Bisimulation Metric Computation for Markov Decision Processes

Sherry Shanshan Ruan Prakash Panangaden Gheorghe Comanici

20th August 2014

1 Asynchronous computation

Given an MDP $\mathcal{M} = \{S, A, P, R, \gamma\}$, we define a functional F on the state metric d. For instance,

$$F(d)(s, s') = \max_{a \in A}(|R_{sa} - R_{s'a}| + \gamma \mathcal{T}(d, P_{sa}, P_{s'a}))$$

where $\mathcal{T}(d, P_{sa}, P_{s'a})$ is the Kantorovich (Wasserstein) metric defined on probability measures:

$$\mathcal{T}(d, P_{sa}, P_{s'a}) = \inf_{\pi \in \Lambda(P_{sa}, P_{s'a})} \sum_{s_i, s_j \in S} \pi(s_i, s_j) d(s_i, s_j)$$

$$\Lambda(P_{sa}, P_{s'a}) = \left\{ \pi : S \times S \to \mathbb{R}_{\geq 0} \mid \sum_{s_i \in S} \pi(s_i, s_j) = P_{s'a}^{s_j}, \sum_{s_j \in S} \pi(s_i, s_j) = P_{sa}^{s_i} \right\}$$

We know that there is a fix point d^* of the functional F such that $d^* = F(d^*)$, and we can compute it by iteratively applying F to an initial metric. The convergence is verified by the theorem below [Ferns et al., 2004]:

$$d^* = \sup_{n \in \mathbb{N}} F^n(0)$$

Updating the entire state space every time can be expensive, so we seek ways to compute them in an "asynchronous" way. We can choose $\{K_i\}_{i=1}^{\infty}$, a sequence of subsets of $S \times S$, and only update metrics related to subset K_i at step i. The convergence of this asynchronous computation can be established provided that the sequence of subsets is carefully designed. We build the proof based on previous work by Comanici et al. [2012], with a more succinct mathematical notation.

Definition 1. Given a set S of states and a subset $K \subseteq S \times S$, we define the operator $F|_K : (S \times S \to \mathbb{R}_{\geq 0}) \to (S \times S \to \mathbb{R}_{\geq 0})$ as

$$F|_K(d)(s, s') = \begin{cases} d(s, s'), & \text{if } (s, s') \notin K \\ F(d)(s, s'), & \text{if } (s, s') \in K \end{cases} \qquad \forall \text{ metric } d, \forall (s, s') \in S \times S$$

Intuitively, the metric remains the same if pairs are not in K, and the metric gets "updated" if pairs lie in K.

Definition 2. Given a sequence of subsets $\{K_i\}_{i=1}^{\infty}$ where $K_i \subseteq S \times S$, we define the function $h_n : S \times S \to \mathbb{R}_{\geq 0}$ for any $n \in \mathbb{N}$ as

$$h_n = \begin{cases} F|_{K_n} \circ F|_{K_{n-1}} \circ \cdots \circ F|_{K_1}(0), & \text{if } n > 0 \\ 0, & \text{if } n = 0 \end{cases}$$

where 0 is the zero metric.

We will show shortly in lemma 6 that h_n is an increasing function, so we can think of h_n as an accumulative metric obtained at step n.

Lemma 3. Given two metrics d_1 and d_2 over S such that $d_1 \leq d_2$ (i.e., $\forall (s, s') \in S \times S$, $d_1(s, s') \leq d_2(s, s')$), then

$$\mathcal{T}(d_1, P, Q) \leq \mathcal{T}(d_2, P, Q) \qquad \text{for any two probability measures } P, Q$$

Proof . Let P, Q be two probability measures and define

$$\Lambda(P, Q) = \left\{ \pi : S \times S \to \mathbb{R}_{\geq 0} \mid \sum_{s_i \in S} \pi(s_i, s_j) = P(s_j), \sum_{s_j \in S} \pi(s_i, s_j) = Q(s_i) \right\}$$

Let $\pi' \in \Lambda(P, Q)$ be an arbitrary transportation strategy, then

$$\begin{aligned} \mathcal{T}(d_1, P, Q) &= \inf_{\pi \in \Lambda(P,Q)} \sum_{s_i, s_j \in S} \pi(s_i, s_j) d_1(s_i, s_j) \\ &\leq \sum_{s_i, s_j \in S} \pi'(s_i, s_j) d_1(s_i, s_j) \\ &\leq \sum_{s_i, s_j \in S} \pi'(s_i, s_j) d_2(s_i, s_j) \end{aligned}$$

Since this holds for all $\pi' \in \Lambda(P, Q)$, we conclude $\mathcal{T}(d_1, P, Q) \leq \mathcal{T}(d_2, P, Q)$. ∎

Lemma 4. Given two metrics d_1 and d_2 over S such that $d_1 \leq d_2$, then

$$F(d_1) \leq F(d_2)$$

Proof . Let $(s, s') \in S \times S$ be given, then for any action $a \in A$:

$$\begin{aligned} |R_{sa} - R_{s'a}| + \mathcal{T}(d_1, P_{sa}, P_{s'a}) &\leq |R_{sa} - R_{s'a}| + \mathcal{T}(d_2, P_{sa}, P_{s'a}) && \text{by lemma 3} \\ &\leq \max_{a' \in A}(|R_{sa'} - R_{s'a'}| + \mathcal{T}(d_2, P_{sa'}, P_{s'a'})) \\ &= F(d_2)(s, s') \end{aligned}$$

Since this is true for all $a \in A$, we conclude that $F(d_1)(s, s') \leq F(d_2)(s, s')$. ∎

Corollary 5.

$$F^n(0) \leq F^m(0) \qquad \text{if } n < m$$

Lemma 6 (Monotonicity)**.** Given a fixed sequence of subsets $\{K_i\}_{i=1}^{\infty}$ and an integer $m \in \mathbb{N}$, then $\forall n < m, n \in \mathbb{N}$,

$$h_n \leq h_m$$

that is, $\forall (s, s') \in S \times S,\ h_n(s, s') \leq h_m(s, s')$

Proof . We prove it by induction on m.

Base case ($m = 1$): Since $n < m$, the only possibility for n is 0, and it clearly holds that $h_m = h_1 \geq 0 = h_n$.

Inductive case: Suppose it holds for all integers $\leq m$. Let $(s, s') \in S \times S$ be a pair of states given. Let n be any integer less than $m + 1$, then $n \leq m$. We proceed by analyzing whether (s, s') is contained in K_{m+1}.

- Case I: $(s, s') \notin K_{m+1}$, then $h_{m+1}(s, s') = h_m(s, s')$, then $h_{m+1}(s, s') = h_m(s, s') \geq h_n(s, s')$ by i.h. on m.
- Case II: $(s, s') \in K_{m+1}$, then $h_{m+1}(s, s') = F(h_m)(s, s')$, then we proceed by a sub-case analysis.
  - Sub-case I: $h_n(s, s') = 0$, then clearly $h_{m+1}(s, s') \geq 0 = h_n(s, s')$.
  - Sub-case II: $h_n(s, s') = F(h_r)(s, s')$ for $0 \leq r < n$, then $h_m \geq h_r$ by i.h. on m, so $h_{m+1}(s, s') = F(h_m)(s, s') \geq F(h_r)(s, s') = h_n(s, s')$ by lemma 4.

Therefore, $h_{m+1}(s, s') \geq h_n(s, s')$ for any $n < m + 1$. ∎

Lemma 7 (Upper bound)**.** Given a sequence of subsets $\{K_i\}_{i=1}^{\infty}$,

$$h_n \leq d^* \qquad \text{for all } n \in \mathbb{N}$$

Proof . We prove it by induction on n.
Base case (n = 0): It is trivial that $h_0 = 0 \leq d^*$

Inductive case: Suppose $h_n \leq d^*$ holds and consider $(s, s') \in S \times S$,

- If $(s, s') \in K_{n+1}$, then $h_{n+1}(s, s') = F(h_n)(s, s') \leq F(d^*)(s, s')$ by i.h. and lemma 4. Note that $F(d^*)(s, s')$ is just $d^*(s, s')$ since d^* is a fixed point of F, so we get $h_{n+1}(s, s') \leq d^*(s, s')$.
- If $(s, s') \notin K_{n+1}$, then $h_{n+1}(s, s') = h_n(s, s') \leq d^*(s, s')$ by i.h..

Therefore, $\forall (s, s') \in S \times S$, $h_{n+1}(s, s') \leq d^*(s, s')$. ∎

Lemma 8 (Lower bound)**.** Suppose $\{K_i\}_{i=1}^{\infty}$ is a sequence of subsets of $S \times S$ satisfying $\bigcap_{n=1}^{\infty} \bigcup_{i=n}^{\infty} K_i = S \times S$, then $\forall n \in \mathbb{N}$, $\exists m \in \mathbb{N}$ such that

$$F^n(0) \leq h_m$$

Proof . We prove it by induction on n.

Base case ($n = 0$): We can simply set $m = 1$ since $F^0(0) = 0 \leq h_1$.

Inductive case: Suppose there exists an integer m such that $F^n(0) \leq h_m$. We want to construct an integer k such that $F^{n+1}(0) \leq h_k$.

Let (s, s') be an arbitrary pair of states in $S \times S$, then by the assumption, we know that $(s, s') \in \bigcup_{i=m+1}^{\infty} K_i$. Without loss of generality, assume $(s, s') \in K_r$ for some $r \geq m + 1$. Then we have

$$\begin{aligned} h_r(s, s') &= F(h_{r-1})(s, s') && \text{since } (s, s') \in K_r \\ &\geq F(h_m)(s, s') && \text{by lemma 6} \\ &\geq F(F^n(0))(s, s') && \text{by i.h. and lemma 4} \\ &= F^{n+1}(0)(s, s') && \text{by definition} \end{aligned}$$

Note that we can find such an r for each pair $(s, s') \in S \times S$. Let

$$k = \sup_{(s,s')} \{r \mid h_r(s, s') \geq F^{n+1}(0)(s, s')\}$$

The supremum exists because the state space is finite. Then by lemma 6, $h_k(s, s') \geq F^{n+1}(0)(s, s')$ for all $(s, s') \in S \times S$, so $h_k \geq F^{n+1}(0)$. ∎

Theorem 9. Suppose $\{K_i\}_{i=1}^{\infty}$ is a sequence of subsets of $S \times S$ such that $\bigcap_{n=1}^{\infty} \bigcup_{i=n}^{\infty} K_i = S \times S$, then

$$\sup_{n \in \mathbb{N}} h_n = d^*$$

Proof . By lemma 7 and lemma 8, $\forall n \in \mathbb{N}, \exists m \in \mathbb{N}$ such that

$$F^n(0) \leq h_m \leq d^*$$

Taking the supremum of them over all natural numbers, we get

$$d^* = \sup_{n \in \mathbb{N}} F^n(0) \leq \sup_{m \in \mathbb{N}} h_m \leq d^*$$

Therefore $\sup_{n \in \mathbb{N}} h_n = d^*$ ∎

Remark 10. To really understand what the assumption $\bigcap_{n=1}^{\infty} \bigcup_{i=n}^{\infty} K_i = S \times S$ means, we can interpret it as follows.

Definition 11. Given a sequence of subsets $\{K_i\}_{i=1}^{\infty}$ where $K_i \subseteq S \times S$, we define the function $\mathcal{X}_i : S \times S \to \{0, 1\}$ for any $i \in \mathbb{N}$ as follows

$$\mathcal{X}_i(s, s') = \begin{cases} 0, & \text{if } (s, s') \notin K_i \\ 1, & \text{if } (s, s') \in K_i \end{cases} \qquad \forall (s, s') \in S \times S$$

We also define the function $a_n : S \times S \to \mathbb{N}$ for any $n \in \mathbb{N}$ as follows

$$a_n(s, s') = \sum_{i=1}^{n} \mathcal{X}_i(s, s') \qquad \forall (s, s') \in S \times S$$

$a_n(s, s')$ is the total number of updates actually performed on the pair (s, s') up to step n.

Definition 12. Following previous assumptions, we define $\eta(n)$ as the least number of updates over all pairs of states in $S \times S$ up to step n.

$$\eta(n) = \min_{(s,s')} a_n(s, s') \qquad \text{for any } n \in \mathbb{N}$$

In other words, all pairs in $S \times S$ are guaranteed to get updated for at least $\eta(n)$ times up to step n.

Theorem 13. Suppose $\{K_i\}_{i=1}^{\infty}$ is a sequence of subsets of $S \times S$ ($K_i \subseteq S \times S$, for $i = 1, 2, 3, \cdots$), then the following holds

$$\bigcap_{n=1}^{\infty} \bigcup_{i=n}^{\infty} K_i = S \times S \iff \lim_{n\to\infty} \eta(n) = \infty$$

Proof . ($\Rightarrow$) Assume $\bigcap_{n=1}^{\infty} \bigcup_{i=n}^{\infty} K_i = S \times S$ holds. Suppose $\lim\limits_{n\to\infty} \eta(n) < \infty$, then we can find a pair $(s, s') \in S \times S$ such that $\sum_{i=1}^{\infty} \delta_i(s, s') < \infty$. Since $\delta_i(s, s')$ is defined to be either 0 or 1, there must exist an $m \in \mathbb{N}$ such that $\delta_i(s, s') = 0$ for all $i > m$, that is, $(s, s') \notin K_i$ for all $i > m$. Then it follows that $(s, s') \notin \bigcup_{i=m+1}^{\infty} K_i$, which clearly contradicts the assumption.

($\Leftarrow$) Assume $\lim_{n\to\infty} \eta(n) = \infty$. Then $\sum_{i=1}^{\infty} \delta_i(s, s') = \infty$ for all $(s, s') \in S \times S$. Since $\delta_i(s, s')$ is defined to be either 0 or 1, there are infinitely many $i \in \mathbb{N}$ such that $\delta_i(s, s') = 1$, that is, for any $m \in \mathbb{N}$, there exists an $i \geq m$ such that $(s, s') \in K_i$. Note that this is true for every $(s, s') \in S \times S$, so we get $\forall (s, s') \in S \times S$, $\forall m \in \mathbb{N}$, $(s, s') \in \bigcup_{i=m}^{\infty} K_i$. Thus, we conclude that $S \times S = \bigcap_{m=1}^{\infty} \bigcup_{i=m}^{\infty} K_i$. ∎

2 Synchronous partitions and metrics

In this section, we establish connections between partitions (relations) and metrics. First, we introduce the concept of partitions and some related definitions which will be used to characterize bisimulation metrics.

Definition 14. Given a set S, a *partition* of S is a collection of subsets $\Pi = \{C_i\}_{i \in I}$ (I is a countable index set) such that $\forall C_1 \neq C_2 \in \Pi, C_1 \bigcap C_2 = \emptyset$ and $\bigcup_{i \in I} C_i = S$. We call each C_i a *block* of Π.

Clearly, every partition defines an equivalence relation, and every equivalence relation defines a partition. We use Π to denote the partition and use R_Π to represent the corresponding equivalence relation.

Definition 15. We say Π' *refines* Π if $R_{\Pi'} \subseteq R_\Pi$, where the subset relation $\subseteq$ is defined as

$$R_{\Pi'} \subseteq R_\Pi \iff sR_\Pi t \text{ whenever } sR_{\Pi'}t$$

i.e., blocks of Π' sit inside blocks of Π

We proceed to consider a sequence of refined partitions paired with metrics. Given a set S, let $\{(\Pi_n, d_n)\}_{n=1}^{\infty}$ be a sequence of metric spaces, where Π_n's are partitions of S such that Π_n refines Π_{n-1} (i.e., $R_{\Pi_n} \subseteq R_{\Pi_{n-1}}$) and d_n is a metric defined on blocks of Π_n. One way to define such a sequence is as follows:

Definition 16. We define $\{(\Pi_n, d_n)\}_{n=1}^{\infty}$ inductively as follows:

(1) $\Pi_0 = \{S\}$, $d_0 = 0$

(2) The equivalence relation R_n is defined as

$$sR_ns' \iff \forall a \in A, R_{sa} = R_{s'a} \text{ and } \forall C \in \Pi_{n-1}, P_{sa}^C = P_{s'a}^C$$

(3) $\Pi_n = \bigcup_{C \in \Pi_{n-1}} D_n(C)$ where $D_n : \mathcal{P}(S) \to \mathcal{P}(\mathcal{P}(S))$ is defined as

$$D_n(C) = C/R_n$$

(4) The metric $d_n : \Pi_n \times \Pi_n \to \mathbb{R}_{\geq 0}$ is defined as

$$d_n(C_1, C_2) = \max_{a \in A}(|R_{s_1a} - R_{s_2a}| + \mathcal{T}(d_{n-1}, P_{s_1a}, P_{s_2a})),\ s_1 \in C_1, s_2 \in C_2$$

Note that d_{n-1} is defined on $\Pi_{n-1} \times \Pi_{n-1}$, so the Kantorovich (Wasserstein) metric $\mathcal{T}(d_{n-1}, P_{s_1a}, P_{s_2a})$ is actually over Π_{n-1}. It is easy to check that the metric is well-defined.

The connection between relations R_{Π_n} and metrics d_n can be characterized as

$$d_n(s, s') = 0 \iff sR_{\Pi_n}s'$$

Theorem 17. Π_n is a partition for all $n \in \mathbb{N}$

Proof . We prove it by induction on n.

Base case: $\Pi_0 = \{S\}$ is obviously well-defined.

Inductive case: Suppose $\Pi_n = \{C_i\}_{i \in I}$ is a partition, then $\Pi_{n+1} = \bigcup_{C \in \Pi_n} D_{n+1}(C)$.

Clearly, $\bigcup_{C' \in \Pi_{n+1}} C' = \bigcup_{C \in \Pi_n} C = S$.

Let $C'_j \neq C'_k \in \Pi_{n+1}$, and suppose $C'_j \in C_j/R_n$ and $C'_k \in C_k/R_n$ where $C_j, C_k \in \Pi_n$.

- If $C_j \neq C_k$, then $C_j \bigcap C_k = \emptyset$ by i.h., so $C'_j \bigcap C'_k = \emptyset$.
- If $C_j = C_k$, then $C'_j \bigcap C'_k = \emptyset$ since they are different equivalence classes of C_j (or C_k).

∎

Definition 18. Define $\hat{d_n} : S \times S \to \mathbb{R}_{\geq 0}$ as

$$\hat{d_n}(s, s') = d_n(C, C'), \quad \forall s \in C, s' \in C'$$

Lemma 19.

$$\mathcal{T}(d_n, P_{sa}, P_{s'a}) = \mathcal{T}(\hat{d_n}, P_{sa}, P_{s'a})$$

Proof . Define

$$\Lambda(P_{sa}, P_{s'a}) = \left\{ \pi : \Pi_n \times \Pi_n \to \mathbb{R}_{\geq 0} \mid \sum_{C_i \in \Pi_n} \pi(C_i, C_j) = P_{s'a}^{C_j}, \sum_{C_j \in \Pi_n} \pi(C_i, C_j) = P_{sa}^{C_i} \right\}$$

$$\hat{\Lambda}(P_{sa}, P_{s'a}) = \left\{ \hat{\pi} : S \times S \to \mathbb{R}_{\geq 0} \mid \sum_{s_i \in S} \hat{\pi}(s_i, s_j) = P_{s'a}^{s_j}, \sum_{s_j \in S} \hat{\pi}(s_i, s_j) = P_{sa}^{s_i} \right\}$$

We first show that $\mathcal{T}(d_n, P_{sa}, P_{s'a}) \leq \mathcal{T}(\hat{d_n}, P_{sa}, P_{s'a})$. We prove it by transforming each strategy $\hat{\pi} \in \hat{\Lambda}(P_{sa}, P_{s'a})$ to a strategy $\underline{\pi} \in \Lambda(P_{sa}, P_{s'a})$.

For every $\hat{\pi} \in \hat{\Lambda}(P_{sa}, P_{s'a})$, we define a corresponding $\underline{\pi} : \Pi_n \times \Pi_n \to \mathbb{R}_{\geq 0}$ as

$$\underline{\pi}(C_i, C_j) = \sum_{s_i \in C_i} \sum_{s_j \in C_j} \hat{\pi}(s_i, s_j)$$

Then we claim that $\underline{\pi} \in \Lambda(P_{sa}, P_{s'a})$ because

$$\begin{aligned}
\sum_{C_i \in \Pi_n} \underline{\pi}(C_i, C_j) &= \sum_{C_i \in \Pi_n} \sum_{s_i \in C_i} \sum_{s_j \in C_j} \hat{\pi}(s_i, s_j) \\
&= \sum_{s_j \in C_j} \sum_{C_i \in \Pi_n} \sum_{s_i \in C_i} \hat{\pi}(s_i, s_j) \\
&= \sum_{s_j \in C_j} \sum_{s_i \in S} \hat{\pi}(s_i, s_j) \\
&= \sum_{s_j \in C_j} P_{s'a}^{s_j} && \text{since } \hat{\pi} \in \hat{\Lambda}(P_{sa}, P_{s'a}) \\
&= P_{s'a}^{C_j}
\end{aligned}$$

Similarly we can show that $\sum_{C_j \in \Pi_n} \underline{\pi}(C_i, C_j) = P_{s'a}^{C_i}$. Hence, $\underline{\pi}$ is indeed in $\Lambda(P_{sa}, P_{s'a})$.

Then we can show that

$$\begin{aligned}
\mathcal{T}(\hat{d}_n, P_{sa}, P_{s'a}) &= \inf_{\hat{\pi}\in\hat{\Lambda}(P_{sa},P_{s'a})} \sum_{s_i,s_j\in S} \hat{\pi}(s_i,s_j)\hat{d}_n(s_i,s_j) \\
&= \inf_{\hat{\pi}\in\hat{\Lambda}(P_{sa},P_{s'a})} \sum_{C_i,C_j\in\Pi_n} \sum_{s_i\in C_i} \sum_{s_j\in C_j} \hat{\pi}(s_i,s_j)\hat{d}_n(s_i,s_j) \\
&= \inf_{\hat{\pi}\in\hat{\Lambda}(P_{sa},P_{s'a})} \sum_{C_i,C_j\in\Pi_n} d_n(C_i,C_j) \sum_{s_i\in C_i} \sum_{s_j\in C_j} \hat{\pi}(s_i,s_j) \\
&= \inf_{\hat{\pi}\in\hat{\Lambda}(P_{sa},P_{s'a})} \sum_{C_i,C_j\in\Pi_n} d_n(C_i,C_j)\underline{\pi}(C_i,C_j) && \text{by transformation} \\
&\geq \inf_{\pi\in\Lambda(P_{sa},P_{s'a})} \sum_{C_i,C_j\in\Pi_n} d_n(C_i,C_j)\pi(C_i,C_j) \\
&= \mathcal{T}(d_n, P_{sa}, P_{s'a})
\end{aligned}$$

Next, we show that $\mathcal{T}(d_n, P_{sa}, P_{s'a}) \geq \mathcal{T}(\hat{d}_n, P_{sa}, P_{s'a})$. We prove it by transforming each strategy $\pi \in \Lambda(P_{sa}, P_{s'a})$ to a strategy $\overline{\pi} \in \hat{\Lambda}(P_{sa}, P_{s'a})$.

For every $\pi \in \Lambda(P_{sa}, P_{s'a})$, we define a corresponding $\overline{\pi}$ as

$$\overline{\pi}(s_i,s_j) = \frac{P_{sa}^{s_i}}{P_{sa}^{C_i}} \frac{P_{s'a}^{s_j}}{P_{s'a}^{C_j}} \pi(C_i,C_j), \text{ where } s_i \in C_i, s_j \in C_j$$

We claim that $\overline{\pi} \in \hat{\Lambda}(P_{sa}, P_{s'a})$ because

$$\begin{aligned}
\sum_{s_i\in S} \overline{\pi}(s_i,s_j) &= \sum_{s_i\in S} \frac{P_{sa}^{s_i}}{P_{sa}^{C_i}} \frac{P_{s'a}^{s_j}}{P_{s'a}^{C_j}} \pi(C_i,C_j) && \text{where } s_i \in C_i, s_j \in C_j \\
&= \sum_{C_i\in\Pi_n} \sum_{s_i\in C_i} \frac{P_{sa}^{s_i}}{P_{sa}^{C_i}} \frac{P_{s'a}^{s_j}}{P_{s'a}^{C_j}} \pi(C_i,C_j) \\
&= \sum_{C_i\in\Pi_n} \frac{P_{s'a}^{s_j}}{P_{s'a}^{C_j}} \pi(C_i,C_j) \\
&= \frac{P_{s'a}^{s_j}}{P_{s'a}^{C_j}} \sum_{C_i\in\Pi_n} \pi(C_i,C_j) \\
&= \frac{P_{s'a}^{s_j}}{P_{s'a}^{C_j}} P_{s'a}^{C_j} && \text{since } \pi \in \Lambda(P_{sa}, P_{s'a}) \\
&= P_{s'a}^{s_j}
\end{aligned}$$

Similarly we can show that $\sum_{s_j \in S} \overline{\pi}(s_i, s_j) = P_{s'a}^{s_i}$. Hence, $\overline{\pi}$ is indeed in $\hat{\Lambda}(P_{sa}, P_{s'a})$.

Then we can show that

$$\begin{aligned}
\mathcal{T}(d_n, P_{sa}, P_{s'a}) &= \inf_{\pi \in \Lambda(P_{sa}, P_{s'a})} \sum_{C_i, C_j \in \Pi_n} \pi(C_i, C_j) d_n(C_i, C_j) \\
&= \inf_{\pi \in \Lambda(P_{sa}, P_{s'a})} \sum_{C_i, C_j \in \Pi_n} \sum_{s_i \in C_i} \sum_{s_j \in C_j} \frac{P_{sa}^{s_i}}{P_{sa}^{C_i}} \frac{P_{s'a}^{s_j}}{P_{s'a}^{C_j}} \pi(C_i, C_j) d_n(s_i, s_j) \\
&= \inf_{\pi \in \Lambda(P_{sa}, P_{s'a})} \sum_{C_i, C_j \in \Pi_n} \sum_{s_i \in C_i} \sum_{s_j \in C_j} \overline{\pi}(s_i, s_j) d_n(s_i, s_j) && \text{by transformation} \\
&= \inf_{\pi \in \Lambda(P_{sa}, P_{s'a})} \sum_{s_i, s_j \in S} \overline{\pi}(s_i, s_j) d_n(s_i, s_j) \\
&\geq \inf_{\hat{\pi} \in \hat{\Lambda}(P_{sa}, P_{s'a})} \sum_{s_i, s_j \in S} \hat{\pi}(s_i, s_j) d_n(s_i, s_j) \\
&= \mathcal{T}(\hat{d}_n, P_{sa}, P_{s'a})
\end{aligned}$$

Therefore, $\mathcal{T}(d_n, P_{sa}, P_{s'a}) = \mathcal{T}(\hat{d}_n, P_{sa}, P_{s'a})$. ∎

Lemma 20.

$$\hat{d}_n = F^n(0)$$

Proof . We show it by induction on n.

Base case ($n = 0$): $\hat{d}_0 = d_0 = 0 = F^0(0)$

Inductive case: Suppose $\hat{d}_n = F^n(0)$. Let $s \in C, s' \in C'$, then

$$\begin{aligned}
\hat{d}_{n+1}(s, s') &= d_{n+1}(C, C') && \text{by definition} \\
&= \max_{a \in A}(|R_{sa} - R_{s'a}| + \mathcal{T}(d_n, P_{sa}, P_{s'a})) && \text{by definition} \\
&= \max_{a \in A}(|R_{sa} - R_{s'a}| + \mathcal{T}(\hat{d}_n, P_{sa}, P_{s'a})) && \text{by lemma 19} \\
&= \max_{a \in A}(|R_{sa} - R_{s'a}| + \mathcal{T}(F^n(0), P_{sa}, P_{s'a})) && \text{by i.h.} \\
&= F^{n+1}(0)(s, s') && \text{by definition}
\end{aligned}$$

This holds for every $(s, s') \in S \times S$, so $\hat{d}_{n+1} = F^{n+1}(0)$ ∎

Theorem 21. The fix point of F, namely d^*, can be constructed by iterating through partition algorithm described in definition 16.

Proof . Directly from lemma 20. ∎

This corollary establishes the correspondence between the synchronous partition procedure described above and the theoretical characterization of synchronous computation.

3 Asynchronous partitions

In this section, we present an asynchronous partition algorithm and then build the connection between the asynchronous algorithm and the theoretical characterization of asynchronous computation proved in section 1.

Definition 22. Following same notations in definition 16, we define $\{(\Pi_n, d_n)\}_{n=1}^{\infty}$ inductively as follows:

(1) $\Pi_0 = \{S\}$, $d_0 = 0$

(2) $\Pi_n = \left(\Pi_{n-1} \setminus \{C^*_{n-1}, C^{**}_{n-1}\}\right) \bigcup D_n(C^*_{n-1}) \bigcup D_n(C^{**}_{n-1})$ where C^*_{n-1} and C^{**}_{n-1} are two different block selected from Π_{n-1}

(3) The metric $d_n : \Pi_n \times \Pi_n \to \mathbb{R}_{\geq 0}$ is defined as

- If $C_1, C_2 \in D_n(C^*_{n-1}) \bigcup D_n(C^{**}_{n-1})$,

$$d_n(C_1, C_2) = \max_{a \in A}(|R_{s_1 a} - R_{s_2 a}| + \mathcal{T}(d_{n-1}, P_{s_1 a}, P_{s_2 a})),\ s_1 \in C_1, s_2 \in C_2$$

- Otherwise,

$$d_n(C_1, C_2) = d_{n-1}(C_1, C_2)$$

It is straightforward to verify that the metric is well-defined.

Definition 23. Define $\hat{d_n} : S \times S \to \mathbb{R}_{\geq 0}$ as

$$\hat{d_n}(s, s') = d_n(C, C'), \quad \forall s \in C, s' \in C'$$

Lemma 24. Given $\{K_n\}_{n=1}^{\infty}$, a sequence of subsets of $S \times S$ such that $K_{n+1} = (C^*_n \bigcup C^{**}_n) \times (C^*_n \bigcup C^{**}_n)$, then

$$\hat{d_n} = h_n$$

Note that C^*, C^{**} are given in definition 22 and h_n is defined in definition 2.

Proof . We show it by induction on n.

Base case ($n = 0$): $\hat{d}_0 = 0 = h_0$

Inductive case: Suppose $\hat{d}_n = h_n$. We want to show that $\hat{d}_{n+1} = h_{n+1}$

If $(s, s') \in K_{n+1}$, then $s, s' \in C_n^* \bigcup C_n^{**}$. Suppose $s \in C, s' \in C'$ and $C, C' \in D_{n+1}(C_n^*) \bigcup D_{n+1}(C_n^{**})$.

$$
\begin{aligned}
h_{n+1}(s, s') &= F(h_n)(s, s') && \text{by definition of } h \\
&= F(\hat{d}_n)(s, s') && \text{by i.h.} \\
&= \max_{a \in A}(|R_{sa} - R_{s'a}| + \mathcal{T}(\hat{d}_n, P_{sa}, P_{s'a})) && \text{by definition of } F \\
&= \max_{a \in A}(|R_{sa} - R_{s'a}| + \mathcal{T}(d_n, P_{sa}, P_{s'a})) && \text{by lemma 19} \\
&= d_{n+1}(C, C') && \text{by definition of } d_{n+1} \\
&= \hat{d}_{n+1}(s, s') && \text{by assumption}
\end{aligned}
$$

If $(s, s') \notin K_{n+1}$, then without loss of generality suppose $s' \in C'$ where $C' \in \Pi_n$ and $C' \neq C_n^*, C' \neq C_n^{**}$.

$$
\begin{aligned}
h_{n+1}(s, s') &= h_n(s, s') && \text{by definition of } h \\
&= \hat{d}_n(s, s') && \text{by i.h.} \\
&= d_n(C, C') && \text{by assumption} \\
&= d_{n+1}(C, C') && \text{by definition of } d_{n+1} \\
&= \hat{d}_{n+1}(s, s') && \text{by assumption}
\end{aligned}
$$

This holds for every $(s, s') \in S \times S$, so $\hat{d}_{n+1} = h_{n+1}$ ∎

Theorem 25. The fix point of F, namely d^*, can be constructed by iterating through asynchronous partition algorithm described in definition 22 provided that the sequence of blocks $\{C_n^*, C_n^{**}\}_{n=1}^{\infty}$ satisfies the condition given in theorem 9 .

Proof . Directly from lemma 24 and the same argument in theorem 9. ∎

References

Gheorghe Comanici, Prakash Panangaden, and Doina Precup. On-the-fly algorithms for bisimulation metrics. In *QEST*, pages 94–103, 2012.

Norm Ferns, Prakash Panangaden, and Doina Precup. Metrics for finite markov decision processes. In David Maxwell Chickering and Joseph Y. Halpern, editors, *UAI*, pages 162–169. AUAI Press, 2004. ISBN 0-9749039-0-6.