

Sudoku Creation and Grading

Andrew C. Stuart
3rd February 2007

Introduction

The purpose of this paper is to explain how my sudoku puzzles are created and how they are graded, the two most common questions asked of Japanese puzzle compilers. A great deal of puzzle information is available on the Internet related to solving sudokus but you'll always see in various forums questions on production and grading, but few concrete replies. Anyone who has attempted both tasks soon realizes that it takes a huge amount of thought and programming and such a project is normally a commercial one so few will divulge the secrets. I can't give algorithms or recipes in this paper for that reason but I can show broadly how and why I have arrived at the system I have and I hope it will answer those people who have queried a particular puzzle. More generally I want to show that my grading spectrum is sophisticated and defensible overall.

The Standards

Three very important standards must be met when creating a good sudoku puzzle. The first and most obvious is that the puzzle should have one solution. In the early days of sudoku many people were convinced that a good number of puzzles published by Michael Mepham and myself had more than one solution and our mail bags were full of duplicate solutions. Michael had a £100 prize for the first person to find a double solution but it has never been claimed. We both wrote many replies showing that the person had transposed a clue incorrectly or missed a clue on their worksheet. We have always guaranteed a single solution and by 2007 this issue has almost died off. Sky famously carved a sudoku into a hillside and offered a £5,000 prize and didn't know that it had 1,905 solutions (see <http://www.sudoku.org.uk/blunder.htm>).

The second standard we meet is that there must be a logical solution to all the puzzles we publish. Most people don't like the idea that they have to guess, it seems to undermine the point of a logic puzzle and I have always agreed with this notion. There is a tricky aspect to this standard though, and it's a problem with language and words. Guessing is actually algorithmic and therefore deterministic which makes it logical in one sense. Turning the question around, is it possible to come up with an illogical strategy? I can't think of one. A mathematician will say that logical strategies are elegant or inelegant and that guessing is strictly inelegant – it's slow and you'll get a great deal of false paths. What the puzzle solver wants are methods that he or she can use which tell them a deduction is correct and they can place a number or remove a candidate with certainty. The puzzles that I produce have this aspect, but in the diabolical or extreme puzzles, the exact combination of 'elegant' logical strategies can be obscure to say the least. And applying strategies and finding new ones are the big attraction of sudoku puzzles.

There is a class of sudoku puzzle where no known logical solution exists – *as far as I know*. Other people may be able to solve it elegantly but no-one has yet proved that all sudoku puzzles can be solved logically without guessing. Currently about 2 in every thousand randomly produced puzzles I make are in this category and I publish some of

these on various web sites. They are important for puzzle aficionados as they are the real coal-face of strategy development. Unless specifically requested, these are not given to newspapers or puzzle book publishers.

The third standard is an aesthetic one. All the puzzles should have some kind of symmetry, normally around the center. Clues should be distributed in a pleasing pattern. It is the case with my puzzles that occasionally one or two extra numbers are removed to make the puzzle ever so slightly harder – but ensuring there is still one solution. Some therefore have a slight asymmetry. I also insist that the average number of clues hangs around 28 and is never more than 32. I have seen some bad sudoku puzzles with half the grid filled in – and labeled ‘hard’. It is possible to reduce the number of clues to 17 and still provide a spectrum of difficulty.

An interesting aside: A question on a forum asked whether it was essential that all the numbers be present in the initial clues. The answer is no, but at least eight of the nine numbers must be present. If there were only seven, for example the numbers 3 to 9, then all the 1s and all the 2s could be swapped around and you’d have a double solution. In a search of my library I discovered that 8% of puzzles had the characteristic of only having eight different numbers in the clues.

Creating a Sudoku Puzzle

To create a puzzle one has to know what the solution is first. That means creating a filled in grid of numbers such that each number 1 to 9 occupies each row, column and box just once. There are a number of ways to fill a sudoku board. Some of the information on the Internet refers to the rules of Chinese Chess and how the King is pinned in a three by three box. I experimented with that early on. Since I was working on logical solving strategies from the start I applied solving methods to board filling. One simply seeds nine cells with 1 to 9 randomly. Then one solves the board – which at this initial stage means removing the candidates that can be seen by the first nine numbers. By iteratively placing a solution to a cell randomly on the board (provided it is permitted by the list of remaining possibles on that cell) and then re-solving from that point, one can quickly fill the board. A randomly placed solution will often cause a solve failure so a note is kept of failed numbers and one backtracks. Nine times out of ten a filled board can be quickly created by placing and solving.

This method works if one has a large set of logical strategies to help remove candidates and prevent unworkable numbers being placed. I don’t include so-called ‘uniqueness’ strategies that depend on a single solution in the tests since the board won’t have a single solution until it is near-filled.

Given a filled board I then start subtracting numbers to make the puzzle. To maintain symmetry either two or four numbers that are diagonally opposite each other must be removed at the same time. For the first twenty or so subtractions four numbers can be removed. Beyond that the chance of four numbers leaving a single solution puzzle get slimmer so two at a time are subtracted. A low target of 20 clues is set and by 30 the remaining numbers are tested individually to see if it can be removed safely. After each subtraction the puzzle is tested to see if it retains a single solution. If this fails the numbers are restored and a different quad, pair or single subtraction is tried. A cut off

limits these tests and over a large run I get a set of puzzles where the number of clues is a bell-curve centered on 28.

Testing for a single solution after each subtraction is that hard part. I use a brute force method based on an algorithm published on the Internet. I've improved this, swapping arrays for bit-wise integers, for example and increasing the run speed. It is very important that this test really does try and find every possible solution to a partial grid in a very short time. You can try this test on my web site – look for the yellow button that says “Solution Count” on my web site: <http://www.scanraid.com/sudoku.htm>.

After a puzzle has been created with the minimum number of clues and it meets the first and third standard described above, it has to be graded.

Grading a Sudoku

Grading a sudoku is the greatest concern of the puzzle maker. If too many people disagree with your grades then you are clearly going to lose your audience. Everyone has different talents and different degrees of each talent so some puzzles will always be easier or harder for any two people. And some solvers might have a talent for pattern matching or guessing which short-cuts the logical method.

However, there are a number of useful pointers that help one to tackle the grading issue. Firstly, if, for example, there are ten squares which can be solved quite independently of each other then this puzzle is clearly easier – at that point, than a different puzzle where each solution relies on you getting the previous ones in a strict sequence. There is a metric of difficulty, therefore, to be gleaned purely by counting the opportunities to solve at all stages of the game. The eye and the mind can only cope with one opportunity and if it is seized, a number is placed, then the board needs to be re-checked to see the knock on effect. ‘Bottlenecks’ occur if there are few or only one chance to make a correct deduction and these make a more difficult puzzle.

The second metric is what kind of necessary strategy is required to identify an opportunity. A gentle puzzle, will for example, merely require the so-called ‘eye balling’ technique – simple looking for cells where only one number is possible. If you have to start jotting down notes to see where that a number might go then it is clearly a more difficult puzzle. Many strategies require you to know all the remaining candidates.

I have tested strategies against very large libraries of sudoku puzzles to find out which ones are frequent and which are obscure – last resort techniques if you will. I've also tested to see how many strategies will break through a bottleneck. Each strategy can be ranked by how hard they are to spot in real life, how often they are needed and how much damage they do (number of eliminations). There are many ways to gauge the usefulness of a given strategy. If scores and weights are carefully given to them we can get a tally from the whole solve route – a new metric of difficulty.

Combining the frequency of opportunity with the necessary strategies gives us a score. Exactly how these weights and factors are set and combined has been a matter of much work and some subjectivity. Is a type 3 Unique Rectangle easier to spot than a Type 1? How much should the weighting difference really be compared to their relative

usefulness? These are difficult questions and there are many of them and it requires a lot of statistical analysis.

Over a great number of puzzles a spectrum of difficulty is built up. Then it is a question of dividing that spectrum into grade bands. I currently have six bands. The pie graph in Figure 1 shows how any given set (normally in the tens of thousands) would be broken up into grades (1=easiest, 6=hardest). It reflects the notion that most randomly produced puzzles will be easy. There is a “long tail” of difficulty at the extreme end of the spectrum where rare sudoku puzzles will be extraordinarily difficult and have very high scores. Most puzzles will be clumped in the easier sextiles.




Figure 1

Some additional rules apply to grades as well. To be a Kids grade the puzzle cannot require any note taking – that is, simple eyeballing is all that is needed to solve the puzzle and there will be a high degree of ‘opportunity’. To be a gentle all solutions can be shown to only require ‘slice and dice’ – simply that in one row or column or box there is only one solution. Moderates may require simple strategies as Naked and Hidden Pairs and Triples. Higher grades will entail more sophisticated strategies. Some sudokus are discarded because of certain unwanted features. For example, if a sudoku is plain sailing but then requires just one very hard strategy so get through a bottleneck it might have a high score and a high grade. But it would not be a satisfying puzzle as most of the board could be filled in trivially yet it could be labeled difficult.

	1	2	3	4	5	6	7	8	9
A	4			7		1			3
B			5				2		
C		6				3		4	
D		7	8	6			9		
E					5				
F			4			2	1	8	
G		1		8				2	
H			2				3		
J	8			2		5			4

Figure 2

Figure 2 is an example diabolical sudoku. To give a rough idea of how such a puzzle is graded consider the solve route in Appendix A. Each solution line – if it solves a whole cell, has two numbers in square brackets. The first number is the number of solved cells at that point (so it starts at 28 and finishes with 81). The second number is the ‘game round’. Identical numbers here mean simultaneous solutions in that round. The lower the final number the more opportunities there were to solve. The calculations for this sudoku give me a score of 595 for the ‘opportunities’ to solve with an average of 2.3 solutions per round.

Score for Order: 595

Average Solving Rate: 2.34 per round

Points where candidates are removed:

Points	Cand	Sol:	points:	
Human Strategy	: 7,	7,	points:	0 7
Naked Singles	: 381,	27,	points:	0 27
Hidden Singles	: 34,	17,	points:	68 34
Naked Pairs	: 6,	1,	points:	12 2
Hidden Pairs	: 0,	0,	points:	0 0
Naked Triples	: 0,	0,	points:	0 0
Hidden Triples	: 0,	0,	points:	0 0
Naked Quads	: 0,	0,	points:	0 0
Hidden Quads	: 0,	0,	points:	0 0
Intersection Removal	: 1,	0,	points:	5 0
X-Wing	: 0,	0,	points:	0 0
Sword-Fish	: 0,	0,	points:	0 0
Simple Colouring	: 1,	1,	points:	20 20
Y-Wings	: 1,	0,	points:	25 0
Avoidable Rectangle	: 0,	0,	points:	0 0
Multivalued X-Wing	: 0,	0,	points:	0 0
Jelly-Fish	: 0,	0,	points:	0 0
X-Cycle	: 0,	0,	points:	0 0
Unique Rectangles	: 0,	0,	points:	0 0
XYZ Wing	: 0,	0,	points:	0 0
XY-Chain	: 1,	1,	points:	50 50
Aligned Pair Exclusion	: 0,	0,	points:	0 0
BUG	: 0,	0,	points:	0 0
Sue-de-Coq	: 0,	0,	points:	0 0
Grouped X-Cycle	: 0,	0,	points:	0 0
Multi Colouring	: 0,	0,	points:	0 0
Empty Rectangles	: 0,	0,	points:	0 0
Forcing Chains	: 0,	0,	points:	0 0
Finned X-Wing	: 0,	0,	points:	0 0
Finned Sword-Fish	: 0,	0,	points:	0 0
Guardian	: 0,	0,	points:	0 0
Almost Locked Sets	: 0,	0,	points:	0 0
Death Blossom	: 0,	0,	points:	0 0
Altern. Inference Chains	: 0,	0,	points:	0 0
AIC with ALS	: 0,	0,	points:	0 0
Pattern Overlay	: 0,	0,	points:	0 0
Bowman Bingo	: 0,	0,	points:	0 0
Nishio	: 0,	0,	points:	0 0

Final Score: $320 * (595 / 1000) = 190$

Scores are allotted to strategies used according to candidates removed and cell solved.

All of these strategies are covered in my book “**Logic of Sudoku**”. I have ordered them in what I perceive to be their rough order of complexity. This is subjective given any two strategies – especially when two strategies could perfectly well solve the same problem, but overall the complexity does increase as one goes down the list. This diabolical requires Simple Colouring, Y-Wings and XY-Chains – minimally, but I don’t claim it is the only way to solve the puzzle but it is near optimal. Provided the rules are applied consistently an overall grading pattern can emerge.

I have ignored guessing as a strategy. This is because it is important to have a bench mark and guessing might short-cut a problem or it might hopelessly confuse a potential solution. My suspicion is that many puzzles which are accused of being easier or harder than the published grade have skipped some logic steps and good or bad hunches have been used. This will effect the outcome of the perceived difficulty.

A Statistical Measurement of Grading

On the Daily Competition Sudoku on www.sudoku.org.uk - where we get 2000 to 3000 correct submissions each day, we’re been very pleased that submitters have given us their times for solving. This helps us calibrate the puzzles. The results of this can be seen here:

<http://www.sudoku.org.uk/solvetimes.htm> (which has been updated to include about 330 days of puzzles).




Figure 3

Submitters are allowed to choose from the following time bands to say how long it took to solve the puzzle:

<=5, <=10, <=15, <=20, <=30, <=40, <=50, <=60, <=120, >120 minutes

Note we also have a “Don’t know” since many people stop and start or don’t want to say and we can ignore such answers. Figure 3 shows the number of correct submissions for the 330 puzzles against each time period (1 to 10) and for each grade. We publish more moderates than others in a week so that line is higher.

Note that the daily competition uses only Gentle, Moderate, Tough and Diabolical so we have no statistics for Kids or Extremes.






Figure 4

If we plot the average time to solve each puzzle (remember there are about 2000-3000 to get an average from) we can form a histogram. The Figure to the left shows each grade coloured and the number of them for “minutes to solve” between 16 minutes and over 1 hour. Clearly the gentles are all being done in a short time and the diabolicals taking longer but over a larger time scale. Only the Tough’s seem to have a spread suggesting they vary in difficulty the most.

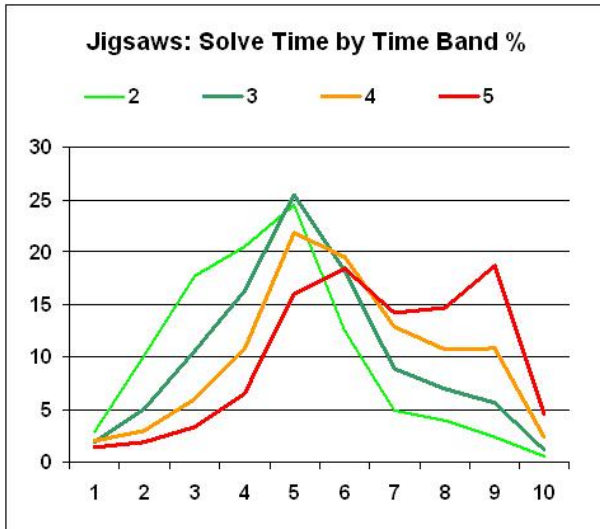


If we normalize the grades and plot the first graph against time periods we can see a shift in difficulty against time band.




Figures 5 and 6

The same information but plotted differently shows that about 10% of solvers are struggling with gentles and 10% are finding some diabolicals easy. This is expected if we consider the wide range of puzzle solving talent possessed by the public at large.



Jigsaw puzzles are derived from normal sudoku puzzles using the same techniques and we'd expect a similar spread of solve times. There are less submissions for these but more than enough to sample.

Jigsaw sudoku has an additional strategy not available to normal sudoku – the Law of Leftovers. It is weighted and factored into the game as part of the grading process. The programs used to make these puzzles in fact consider the normal sudoku to be a special, if boring, jigsaw shape so any changes to production or grading effect both variants equally.



Figures 6 and 7




Figure 8

Figure 8 shows the grade distribution for 6 x 6 sudokus. This interesting variant looks unassuming and if published it is normally aimed at children as an introduction to sudoku. Certainly, when making these 69% are trivial. But the “long tail” effect means that a few in a thousand are extremely difficult and require some of the most advanced techniques known. The 0% for the extremes belies the fact that several hundred extreme grade 6x6 were produced in a batch of 150,000.

Killer Sudoku

Killers are immensely fun and often very difficult puzzles. The same applies to their creation. To begin with a filled sudoku board must be created. If that is done then a set of cages are created and overlaid on the board. Cage creation is an art in its own right and my method is sufficiently efficient that I can create different random cage grids for every puzzle, never repeating one. Some cage grids are discarded – for example if there are too many 2-cages (pairs) or if there are too many single cages (which are useful for gentle and moderate Killers but not suitable for higher grades).

I adhere to the convention that cages cannot contain the same solution number. Imagine a dog-legged cage that spanned three boxes. It is possible for cells at both ends of the cage to be the same. If this occurs then the cage is discarded as well. Given a cage grid we automatically get the clues which are merely the sum of the solutions for each cage. Now, a Killer does not have any starting clues (unless it is an easy one) so the scoring has to be derived from all the normal sudoku rules and strategies plus the additional 'opportunities' and strategies specific to Killer cages. Specifically we are interested in cage combinations (such as a 2-cage with the value of 3 – the two cells must contain 1 and 2) and how this restricts the cell candidates. I distinguish between cages that have only one combination (such as a 4-cage with value 10 – it can only be 1/2/3/4) from others that have more possibilities that cells in the cage. A human solver will always identify those cages that have a restricted set that matches the number of cells. Other strategies that whittle down the candidates are "innies and outies" and cage splitting. These have been weighted and factored into the grading process.

Conclusion

I hope I've given a flavour of what goes into sudoku puzzle production and how much care goes into calibrating them and keeping them up to date. Since 2005 when I could solve only 80% of puzzles to the point now where I can solve 99.98% the puzzles have got steadily harder. Purely because more is understood about them and new strategies have opened doors to myself, as a creator, as well as the solver. Difficulty creep also keeps up with the pace of demand as regulars expect a greater challenge.

If you were skeptical I hope I've addressed your concerns. I'm always interested in solvers feedback and I can be emailed at andrew@scanraid.com (although I can't promise to reply to every email but I'll do my best).

For a demonstration of logical strategies I suggest starting with my own solver at <http://www.scanraid.com/sudoku.htm> which has links to strategy documentation. There is also a Jigsaw Solver as well. On that solver I have ordered the strategies into groups roughly suggesting the grade where they could be called upon. A great deal of interesting discussion can be read on the main sudoku forums as www.sudoku.org.uk and www.sudoku.com. Do feel free to contribute your ideas and experiences as well.

Appendix 1

Solve Route for Diabolical Example

[28, 1] SINGLE candidate 9 changed to SOLUTION at A3
[29, 1] SINGLE candidate 4 changed to SOLUTION at D6
Single occurrence of 4 found at E7 in row/column/box, 2 candidates removed
Single occurrence of 4 found at G5 in row, 4 candidates removed
Single occurrence of 4 found at H2 in column/box, 2 candidates removed
Single occurrence of 5 found at C4 in column/box, 1 candidates removed
Single occurrence of 5 found at F2 in column, 2 candidates removed
Single occurrence of 8 found at E6 in row/box, 2 candidates removed
Single occurrence of 8 found at H9 in row/box, 5 candidates removed
[30, 2] SINGLE candidate 5 changed to SOLUTION at C4
[31, 2] SINGLE candidate 8 changed to SOLUTION at E6
[32, 2] SINGLE candidate 4 changed to SOLUTION at E7
[33, 2] SINGLE candidate 5 changed to SOLUTION at F2
[34, 2] SINGLE candidate 4 changed to SOLUTION at G5
[35, 2] SINGLE candidate 4 changed to SOLUTION at H2
[36, 2] SINGLE candidate 8 changed to SOLUTION at H9
Single occurrence of 3 found at J5 in box, 4 candidates removed
Single occurrence of 4 found at B4 in row/column/box, 1 candidates removed
Single occurrence of 7 found at F5 in box, 2 candidates removed
[37, 3] SINGLE candidate 4 changed to SOLUTION at B4
[38, 3] SINGLE candidate 7 changed to SOLUTION at F5
[39, 3] SINGLE candidate 3 changed to SOLUTION at J5
[40, 4] SINGLE candidate 1 changed to SOLUTION at D5
[41, 4] SINGLE candidate 6 changed to SOLUTION at F9
[42, 4] SINGLE candidate 9 changed to SOLUTION at J2
Single occurrence of 1 found at H4 in column/box, 1 candidates removed
Single occurrence of 1 found at J8 in row, 2 candidates removed
[43, 5] SINGLE candidate 1 changed to SOLUTION at H4
[44, 5] SINGLE candidate 1 changed to SOLUTION at J8
NAKED PAIR (Box 4,D1/E2 2/3), removing 2/3 from E1
NAKED PAIR (Box 4,D1/E2 2/3), removing 3 from E3
NAKED PAIR (Box 4,D1/E2 2/3), removing 3 from F1
[45, 6] SINGLE candidate 9 changed to SOLUTION at F1
[46, 7] SINGLE candidate 3 changed to SOLUTION at F4
[47, 8] 9 is the only possible number in E4
Single occurrence of 3 found at G3 in column, 2 candidates removed
[48, 9] SINGLE candidate 3 changed to SOLUTION at G3
SINGLES CHAIN (Type 1): Removing 9 from B6
[49,10] SINGLE candidate 6 changed to SOLUTION at B6
Single occurrence of 6 found at H5 in column/box, 1 candidates removed
[50,11] SINGLE candidate 6 changed to SOLUTION at H5
Single occurrence of 6 found at A8 in column, 1 candidates removed
[51,12] SINGLE candidate 6 changed to SOLUTION at A8
Single occurrence of 5 found at A7 in row/box, 1 candidates removed
Single occurrence of 8 found at C7 in column/box, 1 candidates removed
[52,13] SINGLE candidate 5 changed to SOLUTION at A7
[53,13] SINGLE candidate 8 changed to SOLUTION at C7
NAKED PAIR (Box 9,G7/J7 6/7), removing 7 from G9
NAKED PAIR (Box 9,G7/J7 6/7), removing 7 from H8
Y-WING 7 taken off B1 - using B8 H8 H1
POINTING PAIR: Between Box=1 and Row=3: 7 taken off C9
XY CHAIN length=4, 3 taken off E2, chain ends: B2 and E8
[54,14] SINGLE candidate 2 changed to SOLUTION at E2
[55,15] SINGLE candidate 8 changed to SOLUTION at A2
[56,15] SINGLE candidate 3 changed to SOLUTION at D1
[57,15] SINGLE candidate 7 changed to SOLUTION at E9
[58,16] SINGLE candidate 2 changed to SOLUTION at A5
[59,16] SINGLE candidate 1 changed to SOLUTION at B1

[60,16] SINGLE candidate 3 changed to SOLUTION at B2
[61,16] SINGLE candidate 5 changed to SOLUTION at D8
[62,16] SINGLE candidate 3 changed to SOLUTION at E8
[63,17] 2 is the only possible number in D9
[64,18] SINGLE candidate 9 changed to SOLUTION at B9
[65,18] SINGLE candidate 7 changed to SOLUTION at C3
[66,18] SINGLE candidate 9 changed to SOLUTION at C5
[67,18] SINGLE candidate 6 changed to SOLUTION at E1
[68,18] SINGLE candidate 9 changed to SOLUTION at H8
[69,19] 2 is the only possible number in C1
[70,19] 8 is the only possible number in B5
[71,19] 1 is the only possible number in E3
[72,20] SINGLE candidate 7 changed to SOLUTION at B8
[73,20] SINGLE candidate 1 changed to SOLUTION at C9
[74,20] SINGLE candidate 5 changed to SOLUTION at G9
[75,20] SINGLE candidate 7 changed to SOLUTION at H6
[76,20] SINGLE candidate 6 changed to SOLUTION at J3
[77,21] 9 is the only possible number in G6
[78,22] SINGLE candidate 7 changed to SOLUTION at G1
[79,22] SINGLE candidate 5 changed to SOLUTION at H1
[80,22] SINGLE candidate 7 changed to SOLUTION at J7
[81,23] 6 is the only possible number in G7