# The Nature of Puzzles

**Cameron Browne,** QUT

*This paper explores the underlying nature of puzzles, and how they relate to games. The discussion focuses on pure deduction puzzles, but with reference to other types of puzzles where appropriate, with examples to support the concepts put forward. These include the notion of puzzles as two-player games between the setter and the solver, the addictiveness of puzzles, and ways in which the setter can exert authorial control, to make challenges more interesting and engaging for the solver.*

## 1 Introduction

PUZZLES come in many forms; there are word puzzles, jigsaw puzzles, logic puzzles, dexterity puzzles, physical puzzles, physics based puzzles, to name just a few. While most readers will have an understanding of what the term 'puzzle' means to them, the genre as a whole has so far defied exact definition, despite many attempts to do so. But perhaps a precise definition is not all that useful – or even possible – given the variety of puzzles that exist. In this paper, I will look instead at the underlying nature of puzzles rather than attempting to provide yet another definition.

The central thesis of this paper is that most puzzles are games played between the setter and the solver, and that their inherent nature allows sufficient authorial control for the setter to impart their personality upon a well designed challenge, in order to challenge, tease and engage the solver. Several examples are presented in support of this argument, which are mostly taken from actual examples of pure deduction puzzles known as *Japanese logic puzzles* [1]. These are characterised by having simple rules, a single (deducible) solution, and no language-dependent content. They are not only my favourite type of puzzle, but also illustrate the principles being discussed as clearly and simply as possible.

To avoid confusion in the following discussion, the term *puzzle* will refer to the actual puzzle game itself, while each instance of a puzzle game presented to the solver will be called a *challenge*.

## 2 Puzzles As Games

A *puzzle* can be defined simply as a task that is *fun and has a right answer* [2], or more precisely: *a question which challenges people to solve, requires their deduction based on its rules to win, and doesn't depend on chance or other people's action.* [3]

Schuh presents a classification scheme for puzzles, and observes that puzzles *can be solved by pure reasoning alone*, must *have a complete analysis* and that *you are your own opponent, in the end* [4]. While the first two observations are true in most cases and they agree with most other definitions of puzzles, I take issue with the third observation that puzzles are a solitary pursuit undertaken without an opponent.

In his classic paper 'Defining the Abstract' (republished in this issue [5]), Thompson makes the astute observation that two-player abstract games may be described as a series of puzzles that the players present to each other. Conversely, I believe that a puzzle may be described as a two-player game played between the *setter* and the *solver*. The task of the setter is to produce a challenge that engages and entertains the solver, while the task of the solver is to avoid the traps laid by the setter to complete the challenge.

It is worth emphasising that, unlike a player in a traditional adversarial game, the setter is not trying to 'win' against the solver. They are instead trying to provide the most entertaining playing experience, in a role not dissimilar to that of the gamemaster in a role playing game. In the language of game theory, this is not a *zero-sum game*, as both players can win.

Puzzles are indeed solitary pursuits in a strictly mathematical sense, as the solver is the only agent making actions towards a solution. However, from a strategic or adversarial viewpoint, these actions are directed by the information encoded in the challenge by the setter, which is revealed as the challenge unfolds. The solver may not be an active player during the solution of a challenge, but participates in absentia by influencing the solver's decisions and actions.

A well designed challenge will include traps and deceptions posed by the setter, which the solver must detect and avoid. In order to see how this works, let's first look at the concepts of dependency and authorial control in puzzle design.

Browne, C., 'The Nature of Puzzles', *Game & Puzzle Design*, vol. 1, no. 1, 2015, pp. 23–34. © 2015

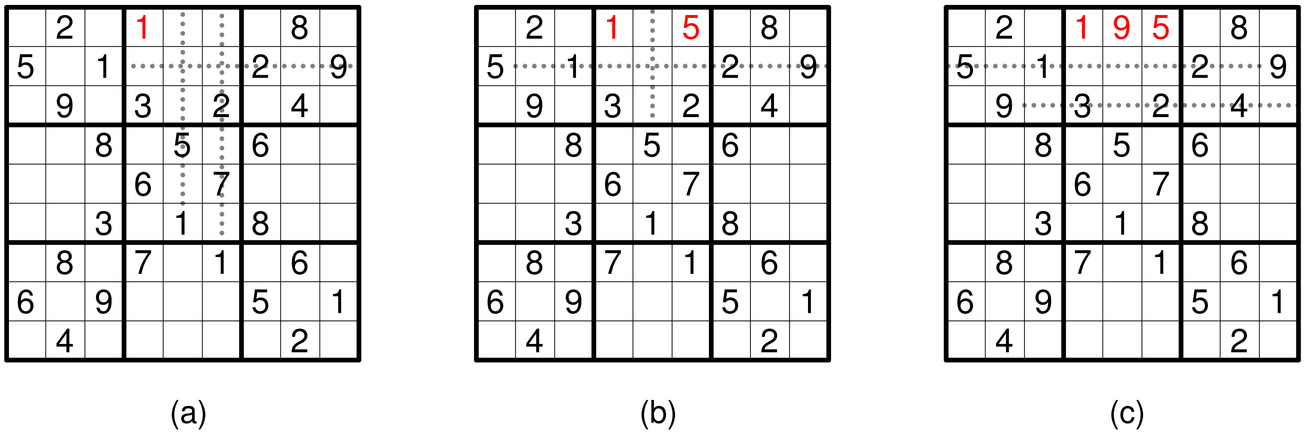


**Figure 1.** Dependent Sudoku hints progressively reveal enough information to make further progress.

### 2.1 Dependency

*Dependency* in this context refers to the degree to which the steps required to solve a given challenge are dependent on prior steps. A challenge in which progress can immediately be made at many places on the board shows low dependency, while a challenge that only exposes enough information for the solver to make progress at one particular point shows high dependency. The solver exploits that piece of information... which reveals further information... which reveals further information... until the solution is reached.

Figure 1 shows a simple example of this process in action, based on Sudoku challenge #31 from [6] (I assume that readers are familiar with the rules of Sudoku). The 1-hints provide enough information to instantiate another 1 on the top row (a). This additional information allows a 5 to be instantiated in the same row (b), which in turn provides enough information to allow a 9 to be instantiated next to it (c).

The required information is meted out in installments, in a self-perpetuating manner such that each action reveals further information to be acted upon. I have heard this process described as the setter leaving a trail of (informational) bread crumbs to follow. However, I prefer to think of the situation as a tapestry with a loose thread or two, in which the majority of the position is impenetrable except for certain weak points, whose exercise unravels further weak points to follow.

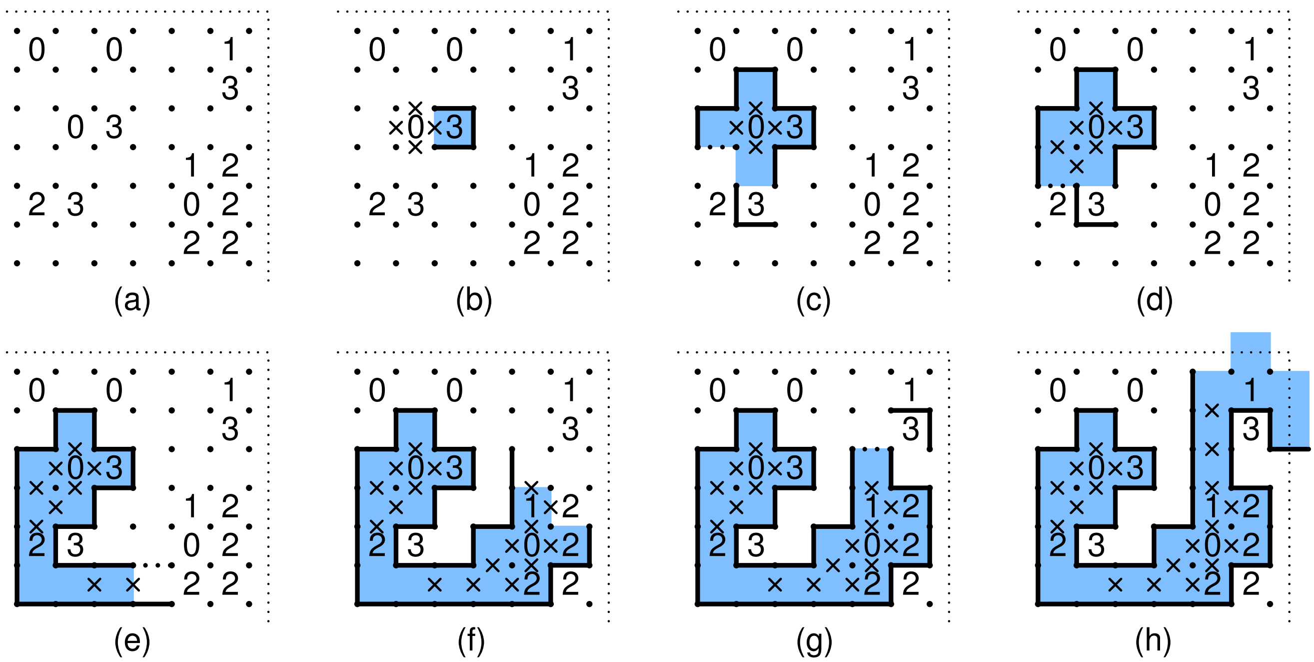


**Figure 2.** An interesting section of a Slitherlink challenge.

This Sudoku example only provides a superficial instance of this process, as it is an easy challenge with several loose threads to follow. For example, the 8-hints immediately dictate that the central cell must take the number 8. Figure 2 shows a much more striking example, with Slitherlink challenge #80 from [7].

Slitherlink is a pure deduction puzzle in which a simple closed path must be traced through orthogonal vertices of a square grid, to visit the number of sides indicated on each numbered *hint* cell [1]. Each adjacent pair of vertices therefore constitutes a *move* whose value can either be an edge (| or –) or no edge (×).

Figure 2 (a) shows the lower left corner of the initial challenge, and (b) shows three edges that must exist where a 3-hint meets a 0-hint. The path thus initiated then bounces off the 0-hints that it encounters (since the path can never visit the side of a zero hint) and moves along the wall to give the position shown in Figure 2 (c). A deduction is required at this point; the dotted move can not be an edge as that would cause the path to close prematurely in a cross shape, so it must be no edge × (d). This allows further progress until a deduction is required at position (e), which allows further progress (f) until another deduction is required at position (g), which leads to the completion of the section, as shown in Figure 2 (h).

These examples demonstrate how puzzles can hide their own solutions in plain sight, only revealing required information as needed in a self-perpetuating way. Each subproblem requires a solution that provides the next subproblem, and so on. In a well designed puzzle, the solver can almost feel the hand of the setter drip-feeding them information and leading them along by the nose along certain avenues to solution.

Pelánek [8] describes the use of dependency as a metric for automatically measuring the difficulty of given Sudoku challenges, based on whether the hints provide enough information to solve the challenge in parallel (i.e. multiple loose threads to follow) or in series (i.e. narrow chain of dependent loose threads). I believe that dependency is a fundamental property of well designed puzzles that runs deeper than just affecting difficulty, as it allows the setter to exert authorial control over the challenges they construct.

## 2.2 Authorial Control

*Authorial control* refers to the degree to which the setter can influence the solver's progress through a given challenge, and manipulate their move choices in absentia. This is the property that makes the setter a second player, in opposition to the solver.

Consider the Killer Sudoku example shown in Figure 3 [9]. Killer Sudoku is played according to the rules of Sudoku, except that no hints are provided initially apart from shaded subregions whose component digits must sum to the value shown on each.

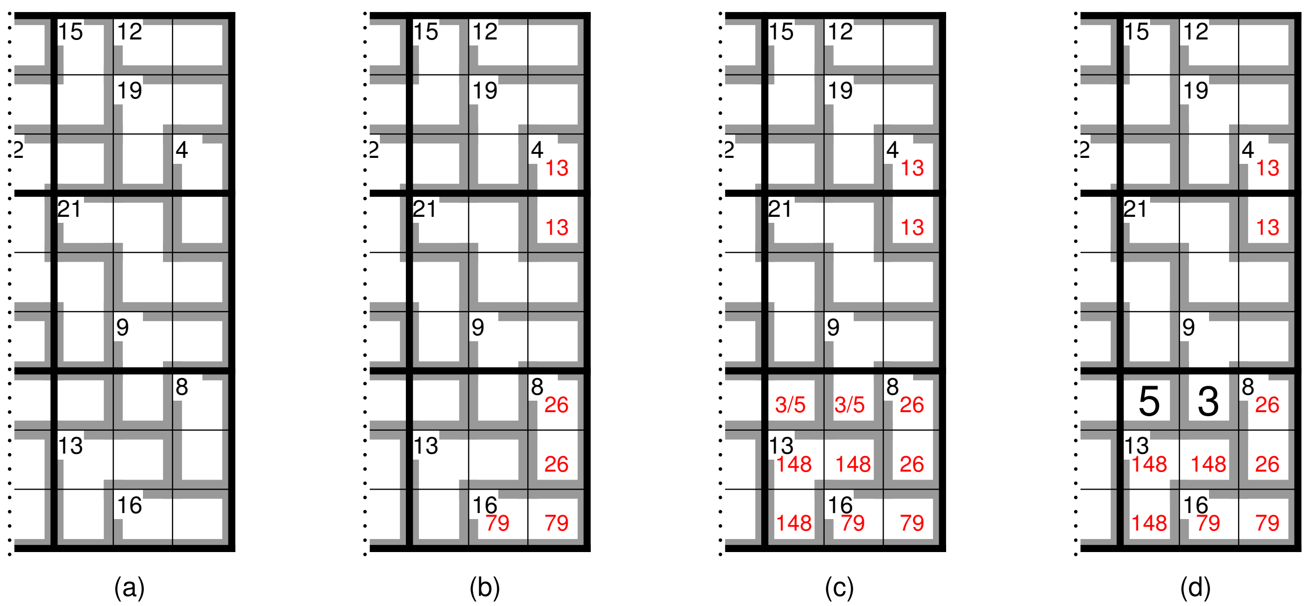


**Figure 3.** An efficient Killer Sudoku sequence that suggests authorial control.

Starting with the rightmost column in (a), the 16-region can only contain {7, 9} and the 4-region can only contain {1, 3}. The latter implies that the 8-region must contain {2, 6}, as shown (b). The only combination that satisfies the 13-region is then {1, 4, 8}, hence the last two remaining cells of the lower right 3×3 subgrid must have values 3 and 5 (c). The 5 cannot occur in the lower cell of the 9-region, since the only possible completion of this region {1, 3} would conflict with the vertical {1, 3} just above it, hence this cell must resolve to 3 and its neighbour to 5 (d).

The lower right 3×3 subgrid of this example resolves itself neatly and efficiently, using a minimum amount of information that self-referentially builds upon information released by prior steps. This pattern is unlikely to have occurred by chance, and the solver has the strong sense of an intelligent hand behind its design.[1]

Expert Sudoku solvers can generally tell whether a given challenge is handcrafted by a human designer or generated by a computer algorithm. Nobuhiko Kanamoto, Chief Editor for Japanese publisher Nikoli, observes that:

> *Computer-generated Sudoku puzzles are lacking a vital ingredient that makes puzzles enjoyable – the sense of communication between solver and author.* [10]

Nikoli have a policy of only publishing handcrafted challenges for their popular line of Japanese logic puzzles, and are sceptical of computer-generated content due to its potential to flood the market with inferior mass product. Challenges may be submitted by amateur fans or experienced designers, but all are hand tested before being approved for publication [11, p. 2].

This communication between setter and solver can only occur if the setter exercises a strong sense of authorial control in their design. For example, consider the computer-generated 6×6 Slitherlink challenge shown in Figure 4 (a).[2] Figure 4 (b) shows obvious simplifications that an experienced player would immediately spot and complete, while (c) shows the number of obvious simplifications arising from each hint and (d) shows these natural directions of progress for this challenge. This example has multiple starting points and no focused solution path.

Compare this with the handcrafted[3] 6×6 Slitherlink challenge shown in Figure 5 (a). This example has only one obvious starting simplification (at the 2 between the two 0s), but it triggers a chain reaction of 84 further simplifications that lead to a complete solution (b) along a few strongly defined directions of progress (c) and (d).

The first challenge may be superficially interesting, as its hints are rotationally symmetrical and it is the more difficult of the two. However, it lacks any underlying strategic structure, and the deductions leading to its solution are homogeneously spread across the board.

The second challenge, on the other hand, has a highly structured solution that unfolds elegantly with each simplification perpetuating the next. It is not symmetrical, nor as difficult to solve, but exhibits a strongly focused sense of authorial control; the solver can feel the hand of the setter and appreciate the craft of the design. This sense of structure tends to be missing from computer-generated designs.

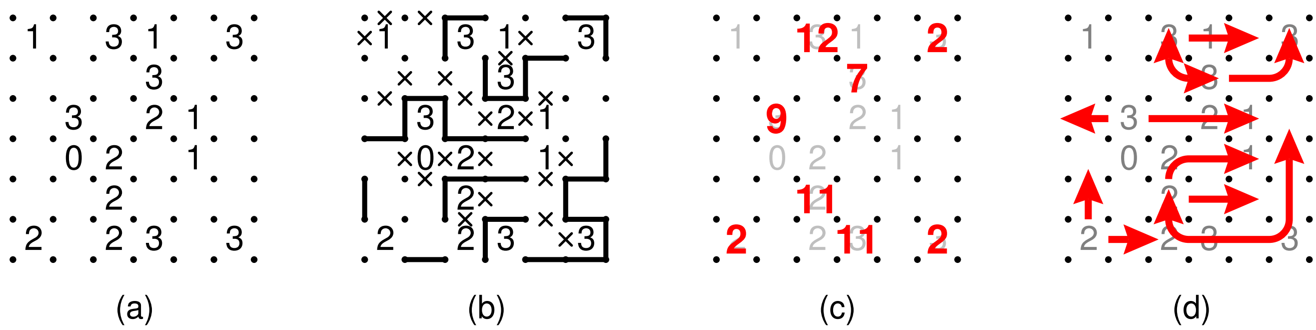

**Figure 4.** Computer-generated 6×6 Slitherlink challenge, showing simplifications and paths to solution.

[1] I would use the term 'intelligent design' if it had not been appropriated for another use recently.
[2] From: http://www.conceptispuzzles.com/ja/index.aspx?uri=puzzle/slitherlink
[3] Handcrafted by the author to illustrate this particular point.

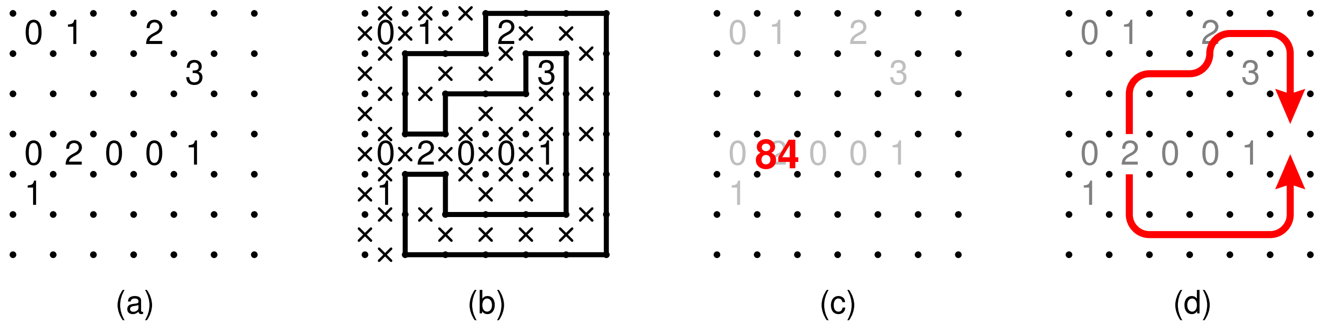


**Figure 5.** Handcrafted 6×6 Slitherlink challenge, showing simplifications and paths to solution.

### 2.3 Addictiveness

These mechanisms of dependency and authorial control could go some way to explaining why many players find solving puzzles so addictive. Drip-feeding subproblems to the solver in this manner makes challenges engaging and addictive, as the satisfaction felt at solving each subproblem is a reward that spurs the player on to solve the next, which itself creates more subproblems to be solved.

Stafford explains this effect in terms of a psychological phenomenon known as the Ziegarnik Effect, which refers to the human brain's tendency to latch onto unsolved problems until they are resolved, with respect to the video puzzle game Tetris [12]. Successful video puzzle games such as Tetris and 2048[4] typically 'hook' the player with such cycles of challenge and reward, as they are presented with continuous sequences of interesting subproblems to solve, each of which feeds the next, until the solution is achieved. This effect may also be described in terms of Gestalt psychology, as the brain's natural tendency to mentally complete incomplete patterns [13].

In both of these games, Tetris and 2048, pieces are added to the board in a nondeterministic manner, and the players' immediate subproblem is where to place those pieces to best effect, given the limited movement options available. These games also tap into our natural betting and risk assessment instincts – *what happens if I put that piece here? or there?* – which builds a sense of anticipation to see whether the next piece will fit the current plan. This gives players a double incentive to continue playing; the satisfaction of completing the immediate subproblem and the revelation of the next piece of hidden information. Note that the same addictive principle is relevant, even though these video puzzle games do not converge to a 'solution' as such.

[4]http://gabrielecirulli.github.io/2048

In Japanese logic puzzles, the subproblems to be solved are the necessary deductions, and the rewards are the simplifications that follow each deduction to reveal new information. If you have any doubt that such puzzles are in fact addictive, then next time you solve one, note the urge to complete just one more item… then one more item… then one more item…

Andrews [14] suggests a more direct causal explanation of why people often find the activity of solving puzzles so emotionally rewarding. He explains that MRI brain scans indicate a relationship between a 'satisfaction centre' in the brain called the *striatum*, which is activated by stimuli associated with reward, and areas of the frontal cortex that are involved with logical thought and planning towards goals. He posits that it is this connection between the 'intellectual' cortex and the 'emotional' striatum that gives us pleasure in response to solving problems, and drives us on to seek further problems to solve.

This addiction for solving puzzles may not even be confined to the human brain. Recent research at the UK's Whipsnade Zoo [15] found that chimpanzees given particular dexterity challenges appeared 'keen to complete the puzzle' for its own sake, regardless of whether those challenges were associated with a food reward or not.

The following section explores the ramifications of these ideas on the form of puzzles.

## 3 Form

In this context, the *form* of a puzzle refers to the degrees of freedom that the setter can manipulate, in order to make challenges more interesting, engaging and aesthetically pleasing for the solver. The *function* of a puzzle refers to the conceptual framework within which such forms exist, i.e. those necessary conditions for the puzzle to work.

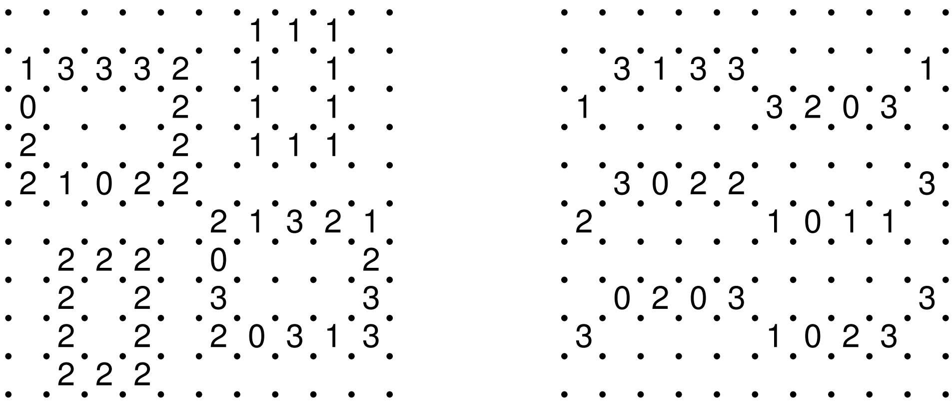

**Figure 6.** Slitherlink examples with symmetrical hint placement.

Authorial control allows the setter to impart some structure on their designs, in order to impart some of their personality on the challenges they produce. This section examines some relevant aspects of form that puzzle setters can manipulate.

## 3.1 Symmetry

An obvious way to inject structure into a design is through symmetry. However, it is important to realise the difference between *visual* (superficial) and *strategic* (underlying) symmetry.

### 3.1.1 Visual Symmetry

*Visual symmetry* is achieved through the symmetrical placement of hints defining each challenge. For example, Figure 6 shows two Slitherlink challenges (#6 and #21 from [1]) with rotationally symmetric hint placement. Many publishers, including Nikoli, have a policy of only publishing symmetrical challenges for most of their puzzles.

To see the reason for this requirement, consider the pair of challenges shown in Figure 7, from a recent study in automated puzzle design [16] involving a new puzzle game called Hour Maze.[5] Both challenges were generated by computer, and both describe valid challenges of similar difficulty on the same background maze, but notice how the symmetrical hint set on the right imposes a sense of order that hints at non-random generation.

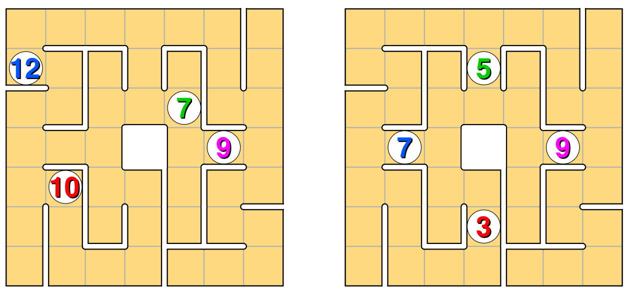

**Figure 7.** An Hour Maze example with asymmetrical (left) and symmetrical (right) hint placement.

Visual symmetry offers the superficial appearance of structure; symmetrical challenges look neater and more elegant but are not necessarily more interesting to solve. However, there is no reason to preclude symmetry as a design constraint, if it pleases the setter or solver, and helps elevate puzzle design to an art form.

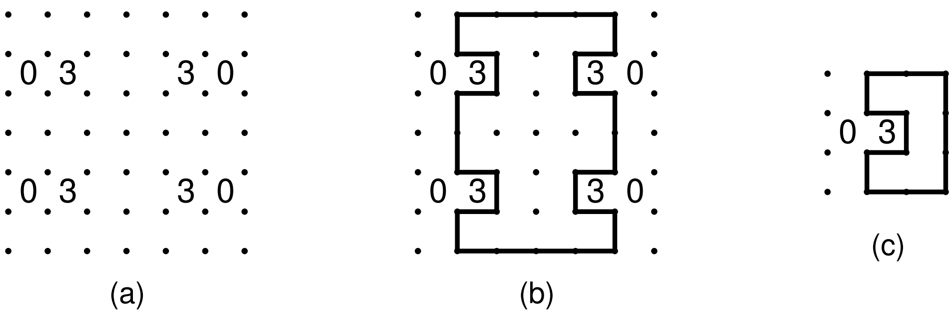


**Figure 8.** Strategic symmetry makes this Slitherlink example a trivial repetition of pattern (c).

[5]The aim in Hour Maze is to fill the grid with coloured number sets 1–12, such that adjacent numbers differ by $\pm 1$.

### 3.1.2 Strategic Symmetry

*Strategic symmetry* refers to pattern or repetition inherent in the solution process itself. This is typically more important than visual symmetry, as it reflects the solution process directly, and can lead to bad designs unless used judiciously.

For example, the Slitherlink challenge shown in Figure 8 is highly symmetric both in its visual design (a) and in its solution (b), which is essentially a repetition of the pattern (c) four times. This challenge is highly redundant and boring; the solver typically wants to be presented with novel subproblems to solve within each challenge.

Similarly, consider the Kakuro challenge shown in Figure 9 [17]. The aim of Kakuro is to fill the grid with digits {1, 2, 3, ..., 9}, such that each consecutive run totals the number shown and does not contain duplicate digits. This challenge exhibits a high degree of strategic symmetry, with several immediate simplifications (small text) being reflected on opposite sides of the grid in the same combinations, creating a high degree of redundancy.

This attempt by the setter to inject some structure into the design may well backfire, unless the solver is happy repeating the same operations in different parts of the board. Indeed, participants in the Hour Maze experiment exhibited a slight *negative* correlation between wall symmetry and puzzle enjoyment [16], perhaps due to the fact that symmetrical mazes tend to produce strategically symmetric (i.e. redundant) solutions.

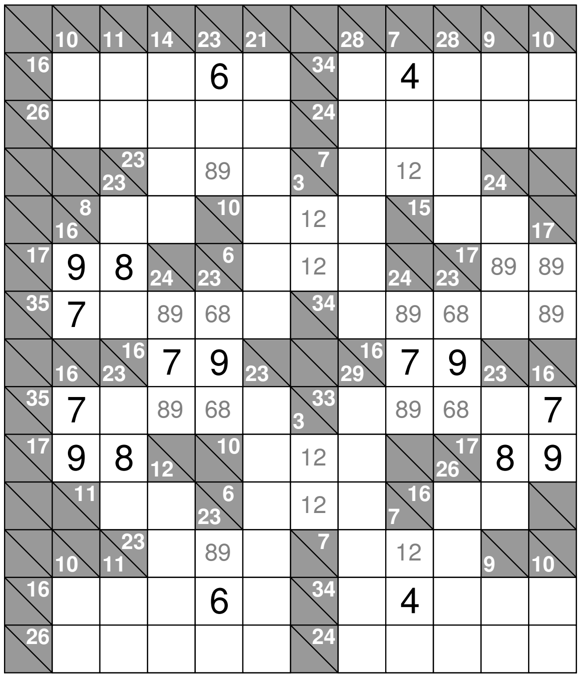

**Figure 9.** A Kakuro challenge showing redundant strategic symmetry.

Slitherlink challenge #19 from [18] shown in Figure 10, on the other hand, demonstrates a positive example of strategic symmetry. The natural solution path for this challenge, once obvious simplifications have been performed, is as follows:

1. a cascade down the left hand side (a),
2. a cascade up the right hand side (b), and
3. a cascade up the centre connecting them (c).

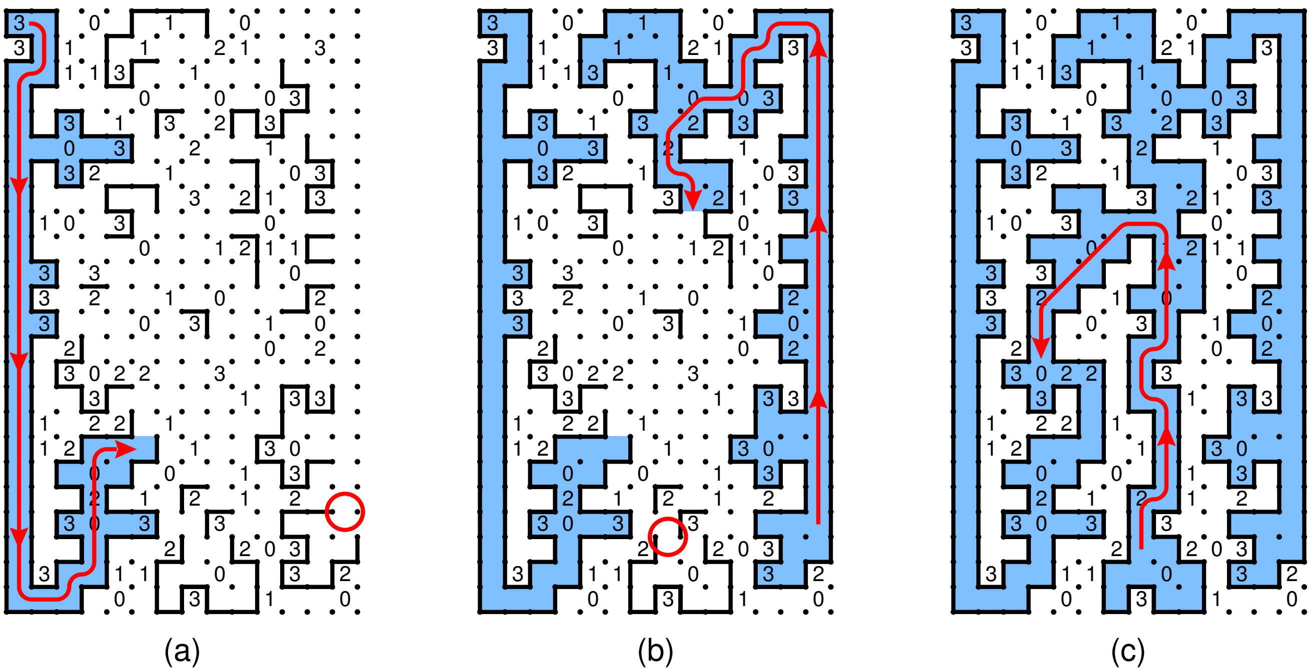


**Figure 10.** Slitherlink challenge with rotationally similar solution paths.

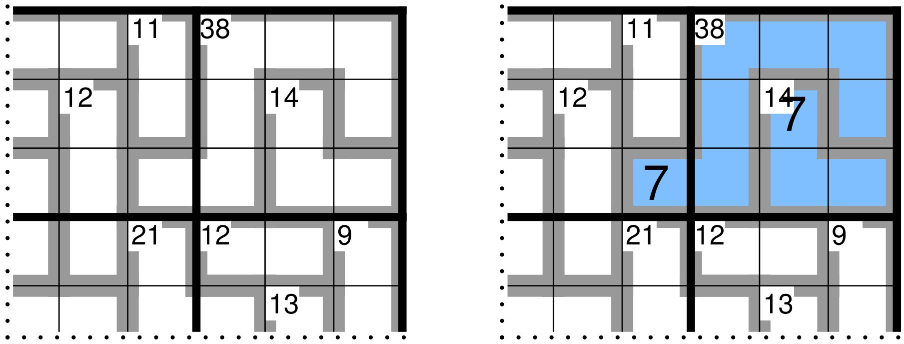


**Figure 11.** A self-resolving 'snail' pattern in Killer Sudoku.

The long cascades along each side are rotationally similar to each other in a global sense. However, each involves the solution of different local subproblems during their propagation, imparting structure on the solution without redundancy. As a general rule of thumb, *global symmetry should be accompanied by local asymmetry*, and vice versa.

The duelling cascades in this challenge are unlikely to have occurred by chance, and the solver has a real feeling of an intelligent hand at work in the design, who is perhaps having a bit of fun with them. This challenge is a nice example of authorial control in action.

## 3.2 Pattern

Handcrafted puzzle challenges often include repeated patterns or *motifs* as an expression of the setter's personality. In spatial puzzles such as Slitherlink, such motifs might involve: matching cascades as shown in Figure 10; letter, number or animal shapes in the solution path; or any other interesting nonrandom patterns. In fact, Nonograms, another type of Japanese logic puzzle, actually produce works of art (or at least pictures) as the solver colours in the cells of a grid according to certain rules.[6]

Motifs may also occur in number puzzles, such as the 'snail' pattern in the Killer Sudoku example shown in Figure 11 [19]. The top right 3×3 subgrid contains two regions, totalling 38 and 14 respectively, with one cell on the 38-region exceeding the subgrid boundary (left). This rogue cell must resolve to 7, which is the difference between the sum of the two regions (38+14 = 52) and the disjoint sum of all digits (1+2+3+4+5+6+7+8+9 = 45) that the 3×3 subgrid must total (right). The central cell can then also be resolved to a 7, as digits cannot be repeated in any row or region.

This snail pattern provides an elegant self-resolving starting point for this challenge. This particular challenge contains a rotated variation of it in each corner, giving the strong impression of carefully structured design. But again, there can be a fine line between amusing the solver through pattern and annoying them through redundancy.

## 3.3 Ladders

*Ladders*, i.e. forced sequences of moves typically in a repeated pattern, are another indicator of authorial control. They are different to the patterns described above, as they are implicit in the design and manifest themselves during its solution.

For example, consider the section of a hypothetical Slitherlink challenge shown in Figure 12 (left). The two 0 hints at the top left dictate which two edges of the adjacent 2 hint are 'on', which in turn dictate which two edges of the adjacent 2 hint are 'on', which in turn dictate which the two edges of the adjacent 2 hint are 'on', and so on. This is a contrived example, but such self-perpetuating ladders are often found in actual challenges.

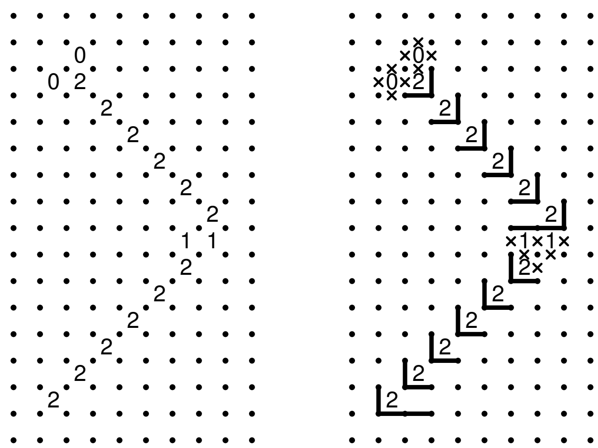

**Figure 12.** A self-perpetuating Slitherlink ladder.

[6]http://www.nonograms.org

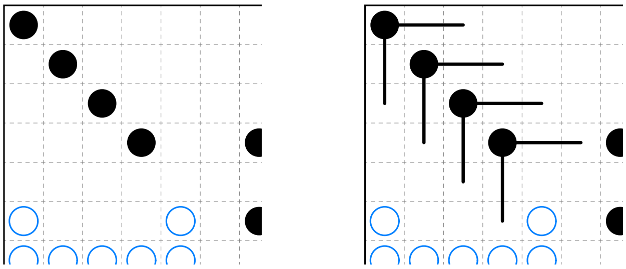

**Figure 13.** A Masyu ladder.

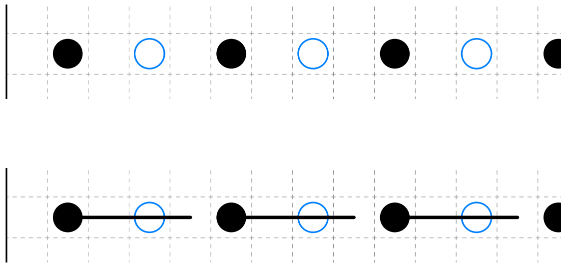

**Figure 14.** Another Masyu ladder.

Figure 13 shows a section of a Japanese logic puzzle called Masyu, in which the solver must draw a single non-self-intersecting path through every circle on the board, such that the path turns within each black circle but passes straight through each white circle. See the 'Masyu' article in this issue (which includes this example) for more details [20]. In this case, the line of black dots triggers a self-perpetuating ladder from the top left corner inwards 13 (right).

Figure 14 shows another example of a ladder, in another Masyu challenge from [20]. In this case, the alternating sequence of black and white circles (top row) forces the self-perpetuating ladder of edges shown (bottom row), as the path must extend straight for two cells from each black circle and also pass straight through each white circle.

Such ladders can also have strategic value, as they can often be clumped into a single unit of information when they are recognised, reducing the solver's mental workload through *modularity* [21] or *chunking*. For example, if a Slitherlink solver notices a diagonal line of 2s as in Figure 12, then the effect of a deduction at one end of the line can often be seen immediately at the other end of the line, without having to think about the intervening items.

### 3.4 Surprise

An element of surprise can keep a challenge interesting and impart the impression of authorial control, typically by establishing a pattern in the solution process and then suddenly disrupting it.

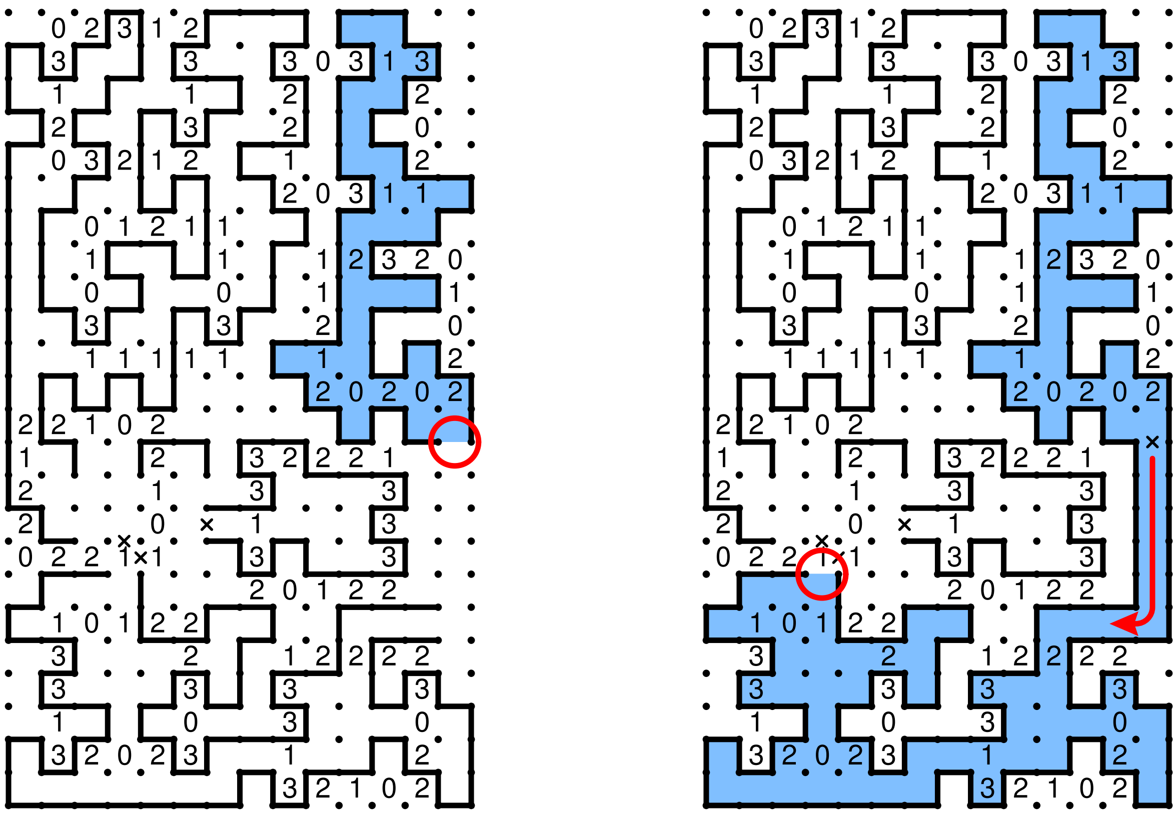

**Figure 15.** A discontinuous jump propagates this Slitherlink solution.

Figure 15 shows a Slitherlink challenge whose obvious progress point is circled on the shaded region in the top right corner (left). This region can only expand downwards, connecting to its neighbouring group (right), but the obvious progress point now jumps to the left side of this extended group (circled). The solver, after following an orderly cascade down the right hand side, must suddenly switch to the other side of the grid.

While the shading in the figure makes this discontinuity easy to spot, larger jumps in more complex situations can be confusing for the solver. Deductions that trigger key information in distant parts of the grid can suggest an intelligent setter actively trying to keep the solver on their toes.

Famous Chess puzzle setter Sam Loyd recognised the importance of surprise, stating that his goal was to compose puzzles whose solutions require a first move that is contrary to what 999 players out of 1,000 would propose [22].

### 3.5 Perversity

When it comes to expressing personality, it is hard to beat sheer perversity. Consider the Slitherlink challenge shown in Figure 16 (a).[7] While this challenge has an obvious solution (b), this can be difficult for experienced Slitherlink solvers to spot. The problem is that two adjacent 3-hints form a common pattern in Slitherlink that invariably implies three parallel edges in a normal context (c).

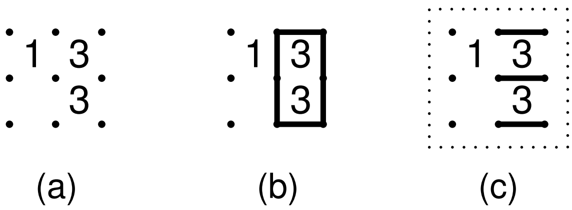


**Figure 16.** A Slitherlink joke.

Experienced solvers will fixate on this learnt pattern and simply not see the obvious solution; they must unlearn habits ingrained over many hours of reinforcement. Such blatant disregard for tradition allows the innovative setter to subvert the solver's expectations for a bit of mischief.

Sudoku challenge #99 from [6], shown in Figure 17, is another case in point. Three values can be immediately resolved from the initial hint set with little effort (highlighted), leading the solver to think that this challenge is not so difficult.

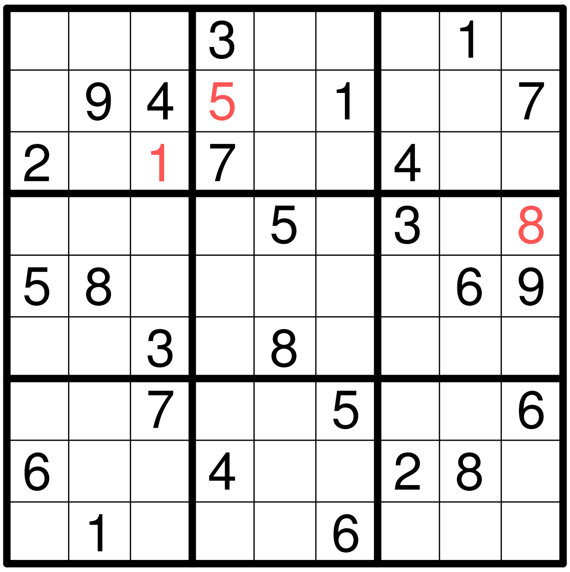

**Figure 17.** An easy Sudoku challenge… or is it?

However, the information soon dries up and its true difficulty becomes apparent; this is actually the most difficult challenge in its collection. Such deception is common in deduction puzzles. If a challenge starts off as being particularly easy, then the solver may be lulled into a false sense of security, but can expect tough times ahead.[8]

The Killer Sudoku challenge shown in Figure 18 [23] has two points of interest. Firstly, the top right 3×3 subgrid contains three regions that fit exactly within the subgrid, whereas typically at least one region would overlap its boundary to provide some information for the solver; this is *almost* a standard solution pattern, but not quite.

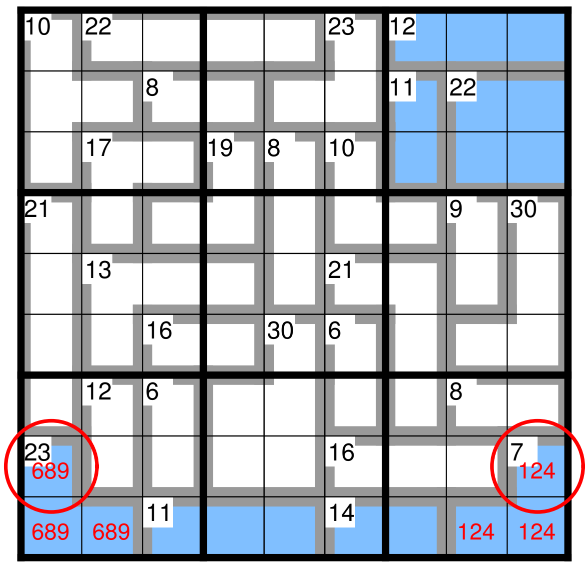

**Figure 18.** Patterns that yield little information.

[7] Provided by Jimmy Goto from Nikoli.

[8] Unless, of course, the challenge is actually rated as 'easy'.

Secondly, the shaded 23, 11, 14 and 7-regions along the bottom row sum to 23+11+14+7 = 55, hence the two circled cells must sum to 10 (since the nine cells along the bottom row must add to the disjoint sum of all digits 1+2+3+4+5+6+7+8+9 = 45). However, all of the values available for these two cells, {6, 8, 9} and {1, 2, 4} respectively, all have pairings that yield 10, hence none can be eliminated and this potential line of enquiry fails. Hence, a normally fruitful solution pattern yields surprisingly little information.

The fact that this challenge immediately thwarts an experienced solver's expectations in two such blatant ways suggests that this challenge was set by a designer wanting to tease the solver. The challenge shown in Figure 19 [24] also demonstrates some unusual design features:

- The hint subregion layout is horizontally, vertically and rotationally symmetrical.
- Single cell subregions exist, which can be immediately instantiated to their shown value.
- Each 3×3 subgrid is dominated by a plus-shaped subregion that separates its four corners.
- Every subregion can be consistently classified as either: 1) corner, 2) edge, 3) plus-shape, or 4) 2×2 junction.

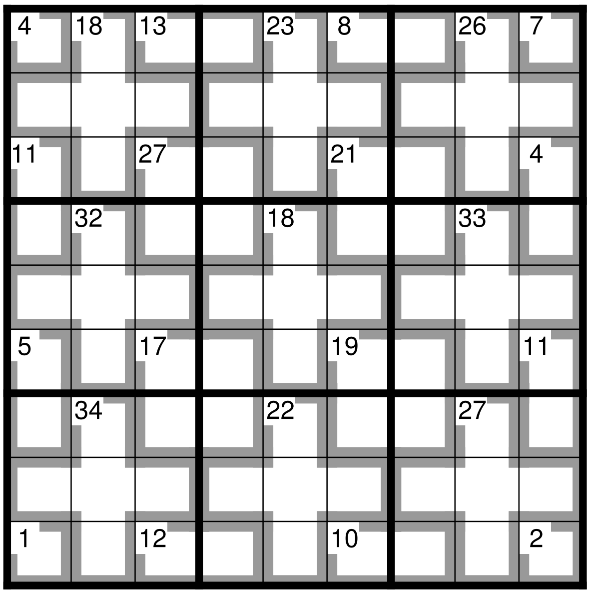


**Figure 19.** An unusually symmetrical challenge.

Killer Sudoku challenges typically involve asymmetrical hint subregion layouts – such as Figure 18 – to allow maximum diversity and strategic depth. To find any of these types of symmetry in a challenge is a surprise, but to find all four speaks of a very eccentric design that pushes this puzzle's design constraints to the limit.

The single cell subregions that immediately instantiate to their shown value, in particular, are unusual and violate the conventions of this puzzle. But this challenge is also engaging and interesting to solve, which indicates a meticulous design carefully constructed by its setter.

## 4 Conclusion

While this paper only touches on a small number of puzzle types, mostly Japanese logic puzzles, I believe that the principles described have relevance to puzzle design in general. Most important is the way in which the setter can exert authorial control on their designs, in order to impart some structure and personality – manipulating the solver's actions in absentia – to make challenges interesting, engaging and aesthetically pleasing.

Phrasing puzzles as two-player games played between the setter and the solver helps explain these ideas, but this is not to say that puzzles are zero-sum games.[9] The setter does not win by defeating the solver as often as possible, but by providing the most interesting, engaging and aesthetically pleasing challenges as often as possible. Both win if the setter provides interesting challenges and the solver enjoys completing them.

I hope that this discussion is useful for puzzle setters, and interesting to solvers, and might suggest concrete approaches for improving the quality of computer-generated designs.

## Acknowledgements

Thanks to Raf Peeters and Stephen Tavener for discussions on puzzle design. Thanks also to Jimmy Goto for insights into Nikoli's design process, and permission to use Nikoli challenges in this article.

## References

[1] *The Times Japanese Logic Puzzles: Hitori, Hashi, Slitherlink and Mosaic*, London, Harper Collins,

[9] A zero-sum game is one in which a win for one player implies a loss for the other participant(s).

2006.

[2] Kim, S., 'What is a Puzzle?', 2008. http://www.scottkim.com.previewc40.carrierzone.com/webgames.html

[3] Higashida, H., 'Machine-Made Puzzles and Hand-Made Puzzles', in *IFIP Advances in Information and Communication Technology (AICT)*, vol. 333, 2010, pp. 214–222.

[4] Schuh, F., *The Master Book of Mathematical Recreations*, New York, Dover, 1968.

[5] Thompson, M. J., 'Defining the Abstract', *Game & Puzzle Design*, vol. 1, no. 1, 2015, pp. 83–86.

[6] *Sudoku 1*, Tokyo, Nikoli, 1988.

[7] *Fresh Slitherlink 3*, Tokyo, Nikoli, 2012.

[8] Pelánek, R., 'Difficulty Rating of Sudoku Puzzles: An Overview and Evaluation', *arXiv*, 2014. http://arxiv.org/abs/1403.7373

[9] Anonymous, 'Killer Sudoku', *The Observer*, 11 August 2013.

[10] Kanamoto, N., 'A well-made Sudoku is a pleasure to solve', *Nikoli*, 2001. http://www.nikoli.co.jp/en/puzzles/why_hand_made.html

[11] *The Nikoli Times*, 26 July 2013.

[12] Stafford, T., 'The Psychology of Tetris', *BBC Future*, 23 October 2012.

[13] Humphrey, G., 'The Psychology of the Gestalt', *Journal of Educational Psychology*, vol. 15, no. 7, 1924, pp. 401–412.

[14] Andrews, M., 'Ask the Brains', *Scientific American: Mind*, August/September, 2007, p. 86.

[15] Clark, F. E. and Smith, L. J., 'Effect of a Cognitive Challenge Device Containing Food and Non-Food Rewards on Chimpanzee Well-Being', *American Journal of Primatology*, vol. 75, no. 8, 2007, pp. 807–816.

[16] Browne, C., 'Metrics for Better Puzzles', in El-Nasr, M. S., Drachen, A., Canossa, A. and Isbister, K. (eds.), *Game Analytics: Maximizing the Value of Player Data*, Berlin, Springer, 2013, pp. 769–800.

[17] Anonymous, 'Kakuro', *The Guardian*, 9 August 2013.

[18] *Fresh Slitherlink 2*, Tokyo, Nikoli, 2011.

[19] Anonymous, 'Killer Sudoku', *The Observer*, 17 November 2013.

[20] Goto, J., 'Masyu', *Game & Puzzle Design*, vol. 1, no. 1, 2015, pp. 19–23.

[21] Karmiloff-Smith, A., *Beyond Modularity*, Massachusetts, MIT Press, 1992.

[22] Hayward, R. and Van Rijswijck, J., 'Hex and Combinatorics', *Discrete Math*, no. 306, 2006, pp. 2515–2528.

[23] Anonymous, 'Killer Sudoku', *The Guardian*, 31 August 2013.

[24] Anonymous, 'Killer Sudoku', *The Observer*, 5 January 2014.

**Cameron Browne** is a Vice-Chancellor's Senior Research Fellow at QUT, Brisbane, Australia, whose research interests include artificial intelligence and automated game design.
**Address:** School of EECS, Science and Engineering Faculty, QUT, Brisbane, 4001, Australia.
**Email:** c.browne@qut.edu.au

## Heptalion Challenge #2

Use all tiles to cover all symbols. See 'Heptalion' (p. 17) for details.