



The hidden logic of Sudoku.- 2nd edition. Conclusion

Denis Berthier

► To cite this version:

Denis Berthier. The hidden logic of Sudoku.- 2nd edition. Conclusion. Denis Berthier. The hidden logic of Sudoku.- 2nd edition, Lulu.com, 2007, 978-1-84799-214-7. hal-00641973

HAL Id: hal-00641973

<https://hal.science/hal-00641973>

Submitted on 17 Nov 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The Hidden Logic of Sudoku (Second Edition):

Conclusion

Denis Berthier

Institut Telecom

<http://www.carva.org/denis.berthier>

Conclusion

What has been achieved

In this conclusion, I'd like first to highlight a few facets of what has been achieved in this book, from four complementary overlapping points of view.

1) ***From the point of view of the Sudoku addict***, the most striking results should be the following.

By formalising the familiar resolution rules, we have eliminated from some of them the ambiguities that plagued their usual presentations and we have clarified their scopes. For instance, we have shown that none of the two usual formulations of Triplets (or Quadruplets) – neither the "strict" nor the "comprehensive" – was the best; our reformulation shows how close Triplets (and Quadruplets) are to xy-chains but also why they cannot be reduced to them (nor to our stronger xyt- or xyzt-chains).

We have fully clarified the symmetry relationships that exist between Naked, Hidden and Super-Hidden Subset rules (the latter being classically known as X-Wing, Swordfish, Jellyfish, Squirmbag and other existent or non existent "fishy patterns"). Such relationships had already been partly mentioned on some Web pages, but never in the systematic way we have dealt with them here. As a result, we have naturally found the proper formulations for the Hidden and Super-Hidden Subset rules and we have proven that no other subset rule of such type can exist.

More generally, we have proven three meta-theorems that automatically produce new resolution rules (their Hidden or Super-Hidden counterparts) from existing ones.

With the introduction of new graphical representations of the problem in auxiliary 2D spaces (mainly row-number and column-number spaces), we have shown that the Hidden or Super-Hidden counterparts of well-known patterns (subsets or chains) defined in natural row-column space can be detected as easily as their originals, because in these new spaces they look like the originals do in the standard representation. We have thus extended the resolution power of such patterns.

We have defined a (non strict) complexity hierarchy between the resolution rules, compatible with their logical symmetry relationships.

We have proven certain theorems showing that particular resolution rules can be formally reduced to simpler ones in the above hierarchy (e.g. "xy-chains of length 3 are subsumed by Naked-Triplets plus XY-Wing").

We have evaluated the strength of each rule by the proportion of new puzzles ("new" according to the above hierarchy) its introduction allows to solve and, for the first time, such an evaluation has been based on a large collection of puzzles (more than 56,000), with detailed results available online.

We have given chain rules a major place in this book (they occupy more than half of it), because they are the main tool for dealing with hard puzzles but they remain the subject of much confusion. We have introduced a general conceptual framework (including the notion of a target not belonging to the chain) for dealing with all conceivable types of chains and we have applied it systematically to all the types we have defined. In particular, we have introduced an intuitive graphical language of patterns for specifying chains and their targets, abstracting them from any irrelevant facet (such as a link being a row or a column or a block), and we have shown that these patterns are equivalent to logical formulæ. In our framework, the confusing notions of "chain of inferences", "weak link" and "strong link" are never used; our chains are well defined patterns of candidates, cells and links.

We have chosen the simplest kind of homogeneous chains, the xy-chains, as our main type of chains, with all the other chains being intuitive generalisations of them.

We have proven that xy-chains and c-chains should have no loops (with the practical consequence that searching for these chains becomes simpler for both a human and a computer).

By pushing the underlying logic of xy-chains to its limits, we have introduced xyt-chains and shown that they are a very natural and powerful generalisation of xy-chains. We have also defined another independent generalisation, the xyz-chains, and combined it with the previous one to get the xyzt-chains.

With each type of chain in natural row-column space, we have associated two new types of chains, their hidden counterparts in row-number and column-number spaces, and, in particular, we have shown the unifying power of the hidden xy-chains. All these chains can be spotted in either of the 2D representations, using our Extended Sudoku Board.

We have also proven theorems allowing to combine our various types of homogeneous chains to build heterogeneous, more complex, ones.

In this second edition, we have generalised the above 2D chains, introduced their fully super-symmetric 3D versions (the nrc-, nrct-, nrcz- and nrczt- chains) and shown that all these types of chains can still be considered in a natural way as various generalisations of the basic xy-chains.

We have produced a multiplicity of well chosen examples proving that particular resolution rules cannot be reduced to simpler ones (in the above hierarchy).

In particular, we have proven that, using only the types of chain rules introduced in the first edition, it is necessary to consider chains of length more than thirty if we want to have a chance of solving all the randomly generated puzzles without resorting to Trial and Error or assuming the uniqueness of a solution.

In the first edition, we had exhibited a set of resolution rules (L13) based on only 2D chains that could solve 97% of the randomly generated minimal puzzles and 99,67% of the 36,628 17-minimal puzzles in the famous Royle database (without resorting to Trial and Error or to an assumption of uniqueness).

In this second edition, we have also exhibited a set of resolution rules (M5) that can solve more than 99% of the randomly generated minimal puzzles using only 3D chains of length no more than five and a set of resolution rules (M7) that can solve 99.9% of these puzzles using only 3D chains of length no more than seven. As psychologists consider that human short term memory has size seven plus or minus two, this means that a human being using these rules should be able to solve almost any puzzle without any computer assistance (but still with some patience). It should be noticed that these chains do not include subsets (Hinges or Almost Locked Sets or grouped chains), contrary to the currently popular chains (Nice Loops or Alternating Inference Chains), thus avoiding a potential source of exponential behaviour.

2) *From the point of view of mathematical logic*, our most obvious result is the introduction of a strict formalism allowing a clear distinction between the straightforward *Sudoku Theory* (that merely expresses the constraints defining the game) and all the possible *Sudoku Resolution Theories* formulated in terms of condition-action rules (that may be put to practical use for solving puzzles). We have given a clear logical definition of what a "resolution rule" is, as opposed to any logically valid formula. With the notions of a resolution theory T and a resolution path (which is merely a proof of the solution within T , in the sense of mathematical logic), we have given a precise meaning to the widespread but as yet informal idea that one wants a "pure logic solution". This leads to both sound foundations and intuitive justifications for our resolution theories, exhibiting the following facets and consequences.

We have established a clear logical (epistemic) status for the notion of a candidate – a notion that is quasi universally introduced for stating the resolution rules but that does not pertain *a priori* to Sudoku Theory and that is usually used only from an intuitive standpoint. Moreover, we have shown that the epistemic operator that must appear in any proper formal definition of this notion can be "forgotten" in practice when we state the resolution rules and that this notion can be considered as primary, provided that we work with intuitionistic (or constructive) logic instead of standard logic (this is not a restriction in practice). Notice that this whole approach can be extended to any game that is based on techniques of progressive elimination of candidates.

We have also defined what a "resolution method" based on a resolution theory is and we have shown that all the resolution theories introduced in this book have the very important *confluence property*, allowing any ordering to be superimposed on their resolution rules without changing their overall resolution capacity. As a major practical consequence, in any software implementation of a resolution theory into a resolution method (e.g. in our SudoRules solver), we may take advantage of any convenient ordering of the rules.

The natural symmetries of the Sudoku problem have been expressed as three general meta-theorems asserting the validity of resolution rules obtained by some simple transformations of those already proven. These meta-theorems have been stated and proven both intuitively and formally. As a first example of how these meta-theorems can be used in practice, we have exhibited a precise relationship between well known (Naked and Hidden) Subset rules with what we call their Super-Hidden counterparts (the famous "fishy patterns") and we have proven some form of completeness of the set of known Subset rules. As a second example of how these meta-theorems can be used in practice, we have defined entirely new types of chain rules, hidden chains of various types, and shown their unifying power.

We have also devised a direct proof of the existence of a simple and striking relationship between Sudoku and Latin Squares: *a block-free resolution rule (i.e. a rule that does not mention blocks or squares) is valid for Sudoku if and only if it is already valid for Latin Squares*. Notice that it does not seem one can prove this result by using only the general methods one would expect to see used in such cases: either the interpolation theorem or the techniques of Gentzen's sequent calculus.

Finally, we have provided a very intuitive example of how difficult it may be to transform a theory formulated in terms of (a few and simple) constraints into a set of "production rules" (or condition-action rules). This also shows that, although the given constraints on rows and columns on the one hand and the constraint on blocks on the other hand can be formulated as axioms with independent predicates, many of the condition-action rules necessary to operationalise them do mix these predicates. This mixture is most visible in the 3D (or fully super-symmetric) chain rules.

3) From the point of view of Artificial Intelligence (AI), the following should be stressed.

Sudoku is a wonderful example for AI teachers. It has simpler rules and is more accessible for student projects than games such as chess or go, but it is much more complex and exciting than the usual examples one can find in AI textbooks (Tic-Tac-Toe, Hanoi Towers, Monkey and Bananas, Bricks World, ...). It easily suggests lots of projects based on the introduction and the formalisation of new types of rules since no complete set of resolution rules is known (see below).

Sudoku provides a very good illustration of a basic software engineering principle: never start writing a program (or a knowledge base for an inference engine) unless you have a non-ambiguous specification for it. The logic of certain resolution rules is so subtle that the least deviation from it (e.g. forgetting that a target of a chain may not belong to it or that loops are not allowed in a chain) has "catastrophic" consequences (typically some puzzles being falsely claimed to have no solution). This can also be considered a nice illustration of Newell's classical distinction (introduced in [New 82]) between the *knowledge level*, here assimilated to the logical formulation, in which knowledge must be formulated in a "declarative" form independent of any processes that might be applied to it, and the *symbol level*, here assimilated to the set of rules in the CLIPS or the JESS language, in which the knowledge level is *operationalised* in a form that may depend (and does largely depend in practice) on the specificities of the knowledge representation and processing tools used to implement it. Although this book does not tackle this point, there are many ways a given resolution rule (as formulated in logic) can be implemented as a production rule in an inference engine, which are more or less related to the different ways it may be used in practice.

Sudoku is also a wonderful testbed for the inference engine chosen to run the knowledge base of resolution rules. The logic of the set of rules is so intricate that many functionalities of the inference engine are stringently tested, which is how we discovered a longstanding undocumented bug in the management of saliences in JESS. With long series of puzzles to solve, memory management can also be a problem (as it is in CLIPS).

The previous topic is related to a crucial problem of AI, both practical and epistemological: how can one be sure that the system does what it is intended to do? Although this is already a very difficult question for classical software (i.e. mainly procedural, notwithstanding the object oriented refinements), it is much worse for AI, for two main reasons. Firstly, the logic underlying a knowledge base is generally much more complex than a set of procedures is (otherwise it would probably be much better to solve the problem with procedural techniques) and secondly an inference engine is a very complex piece of software and debugging it is very difficult.

As a general result, using AI to prove theorems (although this has always been and remains one of the key subtopics of the domain) may make mathematicians and logicians very suspicious.

As a specific result, all the independence theorems that have been proven in this book rely on the correctness of the inference engine we used for our computations. They do not depend on this correctness when we assert that "theory T allows the following resolution path for puzzle P", since the validity of any particular path can easily be checked "by hand", whichever method was used to generate it. But these independence results depend on this correctness any time we state that a particular theory is not able to find a solution for some puzzle P. This is why all our computations have finally been done with CLIPS instead of JESS despite the fact that CLIPS regularly gets lost in memory management problems and computation times on long series of puzzles grow astronomical. But the fact that, due to its problem with the management of saliences¹, JESS misses some inferences on simpler rules, even though this may be infrequent, disqualifies it as a theorem prover in our case (so much so that missed inferences also vary from one version to the next). Obviously, this does not prove that CLIPS is bug free. The only certain conclusions are that, using the same knowledge base, CLIPS solved (a few) more puzzles than JESS and never returned any spurious message of type "this puzzle has no solution".

Of course, these independence theorems also rely on the correctness of the knowledge base we have developed to implement all the rules defined in this book.

¹ It seems that this bug has been corrected in the latest release of JESS (71b1) available as of the publication of this second edition. But we haven't carried out systematic tests.

In this respect, I would like to make three points:

- firstly, thanks to the monotonicity of the facts base (each rule can only add values or not-candidates), a confluence theorem has been proven, which guarantees that there cannot be unwanted interactions between the rules;

- secondly, each individual rule has received the same systematic treatment: it has been stated in plain English, with great care being taken for not forgetting any conditions; it has then been proven, still in plain English, in rather straightforward steps that can be checked by anybody; its formulation in multi-sorted logic (or, for the chain rules, in our equivalent graphical formalism) has been shown to be the direct transliteration of the English one; in turn, the CLIPS or the JESS formulation is the direct transliteration of the logical one, thus minimising the possibilities of discrepancies between successive steps in the development;

- thirdly, the current release of our solver (SudoRules 13) has been tested on more than 56,000 puzzles known to have a unique solution (and this produced the classification results in chapters XXI and XXIII); whereas an error in a rule that would illegitimately eliminate candidates leads very rapidly to the claim that a puzzle has no solution (this allowed the detection of a few subtle bugs in the first version of SudoRules), it is noticeable that SudoRules has not yet produced an incorrect result in the CLIPS environment.

4) *Finally, considering the (currently very fashionable) notion of complexity*, problems that can be stated in simple terms but that need a complex solution are not anything new, although the general idea may remain obscure for the novice thinker. Among the most famous problems of this type, you have certainly heard of the four-colour problem in graph theory or Fermat's conjecture in arithmetic (now known as the "Fermat-Wiles Theorem", since it has been proven recently by Andrew Wiles, more than three centuries after its formulation by Fermat). But proofs of these theorems are not really accessible to the non-mathematician and the type of complexity hidden behind the problem statement therefore remains very abstract. With the notion of deterministic chaos, the second part of the twentieth century has uncovered a new type of complexity: some dynamical systems ruled by very simple equations may have very complex trajectories and two neighbouring points may follow quickly diverging paths – but this also remains a little mysterious if you do not have a minimum mathematical background.

On the contrary, with the popular game of Sudoku, you can get a feeling of another type of complexity, computational complexity (how this is related to the previous ones remains an interesting but very difficult question). Sudoku is known to be NP-complete, i.e., to state it very informally (and somewhat incorrectly), when

we consider grids of increasing sizes, resolution times grow faster than any deterministic polynomial algorithm. As you will never try to solve a Sudoku puzzle on a 100x100 grid (unless you have unlimited free access to a psychoanalyst), this may also remain an abstract definition. There is nevertheless a difference with the previous examples: you can already get a foretaste of the underlying complexity with the standard 9x9 puzzles (e.g. by comparing them to their homologues on 4x4 grids, the so-called Sudokids).

The Sudoku problem is defined by four very simple constraints, immediately understood by everybody in terms of "single occupancy" of a cell and of "mutual exclusion" in a row, a column or a block. For a classically formatted mind, it is therefore natural to think that any puzzle can easily be proven to have no solution or be solved by a finite set of simple operational resolution rules of the condition-action type: "in such a situation carry out such an action (assert a value or eliminate a candidate)". And this idea can only be reinforced if you consider the majority of puzzles published in newspapers. But the independence results proven in this book through a multiplicity of examples have shown that very complex resolution rules are indeed needed.

What this book has shown then, in both intuitive and logically grounded ways, is that writing a set of operational rules for solving an apparently very simple constraints propagation problem may be a very complex task. (Indeed, notwithstanding their overall complexity, the rules that have been defined in this book do not even form a complete resolution theory.) Moreover, as all the NP-complete problems are equivalent (through polynomial transformations) and some of them have lots of practical applications, such as the famous travelling salesman, dealing with the apparently futile example of Sudoku may provide intuitions on problems that seem to be unrelated.

What has been partly achieved (from the point of view of AI)

In the introduction, we said we wanted a set of rules that would simulate a human solver and that could explain each of the resolution steps. The explanations produced by SudoRules are largely illustrated by the listings given in this book; they are sufficiently explicit once you know the definitions of our rules and it would be easy work to make them still more explicit for those who do not know them; but we do not consider this as a very exciting topic. As for the solver facet, SudoRules does simulate a human solver, a particular kind of player who would try all our rules systematically (ordered according to their complexity) on all of their potential instantiations.

Is it likely that any human solver would proceed in such a systematic way? He may prefer to concentrate on a part of the puzzle and try to eliminate a candidate from a chosen cell (or group of cells). What may be missing then in our system is a "strategic" knowledge level: when should one look for such or such pattern? But I have no idea of which criteria could constitute a basis for such strategic knowledge; moreover, as far as I know, whereas there is a plethora of literature on resolution techniques (often misleadingly called strategies), nothing has ever been written on the ways they should be used, i.e. on what might legitimately be called strategies.

To say it otherwise, we do have a strategic level: searching for the different patterns in their order of increasing complexity. Notice that there is already more strategy in this than proposed by most of the books or Web pages on the subject. The question then is: can one define a better (or at least another) strategy? Well, the rules in this book (and the corresponding software SudoRules) are there; you can use them as a basis for further analysis of alternative strategies. One of the simplest ways to do so is to modify the complexity measure and the ordering we have defined on the set of rules. For instance, using psychological analyses, one could break or relax the symmetry constraints we have adopted.

What was not our purpose and has not been achieved; open questions

The first thing that has not been done in this book is a review of all the advanced rules that have been proposed, mainly on the Web (under varying names, in more or less clear formulations, with more or less defined scopes). The list would be too long and moreover it is regularly increasing. The best place to get an idea on this topic is in the recent book by Andrew C. Stuart ([STU 07]) or on the Web, in particular in the Sudoku Players Forum:

<http://www.sudoku.com/forums/viewtopic.php?t=3315>

(with the problem that chains are often considered as "chains of inferences" instead of patterns and they are sometimes classified according to absurd criteria).

Instead, our two main purposes in this regard were to take advantage of the symmetries of the game in a systematic way and to isolate a limited number of rule types, with rules definitions extended as far as the arguments used in their proofs allowed: this is how we introduced xyt-, xyz- and xzyt- chain rules (on the basis of xy-chain rules) and their hidden counterparts (on the basis of our general meta-theorems); this is also how this second edition introduced the 3D chain rules. Of course, we do not claim that there may not be another general type of rules that should be added to ours. For instance, if you admit uniqueness of the solution (i.e. add the axiom of uniqueness), much work remains to be done in order to clarify all the techniques that have been proposed to express it in the form of U-resolution

rules. But one of the main questions in this regard is: should we accept rules for nets or for chains of subsets? In a sense, AICs based on subsets appear to be nets when we try to re-formulate them as chains of candidates; but they are mild kinds of nets. nrczt-chains, which have approximately the same solving power as the most complex AICs, prove that including subsets (Hinges, Almost Locked Sets) in chains is not necessary. On the other hand, general tagging algorithms that can solve anything correspond to unrestricted kinds of nets and they are not of much help for answering the question: which (mild) kinds of nets should we accept?

Viewed from the methodological standpoint, more than proposing a final set of resolution rules, our purpose was to set some minimal standard in the way one should systematise the definition of rules in natural language, formalise them in logic (or in equivalent graphical representations), implement them (as rules for an inference engine or in any other formalism that can be run on a computer, e.g. as strictly defined colouring or tagging resolution techniques) and test their validity and efficiency through the treatment of large collections of examples. It is our opinion that only this complete cycle may bring some clarity into the subject.

The second thing that has not been achieved in this book is the discovery of a *complete* resolution theory that would make no uniqueness assumption and that would not use Trial and Error. Our strongest resolution theory (L16 in the first edition, M28 in this second edition) cannot solve all the minimal puzzles that have a single solution. It can solve *almost all* these puzzles, but not *all* these puzzles, and increasing the maximal length of the chains would not help. Indeed, no set of resolution rules is known that would allow to solve such exceptionally complex puzzles as Easter Monster. Some kind of nets may be necessary. Defining very complex types of nrczt-nets is very easy; defining useful but mild ones is more difficult.

Finally, another related question is: does our strongest resolution theory (L13, or its weak extension L16 or the 3D theory M28) detect all the puzzles that have no solution? We have found no example that could not be detected. But this nevertheless leaves the question open. Underlying this question, there is a more general informal one, still open: is formulating *a priori* necessary and sufficient criteria on the existence of a solution (criteria that would only bear on the entries of a puzzle) easier than finding a complete resolution theory?

References

- [ANG 05-07] ANGUS J., Simple Sudoku, <http://www.angusj.com/sudoku/>, 2005-2007.
- [ARM 00-07] ARMSTRONG S., Sadman Software Sudoku, Solving Techniques, <http://www.sadmansoftware.com/sudoku/techniques.htm>, 2000-2007.
- [BAR 06] BARKER M., Sudoku Players Forum, Advanced solving techniques, post 362, *in* <http://www.sudoku.com/forums/viewtopic.php?t=3315>
- [BER xx] BERTHIER D., Solving Sudoku with the CLIPS or the JESS Inference Engine, forthcoming.
- [BER www] BERTHIER's Web pages: <http://www.carva.org/denis.berthier>
- [BRI 06] BRIDGES D. & VITA L., *Techniques of Constructive Analysis*, Springer, 2006.
- [BRO 06] BROUWER A., Solving Sudokus, <http://homepages.cwi.nl/~aeb/games/sudoku/>, 2006.
- [DAV 06] DAVIS T., The Mathematics of Sudoku, www.geometer.org/mathcircles/sudoku.pdf, 2006.
- [FEL 05] FELGENHAUER B. & JARVIS F., Enumerating possible Sudoku grids, <http://www.afjarvis.staff.shef.ac.uk/sudoku/sudgroup.html>, 2005.
- [FIT 69] FITTING M.C., *Intuitionistic Logic, Model Theory and Forcing*, North Holland, 1969.
- [GAR 03] GARSON J., Modal Logic, Stanford Encyclopedia of Philosophy, <http://plato.stanford.edu/entries/logic-modal/>, 2003.
- [HAN xx] HANSSEN V., <http://www.menneske.no/sudoku/eng/index.html>
- [HEN 06] HENDRICKS V. & SYMONS J., Epistemic Logic, Stanford Encyclopedia of Philosophy, <http://plato.stanford.edu/entries/logic-epistemic/>, 2006.
- [HIN 62] HINTIKKA J., *Knowledge and Belief: An Introduction to the Logic of the Two Notions*, Cornell University Press, 1962.

- [JAR 06] JARVIS F., Sudoku enumeration problems, <http://www.afjarvis.staff.shef.ac.uk/sudoku/>, 2006.
- [KRI 63] KRIPKE S.: Semantical Analysis of Modal Logic, *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, Vol. 9, pp. 67-96, 1963.
- [MEI 93] MEINKE K. & TUCKER J., eds., *Many-Sorted Logic and its Applications*, Wiley, 1993.
- [MEP xx] MEPHAM M., <http://www.sudoku.org.uk/>
- [MOS 07] MOSCHOVAKIS J., Intuitionistic Logic, Stanford Encyclopedia of Philosophy, <http://plato.stanford.edu/entries/logic-intuitionistic/>, 2006.
- [NEW 82] NEWELL A., The Knowledge Level, *Artificial Intelligence*, Vol. 59, pp 87-127, 1982.
- [RUS 05] RUSSELL E. & JARVIS F., There are 5472730538 essentially different Sudoku grids ... and the Sudoku symmetry group, http://www.afjarvis.staff.shef.ac.uk/sudoku/sudoku_group.html, 2005.
- [ROY xx] ROYLE G., Minimum Sudoku, <http://www.csse.uwa.edu.au/~gordon/sudokumin.php>
- [SPIF] Sudoku Players Forums, <http://www.sudoku.com/forums/index.php>
- [SPrF] Sudoku Programmers Forums, <http://www.setbb.com/sudoku/index.php?mforum=sudoku>
- [STE] STERTEN, <http://magictour.free.fr/sudoku.htm>
- [STU 06] STUART A., <http://www.scanraid.com>, 2006.
- [STU 07] STUART A., *The Logic of Sudoku*, Michael Mepham Publishing, 2007.
- [SUD] Sudopedia, http://www.sudopedia.org/wiki/Main_Page
- [SUF] Sudoku UK Forums, <http://www.sudoku.org.uk/cgi-bin/discus/discus.cgi?pg=topics>
- [WER 05-07] van der WERF R., Sudocue, Sudoku Solving Guide, <http://www.sudocue.net/guide.php>, 2005-2007.
- [YAT 02] YATO T. & SETA T., Complexity and completeness of finding another solution and its application to puzzles, IPSG SIG Notes 2002-AL-87-2, <http://www.imai.is.s.u-tokyo.ac.jp/~yato/data2/SIGAL87-2.pdf>, 2002.

Introduction

1. The Sudoku problem and the resolution methods

1.1. Statement of the Sudoku problem

Given a 9x9 *grid*, partially filled with *numbers* from 1 to 9 (the "entries" of the problem, also called the "clues" or the "givens"), complete it with numbers from 1 to 9 so that in every of the nine *rows*, in every of the nine *columns* and in every of the nine disjoint *blocks* of 3x3 contiguous *cells*, the following property holds:

- there is at most one occurrence of each of these numbers.

Although this defining property can be replaced by either of the following two, that are obviously equivalent to it, we shall stick to the first formulation, for reasons that will appear later (in chapter IV, section 1.2):

- there is at least one occurrence of each of these numbers,
- there is exactly one occurrence of each of these numbers.

Since rows, columns and blocks play similar roles in the defining constraints, they will naturally appear to do so in many other places and it is convenient to introduce a word that makes no difference between them: a *unit* is either a row or a column or a block. And we say that two cells *share a unit* if they are either in the same row or in the same column or in the same block (where "or" is non exclusive). We also say that these two cells are *linked*, or that they *see* each other. It should be noticed that this (symmetric) relation between two cells, whichever of the three equivalent names it is given, does not depend in any way on the content of these cells but only on their place in the grid; it is therefore a straightforward and quasi physical notion.

As can be seen from the definition, a Sudoku grid is a special case of a Latin Square. Latin Squares must satisfy the same constraints as Sudoku, except the condition on blocks. The practical consequences of this relationship between Sudoku and Latin Squares will appear throughout this book (and the logical relationship between the two theories will be fully clarified in chapter IV).

Figure 1 below shows the standard presentations of a *problem grid* (also called a *puzzle*) and of a *solution grid* (also called a *complete Sudoku grid*).

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

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

Figure 1. A puzzle (Royle17-3) and its solution

1.2. Resolution methods, candidates

The problem statement lists the constraints a solution grid must satisfy, i.e. it says *what* we want. It does not say anything about *how* we can obtain it: this is the job of the *resolution methods* and the *resolution rules* on which they are based (two notions that will be progressively refined in this introduction, until the final definition of a resolution rule can be given in chapter IV).

Different kinds of resolution methods can be used, depending mainly on whether they are primarily intended for a human solver or for a machine. For instance, one may be interested in efficient machine solving techniques; one will then choose machine representations of the problem specially adapted to efficient processing but that may be rather obscure for most human Sudoku players; one well-known technique of this kind relies on graph theory and maximal cliques. Although we do not neglect efficiency matters and have developed an automatised solver (SudoRules) implementing all the resolution rules described later in this book, we want to make it clear from the start that such questions are not our primary concern. Instead, our

approach is totally player oriented and we shall concentrate on formalising resolution rules that can be applied by a human solver equipped only with a sheet of paper and a pen (and quite a lot of patience) and on simulating them with the (most classical, i.e. rule based) techniques of Artificial Intelligence (AI).

	c1	c2	c3	c4	c5	c6	c7	c8	c9	
r1	<div>3 4 5 6 8 9</div>	<div>3 4 6 7 9</div>	<div>3 4 5 6 7 8 9</div>	<div>7 8 9</div>	<div>4 7 8 9</div>	<div>4 7 8 9</div>	<div>4 5 9</div>	1	2	r1
r2	<div>2 4 6 8 9</div>	<div>1 2 4 6 7 9</div>	<div>1 2 4 6 7 8 9</div>	<div>2 7 8 9</div>	3	5	<div>4 9</div>	<div>6 8 9</div>	<div>4 6 8 9</div>	r2
r3	<div>2 3 4 5 8 9</div>	<div>1 2 3 4 9</div>	<div>1 2 3 4 5 8 9</div>	6	<div>1 4 8 9</div>	<div>1 2 4 8 9</div>	<div>4 5 9</div>	7	<div>3 4 5 8 9</div>	r3
r4	7	<div>2 4 6 9</div>	<div>2 4 5 6 8 9</div>	<div>2 5 8 9</div>	<div>1 5 6 8 9</div>	<div>1 2 6 8 9</div>	3	<div>2 5 6 9</div>	<div>1 4 5 6 9</div>	r4
r5	<div>2 3 5 6 9</div>	<div>2 3 6 9</div>	<div>2 3 5 6 9</div>	4	<div>1 5 6 7 9</div>	<div>1 2 3 6 7 9</div>	8	<div>2 5 6 9</div>	<div>1 5 6 7 9</div>	r5
r6	1	<div>2 3 4 6 9</div>	<div>2 3 4 5 6 8 9</div>	<div>2 3 5 7 8 9</div>	<div>5 6 7 8 9</div>	<div>2 3 6 7 8 9</div>	<div>2 4 5 7 9</div>	<div>2 5 6 9</div>	<div>4 5 6 7 9</div>	r6
r7	<div>3 4 6 9</div>	<div>3 4 6 7 9</div>	<div>3 4 6 7 9</div>	1	2	<div>3 4 6 7 8 9</div>	<div>5 7 9</div>	<div>3 5 8 9</div>	<div>3 5 7 8 9</div>	r7
r8	<div>2 3 6 9</div>	8	<div>1 2 3 6 7 9</div>	<div>3 5 7 9</div>	<div>5 6 7 9</div>	<div>3 6 7 9</div>	<div>1 2 5 7 9</div>	4	<div>1 3 5 7 9</div>	r8
r9	<div>2 3 4 9</div>	5	<div>1 2 3 4 7 9</div>	<div>3 7 8 9</div>	<div>4 7 8 9</div>	<div>3 4 7 8 9</div>	6	<div>2 3 8 9</div>	<div>1 3 7 8 9</div>	r9
	c1	c2	c3	c4	c5	c6	c7	c8	c9	

Figure 2. Grid Royle17-3 of Figure 1, with the candidates remaining after the elementary constraints have been propagated

Given this choice, the process of solving a grid "by hand" is generally initialised by defining the "candidates" for each cell. For later formalisation, one must give a careful definition of this notion: *at any stage of the resolution process, candidates for a cell are the numbers that are not yet explicitly known to be impossible values for this cell*. At the start of the game, one possibility is to consider that any cell with

no input value admits all numbers from 1 to 9 as candidates (but more subtle initialisations can be considered).

Usually candidates for a cell are displayed in the grid as smaller and/or clearer letters in this cell (as shown in Figure 2); for better readability of such representations, the nine blocks will be marked by thick borders and each of the possible values will always be represented at the same relative place in each of the cells.

Then, the resolution process is a sequence of steps consisting of repeatedly applying "resolution rules" (some of which have become very classical and some of which may be very complex) of the general condition-action type: if some pattern (i.e. configuration) of cells, links, values and candidates for these cells is present on the grid, then carry out the action specified by the rule. Notice that any such pattern always has a purely "physical" part (which may be called its "physical" support), defined by the conditions on the cells and links between them, and an additional part, depending on the conditions put on the values and candidates in these cells.

According to the type of their action part, such rules can be classified into three categories:

- either assert the final value of a cell (when it is proven there is only one possibility left for it); there are very few rules of this type;
- or delete some candidate(s) (which we call the target values of the pattern) from some cell(s) (which we call the target cells of the pattern); as appears from a quick browsing of the available literature and as will be confirmed by this book, most resolution rules are of this type; they express specific forms of constraints propagation; their general form is: if such a pattern is present, then it is impossible for some value(s) to be in some cell(s) and the corresponding candidates must be deleted from them;
- or, for some very difficult grids, recursively make a hypothesis on the value of a cell, analyse its consequences and apply the eliminations induced by the contradictions thus discovered; techniques of this kind (named "recursive Trial and Error" or "recursive guess"), do not fit our condition-action form and are proscribed by purists (for the reason that, most of the time, they make solving the puzzle totally uninteresting); this book will show that they are very rarely needed if one admits complex chain rules.

It should be noted that all of the above resolution rules, whatever their type, do not assert that there is a solution. But for recursive Trial and Error, they may be interpreted from an operational point of view as: "from what is known in the current situation, do conclude that any solution, if there is any, will satisfy the following".

As one proceeds with resolution, candidates for each cell form a monotone decreasing set. With a little care, this remains true even when making hypotheses (i.e. resorting to recursive Trial and Error) cannot be avoided; in our "SudoRules" solver, for instance, in this case all candidates are explicitly relativised to finite sets of hypotheses (called "contexts") and monotonicity is thus maintained.

1.3. Elementary rules, Trial and Error, and their limitations

The four simpler constraints propagation rules (obviously valid) are the direct translation of the initial problem formulation into operational rules for managing candidates. We call them "the (four) elementary constraints propagation rules" (ECP):

- ECP(cell): "if a value is asserted for a cell (as is the case for the initial values), then remove all the other candidates for this cell";
- ECP(row): "if a value is asserted for a cell (as is the case for the initial values), then remove this value from the candidates for any other cell in the same row";
- ECP(col): "if a value is asserted for a cell (as is the case for the initial values), then remove this value from the candidates for any other cell in the same column";
- ECP(blk): "if a value is asserted for a cell (as is the case for the initial values), then remove this value from the candidates for any other cell in the same block".

The simpler assertion rule (also obviously valid) is called Naked-Single:

- NS: "if a cell has only one candidate left, then assert it as the only possible value of the cell".

Together with NS, the four elementary constraints propagation rules constitute "the (five) elementary rules".

A novice player may think that these five elementary rules express the whole problem and that applying them repeatedly is therefore enough to solve any puzzle. If such were the case, you'd probably never have heard of Sudoku, because it would amount to mere paper scratching. Anyway, as he gets stuck in situations where none of these rules remains applicable, he soon discovers that, except for the simplest grids, this is very far from being sufficient. The puzzle in Figure 1 is a simple illustration of how you get stuck if you only know and use the five elementary rules: the resulting situation is shown in Figure 2, in which none of these rules can be applied. As we shall see later (in chapter V), for this particular puzzle, there is an easy way to unblock it. But, as we shall also see, there are lots of puzzles that need rules of a much higher complexity in order to be solved. And this is why Sudoku has become

so popular: all but the easiest puzzles require a particular combination of neuron-titillating techniques and may even suggest the discovery of as yet unknown ones.

One general way out of the blocked situation described above is recursive Trial and Error: when stuck, one can start a systematic (depth first) exploration of the tree of possibilities, duly pruned by the propagation of elementary constraints (thus avoiding the exploration of obviously contradictory possibilities). Simplistic as this method may be, it has a major theoretical advantage, justifying that we keep it in our arsenal of techniques to solve a grid: it is guaranteed either to find a solution if there is (at least) one or to prove there is none.

The technical drawback is a great variance in the computation times (be it by a human or a machine), and this variability is unrelated to any sensible notion of difficulty of the grid (it depends mainly on the order chosen to explore the tree of possibilities, i.e. on chance). But the major drawback is its unrealistic character by any human standards: some puzzles would require exploring thousands of hypotheses, sometimes more than twenty levels deep. This is one of the reasons why this technique is anathemised by purists, the second being that using it generally makes the puzzle totally uninteresting.

Finally, we keep recursive Trial and Error in our arsenal, but we keep it as a last resort weapon, to be used when nothing else can be done; we shall see that with all the rules defined later in this book, such a strategy guarantees that, most of the time (in 99.7% of Royle's 36,628 reference cases defined below; in 97% of the randomly generated puzzles), this technique is not needed; and, when it is needed, we have found no case for which one level of hypothesis was not enough. With the rules for 3D chains introduced in this second edition, these percentages rise to over 99.99%.

1.4. Resolution rules and guiding principles for their formulation

Because the five elementary rules are not enough to solve any puzzle and recursive Trial and Error is not realistic from a human solver point of view, other resolution techniques must be devised.

Since Sudoku was invented, more or less complex resolution rules have been defined. They are based on the eliciting of various types of additional constraints, some of which may be non-obvious consequences of the problem statement.

Unfortunately, very often in the available literature, these rules, especially the most complex ones, are only illustrated by examples and their definitions remain

rather vague – which incurs both redundancy in the rules proposed by various authors and much uncertainty regarding their scopes of application. To check this, just look at some of the innumerable Web sites dedicated to Sudoku (more than sixty millions, only a few of which are listed in the bibliography at the end of this book).

It appears that this vagueness is due to the lack of a general guiding principle for stating the rules, and this in turn is due to the lack of a clear notion of the complexity of a rule.

Later in this book, we shall provide a precise and sometimes unusual rephrasing of most of the familiar rules. Besides the constraint of non-ambiguity, the general guiding principle we adopt can be considered a version of Occam's Razor. One can easily find some logical and some psychological support for it. It can be viewed from two complementary, but essentially equivalent, points of view:

- from the point of view of the preconditions of a rule: a rule should apply only in cases when simpler rules do not, i.e. its preconditions must be so specific as not to subsume those of simpler rules; but they must also be so general as to cover as many cases as possible; said otherwise, the scope of a rule must be extended as far as the logic underlying it allows;
- from the point of view of the conclusions of a rule (its action part): a rule should produce effective results, i.e. its conclusions should not be obtainable by simpler rules.

Of course, with the same reference to "simpler rules" in the two points of view, this principle relies on a definition of the complexity of a rule (or at least of the relative complexities of two rules). In this book, we build a hierarchy of rules progressively, based on:

- a distinction between three general classes of rules: subset rules, interaction rules and chain rules;
- a generalised notion of logical symmetry and associated representations;
- a second guiding principle: a rule obtained from another by some (generalised or not) logical symmetry must be granted the same logical complexity.

Given our objective of formalising the methods applied by a human solver, our second principle is highly debatable. There may be a great gap between abstract logical complexity and psychological complexity for the human solver. But the fact is that, in most cases, we have no idea of how psychological complexity can be measured. It is even doubtful that a given resolution rule could be given a psychological complexity measure independent of the "geographical" situation on the grid of the cells it applies to, i.e. independent of the most elementary symmetries inherent to

Sudoku (see chapter I); for instance, identical patterns of candidates on adjacent cells may be easier to see than the same patterns in distant cells; and this may also depend on individual psychological specificities. On the other hand, it is our hope that a partial relative ordering of our rules, based on their logical formulation and consistent with all the logical symmetries of the game, will serve as a reference for future measures of the psychological deviations from it. Moreover, there is a strong argument in favour of this principle, if one adopts the graphical representations and the extended Sudoku board (defined in chapter II) that makes obvious the equivalences associated to generalised symmetries.

Notice that we are looking for a partial complexity order relation on the set of resolution rules and that this is a very different task from trying to rank the puzzles based on some definition of the complexity of their resolution path (unless one defines the ranking of a puzzle as the complexity of the most complex rule necessary to solve it – not a very realistic ranking). Of course, there must be some relationship between a ranking of the puzzles and a partial complexity order on the set of resolution rules. Nevertheless, given a fixed set of rules, we shall see through examples that it can solve puzzles whose solution paths vary largely in complexity (whatever intuitive notion of complexity one adopts for the paths). In this book, we shall not tackle the problem of ranking the puzzles.

One last point can now be clarified. Everywhere in this book, a *resolution method* must be understood strictly as:

- a set of *resolution rules*,
- a *non-strict precedence ordering* among them. Non-strict means that two rules can have the same precedence (for instance, there is no reason to give a rule higher precedence than that obtained from it by transposing rows and columns or by any generalised symmetry).

As a consequence of this definition, several resolution methods can be based on the same set of rules with different partial orderings.

Moreover, to every resolution method one can associate a simple systematic procedure for solving a puzzle:

List all the resolution rules in a way compatible with their precedence ordering (i.e. among the different possibilities of doing so choose one)

Loop until a solution is found (or until it is proven there can be no solution)

- | Do until a rule applies effectively
 - | | Take the first rule not yet tried in the list
 - | | Do until its conditions pattern effectively maps to the grid
 - | | Try all possible mappings of the conditions pattern

```
|           |           End do
|           End do
|           Apply rule on selected matching pattern
End loop
```

In this context, a natural question arises: given a set of resolution rules, can different orderings lead to different puzzles being solved or unsolved? The answer is in the notion of confluence, to be explained in chapter XXII, where it will be shown that all the sets of rules introduced in this book have the *confluence property* and that the ordering of the rules is therefore irrelevant as long as we are only interested in solving puzzles; but it is of course very relevant when we also consider the efficiency of the associated method, e.g. the simplicity of the solution paths.

This abstract property has a very practical meaning for the player: it allows him/her not to be as systematic in the application of the rules as a machine would be, without running the risk of being blocked because of missing an elimination he could have done earlier in the resolution process.

2. The roles of logic and AI in this book

As its organisation shows, this book is centred on the Sudoku problem itself. Nevertheless, from the points of view of logic or AI, it can also be considered as a long exercise in either of these disciplines. So let us clarify the roles we grant them.

2.1. The role of logic

Throughout this book, the primary function of logic will be that of a compact notation tool for expressing the resolution rules in a non ambiguous way and expliciting the symmetry relationships between them (the simplest and most striking example of this is the set of rules for Singles in section V.2).

For better readability, the rules we introduce will always be formulated first in plain English and their validity will only be established by elementary non-formal means. The non mathematically oriented reader should therefore not be discouraged by the logical formalism. He can even skip chapters III and IV and the formal version of each rule that will usually follow its intuitive definition.

Moreover, in the very important case of the various types of chains we shall consider, the associated rules will always be expressed in an intuitive graphical formalism (partly inspired from existing informal representations one can find on Web forums, but also resolutely diverging from them when necessary); it will be

shown to be strictly equivalent to logical formulæ – so that the explicit writing of the corresponding logical formulæ will not even be needed.

The formalism we use relies effectively on the strictest formal logic and it would not be very difficult to use it as a basis for formal proofs. From a logical point of view and given the basic definitions of chapters III and IV, we consider that these formal proofs are no more than easy exercises for students in logic and we should not overload this book with them.

As a fundamental and practical consequence of our strict logical foundations, the natural symmetry properties of the Sudoku problem can be transposed into three formal meta-theorems allowing one to deduce systematically new rules from given ones (see chapters I and IV). This will allow us to introduce chain rules of completely new types ("hidden chains").

Finally, the other role assigned to logic is that of a mediator between the intuitive formulation of the resolution rules and their implementation in our AI program (SudoRules, or any other). This is a methodological point for AI (or software engineering in general): no program development should ever be started before precise definitions of its components are given (though not necessarily in strict logical form) – a common sense principle that is very often violated, even by those who consider it as obvious (this is the teacher speaking)!

2.2. The role of AI in this book

What role do we impart to AI in this book?

The resolution of each puzzle by a human solver needs a significant amount of time. Therefore, the number of puzzles that can be tested "by hand" against any resolution method is very limited. Simulating human solvers by AI will allow us to test tens of thousands of puzzles (see section 3.1 below). This will give us indications of the relative efficiency of different rules. It is not mere chance that the writing of this book and the development of our SudoRules solver occurred in parallel. Abstract definitions of relative complexities of rules were checked against our puzzle collections for their resolution times.

Productivity of new rules was tested as soon as they were introduced. Sometimes, it was very hard to find an example for a rule (such as rules for Naked-, Hidden- or Super-Hidden- Quadruplets). And sometimes, when no example could be found, it led to the conjecture, and then to the proof, that the supposedly new rule was subsumed by (i.e. could be reduced to) simpler ones.

As I said above, this book can also be considered as a (long) exercise in AI. Many computer science departments in universities have taken Sudoku as a basis for various projects. My personal experience is that it is a most welcome topic for a project in computer science or AI.

Actually, this is how my work on Sudoku started. But, on second thoughts, I realised that there might be something original in the point of view I had developed in the meanwhile and that it might concern a wider audience of "sudoku-ka". I therefore decided to soften the mathematical content (without sacrificing its logical foundations), in the hope that the result would be of interest to the union of the three populations instead of their intersection.

3. Examples and classification results

As can be seen from a fast browsing of this book, many examples are scattered in every chapter, making nearly a third of the content. This is not only because a book on Sudoku without a lot of examples would be like a French lunch without cheese. All our examples satisfy precise functions and their choice is anything but arbitrary. We decided that each example should:

- be as short as possible,
- illustrate a precise rule,
- prove that the rule it illustrates cannot be reduced to simpler ones (in this sense, *the detailed resolution paths given for all the examples must be considered as proofs of independence theorems*),
- originate from a real puzzle (this may seem an obvious constraint, but one can find examples on the Web where a partial situation is displayed with no indication as to its origin; for instance, one can find an example of an xy-chain of length 20; but I have never seen any real puzzle whose resolution needed to consider such a long xy-chain).

3.1. The origin of our examples

All our examples rely on three large puzzle collections:

- the first, hereafter named the Royle17 collection, has been assembled by the graph theorist Gordon Royle; it consists of the 36,628 known (non essentially equivalent) minimal grids with a unique solution; in this context, a grid is called minimal if it has seventeen entries and it has a unique solution; it is termed minimal because there is no known example of a grid with less than seventeen entries and a unique solution (it might be called absolutely minimal, but, as of the writing of this book, it

has not been proven that grids with fewer than seventeen entries cannot have a unique solution); in order to avoid confusion with the broader notion of minimality defined below, we call this case 17-minimal; grid number n in this collection is always named Royle17- n ;

- the second, hereafter named the Sudogen0 collection, consists of 10,000 puzzles randomly generated with the C generator `suexg` (see <http://magictour.free.fr/suexco.txt> for a description of the generation principles), with seed 0 for the random numbers generator; grid number n in this collection is always named Sudogen0- n ;

- the third, hereafter named the Sudogen17 collection, consists of 10,000 puzzles randomly generated with the same software as above, but using a different seed (17); grid number n in this collection is always named Sudogen17- n .

All the puzzles in our three test databases are *minimal* in the following sense (broader than the one used by Gordon Royle, in that they may have more than seventeen entries): they have a unique solution and any puzzle obtained from them by eliminating any one of their entries has more than one solution. For the Royle-17 case, this property results from the assembling choices of the collection; for the two Sudogen cases, the property is included in the principles of the generating software.

As for the specific examples chosen in this book to illustrate our rules, most of them draw upon the Royle17 collection. Occasionally, we also take examples from the Sudogen0 and Sudogen17 collections. The main reason for preferring the Royle-17 puzzle database is that showing that there is a 17-minimal puzzle for which a rule applies is a stronger result than just showing that this rule applies to some grid with no specific property (but this is not really important for the purposes of this book). And the main reason for using also randomly generated puzzles is for not relying on biased databases when we study global classification results.

3.2. *Uniform presentation of our examples*

If we displayed the full trace of the resolution process of an example puzzle, it would take several pages, most of which would describe obvious or uninteresting steps. Indeed, in the worst cases, starting from a 17-minimal puzzle, there are 64 (81-17) unknown values, which makes 576 (64x9) candidates to be eliminated and 64 values to be asserted, that is 640 steps. In order to skip some of these steps, we shall use the following five conventions.

Convention 1: obvious elementary constraint propagation rules ECP(cell), ECP(row), ECP(col) and ECP(blk) will never be displayed.

Convention 2: let us define the theory (i.e. the set of rules) *L1_0* as the union of all the above elementary (ECP) rules and the semi-elementary rules (Naked Singles and Hidden Singles – NS and HS) defined in chapter V. It can easily be checked that the final rules that apply to a puzzle always belong to *L1_0*, at least when these rules are given higher priority than more complex ones. Except in chapter V, where they are introduced, they will always be omitted from the end of the listing of the resolution path.

Convention 3: it can easily be checked that for most puzzles (due to the fact that they are minimal), the first rules that can be applied (not mentioning ECP) are semi-elementary propagation rules (NS and HS) whose direct effect is to add values. Except in chapter V, we shall replace the initial puzzle *P* by its "*L1_0* elaboration", i.e. by the puzzle obtained by adding to *P* all the values asserted by these first semi-elementary rules. Of course, this puzzle is no longer minimal (but it originates in a clearly defined minimal one).

Convention 4: since, for complex puzzles, this is sometimes still not enough for concentrating the listing on the rule we want to illustrate, when necessary we shall replace the original puzzle by the one obtained after applying rules of higher complexity than the semi-elementary ones (but of lower complexity than the one the example intends to illustrate). If *P* is a puzzle and *T* is a Sudoku Resolution Theory (i.e. a set of resolution rules), the puzzle obtained from *P* by applying repeatedly all the rules in *T* until none of them can be applied and keeping only the values thus asserted (i.e. discarding any information on the candidates eliminated) is called *the T elaboration of P*. Notice that the *T* elaboration of *P* is a real puzzle (although not minimal). It includes all the consequences of *T* on *P* that can be expressed as values.

Convention 5: nevertheless, since the *T* elaboration of *P* discards information that can be expressed only in terms of candidates, when it is taken as the new starting point and it is submitted to a new resolution process using an extension *T'* of *T* with rules more complex than those in *T*, the first rules that apply may still be rules from *T*. The worst situation is when the *T* elaboration of *P* coincides with *P* (i.e. no new value is produced by *T*). Such an extreme situation seems nearly impossible when we start with a minimal *P*, but it is often the case that the *T* elaboration of *P* coincides with the elaboration of *T* by a much simpler theory than *T* (*L1_0* for instance), i.e. all the rules in *T* do not produce more values than the rules in *L1_0*. A situation that arises rather frequently is when many candidates are deleted by *T* but few values are asserted; this makes the strategy described in convention 4 more or less inefficient. Conversely, the ideal case is when all results produced by the elaboration process are subsumed by the values asserted (this is of course the case when *T* is *L1_0*, but this condition is not necessary). In this case, we may be certain that the first rule applicable in *T'* will not be in *T* (still not mentio-

ning ECP). Therefore the T elaboration of P illustrates a situation in which the patterns used by the rules added to T in order to obtain T' are immediately visible (after rules from ECP have been applied). Whenever possible, our examples will be chosen from such ideal situations.

With the above conventions, only the interesting part of the resolution process of the initial minimal puzzle will be displayed. But, even with the economy resulting from the above conventions, some traces of resolution processes may be very long. Most of the time, we shall select examples with short traces, but this will not always be possible and, in order to help you keep in mind that these examples are somewhat exceptional and the possibilities are much more varied, especially for rules relative to long chains, longer examples will also be given from time to time (see for instance chapters XV or XVIII).

All our examples will respect the following uniform format (Figure 3), except that, due to page setting constraints, the figure displaying the puzzle, the introductory text and/or the listing may be inverted.

After an introductory text, explaining the purpose of the example and/or commenting on some particular point, a row of three grids is displayed: the original puzzle (always a minimal puzzle taken from one of our three collections), its indicated elaborated version and its solution. This is followed by a listing of the resolution path.

The first line of the resolution path indicates by which resolution theory T1 (L3_0 in the above example) the elaborated puzzle displayed in the second grid was obtained and (inside parenthesis) to which simpler elaboration T0 (L1 in the above example) it is equivalent. Both statements are useful: the second indicates the minimal theory T0 necessary (which is also the part of T1 effectively used) to get the elaborated version of P; the first indicates that P cannot be solved in (the stronger) T1 alone.

This first line of the resolution path also indicates in which theory T (stronger than T1) the resolution path that follows is obtained (L3_0+XY3 in the example). Most of the time, T will be of type T1+R, i.e. will be obtained from T1 by adding a single rule, R. Thus, *by showing that P can be solved in T1+R but not in T1 alone, the example proves that the R rule is not subsumed by (i.e. cannot be reduced to) the set of rules in T1*. This is a very important property, because the converse would mean that, given T1, R is useless. To express it, *we write that P belongs to [T1]+R*. Every example of this form can thus be considered as an *independence theorem*.

<Introductory text>




Figure x. Puzzle Royle17-186, its L1 elaboration and its solution

Resolution path in L3_0+XY3 for the L3_0 (or L1) elaboration of Royle17-186:
xy3-chain {n8 n5}r2c8 — {n5 n4}r3c7 — {n4 n8}r4c7 ==> r4c8 ≠ 8
... (Naked-Singles and Hidden-Singles)

Figure 3. Uniform format of our examples

Then comes the resolution path proper. Each step in the resolution path is the application of a well defined resolution rule in *T* to the precisely decribed and purely factual situation resulting from the previous rule applications; the resolution path is thus a proof of the solution within theory *T* (where "proof" is meant in the strict mathematical logic sense).

Starting from the elaborated version of the puzzle, only the sequence of non-obvious resolution steps will be displayed. Each line in the sequence consists of the name of the rule applied, followed in order by: the description of how the condition part is satisfied (how the rule is "instantiated"), the "==" sign, the conclusion(s) allowed by the "action" part. Details of the "nrc notation" used for the condition part will be described progressively with each rule we study. The conclusion part is always either that a candidate can be eliminated, symbolically written as here: r4c8 ≠ 8, or that a value must be asserted, written symbolically as e.g. r4c8 = 8. When the same rule instantiation justifies several conclusions, they will be written on the same line, separated by commas: e.g. r4c8 ≠ 8, r5c8 ≠ 8.

The rule(s) of interest in the path will be displayed in bold characters. In the above example, there is only one step, the application of the XY3 rule to some clearly described pattern of cells and values.

The trace of a resolution path will always end with the line "... (Naked-Singles and Hidden-Singles)" or something similar to remind you of convention C2.

The above conventions present the following advantages for you reader, if you want to try the examples. First, you may skip the uninteresting parts and start from the central puzzle; it is not minimal, but it is a real puzzle. Then, most of the time, the first rule you will have to apply (after the obvious ECP) will be the one studied in the chapter of the example; and, when this is not the case, the steps you will have to apply before you reach this rule will be clearly indicated so that you can easily reproduce them until you reach the pattern of interest. Our examples are designed to help you detect these patterns but they suppose an active participation on your part: only the initial values are displayed; it is left to you to apply ECP and the other rules of the resolution path to reach the desired situation. Occasionally, the detailed situation at some point in the resolution path (i.e. all the values and candidates present at this point) will be displayed so that you can directly check the presence of the pattern under discussion, but, due to place constraints, this cannot be systematic.

Finally, note that all the traces of resolution processes given in this book were obtained with version 13 of our SudoRules solver (with some hand editing for a shorter and cleaner appearance), run in the CLIPS 6.24 environment (more on this in chapter XXI).

3.3. Classification results

The available literature on resolution rules has concentrated on isolated examples illustrating specific rules but systematic studies on large collections of puzzles are lacking. To palliate this deficiency, and as a concrete counterpart to our abstract ideas about rules classification, detailed numerical results about the number of grids solved by each type of rule will be given in chapter XXI. As a justification of these results (that would need far too many pages to be published in paper form), detailed lists of the corresponding grids are available on the author's Web pages (permanent address: <http://carva.org/denis.berthier>), together with lots of additional material.

Conclusion

What has been achieved

In this conclusion, I'd like first to highlight a few facets of what has been achieved in this book, from four complementary overlapping points of view.

1) ***From the point of view of the Sudoku addict***, the most striking results should be the following.

By formalising the familiar resolution rules, we have eliminated from some of them the ambiguities that plagued their usual presentations and we have clarified their scopes. For instance, we have shown that none of the two usual formulations of Triplets (or Quadruplets) – neither the "strict" nor the "comprehensive" – was the best; our reformulation shows how close Triplets (and Quadruplets) are to xy-chains but also why they cannot be reduced to them (nor to our stronger xyt- or xyzt-chains).

We have fully clarified the symmetry relationships that exist between Naked, Hidden and Super-Hidden Subset rules (the latter being classically known as X-Wing, Swordfish, Jellyfish, Squirmbag and other existent or non existent "fishy patterns"). Such relationships had already been partly mentioned on some Web pages, but never in the systematic way we have dealt with them here. As a result, we have naturally found the proper formulations for the Hidden and Super-Hidden Subset rules and we have proven that no other subset rule of such type can exist.

More generally, we have proven three meta-theorems that automatically produce new resolution rules (their Hidden or Super-Hidden counterparts) from existing ones.

With the introduction of new graphical representations of the problem in auxiliary 2D spaces (mainly row-number and column-number spaces), we have shown that the Hidden or Super-Hidden counterparts of well-known patterns (subsets or chains) defined in natural row-column space can be detected as easily as their originals, because in these new spaces they look like the originals do in the standard representation. We have thus extended the resolution power of such patterns.

We have defined a (non strict) complexity hierarchy between the resolution rules, compatible with their logical symmetry relationships.

We have proven certain theorems showing that particular resolution rules can be formally reduced to simpler ones in the above hierarchy (e.g. "xy-chains of length 3 are subsumed by Naked-Triplets plus XY-Wing").

We have evaluated the strength of each rule by the proportion of new puzzles ("new" according to the above hierarchy) its introduction allows to solve and, for the first time, such an evaluation has been based on a large collection of puzzles (more than 56,000), with detailed results available online.

We have given chain rules a major place in this book (they occupy more than half of it), because they are the main tool for dealing with hard puzzles but they remain the subject of much confusion. We have introduced a general conceptual framework (including the notion of a target not belonging to the chain) for dealing with all conceivable types of chains and we have applied it systematically to all the types we have defined. In particular, we have introduced an intuitive graphical language of patterns for specifying chains and their targets, abstracting them from any irrelevant facet (such as a link being a row or a column or a block), and we have shown that these patterns are equivalent to logical formulæ. In our framework, the confusing notions of "chain of inferences", "weak link" and "strong link" are never used; our chains are well defined patterns of candidates, cells and links.

We have chosen the simplest kind of homogeneous chains, the xy-chains, as our main type of chains, with all the other chains being intuitive generalisations of them.

We have proven that xy-chains and c-chains should have no loops (with the practical consequence that searching for these chains becomes simpler for both a human and a computer).

By pushing the underlying logic of xy-chains to its limits, we have introduced xyt-chains and shown that they are a very natural and powerful generalisation of xy-chains. We have also defined another independent generalisation, the xyz-chains, and combined it with the previous one to get the xyzt-chains.

With each type of chain in natural row-column space, we have associated two new types of chains, their hidden counterparts in row-number and column-number spaces, and, in particular, we have shown the unifying power of the hidden xy-chains. All these chains can be spotted in either of the 2D representations, using our Extended Sudoku Board.

We have also proven theorems allowing to combine our various types of homogeneous chains to build heterogeneous, more complex, ones.

In this second edition, we have generalised the above 2D chains, introduced their fully super-symmetric 3D versions (the nrc-, nrct-, nrcz- and nrczt- chains) and shown that all these types of chains can still be considered in a natural way as various generalisations of the basic xy-chains.

We have produced a multiplicity of well chosen examples proving that particular resolution rules cannot be reduced to simpler ones (in the above hierarchy).

In particular, we have proven that, using only the types of chain rules introduced in the first edition, it is necessary to consider chains of length more than thirty if we want to have a chance of solving all the randomly generated puzzles without resorting to Trial and Error or assuming the uniqueness of a solution.

In the first edition, we had exhibited a set of resolution rules (L13) based on only 2D chains that could solve 97% of the randomly generated minimal puzzles and 99,67% of the 36,628 17-minimal puzzles in the famous Royle database (without resorting to Trial and Error or to an assumption of uniqueness).

In this second edition, we have also exhibited a set of resolution rules (M5) that can solve more than 99% of the randomly generated minimal puzzles using only 3D chains of length no more than five and a set of resolution rules (M7) that can solve 99.9% of these puzzles using only 3D chains of length no more than seven. As psychologists consider that human short term memory has size seven plus or minus two, this means that a human being using these rules should be able to solve almost any puzzle without any computer assistance (but still with some patience). It should be noticed that these chains do not include subsets (Hinges or Almost Locked Sets or grouped chains), contrary to the currently popular chains (Nice Loops or Alternating Inference Chains), thus avoiding a potential source of exponential behaviour.

2) *From the point of view of mathematical logic*, our most obvious result is the introduction of a strict formalism allowing a clear distinction between the straightforward *Sudoku Theory* (that merely expresses the constraints defining the game) and all the possible *Sudoku Resolution Theories* formulated in terms of condition-action rules (that may be put to practical use for solving puzzles). We have given a clear logical definition of what a "resolution rule" is, as opposed to any logically valid formula. With the notions of a resolution theory T and a resolution path (which is merely a proof of the solution within T , in the sense of mathematical logic), we have given a precise meaning to the widespread but as yet informal idea that one wants a "pure logic solution". This leads to both sound foundations and intuitive justifications for our resolution theories, exhibiting the following facets and consequences.

We have established a clear logical (epistemic) status for the notion of a candidate – a notion that is quasi universally introduced for stating the resolution rules but that does not pertain *a priori* to Sudoku Theory and that is usually used only from an intuitive standpoint. Moreover, we have shown that the epistemic operator that must appear in any proper formal definition of this notion can be "forgotten" in practice when we state the resolution rules and that this notion can be considered as primary, provided that we work with intuitionistic (or constructive) logic instead of standard logic (this is not a restriction in practice). Notice that this whole approach can be extended to any game that is based on techniques of progressive elimination of candidates.

We have also defined what a "resolution method" based on a resolution theory is and we have shown that all the resolution theories introduced in this book have the very important *confluence property*, allowing any ordering to be superimposed on their resolution rules without changing their overall resolution capacity. As a major practical consequence, in any software implementation of a resolution theory into a resolution method (e.g. in our SudoRules solver), we may take advantage of any convenient ordering of the rules.

The natural symmetries of the Sudoku problem have been expressed as three general meta-theorems asserting the validity of resolution rules obtained by some simple transformations of those already proven. These meta-theorems have been stated and proven both intuitively and formally. As a first example of how these meta-theorems can be used in practice, we have exhibited a precise relationship between well known (Naked and Hidden) Subset rules with what we call their Super-Hidden counterparts (the famous "fishy patterns") and we have proven some form of completeness of the set of known Subset rules. As a second example of how these meta-theorems can be used in practice, we have defined entirely new types of chain rules, hidden chains of various types, and shown their unifying power.

We have also devised a direct proof of the existence of a simple and striking relationship between Sudoku and Latin Squares: *a block-free resolution rule (i.e. a rule that does not mention blocks or squares) is valid for Sudoku if and only if it is already valid for Latin Squares*. Notice that it does not seem one can prove this result by using only the general methods one would expect to see used in such cases: either the interpolation theorem or the techniques of Gentzen's sequent calculus.

Finally, we have provided a very intuitive example of how difficult it may be to transform a theory formulated in terms of (a few and simple) constraints into a set of "production rules" (or condition-action rules). This also shows that, although the given constraints on rows and columns on the one hand and the constraint on blocks on the other hand can be formulated as axioms with independent predicates, many of the condition-action rules necessary to operationalise them do mix these predicates. This mixture is most visible in the 3D (or fully super-symmetric) chain rules.

3) *From the point of view of Artificial Intelligence* (AI), the following should be stressed.

Sudoku is a wonderful example for AI teachers. It has simpler rules and is more accessible for student projects than games such as chess or go, but it is much more complex and exciting than the usual examples one can find in AI textbooks (Tic-Tac-Toe, Hanoi Towers, Monkey and Bananas, Bricks World, ...). It easily suggests lots of projects based on the introduction and the formalisation of new types of rules since no complete set of resolution rules is known (see below).

Sudoku provides a very good illustration of a basic software engineering principle: never start writing a program (or a knowledge base for an inference engine) unless you have a non-ambiguous specification for it. The logic of certain resolution rules is so subtle that the least deviation from it (e.g. forgetting that a target of a chain may not belong to it or that loops are not allowed in a chain) has "catastrophic" consequences (typically some puzzles being falsely claimed to have no solution). This can also be considered a nice illustration of Newell's classical distinction (introduced in [New 82]) between the *knowledge level*, here assimilated to the logical formulation, in which knowledge must be formulated in a "declarative" form independent of any processes that might be applied to it, and the *symbol level*, here assimilated to the set of rules in the CLIPS or the JESS language, in which the knowledge level is *operationalised* in a form that may depend (and does largely depend in practice) on the specificities of the knowledge representation and processing tools used to implement it. Although this book does not tackle this point, there are many ways a given resolution rule (as formulated in logic) can be implemented as a production rule in an inference engine, which are more or less related to the different ways it may be used in practice.

Sudoku is also a wonderful testbed for the inference engine chosen to run the knowledge base of resolution rules. The logic of the set of rules is so intricate that many functionalities of the inference engine are stringently tested, which is how we discovered a longstanding undocumented bug in the management of saliences in JESS. With long series of puzzles to solve, memory management can also be a problem (as it is in CLIPS).

The previous topic is related to a crucial problem of AI, both practical and epistemological: how can one be sure that the system does what it is intended to do? Although this is already a very difficult question for classical software (i.e. mainly procedural, notwithstanding the object oriented refinements), it is much worse for AI, for two main reasons. Firstly, the logic underlying a knowledge base is generally much more complex than a set of procedures is (otherwise it would probably be much better to solve the problem with procedural techniques) and secondly an inference engine is a very complex piece of software and debugging it is very difficult.

As a general result, using AI to prove theorems (although this has always been and remains one of the key subtopics of the domain) may make mathematicians and logicians very suspicious.

As a specific result, all the independence theorems that have been proven in this book rely on the correctness of the inference engine we used for our computations. They do not depend on this correctness when we assert that "theory T allows the following resolution path for puzzle P", since the validity of any particular path can easily be checked "by hand", whichever method was used to generate it. But these independence results depend on this correctness any time we state that a particular theory is not able to find a solution for some puzzle P. This is why all our computations have finally been done with CLIPS instead of JESS despite the fact that CLIPS regularly gets lost in memory management problems and computation times on long series of puzzles grow astronomical. But the fact that, due to its problem with the management of saliences², JESS misses some inferences on simpler rules, even though this may be infrequent, disqualifies it as a theorem prover in our case (so much so that missed inferences also vary from one version to the next). Obviously, this does not prove that CLIPS is bug free. The only certain conclusions are that, using the same knowledge base, CLIPS solved (a few) more puzzles than JESS and never returned any spurious message of type "this puzzle has no solution".

Of course, these independence theorems also rely on the correctness of the knowledge base we have developed to implement all the rules defined in this book.

² It seems that this bug has been corrected in the latest release of JESS (71b1) available as of the publication of this second edition. But we haven't carried out systematic tests.

In this respect, I would like to make three points:

- firstly, thanks to the monotonicity of the facts base (each rule can only add values or not-candidates), a confluence theorem has been proven, which guarantees that there cannot be unwanted interactions between the rules;

- secondly, each individual rule has received the same systematic treatment: it has been stated in plain English, with great care being taken for not forgetting any conditions; it has then been proven, still in plain English, in rather straightforward steps that can be checked by anybody; its formulation in multi-sorted logic (or, for the chain rules, in our equivalent graphical formalism) has been shown to be the direct transliteration of the English one; in turn, the CLIPS or the JESS formulation is the direct transliteration of the logical one, thus minimising the possibilities of discrepancies between successive steps in the development;

- thirdly, the current release of our solver (SudoRules 13) has been tested on more than 56,000 puzzles known to have a unique solution (and this produced the classification results in chapters XXI and XXIII); whereas an error in a rule that would illegitimately eliminate candidates leads very rapidly to the claim that a puzzle has no solution (this allowed the detection of a few subtle bugs in the first version of SudoRules), it is noticeable that SudoRules has not yet produced an incorrect result in the CLIPS environment.

4) *Finally, considering the (currently very fashionable) notion of complexity*, problems that can be stated in simple terms but that need a complex solution are not anything new, although the general idea may remain obscure for the novice thinker. Among the most famous problems of this type, you have certainly heard of the four-colour problem in graph theory or Fermat's conjecture in arithmetic (now known as the "Fermat-Wiles Theorem", since it has been proven recently by Andrew Wiles, more than three centuries after its formulation by Fermat). But proofs of these theorems are not really accessible to the non-mathematician and the type of complexity hidden behind the problem statement therefore remains very abstract. With the notion of deterministic chaos, the second part of the twentieth century has uncovered a new type of complexity: some dynamical systems ruled by very simple equations may have very complex trajectories and two neighbouring points may follow quickly diverging paths – but this also remains a little mysterious if you do not have a minimum mathematical background.

On the contrary, with the popular game of Sudoku, you can get a feeling of another type of complexity, computational complexity (how this is related to the previous ones remains an interesting but very difficult question). Sudoku is known to be NP-complete, i.e., to state it very informally (and somewhat incorrectly), when

we consider grids of increasing sizes, resolution times grow faster than any deterministic polynomial algorithm. As you will never try to solve a Sudoku puzzle on a 100x100 grid (unless you have unlimited free access to a psychoanalyst), this may also remain an abstract definition. There is nevertheless a difference with the previous examples: you can already get a foretaste of the underlying complexity with the standard 9x9 puzzles (e.g. by comparing them to their homologues on 4x4 grids, the so-called Sudokids).

The Sudoku problem is defined by four very simple constraints, immediately understood by everybody in terms of "single occupancy" of a cell and of "mutual exclusion" in a row, a column or a block. For a classically formatted mind, it is therefore natural to think that any puzzle can easily be proven to have no solution or be solved by a finite set of simple operational resolution rules of the condition-action type: "in such a situation carry out such an action (assert a value or eliminate a candidate)". And this idea can only be reinforced if you consider the majority of puzzles published in newspapers. But the independence results proven in this book through a multiplicity of examples have shown that very complex resolution rules are indeed needed.

What this book has shown then, in both intuitive and logically grounded ways, is that writing a set of operational rules for solving an apparently very simple constraints propagation problem may be a very complex task. (Indeed, notwithstanding their overall complexity, the rules that have been defined in this book do not even form a complete resolution theory.) Moreover, as all the NP-complete problems are equivalent (through polynomial transformations) and some of them have lots of practical applications, such as the famous travelling salesman, dealing with the apparently futile example of Sudoku may provide intuitions on problems that seem to be unrelated.

What has been partly achieved (from the point of view of AI)

In the introduction, we said we wanted a set of rules that would simulate a human solver and that could explain each of the resolution steps. The explanations produced by SudoRules are largely illustrated by the listings given in this book; they are sufficiently explicit once you know the definitions of our rules and it would be easy work to make them still more explicit for those who do not know them; but we do not consider this as a very exciting topic. As for the solver facet, SudoRules does simulate a human solver, a particular kind of player who would try all our rules systematically (ordered according to their complexity) on all of their potential instantiations.

Is it likely that any human solver would proceed in such a systematic way? He may prefer to concentrate on a part of the puzzle and try to eliminate a candidate from a chosen cell (or group of cells). What may be missing then in our system is a "strategic" knowledge level: when should one look for such or such pattern? But I have no idea of which criteria could constitute a basis for such strategic knowledge; moreover, as far as I know, whereas there is a plethora of literature on resolution techniques (often misleadingly called strategies), nothing has ever been written on the ways they should be used, i.e. on what might legitimately be called strategies.

To say it otherwise, we do have a strategic level: searching for the different patterns in their order of increasing complexity. Notice that there is already more strategy in this than proposed by most of the books or Web pages on the subject. The question then is: can one define a better (or at least another) strategy? Well, the rules in this book (and the corresponding software SudoRules) are there; you can use them as a basis for further analysis of alternative strategies. One of the simplest ways to do so is to modify the complexity measure and the ordering we have defined on the set of rules. For instance, using psychological analyses, one could break or relax the symmetry constraints we have adopted.

What was not our purpose and has not been achieved; open questions

The first thing that has not been done in this book is a review of all the advanced rules that have been proposed, mainly on the Web (under varying names, in more or less clear formulations, with more or less defined scopes). The list would be too long and moreover it is regularly increasing. The best place to get an idea on this topic is in the recent book by Andrew C. Stuart ([STU 07]) or on the Web, in particular in the Sudoku Players Forum:

<http://www.sudoku.com/forums/viewtopic.php?t=3315>

(with the problem that chains are often considered as "chains of inferences" instead of patterns and they are sometimes classified according to absurd criteria).

Instead, our two main purposes in this regard were to take advantage of the symmetries of the game in a systematic way and to isolate a limited number of rule types, with rules definitions extended as far as the arguments used in their proofs allowed: this is how we introduced xyt-, xyz- and xzyt- chain rules (on the basis of xy-chain rules) and their hidden counterparts (on the basis of our general meta-theorems); this is also how this second edition introduced the 3D chain rules. Of course, we do not claim that there may not be another general type of rules that should be added to ours. For instance, if you admit uniqueness of the solution (i.e. add the axiom of uniqueness), much work remains to be done in order to clarify all the techniques that have been proposed to express it in the form of U-resolution

rules. But one of the main questions in this regard is: should we accept rules for nets or for chains of subsets? In a sense, AICs based on subsets appear to be nets when we try to re-formulate them as chains of candidates; but they are mild kinds of nets. nrczt-chains, which have approximately the same solving power as the most complex AICs, prove that including subsets (Hinges, Almost Locked Sets) in chains is not necessary. On the other hand, general tagging algorithms that can solve anything correspond to unrestricted kinds of nets and they are not of much help for answering the question: which (mild) kinds of nets should we accept?

Viewed from the methodological standpoint, more than proposing a final set of resolution rules, our purpose was to set some minimal standard in the way one should systematise the definition of rules in natural language, formalise them in logic (or in equivalent graphical representations), implement them (as rules for an inference engine or in any other formalism that can be run on a computer, e.g. as strictly defined colouring or tagging resolution techniques) and test their validity and efficiency through the treatment of large collections of examples. It is our opinion that only this complete cycle may bring some clarity into the subject.

The second thing that has not been achieved in this book is the discovery of a *complete* resolution theory that would make no uniqueness assumption and that would not use Trial and Error. Our strongest resolution theory (L16 in the first edition, M28 in this second edition) cannot solve all the minimal puzzles that have a single solution. It can solve *almost all* these puzzles, but not *all* these puzzles, and increasing the maximal length of the chains would not help. Indeed, no set of resolution rules is known that would allow to solve such exceptionally complex puzzles as Easter Monster. Some kind of nets may be necessary. Defining very complex types of nrczt-nets is very easy; defining useful but mild ones is more difficult.

Finally, another related question is: does our strongest resolution theory (L13, or its weak extension L16 or the 3D theory M28) detect all the puzzles that have no solution? We have found no example that could not be detected. But this nevertheless leaves the question open. Underlying this question, there is a more general informal one, still open: is formulating *a priori* necessary and sufficient criteria on the existence of a solution (criteria that would only bear on the entries of a puzzle) easier than finding a complete resolution theory?

References

- [ANG 05-07] ANGUS J., Simple Sudoku, <http://www.angusj.com/sudoku/>, 2005-2007.
- [ARM 00-07] ARMSTRONG S., Sadman Software Sudoku, Solving Techniques, <http://www.sadmansoftware.com/sudoku/techniques.htm>, 2000-2007.
- [BAR 06] BARKER M., Sudoku Players Forum, Advanced solving techniques, post 362, *in* <http://www.sudoku.com/forums/viewtopic.php?t=3315>
- [BER xx] BERTHIER D., Solving Sudoku with the CLIPS or the JESS Inference Engine, forthcoming.
- [BER www] BERTHIER's Web pages: <http://www.carva.org/denis.berthier>
- [BRI 06] BRIDGES D. & VITA L., *Techniques of Constructive Analysis*, Springer, 2006.
- [BRO 06] BROUWER A., Solving Sudokus, <http://homepages.cwi.nl/~aeb/games/sudoku/>, 2006.
- [DAV 06] DAVIS T., The Mathematics of Sudoku, www.geometer.org/mathcircles/sudoku.pdf, 2006.
- [FEL 05] FELGENHAUER B. & JARVIS F., Enumerating possible Sudoku grids, <http://www.afjarvis.staff.shef.ac.uk/sudoku/sudgroup.html>, 2005.
- [FIT 69] FITTING M.C., *Intuitionistic Logic, Model Theory and Forcing*, North Holland, 1969.
- [GAR 03] GARSON J., Modal Logic, Stanford Encyclopedia of Philosophy, <http://plato.stanford.edu/entries/logic-modal/>, 2003.
- [HAN xx] HANSSEN V., <http://www.menneske.no/sudoku/eng/index.html>
- [HEN 06] HENDRICKS V. & SYMONS J., Epistemic Logic, Stanford Encyclopedia of Philosophy, <http://plato.stanford.edu/entries/logic-epistemic/>, 2006.
- [HIN 62] HINTIKKA J., *Knowledge and Belief: An Introduction to the Logic of the Two Notions*, Cornell University Press, 1962.

- [JAR 06] JARVIS F., Sudoku enumeration problems, <http://www.afjarvis.staff.shef.ac.uk/sudoku/>, 2006.
- [KRI 63] KRIPKE S.: Semantical Analysis of Modal Logic, *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, Vol. 9, pp. 67-96, 1963.
- [MEI 93] MEINKE K. & TUCKER J., eds., *Many-Sorted Logic and its Applications*, Wiley, 1993.
- [MEP xx] MEPHAM M., <http://www.sudoku.org.uk/>
- [MOS 07] MOSCHOVAKIS J., Intuitionistic Logic, Stanford Encyclopedia of Philosophy, <http://plato.stanford.edu/entries/logic-intuitionistic/>, 2006.
- [NEW 82] NEWELL A., The Knowledge Level, *Artificial Intelligence*, Vol. 59, pp 87-127, 1982.
- [RUS 05] RUSSELL E. & JARVIS F., There are 5472730538 essentially different Sudoku grids ... and the Sudoku symmetry group, http://www.afjarvis.staff.shef.ac.uk/sudoku/sudoku_group.html, 2005.
- [ROY xx] ROYLE G., Minimum Sudoku, <http://www.csse.uwa.edu.au/~gordon/sudokumin.php>
- [SPIF] Sudoku Players Forums, <http://www.sudoku.com/forums/index.php>
- [SPrF] Sudoku Programmers Forums, <http://www.setbb.com/sudoku/index.php?mforum=sudoku>
- [STE] STERTEN, <http://magictour.free.fr/sudoku.htm>
- [STU 06] STUART A., <http://www.scanraid.com>, 2006.
- [STU 07] STUART A., *The Logic of Sudoku*, Michael Mepham Publishing, 2007.
- [SUD] Sudopedia, http://www.sudopedia.org/wiki/Main_Page
- [SUF] Sudoku UK Forums, <http://www.sudoku.org.uk/cgi-bin/discus/discus.cgi?pg=topics>
- [WER 05-07] van der WERF R., Sudocue, Sudoku Solving Guide, <http://www.sudocue.net/guide.php>, 2005-2007.
- [YAT 02] YATO T. & SETA T., Complexity and completeness of finding another solution and its application to puzzles, IPSG SIG Notes 2002-AL-87-2, <http://www.imai.is.s.u-tokyo.ac.jp/~yato/data2/SIGAL87-2.pdf>, 2002.