Planning and Optimization

E6. Optimal Cost-Partitioning

Malte Helmert and Thomas Keller

Universität Basel

November 18, 2019

Content of this Course

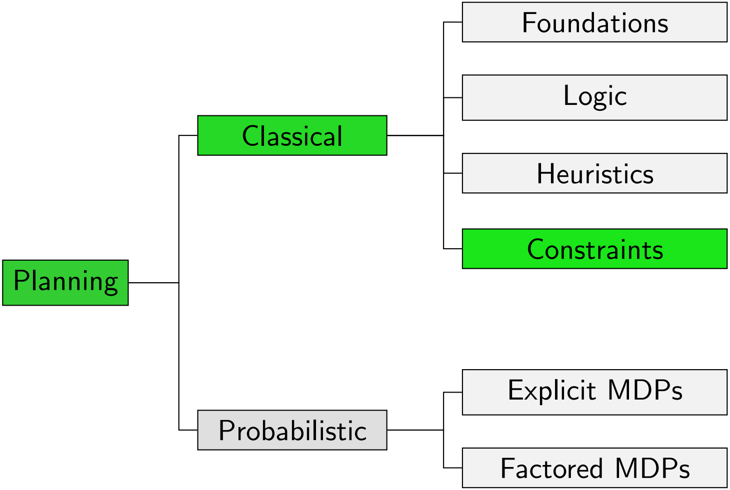

Content of this Course: Constraints

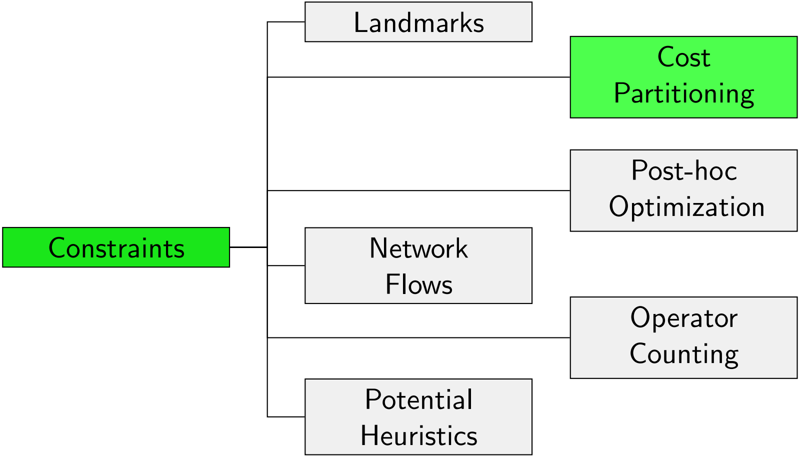

Optimal Cost Partitioning

Optimal Cost Partitioning with LPs

- Use variables for cost of each operator in each task copy
- Express heuristic values with linear constraints
- Maximize sum of heuristic values subject to these constraints

LPs known for

- abstraction heuristics
- disjunctive action landmarks

Abstractions

LP for Shortest Path in State Space

Variables

Non-negative variable Distance_s for each state s

Objective

Maximize Distance_{s_I}

Subject to

$$\text{Distance}_{s_\star} = 0 \quad \text{for all goal states } s_\star$$

$$\text{Distance}_s \leq \text{Distance}_{s'} + cost(o) \quad \text{for all transitions } s \xrightarrow{o} s'$$

Optimal Cost Partitioning for Abstractions I

Variables

For each abstraction α:

Non-negative variable Distance_s^α for each abstract state s,

Non-negative variable Cost_o^α for each operator o

Objective

Maximize $\sum_\alpha \text{Distance}_{\alpha(s_I)}^\alpha$

...

Optimal Cost Partitioning for Abstractions II

Subject to

$\sum_\alpha \text{Cost}_o^\alpha \leq cost(o)$ for all operators o

and for all abstractions α

$\text{Distance}_{s_\star}^\alpha = 0$ for all abstract goal states $s_\star$

$\text{Distance}_s^\alpha \leq \text{Distance}_{s'}^\alpha + \text{Cost}_o^\alpha$ for all transition $s \xrightarrow{o} s'$

Example (1)

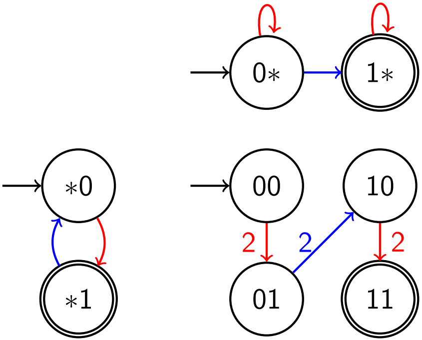

Example (2)

Maximize $\text{Distance}^1_0 + \text{Distance}^2_0$ subject to

$$
\begin{aligned}
\text{Cost}^1_{\text{red}} + \text{Cost}^2_{\text{red}} &\leq 2\\
\text{Cost}^1_{\text{blue}} + \text{Cost}^2_{\text{blue}} &\leq 2\\
\text{Distance}^1_1 &= 0\\
\text{Distance}^1_0 &\leq \text{Distance}^1_0 + \text{Cost}^1_{\text{red}}\\
\text{Distance}^1_0 &\leq \text{Distance}^1_1 + \text{Cost}^1_{\text{blue}}\\
\text{Distance}^1_1 &\leq \text{Distance}^1_1 + \text{Cost}^1_{\text{red}}\\
\text{Distance}^2_1 &= 0\\
\text{Distance}^2_0 &\leq \text{Distance}^2_1 + \text{Cost}^2_{\text{red}}\\
\text{Distance}^2_1 &\leq \text{Distance}^2_0 + \text{Cost}^2_{\text{blue}}\\
\text{Distance}^\alpha_s &\geq 0 \quad \text{for } \alpha \in \{1,2\}, s \in \{0,1\}\\
\text{Cost}^\alpha_o &\geq 0 \quad \text{for } \alpha \in \{1,2\}, o \in \{\text{red}, \text{blue}\}
\end{aligned}
$$

Caution

A word of warning

- optimization for every state gives best-possible cost partitioning
- but takes time

Better heuristic guidance often does not outweigh the overhead.

Landmarks

Optimal Cost Partitioning for Landmarks

- Use again LP that covers heuristic computation and cost partitioning.
- LP variable Cost_L for cost of landmark L in induced task
- Explicit variables for cost partitioning not necessary. Use implicitly $cost_L(o) = \mathrm{Cost}_L$ for all $o \in L$ and 0 otherwise.

Optimal Cost Partitioning for Landmarks: LP

Variables

Non-negative variable Cost_L for each disj. action landmark $L \in \mathcal{L}$

Objective

Maximize $\sum_{L \in \mathcal{L}} \mathrm{Cost}_L$

Subject to

$$\sum_{L \in \mathcal{L} : o \in L} \mathrm{Cost}_L \leq cost(o) \quad \text{for all operators } o$$

Example (1)

Example

Let Π be a planning task with operators $o_1, \ldots, o_4$ and $cost(o_1) = 3, cost(o_2) = 4, cost(o_3) = 5$ and $cost(o_4) = 0$.
Let the following be disjunctive action landmarks for Π:

$$\mathcal{L}_1 = \{o_4\}$$
$$\mathcal{L}_2 = \{o_1, o_2\}$$
$$\mathcal{L}_3 = \{o_1, o_3\}$$
$$\mathcal{L}_4 = \{o_2, o_3\}$$

Example (2)

Example

Maximize $\mathrm{Cost}_{\mathcal{L}_1} + \mathrm{Cost}_{\mathcal{L}_2} + \mathrm{Cost}_{\mathcal{L}_3} + \mathrm{Cost}_{\mathcal{L}_4}$ subject to

$$
\begin{array}{lrl}
[o_1] & \mathrm{Cost}_{\mathcal{L}_2} + \mathrm{Cost}_{\mathcal{L}_3} & \leq 3 \\
[o_2] & \mathrm{Cost}_{\mathcal{L}_2} + \mathrm{Cost}_{\mathcal{L}_4} & \leq 4 \\
[o_3] & \mathrm{Cost}_{\mathcal{L}_3} + \mathrm{Cost}_{\mathcal{L}_4} & \leq 5 \\
[o_4] & \mathrm{Cost}_{\mathcal{L}_1} & \leq 0 \\
 & \mathrm{Cost}_{\mathcal{L}_i} & \geq 0 \quad \text{for } i \in \{1, 2, 3, 4\}
\end{array}
$$

Optimal Cost Partitioning for Landmarks (Dual view)

Variables

Non-negative variable Applied_o for each operator o

Objective

Minimize $\sum_o \text{Applied}_o \cdot cost(o)$

Subject to

$$\sum_{o \in L} \text{Applied}_o \geq 1 \text{ for all landmarks } L$$

Minimize "plan cost" with all landmarks satisfied.

Example: Dual View

Example (Optimal Cost Partitioning: Dual View)

Minimize $3\text{Applied}_{o_1} + 4\text{Applied}_{o_2} + 5\text{Applied}_{o_3}$ subject to

$$\begin{aligned}
\text{Applied}_{o_4} &\geq 1 \\
\text{Applied}_{o_1} + \text{Applied}_{o_2} &\geq 1 \\
\text{Applied}_{o_1} + \text{Applied}_{o_3} &\geq 1 \\
\text{Applied}_{o_2} + \text{Applied}_{o_3} &\geq 1 \\
\text{Applied}_{o_i} &\geq 0 \quad \text{for } i \in \{1, 2, 3, 4\}
\end{aligned}$$

Example: Dual View

Example (Optimal Cost Partitioning: Dual View)

Minimize $\quad 3\text{Applied}_{o_1} + 4\text{Applied}_{o_2} + 5\text{Applied}_{o_3} \quad$ subject to

$$
\begin{aligned}
\text{Applied}_{o_4} &\geq 1 \\
\text{Applied}_{o_1} + \text{Applied}_{o_2} &\geq 1 \\
\text{Applied}_{o_1} + \text{Applied}_{o_3} &\geq 1 \\
\text{Applied}_{o_2} + \text{Applied}_{o_3} &\geq 1 \\
\text{Applied}_{o_i} &\geq 0 \quad \text{for } i \in \{1, 2, 3, 4\}
\end{aligned}
$$

This is equal to the LP relaxation of MHS heuristic

Reminder: LP Relaxation of MHS heuristic

Example (Minimum Hitting Set)

minimize $3X_{o_1} + 4X_{o_2} + 5X_{o_3}$ subject to

$$
\begin{aligned}
X_{o_4} &\geq 1 \\
X_{o_1} + X_{o_2} &\geq 1 \\
X_{o_1} + X_{o_3} &\geq 1 \\
X_{o_2} + X_{o_3} &\geq 1
\end{aligned}
$$

$$X_{o_1} \geq 0, \quad X_{o_2} \geq 0, \quad X_{o_3} \geq 0, \quad X_{o_4} \geq 0$$

- ⇝ optimal solution of LP relaxation: $X_{o_4} = 1$ and $X_{o_1} = X_{o_2} = X_{o_3} = 0.5$ with objective value 6
- ⇝ LP relaxation of MHS heuristic is admissible and can be computed polynomial time

General Cost Partitioning

General Cost Partitioning

Cost functions usually non-negative

- We tacitly also required this for task copies
- Makes intuitively sense: original costs are non-negative
- But: not necessary for cost-partitioning!

General Cost Partitioning

Definition (General Cost Partitioning)

Let Π be a planning task with operators O.

A general cost partitioning for Π is a tuple $\langle cost_1, \dots, cost_n \rangle$, where

- $cost_i : O \to \mathbb{R}$ for $1 \leq i \leq n$ and
- $\sum_{i=1}^{n} cost_i(o) \leq cost(o)$ for all $o \in O$.

General Cost Partitioning: Admissibility

Theorem (Sum of Solution Costs is Admissible)

Let Π be a planning task, $\langle cost_1, \dots, cost_n \rangle$ be a general cost partitioning and $\langle \Pi_1, \dots, \Pi_n \rangle$ be the tuple of induced tasks.

Then the sum of the solution costs of the induced tasks is an admissible heuristic for Π, i.e., $\sum_{i=1}^{n} h^*_{\Pi_i} \le h^*_{\Pi}$.

(Proof omitted.)

General Cost Partitioning: Example

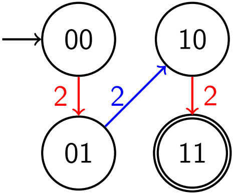

General Cost Partitioning: Example

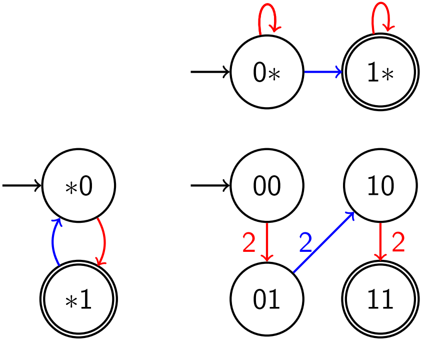

General Cost Partitioning: Example

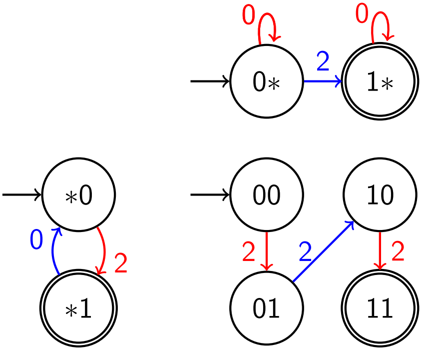

Heuristic value: $2 + 2 = 4$

General Cost Partitioning: Example

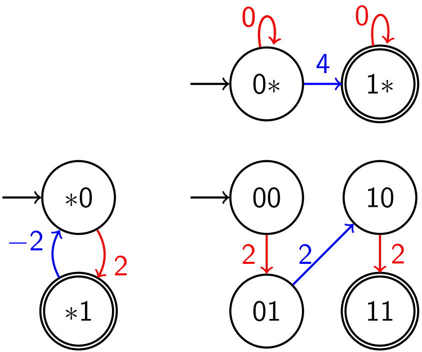

Heuristic value: $4 + 2 = 6$

General Cost Partitioning: Example

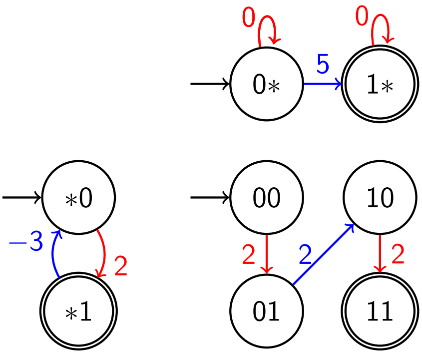

Heuristic value: $-\infty + 5 = -\infty$

LP for Shortest Path in State Space with Negative Costs

Variables

General variable Distance_s for each state s

Objective

Maximize Distance_{s_I}

Subject to

$$\text{Distance}_{s_\star} \leq 0 \quad \text{for all goal states } s_\star$$

$$\text{Distance}_s \leq \text{Distance}_{s'} + cost(o) \quad \text{for all alive transitions } s \xrightarrow{o} s'$$

alive: on any path from initial state to goal state

Modifications also correct (but unnecessary) for non-negative costs

Optimal General Cost Partitioning for Abstractions I

Variables

For each abstraction α:

General variable Distance_s^α for each abstract state s,

General variable Cost_o^α for each operator o

Objective

Maximize $\sum_\alpha \text{Distance}_{\alpha(s_I)}^\alpha$

...

Optimal Cost Partitioning for Abstractions II

Subject to

$\sum_\alpha \text{Cost}_o^\alpha \leq cost(o)$ for all operators o

and for all abstractions α

$\text{Distance}_{s_\star}^\alpha \leq 0$ for all abstract goal states $s_\star$

$\text{Distance}_s^\alpha \leq \text{Distance}_{s'}^\alpha + \text{Cost}_o^\alpha$ for all alive transition $s \xrightarrow{o} s'$

Summary

Summary

- For abstraction heuristics and disjunctive action landmarks, we know how to determine an optimal cost partitioning, using linear programming.
- Although solving a linear program is possible in polynomial time, the better heuristic guidance often does not outweigh the overhead.
- In constrast to standard (non-negative) cost partitioning, general cost partitioning allows negative operators costs.
- General cost partitioning has the same relevant properties as non-negative cost partitioning but is more powerful.