Graph Coloring and Sudoku

The game of Sudoku has become, in the last few years, the rage among puzzle and game enthusiasts looking for a more intellectual (and cheaper) challenge than the one provided by an X-Box. Sudoku is addictive, and even ordinary people that are not drawn to video games are hooked on it. These days practically every major newspaper carries a daily Sudoku puzzle.

If you haven't played Sudoku yet, the rules are quite simple: You start with a 9×9 grid of 81 squares called *cells*. The grid is also subdivided into nine 3×3 subgrids called *boxes*. Some of the cells are already filled with the numbers 1 through 9. These are called the *givens*. The challenge of the game is to complete the grid by filling the remaining cells with the numbers 1 through 9. The requirements are: (i) every row and every column of the grid must have the numbers 1 through 9 appear once; (ii) each of the nine boxes must have the numbers 1 through 9 appear once.

	1	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*
a			4	8					
b		9		4	6			7	
c		5					6	1	4
d	2	1		6			5		
e	5	8		7		9		4	1
f			7			8		6	9
g	3	4	5					9	
h		6			3	7		2	
i						4	1		

FIGURE 2-10

A typical Sudoku puzzle may have somewhere between 25 and 40 givens, depending on the level of difficulty. Figure 2-10 is an example of a moderately easy Sudoku puzzle. (*Source*: *http://www.geometer.org/mathcircles*). The labels 1 through 9 on the columns and *a* through *i* on the rows are not part of the puzzle, but they provide a convenient way to refer to the cells. Just for fun, you may want to try this one out before you read on. (*Hint*: Try to figure out what number should go in cell *f*2. Once you have that one figured, go to cell *d*3. That's enough help for now. The solution is shown after the References and Further Readings.)

The Sudoku Graph

To see the connection between Sudoku and graph coloring, we will first describe the **Sudoku graph**, which for convenience we will refer to as *S*. The graph *S* has 81 vertices, with each vertex representing a cell. When two cells *cannot* have the same number (either because they are in the same row, in the same column, or in the same box) we put an edge connecting the corresponding vertices of the Sudoku graph *S*. For example, since cells *a*3 and *a*7 are in the same row, there is an edge joining their corresponding vertices; there is also an edge connecting *a*1 and *b*3 (they are in the same box), and so on. When everything is said and done, each vertex of the Sudoku graph has degree 20, and the graph has a total of 810 edges. *S* is too large to draw, but we can get a sense of the structure of *S* by looking at a partial drawing such as the one in Fig. 2-11. The drawing shows all 81 vertices of *S*, but only two (*a*1 and *e*5) have their full set of incident edges showing.

See Exercise 22.

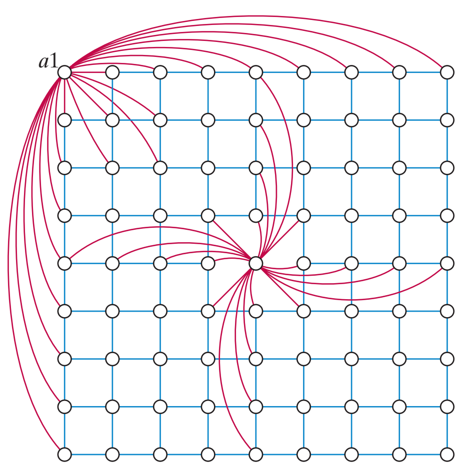

FIGURE 2-11 A partial drawing of the Sudoku graph

The second step in converting a Sudoku puzzle into a graph coloring problem is to assign colors to the numbers 1 through 9. This assignment is arbitrary, and is not a priority ordering of the colors as in the greedy algorithm —it's just a simple correspondence between numbers and colors. Figure 2-12 shows one such assignment.

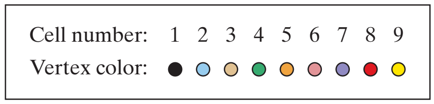

FIGURE 2-12

Once we have the Sudoku graph and an assignment of colors to the numbers 1 through 9, any Sudoku puzzle can be described by a Sudoku graph where some of the vertices are already colored (the ones corresponding to the givens). For example, the Sudoku puzzle shown in Fig. 2-10 is equivalent to the partial coloring shown in Fig. 2-13. To solve the Sudoku puzzle all we have to do is color the rest of the vertices using the nine colors in Fig. 2-12.

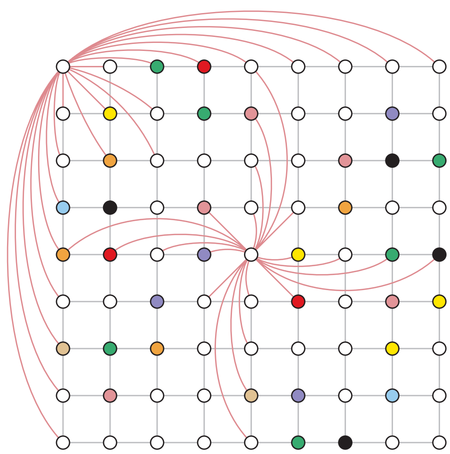

FIGURE 2-13 A Sudoku puzzle as a graph coloring problem

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