

Planning and Optimization

M. Helmert, G. Röger
F. Pommerening

University of Basel
Fall Semester 2017

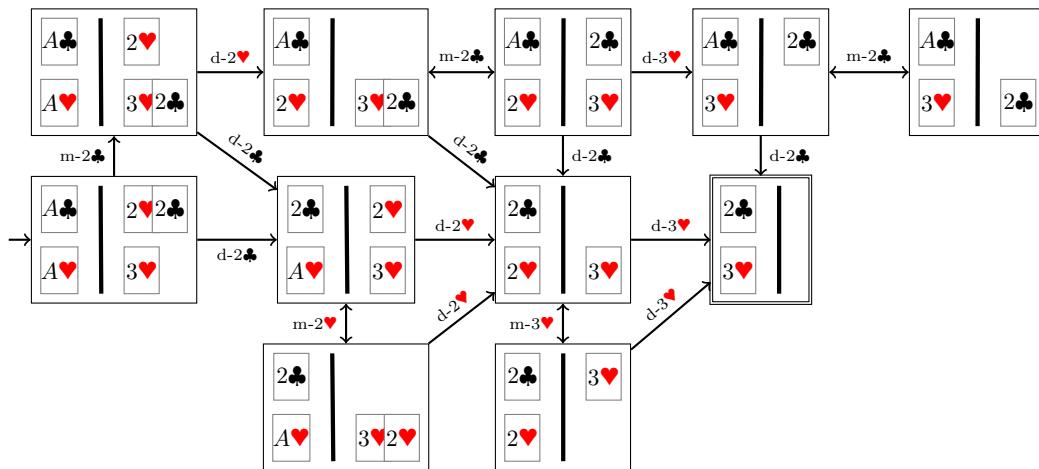
Exercise Sheet D

Due: November 21, 2017

The files required for this exercise are in the directory `exercise-d` of the course repository (<https://bitbucket.org/aibasel/planopt-hs17>). All paths are relative to this directory. Update your clone of the repository with `hg pull -u` to see the files. For the runs with Fast Downward, set a time limit of 1 minute and a memory limit of 2 GB. Using Linux, such limits can be set with `ulimit -t 60` and `ulimit -v 2000000`, respectively.

Exercise D.1 (4+2+1+4+2+2 marks)

(a) Consider the following graph G showing a simplified version of the reachable state space of the beleaguered castle instance from exercise A.1 (f). Note that some operators that were different in exercise A.1 (f) are now the same. Assume for this exercise that discarding a card has a cost of 5 while moving a card has a cost of 1:



Provide the following graphs:

- a graph G_1 which is isomorphic to G but not the same.
- a graph G_2 which is graph equivalent to G but not isomorphic to it.
- a graph G_3 which is a strict homomorphism of G but not graph equivalent to it.
- a graph G_4 which is a non-strict homomorphism of G but not graph equivalent to it.
- a graph G_5 that is the transition system induced by the abstraction α that maps states that are in the column i in the image above to the abstract state s_i . For example, the two states in the first column are mapped to an abstract state s_1 and the three states in the second column to an abstract state s_2 .
- a graph G_6 that is the induced transition system of an abstraction β that is a non-trivial coarsening of α .
- a graph G_7 that is the induced transition system of an abstraction γ that is a non-trivial refinement of β but different from α .

In all cases, highlight an optimal path and compute its cost. For graphs G_1 – G_4 also briefly argue (one sentence) why they don't have the property they are not supposed to have, for example, why G_2 is not isomorphic to G .

(b) Point out the problems with the following ideas for abstraction mappings in the beleaguered castle domain:

- α_1 : For each card value v there is one abstract state representing all world states where v is the highest undiscarded value.
- α_2 : A state is mapped to an abstract state by ignoring the suit of the top card on each tableau pile.
- α_3 : There are up to $n = 10^6$ abstract states s_0, \dots, s_n . A world state s is mapped to the abstract state s_k , where k is the MD5 hash of s modulo 10^6 .
- α_4 : All states s with $0 \leq h^*(s) < 5$ are mapped to the first abstract state, all states s with $5 \leq h^*(s) < 10$ are mapped to the second abstract state, and so on.

(c) Prove the following claim from the lecture: let α_1 and α_2 be abstractions of a transition system \mathcal{T} . If no label of \mathcal{T} affects both \mathcal{T}^{α_1} and \mathcal{T}^{α_2} , then α_1 and α_2 are orthogonal.

(d) Let Π be a SAS⁺ planning task that is not trivially unsolvable and does not contain trivially inapplicable operators, and let P be a pattern for Π . Prove that $\mathcal{T}(\Pi|_P) \xrightarrow{G} \mathcal{T}(\Pi)^{\pi_P}$, i.e., $\mathcal{T}(\Pi|_P)$ is graph-equivalent to $\mathcal{T}(\Pi)^{\pi_P}$.

(e) Discuss the theorem from exercise (d). What is it good for? Why is it important to exclude trivially unsolvable tasks or trivially inapplicable operators?

(f) Provide an FDR task that is not trivially unsolvable and does not contain trivially inapplicable operators, and a pattern P for Π such that $\mathcal{T}(\Pi|_P) \not\xrightarrow{G} \mathcal{T}(\Pi)^{\pi_P}$. Explain your solution.

Exercise D.2 (4+3+2+3+2+4+2 marks)

(a) In the files `fast-downward/src/search/planopt_heuristics/projection.*` you can find an incomplete implementation of a class projecting a TNF task to a given pattern. Complete the implementation by projecting the initial state, the goal state and the operators.
The example task from the lecture and two of its projections are implemented in the method `test_projections`. You can use them to test and debug your implementation by calling Fast Downward as `./fast-downward.py --test-projections`.

(b) In the files `fast-downward/src/search/planopt_heuristics/pdb.*` you can find an incomplete implementation of a pattern database. Complete the implementation by computing the distances for all abstract states as described in the code comments.
You can use the built-in implementation of Fast Downward to debug your code as explained in exercise (c).

(c) Use the heuristic `pdb(pattern=greedy(1000))` to find a good pattern with at most 1000 abstract states for each instance in the directory `castle`. Then run your implementation from exercise (b) using the heuristic `planopt_pdb(pattern=P)`. For each instance use the same pattern P used by the built-in implementation.
 Compare the two implementations and discuss the preprocessing time, the search time, and the number of expanded states excluding the last f -layer (printed as “Expanded until last jump”). Repeat the experiment for 100000 abstract states and compare the results.

(d) In the files `fast-downward/src/search/planopt_heuristics/canonical_pdb.*` you can find an incomplete implementation of the canonical pattern database heuristic. Complete the implementation in the methods `build_compatibility_graph` and `compute_heuristic` to create the compatibility graph for a given pattern collection and for computing the heuristic value given the maximal cliques of that graph.
You can use the built-in implementation of Fast Downward to debug your code as explained in exercise (e).

(e) Use the heuristic `cpdbs(patterns=combo(1000))` to find a good pattern collection with at most 1000 abstract states for each instance in the directory `nomystery-opt11-strips`. Then run your implementation from exercise (d) using the heuristic `planopt_cpdbs(patterns=C)`. For each instance use the same pattern collection C used by the built-in implementation.

Compare the two implementations and discuss the total time, and the number of expanded states excluding the last f -layer (printed as “Expanded until last jump”). Also compare your results to using a single pattern database heuristic with up to 1000 abstract states as in exercise (c).

(f) In the files `fast-downward/src/search/planopt_heuristics/pattern_hillclimbing.*` you can find an incomplete implementation of the hill-climbing search for iPDB. Complete the implementation in the methods `fits_size_bound`, `compute_initial_collection`, `compute_neighbors`, and `run` by following the comments in the code.

The built-in implementation of Fast Downward is different in some details so its results are not guaranteed to match those of your implementation.

(g) Use the heuristic `ipdbs(collection_max_size=1000, min_improvement=1)` on the instances in the directory `nomystery-opt11-strips`. This will use hill climbing to find a pattern collection with at most 1000 abstract states and continues as long as there is a state with an improved heuristic value. Then run your implementation from exercise (f) using the heuristic `planopt_ipdb(size_bound=1000)`.

Compare the two implementations and discuss the total time and the number of expanded states excluding the last f -layer (printed as “Expanded until last jump”). Also compare your results to the results from exercise (e).

The bash scripts in the directory `scripts` can be extended to run your experiments.

Exercise D.3 (7+5+3 marks)

(a) Consider the planning domain “Gripper” where one robot with two arms (“grippers”) moves n balls from a room A to a room B . It consists of the planning tasks $\Pi_n = \langle V, I, O, \gamma \rangle$ with

- Variables $V = \{at_r, empty_L, empty_R\} \cup \{at_{b_i} \mid 1 \leq i \leq n\}$ with $dom(at_r) = \{A, B\}$, $dom(empty_L) = dom(empty_R) = \{\mathbf{T}, \mathbf{F}\}$, $dom(at_{b_i}) = \{A, B, L, R\}$ for $1 \leq i \leq n$.
- Initial state $I = \{at_r \mapsto B, empty_L \mapsto \mathbf{T}, empty_R \mapsto \mathbf{T}\} \cup \{at_{b_i} \mapsto A \mid 1 \leq i \leq n\}$

- Operators $O = \{move_{A,B}, move_{B,A}\}$
 - $\cup \{pick-up_{i,x,g} \mid 1 \leq i \leq n, x \in \{A, B\}, g \in \{L, R\}\}$
 - $\cup \{drop_{i,x,g} \mid 1 \leq i \leq n, x \in \{A, B\}, g \in \{L, R\}\}$

where

- $move_{x,y} = \langle (at_r = x), (at_r := y), 1 \rangle$
- $pick-up_{i,x,g} = \langle (at_r = x) \wedge (at_{b_i} = x) \wedge (empty_g = \mathbf{T}), (at_{b_i} := g) \wedge (empty_g := \mathbf{F}), 1 \rangle$
- $drop_{i,x,g} = \langle (at_r = x) \wedge (at_{b_i} = g), (at_{b_i} := x) \wedge (empty_g := \mathbf{T}), 1 \rangle$

- Goal $\gamma = \bigwedge_{1 \leq i \leq n} (at_{b_i} = B)$

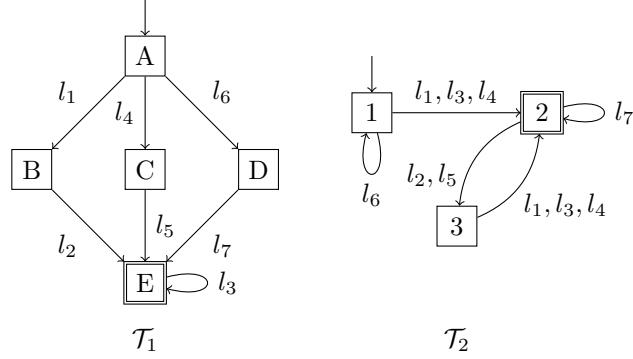
For simplicity, we assume that n is even. Then the optimal solution of task Π_n has cost $3n$ (each ball is picked up and dropped once, and for every two balls the robot has to move to A once and then back to B).

Show that polynomially sized merge&shrink heuristics can lead to an exponentially smaller search effort than polynomially sized pattern database heuristics in this domain. To do so, first show that every pattern database heuristic that uses a pattern $P \subset V$ must have an imperfect heuristic value for all states with a certain property. (We exclude the case $P = V$ here because a PDB for all variables no longer has a polynomial size.) Show that there is an exponential number of states with this property that can be reached with a plan cheaper

than $3n$. An A* search with such a heuristic will expand all of these states before expanding the goal state.

Next, show that there is a merge strategy and a shrink strategy that will result in a perfect heuristic while maintaining polynomial sized transition systems. An A* search with the perfect heuristic (and sensible tie-breaking) will expand only $3n$ states.

(b) Consider a set $X = \{\mathcal{T}_1, \mathcal{T}_2\}$ of abstract transition systems with identical label set $L = \{l_1, \dots, l_7\}$ and cost function c such that $c(l_1) = c(l_4) = c(l_6) = 1$ and $c(l_2) = c(l_3) = c(l_5) = c(l_7) = 2$. \mathcal{T}_1 and \mathcal{T}_2 are depicted graphically below. As usual, an incoming arrow indicates the initial state, and goal states are marked by a double rectangle.



- Determine a mapping $\lambda : L \mapsto L'$ that maps all \mathcal{T}_1 -combinable labels with identical cost to the same (new) label and all labels l that are not \mathcal{T}_1 -combinable with another label to l . Let c' be the cost function that allows exact label reduction with $\langle \lambda, c' \rangle$. Graphically provide $\mathcal{T}_1^{\langle \lambda, c' \rangle}$ and $\mathcal{T}_2^{\langle \lambda, c' \rangle}$.
- Graphically provide the transition systems \mathcal{T}_1' and \mathcal{T}_1'' that result from shrinking $\mathcal{T}_1^{\langle \lambda, c' \rangle}$ with the following shrinking strategies:
 - \mathcal{T}_1' results from applying f -preserving shrinking, and
 - \mathcal{T}_1'' results from applying bisimulation-based shrinking.
- Graphically provide the transition systems $\mathcal{T}_1' \otimes \mathcal{T}_2^{\langle \lambda, c' \rangle}$, $\mathcal{T}_1'' \otimes \mathcal{T}_2^{\langle \lambda, c' \rangle}$, and $\mathcal{T}_1 \otimes \mathcal{T}_2$. How do they compare with respect to size and heuristic value of the initial state?

(c) Let X and X' be collections of transition systems. Why is $h(s) = h_{\mathcal{T}_{X'}}^*(\sigma(s))$ not necessarily an admissible heuristic for \mathcal{T}_X if the transformation from X to X' is not safe? Discuss the question for each of the following reasons why a transformation with functions σ and λ can be unsafe:

- $c'(\lambda(l)) > c(l)$ for at least one $l \in L$
- there is a transition $\langle s, l, t \rangle$ of \mathcal{T}_X such that $\langle \sigma(s), \lambda(l), \sigma(t) \rangle$ is not a transition of $\mathcal{T}_{X'}$, or
- there is a goal state s of \mathcal{T}_X such that $\sigma(s)$ is not a goal state of $\mathcal{T}_{X'}$.

The exercise sheets can be submitted in groups of two students. Please provide both student names on the submission.