Supplementary Material

MvPOP Validity

Lemma 2. *If $l \in \{S_0 \to g\}_{\mathcal{A}}$, then:*

$$\mathrm{S}_0 + \mathfrak{P}_l \Sigma \geq \Pi$$

where $\geq$ applies to each element of the matrices.

Proof. Let $\mathrm{S} := \mathrm{S}_0 + \mathfrak{P}_l \Sigma$, i.e.:

$$\mathrm{S}[b, x] = S_0(x) + \sum_{a \in \mathbf{A}} \mathfrak{P}_l[b, a]\sigma(a, x)$$

Then, $\mathrm{S}[b, x]$ represents the state of variable $x \in V$ before one of the occurrences of action b in l, because $\mathfrak{P}_l[b, a] = 1$ iff a is before b in l. This holds for all occurrences of b in l because there is a row in S for each such occurrence. Since $\mathrm{S}_0 + \mathfrak{P}_l \Sigma \geq \Pi$, then $\mathrm{S}[b, x] \geq \pi(b, x)$ for all b occurring in l and all variables $x \in V$. □

Theorem 2. *For an NAP $\{S_0 \to g\}_{\mathcal{A}}$, a multi-set (A, μ), and an MvPOP $\mathfrak{P}$ over (A, μ) with max. g:*

$$\begin{gathered} \textit{If for all } x \in V : \mathrm{S}_0 + (\mathfrak{P} + \mathfrak{V}_x^{\mathfrak{P}})\Sigma \geq \Pi \\ \textit{then, } \operatorname{lin}(\mathfrak{P}) \subseteq \{S_0 \to g\}_{\mathcal{A}} \end{gathered} \tag{1}$$

Proof. Let $\mathrm{S}^+ := \mathrm{S}_0 + \mathfrak{P}\Sigma$. We have proven in Lemma 2 that in any linearization l where all actions incomparable to b occur after b, we get $\mathrm{S}^+[b, x] \geq \pi(b, x)$. However, we need to ensure that the same holds for linearizations where part of the incomparable actions to b occur before b.
Let $\mathrm{S} := \mathrm{S}_0 + (\mathfrak{P} + \mathfrak{V}_x^{\mathfrak{P}})\Sigma = \mathrm{S}^+ + \mathfrak{V}_x^{\mathfrak{P}}\Sigma$. Then, for all $b \in \mathbf{A}, x \in V$:

$$\mathrm{S}[b, x] = \mathrm{S}^+[b, x] + \sum_{a \in \mathbf{A}} \mathfrak{V}_x^{\mathfrak{P}}[b, a]\sigma(a, x) \tag{2}$$

There are two cases for $\mathfrak{I}_x^{\mathfrak{P}}[b, a]$:

- If $a \not\sim_x b$ and $\pi(b, x) \neq -\infty$, then, $\sigma(a, x) > 0$, i.e., none of the incomparable occurrences of a decrements the value of x. Therefore, if $\mathrm{S}^+[b, x] \geq \pi(b, x)$, it also holds after adding any number of a occurrences. This is the case where $\mathfrak{V}_x^{\mathfrak{P}}[b, a] = 0$.
- If $a \sim_x b$, we need to consider the worst-case linearization, i.e., where all incomparable a's occur before b. The value $\mathfrak{V}_x^{\mathfrak{P}}[b, a]$ in this case is the maximal number of incomparable a's occurrences to b.

From (1), (2), we know that

$$\mathrm{S}^+[b, x] + \sum_{a \in \mathbf{A}} \mathfrak{V}_x^{\mathfrak{P}}[b, a]\sigma(a, x) \geq \pi(b, x)$$

Therefore, even the worst-case linearizations w.r.t. variable $x \in V$, i.e., where the effects of all violating, incomparable actions $\sum_{a \in \mathbf{A}} \mathfrak{V}_x^{\mathfrak{P}}[b, a]\sigma(a, x)$ are counted before b, the precondition of b w.r.t. x still holds. Since (1) is quantified over all variables $x \in V$, thus, any linearization $l \in \operatorname{lin}(\mathfrak{P})$ is in $\{S_0 \to g\}_{\mathcal{A}}$. □

Further Preliminaries

For a relaxed solution $\mathfrak{P}$ over a multi-set $\mathbf{A} = (A, \mu)$, let $S_{<a}^{\mathfrak{P}}$ be *the state before applying* $a \in \mathbf{A}$, i.e., for all $x \in V$:

$$S_{<a}^{\mathfrak{P}}(x) := S_0(x) + \sum_{b \in \mathbf{A}} \mathfrak{P}[a, b]\sigma(b, x)$$

Then, the number of *valid repetitions* of a in $\mathfrak{P}$ is:

$$u^{\mathfrak{P}}(a) := \max\{i \in \mathbb{N} : i \leq \mathfrak{P}[g, a], S_{<a}^{\mathfrak{P}} + i \cdot a \vDash a\}$$

where $S + 0 \cdot a := S$, and $S + i \cdot a := (S + (i-1) \cdot a) + a$ for all $i \in \mathbb{N}_{\geq 1}$ and states $S : V \to \mathbb{Z}$. Now, the *maximal* (w.r.t. number of a repetitions) *validly reachable state* in $\mathfrak{P}$ is denoted $S_{>a}^{\mathfrak{P}}$, where for all $x \in V$:

$$S_{>a}^{\mathfrak{P}}(x) := S_{<a}^{\mathfrak{P}}(x) + u^{\mathfrak{P}}(a)\sigma(a, x)$$

And the number of *violated repetitions* of $a \in \mathbf{A}$ in a relaxed solution $\mathfrak{P}$ is $v^{\mathfrak{P}}(a) := \mathfrak{P}[g, a] - u^{\mathfrak{P}}(a)$.

For $a \in A$, let $U_a := \{x \in V : \sigma(a, x) > 0\}$. Additionally, for a multi-set $\mathbf{A} = (A, \mu)$, let $\mathbf{A} \cup \{a\}$ be the multi-set with one additional sequential repetition of a, i.e., $\mathbf{A} \cup \{a\} = (A, \mu')$ with $\mu'(b) = \mu(b)$ for all $b \in A \setminus \{a\}$, and $\mu'(a) = \mu(a) + 1$. Finally, with $\mathcal{A} \setminus \{a\}$, we denote the NAD $\mathcal{A}$ with the action a removed.

For $a \in \mathbf{A}$, we call an MvPOP $\mathfrak{P}$ over $\mathbf{A}$ an *a-relaxed* solution if (a, a) is the only pair that causes a threat in $\mathfrak{P}$.

Maintainable NAP

Theorem 3. *For a* NAP $\{S_0 \to g\}_{\mathcal{A}}$*, a maintainable action $a \in A$, and a function $\mu^* : T_{\mathfrak{P}}(a) \to \mathbb{N}$ that defines an upper bound for the number of sequential b repetitions needed for all $b \in T_{\mathfrak{P}}(a)$. We can extend the definition set of μ^* to a by:*

$$\mu^*(a) := 1 + \sum_{b \in T_{\mathfrak{P}}(a)} \mu^*(b)$$

Proof. For any two actions $a, b \in A$, if $a \not\sim b$, the number of sequential repetitions needed of a does not depend on b. If $a \sim b$ and $\mathfrak{P}[b, a] = 0$, then, either $\mathfrak{I}^{\mathfrak{P}}[b, a] = \mathfrak{P}[g, a]$ or $\mathfrak{I}^{\mathfrak{P}}[b, a] = 0$. In the first case, all occurrences of a are incomparable to b in $\mathfrak{P}$. Therefore, all a occurrences can be promoted over all b occurrences to solve a threat (a, b). Finally, for $\mathfrak{I}^{\mathfrak{P}}[b, a] = 0$, no ocurrences of a are incomparable to b in $\mathfrak{P}$. Therefore, since a is maintainable, i.e., whenever a is applicable once, it can be repeated arbitrarily often with one sequential repetition, we conclude that sequential repetitions of a are required only to differentiate between a occurrences after and before each sequential repetition of an action b only if $a \precsim_{\mathfrak{P}} b$. □

Beneficially Maintainable NAP

Lemma 4. *There exists a polynomial time reduction of* PP *to* bmNAP.

Proof. Remember that we can translate any PP domain to a NAD in polynomial time. Additionally, notice that, in propositional planning, any repetition is beneficial. We can intuitively validate that from our understanding of PP because an activation effect p of a cannot be reversed by an effect $\neg p$ even if it occurs many times before that a. E.g., if $\{p\} = \mathit{eff}(a) = \mathit{pre}(g)$, and $\{\neg p\} = \mathit{eff}(b)$, then, after applying the propositional plan $[b, b, b, a, g]$, the effect p of a still satisfies the precondition for g even if b deactivates it arbitrarily many times before. In NAP, this can be done by using the correcting actions $c_p, c_{\neg p}$, e.g., $[b, b, b, a, g]$ in PP translates to $[b, b, c_{\neg p}, b, c_{\neg p}, a, g]$ in NAP. □

In the next theorem, we prove that knowing if a mNAP problem is unsolvable is decidable by proving that the NAP as Search algorithm must terminate after a finite number of steps.

Lemma 5. *For an action $a \in A$ with a beneficial repetition policy and a multi-set $\mathbf{A}_0 = (A, \mu_0)$. Let $\mathfrak{P}_0$ be an a-relaxed solution over $\mathbf{A}_0$, and let $\mathbf{A}_i \supseteq \mathbf{A}_{i-1} \cup \{a\}$ for all $i \in \mathbb{N}$. If $\{S_0 \to g\}_\mathcal{A} = \emptyset$, then, there exists $k \in \mathbb{N}$ s.t. no a-relaxed solutions $\mathfrak{P}_k$ over $\mathbf{A}_k$ exist.*

Proof. Assume that for all $k \in \mathbb{N}$, there exists an a-relaxed solution $\mathfrak{P}_k$ over $\mathbf{A}_k$. Notice that if $\{S_0 \to g\}_\mathcal{A} = \emptyset$, then, for all $k \in \mathbb{N}$, $\mathfrak{P}_k$ is not valid; however, since it is a-relaxed, there exists $b_k \in \mathbf{A}_0$ and $x_k \in U_a$ s.t.

$$S_{>a}^{\mathfrak{P}_k}(x_k) < \pi(b_k, x_k) \leq \max_{b \in A, x \in U_a} \pi(b, x) =: m$$

Remember that a has a beneficial repetition policy, i.e., we can ensure that the states $S_{>a}^{\mathfrak{P}_k}$ are increasing with respect to all $x \in U_a$, which contradicts that m is constant w.r.t. $k \in \mathbb{N}$. □

With this last result, we now how beneficial repetition policies behave and can ensure termination of the search.

Theorem 4. *bmNAP is* PSPACE*-complete.*

Proof. Inclusion of PP proves the PSPACE-hardness. We prove by induction over the number of non-maintainable actions that a polynomial space algorithm solves bmNAP. First, we proved that a polynomial space algorithm exists if all actions are maintainable. Let a be an additional non-maintainable action for the induction step. Given a relaxed solution, we can transform it into an a-relaxed solution by the induction hypothesis. It suffices to prove that a NAP problem $\{S_0 \to g\}_\mathcal{A}$ can be decomposed into problems that can be solved with a polynomial number of sequential a repetitions. Consider the following algorithm: For each $i \in \mathbb{N}$, if $\{S_i \to g\}_{\mathcal{A}\setminus\{a\}} = \emptyset$ but $\{S_i \to g\}_\mathcal{A}$ has an a-relaxed solution, repeat a beneficially to reach S_{i+1} and continue, else terminate. This search terminates if the NAP is solvable after finding the minimal number of a repetitions needed. If the NAP is unsolvable, the search terminates as shown in the last lemma. □

Finitely Maintainable NAP

In this subsection, we deal with plans as words over A. Therefore, it is important to define validity of these words.

Definition 19. *For any NAD $\mathcal{A} = (A, V\sigma, \pi)$, and any actions $a, b \in A$, the NAD can be extended by the action ab, denoted by $\mathcal{A} \cup \{ab\} := (A \cup \{ab\}, V, \sigma', \pi')$, where σ', π' are extension of σ, π to the new action ab, and are defined for every variable $x \in V$ as:*

- *The effects are summed:* $\sigma'(ab, x) = \sigma(a, x) + \sigma(b, x)$.
- *The preconditions are backwardly propagated:* $\pi'(ab, x) = \max\{\pi(a, x), \pi(b, x) - \sigma(a, x)\}$.

Notice that if $S_0 \vDash ab$, then $S_0 \vDash a$, and $S_0 + a \vDash b$, meaning that the plan ab is valid from S_0.

Let us focus on the number of violated repetitions $v^{\mathfrak{P}}(a)$ that are produced during the search for the number of sequential repetitions of the non-maintainable action $a \in A$.

Lemma 6. *For non maintainable action $a \in A$, let $\mathfrak{P}_0, \mathfrak{P}_1$ be two successive a-relaxed solutions of $\{S_0 \to g\}_\mathcal{A}$ over $\mathbf{A} = (A, \mu), \mathbf{A} \cup \{a\}$, respectively. If $v^{\mathfrak{P}_1}(a) > v^{\mathfrak{P}_0}(a)$ and $[S_0 \to g]_\mathcal{A}^{\mathbf{A}\cup\{a\}} = \emptyset$, then, $\{S_0 \to g\}_\mathcal{A} = \emptyset$.*

Proof. Let $i_j := u^{\mathfrak{P}_j}(a); j \in \{0, 1\}$. Notice that there exist plans $l, w_0, u \in A^*$, s.t. $\#_a(w_0) = 0$, $l := w_0 a^{\mathfrak{P}_0[g,a]} ug$, and $l \in \mathrm{lin}(\mathfrak{P}_0)$. If $l \in \{S_0 \to g\}_\mathcal{A}$, we have a contradiction $l \in [S_0 \to g]_\mathcal{A}^{\mathbf{A}\cup\{a\}} = \emptyset$. Otherwise, assuming that there exists a minimal solution $l' \in \{S_0 + w_0 a^{i_0} \to g\}_\mathcal{A}$ w.r.t. a, then, there exists plans $w_1, u' \in A^*$ s.t. $l' = w_1 a^{i_1} u' g$, with $\#_a(w_1) = 0$. Therefore, for all $x \in V$:

$$S_0(x) + \sigma(w_0, x) + \sigma(w_1, x) + (i_0 + i_1)\sigma(a, x) \geq \pi(u'g, x)$$

If $l'' = w_0 a^{i_0} u' g \in \{S_0 \to g\}_A$, then, l' is not minimal, therefore, $l'' \notin \{S_0 \to g\}_A$, i.e., there exists $x \in V$ s.t.:

$$S_0(x) + \sigma(w_0, x) + i_0\sigma(a, x) < \pi(u'g, x)$$

By combining the last two inequalities:

$$i_1\sigma(a, x) + \sigma(w_1, x) > 0$$

Since $\mathfrak{P}_0$ is an a-relaxed solution:

$$S_0(x) + \mathfrak{P}_0[g, a]\sigma(a, x) + \sigma(w_0, x) \geq \pi(ug, x)$$

By summing the last two inequalities:

$$S_0(x) + (i_1 + \mathfrak{P}_0[g, a])\sigma(a, x) + \sigma(w_0, x) + \sigma(w_1, x) > \pi(ug, x)$$

Notice that if $l''' := w_0 a^{i_0} w_1 a^{\mathfrak{P}_1[g,a]} ug \in \mathrm{lin}(\mathfrak{P}_1)$, we have a contradiction for $l''' \in \{S_0 \to g\}_\mathcal{A}$, because in that case $l''' \in [S_0 \to g]_\mathcal{A}^{\mathbf{A}\cup\{a\}} = \emptyset$. Therefore, since $\mathfrak{P}_1$ is minimal w.r.t. a:

$$S_0(x) + (i_0 + \mathfrak{P}_1[g, a] - 1)\sigma(a, x) + \sigma(w_0, x) + \sigma(w_1, x) < \pi(ug, x)$$

By combining the last two inequalities, we get

$$v^{\mathfrak{P}_1}(a) - 1 = \mathfrak{P}_1[g, a] - i_1 - 1 < \mathfrak{P}_0[g, a] - i_0 = v^{\mathfrak{P}_0}(a)$$

which contradicts $v^{\mathfrak{P}_1}(a) > v^{\mathfrak{P}_0}(a)$. □

For this reason, we can abort the search whenever the number of violated a repetitions in the a-relaxed solution increases after adding a sequential repetition of a.

Theorem 5. *fmNAP is decidable.*

Proof. Notice that a de-violating repetition policy guarantees termination because the number of violated copies eventually decreases. Since we can prune if that number increases as proven in the last lemma, we ensure termination of the search. □