In *ICALP 04: 31st International Colloquium on Automata, Languages, and Programming,* LNCS, Springer-Verlag, 2004.

Linear and Branching Metrics for Quantitative Transition Systems*

Luca de Alfaro, Marco Faella, and Mariëlle Stoelinga

Department of Computer Engineering, Universitity of California, Santa Cruz, USA

Abstract. We extend the basic system relations of trace inclusion, trace equivalence, simulation, and bisimulation to a quantitative setting in which propositions are interpreted not as boolean values, but as real values in the interval $[0,1]$. Trace inclusion and equivalence give rise to asymmetrical and symmetrical *linear distances,* while simulation and bisimulation give rise to asymmetrical and symmetrical *branching distances.* We study the relationships among these distances, and we provide a full logical characterization of the distances in terms of quantitative versions of LTL and μ-calculus. We show that, while trace inclusion (resp. equivalence) coincides with simulation (resp. bisimulation) for deterministic boolean transition systems, linear and branching distances do not coincide for deterministic quantitative transition systems. Finally, we provide algorithms for computing the distances, together with matching lower and upper complexity bounds.

1 Introduction

Quantitative transition systems extend the usual transition systems, by interpreting propositions as numbers in [0,1], rather than as truth values. Quantitative transition systems arise in a wide range of contexts. They provide models for optimization problems, where the propositions can be interpreted as rewards, costs, or as the use of resources such as power and memory. They also provide models for discrete-time samplings of continuous systems, where the propositions represent the values of continuous variables at discrete instants of time. We extend the classical relations of trace inclusion, trace equivalence, simulation, and bisimulation to a quantitative setting, by defining linear and branching *distances*[1]. Considering distances, rather than relations, is particularly useful in the quantitative setting, as it leads to a theory of system approximations [5, 16, 1], enabling the quantification of how closely a concrete system implements a specification.

We define two families of distances: *linear distances*, which generalize trace inclusion and equivalence, and *branching distances*, which generalize (bi)simulation. We relate these distances to the quantitative version of the two well-known specification languages LTL and μ-calculus, showing that the distances measure to what extent the logic can tell one system from the other.

Our starting point for linear distances is the distance $\|\sigma-\rho\|_\infty$ between two traces σ and ρ, which measures the supremum of the difference in predicate valuations at

* This research was supported in part by the NSF CAREER grant CCR-0132780, the NSF grant CCR-0234690, and the ONR grant N00014-02-1-0671.

[1] In this paper, we use the term “distance” in a generic way, applying it to quantities that are traditionally called pseudo-metrics and quasi-pseudo-metrics [7].

corresponding positions of σ and ρ. To lift this trace distance to a distance over states, we define $ld^{\text{s}}(s,t) = \sup_{\sigma \in Tr(s)} \inf_{\rho \in Tr(t)} \|\sigma - \rho\|_\infty$, where $Tr(s)$ and $Tr(t)$ are the set of traces from s and t, respectively. The distance $ld^{\text{s}}(s,t)$ is asymmetrical, and is a quantitative extension of trace containment: if $ld^{\text{s}}(s,t) = b$, then for all traces σ from s, there is a trace ρ from t such that $\|\sigma - \rho\|_\infty \leq b$. In particular, $Tr(s) \subseteq Tr(t)$ iff $ld^{\text{s}}(s,t) = 0$. We define a symmetrical version of this distance by $\overline{ld}^{\text{s}}(s,t) = \max\{ld^{\text{s}}(s,t), ld^{\text{s}}(t,s)\}$, yielding a distance that generalizes trace equivalence; thus, $\overline{ld}^{\text{s}}(s,t)$ is the Hausdorff distance between $Tr(s)$ and $Tr(t)$.

We relate the linear distance to the logic QLTL, a quantitative version of LTL [12]. When interpreted on a quantitative transition system, QLTL formulas yield a real value in the interval [0,1]. The formula "next p" returns the (quantitative) value of p in the next step of a trace, while "eventually p" seeks the maximum value attained by p throughout the trace. The logical connectives "and" and "or" are interpreted as "min" and "max", and "not x" is interpreted as $1 - x$. Furthermore, QLTL has a bounded difference operator $\dot{-}$, defined as $x \dot{-} y = \max\{x - y, 0\}$.

In the boolean setting, for a relation to characterize a logic, two states must be related if and only if all formulas from the logic have the same truth value on them. In the quantitative framework, we can achieve a finer characterization: in addition to relating those states that formulas cannot distinguish, we can also *measure* to what extent the logic can tell one state from the other. We show that the linear distances provide such a measure for QLTL: for all states s,t we have $\overline{ld}^{\text{s}}(s,t) = \sup_{\varphi \in \text{QLTL}} |\varphi(s) - \varphi(t)|$ and $ld^{\text{s}}(s,t) = \sup_{\varphi \in \text{QLTL}} (\varphi(s) \dot{-} \varphi(t))$. We investigate what syntactic fragment of QLTL is necessary for such a characterization, showing in particular that the fragment must include the operator $\dot{-}$, in line with the results of [5, 11]. We also consider linear distances based on the asymmetric trace distance $\|\sigma \dot{-} \rho\|_\infty$ for traces σ and ρ. Intuitively, if $\|\sigma \dot{-} \rho\|_\infty = b$, then all predicate valuations along ρ are no more than b below the corresponding valuations in σ. Such asymmetrical distances are useful in optimization and control problems, where it is desired to approximate a given quantity from above or below. We show that these distances are characterized by the *positive* fragment of QLTL, in which all propositions occur with positive polarity.

We then study the branching distances that are the analogous of simulation and bisimulation on quantitative systems. A state s simulates a state t via R if the proposition valuations at s and t coincide, and if every successor of s is related via R to some successor of t. We generalize simulation to a distance bd^{As} over states. If $bd^{\text{As}}(s,t) = b$, then $\|s - t\|_\infty < b$, and every successor of s can be matched by a successor of t within bd^{As}-distance b. In a similar fashion, we can define a distance bd^{Ss} that is a quantitative analogous of bisimulation; such a distance has been studied in [5, 16]. We relate these distances to QMU, a quantitative fixpoint calculus that essentially coincides with the μ–calculus of [2], and is related to the calculi of [9, 3] (see also [8, 13]). In particular, we show that $bd^{\text{Ss}}(s,t) = \sup_{\varphi \in \text{QMU}} |\varphi(s) - \varphi(t)|$ and $bd^{\text{As}}(s,t) = \sup_{\varphi \in \exists\text{QMU}} (\varphi(s) \dot{-} \varphi(t))$, where $\exists$QMU is the fragment of QMU in which only existential predecessor operators occur. Similarly, starting from the asymmetrical state distance $\|s \dot{-} t\|_\infty$, we obtain branching distances that are characterized by the corresponding positive fragments of QMU. As before, these characterizations require the presence of the $\dot{-}$ operator in the calculus.

We relate linear and branching distances, showing that just as simulation implies trace containment, so the branching distances are greater than or equal to the corresponding linear distances. However, while trace inclusion (resp. equivalence) coincides with simulation (resp. bisimulation) for deterministic boolean transition systems, we show that linear and branching distances do not coincide for deterministic quantitative transition systems. Finally, we present algorithms for computing linear and branching distances over quantitative transition systems. We show that the problem of computing the linear distances is PSPACE-complete, and it remains PSPACE-complete even over deterministic systems, showing once more that determinism plays a lesser role in quantitative transition systems. The branching distances can be computed in polynomial time using standard fixpoint algorithms [2].

We also present our results in a *discounted* version, in which distances occurring i steps in the future are multiplied by α^i, where α is a discount factor in $[0,1]$. This discounted setting is common in the theory of games (see e.g. [6]) and optimal control (see e.g. [4]), and it leads to robust theories of quantitative systems [2].

2 Preliminaries

For two numbers $x,y \in [0,1]$, we write $x \sqcup y = \max(x,y)$, $x \sqcap y = \min(x,y)$, $x \dot{+} y = 1 \sqcap (x+y)$ and $x \dot{-} y = 0 \sqcup (x-y)$. We lift the operators $\sqcup$ and $\sqcap$, and the relations $<$, $\leq$ to functions via their pointwise extensions. Given a function $d : X^2 \mapsto \mathbb{R}^{\geq 0}$, we denote by $\text{Zero}(d) = \{(x,y) \in X^2 \mid d(x,y) = 0\}$ its zero set.

Quantitative transition systems. A *quantitative transition system* (QTS) $\mathscr{S} = (S, \tau, \Sigma, [\cdot])$ consists of a set S of states, a transition relation $\tau \subseteq S \times S$, a finite set Σ of propositions, and a function $[\cdot]$: $S \to (\Sigma \to [0,1])$ which assigns to each state $s \in S$ and proposition $r \in \Sigma$ a value $[s](r)$. For a state $s \in S$, we write $\tau(s)$ for $\{t \in S \mid (s,t) \in \tau\}$. We require that $\mathscr{S}$ is finite-branching and non-blocking: for all $s \in S$, the set $\tau(s)$ is finite and non-empty. We call a function $u : \Sigma \to [0,1]$ a Σ-*valuation*, and we denote by $\mathscr{U}$ the set of all Σ-valuations. A QTS $\mathscr{S}$ is *boolean* if for all $s \in S$ and all $r \in \Sigma$, we have $[s](r) \in \{0,1\}$. A QTS $\mathscr{S}$ is *deterministic* if for all states $s \in S$ and $t,t' \in \tau(s)$ with $t \neq t'$, there is $r \in \Sigma$ such that $[t](r) \neq [t'](r)$. When discussing algorithmic complexity, we assume that values $x \in [0,1]$ are encoded as fixed-point binary numbers, and we denote by $|x|_b$ the number of bits their encoding. We define the size of a (finite) QTS $\mathscr{S} = (S, \tau, \Sigma, [\cdot])$ by $|\mathscr{S}| = \sum_{s \in S} \sum_{r \in \Sigma} |[s](r)|_b + \sum_{s \in S} |\tau(s)|$.

Paths and traces. Given a set A and a sequence $\pi = a_0 a_1 a_2 \cdots \in A^\omega$, we write π_i for the i-th element a_i of π, and we write $\pi^i = a_i a_{i+1} a_{i+2} \cdots$ for the (infinite) suffix of π starting from π_i. A *path* of $\mathscr{S}$ is an infinite sequence $\pi = s_0 s_1 s_2 \cdots$ of states such that $(s_i, s_{i+1}) \in \tau$ for all $i \in \mathbb{N}$. Given a state $s \in S$, we write $Pts(s)$ for the set of all paths starting in s. A Σ-*trace* is an infinite sequence $\sigma = u_0 u_1 u_2 \cdots \in \mathscr{U}^\omega$; we call a Σ-trace simply a trace when Σ is clear from the context. Every path π of $\mathscr{S}$ induces a Σ-trace $[\pi] = [\pi_0][\pi_1][\pi_2] \cdots$; we write $Tr(s) = \{[\pi] \mid \pi \in Pts(s)\}$ for the set of traces from $s \in S$.

We define simulation, bisimulation, and trace containment for QTS as usual. Specifically, for a QTS $\mathscr{S} = (S, \tau, \Sigma, [\cdot])$, the simulation relation $\preceq_{sim}$ (resp. the bisimulation relation $\approx_{bis}$) is the largest relation $R \subseteq S \times S$ such that, for all sRt, the following conditions (i) and (ii) (resp. (i), (ii), and (iii)) hold: (i) $[s] = [t]$; (ii) for all $s' \in \tau(s)$, there is $t' \in \tau(t)$ with $s'Rt'$; (iii) for all $t' \in \tau(t)$, there is $s' \in \tau(s)$ with $s'Rt'$. For $s,t \in S$, we write $s \sqsubseteq_{tr} t$ if $Tr(s) \subseteq Tr(t)$, and $s \equiv_{tr} t$ if $Tr(s) = Tr(t)$.

Directed metrics and pseudometrics. A *directed metric* on X is a function $d : X \times X \to \mathbb{R}^{\geq 0}$ that satisfies $d(x,x) = 0$ for all $x \in X$ and the triangle inequality: $d(x,z) \leq d(x,y) + d(y,z)$ for all $x,y,z \in X$. A *pseudometric* d is a directed metric that is symmetric, i.e. $d(x,y) = d(y,x)$ for all $x,y \in X$. Given a directed metric, we denote by $\overline{d}$ its *symmetrization*, defined by $\overline{d}(s,t) = d(s,t) \sqcup d(t,s)$.

We develop our definitions in terms of directed metrics. Given a directed metric d on X and a mapping $q : X \to [0,1]$, the "directed" bound $d(x,y) \geq q(x) \dot{-} q(y)$ for all $x,y \in X$ immediately yields the "symmetrical" bound $\overline{d}(x,y) \geq |q(x) - q(y)|$ for all $x,y \in X$. Hence, we focus on directed metrics and directed bounds, deriving the symmetrical results through the above observation.

3 Linear Distances and Logics

Throughout this paper, unless specifically noted, we consider a fixed a QTS $\mathscr{S} = (S, \tau, \Sigma, [\cdot])$. The propositional distance between two states measures the maximum difference in their proposition evaluations.

Definition 1 (propositional distance) We define the *propositional distance* $pd : \mathscr{U}^2 \to [0,1]$, for all $u,v \in \mathscr{U}$, as $pd(u,v) = \max_{r \in \Sigma} (u(r) \dot{-} v(r))$. ∎

For ease of notation, we will write $pd(s,t)$ for $pd([s],[t])$. For $u,v \in \mathscr{U}$ we have $(u,v) \in \text{Zero}(\overline{pd})$ iff $u(r) = v(r)$ for all $r \in \Sigma$, and $(u,v) \in \text{Zero}(pd)$ iff $u(r) \leq v(r)$ for all $r \in \Sigma$. The definition of trace distance discounts the propositional distance at positions i of the trace by multiplying it by α^i, for $\alpha \in [0,1]$.

Definition 2 (trace distance) We define the *trace distance* $td_\alpha : \mathscr{U}^\omega \to [0,1]$ by letting, for $\sigma, \rho \in \mathscr{U}^\omega$ and $\alpha \in [0,1]$, $td_\alpha(\sigma,\rho) = \sup_{i \in \mathbb{N}} \alpha^i pd(\sigma_i, \rho_i)$. ∎

For $\alpha = 1$, the definitions reduce to the classical notions of trace distance: $td_1(\sigma,\rho) = \|\sigma \dot{-} \rho\|_\infty$, and $\overline{td}_1(\sigma,\rho) = \|\sigma - \rho\|_\infty$. We note that $\overline{td}_\alpha$ is a generalization of the Cantor metric, which equals $\overline{td}_{1/2}$. Intuitively, td (resp. $\overline{td}$) [2] corresponds to implication (resp. equivalence) along all the trace. Indeed, lifting $\leq$ and $=$ to traces in a pointwise way, for all $\sigma, \rho \in \mathscr{U}^\omega$ and $\alpha \in (0,1]$ we have that $(\sigma,\rho) \in \text{Zero}(\overline{td}_\alpha)$ iff $\sigma = \rho$, and $(\sigma,\rho) \in \text{Zero}(td_\alpha)$ iff $\sigma \leq \rho$. The linear distances are obtained by lifting trace distances to the set of all outgoing traces in two states, as in the Hausdorff distance.

Definition 3 (linear distance) We define the two *linear distances* ld^{a} and ld^{s} over S as follows, for $s,t \in S$ and $\alpha \in [0,1]$:

$$ld^{\text{a}}_\alpha(s,t) = \sup_{\sigma \in Tr(s)} \inf_{\rho \in Tr(t)} td_\alpha(\sigma,\rho) \qquad ld^{\text{s}}_\alpha(s,t) = \sup_{\sigma \in Tr(s)} \inf_{\rho \in Tr(t)} \overline{td}_\alpha(\sigma,\rho)$$ ∎

One can easily check that, for all $\alpha \in [0,1]$, the functions ld^{a}_α, ld^{s}_α are directed metrics and $\overline{ld}^{\text{a}}_\alpha$, $\overline{ld}^{\text{s}}_\alpha$ are pseudometrics. Intuitively, the distance ld^{s} is a quantitative

[2] When discussing properties that are independent of the discount factor, we sometimes omit the α subscript from distance names.

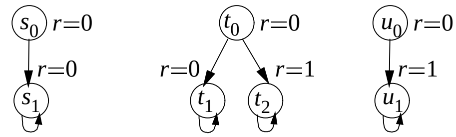

Fig. 1. A QTS showing the difference between $ld^{\rm a}_\alpha$, $ld^{\rm s}_\alpha$, $\overline{ld}^{\rm a}_\alpha$, and $\overline{ld}^{\rm s}_\alpha$.

extension of trace containment: for $s,t \in S$, the distance $ld^{\rm s}(s,t)$ measures how closely (in a quantitative sense) can a trace from t simulate a trace from s. The symmetrization of $ld^{\rm s}$ is $\overline{ld}^{\rm s}$, which is related to trace equivalence. The following result makes this observation precise.

Theorem 1 *For all* $\alpha \in (0,1]$, *we have* $\sqsubseteq_{tr} = \mathrm{Zero}(ld^{\rm s}_\alpha)$ *and* $\equiv_{tr} = \mathrm{Zero}(\overline{ld}^{\rm s}_\alpha)$.

We will see that the valuation of QLTL formulas at s and t can differ by at most $\overline{ld}^{\rm s}(s,t)$, and similarly, the valuation of any QLTL formula at t is at most $ld^{\rm s}(s,t)$ below the valuation at s. For $\alpha = 1$, the distances $ld^{\rm a}$ and $\overline{ld}^{\rm a}$ have the following intuitive characterization. For a trace $\sigma \in \mathscr{U}^\omega$ and $c \in \mathbb{R}$, denote by $\sigma \dot{-} c$ the trace defined by $(\sigma \dot{-} c)_k(r) = \sigma_k(r) \dot{-} c$ for all $k \in \mathbb{N}$ and $r \in \Sigma$: in other words, $\sigma \dot{-} c$ is obtained from σ by decreasing all proposition valuations by c. For all $s,t \in S$, if $ld^{\rm a}_1(s,t) = c$ then for every trace σ from s there is a trace ρ from t such that $\rho \geq \sigma \dot{-} c$. This means that $ld^{\rm a}_1(s,t)$ is a "positive" version of trace containment: for each trace σ of s, the goal of a trace ρ from t is not that of being close to σ, but rather, that of not being below $\sigma \dot{-} c$. This version of trace containment will preserve within c the valuation of QLTL formulas with only positive occurrences of propositions (called positive QLTL formulas). The relations among linear distances are summarized by the following theorem.

Theorem 2 *The relations in Figure 4(a) hold for all* $\alpha \in [0,1]$. *Moreover, for* $\alpha \in (0,1]$ *the inequalities cannot be replaced by equalities.*

Proof. The inequalities are immediate. For $\alpha \in (0,1]$ and the QTS in Figure 1, we have

$$\begin{array}{rcc} ld^{\rm a}_\alpha(s_0,t_0) = 0 & ld^{\rm a}_\alpha(t_0,u_0) = 0 & ld^{\rm a}_\alpha(u_0,t_0) = 0 \\ ld^{\rm s}_\alpha(s_0,t_0) = 0 & ld^{\rm s}_\alpha(t_0,u_0) = \alpha & ld^{\rm s}_\alpha(u_0,t_0) = 0 \\ \overline{ld}^{\rm a}_\alpha(s_0,t_0) = \alpha & \overline{ld}^{\rm a}_\alpha(t_0,u_0) = 0 & \overline{ld}^{\rm a}_\alpha(u_0,t_0) = 0 \\ \overline{ld}^{\rm s}_\alpha(s_0,t_0) = \alpha & \overline{ld}^{\rm s}_\alpha(t_0,u_0) = \alpha & \overline{ld}^{\rm s}_\alpha(u_0,t_0) = \alpha \end{array}$$

Thus, we have an example where $ld^{\rm a}_\alpha \neq ld^{\rm s}_\alpha$, $ld^{\rm a}_\alpha \neq \overline{ld}^{\rm a}_\alpha$, $ld^{\rm s}_\alpha \neq \overline{ld}^{\rm s}_\alpha$, $\overline{ld}^{\rm a}_\alpha \neq \overline{ld}^{\rm s}_\alpha$, and neither $ld^{\rm s}_\alpha \leq \overline{ld}^{\rm a}_\alpha$ nor $ld^{\rm s}_\alpha \geq \overline{ld}^{\rm a}_\alpha$. ∎

3.1 Quantitative Linear-Time Temporal Logic

The linear distances introduced above are closely connected to a quantitative extension of linear-time temporal logic which we call *quantitative linear-time temporal logic* (QLTL). The logic QLTL includes quantitative versions of the temporal operators and logic connectives. Following [5], QLTL also has a "threshold" operator, enabling the comparison of a formula against a constant in the interval $[0,1]$. The QLTL formulas over Σ are generated by the following grammar:

$$\varphi ::= r \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \neg\varphi \mid c \dot{+} \varphi \mid c \dot{-} \varphi \mid \bigcirc_\alpha \varphi \mid \ominus_\alpha \varphi \mid \Diamond_\alpha \varphi \mid \Box_\alpha \varphi$$

Here $r \in \Sigma$ is a proposition, $c \in [0,1]$ a constant and $\alpha \in [0,1]$ a discount factor. A formula φ assigns a value $[\![\varphi]\!](\rho) \in [0,1]$ to each trace $\sigma \subseteq \mathscr{U}^\omega$.

$$
\begin{aligned}
[\![r]\!](\sigma) &= \sigma_0(r) \\
[\![\neg\varphi]\!](\sigma) &= 1 - [\![\varphi]\!](\sigma) & [\![\varphi_1 \wedge \varphi_2]\!](\sigma) &= [\![\varphi_1]\!](\sigma) \sqcap [\![\varphi_2]\!](\sigma) \\
[\![c \dotplus \varphi]\!](\sigma) &= c \dotplus [\![\varphi]\!](\sigma) & [\![\varphi_1 \vee \varphi_2]\!](\sigma) &= [\![\varphi_1]\!](\sigma) \sqcup [\![\varphi_2]\!](\sigma) \\
[\![c \dot{-} \varphi]\!](\sigma) &= c \dot{-} [\![\varphi]\!](\sigma) & [\![\Diamond_\alpha\varphi]\!](\sigma) &= \sup\{\alpha^i \cdot [\![\varphi]\!](\sigma^i) \mid i \geq 0\} \\
[\![\bigcirc_\alpha\varphi]\!](\sigma) &= \alpha \cdot [\![\varphi]\!](\sigma^1) & [\![\Box_\alpha\varphi]\!](\sigma) &= \inf\{1 - \alpha^i \cdot (1 - [\![\varphi]\!](\sigma^i)) \mid i \geq 0\} \\
[\![\ominus_\alpha\varphi]\!](\sigma) &= 1 - \alpha + \alpha \cdot [\![\varphi]\!](\sigma^1)
\end{aligned}
$$

A QLTL formula φ assings a real value $[\![\varphi]\!](s) \in [0,1]$ to each state s of a given an QTS, according to the rule[3] $[\![\varphi]\!](s) = \sup\{[\![\varphi]\!](\rho) \mid \rho \in Tr(s)\}$. Thanks to the equivalences $\neg\bigcirc_\alpha\varphi \equiv \ominus_\alpha\neg\varphi$, $\neg(c \dotplus \varphi) \equiv ((1-c) \dot{-} \varphi)$, $\neg(c \dot{-} \varphi) \equiv ((1-c) \dotplus \varphi)$, $\neg(\Diamond_\alpha\varphi) \equiv \Box_\alpha\neg\varphi$, and the classical dualities between $\wedge$, $\vee$, μ, and ν, the syntax of QLTL allows negations to be pushed to the atomic propositions without affecting the value of a formula. For $\alpha \in [0,1]$, we denote by QLTL_α the set of formulas containing only discount factors smaller than or equal to α. All QLTL operators are *positive*, with the exception of $\neg$ and $c \dot{-}$ for $c \in [0,1]$, which are *negative*. We say that a QLTL formula is *positive* if all propositions occur with positive polarity, that is, within an even number of negative operators; we denote by QLTL^+_α the positive fragment of QLTL_α. Furthermore, for $ops \subseteq \{\bigcirc, \ominus, \Diamond, \Box, \dotplus, \dot{-}\}$, we denote by $\text{QLTL}_\alpha(ops)$ the set of formulas which only contain boolean connectives and operators in *ops*. We denote by $\text{QLTL}^+_\alpha(ops)$ the restrictions of these sets to positive formulas. Notice that for $\alpha = 1$, $\bigcirc_\alpha$ and $\ominus_\alpha$ coincide with the usual $\bigcirc$ operator of LTL. Thus, if we forbid the use of $\dotplus$ and $\dot{-}$ and we take all discount factors to be 1, the semantics of QLTL on boolean systems coincides with the one of LTL.

3.2 Logical Characterization of Linear Distances

Linear distances provide a bound for the difference in valuation of QLTL formulas. We begin by relating distances and logics over traces.

Lemma 1 *For all $\alpha \in [0,1]$ and all traces $\sigma, \rho \in \mathscr{U}^\omega$, the following holds.*

$$
\begin{aligned}
&\textit{For all } \varphi \in \text{QLTL}^+_\alpha: \quad td_\alpha(\sigma,\rho) \geq [\![\varphi]\!](\sigma) \dot{-} [\![\varphi]\!](\rho); \\
&\textit{For all } \varphi \in \text{QLTL}_\alpha: \quad \overline{td}_\alpha(\sigma,\rho) \geq |[\![\varphi]\!](\sigma) - [\![\varphi]\!](\rho)|.
\end{aligned}
$$

The following theorem uses the linear distances to provide the desired bounds for QLTL.

Theorem 3 *For all $\alpha \in [0,1]$ and $s,t \in S$, we have:*

$$
\begin{aligned}
&\textit{For all } \varphi \in \text{QLTL}^+_\alpha: ld^{\text{a}}_\alpha(s,t) \geq [\![\varphi]\!](s) \dot{-} [\![\varphi]\!](t) \textit{ and } \overline{ld}^{\text{a}}_\alpha(s,t) \geq |[\![\varphi]\!](s) - [\![\varphi]\!](t)|, \\
&\textit{For all } \varphi \in \text{QLTL}_\alpha: ld^{\text{s}}_\alpha(s,t) \geq [\![\varphi]\!](s) \dot{-} [\![\varphi]\!](t) \textit{ and } \overline{ld}^{\text{s}}_\alpha(s,t) \geq |[\![\varphi]\!](s) - [\![\varphi]\!](t)|.
\end{aligned}
$$

[3] We chose to give the existential interpretation of QLTL. Obviously, the minimum value of φ from s is obtained by one minus the maximum value of $\neg\varphi$ in s.

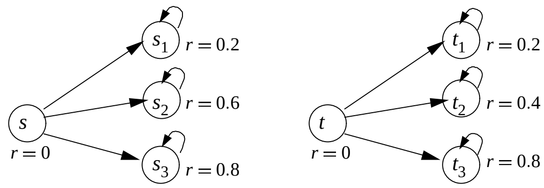

Fig. 2. QLTL cannot distinguish between s and t.

The results for ld^{s} and $\overline{ld}^{\mathrm{s}}$ are the quantitative analogous of the standard connection between trace containment and trace equivalence, and LTL. For instance, the result about ld^{s} states that, if $ld^{\mathrm{s}}_{\alpha}(s,t) = c$, then for every QLTL formula φ and every trace σ from s, there is a trace ρ from t such that $[\![\varphi]\!](\rho) \geq [\![\varphi]\!](\sigma) - c$.

The following theorem states that the linear distances can be characterized by a syntactic subset of the logics that includes only the $\bigcirc$ and $\dot{+}$ operators, in addition to boolean connectives. Together with Theorem 3, this result constitutes a full characterization of linear distances in terms of QLTL.

Theorem 4 *For all $\alpha \in [0,1]$ and $s,t \in S$,*

$$ld^{\mathrm{a}}_{\alpha}(s,t) = \sup_{\varphi \in \mathrm{QLTL}^{+}_{\alpha}(\bigcirc,\dot{+})} [\![\varphi]\!](s) \dot{-} [\![\varphi]\!](t) \qquad \overline{ld}^{\mathrm{a}}_{\alpha}(s,t) = \sup_{\varphi \in \mathrm{QLTL}^{+}_{\alpha}(\bigcirc,\dot{+})} |[\![\varphi]\!](s) - [\![\varphi]\!](t)|$$

$$ld^{\mathrm{s}}_{\alpha}(s,t) = \sup_{\varphi \in \mathrm{QLTL}_{\alpha}(\bigcirc,\dot{+})} [\![\varphi]\!](s) \dot{-} [\![\varphi]\!](t) \qquad \overline{ld}^{\mathrm{s}}_{\alpha}(s,t) = \sup_{\varphi \in \mathrm{QLTL}_{\alpha}(\bigcirc,\dot{+})} |[\![\varphi]\!](s) - [\![\varphi]\!](t)|$$

The next result shows that the operator $\dot{+}$ is indeed necessary to obtain such a characterization ($\bigcirc$ is also trivially necessary). This result is reminiscent of a result by [5] for Markov systems.

Theorem 5 *There is a finite QTS and two states s and t such that, for all $\alpha \in (0,1]$, $\overline{ld}^{\mathrm{s}}_{\alpha}(s,t) = ld^{\mathrm{s}}_{\alpha}(s,t) > 0$, and $\sup_{\varphi \in \mathrm{QLTL}_{\alpha}(\bigcirc,\circleddash,\diamond,\square)} |[\![\varphi]\!](s) - [\![\varphi]\!](t)| = 0$.*

As an example, consider the QTS in Figure 2, and assume $\alpha = 1$. It holds that $\overline{ld}^{\mathrm{s}}_{\alpha}(s,t) = ld^{\mathrm{s}}_{\alpha}(s,t) = 0.2$. A suitable formula for distinguishing s and t is $\varphi : \bigcirc[(0.6 \dot{+} \neg r) \wedge (0.4 \dot{+} r)]$; we have $\varphi(s) = 1$ and $\varphi(t) = 0.8$. On the other hand, it can be proved by induction on the structure of the formula that, if $\dot{+}$ and $\dot{-}$ are not used, there is no QLTL formula that distinguishes between s and t.

3.3 Computing the Linear Distance

Given a finite QTS $\mathscr{S} = (S, \tau, \Sigma, [\cdot])$ we wish to compute $ld^{x}_{\alpha}(s_0, t_0)$, for all $s_0, t_0 \in S$, all $x \in \{\mathrm{a}, \mathrm{s}\}$, and all $\alpha \in (0,1]$ (the case $\alpha = 0$ is trivial). We describe the computation of ld^{a}, as the computation of ld^{s} is analogous. We can read the definition of ld^{a} as a two-player game. Player 1 chooses a path $\pi = s_0 s_1 s_2 \cdots$ from s_0; Player 2 chooses a path $\pi' = t_0 t_1 t_2 \cdots$ from t_0; the goal of Player 1 (resp. Player 2) is to maximize (resp. minimize) $\sup_k \alpha^k pd(\pi_k, \pi'_k)$. The game is played with partial information: after $s_0 \cdots s_n$, Player 1 must choose s_{n+1} without knowledge[4] of $t_0 \cdots t_n$. Such

[4] Indeed, if the game were played with total information, we would obtain the branching distances of the next section.

a game can be solved via a variation of subset construction [14]. The key idea is to associate with each final state s_n of a finite path $s_0 s_1 \cdots s_n$ chosen by Player 1, all final states t_n of finite paths $t_0 t_1 \cdots t_n$ chosen by Player 2, each labeled by the distance $v(s_0 \cdots s_n, t_0 \cdots t_n) = \max_{0 \leq k \leq n} \alpha^{k-n} pd(s_k, t_k)$.

From $\mathscr{S}$, we construct another QTS $\mathscr{S}' = (S', \tau', \{r\}, [\cdot]')$, having set of states $S' = S \times 2^{S \times \mathbb{D}}$. If $\alpha = 1$ we can take $\mathbb{D} = \{pd(s,t) \mid s,t \in S\}$, so that $|\mathbb{D}| \leq |S|^2$. For $\alpha \in (0,1]$, we take $\mathbb{D} = \{pd(s,t)/\alpha^k \mid s,t \in S \wedge k \in \mathbb{N} \wedge pd(s,t) \leq \alpha^k\} \cup \{1\}$, so that $|\mathbb{D}| \leq |S|^2 \cdot \lceil \log_\alpha \min\{pd(s,t) \mid s,t \in S \wedge pd(s,t) > 0\} \rceil + 1$. The transition relation τ' consists of all pairs $(\langle s, C\rangle, \langle s', C'\rangle)$ such that $s' \in \tau(s)$ and $C' = \{\langle t', v'\rangle \mid \exists \langle t, v\rangle \in C \,.\, t' \in \tau(t) \wedge v' = (v/\alpha \sqcup pd(s', t')) \sqcap 1\}$. Note that only Player 1 has a choice of moves in this game, since the moves of Player 2 are accounted for by the subset construction. Finally, the interpretation $[\cdot]'$ is given by $[\langle s, C\rangle]'(r) = \min\{v \mid \langle t, v\rangle \in C\}$, so that r indicates the minimum distance achievable by Player 2 while trying to match a path to $\langle s, C\rangle$ chosen by Player 1. The goal of the game, for Player 1, consists in reaching a state of $\mathscr{S}'$ with the highest possible (discounted) value or r. Thus, for all $s,t \in S$, we have $ld^x_\alpha(s,t) = [\![\exists \Diamond_\alpha r]\!]_{\mathscr{S}'}(\langle s, \{\langle t, pd(s,t)\rangle\}\rangle)$, where the right-hand side is to be computed on $\mathscr{S}'$. This expression can be evaluated by a depth-first traversal of the state space of $\mathscr{S}'$, noting that no state of $\mathscr{S}'$ needs to be visited twice, as subsequent visits do not increase the value of $\Diamond_\alpha r$.

Theorem 6 *For all $x \in \{\mathrm{a}, \mathrm{s}\}$, the following assertions hold:*

1. *Computing ld^x_α for $\alpha \in [0,1]$ and QTS $\mathscr{S}$ is PSPACE-complete in $|\mathscr{S}| + |\alpha|_b$.*
2. *Computing ld^x_α for $\alpha \in [0,1]$ and deterministic QTS $\mathscr{S}$ is PSPACE-complete in $|\mathscr{S}| + |\alpha|_b$.*
3. *Computing ld^x_α for $\alpha \in [0,1]$ and boolean, deterministic QTS $\mathscr{S}$ is in time $O(|\mathscr{S}|^4)$.*

The upper complexity bound for part 1 comes from the above algorithm; the lower bound comes from a reduction from the corresponding result for trace inclusion [15]. Part 2 states that, unlike in the boolean case, the problem remains PSPACE-complete even for deterministic QTSs. This result is proved by a reduction to the nondeterministic case: by introducing perturbations in the valuations, we can tranform a nondeterministic QTS into a deterministic one; for appropriately small perturbations, the distances computed on the derived deterministic QTS enable the determination of the distances over the nondeterministic QTS. Finally, part 3 is a consequence of Theorems 13 and 12.

4 Branching Distances and Logics

Definition 4 (branching distances) Consider the following four equations involving the function $d : S^2 \to [0,1]$ and the parameter $\alpha \in [0,1]$.

$$\text{(Aa)} \qquad d(s,t) = pd(s,t) \sqcup \alpha \cdot \max_{s' \in \tau(s)} \min_{t' \in \tau(t)} d(s',t')$$

$$\text{(As)} \qquad d(s,t) = \overline{pd}(s,t) \sqcup \alpha \cdot \max_{s' \in \tau(s)} \min_{t' \in \tau(t)} d(s',t')$$

$$\text{(Sa)} \qquad d(s,t) = pd(s,t) \sqcup \alpha \cdot \max_{s' \in \tau(s)} \min_{t' \in \tau(t)} d(s',t') \sqcup \alpha \cdot \max_{t' \in \tau(t)} \min_{s' \in \tau(s)} d(s',t')$$

$$\text{(Ss)} \qquad d(s,t) = \overline{pd}(s,t) \sqcup \alpha \cdot \max_{s' \in \tau(s)} \min_{t' \in \tau(t)} d(s',t') \sqcup \alpha \cdot \max_{t' \in \tau(t)} \min_{s' \in \tau(s)} d(s',t')$$

For $x \in \{\mathrm{Aa}, \mathrm{As}, \mathrm{Sa}, \mathrm{Ss}\}$, we define the branching distance bd_α^x as the smallest function $d : S^2 \to [0,1]$ satisfying the equation (x). ∎

The distance bd^{Ss} is related to the metrics of [5, 16, 2]. Clearly $\overline{bd}^{\mathrm{Ss}} = bd^{\mathrm{Ss}}$, so we obtain three symmetrical versions $\overline{bd}^{\mathrm{Aa}}$, $\overline{bd}^{\mathrm{As}}$, and $\overline{bd}^{\mathrm{Sa}}$. For all $\alpha \in [0,1]$, the functions bd_α^{Aa}, bd_α^{As}, and bd_α^{Sa} are directed metrics, and the functions bd_α^{Ss}, $\overline{bd}_\alpha^{\mathrm{Aa}}$, $\overline{bd}_\alpha^{\mathrm{As}}$, and $\overline{bd}_\alpha^{\mathrm{Sa}}$ are pseudometrics.

For $\alpha \in (0,1]$, bd_α^{As} characterizes similarity and bd_α^{Ss} characterizes bisimilarity.

Theorem 7 *For all $\alpha \in (0,1]$, we have $\preceq_{sim} = \mathrm{Zero}(bd_\alpha^{\mathrm{As}})$ and $\approx_{bis} = \mathrm{Zero}(bd_\alpha^{\mathrm{Ss}})$.*

The distance bd^{Aa} corresponds to a variant of simulation where, if $bd_1^{\mathrm{Aa}}(s,t) = 0$ (that is, if s is related to t), then $[s] \leq [t]$. This notion is the quantitative equivalent of a boolean notion of simulation proposed in [10] for the preservation of *positive* ACTL formulas, that is, ACTL formulas where all propositions occur with positive polarity. Indeed, Theorem 8 states that a similar characterization holds for bd^{Aa} in the quantitative setting. Just as similarity in both directions does not imply bisimulation, $\overline{bd}^{\mathrm{As}}$ can be strictly smaller than bd^{Ss}, and $\overline{bd}^{\mathrm{Aa}}$ can be strictly smaller than bd^{Sa}.

Theorem 8 *The relations in Figure 4(b) hold for all QTS and for all $\alpha \in [0,1]$. For $\alpha \in (0,1]$, no other inequalities hold on all QTSs.*

4.1 Quantitative μ-Calculus

We define quantitative μ-calculus after [2]. Given a set of variables X and a set of atomic propositions Σ, the formulas of the *quantitative μ-calculus* are generated by the grammar

$$\begin{aligned}\varphi \ ::= \ & r \mid x \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \neg \varphi \mid c \dot{+} \varphi \mid c \dot{-} \varphi \mid \\ & \exists \bigcirc_\alpha \varphi \mid \exists \ominus_\alpha \varphi \mid \forall \bigcirc_\alpha \varphi \mid \forall \ominus_\alpha \varphi \mid \mu x \,.\, \varphi \mid \nu x \,.\, \varphi\end{aligned}$$

for propositions $r \in \Sigma$, variables $x \in X$, and discount factors $\alpha \in [0,1]$. Denoting by $\mathscr{F} = (S \to [0,1])$, a (variable) interpretation is a function $\mathscr{E} : X \to \mathscr{F}$. Given an interpretation $\mathscr{E}$, a variable $x \in X$ and a function $f \in \mathscr{F}$, we denote by $\mathscr{E}[x := f]$ the interpretation $\mathscr{E}'$ such that $\mathscr{E}'(x) = f$ and, for all $y \neq x$, $\mathscr{E}'(y) = \mathscr{E}(y)$. Given a QTS and an interpretation $\mathscr{E}$, every formula φ of the quantitative μ-calculus defines a valuation $[\![\varphi]\!]_{\mathscr{E}} : S \to [0,1]$:

$$\begin{aligned}
[\![r]\!]_{\mathscr{E}}(s) &= [s](r) \\
[\![x]\!]_{\mathscr{E}} &= \mathscr{E}(x) \\
[\![\varphi_1 \wedge \varphi_2]\!]_{\mathscr{E}} &= [\![\varphi_1]\!]_{\mathscr{E}} \sqcap [\![\varphi_2]\!]_{\mathscr{E}} \\
[\![\varphi_1 \vee \varphi_2]\!]_{\mathscr{E}} &= [\![\varphi_1]\!]_{\mathscr{E}} \sqcup [\![\varphi_2]\!]_{\mathscr{E}} \\
[\![\neg \varphi]\!]_{\mathscr{E}}(s) &= 1 - [\![\varphi]\!]_{\mathscr{E}}(s) \\
[\![c \dot{+} \varphi]\!](s) &= c \dot{+} [\![\varphi]\!](s) \\
[\![c \dot{-} \varphi]\!](s) &= c \dot{-} [\![\varphi]\!](s)
\end{aligned}
\qquad
\begin{aligned}
[\![\exists \bigcirc_\alpha \varphi]\!]_{\mathscr{E}}(s) &= \alpha \cdot \max_{s' \in \tau(s)} [\![\varphi]\!]_{\mathscr{E}}(s') \\
[\![\exists \ominus_\alpha \varphi]\!]_{\mathscr{E}}(s) &= 1 - \alpha + \alpha \cdot \max_{s' \in \tau(s)} [\![\varphi]\!]_{\mathscr{E}}(s') \\
[\![\forall \bigcirc_\alpha \varphi]\!]_{\mathscr{E}}(s) &= \alpha \cdot \min_{s' \in \tau(s)} [\![\varphi]\!]_{\mathscr{E}}(s') \\
[\![\forall \ominus_\alpha \varphi]\!]_{\mathscr{E}}(s) &= 1 - \alpha + \alpha \cdot \min_{s' \in \tau(s)} [\![\varphi]\!]_{\mathscr{E}}(s') \\
[\![\mu x \,.\, \varphi]\!]_{\mathscr{E}} &= \inf\{ f \in \mathscr{F} \mid f = [\![\varphi]\!]_{\mathscr{E}[x:=f]} \} \\
[\![\nu x \,.\, \varphi]\!]_{\mathscr{E}} &= \sup\{ f \in \mathscr{F} \mid f = [\![\varphi]\!]_{\mathscr{E}[x:=f]} \}.
\end{aligned}$$

The existence of the required fixpoints is guaranteed by the monotonicity and continuity of all operators. If φ is closed, we write $[\![\varphi]\!]$ for $[\![\varphi]\!]_{\mathscr{E}}$. A formula is *positive* if all atomic propositions occur in the scope of an even number of negations. For all $\alpha \in [0,1]$, we call $\textsc{ClMuCalc}_\alpha$ the set of closed μ–calculus formulas where all discount factors are smaller than or equal to α and $\textsc{ClMuCalc}^+_\alpha$ the subset of $\textsc{ClMuCalc}_\alpha$ that only contains positive formulas. We denote by $\exists\textsc{ClMuCalc}_\alpha$, $\exists\textsc{ClMuCalc}^+_\alpha$ the respective subsets with no occurrences of $\forall$. For $ops \subseteq \{\bigcirc, \ominus, \Diamond, \Box, \dot{+}, \dot{-}, \mu, \nu, \exists, \forall\}$, we denote by $\textsc{ClMuCalc}_\alpha(ops)$ the set of formulas that only contain boolean connectives and operators in *ops*. Notice that, if we omit the operators $\dot{+}$ and $\dot{-}$ and we take all discount factors to be 1, then the semantics of the quantitative μ-calculus on boolean systems coincides with the one of the classical μ-calculus.

4.2 Logical Characterizations of Branching Distances

The following result shows that the branching distances provide bounds for the corresponding fragments of the μ-calculus.

Theorem 9 *For all QTSs, states s and t, and $\alpha \in [0,1]$, we have*

$$\begin{array}{ll} \text{for all } \varphi \in \exists\textsc{ClMuCalc}^+_\alpha & bd^{\text{Aa}}_\alpha(s,t) \geq [\![\varphi]\!](s) \dot{-} [\![\varphi]\!](t) \\ \text{for all } \varphi \in \exists\textsc{ClMuCalc}_\alpha & bd^{\text{As}}_\alpha(s,t) \geq [\![\varphi]\!](s) \dot{-} [\![\varphi]\!](t) \\ \text{for all } \varphi \in \textsc{ClMuCalc}^+_\alpha & bd^{\text{Sa}}_\alpha(s,t) \geq [\![\varphi]\!](s) \dot{-} [\![\varphi]\!](t) \\ \text{for all } \varphi \in \textsc{ClMuCalc}_\alpha & bd^{\text{Ss}}_\alpha(s,t) \geq |[\![\varphi]\!](s) - [\![\varphi]\!](t)| \end{array}$$

As noted before, each bound of the form $d(s,t) \geq [\![\varphi]\!](s) \dot{-} [\![\varphi]\!](t)$, trivially leads to a bound of the form $\overline{d}(s,t) \geq |[\![\varphi]\!](s) - [\![\varphi]\!](t)|$. The bounds are tight, and the following theorem identifies which fragments of quantitative μ-calculus suffice for characterizing each branching distance.

Theorem 10 *For all QTSs, states s and t, and $\alpha \in [0,1]$, we have*

$$\begin{aligned} bd^{\text{Aa}}_\alpha(s,t) &= \sup\nolimits_{\varphi \in \textsc{ClMuCalc}^+_\alpha(\exists,\ominus,\dot{+})} [\![\varphi]\!](s) \dot{-} [\![\varphi]\!](t), \\ bd^{\text{As}}_\alpha(s,t) &= \sup\nolimits_{\varphi \in \textsc{ClMuCalc}_\alpha(\exists,\ominus,\dot{+})} [\![\varphi]\!](s) \dot{-} [\![\varphi]\!](t), \\ bd^{\text{Sa}}_\alpha(s,t) &= \sup\nolimits_{\varphi \in \textsc{ClMuCalc}^+_\alpha(\exists,\forall,\ominus,\dot{+})} [\![\varphi]\!](s) \dot{-} [\![\varphi]\!](t), \\ bd^{\text{Ss}}_\alpha(s,t) &= \sup\nolimits_{\varphi \in \textsc{ClMuCalc}_\alpha(\exists,\forall,\ominus,\dot{+})} [\![\varphi]\!](s) \dot{-} [\![\varphi]\!](t). \end{aligned}$$

The next result shows that the operator $\dot{+}$ (or $\dot{-}$), which is not present in the ordinary μ-calculus, is necessary to characterize the branching distances. This parallels a result of [5] for a metric related to bd^{Ss} on labeled Markov chains, and a result of [11] for Markov decision processes and games.

Theorem 11 *There is a finite QTS and two states s and t such that, for all $\alpha \in (0,1]$, $bd^{\text{Ss}}_\alpha(s,t) = bd^{\text{As}}_\alpha(s,t) > 0$ and for all $\varphi \in \textsc{ClMuCalc}$ that do not contain $\dot{+}$ and $\dot{-}$, we have $[\![\varphi]\!](s) = [\![\varphi]\!](t)$.*

Proof (sketch). Consider again the QTS in Figure 2 and take $\alpha = 1$. Then $bd^{\text{Ss}}(s,t) = bd^{\text{As}}(s,t) = 0.2$. Theorem 5 states that formulas from $\textsc{Qltl}(\bigcirc, \Diamond)$ are not sufficient for distinguishing s from t. Compared to $\textsc{Qltl}$, the μ–calculus allows to specify branching formulas and take fixpoints of expressions. However, in the example here, these capabilities do not help, since, starting from s or t, the only branching points occurs in the first state. ∎

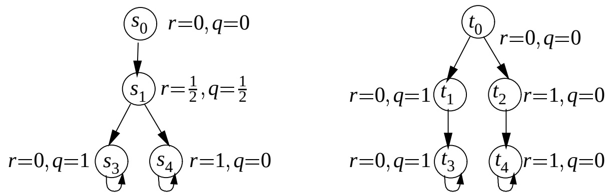

Fig. 3. Linear versus branching distances on a deterministic QTS.

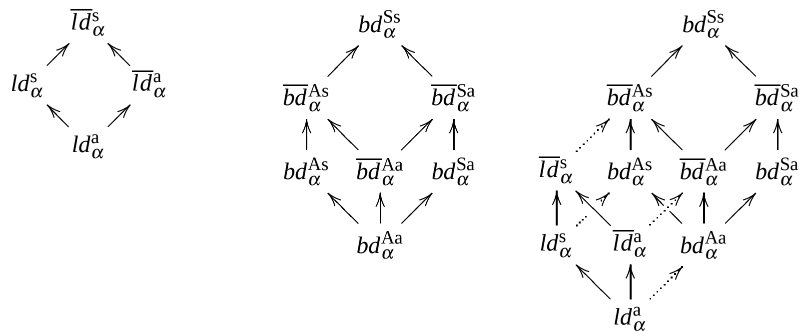

(a) Linear distances. (b) Branching distances. (c) All distances.

Fig. 4. Relations between distances, where $f \to g$ means $f \leq g$. In (c), the dotted arrows collapse to equality for boolean, deterministic QTSs.

4.3 Computing the branching distances

Given a finite QTS $\mathscr{S} = (S, \tau, \Sigma, [\cdot])$ a rational number $\alpha \in [0,1]$, and $x \in \{\mathrm{Ss}, \mathrm{Sa}, \mathrm{As}, \mathrm{Aa}\}$, we can compute $bd^x_\alpha(s,t)$ for all states $s,t \in S$ by computing in an iterative fashion the fixpoints of Definition 4. For instance, bd^{Aa}_α can be computed by letting $d^0(s,t) = 0$ for all $s,t \in S$ and, for $k \in \mathbb{N}$, by letting $d^{k+1}(s,t) = pd(s,t) \sqcup \alpha \cdot \max_{s' \in \tau(s)} \min_{t' \in \tau(t)} d^k(s',t')$, for all $s,t \in S$. Then $bd^x_\alpha = \lim_{k \to \infty} d^k$, and it can be shown that this and the other computations terminate in at most $|S|^2$ iterations. This gives the following complexity result.

Theorem 12 *Computing bd^x_α for $x \in \{\mathrm{Ss}, \mathrm{Sa}, \mathrm{As}, \mathrm{Aa}\}$, $\alpha \in [0,1]$ and a QTS $\mathscr{S}$ can be done in time $O(|\mathscr{S}|^4)$.*

5 Comparing the Linear and Branching Distances

Just as similarity implies trace inclusion, we have both $ld^{\mathrm{a}} \leq bd^{\mathrm{Aa}}$ and $ld^{\mathrm{s}} \leq \overline{bd}^{\mathrm{As}}$; just as bisimilarity implies trace equivalence, we have $\overline{ld}^{\mathrm{s}} \leq bd^{\mathrm{Ss}}$ and $\overline{ld}^{\mathrm{a}} \leq \overline{bd}^{\mathrm{Sa}}$. Moreover, in the non-quantitative setting, trace inclusion (resp. trace equivalence) coincides with (bi-)similarity on deterministic systems. This result generalizes to distances over QTSs that are both deterministic and boolean, but not to distances over QTSs that are just deterministic.

Theorem 13 *The following properties hold.*

1. The relations in Figure 4(c) hold for all $\alpha \in [0,1]$*. Moreover, for* $\alpha \in (0,1]$*, the inequalities cannot be replaced by equalities.*
2. For all boolean, deterministic QTSs, all $\alpha \in [0,1]$*, we have*

$$ld^{\mathrm{a}}_{\alpha} = bd^{\mathrm{Aa}}_{\alpha} \quad ld^{\mathrm{s}}_{\alpha} = bd^{\mathrm{As}}_{\alpha} \quad \overline{ld}^{\mathrm{a}}_{\alpha} = \overline{bd}^{\mathrm{Aa}}_{\alpha} \quad \overline{ld}^{\mathrm{s}}_{\alpha} = \overline{bd}^{\mathrm{As}}_{\alpha}.$$

These equalities need not to hold for non-boolean, deterministic QTSs.

To see that on deterministic, non-boolean QTSs, the linear distances between states can be strictly smaller than the corresponding branching ones, consider the QTS in Figure 3. We assume that $\alpha > \frac{1}{2}$; a similar example works if $\alpha \leq \frac{1}{2}$. Then $ld^{\mathrm{a}}_{\alpha}(s,t) = ld^{\mathrm{s}}_{\alpha}(s,t) = \overline{ld}^{\mathrm{a}}_{\alpha}(s,t) = \overline{ld}^{\mathrm{s}}_{\alpha}(s,t) = \frac{1}{2}\alpha$, while $bd^{\mathrm{Aa}}_{\alpha}(s,t) = bd^{\mathrm{As}}_{\alpha}(s,t) = \overline{bd}^{\mathrm{Aa}}_{\alpha}(s,t) = \overline{bd}^{\mathrm{As}}_{\alpha}(s,t) = \alpha^2$.

References

1. P. Caspi and A. Benveniste. Toward an approximation theory for computerized control. In *EMSOFT 02: 2nd Intl. Workshop on Embedded Software*, volume 2491 of *Lect. Notes in Comp. Sci.*, pages 294–304, 2002.
2. L. de Alfaro, T.A. Henzinger, and R. Majumdar. Discounting the future in systems theory. In *Proc. 30th Int. Colloq. Aut. Lang. Prog.*, volume 2719 of *Lect. Notes in Comp. Sci.*, pages 1022–1037. Springer-Verlag, 2003.
3. L. de Alfaro and R. Majumdar. Quantitative solution of omega-regular games. In *Proc. 33rd ACM Symp. Theory of Comp.*, pages 675–683. ACM Press, 2001.
4. C. Derman. *Finite State Markovian Decision Processes*. Academic Press, 1970.
5. J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden. Metrics for labelled markov systems. In *CONCUR'99: Concurrency Theory. 10th Int. Conf.*, volume 1664 of *Lect. Notes in Comp. Sci.*, pages 258–273. Springer, 1999.
6. J. Filar and K. Vrieze. *Competitive Markov Decision Processes*. Springer-Verlag, 1997.
7. Peter Fletcher and William F. Lindgren. *Quasi-uniform spaces*, volume 77 of *Lecture Notes in Pure and Applied Mathematics*. Marcel Dekker Inc., New York, 1982.
8. M. Huth and M. Kwiatkowska. Quantitative analysis and model checking. In *Proc. 12th IEEE Symp. Logic in Comp. Sci.*, pages 111–122, 1997.
9. D. Kozen. A probabilistic PDL. In *Proc. 15th ACM Symp. Theory of Comp.*, pages 291–297, 1983.
10. C. Loiseaux, S. Graf, J. Sifakis, A. Bouajjani, and S. Bensalem. Property preserving abstractions for the verification of concurrent systems. *Formal Methods in System Design: An International Journal*, 6(1):11–44, January 1995.
11. R. Majumdar. *Symbolic algorithms for verification and control*. PhD thesis, University of California, Berkeley, 2003.
12. Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems: Specification*. Springer-Verlag, New York, 1991.
13. A. McIver and Carroll Morgan. Games, probability, and the quantitative μ-calculus qMμ. In *Proc. of LPAR*, pages 292–310, 2002.
14. J.H. Reif. Universal games of incomplete information. In *11th Annual ACM Symposium on Theory of Computing*, pages 288–308, April, Atlanta, Georgia 1979.
15. L.J. Stockmeyer and A.R. Meyer. Word problems requiring exponential time. In *Proc. 5th ACM Symp. Theory of Comp.*, pages 1–9. ACM Press, 1973.
16. F. van Breugel and J. Worrel. An algorithm for quantitative verification of probabilistic transition systems. In *CONCUR 01: Concurrency Theory. 12th Int. Conf.*, volume 2154 of *Lect. Notes in Comp. Sci.*, pages 336–350, 2001.