Online Saturated Cost Partitioning for Classical Planning

Jendrik Seipp

Linköping University and University of Basel

Setting and Motivation

- optimal classical planning
- A^* search + multiple abstraction heuristics
- (saturated) cost partitioning
- different states need different cost partitionings
- precompute cost partitionings

→ no good stopping criterion, search starts late

- compute cost partitioning for each state

→ too expensive

Cost partitioning

- split action costs among heuristics
- ensure that sum of costs $\leq$ original cost

Saturated cost partitioning

- order heuristics, then for each heuristic h:
  - use minimum costs preserving all estimates of h
  - use remaining costs for subsequent heuristics

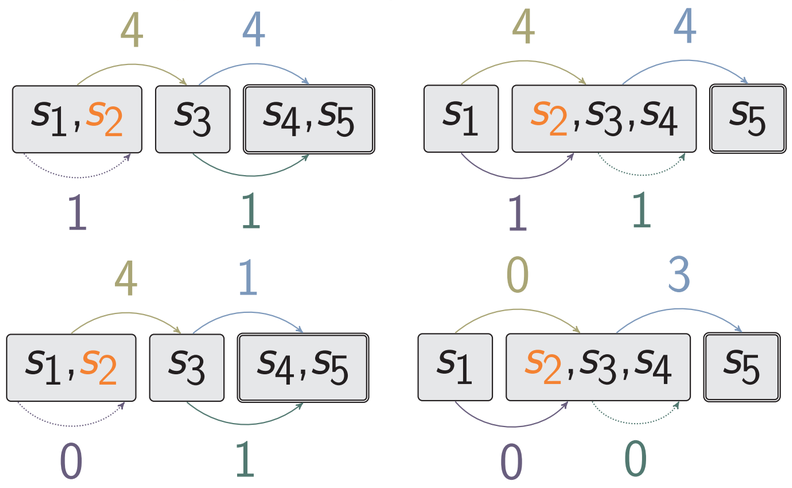

Offline diversification

- sample 1000 states
- start with empty set of orders
- until time limit is reached:
  - compute order for new sample
  - store order if a sample profits from it

Online diversification: ComputeHeuristic(s)

- if SELECT(s) and not time limit reached
  - compute order for s
  - store order if s profits from it
- return maximum over all stored orders for s

Selection strategies

- Bellman (Eifler and Fickert 2018)
- Novelty (Lipovetzky and Geffner 2012)
- Interval

Offline vs. online diversification

Offline

- compute orders for samples for T seconds
- store order if one of 1000 samples profits from it

Online

- compute orders for subset of evaluated states for at most T seconds
- store order if single evaluated state profits from it

Results