Merge-and-Shrink for Stochastic Shortest-Path Problems with Pruning Transformations — Technical Report

Thorsten Klößner[1], **Álvaro Torralba**[2], **Marcel Steinmetz**[1], **Silvan Sievers**[3]

[1]Saarland University, Germany, [2]Aalborg University, Denmark, [3]LAAS-CNRS, France, [4]Basel University, Switzerland
kloessner@cs.uni-saarland.de, marcel.steinmetz@laas.fr, alto@cs.aau.dk, silvan.sievers@unibas.ch

This techincal report contains the full versions of all sketched or omitted proofs of our HSDIP 2024 paper "Merge-and-Shrink for Stochastic Shortest-Path Problems with Pruning Transformations". We follow the same notation as originally introduced in the paper.

For the proofs below, recall again the definition of the transformed policy $\pi_{\tau,s}$ as given in the paper:

Definition 1. *Let $\langle \Theta, \Theta', \sigma, \lambda \rangle$ be a transformation, let π be a policy for Θ and let $s \in S_\Theta$ be some starting state. The transformed policy $\pi_{\tau,s}$ of π for s is defined by*

$$\pi_{\tau,s}(h')(\mathfrak{t}') := \Pr_s^\pi[\dot{\bigcup}_{\substack{h \in ind_\tau^{-1}(h') \\ \mathfrak{t} \in ind_\tau^{-1}(\mathfrak{t}')}} Cyl(h\mathfrak{t}) \mid \dot{\bigcup}_{h \in ind_\tau^{-1}(h')} Cyl(h)] = \frac{\Pr_s^\pi[\dot{\bigcup}_{\langle h,\mathfrak{t}\rangle \in ind_\tau^{-1}(h') \times ind_\tau^{-1}(\mathfrak{t}')} Cyl(h\mathfrak{t})]}{\Pr_s^\pi[\dot{\bigcup}_{h \in ind_\tau^{-1}(h')} Cyl(h)]}$$

if $\Pr_s^\pi[\dot{\bigcup}_{h \in ind_\tau^{-1}(h')} Cyl(h)] > 0$*, and* $\pi_{\tau,s}(h')(\mathfrak{t}') := 0$ *otherwise.*

Proof of Lemma 1

Recall the statement made by Lemma 1.

Lemma 1. *Let $\tau = \langle \Theta, \Theta', \sigma, \lambda \rangle$ be a transformation, let π be a policy and let $s \in S_\Theta$ be a starting state with $s \in dom(\sigma)$. For every abstract history $h' \in Hist(\Theta')$:*

$$\Pr_{\sigma(s)}^{\pi_{\tau,s}}[Cyl(h')] = \Pr_s^\pi\left[\dot{\bigcup}_{h \in ind_\tau^{-1}(h')} Cyl(h)\right] \tag{1}$$

$$\Pr_{\sigma(s)}^{\pi_{\tau,s}}[h'] = \Pr_s^\pi\left[ind_\tau^{-1}(h')\right] + \Pr_s^\pi\left[\dot{\bigcup}_{\langle h,\mathfrak{t}\rangle \in ind_\tau^{-1}(h') \times (T_\Theta \setminus ind_\tau^{-1}(T_{\Theta'}))} Cyl(h\mathfrak{t})\right] \tag{2}$$

Proof. Let $h' = s'_0\mathfrak{t}'_0 \dots s'_n$ in the following.

Proof of Equation (1) First, we show Equation (1). By definition of the probability measure, we have

$$\Pr_{\sigma(s)}^{\pi_{\tau,s}}[Cyl(s'_0\mathfrak{t}'_0 \dots s'_n)] = [s'_0 = s'] \cdot \prod_{i=0}^{n-1} \pi_{\tau,s}(s'_0\mathfrak{t}'_0 \dots s'_i)(\mathfrak{t}'_i) \cdot \delta_{\mathfrak{t}'_i}(s'_{i+1}).$$

If $\Pr_s^\pi\left[\bigcup_{h \in ind_\tau^{-1}(s'_0\mathfrak{t}'_0 \dots s'_i)} Cyl(h)\right] = 0$ for some prefix $s'_0\mathfrak{t}'_0 \dots s'_i$ of h', then the right hand side of Equation (1) is $\Pr_s^\pi\left[\bigcup_{h \in ind_\tau^{-1}(h')} Cyl(h)\right] = 0$ and the equation above has a factor $\pi_{\tau,s}(s'_0\mathfrak{t}'_0 \dots s'_i) = 0$ and thus also evaluates to zero. Otherwise we apply the definition of $\pi_{\tau,s}$ and exploit countable additivity of the probability measure to obtain:

$$\dots = [s'_0 = s'] \cdot \prod_{i=0}^{n-1} \frac{\sum_{\langle h,\mathfrak{t}\rangle \in ind_\tau^{-1}(s'_0\mathfrak{t}'_0 \dots s'_i) \times ind_\tau^{-1}(\mathfrak{t}'_i)} \Pr_s^\pi[Cyl(h\mathfrak{t})] \cdot \delta_{\mathfrak{t}'_i}(s'_{i+1})}{\Pr_s^\pi[\bigcup_{h \in ind_\tau^{-1}(s'_0\mathfrak{t}'_0 \dots s'_i)} Cyl(h)]}$$

In the next step, we use the fact that $\Pr_s^\pi[Cyl(h\mathfrak{t})] = \Pr_s^\pi[Cyl(h)] \cdot \pi(h)(\mathfrak{t})$, and exploit that $\mathfrak{t} \in ind_\tau^{-1}(\mathfrak{t}'_i)$ in the sum, from which it follows that $\delta_{\mathfrak{t}'_i} = \sigma_{Dist}(\delta_\mathfrak{t})$.

$$\dots = [s'_0 = s'] \cdot \prod_{i=0}^{n-1} \frac{\sum_{\langle h,\mathfrak{t}\rangle \in ind_\tau^{-1}(s'_0\mathfrak{t}'_0 \dots s'_i) \times ind_\tau^{-1}(\mathfrak{t}'_i)} \Pr_s^\pi[Cyl(h)] \cdot \pi(h)(\mathfrak{t}) \cdot \sigma_{Dist}(\delta_\mathfrak{t})(s'_{i+1})}{\Pr_s^\pi[\bigcup_{h \in ind_\tau^{-1}(s'_0\mathfrak{t}'_0 \dots s'_i)} Cyl(h)]}$$

We now apply the definition of σ_{Dist}, leading to:

$$\ldots = [s'_0 = s'] \cdot \prod_{i=0}^{n-1} \frac{\sum_{\langle h,\mathfrak{t}\rangle \in ind_\tau^{-1}(s'_0\mathfrak{t}'_0 \ldots s'_i) \times ind_\tau^{-1}(\mathfrak{t}'_i)} \Pr_s^\pi[Cyl(h)] \cdot \pi(h)(\mathfrak{t}) \cdot \sum_{s \in \sigma^{-1}(s'_{i+1})} \delta_\mathfrak{t}(s)}{\Pr_s^\pi[\bigcup_{h \in ind_\tau^{-1}(s'_0\mathfrak{t}'_0 \ldots s'_i)} Cyl(h)]}$$

$$= [s'_0 = s'] \cdot \prod_{i=0}^{n-1} \frac{\sum_{\langle h,\mathfrak{t},s\rangle \in ind_\tau^{-1}(s'_0\mathfrak{t}'_0 \ldots s'_i) \times ind_\tau^{-1}(\mathfrak{t}'_i) \times \sigma^{-1}(s'_{i+1})} \Pr_s^\pi[Cyl(h)] \cdot \pi(h)(\mathfrak{t}) \cdot \delta_\mathfrak{t}(s)}{\Pr_s^\pi[\bigcup_{h \in ind_\tau^{-1}(s'_0\mathfrak{t}'_0 \ldots s'_i)} Cyl(h)]}$$

Finally, we make use of the fact $\Pr_s^\pi[Cyl(h)] \cdot \pi(h)(\mathfrak{t}) \cdot \delta_\mathfrak{t}(s) = \Pr_s^\pi[Cyl(h\mathfrak{t}s)]$ and simplify:

$$= [s'_0 = s'] \cdot \prod_{i=0}^{n-1} \frac{\sum_{\langle h,\mathfrak{t},s\rangle \in ind_\tau^{-1}(s'_0\mathfrak{t}'_0 \ldots s'_i) \times ind_\tau^{-1}(\mathfrak{t}'_i) \times \sigma^{-1}(s'_{i+1})} \Pr_s^\pi[Cyl(h\mathfrak{t}s)]}{\Pr_s^\pi[\bigcup_{h \in ind_\tau^{-1}(s'_0\mathfrak{t}'_0 \ldots s'_i)} Cyl(h)]}$$

$$= [s'_0 = s'] \cdot \prod_{i=0}^{n-1} \frac{\sum_{h \in ind_\tau^{-1}(s'_0\mathfrak{t}'_0 \ldots s'_i\mathfrak{t}'_i s'_{i+1})} \Pr_s^\pi[Cyl(h)]}{\Pr_s^\pi[\bigcup_{h \in ind_\tau^{-1}(s'_0\mathfrak{t}'_0 \ldots s'_i)} Cyl(h)]}$$

$$= [s'_0 = s'] \cdot \prod_{i=0}^{n-1} \frac{\Pr_s^\pi[\bigcup_{h \in ind_\tau^{-1}(s'_0\mathfrak{t}'_0 \ldots s'_{i+1})} Cyl(h)]}{\Pr_s^\pi[\bigcup_{h \in ind_\tau^{-1}(s'_0\mathfrak{t}'_0 \ldots s'_i)} Cyl(h)]}$$

This is a telescoping product, in which every numerator cancels with the following denominator, and only the final numerator for $i = n-1$ remains. Note in particular that $[s'_0 = s'] = \Pr_s^\pi[\bigcup_{s_0 \in \sigma^{-1}(s'_0)} Cyl(s_0)]$. Therefore:

$$\ldots = \Pr_s^\pi[\bigcup_{h \in ind_\tau^{-1}(h')} Cyl(h)]$$

This concludes the proof of Equation (1).

Proof of Equation (2) Before we show Equation (2), note that by definition of $\pi_{\tau,s}$ and Equation (1):

$$\Pr_{\sigma(s)}^{\pi_{\tau,s}}[Cyl(h'\mathfrak{t}')] = \Pr_{\sigma(s)}^{\pi_{\tau,s}}[Cyl(h')] \cdot \pi_{\tau,s}(h')(\mathfrak{t}') = \Pr_s^\pi\big[\bigcup_{\langle h,\mathfrak{t}\rangle \in ind_\tau^{-1}(h') \times ind_\tau^{-1}(\mathfrak{t}')} Cyl(h\mathfrak{t})\big] \tag{3}$$

To prove Equation (2), we first express the event $\{h'\}$ as the event of all executions with prefix h', without those which continue with some transition:

$$\Pr_{\sigma(s)}^{\pi_{\tau,s}}[h'] = \Pr_{\sigma(s)}^{\pi_{\tau,s}}\big[Cyl(h') \setminus \big(\bigcup_{\mathfrak{t}' \in T_{\Theta'}} Cyl(h'\mathfrak{t}')\big)\big]$$

$$= \Pr_{\sigma(s)}^{\pi_{\tau,s}}[Cyl(h')] - \Pr_{\sigma(s)}^{\pi_{\tau,s}}\big[\bigcup_{\mathfrak{t}' \in T_{\Theta'}} Cyl(h'\mathfrak{t}')\big]$$

Next, we apply Equation (1) and Equation (3), and simplify.

$$= \Pr_s^\pi[\bigcup_{h \in ind_\tau^{-1}(h')} Cyl(h)] - \Pr_s^\pi\big[\bigcup_{\mathfrak{t}' \in T_{\Theta'}} \bigcup_{\langle h,\mathfrak{t}\rangle \in ind_\tau^{-1}(h') \times ind_\tau^{-1}(\mathfrak{t}')} Cyl(h\mathfrak{t})\big]$$

$$= \Pr_s^\pi[\bigcup_{h \in ind_\tau^{-1}(h')} Cyl(h)] - \Pr_s^\pi\big[\bigcup_{\mathfrak{t} \in ind_\tau^{-1}(T_{\Theta'})} \bigcup_{h \in ind_\tau^{-1}(h')} Cyl(h\mathfrak{t})\big]$$

$$= \Pr_s^\pi\big[\bigcup_{h \in ind_\tau^{-1}(h')} Cyl(h) \setminus \big(\bigcup_{\mathfrak{t} \in ind_\tau^{-1}(T_{\Theta'})} \bigcup_{h \in ind_\tau^{-1}(h')} Cyl(h\mathfrak{t})\big)\big]$$

The event above considers all executions with a prefix $h \in ind_\tau^{-1}(h')$, excluding those continuing with an induced transition $\mathfrak{t} \in ind_\tau^{-1}(T_{\Theta'})$. In other words, we consider exactly terminating executions $h \in ind_\tau^{-1}(h')$, and the executions that start with a prefix $h \in ind_\tau^{-1}(h')$ and then continue with a transition $\mathfrak{t} \notin ind_\tau^{-1}(T_{\Theta'})$. Rewriting the event accordingly:

$$= \Pr_s^\pi\big[ind_\tau^{-1}(h') \mathbin{\dot\cup} \big(\bigcup_{\substack{h \in ind_\tau^{-1}(h') \\ \mathfrak{t} \notin ind_\tau^{-1}(T_{\Theta'})}} Cyl(h\mathfrak{t})\big)\big]$$

$$= \Pr_s^\pi[ind_\tau^{-1}(h')] + \Pr_s^\pi\big[\big(\bigcup_{\substack{h \in ind_\tau^{-1}(h') \\ \mathfrak{t} \notin ind_\tau^{-1}(T_{\Theta'})}} Cyl(h\mathfrak{t})\big)\big]$$

This shows the claim. □

Proof of Proposition 1

Proposition 1 made the following statement.

Proposition 1. *Let $\tau = \langle \Theta, \Theta', \sigma, \lambda \rangle \in \mathbf{CONS_{L+T}}$ and let π be a policy. For all states $s \in S_\Theta$ with $\mathit{Reach}^{\rightarrow}_{\Theta,\pi}(s) \subseteq \mathit{dom}(\sigma)$ and every set of target states $T \subseteq S_\Theta$:*

(i) $\mathit{Reach}^{\rightarrow}_{\Theta,\pi_{\tau,s}}(\sigma(s)) = \sigma(\mathit{Reach}^{\rightarrow}_{\Theta,\pi}(s))$ and

(ii) If $\pi \in \mathit{Sols}_\Theta(s,T)$, then $\pi_{\tau,s} \in \mathit{Sols}_{\Theta'}(\sigma(s), \sigma(T))$.

Proof. To show claim (i), recall Equation (1):

$$\Pr^{\pi_{\tau,s}}_{\sigma(s)}[\mathit{Cyl}(h')] = \Pr^{\pi}_{s}\left[\dot\bigcup_{h \in \mathit{ind}_\tau^{-1}(h')} \mathit{Cyl}(h)\right]$$

The inclusion $\mathit{Reach}^{\rightarrow}_{\Theta,\pi_{\tau,s}}(\sigma(s)) \subseteq \sigma(\mathit{Reach}^{\rightarrow}_{\Theta,\pi}(s))$ holds without any assumptions, since by this equation, $\Pr^{\pi_{\tau,s}}_{\sigma(s)}[\mathit{Cyl}(h')] > 0$ for some history $h' \in \mathit{Hist}(\Theta')$ implies that there is a concrete history $h \in \mathit{ind}_\tau^{-1}(h')$ with $\Pr^{\pi}_{s}[\mathit{Cyl}(h)] > 0$, and we have $\mathit{last}(h') = \sigma(\mathit{last}(h))$ in particular.

For the other direction, acknowledge that under the assumptions $\mathit{Reach}^{\rightarrow}_{\Theta,\pi}(s) \subseteq \mathit{dom}(\sigma)$ and $\tau \in \mathbf{CONS_L}$, every possible history $h \in \mathit{Hist}(\Theta)$ with $\Pr^{\pi}_{s}[\mathit{Cyl}(h)] > 0$ induces an abstract history, i.e., $h \in \mathit{dom}(\mathit{ind}_\tau)$. With $\tau \in \mathbf{CONS_T}$ we even have $\mathit{ind}_\tau(h) \in \mathit{Hist}(\Theta')$. Therefore, $0 < \Pr^{\pi}_{s}[\mathit{Cyl}(h)] \leq \Pr^{\pi}_{s}\left[\dot\bigcup_{h \in \mathit{ind}_\tau^{-1}(\mathit{ind}_\tau(h))} \mathit{Cyl}(h)\right] = \Pr^{\pi_{\tau,s}}_{\sigma(s)}[\mathit{Cyl}(\mathit{ind}_\tau(h))]$ and in particular, $\sigma(\mathit{last}(h)) = \mathit{last}(\mathit{ind}_\tau(h))$.

For the claim (ii), consider Equation (2) for a history $h' \in \mathit{Finish}_{\Theta'}(\sigma(T))$. In the event of the right summand, since $\mathfrak{t} \notin \mathit{ind}_\tau^{-1}(T_{\Theta'})$, we also have $h\mathfrak{t}u \notin \mathit{ind}_\tau^{-1}(\mathit{Hist}(\Theta'))$ and in particular $\Pr^{\pi}_{s}[\mathit{Cyl}(h\mathfrak{t}u)] = 0$ by contraposition of the argument above. Hence, the right summand vanishes and we obtain $\Pr^{\pi_{\tau,s}}_{\sigma(s)}[h'] = \Pr^{\pi}_{s}[\mathit{ind}_\tau^{-1}(h')]$. Concludingly:

$$\begin{aligned}\Pr^{\pi_{\tau,s}}_{\sigma(s)}[\mathit{Finish}_{\Theta'}(\sigma(T))] &= \Pr^{\pi}_{s}[\mathit{ind}_\tau^{-1}(\mathit{Finish}_{\Theta'}(\sigma(T)))] \\ &\geq \Pr^{\pi}_{s}[\mathit{Finish}_{\Theta}(T)] = 1.\end{aligned}$$

Note here that $\mathit{Finish}_\Theta(T) \subseteq \mathit{ind}_\tau^{-1}(\mathit{Finish}_{\Theta'}(\sigma(T)))$ if we only consider possible histories $h \in \mathit{Finish}_\Theta(T)$, as then $h \in \mathit{dom}(\mathit{ind}_\tau)$ and $\mathit{ind}_\tau(h) \in \mathit{Finish}_{\Theta'}(\sigma(T))$. $\square$