Dantzig-Wolfe Decomposition for Cost Partitioning: Technical Report

Florian Pommerening,[1] Thomas Keller,[1] Valentina Halasi,
Jendrik Seipp,[1,2] Silvan Sievers,[1] Malte Helmert[1]
[1] University of Basel, Switzerland
[2] Linköping University, Sweden
florian.pommerening@unibas.ch, tho.keller@unibas.ch, mail@valentinahalasi.com
jendrik.seipp@liu.se, silvan.sievers@unibas.ch, malte.helmert@unibas.ch

Technical Report CS-2021-001

Abstract

This technical report contains full proofs for the claim on redundant patterns for general cost partitioning of section "Restricting the Considered Patterns" of the main paper.

Definitions

We make some of the terminology used in the main paper formal definitions in the following.

Definition 1. *Let Π be a planning task with variables V, states S, and cost function c. A pattern P is* redundant in Π for general cost partitioning *if there are two patterns $P_1, P_2 \subset P$, such that for every collection of patterns $C \subset 2^V$*

$$h^{\text{gOCP}}_{\{P_1,P_2\}\cup C}(s, c) \geq h^{\text{gOCP}}_{\{P\}\cup C}(s, c) \quad \textit{for all } s \in S.$$

Definition 2. *A pattern is* interesting *for a task for general cost partitioning if it is not redundant. It is* interesting for a set of tasks *if it is interesting in one of them.*

Definition 3. *A pattern P is* causally connected *if the subgraph of the causal graph induced by P is weakly connected.*

It is causally relevant *if the subgraph of the causal graph induced by P contains a directed path via precondition edges from each node to some goal variable node.*

Interesting Patterns for General Cost Partitioning

When talking about general cost partitioning, it is easier to consider its dual view which is operator counting. A flow in a graph is a function mapping edges to non-negative real numbers such that the total incoming and outgoing flow is balanced at each node (with exceptions for initial and goal nodes). The *operator count* of a flow maps each operator to the total flow along edges labeled with this operator. A pattern P is then associated with a set of operator counts for all valid flows in the abstract transition system of the projection to P after removing all dead states. Let us call this set *counts*(P, s). The cost of a count x under cost function c is simply $c(x) = \sum_{o\in O} x(o)c(o)$. When considering a collection of patterns C, the operator-counting heuristic with the flow constraints for all patterns selects the cheapest operator count that has a valid flow in each projection, i.e., it considers the cheapest flow from the set *counts*$(C, s) = \bigcap_{P\in C}$ *counts*(P, s). From the relation of general cost partitioning to operator counting, we know that

$$h^{\text{gOCP}}_{C}(s, c) = \min_{x\in \textit{counts}(C,s)} c(x)$$

We can use this connection to show that $h^{\text{gOCP}}_{C}(s, c) \geq h^{\text{gOCP}}_{C'}(s, c)$ by showing that *counts*$(C, s) \subseteq$ *counts*(C', s).

We show that the two conditions of Definition 1 of the main paper/Definition 3 of this technical report lead to redundant patterns for general cost partitioning in two separate theorems.

Theorem 1. *A pattern P that is not causally connected is redundant for general cost partitioning.*

Proof. Let P_1, P_2 be a partition of P into non-empty subsets such that the subgraph of the causal graph induced by P contains no arc between the two sets. Let $\alpha, \alpha_1, \alpha_2$ be the projections to P, P_1, P_2. It is sufficient to show that *counts*$(P_1, s) \cap$ *counts*$(P_2, s) \subseteq$ *counts*(P, s) for a state s.

Consider a count $x \in$ *counts*$(P_1, s) \cap$ *counts*(P_2, s). There have to be flows f_1, f_2 in α_1 and α_2 with count x. We define the following flow that can be seen as a way of concatenating abstract plans.

$$f(\langle\langle s, s'\rangle, o, \langle t, t'\rangle\rangle) = \begin{cases} f_1(\langle s, o, t\rangle) & \text{if } s' = t' = \alpha_2(I) \\ w_s f_2(\langle s', o, t'\rangle) & \text{if } s = t \text{ and } s' \neq t' \end{cases}$$

where w_s is the outgoing flow in s according to f_1. If there is a single goal state g, then $w_g = 1$ and $w_s = 0$ for $s \neq g$, otherwise the total outgoing flow of 1 can be split among all goal states.

As no operator can affect both P_1 and P_2, we can partition the operators into three sets $O_1, O_2, O_\emptyset$ of operators affecting P_1, operators affecting P_2 and operators affecting neither pattern.

Consider an operator $o \in O_1 \cup O_\emptyset$. It induces only self loops in α_2. For every transition $\langle s, o, t\rangle$ in α_1, there is exactly one transition $\langle\langle s, \alpha_2(I)\rangle, o, \langle t, \alpha_2(I)\rangle\rangle$ in α that has the same flow and all other transitions induced by o have a flow of 0. The count of o thus matches $x(o)$.

Similarly, operators $o \in O_2$ match $x(o)$ because each is counted with a weight of w_s for every state s in α_1. These weights have to sum to 1, so the values sum to the count of o in f_2, which is $x(o)$.

Finally, we have to check that f is a flow in α. When considering only f_1, all states $\langle s, s' \rangle$ with $s' = \alpha_2(I)$ have an outgoing flow of w_s. The flow $w_s f_2$ then moves this flow to a goal state while remaining within states with the first component s.

We have shown that f is a flow in α with count x, so $x \in$ *counts*(P). Since x was an arbitrary count in *counts*$(P_1, s) \cap$ *counts*(P_2, s), we have *counts*$(P_1, s) \cap$ *counts*$(P_2, s) \subseteq$ *counts*(P, s). □

Theorem 2. *A pattern P is redundant for general cost partitioning if it contains a variable v such that the causal graph does not contain a directed path along precondition arcs from v to a goal variable.*

Proof. Let P be a pattern containing a variable v as described above and let C be a collection of patterns. Further let dep(v) be the set of variables than can be reached from v via precondition-effect arcs in the causal graph. Note that this set cannot contain a goal variable. Let O_{irr} be the subset of all operators that have a precondition in dep(v). Operators in O_{irr} cannot set goal variables nor any variables outside of dep(v). Any plan (and in particular any abstract plan) is still a plan once all of these operators are removed because the operators reaching the goal conditions remain and all operators that remain can only have preconditions on variables outside of dep(v), which were not modified by the removed operators.

We first show that an optimal cost partitioning $\mathcal{C} = \langle c_1, \dots, c_n \rangle$ for any pattern collection $\{P_1, \dots, P_n\}$ remains optimal if the cost of all operators in O_{irr} is changed to 0 in all cost functions. To see this, consider a cheapest abstract plan π in the projection to a pattern P_i under the corresponding cost function c_i. The total contribution of operators in O_{irr} to the cost of π cannot be positive. If it were, then the plan π' that is like π but with all operators in O_{irr} removed would be cheaper under the same cost function. Now consider the cost partitioning $\mathcal{C}' = \langle c'_1, \dots, c'_n \rangle$ that results from changing the cost of all operators in O_{irr} to 0 in all cost functions. Under such a cost function the total contribution of operators in O_{irr} is always 0, while the contribution of the remaining operators remains the same. Such a cost function can only increase the cost of plans with operators in O_{irr} and leaves the cost of all other plans the same. So the total heuristic value under $\mathcal{C}'$ cannot be lower than that under $\mathcal{C}$, and thus $\mathcal{C}'$ is also an optimal cost partitioning.

So, when considering the pattern collection $\{P\} \cup C$ we can assume that an optimal cost partitioning has a cost of 0 for all operators in O_{irr} and that there is an optimal plan π without operators from O_{irr} in the projection to P. This plan remains optimal in the projection to $P' = P \setminus \text{dep}(v)$ under the same cost function. There cannot be a cheaper plan in the projection to P' because it could only consist of operators outside of O_{irr} and thus also be applicable in the projection to P. Thus h^P and $h^{P'}$ have the same values under this cost function and pattern P can be replaced by P' without lowering the value of the optimal cost partitioning. (Technically P is replaced by $\{P', \emptyset\}$ because according to our definition a pattern is redundant if it can be replaced by two other patterns.) □