

Multi-Agent Systems

B. Nebel, F. Lindner, T. Engesser
 Winter Semester 2018/19

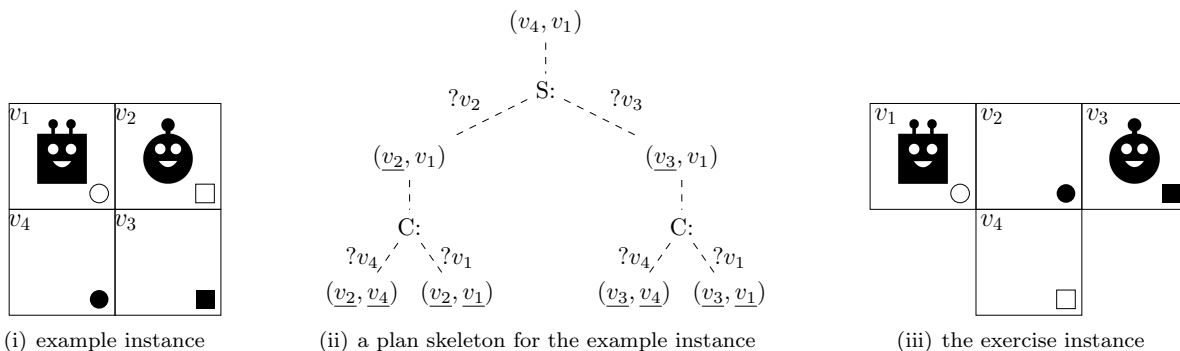
University of Freiburg
 Department of Computer Science

Exercise Sheet 8

Due: December 19, 2018, 16:00

Exercise 8.1 (Multi-Agent Pathfinding with Destination Uncertainty, 2+2+2+1)

In this exercise, you are supposed to systematically construct an i -strong plan for a multi-agent pathfinding instance with destination uncertainty. To simplify things, we will make use of *plan skeletons*, which we define as transition systems alternating between stepping stone configurations and branching points. Intermediate moves and success announcements are left out, but we mark in each configuration the agents that already have announced success and cannot move. Each leaf of the tree has to be a goal configuration. Consider the example instance (i) below. The plan skeleton shown in (ii) describes the following class of plans in which at first, both agents create the stepping stone (v_4, v_1) , where S is on v_4 and C is on v_1 . In the case where the actual destination of S is v_2 , after the agent S goes there and announces success, the stepping stone (v_2, v_1) is created from which agent C can move to v_1 or v_4 depending on his actual destination. The branch where the actual destination of S is v_3 works analogously.



- Identify all initially reachable stepping stones in the exercise instance (iii).
- Choose one of the stepping stones and find a plan skeleton starting with that stepping stone.
- Use the plan skeleton to construct a full plan containing all movement and announcement actions.
- Determine the execution cost of your plan.

Exercise 8.2 (Subjective and Objective Strength, 2)

What is the difference between an i -strong and an objectively strong plan? Explain in your own words using an example that is as simple as possible (the example from the lecture doesn't count). *Hint:* There are some very trivial examples with only two agents.

Exercise 8.3 (CA*, 1+2)

Consider the following MAPF instance on a graph $G = (V, E)$ where $G = \{v_1, v_2, v_3, v_4, v_5, v_6, v_7\}$ and $E = \{\{v_1, v_2\}, \{v_2, v_3\}, \{v_2, v_4\}, \{v_4, v_6\}, \{v_5, v_6\}, \{v_6, v_7\}\}$, with agents $A = \{S, C, T\}$, initial state $\alpha_0 = \{S \mapsto v_1, C \mapsto v_5, T \mapsto v_2\}$ and goal state $\alpha_* = \{S \mapsto v_7, C \mapsto v_3, T \mapsto v_6\}$.

- Decide (without proof) for all agent orderings whether or not the CA* algorithm finds a plan.
- Generate a plan and the corresponding reservation table for one ordering where CA* is successful.

Hint: Since the planning for each agent is done in *space-time*, it is possible for agents to actively go out of the way of other agents who precede them in the agent ordering (and whose plans can thus already be considered fixed).