

Andries E. Brouwer

Technische Universiteit Eindhoven
Postbus 513, 5600 MB Eindhoven
aeb@win.tue.nl

Recreational Mathematics

Sudoku Puzzles and How to Solve Them

The Sudoku, a Latin square completion puzzle, has already conquered The Netherlands some time ago, surpassing in popularity the so-called griddlers ('Japanese puzzles' in Dutch). The puzzle was invented in Indianapolis in 1979 by Howard Garns. Garns contributed his puzzles to Dell Magazines, under the name of 'Number Place'. Interest in Sudoku surged due to a revival in Japan in 1986, when puzzle publisher Nikoli rediscovered the game. The name 'Sudoku' is the Japanese abbreviation of a longer phrase, "Suuji wa dokushin ni kagiru", meaning 'the digits must remain single' (source: Wikipedia). Although solving the puzzle is a trivial task for a computer, the puzzle has many interesting mathematical properties. There are even some open problems attached to it, as Andries Brouwer, professor in graph theory at the Technische Universiteit Eindhoven, shows in this article.

A Sudoku puzzle (of 'classical type') consists of a 9-by-9 matrix partitioned into nine 3-by-3 sub-matrices ('boxes'). Some of the entries are given, and the goal of the puzzle is to find the remaining entries, under the condition that the nine rows, the nine columns, and the nine boxes all contain a permutation of the symbols of some given alphabet of size 9, usually the digits 1–9, or the letters A–I.

Some mathematicians will claim that since this is a finite problem, it is trivial. The time needed to solve a Sudoku puzzle is $O(1)$ – indeed, one can always try the 9^{81} possible ways of filling the grid. But one can still ask for efficient ways of finding a solution. Or, if one already knows the solution, one can ask for a sequence of logical arguments one can use to convince someone else of the fact that this really is the unique solution.

Backtrack and elegance

It is very easy to write an efficient computer solver. Straightforward backtrack search suffices, and Knuth's 'dancing links' formulation of the backtrack search for an exact covering problem takes a few microseconds per puzzle on common hardware today.

For a human solver, backtracking is the last resort. If all attempts at further progress fail, one can always select an open square, preferably with only a few possibilities, and try these possibilities one by one—maybe using pencil and eraser, or maybe copying the partially filled diagram to several auxiliary sheets of paper and trying each possibility on a separate sheet of paper.

For very difficult Sudoku puzzles, this is the fastest way to solve them, both for computers and for humans.

However, one solves puzzles not because the answer is needed, but for fun, in order to exercise one's capabilities in logical reasoning. And solving by backtracking is dull, boring, mindless; it requires no thinking, is no fun at all, and should be left to computers.

So, Backtrack, or Trial & Error, is taboo. And if it cannot be avoided one prefers some limited form. Maybe whatever can be done entirely in one's head.

Grading

Most Sudoku puzzles one meets are computer-produced, and it is necessary to have a reasonable estimate of the difficulty of these puzzles. To this end one needs computer solvers that mimic human solvers. Thus, one would also implement the solving steps described below in a Sudoku solving program, not in order to find the solution as quickly as possible, but in order to judge the difficulty of the puzzle, or in order to be able to give hints to a human player. Such AI-type Sudoku solving programs tend to be a thousand to a million times slower than straightforward backtrack.

Generating

Given the backtrack solver, generating Sudoku puzzles is easy: start with an empty grid, and each time the backtrack solver says

that the solution is not unique throw in one more digit. (If at that point there is no solution anymore, try a different digit in the same place.) To generate a puzzle in this way requires maybe thirty calls to the backtrack solver, less than a millisecond. One can polish the puzzle a little by checking that none of the givens is superfluous.

Afterwards one feeds the puzzles that were generated to a grader. Maybe half will turn out to be very easy, and most will be rather easy ('humanly solvable'). It is very difficult to generate very difficult puzzles, puzzles that are too difficult even for very experienced humans.

Solving

Below we sketch a possible approach for a human solver. The goal is to be efficient. In particular, the boring and time-consuming action of writing all possibilities in every empty square is postponed as much as possible. On the other hand, some form of markup helps.

Baby steps

When eight digits in some row or column or box are known, one can find the last missing digit.

Singles

When there is only one place for a given digit in a given row or column or box, write it there. If there is only one digit that can go in a given square, write it there.

	1							
						3		
				1				
								1

Singles: the 1 in the center right box must be placed in the yellow square.

5						2	3	
			6					
								6
	6							

Singles: the 6 in the top row must be put be in the yellow square.

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

Exercise (i) Solve this puzzle using baby steps only. (ii) Show that if a puzzle can be solved using baby steps only, it has at most 21 open squares.

Baby steps are particularly easy cases of singles. Checking for singles requires 324 steps. Knuth's 'dancing links' backtracker will take 324 steps if and only if the puzzle can be solved by singles only. It is unknown how

many open squares a puzzle can have and be solvable by singles only. There are examples with 17 givens. It is unknown whether any Sudoku puzzles exist with 16 givens and a unique solution. (There exists an example with 16 givens and exactly two solutions.)

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

Exercise Solve the above diagram, using singles only.



Illustrator: Ryu Tajiri

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

The subset principle: one may remove a candidate for a cell outside S if its presence would force $\sum n_d < |S|$.

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

at most n_d . If $\sum n_d = |S|$, then the situation is tight: each digit d must occur precisely n_d times in S . In this case we can eliminate a digit d from the candidates of any cell C such that the presence of a d in C would force the number of d 's in S to be less than n_d .

Matchings (3): For digit 7, the only possibilities in columns 2, 5, 6 occur in rows 1, 6, 7. Therefore, digit 7 cannot occur outside columns 2, 5, 6 in these rows.

Exercise Complete this Sudoku.

More generally, one may remove a candidate for a cell outside S if its presence would force $\sum n_d < |S|$.

and columns interchanged. This argument is called *X-wing* for $n = 2$, *Swordfish* for $n = 3$, *Jellyfish* for $n = 4$.

Hinge

The previous subsection used that each cell contains at least one digit. Conversely, each digit is in at least one cell in any given row, column or box.

The subset principle

Let S be a subset of the set of cells of a partially filled Sudoku diagram, and let for each digit d the number of occurrences of d in S be

Forcing Chains

Consider proposition $(i, j)d$: 'cell (i, j) has value d ' and proposition $(i, j)!d$: 'cell (i, j) has a value different from d '. By observing the grid one finds implications among such propositions.

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

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

The subset principle: In the five colored squares, the five digits 2,3,4,5,9 each occur at most once (since all occurrences of 3,4,5 are in a single box and all occurrences of 2,5,9 in a single column). Since the situation is tight, digits 3,4,5 do not occur elsewhere in this box, and digits 2,5,9 do not occur elsewhere in this column.

lie in the same row, column or box, so that they must contain different digits. This gives the first type (denoted by I): If (i, j) is adjacent to (k, l) then $(i, j)d > (k, l)!d$, where $>$ denotes implication.

4		2	8			1	6	
8	b		2	c	9	5	4	
			4		2	8		
9	4	1		2		8	3	5
2			3	8	4	6	9	1
6	8	3		9		4	7	2
7	a	8				9	2	4
3	2	4	9		8	7		6
			4	7	2	3		8

Hinge: if cell a has a 1, then cell b does not have a 1, and then the 1 in row 2 must be in cell c . But then the yellow area cannot contain a 1, which is impossible. (So, cell a has a 5.)

In case (i, j) is adjacent to (k, l) and (k, l) only has the two possibilities d and e , then $(i, j)d > (k, l)e$. This is the second type of implication (denoted by II).

Finally, if some row, column or box has only two possible positions (i, j) and (k, l) for some digit d , then $(i, j)!d > (k, l)d$. This is the third type (denoted by III).

Let us now consider chains of implications. If $(i, j)d > \dots > (i, j)!d$ then $(i, j)!d$. For chains that only involve a single digit, and where the implication types alternate between I and III , one often uses a simplified notation like

$1 : (8, 2) - (6, 2) = (6, 9) - (4, 7) = (8, 7) - (8, 2)$.

Here ' $-$ ' denotes that at most one is true and ' $=$ ' that at least one is true.

	56					68		
		56						
		26				68		

Finding useful chains may be nontrivial, and there are various techniques such as 'colouring' that help.

Uniqueness

A properly formulated Sudoku puzzle has a unique solution. One can assume that a given puzzle actually is properly formulated and use that in the reasoning, to exclude branches that would not lead to a unique solution.

		124		124	29			
			14		124			

	56					68		
		56						
		2				68		

			14		14			
			14		14			

A chain of implications. If $(i, j)d > \dots > (i, j)!d$ then $(i, j)!d$. For example, the sequence $(8, 2)6 > (8, 6)8 > (5, 6)6 > (5, 1)5 > (6, 2)6 > (8, 2)2$ is used to conclude that $(8, 2)2$.

For example, the above can be completed in at least two ways, violating the uniqueness assumption. This can be avoided: for instance in the diagram above at least one of the corners of the rectangle is 2.

Uniqueness: at least one of the corners of the rectangle must be 2.

2	8		46	5	1		46	3	9	7
136	5		136	36	7	9	8	4	2	
9	7		34	34	8	2	5		16	16
7	9	2		8	5		146	16	3	146
1346		146	136	7	2		146	9	5	8
5		8		46	9	3	7	2		146
	146									
146	3	7	9		46	8	2		16	5
8		146	9	2	46	5		146	7	3
	46	2	5	1	3	7		46	8	9

2	8		46	5	1		46	3	9	7
136	5		136	36	7	9	8	4	2	
9	7		34	34	8	2	5		16	16
7	9	2		8	5		146	16	3	146
1346		146	136	7	2		146	9	5	8
5		8		46	9	3	7	2		146
	146									
1	3	7	9		46	8	2		16	5
8		46	9	2	46	5		146	7	3
	46	2	5	1	3	7		46	8	9

Another chain of implications. (Check the given possibilities in red!) Here $(8, 2)1 > (6, 2)1 > (6, 9)1 > (4, 7)1 > (8, 7)1 > (8, 2)1$ is used to conclude $(8, 2)1$.

5	1	3	9			7	2	8
6	9	7	5	2	8	4	3	1
4	8	2	1	$\frac{3}{7}$		5		
3		5		1			8	
9		1	8			6		$\frac{3}{}$
2		8				1		$\frac{3}{}$
1	5		4	$\frac{8}{}$	$\frac{8}{}$			
7			6	$\frac{38}{}$	1	$\frac{238}{}$		
8	$\frac{23}{}$	$\frac{46}{}$						$\frac{46}{}$

Uniqueness: if (7, 7)3, then (7, 5)8 and (8, 7)8 so that (8, 5)3 and we have a forbidden rectangle with pattern 83-38. So, (7, 7)!3, which means that we have an X-wing: digit 3 in columns 2,7 can only be in rows 8,9 and does not occur elsewhere in these rows. In particular (9, 4)!3 so that (6, 4)3, and (8, 5)!3 so that (8, 5)8.

More generally the following theorem holds.

Theorem Suppose one writes some (more than 0) candidate numbers in some of the initially open cells of a given Sudoku diagram, 0 or 2 in each cell, such that each value occurs 0 or 2 times in any row, column, or box. Then

this Sudoku diagram has an even number of completions that agree with at least one of the candidates in each cell where candidates were given. In particular, if the Sudoku diagram has a unique solution, then that unique solution differs from both candidates in at least one cell.

Digit patterns and jigsaw puzzles

A more global approach was described by *Myth Jellies* (a pseudonym of a user of the Sudoku forum www.sudocue.net), the inventor of the *Pattern Overlay Method* (PDM), that is, jigsaw technique. Solve a puzzle until no further progress is made. Then, for each of the nine digits, write down all possible solution patterns for that digit. One hopes to find no more than a few dozen patterns in all. Now the actual solution has one pattern for each digit, where these 9 patterns partition the grid.

Regard each digit pattern ('jigsaw piece') as a boolean formula ('this pattern occurs in the solution'). Write down the formulas that express that for each of the nine digits exactly one pattern occurs, and that overlapping pieces cannot both be true. Solve the resulting system of propositional formulas.

This approach allows one to solve some otherwise unapproachable puzzles.

Remark This note is a condensed version of <http://homepages.cwi.nl/~aeb/games/sudoku>



Illustratie: Ryu Tajiri