

Approaching the minimum number of clues Sudoku problem via the polynomial method

Aden Forrow John R. Schmitt

Department of Mathematics
Middlebury College
Middlebury, Vermont 05753, U.S.A.

aforrow@mit.edu jschmitt@middlebury.edu

Submitted: July 24, 2013; Accepted: YY; Published: XX
Mathematics Subject Classification: 05C35, 05C31, 05C15

Abstract

Determining the minimum number of clues that must be present in a Sudoku puzzle in order to uniquely complete the puzzle is known as the *minimum number of clues problem*. For a 9×9 Sudoku board, it has been conjectured that one needs 17 clues. We apply the polynomial method to the analogous problem for the 4×4 Shidoku board to illustrate how one might approach the more general problem.

Keywords: polynomial method, Schauz's Coefficient Formula, Sudoku, minimum number of clues

1 Introduction

Shortly after the introduction of the popular number-puzzle Sudoku, enthusiasts and mathematicians alike began asking the following question: how difficult can a Sudoku puzzle be? One measure of difficulty is the number of clues (or hints) initially given in the puzzle – the fewer clues given, the harder the puzzle. But just how few clues could be given and still yield a puzzle with a unique completion? The conjectured value is 17. (It is unknown to the authors who initially made the conjecture.) That is, it is conjectured that any puzzle with 16 or fewer clues that has a completion has at least two distinct completions. Determining the minimum number of clues necessary is known as the *minimum number of clues problem*.

As evidence that many do believe that 17 is the correct answer to the minimum number of clues problem for the 9×9 Sudoku board, one should consult the webpage¹ of G. Royle of the University of Western Australia. The webpage contains his collection of 49,151

¹See <http://school.maths.uwa.edu.au/~gordon/sudokumin.php>.

distinct Sudoku puzzles each with 17 clues (one of which is reproduced in Figure 1). Royle notes that the members of his set of 49,151 puzzles are *mathematically inequivalent*. That is, any permissible permuting of the symbols, rows, columns, boxes or any transposing of a given puzzle will not yield any of the other 49,150.

Recently, the team of G. McGuire, B. Tugemann, and G. Civario [8] claimed to have solved the problem, confirming that 17 is the answer. Their approach involves an exhaustive (!) computer search (that included 7.1 million core hours on a 320-node cluster). A brief critique of their approach, with which we mostly agree and do not reproduce here, is given by J. Cooper and A. Kirkpatrick [3]. It suffices to say that a “human-readable” mathematical approach is still wanting.

While the puzzle in Figure 1 provides an upper bound of 17 to the minimum number of clues Sudoku problem, the best known lower bound that can be arrived at via a purely mathematical argument is 8. The argument is rather trivial. Suppose that only seven clues are given. Then there are at least two labels not used, say x and y , and from a completion of the puzzle we may replace all instances of x with y and vice versa. The result of this is a distinct completion.

The purpose of this paper is to suggest a mathematical approach that might prove successful in yielding a human-readable solution to the minimum number of clues problem or at least provide a computational alternative to that of McGuire et al.[8]. We use the polynomial method – this is a method that encodes a combinatorial problem with a polynomial, the nonzero values of which correspond to some interesting property of the discrete system, and then uses algebraic results about polynomials to yield information about the combinatorial problem by showing that some nonzero value exists or that none do. The results we use are due to U. Schauz [9]. We illustrate an application of these results by solving the minimum number of clues problem for the 4×4 Shidoku board, showing that one needs at least four clues.

Theorem 1 *The minimum number of clues in a Shidoku puzzle with a unique completion is 4.*

The above statement is by no means new (nor surprising). The statement has been given by A. Herzberg and M. Ram Murty [6], but their proof still requires us to “(check) the cases that arise one by one.” L. Taalman [10] gives a short proof which uses ideas found in the McGuire et al. approach to the larger problem. However, the reader should not assume that the objective of this paper is to prove this theorem; it is not. Rather, our objective is to illustrate the aspects of the polynomial method that are given below.

This paper is organized as follows. In Section 2 we discuss the aspects of the polynomial method that we require, including results of U. Schauz [9], though we adopt the notation of M. Lason [7], who gave simpler proofs of some of Schauz’s results. In Section 3 we give definitions and notation relevant to the minimum number of clues problem. In Section 4 we prove Theorem 1. In Section 5 we comment on the proof of Theorem 1 given in Section 4 and discuss how this approach might be extended to the minimum number of clues Sudoku problem.

							1	5
7	9							
			2					
					8	7		6
		1						
						9		
	7					8	3	
4			1	5				
			3					

Figure 1: A 17-clue Sudoku puzzle

2 The polynomial method and Schauz's Coefficient Formula

Let \mathbb{F} be an arbitrary field. The Fundamental Theorem of Algebra tells us that a degree- t polynomial $f(x)$ contained in the polynomial ring $\mathbb{F}[x]$ has at most t zeros. Said another way, for any set A of cardinality greater than t and contained in \mathbb{F} , there is an element $a \in A$ such that $f(a)$ is nonzero. One may think of this as saying, either a polynomial is zero everywhere or it is zero in very few places. The second of two theorems collectively known as the Combinatorial Nullstellensatz generalizes this fact to polynomials of several variables – it is due to N. Alon [1, Theorem 1.2]. We may think of it as saying that a multivariable polynomial that isn't zero everywhere has a non-root in a box of large enough volume. To reach this conclusion, Alon's theorem requires the coefficient of an appropriate monomial of maximum degree to be nonzero. Alon's assumption can be relaxed somewhat; it suffices to find a monomial with nonzero coefficient that does not divide any other monomial in the polynomial.

To state the generalization, we require the following definition. Given a polynomial $f \in \mathbb{F}[x_1, \dots, x_n]$, define the *support of f* , $\text{Supp}(f)$, as the set of all $(\alpha_1, \dots, \alpha_n)$ such that the coefficient of $x_1^{\alpha_1} \dots x_n^{\alpha_n}$ in f is nonzero. We say $(\alpha_1, \dots, \alpha_n) \geq (\beta_1, \dots, \beta_n)$ if $\alpha_i \geq \beta_i$ for all i ; this gives us a partial ordering of the elements of $\text{Supp}(f)$.

Theorem 2 [Generalized Combinatorial Nullstellensatz, U. Schauz [9] Theorem 7.3, M. Lason [7] Theorem 2] *Let \mathbb{F} be an arbitrary field, and let f be a polynomial in $\mathbb{F}[x_1, \dots, x_n]$. Suppose that $(\alpha_1, \dots, \alpha_n)$ is maximal in $\text{Supp}(f)$. Then for any subsets A_1, \dots, A_n of \mathbb{F} satisfying $|A_i| \geq \alpha_i + 1$, there are $a_1 \in A_1, \dots, a_n \in A_n$ so that $f(a_1, \dots, a_n) \neq 0$.*

Schauz also discovered a general method for finding coefficients of maximal monomials in terms of the values of a polynomial. The following function will be useful in writing the formula:

$$N(a_1, \dots, a_n) = \prod_{i=1}^n \prod_{b \in A_i \setminus \{a_i\}} (a_i - b).$$

The function $N(a_1, \dots, a_n)$ may be described as a normalizing factor for

$$\chi_{(a_1, \dots, a_n)}(x_1, \dots, x_n) = N(a_1, \dots, a_n)^{-1} \cdot \prod_{i=1}^n \prod_{b \in A_i \setminus \{a_i\}} (x_i - b).$$

The function $\chi_{(a_1, \dots, a_n)}$ selects out the point (a_1, \dots, a_n) ; it equals 1 at this point and zero at every other $(x_1, \dots, x_n) \in A_1 \times \dots \times A_n$.

With this we can now write an expression for the coefficient of any monomial in f whose exponent vector is maximal in $\text{Supp}(f)$ or greater than a maximal element of $\text{Supp}(f)$. Note that when we first encountered polynomials (say, in elementary school), we knew the coefficients, plugged in values for the variables, and got out a value of the polynomial; here we do that backwards, using values of the polynomial to get to coefficients. This will prove useful, allowing us to learn more about the polynomial values from one coefficient, even if that coefficient is zero.

Theorem 3 [*Coefficient Formula, U. Schauz [9] Theorem 3.2, M. Lasoń [7] Theorem 3*] *Let f be a polynomial in $\mathbb{F}[x_1, \dots, x_n]$ and let $f_{\alpha_1, \dots, \alpha_n}$ denote the coefficient of $x_1^{\alpha_1} \dots x_n^{\alpha_n}$ in f . Suppose that there is no greater element than $(\alpha_1, \dots, \alpha_n)$ in $\text{Supp}(f)$. Then for any sets A_1, \dots, A_n in \mathbb{F} such that $|A_i| = \alpha_i + 1$ we have*

$$f_{\alpha_1, \dots, \alpha_n} = \sum_{(a_1, \dots, a_n) \in A_1 \times \dots \times A_n} \frac{f(a_1, \dots, a_n)}{N(a_1, \dots, a_n)}. \quad (1)$$

Note that in Theorem 3 we do not say $(\alpha_1, \dots, \alpha_n)$ is maximal in the support of f because we do not require it to be in $\text{Supp}(f)$. Theorem 2 follows immediately from Theorem 3, as does the next corollary.

Corollary 4 [*U. Schauz [9] Corollary 3.4*] *If $f_{\alpha_1, \dots, \alpha_n} = 0$, then either f vanishes over $A_1 \times \dots \times A_n$ or f has at least two nonzero values over $A_1 \times \dots \times A_n$.*

PROOF: If f is nonzero for exactly one element $(a_1, \dots, a_n) \in A_1 \times \dots \times A_n$, Equation (1) becomes $f_{\alpha_1, \dots, \alpha_n} = \frac{f(a_1, \dots, a_n)}{N(a_1, \dots, a_n)} \neq 0$, as all other terms in the sum are zero. \square

The following corollary is a direct consequence of Corollary 4 and these two corollaries will be our main tools going forward.

Corollary 5 [*U. Schauz [9] Corollary 3.4*] *Let f be a polynomial of degree d in $\mathbb{F}[x_1, \dots, x_n]$. Then for any subsets A_1, \dots, A_n of \mathbb{F} satisfying $\sum_{i=1}^n (|A_i| - 1) > d$, f either vanishes over $A_1 \times \dots \times A_n$ or f has at least two nonzero values over $A_1 \times \dots \times A_n$.*

PROOF: Consider the monomial $x_1^{|A_1|-1} \cdots x_n^{|A_n|-1}$ in f . This is a monomial of degree greater than d , so its coefficient is zero. Applying Corollary 4, the conclusion follows immediately. \square

As pointed out in T. Tao and V. Vu's text on additive combinatorics [11], the polynomial method is useful in providing lower bounds concerning cardinalities of sets. This is precisely how we use the method below. To do so, we will most frequently make use of Corollary 5 since the use of Corollary 4 requires us to explicitly compute a coefficient, at times a challenging task. (For other applications of these corollaries, see Section 4 of [9].) When we do need to compute a coefficient, we will make use of a result of N. Alon and M. Tarsi [2].

The following definitions are given in [2]. The *graph polynomial* $f_G = f_G(x_1, \dots, x_n)$ of an undirected graph $G = (V, E)$ on a set $V = \{v_1, \dots, v_n\}$ of n vertices is defined by $f_G(x_1, \dots, x_n) = \prod \{(x_i - x_j) : i < j, \{v_i, v_j\} \in E\}$. An oriented edge (v_i, v_j) of G is *decreasing* if $i > j$. An orientation D of G is *even* (*odd*) if it has an even (odd) number of decreasing edges. For non-negative integers d_1, \dots, d_n , let $DE(d_1, \dots, d_n)$ and $DO(d_1, \dots, d_n)$ denote, respectively, the sets of all even and odd orientations of G in which the out degree of the vertex v_i is d_i for $1 \leq i \leq n$.

Lemma 6 [N. Alon, M. Tarsi [2]] *In the above notation*

$$f_G(x_1, \dots, x_n) = \sum_{d_1, \dots, d_n \geq 0} (|DE(d_1, \dots, d_n)| - |DO(d_1, \dots, d_n)|) \prod_{i=1}^n x_i^{d_i}.$$

This result follows simply from the observation that each term in the expansion of the product f_G corresponds to an orientation of the edges of G .

3 Notation and terminology

Sudoku is a single-player game (or puzzle) played on a 9×9 matrix with a given subset of cells filled with labels called *clues* – these labels are usually drawn from the set $\{1, \dots, 9\}$ – and one must fill in the remaining unfilled cells subject to the following rule, no two cells from the same row, column or one of the nine 3×3 sub-matrices – see the darkened lines in Figure 1 – share the same label. If there is a labeling that adheres to these rules, such a labeling is called a *completion*. We are interested in puzzles that have a unique completion. We also define a more general version. We define *Sudoku*(n) to be the single-player game played on an $n^2 \times n^2$ matrix with a given subset of cells filled with labels called clues – say these are drawn from the set $\{1, \dots, n^2\}$ – and one must fill in the remaining unfilled cells subject to the following rule, no two cells from the same row, column or one of the n^2 $n \times n$ sub-matrices called *boxes* share the same label. (The term *chute* will be used to refer to a set of n vertical or horizontal collinear boxes.) Thus, *Sudoku*(3) is simply *Sudoku*. *Sudoku*(2) is also referred to as *Shidoku*, an example of which is given below. The example below is a *Shidoku* puzzle with a unique completion; we will prove that this 4-clue puzzle is of minimum size.

The minimum number of clues *Sudoku* conjecture may now be stated as follows.

Conjecture 7 [Folklore] *The minimum number of clues in a Sudoku puzzle with a unique completion is 17.*

Notice that the rule set for $\text{Sudoku}(n)$ leads naturally to a graph. To each cell of the matrix we identify a vertex and a pair of vertices is joined by an edge if the corresponding cells lie in the same row, column or one of the n^2 boxes. Each edge of this graph may be thought of as one of the many rules defining play on the puzzle. That is, an edge corresponds to the fact that we insist that the two labels associated to the vertices at opposite ends of the edge to be different. One can easily check that each of the n^4 vertices is incident with same number of edges, namely $(n^2-1)+(n^2-1)+(n^2-(n-1)-(n-1)-1) = 3n^2 - 2n - 1$. Thus, the graph has $\frac{n^4(3n^2-2n-1)}{2}$ edges. We call the graph that arises from the rule set of $\text{Sudoku}(n)$ the *Sudoku*(n) graph and denote it by SUD_n .

	1	2	
1			3

To each vertex in SUD_n we assign a name (i, j) which corresponds to the i^{th} row and j^{th} column of the associated cell. We label rows from top to bottom and columns from left to right. Thus, the $(2, 3)$ -cell in the above Shidoku puzzle, which contains the label 2, shares its name and label with a vertex in SUD_2 .

It is important to note that, in the above puzzle one has more information than is needed, that is, some of the rules are *redundant*. For example, consider the edge joining vertices $(4, 1)$ and $(4, 2)$ and the edge joining the vertices $(2, 2)$ and $(4, 2)$. As vertices $(4, 1)$ and $(2, 2)$ are both adjacent to the same vertex and each has the label 1, we only need one of these two edges to know that the label for vertex $(4, 2)$ cannot be a 1. We may consider one of the edges to be a *redundant edge* and it can be safely removed from the graph as the information it contains is entailed by others. *Finding redundant edges will be key to our work.* A recent paper by B. Demoen and M.G. de la Banda [4], aptly titled Redundant Sudoku Rules, discusses how one might do this starting with a Sudoku board with no clues present. A subsequent technical report [5] by the same authors does the same for the Shidoku board.

Figure 2, given in [5], shows some redundant sets of edges that can be removed from the Shidoku graph SUD_2 while still supplying enough rules to know whether or not a completion satisfies all of the initial rules. Note that in each of the graphs there are precisely 16 edges present, meaning that we have kept 40 of the original 56 edges. We will refer to the graphs from left to right and top to bottom as G_1, \dots, G_8 . Each of these graphs has multiple symmetries, obtained by swapping two rows or two columns

of vertices from the same chute, swapping two chutes or reflecting across either of the long diagonals of the matrix. In regards to the computational challenges of the minimum number of clues problem, it is important to note that these graphs are the result of a short combinatorial argument, see [5].

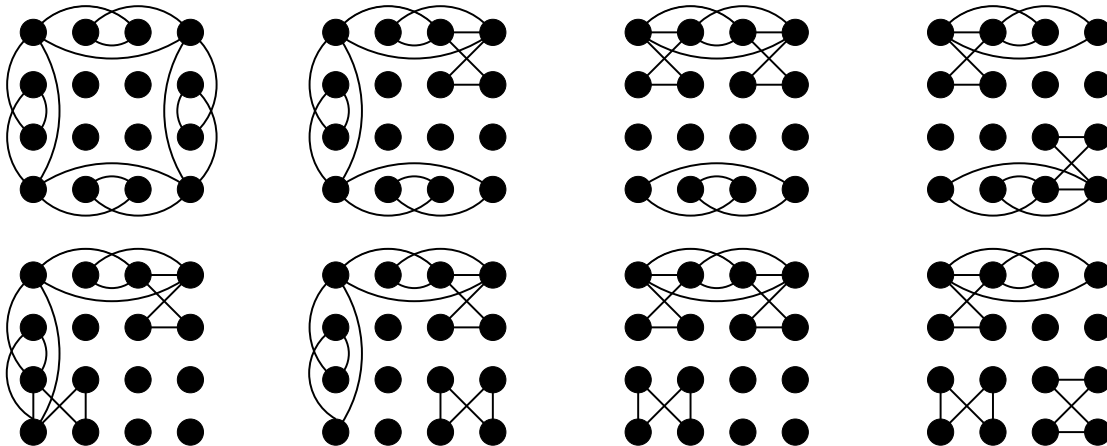


Figure 2: G_1, \dots, G_8 , from left to right, top to bottom

4 Solving the minimum number of clues Shidoku problem

If one were to take a brute-force approach to proving the lower bound of 4 for the minimum number of clues Shidoku problem, then one would have to consider at most $\binom{16}{3} \cdot 4^3 = 35,840$ sets of initial clues as there are $\binom{16}{3}$ ways to choose the cells and for each cell chosen there are 4 possible labels. Of course one can find symmetries to reduce this number, but it gives us a sense of the size of the problem. Furthermore, searching this space via a computer is much more manageable than the computation that McGuire et al. [8] took on and we believe that many of the objections pointed out by Cooper and Kirkpatrick [3] would dissolve. *However*, our intention is to highlight what we believe to be a beautiful algebraic approach, one that might lead to a human-readable solution to the minimum number of clues Sudoku problem.

PROOF OF THEOREM 1

Note that the Shidoku puzzle given above is of size 4 and has a unique completion. This establishes the upper bound.

We now show the lower bound. We begin by noting that any $\text{Sudoku}(n)$ puzzle that omits two distinct labels from the clue-set, say x and y are omitted, will allow any completion of the puzzle to yield a distinct completion upon a permutation of labels that interchanges x and y . Thus, we need only consider 3-clue puzzles with three distinct labels.

The approach is as follows. Let a 3-clue set with three distinct labels be given and consider SUD_2 . If necessary, we extend to a partial completion, where the additional clues determined are forced by the given three. For each unlabeled vertex u in SUD_2 assign a variable x_u . Let n count the number of such variables assigned. (Note that n will be at most $16 - 3 = 13$, and will be less only when we extend to a partial completion.) For each labeled vertex v , let c_v denote the label the corresponding cell has been assigned. Depending upon the relative placement of the clues and the relative labeling, we will choose one of G_1, \dots, G_8 or the empty graph and delete the edges contained in this graph and possibly some additional edges from SUD_2 . For each edge $\{u, v\}$ in the resulting graph G associate a linear factor of the form $(x_u - x_v)$, $(x_u - c_v)$ or $(c_u - c_v)$, where the factor corresponds to the variable-variable, variable-label or label-label assignment, respectively, of the endpoints of the edge. Now consider the polynomial that is the product of all such linear factors; call it f_G , the graph polynomial. We will consider f_G as a polynomial in the x_u 's. That is, $f_G = f_G(x_1, \dots, x_n) \in \mathbb{R}[x_1, \dots, x_n]$. This polynomial is zero when any one of these factors is zero and is nonzero otherwise. When we evaluate this polynomial at a point, the point will correspond to a labeling of the graph and the polynomial evaluates as zero if and only if any pair of labels violates the rules of the puzzle. A nonzero corresponds to a completion of the puzzle. We seek points that belong to $A_1 \times \dots \times A_n$, where $A_i = \{1, 2, 3, 4\}$ for $1 \leq i \leq n$. If we can show that there is no greater element than the n -length exponent vector of the form $(3, 3, \dots, 3)$ in $\text{Supp}(f_G)$ and the coefficient of the corresponding monomial is 0 or the degree of f_G is 'small enough' (i.e., strictly less than $3n$), then either by Corollary 4 or Corollary 5, respectively, either f_G vanishes over $A_1 \times \dots \times A_n$ or f_G has at least two nonzero values over $A_1 \times \dots \times A_n$. If the former occurs, then this is saying that there is no completion and if the latter occurs, then this is saying that there is more than one completion. In either case, the puzzle does not have a unique completion.

Below, when we say we "apply G_i ", we mean we delete the edges from SUD_2 that are also in G_i or the described symmetry of G_i .

Swapping two rows or two columns in the same chute does not change the adjacency graph, nor does swapping two chutes or reflecting across the diagonal, so puzzles that can be transformed into each other by such operations have the same number of completions. Also, the identity of the labels does not matter: the puzzle is the same whether the three given clues are $(1, 2, 3)$, $(3, 1, 2)$, or (p, q, r) .

The proof now breaks into various cases, depending upon the relative placement of the clues and the relative labeling.

We first consider the case where one of the following holds: three clues in the same box; three in the same row/column (which we illustrate below in Figure 3); two in the same box or in the same row/column and the third in the same row/column as one of the first two. (This case corresponds to the existence of at least two edges "between" clues that may be dropped.) In each instance, we apply an appropriate symmetry of G_7 . For each instance, there are at least two edges "between" clues that are not in the model and may be deleted from SUD_2 . We obtain a polynomial with 13 variables and of degree at most $40 - 2 = 38$. We apply Corollary 5 and we are done.

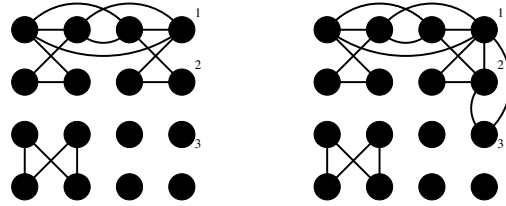
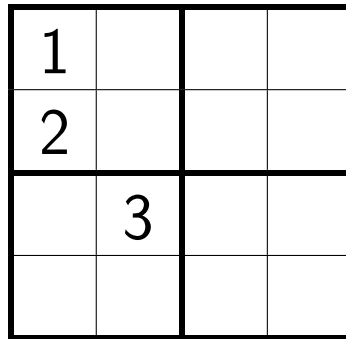


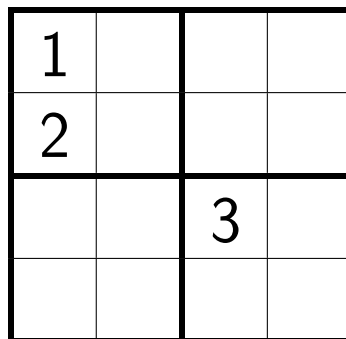
Figure 3: From left to right: G_7 with 3 vertices labeled in the fourth column; and with the three additional edges that may also be deleted from SUD_2 , for a total of 19 edges deleted

There are five cases with three distinct clues with at most one edge between them; up to isomorphism (and with an accompanying board) these are:

(1) two clues in the same row/column and box, with the third clue in the same chute as the first two;



(2) two clues in the same row/column and box, with the third clue in a different chute to the first two;



(3) two clues in the same row/column, different box;

1			
2			
		3	

(4) two clues in the same box, different row/column;

1			
	2		
		3	

and, (5) no two clues in the same row, column, or box.

			1
	2		
		3	

For Case (1) we uniquely complete the first column (as shown below, with new clues given in italicized font) and need not apply any model. With the double appearance of the 3-label in the lower left box, the polynomial f_G has a factor of 0 and so is identically 0. We apply Corollary 5 and reach the desired conclusion.

1			
2			
<i>4</i>	3		
<i>3</i>			

For Case (2) we uniquely complete the first column (as shown below, with new clues given in italicized font) and apply a 90-degree clockwise-rotation of model G_7 . We delete six edges in the first column and one edge in the third row, none of which are present in the model. With the double appearance of the 3-label, we delete one more edge, an edge that is incident with the cell $(3, 2)$. We now have 11 variables and $40 - 8 = 32$ edges. We apply Corollary 5 and reach the desired conclusion.

1			
2			
<i>4</i>		3	
<i>3</i>			

For Case (3) we uniquely complete the first column and the $(3, 2)$ -entry (as shown below, with new clues given in italicized font) and apply a 90-degree clockwise-rotation of model G_7 . We delete nine edges between the known clues, none of which are present in the model. With the triple appearance of the 3-label, we delete two more edges (neither of which is present in the considered symmetry of G_7) – one of the pair incident with cell $(4, 2)$ and one of the pair incident with cell $(2, 3)$. We now have 10 variables and $40 - 9 - 2 = 29$ edges. We apply Corollary 5 and reach the desired conclusion.

1			
<i>3</i>			
2	<i>3</i>		
<i>4</i>		3	

For Case (4) we show two distinct completions.

1	3	4	2
4	2	1	3
2	4	3	1
3	1	2	4

1	3	4	2
4	2	1	3
2	1	3	4
3	4	2	1

For Case (5) we uniquely complete the second row, third column, and upper-right box uniquely (as shown below, with new clues given in italicized font) and apply a 180-degree clockwise rotation of model G_2 . We delete 16 edges between the known clues, none of which are present in the model. We now have 8 variables and $40 - 16 = 24$ edges. As the degree of f is 24, we know that there is no greater element than the 8-length exponent vector of the form $(3, 3, \dots, 3)$ in $\text{Supp}(f_G)$. To show that the coefficient of the monomial that corresponds to this exponent vector is 0, we give a visual proof via Figure 4 to help illustrate the application of Lemma 6.

		<i>2</i>	<i>1</i>
<i>1</i>	<i>2</i>	<i>4</i>	<i>3</i>
		<i>3</i>	
		<i>1</i>	

As we seek the monomial corresponding to the 8-length exponent vector of the form $(3, 3, \dots, 3)$, we are interested in finding all orientations that admit 3 outgoing arcs from each of the vertices that correspond to an x -variable. Due to the degree of the polynomial, each of the c -labeled vertices must have all arcs as incoming. Those x -variable vertices with degree 3 obviously must have all arcs directed outwards. These two facts force the direction of most of the edges. The remaining edges that are not forced are highlighted in the right-most graph of Figure 4. As seen, these edges form a 5-cycle through the vertices associated to x_a -variables. Each of these vertices has an out-degree of two as given by the ‘forcing’. Thus, each requires one more outgoing arc, and so the other must be incoming. This allows for two possible orientations. As the parity of the length of the cycle is odd, one of these orientations is even and the other odd. Lemma 6 implies that the coefficient of the monomial we seek is 0. We now apply Corollary 4 and reach the desired conclusion.

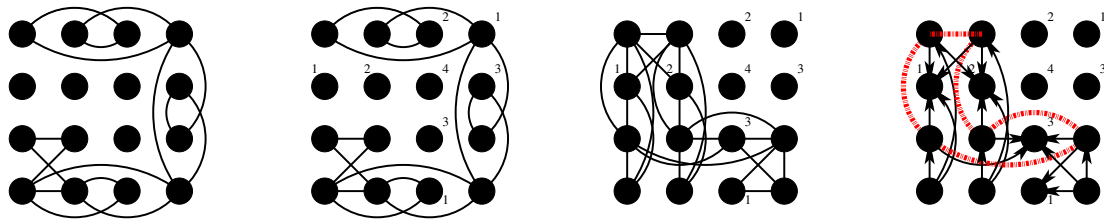


Figure 4: From left to right: G_2 rotated 180 degrees; and with the labeling; the built graph G ; and desired orientations of built graph

This completes the proof. \square

5 Concluding remarks

5.1 Remarks on the proof of Theorem 1

The reader should have noted that, although the conclusion of each of the corollaries that we used in the proof allows for two possibilities, we needn't distinguish between the two. With either possibility, the puzzle is not valid!

The reader should note that in the proof of Theorem 1 we sought to use Corollary 5 whenever we could. The reason for this was that we then didn't need to compute the coefficient of a leading monomial, which an application of Corollary 4 would have required. Our work with the polynomial method has confirmed for us what others have found – extracting coefficients of a (leading) monomial is challenging. So, this is one reason we were so happy to 'find' Corollary 5 – we simply needed to check that the degree of the polynomial was 'small enough', we didn't need to find any coefficients!

It seems to us that Case (4), which was handled by providing two distinct completions, could not have been approached via Corollary 5, but could have been handled using Corollary 4. However, doing so would have been a little bit more tedious than simply presenting the two completions. Further, our goal has been to demonstrate how the polynomial method can be used and we believe we have accomplished doing so.

5.2 Remarks on the Sudoku problem

We believe that the approach we used for the minimum number of clues Shidoku problem can also be used for the minimum number of clues Sudoku problem, though this is not without challenge.

One of the first obstacles has already been overcome – Demoen and de la Banda [4] have shown how to eliminate some of the 810 edges in SUD_3 . They have given various models that show which sets of redundant edges can be removed from the Sudoku graph while still supplying enough rules to know whether or not a completion satisfies all of the initial rules. The best known model allows one to consider an 81-vertex graph with 648 edges – others leave the number of edges in the high 600s.

If one were to try to prove Conjecture 7 via the method above, one would begin by filling in 16 cells and writing down a polynomial in at most 65 variables. The arithmetic condition called for by this application of the polynomial method is a polynomial of degree at most $8 \cdot 65 = 520$. The results of Demoen and de la Banda leave us short of this goal and so we must find many redundant edges by examining the clues given. This is particularly challenging, especially when one discovers that a redundant edge need not share an endpoint with a vertex whose corresponding cell has been labeled.

Of course one might try for a weaker lower bound than the conjectured one. Doing so has the advantage of not having to find as many redundant edges based upon the clues, but there is a trade-off as with fewer clues there will be less redundancy to find.

Finally, we point the reader to a question of Herzberg and Murty [6] who asked, is it true that the answer for the minimum number of clues $\text{Sudoku}(n)$ problem is $o(n^4)$?

6 Acknowledgments

The authors wish to thank Pete Schumer who read and commented on an early draft of our work. We also wish to thank Jeff Dinitz for his encouragement.

References

- [1] N. Alon. Combinatorial Nullstellensatz. Recent trends in combinatorics (Mátraháza, 1995). *Combin. Probab. Comput.*, **8** (1999), no. 1–2, 7–29.
- [2] N. Alon and M. Tarsi. Colorings and orientations of graphs. *Combinatorica* **12** (1992), no. 2, 125–134.
- [3] J. Cooper and A. Kirkpatrick. Critical sets for Sudoku and general graphs. *Discrete Math.*, **315–316** (2014), 112–119.
- [4] B. Demoen and M.G. de la Banda. Redundant Sudoku rules. *Theory Pract. Log. Program.*, to appear.
- [5] B. Demoen and M.G. de la Banda. Inequalities in 2×2 –Sudoku and beyond. CW Reports, volume CW636, 19 pages, Department of Computer Science, KU Leuven, Leuven, Belgium, February 2013.
- [6] A. Herzberg and M.R. Murty. Sudoku squares and chromatic polynomials. *Notices Amer. Math. Soc.* **54** (2007), no. 6, 708–717.
- [7] M. Lasoń. A generalization of combinatorial Nullstellensatz. *Electron. J. Combin.* **17** (2010), no. 1, Note 32, 6pp.
- [8] G. McGuire, B. Tugemann, G. Civario. There is no 16-clue Sudoku: Solving the Sudoku minimum number of clues problem. [arXiv:1201.0749](https://arxiv.org/abs/1201.0749).
- [9] U. Schauz. Algebraically solvable problems: describing polynomials as equivalent to explicit solutions. *Electron. J. Combin.* **15** (2008), no. 1, Research Paper 10, 35 pp.
- [10] L. Taalman. Taking Sudoku seriously. *Math Horizons* **15** (2007), no. 1, 5–9.
- [11] T. Tao and V. Vu. Additive combinatorics. Cambridge Studies in Advanced Mathematics, 105. Cambridge University Press, Cambridge, 2010.