From Simple Stochastic Games to Bisimulation Pseudometrics on Markov Decision Processes

Franck van Breugel[1,2] and James Worrell[1]

[1] Department of Computer Science, University of Oxford, UK
[2] Department of Computer Science, York University, Canada

Abstract. In this paper we investigate the complexity of computing bisimulation pseudometrics on Markov decision processes (MDPs). Our first main result is that such pseudometrics can be computed in the complexity class **PPAD**. We show that another well-known problem in **PPAD**—computing the value of a simple stochastic game (SSG)—can be reduced in logarithmic space to the problem of computing the bisimulation pseudometric on a given MDP. In the other direction, we reduce the problem of computing the bisimulation pseudometric to that of computing the value of an SSG. This reduction uses a construction similar to the classical attacker-defender game for bisimulation in the non-probabilistic case, and works in polynomial time for MDPs of a fixed branching degree. Finally, we investigate whether the above bound on the branching degree can be dropped, relating it to the question of whether there is a family of polynomial size SSGs that solve the linear assignment problem.

1 Introduction

Recently, a family of *pseudometrics* on labelled Markov chains, *Markov decision processes* (MDPs for short), and stochastic games has been proposed to measure behavioural equivalence of states (see, for example, [8,9,14,19]). These pseudometrics can be seen as quantitative generalizations of *probabilistic bisimilarity*, a fundamental notion of equivalence for probabilistic systems [23,27] that generalises lumpability for Markov chains. Specifically, states have zero distance under the pseudometric if and only if they are bisimilar. A major motivation for such a quantitative relaxation of the notion of probabilistic bisimilarity is robustness: the resulting notion is less sensitive to numerical imprecision in modelling or computation [21]. The price to pay for this is that efficient computation of the pseudometrics becomes more challenging.

In the simplest setting, labelled Markov chains, the authors in collaboration with Chen [9] have shown that the pseudometric is rational and computable in polynomial time using the ellipsoid algorithm. In the most complex setting, stochastic games, the pseudometric has been shown to be irrational, and as hard as the sum-of-square-roots problem by Chatterjee, de Alfaro, Majumdar and Raman [8]. In this paper, we consider a pseudometric on MDPs, i.e., we are in the intermediate setting. De Alfaro, Majumdar, Raman and Stoelinga have provided a logical characterization of this pseudometric in [13]. Chatterjee, de Alfaro, Majumdar and Raman have considered a closely related pseudometric in [8]. They have shown that deciding whether the distance of two states is smaller than or equal to some given rational is in **PSPACE**. This decision procedure can be used to approximate the distance of two states by exploiting binary search. Fu [19] has shown that the pseudometric is rational and computable in **NP** $\cap$ **coNP**. It is even in **UP** $\cap$ **coUP** [20]. His proofs can easily be adapted to our setting [20]. Ferns, Panangaden and Precup [17,18] have considered a pseudometric on MDPs that discounts the future. They have shown that their pseudometric can be approximated in polynomial time.

In this paper, we investigate the computational complexity of the above-mentioned bisimulation pseudometric on MDPs in the undiscounted case. We show that the value of the pseudometric is rational and can be computed in **PPAD**, thus improving the bounds given in [8,20]. Our proof relies on a connection between bisimulation pseudometrics and *simple stochastic games* (SSGs for short). This generalises the classical game-theoretic formulation of bisimulation on labelled transition systems as a two-player turn-based finite game between "Attacker", who is trying to prove the systems are not bisimilar, and "Defender", who is trying to prove that the systems are bisimilar [28].

Recall that an SSG consists of a directed graph with a designated start vertex and two distinguished vertices, known as 0-sink and 1-sink, that have no outgoing edges. All other vertices have outdegree two and are classified either as max-vertices, min-vertices or as average-vertices. The game is played by two players, Player 0 and Player 1, with a single token. Initially, the token is placed on the start vertex of the graph. At each step of the game, the token is moved from a vertex to one of its two successors. At a min-vertex Player 0 chooses the successor, at a max-vertex Player 1 chooses the successor, and at an average-vertex the successor is chosen randomly. Player 1 wins the game if the token reaches the 1-sink. If the token reaches the 0-sink or the token never reaches a sink, Player 0 wins. As we will discuss below, one can assign a *value*, a real number in the unit interval, to each vertex of the graph. For the start vertex, this value captures the probability that Player 1 wins if both play optimally.

Computing the value of an SSG is a natural extension of linear programming. The decision version of the problem also generalises the problem of computing the winner of a parity game. It was shown to be in **NP** $\cap$ **coNP** by Condon [11] and, in fact, is known to be in **UP** $\cap$ **coUP** (see, for example, the paper [30] by Yannakakis). However, after 20 years, the exact complexity remains unknown. Computing the value of an SSG was shown to be in **PLS** by Yannakakis [29], **PPAD** by Etessami and Yannakakis [16], and **CLS** and **CCLS** by Daskalakis and Papadimitriou [12]. It is still not known whether this problem is complete for any complexity class.

We reduce the problem of computing the bisimulation pseudometric on an MDP to that of computing the value of an SSG derived from the MDP. The size of this game depends exponentially on the branching degree of the MDP and so our reduction is not polynomial-time in general. Nevertheless we are still able to inherit the **PPAD** complexity from SSGs. Here we use the result of Etessami and Yannakakis [16] that a fixed point of *polynomial piecewise linear* functional on $\mathbb{R}^n$ can be computed in **PPAD**.

The exponential blow up in the reduction from bisimulation pseudometrics to SSGs can be avoided if one can encode an arithmetic circuit to solve the *linear assignment problem* as an SSG. Recall that in the linear assignment problem there are n agents and n tasks and a cost matrix specifying the cost for each agent to perform each task. One has to assign exactly one agent to each task while minimizing the total cost (see, for example, [7] for a detailed discussion). Note that this is a quantitative generalization of the *perfect matching problem*.

We show that one can construct an MDP from a cost matrix such that the bisimulation distance of two designated states of the MDP coincides with the minimal total cost. Thus a polynomial transformation of bisimulation pseudometrics on MDPs to SSGs would imply that there are polynomial size SSGs that solve the linear assignment problem (i.e., for each n there would be an SSG $\mathcal{G}_n$ such that if the cost matrix of a linear assignment problem with n agents and n tasks is represented in the values of the sink nodes of $\mathcal{G}_n$, then the value of the start vertex gives the value of the optimal assignment). Note here the restriction that all vertices of an SSG compute monotone functions (min, max, or average). As shown by Razborov [25], there is a superpolynomial gap between the size of monotone and non-monotone (boolean-valued) circuits for computing a perfect matching. If there also exists a superpolynomial gap between the size of monotone and non-monotone (real-valued) circuits for solving the linear assignment problem (a problem that we leave for future research), then this may shed more light on the relative complexity of computing the bisimulation pseudometric on an MDP and computing the value of an SSG. Schwiegelshohn and Thiele [26] have presented an arithmetic circuit to solve the linear assignment problem. However, the circuit contains subtraction, which is a non-monotone operator.

In summary, the main contributions of this paper are the following.

- We prove that computing the bisimulation pseudometric on an MDP is in **PPAD**, thus improving the bounds given in [8,20].
- We show that computing the value of an SSG can be reduced in logarithmic space to computing the bisimulation pseudometric on an MDP.
- We prove that the above two problems are polynomial time equivalent if there is an absolute bound on the branching degree of the MDP.
- We show that the linear assignment problem can be reduced in logarithmic space to computing the bisimulation pseudometric on an MDP.

Due to lack of space, no proofs have been included. These can be found in [6].

Acknowledgements We would like to thank Rainer Burkard for discussions about the linear assignment problem and for pointing us to [26]. We would also like to thank Taolue Chen and Hongfei Fu for helpful discussions.

2 Orders and Metrics

Both value assignments to the vertices of an SSG and distance functions on the states of an MDP are functions mapping tuples to the unit interval. Such functions carry a natural order and metric which will play a key role in our technical development.

Let X be a set and $n \in \mathbb{N}$. The relation $\sqsubseteq \subseteq [0,1]^{X^n} \times [0,1]^{X^n}$ is defined by

$$f_1 \sqsubseteq f_2 \text{ iff } f_1(x_1, \ldots, x_n) \leq f_2(x_1, \ldots, x_n) \text{ for all } x_1, \ldots, x_n \in X.$$

One can easily verify that $\langle [0,1]^{X^n}, \sqsubseteq \rangle$ is a complete lattice.

The function $\| \cdot - \cdot \| : [0,1]^{X^n} \times [0,1]^{X^n} \to [0,1]$ is defined by

$$\|f_1 - f_2\| = \sup_{x_1, \ldots, x_n \in X} |f_1(x_1, \ldots, x_n) - f_2(x_1, \ldots, x_n)|.$$

One can also easily check that $\langle [0,1]^{X^n}, \| \cdot - \cdot \| \rangle$ is a nonempty complete metric space.

To define the bisimulation pseudometric on the states of an MDP, we will use two key ingredients: a distance function on finite sets and a distance function on probability distributions. The former captures the nondeterministic choices in an MDP, whereas the latter deals with the probabilistic choices in an MDP.

We denote the set of finite subsets of a set X by $\mathcal{P}(X)$. We define $\max \emptyset = 0$ and $\min \emptyset = 1$. We lift a distance function on X to a distance function on $\mathcal{P}(X)$ as follows.

Definition 1. *Let X be a set. The function $\mathcal{P} : [0,1]^{X \times X} \to [0,1]^{\mathcal{P}(X) \times \mathcal{P}(X)}$ is defined by*

$$\mathcal{P}(d)(A_1, A_2) = \max \left\{ \max_{x_1 \in A_1} \min_{x_2 \in A_2} d(x_1, x_2), \max_{x_2 \in A_2} \min_{x_1 \in A_1} d(x_2, x_1) \right\}.$$

The above is known as the Hausdorff metric (in case d is a metric). One can show that the singleton sets isometrically embed X into $\mathcal{P}(X)$, that is, for all x_1, $x_2 \in X$, $\mathcal{P}(d)(\{x_1\}, \{x_2\}) = d(x_1, x_2)$.

One can show that $\mathcal{P}$ preserves both the order and the metric: for all d_1, $d_2 \in [0,1]^{X \times X}$, if $d_1 \sqsubseteq d_2$ then $\mathcal{P}(d_1) \sqsubseteq \mathcal{P}(d_2)$, that is, $\mathcal{P}$ is monotone, and $\|\mathcal{P}(d_1) - \mathcal{P}(d_2)\| \leq \|d_1 - d_2\|$, that is, $\mathcal{P}$ is nonexpansive.

Recall that $d \in [0,1]^{X \times X}$ is a pseudometric if for all x_1, x_2, $x_3 \in X$, $d(x_1, x_1) = 0$, $d(x_1, x_2) = d(x_2, x_1)$, and $d(x_1, x_3) \leq d(x_1, x_2) + d(x_2, x_3)$. If d is a pseudometric then $\mathcal{P}(d)$ is a pseudometric as well (see, for example, [3, Proposition A.25]).

We denote the set of probability distributions on a set X by $\mathcal{D}(X)$. To lift a distance function on X to a distance function on $\mathcal{D}(X)$, we use the set of nonexpansive functions from the set X endowed with the distance function d to the unit interval, which we denote by $(X, d) \xrightarrow{1} [0,1]$.

Definition 2. *Let X be a set. The function $\mathcal{D} : [0,1]^{X \times X} \to [0,1]^{\mathcal{D}(X) \times \mathcal{D}(X)}$ is defined by*

$$\mathcal{D}(d)(\mu_1, \mu_2) = \sup \left\{ \sum_{x \in X} f(x)(\mu_1(x) - \mu_2(x)) \;\middle|\; f \in (X, d) \nrightarrow [0,1] \right\}.$$

The above is known as the Kantorovich metric (in case d is a metric). We denote the Dirac distribution centered on x by δ_x. Recall that this probability distribution maps x to one and all other elements to zero. One can show that these Dirac distributions isometrically embed X into $\mathcal{D}(X)$, that is, for all x_1, $x_2 \in X$, $\mathcal{D}(d)(\delta_{x_1}, \delta_{x_2}) = d(x_1, x_2)$ (see, for example, [15, page 108]). One can show that $\mathcal{D}$ is monotone (see, for example, [5, Proposition 38]) and nonexpansive (see, for example, [4, Section 3]). One can also prove that $\mathcal{D}(d)$ is a pseudometric if d is a pseudometric (see, for example, [15, Proposition 2.5.14]).

Definition 3. *Let X be a set. Let μ_1, $\mu_2 \in \mathcal{D}(X)$. Then $\omega \in \mathcal{D}(X \times X)$ is a coupling of μ_1 and μ_2 if for all x_1 $x_2 \in X$,*

$$\sum_{x_2 \in X} \omega(x_1, x_2) = \mu_1(x_1) \text{ and } \sum_{x_1 \in X} \omega(x_1, x_2) = \mu_2(x_2).$$

In other words, ω is a joint probability distribution whose marginals are μ_1 and μ_2. We denote the set of couplings of μ_1 and μ_2 by Ω_{μ_1,μ_2}. Using the duality theorem of linear programming (see, for example, [10, Theorem 5.1]) we can characterise $\mathcal{D}$ as follows:

$$\mathcal{D}(d)(\mu_1, \mu_2) = \min \left\{ \sum_{x_1, x_2 \in X} \omega(x_1, x_2) d(x_1, x_2) \;\middle|\; \omega \in \Omega_{\mu_1,\mu_2} \right\}.$$

In [9, Lemma 1], a characterization of probabilistic bisimilarity (to be defined in Definition 8) in terms of permutations is given. The following proposition can be viewed as a quantitative analogue of that result. We denote the set of permutations of the set $\{1, \dots, n\}$ by $\mathcal{S}_n$.

Proposition 4. *Let $n \in \mathbb{N}$. Let X be a set with x_i, $y_i \in X$ for all $1 \leq i \leq n$. Let $d \in [0,1]^{X \times X}$ be a pseudometric. Then*

$$\mathcal{D}(d) \left(\sum_{1 \leq i \leq n} \tfrac{1}{n} \delta_{x_i}, \sum_{1 \leq i \leq n} \tfrac{1}{n} \delta_{y_i} \right) = \min \left\{ \sum_{1 \leq i \leq n} \tfrac{1}{n} \cdot d(x_i, y_{\sigma(i)}) \;\middle|\; \sigma \in \mathcal{S}_n \right\}.$$

3 MDPs and SSGs

Next, we formally introduce the two main players in this paper: MDPs and SSGs. Furthermore, we define a pseudometric on the states of an MDP and show that it indeed generalizes probabilistic bisimilarity and, therefore, deserves the name bisimulation pseudometric. Finally, we define a value assignment of the vertices of an SSG.

Definition 5. *A Markov decision process (MDP) is a tuple* $(S, L, \rightarrow, \ell)$ *consisting of*

- *a finite set* S *of states,*
- *a finite set* L *of labels,*
- *a finite transition relation* $\rightarrow \subseteq S \times \mathcal{D}(S)$*, and*
- *a labelling function* $\ell : S \rightarrow L$*.*

For the remainder of this section, we fix an MDP $(S, L, \rightarrow, \ell)$. As we will see below, the bisimulation pseudometric is defined as the least fixed point of the following function.

Definition 6. *The function* $\Delta : [0,1]^{S \times S} \rightarrow [0,1]^{S \times S}$ *is defined as follows. If* $\ell(s_1) \neq \ell(s_2)$ *then* $\Delta(d)(s_1, s_2) = 1$*. Otherwise*

$$\Delta(d)(s_1, s_2) = \mathcal{P}(\mathcal{D}(d))(\{\, \mu_1 \mid s_1 \rightarrow \mu_1 \,\}, \{\, \mu_2 \mid s_2 \rightarrow \mu_2 \,\}).$$

Fu [19] also defines his pseudometric as a least fixed point. He considers sets of convex combinations of transitions, also known as combined transitions, rather than just sets of transitions as we do. For a discussion of the difference between transitions and combined transitions we refer the reader to, for example, [27].

From the facts that $\mathcal{P}$ and $\mathcal{D}$ are monotone, we can conclude that Δ is monotone as well. According to Tarski's fixed point theorem, a monotone function on a complete lattice has a least fixed point. Hence, Δ has a least fixed point, which we denote by δ. This is our bisimulation pseudometric.

From the facts that $\mathcal{P}$ and $\mathcal{D}$ are nonexpansive, we can conclude that Δ is nonexpansive as well. Since Δ is monotone and nonexpansive, we can conclude from [4, Corollary 1] that the closure ordinal of Δ is ω, that is, $\delta = \bigsqcup_{n \in \mathbb{N}} \Delta^n(\mathbf{0})$, where the distance function $\mathbf{0}$ maps every pair of states to zero. The latter characterization of δ allows for inductive proofs. For example, to conclude that δ is a pseudometric, it suffices to prove by induction on n that $\Delta^n(\mathbf{0})$ is a pseudometric.

Next, we introduce the notion of probabilistic bisimilarity. To define this notion, we first show how to lift a relation on states to a relation on probability distributions on states.

Definition 7. *The lifting of a relation* $R \subseteq S \times S$ *is the relation* $\bar{R} \subseteq \mathcal{D}(S) \times \mathcal{D}(S)$ *defined by* $\mu_1 \bar{R} \mu_2$ *if there exists a coupling* $\omega \in \Omega_{\mu_1, \mu_2}$ *such that* $\omega(s_1, s_2) > 0$ *implies* $s_1 \mathrel{R} s_2$ *for all* s_1*,* $s_2 \in S$*.*

Definition 8. $R \subseteq S \times S$ *is a probabilistic bisimulation if* $s_1 \mathrel{R} s_2$ *implies*

- $\ell(s_1) = \ell(s_2)$*,*
- *if* $s_1 \rightarrow \mu_1$ *then there exists* $s_2 \rightarrow \mu_2$ *such that* $\mu_1 \mathrel{\bar{R}} \mu_2$*,*
- *if* $s_2 \rightarrow \mu_2$ *then there exists* $s_1 \rightarrow \mu_1$ *such that* $\mu_1 \mathrel{\bar{R}} \mu_2$*.*

States s_1 *and* s_2 *are probabilistic bisimilar, denoted* $s_1 \sim s_2$*, if* $s_1 \mathrel{R} s_2$ *for some probabilistic bisimulation* R*.*

One can show that distance zero captures probabilistic bisimilarity, that is, for all s_1, $s_2 \in S$, $\delta(s_1, s_2) = 0$ if and only if $s_1 \sim s_2$.

Definition 9. *A simple stochastic game (SSG) is a tuple (V, E, v_s) such that*

- *(V, E) is a finite directed graph,*
- *$v_s \in V$,*
- *V is the disjoint union of the sets $\{v_0, v_1\}$, $V_{\max}$, $V_{\min}$, and V_{avg},*
- *the indegree of v_0 and v_1 is zero, and*
- *the indegree of v is two for all $v \in V \setminus \{v_0, v_1\}$.*

Whenever the start vertex v_s does not play a role, we write (V, E) instead of (V, E, v_s). We denote the predecessors of vertex $v \in V \setminus \{v_0, v_1\}$ by v_ℓ and v_r. For the remainder of this section, we fix an SSG (V, E, v_s). It is a bit of folklore that the value of an SSG, as described in the introduction, can be defined as the least fixed point of the following function (see, for example, [22]).

Definition 10. *The function $\Phi : [0,1]^V \to [0,1]^V$ is defined by*

$$\Phi(f)(v) = \begin{cases} 0 & \text{if } v = v_0 \\ 1 & \text{if } v = v_1 \\ \max\{f(v_\ell), f(v_r)\} & \text{if } v \in V_{\max} \\ \min\{f(v_\ell), f(v_r)\} & \text{if } v \in V_{\min} \\ \frac{1}{2} f(v_\ell) + \frac{1}{2} f(v_r) & \text{if } v \in V_{\mathrm{avg}} \end{cases}$$

It follows immediately from the definitions that Φ is monotone. Again using Tarski's fixed point theorem, we can conclude that Φ has a least fixed point which we denote by ϕ. This function assigns to each vertex of the SSG its value. Hence, $\phi(v_s)$ is the value of the SSG.

One can prove that Φ is nonexpansive. Since Φ is also monotone, we can conclude from [4, Corollary 1] that the closure ordinal of Φ is ω, that is, $\phi = \bigsqcup_{n \in \mathbb{N}} \Phi^n(\mathbf{0})$, where the function $\mathbf{0}$ maps every vertex to zero. Again, this characterization allows for inductive proofs.

4 Computing Bisimulation Pseudometrics is in PPAD

In this section we show that the least fixed point of the functional Δ can be computed in **PPAD**.

Proposition 11. *Δ is piecewise linear.*

Proof. Writing out the definition of Δ explicitly, we have that if $\ell(s_1) = \ell(s_2)$ then

$$\begin{aligned} \Delta(d)(s_1, s_2) = \max \Big\{ & \max_{s_1 \to \mu_1} \min_{s_2 \to \mu_2} \min_{\omega \in \Omega_{\mu_1, \mu_2}} \sum_{u,v \in S} \omega(u,v) d(u,v), \\ & \max_{s_2 \to \mu_2} \min_{s_1 \to \mu_1} \min_{\omega \in \Omega_{\mu_1, \mu_2}} \sum_{u,v \in S} \omega(u,v) d(u,v) \Big\} . \end{aligned} \qquad (1)$$

Observe, moreover, that the innermost minima in (3) are achieved at the vertices of each polytope Ω_{μ_1,μ_2}.

Next we reformat the definition of Δ, essentially by currying. To this end, define the set of *attacker strategies* $\mathbb{AS}$ to comprise all functions

$$\sigma : S \times S \to \mathcal{D}(S) \times S$$

such that either $\sigma(s_1, s_2) = (\mu_1, s_2)$, where $s_1 \to \mu_1$, or $\sigma(s_1, s_2) = (\mu_2, s_1)$, where $s_2 \to \mu_2$. Define the set of *defender strategies* $\mathbb{DS}$ to comprise all functions

$$\sigma : \mathcal{D}(S) \times S \to \mathcal{D}(S) \times \mathcal{D}(S)$$

such that $\sigma(\mu_1, s_2) = (\mu_1, \mu_2)$, where $s_2 \to \mu_2$. Finally, the set of *coupling strategies* $\mathbb{CS}$ comprises all functions $\sigma : \mathcal{D}(S) \times \mathcal{D}(S) \to \mathcal{D}(S \times S)$ such that $\sigma(\mu_1, \mu_2)$ is a vertex of the set of couplings Ω_{μ_1,μ_2}.

Given an attacker strategy $\sigma_1 \in \mathbb{AS}$, defender strategy $\sigma_2 \in \mathbb{DS}$, and coupling strategy $\sigma_3 \in \mathbb{CS}$, we define $\Delta_{\sigma_1,\sigma_2,\sigma_3} : [0,1]^{S\times S} \to [0,1]^{S\times S}$ by $\Delta_{\sigma_1,\sigma_2,\sigma_3}(d)(s_1, s_2) = 1$ if $\ell(s_1) \neq \ell(s_2)$, and otherwise

$$\Delta_{\sigma_1,\sigma_2,\sigma_3}(d)(s_1, s_2) = \sum_{u,v \in S} \omega(u, v) d(u, v)\,,$$

where $\omega = (\sigma_3 \circ \sigma_2 \circ \sigma_1)(s_1, s_2)$.

Currying the definition of Δ in (3) we have that

$$\Delta(d) = \max_{\sigma_1 \in \mathbb{AS}} \min_{\sigma_2 \in \mathbb{DS}} \min_{\sigma_3 \in \mathbb{CS}} \Delta_{\sigma_1,\sigma_2,\sigma_3}\,. \tag{2}$$

But each function $\Delta_{\sigma_1,\sigma_2,\sigma_3}$ is clearly piecewise linear and the class of piecewise linear functions $[0,1]^{S\times S}$ is closed under finite pointwise maxima and minima. It follows that Δ is piecewise linear. □

In fact we have a stronger result than Proposition 11: the functional Δ is *polynomial* piecewise linear in the sense of [16]. This means that Δ's domain can be divided into polyhedral cells, and given $d \in [0,1]^{S\times S}$ we can output in polynomial time (in the representation of d and the underlying MDP) a system of linear inequalities defining the cell $C_d \subseteq [0,1]^{S\times S}$ containing d and a linear function coinciding with $\Delta \restriction_{C_d}$.

Proposition 12. *Δ is polynomial piecewise linear.*

It is shown in [16] that one can compute *a* fixed point of a polynomial piecewise linear map in **PPAD**. To compute specifically the least fixed point δ of Δ we use a similar strategy to that employed for SSGs in [16]. We consider a contractive map $\Delta_\varepsilon(d) = \Delta(\varepsilon \cdot d)$ for some ε sufficiently close to 1. We compute the (unique) fixed point of Δ_ε in **PPAD** and round using continued fractions to obtain δ. This relies on the fact that the unique fixed point of Δ_ε converges to δ as ε converges to 1, as well as the fact that there is a polynomial bound on the bit size of the rational numbers occurring in δ. Thus we obtain the following.

Theorem 13. *The problem of computing δ on a given MDP is in* ***PPAD****.*

5 From SSGs to MDPs

For the remainder of this section, we fix an SSG (V, E). Next, we will construct a corresponding MDP. Before formally defining this MDP, let us provide the intuition behind its construction. For each vertex v and its outgoing edges of the SSG, we include a gadget consisting of states and their outgoing transitions in the MDP. Each gadget contains states v^0 and v^1 such that $\delta(v^0, v^1) = \phi(v)$. Let us informally describe the MDP by giving gadgets for each type of vertex. The gadget for vertex v_0 is as follows.

The gadget for vertex v_1 is shown below.

The gadget for a min vertex is as follows.

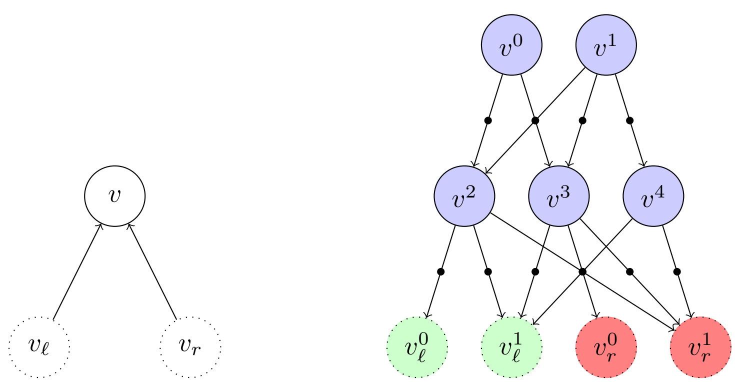

Next, we consider the gadget for a max vertex.

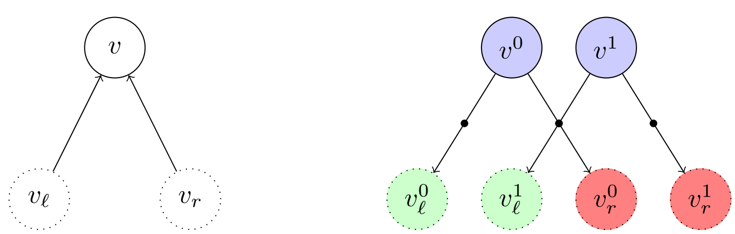

Finally, we present the gadget for an average vertex.

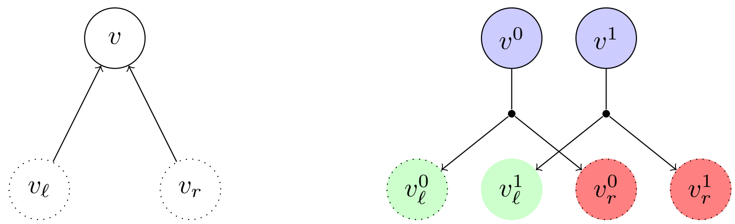

This completes the description of the gadgets. The MDP $\mathcal{M}(V, E)$ is obtained by composing the gadgets for each vertex of the SSG (V, E). The transduction of an SSG to an MDP is done vertex by vertex. To produce the gadget corresponding to each vertex one only needs to store the vertex and its two predecessors, and the states of the gadget. Thus the reduction can be done in deterministic logarithmic space.

Definition 14. *The MDP $\mathcal{M}(V, E)$ is defined by $\mathcal{M}(V, E) = (S, L, \rightarrow, \ell)$ where*

- *the set S of states is defined by*
$$S = \{\, v^0, v^1 \mid v \in V \,\} \cup \{\, v^2, v^3, v^4 \mid V \in V_{\min} \,\},$$
- *the set L of labels is defined by*
$$L = V,$$
- *the transition relation $\rightarrow$ is defined by*
$$\rightarrow = \bigcup \{\, \rightarrow_v \mid v \in V \,\},$$
where
  - $\rightarrow_{v_0} = \emptyset$,
  - $\rightarrow_{v_1} = \{(v_1^0, \delta_{v_1^0})\}$,
  - *for each $v \in V_{\min}$,*
  $\rightarrow_v = \{(v^0, \delta_{v^2}), (v^0, \delta_{v^3}), (v^1, \delta_{v^2}), (v^1, \delta_{v^3}), (v^1, \delta_{v^4}), (v^2, \delta_{v_\ell^0}), (v^2, \delta_{v_\ell^1}),$
  $(v^2, \delta_{v_r^1}), (v^3, \delta_{v_\ell^1}), (v^3, \delta_{v_r^0}), (v^3, \delta_{v_r^1}), (v^4, \delta_{v_\ell^1}), (v^4, \delta_{v_r^1})\}$,
  - *for each $v \in V_{\max}$,*
  $\rightarrow_v = \{(v^0, \delta_{v_\ell^0}), (v^0, \delta_{v_r^0}), (v^1, \delta_{v_\ell^1}), (v^1, \delta_{v_r^1})\}$,
  - *for each $v \in V_{\text{avg}}$,*
  $\rightarrow_v = \{(v^0, \frac{1}{2}\delta_{v_\ell^0} + \frac{1}{2}\delta_{v_r^0}), (v^1, \frac{1}{2}\delta_{v_\ell^1} + \frac{1}{2}\delta_{v_r^1})\}$,
- *the labelling function $\ell : S \rightarrow L$ is defined by*
$$\ell(v^i) = v.$$

In the remainder of this section, $\varPhi$ and ϕ correspond to the SSG (V, E), and $\varDelta$ and δ correspond to the MDP $\mathcal{M}(V, E)$. $\varDelta$ and $\varPhi$ are related as follows.

Lemma 15. *For all $n \in \mathbb{N}$ and $v \in V$,*
$$\varDelta^n(\mathbf{0})(v^0, v^1) \leq \varPhi^n(\mathbf{0})(v) \leq \varDelta^{2n}(\mathbf{0})(v^0, v^1).$$

From the above lemma and the fact that for all $n \in \mathbb{N}$, $\Delta^n(\mathbf{0}) \sqsubseteq \Delta^{n+1}(\mathbf{0})$ and $\Phi^n(\mathbf{0}) \sqsubseteq \Phi^{n+1}(\mathbf{0})$, we can conclude the following result.

Theorem 16. *For all $v \in V$, $\delta(v^0, v^1) = \phi(v)$.*

Hence, computing the value of an SSG can be reduced in deterministic logarithmic space to computing the bisimulation pseudometric on an MDP.

6 From MDPs to SSGs

For the remainder of this section, we fix an MDP $(S, L, \rightarrow, \ell)$. Next, we will construct an SSG that corresponds to the MDP. Before formally defining this SSG, let us provide the intuition behind its construction. For each pair of states s and t there is a max vertex $v_{s,t}$, with $\phi(v_{s,t}) = \delta(s,t)$. For each state s and transition $t \rightarrow \nu$, there is a min vertex $v_{s,\nu}$ with $\phi(v_{s,\nu}) = \min_{s\rightarrow\mu} \mathcal{D}(\delta)(\mu,\nu)$ and for each pair of transitions $s \rightarrow \mu$ and $t \rightarrow \nu$, there is a min vertex $v_{\mu,\nu}$ with $\phi(v_{\mu,\nu}) = \mathcal{D}(\delta)(\mu,\nu)$.

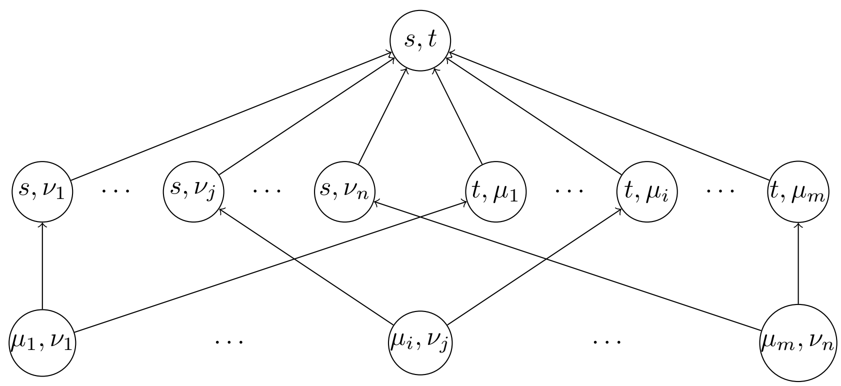

With each each vertex $v_{\nu,\mu}$ we associate the set of couplings $\Omega_{\mu,\nu}$. This set of couplings forms a convex polytope. Let $V_{\mu,\nu}$ be the set of its vertices. For each $\omega \in V_{\mu,\nu}$, we can construct an SSG (V_ω, E_ω) along the lines of the constructions in [2, Section 4.1].

Proposition 17. *For each $\omega \in V_{\mu,\nu}$ there exists an SSG (V_ω, E_ω) with*

- $v_\omega \in V_\omega$,
- $\{ v_{s_1,s_2} \mid (s_1, s_2) \in \mathit{support}(\omega) \} \subseteq V_\omega$, *and*
- $\phi(v_\omega) = \sum_{s_1,s_2 \in S} \omega(s_1,s_2)\phi(v_{s_1,s_2})$

of size polynomial in ω.

Combining the above we arrive at the following construction.

Definition 18. *The SSG $\mathcal{G}(S, L, \rightarrow, \ell)$ is defined by $\mathcal{G}(S, L, \rightarrow, \ell) = (V, E)$ where the set V of vertices is defined by*

$$V = \bigcup \left\{ V_\omega \;\middle|\; \exists s \rightarrow \mu : \exists t \rightarrow \nu : \omega \in V_{\mu,\nu} \right\} \cup V_{\min} \cup V_{\max}$$

with

$$\begin{aligned} V_{\min} &= \{ v_{s,\nu} \mid s \in S \wedge \exists t \in S : t \rightarrow \nu \} \cup \\ &\quad\; \{ v_{\mu,\nu} \mid \exists s \in S : s \rightarrow \mu \wedge \exists t \in S : t \rightarrow \nu \} \\ V_{\max} &= \{ v_{s,t} \mid s, t \in S \} \end{aligned}$$

and the set E of edges is defined by

$$\begin{aligned} E = \bigcup & \left\{ E_\omega \;\middle|\; \exists s \rightarrow \mu : \exists t \rightarrow \nu : \omega \in V_{\mu,\nu} \right\} \cup \\ & \{ (v_{s,t}, v_{t,\mu}) \mid s \rightarrow \mu \wedge t \in S \wedge \ell(s) = \ell(t) \} \cup \\ & \{ (v_{s,t}, v_{s,\nu}) \mid s \in S \wedge t \rightarrow \nu \wedge \ell(s) = \ell(t) \} \cup \\ & \{ (v_{s,t}, v_1) \mid s \in S \wedge t \in S \wedge \ell(s) \neq \ell(t) \} \cup \\ & \{ (v_{s,\nu}, v_{\mu,\nu}) \mid s \rightarrow \mu \wedge \exists t \in S : t \rightarrow \nu \} \cup \\ & \{ (v_{\mu,\nu}, v_\omega) \mid v_{\mu,\nu} \in V_{\min} \wedge \omega \in V_{\mu,\nu} \} \end{aligned}$$

Let δ be the bisimulation pseudometric of the MDP $(S, L, \rightarrow, \ell)$ and let ϕ be the value assignment of the SSG $\mathcal{G}(S, L, \rightarrow, \ell)$. These are related as follows.

Theorem 19. *For all s_1, $s_2 \in S$, $\phi(v_{s_1,s_2}) = \delta(s_1, s_2)$.*

Note that, in general, the size of $V_{\mu,\nu}$ may be exponential in the branching degree of the MDP. To obtain a polynomial time reduction we restrict ourselves to the following class of MDPs.

Definition 20. *Let $n \in \mathbb{N}$. An MDP $(S, L, \rightarrow, \ell)$ has n-bounded support if for all $(s, \mu) \in \rightarrow$, $|support(\mu)| \leq n$.*

If we restrict ourselves to MDPs with n-bounded support for some fixed $n \in \mathbb{N}$, the size of V_{μ_1,μ_2} is polynomial in the size of the MDP and, hence, we obtain a polynomial time reduction.

7 From Linear Assignment Problems to MDPs

Given a cost matrix $c : \mathbb{Q}^{\{1,\ldots,n\} \times \{1,\ldots,n\}}$, the linear assignment problem boils down to computing

$$\min \left\{ \sum_{1 \leq i \leq n} c(i, \sigma(i)) \;\middle|\; \sigma \in \mathcal{S}_n \right\}.$$

Without loss of generality we may assume that the costs are rationals in the unit interval. For the remainder of this section, we fix $c \in (\mathbb{Q} \cap [0,1])^{\{1,\ldots,n\} \times \{1,\ldots,n\}}$.

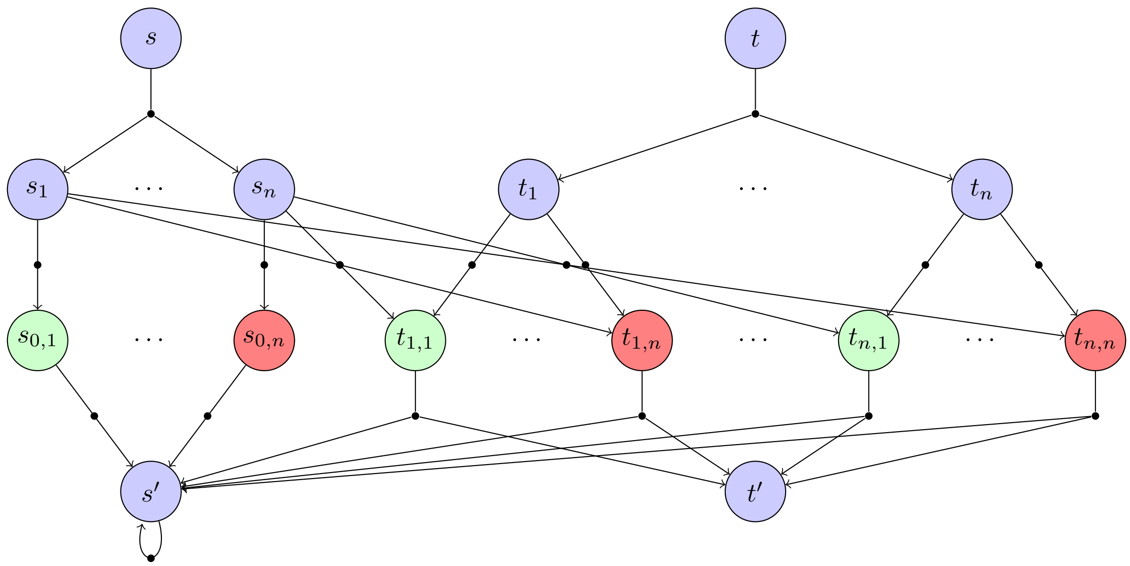

Definition 21. *The MDP* $\mathcal{M}(c)$ *is defined by* $\mathcal{M}(c) = (S, L, \rightarrow, \ell)$ *where*

– *the set* S *is states is defined by*

$$\begin{aligned} S = \{s, t, s', t'\} \cup \\ \{\, s_i \mid 1 \leq i \leq n \,\} \cup \{\, t_i \mid 1 \leq i \leq n \,\} \cup \\ \{\, s_{0,j} \mid 1 \leq j \leq n \,\} \cup \{\, t_{i,j} \mid 1 \leq i, j \leq n \,\} \end{aligned}$$

– *the set* L *of labels is defined by*

$$L = \{l\} \cup \{l_i \mid 1 \leq i \leq n\},$$

– *the transition relation* $\rightarrow$ *is defined by*

$$\begin{aligned} \rightarrow = & \left\{ \left(s, \sum_{1 \leq i \leq n} \tfrac{1}{n} \delta_{s_i} \right), \left(t, \sum_{1 \leq i \leq n} \tfrac{1}{n} \delta_{t_i} \right) \right\} \cup \\ & \{\, (s_j, \delta_{s_{0,j}}) \mid 1 \leq j \leq n \,\} \cup \{\, (s_h, \delta_{t_{i,j}}) \mid 1 \leq h, i, j \leq n \wedge h \neq j \,\} \cup \\ & \{\, (t_i, \delta_{t_{i,j}}) \mid 1 \leq i, j \leq n \,\} \cup \\ & \{\, (s_{0,j}, \delta_{s'}) \mid 1 \leq j \leq n \,\} \cup \\ & \{\, (t_{i,j}, c(j,i)\delta_{t'} + (1 - c(j,i))\delta_{s'}) \mid 1 \leq i, j \leq n \,\} \cup \\ & \{\, (s', \delta_{s'}) \mid 1 \leq i \leq n \,\}, \end{aligned}$$

– *the labelling function* $\ell : S \rightarrow L$ *is defined by*

$$\ell(s'') = \begin{cases} l_j & \text{if } s'' = s_{0,j} \text{ or } s'' = t_{i,j} \\ l & \text{otherwise} \end{cases}$$

Let δ be the bisimulation pseudometric of the MDP $\mathcal{M}(c)$. Using Proposition 4, we can show the following.

Theorem 22. $\delta(s,t) = \frac{1}{n} \min\{ \sum_{1 \leq i \leq n} c(i, \sigma(i)) \mid \sigma \in \mathcal{S}_n \}$.

Hence, the linear assignment problem can be reduced in deterministic logarithmic space to computing a bisimulation pseudometric.

References

1. Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network flows – theory, algorithms and applications.* Prentice Hall, 1993.
2. Manuel Blum, Rachel Rue, and Ke Yang. On the complexity of max/min/avrg circuits. Technical Report CMU-CS-02-110, Carnegie Mellon University, Pittsburgh, PA, USA, March 2002.
3. Franck van Breugel. An introduction to metric semantics: operational and denotational models for programming and specification languages. *Theoretical Computer Science*, 258(1/2):1–98, May 2001.
4. Franck van Breugel. On behavioural pseudometrics and closure ordinals. *Information Processing Letters*, 112(19):715–718, October 2012.
5. Franck van Breugel, Claudio Hermida, Michael Makkai, and James Worrell. Recursively defined metric spaces without contraction. *Theoretical Computer Science*, 380(1/2):143–163, June 2007.
6. Franck van Breugel and James Worrell. From simple stochastic games to bisimulation pseudometrics on Markov decision processes. Available at `www.cse.yorku.ca/~franck/research/drafts`, April 2013.
7. Rainer Burkard, Mauro Dell' Amico, and Silvano Martello. *Assignment Problems.* SIAM, Philadelphia, PA, USA, 2012.
8. Krishnendu Chatterjee, Luca de Alfaro, Rupak Majumdar, and Vishwanath Raman. Algorithms for game metrics. In *Proceedings of the IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science*, volume 2 of *LIPIcs*, pages 107–118, Bangalore, India, December 2008. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik.
9. Di Chen, Franck van Breugel, and James Worrell. On the complexity of computing probabilistic bisimilarity. In *Proceedings of the 15th International Conference on Foundations of Software Science and Computational Structures*, volume 7213 of *Lecture Notes in Computer Science*, pages 437–451, Tallinn, Estonia, March 2012. Springer-Verlag.
10. Vašek Chvátal. *Linear Programming.* W.H. Freeman and Company, New York, NY, USA, 1983.
11. Anne Condon. The complexity of stochastic games. *Information and Computation*, 96(2):203–224, February 1992.
12. Constantinos Daskalakis and Christos Papadimitriou. Continuous local search. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 790–804, San Francisco, CA, USA, January 2011. SIAM.
13. Luca de Alfaro, Rupak Majumdar, Vishwanath Raman, and Mariëlle Stoelinga. Game relations and metrics. In *Proceedings of the 22nd IEEE Symposium on Logic in Computer Science*, pages 99–108, Wroclaw, Poland, July 2007. IEEE.

14. Josee Desharnais, Vineet Gupta, Radha Jagadeesan, and Prakash Panangaden. Metrics for labeled Markov systems. In *Proceeding of the 10th International Conference on Concurrency Theory*, volume 1664 of *Lecture Notes in Computer Science*, pages 258–273, Eindhoven, The Netherlands, August 1999. Springer-Verlag.
15. Gerald A. Edgar. *Integral, Probability, and Fractal Measures.* Springer-Verlag, New York, NY, USA, 1998.
16. Kousha Etessami and Mihalis Yannakakis. On the complexity of Nash equilibria and other fixed points (extended abstract). In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*, pages 113–123, Providence, RI, USA, October 2007. IEEE.
17. Norm Ferns, Prakash Panangaden, and Doina Precup. Metrics for finite Markov decision processes. In *Proceedings of the 20th Conference in Uncertainty in Artificial Intelligence*, pages 162–169, Banff, Canada, July 2004. AUAI Press.
18. Norm Ferns, Prakash Panangaden, and Doina Precup. Metrics for finite Markov decision processes. In *Proceedings of the 19th National Conference on Artificial Intelligence*, pages 950–951, San Jose, CA, USA, July 2004. AAAI Press.
19. Hongfei Fu. Computing game metrics on Markov decision processes. In *Proceedings of the 39th International Colloquium on Automata, Languages, and Programming*, volume 7392 of *Lecture Notes in Computer Science*, pages 227–238, Warwick, UK, July 2012. Springer-Verlag.
20. Hongfei Fu. Personal communication, January 2013.
21. Alessandro Giacalone, Chi-Chang Jou, and Scott A. Smolka. Algebraic reasoning for probabilistic concurrent systems. In *Proceedings of the IFIP WG 2.2/2.3 Working Conference on Programming Concepts and Methods*, pages 443–458, Sea of Gallilee, April 1990. North-Holland.
22. Brendan Juba. On the hardness of simple stochastic games. Master's thesis, Carnegie Mellon University, Pittsburgh, PA, USA, May 2005.
23. Kim Guldstrand Larsen and Arne Skou. Bisimulation through probabilistic testing. In *Proceedings of the 16th Annual ACM Symposium on Principles of Programming Languages*, pages 344–352, Austin, TX, USA, January 1989. ACM.
24. William Lawvere. Metric spaces, generalized logic, and closed categories. *Rendiconti del Seminario Matematico e Fisico di Milano*, 43:135–166, 1973.
25. Alexander A. Razborov. Lower bounds on monotone complexity of the logical permanent. *Mathematical Notes of the Academy of Sciences of the USSR*, 37(6):485–493, June 1985.
26. Uwe Schwiegelshohn and Lothar Thiele. A systolic array for the assignment problem. *IEEE Transactions on Computers*, 37(11):1422–1425, November 1988.
27. Roberto Segala and Nancy Lynch. Probabilistic simulations for probabilistic processes. In *Proceedings of the 5th International Conference on Concurrency Theory*, volume 836 of *Lecture Notes in Computer Science*, pages 481–496, Uppsala, Sweden, August 1994. Springer-Verlag.
28. Colin Stirling. Bisimulation, modal logic and model checking games. *Logic Journal of the IGPL*, 7(1):103–124, 1999.
29. Mihalis Yannakakis. The analysis of local search problems and their heuristics. In *Proceedings of the 7th Annual Symposium on Theoretical Aspects of Computer Science*, volume 415 of *Lecture Notes in Computer Science*, pages 298–311, Rouen, France, February 1990. Springer-Verlag.
30. Mihalis Yannakakis. Compuational aspects of equilibria. In *Proceedings of the 2nd International Symposium on Algorithmic game Theory*, volume 5814 of *Lecture Notes in Computer Science*, pages 2–13, Paphos, Cyprus, October 2009. Springer-Verlag.

A Proofs of Section 2

Proposition 23. *Let X be a set and $n \in \omega$. $\langle [0,1]^{X^n}, \sqsubseteq \rangle$ is a complete lattice.*

Proof. Obviously, $\sqsubseteq$ is a partial order. The bottom element assigns to every tuple zero and the top element assigns to every tuple one. Let F be a nonempty subset of $[0,1]^{X^n}$. The meet of F is defined by $(\bigsqcap F)(x_1, \ldots, x_n) = \inf_{f \in F} f(x_1, \ldots, x_n)$ and the join of F is defined by $(\bigsqcup F)(x_1, \ldots, x_n) = \sup_{f \in F} f(x_1, \ldots, x_n)$. □

Proposition 24. *Let X be a set. For all $d \in [0,1]^{X \times X}$, A_1, $A_2 \in \mathcal{P}(X)$, if $A_1 \subseteq A_2$ then $\mathcal{P}(d)(A_1, A_2) = \max_{x_2 \in A_2 \setminus A_1} \min_{x_1 \in A_1} d(x_1, x_2)$.*

Proof.

$$\begin{aligned}
\mathcal{P}(d)(A_1, A_2) &= \max\left\{\max_{x_1 \in A} \min_{x_2 \in A_2} d(x_1, x_2), \max_{x_2 \in A_2} \min_{x_1 \in A_1} d(x_1, x_2)\right\} \\
&= \max_{x_2 \in A_2} \inf_{x_1 \in A_1} d(x_1, x_2) \\
&= \max_{x_2 \in A_2 \setminus A_1} \min_{x_1 \in A_1} d(x_1, x_2).
\end{aligned}$$

□

Proposition 25. *$\mathcal{P}$ is nonexpansive.*

Proof. Let d_1, $d_2 \in [0,1]^{X \times X}$ and let A_1, $A_2 \in \mathcal{P}(X)$. Then

$$\begin{aligned}
&\mathcal{P}(d_1)(A_1, A_2) - \mathcal{P}(d_2)(A_1, A_2) \\
&= \max\left\{\max_{x_1 \in A_1} \min_{x_2 \in A_2} d_1(x_1, x_2), \max_{x_2 \in A_2} \min_{x_1 \in A_1} d_1(x_2, x_1)\right\} - \\
&\quad \max\left\{\max_{x_1 \in A_1} \min_{x_2 \in A_2} d_2(x_1, x_2), \max_{x_2 \in A_2} \min_{x_1 \in A_1} d_2(x_2, x_1)\right\} \\
&\leq \max\left\{\max_{x_1 \in A_1} \min_{x_2 \in A_2} d_1(x_1, x_2) - \max_{x_1 \in A_1} \min_{x_2 \in A_2} d_2(x_1, x_2),\right. \\
&\qquad \left.\max_{x_2 \in A_2} \min_{x_1 \in A_1} d_1(x_2, x_1) - \max_{x_2 \in A_2} \min_{x_1 \in A_1} d_2(x_2, x_1)\right\}
\end{aligned}$$

Furthermore, there exist x_1', $x_2' \in X$,

$$\begin{aligned}
&\max_{x_1 \in A_1} \min_{x_2 \in A_2} d_1(x_1, x_2) - \max_{x_1 \in A_1} \min_{x_2 \in A_2} d_2(x_1, x_2) \\
&= \max_{x_1 \in A_1} \min_{x_2 \in A_2} d_1(x_1, x_2) - \min_{x_2 \in A_2} d_2(x_1', x_2) \\
&\leq \min_{x_2 \in A_2} d_1(x_1', x_2) - \min_{x_2 \in A_2} d_2(x_1', x_2) \\
&= d_1(x_1', x_2') - \min_{x_2 \in A_2} d_2(x_1', x_2) \\
&\leq d_1(x_1', x_2') - d_2(x_1', x_2') \\
&\leq \|d_1 - d_2\|.
\end{aligned}$$

By symmetry, we can conclude that

$$\mathcal{P}(d_1)(A_1, A_2) - \mathcal{P}(d_2)(A_1, A_2) \leq \|d_1 - d_2\|.$$

$\square$

Proposition 26. *Let X be a set. For all $d \in [0,1]^{X \times X}$, x_1, x_2, y_1, $y_2 \in X$, if d is a pseudometric and*

$$\begin{aligned} d(x_1, x_2) &= 1 \\ d(x_1, y_2) &= 1 \\ d(y_1, x_2) &= 1 \\ d(y_1, y_2) &= 1 \end{aligned}$$

then $\mathcal{D}(d)(\frac{1}{2}\delta_{x_1} + \frac{1}{2}\delta_{x_2}, \frac{1}{2}\delta_{y_1} + \frac{1}{2}\delta_{y_2}) = \frac{1}{2}(d(x_1, y_1) + d_X(x_2, y_2))$.

Proof. We have that

$$\begin{aligned} &\mathcal{D}(d)(\tfrac{1}{2}\delta_{x_1} + \tfrac{1}{2}\delta_{x_2}, \tfrac{1}{2}\delta_{y_1} + \tfrac{1}{2}\delta_{y_2}) \\ &= \sup_{f \in X \nrightarrow [0,1]} (\tfrac{1}{2}f(x_1) + \tfrac{1}{2}f(x_2)) - (\tfrac{1}{2}f(y_1) + \tfrac{1}{2}f(y_2)) \\ &= \tfrac{1}{2} \sup_{f \in X \nrightarrow [0,1]} (f(x_1) - f(y_1)) + (f(x_2) - f(y_2)) \\ &= \tfrac{1}{2}(d(x_1, y_1) + d(x_2, y_2)), \end{aligned}$$

since $f(x_1) - f(y_1) \leq d(x_1, y_1)$ and $f(x_2) - f(y_2) \leq d(x_2, y_2)$ for all $f \in X \nrightarrow [0,1]$, and function $g : \{x_1, x_2, y_1, y_2\} \to [0,1]$ defined by

$$g(z) = \begin{cases} d(x_1, y_1) & \text{if } z = x_1 \\ d(x_2, y_2) & \text{if } z = x_2 \\ 0 & \text{otherwise} \end{cases}$$

is nonexpansive and $g(x_1) - g(y_1) = d(x_1, y_1)$ and $g(x_2) - g(y_2) = d(x_2, y_2)$. Note that every nonexpansive function $\{x_1, x_2, y_1, y_2\}$ to $[0,1]$ can be extended to a nonexpansive function from X to $[0,1]$ (see, for example, [24, page 162] for such a McShane-Whitney extension theorem). $\square$

Proposition 27. *Let X be a set. For all $q \in [0,1]$, $d \in [0,1]^{X \times X}$, x, $y \in X$, if d is a pseudometric then $\mathcal{D}(d)(\delta_x, (1-q)\delta_x + q\delta_y) = qd(x, y)$.*

Proof. We have that

$$\begin{aligned} &\mathcal{D}(d)(\delta_x, (1-q)\delta_x + q\delta_y) \\ &= \sup_{f \in X \nrightarrow [0,1]} f(x) - ((1-q)f(x) + qf(y)) \\ &= q \sup_{f \in X \nrightarrow [0,1]} f(x) - f(y) \\ &= qd(x, y), \end{aligned}$$

since $f(x) - f(y) \leq d(x, y)$ for all $f \in X \nrightarrow [0,1]$, and function $g : \{x, y\} \to [0,1]$, defined by $g(x) = d(x, y)$ and $g(y) = 0$, is nonexpansive and $g(x) - g(y) = d(x, y)$ (again using a McShane-Whitney extension theorem). $\square$

Proposition 4. *Let $m \in \omega$. Let X be a set with x_i, $y_i \in X$ for all $1 \leq i \leq m$. Let $d \in [0,1]^{X \times X}$ be a pseudometric. Then*

$$\mathcal{D}(d)\left(\sum_{1 \leq i \leq m} \tfrac{1}{m}\delta_{x_i}, \sum_{1 \leq i \leq m} \tfrac{1}{m}\delta_{y_i}\right) = \min\left\{ \sum_{1 \leq i \leq m} \tfrac{1}{m} \cdot d(x_i, y_{\sigma(i)}) \;\middle|\; \sigma \in S_m \right\}.$$

Proof. By definition,

$$\begin{aligned}
&\mathcal{D}(d)\left(\sum_{1 \leq i \leq m} \tfrac{1}{m}\delta_{x_i}, \sum_{1 \leq i \leq m} \tfrac{1}{m}\delta_{y_i}\right) \\
&= \sup\left\{ \sum_{1 \leq i \leq m} \tfrac{1}{m} f(x_i) - \sum_{1 \leq i \leq m} \tfrac{1}{m} f(y_i) \;\middle|\; f \in X \nrightarrow [0,1] \right\} \\
&= \tfrac{1}{m} \sup\left\{ \sum_{1 \leq i \leq m} f(x_i) - f(y_i) \;\middle|\; f \in X \nrightarrow [0,1] \right\}
\end{aligned}$$

By the above mentioned McShane-Whitney extension theorem, every nonexpansive function from $\{x_i, y_i \mid 1 \leq i \leq m\}$ to $[0,1]$ can be extended to a nonexpansive function from X to $[0,1]$ and, hence, we can restrict our attention to the former. The above is equal to $\frac{1}{m}$ times the value of the following linear programming problem. In the problem we use the variables f_i and g_i for the values of $f(x_i)$ and $f(y_i)$, respectively.

$$\begin{array}{lll}
\text{maximize} & \sum_{1 \leq i \leq m} f_i - g_i & \\
\text{such that} & f_i - f_j \leq d(x_i, x_j) & 1 \leq i, j \leq m \\
& g_i - g_j \leq d(y_i, y_j) & 1 \leq i, j \leq m \\
& f_i - g_j \leq d(x_i, y_j) & 1 \leq i, j \leq m \\
& g_i - f_j \leq d(y_i, x_j) & 1 \leq i, j \leq m \\
& f_i \geq 0 & 1 \leq i \leq m \\
& g_i \geq 0 & 1 \leq i \leq m
\end{array}$$

If all f_i's and g_i's are zero then the constraints are satisfied and, hence, the linear programming problem has a feasible origin. Since the distance function d is bounded by one, all f_i's and g_i's are bounded by one. Therefore, $\sum_{1 \leq i \leq m} f_i - g_i$ is bounded. According to the fundamental theorem of linear programming (see, for example, [10, Theorem 3.4]), the linear programming problem has an optimal solution.

Next, we show that the constraints $f_i - f_j \leq d(x_i, x_j)$ and $g_i - g_j \leq d(y_i, y_j)$ are redundant. For example,

$$\begin{aligned}
f_i - f_j &= (f_i - g_j) - (f_j - g_j) \\
&\leq d(x_i, y_j) - d(x_j, y_j) \\
&\leq d(x_i, x_j) \quad [d \text{ satisfies the triangle inequality}]
\end{aligned}$$

The dual of this simplified linear programming problem is

$$\begin{array}{lll} \text{minimize} & \displaystyle\sum_{1\leq i,j\leq m} d(x_i,y_j)v_{ij} & \\ \text{such that} & \displaystyle\sum_{1\leq j\leq m} v_{ij} \leq 1 & 1\leq i\leq m \\ & \displaystyle\sum_{1\leq i\leq m} v_{ij} \geq 1 & 1\leq j\leq m \\ & v_{ij}\geq 0 & 1\leq i,j\leq m \end{array}$$

According to the duality theorem of linear programming (see, for example, [10, Theorem 5.1]), this dual problem also has an optimal solution with the same optimal value as the original problem.

From the above constraints we can conclude that

$$\sum_{1\leq i,j\leq m} v_{ij} \leq m \text{ and } \sum_{1\leq i,j\leq m} v_{ij} \geq m.$$

Hence, we can replace the inequalities in those constraints by equalities resulting in the following.

$$\begin{array}{lll} \text{minimize} & \displaystyle\sum_{1\leq i,j\leq m} d(x_i,y_j)v_{ij} & \\ \text{such that} & \displaystyle\sum_{1\leq j\leq m} v_{ij} = 1 & 1\leq i\leq m \\ & \displaystyle\sum_{1\leq i\leq m} v_{ij} = 1 & 1\leq j\leq m \\ & v_{ij}\geq 0 & 1\leq i,j\leq m \end{array}$$

The value of the above problem does not change if we replace the constraints $v_{ij}\geq 0$ with $v_{ij}\in\{0,1\}$ (see, for example, [7, Chapter 4]). This modified linear programming problem is known as the linear assignment problem. According to, for example, [7, Chapter 4], its value is $\min\{\sum_{1\leq i\leq m} d(x_i,y_{\sigma(i)}) \mid \sigma\in S_m\}$. □

B Proofs of Section 3

Proposition 28. Δ *is monotone.*

Proof. Let d_1, $d_2\in[0,1]^{S\times S}$ with $d_1\sqsubseteq d_2$. Let s_1, $s_2\in S$. We distinguish two cases. If $\ell(s_1)\neq\ell(s_2)$ then

$$\Delta(d_1)(s_1,s_2) = 1 = \Delta(d_2)(s_1,s_2).$$

Otherwise,

$$\begin{aligned} &\Delta(d_1)(s_1,s_2)\\ &= \mathcal{P}(\mathcal{D}(d_1))(\{\,\mu_1 \mid s_1\to\mu_1\,\},\{\,\mu_2\mid s_2\to\mu_2\,\})\\ &\leq \mathcal{P}(\mathcal{D}(d_2))(\{\,\mu_1 \mid s_1\to\mu_1\,\},\{\,\mu_2\mid s_2\to\mu_2\,\}) \quad \text{[Proposition \textbf{??} and \textbf{??}]}\\ &= \Delta(d_2)(s_1,s_2). \end{aligned}$$

□

Proposition 29. *Δ is nonexpansive.*

Proof. Let d_1, $d_2 \in [0,1]^{S \times S}$. Let s_1, $s_2 \in S$. We distinguish two cases. If $\ell(s_1) \neq \ell(s_2)$ then

$$|\Delta(d_1)(s_1, s_2) - \Delta(d_2)(s_1, s_2)| = |1 - 1| = 0 \leq \|d_1 - d_2\|.$$

Otherwise,

$$\begin{aligned}
&|\Delta(d_1)(s_1, s_2) - \Delta(d_2)(s_1, s_2)| \\
&= |\mathcal{P}(\mathcal{D}(d_1))(\{\, \mu_1 \mid s_1 \to \mu_1 \,\}, \{\, \mu_2 \mid s_2 \to \mu_2 \,\}) - \mathcal{P}(\mathcal{D}(d_1))(\{\, \mu_1 \mid s_1 \to \mu_1 \,\}, \{\, \mu_2 \mid s_2 \to \mu_2 \,\})| \\
&\leq \epsilon \|d_1 - d_2\| \quad \text{[Proposition 25 and \textbf{??}]}
\end{aligned}$$

□

Proposition 30. *δ is a pseudometric.*

Proof. Obviously, $\mathbf{0}$ is a pseudometric. Using Proposition **??**, we can show that for all $n \in \omega$, $\Delta^n(\mathbf{0})$ is a pseudometric by induction on n. By Proposition **??**, $\delta = \bigsqcup \{\, \Delta^n(\mathbf{0}) \mid n \in \omega \,\}$. From this, we easily derive that δ is a pseudometric as well. □

Proposition 31. *Φ is nonexpansive.*

Proof. Let f_1, $f_2 \in [0,1]^V$. Let $v \in V$. distinguish five cases.

- Let $v = v_0$. Then
$$|\Phi(f_1)(v) - \Phi(f_2)(v)| = |0 - 0| = 0 \leq \|f_1 - f_2\|.$$
- Let $v = v_1$. Then
$$|\Phi(f_1)(v) - \Phi(f_2)(v)| = |1 - 1| = 0 \leq \|f_1 - f_2\|.$$
- Let v be a min vertex.
$$\begin{aligned}
|\Phi(f_1)(v) - \Phi(f_2)(v)| &= |\min\{f_1(v_\ell), f_1(v_r)\} - \min\{f_2(v_\ell), f_2(v_r)\}| \\
&\leq \max\{|f_1(v_\ell) - f_2(v_\ell)|, |f_1(v_r) - f_2(v_r)|\} \\
&\leq \|f_1 - f_2\|.
\end{aligned}$$
- Let v be a max vertex.
$$\begin{aligned}
|\Phi(f_1)(v) - \Phi(f_2)(v)| &= |\max\{f_1(v_\ell), f_1(v_r)\} - \max\{f_2(v_\ell), f_2(v_r)\}| \\
&\leq \max\{|f_1(v_\ell) - f_2(v_\ell)|, |f_1(v_r) - f_2(v_r)|\} \\
&\leq \|f_1 - f_2\|.
\end{aligned}$$
- Let v be an avg vertex.
$$\begin{aligned}
|\Phi(f_1)(v) - \Phi(f_2)(v)| &= |(\tfrac{1}{2} f_1(v_\ell) + \tfrac{1}{2} f_1(v_r)) - (\tfrac{1}{2} f_2(v_\ell) + \tfrac{1}{2} f_2(v_r))| \\
&\leq |\tfrac{1}{2}(f_1(v_\ell) - f_2(v_\ell)) + \tfrac{1}{2}(f_1(v_r) - f_2(v_r))| \\
&\leq \|f_1 - f_2\|.
\end{aligned}$$

□

C Proofs of Section 4

Proposition 11. Δ *is piecewise linear.*

Proof. Writing out the definition of Δ explicitly, we have that if $\ell(s_1) = \ell(s_2)$ then

$$\Delta(d)(s_1, s_2) = \max\Big\{ \max_{s_1 \to \mu_1} \min_{s_2 \to \mu_2} \min_{\omega \in \Omega_{\mu_1,\mu_2}} \sum_{u,v \in S} d(u,v) \cdot \omega(u,v), \\ \max_{s_2 \to \mu_2} \min_{s_1 \to \mu_1} \min_{\omega \in \Omega_{\mu_1,\mu_2}} \sum_{u,v \in S} d(u,v) \cdot \omega(u,v) \Big\}. \tag{3}$$

Observe, moreover, that the innermost minimum in (3) are achieved at the vertices of each polytope Ω_{μ_1,μ_2}.

Next we reformat the definition of Δ, essentially by currying. To this end, define the set of *attacker strategies* $\mathbb{AS}$ to comprise all functions

$$\sigma : S \times S \to \mathcal{D}(S) \times S$$

such that either $\sigma(s_1, s_2) = (\mu_1, s_2)$, where $s_1 \to \mu_1$, or $\sigma(s_1, s_2) = (\mu_2, s_1)$, where $s_2 \to \mu_2$. Define the set of *defender strategies* $\mathbb{DS}$ to comprise all functions

$$\sigma : \mathcal{D}(S) \times S \to \mathcal{D}(S) \times \mathcal{D}(S)$$

such that $\sigma(\mu_1, s_2) = (\mu_1, \mu_2)$, where $s_2 \to \mu_2$. Finally, the set of *coupling strategies* $\mathbb{CS}$ comprises all functions $\sigma : \mathcal{D}(S) \times \mathcal{D}(S) \to \mathcal{D}(S \times S)$ such that $\sigma(\mu_1, \mu_2)$ is a vertex of the set of couplings Ω_{μ_1,μ_2}.

Given an attacker strategy $\sigma_1 \in \mathbb{AS}$, defender strategy $\sigma_2 \in \mathbb{DS}$, and coupling strategy $\sigma_3 \in \mathbb{CS}$, we define $\Delta_{\sigma_1,\sigma_2,\sigma_3} : [0,1]^{S \times S} \to [0,1]^{S \times S}$ by $\Delta_{\sigma_1,\sigma_2,\sigma_3}(d)(s_1, s_2) = 1$ if $\ell(s_1) \neq \ell(s_2)$, and otherwise

$$\Delta_{\sigma_1,\sigma_2,\sigma_3}(d)(s_1, s_2) = \sum_{u,v \in S} d(u,v) \cdot \omega(u,v)\,,$$

where $\omega = (\sigma_3 \circ \sigma_2 \circ \sigma_1)(s_1, s_2)$.

Currying the definition of Δ in (3) we have that

$$\Delta(d) = \max_{\sigma_1 \in \mathbb{AS}} \min_{\sigma_2 \in \mathbb{DS}} \min_{\sigma_3 \in \mathbb{CS}} \Delta_{\sigma_1,\sigma_2,\sigma_3}\,. \tag{4}$$

But each function $\Delta_{\sigma_1,\sigma_2,\sigma_3}$ is clearly piecewise linear and the class of piecewise linear functions $[0,1]^{S \times S}$ is closed under finite pointwise maxima and minima. It follows that Δ is piecewise linear. □

Proposition 12. Δ *is polynomial piecewise linear.*

Proof. Suppose we are given a pseudometric $d \in [0,1]^{S \times S}$. Given distributions $\mu_1, \mu_2 \in \mathcal{D}(S)$, using the Network Simplex Algorithm [1] we can compute in polynomial time (in the representation of d and the MDP $\mathcal{M}$) a coupling $\omega^* \in \Omega_{\mu_1,\mu_2}$ that minimises

$$\sum_{u,v \in S} d(u,v) \cdot \omega(u,v)\,. \tag{5}$$

Without loss of generality we can assume that ω^* is a vertex of Ω_{μ_1,μ_2}. Let $\omega_1, \ldots, \omega_k$ be the adjacent vertices of ω^* in Ω_{μ_1,μ_2} (also computable in polynomial time). Define

$$C_d^{\mu_1,\mu_2} := \left\{ d' \in [0,1]^S : \sum_{u,v \in S} (\omega^*(u,v) - \omega_i(u,v)) \cdot d'(u,v) \geq 0, i = 1, \ldots, k \right\}.$$

By convexity of Ω_{μ_1,μ_2}, for all $d' \in C_d$ we have that the linear function (5) is maximised on Ω_{μ_1,μ_2} at the vertex ω^*.

Write

$$C_d = \bigcap_{\substack{s_1 \in S \\ s_2 \in S}} \bigcap_{\substack{s_1 \to \mu_1 \\ s_2 \to \mu_2}} C_d^{\mu_1,\mu_2}\,.$$

Then C_d is a polyhedral cell containing d, computable in polynomial time, on which Δ restricts to a linear functional (also computable in polynomial time).

D Proofs of Section 5

Proposition 32. *For all $n \in \omega$,*

$$\Delta^n(\mathbf{0}) \sqsubseteq \Delta^{n+1}(\mathbf{0}).$$

Proof. We prove this proposition by induction on n. Obviously, it holds in the case that $n = 0$. Let $n{>}0$. By induction, $\Delta^{n-1}(\mathbf{0}) \sqsubseteq \Delta^n(\mathbf{0})$. Since Δ is monotone (Proposition 28), we can conclude that $\Delta^n(\mathbf{0}) \sqsubseteq \Delta^{n+1}(\mathbf{0})$. □

Lemma 15. *For all $n \in \omega$ and $v \in V$,*

$$\Delta^n(\mathbf{0})(v^0, v^1) \leq \Phi^n(\mathbf{0})(v) \leq \Delta^{2n}(\mathbf{0})(v^0, v^1).$$

Proof. We prove this lemma by induction on n. The result obviously holds for $n = 0$. Let $n > 0$. We distinguish the following cases.

– Let $v = v_0$. Then

$$\begin{aligned}
\Delta^n(\mathbf{0})(v_0^0, v_0^1) &= \mathcal{P}(\mathcal{D}(\Delta^{n-1}(\mathbf{0})))(\emptyset, \emptyset) \\
&= 0 \\
&= \Phi^n(\mathbf{0})(v_0) \\
&= 0 \\
&= \mathcal{P}(\mathcal{D}(\Delta^{2n-1}(\mathbf{0})))(\emptyset, \emptyset) \\
&= \Delta^{2n}(\mathbf{0})(v_0^0, v_0^1).
\end{aligned}$$

Let $v = v_1$. Then

$$\begin{aligned}
\Delta^n(\mathbf{0})(v_1^0, v_1^1) &= \mathcal{P}(\mathcal{D}(\Delta^{n-1}(\mathbf{0})))(\{\delta_{v_1^0}\}, \emptyset) \\
&= 1 \\
&= \Phi^n(\mathbf{0})(v_1) \\
&= 1 \\
&= \mathcal{P}(\mathcal{D}(\Delta^{2n-1}(\mathbf{0})))(\{\delta_{v_1^0}\}, \emptyset) \\
&= \Delta^{2n}(\mathbf{0})(v_1^0, v_1^1).
\end{aligned}$$

– Let v be a min vertex. Then

$$\begin{aligned}
&\Delta^n(\mathbf{0})(v^0, v^1) \\
&= \mathcal{P}(\mathcal{D}(\Delta^{n-1}(\mathbf{0})))(\{\delta_{v^2}, \delta_{v^3}\}, \{\delta_{v^2}, \delta_{v^3}, \delta_{v^4}\}) \\
&= \min\{\mathcal{D}(\Delta^{n-1}(\mathbf{0}))(\delta_{v^2}, \delta_{v^4}), \mathcal{D}(\Delta^{n-1}(\mathbf{0}))(\delta_{v^3}, \delta_{v^4})\} \quad \text{[Proposition 24]} \\
&= \min\{\Delta^{n-1}(\mathbf{0})(v^2, v^4), \Delta^{n-1}(\mathbf{0})(v^3, v^4)\} \quad \text{[Proposition \textbf{??}]}
\end{aligned}$$

Hence, if $n = 1$ then

$$\begin{aligned}
\Delta^n(\mathbf{0})(v^0, v^1) &= 0 \\
&\leq \Phi^n(\mathbf{0})(v) \\
&= \min\{\Phi^{n-1}(\mathbf{0})(v_\ell), \Phi^{n-1}(\mathbf{0})(v_r)\} \\
&= 0 \\
&\leq \Delta^{2n}(\mathbf{0})(v^0, v^1).
\end{aligned}$$

Let $n > 1$.

$$\begin{aligned}
&\Delta^{n-1}(\mathbf{0})(v^2, v^4) \\
&= \mathcal{P}(\mathcal{D}(\Delta^{n-2}(\mathbf{0})))(\{\delta_{v_\ell^0}, \delta_{v_\ell^1}, \delta_{v_r^1}\}, \{\delta_{v_\ell^1}, \delta_{v_r^1}\}) \\
&= \min\{\mathcal{D}(\Delta^{n-2}(\mathbf{0}))(\delta_{v_\ell^0}, \delta_{v_\ell^1}), \mathcal{D}(\Delta^{n-2}(\mathbf{0}))(\delta_{v_\ell^0}, \delta_{v_r^1})\} \quad \text{[Proposition 24]} \\
&= \min\{\Delta^{n-2}(\mathbf{0})(v_\ell^0, v_\ell^1), \Delta^{n-2}(\mathbf{0})(v_\ell^0, v_r^1)\} \quad \text{[Proposition \textbf{??}]}
\end{aligned}$$

and

$$\begin{aligned}
&\Delta^{n-1}(\mathbf{0})(v^3, v^4) \\
&= \mathcal{P}(\mathcal{D}(\Delta^{n-2}(\mathbf{0})))(\{\delta_{v_\ell^1}, \delta_{v_r^0}, \delta_{v_r^1}\}, \{\delta_{v_\ell^1}, \delta_{v_\ell^1}, \delta_{v_r^1}\}) \\
&= \min\{\mathcal{D}(\Delta^{n-2}(\mathbf{0}))(\delta_{v_r^0}, \delta_{v_\ell^1}), \mathcal{D}(\Delta^{n-2}(\mathbf{0}))(\delta_{v_r^0}, \delta_{v_r^1})\} \quad \text{[Proposition 24]} \\
&= \min\{\Delta^{n-2}(\mathbf{0})(v_r^0, v_\ell^1), \Delta^{n-2}(\mathbf{0})(v_r^0, v_r^1)\} \quad \text{[Proposition \textbf{??}]}
\end{aligned}$$

Combining the above, we arrive at

$$\Delta^n(\mathbf{0})(v^0, v^1) = \min\{\Delta^{n-2}(\mathbf{0})(v_\ell^0, v_\ell^1), \Delta^{n-2}(\mathbf{0})(v_\ell^0, v_r^1), \Delta^{n-2}(\mathbf{0})(v_r^0, v_\ell^1), \Delta^{n-2}(\mathbf{0})(v_r^0, v_r^1)\}.$$

If $n = 2$, then

$$\Delta^n(\mathbf{0})(v^0, v^1) = 0 \leq \Phi^n(\mathbf{0})(v).$$

Let $n > 2$. Due to the labelling, we have that $\Delta^{n-2}(\mathbf{0})(v_\ell^0, v_r^1) = 1$ and $\Delta^{n-2}(\mathbf{0})(v_r^0, v_\ell^1) = 1$. Hence,

$$\Delta^n(\mathbf{0})(v_1^0, v_1^1) = \min\{\Delta^{n-2}(\mathbf{0})(v_\ell^0, v_\ell^1), \Delta^{n-2}(\mathbf{0})(v_r^0, v_r^1)\}.$$

Therefore,

$$\begin{aligned}
&\Delta^n(\mathbf{0})(v_1^0, v_1^1)\\
&= \min\{\Delta^{n-2}(\mathbf{0})(v_\ell^0, v_\ell^1), \Delta^{n-2}(\mathbf{0})(v_r^0, v_r^1)\}\\
&\le \min\{\Phi^{n-2}(\mathbf{0})(v_\ell), \Phi^{n-2}(\mathbf{0})(v_r)\} \quad \text{[induction]}\\
&= \Phi^{n-1}(\mathbf{0})(v)\\
&\le \Phi^n(\mathbf{0})(v) \quad \text{[Proposition \textbf{??}]}
\end{aligned}$$

Furthermore, for $n > 1$ we have that

$$\begin{aligned}
\Phi^n(\mathbf{0})(v) &= \min\{\Phi^{n-1}(\mathbf{0})(v_\ell), \Phi^{n-1}(\mathbf{0})(v_r)\}\\
&\le \min\{\Delta^{2n-2}(\mathbf{0})(v_\ell^0, v_\ell^1), \Delta^{2n-2}(\mathbf{0})(v_r^0, v_r^1)\} \quad \text{[induction]}\\
&= \Delta^{2n}(\mathbf{0})(v^0, v^1).
\end{aligned}$$

– Let v be a max vertex. Then

$$\begin{aligned}
&\Delta^n(\mathbf{0})(v^0, v^1)\\
&= \mathcal{P}(\mathcal{D}(\Delta^{n-1}(\mathbf{0})))(\{\delta_{v_\ell^0}, \delta_{v_r^0}\}, \{\delta_{v_\ell^1}, \delta_{v_r^1}\})\\
&= \max\{\max\{\min\{\mathcal{D}(\Delta^{n-1}(\mathbf{0}))(\delta_{v_\ell^0}, \delta_{v_\ell^1}), \mathcal{D}(\Delta^{n-1}(\mathbf{0}))(\delta_{v_\ell^0}, \delta_{v_r^1})\},\\
&\qquad\qquad \min\{\mathcal{D}(\Delta^{n-1}(\mathbf{0}))(\delta_{v_r^0}, \delta_{v_\ell^1}), \mathcal{D}(\Delta^{n-1}(\mathbf{0}))(\delta_{v_r^0}, \delta_{v_r^1})\}\}\\
&\qquad \max\{\min\{\mathcal{D}(\Delta^{n-1}(\mathbf{0}))(\delta_{v_r^0}, \delta_{v_\ell^1}), \mathcal{D}(\Delta^{n-1}(\mathbf{0}))(\delta_{v_\ell^0}, \delta_{v_\ell^1})\},\\
&\qquad\qquad \min\{\mathcal{D}(\Delta^{n-1}(\mathbf{0}))(\delta_{v_\ell^0}, \delta_{v_r^1}), \mathcal{D}(\Delta^{n-1}(\mathbf{0}))(\delta_{v_r^0}, \delta_{v_r^1})\}\}\}\\
&= \max\{\max\{\min\{\Delta^{n-1}(\mathbf{0})(v_\ell^0, v_\ell^1), \Delta^{n-1}(\mathbf{0})(v_\ell^0, v_r^1)\},\\
&\qquad\qquad \min\{\Delta^{n-1}(\mathbf{0})(v_r^0, v_\ell^1), \Delta^{n-1}(\mathbf{0})(v_r^0, v_r^1)\}\}\\
&\qquad \max\{\min\{\Delta^{n-1}(\mathbf{0})(v_r^0, v_\ell^1), \Delta^{n-1}(\mathbf{0})(v_\ell^0, v_\ell^1)\},\\
&\qquad\qquad \min\{\Delta^{n-1}(\mathbf{0})(v_\ell^0, v_r^1), \Delta^{n-1}(\mathbf{0})(v_r^0, v_r^1)\}\}\} \quad \text{[Proposition \textbf{??}]}
\end{aligned}$$

If $n = 1$, then

$$\begin{aligned}
\Delta^n(\mathbf{0})(v^0, v^1) &= 0\\
&\le \Phi^n(\mathbf{0})(v)\\
&= \max\{\Phi^{n-1}(\mathbf{0})(v_\ell), \Phi^{n-1}(\mathbf{0})(v_r)\}\\
&= 0\\
&\le \Delta^{2n}(\mathbf{0})(v^0, v^1).
\end{aligned}$$

Let $n > 1$. Due to the labelling, we have that $\Delta^{n-1}(\mathbf{0})(v_\ell^1, v_r^1) = 1$ and $\Delta^{n-1}(\mathbf{0})(v_r^0, v_\ell^1) = 1$. Hence,

$$\begin{aligned}
\Delta^n(\mathbf{0})(v^0, v^1) &= \max\{\Delta^{n-1}(\mathbf{0})(v_\ell^0, v_\ell^1), \Delta^{n-1}(\mathbf{0})(v_r^0, v_r^1)\} \\
&\leq \max\{\Phi^{n-1}(\mathbf{0})(v_\ell), \Phi^{n-1}(\mathbf{0})(v_r)\} \quad \text{[induction]} \\
&= \Phi^n(\mathbf{0})(v) \\
&= \max\{\Phi^{n-1}(\mathbf{0})(v_\ell), \Phi^{n-1}(\mathbf{0})(v_r)\} \\
&\leq \max\{\Delta^{2n-2}(\mathbf{0})(v_\ell^0, v_\ell^1), \Delta^{2n-2}(\mathbf{0})(v_r^0, v_r^1)\} \quad \text{[induction]} \\
&= \Delta^{2n-1}(\mathbf{0})(v^0, v^1) \\
&\leq \Delta^{2n}(\mathbf{0})(v^0, v^1).
\end{aligned}$$

– Let v be an avg vertex. Then

$$\begin{aligned}
\Delta^n(\mathbf{0})(v^0, v^1) &= \mathcal{P}(\mathcal{D}(\Delta^{n-1}(\mathbf{0})))(\{\tfrac{1}{2}\delta_{v_\ell^0} + \tfrac{1}{2}\delta_{v_r^0}\}, \{\tfrac{1}{2}\delta_{v_\ell^1} + \tfrac{1}{2}\delta_{v_r^1}\}) \\
&= \mathcal{D}(\Delta^{n-1}(\mathbf{0}))(\tfrac{1}{2}\delta_{v_\ell^0} + \tfrac{1}{2}\delta_{v_r^0}, \tfrac{1}{2}\delta_{v_\ell^1} + \tfrac{1}{2}\delta_{v_r^1})
\end{aligned}$$

Let $n = 1$. Note that $\mathbf{0} \nrightarrow [0,1]$ are the constant functions. Hence,

$$\begin{aligned}
\Delta^n(\mathbf{0})(v^0, v^1) &= \mathcal{D}(\Delta^{n-1}(\mathbf{0}))(\tfrac{1}{2}\delta_{v_\ell^0} + \tfrac{1}{2}\delta_{v_r^0}, \tfrac{1}{2}\delta_{v_\ell^1} + \tfrac{1}{2}\delta_{v_r^1}) \\
&= \sup_{f \in \mathbf{0} \nrightarrow [0,1]} (\tfrac{1}{2} f(v_\ell^0) + f(v_r^0)) - (\tfrac{1}{2} f(v_\ell^1) + f(v_r^1)) \\
&= 0 \\
&\leq \Phi^n(\mathbf{0})(v_0) \\
&= \max\{\Phi^{n-1}(\mathbf{0})(v_\ell), \Phi^{n-1}(\mathbf{0})(v_r)\} \\
&= 0 \\
&\leq \Delta^{2n}(\mathbf{0})(v^0, v^1).
\end{aligned}$$

Let $n > 1$. Due to the labelling, we have that $\Delta^{n-1}(\mathbf{0})(v_\ell^1, v_r^1) = 1$ and $\Delta^{n-1}(\mathbf{0})(v_r^0, v_\ell^1) = 1$. Hence,

$$\begin{aligned}
\Delta^n(\mathbf{0})(v^0, v^1) &= \mathcal{D}(\Delta^{n-1}(\mathbf{0}))(\tfrac{1}{2}\delta_{v_\ell^0} + \tfrac{1}{2}\delta_{v_r^0}, \tfrac{1}{2}\delta_{v_\ell^1} + \tfrac{1}{2}\delta_{v_r^1}) \\
&= \tfrac{1}{2}\Delta^{n-1}(\mathbf{0})(v_\ell^0, v_\ell^1) + \tfrac{1}{2}\Delta^{n-1}(\mathbf{0})(v_r^0, v_r^1) \quad \text{[Proposition 26]} \\
&\leq \tfrac{1}{2}\Phi^{n-1}(\mathbf{0})(v_\ell) + \tfrac{1}{2}\Phi^{n-1}(\mathbf{0})(v_r) \quad \text{[induction]} \\
&= \Phi^n(\mathbf{0})(v) \\
&= \tfrac{1}{2}\Phi^{n-1}(\mathbf{0})(v_\ell) + \tfrac{1}{2}\Phi^{n-1}(\mathbf{0})(v_r) \\
&\leq \tfrac{1}{2}\Delta^{2n-2}(\mathbf{0})(v_\ell^0, v_\ell^1) + \tfrac{1}{2}\Delta^{2n-2}(\mathbf{0})(v_r^0, v_r^1) \quad \text{[induction]} \\
&= \Delta^{2n-1}(\mathbf{0})(v^0, v^1) \\
&\leq \Delta^{2n}(\mathbf{0})(v^0, v^1).
\end{aligned}$$

□

E Proofs of Section 6

Theorem 19. *For all s_1, $s_2 \in S$, $\delta(s_1, s_2) = \phi(v_{s_1,s_2})$.*

Proof. We first show that for all s_1, $s_2 \in S$, $\delta(s_1, s_2) \leq \phi(v_{s_1,s_2})$.

We define the distance function $d_\phi : S \times S \to [0, 1]$ by

$$d_\phi(s_1, s_2) = \phi(v_{s_1,s_2}).$$

It suffices to show that d_ϕ is a fixed point of Δ. Let s_1, $s_2 \in S$. We distinguish two cases. If $\ell(s_1) \neq \ell(s_2)$ then

$$\Delta(d_\phi)(s_1, s_2) = 1 = \phi(v_{s_1,s_2}) = d_\phi(s_1, s_2).$$

According to Proposition **??**,

$$\mathcal{D}(d)(\mu_1, \mu_2) = \min_{\omega \in \Omega_{\mu_1,\mu_2}} \sum_{s'_1,s'_2 \in S} \omega(s'_1, s'_2) d(s'_1, s'_2).$$

Since a linear function on a convex polytope attains its minimum at a vertex of the polytope, the above is equal to

$$\mathcal{D}(d)(\mu_1, \mu_2) = \min_{\omega \in V_{\mu_1,\mu_2}} \sum_{s'_1,s'_2 \in S} \omega(s'_1, s'_2) d(s'_1, s'_2).$$

Now we consider the second case, that is, $\ell(s_1) = \ell(s_2)$. Then

$$
\begin{aligned}
&\Delta(d_\phi)(s_1, s_2)\\
&= \mathcal{P}(\mathcal{D}(d_\phi))(\{\,\mu_1 \mid s_1 \to \mu_1\,\}, \{\,\mu_2 \mid s_2 \to \mu_2\,\})\\
&= \max\{\max_{s_1\to\mu_1}\min_{s_2\to\mu_2}\mathcal{D}(d_\phi)(\mu_1,\mu_2), \max_{s_2\to\mu_2}\min_{s_1\to\mu_1}\mathcal{D}(d_\phi)(\mu_1,\mu_2)\}\\
&= \max\left\{\max_{s_1\to\mu_1}\min_{s_2\to\mu_2}\min_{\omega\in V_{\mu_1,\mu_2}}\sum_{s_1',s_2'\in S}\omega(s_1',s_2')d_\phi(s_1',s_2'),\right.\\
&\qquad\left.\max_{s_2\to\mu_2}\min_{s_1\to\mu_1}\min_{\omega\in V_{\mu_1,\mu_2}}\sum_{s_1',s_2'\in S}\omega(s_1',s_2')d_\phi(s_1',s_2')\right\}\\
&= \max\left\{\max_{s_1\to\mu_1}\min_{s_2\to\mu_2}\min_{\omega\in V_{\mu_1,\mu_2}}\sum_{s_1',s_2'\in S}\omega(s_1',s_2')\phi(v_{s_1',s_2'}),\right.\\
&\qquad\left.\max_{s_2\to\mu_2}\min_{s_1\to\mu_1}\min_{\omega\in V_{\mu_1,\mu_2}}\sum_{s_1',s_2'\in S}\omega(s_1',s_2')\phi(v_{s_1',s_2'})\right\}\\
&= \max\left\{\max_{s_1\to\mu_1}\min_{s_2\to\mu_2}\min_{\omega\in V_{\mu_1,\mu_2}}\phi(v_\omega), \max_{s_2\to\mu_2}\min_{s_1\to\mu_1}\min_{\omega\in V_{\mu_1,\mu_2}}\phi(v_\omega)\right\}\\
&= \max\left\{\max_{s_1\to\mu_1}\min_{s_2\to\mu_2}\Phi(\phi)(v_{\mu_1,\mu_2}), \max_{s_2\to\mu_2}\min_{s_1\to\mu_1}\Phi(\phi)(v_{\mu_1,\mu_2})\right\}\\
&= \max\left\{\max_{s_1\to\mu_1}\min_{s_2\to\mu_2}\phi(v_{\mu_1,\mu_2}), \max_{s_2\to\mu_2}\min_{s_1\to\mu_1}\phi(v_{\mu_1,\mu_2})\right\}\\
&= \max\left\{\max_{s_1\to\mu_1}\Phi(\phi)(v_{s_2,\mu_1}), \max_{s_2\to\mu_2}\Phi(\phi)(v_{s_1,\mu_2})\right\}\\
&= \max\left\{\max_{s_1\to\mu_1}\phi(v_{s_2,\mu_1}), \max_{s_2\to\mu_2}\phi(v_{s_1,\mu_2})\right\}\\
&= \Phi(\phi)(v_{s_1,s_2})\\
&= \phi(v_{s_1,s_2})\\
&= d_\phi(s_1, s_2).
\end{aligned}
$$

Next, we conclude that for all s_1, $s_2 \in S$, $\delta(s_1, s_2) \geq \phi(v_{s_1,s_2})$ by proving $\Delta^n(\mathbf{0})(s_1, s_2) \geq \Phi^n(\mathbf{0})(v_{s_1,s_2})$ for all $n \in \omega$ by induction. Obviously, the result holds when $n = 0$. Let $n > 0$. Then

$$
\begin{aligned}
&\Delta^n(\mathbf{0})(s_1, s_2)\\
&= \max \left\{ \begin{array}{l} \max\limits_{s_1 \to \mu_1} \min\limits_{s_2 \to \mu_2} \min\limits_{\omega \in V_{\mu_1,\mu_2}} \sum\limits_{s_1', s_2' \in S} \omega(s_1', s_2') \Delta^{n-1}(\mathbf{0})(s_1', s_2'), \\ \max\limits_{s_2 \to \mu_2} \min\limits_{s_1 \to \mu_1} \min\limits_{\omega \in V_{\mu_1,\mu_2}} \sum\limits_{s_1', s_2' \in S} \omega(s_1', s_2') \Delta^{n-1}(\mathbf{0})(s_1', s_2') \end{array} \right\}\\
&\geq \max \left\{ \begin{array}{l} \max\limits_{s_1 \to \mu_1} \min\limits_{s_2 \to \mu_2} \min\limits_{\omega \in V_{\mu_1,\mu_2}} \sum\limits_{s_1', s_2' \in S} \omega(s_1', s_2') \Phi^{n-1}(\mathbf{0})(v_{s_1', s_2'}), \\ \max\limits_{s_2 \to \mu_2} \min\limits_{s_1 \to \mu_1} \min\limits_{\omega \in V_{\mu_1,\mu_2}} \sum\limits_{s_1', s_2' \in S} \omega(s_1', s_2') \Phi^{n-1}(\mathbf{0})(v_{s_1', s_2'}) \end{array} \right\} \quad \text{[induction]}\\
&\geq \max \left\{ \begin{array}{l} \max\limits_{s_1 \to \mu_1} \min\limits_{s_2 \to \mu_2} \Phi^{n}(\mathbf{0})(v_{\mu_1,\mu_2}), \\ \max\limits_{s_2 \to \mu_2} \min\limits_{s_1 \to \mu_1} \Phi^{n}(\mathbf{0})(v_{\mu_1,\mu_2}) \end{array} \right\}\\
&\geq \max \left\{ \begin{array}{l} \max\limits_{s_1 \to \mu_1} \min\limits_{s_2 \to \mu_2} \Phi^{n-1}(\mathbf{0})(v_{\mu_1,\mu_2}), \\ \max\limits_{s_2 \to \mu_2} \min\limits_{s_1 \to \mu_1} \Phi^{n-1}(\mathbf{0})(v_{\mu_1,\mu_2}) \end{array} \right\} \quad \text{[Proposition \textbf{??}]}\\
&= \max \left\{ \max\limits_{s_1 \to \mu_1} \Phi^{n-1}(\mathbf{0})(v_{s_2,\mu_1}), \max\limits_{s_2 \to \mu_2} \Phi^{n-1}(\mathbf{0})(v_{s_1,\mu_2}) \right\}\\
&\geq \max \left\{ \max\limits_{s_1 \to \mu_1} \Phi^{n}(\mathbf{0})(v_{s_2,\mu_1}), \max\limits_{s_2 \to \mu_2} \Phi^{n}(\mathbf{0})(v_{s_1,\mu_2}) \right\} \quad \text{[Proposition \textbf{??}]}\\
&= \Phi^n(\mathbf{0})(v_{s_1,s_2})
\end{aligned}
$$

□

F Proofs of Section 7

Theorem 22. $\delta(s,t) = \frac{1}{n} \min_{\sigma \in \mathcal{S}_n} \sum_{1 \leq i \leq n} c(i, \sigma(i))$.

Proof. We have that

$$\delta(s', t') = \Delta(\delta)(s', t') = \mathcal{P}(\mathcal{D}(\delta))(\{\delta_{s'}\}, \emptyset) = 1. \tag{6}$$

Furthermore, for all $1 \leq i, j \leq n$,

$$\begin{aligned}
&\delta(s_{0,i}, t_{j,i}) \\
&= \Delta(\delta)(s_{0,i}, t_{j,i}) \\
&= \mathcal{P}(\mathcal{D}(\delta))(\{\delta_{s'}\}, \{c(i,j)\delta_{t'} + (1 - c(i,j))\delta_{s'}\}) \\
&= \mathcal{D}(\delta)(\delta_{s'}, c(i,j)\delta_{t'} + (1 - c(i,j))\delta_{s'}) \quad \text{[Proposition ??]} \\
&= c(i,j)\delta(s', t') \quad \text{[Proposition 27]} \\
&= c(i,j) \quad \text{[Equation (6)]}
\end{aligned} \tag{7}$$

From the labelling, we can conclude that for all $1 \leq i, j, k \leq n$ with $i \neq k$,

$$\delta(s_{0,i}, t_{j,k}) = \Delta(\delta)(s_{0,i}, t_{j,k}) = 1. \tag{8}$$

Also, for all $1 \leq i, j \leq n$,

$$\begin{aligned}
&\delta(s_i, t_j) \\
&= \Delta(\delta)(s_i, t_j) \\
&= \mathcal{P}(\mathcal{D}(\delta))(\{\delta_{s_{0,i}}\} \cup \{\, \delta_{t_{j,k}} \mid 1 \leq j, k \leq n \wedge k \neq i \,\}, \{\, \delta_{t_{j,k}} \mid 1 \leq k \leq n \,\}) \\
&= \min_{1 \leq k \leq n} \mathcal{D}(\delta)(\delta_{s_{0,i}}, \delta_{t_{j,k}}) \\
&= \min_{1 \leq k \leq n} \delta(s_{0,i}, t_{j,k}) \quad \text{[Proposition ??]} \\
&= \delta(s_{0,i}, t_{j,i}) \quad \text{[Equation (8)]} \\
&= c(i,j) \quad \text{[Equation (7)]}
\end{aligned} \tag{9}$$

Finally,

$$\begin{aligned}
&\delta(s, t) \\
&= \Delta(\delta)(s, t) \\
&= \mathcal{P}(\mathcal{D}(\delta)) \left(\left\{ \sum_{1 \leq i \leq n} \tfrac{1}{n} \delta_{s_i} \right\}, \left\{ \sum_{1 \leq i \leq n} \tfrac{1}{n} \delta_{t_i} \right\} \right) \\
&= \mathcal{P}(\mathcal{D}(\delta)) \left(\sum_{1 \leq i \leq n} \tfrac{1}{n} \delta_{s_i}, \sum_{1 \leq i \leq n} \tfrac{1}{n} \delta_{t_i} \right) \quad \text{[Proposition ??]} \\
&= \min_{\sigma \in S_n} \sum_{1 \leq i \leq n} \tfrac{1}{n} \delta(s_i, t_{\sigma(i)}) \quad \text{[Proposition 4]} \\
&= \min_{\sigma \in S_n} \sum_{1 \leq i \leq n} \tfrac{1}{n} c(i, \sigma(i)) \quad \text{[Equation (9)]} \\
&= \tfrac{1}{n} \min_{\sigma \in S_n} \sum_{1 \leq i \leq n} c(i, \sigma(i))
\end{aligned}$$

□

Theorem 33. *For all s_1, $s_2 \in S$,*

$$\delta(s_1, s_2) = 0 \textit{ iff } s_1 \sim s_2.$$

Proof. We prove two implications.

Let $R = \{ (s_1, s_2) \mid \delta(s_1, s_2) = 0 \}$. It suffices to prove that R is a probabilistic bisimulation. Let s_1, $s_2 \in S$. Since $\delta(s_1, s_2) = 0$, we can conclude that $\ell(s_1) = \ell(s_2)$. Assume that if $s_1 \to \mu_1$. Since $\delta(s_1, s_2) = 0$, there exists a $s_2 \to \mu_2$ such that $\mathcal{D}(\delta)(\mu_1, \mu_2) = 0$. According to Proposition **??**, there exists a coupling ω of μ_1, μ_2 such that $\sum_{s'_1, s'_2 \in S} \omega(s'_1, s'_2)\delta(s'_1, s'_2)$. Hence, $\omega(s'_1, s'_2) > 0$ implies $\delta(s'_1, s'_2) = 0$ and, hence, $\mu_1 \bar{R} \mu_2$.

Let s_1, $s_2 \in S$ with $s_1 \sim s_2$. From Proposition **??** and **??** we can conclude that it suffices to show that

$$\forall n \in \omega : \Delta^n(\mathbf{0})(s_1, s_2) = 0.$$

We prove this by induction on n. The above obviously holds for $n = 0$. Let $n > 0$. Then for each $s_1 \to \mu_1$ there exists a $s_2 \to \mu_2$ and a coupling ω of μ_1, μ_2 such that for all s'_1, $s'_2 \in S$, $\omega(s'_1, s'_2) > 0$ implies $s'_1 \sim s'_2$. By induction, $\Delta^{n-1}(\mathbf{0})(s'_1, s'_2) = 0$. From Proposition **??** we can conclude that $\Delta^{n-1}(\mathbf{0})(s_1, s_2) = 0$. □