

**Figure 3.32** The track pieces in a wooden railway set; each is labeled with the number of copies in the set. Note that curved pieces and “fork” pieces (“switches” or “points”) can be flipped over so they can curve in either direction. Each curve subtends 45 degrees.

- 3.14** Which of the following are true and which are false? Explain your answers.
- Depth-first search always expands at least as many nodes as A\* search with an admissible heuristic.
  - $h(n) = 0$  is an admissible heuristic for the 8-puzzle.
  - A\* is of no use in robotics because percepts, states, and actions are continuous.
  - Breadth-first search is complete even if zero step costs are allowed.
  - Assume that a rook can move on a chessboard any number of squares in a straight line, vertically or horizontally, but cannot jump over other pieces. Manhattan distance is an admissible heuristic for the problem of moving the rook from square A to square B in the smallest number of moves.
- 3.15** Consider a state space where the start state is number 1 and each state  $k$  has two successors: numbers  $2k$  and  $2k + 1$ .
- Draw the portion of the state space for states 1 to 15.
  - Suppose the goal state is 11. List the order in which nodes will be visited for breadth-first search, depth-limited search with limit 3, and iterative deepening search.
  - How well would bidirectional search work on this problem? What is the branching factor in each direction of the bidirectional search?
  - Does the answer to (c) suggest a reformulation of the problem that would allow you to solve the problem of getting from state 1 to a given goal state with almost no search?
  - Call the action going from  $k$  to  $2k$  Left, and the action going to  $2k + 1$  Right. Can you find an algorithm that outputs the solution to this problem without any search at all?
- 3.16** A basic wooden railway set contains the pieces shown in Figure 3.32. The task is to connect these pieces into a railway that has no overlapping tracks and no loose ends where a train could run off onto the floor.
- Suppose that the pieces fit together *exactly* with no slack. Give a precise formulation of the task as a search problem.
  - Identify a suitable uninformed search algorithm for this task and explain your choice.
  - Explain why removing any one of the “fork” pieces makes the problem unsolvable.

- Give an upper bound on the total size of the state space defined by your formulation. (*Hint*: think about the maximum branching factor for the construction process and the maximum depth, ignoring the problem of overlapping pieces and loose ends. Begin by pretending that every piece is unique.)



**3.17** On page 90, we mentioned **iterative lengthening search**, an iterative analog of uniform cost search. The idea is to use increasing limits on path cost. If a node is generated whose path cost exceeds the current limit, it is immediately discarded. For each new iteration, the limit is set to the lowest path cost of any node discarded in the previous iteration.

- Show that this algorithm is optimal for general path costs.
- Consider a uniform tree with branching factor  $b$ , solution depth  $d$ , and unit step costs. How many iterations will iterative lengthening require?
- Now consider step costs drawn from the continuous range  $[\epsilon, 1]$ , where  $0 < \epsilon < 1$ . How many iterations are required in the worst case?
- Implement the algorithm and apply it to instances of the 8-puzzle and traveling salesman problems. Compare the algorithm’s performance to that of uniform-cost search, and comment on your results.

**3.18** Describe a state space in which iterative deepening search performs much worse than depth-first search (for example,  $O(n^2)$  vs.  $O(n)$ ).



**3.19** Write a program that will take as input two Web page URLs and find a path of links from one to the other. What is an appropriate search strategy? Is bidirectional search a good idea? Could a search engine be used to implement a predecessor function?



**3.20** Consider the vacuum-world problem defined in Figure 2.2.

- Which of the algorithms defined in this chapter would be appropriate for this problem? Should the algorithm use tree search or graph search?
- Apply your chosen algorithm to compute an optimal sequence of actions for a  $3 \times 3$  world whose initial state has dirt in the three top squares and the agent in the center.
- Construct a search agent for the vacuum world, and evaluate its performance in a set of  $3 \times 3$  worlds with probability 0.2 of dirt in each square. Include the search cost as well as path cost in the performance measure, using a reasonable exchange rate.
- Compare your best search agent with a simple randomized reflex agent that sucks if there is dirt and otherwise moves randomly.
- Consider what would happen if the world were enlarged to  $n \times n$ . How does the performance of the search agent and of the reflex agent vary with  $n$ ?

**3.21** Prove each of the following statements, or give a counterexample:

- Breadth-first search is a special case of uniform-cost search.
- Depth-first search is a special case of best-first tree search.
- Uniform-cost search is a special case of A\* search.

**3.22** Compare the performance of A\* and RBFS on a set of randomly generated problems in the 8-puzzle (with Manhattan distance) and TSP (with MST—see Exercise 3.30) domains. Discuss your results. What happens to the performance of RBFS when a small random number is added to the heuristic values in the 8-puzzle domain?

**3.23** Trace the operation of A\* search applied to the problem of getting to Bucharest from Lugoj using the straight-line distance heuristic. That is, show the sequence of nodes that the algorithm will consider and the  $f$ ,  $g$ , and  $h$  score for each node.

**3.24** Devise a state space in which A\* using GRAPH-SEARCH returns a suboptimal solution with an  $h(n)$  function that is admissible but inconsistent.

**3.25** The **heuristic path algorithm** (Pohl, 1977) is a best-first search in which the evaluation function is  $f(n) = (2 - w)g(n) + wh(n)$ . For what values of  $w$  is this complete? For what values is it optimal, assuming that  $h$  is admissible? What kind of search does this perform for  $w = 0$ ,  $w = 1$ , and  $w = 2$ ?

**3.26** Consider the unbounded version of the regular 2D grid shown in Figure 3.9. The start state is at the origin,  $(0,0)$ , and the goal state is at  $(x, y)$ .

- What is the branching factor  $b$  in this state space?
- How many distinct states are there at depth  $k$  (for  $k > 0$ )?
- What is the maximum number of nodes expanded by breadth-first tree search?
- What is the maximum number of nodes expanded by breadth-first graph search?
- Is  $h = |u - x| + |v - y|$  an admissible heuristic for a state at  $(u, v)$ ? Explain.
- How many nodes are expanded by A\* graph search using  $h$ ?
- Does  $h$  remain admissible if some links are removed?
- Does  $h$  remain admissible if some links are added between nonadjacent states?

**3.27**  $n$  vehicles occupy squares  $(1, 1)$  through  $(n, 1)$  (i.e., the bottom row) of an  $n \times n$  grid. The vehicles must be moved to the top row but in reverse order; so the vehicle  $i$  that starts in  $(i, 1)$  must end up in  $(n - i + 1, n)$ . On each time step, every one of the  $n$  vehicles can move one square up, down, left, or right, or stay put; but if a vehicle stays put, one other adjacent vehicle (but not more than one) can hop over it. Two vehicles cannot occupy the same square.

- Calculate the size of the state space as a function of  $n$ .
- Calculate the branching factor as a function of  $n$ .
- Suppose that vehicle  $i$  is at  $(x_i, y_i)$ ; write a nontrivial admissible heuristic  $h_i$  for the number of moves it will require to get to its goal location  $(n - i + 1, n)$ , assuming no other vehicles are on the grid.
- Which of the following heuristics are admissible for the problem of moving all  $n$  vehicles to their destinations? Explain.
  - $\sum_{i=1}^n h_i$ .
  - $\max\{h_1, \dots, h_n\}$ .
  - $\min\{h_1, \dots, h_n\}$ .

**3.28** Invent a heuristic function for the 8-puzzle that sometimes overestimates, and show how it can lead to a suboptimal solution on a particular problem. (You can use a computer to help if you want.) Prove that if  $h$  never overestimates by more than  $c$ , A\* using  $h$  returns a solution whose cost exceeds that of the optimal solution by no more than  $c$ .

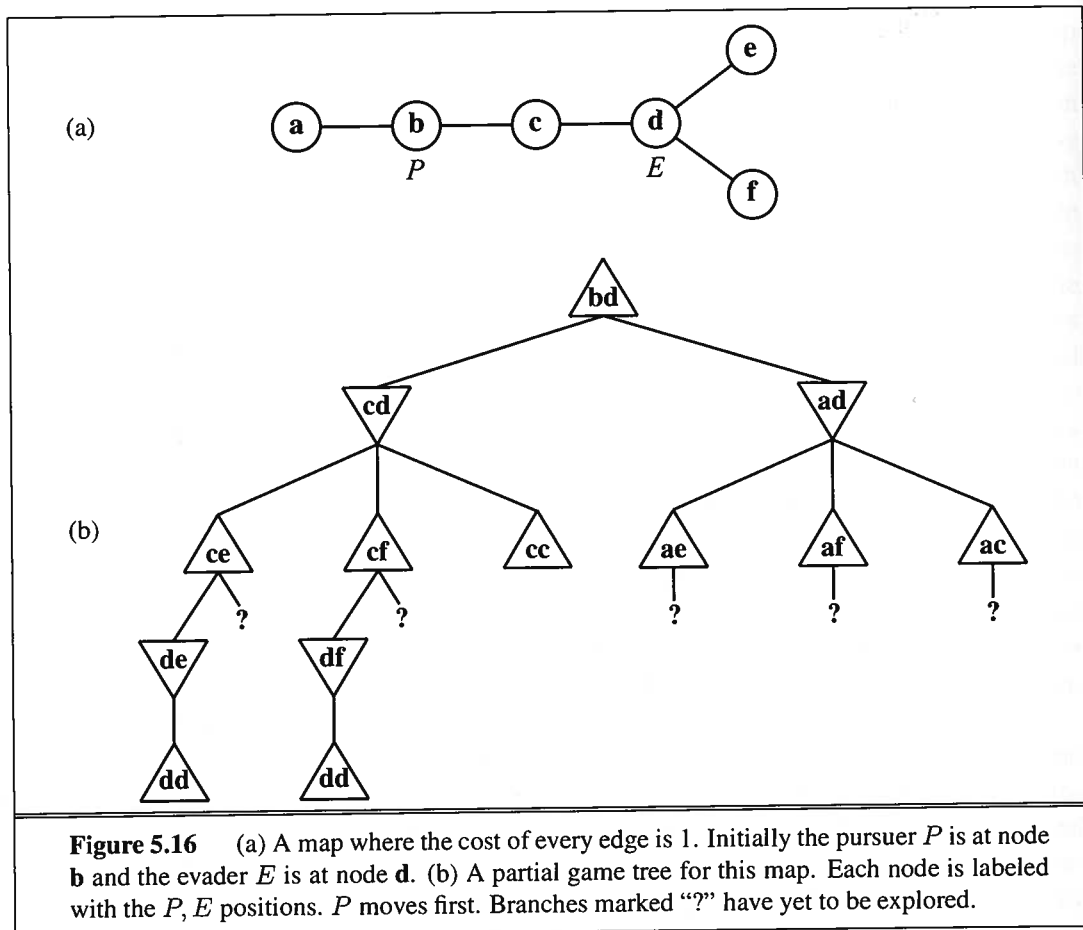
**3.29** Prove that if a heuristic is consistent, it must be admissible. Construct an admissible heuristic that is not consistent.

**3.30** The traveling salesperson problem (TSP) can be solved with the minimum-spanning-tree (MST) heuristic, which estimates the cost of completing a tour, given that a partial tour has already been constructed. The MST cost of a set of cities is the smallest sum of the link costs of any tree that connects all the cities.

- Show how this heuristic can be derived from a relaxed version of the TSP.
- Show that the MST heuristic dominates straight-line distance.
- Write a problem generator for instances of the TSP where cities are represented by random points in the unit square.
- Find an efficient algorithm in the literature for constructing the MST, and use it with A\* graph search to solve instances of the TSP.

**3.31** On page 105, we defined the relaxation of the 8-puzzle in which a tile can move from square A to square B if B is blank. The exact solution of this problem defines **Gaschnig's heuristic** (Gaschnig, 1979). Explain why Gaschnig's heuristic is at least as accurate as  $h_1$  (misplaced tiles), and show cases where it is more accurate than both  $h_1$  and  $h_2$  (Manhattan distance). Explain how to calculate Gaschnig's heuristic efficiently.

**3.32** We gave two simple heuristics for the 8-puzzle: Manhattan distance and misplaced tiles. Several heuristics in the literature purport to improve on this—see, for example, Nilsson (1971), Mostow and Prieditis (1989), and Hansson *et al.* (1992). Test these claims by implementing the heuristics and comparing the performance of the resulting algorithms.



**Figure 5.16** (a) A map where the cost of every edge is 1. Initially the pursuer  $P$  is at node  $b$  and the evader  $E$  is at node  $d$ . (b) A partial game tree for this map. Each node is labeled with the  $P, E$  positions.  $P$  moves first. Branches marked “?” have yet to be explored.

**5.3** Imagine that, in Exercise 3.3, one of the friends wants to avoid the other. The problem then becomes a two-player **pursuit–evasion** game. We assume now that the players take turns moving. The game ends only when the players are on the same node; the terminal payoff to the pursuer is minus the total time taken. (The evader “wins” by never losing.) An example is shown in Figure 5.16.


- Copy the game tree and mark the values of the terminal nodes.
- Next to each internal node, write the strongest fact you can infer about its value (a number, one or more inequalities such as “ $\geq 14$ ”, or a “?”).
- Beneath each question mark, write the name of the node reached by that branch.
- Explain how a bound on the value of the nodes in (c) can be derived from consideration of shortest-path lengths on the map, and derive such bounds for these nodes. Remember the cost to get to each leaf as well as the cost to solve it.
- Now suppose that the tree as given, with the leaf bounds from (d), is evaluated from left to right. Circle those “?” nodes that would *not* need to be expanded further, given the bounds from part (d), and cross out those that need not be considered at all.
- Can you prove anything in general about who wins the game on a map that is a tree?

**5.4** Describe and implement state descriptions, move generators, terminal tests, utility functions, and evaluation functions for one or more of the following stochastic games: Monopoly, Scrabble, bridge play with a given contract, or Texas hold’em poker.

**5.5** Describe and implement a *real-time, multiplayer* game-playing environment, where time is part of the environment state and players are given fixed time allocations.

**5.6** Discuss how well the standard approach to game playing would apply to games such as tennis, pool, and croquet, which take place in a continuous physical state space.

**5.7** Prove the following assertion: For every game tree, the utility obtained by MAX using minimax decisions against a suboptimal MIN will be never be lower than the utility obtained playing against an optimal MIN. Can you come up with a game tree in which MAX can do still better using a *suboptimal* strategy against a suboptimal MIN?



**Figure 5.17** The starting position of a simple game. Player  $A$  moves first. The two players take turns moving, and each player must move his token to an open adjacent space in either direction. If the opponent occupies an adjacent space, then a player may jump over the opponent to the next open space if any. (For example, if  $A$  is on 3 and  $B$  is on 2, then  $A$  may move back to 1.) The game ends when one player reaches the opposite end of the board. If player  $A$  reaches space 4 first, then the value of the game to  $A$  is  $+1$ ; if player  $B$  reaches space 1 first, then the value of the game to  $A$  is  $-1$ .

**5.8** Consider the two-player game described in Figure 5.17.

- Draw the complete game tree, using the following conventions:
  - Write each state as  $(s_A, s_B)$ , where  $s_A$  and  $s_B$  denote the token locations.
  - Put each terminal state in a square box and write its game value in a circle.
  - Put *loop states* (states that already appear on the path to the root) in double square boxes. Since their value is unclear, annotate each with a “?” in a circle.
- Now mark each node with its backed-up minimax value (also in a circle). Explain how you handled the “?” values and why.
- Explain why the standard minimax algorithm would fail on this game tree and briefly sketch how you might fix it, drawing on your answer to (b). Does your modified algorithm give optimal decisions for all games with loops?
- This 4-square game can be generalized to  $n$  squares for any  $n > 2$ . Prove that  $A$  wins if  $n$  is even and loses if  $n$  is odd.

**5.9** This problem exercises the basic concepts of game playing, using tic-tac-toe (noughts and crosses) as an example. We define  $X_n$  as the number of rows, columns, or diagonals

with exactly  $n$   $X$ 's and no  $O$ 's. Similarly,  $O_n$  is the number of rows, columns, or diagonals with just  $n$   $O$ 's. The utility function assigns  $+1$  to any position with  $X_3 = 1$  and  $-1$  to any position with  $O_3 = 1$ . All other terminal positions have utility 0. For nonterminal positions, we use a linear evaluation function defined as  $Eval(s) = 3X_2(s) + X_1(s) - (3O_2(s) + O_1(s))$ .


- Approximately how many possible games of tic-tac-toe are there?
- Show the whole game tree starting from an empty board down to depth 2 (i.e., one  $X$  and one  $O$  on the board), taking symmetry into account.
- Mark on your tree the evaluations of all the positions at depth 2.
- Using the minimax algorithm, mark on your tree the backed-up values for the positions at depths 1 and 0, and use those values to choose the best starting move.
- Circle the nodes at depth 2 that would *not* be evaluated if alpha-beta pruning were applied, assuming the nodes are generated in the optimal order for alpha-beta pruning.

**5.10** Consider the family of generalized tic-tac-toe games, defined as follows. Each particular game is specified by a set  $\mathcal{S}$  of squares and a collection  $\mathcal{W}$  of winning positions. Each winning position is a subset of  $\mathcal{S}$ . For example, in standard tic-tac-toe,  $\mathcal{S}$  is a set of 9 squares and  $\mathcal{W}$  is a collection of 8 subsets of  $\mathcal{W}$ : the three rows, the three columns, and the two diagonals. In other respects, the game is identical to standard tic-tac-toe. Starting from an empty board, players alternate placing their marks on an empty square. A player who marks every square in a winning position wins the game. It is a tie if all squares are marked and neither player has won.

- Let  $N = |\mathcal{S}|$ , the number of squares. Give an upper bound on the number of nodes in the complete game tree for generalized tic-tac-toe as a function of  $N$ .
- Give a lower bound on the size of the game tree for the worst case, where  $\mathcal{W} = \{\}$ .
- Propose a plausible evaluation function that can be used for any instance of generalized tic-tac-toe. The function may depend on  $\mathcal{S}$  and  $\mathcal{W}$ .
- Assume that it is possible to generate a new board and check whether it is a winning position in  $100N$  machine instructions and assume a 2 gigahertz processor. Ignore memory limitations. Using your estimate in (a), roughly how large a game tree can be completely solved by alpha-beta in a second of CPU time? a minute? an hour?

**5.11** Develop a general game-playing program, capable of playing a variety of games.

- Implement move generators and evaluation functions for one or more of the following games: Kalah, Othello, checkers, and chess.
- Construct a general alpha-beta game-playing agent.
- Compare the effect of increasing search depth, improving move ordering, and improving the evaluation function. How close does your effective branching factor come to the ideal case of perfect move ordering?
- Implement a selective search algorithm, such as B\* (Berliner, 1979), conspiracy number search (McAllester, 1988), or MGSS\* (Russell and Wefald, 1989) and compare its performance to A\*.



**Figure 5.18** Situation when considering whether to prune node  $n_j$ .

**5.12** Describe how the minimax and alpha-beta algorithms change for two-player, non-zero-sum games in which each player has a distinct utility function and both utility functions are known to both players. If there are no constraints on the two terminal utilities, is it possible for any node to be pruned by alpha-beta? What if the player's utility functions on any state differ by at most a constant  $k$ , making the game almost cooperative?

**5.13** Develop a formal proof of correctness for alpha-beta pruning. To do this, consider the situation shown in Figure 5.18. The question is whether to prune node  $n_j$ , which is a max-node and a descendant of node  $n_1$ . The basic idea is to prune it if and only if the minimax value of  $n_1$  can be shown to be independent of the value of  $n_j$ .

- Node  $n_1$  takes on the minimum value among its children:  $n_1 = \min(n_2, n_{21}, \dots, n_{2b_2})$ . Find a similar expression for  $n_2$  and hence an expression for  $n_1$  in terms of  $n_j$ .
- Let  $l_i$  be the minimum (or maximum) value of the nodes to the left of node  $n_i$  at depth  $i$ , whose minimax value is already known. Similarly, let  $r_i$  be the minimum (or maximum) value of the unexplored nodes to the right of  $n_i$  at depth  $i$ . Rewrite your expression for  $n_1$  in terms of the  $l_i$  and  $r_i$  values.
- Now reformulate the expression to show that in order to affect  $n_1$ ,  $n_j$  must not exceed a certain bound derived from the  $l_i$  values.
- Repeat the process for the case where  $n_j$  is a min-node.

**5.14** Prove that alpha-beta pruning takes time  $O(2^{m/2})$  with optimal move ordering, where  $m$  is the maximum depth of the game tree.

**5.15** Suppose you have a chess program that can evaluate 10 million nodes per second. Decide on a compact representation of a game state for storage in a transposition table. About how many entries can you fit in a 2-gigabyte in-memory table? Will that be enough for the

from database theory, Gottlob *et al.* (1999a, 1999b) developed a notion, **hypertree width**, that is based on the characterization of the CSP as a hypergraph. In addition to showing that any CSP with hypertree width  $w$  can be solved in time  $O(n^{w+1} \log n)$ , they also showed that hypertree width subsumes all previously defined measures of “width” in the sense that there are cases where the hypertree width is bounded and the other measures are unbounded.

Interest in look-back approaches to backtracking was rekindled by the work of Bayardo and Schrag (1997), whose RELSAT algorithm combined constraint learning and backjumping and was shown to outperform many other algorithms of the time. This led to AND/OR search algorithms applicable to both CSPs and probabilistic reasoning (Dechter and Mateescu, 2007). Brown *et al.* (1988) introduce the idea of symmetry breaking in CSPs, and Gent *et al.* (2006) give a recent survey.

The field of **distributed constraint satisfaction** looks at solving CSPs when there is a collection of agents, each of which controls a subset of the constraint variables. There have been annual workshops on this problem since 2000, and good coverage elsewhere (Collin *et al.*, 1999; Pearce *et al.*, 2008; Shoham and Leyton-Brown, 2009).

Comparing CSP algorithms is mostly an empirical science: few theoretical results show that one algorithm dominates another on all problems; instead, we need to run experiments to see which algorithms perform better on typical instances of problems. As Hooker (1995) points out, we need to be careful to distinguish between competitive testing—as occurs in competitions among algorithms based on run time—and scientific testing, whose goal is to identify the properties of an algorithm that determine its efficacy on a class of problems.

The recent textbooks by Apt (2003) and Dechter (2003), and the collection by Rossi *et al.* (2006) are excellent resources on constraint processing. There are several good earlier surveys, including those by Kumar (1992), Dechter and Frost (2002), and Bartak (2001); and the encyclopedia articles by Dechter (1992) and Mackworth (1992). Pearson and Jeavons (1997) survey tractable classes of CSPs, covering both structural decomposition methods and methods that rely on properties of the domains or constraints themselves. Kondrak and van Beek (1997) give an analytical survey of backtracking search algorithms, and Bacchus and van Run (1995) give a more empirical survey. Constraint programming is covered in the books by Apt (2003) and Fruhwirth and Abdennadher (2003). Several interesting applications are described in the collection edited by Freuder and Mackworth (1994). Papers on constraint satisfaction appear regularly in *Artificial Intelligence* and in the specialist journal *Constraints*. The primary conference venue is the International Conference on Principles and Practice of Constraint Programming, often called *CP*.

## EXERCISES

**6.1** How many solutions are there for the map-coloring problem in Figure 6.1? How many solutions if four colors are allowed? Two colors?

**6.2** Consider the problem of placing  $k$  knights on an  $n \times n$  chessboard such that no two knights are attacking each other, where  $k$  is given and  $k \leq n^2$ .

- Choose a CSP formulation. In your formulation, what are the variables?
- What are the possible values of each variable?
- What sets of variables are constrained, and how?
- Now consider the problem of putting *as many knights as possible* on the board without any attacks. Explain how to solve this with local search by defining appropriate ACTIONS and RESULT functions and a sensible objective function.

**6.3** Consider the problem of constructing (not solving) crossword puzzles:<sup>5</sup> fitting words into a rectangular grid. The grid, which is given as part of the problem, specifies which squares are blank and which are shaded. Assume that a list of words (i.e., a dictionary) is provided and that the task is to fill in the blank squares by using any subset of the list. Formulate this problem precisely in two ways:

- As a general search problem. Choose an appropriate search algorithm and specify a heuristic function. Is it better to fill in blanks one letter at a time or one word at a time?
- As a constraint satisfaction problem. Should the variables be words or letters?

Which formulation do you think will be better? Why?

**6.4** Give precise formulations for each of the following as constraint satisfaction problems:

- Rectilinear floor-planning: find non-overlapping places in a large rectangle for a number of smaller rectangles.
- Class scheduling: There is a fixed number of professors and classrooms, a list of classes to be offered, and a list of possible time slots for classes. Each professor has a set of classes that he or she can teach.
- Hamiltonian tour: given a network of cities connected by roads, choose an order to visit all cities in a country without repeating any.

**6.5** Solve the cryptarithmic problem in Figure 6.2 by hand, using the strategy of backtracking with forward checking and the MRV and least-constraining-value heuristics.

**6.6** Show how a single ternary constraint such as “ $A + B = C$ ” can be turned into three binary constraints by using an auxiliary variable. You may assume finite domains. (*Hint:* Consider a new variable that takes on values that are pairs of other values, and consider constraints such as “ $X$  is the first element of the pair  $Y$ .”) Next, show how constraints with more than three variables can be treated similarly. Finally, show how unary constraints can be eliminated by altering the domains of variables. This completes the demonstration that any CSP can be transformed into a CSP with only binary constraints.

**6.7** Consider the following logic puzzle: In five houses, each with a different color, live five persons of different nationalities, each of whom prefers a different brand of candy, a different drink, and a different pet. Given the following facts, the questions to answer are “Where does the zebra live, and in which house do they drink water?”

<sup>5</sup> Ginsberg *et al.* (1990) discuss several methods for constructing crossword puzzles. Littman *et al.* (1999) tackle the harder problem of solving them.