

Game Relations and Metrics*

Luca de Alfaro
Computer Engineering Department
University of California, Santa Cruz, USA

Vishwanath Raman
Computer Science Department
University of California, Santa Cruz, USA
and
Synopsis Inc, Mountain View, USA

Rupak Majumdar
Department of Computer Science
University of California, Los Angeles, USA

Mariëlle Stoelinga
FMT Group
University of Twente, the Netherlands

Abstract

We consider two-player games played over finite state spaces for an infinite number of rounds. At each state, the players simultaneously choose moves; the moves determine a successor state. It is often advantageous for players to choose probability distributions over moves, rather than single moves. Given a goal (e.g., “reach a target state”), the question of winning is thus a probabilistic one: “what is the maximal probability of winning from a given state?”.

On these game structures, two fundamental notions are those of equivalences and metrics. Given a set of winning conditions, two states are equivalent if the players can win the same games with the same probability from both states. Metrics provide a bound on the difference in the probabilities of winning across states, capturing a quantitative notion of state “similarity”.

We introduce equivalences and metrics for two-player game structures, and we show that they characterize the difference in probability of winning games whose goals are expressed in the quantitative μ -calculus. The quantitative μ -calculus can express a large set of goals, including reachability, safety, and ω -regular properties. Thus, we claim that our relations and metrics provide the canonical extensions to games, of the classical notion of bisimulation for transition systems. We develop our results both for equivalences and metrics, which generalize bisimulation, and for asymmetrical versions, which generalize simulation.

*This research was sponsored in part by the grants NSF-CCF-0427202, NSF-CCF-0546170, and NSF-CCR-0132780.

1. Introduction

We consider two-player games played for an infinite number of rounds over finite state spaces. At each round, the players simultaneously and independently select moves; the moves then determine a probability distribution over successor states. These games, known variously as *stochastic games* [24] or *concurrent games* [3, 1, 5], generalize many common structures in computer science, from transition systems, to Markov chains [12] and Markov decision processes [6]. The games are *turn-based* if, at each state, at most one of the players has a choice of moves, and *deterministic* if the successor state is uniquely determined by the current state, and by the moves chosen by the players.

It is well-known that in such games with simultaneous moves it is often advantageous for the players to randomize their moves, so that at each round, they play not a single “pure” move, but rather, a probability distribution over the available moves. These probability distributions over moves, called *mixed moves* [20], lead to various notions of equilibria [29, 20], such as the equilibrium result expressed by the minimax theorem [29]. Intuitively, the benefit of playing mixed, rather than pure, moves lies in preventing the adversary from tailoring a response to the individual move played. Even for simple reachability games, the use of mixed moves may allow players to win, with probability 1, games that they would lose (i.e., win with probability 0) if restricted to playing pure moves [3]. With mixed moves, the question of winning a game with respect to a goal is thus a probabilistic one: what is the maximal probability a player can be guaranteed of winning, regardless of how the other player plays? This probability is known, in brief, as the *winning probability*.

In structures ranging from transition systems to Markov decision processes and games, a fundamental question is the one of *equivalence* of states. Given a suitably large class Φ of properties, containing all properties of interest to the modeler, two states are equivalent if the same properties hold in both states. For a property ϕ , denote the value of ϕ at s by $\phi(s)$: in the case of games, this might represent the maximal probability of a player winning with respect to a goal expressed by ϕ . Two states s and t are equivalent if $\phi(s) = \phi(t)$ for all $\phi \in \Phi$. For (finite-branching) transition systems, and for the class of properties Φ expressible in the μ -calculus [14], state equivalence is captured by bisimulation [19]; for Markov decision processes, it is captured by probabilistic bisimulation [22]. For quantitative properties, a notion related to equivalence is that of a *metric*: a metric provides a tight bound for how much the value of a property can differ at states of the system, and provides thus a quantitative notion of similarity between states. Given a set Φ of properties, the metric distance of two states s and t can be defined as $\sup_{\phi \in \Phi} |\phi(s) - \phi(t)|$. Metrics for Markov decision processes have been studied in [7, 27, 28, 8, 9]. Obviously, the metrics and relations are connected, in the sense that the relations are the *kernels* of the metrics (the pairs of states having metric distance 0). The metrics and relations are at the heart of many verification techniques, from approximate reasoning (one can substitute states that are close in the metric) to system reductions (one can collapse equivalent states) to compositional reasoning and refinement (providing a notion of substitutivity of equivalents).

We introduce metrics and equivalence relations for concurrent games, with respect to the class of properties Φ expressible in the quantitative μ -calculus [5, 18]. We claim that these metrics and relations represent the canonical extension of bisimulation to games. We also introduce asymmetrical versions of these metrics and equivalences, which constitute the canonical extension of simulation.

An equivalence relation for deterministic games that are either turn-based, or where the players are constrained to playing pure moves, has been introduced in [2] and called *alternating bisimulation*. Relations and metrics for the general case of concurrent games have so far proved elusive, with some previous attempts at their definition by a subset of the authors following a subtly flawed approach [4, 16]. The cause of the difficulty goes to the heart of the definition of bisimulation. In the definition of bisimulation for transition systems, for every pair s, t of bisimilar states, we require that if s can go to a state s' , then t should be able to go to t' , such that s' and t' are again bisimilar (we also ask that s, t have an equivalent predicate valuation). This definition has been extended to Markov decision processes by requiring that for every mixed move from s , there is a mixed move from t , such that the moves induce probability distributions over successor states that are equivalent modulo the under-

lying bisimulation [22, 21]. Unfortunately, the generalization of this appealing definition to games fails. It turns out, as we prove in this paper, that requiring players to be able to replicate probability distributions over successors (modulo the underlying equivalence) leads to an equivalence that is too fine, and that may fail to relate states at which the same quantitative μ -calculus formulas hold. We show that phrasing the definition in terms of distributions over successor states is the wrong approach for games; rather, the definition should be phrased in terms of expectations of certain metric-bounded quantities.

Our starting point is a closer look at the definition of metrics for Markov decision processes. We observe that we can manipulate the definition of metrics given in [28], obtaining an alternative form, which we call the *a priori* form, in contrast with the original form of [28], which we call the *a posteriori* form. Informally, the a posteriori form is the traditional definition, phrased in terms of similarity of probability distributions; the a priori form is instead phrased in terms of expectations. We show that, while on Markov decision processes these two forms coincide, this is not the case for games; moreover, we show that it is the a priori form that provides the canonical metrics for games.

We prove that the a priori metric distance between two states s and t of a concurrent game is equal to $\sup_{\phi \in \Phi} |\phi(s) - \phi(t)|$, where Φ is the set of properties expressible via the quantitative μ -calculus. This result can be summarized by saying that the quantitative μ -calculus provides a *logical characterization* for the a priori metrics, similar to the way the ordinary μ -calculus provides a logical characterization of bisimulation. Furthermore, we prove that a priori metrics — and their kernels, the a priori relations — satisfy a *reciprocity* property, stating that properties expressed in terms of player-1 and player-2 winning conditions have the same distinguishing power. This property is intimately connected to the fact that concurrent games, played with mixed moves, are *determined* for ω -regular goals [17, 5]: the probability that player 1 achieves a goal ψ is one minus the probability that player 2 achieves the goal $\neg\psi$. Reciprocity ensures that there is one, canonical, notion of game equivalence. This is in contrast to the case of alternating bisimulation of [2], in which there are distinct player-1 and player-2 versions, as a consequence of the fact that concurrent games, when played with pure moves, are not determined. The logical characterization and reciprocity result justify our claim that a priori metrics and relations are the canonical notion of metrics, and equivalence, for concurrent games. Neither the logical characterization nor the reciprocity result hold for the a posteriori metrics and relations.

While this introduction focused mostly on metrics and equivalence relations, we also develop results for the asymmetrical versions of these notions, related to simulation.

2. Games and Goals

We will develop metrics for game structures over a set S of states. We start with some preliminary definitions. For a finite set A , let $\text{Dist}(A)$ denote the set of probability distributions over A . We say that $p \in \text{Dist}(A)$ is *deterministic* if there is $a \in A$ such that $p(a) = 1$.

For a set S , a *valuation over S* is a function $f : S \mapsto [0, 1]$ associating with every element $s \in S$ a value $0 \leq f(s) \leq 1$; we let \mathcal{F} be the set of all valuations. For $c \in [0, 1]$, we denote by \mathbf{c} the constant valuation such that $\mathbf{c}(s) = c$ at all $s \in S$. We order valuations pointwise: for $f, g \in \mathcal{F}$, we write $f \leq g$ iff $f(s) \leq g(s)$ at all $s \in S$; we remark that \mathcal{F} , under \leq , forms a lattice.

Given $a, b \in \mathbb{R}$, we write $a \sqcup b = \max\{a, b\}$, and $a \sqcap b = \min\{a, b\}$; we also let $a \oplus b = \min\{1, \max\{0, a + b\}\}$ and $a \ominus b = \max\{0, \min\{1, a - b\}\}$. We extend $\sqcap, \sqcup, +, -, \oplus, \ominus$ to valuations by interpreting them in pointwise fashion.

A *directed metric* is a function $d : S^2 \mapsto \mathbb{R}_{\geq 0}$ which satisfies $d(s, s) = 0$ and $d(s, t) \leq d(s, u) + d(u, t)$ for all $s, t, u \in S$. We denote by $\mathcal{M} \subseteq S^2 \mapsto \mathbb{R}$ the space of all metrics; this space, ordered pointwise, forms a lattice which we indicate with (\mathcal{M}, \leq) . Given a metric $d \in \mathcal{M}$, we denote by \check{d} its *opposite* version, defined by $\check{d}(s, t) = d(t, s)$ for all $s, t \in S$; we say that d is symmetrical if $d = \check{d}$.

2.1. Game Structures

We assume a fixed, finite set \mathcal{V} of *observation variables*. A (two-player, concurrent) *game structure* $G = \langle S, [\cdot], \text{Moves}, \Gamma_1, \Gamma_2, \delta \rangle$ consists of the following components [1, 3]:

- A finite set S of states.
- A variable interpretation $[\cdot] : \mathcal{V} \times S \mapsto [0, 1]$, which associates with each variable $v \in \mathcal{V}$ a valuation $[v]$.
- A finite set Moves of moves.
- Two move assignments $\Gamma_1, \Gamma_2 : S \mapsto 2^{\text{Moves}} \setminus \emptyset$. For $i \in \{1, 2\}$, the assignment Γ_i associates with each state $s \in S$ the nonempty set $\Gamma_i(s) \subseteq \text{Moves}$ of moves available to player- i at state s .
- A probabilistic transition function $\delta : S \times \text{Moves} \times \text{Moves} \mapsto \text{Dist}(S)$, that gives the probability $\delta(s, a_1, a_2)(t)$ of a transition from s to t when player-1 plays move a_1 and player-2 plays move a_2 .

At every state $s \in S$, player 1 chooses a move $a_1 \in \Gamma_1(s)$, and simultaneously and independently player 2 chooses a move $a_2 \in \Gamma_2(s)$. The game then proceeds to the successor state $t \in S$ with probability $\delta(s, a_1, a_2)(t)$. We denote by $\text{Dest}(s, a_1, a_2) = \{t \in S \mid \delta(s, a_1, a_2)(t) > 0\}$ the set of *destination states* when actions a_1, a_2 are chosen at

s . The variables in \mathcal{V} naturally induce an equivalence on states: for states s, t , define $s \equiv t$ if for all $v \in \mathcal{V}$ we have $[v](s) = [v](t)$. In the following, unless otherwise noted, the definitions refer to a game structure with components $G = \langle S, [\cdot], \text{Moves}, \Gamma_1, \Gamma_2, \delta \rangle$. For player $i \in \{1, 2\}$, we write $\sim i = 3 - i$ for the opponent. We also consider the following subclasses of game structures.

- *Turn-based game structures*. A game structure G is *turn-based* if we can write S as the disjoint union of two sets: the set S_1 of *player-1 states*, and the set S_2 of *player-2 states*, such that $s \in S_1$ implies $|\Gamma_2(s)| = 1$, and $s \in S_2$ implies $|\Gamma_1(s)| = 1$, and further, there is a special variable $\text{turn} \in \mathcal{V}$, such that $[\text{turn}]s = 1$ iff $s \in S_1$, and $[\text{turn}]s = 0$ iff $s \in S_2$: thus, the variable turn indicates whose turn it is to play at a state. Turn-based games are often called *perfect information games* [20].
- *Markov decision processes*. A game structure G is a *Markov decision process* (MDP) [6] if only one of the two players has a choice of moves. For $i \in \{1, 2\}$, we say that a structure is an i -MDP if $\forall s \in S, |\Gamma_{\sim i}(s)| = 1$. For MDPs, we omit the (single) move of the player without a choice of moves, and write $\delta(s, a)$ for the transition function.
- *Deterministic game structures*. A game structure G is *deterministic* if, for all $s \in S, a_1 \in \text{Moves}$, and $a_2 \in \text{Moves}$, there exists a $t \in S$ such that $\delta(s, a_1, a_2)(t) = 1$; we denote such t by $\tau(s, a_1, a_2)$. We sometimes call *probabilistic* a general game structure, to emphasize the fact that it is not necessarily deterministic.

Pure and mixed moves. A *mixed move* is a probability distribution over the moves available to a player at a state. We denote by $\mathcal{D}_i(s) = \text{Dist}(\Gamma_i(s))$ the set of mixed moves available to player $i \in \{1, 2\}$ at $s \in S$. The moves in Moves are called *pure moves*, in contrast to mixed moves. We extend the transition function to mixed moves. For $s \in S$ and $x_1 \in \mathcal{D}_1(s), x_2 \in \mathcal{D}_2(s)$, we write $\delta(s, x_1, x_2)$ for the next-state probability distribution induced by the mixed moves x_1 and x_2 , defined for all $t \in S$ by

$$\delta(s, x_1, x_2)(t) = \sum_{a_1 \in \Gamma_1(s)} \sum_{a_2 \in \Gamma_2(s)} \delta(s, a_1, a_2)(t) x_1(a_1) x_2(a_2).$$

In the following, we sometimes restrict the moves of the players to pure moves. We identify a pure move $a \in \Gamma_i(s)$ available to player $i \in \{1, 2\}$ at a state s with a deterministic distribution that plays a with probability 1.

The deterministic setting. The *deterministic setting* consists in considering deterministic game structures, and players restricted to play pure moves.

Predecessor operators. Given a valuation $f \in \mathcal{F}$, a state $s \in S$, and two mixed moves $x_1 \in \mathcal{D}_1(s)$ and $x_2 \in \mathcal{D}_2(s)$, we define the expectation of f from s under x_1, x_2 :

$$\mathbb{E}_s^{x_1, x_2}(f) = \sum_{t \in S} \delta(s, x_1, x_2)(t) f(t)$$

For a game structure G , for $i \in \{1, 2\}$ we define the *valuation transformer* $\text{Pre}_i : \mathcal{F} \mapsto \mathcal{F}$ by, for all $f \in \mathcal{F}$ and $s \in S$,

$$\text{Pre}_i(f)(s) = \sup_{x_i \in \mathcal{D}_i(s)} \inf_{x_{\sim i} \in \mathcal{D}_{\sim i}(s)} \mathbb{E}_s^{x_1, x_2}(f).$$

Intuitively, $\text{Pre}_i(f)(s)$ is the maximal expectation player i can achieve of f after one step from s : this is the classical “one-day” or “next-stage” operator of the theory of repeated games [10]. We also define a *deterministic* version of this operator, in which players are forced to play pure moves:

$$\text{Pre}_i^\Gamma(f)(s) = \max_{x_i \in \Gamma_i(s)} \min_{x_{\sim i} \in \Gamma_{\sim i}(s)} \mathbb{E}_s^{x_1, x_2}(f).$$

2.2. Quantitative μ -calculus

We consider the set of properties expressed by the *quantitative μ -calculus* ($q\mu$). As discussed in [13, 5, 18], a large set of properties can be encoded in $q\mu$, spanning from basic properties such as maximal reachability and safety probability, to the maximal probability of satisfying a general ω -regular specification.

Syntax. The syntax of quantitative μ -calculus is defined with respect to the set of observation variables \mathcal{V} as well as a set $MVars$ of *calculus variables*, which are distinct from the observation variables in \mathcal{V} . The syntax is given as follows:

$$\begin{aligned} \phi ::= & c \mid v \mid Z \mid \neg\phi \mid \phi \vee \phi \mid \phi \wedge \phi \mid \phi \oplus c \mid \phi \ominus c \mid \\ & \text{pre}_1(\phi) \mid \text{pre}_2(\phi) \mid \mu Z. \phi \mid \nu Z. \phi \end{aligned}$$

for constants $c \in [0, 1]$, observation variables $v \in \mathcal{V}$, and calculus variables $Z \in MVars$. In the formulas $\mu Z. \phi$ and $\nu Z. \phi$, we furthermore require that all occurrences of the bound variable Z in ϕ occur in the scope of an even number of occurrences of the complement operator \neg . A formula ϕ is *closed* if every calculus variable Z in ϕ occurs in the scope of a quantifier μZ or νZ . From now on, with abuse of notation, we denote by $q\mu$ the set of closed formulas of $q\mu$. A formula is a *player- i formula*, for $i \in \{1, 2\}$, if ϕ does not contain the $\text{pre}_{\sim i}$ operator; we denote with $q\mu_i$ the syntactic subset of $q\mu$ consisting only of closed player- i formulas. A formula is in *positive form* if the negation appears only in front of game variables, i.e., in the context $\neg v$; we denote with $q\mu^+$ and $q\mu_i^+$ the subsets of $q\mu$ and $q\mu_i$ consisting only of positive formulas.

We remark that the fixpoint operators μ and ν will not be needed to achieve our results on the logical characterization

of game relations. They have been included in the calculus because they allow the expression of many interesting properties, such as safety, reachability, and in general, ω -regular properties.

Semantics. A variable valuation $\xi: MVars \mapsto \mathcal{F}$ is a function that maps every variable $Z \in MVars$ to a valuation in \mathcal{F} . We write $\xi[Z \mapsto f]$ for the valuation that agrees with ξ on all variables, except that Z is mapped to f . Given a game structure G and a variable valuation ξ , every formula ϕ of the quantitative μ -calculus defines a valuation $\llbracket \phi \rrbracket_\xi^G \in \mathcal{F}$ (the superscript G is omitted if the game structure is clear from the context):

$$\begin{aligned} \llbracket c \rrbracket_\xi &= \mathbf{c} & \llbracket v \rrbracket_\xi &= [v] & \llbracket Z \rrbracket_\xi &= \xi(Z) \\ \llbracket \neg\phi \rrbracket_\xi &= \mathbf{1} - \llbracket \phi \rrbracket_\xi & \llbracket \phi \left\{ \begin{smallmatrix} \oplus \\ \ominus \end{smallmatrix} \right\} c \rrbracket_\xi &= \llbracket \phi \rrbracket_\xi \left\{ \begin{smallmatrix} \oplus \\ \ominus \end{smallmatrix} \right\} \mathbf{c} \\ \llbracket \phi_1 \left\{ \begin{smallmatrix} \vee \\ \wedge \end{smallmatrix} \right\} \phi_2 \rrbracket_\xi &= \llbracket \phi_1 \rrbracket_\xi \left\{ \begin{smallmatrix} \vee \\ \wedge \end{smallmatrix} \right\} \llbracket \phi_2 \rrbracket_\xi \\ \llbracket \text{pre}_i(\phi) \rrbracket_\xi &= \text{Pre}_i(\llbracket \phi \rrbracket_\xi) \\ \llbracket \left\{ \begin{smallmatrix} \mu \\ \nu \end{smallmatrix} \right\} Z. \phi \rrbracket_\xi &= \left\{ \begin{smallmatrix} \inf \\ \sup \end{smallmatrix} \right\} \{ f \in \mathcal{F} \mid f = \llbracket \phi \rrbracket_{\xi[Z \mapsto f]} \} \end{aligned}$$

where $i \in \{1, 2\}$. The existence of the fixpoints is guaranteed by the monotonicity and continuity of all operators and can be computed by Picard iteration [5]. If ϕ is closed, $\llbracket \phi \rrbracket_\xi$ is independent of ξ , and we write simply $\llbracket \phi \rrbracket$.

We also define a *deterministic* semantics $\llbracket \cdot \rrbracket^\Gamma$ for $q\mu$, in which players can select only pure moves in the operators $\text{pre}_1, \text{pre}_2$. $\llbracket \cdot \rrbracket^\Gamma$ is defined as $\llbracket \cdot \rrbracket$, except for the clause

$$\llbracket \text{pre}_i(\phi) \rrbracket_\xi^\Gamma = \text{Pre}_i^\Gamma(\llbracket \phi \rrbracket_\xi^\Gamma).$$

Example 1 Given a set $T \subseteq S$, the *characteristic valuation* \mathbf{T} of T is defined by $\mathbf{T}(s) = 1$ if $s \in T$, and $\mathbf{T}(s) = 0$ otherwise. With this notation, the maximal probability with which player $i \in \{1, 2\}$ can ensure eventually reaching $T \subseteq S$ is given by $\llbracket \mu Z. (\mathbf{T} \vee \text{pre}_i(Z)) \rrbracket$, and the maximal probability with which player i can guarantee staying in T forever is given by $\llbracket \nu Z. (\mathbf{T} \wedge \text{pre}_i(Z)) \rrbracket$ [5].

3. Metrics

We are interested in developing a *metric* on states of a game structure that captures an approximate notion of equivalence: states close in the metric should yield similar values to the players for any winning objective. Specifically, we are interested in defining a bisimulation metric $[\simeq_g] \in \mathcal{M}$ such that for any game structure G and states s, t of G , the following continuity property holds:

$$[\simeq_g](s, t) = \sup_{\phi \in q\mu} |\llbracket \phi \rrbracket(s) - \llbracket \phi \rrbracket(t)|. \quad (1)$$

In particular, the kernel of the metric, that is, states at distance 0, are equivalent: each player can get exactly the same value from either state for any objective. Notice that

in defining the metric independent of a player, we are expecting our metrics to be *reciprocal*, that is, invariant under a change of player. Reciprocity is expected to hold since the underlying games we consider are determined — for any game, the value obtained by player 2 is 1 minus the value obtained by player 1— and yields canonical metrics on games.

Thus, our metrics will generalize equivalence and refinement relations that have been studied on MDPs and in the deterministic setting. To underline the connection between classical equivalences and the metrics we develop, we write $[s \simeq_g t]$ for $[\simeq_g](s, t)$, so that the desired property of the bisimulation metric can be stated as $[s \simeq_g t] = \sup_{\phi \in q\mu} |[\phi](s) - [\phi](t)|$. Metrics of this type have already been developed for Markov decision processes (MDPs) [27, 8]. Our construction of metrics for games starts from an analysis of these constructions.

3.1. Metrics for MDPs

We consider the case of 1-MDPs; the case for 2-MDPs is symmetrical. Throughout this subsection, we fix a 1-MDP $\langle S, [\cdot], \text{Moves}, \Gamma_1, \Gamma_2, \delta \rangle$. Before we present the metric correspondent of probabilistic simulation, we first rephrase classical probabilistic (bi)simulation on MDPs [15, 11, 22, 23] as a fixpoint of a *relation transformer* $F : 2^{S \times S} \mapsto 2^{S \times S}$. For all relations $R \subseteq S \times S$ and $s, t \in S$, let $(s, t) \in F(R)$ iff

$$s \equiv t \wedge \forall x_1 \in M_1(s) . \exists y_1 \in M_1(t) . \delta(s, x_1) \sqsubseteq_R \delta(t, y_1), \quad (2)$$

for all states $s, t \in S$. In (2), \equiv is the predicate equivalence relation: $s \equiv t$ if the predicates have the same value at s and t . The relation \sqsubseteq_R lifts the relation R on states to a relation on distributions. Precisely, for a relation $R \subseteq S \times S$ and two distributions $p, q \in \text{Dist}(S)$, we let $p \sqsubseteq_R q$ if there is a function $\Delta : S \times S \rightarrow [0, 1]$ such that (i) $\Delta(s, s') > 0$ implies $(s, s') \in R$, (ii) $p(s) = \sum_{s' \in S} \Delta(s, s')$ for any $s \in S$, and (iii) $q(s') = \sum_{s \in S} \Delta(s, s')$ for any $s' \in S$. Probabilistic simulation is the greatest fixpoint of (2); probabilistic bisimulation is the greatest symmetrical fixpoint of (2).

To obtain a metric equivalent of probabilistic simulation, we lift the above fixpoint from relations (subsets of S^2) to metrics (maps $S^2 \mapsto \mathbb{R}$). First, we define $[\equiv] \in \mathcal{M}$ for all $s, t \in S$ by $[s \equiv t] = \max_{v \in \mathcal{V}} |[v](s) - [v](t)|$. Second, we lift (2) to metrics, defining a metric transformer $H_{\text{post}}^{\text{IMDP}} : \mathcal{M} \mapsto \mathcal{M}$. For all $d \in \mathcal{M}$, let $D(\delta(s, x_1), \delta(t, y_1))(d)$ be the *distribution distance* between $\delta(s, x_1)$ and $\delta(t, y_1)$ with respect to the metric d . We will show later how to define such a distribution distance. For $s, t \in S$, we let

$$\begin{aligned} H_{\text{post}}^{\text{IMDP}}(d)(s, t) &= [s \equiv t] \sqcup \sup_{x_1 \in \Gamma_1(s)} \inf_{y_1 \in \Gamma_1(t)} D(\delta(s, x_1), \delta(t, y_1))(d). \end{aligned} \quad (3)$$

In this definition, the \forall and \exists of (2) have been replaced by \sup and \inf , respectively. Since equivalent states should have distance 0, the simulation metric in MDPs is defined as the *least* (rather than greatest) fixpoint of (3) [27, 8]. Similarly, the bisimulation metric is defined as the least symmetrical fixpoint of (3).

For a distance $d \in \mathcal{M}$ and two distributions $p, q \in \text{Dist}(S)$, the *distribution distance* $D(p, q)(d)$ between p and q with respect to d , is a measure of how much “work” we have to do to make p look like q , given that moving a unit of probability mass from $s \in S$ to $t \in S$ has cost $d(s, t)$. More precisely, $D(p, q)(d)$ is defined via the *trans-shipping problem*, as the minimum cost of shipping the distribution p into q , with edge costs d . Thus, $D(p, q)(d)$ is the solution of the following linear programming (LP) problem over the set of variables $\{\lambda_{s,t}\}_{s,t \in S}$ [27]:

$$\begin{aligned} \text{Minimize} \quad & \sum_{s,t \in S} d(s, t) \lambda_{s,t} \quad \text{subject to} \\ & \sum_{s \in S} \lambda_{s,t} = p(s), \quad \sum_{t \in S} \lambda_{s,t} = q(t), \quad \lambda_{s,t} \geq 0. \end{aligned}$$

Equivalently, we can define $D(p, q)(d)$ via the dual of the above LP problem. Given a metric $d \in \mathcal{M}$, let $C(d) \subseteq \mathcal{F}$ be the subset of valuations $k \in \mathcal{F}$ such that $k(s) - k(t) \leq d(s, t)$ for all $s, t \in S$. Then the dual formulation is:

$$\begin{aligned} \text{Maximize} \quad & \sum_{s \in S} p(s) k(s) - \sum_{s \in S} q(s) k(s) \quad (4) \\ & \text{subject to } k \in C(d). \end{aligned}$$

The constraint $C(d)$ on the valuation k , states the value of k across states cannot differ by more than d . This means, intuitively, that k behaves like the valuation of a $q\mu$ formula: as we will see, the logical characterization implies that d is a bound for the difference in valuation of $q\mu$ formulas across states. Indeed, the logical characterization of the metrics is proved by constructing formulas whose valuation approximate that of the optimal k . Plugging (4) into (3), we obtain:

$$\begin{aligned} H_{\text{post}}^{\text{IMDP}}(d)(s, t) &= [s \equiv t] \sqcup \\ & \sup_{x_1 \in \Gamma_1(s)} \inf_{y_1 \in \Gamma_1(t)} \sup_{k \in C(d)} (\mathbb{E}_s^{x_1}(k) - \mathbb{E}_t^{y_1}(k)) \end{aligned} \quad (5)$$

We can interpret this definition as follows. State t is trying to simulate state s (this is a definition of a simulation metric). First, state s chooses a mixed move x_1 , attempting to make simulation as hard as possible; then, state t chooses a mixed move y_1 , trying to match the effect of x_1 . Once x_1 and y_1 have been chosen, the resulting distance between s and t is equal to the maximal difference in expectation, for moves x_1 and y_1 , of a valuation $k \in C(d)$. We call the metric transformer $H_{\text{post}}^{\text{IMDP}}$ the *a posteriori* metric transformer: the valuation k in (5) is chosen *after* the moves x_1 and y_1 are chosen. We can define an *a priori* metric transformer,

where k is chosen before x_1 and y_1 :

$$H_{prio}^{IMDP}(d)(s, t) = [s \equiv t] \sqcup \sup_{k \in C(d)} \sup_{x_1 \in \Gamma_1(s)} \inf_{y_1 \in \Gamma_1(t)} (\mathbb{E}_s^{x_1}(k) - \mathbb{E}_t^{y_1}(k)) \quad (6)$$

Intuitively, in the a priori transformer, first a valuation $k \in C(d)$ is chosen. Then, state t must simulate state s with respect to the expectation of k . State s chooses a move x_1 , trying to maximize the difference in expectations, and state t chooses a move y_1 , trying to minimize it. The distance between s and t is then equal to the difference in the resulting expectations of k .

Theorem 1 below states that for MDPs, a priori and a posteriori simulation metrics coincide. In the next section, we will see that this is not the case for games.

Theorem 1 For all MDPs, $H_{post}^{IMDP} = H_{prio}^{IMDP}$.

Proof Consider two states $s, t \in S$, and a metric $d \in \mathcal{M}$. We have to prove that $\sup_k \sup_{x_1} \inf_{y_1} \mathbb{E}_s^{x_1}(k) - \mathbb{E}_t^{y_1}(k) = \sup_{x_1} \inf_{y_1} \sup_k \mathbb{E}_s^{x_1}(k) - \mathbb{E}_t^{y_1}(k)$. In the left-hand side, we can exchange the two outer sups. Then, noticing that the difference in expectation is bilinear in k and y_1 for a fixed x_1 , that y_1 is a probability distribution, and that k is chosen from a compact convex subset, we apply the generalized minimax theorem [25] to exchange $\sup_k \inf_{y_1}$ into $\inf_{y_1} \sup_k$, thus obtaining the right-hand side. ■

The metrics defined above are logically characterized by $q\mu$. Precisely, let $[\sim] \in \mathcal{M}$ be the least symmetrical fixpoint of $H_{prio}^{IMDP} = H_{post}^{IMDP}$. Then, the results of [8] (originally stated for H_{post}^{IMDP}) state that for all states s, t of a 1-MDP, we have $[s \sim t] = \sup_{\phi \in q\mu} |\llbracket \phi \rrbracket(s) - \llbracket \phi \rrbracket(t)|$.

3.2. Metrics for Concurrent Games

We now extend the simulation and bisimulation metrics from MDPs to general game structures. As we shall see, unlike for MDPs, the a priori and the a posteriori metrics do not coincide over games. In particular, we show that the a priori formulation satisfies both a tight logical characterization as well as reciprocity while, perhaps surprisingly, the more natural a posteriori version does not.

A posteriori metrics are defined via the metric transformer $H_{\sqsubseteq_1} : \mathcal{M} \mapsto \mathcal{M}$ as follows, for all $d \in \mathcal{M}$ and $s, t \in S$:

$$\begin{aligned} H_{\sqsubseteq_1}(d)(s, t) &= [s \equiv t] \sqcup \sup_{x_1 \in \mathcal{D}_1(s)} \inf_{y_1 \in \mathcal{D}_1(t)} \sup_{y_2 \in \mathcal{D}_2(t)} \inf_{x_2 \in \mathcal{D}_2(s)} \\ &\quad D(\delta(s, x_1, x_2), \delta(t, y_1, y_2), d) \\ &= [s \equiv t] \sqcup \sup_{x_1 \in \mathcal{D}_1(s)} \inf_{y_1 \in \mathcal{D}_1(t)} \sup_{y_2 \in \mathcal{D}_2(t)} \inf_{x_2 \in \mathcal{D}_2(s)} \\ &\quad \sup_{k \in C(d)} (\mathbb{E}_s^{x_1, x_2}(k) - \mathbb{E}_t^{y_1, y_2}(k)). \quad (7) \end{aligned}$$

We define a *posteriori game simulation metric* $[\sqsubseteq_1]$ as the least fixpoint of H_{\sqsubseteq_1} , and we define a *posteriori game bisimulation metric* $[\cong_1]$ as the least symmetrical fixpoint of H_{\sqsubseteq_1} . The a posteriori simulation metric $[\sqsubseteq_1]$ has been introduced in [4, 16].

A priori metrics are defined by bringing the \sup_k outside. Precisely, we define a metric transformer $H_{\preceq_1} : \mathcal{M} \mapsto \mathcal{M}$ as follows, for all $d \in \mathcal{M}$ and $s, t \in S$:

$$\begin{aligned} H_{\preceq_1}(d)(s, t) &= [s \equiv t] \sqcup \sup_{k \in C(d)} \sup_{x_1 \in \mathcal{D}_1(s)} \inf_{y_1 \in \mathcal{D}_1(t)} \sup_{y_2 \in \mathcal{D}_2(t)} \inf_{x_2 \in \mathcal{D}_2(s)} \\ &\quad (\mathbb{E}_s^{x_1, x_2}(k) - \mathbb{E}_t^{y_1, y_2}(k)) \\ &= [s \equiv t] \sqcup \sup_{k \in C(d)} \left(\sup_{x_1 \in \mathcal{D}_1(s)} \inf_{x_2 \in \mathcal{D}_2(s)} \mathbb{E}_s^{x_1, x_2}(k) \right. \\ &\quad \left. - \sup_{y_1 \in \mathcal{D}_1(t)} \inf_{y_2 \in \mathcal{D}_2(t)} \mathbb{E}_t^{y_1, y_2}(k) \right) \\ &= [s \equiv t] \sqcup \sup_{k \in C(d)} (\text{Pre}_1(k)(s) - \text{Pre}_1(k)(t)). \quad (8) \end{aligned}$$

The *a priori simulation metric* $[\preceq_1]$ is the least fixpoint of H_{\preceq_1} , and the *a priori bisimulation metric* $[\simeq_1]$ is the least symmetrical fixpoint of H_{\preceq_1} .

We now show some basic properties of these metrics. We show that a priori fixpoints give a (directed) metric. We show that a priori and a posteriori metrics are distinct. We then focus on the a priori metrics, and show through our results that they are the natural metrics for concurrent games.

Theorem 2 For all game structures G , and all states s, t, u of G , we have (1) $[s \preceq_1 t] \geq 0$, and $[s \sqsubseteq_1 t] \geq 0$, and (2) $[s \preceq_1 u] \leq [s \preceq_1 t] + [t \preceq_1 u]$.

Proof The first part follows by considering $k(s) = 0$ for all $s \in S$. The second part is proved by induction on the iterations used to compute the fixpoint of H_{\preceq_1} , as given in (8). The crucial step consists in observing that, for $d \in \mathcal{M}$, we have

$$\begin{aligned} &\sup_{k \in C(d)} (\text{Pre}_1(k)(s) - \text{Pre}_1(k)(t)) \\ &+ \sup_{k \in C(d)} (\text{Pre}_1(k)(t) - \text{Pre}_1(k)(u)) \\ &\geq \sup_{k \in C(d)} (\text{Pre}_1(k)(s) - \text{Pre}_1(k)(u)). \quad \blacksquare \end{aligned}$$

A priori and a posteriori metrics are distinct. First, we show that a priori and a posteriori metrics are distinct in general: the a priori metric never exceeds the a posteriori one, and there are concurrent games where it is strictly smaller. Intuitively, this can be explained as follows. Simulation entails trying to simulate the expectation of a valuation k , as we see from (7), (8). It is easier to simulate a state




Figure 1. Transition probabilities from states s, t to states u, w .

s from a state t if the valuation is known in advance, as in a priori metrics (8), than if the valuation k is chosen after all the moves have been chosen, as in a posteriori metrics (7).

As a special case, we shall see that equality holds for turn-based game structures, in addition to MDPs as we have seen in the previous subsection.

Theorem 3 *The following assertions hold.*

1. For all game structures G , and for all states s, t of G , we have $[s \preceq_1 t] \leq [s \sqsubseteq_1 t]$.
2. There is a game structure G , and states s, t of G , such that $[s \preceq_1 t] = 0$ and $[s \sqsubseteq_1 t] > 0$.
3. For all turn-based game structures, we have $\preceq_1 = \sqsubseteq_1$.

Proof The first assertion is a consequence of the fact that, for all functions $f : \mathbb{R}^2 \mapsto \mathbb{R}$, we have $\sup_x \inf_y f(x, y) \leq \inf_y \sup_x f(x, y)$ (intuitively, it is easier to maximize if we can choose x after y). Repeated applications of this allows us to show that, for all $d \in \mathcal{M}$, we have $H_{\preceq}(d) \leq H_{\sqsubseteq}(d)$ (with pointwise ordering). The result then follows from the monotonicity of H_{\preceq} and H_{\sqsubseteq} .

For the second assertion, we give an example where a priori distances are strictly less than a posteriori distances. Consider a game with states $S = \{s, t, u, w\}$. States u and w are sink states with $[u \equiv w] = 1$; states s and t are such that $[s \equiv t] = 0$. At states s and t , player-2 has moves $\{f, g\}$. Player-1 has a single move $\{a\}$ at state s , and moves $\{b, c\}$ at state t . The moves from s and t lead to u and w with transition probabilities indicated in Figure 1. In the figure, the point b, f indicates the probability of going to u and w when the move pair (b, f) is played, with $\delta(s, b, f)(u) + \delta(s, b, f)(w) = 1$; similarly for the other move pairs. The thick line segment between the points a, f and a, g represents the transition probabilities arising when player 1 plays move a , and player 2 plays a mixed move (a mix of f and g).

We show that, in this game, we have $[s \sqsubseteq_1 t] > 0$. Consider the metric d where $d(u, w) = 1$ (recall that $[u \equiv w] = 1$, and note the other distances do not matter,

since u, w are the only two destinations). We need to show

$$\forall y_1 \in \mathcal{D}_1(t). \exists y_2 \in \mathcal{D}_2(t). \forall x_2 \in \mathcal{D}_2(s). \exists k \in C(d). \\ (\mathbb{E}_s^{a, x_2}(k) - \mathbb{E}_t^{y_1, y_2}(k)) > 0 \quad (9)$$

Consider any mixed move $y_1 = \alpha b + (1 - \alpha)c$, where b, c are the moves available to player 1 at t , and $0 \leq \alpha \leq 1$. If $\alpha < \frac{1}{2}$, choose move f from t as y_2 , and choose $k(w) = 1$, $k(u) = 0$. Otherwise, choose move g from t as y_2 , and choose $k(w) = 0$, $k(u) = 1$. With these choices, the transition probability $\delta(t, y_1, y_2)$ will fall outside of the segment $[(a, f), (a, g)]$ in Figure 1. Thus, with the choice of k above, we ensure that the difference in (9) is always positive.

To show that in the game we have $[s \preceq_1 t] = 0$, it suffices to show (given that $[s \preceq_1 t] \geq 0$) that

$$\forall k \in C(d). \exists y_1 \in \mathcal{D}_1(t). \forall y_2 \in \mathcal{D}_2(t). \exists x_2 \in \mathcal{D}_2(s). \\ (\mathbb{E}_s^{a, x_2}(k) - \mathbb{E}_t^{y_1, y_2}(k)) \leq 0.$$

If $k(u) = k(w)$, the result is immediate. Assume otherwise, without loss of generality, that $k(u) < k(w)$, and choose $y_1 = c$. For every y_2 , the distribution of successor states (and of k -expectations) will be in the interval $[(c, f), (c, g)]$ in Figure 1. By choosing $x_2 = f$, we have that $\mathbb{E}_s^{a, f}(k) < \mathbb{E}_t^{c, y_2}(k)$ for all $y_2 \in \mathcal{D}_2(t)$, leading to the result.

The last assertion of the theorem is proven in the same way as Theorem 1. \blacksquare

Reciprocity of a priori metric. The previous theorem establishes that the a priori and a posteriori metrics are in general distinct. We now prove that it is the a priori metric, rather than the a posteriori one, that enjoys reciprocity, and that provides a (quantitative) logical characterization of $\llbracket q\mu \rrbracket$. We begin by considering reciprocity.

Theorem 4 *The following assertions hold.*

1. For all game structures G , we have $[\preceq_1] = [\succeq_2]$, and $[\simeq_1] = [\simeq_2]$.
2. There is a concurrent game structure G , with states s and t , where $[\sqsubseteq_1] \neq [\sqsupset_2]$.

Proof For the first assertion, it suffices to show that, for all $d \in \mathcal{M}$, and states $s, t \in S$, we have $H_{\preceq_1}(d)(s, t) = H_{\succeq_2}(d)(t, s)$. We proceed as follows:

$$\sup_{k \in C(d)} (\text{Pre}_1(k)(s) - \text{Pre}_1(k)(t)) \quad (10)$$

$$= \sup_{k \in C(d)} (-\text{Pre}_2(1 - k)(s) + \text{Pre}_2(1 - k)(t)) \quad (11)$$

$$= \sup_{k \in C(\check{d})} (\text{Pre}_2(k)(t) - \text{Pre}_2(k)(s)) \quad (12)$$

The step from (10) to (11) uses $\text{Pre}_1(k)(s) = 1 - \text{Pre}_2(1 - k)(s)$ [5], and the step from (11) to (12) uses the change of variables $k \rightarrow 1 - k$.

To show that reciprocity fails for a posteriori simulation metrics, consider the game from Figure 1. We add two new moves to player 2 at state t , namely $\{e, h\}$, such that for any $\alpha \in [0, 1]$, $\delta(t, *, \alpha e + (1 - \alpha)h)(u) = \delta(s, a, \alpha f + (1 - \alpha)g)(u)$: using move $\{e, h\}$ player 2 can ensure that transition probabilities are on the solid line in Figure 1. Reasoning as in Theorem 3, it still is the case that $[s \sqsubseteq_1 t] > 0$. Since the player-2 moves $\{e, h\}$ from state t are identical in induced transition probabilities to player-2 moves $\{f, g\}$ from state s , player 2 can restrict herself to moves $\{e, h\}$ at state t and ensure that $[s \sqsubseteq_2 t] = 0$. ■

As a consequence of this theorem, we write $[\simeq_g]$ in place of $[\simeq_1] = [\simeq_2]$, to emphasize that the player-1 and player-2 versions of game equivalence metrics coincide.

Logical characterization of a priori metric. The following theorem expresses the fact that $\llbracket q\mu \rrbracket$ provides a logical characterization for the a priori metrics. The proof of the theorem uses ideas from [16] and [8].

Theorem 5 *The following assertions hold for all game structures G , and for all states s, t of G :*

$$[s \preceq_1 t] = \sup_{\phi \in q\mu_1^+} (\llbracket \phi \rrbracket(s) - \llbracket \phi \rrbracket(t))$$

$$[s \simeq_g t] = \sup_{\phi \in q\mu} |\llbracket \phi \rrbracket(s) - \llbracket \phi \rrbracket(t)|$$

We note that, due to Theorem 3, an analogous result does not hold for the a posteriori metrics. Together with the lack of reciprocity of the a posteriori metrics, this is a strong indication that the a priori metrics, and not the a posteriori ones, are the “natural” metrics on concurrent games.

The Kernel. The kernel of the metric $[\simeq_g]$ defines an equivalence relation \simeq_g on the states of a game structure: $s \simeq_g t$ iff $[s \simeq_g t] = 0$. We call this the *game bisimulation* relation. Notice that by the reciprocity property of \simeq_g , the game bisimulation relation is canonical: $\simeq_1 = \simeq_2 = \simeq_g$. Similarly, we define the *game simulation* preorder $s \preceq_1 t$ as the kernel of the directed metric $[\preceq_1]$, that is, $s \preceq_1 t$ iff $[s \preceq_1 t] = 0$. Alternatively, it is possible to define \preceq_1 and \simeq_g directly. Given a relation $R \subseteq S \times S$, let $B(R) \subseteq \mathcal{F}$ consist of all valuations $k \in \mathcal{F}$ such that, for all $s, t \in S$, if sRt then $k(s) \leq k(t)$. We have the following result.

Theorem 6 *Given a game structure G , \preceq_1 (resp. \simeq_1) can be characterized as the largest (resp. largest symmetrical) relation R such that, for all states s, t with sRt , we have $s \equiv t$ and*

$$\forall k \in B(R). \forall x_1 \in \mathcal{D}_1(s). \exists y_1 \in \mathcal{D}_1(t). \forall y_2 \in \mathcal{D}_2(t).$$

$$\exists x_2 \in \mathcal{D}_2(s). (\mathbb{E}_t^{y_1, y_2}(k) \geq \mathbb{E}_s^{x_1, x_2}(k)). \quad (13)$$

We note that the above theorem allows the computation of \simeq_1 via a partition-refinement scheme. From the logical characterization theorem, we obtain the following corollary.

Corollary 1 *For any game structure G and states s, t of G , we have $s \simeq_g t$ iff $\llbracket \phi \rrbracket(s) = \llbracket \phi \rrbracket(t)$ holds for every $\phi \in q\mu$ and $s \preceq_1 t$ iff $\llbracket \phi \rrbracket(s) \leq \llbracket \phi \rrbracket(t)$ holds for every $\phi \in q\mu_1^+$.*

Computation. The next theorem states that the metrics are computable to any degree of precision. This follows, since the definition of the distance between two states of a given game, as the least fixpoint of the metric transformer (8), can be written as a formula in the theory of reals, which is decidable [26]. Since the distance between two states may not be rational, we can only guarantee an approximate computation in general.

Theorem 7 *For any game structure G , and states s, t of G , the following assertions hold.*

1. *For all rational v , and all $\epsilon > 0$, it is decidable if $|[s \preceq_1 t] - v| < \epsilon$ and if $|[s \simeq_g t] - v| < \epsilon$.*
2. *It is decidable if $s \preceq_1 t$ and if $s \simeq_g t$.*

Game Metrics and (Bi-)simulation Metrics. The a priori metrics assume an adversarial relationship between the players. We show that, on turn-based games, the a priori bisimulation metric coincides with the classical bisimulation metric where the players cooperate. Moreover, on 1-MDPs, the player-1 a priori simulation metric coincides with the cooperative simulation metric. The metric analog of classical (bi)simulation [19, 22] is obtained through the metric transformers $H_{\preceq_{12}} : \mathcal{M} \mapsto \mathcal{M}$ and $H_{\simeq_{12}} : \mathcal{M} \mapsto \mathcal{M}$ given by

$$H_{\preceq_{12}}(d)(s, t)$$

$$= [s \equiv t] \sqcup \sup_{k \in C(d)} \sup_{x_1 \in \mathcal{D}_1(s)} \sup_{x_2 \in \mathcal{D}_1(s)} \inf_{y_1 \in \mathcal{D}_1(t)} \inf_{y_2 \in \mathcal{D}_1(t)}$$

$$\{\mathbb{E}_s^{x_1, x_2}(k) - \mathbb{E}_t^{y_1, y_2}(k)\}$$

$$H_{\simeq_{12}}(d)(s, t) = H_{\preceq_{12}}(d)(s, t) \sqcup H_{\preceq_{12}}(d)(t, s)$$

The metrics $[\preceq_{12}]$ and $[\simeq_{12}]$ are defined as the least fixed points of $H_{\preceq_{12}}$ and $H_{\simeq_{12}}$ respectively. The kernel of these metrics define the classical probabilistic simulation and bisimulation relations.

Theorem 8 *The following assertions hold.*

1. *On turn-based game structures, $[\simeq_g] = [\simeq_{12}]$.*
2. *There is a deterministic game structure G and states s, t in G such that $[s \simeq_g t] > [s \simeq_{12} t]$.*
3. *There is a deterministic game structure G and states s, t in G such that $[s \simeq_g t] < [s \simeq_{12} t]$.*

The first part follows easily from the definition. The second part is proven by the game in Figure 2, where it holds that $[s \simeq_g t] = \frac{1}{2}$ and $[s \simeq_{12} t] = 0$. In this figure, as in all subsequent ones, different state colors denote that observation variables have different values at the states, so that the




Figure 2. $[s \simeq_g t] = \frac{1}{2}$ and $[s \simeq_{12} t] = 0$




Figure 3. $[s \simeq_g t] = 0$ but $[s \simeq_{12} t] = 1$.

states are distinguishable in $q\mu$. The third part is proven by the game in Figure 3 where $[s \equiv t] = 0$ and $[u \equiv v] = 1$.

Theorem 9 *The following assertions hold.*

1. For i -MDPs we have $[\preceq_i] = [\preceq_{12}]$.
2. There is a deterministic 2-MDP G with states s, t such that $[s \preceq_1 t] < [s \preceq_{12} t]$.
3. There is a deterministic 2-MDP G with states s, t such that $[t \preceq_1 s] > [t \preceq_{12} s]$.

Again, the first statement follows easily from the definition. The second and third are proven by the deterministic 2-MDP in Figure 4, where again $[s \equiv t] = 0$ and $[u \equiv v] = 1$.

4. Discussion

Our derivation of \preceq_i and \simeq_g , for $i \in \{1, 2\}$, as kernels of metrics, seems somewhat abstruse: most equivalence or similarity relations have been defined, after all, without resorting to metrics. We now point out how a generalization of the usual definitions [22, 2, 7, 8], suggested in [4, 16], fails to produce the “right” relations. Furthermore, the flawed relations obtained as a generalization of [22, 2, 7, 8] are no simpler than our definitions, based on kernel metrics. Thus, our study of game relations as kernels of metrics carries no drawbacks in terms of leading to more complicated definitions. Indeed, we believe that the metric approach is the superior one for the study of game relations.

We outline the flawed generalization of [22, 2, 7, 8] as proposed in [4, 16], explaining why it would seem a natural generalization. The alternating simulation of [2] is defined over deterministic game structures. Player- i alternating simulation, for $i \in \{1, 2\}$, is the largest relation R satisfying the following conditions, for all states $s, t \in S$: sRt implies $s \equiv t$ and $\forall a_i \in \Gamma_i(s) . \exists y_i \in \Gamma_i(t) . \forall y_{\sim i} \in \Gamma_{\sim i}(t) . \exists x_{\sim i} \in \Gamma_{\sim i}(s) . \tau(s, x_1, x_2)R\tau(t, y_1, y_2)$.




Figure 4. $[s \preceq_1 t] = 0$ and $[s \preceq_{12} t] = 1$. Also, $[t \preceq_1 s] = 1$ and $[t \preceq_{12} s] = 0$.

The MDP relations of [22], later extended to metrics by [7, 8], rely on the fixpoint (2), where \sup plays the role of \forall , \inf plays the role of \exists , and R is replaced by distribution equality modulo R , or \sqsubseteq_R . This strongly suggests — incorrectly — that equivalences for general games (probabilistic, concurrent games) can be obtained by taking the double alternating of $\forall\exists\forall\exists$ in the definition of alternating simulation, changing all \forall into \sup , all \exists into \inf , and replacing R by \sqsubseteq_R . The definition that would result is as follows. We parameterize the new relations by a player $i \in \{1, 2\}$, as well as by whether mixed moves or only pure moves are allowed. For a relation $R \subseteq S \times S$, for $M \in \{\Gamma, \mathcal{D}\}$, for all $s, t \in S$ and $i \in \{1, 2\}$ consider the following conditions:

- **(loc)** $s R t$ implies $s \equiv t$.
- **(M - i -altsim)** $s R t$ implies $\forall x_i \in M_i(s) . \exists y_i \in M_i(t) . \forall y_{\sim i} \in M_{\sim i}(t) . \exists x_{\sim i} \in M_{\sim i}(s) . \delta(s, x_1, x_2) \sqsubseteq_R \delta(t, y_1, y_2)$;

We then define the following relations:

- For $i \in \{1, 2\}$ and $M \in \{\Gamma, \mathcal{D}\}$, *player- i M -alternating simulation* \sqsubseteq_i^M is the largest relation that satisfies (loc) and (M - i -altsim).
- For $i \in \{1, 2\}$ and $M \in \{\Gamma, \mathcal{D}\}$, *player- i M -alternating bisimulation* \cong_i^M is the largest symmetrical relation that satisfies (loc) and (M - i -altsim).

Over deterministic game structures, the definitions of \sqsubseteq_i^Γ and \cong_i^Γ coincide with the alternating simulation and bisimulation relations of [2]. In fact, \sqsubseteq_i^Γ and \cong_i^Γ capture the *deterministic semantics* of $q\mu$, and thus in some sense generalize the results of [2] to probabilistic game structures.

Theorem 10 *For any game structure G and states s, t of G , the following assertions hold:*

1. $s \cong_i^\Gamma t$ iff $\llbracket \phi \rrbracket^\Gamma(s) = \llbracket \phi \rrbracket^\Gamma(t)$ holds for every $\phi \in q\mu_i$.
2. $s \sqsubseteq_i^\Gamma t$ iff $\llbracket \phi \rrbracket^\Gamma(s) \leq \llbracket \phi \rrbracket^\Gamma(t)$ holds for every $\phi \in q\mu_i^+$.

The following lemma states that $\sqsubseteq_i^{\mathcal{D}}$ and $\cong_i^{\mathcal{D}}$ are the kernels of $\llbracket \cdot \rrbracket_i$ and $\llbracket \cdot \rrbracket_i^+$, connecting thus the result of combining the definitions of [22] and [2] with a posteriori metrics.

Lemma 1 *For all game structures G , all players $i \in \{1, 2\}$, and all states s, t of G , we have $s \sqsubseteq_i^{\mathcal{D}} t$ iff $[s \sqsubseteq_i] = 0$, and $s \cong_i^{\mathcal{D}} t$ iff $[s \cong_i t] = 0$.*

We are now in a position to prove that neither the Γ -relations nor the \mathcal{D} -relations are the “correct” relations on general concurrent games, since neither characterizes $\llbracket q\mu \rrbracket$. In particular, the \mathcal{D} -relations are too fine, and the Γ -relations are incomparable with the relations \preceq_i and \simeq_g , for $i \in \{1, 2\}$. We prove these negative results first for the \mathcal{D} -relations. They follow from Theorem 3 and 5.

Theorem 11 *The following assertions hold:*

1. For all game structures G , all states s, t of G , and all $i \in \{1, 2\}$, we have that $s \sqsubseteq_i^{\mathcal{D}} t$ implies $s \preceq_i t$, and $s \cong_i^{\mathcal{D}} t$ implies $s \simeq_i t$.
2. There is a game structure G , and states s, t of G , such that $s \preceq_i t$ but $s \not\sqsubseteq_i^{\mathcal{D}} t$.
3. There is a game structure G , and states s, t of G , such that $\llbracket \phi \rrbracket(s) = \llbracket \phi \rrbracket(t)$ for all $\phi \in q\mu$, but $s \not\cong_i^{\mathcal{D}} t$ for some $i \in \{1, 2\}$.

We now turn our attention to the Γ -relations, showing that they are incomparable with \preceq_i and \simeq_g , for $i \in \{1, 2\}$.

Theorem 12 *The following assertions hold:*

1. There exists a deterministic game structure G and states s, t of G such that $s \sqsubseteq_1^{\Gamma} t$ but $s \not\preceq_1 t$, and $s \cong_1^{\Gamma} t$ but $s \not\simeq_g t$.
2. There exists a turn-based game structure G and states s, t of G such that $s \preceq_1 t$ but $s \not\sqsubseteq_1^{\Gamma} t$, and $s \simeq_g t$ but $s \not\cong_1^{\Gamma} t$.

Finally, we remark that, in view of Theorem 6, the definitions of the relations \preceq_i and \simeq_g for $i \in \{1, 2\}$ are no more complex than the definitions of $\sqsubseteq_1^{\mathcal{D}}$, \sqsubseteq_1^{Γ} , $\cong_1^{\mathcal{D}}$, and \cong_1^{Γ} .

References

- [1] R. Alur, T.A. Henzinger, and O. Kupferman. Alternating time temporal logic. *J. ACM*, 49:672–713, 2002.
- [2] R. Alur, T.A. Henzinger, O. Kupferman, and M.Y. Vardi. Alternating refinement relations. In *CONCUR 98*, LNCS 1466, pp. 163–178. Springer, 1998.
- [3] L. de Alfaro, T.A. Henzinger, and O. Kupferman. Concurrent reachability games. In *FOCS 98*, pp. 564–575. IEEE, 1998.
- [4] L. de Alfaro, T.A. Henzinger, and R. Majumdar. Discounting the future in systems theory. In *ICALP 03*, LNCS 2719, pp. 1022–1037. Springer, 2003.
- [5] L. de Alfaro and R. Majumdar. Quantitative solution of omega-regular games. *Journal of Computer and Systems Sciences*, 68:374–397, 2004.
- [6] C. Derman. *Finite State Markovian Decision Processes*. Academic Press, 1970.
- [7] J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden. Metrics for labelled markov systems. In *CONCUR 99*, LNCS 1664, pp. 258–273. Springer, 1999.
- [8] J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden. Approximating labelled markov processes. *Information and Computation*, 184(1):160–200, 2003.
- [9] J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden. The metric analogue of weak bisimulation for probabilistic processes. In *LICS 02*, pp. 413–422, 2002.
- [10] J. Filar and K. Vrieze. *Competitive Markov Decision Processes*. Springer-Verlag, 1997.
- [11] C.-C. Jou and S.A. Smolka. Equivalences, congruences and complete axiomatizations for probabilistic processes. In *CONCUR 90*, LNCS 458, pp. 367–383. Springer, 1990.
- [12] J.G. Kemeny, J.L. Snell, and A.W. Knapp. *Denumerable Markov Chains*. D. Van Nostrand Company, 1966.
- [13] D. Kozen. A probabilistic PDL. In *STOC 83*, pp. 291–297, 1983.
- [14] D. Kozen. Results on the propositional μ -calculus. *Theoretical Computer Science*, 27(3):333–354, 1983.
- [15] K.G. Larsen and A. Skou. Compositional Verification of Probabilistic Processes. In *CONCUR 92*, LNCS 630. Springer, 1992.
- [16] R. Majumdar. *Symbolic algorithms for verification and control*. PhD thesis, University of California, Berkeley, 2003.
- [17] D.A. Martin. The determinacy of Blackwell games. *J. of Symbolic Logic*, 63(4):1565–1581, 1998.
- [18] A. McIver and C. Morgan. *Abstraction, Refinement, and Proof for Probabilistic Systems*. Springer, 2004.
- [19] R. Milner. Operational and algebraic semantics of concurrent processes. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, vol B, pp. 1202–1242. Elsevier, 1990.
- [20] M.J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
- [21] R. Segala. *Modeling and Verification of Randomized Distributed Real-Time Systems*. PhD thesis, MIT, 1995. Technical Report MIT/LCS/TR-676.
- [22] R. Segala and N.A. Lynch. Probabilistic simulations for probabilistic processes. In *CONCUR 94*, LNCS 836, pp. 481–496. Springer, 1994.
- [23] R. Segala and N.A. Lynch. Probabilistic simulations for probabilistic processes. *Nordic Journal of Computing*, 2(2):250–273, 1995.
- [24] L.S. Shapley. Stochastic games. *Proceedings of the National Academy of Sciences, USA*, 39:1095–1100, 1953.
- [25] M. Sion. On general minimax theorems. *Pacific J. Math.*, 8:171–176, 1958.
- [26] A. Tarski. *A Decision Method for Elementary Algebra and Geometry*. University of California Press, 1951.
- [27] F. van Breugel and J. Worrel. An algorithm for quantitative verification of probabilistic transition systems. In *CONCUR 01*, LNCS 2154, pp. 336–350. Springer, 2001.
- [28] F. van Breugel and J. Worrel. Towards quantitative verification of probabilistic systems. In *ICALP 01*, LNCS 2976, pp. 421–432. Springer, 2001.
- [29] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. New York: John Wiley and Sons, 1944.