16.410-13: Principles of Autonomy and Decision Making
Sample Final Exam

December 3rd, 2004

Name	
E-mail	

<u>Note:</u> Budget your time wisely. Some parts of this exam could take you much longer than others. Move on if you are stuck and return to the problem later.

Problem Number	Max	Score	Grader
Problem 1	20		
Problem 2	20		
Problem 3	20		
Problem 4	27		
Problem 5	30		
Problem 6	30		
Problem 7	14		
Problem 8	15		
Problem 9	15		
Total	191		

Problem 1 – Search (20 points)

Part A – Dijkstra's Algorithm (6 points)

Consider the following graph, with node 1 as the start, and node 9 as the goal.

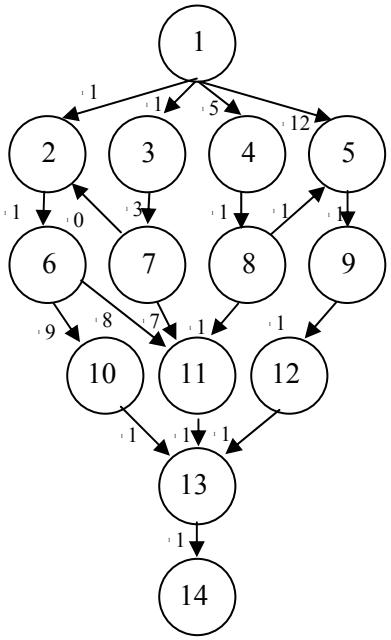

Part A-1 – (2 points)

How many times does the value at node 11 change?

3 – Once for each of node 11's parents - nodes 5, 6 and 7.

Part A-2 – (2 points)

What is the length of the shortest route from start to goal?

8, following along the path 1-4-7-8-9.

Part A-3 – (2 points)

In what order are the nodes expanded?

1, 2, 3, 5, 6, 4, 7, 8, and then 9.

Part B – A* Search (8 points)

Consider the following maze. Actions are moves to the 8 neighboring squares. Each such move involves a move over a distance, which is the cost of the move. Let's say that the distance to the left and right neighbors, and to the upper and lower neighbors is 1, and that the distance to the corner neighbors is $\sqrt{2}$.

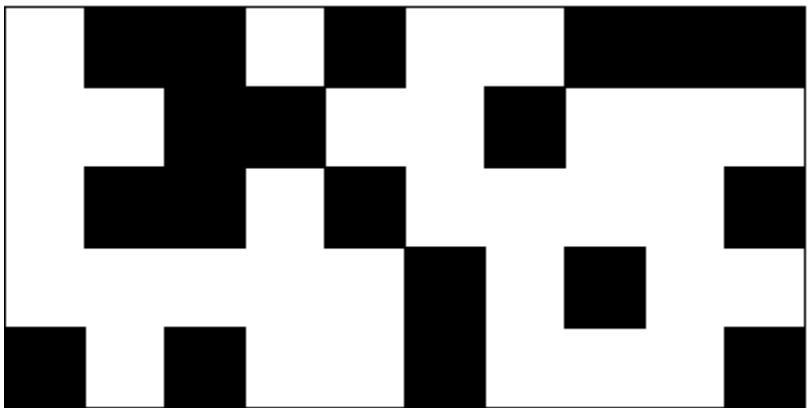

Part B-1 – (4 points)

Is Euclidean distance an admissible heuristic? Why or why not? (The Euclidean distance between two points <x1,y1> and <x2,y2> is $[(x2 - x1\})^2 + (y2 - y1)^2]^{1/2}$).

Yes. Euclidean distance is the straight line distance between two points, which is the shortest distance between those points. Hence Euclidean distance is always less than or equal to the true distance of any path between those two points within the maze.

Part B-2 – (4 points)

Is Manhattan distance an admissible heuristic? Why or why not? (Manhattan distance between two points <x1,y1> and <x2,y2> is (|x2 – x1| + |y2 – y1|)).

No. The Manhattan distance may be greater than the true distance. For example, the Manhattan distance to the corner neighbors will be greater than the true distance.

Part C – Properties of Search (6 points)

Part C-1 – (1 point)

How can A* be made to behave just like breadth-first search?

Set h = 0

Part C-2 – (1 point)

How can depth-first search be made to behave just like breadth-first search?

Use iterative deepening

Part C-3 – (2 points)

Is A* always the fastest search method? Explain your answer.

No. Hill climbing w/o backtracking, for example, could find a solution much faster (if it is lucky).

Part C-4 – (2 points)

Is depth first always slower than breadth first search? Explain your answer.

No. If the goal happens to be on an early expansion of the depth-first algorithm, then dfs could be faster. For example, if dfs expands from left to right, and the goal is to the farthest left of the tree.

Problem 2 – Planning with Graphplan (20 points)

Consider the following set of initial facts and operators.

Initial facts:

```
(Item Brian)
(Item Laptop)
(City Boston)
(City PaloAlto)
(City Ames)
(Plane USAir-1)
(Plane USAir-2)

(in Boston Brian)
(in PaloAlto Laptop)
(in Boston USAir-1)
(in PaloAlto USAir-2)
```

Operators:

```
(operator board
        (parameters (Item x) (Plane y) (City z))
        (preconditions (in z x) (in z y))
        (effects (on y x))
)

(operator fly
        (parameters (Plane x) (City y) (City z))
        (preconditions (in y x))
        (effects (in z x))
)

(operator deplane
        (parameters (Item x) (Plane y) (City z))
        (preconditions (on y x) (in z y))
        (effects (in z x))
)
```

Part A – Get Brian from Boston to his laptop in Palo Alto (5 points)

What is the goal fact?

(in Boston Laptop)

How many different propositional symbols result from instantiating fly(x, y, z) for the first ten levels of the plan graph, given that the initial facts have already been instantiated?

Instantiates (in z x) for each of 10 layers, with z = 2 planes and x = three cities
10 x 2 x 3 = 60

Part B – Get Brian's laptop from Palo Alto to Boston using Graphplan (15 points)

Fill in the plan graph, showing level 1 operators and level 1 and 2 facts (but no mutexes).

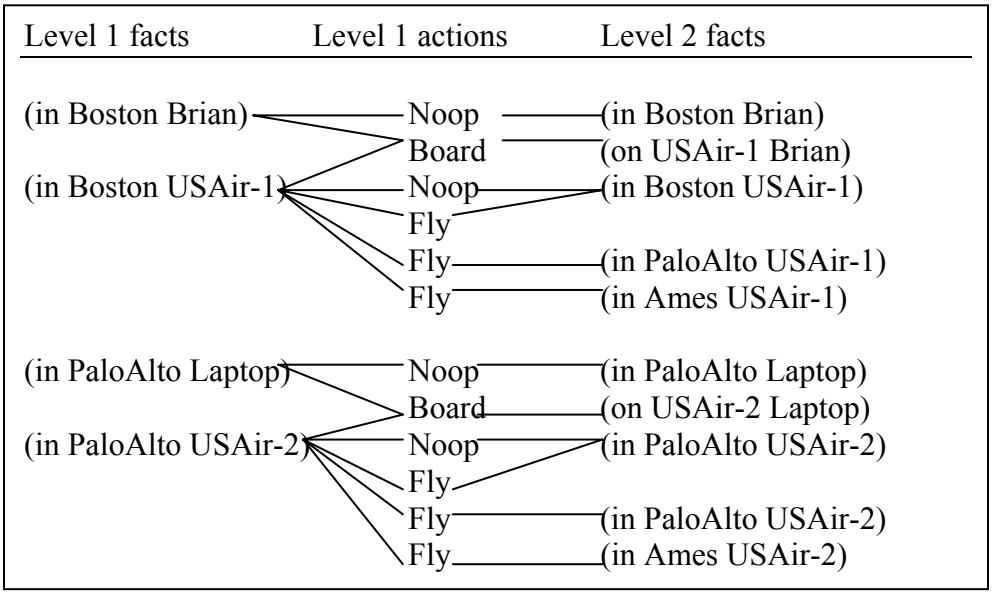

Problem 3 – Propositional Logic and Inference (20 points)

Consider a theory comprised of the following six clauses:

not A or B;
not B or C;
not C or not D or E;
not D or not E;
not F or not G;
F.

Part A – Satisfiability Using DPLL (15 points)

Use the DPLL algorithm (backtrack search plus unit propagation) to **find a truth assignment** to propositions A, B, C, D, E, F and G, that is consistent with the theory. **Fill out the search tree supplied below, stopping at the first consistent assignment found.**

- **Search** the propositions **in alphabetical order** (no other order please!).
- For each proposition P, **assign** the value **True before** trying **False**.
- On the line next to each node in the tree, **write the proposition being assigned** a truth value at that point in the search.
- In the box next to each branch**, list all propositions** whose truth value is **determined by unit propagation** based on the assignment to the proposition at that branch.
- **Indicate** the **truth value derived** for each of these propositions.
- **Draw an X** at each node that is immediately below the branch **where at least one clause becomes false**; this is where the search backtracks.
- **Circle** the node that denotes the **first complete and consistent assignment** to the propositions A **–** G.

We filled out the result of the initial propagation in the box above the tree. In addition, we filled out the first variable to be assigned (A), next to the root.

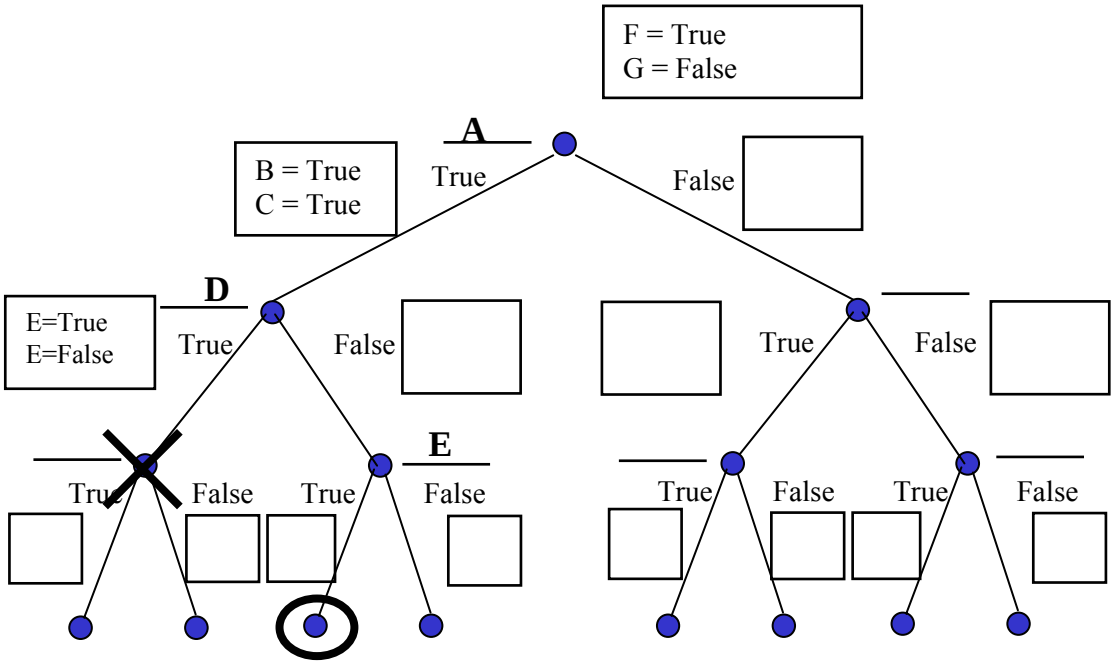

Part B – Satisfiability with Backtracking + Forward Checking (5 points)

Consider solving the same problem (finding the first consistent assignment for the theory) using **backtrack search plus forward checking**. Would backtrack search plus forward checking expand the same number of search nodes, up to the first solution, **fewer** nodes, or **more** nodes than the DPLL algorithm? Give a brief explanation for your answer.

Circle one of Fewer, Same or More

Why? DPLL deduces truth assignments to propositions at the root node by using unit propagation. BT+FC may need to assign these propositions by explicit splitting. In addition, DPLL will use unit propagation at a node to detect inconsistencies involving two or unassigned variables. FC would need to expand the node to its descendants in order to detect the same inconsistency.

Problem 4 – Model-based Diagnosis (27 points)

You are having trouble with your car. Each time you brake, the car drags to the right (the left wheel underbrakes and the right wheel overbrakes), while the brake pedal feels harder than normal. Let's see if what you have learned in 16.410/13 can help you find the cause.

The reference manual of the car states that the hydraulic circuit consists of the pedal brake cylinder B, the left and right wheel brake cylinders LW and RW, and two valves LV and RV:

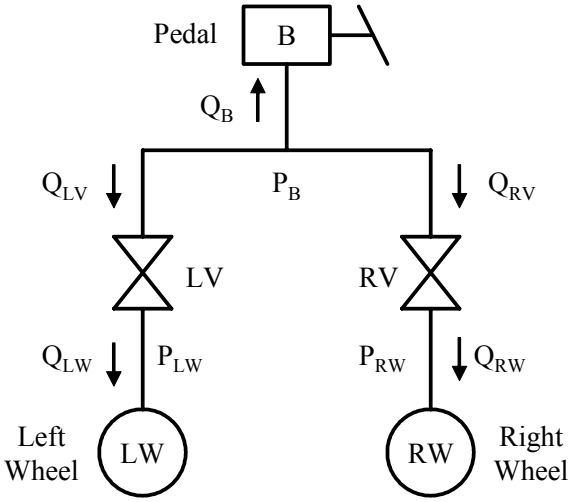

The fluid flow in the pipes (in the direction indicated by the arrows) is denoted Q_B, Q_{LV}, Q_{RV}, Q_{LW}, and Q_{RW}. The pressure in the pipes is denoted P_B, P_{LW}, and P_{RW}.

If a valve is working correctly, then the flow across it is proportional to the pressure difference in the adjacent pipes. If a pedal or wheel brake cylinder is working correctly, then the flow into it is proportional to the pressure in the adjacent pipe. The brake fluid is incompressible, so the flow sums up to zero at the junction of the pipes to the left and right branches.

The next page shows a constraint-based model of the hydraulic circuit as described above. Each variable can assume one of the three values "low", "nominal", and "high", abbreviated as "-", "0", and "+". Each constraint lists the possible combinations of values if the component works correctly (denoted "G"). In this representation, the observations of the car's strange behavior can be expressed as P_B="+", P_{LW}="-", P_{RW}="+".

B:

P_B	Q_B
-	-
0	0
+	+

LW:

P_{LW}	Q_{LW}
-	-
0	0
+	+

RW:

P_{RW}	Q_{RW}
-	-
0	0
+	+

LV:

Q_{LV}	Q_{LW}	P_B	P_{LW}
-	-	-	-
-	-	-	0
-	-	-	+
-	-	0	+
-	-	+	+
0	0	-	-
0	0	0	0
0	0	+	+
+	+	-	-
+	+	0	-
+	+	+	-
+	+	+	0
+	+	+	+

RV:

Q_{RV}	Q_{RW}	P_B	P_{RW}
-	-	-	-
-	-	-	0
-	-	-	+
-	-	0	+
-	-	+	+
0	0	-	-
0	0	0	0
0	0	+	+
+	+	-	-
+	+	0	-
+	+	+	-
+	+	+	0
+	+	+	+

Pipe Junction:

Q_{LV}	Q_B	Q_{RV}
+	-	-
-	+	-
+	0	-
0	+	-
+	+	-
+	-	0
0	0	0
-	+	0
+	-	+
-	+	+
-	0	+
0	-	+
-	-	+

In the lecture, you learned how to find symptoms by applying unit propagation to a set of clauses, and how to extract conflicts from inconsistent clauses by tracing support. Now you will **generalize** propagation and conflict extraction **from clauses to constraints**.

Recall that in unit propagation, we look for clauses where **all literals except one** are false. We then **assign** true to the remaining literal, and **record the clause as a support** for this literal.

Constraint propagation generalizes this in the following way: In constraint propagation we look for constraints where **all values of a variable except one** are excluded. We then **assign** this remaining value to the variable, and **record the constraint as a support** for this assignment.

For example, consider the constraint for LV. If P_B = "+" and P_{LW} = "-", then Q_{LV} is restricted to the single value "+" and Q_{LW} is restricted to the single value "+". We assign these values to the variables and record "LV=G" as a support for Q_{LV} = "+" and as a support for Q_{LW} = "+".

In the diagram below, we applied constraint propagation to the constraint model above, given the observations P_B="+", P_{LW}="-", P_{RW}="+". The support for each predicted value is shown next to that value. The pipes are assumed to be ok, so the pipe junction is not included as a support (for Q_{LV}).

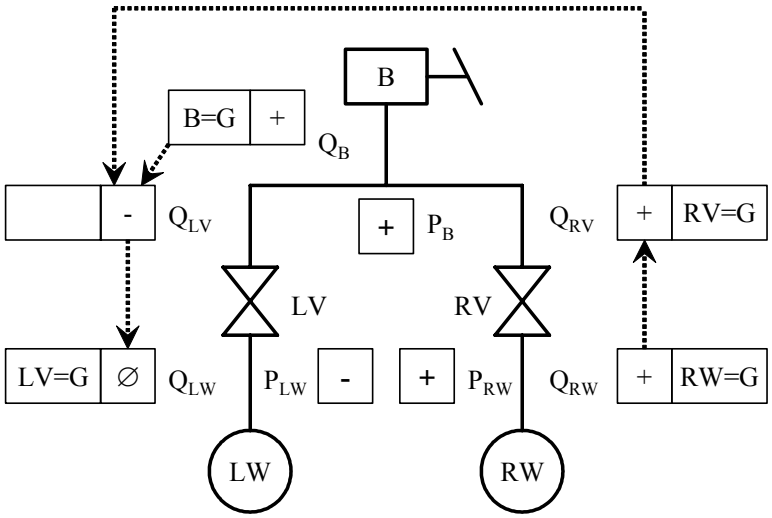

Part A – Conflict Extraction from Support [4 points]

As shown in the diagram above, the model and the observations are **inconsistent** with the assertion that every component is correct: constraint LV has become empty, as there is no tuple in the constraint LV that allows for Q_{LV} ="-", P_B = "+", and P_{LW} = "-". **Extract the conflict** by tracing back the support for the predictions (you should be able to do this by inspection):

Not (LV=G and B=G and RV=G and RW=G) Note: pipe excluded, because it is given that it is okay.

Part B – Constraint Suspension for Fault Localization [20 points]

Let's now see what single failures can account for the symptoms you observed in your car. Find this out by successively **suspending constraints**.

Show your results on the following diagrams, one for each candidate. For each diagram, write at the top if the candidate **is subsumed by (can be pruned by) conflicts** that have been discovered so far. **If the candidate is not subsumed by conflicts discovered so far**, then:

- **Cross out** the suspended constraint.
- Perform constraint propagation and write the **predicted values** in the box next to the corresponding variable. You can stop propagation as soon as you have found that the candidate is inconsistent.
- Write the **support** in the box next to the predicted value. Do not record the pipe junction as support, we assume it to be fault-free.
- Write at the top of each diagram whether or not the candidate is **consistent**.
- If the candidate is inconsistent, extract a **conflict** and write it on top of the diagram.

Recall, the observations are P_B="+", P_{LW}="-", P_{RW}="+".

Candidate B:

Subsumed by conflicts (yes/no)? __No__ Consistent (yes/no)? ___Yes____
Conflict: ____None_____

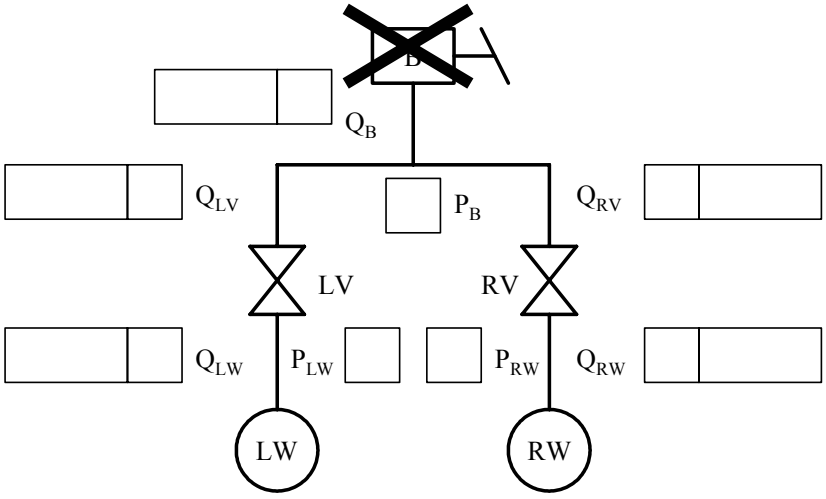

Assignment	Support
$P_B = +$, $P_{LW} = -$, $P_{RW} = +$	Observed
$Q_{LV} = +$	LV=G, $P_B = +$, $P_{LW} = -$
$Q_{LW} = +$	LV=G, $P_B = +$, $P_{LW} = -$
$Q_{RW} = +$	RW=G, $P_{RW} = +$
$Q_{RV} = +$	RV=G, $Q_{RW} = +$, $P_{RW} = +$
$Q_B = -$	(Pipe), $Q_{LV} = +$, $Q_{RV} = +$

Candidate LV:

Subsumed by conflicts (yes/no)? ___No_ Consistent (yes/no)? __Yes__
Conflict: _____

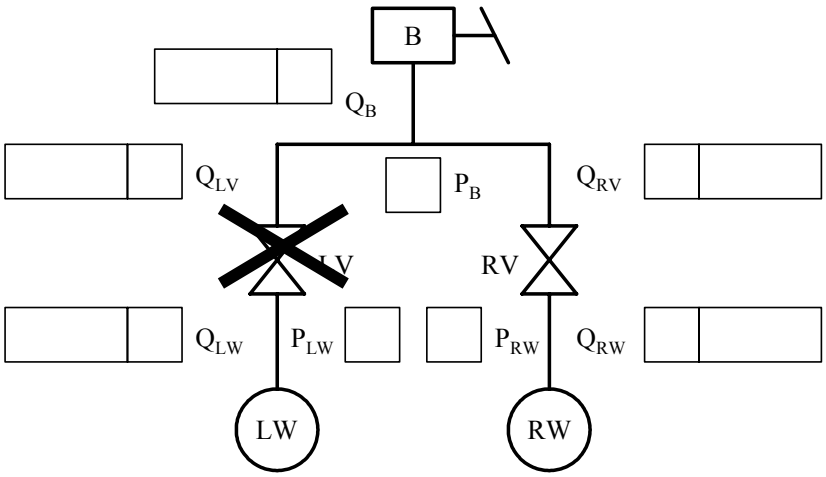

Assignment	Support
$P_B = +$, $P_{LW} = -$, $P_{RW} = +$	Observed
$Q_{RW} = +$	RW=G, $P_{RW} = +$
$Q_{RV} = +$	RV=G, $Q_{RW} = +$, $P_{RW} = +$, $P_B = +$
$Q_{LW} = +$	LW=G, $P_{LW} = +$
$Q_B = +$	B=G, $P_B = +$
$Q_{LV} = -$	(Pipe), $Q_B = +$, $Q_{RV} = +$

Candidate RV:

Subsumed by conflicts (yes/no)? __No____ Consistent (yes/no)? _Yes_____
Conflict: _____

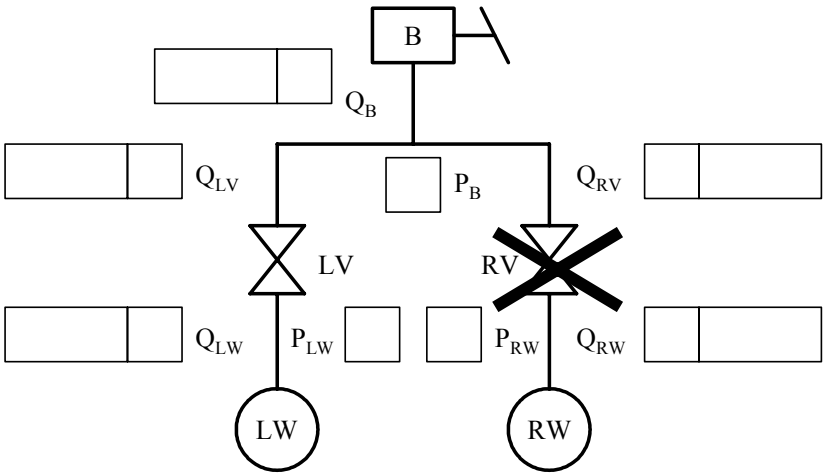

Assignment	Support
$P_B = +$, $P_{LW} = -$, $P_{RW} = +$	Observed
$Q_{LV} = +$	LV=G, $P_B = +$, $P_{LW} = -$
$Q_{LW} = +$	LV=G, $P_B = +$, $P_{LW} = -$
$Q_{RW} = +$	RW=G, $P_{RW} = +$
$Q_B = +$	B=G, $P_B = +$
$Q_{RV} = -$	(Pipe), $Q_B = +$, $Q_{LV} = +$

Candidate LW:

Subsumed by conflicts (yes/no)? __Yes____ Consistent (yes/no)? _____
Conflict: _____

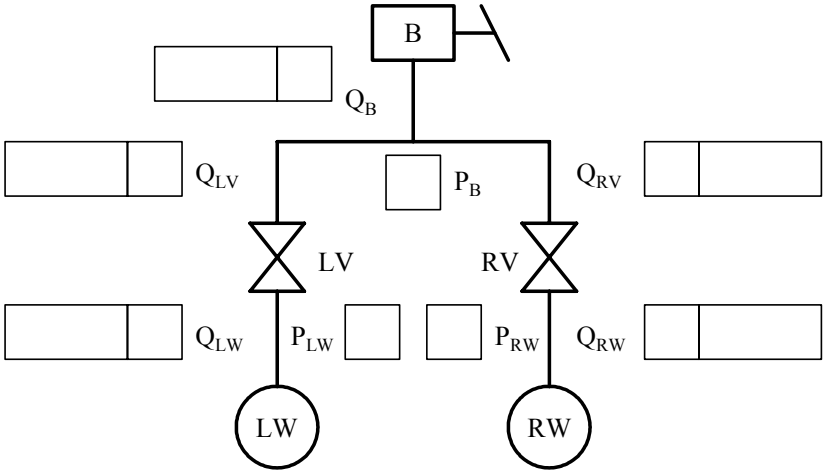

Candidate RW:

Subsumed by conflicts (yes/no)? __No____ Consistent (yes/no)? ___No____
Conflict: __Not (B=G and LV=G and RV=G)_____

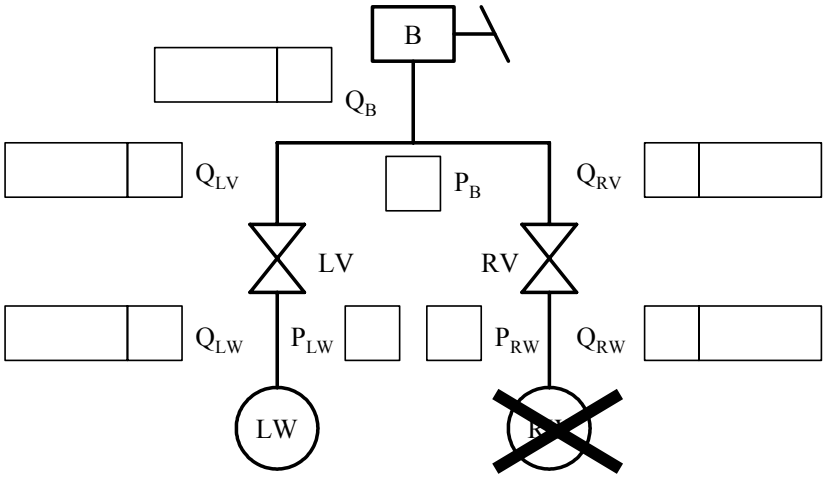

Assignment	Support
$P_B = +$, $P_{LW} = -$, $P_{RW} = +$	Observed
$Q_B = +$	B=G, $P_B = +$
$Q_{LV} = +$	LV=G, $P_B = +$, $P_{LW} = -$
$Q_{LW} = +$	LV=G, $P_B = +$, $P_{LW} = -$
$Q_{RV} = -$	(Pipe), $Q_{LV} = +$, $Q_B = +$
$Q_{RW} = \phi$	RV=G, $Q_{RW} = +$, $P_{RW} = +$, $Q_{RV} = -$

Problem 5 – Navigation by MDP (30 points)

Captain Jack Sparrow, infamous pirate, has sailed his ship to the eastern side of the island of Tortuga (see chart below).

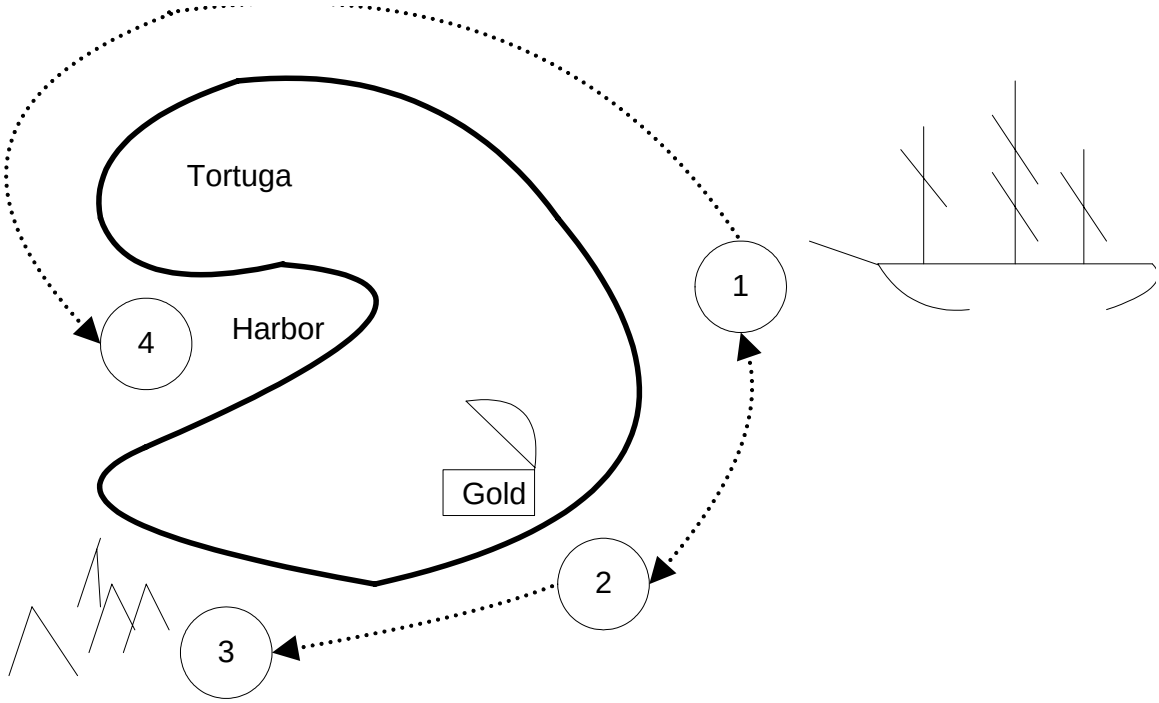

Captain Jack would like to anchor in the harbor on the western side. Let's help him by using an ancient navigation technique that is known to all sailors worth their salt: value iteration.

First, let's consider some details shown on the chart. There are four locations, with the dotted arrows indicating valid moves between them. We will assume that all moves are deterministic. Location 3 represents rocks that will sink the ship, so there are no actions that lead out of this state. Similarly, location 4 represents the goal, and there are no actions that lead out of location 4 (Captain Jack wants to relax after he has anchored). In order to relax, Captain Jack needs some gold. Fortunately, he remembers that he has previously stashed some at location 2.

Let's assume that any time the ship reaches location 2, and the ship is not carrying the gold, the gold is automatically loaded onto the ship.

To use an MDP formulation, we need some notion of reward. Let's assume that location 4, the goal, has a reward of 1000. Also, let's assume that location 3 has a reward of

-1000 (because the ship sinks). Finally, let's assume that location 2 has a reward of 200 if the ship has not yet picked up the gold. Rewards for actions in all other states are 0.

Part A - Modeling

In this part, you will design an MDP model for this problem. Assume that the state vector consists of the following two variables (with corresponding possible values):

Location (1, 2, 3, 4)
Ship-has-gold (true, false)

Part A-1 (5 points)

Write the transition function for this problem. Note that the system is deterministic.

Current state	Action	Next state
1, no gold	To 2	2, no gold
	To 4	4, no gold
1, gold	To 2	2, gold
	To 4	4, gold
2, no gold	To 1	1, gold
	To 3	3, gold
2, gold	To 1	1, gold
	To 3	3, gold

States 3 and 4 have no actions.

Part A-2 (5 points)

Write the reward function for this problem.

Current state	Action	Reward
2, no gold	Don't care	200
3 (gold or no gold)	Don't care	-1000
4 (gold or no gold)	Don't care	1000

For all other combinations of states and actions, the reward is 0.

Part B – Value Iteration (10 points)

Use value iteration to determine the values of each state. Perform two iterations using a discount factor of $\gamma = 0.9$. Assume that the initial value for all states corresponding to location 3 is -1000, the initial value for all states corresponding to location 4 is 1000, and the initial value for all other states is 0.

State	Initial value	Iteration 1 value	Iteration 2 value
1, no gold	0	900	900
1, gold	0	900	900
2, no gold	0	200	1010
2, gold	0	0	810

States for location 3 have a constant value of –1000 since this is the initial value, and since there are no actions leading out of this state. Similarly, states for location 4 have a constant value of 1000.

Part C – Optimal Policy (5 points)

What is the optimal policy for each state?

Assuming that V_2 is a reasonable approximation of V* after two iterations, then:

State	Action
1, no gold	to 2
1, gold	to 4
2, no gold	to 1
2, gold	to 1

Part D – Effect of discount factor (5 points)

What would be the effect of changing the discount factor from 0.9 to 0.5?

The algorithm would become more greedy, and would go directly to the harbor without picking up the gold.

Problem 6 - Scheming with CNF (30 points)

For this problem, we'll use a subset of propositional logic. In particular, we don't worry about negation (`not`), and the `and` and `or` operators are binary (as opposed to n-ary). Well-formed formulas (WFF's) are either propositional symbols (e.g. `a`), or a conjunction of two WFF's, (e.g. (`and` ... ...)), or a disjunction of two WFF's, (e.g. (`or` ... ...)). We can write the grammar for WFF's as:

WFF :: = Symbol
| (`or` WFF WFF)
| (`and` WFF WFF)

The following are WFF's:

```
1. a
2. (or a b)
3. (and a b)
4. (and (or a b) c)
5. (and (or a b) (and c (or d e)))
6. (and a (or (and b c) d))
7. (or (or a (and b c)) (and (or d e) f)))
```

A WFF is a *Disjunctive Normal Clause(DNC)* if it is a disjunction of terms, where each term is either a propositional symbol or a disjunctive normal clause. A single symbol can be considered a degenerate disjunctive normal clause. For example, formulas 1 and 2, above, are DNC. We can write the grammar for DNC as:

DNC ::= Symbol
| (`or` DNC DNC)

A WFF in *Conjunctive Normal Form (CNF)* consists of a conjunction of WFF's such that each conjunct is either in CNF itself or is a DNC. In the list of formulas above, WFF's 1 through 5 are in CNF. Note that formula 1 and 2 can be considered to consist of a single conjunct. A grammar for CNF formula is:

CNF ::= Symbol
| (`and` CNF CNF)
| (`or` DNC DNC)

An important step in converting logical formulas to CNF is to distribute `or` over `and` whenever possible. For example, we can convert (`or` A (`and` B C)) to (`and` (`or` A B) (`or` A C)), and we can convert (`or` (`and` A B) C) to (`and` (`or` A C) (`or` B C))

Part A – Convert to CNF Manually (10 points)

Convert the following WFF's to CNF by repeatedly applying `or` distribution.
Show intermediate steps, so that if you make a mistake we can give you partial credit.

```
(and a (or (and b c) d))

Recall:(or x (and y z))=>(and (or x y)(or x z)-distribution

Soln: (and a(and(or b d)(or c d))) by distribution
```

```
(or (or a b) (and c d))

Soln: (and (or (or a b) c)(or (or a b) d)) by distribution
```

```
(or (or a (and b c)) (and (or d e) f))

recall:(or(and A B)(and C D)
=>(and (and (or A C)(or A D))(and(or B C)(or B D)))

Soln: (or (and (or a b)(or a c)) (and (or d e) f))
=>
(and (and (or (or a b)(or d e))
          (or (or a b) f))
     (and (or (or a c)(or d e))
          (or (or a c) f)))
```

Part B – Scheme Code to Convert to CNF (20 points)

Following is the skeleton of a function to convert a WFF into CNF by distributing `or` over `and`:

```
(define first car)
(define second cadr)
(define third caddr)
;; Assume formulas are well-formed, and conform to the following grammar:
;;   wff ::= symbol | (and wff wff) | (or wff wff)
(define (and-clause? wff) (and (list? wff) (eq? 'and (first wff))))
(define (or-clause? wff) (and (list? wff) (eq? 'or (first wff))))
;; Distribute OR over AND in a well-formed formula (wff).
;; 1. (or (and A B) C) => (and (or A C) (or B C))
;; 2. (or A (and B C)) => (and (or A B) (or A C))
;; Note that in the above, if A, B, or C are not symbols or simple
;; disjunctive clauses, they will need to be further converted, so
;; that we wind up with, in general, wff's of the form
;; cnf  ::= symbol | (and cnf cnf) | (or dnc dnc)
;; dnc  ::= symbol | (or dnc dnc)
(define (distribute-or wff)           ; => CNF
  (cond
   ((symbol? wff)                     ; symbol
    wff)
   ((and-clause? wff)                 ; an AND clause
    (list 'and
        ;; [1] construct 1st conjunct here
        ;; [2] construct 2nd conjunct here
   ((or-clause? wff)                  ; an OR clause
    (let ((disj1 (distribute-or (second wff))) ; get disjuncts
          (disj2 (distribute-or (third wff)))) ; into CNF
      (cond
       ((and-clause? disj1)           ; case 1, above
       (let ((a (second disj1))
             (b (third disj1))
             (c disj2))
         (list 'and
             ;; [3] construct 1st conjunct here
             ;; [4] construct 2nd conjunct here
       ((and-clause? disj2)           ; case 2, above
       (let ((a disj1)
             (b (second disj2))
             (c (third disj2)))
         (list 'and
             ;; [5] construct 1st conjunct here
             ;; [6] construct 2nd conjunct here
       (else
       (list 'or
             ;; [7] construct 1st disjunct here
             ;; [8] construct 2nd disjunct here
             )))))))
```

Provide the eight Scheme expressions that should replace the eight commented lines in the code of the form "[n] construct ... here".

1. `(distribute-or (second wff))`
2. `(distribute-or (third wff))`
3. (list 'or a c)
4. (list 'or b c)
5. (list 'or a b)
6. (list 'or a c)
7. disj1
8. disj2

Problem 7 - Mixed Integer Programming (14 points)

Part A. Formulation with fixed costs (7 points)

Formulate the following mixed integer linear program. Stacy's Subway Company needs to supply 5 subway cars to Manhattan. The alternatives for shipping include transporting them by truck over a bridge (at a cost of $2000 each) or transporting them by ferry (at a cost of $1000 each). However, the permit process in Manhattan is absurd and so there is a one time red-tape cost of $6000 for transporting any subway cars in trucks over a bridge and a one time red-tape cost of $10,000 for transporting any subway cars by ferry.

Formulate a MILP to minimize the transportation costs.

Minimize $z = x_1 + 2x_2 + 10b_1 + 6b_2$

subject to $x_1 + x_2 \geq 5$

$x_1 \leq Mb_1$ [Here, accept M or any numerical value ≥ 5]

$x_2 \leq Mb_2$ [Here, accept M or any numerical value ≥ 5]

b_1, b_2 are binary (i.e., $b_i \in \{1,0\}$).

Part B. Branch and Bound Search Tree (7 points)

Solve the following mixed integer linear problem using Branch and Bound.

Minimize $z = 10x_1 + 4x_2 + 20b_1 + 25b_2$

subject to
$2x_1 + x_2 \geq 10$
$x_1 \leq 10b_1$
$x_2 \leq 10b_2$
$x1 \geq 0, x2 \geq 0$

$b_1 + b_2 \leq 1$
b_1, b_2 are binary (i.e., $b_i \in \{1,0\}$).

Fill in the branch and bound search tree below. Branch on the binary variables in the **following order**: b_1, b_2. Evaluate the **0** branch **before** the **1** branch. Cross off each node that is infeasible or fathomed. For feasible, non-fathomed nodes, give the relaxed solution and the value of Z. For fathomed nodes, give the solution and value of Z.

Derivation of Answer:

You can solve each relaxed problem using the simplex method. This problem is simple enough that you can solve it by inspection. The solution below is based on the use of monotonicity arguments to determine that particular constraints are active (this is called monotonicity analysis, or activity analysis).

We start by initializing the incumbent to:

z*= infinity
b1=?, b2=?, x1=?, x2=?,

We then relax the binary variables at the Root Node and solve:

Root Node:

Minimize $z = 10x_1 + 4x_2 + 20b_1 + 25b_2$

subject to $2x_1 + x_2 \geq 10$
$x_1 \leq 10b_1$
$x_2 \leq 10b_2$
$x1 \geq 0$, $x2 \geq 0$
$b_1 + b_2 \leq 1$
$b_1 \leq 1$, $b_2 \leq 1$
$b1 \geq 0$, $b2 \geq 0$

z monotonically decreases as x1 and x2 monotonically decrease, hence z is a minimum when
$2x_1 + x_2 \geq 10 \rightarrow 2x_1 + x_2 = 10$ (i.e., the constraint is said to be active)

Note that this argument holds unless x1 and x2 reach another constraint boundary first. We can test this at the end by checking feasibility of the solution.

Next, solving for x_2 produces:

$x_2 = 10 - 2x_1$

Substituting for x_2 into the problem and simplifying produces:

Minimize $z = 2x_1 + 20b_1 + 25b_2 + 40$
subject to $x_1 \leq 10b_1$
$10 \leq 10b_2 + 2x_1$
$x_1 \geq 0$, $5 \geq x_1$
$b_1 + b_2 \leq 1$
$b_1 \leq 1$, $b_2 \leq 1$, $b_1 \geq 0$, $b_2 \geq 0$

z monotonically decreases as x_1 and b_2 monotonically decrease, hence z is a minimum when
$10 \leq 10b_2 + 2x_1 \rightarrow 10 = 10b_2 + 2x_1$ (i.e., active constraint).

Solving for x_1:

$x_1 = 5 - 5b_2$

Substituting for x_1 into the problem and simplifying produces:

Minimize $z = 20b_1 + 15b_2 + 50$
subject to $5 \leq 10b_1 + 5\ b_2$
$1 \geq b_2,\ b_2 \geq 0$
$b_1 + b_2 \leq 1$
$b_1 \leq 1,\ b_2 \leq 1,\ b_1 \geq 0,\ b_2 \geq 0$

z monotonically decreases as b_1 and b_2 monotonically decrease, hence z is a minimum when
$5 \leq 10b_1 + 5\ b_2 \rightarrow 5 = 10b_1 + 5\ b_2$ (i.e., active constraint).

Solving for b_2:
$b_2 = 1 - 2b_1$

Substituting for b_2 into the problem and simplifying produces:
Minimize $z = -\ 10b_1 + 65$
subject to $b_1 \geq 0,\ b_1 \leq .5$

z monotonically decreases as b_1 monotonically increases, hence z is a minimum when
$b_1 \leq .5 \rightarrow b_1 = .5$ (i.e., active constraint).
Substituting for b_1 into the problem and simplifying produces:
$z = -\ 10(.5) + 65 = 60$
Substituting b_1 into the active constraints produces:
$b_2 = 1 - 2b_1 = 1 - 2(.5) = 0$
$x_1 = 5 - 5b_2 = 5 - 5(0) = 5$
$x_2 = 10 - 2(5) = 0$
To summarize, the relaxed solution is:
$b_1 = .5,\ b_2 = 0,\ x_1 = 5,\ x_2 = 0,\ z = 60$

This relaxed solution is feasible (substituting into the inequality constraints to confirm). It is not integer, and the relaxed solution is better than the incumbent. Hence we branch on the next variable, b_1.

Node $b_1 = 0$:
Substituting for b_1 induces an assignment for x_1:
$x_1 \leq 10(0),\ x_1 \geq 0 \quad \rightarrow x_1 = 0$

Eliminating b_1 and x_1, the problem simplifies to:
Minimize $z = 4x_2 + 25b_2$
subject to $x_2 \geq 10$
$x_2 \leq 10b_2$
$x2 \geq 0$
$b_2 \leq 1,\ b2 \geq 0$

z decreases monotonically with x2, reaching a minimum when:
$x_2 \geq 10 \quad \rightarrow x_2 = 10$ (active constraint)

z also decreases monotonically with b2, reaching a minimum when:
$10 \leq 10b_2 \quad \rightarrow 10 = 10b_2 \rightarrow b_2 = 1$ (active constraint)

substituting x_2 and b_2 into z produces:

$z = 4(10) + 25(1) = 65$

Thus, the relaxed solution is:

b1=0, b2=1, x1=0, x2=10, z=65

All integer variables are integer in the relaxed solution, hence the node is fathomed.
The solution is better than the current incumbent, thus the incumbent is changed to:

$z^* = 65$

b1=0, b2=1, x1=0, x2=10

Node b1 = 1:

Substituting b1 into the constraints induces two additional assignments:

$1 + b_2 \leq 1,\ b_2 > 0 \quad \rightarrow b_2 = 0$

$x_2 \leq 10(0),\ x2 \geq 0 \quad \rightarrow x_2 = 0$

Substitutuing these three assignments into the original problem produces:

Minimize $z = 10x_1 + 20$

subject to $2x_1 \geq 10$

$x_1 \leq 10,\ x1 \geq 0$

z monotonically decreases with x1. Reaching a minimum when:

$2x_1 \geq 10 \rightarrow 2x1 = 10 \rightarrow x1 = 5$ (active constraint)

Substituting this assignment into z produces:

$z = 10(5) + 20 = 70$

Hence the solution to the relaxed problem is:

b1=1, b2=0, x1=5, x2=0, z=70

This is worse than the incumbent, hence we fathom this node, b1 = 1.

The tree is fully explored, hence the optimal solution is the incumbent:

$z^* = 65$, b1=0, b2=1, x1=0, x2=10

Problem 8 - Decision Trees (15 points)

City planners have accumulated the following data on a variety of threats to urban life. The purpose of this data is to predict, based upon a monster's origin, appearance, and breath attack (if any), whether or not it will practically raze a city before it retires for mid-afternoon tea and crumpets.

Origin	**Appearance**	**Foul Breath**	**City-Destroyer**
Outer Space	Blob	Acid	Yes
Outer Space	Reptile	None	Yes
Outer Space	Reptile	Acid	Yes
Outer Space	Blob	Fire	No
Pacific Ocean	Blob	None	No
Pacific Ocean	Blob	Acid	No
Pacific Ocean	Reptile	Fire	No
Pacific Ocean	Reptile	Acid	No

1. (5 points) With no other information, how many bits on average would you need to transmit City-Destroyer ?

$$H(3/8, 5/8) = -3/8 \times \log_2(3/8) - 5/8 \log_2(5/8)$$

2. (1 point) Which of the three other attributes gives you the highest information gain with respect to City-Destroyer? Show your reasoning.

By inspection, the Origin attribute is the best since it only makes one mistake when used as a classifier and has high entropy. The others make more mistakes.

3. (2 points) Using information gain as your splitting criteria, what decision tree would you get? Do not prune the tree.

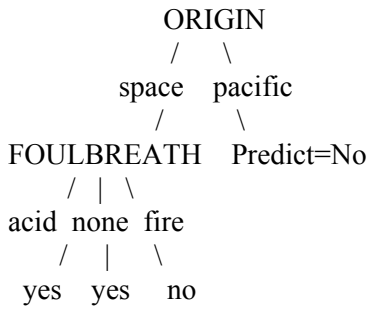

4. (2 points) Fill in the following table with your predictions.

Origin	**Appearance**	**Foul Breath**	**City-Destroyer**
Outer Space	Blob	Fire	No
Pacific Ocean	Reptile	Acid	No

5. (5 points) How might you deal with missing values? For example, what would you predict if all you knew about a visitor was that it was blob-shaped, and why?

Any of the following answers:
1. Use a Bayes Classifier instead
2. Retrain the decision tree using only the attributes that are present in the query
3. Guess the most common value as the missing value.

Problem 9- Hidden Markov Models (15 points)

Consider an HMM defined by the traditional set of parameters:

$$\lambda = (S,Z,T,O,p_0)$$

Recall that the following array of probabilities can be obtained by dynamic programming in the forward algorithm:

$$\alpha_t(i) = p(z_1,..., z_t \wedge q_t = s_i|\lambda)$$

a) We can compute a similar array of probabilities using dynamic programming in a backward algorithm:

$$\beta_t(i) = p(z_t,..., z_T \wedge q_t = s_i|\lambda)$$

Give a mathematical expression for $\beta_\tau(\iota)$ that can be computed using dynamic programming.

$$\beta_t(i) = \Sigma^S_{j=1} T(s_i,s_j)O(s_j,z_{t+1})\beta_{t+1}(j)$$

b) In the process of learning an HMM from data we need the following array of probabilities:

$$\gamma_t(i) = p(q_t = s_i|z_1, ..., z_t, z_{t+1}, ..., z_T, \lambda)$$

Your job is to define $\gamma_t(i)$ in a simple finite expression in terms of $\alpha_t(i)$, $\beta_t(i)$ and $p(z_1, ..., z_t, z_{t+1}, ..., z_T, \lambda)$.

$$\gamma_t(i) = \alpha_t(i)\beta_t(i)/p(z_1, ..., z_t, z_{t+1}, ..., z_T, \lambda)$$