

Assignment 1: Vacuum World Planner
CS 730/730W/830, Spring 2012

Electronic submission of milestone due at **11:30pm on Sun, Feb 5**

Full electronic submission due at **11:30pm on Sun, Feb 12**

Hardcopy submission due at **1:10pm in class on Mon, Feb 13**

Overview

You will write a program to solve robot planning problems in Vacuum World, a simple grid environment where some cells are contaminated with dirt. This problem is fundamentally similar to game AI path-finding and NASA rover surface operations. Given a world description, your program will return a sequence of actions that will move the robot around and vacuum up all the dirt. The vacuum robot can only move in the four cardinal directions. All actions have the same cost and the objective is to achieve a clean world with a plan of minimum cost.

Input

Your program should read a world description from standard input, such as:

```
4
3
_ * _ _
_ _ # *
_ @ * #
```

The format is:

1. The number of columns and the number of rows, on separate lines. Worlds might be hundreds or thousands of cells in each dimension.
2. Rows of characters, one character per cell and one world row per line of input. An underscore (`_`) means a blank cell, a hash (`#`) means a blocked cell that the robot cannot enter, an asterisk (`*`) means a dirty cell. An at-sign (`@`) is the starting position of the robot.

You can assume that the world is correctly formatted and solvable.

Your program should accept up to two command-line arguments:

algorithm one of `depth-first`, `depth-first-id` (depth-first iterative deepening), `uniform-cost`, or `a-star`

heuristic if the algorithm is `a-star`, this second argument will specify either `h0` ($h(n) = 0$), `h1` (a heuristic you design), or `h2` (an even better heuristic you design). Your heuristics must be admissible.

Output

Your program should write only the following to standard output: a list of actions (`N`, `S`, `E`, `W`, or `V` (vacuum), one per line) followed by the number of search nodes expanded and generated (one per line, in that order). For example:

```
katsura(1153)> ./run.sh uniform-cost < tiny-2.vw
E
V
W
N
N
V
E
E
S
```

85 nodes generated

39 nodes expanded

Execution

Read the submissions instructions on the course web page for information on the required `make.sh` and `run.sh` scripts. Basically, if you are using a compiled language, you should supply a `make.sh` script that builds your code and the final executable should be named `run.sh`. If you are using an interpreted or JIT compiled language, your `make.sh` script can do nothing and your `run.sh` should load and run your program with the arguments supplied to the script.

On the course web page, we supply some utility programs and input files. Most of the programs we distribute in this class will tell you their command-line arguments if you run them with the `-help` option.

`*.vw` a few example benchmark problems.

`make-vw` takes four command line arguments (the number of rows, number of columns, percentage of blocked cells (eg, 0.15), and number of dirty cells) and prints a new random world to standard output. There are some optional arguments too — try `-help`.

`vacuum-plan-reference` a sample solution. It actually has three informed heuristics (`h1-h3`) and then two additional ones (`h4` and `h5`) for the graduate extension (described below).

`vw-validator` runs your program, validates its output, and displays the solution plan (using digits to show progression from start to end). For example:

```
vw-validator -time 5 -- run.sh a-star h1 < tiny-1.vw
```

will run A* on `tiny-1.vw` using your first heuristic and a 5 second time limit.

As the example illustrates, the validator takes two groups of arguments, separated by `--`. The first group are optional flags to control the validator's behavior: `-novis` turns off display of the solution plan (useful for large problems), `-time float` specifies a time bound in seconds (default 60), and `-mem float` specifies a memory bound in gigabytes (default 1). The second group of arguments are the program to run and the arguments to pass to it. The validator expects a world on standard input and will pass it to your program.

Intermediate Milestone

Your milestone solution should correctly implement the `depth-first` search algorithm. You should use only linear memory, but be complete in finite domains like vacuum world. Electronically submit your source code as well as a transcript of your program running with the validator on the supplied example problems.

Submission

Electronically submit your milestone and final solution using the instructions on the course web page. Submit a brief write-up with your final hardcopy solution answering the following questions:

1. What is the size of the state space for this problem?
2. Explain the heuristic functions you devised and why they are admissible.
3. Describe any implementation choices you made that you felt were important. Mention anything else that we should know when evaluating your program.
4. What is the time and space complexity of each algorithm you implemented? Which algorithms are admissible?
5. Provide empirical results confirming your answers to the previous question.

6. What suggestions do you have for improving this assignment in the future?

Attach to the write-up a transcript of your code solving the example cases using the validator, and a listing of your source code (2 pages per page, as with `a2ps -2`).

Graduate Extensions

Those in 830 must extend the base assignment in two ways. First, you must support two additional algorithms `ida-star` (iterative-deepening A*) and `greedy` ($f(n) = h(n)$). Second, your program must accept as an optional last command-line argument the flag `-battery` that indicates, when present, that you must model the battery usage of the vacuum robot. Special cells (marked `:`) are recharging stations and a special action (`R`) completely recharges the battery when in those cells. Each action (except `R`) uses one unit of power. The robot starts with a full charge, which is defined to be exactly enough to move from one corner of the world to the opposite one, vacuum the cell, and move back, assuming no obstacles. You should implement an additional admissible heuristic function `h3` that takes the battery into account somehow.

Evaluation

If you are in 730W, 2 points of your grade will depend on how clean, elegant, and readable your code is. Otherwise, we will look at your code mainly to try to give you partial credit. For the milestone:

0 does not build or does not work with validator

1 works with validator but does not find proper results

2 produces a correct result

For the final submission:

0 nothing

1 something but basically nothing

3 write up is correct but no code works

6 some basic algorithm works

9 a heuristic search algorithm works too

12 One of the algorithms doesn't really work. Eg, no handling of duplicates, or takes gobs of memory, or code is a mess. Write-up is ambiguous and not convincing.

15 Very slow on large problems, but code looks nice. Write-up has an error in reasoning.

18 A* is quick even on large problems. Write-up is clear and convincing, with no errors.

Extra credit will be given to the best-performing undergraduate and graduate student solutions, meaning fastest `a-star` with `h2` on large worlds with many obstacles and dirty cells (`h3` and battery for grad students).

Design Suggestions

Make the state representation (stuff like current location, remaining dirt, actions done so far) separate from the node representation (stuff like g and h values). You probably want to make both of those separate from the static problem information (eg, which cells are blocked) that doesn't change from node to node. Ideally, these should all be independent from the search algorithms (we will be using A* again in assignment 2). One way to do this is to have each search algorithm take an initial state, state expansion function, and goal test function and do all its work in terms of those functions, with no other knowledge about the problem being solved. (The heuristic search algorithms will obviously take the heuristic as an additional argument.)

Start by making sure you can read in and print out a world state. Use very small examples until you are sure your basic functions work correctly.

Detecting and handling duplicate states is very important in this domain.

Use appropriate data structures for the state representation, the open list, and the closed list. Consider such options as: linked list, array, hash table, binary heap.