Small defining sets in $n \times n$ Sudoku squares

†**Mohammad Mahdian**[1] **and** ***Ebadollah S. Mahmoodian**[2]

[1]Google Research, Mountain View, CA, USA
[2]Department of Mathematical Sciences, Sharif University of Technology, Tehran, I. R. Iran

*Presenting and Corresponding author: emahmood@sharif.edu

Abstract

Over the last decade, Sudoku, a combinatorial number-placement puzzle, has become a favorite pastimes of many all around the world. Recently it is shown that this concept has many mathematical and computational relations and applications. In this puzzle, the task is to complete a partially filled 9 by 9 square with numbers 1 through 9, subject to the constraint that each number must appear once in each row, each column, and each of the nine 3 by 3 blocks. Sudoku squares can be considered a subclass of the well-studied class of Latin squares. Actually a Sudoku square of order $n = k^2$ is a Latin square of order n such that every element in $[n] = \{1, \ldots, n\}$ appears exactly once in each block. A partial Sudoku square P is a defining set for a Sudoku square S if S is the unique Sudoku square that is an extension of P. A central problem is to determine the size of the smallest defining set for Sudoku squares of order n. For $n = 9$ (regular Sudoku) extensive computer search showed that this number is 17 (McGuire et al, 2014), but the asymptotics of this value is unknown. For Latin squares, this number is conjectured to be $\lfloor n^2/4 \rfloor$ (Mahmoodian 1995, Van Rees and Bates 1999). A construction based on back-circulant Latin squares shows that this number is at most $\lfloor n^2/4 \rfloor$, but the best proven lower bound is just slightly superlinear. Also, the $\lfloor n^2/4 \rfloor$ conjecture is proved if "defining set" is replaced by a more strict notion called "forcing set".

For Sudoku squares, we show that the same construction (with a permutation on the rows of the matrix) works, giving an upper bound of $\lfloor n^2/4 \rfloor$. We also show that the size of the smallest forcing set for Sudoku squares of order n is at least $\Theta(n^2)$. Our conjecture is that the size of the smallest defining set for Sudoku squares of order n is also $\Theta(n^2)$. Finally, we discuss open problems related to Sudoku squares, their defining sets, and the computational complexity of Sudoku completion.

Keywords: Computation, Sudoku, Latin squares, defining set, forcing set, extension.

Introduction

A Latin square of order n is an $n \times n$ matrix with entries from $[n] = \{1, \ldots, n\}$ such that every element in $[n]$ appears exactly once in each row and in each column.

A partial Latin square of order n is an $n \times n$ matrix with entries from $[n] \cup \{*\}$ such that every element in $[n]$ appears at most once in each row and in each column. A partial Latin square P_1 is an extension of a partial Latin square P_2 if for every $(i, j) \in [n]^2$, if $P_2(i, j) \neq *$, then $P_1(i, j) = P_2(i, j)$.

A partial Latin square P is a defining set for a Latin square L if L is the unique Latin square that is an extension of P. A critical set is a minimal defining set. A forcing set (also called a strong critical set) is a partial Latin square P such that there is a sequence $P = P_0, P_1, \ldots, P_\ell$ such that P_ℓ is a Latin square and for every r,

- P_r is a partial Latin square and an extension of P_{r-1},

- the difference between P_r and P_{r-1} is in precisely one entry, i.e., there is $(x, y) \in [n]^2$ such that $P_r(i, j) = P_{r-1}(i, j)$ for every $(i, j) \neq (x, y)$ and $P_{r-1}(x, y) = *$ and $P_r(x, y) \neq *$, and
- for every $z \in [n]$ and $z \neq P_r(x, y)$, the matrix obtained from P_r by setting $P_r(x, y)$ to z is not a partial Latin square.

In a Latin square of order $n = k^2$, the (i, j)'th block (for $i, j \in [k]$) is the set of entries with coordinates in $((i-1)k+x, (j-1)k+y)$ for $x, y \in [k]$. We say that (i, j) are the coordinates of this block. These blocks partitions the set of entries in the matrix into n blocks, each containing n entries. A Sudoku square of order $n = k^2$ is a Latin square of order n such that every element in $[n]$ appears exactly once in each block. We say that the (i, j)'th block belongs to the i'th row block and the j'th column block.

Notions of partial Sudoku square, extensions of a partial Sudoku square, defining sets, critical sets, and forcing sets for Sudoku squares can be defined similarly.

A central problem is to determine the size of the smallest defining set for Sudoku squares of order n. For $n = 9$ (regular Sudoku) extensive computer search showed that this number is 17 (McGuire et al [6]), but the asymptotics of this value is unknown. For Latin squares, this number is conjectured to be $\lfloor n^2/4 \rfloor$ (Mahmoodian [5], Bate and Van Rees [1]). A construction based on back-circulant Latin squares shows that this number is at most $\lfloor n^2/4 \rfloor$, but the best proven lower bound is just slightly superlinear. Also, the $\lfloor n^2/4 \rfloor$ conjecture is proved if "defining set" is replaced by "forcing set". For Sudoku square, we show that the same construction (with a permutation on the rows of the matrix) works, giving an upper bound of $\lfloor n^2/4 \rfloor$. We also show that the size of the smallest forcing set for Sudoku squares of order n is at least $\Theta(n^2)$. Our conjecture is that the size of the smallest defining set for Sudoku squares of order n is also $\Theta(n^2)$. We conclude with the discussion of many Sudoku-related problems that remain open.

Lower bound on the size of forcing sets

In this section, we prove the main result of this paper, which is the following lower bound on the size of the smallest forcing set in Sudoku squares. This result, combined with the observation that essentially the same construction as the one for back-circulant Latin squares gives us a forcing set of size $\lfloor n^2/4 \rfloor$ for an equivalent Sudoku square, shows that the smallest forcing set of Sudoku squares of order n is precisely $\Theta(n^2)$.

Theorem 1 *For every n, the size of the smallest forcing set for Sudoku squares of order $n = k^2$ is at least $\Omega(n^2)$.*

Proof. Let F be a partial Sudoku square that is a forcing set, and consider the forcing order on the entries not specified by F. Let S denote this ordering, i.e., S_1 is an entry that is forced by F, S_2 is an entry that is forced by $F \cup \{S_1\}$, and so on.

We start by defining a subsequence S' of S as follows: $S'_1 = S_1$, and for every $i > 1$, S'_i is the first element in S after S'_{i-1} that is not in the same row, the same column, or the same block as any of $S'_1, S'_2, \ldots, S'_{i-1}$. In other words, S' is obtained from S by removing elements that are in the same row, same column, or same block. Therefore, the sequence S' has at most n elements, and contains at most one element from each row, each column, and each block of the Sudoku square.

We now transform S' into an ordering of a $k \times k$ square. More formally, we define a permutation π of the set $[k]^2$ as follows: for every i where S'_i is defined, π_i is the coordinates of the block

containing S'_i. Since S' contains at most one element from each block, the π_i's defined based on S'_i's are distinct. There can be blocks with no element present in S'; we add the coordinates of such blocks in an arbitrary order to the end of π. This completes the definition of the permutation π of $[k]^2$. The proof of the theorem is based on two lemmas. The first lemma bounds the size of the forcing set in terms of a quantity associated with the permutation π, and the second lemma bounds this quantity for every such permutation.

To state the first lemma, we need a few notations. For every permutation π of $[k]^2$ and every $u, v \in [k]^2$ ($u \neq v$), we say $u \prec_\pi v$ if u comes before v in π. Let $B^r_\pi(v)$ denote the number of $u \in [k]^2$ such that $u \prec_\pi v$ and u and v are in the same row (i.e., $u = (i, j)$ and $v = (i, j')$ for $i, j, j' \in [k]$). Similarly, let $B^c_\pi(v)$ denote the number of $u \prec_\pi v$ that are in the same column as v. Finally, let $B_\pi(v) = B^r_\pi(v) + B^c_\pi(v)$. We are now ready to state the first lemma.

Lemma 1 *Let π be the permutation defined based on a forcing set F using the above procedure. Then,*

$$|F| \geq \sum_{i=1}^{n} \max(0, n + 1 - 2i - (2k - 2)B_\pi(\pi_i)).$$

Let $L(\pi)$ denote the quantity on the right-hand side of the inequality in Lemma 1. The second lemma bounds this quantity for every permutation π.

Lemma 2 *There is a constant c such that for every permutation π of $[k]^2$, we have $L(\pi) \geq cn^2$.*

We start by proving the first lemma.

[Proof of Lemma 1] Let $i \in [n]$ be an index for which S'_i exists. Therefore, π_i is the coordinates of the block containing S'_i. We argue that to uniquely force S'_i, we need at least $n + 1 - 2i - (2k - 2)B_\pi(\pi_i)$ *new* elements in F (i.e., elements other than the ones needed to force S'_j for $j < i$).

Let A_i denote the set of entries of the Sudoku square that are in the same row, same column, or the same block as S'_i. The following lemma bounds the cardinality of the intersection of these sets.

Lemma 3 *For every i, j, $j \neq i$, if S'_i and S'_j are not in the same row block or the same column block, then $|A_i \cap A_j| = 2$. If they are on the same row block or same column block, then $|A_i \cap A_j| = 2k$.*

Proof. Proof is easy. Omitted for now.

Since F is a forcing set, by the time S'_i is forced, there must be at least $n-1$ entries in A_i whose values are uniquely specified. We argue that out of these $n-1$, at most $2(i-1)+(2k-2)B_\pi(\pi_i)$ are either forced in previous steps or already counted in F, and therefore there must be at least $n + 1 - 2i - (2k - 2)B_\pi(\pi_i)$ of them that are in F and are not previously counted in F. Note that any entry that is either forced in previous steps or already counted in F must be in A_j for a $j < i$. This is because any entry that is forced before S'_i must either be present in the sequence $S'_1, \ldots, S'_{i-1}$, or be in the same row, same column, or same block as one of the elements of this sequence. Either way, this element belongs to $\bigcup_{j<i} A_j$. Also, in each step $j < i$, we

count elements of F that are used to force S'_j, and these elements belong to A_j. Therefore, the number of elements in A_i that either forced before S'_i or are already counted in F is at most $|A_i \cap (\bigcup_{j<i} A_j)|$. To bound this cardinality, we use Lemma 3. By this lemma and the definition of $B_\pi(\pi_i)$, the value of $|A_i \cap A_j|$ is equal to $2k$ for precisely $B_\pi(\pi_i)$ values of j and is equal to 2 for the remaining $i - 1 - B_\pi(\pi_i)$. Therefore,

$$|A_i \cap (\bigcup_{j<i} A_j)| \leq \sum_{j<i} |A_i \cap A_j| = 2kB_\pi(\pi_i) + 2(i - 1 - B_\pi(\pi_i)).$$

Therefore, there must be at least $\max(0, n + 1 - 2i - (2k - 2)B_\pi(\pi_i))$ elements in F that are used to force S'_i and are not counted in previous steps.

Next, we consider i's for which S'_i does not exist. Recall that when the length of S' is less than n, we append a list of block coordinates that contain no element of S' at the end of π in an arbitrary order. Therefore π_i is the coordinate of a block none of whose elements appears in S'. This means that all of the n elements of the block at coordinates π_i must either be in F, or in the same row, column, or block as an element of S', since otherwise they would have been included in S'. We can now repeat the same argument with A_i replaced by the set of entries in the block at coordinates π_i.

Putting these cases together, we get that in total F must contain at least $\sum_{i=1}^{n} \max(0, n + 1 - 2i - (2k - 2)B_\pi(\pi_i))$ elements.

Next, we prove Lemma 2, which gives a bound on the quantity $L(\pi)$ for every permutation π of $[k]^2$.

[Proof of Lemma 2] Let $\alpha \in [0, 1]$ be a parameter that will be fixed later. For convenience we assume that $(1 - \alpha)k/2$ (and therefore $(1 - \alpha)n/2$) is an integer. We use the following lower bound on $L(\pi)$:

$$L(\pi) \geq \sum_{i=1}^{(1-\alpha)n/2} \max(0, n + 1 - 2i - (2k - 2)B_\pi(\pi_i)).$$

Since for every $i \leq (1 - \alpha)n/2$, we have $n + 1 - 2i > \alpha n$, the above inequality implies:

$$L(\pi) \geq \sum_{i=1}^{(1-\alpha)n/2} \max(0, \alpha n - (2k - 2)B_\pi(\pi_i)).$$

For every i, we define

$$L(\pi, i) = \begin{cases} 0 & \text{if } \max\{B^r_\pi(\pi_i), B^c_\pi(\pi_i)\} > \frac{\alpha n}{4k-4} \\ \alpha n - (2k - 2)B_\pi(\pi_i) & \text{otherwise.} \end{cases}$$

It is easy to see that $\max(0, \alpha n - (2k - 2)B_\pi(\pi_i)) \geq L(\pi, i)$ for every i. Therefore,

$$L(\pi) \geq \sum_{i=1}^{(1-\alpha)n/2} L(\pi, i).$$

Let $L'(\pi)$ denote the right-hand side of the above inequality. We will show how the permuta-

tion π can be transformed into a structurally simpler permutation π' such $L'(\pi) \geq L'(\pi')$. Let $t = \lfloor \frac{\alpha n}{4k-4} \rfloor$. Consider the smallest index i such that $\max\{B^r_\pi(\pi_i), B^c_\pi(\pi_i)\} = t$, and assume, without loss of generality, that $B^r_\pi(\pi_i) = t$. This means that there are t indices $i_1, i_2, \ldots, i_t = i$ such that π_{i_ℓ}'s, for all $\ell = 1, \ldots, t$, are on the same row in $[k]^2$. It is not hard to see that moving all these π_{i_ℓ}'s to the beginning of the permutation does not change the value of $L'(\pi)$. Furthermore, all other elements of the same row can be added after these elements without increasing $L'(\pi)$. Therefore, by moving all entries that are on the same row as π_i to the beginning of the permutation, we obtain another permutation whose L' value is not more than the L' value of the original permutation. We can continue this process, by finding the first index i' such that $\max\{B^r_\pi(\pi_{i'}), B^c_\pi(\pi_{i'})\} = t$ and $\pi_{i'}$ is not on the same row as π_i. Using the same argument, depending on whether $B^r_\pi(\pi_{i'}) = t$ or $B^c_\pi(\pi_{i'}) = t$, elements of the row or column of $\pi_{i'}$ (except possibly the ones that were on the same row as π_i) can be moved right after the elements of the row of π_i. Continuing with this process, we can build a permutation π' such that $L'(\pi) \geq L'(\pi)$, and π' has the following structure: it starts with the list of all elements of a row/column of $[k]^2$, then all elements of another row/column of $[k]^2$ except the ones that have appeared before, and so on.

What remains is to prove that for a permutation π' that has the above structure, $L'(\pi') = \Omega(n^2)$. Using the structure of π', we can decompose it into segments, where each segment lists all elements of a row/column of $[k]^2$ except the ones that are listed that are listed in previous segments. We call a segment a row/column segment, depending on whether it is a list of elements in a row or a column of $[k]^2$. The value of a segment is the sum of $L(\pi', i)$ for all i that belong to that segment. Let ℓ^r_j (ℓ^c_j, respectively) denote the number of row (column, respectively) segments *before* the j'th segment. Therefore, if the j'th segment is a column segment, its value can be written as:

$$\begin{aligned} V_j &= (\alpha n - (2k-2)\ell^r_j) + (\alpha n - (2k-2)(\ell^r_j + 1)) + \cdots + (\alpha n - (2k-2)t) \\ &= (t - \ell^r_j + 1)(\alpha n - (k-1)(t + \ell^r_j)), \end{aligned} \tag{1}$$

if $\ell^r_j \leq t$. We also have $V_j = 0$ if $\ell^r_j > t$. If the j'th segment is a row segment, we get a similar expression for V_j, with ℓ^r_j replaced by ℓ^c_j.

Since each segment contains at most k elements, there are at least $\frac{(1-\alpha)n}{2k} = (1-\alpha)k/2$ segments that are entirely contained in the first $(1-\alpha)n/2$ elements of π. Therefore, $L'(\pi')$ is at least the sum of the values of the first $(1-\alpha)k/2$ segments, i.e., $L'(\pi') \geq \sum_{j=1}^{(1-\alpha)k/2} V_j$. We let $L''(\pi') := \sum_{j=1}^{(1-\alpha)k/2} V_j$.

The final step is to change π' to another permutation π'' (with a similar segmented structure) such that $L''(\pi') \geq L''(\pi'')$. We do this as follows: assume, for some j, $\ell^r_j > \ell^c_j$ and the j'th segment is a row segment. Find the smallest index $j' \in [j, (1-\alpha)k/2]$ such that the j''th segment is a column segment, if such an index exists. We can write down the difference in the total L'' value if we replace the order of the segments j' and $j'-1$ (i.e., first list all elements in the column corresponding to segment j' and then list all elements in the row corresponding to segment $j'-1$). It is easy to see that the inequality $\ell^r_j > \ell^c_j$ implies that this swap cannot increase the L'' value of the permutation. If such an index j' does not exist, we can change the last segment to a column segment. Again, it is not hard to see that the assumption $\ell^r_j > \ell^c_j$ implies that this change does not increase the L'' value of the permutation. Similar statements hold if we switch the role of row segments and the column segments. By repeatedly using this

procedure, we get a permuation π'' that consists of alternating row and column segments, and satisfies $L''(\pi') \geq L''(\pi'')$.

All that remains is to write down the value of $L''(\pi'')$. This permutation satisfies $\ell_j^r = \ell_j^c = \lfloor j/2 \rfloor$ for j odd and $\ell_j^r = \lfloor j/2 \rfloor = \ell_j^c + 1$ for j even. Using Equation (1), the value of $L''(\pi'')$ can be written as follows:

$$\begin{aligned} L''(\pi'') &= \sum_{s=0}^{p}(t-s+1)(\alpha n-(k-1)(t+s))+ \\ &\quad \sum_{s=0}^{p}(t-s+1)(\alpha n-(k-1)(t+s)) \\ &\geq 2\sum_{s=0}^{p}(t-s+1)(\alpha n-(k-1)(t+s)), \end{aligned}$$

where $p = \min\{t, \lfloor (1-\alpha)k/4 \rfloor\}$. Recall that $t = \lfloor \frac{\alpha n}{4k-4} \rfloor$. Therefore,

$$\begin{aligned} L''(\pi'') &\geq 2(k-1)\sum_{s=0}^{p}(t-s)\left(\frac{\alpha n}{k-1}-t-s\right) \\ &\geq 2(k-1)\sum_{s=0}^{p}(t-s)^2. \end{aligned}$$

If we pick α in such a way that $t \leq \lfloor (1-\alpha)k/4 \rfloor$, we have $p = t$ and therefore,

$$L''(\pi'') \geq \frac{2(k-1)t^3}{3} \geq \frac{2\alpha^3 k^4}{3 \cdot 4^3}.$$

Now, it suffices to pick any $\alpha < 1/2$. It is easy to see that this satisfies the inequality $t \leq \lfloor (1-\alpha)k/4 \rfloor$, and gives us $L''(\pi'') \geq \frac{1}{3.4^4} n^2$.

The theorem follows by putting Lemmas 1 and 2 together.

Conclusions (Open Problems and Future Directions)

Sudoku is a fascinating source of new interesting open questions in combinatorics. The obvious open question is whether the result in this paper can be strengthened to defining sets. Our conjecture is that this is true, i.e., the size of the smallest defining set of Sudoku squares of order n is $\Theta(n^2)$. If true, this is probably a difficult problem, since the similar question for Latin squares has been open for years.

A simpler problem is to strengthen the result to a notion like "semi-strong critical set" ([1]), as defined similarly to Latin squares. Also, finding any super-linear lower bound is an interesting open question. Note that in the case of Latin squares, the best lower bounds we know are just barely superlinear.

As mentioned earlier in the paper, for Latin squares, there is a construction for a defining set of size $\lfloor n^2/4 \rfloor$. This defining set has a unique extension to a back-circulant Latin square. In fact, it is proved that for even n, this is the smallest defining set of a back-circulant Latin square. It is not hard to show that by permuting rows and columns of a back-circulant Latin square, one can obtain a Sudoku square. This gives a construction for a defining set of size $\lfloor n^2/4 \rfloor$ for Sudoku.

Two questions remain open: Are there Sudoku squares with smaller defining sets, and are there smaller defining sets for this particular Sudoku squares. The answer to both of these questions are conjectured to be negative in the case of Latin squares (and proved to be so in the case of the second question for n even). For Sudoku, however, these conjectures might not be true, since the block constraint could reduce the size of the smallest defining set.

There are also many computational open questions arising from the Sudoku puzzle. The first question is whether the problem of Sudoku completion (given a partial Sudoku square, is there a completion to a Sudoku square) is NP-hard. Our conjecture, of course, is that it is. A more difficult problem is the complexity of completing a defining set (i.e., a set that is guaranteed to have a unique completion) to a full Sudoku square. As a less mathematical problem, it would be interesting if one can define a measure of difficulty for Sudoku puzzles that roughly correspond to how hard the puzzle is for humans. An online search reveals many 9×9 Sudoku puzzles that are claimed to be the hardest Sudoku puzzle. It would be interesting to have a quantitative measure of such puzzles.

Finally, there are many open combinatorial conjectures for Latin squares for which the corresponding Sudoku problem might be more approachable. Two example are two long-standing conjectures of Brualdi-Stein and Ryser.

Conjecture 1 ([3, 8]) *Every Latin square of even order n contains a partial transversal of length $n - 1$.*

Conjecture 2 ([7]) *Every Latin square of odd order contains a transversal.*

Another interesting question is whether Galvin's theorem about list colorability of the Latin squares ([4]) to Sudoku squares.

References

[1] J. A. Bate and G. H. J. van Rees. The size of the smallest strong critical set in a Latin square. *Ars Combin.*, 53:73–83, 1999.

[2] Richard A. Brualdi and Herbert J. Ryser. *Combinatorial matrix theory*, volume 39 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, 1991.

[3] J. Dénes and A. D. Keedwell. *Latin squares and their applications*. Academic Press, New York, 1974.

[4] Fred Galvin. The list chromatic index of a bipartite multigraph. *J. Combin. Theory Ser. B*, 63(1):153–158, 1995.

[5] E. S. Mahmoodian. Some problems in graph colorings. In *Proceedings of the 26th Annual Iranian Mathematics Conference, Vol. 2 (Kerman, 1995)*, pages 215–218. Shahid Bahonar Univ. Kerman, Kerman, 1995.

[6] Gary McGuire, Bastian Tugemann, and Gilles Civario. There is no 16-clue Sudoku: solving the Sudoku minimum number of clues problem via hitting set enumeration. *Exp. Math.*, 23(2):190–217, 2014.

[7] H. J. Ryser. Neuere Probleme in der Kombinatorik (prepared by D.W. Miller). *Vortrage uber Kombinatorik*, pages 69–91, 1967. (cited in [2]).

[8] S. K. Stein. Transversals of Latin squares and their generalizations. *Pacific J. Math.*, 59(2):567–575, 1975.