

# TECHNIQUES FOR SOLVING SUDOKU PUZZLES

ERIC C. CHI\* AND KENNETH LANGE\*†

**Abstract.** Solving Sudoku puzzles is one of the most popular pastimes in the world. Puzzles range in difficulty from easy to very challenging; the hardest puzzles tend to have the most empty cells. The current paper explains and compares three algorithms for solving Sudoku puzzles. Backtracking, simulated annealing, and alternating projections are generic methods for attacking combinatorial optimization problems. Our results favor backtracking. It infallibly solves a Sudoku puzzle or deduces that a unique solution does not exist. However, backtracking does not scale well in high-dimensional combinatorial optimization. Hence, it is useful to expose students in the mathematical sciences to the other two solution techniques in a concrete setting. Simulated annealing shares a common structure with MCMC (Markov chain Monte Carlo) and enjoys wide applicability. The method of alternating projections solves the feasibility problem in convex programming. Converting a discrete optimization problem into a continuous optimization problem opens up the possibility of handling combinatorial problems of much higher dimensionality.

**Key words.** Backtracking, Simulated Annealing, Alternating Projections, NP-complete, Satisfiability

**1. Introduction.** As all good mathematical scientists know, a broad community has contributed to the invention of modern algorithms. Computer scientists, applied mathematicians, statisticians, economists, and physicists, to name just a few, have made lasting contributions. Exposing students to a variety of perspectives outside the realm of their own disciplines sharpens their instincts for modeling and arms them with invaluable tools. In this spirit, the current paper discusses techniques for solving Sudoku puzzles, one of the most popular pastimes in the world. One could make the same points with more serious applications, but it is hard to imagine a more beguiling introduction to the algorithms featured here. Sudoku diagrams are special cases of the Latin squares long familiar in experimental design and, as such, enjoy interesting mathematical and statistical properties [1]. The complicated constraints encountered in solving Sudoku puzzles have elicited many clever heuristics that amateurs use to good effect. Here we examine three generic methods with broader scientific and societal applications. The fact that one of these methods outperforms the other two is mostly irrelevant. No two problem categories are completely alike, and it is best to try many techniques before declaring a winner.

The three algorithms tested here are simulated annealing, alternating projections, and backtracking. Simulating annealing is perhaps the most familiar to statisticians. It is the optimization analog of MCMC (Markov chain Monte Carlo) and has been employed to solve a host of combinatorial problems. The method of alternating projections was first proposed by von Neumann [20] to find a feasible point in the intersection of a family of hyperplanes. Modern versions of alternating projections more generally seek a point in the intersection of a family of closed convex sets. Backtracking is a standard technique taken from the toolkits of applied mathematics and computer science. Backtracking infallibly finds all solutions of a Sudoku puzzle or determines that no solution exists. Its Achilles heel of excessive computational complexity does not come into play with Sudoku puzzles because they are, despite appearances, relatively benign computationally. Sudoku puzzles are instances of the satisfiability problem in computer science. As problem size increases, such problems are combinatorially hard

---

\*Dept. Human Genetics, University of California, Los Angeles, CA. Email: ecchi@ucla.edu

†Depts. of Biomathematics and Statistics, University of California, Los Angeles, CA. Email: klange@ucla.edu

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

Fig. 1.1: Sample Puzzle

and often defy backtracking. For this reason alone, it is useful to examine alternative strategies.

In a typical Sudoku puzzle, there are 81 cells arranged in a 9-by-9 grid, some of which are occupied by numerical clues. See Figure 1.1. The goal is to fill in the remaining cells subject to the following three rules:

1. Each integer between 1 and 9 must appear exactly once in a row,
2. Each integer between 1 and 9 must appear exactly once in a column,
3. Each integer between 1 and 9 must appear exactly once in each of the 3-by-3 subgrids.

Solving a Sudoku game is a combinatorial task of intermediate complexity. The general problem of filling in an incomplete  $n^2 \times n^2$  grid with  $n \times n$  subgrids belongs to the class of NP-complete problems [22]. These problems are conjectured to increase in computational complexity at an exponential rate in  $n$ . Nonetheless, a well planned exhaustive search can work quite well for a low value of  $n$  such as 9. For larger values of  $n$ , brute force, no matter how cleverly executed, is simply not an option. In contrast, simulated annealing and alternating projections may yield good approximate solutions and partially salvage the situation.

In the rest of this paper, we describe the three methods for solving Sudoku puzzles and compare them on a battery of puzzles. The puzzles range in difficulty from pencil and paper exercises to hard benchmark tests that often defeat the two approximate methods. Our discussion reiterates the rationale for equipping students with the best computational tools.

## 2. Three methods for solving Sudoku.

**2.1. Backtracking.** Backtracking systematically grows a partial solution until it becomes a full solution or violates a constraint [18]. In the latter case it backtracks to the next permissible partial solution and begins the growing process anew. The advantage of backtracking is that a block of potential solutions can be discarded en masse. Backtracking starts by constructing for each empty Sudoku cell  $(i, j)$  a

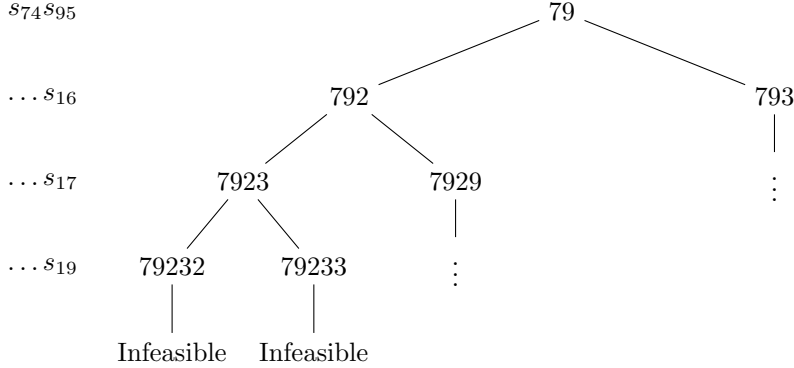


Fig. 2.1: Backtracking on the puzzle shown in Figure 1.1. Starting from  $s_{74}s_{95} = 79$ , the algorithm attempts and fails to grow the solution beyond  $s_{74}s_{95}s_{16}s_{17}s_{19} = 79232$ . After failing to grow the solution beyond  $s_{74}s_{95}s_{16}s_{17}s_{19} = 79233$ , all partial solutions beginning with 7923 are eliminated from further consideration. The algorithm starts anew by attempting to grow  $s_{74}s_{95}s_{16}s_{17} = 7929$ .

list  $L_{ij}$  of compatible digits. This is done by scanning the cell's row, column, and subgrid. The empty cells are then ordered by the cardinalities of the lists  $|L_{ij}|$ . For example in Figure 1.1, two cells (7, 4) and (9, 5) possess lists  $L_{74} = \{7\}$  and  $L_{95} = \{9\}$  with cardinality 1 and come first. Next come cells such as (1, 6) with  $L_{16} = \{2, 3\}$ , (1, 7) with  $L_{17} = \{3, 9\}$ , and (1, 9) with  $L_{19} = \{2, 3\}$  whose lists have cardinality 2. Finally come cells such as (2, 9) with  $L_{29} = \{2, 3, 5, 6, 7\}$  whose lists have maximum cardinality 5. Partial solutions are character strings such as  $s_{74}s_{95}s_{16}s_{17}s_{19}$  taken in dictionary order with the alphabet at cell  $(i, j)$  limited to the list  $L_{ij}$ . In dictionary order a string such as 7939 is treated as coming after a string such as 79232.

Backtracking starts with the string 7 by taking the only element of  $L_{74}$ , grows it to 79 by taking the only element of  $L_{95}$ , grows it to 792 by taking the first element of  $L_{16}$ , grows it to 7923 by taking the first element of  $L_{17}$ , and finally grows it to 79232 by taking the first element of  $L_{19}$ . At this juncture a row violation occurs, namely a 2 in both cells (1, 6) and (1, 9). Backtracking discards all strings beginning with 79232 and moves on to the string 79233 by replacing the first element of  $L_{19}$  by the second element of  $L_{19}$ . This leads to another row violation with a 3 in both cells (1, 7) and (1, 9). Backtracking moves back to the string 7929 by discarding the fifth character of 79239 and replacing the first element of  $L_{17}$  by its second element. This sets the stage for another round of growing.

Backtracking is also known as depth first search. In this setting the strings are viewed as nodes of a tree as depicted in Figure 2.1. Generating strings in dictionary order constitutes a tree traversal that systematically eliminates subtrees and moves down and backs up along branches. Because pruning large subtrees is more efficient than pruning small subtrees, ordering of cells by cardinality compels the decision tree to have fewer branches at the top. We use the C code from Skiena and Revilla [17] implementing backtracking on Sudoku puzzles. Backtracking has the virtue of finding all solutions when multiple solutions exist. Thus, it provides a mechanism for validating the correctness of puzzles.

**2.2. Simulated Annealing.** Simulated annealing [2, 10, 16] attacks a combinatorial optimization problem by defining a state space of possible solutions, a cost function quantifying departures from the solution ideal, and a positive temperature parameter. For a satisfiability problem, it is sensible to equate cost to the number of constraint violations. Solutions then correspond to states of zero cost. Each step of annealing operates by proposing a move to a new randomly chosen state. Proposals are Markovian in the sense that they depend only on the current state of the process, not on its past history. Proposed steps that decrease cost are always accepted. Proposed steps that increase cost are taken with high probability in the early stages of annealing when temperature is high and with low probability in the late stages of annealing when temperature is low. Inspired by models from statistical physics, simulated annealing is designed to sample the state space broadly before settling down at a local minimum of the cost function.

For the Sudoku problem, a state is a  $9 \times 9$  matrix (board) of integers drawn from the set  $\{1, \dots, 9\}$ . Each integer appears nine times, and all numerical clues are respected. Annealing starts from any feasible board. The proposal stage randomly selects two different cells without clues. The corresponding move swaps the contents of the cells, thus preserving all digit counts. To ensure that the most troublesome cells are more likely to be chosen for swapping, we select cells non-uniformly with probability proportional to  $\exp(i)$  for a cell involved in  $i$  constraint violations. Let  $\mathbf{B}$  denote a typical board,  $c(\mathbf{B})$  its associated cost, and  $n$  the current iteration index. At temperature  $\tau$ , we decide whether to accept a proposed neighboring board  $\mathbf{B}$  by drawing a random deviate  $U$  uniformly from  $[0, 1]$ . If  $U$  satisfies

$$U \leq \min \{ \exp([c(\mathbf{B}_n) - c(\mathbf{B})] / \tau_n), 1 \},$$

then we accept the proposed move and set  $\mathbf{B}_{n+1} = \mathbf{B}$ . Otherwise, we reject the move and set  $\mathbf{B}_{n+1} = \mathbf{B}_n$ . Thus, the greater the increase in the number of constraint violations, the less likely the move is made to a proposed state. Also, the higher the temperature, the more likely a move is made to an unfavorable state. The final ingredient of simulated annealing is the cooling schedule. In general, the temperature parameter  $\tau$  starts high and slowly declines to 0, where only favorable or cost neutral moves are taken. Typically temperature is lowered at a slow geometric rate.

**2.3. Alternating Projections.** The method of alternating projections relies on projection operators. In the projection problem, one seeks the closest point  $\mathbf{x}$  in a set  $C \subset \mathbb{R}^d$  to a point  $\mathbf{y} \in \mathbb{R}^d$ . Distance is quantified by the usual Euclidean norm  $\|\mathbf{x} - \mathbf{y}\|$ . If  $\mathbf{y}$  already lies in  $C$ , then the problem is trivially solved by setting  $\mathbf{x} = \mathbf{y}$ . It is well known that a unique minimizer exists whenever the set  $C$  is closed and convex [11]. We will denote the projection operator taking  $\mathbf{y}$  to  $\mathbf{x}$  by  $P_C(\mathbf{y}) = \mathbf{x}$ .

Given a finite collection of closed convex sets with a nonempty intersection, the alternating projection algorithm finds a point in that intersection. Consider the case of two closed convex sets  $A$  and  $B$ . The method recursively generates a sequence  $\mathbf{y}_n$  by taking  $\mathbf{y}_0 = \mathbf{y}$  and  $\mathbf{y}_{n+1} = P_A(\mathbf{y}_n)$  for  $n$  even and  $\mathbf{y}_{n+1} = P_B(\mathbf{y}_n)$  for  $n$  odd. Figure 2.2 illustrates a few iterations of the algorithm. As suggested by the picture, the algorithm does indeed converge to a point in  $A \cap B$  [3]. For more than two closed convex sets with nonempty intersection, the method of alternating projections cycles through the projections in some fixed order. Convergence occurs in this more general case as well based on some simple theory involving paracontractive operators [7]. The limit is not guaranteed to be the closest point in the intersection to the original point

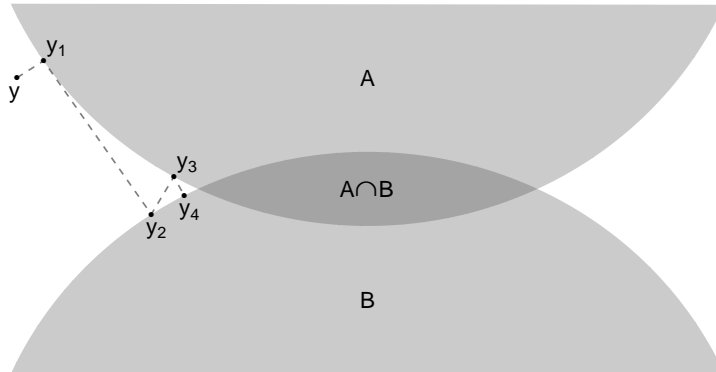


Fig. 2.2: Alternating projections find a point in  $A \cap B$ , where  $A$  and  $B$  are closed convex sets. The initial point is  $\mathbf{y}$ . The sequence of points  $\mathbf{y}_n$  is generated by alternating projection onto  $A$  with projection onto  $B$ .

$\mathbf{y}$ . The related but more complicated procedure known as Dykstra's algorithm [6] finds this point.

It is easy to construct some basic projection operators. For instance, projection onto the rectangle  $R = \{\mathbf{x} \in \mathbb{R}^d : a_i \leq x_i \leq b_i \text{ for all } i\}$  is achieved by defining  $\mathbf{x} = P_R(\mathbf{y})$  to have components  $x_i = \min\{\max\{a_i, y_i\}, b_i\}$ . This example illustrates a more general rule; namely, if  $A$  and  $B$  are two closed convex sets, then projection onto the Cartesian product  $A \times B$  is effected by the Cartesian product operator  $(\mathbf{x}, \mathbf{y}) \mapsto [P_A(\mathbf{x}), P_B(\mathbf{y})]$ . When  $A$  is an entire Euclidean space,  $P_A(\mathbf{x})$  is just the identity map. Projection onto the hyperplane

$$H = \{\mathbf{y} \in \mathbb{R}^d : \mathbf{v}^T \mathbf{y} = c\}$$

is implemented by the operator

$$P_H(\mathbf{x}) = \mathbf{x} - \frac{\mathbf{v}^T \mathbf{x} - c}{\|\mathbf{v}\|^2} \mathbf{v}.$$

Projection onto the unit simplex  $U = \{\mathbf{x} \in \mathbb{R}^d : \sum_{i=1}^d x_i = 1, x_i \geq 0 \forall i\}$  is more subtle. Fortunately there exist fast algorithms for this purpose [5, 13].

In either the alternating projection algorithm or Dykstra's algorithm, it is advantageous to reduce the number of participating convex sets to the minimum possible consistent with fast projection. For instance, it is better to take the unit simplex  $U$  as a whole rather than as an intersection of the halfspaces  $\{\mathbf{x} : x_i \geq 0\}$  and the affine subspace  $\{\mathbf{x} : \sum_{i=1}^d x_i = 1\}$ . Because our alternating projection algorithm for solving Sudoku puzzles relies on projecting onto several simplexes, it is instructive to derive the Duchi et al [5] projection algorithm. Consider minimization of a convex smooth function  $f(\mathbf{x})$  over  $U$ . The Karush-Kuhn-Tucker stationarity condition

involves setting the gradient of the Lagrangian

$$\mathcal{L}(\mathbf{x}, \lambda, \boldsymbol{\mu}) = f(\mathbf{x}) + \lambda \left( \sum_{i=1}^d x_i - 1 \right) - \sum_{i=1}^d \mu_i x_i$$

equal to  $\mathbf{0}$ . This is stated in components as the Gibbs criterion

$$0 = \frac{\partial}{\partial x_i} f(\mathbf{x}) + \lambda - \mu_i$$

for multipliers  $\mu_i \geq 0$  obeying the complementary slackness conditions  $\mu_i x_i = 0$ . For the choice  $f(\mathbf{x}) = \frac{1}{2} \|\mathbf{x} - \mathbf{y}\|^2$ , the Gibbs condition can be solved in the form

$$x_i = \begin{cases} y_i - \lambda & x_i > 0 \\ y_i - \lambda + \mu_i & x_i = 0. \end{cases}$$

If we let  $I_+ = \{i : x_i > 0\}$ , then the equality constraint

$$1 = \sum_{i \in I_+} x_i = \sum_{i \in I_+} y_i - |I_+| \lambda$$

implies

$$\lambda = \frac{1}{|I_+|} \left( \sum_{i \in I_+} y_i - 1 \right).$$

The catch, of course, is that we do not know  $I_+$ .

The key to avoid searching over all  $2^d$  subsets is the simple observation that the  $x_i$  and  $y_i$  are consistently ordered. Suppose on the contrary that  $y_i < y_j$  and  $x_j < x_i$ . For small  $s > 0$  substitute  $x_j + s$  for  $x_j$  and  $x_i - s$  for  $x_i$ . The objective function  $f(\mathbf{x}) = \frac{1}{2} \|\mathbf{x} - \mathbf{y}\|^2$  then changes by the amount

$$\frac{1}{2} \left[ (x_i - s - y_i)^2 + (x_j + s - y_j)^2 - (x_i - y_i)^2 - (x_j - y_j)^2 \right] = s(y_i - y_j + x_j - x_i) + s^2,$$

which is negative for  $s$  small. Let  $w_i$  denote the  $i$ th largest entry of  $\mathbf{y}$ . Then the Gibbs condition implies that  $w_1 \geq w_2 \geq \dots \geq w_{|I_+|} > \lambda$ . Thus, to determine  $\lambda$  we seek the largest  $k$  such that

$$w_k > \frac{1}{k} \left( \sum_{i=1}^k w_i - 1 \right)$$

and set  $\lambda$  equal to the right hand side of this inequality. With  $\lambda$  in hand, the Gibbs condition implies that  $x_i = \max\{y_i - \lambda, 0\}$ . It follows that projection onto  $U$  can be accomplished in  $O(d \log d)$  operations dominated by sorting. Algorithm 1 displays pseudocode for projection onto  $U$ .

Armed with these results, we now describe how to solve a continuous relaxation of Sudoku by the method of alternating projections. In the relaxed version of the problem, we imagine generating candidate solutions by random sampling. Each cell  $(i, j)$  is assigned a sampling distribution  $p_{ijk} = \Pr(S_{ij} = k)$  for choosing a random deviate  $S_{ij} \in \{1, \dots, 9\}$  to populate the cell. If a numerical clue  $k$  occupies cell  $(i, j)$ ,

---

**Algorithm 1** PROJECTION ONTO SIMPLEX

---

$\mathbf{w} \leftarrow \text{SORT\_DESCENDING}(\mathbf{y}).$   
 $k \leftarrow \max \left\{ j : w_j > \frac{1}{j} \left( \sum_{i=1}^j w_i \right) \right\}$   
 $\lambda \leftarrow \frac{1}{k} \left( \sum_{i=1}^k w_i \right)$   
 $x_i \leftarrow \max \{ y_i - \lambda, 0 \}.$

---

then we set  $p_{ijl} = 1$  for  $l = k$  and 0 otherwise. A matrix of sampled deviates  $\mathbf{S}$  constitutes a candidate solution. It seems reasonable to demand that the average puzzle obey the constraints. Once we find a feasible 3-dimensional tensor  $\mathcal{P} = (p_{ijk})$  obeying the constraints, a good heuristic for generating an integer solution  $\hat{\mathbf{S}}$  is to put

$$\hat{s}_{ij} = \max_{k \in \{1, \dots, 9\}} p_{ijk}.$$

In other words, we impute the most probable integer to each unknown cell  $(i, j)$ . It is easy to construct counterexamples where imputation of the most probable integer from a feasible tensor  $\mathcal{P}$  of the relaxed problem fails to solve the Sudoku puzzle.

In any case, the remaining agenda is to specify the constraints and the corresponding projection operators. The requirement that each digit appear in each row on average once amounts to the constraint  $\sum_{j=1}^9 p_{ijk} = 1$  for all  $i$  and  $k$  between 1 and 9. There are 81 such constraints. The requirement that each digit appear in each column on average once amounts to the constraint  $\sum_{i=1}^9 p_{ijk} = 1$  for all  $j$  and  $k$  between 1 and 9. Again, there are 81 such constraints. The requirement that each digit appear in each subgrid on average once amounts to the constraint  $\sum_{j=1}^3 \sum_{k=1}^3 p_{a+i, b+j, k} = 1$  for all  $k$  between 1 and 9 and all  $a$  and  $b$  chosen from the set  $\{0, 3, 6\}$ . This contributes another 81 constraints. Finally, the probability constraints  $\sum_{k=1}^9 p_{ijk} = 1$  for all  $i$  and  $j$  between 1 and 9 contribute 81 more affine constraints. Hence, there are a total of 324 affine constraints on the  $9^3 = 729$  parameters. In addition there are 729 nonnegativity constraints  $p_{ijk} \geq 0$ .

Every numerical clue voids several constraints. For example, if the digit 7 is mandated for cell  $(9,2)$ , then we must take  $p_{927} = 1$ ,  $p_{92k} = 0$  for  $k \neq 7$ ,  $p_{i27} = 0$  for all  $i \neq 9$ ,  $p_{9j7} = 0$  for all  $j \neq 2$ , and  $p_{ij7} = 0$  for all other pairs  $(i, j)$  in the  $(3,1)$  subgrid. In carrying out alternating projection, we eliminate the corresponding variables. With this proviso, we cycle through the simplex projections summarized in Algorithm 1. The process is very efficient but slightly tedious to code. For the sake of brevity we omit the remaining details. All code used to generate the subsequent results are available at <https://github.com/echi/Sudoku>, and we direct the interested reader there.

**3. Comparisons.** We generated test puzzles from code available online [21] and discarded puzzles that could be completely solved by filling in entries directly implied by the initial clues. This left 87 easy puzzles, 130 medium puzzles, and 100 hard puzzles. We also downloaded an additional 95 very hard benchmark puzzles [4, 19]. In simulated annealing, the temperature  $\tau$  was initialized to 200 and lowered by a factor of 0.99 after every 50 steps. We allowed at most  $2 \times 10^5$  iterations and reset the temperature to 200 if a solution had not been found after  $10^5$  iterations. For the alternating projection algorithm, we commenced projecting from the origin  $\mathbf{0}$ .

Backtracking successfully solved all puzzles. Table 3.1 shows the fraction of puzzles the two heuristics were able to successfully complete. Table 3.2 records summary

	Alt. Projection	Sim. Annealing	Backtracking	Number of Puzzles
Easy	0.85	1.00	1.00	87
Medium	0.89	1.00	1.00	130
Hard	0.72	0.97	1.00	100
Top 95	0.41	0.03	1.00	95

Table 3.1: Success rates for solving puzzles of varying difficulty.

		CPU Time (sec)		
		Alt. Projection	Sim. Annealing	Backtracking
Easy	Minimum	0.032	0.006	0.007
	Median	0.041	0.021	0.008
	Mean	0.052	0.112	0.008
	Maximum	0.237	0.970	0.009
Medium	Minimum	0.032	0.007	0.007
	Median	0.051	0.037	0.008
	Mean	0.062	0.231	0.008
	Maximum	0.269	3.36	0.010
Hard	Minimum	0.033	0.008	0.008
	Median	0.110	0.753	0.008
	Mean	0.159	1.104	0.009
	Maximum	0.525	7.204	0.031

Table 3.2: Summary statistics on the run times for different methods on puzzles of varying difficulty. For the alternating projection and simulated annealing techniques, only successfully solved puzzles are included in the statistics.

statistics for the CPU time taken by each method for each puzzle category. All computations were done on an iMac computer with a 3.4 GHz Intel Core i7 processor and 8 GB of RAM. We implemented the alternating projection and simulated annealing algorithms in Fortran 95. For backtracking we relied on the existing implementation in C.

The comparisons show that backtracking performs best, and for the vast majority of  $9 \times 9$  Sudoku problems it is probably going to be hard to beat. Simulated annealing finds the solution except for a handful of the most challenging paper and pencil problems, but its maximum run times are unimpressive. While alternating projection does not perform as well on the pencil and paper problems compared to the other two algorithms, it does not do terribly either. Moreover, we see hints of the tables turning on the hard puzzles.

Simulated annealing struggles mightily on the 95 benchmark puzzles. Closer inspection of individual puzzles reveals that these very hard puzzles admit many local minima with just a few constraint violations. Figure 3.1 shows a typical local minimum that traps the simulated annealing algorithm. Additionally, something curious is happening in Figure 3.2, which plots CPU solution times for alternating projection



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

Fig. 3.1: A typical local minimum that traps simulated annealing in a top 95 puzzle. Clues are shaded light gray. There are two column constraint violations caused by the cells shaded dark gray. The local minimum is deep in the sense that all one-step swaps result in further constraint violations.

versus backtracking. Points below the dashed line indicate puzzles that the method of alternating projection solves more efficiently than backtracking. It appears that when the method of alternating projections finds correct solutions to very hard problems, it tends to find them more quickly than backtracking.

**4. Discussion.** It goes almost without saying that students of the mathematical sciences should be taught a variety of solution techniques for combinatorial problems. Because Sudoku puzzles are easy to state and culturally neutral, they furnish a good starting point for the educational task ahead. It is important to stress the contrast between exact strategies that scale poorly with problem size and approximate strategies that adapt more gracefully. The performance of the alternating projection algorithm on the benchmark tests suggest it may have a role in solving much harder combinatorial problems. Certainly, the electrical engineering community takes this attitude, given the close kinship of Sudoku puzzles to problems in coding theory [8, 9, 14, 15].

One can argue that algorithm development has assumed a dominant role within the mathematical sciences. Three inter-related trends are feeding this phenomenon. First, computing power continues to grow. Execution times are dropping, and computer memory is getting cheaper. Second, good computing simply tempts scientists to tackle larger data sets. Third, certain fields, notably communications, imaging, genomics, and economics generate enormous amounts of data. All of these fields create problems in combinatorial optimization. For instance, modern DNA sequencing is still challenged by the phase problem of discerning the maternal and paternal origin of genetic variants. Computation is also being adopted more often as a means of proving propositions. The claim that at least 17 numerical clues are needed to ensure uniqueness of a Sudoku solution has apparently been proved using intelligent brute force [12]. Mathematical scientists need to be aware of computational developments outside their classical application silos. Importing algorithms from outside fields is

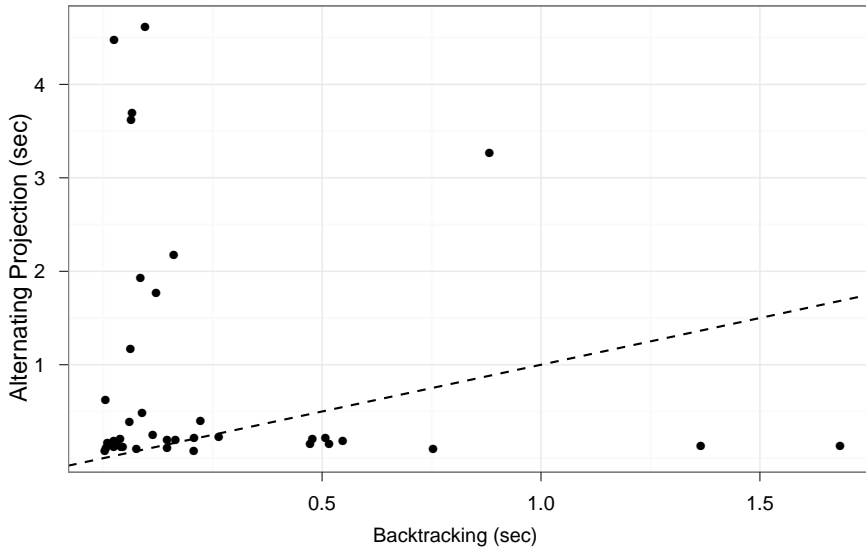


Fig. 3.2: Scatterplot of solution times for the top 95 benchmark problems.

one of the quickest means of refreshing an existing field.

**Acknowledgments.** We thank Peter Cock for helpful comments and pointing us to the source and author of the “Top 95” puzzles.

#### REFERENCES

- [1] R. A. BAILEY, P. J. CAMERON, AND R. CONNELLY, *Sudoku, gerechte designs, resolutions, affine space, spreads, reguli, and Hamming codes*, The American Mathematical Monthly, 115 (2008), pp. 383–404.
- [2] V. CERNY, *Thermodynamical approach to the traveling salesman problem: An efficient simulation algorithm*, J. Opt. Theory Appl., 45 (1985), pp. 41–51.
- [3] W. CHENEY AND A. A. GOLDSTEIN, *Proximity maps for convex sets*, Proceedings of the American Mathematical Society, 10 (1959), pp. 448–450.
- [4] P. COCK, *Solving Sudoku puzzles with Python*. Available at [http://www2.warwick.ac.uk/fac/sci/moac/people/students/peter\\_cock/python/sudoku](http://www2.warwick.ac.uk/fac/sci/moac/people/students/peter_cock/python/sudoku), 2012.
- [5] J. DUCHI, S. SHALEV-SHWARTZ, Y. SINGER, AND T. CHANDRA, *Efficient projections onto the  $\ell_1$ -ball for learning in high dimensions*, in Proceedings of the International Conference on Machine Learning, 2008.
- [6] R. L. DYKSTRA, *An algorithm for restricted least squares regression*, Journal of the American Statistical Association, 78 (1983), pp. 837–842.
- [7] L. ELSNER, I. KOLTRACHT, AND M. NEUMANN, *Convergence of sequential and asynchronous nonlinear paracontractions*, Numerische Mathematik, 62 (1992), pp. 305–319.
- [8] Y. ERLICH, K. CHANG, A. GORDON, R. RONEN, O. NAVON, M. ROOKS, AND G. J. HANNON, *DNA Sudoku - Harnessing high-throughput sequencing for multiplexed specimen analysis*, Genome Research, 19 (2009), pp. 1243–1253.
- [9] J. GUNTHER AND T. MOON, *Entropy minimization for solving Sudoku*, Signal Processing, IEEE Transactions on, 60 (2012), pp. 508–513.
- [10] S. KIRKPATRICK, C. D. GELATT, AND M. P. VECCHI, *Optimization by simulated annealing*, Science, 220 (1983), pp. 671–680.
- [11] K. LANGE, *Optimization*, Springer-Verlag, New York, 2004.
- [12] G. MCGUIRE, B. TUGEMANN, AND G. CIVARIO, *There is no 16-clue Sudoku: Solving the Sudoku minimum number of clues problem*. arXiv:1201.0749v1 [cs.DS], January 2012. Available

- at <http://arxiv.org/abs/1201.0749v1>.
- [13] C. MICHELOT, *A finite algorithm for finding the projection of a point onto the canonical simplex of  $\mathbb{R}^n$* , Journal of Optimization Theory and Applications, 50 (1986), pp. 195–200.
  - [14] T. MOON AND J. GUNTHER, *Multiple constraint satisfaction by belief propagation: An example using Sudoku*, in Adaptive and Learning Systems, 2006 IEEE Mountain Workshop on, july 2006, pp. 122–126.
  - [15] T. MOON, J. GUNTHER, AND J. KUPIN, *Sinkhorn solves Sudoku*, Information Theory, IEEE Transactions on, 55 (2009), pp. 1741–1746.
  - [16] W. H. PRESS, S. A. TEUKOLSKY, W. T. VETTERLING, AND B. P. FLANNERY, *Numerical Recipes: The Art of Scientific Computing*, Cambridge University Press, New York, New York, 3rd ed., 2007.
  - [17] S. SKIENA AND M. REVILLA, *Programming Challenges: The Programming Contest Training Manual*, Springer-Verlag, New York, 2003.
  - [18] S. S. SKIENA, *The Algorithm Design Manual*, Springer-Verlag, 2nd ed., 2008.
  - [19] G. STERTENBRINK, *Sudoku*. <http://magictour.free.fr/sudoku.htm>, November 2005.
  - [20] J. VON NEUMANN, *Functional operators. Vol. II.*, vol. 22 of Annals of Math. Studies, Princeton Univ. Press, Princeton, NJ, 1950. [This is a reprint of mimeographed lecture notes first distributed in 1933].
  - [21] H. WANG, *Solve and create Sudoku puzzles for different levels*. <http://www.mathworks.com/matlabcentral/fileexchange/13846-solve-and-create-sudoku-puzzles-for-different-levels>, 2007.
  - [22] T. YATO AND T. SETA, *Complexity and completeness of finding another solution and its application to puzzles*, IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences, 86 (2003), pp. 1052–1060.