

Planning and Optimization

M. Helmert, G. Röger
C. Büchner, T. Keller, S. Sievers

University of Basel
Fall Semester 2021

Exercise Sheet A

Due: October 14, 2021

Important: for submission, consult the rules at the end of the exercise. Non-adherence to the rules will lead to your submission not being corrected.

The files required for this exercise are in the directory `exercise-a` of the course repository (<https://github.com/aibasel-teaching/planopt-hs21>). All paths are relative to this directory. Update your clone of the repository with `git pull` to see the files. In the virtual machine, `/vagrant/plan-opt-hs21` is the repository.

Exercise A.1 (0.5+0.5+0.5+0.5 marks)(Lecture A2)

Characterize the following planning domains by describing if they are

- *fully* or *partially observable*,
- *deterministic* or *stochastic*,
- *instantaneous* or *durative*, and
- whether the planning agent is the *sole source of change* or there are *exogenous events*.

- | | |
|--------------------|--------------------|
| (a) Route planning | (c) Cooking a meal |
| (b) Solitaire | (d) Autonomous car |

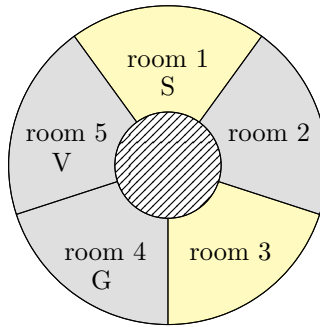
Note: Since these properties depend on the explicit model of each domain, make reasonable assumptions and justify your answer if you feel the need for explanation.

Exercise A.2 (3+2+3+2 marks)(PDDL Intro from Exercise Session 1 for (a) and (b), Lecture A5 for (c) and (d))

Consider the following planning task:

- You are trapped in the cellar of a building with a switch board full of light switches. In the rooms above you there is a vampire (V). Luckily, there also is a vampire slayer (S) in those rooms. To keep things simple, we consider only room layouts that are circular corridors where each room has a clockwise and an anti-clockwise neighbor.
- The vampire avoids the light: whenever the light in the vampire's room is switched on, it moves to a neighboring room. If one of the neighboring rooms is dark, it will move there, preferring the anti-clockwise one if both are dark. If both neighboring rooms are bright, it will move clockwise.
- The slayer tries to stay in the light. If the light in her room is switched off, she moves to a neighboring room. She moves clockwise if that room is bright and anti-clockwise otherwise.
- If the two of them meet in a room they will fight. The vampire wins the fight in a dark room unless there is garlic (G) in that room. In bright rooms or in rooms with garlic, the slayer wins.
- All you can do is use the switch board to toggle lights and watch the fight, when it happens. Your objective is to toggle the lights so that the slayer can win the fight.

Example instance with five rooms:



- (a) There is a partial model of this domain in the directory `vampire`. Complete it by adding the effects of `toggle-light` and `watch-fight`. Do not add new actions or predicates.

The directory also contains instances which you can use for debugging. Their optimal plan costs are 6, 4, 7, 5, 4, 12, 11, 10, 13, and 8.

You can use `INVAL` for debugging your PDDL code: `inval <domain.pddl> <problem.pddl> <plan>` where `problem` and `plan` file are optional.

- (b) PDDL uses first-order predicate logic to model planning tasks. However, the models discussed in the lecture are all based on propositional logic. Most planners convert PDDL into one of the propositional models in a step called *grounding*. The directory `preprocess` contains a Python tool to do this step. The call

```
./preprocess/ground.py vampire/domain.pddl vampire/p01.pddl
```

will create a new domain file `vampire/domain_grounded_for_p01.pddl` and a new task file `vampire/p01_grounded.pddl`. Repeat this for all task files and describe the effect of the grounding procedure.

- (c) In addition to `ground.py` there is an incomplete Python program called `transform.py` in the directory `preprocess` which should transform grounded domains into conflict-free effect normal form. Complete the missing parts and use it to transform your grounded domains from exercise (b) into conflict-free effect normal form. The call

```
./preprocess/transform.py vampire/domain_grounded_for_p01.pddl
```

will create the file `vampire/domain_grounded_for_p01_normalized.pddl`.

- (d) Use Fast Downward to generate plans for all tasks using the domains you created in exercises (a)–(c). Then use `INVAL` to validate each plan against each domain formulation. In which combinations are the plans valid? Discuss the reason for that.

Exercise A.3 (2 marks)(Lecture A3)

Consider the transition system given in Figure ???. Write down true statements about it using all terms from the following list. The goal of this exercise is to confirm that you understand the meaning of these terms.

- *successor*
- *predecessor*
- *initial state*
- *goal state*
- *unreachable state*
- *solution*
- *path length*
- *path cost*




Figure 1: Example transition system with cost function c such that $c(\ell_i) = i$ for all $i \in \{0, 1, 2, 3\}$.

Exercise A.4 (0.5+0.5+0.5+0.5 marks)(Lecture A4)

Let $o = \langle \neg a \wedge \neg b, a \wedge (c \triangleright b) \wedge ((\neg c \wedge d) \triangleright (a \wedge b)) \rangle$ be an operator for a planning task Π . Provide the following effect conditions:

- $effcond(\neg a, eff(o))$
- $effcond(b, eff(o))$
- $effcond(c, eff(o))$
- $effcond(a \wedge b, eff(o))$

It is sufficient to give the final result.

Exercise A.5 (4 marks)(Lecture A4)

In this exercise we consider the solitaire game Beleguered Castle (http://justsolitaire.com/Beleguered_Castle_Solitaire/). It consists of a deck of cards stacked face-up in several tableau piles. For each suit in the deck there is a discard pile consisting only of the ace initially. There are three types of legal moves:

- The top card of a tableau pile can be moved on top of another tableau pile if the top card of the target pile has a value that is one higher. The suit of both cards does not matter for this move. For example, $2\clubsuit$ can be moved on $3\heartsuit$, $10\clubsuit$ on $J\spadesuit$, or $Q\diamondsuit$ on $K\diamondsuit$.
- The top card of a tableau pile can be moved to an empty tableau pile. This is allowed for all cards (not just for kings as in other solitaire games).
- The top card of a tableau pile can be moved to the discard pile for the matching suit if the top card on the discard pile has a value one lower. For example, if $7\heartsuit$ was discarded last, then $8\heartsuit$ can be discarded next. Discarded cards can never be moved again.

We consider a parameterized version of the game with m tableau piles $Tableaus = \{t_1, \dots, t_m\}$ and any set of cards $Cards$. For a given card $c \in Cards$ we use $suit(c)$ and $value(c)$ to refer to its suit and numerical value. The set of discard piles contains one discard pile for each suit: $Discards = \{discard_s \mid s = suit(c) \text{ for a } c \in Cards\}$. The set of all piles is $Piles = Tableaus \cup Discards$.

Model Beleguered Castle as a family of propositional planning tasks. Use the following state variables:

- $c\text{-on-}x$ for all $c \in Cards$ and $x \in Cards \cup Piles$
For cards c_1, c_2 the variable $c_1\text{-on-}c_2$ should be true iff c_1 is directly on top of c_2 ; For a pile p the variable $c_1\text{-on-}p$ should be true iff c_1 is directly on top of the pile, i.e., iff c_1 is the bottom-most card in the pile p .

- *x-clear* for all $x \in Cards \cup Tableaus$
An object should be clear iff there is no card on top of it.
- *c-discarded* for all $c \in Cards$
A card is discarded iff it is on the discard pile.

Exercise A.6 (4+2+2 marks)(Lecture A6 for (a) and (b), Lecture A7 for (c))

Consider the following planning domain: an agent is moving on a map and is trying to reach a specific target location. However, there are locked doors between some locations that can only be unlocked by first picking up their corresponding key from some other location. The agent can hold only one key at a time.

Let $G = \langle N, E \rangle$ be an undirected graph which represents the map (nodes correspond to locations, edges to doors), let $s, t \in N$ be the start and target location of the agent, let $L \subseteq E$ be the set of locked doors, let $Keys$ be the set of keys and let the function $unlocks: Keys \rightarrow L$ denote which door is unlocked by a key and the function $initial-location: Keys \rightarrow N$ denote where each key is located in the beginning.

The domain can then be modeled as a family of FDR tasks $\Pi = \langle V, I, O, \gamma \rangle$ where:

- V consists of the following variables:
 - Variable **at** with domain N denotes the position of the agent.
 - Variable **{a, b}-locked** for each $\{a, b\} \in E$ with domain $\{\top, \perp\}$ denotes whether or not an edge is locked.
 - Variable **holding** with domain $Keys \cup \{none\}$ denotes which key the agent currently holds.
 - Variable **k-location** for each $k \in Keys$ with domain $N \cup \{door\} \cup \{agent\}$ denotes the position of the key (at a location, in an unspecified door or carried by the agent).
- I sets above variables as follows:
 - $I(\mathbf{at}) = s$,
 - $I(\mathbf{\{a, b\}-locked}) = \top$ for all $\{a, b\} \in L$ and $I(\mathbf{\{a, b\}-locked}) = \perp$ otherwise,
 - $I(\mathbf{holding}) = none$,
 - $I(\mathbf{k-location}) = n$ for all $k \in Keys$, where $n = initial-location(k)$.
- O contains the following actions:
 - **move- $\langle a, b \rangle$** and **move- $\langle b, a \rangle$** for all $\{a, b\} \in E$, with precondition $\mathbf{at} = a \wedge \mathbf{\{a, b\}-locked} = \perp$ and effect $\mathbf{at} = b$ for action **move- $\langle a, b \rangle$** and analogous precondition and effect for action **move- $\langle b, a \rangle$** .
 - **unlock- $\{a, b\}$** for all $\{a, b\} \in L$, with precondition $(\mathbf{at} = a \vee \mathbf{at} = b) \wedge \mathbf{holding} = k \wedge \mathbf{k-location} = \mathbf{agent} \wedge \mathbf{\{a, b\}-locked} = \top$ and effect $\mathbf{\{a, b\}-locked} = \perp \wedge \mathbf{holding} = none \wedge \mathbf{k-location} = \mathbf{door}$, where $k \in Keys$ with $unlocks(k) = \{a, b\}$.
 - **pick- k** for all $k \in Keys$, with precondition $\mathbf{at} = x \wedge \mathbf{holding} = none \wedge \mathbf{k-location} = x$ and effect $\mathbf{holding} = k \wedge \mathbf{k-location} = \mathbf{agent}$, where $x = initial-location(k)$.
- $\gamma = (\mathbf{at} = t)$

(a) Transform the model into an equivalent family of STRIPS planning tasks.

(b) Consider the example instance in Figure ?? in the given FDR model and your STRIPS model. Compare the state spaces of the two formulations. How many states do they have? How many of those states are reachable? (Note that we expect actual numbers for these two questions.) Are the state spaces the same? Are they equivalent? What remains the same, what changes?




Figure 2: Example instance with four locations (circles, numbered for easy reference), doors (edges) with specific colors, and keys (numbered k1 to k4) of different colors located initially at the locations as shown. The agent starts at location 1. The goal of the agent does not matter for the analysis of this exercise.

- (c) List the mutex groups over the state variables of your STRIPS task in part (a) and specify which variables in the given FDR task they would induce.

Exercise A.7 (0.5+1+0.5 marks)(Lecture A7)

Consider a propositional planning task Π with state variables $V = \{a, b, c, d, e, f, g, h\}$.

- (a) Let v_1 and v_2 be mutex in Π for all $\{v_1, v_2\} \in \{\{a, b\}, \{c, d\}\}$. Provide the invariant that covers exactly this information.
- (b) Let v_1 and v_2 be mutex in Π for all $\{v_1, v_2\} \in \{\{a, b\}, \{a, e\}, \{b, c\}, \{b, e\}, \{c, d\}, \{c, e\}, \{c, g\}, \{d, e\}, \{d, g\}, \{e, f\}, \{e, g\}, \{f, g\}, \{f, h\}, \{g, h\}\}$.

Provide all mutex groups G for which there is no mutex group G' such that $G \subset G'$.

- (c) Let $\{\{a, b, c, d\}, \{b, d, e\}, \{c, f\}, \{c, g\}\}$ be a set of mutex groups for Π . Provide two different mutex covers for Π .

Submission rules:

- Exercise sheets must be submitted in groups of three students. Please submit a single copy of the exercises per group (only one member of the group does the submission).
- Create a single PDF file (ending .pdf) for all non-programming exercises. Use a file name that does not contain any spaces or special characters other than the underscore “_”. If you want to submit handwritten solutions, include their scans in the single PDF. Make sure it is in a reasonable resolution so that it is readable, but ensure at the same time that the PDF size is not astronomically large. Put the names of all group members on top of the first page. Either use page numbers on all pages or put your names on each page. Make sure your PDF has size A4 (fits the page size if printed on A4).
- For programming exercises, only create those code textfiles required by the exercise. Put your names in a comment on top of each file. Make sure your code compiles and test it. Code that does not compile or which we cannot successfully execute will not be graded.
- For the submission: if the exercise sheet does not include programming exercises, simply upload the single PDF. If the exercise sheet includes programming exercises, upload a ZIP file (ending .zip, .tar.gz or .tgz; *not* .rar or anything else) containing the single PDF and the code textfile(s) and nothing else. Do not use directories within the ZIP, i.e., zip the files directly.

- Do not upload several versions to ADAM, i.e., if you need to resubmit, use the same file name again so that the previous submission is overwritten.