Sudoku Solving Inference Engine

D.P.Playne

daniel.playne@gmail.com

Institute of Information and Mathematical Sciences, Massey University, Albany, North Shore 102-904, Auckland

Introduction

Sudoku is a popular game where the aim is to complete a 9x9 board of tiles which each contain a value in the range 1-9. Each 3x3 square, row and column should contain the values 1-9. Each Sudoku puzzle has only one logical solution; however it is possible for there to be multiple non-logical Solutions for a board. The Sudoku Solving Inference Engine (SuSIE) can solve Sudoku puzzles using logic rules rather than brute force techniques which means it always finds this correct logical solution. Given an input of an incomplete Sudoku board, SuSIE will logically solve the Sudoku and display the solution with each of the tiles filled in with its correct value.

Rule One

Rule One: We know a tile cannot have the same value as a different tile in the same 3x3 square, row or column. Thus if all other values have already been filled, the tile must contain the remaining value. Figure: 1 is an illustration of how Rule One works, the grey square can only be 2, because all the other numbers have already be taken by another tile in the same 3x3 square, row or column (red squares).

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

Figure 1: An example of Rule One.

Rule Two

Rule Two: Each 3x3 square, row and column must contain the values 1-9. So if there is only one tile within one of these constructs that could contain a certain value, it should be assigned that value. Figure: 2 shows Rule Two in action, the grey tile must be 5 because no other tile in that row (red squares) can possibly be 5.

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

Figure 2: An example of Rule Two.

Rule Three

Rule Three: If any 3x3 square, row or column had only two tiles that can have a certain value then no other tile in that 3x3 square, row or column can possibly have that value. The application of Rule Three can be seen in Figure: 3. The grey tile cannot be 2 because one of the two red tiles must be. This information can then be used to work out more details about other tiles.

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

Figure 3: An example of Rule Three.

Rule Four

Rule Four: If a square, row or column has only two tiles that can contain two certain values, then those two tiles must contain either one or the other of those values. The grey tiles must be either 8 or 5 because no other tile in that 3x3 square can be. This information can be used to discover more information about other tiles.

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

Figure 4: An example of Rule Four.

Results and Testing

SuSIE was testing with ten Sudoku puzzles for each difficulty level - Easy, Medium, Hard and Evil. The time taken for SuSIE to solve each puzzle was measured. SuSIE was able to successfully solve every puzzle tested and the results are shown in Table: 1.

Difficulty	Time Taken (seconds)
Easy	0.001
Medium	0.004
Hard	0.008
Evil	0.012

Table 1: Results.

Summary and Conclusions

SuSIE can successfully find the correct logical solution to very hard Sudoku puzzles using the discussed logical inference rules. The logic rule implementation of SuSIE allows it to find the solution to Sudoku puzzles significantly faster than any brute force search method. During testing SuSIE was able to solve every: easy, medium, hard and evil puzzle in a fraction of a second.

Logical inference rules such as these can be studied at Massey in papers such as **159.202** and techniques for implementing this program can be learned in: **159.101, 159.102, 159.201** and **159.234**.