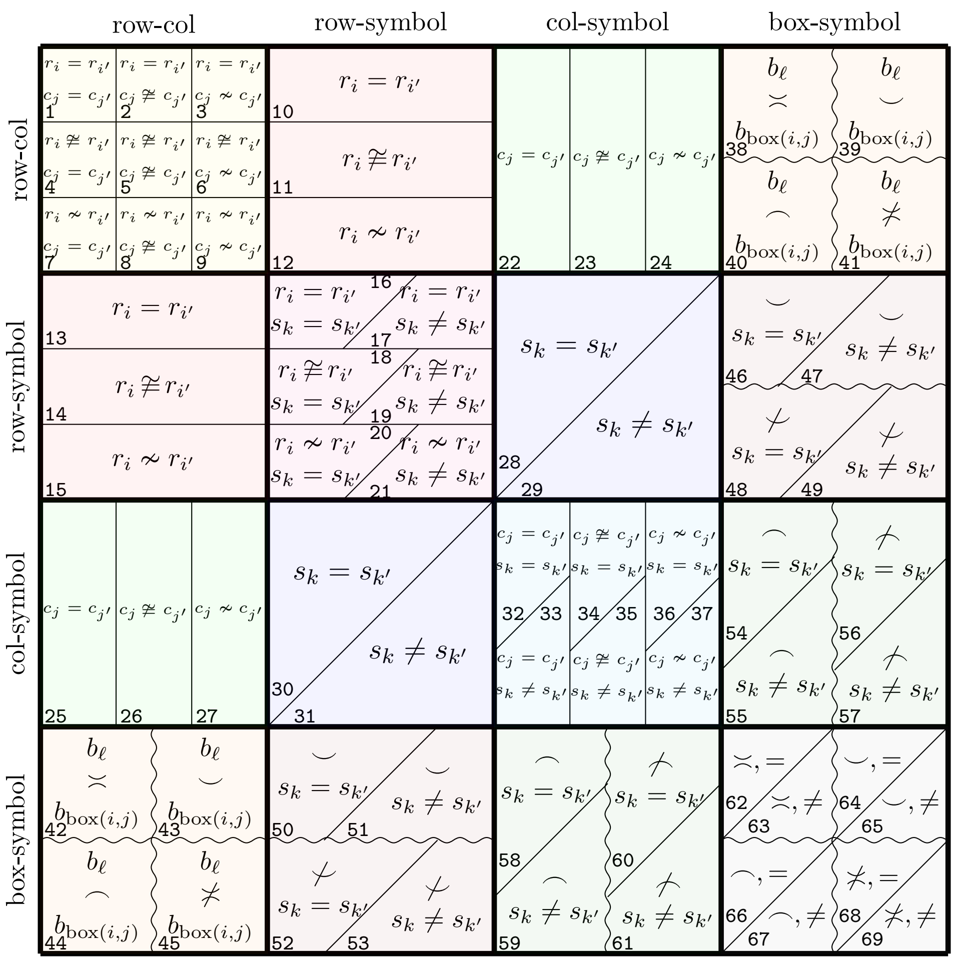

Figure 3. Relations on edges of G_{hw}

We discuss these points in a little more detail.

Proposition 3.2. *In the coherent configuration defined by Γ_{hw}, each structure constant p_{ij}^k is a polynomial of degree at most 2 in each of h and w.*

Proof. Fix two edges $x, z \in E(G_{hw})$ with $(x,z) \in R_k$. The quantity p_{ij}^k counts the edges $y \in E(G_{hw})$ with $(x,y) \in R_i$ and $(y,z) \in R_j$. This quantity is zero unless the indices i and j simultaneously allow one of the four types of edges for y. Given indices i and j which admit a choice of y, we must choose either a row-column pair, a row-symbol pair, a column-symbol pair, or a box-symbol pair. The two components of each pair can be selected separately, leading to a product of choices for the two components. Each factor is easily seen to have degree at most one in both h and w. The number of choices for a row is an element of $\{0, 1, h-2, h-1, h, n-2h, n-h, n\}$. Similarly, the number of choices for a column is an element of $\{0, 1, w-2, w-1, w, n-2w, n-w, n\}$. The number of choices for a symbol is an element of $\{0, 1, n-2, n-1, n\}$. Finally, the number of choices for a box is a product of an element of $\{0, 1, h-2, h-1, h\}$ with an element of $\{0, 1, w-2, w-1, w\}$. $\square$

The choice of nine cases $2 \leqslant h, w \leqslant 4$ suffices because of the degree bound in Proposition 3.2. The computation was carried out on computer by explicitly listing all $4h^2w^2$ edges and counting incidences. This took several minutes for the larger cases. From this, the interpolation in (3) can be performed using a 9×9 Vandermonde matrix based on the terms $1, h, w, h^2, hw, w^2, h^2w, hw^2, h^2w^2$, where $(h, w) \in \{2, 3, 4\}^2$.

For each relation index $i = 1, \ldots, 69$, we consider its corresponding adjacency matrix A_i. Let $\mathfrak{A}$ denote the $\mathbb{R}$-vector space spanned by the A_i. Since the relations form a coherent configuration, $\mathfrak{A}$ is closed under matrix multiplication, and hence forms an algebra.

If we view each relation as a graph, then $\{R_i : i = 1, \ldots, 69\}$ is a decomposition of the line graph of G_{hw} into regular graphs. The degrees of these graphs are the nonzero row sums of the corresponding adjacency matrices. We give the degrees d_i for each of the relations in Table 2. These are arranged into columns according to the four edge types: row-column, row-symbol, column-symbol, and symbol-box. Consider, for example, the degree d_{47}. Given a row-symbol edge, say r_1s_1, this degree counts the symbol-box edges s_kb_ℓ with $s_k \neq s_1$ and b_ℓ in the same row bundle as r_1. There are $n-1$ choices for s_k and h choices for b_ℓ, since every row is incident with exactly $n/w = h$ boxes. So $d_{47} = (n-1)h$.

i	d_i	i	d_i	i	d_i	i	d_i
1	1	13	n	25	n	42	n
2	$w-1$	14	$n(h-1)$	26	$n(w-1)$	43	$n(h-1)$
3	$(h-1)w$	15	$nh(w-1)$	27	$n(h-1)w$	44	$n(w-1)$
4	$h-1$	16	1	30	n	45	$n(h-1)(w-1)$
5	$(h-1)(w-1)$	17	$n-1$	31	$n(n-1)$	50	h
6	$(h-1)^2w$	18	$h-1$	32	1	51	$(n-1)h$
7	$h(w-1)$	19	$(n-1)(h-1)$	33	$n-1$	52	$h(w-1)$
8	$h(w-1)^2$	20	$h(w-1)$	34	$w-1$	53	$(n-1)h(w-1)$
9	$n(h-1)(w-1)$	21	$(n-1)h(w-1)$	35	$(n-1)(w-1)$	58	w
10	n	28	n	36	$(h-1)w$	59	$(n-1)w$
11	$n(h-1)$	29	$n(n-1)$	37	$(n-1)(h-1)w$	60	$(h-1)w$
12	$nh(w-1)$	46	h	54	w	61	$(n-1)(h-1)w$
22	n	47	$(n-1)h$	55	$(n-1)w$	62	1
23	$n(w-1)$	48	$h(w-1)$	56	$(h-1)w$	63	$n-1$
24	$n(h-1)w$	49	$(n-1)h(w-1)$	57	$(n-1)(h-1)w$	64	$h-1$
38	n					65	$(n-1)(h-1)$
39	$n(h-1)$					66	$w-1$
40	$n(w-1)$					67	$(n-1)(w-1)$
41	$n(h-1)(w-1)$					68	$(h-1)(w-1)$
						69	$(n-1)(h-1)(w-1)$

TABLE 2. Relation degrees d_i; alternatively the nonzero row sums of A_i

3.2. THE COEFFICIENT MATRIX M. Recall that W denotes the $\{0, 1\}$ inclusion matrix of edges versus tiles in G_{hw}, and that $M = WW^\top$. We computed the rank and a basis for the kernel of M in Section 2. A key next observation is that M belongs to our adjacency algebra.

PROPOSITION 3.3. *The matrix $M = WW^\top$ lies in $\mathfrak{A}$, with*

$$M = hw(A_1 + A_{16} + A_{32} + A_{62}) + w(A_{46} + A_{50}) + h(A_{54} + A_{58}) + A_{10} + A_{13} + A_{22} + A_{25} + A_{28} + A_{30} + A_{38} + A_{42}. \tag{5}$$

Proof. Given two edges e, f in G_{hw}, the entry $M(e,f)$ equals the number of tiles containing both e and f. Since a tile contains exactly one row, column, box, and symbol, this number is zero whenever $e \cup f$ contains two distinct vertices of the same type. Moreover, since the box in a tile must correspond with the row-column pair, $M(e,f)$ is zero if $e \cup f \supset \{r_i, b_\ell\}$ or $\{c_j, b_\ell\}$ with, respectively $r_i \not\smile b_\ell$ or $c_j \not\frown b_\ell$. It suffices to consider those remaining cases when there exist tiles extending $e \cup f$.

If $e = f$, we claim that there are n such tiles, regardless of the type of edge. For $e = \{r_i, c_j\}$, any of the n symbols extend e to a tile (and there is a unique box involved). For $e = \{r_i, s_k\}$, there are any of n columns (with one corresponding box for each) extending e. This is similar when we exchange the roles of rows and columns. Finally, for $e = \{b_\ell, s_k\}$, any of the $hw = n$ cells (i,j) with $\mathrm{box}(i,j) = \ell$ extend e to a tile. The identity relation in our setup decomposes into the identity on the four types of edges; in terms of matrices,

$$I = A_1 + A_{16} + A_{32} + A_{62}.$$

We have shown that the diagonal entries of M, and hence the coefficient for each of these four adjacency matrices, equals n.

Next, consider $e = \{r_i, s_k\}$ and $f = \{b_\ell, s_k\}$. In the event that $r_i \smile b_\ell$, we obtain w possible columns c_j such that $\mathrm{box}(i,j) = \ell$, and $e \cup f \cup \{c_j\}$ defines a valid tile. The two relations corresponding to this choice of e and f (transposes of each other) have indices 46 and 50 in our labeling. If instead we take $e = \{c_j, s_k\}$ for $c_j \frown b_\ell$, there are likewise exactly h extensions to a tile via some row r_i. This choice corresponds to relations numbered 54 and 58.

Finally, in each of the following possibilities for $\{e, f\}$, there is a unique tile $r_i c_j s_k b_\ell$ extending $e \cup f$, where $\ell = \mathrm{box}(i,j)$:

$$\{r_i c_j, r_i s_k\}, \{r_i c_j, c_j s_k\}, \{r_i s_k, c_j s_k\}, \{r_i c_j, b_\ell s_k\}.$$

The corresponding relation labels are $10, 13, 22, 25, 28, 30, 38, 42$. □

The structure of entries of M is depicted in Figure 4. On the left, we present M as a block matrix, whose block partition corresponds to the four edge types. Each block is an $n^2 \times n^2$ matrix which can be factored as a Kronecker product. It is convenient to slightly abuse the Kronecker product in the following way. In forming $A \otimes B$, each factor will be indexed by one of our four Sudoku objects: rows, columns, symbols, and boxes. The product is then indexed by corresponding pairs of elements. For instance, the $(1,2)$-block of M can be represented as $I_r \otimes J_{cs}$, where I_r is the identity matrix indexed by $\{r_1, \dots, r_n\}$ and J_{cs} is the all-ones matrix whose rows are indexed by $\{c_1, \dots, c_n\}$ and columns are indexed by $\{s_1, \dots, s_n\}$. The latter can be factored as $\mathbf{j}_c \otimes \mathbf{j}_s^\top$, where $\mathbf{j}$ is an $n \times 1$ all-ones vector and the subscript indicates the indexing set. The $(1,2)$-block of M has rows indexed by edges $e = r_i c_j$, columns indexed by edges $f = r_{i'} s_{k'}$, and the (e,f)-entry is 1 if and only if $i = i'$. This exactly recovers the condition for e and f sharing a common tile. Other blocks of M are similar. We use H_{rb} to denote the zero-one matrix indexed by rows versus boxes in which $H_{rb}(r_i, b_\ell) = 1$ if and only if $r_i \smile b_\ell$. We use H_{cb} analogously for columns. Finally, H_{rcb} is $n^2 \times n$, indexed by row-column edges versus boxes, and $H_{rcb}(r_i c_j, b_\ell) = 1$ if and only if $\mathrm{box}(i,j) = \ell$.

On the right of Figure 4, we display the locations of nonzero entries as a graphic, illustrated in the case $h = 2$, $w = 3$. The diagonal has entries $n = hw$. Blocks $(2,4)$ and $(4,2)$ correspond to A_{46} and A_{50}, with coefficient w. Blocks $(3,4)$ and $(4,3)$ correspond to A_{54} and A_{58}, with coefficient h. The other blocks correspond to the remaining terms with coefficient 1.

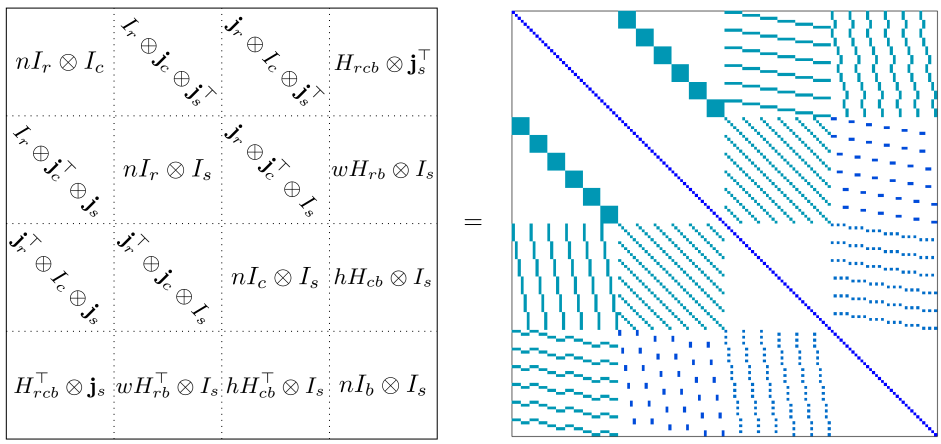

FIGURE 4. Illustration of the block matrix structure of M

4. Spectral decomposition of M

4.1. Eigenvalues and eigenvectors. Since $M = WW^\top$, we know it is symmetric and hence has real eigenvalues. We also have $\operatorname{rank}(M) = \operatorname{rank}(W)$, so from Section 2.3, we know that zero is an eigenvalue of M with multiplicity $3n+(h+w)(n-1)$. Moreover, M has constant row sums equal to $4n$, since every edge belongs to n tiles, and every tile has four edges. This gives an eigenvalue $4n$ corresponding to the one-dimensional eigenspace of constant vectors.

In this section, we compute all other eigenvalues and corresponding eigenvectors for M. By (5), we know that $M \in \mathfrak{A}$. Later, a generalized inverse for M is expressed with a list of coefficients in $\mathfrak{A}$. For the discovery of these coefficients, it is helpful to have a good understanding of the spectral decomposition of M. This is summarized here, with more details and verifications for eigenvectors appearing in the remainder of this subsection.

Proposition 4.1. *The eigenvalues of M are $\theta_j = jn$, $j = 0, 1, \ldots, 4$. Each eigenspace has a basis of eigenvectors consisting of vectors with entries in $\{0, \pm 1\}$.*

We have discussed $\theta_0 = 0$ and $\theta_4 = 4n$ earlier, so we turn our focus to $\theta_1, \theta_2, \theta_3$. Below, we describe different varieties of eigenvectors (A), (B), etc., for each of these θ_j. A basis for each eigenspace can be found by taking a union of linearly independent vectors over the different varieties. Making a selection of linearly independent vectors of the indicated size within each variety can be done using relations as in Section 2.3. More details can be found in [13].

We give an informal description and brief verification for each eigenvector. Checking that $M\mathbf{v} = \theta_j \mathbf{v}$ can be done as follows. Take each edge $f \in E(G_{hw})$ and extend to a tile $t \supset f$ in all possible ways. Then, sum the values of $\mathbf{v}$ on the four edges of t, and check that this total equals $\theta_j \mathbf{v}(f)$. This often equals zero, either from canceling signs or when the support of $\mathbf{v}$ is disjoint from the relevant tiles t. Figures 5, 6 and 7 give diagrams illustrating the eigenvector varieties in the case $(h, w) = (2, 3)$. In these diagrams, the four sections correspond to the four edge types: row-column (top left), row-symbol (top right), symbol-column (bottom left), and box-symbol (bottom right). Symbols $+$ and $-$ denote vector entries 1 and -1, respectively, and blanks represent 0 in the corresponding positions.

- $\theta_1 = n$; eigenspace dimension $4n^2 - (2n-3)(h+w) - 5n - 1$

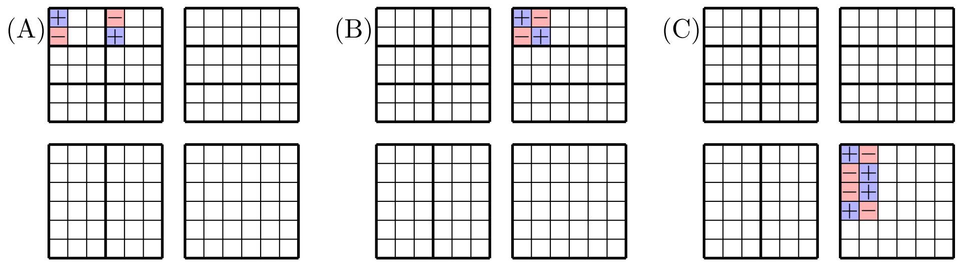

Figure 5. Eigenvectors for $\theta_1 = n$

(A) Opposite signs on two distinct rows and two distinct columns, at least one pair of which is in a common bundle. There are $(n-1)^2 - (h-1)(w-1)$ linearly independent vectors of this kind.

If r_{i_0}, r_{i_1} are the rows and c_{j_0}, c_{j_1} are the columns, then the entries are given explicitly by $\mathbf{v}(r_i c_j) = (-1)^{\alpha+\beta}$ if $(i,j) = (i_\alpha, j_\beta)$, and otherwise $\mathbf{v}(e) = 0$. For an edge $f = r_i c_j$, we have n tiles extending f corresponding to a choice of symbol s_k. Each such tile contains at most one non-vanishing edge, namely that corresponding to f. So $M\mathbf{v}(f) = n\mathbf{v}(f)$. For f of type row-symbol, column-symbol or box-symbol, we have $M\mathbf{v}(f) = 0$, either from cancellation or disjoint supports. Importantly, having either $r_{i_0} \sim r_{i_1}$ or $c_{j_0} \sim c_{j_1}$ ensures cancellation within each box.

(B) Opposite signs on two distinct rows (or columns) in the same bundle and on two distinct symbols. There are $(n-1)(h(w-1) + w(h-1))$ linearly independent vectors of this kind.

If r_{i_0}, r_{i_1} are the rows and s_{k_0}, s_{k_1} are the symbols, then the entries are given explicitly by $\mathbf{v}(r_i s_k) = (-1)^{\alpha+\gamma}$ if $(i,k) = (i_\alpha, k_\gamma)$, and otherwise $\mathbf{v}(e) = 0$. Verification that $M\mathbf{v} = n\mathbf{v}$ is similar to (A).

(C) Alternating signs on a rectangle of boxes and opposite signs on two distinct symbols. There are $(n-1)(h-1)(w-1)$ linearly independent vectors of this kind.

Suppose $\ell_{\alpha\beta}$ are the four box indices, where $\alpha, \beta \in \{0,1\}$ tell us the chosen row/column bundles, respectively. As in (B), let k_γ be the chosen symbol indices, $\gamma \in \{0,1\}$. The entries of the eigenvector are given by $\mathbf{v}(s_k b_\ell) = (-1)^{\alpha+\beta+\gamma}$ when $(k,\ell) = (k_\gamma, \ell_{\alpha\beta})$, and otherwise $\mathbf{v}(e) = 0$. Cancellation occurs if we sum over rows, columns, or symbols. For a symbol-box edge $f = s_k b_\ell$, the n tiles extending f correspond to a choice of entry in box ℓ. This picks up the value of $\mathbf{v}(f)$ with a multiplicity of n. So $M\mathbf{v} = n\mathbf{v}$.

• $\theta_2 = 2n$; eigenspace dimension $(2n-h-w)+(n-1)(h+w-2)+(h-1)(w-1) = (n-3)(h+w-1)+2n$.

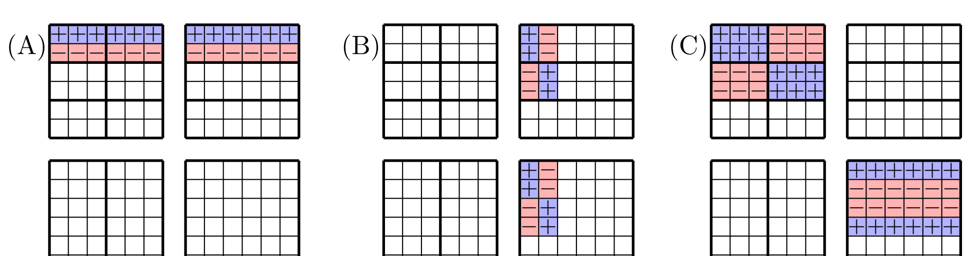

Figure 6. Eigenvectors for $\theta_2 = 2n$

(A) Opposite signs on two distinct rows in the same bundle; constant on all columns and symbols. A similar variety exists with rows and columns swapped. There are $h(w-1)+w(h-1) = 2n-h-w$ linearly independent vectors of this kind.

If r_{i_0}, r_{i_1} are the rows, then the eigenvector entries are $\mathbf{v}(r_i c_j) = \mathbf{v}(r_i s_k) = (-1)^\alpha$ when $i = i_\alpha$, and otherwise $\mathbf{v}(e) = 0$ for all other edges. An edge $f = r_i c_j$ or $r_i s_k$ has exactly n extensions to a tile, each of which has two edges of a common sign. So $M\mathbf{v}(f) = 2n\mathbf{v}(f)$ in those cases. It is easy to see that $M\mathbf{v}(f) = 0$ on all other edges due to cancellation on rows.

(B) Opposite signs on both rows and boxes of two distinct row bundles; opposite signs on symbols. A similar variety exists with rows and columns swapped. There are $(n-1)(h+w-2)$ linearly independent vectors of this kind.

If f is a row-column edge, the cancellation on symbols gives $M\mathbf{v}(f) = 0$. Likewise, if f is a column-symbol edge, the cancellation on rows gives $M\mathbf{v}(f) = 0$. For f of either of the other two edge types, there are n extensions to a tile, and again the nonzero edges (if any) agree in sign.

(C) Alternating signs on a rectangle of boxes; constant on all symbols and on entries within each box. There are $(h-1)(w-1) = n-h-w+1$ linearly independent vectors of this kind.

For row-symbol or column-symbol edges, the extension to a tile leads to cancellation. For a row-column edge f, there n extensions to a tile by selecting a symbol, and each has two matching edges from the entry and box. So $M\mathbf{v}(f) = 2n\mathbf{v}(f)$. Similarly, for a box-symbol edge f, we have $M\mathbf{v}(f) = 2n\mathbf{v}(f)$.

- $\theta_3 = 3n$; eigenspace dimension $n+h+w-3$

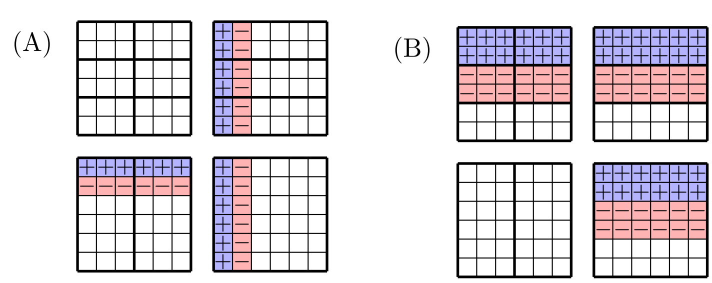

FIGURE 7. Eigenvectors for $\theta_3 = 3n$

(A) Opposite signs on two distinct symbols; constant on all rows, columns, and boxes. There are $n-1$ linearly independent vectors of this kind.

If k_0, k_1 are the two symbols, then the eigenvector entries are $\mathbf{v}(r_i s_k) = \mathbf{v}(c_j s_k) = \mathbf{v}(b_\ell s_k) = (-1)^\gamma$ when $k = k_\gamma$, and otherwise $\mathbf{v}(e) = 0$ for all other edges. If f is any edge involving a symbol s_{k_γ}, the n tiles extending f each have (if any) three nonzero edges of matching sign. So $M\mathbf{v}(f) = 3n\mathbf{v}(f)$. In other cases, it is easy to see that $M\mathbf{v}(f) = 0$ by cancellation.

(B) Opposite signs on both rows and boxes of two distinct row bundles; constant on all columns and symbols. A similar variety exists with rows and columns swapped. There are $(h-1)+(w-1)$ linearly independent vectors of this kind. The verification here is similar to (A), except that row bundles take the role of symbols.

Although it is somewhat cumbersome, one can use Kronecker product to express all of the above eigenvectors, including the kernel vectors from Section 2.3. They can be verified using Figure 4 and block matrix multiplication.

We next consider in more detail the projectors onto the eigenspaces of M.

4.2. Projectors and a generalized inverse for M. Since M is symmetric, the projectors E_j onto eigenspaces for θ_j are pairwise orthogonal idempotents summing to I. Moreover, we have $E_j \in \mathfrak{A}$ for each j as a general fact of coherent configurations; see for instance [9].

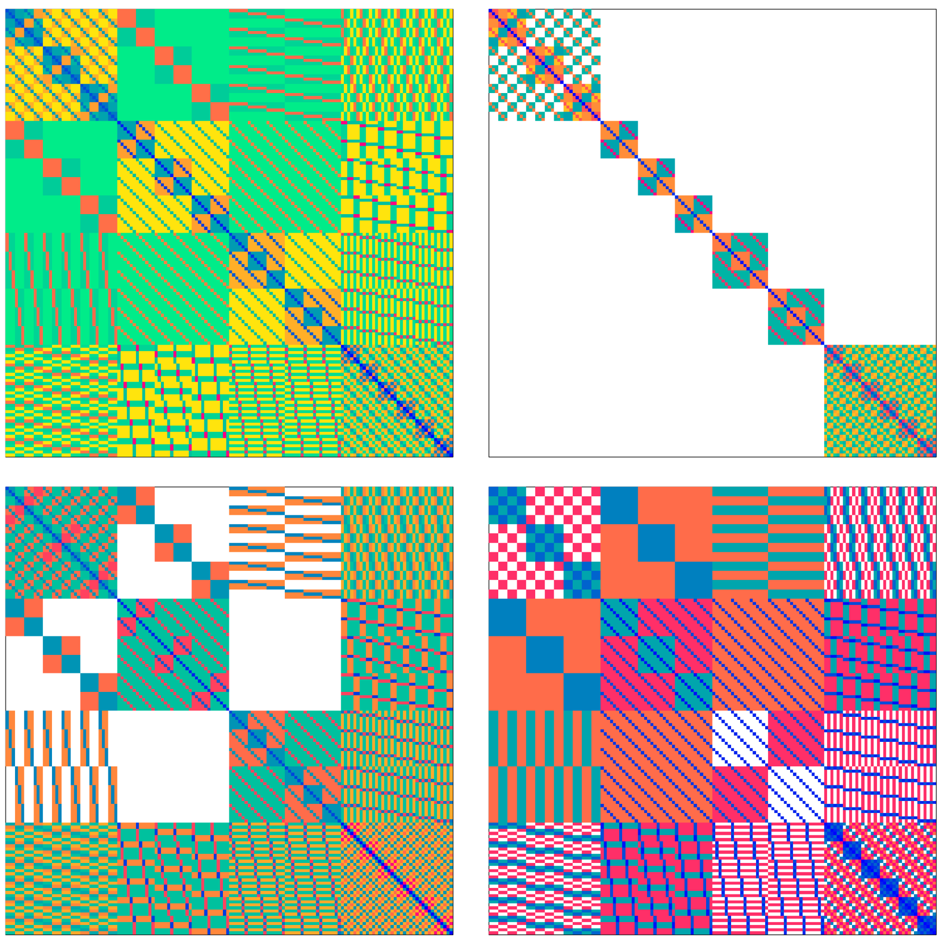

Figure 8. Structure of entries of E_0, E_1 (top), E_2, E_3 (bottom) for $(h, w) = (2, 3)$

The projectors can be computed as $E_i = V_i(V_i^\top V_i)^{-1}V_i^\top$, where V_i is a matrix whose columns are a basis of eigenvectors for θ_i. As a special case, since V_4 is the all-ones vector, we have $E_4 = \frac{1}{4n^2}J$. The structure of entries for each of the other projectors is shown in Figure 8. Intensity of blue/red correspond respectively to extreme positive/negative entries, while shades of green/yellow correspond to positive/negative entries which are smaller in magnitude.

Knowing the eigenvalues and eigenspace projectors for M can be used to compute a generalized inverse M^+ satisfying $MM^+M = M$. We explain this computation in the rest of this section.

The spectral decomposition of M is given by $M = nE_1 + 2nE_2 + 3nE_3 + 4nE_4$. In what follows, E_0 will also be denoted K, since it projects onto the kernel of M. Although M itself is not invertible, if we take $\eta \neq 0$, say $\eta = n/x$, we can invert the additive shift $A = M + \eta K$ as

$$A^{-1} = \frac{1}{n}\left(xK + \sum_{j=1}^{4} \frac{1}{j} E_j\right). \tag{6}$$

This formula results from the E_i being orthogonal idempotents with $E_0 + E_1 + \cdots + E_4 = I$. Later on, to solve our linear system for fractional Sudoku, we make use of a generalized inverse M^+ of the form in (6). It turns out that $x = 3/2$, or $\eta = 2n/3$, is a nice choice. A discussion of this choice is given in the next subsection.

With some computer-assisted algebra, we found coefficients to express A^{-1} in the basis $\{A_i : i = 1, \ldots, 69\}$ for the adjacency algebra $\mathfrak{A}$. These are expressed in Table 3. For convenience, we have cleared a denominator of $9n^3$ and then applied an additive shift of 5/16. Using our Sage [16] worksheet at `https://github.com/pbd345/sudoku`, the interested reader can compute various symbolic products in $\mathfrak{A}$, including a verification that Table 3 does indeed give an inverse of A.

relations	coeffs	relations	coeffs
1	$9n^2 + h + w$	32	$9n^2 + n + h$
$2, 4, 5$	$h + w$	34	$n + h$
$3, 6, 17, 19, 39, 43, 47, 51$	w	$38, 42$	$-9n/2 + h + w$
$7, 8, 33, 35, 40, 44, 55, 59$	h	$46, 50$	$-9nw/2 + n + w$
$10, 13$	$-9n/2 + w + 1$	$54, 58$	$-9nh/2 + n + h$
$11, 14$	$w + 1$	62	$9n^2 + n + h + w - 1$
$12, 15, 24, 27, 29, 31$	1	63	$h + w - 1$
16	$9n^2 + n + w$	64	$n + w - 1$
18	$n + w$	65	$w - 1$
$20, 36, 48, 52, 56, 60$	n	66	$n + h - 1$
$22, 25$	$-9n/2 + h + 1$	67	$h - 1$
$23, 26$	$h + 1$	68	$n - 1$
$28, 30$	$-7n/2 + 1$	69	-1

Table 3. Coefficients of $9n^3A^{-1} + \frac{5}{16}J$

4.3. Norm bounds.

We work with the ∞-norm of vectors $\|\mathbf{x}\|_\infty = \max\{|x_1|, \ldots, |x_n|\}$ and the induced norm on matrices

$$\|A\|_\infty = \max_i \sum_j |A_{ij}|.$$

It is straightforward to obtain a bound on the ∞-norm of (6) using the values in Tables 2 and 3. The triangle inequality gives a crude bound of order $O(n^{-1})$, but we can get an exact value with the help of a computer. First, we store the coefficients of the projectors relative to our coherent configuration basis and make a note of their signs. For each of the four sections corresponding to the edge types, we sum the absolute values of projector coefficients times the section row sums. When we combine these as in (6), the result is a list of three piecewise linear functions (one duplicate occurs for two sections), each multiplied by n^{-1}. These functions are

$$f_1(x) = 3|\tfrac{x}{2} - \tfrac{3}{4}| + 4|\tfrac{x}{6} - \tfrac{5}{36}| + 2|\tfrac{x}{12} - \tfrac{7}{144}| + 2|\tfrac{x}{12} - \tfrac{13}{144}| + 2|\tfrac{x}{3} - \tfrac{11}{18}| + 3|\tfrac{x}{2} - \tfrac{1}{4}| + 1,$$
$$f_2(x) = 2|\tfrac{x}{2} - \tfrac{3}{4}| + 4|\tfrac{x}{6} - \tfrac{5}{36}| + 2|\tfrac{x}{12} - \tfrac{7}{144}| + 2|\tfrac{x}{12} - \tfrac{13}{144}| + |\tfrac{x}{3} - \tfrac{11}{18}| + |\tfrac{x}{3} - \tfrac{1}{9}| + 2|\tfrac{x}{2} - \tfrac{1}{4}| + 1,$$
$$f_3(x) = 3|\tfrac{x}{2} - \tfrac{3}{4}| + |\tfrac{x}{4} - \tfrac{25}{48}| + 6|\tfrac{x}{6} - \tfrac{5}{36}| + 3|\tfrac{x}{12} - \tfrac{13}{144}| + 3|\tfrac{x}{3} - \tfrac{11}{18}| + 3|\tfrac{x}{2} - \tfrac{1}{4}| + 1.$$

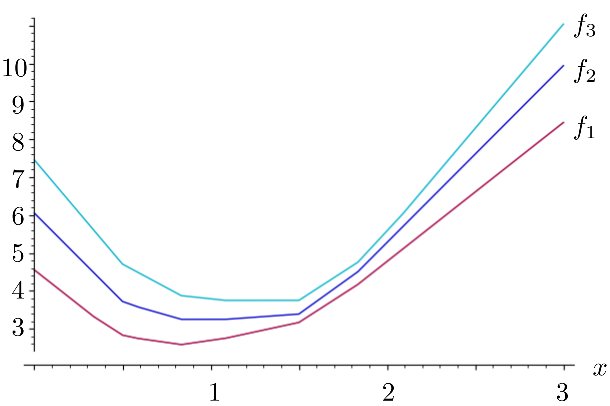

Figure 9. A^{-1} norm plots (f_1, f_2, f_3) as functions of $x := n/\eta$

Graphs for the functions $f_i(x)$ are shown in Figure 9. It turns out that $\max\{f_i(x) : i = 1,2,3\}$ is minimized at $x = 3/2$, yielding the dominant term $15/4n$, also an upper bound for all $h, w \geqslant 2$. The results of this computation are summarized in the lemma below.

Lemma 4.2. *Let* $A = M + \frac{2n}{3}K$. *Then*

$$\|A^{-1}\|_\infty = \frac{15}{4n} - \frac{7(h+w)}{8n^2} - \frac{4}{9n^2} + \frac{31(h+w)-21}{72n^3} < \frac{15}{4n}.$$

We also note the following bound on $K = E_0$.

Lemma 4.3. *We have* $\|K\|_\infty \leqslant \dfrac{11}{2} - \dfrac{17(h+w)}{6n} + O(n^{-1})$.

5. Perturbation

5.1. Changes to M resulting from pre-filled entries. Let S be a partial Sudoku of type (h, w), where $hw = n$. Recall that G_S is the graph obtained from G_{hw} by deleting the edges of tiles corresponding to pre-filled entries in S. Suppose S is ϵ-dense. We know that every edge in G_S is contained in at least $(1-3\epsilon)n$ tiles in G_S.

Let $M = WW^\top$ and $M_S = W_S W_S^\top$, as introduced in Section 2. To set up our perturbation argument, we are interested in quantifying the change in M resulting from pre-filling the entries of S. It makes no sense to subtract M_S from M directly, since these matrices have different sizes. However, we can use a convenient border.

Let $\widetilde{M}$ denote the $4n^2 \times 4n^2$ matrix, indexed by edges of G, whose entries are given by

$$\widetilde{M}(e,f) = \begin{cases} M_S(e,f) & \text{if } e, f \in E(G_S); \\ 0 & \text{if } e \in E(G_S) \text{ and } f \notin E(G_S); \\ M(e,f) & \text{if } e \notin E(G_S). \end{cases}$$

If we sort the rows and columns so that those indexed by $E(G_S)$ come first, then

$$\widetilde{M} = \left[\begin{array}{c|c} M_S & O \\ \hline \multicolumn{2}{c}{\text{as in } M} \end{array}\right]. \tag{7}$$

Put $\Delta M = M - \widetilde{M}$. We next estimate $\|\Delta M\|_\infty$ under our sparseness assumption.

For an edge $e \in E(G_S)$, let $U(e)$ denote the set of unavailable options

$$U(e) = \{t \in T(G_{hw}) : e \in t \text{ and } f \in t \text{ for some } f \in E(G_{hw}) \smallsetminus E(G_S)\}.$$

Put $u(e) = |U(e)|$ and $\mathbf{u} = (u(e) : e \in E(G_S))$. In more detail, if e is an edge of type row-column, say $e = \{r_i, c_j\}$, then $U(e)$ keeps track of those symbols k which are not able to be placed in cell (i, j) because k already appears in row i or column j or box $\text{box}(i, j)$. If e is an edge of type row-symbol, say $e = \{r_i, s_k\}$, then $U(e)$ keeps track of those columns j which are unavailable for symbol k in row i, either because cell (i, j) was pre-filled or k appears somewhere else in column j or box $\text{box}(i, j)$. Note that several columns might be eliminated as options if k appears in a box intersecting row i. Edges of type column-symbol behave in an analogous way. Finally, if e is an edge of type box-symbol, say $e = \{b_\ell, s_k\}$, then $U(e)$ keeps track of those cells (i, j) in box ℓ for which k is not allowed, either because (i, j) was already filled in S, or because k already appears in the row or column bundle for box ℓ.

Lemma 5.1. *We have $\mathbf{0} \leqslant \Delta M \mathbb{1} \leqslant 4\mathbf{u}$ entrywise. In particular, $\|\Delta M\|_\infty \leqslant 4\|\mathbf{u}\|_\infty$.*

Proof. Entry e of $(\Delta M)\mathbb{1}$ equals $\sum_f \Delta M(e, f)$. The summand is the number of unavailable tiles t with $e \in t$ and $f \in t$. Since each copy t contains four edges, this count is at most $4u(e)$. ☐

If S is ϵ-dense, then $\|\mathbf{u}\|_\infty \leqslant 3\epsilon n$. This gives the following bound on ΔM.

Lemma 5.2. *Suppose S is an ϵ-dense partial Sudoku. With ΔM constructed from S as above, we have $\|\Delta M\|_\infty \leqslant 12\epsilon n$.*

5.2. A guarantee on nonnegative solutions. The following can be distilled from [4, Section 3].

Lemma 5.3. *Let A be an $N \times N$ invertible matrix over the reals. Suppose $A - \Delta A$ is a perturbation. Then*

(1) *$A - \Delta A$ is invertible provided $\|A^{-1}\Delta A\|_\infty < 1$; and*
(2) *the solution $\mathbf{x}$ to $(A - \Delta A)\mathbf{x} = A\mathbb{1}$ is entrywise nonnegative provided $\|A^{-1}\Delta A\|_\infty \leqslant \frac{1}{2}$.*

Lemma 5.3 can be proved using the expansion $(A-\Delta A)^{-1} = \sum_{k=0}^{\infty}(A^{-1}\Delta A)^k A^{-1}$. More details on matrix norms and the convergence of this series can be found in Horn and Johnson's book [10].

We would like to apply Lemma 5.3 to the perturbation $M - \Delta M$, but we must take care to handle the nontrivial kernel. Of various possible approaches, one convenient thing to do is to place those columns of $K = E_0$ corresponding to non-edges of G_S in the perturbation. In more detail, let $A = M + \eta K$ and observe that $A\mathbb{1} = 4n\mathbb{1}$. That is, for this choice of A, the right side of the system in Lemma 5.3 is just a scalar multiple of the all-ones vector. Define $\Delta A = \Delta M + \eta K'$, where

$$K'(e, f) = \begin{cases} 0 & \text{if } f \in E(G_S); \\ K(e, f) & \text{otherwise.} \end{cases}$$

Note that

$$A^{-1}(\eta K') = \left(\frac{1}{\eta}K + \sum_{j=1}^{4}\frac{1}{jn}E_j\right)(\eta K') = K', \tag{8}$$

since the columns of K' are orthogonal to each of the other eigenspaces.

Lemma 5.4. *Suppose S is ϵ-dense. Then, for large h and w, $\|K'\|_\infty \leqslant (\epsilon + o(1))\|K\|_\infty$.*

Proof. Write $K = \sum_{i=1}^{m} c_i A_i$. Fix $e \in E(G_{hw})$. Then we have

$$\sum_{f \in E(G_{hw})} |K(e, f)| = \sum_i |c_i| d_i(e),$$

where $d_i(e)$ is the number of edges f with $(e,f) \in R_i$. Recall that $d_i(e)$ is zero unless e is of an edge type corresponding to the first coordinate of R_i, and we may assume a canonical choice r_1c_1, r_1s_1, c_1s_1, or s_1b_1 for e.

Let $\overline{G_S}$ denote the complement of G_S in G_{hw}. Then we have

$$\sum_{f\in E(G_{hw})} |K'(e,f)| = \sum_{f \in E(\overline{G_S})} |K(e,f)| = \sum_{i=1}^{m} |c_i| d_i'(e), \tag{9}$$

where $d_i'(e)$ is the number of edges $f \in E(\overline{G_S})$ with $(e,f) \in R_i$. With the exception of $i \in I := \{1,2,4,16,32,62\}$, each relation R_i has an associated feature which, owing to our ϵ-density assumption, limits the number of missing edges f in G_S with $(e,f) \in R_i$. These features are indicated in Table 4, along with bounds on leading terms of $|c_i|d_i'(e)$. A legend and upper bound on corresponding $d_i'(e)$ are given in Table 5. Terms with $i \in I$ are of lower order. Otherwise, when we compute the sum (9), we obtain the same leading terms as in the computation of $\|K\|_\infty$, each times ϵ. The edge type with largest total coefficient of ϵ is the box-symbol type, or column 4 in Table 4. This results in

$$\|K'\|_\infty \leqslant \left(\frac{11}{2} + O(h^{-1}+w^{-1})\right)\epsilon + \frac{h+w}{2n} + O(h^{-2}+w^{-2}+n^{-1}). \qquad \square$$

i	leading term	sparse feature	i	leading term	sparse feature	i	leading term	sparse feature	i	leading term	sparse feature
1	$3/2n$	-	13	$\epsilon/2$	r	25	$\epsilon/2$	c	42	$\epsilon/2$	b
2	$1/h$	-	14	$\epsilon/6$	rb	26	$\epsilon/6$	cb	43	$\epsilon/6$	rb
3	$\epsilon/2$	r	15	$\epsilon/12$	all	27	$\epsilon/12$	all	44	$\epsilon/6$	cb
4	$1/w$	-	16	$1/2h$	-	30	$\epsilon/3$	s	45	$\epsilon/12$	all
5	$\epsilon/2$	b	17	$\epsilon/2$	r	31	$\epsilon/12$	all	50	$\epsilon/2$	srb
6	$\epsilon/3$	rb	18	$\epsilon/2$	srb	32	$1/2w$	-	51	$\epsilon/6$	rb
7	$\epsilon/2$	c	19	$\epsilon/3$	rb	33	$\epsilon/2$	c	52	$\epsilon/6$	s
8	$\epsilon/3$	cb	20	$\epsilon/6$	s	34	$\epsilon/2$	scb	53	$\epsilon/12$	all
9	$\epsilon/12$	all	21	$\epsilon/12$	all	35	$\epsilon/3$	cb	58	$\epsilon/2$	scb
10	$\epsilon/2$	r	28	$\epsilon/3$	s	36	$\epsilon/6$	s	59	$\epsilon/6$	cb
11	$\epsilon/6$	rb	29	$\epsilon/12$	all	37	$\epsilon/12$	all	60	$\epsilon/6$	s
12	$\epsilon/12$	all	46	$\epsilon/2$	srb	54	$\epsilon/2$	scb	61	$\epsilon/12$	all
22	$\epsilon/2$	c	47	$\epsilon/6$	rb	55	$\epsilon/6$	cb	62	$(h+w)/2n$	-
23	$\epsilon/6$	cb	48	$\epsilon/6$	s	56	$\epsilon/6$	s	63	$\epsilon/2$	b
24	$\epsilon/12$	all	49	$\epsilon/12$	all	57	$\epsilon/12$	all	64	$\epsilon/2$	srb
38	$\epsilon/2$	b							65	$\epsilon/3$	rb
39	$\epsilon/6$	rb							66	$\epsilon/2$	scb
40	$\epsilon/6$	cb							67	$\epsilon/3$	cb
41	$\epsilon/12$	all							68	$\epsilon/3$	s
									69	$\epsilon/4$	all

Table 4. Terms contributing to $\|K\|_\infty$

	sparse feature	bound		sparse feature	bound
r	cells filled in a row	ϵn	b	cells filled in a box	ϵn
c	cells filled in a column	ϵn	s	occurrences of a symbol	ϵn
rb	cells filled in a row bundle	ϵnh	srb	times a symbol is in a row bundle	ϵh
cb	cells filled in a column bundle	ϵnw	scb	times a symbol is in a column bundle	ϵw
all	cells filled in all of S	ϵn^2			

Table 5. Legend for Table 4 and ϵ-density bounds

Putting together Lemmas 4.2, 4.3, 5.2 and 5.4, we obtain a bound on $A^{-1}\Delta A$.

Proposition 5.5. *Suppose S is an ϵ-dense partial Sudoku of type (h, w) where h, w are large. Then $\|A^{-1}\Delta A\|_\infty < 101\epsilon/2 + o(1)$.*

Proof. From (8), submultiplicativity and the triangle inequality,

$$\|A^{-1}\Delta A\|_\infty \leqslant \|A^{-1}\|_\infty \|\Delta M\|_\infty + \|K'\|_\infty < \tfrac{15}{4n} \times 12\epsilon n + \tfrac{11}{2}\epsilon + o(1) = \tfrac{101}{2}\epsilon + o(1). \quad \square$$

5.3. Proof of the main result. We are now ready to prove our result on partial Sudoku completion under the ϵ-dense assumption.

Proof of Theorem 1.1. Apply Lemma 5.3 to A and ΔA constructed as above. Under the assumption $\epsilon < 1/101$, Proposition 5.5 gives $\|A^{-1}\Delta A\|_\infty < 1/2$ for sufficiently large h, w. This implies an entrywise nonnegative solution to $(A - \Delta A)\mathbf{x} = \mathbb{1}$. Let $\mathbf{x}'$ denote the restriction of $\mathbf{x}$ to $E(G_S)$. Since $A - \Delta A$ is block lower-triangular with respect to the partition into edges and non-edges of G_S, it follows that $(M_S + \eta K[S])\mathbf{x}' = \mathbb{1}$. We note that M_S and $K[S]$ are symmetric and satisfy the conditions in Proposition 2.4. Therefore, Lemma 2.5 implies $M_S\mathbf{x}' = \mathbb{1}$. This, in turn, implies a nonnegative solution to the linear system for completing S via the coefficient matrix W_S. $\square$

It is worth a remark that the lower order terms in Lemmas 4.2 and 4.3 are actually negative. This means our hypothesis of large h and w is only really used to control the mild lower-order terms in K'. In general, our method is robust for small partial Sudoku, often succeeding in practice with densities much larger than 1/101. For instance, the completion shown in Figure 1 came from applying the above proof method.

6. Variations and concluding remarks

In the case of Sudoku with 'thin boxes', say with fixed w and height $h = n/w$ tending to infinity, the bundle condition for ϵ-density is seemingly stronger than necessary to ensure completion. That is, by adding symbols before completion, one could ensure that a symbol occurring in some column bundle (of fixed size w) occurs in *each* of the w boxes of that bundle. Then, one only needs to allocate remaining symbols that occur in a small fraction of the rows, columns, and boxes. This setting of thin boxes is discussed in more detail in [13].

Suppose we generalize our setting so that each Sudoku box/cage is an arbitrary polyomino of n cells. Most of our set-up stays the same, except that the n-to-1 function $\text{box}(i, j)$ mapping cells to boxes changes, say to $\text{box}'(i, j)$. If this change is sufficiently small, we can reasonably expect the same perturbation methods to give a fractional completion guarantee for sparse partial Sudoku of this generalized type.

We describe in a little more detail a setting in which this could work. Define a polyomino Sudoku as above to have *α-approximate* type (h, w) if, for each box ℓ, the symmetric difference between $\text{box}^{-1}(\ell)$ and $(\text{box}')^{-1}(\ell)$ can be covered by αh rows and αw columns. The 0-approximate case coincides with our standard setting of rectangular boxes. Let M' denote the $4n^2 \times 4n^2$ matrix for the empty Sudoku with polyomino boxes defined by box', and let M be our usual matrix for the case of $h \times w$ boxes. It is not hard to see that, for an α-approximate type (h, w),

$$\|M - M'\|_\infty \leqslant 2(\alpha h)w + 2(\alpha w)h = 4\alpha n. \tag{10}$$

Applying the triangle inequality to (10) and our earlier estimates, a completion result for the α-approximate setting would follow readily.

It is possibly of interest to consider properties of the matrix M' for specific box arrangements. A 'Pentadoku' is a 5×5 Sudoku-like puzzle whose cages are pentomino shapes. Each cage (in addition to each line) must contain the numbers from 1 to 5 exactly once. Figure 10 shows an example of a completed Pentadoku puzzle.

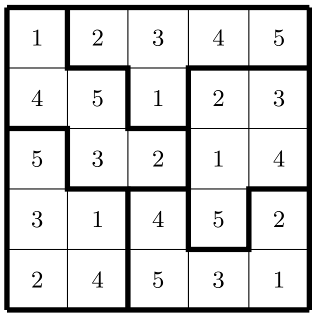

Figure 10. A solved Pentadoku puzzle

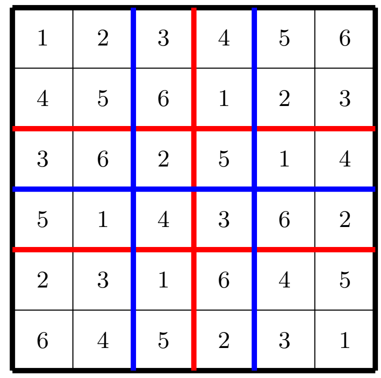

Figure 11. Puzzle with simultaneous box conditions

A generalization which we have not considered could allow two or more simultaneous box patterns. This is natural because the row and column conditions in a latin square can already be viewed as degenerate box conditions. An example for $n = 6$ with both 2×3 and 3×2 boxes is given in Figure 11. Using a generalized notion of tile and a suitably enlarged linear system, similar methods as in this paper could apply, at least in principle.

Methods of algebraic graph theory have been applied to Sudoku before, but in a slightly different way. A *Sudoku graph* has n^2 vertices corresponding to cells, and two vertices are declared adjacent if they share the same row, column or box. Eigenvalues and eigenvectors of Sudoku graphs have been investigated in [1, 11, 15]. Although they have integral eigenvalues and Kronecker-structured $\{\pm 1, 0\}$-valued eigenvectors, as ours, we could see no way to use the Sudoku graph alone to build the linear system for completion. Still, it would be interesting to explore the Sudoku graph in the context of completing partial Sudoku.

As a last remark, our results are only about fractional completion. For partial latin squares, the iterative absorption methods of [2] are able to convert a sparseness guarantee for fractional completion into a guarantee for (exact) completion. We do not know whether these or other methods could work for the Sudoku setting.

Acknowledgements. The authors are grateful to Akihiro Munemasa, who suggested expressing our coherent configuration in terms of a group action.

References

[1] Mohammad Abudayah, Omar Alomari, and Torsten Sander, *Spectrum of free-form Sudoku graphs*, Open Math. **16** (2018), no. 1, 1445–1454, https://doi.org/10.1515/math-2018-0125.

[2] Ben Barber, Daniela Kühn, Allan Lo, Deryk Osthus, and Amelia Taylor, *Clique decompositions of multipartite graphs and completion of Latin squares*, J. Combin. Theory Ser. A **151** (2017), 146–201, https://doi.org/10.1016/j.jcta.2017.04.005.

[3] Padraic Bartlett, *Completions of ϵ-dense partial Latin squares*, J. Combin. Des. **21** (2013), no. 10, 447–463, https://doi.org/10.1002/jcd.21355.

[4] Flora C. Bowditch and Peter J. Dukes, *Fractional triangle decompositions of dense 3-partite graphs*, J. Comb. **10** (2019), no. 2, 255–282, https://doi.org/10.4310/JOC.2019.v10.n2.a5.

[5] Hang Chen and Curtis Cooper, *Solving Sudoku: structures and strategies*, Missouri J. Math. Sci. **29** (2017), no. 1, 12–18, https://doi.org/10.35834/mjms/1488423697.

[6] Amanda Chetwynd and Roland Häggkvist, *Completing partial $n \times n$ latin squares where each row, column and symbol is used at most cn times*, 1985, Reports, Dept. of Mathematics, University of Stockholm.

[7] David E. Daykin and Roland Häggkvist, *Completion of sparse partial Latin squares*, in Graph theory and combinatorics (Cambridge, 1983), Academic Press, London, 1984, pp. 127–132.

[8] T. Gustavsson, *Decompositions of large graphs and digraphs with high minimum degree*, Ph.D. thesis, Stockholm University, 1991.

[9] D. G. Higman, *Coherent configurations. I*, Rend. Sem. Mat. Univ. Padova **44** (1970), 1–25, `http://www.numdam.org/item?id=RSMUP_1970__44__1_0`.

[10] Roger A. Horn and Charles R. Johnson, *Matrix analysis*, Cambridge University Press, Cambridge, 1985, `https://doi.org/10.1017/CBO9780511810817`.

[11] Sho Kubota, Sho Suda, and Akane Urano, *Mutually orthogonal Sudoku Latin squares and their graphs*, Graphs Combin. **39** (2023), no. 6, article no. 122 (26 pages), `https://doi.org/10.1007/s00373-023-02718-9`.

[12] Richard Montgomery, *Fractional clique decompositions of dense partite graphs*, Combin. Probab. Comput. **26** (2017), no. 6, 911–943, `https://doi.org/10.1017/S0963548317000165`.

[13] K.I. Nimegeers, *Pseudoku: A sudoku adjacency algebra and a fractional completion threshold*, Master's thesis, University of Victoria, 2024.

[14] J. Scott Provan, *Sudoku: strategy versus structure*, Amer. Math. Monthly **116** (2009), no. 8, 702–707, `https://doi.org/10.4169/193009709X460822`.

[15] Torsten Sander, *Sudoku graphs are integral*, Electron. J. Combin. **16** (2009), no. 1, article no. 25 (7 pages), `https://doi.org/10.37236/263`.

[16] The Sage Developers, *Sagemath, the Sage Mathematics Software System (Version 9.3)*, 2021, `https://www.sagemath.org`.

Peter J. Dukes, University of Victoria, Dept. of Mathematics and Statistics, Victoria, BC, V8N 5L5 (Canada)
E-mail : `dukes@uvic.ca`

Kate I. Nimegeers, University of Victoria, Dept. of Mathematics and Statistics, Victoria, BC, V8N 5L5 (Canada)
E-mail : `nimegeek@uvic.ca`