Linear and Branching System Metrics ⋆

Luca de Alfaro[1], Marco Faella[1,2], and Mariëlle Stoelinga[3]

[1] CE, University of California, Santa Cruz, USA
[2] Dipartimento di Scienze Fisiche, Università di Napoli “Federico II”, Italy
[3] EWI, University of Twente, the Netherlands

Technical report UCSC-CRL-05-01
April 8, 2005
School of Engineering
University of California, Santa Cruz

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 elements of arbitrary metric spaces. 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 metric transition systems. Finally, we provide algorithms for computing the distances over finite systems, together with matching lower complexity bounds.

⋆ This work is an improved and extended version of the paper “Linear and Branching Metrics for Quantitative Transition Systems” by the same authors, appeared in the *Proceedings of the 31st International Colloquium on Automata, Languages, and Programming* (ICALP), 2004 [3]. This research was supported in part by the NSF CAREER grant CCR-0132780, the NSF grant CCR-0234690, the ONR grant N00014-02-1-0671, and the EU grant IST-004527 (ARTIST2).

1 Introduction

We consider *metric transition systems,* which are transition systems in which the predicates, at each state, are interpreted as elements of generic metric spaces. Many examples of metric transition systems have been studied in the literature. As the set $\mathbb{R}$ of real numbers is a metric space (when equipped, for instance, with the metric $d(x, y) = |x - y|$), hybrid systems (where clocks and hybrid variables are interpreted in $\mathbb{R}$) and priced automata (where a real-valued "price" is associated with each state) are both examples of metric transition systems. Kripke structures are also a special case of metric transition systems, as the set $\{\mathrm{T}, \mathrm{F}\}$ of boolean values can be associated with the metric $d(\mathrm{T}, \mathrm{T}) = d(\mathrm{F}, \mathrm{F}) = 0$, and $d(\mathrm{T}, \mathrm{F}) = d(\mathrm{F}, \mathrm{T}) = 1$. Indeed, it is difficult to think of a class of transition systems that has been proposed in the literature, and that *cannot* be cast as a metric transition system.

Trace inclusion, trace equivalence, simulation, and bisimulation are classical system relations which play a very important role in system specification and verification. These notions are defined in terms of the *equality* of predicate valuations: for example, trace inclusion holds between two states s, t if, for every trace from s, we can find a trace from t with equal predicate valuations. Once the predicate valuations belong to metric spaces, it becomes natural to extend these system relations to metrics, that capture how close the valuations are, rather than requiring equality. For example, trace inclusion can be generalized to a metric, where the distance between s and t provides a bound for how closely the valuations of an arbitrary trace from s can be matched by a trace from t. Following this idea, we extend the classical relations of trace inclusion, trace equivalence, simulation, and bisimulation to a metric setting, by defining linear and branching *distances*[4]. Considering distances, rather than relations, leads to a theory of system approximations [7, 18, 2], enabling the quantification of how closely a concrete system implements a specification. System metrics, rather than relations, are also appropriate when the system structure is derived from experimental observations, so that the predicate valuations are subject to measurement errors. In this case, system metrics provide useful information about the similarity of system behaviors, while relations, relying on equality in predicate valuations, are unnecessarily fine-grained.

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. The distance notions arising as generalizations of trace inclusion and simulation are asymmetrical, just like the relations they generalize: the "simulation distance" between s and t is in general different from the "simulation distance" between t and s. We call these asymmetrical distances *directed metrics,* preferring this term

[4] In this paper, we use the term "distance" in a generic way, applying it to various types of metrics.

to the term *quasi-pseudometrics* used elsewhere in the literature [9]; symmetrical distances will be called *undirected metrics.* Thus, for the sake of generality, we develop our results in the general setting where predicates are evaluated in spaces endowed with directed metrics.

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

We relate the linear distance to the logic QLTL, a quantitative version of LTL [14]. When interpreted on a metric transition system, QLTL formulas yield a value in the positive reals extended with infinity, or $\mathbb{R}_+ \cup \{\infty\}$. The propositional formulas of QLTL are of the form $D(r,c)$ and $D(c,r)$, where r is a predicate, and c a constant. The formula $D(r,c)$, at a state, yields the distance of the valuation of r at the state from the constant c. Both $D(r,c)$ and $D(c,r)$ are present as basic formulas, since in our setting based on directed distances, the distance between the valuation of r and c, and the distance between c and the valuation of r, need not be the same. The formula "next p" returns the (quantitative) value of the subformula 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."

In the standard, relational 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 our metric 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 give two kinds of characterizations. We show that for arbitrary metric transition systems, the distances provide a bound for the difference in value of QLTL formulas: precisely, for all states s,t we have $|\varphi(t) - \varphi(s)| \leq \overline{ld}^{\rm s}(s,t)$ and $\varphi(t) - \varphi(s) \leq ld^{\rm s}(s,t)$. Moreover, we show that for finitely branching metric transition systems, such characterizations are tight: for all states s,t we have $\overline{ld}^{\rm s}(s,t) = \sup_{\varphi \in \text{QLTL}} |\varphi(t) - \varphi(s)|$ and $ld^{\rm s}(s,t) = \sup_{\varphi \in \text{QLTL}} (\varphi(t) - \varphi(s))$. This tightness result does not hold in general for non-finitely-branching metric transition systems.

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 $bd^{\rm As}$

over states. If $bd^{\mathrm{As}}(s,t) = b$, then the valuations of corresponding predicates at s and t differ by at most 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 [7, 18]. We relate these distances to QMU, a quantitative fixpoint calculus that closely resembles the μ-calculus of [4], and is related to the calculi of [11, 5] (see also [10, 15]). Similarly to QLTL, the basic formulas of QMU are of the form $D(r,c)$ and $D(c,r)$, for a predicate r and a valuation c. The modal formulas $\forall\bigcirc p$, $\exists\bigcirc p$ compute respectively the least and greatest value of a subformula p at all successor states; the logical connectives "and" and "or" are interpreted as "min" and "max", and the fixpoints are given a quantitative interpretation.

Again, we provide a twofold logical characterization of the branching distances in terms of QMU. We show that for arbitrary metric transition systems, we have $|\varphi(t) - \varphi(s)| \leq bd^{\mathrm{Ss}}(s,t)$ and $\psi(t) - \psi(s) \leq bd^{\mathrm{As}}(s,t)$, where φ is any QMU-formula, and ψ is any "universal" QMU-formula, i.e., any formula of QMU which does not contain $\exists\bigcirc$. Moreover, if the metric transition system is finitely branching, then we have the stronger result $bd^{\mathrm{Ss}}(s,t) = \sup_{\varphi\in\mathrm{QMU}} |\varphi(t) - \varphi(s)|$ and $bd^{\mathrm{As}}(s,t) = \sup_{\varphi\in\exists\mathrm{QMU}}(\varphi(t) - \varphi(s))$, where $\exists$QMU is the fragment of QMU in which $\exists\bigcirc$ does not occur; these results do not hold in general for non-finitely-branching metric transition systems.

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, we show that determinism plays a lesser role in the quantitative setting than in the standard boolean setting: 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 [4].

We present all 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. [8]) and optimal control (see e.g. [6]), and it leads to robust theories of quantitative systems [4]. In the discouned setting, behavioral differences arising far into the future are given less relative weight than behavioral differences affecting the present or the near future. Hence, the discounted setting leads to notions of "local similarity" that enjoy many pleasant mathematical properties.

2 Preliminaries

We denote by $\mathbb{R}$ the set of real numbers, by $\mathbb{R}_+$ the set of non-negative reals and we set $\widehat{\mathbb{R}}_+ = \mathbb{R}_+ \cup \{\infty\}$. We extend the operations $+, -, \cdot$ to $\widehat{\mathbb{R}}_+$ as usual: namely, $\infty - \infty = 0$, $\infty + \infty = \infty$, and $\infty \pm x = \infty$ for all $x \in \mathbb{R}$, $\infty \cdot x = \infty$ for $x \in \mathbb{R} \setminus \{0\}$. For two numbers $x, y \in \widehat{\mathbb{R}}_+$, we write $x \sqcup y = \max(x, y)$ and $x \sqcap y = \min(x, y)$. We lift the operators $\sqcup$ and $\sqcap$, and the relations $<$, $\leq$ to functions via their pointwise extensions. Precisely, for n-argument functions $f_1, f_2 : A_1 \times \cdots \times A_n \to B$, we write $f_1 \sqcup f_2$ for the function $g : A_1 \times \cdots \times A_n \to B$ defined by $g(x_1, \ldots, x_n) = f_1(x_1, \ldots, x_n) \sqcup f_2(x_1, \ldots, x_n)$, and similarly for $\sqcap$; we write $f_1 \leq f_2$ if $f_1(x_1, \ldots, x_n) \leq f_2(x_1, \ldots, x_n)$ for all $x_1 \in A_1$, $\ldots$, $x_n \in A_n$, and we write $f_1 < f_2$ if $f_1 \leq f_2$ and if there are some $x_1 \in A_1$, $\ldots$, $x_n \in A_n$ for which $f_1(x_1, \ldots, x_n) < f_2(x_1, \ldots, x_n)$. Given a function $d : X^2 \to \widehat{\mathbb{R}}_+$, we denote by $\mathrm{Zero}(d) = \{(x, y) \in X^2 \mid d(x, y) = 0\}$ its zero set. Given a sequence $\{x_i\}_{i \in \mathbb{N}}$, we commonly write $\lim_i x_i$ (resp. $\sup_i X_i$, $\inf_i x_i$) for $\lim_{i \to \infty} x_i$ (resp. $\sup_{i \to \infty} X_i$, $\inf_{i \to \infty} x_i$). The following lemma summarizes some simple facts about sequences of real numbers that will be needed in subsequent proofs.

Lemma 1 *Let $\mathcal{I}$ be a set and $\{x_i\}_{i \in \mathcal{I}}$, $\{y_i\}_{i \in \mathcal{I}}$ be two families of numbers in $\mathbb{R}$. The following assertions hold.*

1. *If $x_i - y_i \leq c$ for all $i \in \mathcal{I}$, then $\sup_i x_i - \sup_i y_i \leq c$ and $\inf_i x_i - \inf_i y_i \leq c$.*
2. *Let X, Y be sets and $f : X \times Y \to \mathbb{R}$ be a function. Then*

$$\sup_{x \in X} \inf_{y \in Y} f(x, y) \leq \inf_{y \in Y} \sup_{x \in X} f(x, y). \quad \square$$

2.1 Metrics and metric spaces

We define *directed and undirected metrics,* where undirected metrics are required to be symmetrical and directed metrics are not. For example, the travel distance between two points in a city with one-way streets is a directed metric. Our directed and undirected metrics generalize the usual metrics, in that elements that have metric 0 are not required to be identical. This terminology, used throughout the paper, differs somewhat from the standard one: directed metrics have been called *generalized pseudometrics* [9]. We prefer the term "directed", as it is more specific, and parallels the distinction between directed and undirected graphs. The definitions are as follows.

Definition 1 We introduce the following terminology.

1. A *directed metric* on a set X is a function $d : X \times X \to \widehat{\mathbb{R}}_+$ that satisfies
   - $d(x, x) = 0$ for all $x \in X$;
   - $d(x, z) \leq d(x, y) + d(y, z)$ for all $x, y, z \in X$ (triangle inequality).
2. An *undirected metric* is a directed metric $d : X \times X \to \widehat{\mathbb{R}}_+$ that is symmetrical, that is, such that $d(x, y) = d(y, x)$ for all $x, y \in X$. Undirected metrics are also called simply *metrics.* $\square$

We will often define a directed metrics, and obtain the corresponding undirected metrics by *symmetrization.*

Definition 2 (symmetrization) Given a directed metric d on a set X, we denote by $\bar{d}$ its *symmetrization,* defined by $\bar{d}(x, y) = d(x, y) \sqcup d(y, x)$ for all $x, y \in X$. Obviously, for all $x, y \in X$, we have $d(x, y) \leq \bar{d}(x, y)$. $\square$

In a Kripke structure, the value of a proposition, at each state, is a member of the truth-value set $\{\text{T}, \text{F}\}$. We extend this setting by evaluating propositions, at each state, to elements of *metric spaces.* A metric space is a set with a metric defined on it; for the sake of generality, we assume only that the metric is a directed metric.

Definition 3 A *directed metric space*, or shortly a metric space, is a pair (X, d), where d is a directed metric on X. $\square$

Example 1 An example of a metric space is the space of RGB-represented colors, where the distance between colors c_1 and c_2 represents the difference in brightness between c_1 and c_2. The space is then $X = [0, 1]^3$, and for $\boldsymbol{x} = \langle x_1, x_2, x_3 \rangle$ and $\boldsymbol{y} = \langle y_1, y_2, y_3 \rangle$ we define $d(\boldsymbol{x}, \boldsymbol{y}) = |\boldsymbol{x} \cdot \boldsymbol{b} - \boldsymbol{y} \cdot \boldsymbol{b}|$, where $\boldsymbol{b}$ is a vector giving the brightness of each basic color, and $\cdot$ is the internal product. It is easy to see that (X, d) is a directed metric space. In particular, d is undirected, and note that different colors may have the same brightness. $\square$

Example 2 Another example of metric space is $\mathbf{X}_{\mathbb{R}} = (\mathbb{R}, d_{\mathbb{R}})$, with $d(x, y) = x \dot{-} y \stackrel{\text{def}}{=} \max\{x - y, 0\}$ for $x, y \in \mathbb{R}$. It is immediate that d is a directed metric. $\square$

Example 3 A particularly simple example of metric space is $\mathbf{X}_{\mathbb{B}} = (X, d_{\mathbb{B}})$ is $X = \{0, 1\}$ and $d(x, y) = |x - y|$ for $x, y \in \{0, 1\}$. This is the usual space of "boolean" valuations; it is immediate that d is an undirected metric. $\square$

When providing logical characterizations for the distances, we will first consider logics in which any element of the metric space can be used as a constant. If the metric space is uncountable, however, this leads to the consideration of logics with uncountably many symbols. If a metric space is *separable,* each element can be approximated by arbitrarily close elements of a *countable basis.* In this case, we will see that logics with countably many symbols (corresponding to the elements of the basis) will suffice.

Definition 4 (separable directed metric space) A directed metric space (X, d) is *separable* if there is a countable *basis* $\mathcal{B} \subseteq X$ such that, for all $x \in X$ and all $\varepsilon > 0$, there is $y \in \mathcal{B}$ with $d(x, y) < \varepsilon$ and $d(y, x) < \varepsilon$. $\square$

2.2 Metric transition systems

A *metric transition system* is a transition system where the value of a proposition, at each state, is an element of a directed metric space. To simplify the notation, we assume throughout the paper an underlying set AP of propositions, where each proposition $r \in \Sigma$ takes values in a metric space (X_r, d_r).

Definition 5 (valuations) A *valuation* u of a set $\Sigma \subseteq AP$ of propositions is a function with domain Σ that assigns to each $r \in \Sigma$ an element $q \in X_r$ of the metric space (X_r, d_r) corresponding to r. We denote by $\mathcal{U}[\Sigma]$ the set of all valuations of Σ. □

Definition 6 (metric transition system) A *metric transition system* (MTS) is a tuple $M = (S, \tau, \Sigma, [\cdot])$ consisting of the following components:

- a set S of states;
- a transition relation $\tau \subseteq S \times S$;
- a finite set $\Sigma \subseteq AP$ of propositions;
- a function $[\cdot]\colon S \to \mathcal{U}[\Sigma]$ which assigns to each state $s \in S$ a valuation.

For a state $s \in S$, we write $\tau(s)$ for $\{t \in S \mid (s,t) \in \tau\}$. We require that M is non-blocking: for all $s \in S$, the set $\tau(s)$ is non-empty. □

We distinguish the special classes of *deterministic* and *finitely branching* MTSs.

Definition 7 (special types of MTSs) Let $M = (S, \tau, [\cdot])$ be a MTS.

- We say that M 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 $[t](r) \neq [t'](r)$.
- We say that M is *finitely branching* if $\tau(s)$ is finite for all $s \in S$.
- We say that M is *separable* if, for all $r \in \Sigma$, the metric space (q_r, d_r) is separable. In this case, we denote by $\mathcal{B}_r$ a countable basis for (q_r, d_r). □

2.3 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.

Definition 8 (paths and traces) Consider an MTS $M = (S, \tau, \Sigma, [\cdot])$. A *path* of M is an infinite sequence of states $\pi \in S^\omega$ such that $(\pi_i, \pi_{i+1}) \in \tau$ for all $i \in \mathbb{N}$. Given a state $s \in S$, we write $\mathit{Paths}_M(s)$ for the set of all paths of M starting from s; we omit the subscript M when clear from the context.

A *trace* is an infinite sequence $\sigma \in \mathcal{U}[\Sigma]^\omega$. Every path π of M induces a trace $[\pi] = [\pi_0][\pi_1][\pi_2] \cdots$. We write $\mathit{Tr}_M(s) = \{[\pi] \mid \pi \in \mathit{Paths}_M(s)\}$ for the set of traces of M starting from the state $s \in S$, and we omit the subscript M when clear from the context. □

2.4 Branching and trace relations

We define simulation, bisimulation, trace containment, and trace equivalence for MTSs as usual.

Definition 9 ((bi)simulation, trace containment/equivalence) For an MTS $M = (S, \tau, \Sigma, [\cdot])$, the simulation relation $\preceq_{sim}$ (resp. the bisimulation relation $\approx_{bis}$) is the largest relation $R \subseteq S \times S$ such that, for all $s\,R\,t$, the following Conditions 1 and 2 (resp. 1, 2, and 3) hold:

1. $[s] = [t]$;
2. for all $s' \in \tau(s)$, there is $t' \in \tau(t)$ with $s' \, R \, t'$;
3. for all $t' \in \tau(t)$, there is $s' \in \tau(s)$ with $s' \, R \, t'$.

For $s, t \in S$, we write $s \sqsubseteq_{tr} t$ if $Tr(s) \subseteq Tr(t)$, and $s \equiv_{tr} t$ if $Tr(s) = Tr(t)$. □

2.5 Discussion

We note that, for some of the results on system metrics, it would have been sufficient to define a metric transition system as a system that maps each state into an element of a metric space, bypassing thus the introduction of a set of predicates, and the related machinery. Such a definition, of course, is a special case of the one we adopt, and corresponds to considering metric transition systems with only one proposition. The main function of predicates is to enable us to develop the connection between system metrics and logics, since the logics refer to quantities via the predicates.

In an MTS $(S, \tau, \Sigma, [\cdot])$, we call each $r \in \Sigma$ a "proposition", rather than "variable", in spite of the fact that r takes values in a generic metric space (X_r, d_r), rather than in the set of truth-values. Our choice of terminology is motivated by the fact that in the system logics we consider, the symbol r plays a (syntactic) role that is analogous to that of ordinary propositions. We reserve instead the term "variable" for the variables used to construct fixpoint expressions in μ-calculus.

3 Linear Distances and Logics

3.1 Linear distances

Throughout the paper, unless specifically noted, we consider a fixed MTS $M = (S, \tau, \Sigma, [\cdot])$. We proceed by defining the linear distances between valuations, then between traces and finally between states. The propositional distance between two valuations is the maximum difference in their proposition evaluations, where differences in the assignments of proposition r are measured by the metric d_r.

Definition 10 (propositional distance) We define the *propositional distance* $pd : \mathcal{U}[\Sigma]^2 \to \widehat{\mathbb{R}}_+$, for all valuations $u, v \in \mathcal{U}[\Sigma]$, as $pd(u, v) = \max_{r \in \Sigma} d_r(u(r), v(r))$. □

For ease of notation, we write $pd(s, t)$ for $pd([s], [t])$. If d_r is a distance for each $r \in \Sigma$, then given $u, v \in \mathcal{U}[\Sigma]$ we have $(u, v) \in \text{Zero}(\overline{pd})$ iff $u = v$, and $(u, v) \in \text{Zero}(pd)$ iff $u \leq v$. The trace distance is the pointwise extension of the propositional distance to infinite sequences of valuations, where the value at position i is discounted by α^i, for a *discount factor* $\alpha \in (0, 1]$.

Definition 11 (trace distance) We define the *trace distance* $td_\alpha : \mathcal{U}[\Sigma]^\omega \times \mathcal{U}[\Sigma]^\omega \to \widehat{\mathbb{R}}_+$ by letting, for $\sigma, \rho \in \mathcal{U}[AP]^\omega$ and $\alpha \in (0, 1]$, $td_\alpha(\sigma, \rho) = \sup_{i \in \mathbb{N}} \alpha^i pd(\sigma_i, \rho_i)$. □

It is easy to show that td_α is a directed metric. The following result states that if we base the notion of trace distance on $\overline{pd}$ instead of on pd (i.e. if we replace pd by $\overline{pd}$ in the definition above), we obtain the symmetrization $\overline{td}_\alpha$ of td_α. Moreover, the kernel of this symmetrization is trace equality.

Lemma 1. *For all sequences $\sigma, \rho \in \mathcal{U}[\Sigma]^\omega$ and all $\alpha \in (0,1]$, we have $\overline{td}_\alpha(\sigma,\rho) = \sup_{i\in\mathbb{N}} \alpha^i \overline{pd}(\sigma_i,\rho_i)$ and $(\sigma,\rho) \in \text{Zero}(\overline{td}_\alpha) \iff \sigma = \rho$.*

The linear distances between two states are obtained by lifting trace distances to the set of all traces from the two states, as in the definition of Hausdorff distance between sets.

Definition 12 (linear distance) We define the two *linear distances* ld^{a} and ld^{s} over S as follows, for $s, t \in S$ and $\alpha \in (0,1]$:

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

One can easily check that, for all $\alpha \in (0,1]$, the functions ld^{a}_α and ld^{s}_α are directed metrics, while $\overline{ld}^{\mathrm{a}}_\alpha$ and $\overline{ld}^{\mathrm{s}}_\alpha$ are undirected ones. Intuitively, the distance ld^{s}_α is a quantitative extension of trace containment: for $s, t \in S$, the distance $ld^{\mathrm{s}}_\alpha(s,t)$ measures how closely (in a quantitative sense) a trace from s can be simulated a trace from t. The symmetrization of ld^{s}_α is $\overline{ld}^{\mathrm{s}}_\alpha$, which is related to trace equivalence. Indeed, we will see in the next section that it is possible to define a quantitative logic QLTL such that the valuation of QLTL formulas at s and t can differ by at most $\overline{ld}^{\mathrm{s}}_\alpha(s,t)$, and similarly, the valuation of any QLTL formula at t is at most $ld^{\mathrm{s}}_\alpha(s,t)$ below the valuation at s.

Example 4 Consider the case where $(X_r, d_r) = \mathbf{X}_{\mathbb{R}}$ for all $r \in \Sigma$, that is, all propositions are interpreted as real numbers, and $d_r(a,b)$ is a measure of how much greater is a than b. In this setting, for $\alpha = 1$ the distances ld^{a}_α and $\overline{ld}^{\mathrm{a}}_\alpha$ have the following intuitive characterization. For a trace $\sigma \in \mathcal{U}[\Sigma]^\omega$ 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 $ld^{\mathrm{a}}_1(s,t) = c$ then for every trace σ from s there is a trace ρ from t such that $\rho \geq \sigma \dot{-} c$. This means that $ld^{\mathrm{a}}_1(s,t)$ 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 $\sigma \dot{-} c$. $\square$

Theorem 1 *For all finitely branching MTSs $(S,\tau,\Sigma,[\cdot])$ and for all $\alpha \in (0,1]$, we have $\sqsubseteq_{tr} = \text{Zero}(ld^{\mathrm{s}}_\alpha)$ and $\equiv_{tr} = \text{Zero}(\overline{ld}^{\mathrm{s}}_\alpha)$.*

Proof. Let $(S,\tau,\Sigma,[\cdot])$ be an MTS with $s,t \in S$ and $\alpha \in (0,1]$. It is easy to see that $s \sqsubseteq_{tr} t$ implies $ld^{\mathrm{s}}_\alpha(s,t) = 0$. To prove the converse, assume that $ld^{\mathrm{s}}_\alpha(s,t) = 0$ and let $\sigma \in Tr(s)$. Then, there are traces $\rho_0, \rho_1, \rho_2 \ldots \in Tr(t)$ such that $\overline{td}_\alpha(\sigma,\rho_i) < \frac{1}{2^i}$ for all i. Due to the finitely branching property, there exists a trace ρ^* such that $\overline{td}_\alpha(\sigma,\rho^*) < \frac{1}{2^i}$ for all i. This means that $\overline{td}_\alpha(\sigma,\rho^*) = 0$, which, by Lemma 1, is the same as $\sigma = \rho$. Now, the result for $\equiv_{tr}$ and $\overline{ld}^{\mathrm{s}}_\alpha$ easily follows. $\square$

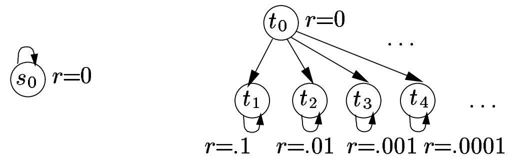

Fig. 1. An MTS showing the difference between $\text{Zero}(ld^{\mathrm{s}}_\alpha)$ and $\sqsubseteq_{tr}$. The proposition r is evaluated in the metric space $\mathbf{X}_{\mathbb{R}}$.

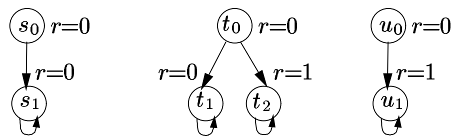

Fig. 2. An MTS showing the difference between ld^{a}_α, ld^{s}_α, $\overline{ld}^{\mathrm{a}}_\alpha$, and $\overline{ld}^{\mathrm{s}}_\alpha$. The proposition r is evaluated in the metric space $\mathbf{X}_{\mathbb{R}}$.

To show that the result above does not hold for infinitely branching systems, consider the MTS in Figure 1, where the proposition r is again evaluated in the metric space $\mathbf{X}_{\mathbb{R}}$. This MTS has infinitely many states s_0, t_0, t_1, $t_2, \ldots$ and transitions (s_0, s_0), (t_0, t_i) and (t_i, t_i) for each $i \in \mathbb{N}$. Moreover, we put $[r](s_0) = [r](t_0) = 0$ and $[r](t_i) = 10^{-i}$ for $i > 0$. Then, we have for all $\alpha \in (0, 1]$ that $(s_0, t_0) \in \text{Zero}(ld^{\mathrm{s}}_\alpha)$, but $s_0 \not\sqsubseteq_{tr} t_0$. To obtain an MTS with $\overline{ld}^{\mathrm{s}}_\alpha(t_0, u_0) = 0$, but $t_0 \not\equiv_{tr} u_0$, we let u_0 be a state that is the exactly same as t_0 (i.e. same valuation and same successor states), except that it has a self-loop (i.e. a transition $(u_0, u_0) \in \tau$).

The relations among linear distances are summarized by the following theorem.

Theorem 2 *The following assertions hold.*

1. *For all MTSs, and for all $\alpha \in (0, 1]$, we have $ld^{\mathrm{a}}_\alpha \leq \overline{ld}^{\mathrm{a}}_\alpha$, $ld^{\mathrm{a}}_\alpha \leq ld^{\mathrm{s}}_\alpha$, $ld^{\mathrm{s}}_\alpha \leq \overline{ld}^{\mathrm{s}}_\alpha$, and $\overline{ld}^{\mathrm{a}}_\alpha \leq \overline{ld}^{\mathrm{s}}_\alpha$; moreover, for $\alpha \in (0, 1]$ the inequalities cannot be replaced by equalities.*
2. *For $\alpha \in (0, 1]$, the distances ld^{s}_α and $\overline{ld}^{\mathrm{a}}_\alpha$ are incomparable: there is an MTS with states $s, t, z \in S$ such that $ld^{\mathrm{s}}_\alpha(s, t) < \overline{ld}^{\mathrm{a}}_\alpha(s, t)$ and $ld^{\mathrm{s}}_\alpha(t, z) > \overline{ld}^{\mathrm{a}}_\alpha(t, z)$.*

Proof. The first and third inequalities of statement (1) are trivial, while the second and fourth follow immediately from the fact that, for all traces σ and ρ, $td(\sigma, \rho) \leq \overline{td}(\sigma, \rho)$. For $\alpha \in (0, 1]$ and the MTS in Figure 2, we have

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

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 $ld^{\rm s}_\alpha \geq \overline{ld}^{\rm a}_\alpha$. □

Next, we show that the linear distances are robust with respect to perturbations in the state valuations: small changes in the proposition valuations causes small changes in the distances. Given two state valuations $[\cdot]_1, [\cdot]_2 : S \to \mathcal{U}[\Sigma]$, we define their directed distance by:

$$d([\cdot]_1, [\cdot]_2) = \sup_{s \in S} \max_{r \in \Sigma} d_r([s]_1(r), [s]_2(r))$$

Moreover, for a state valuation $f : S \to \mathcal{U}[\Sigma]$ and $\alpha \in (0,1]$, we write $ld^{\rm a}_{f,\alpha}$, $ld^{\rm s}_{f,\alpha}$ for the distances defined as in Definition 12, using f as the state valuation.

Theorem 3 (linear distance robustness) *For all $\alpha \in (0,1]$, all predicate valuations $[\cdot]_1, [\cdot]_2$, and all $s, t \in S$, we have*

$$ld^{\rm a}_{[\cdot]_1,\alpha}(s,t) - ld^{\rm a}_{[\cdot]_2,\alpha}(s,t) \leq d([\cdot]_1, [\cdot]_2) + d([\cdot]_2, [\cdot]_1)$$
$$|ld^{\rm s}_{[\cdot]_1,\alpha}(s,t) - ld^{\rm s}_{[\cdot]_2,\alpha}(s,t)| \leq 2 \cdot \overline{d}([\cdot]_1, [\cdot]_2)$$

Proof. The result follows by showing that the trace distance between two traces ρ and σ, measured under $[\cdot]_1$ and $[\cdot]_2$, differs by at most $d([\cdot]_1, [\cdot]_2) + d([\cdot]_2, [\cdot]_1)$. The key step consists in noting that, for any $r \in \Sigma$, from the triangular inequality

$$d_r([s]_1(r), [t]_1(r)) \leq d_r([s]_1(r), [s]_2(r)) + d_r([s]_2(r), [t]_2(r)) + d_r([t]_2(r), [t]_1(r))$$

follows

$$\begin{aligned} d_r([s]_1(r), [t]_1(r)) - d_r([s]_2(r), [t]_2(r)) &\leq d_r([s]_1(r), [s]_2(r)) + d_r([t]_2(r), [t]_1(r)) \\ &\leq d([\cdot]_1, [\cdot]_2) + d([\cdot]_2, [\cdot]_1). \end{aligned}$$

Now the result follows by repetitive application of Lemma 1(1). □

3.2 Quantitative linear-time temporal logic

The linear distances introduced above can be characterized in terms *quantitative linear-time temporal logic* (QLTL), a quantitative extension of linear-time temporal logic [14] which includes quantitative versions of the temporal operators and logic connectives. Following [7], QLTL has a "threshold" operator, enabling the comparison of a formula against a real constant. The QLTL formulas over a set Σ of propositions are generated by the following grammar:

$$\varphi \ ::= \ D(r,c) \mid D(c,r) \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \bigcirc_\alpha \varphi \mid \Diamond_\alpha \varphi$$

Here $r \in \Sigma$ is a proposition, $c \in \bigcup_{r \in AP} X_r$ is a constant and $\alpha \in (0,1]$ a discount factor. We assume that, in a term of the form $D(r,c)$ or $D(c,r)$, we have $c \in X_r$.

A formula φ assigns a value $[\![\varphi]\!](\rho) \in \widehat{\mathbb{R}}_+$ to each trace $\sigma \subseteq \mathcal{U}[\Sigma]^\omega$:

$$\begin{aligned}
[\![D(r,c)]\!](\sigma) &= d_r(\sigma_0(r), c) \\
[\![D(c,r)]\!](\sigma) &= d_r(c, \sigma_0(r)) \\
[\![\varphi_1 \wedge \varphi_2]\!](\sigma) &= [\![\varphi_1]\!](\sigma) \sqcap [\![\varphi_2]\!](\sigma) \\
[\![\varphi_1 \vee \varphi_2]\!](\sigma) &= [\![\varphi_1]\!](\sigma) \sqcup [\![\varphi_2]\!](\sigma) \\
[\![\bigcirc_\alpha \varphi]\!](\sigma) &= \alpha \cdot [\![\varphi]\!](\sigma^1) \\
[\![\diamond_\alpha \varphi]\!](\sigma) &= \sup\{\alpha^i \cdot [\![\varphi]\!](\sigma^i) \mid i \geq 0\}
\end{aligned}$$

A QLTL formula φ assigns a real value $[\![\varphi]\!](s) \in \widehat{\mathbb{R}}_+$ to each state s of a given MTS, by defining

$$[\![\varphi]\!](s) = \inf\{[\![\varphi]\!](\rho) \mid \rho \in Tr(s)\}.$$

We note that the above definition could also be phrased in terms of sup over all traces from s, rather than inf. However, as our setting is based on distances, the inf operator most closely corresponds to the universal quantification over all paths present in the classical definition of LTL semantics.

For $\alpha \in (0,1]$, we denote by QLTL_α the set of formulas containing only discount factors smaller than or equal to α. Furthermore, for $ops \subseteq \{\bigcirc, \diamond, D(c,r), D(r,c)\}$, we denote by $\text{QLTL}_\alpha \setminus ops$ the set of formulas which do not employ the operators in ops.

Notice that QLTL is a proper extension to the fragment of LTL without the Until operator, in the following sense. Consider the metric space $E = (\{0,1\}, \lambda xy.|x-y|)$. Any Kripke structure M has an obvious translation to an MTS M' over E. Moreover, any LTL formula φ in positive normal form can be translated into a QLTL formula φ' by adding the discount factor 1 as a subscript to all temporal operators and replacing r and $\neg r$ with $d(r,0)$ and $d(r,1)$, respectively. Then, φ is true on a Kripke structure M if and only if φ' evaluates to 1 on M'.

3.3 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 2 *For all MTSs $(S, \tau, \Sigma, [\cdot])$, all $\alpha \in (0,1]$ and traces $\sigma, \rho \in \mathcal{U}[\Sigma]^\omega$, the following holds.*

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

Proof. Let us consider the first assertion. We proceed by structural induction on φ. If $\varphi = D(c,r)$, using triangle inequality we get $[\![\varphi]\!](\rho) - [\![\varphi]\!](\sigma) =$

$$d(c, [\rho_0](r)) - d(c, [\sigma_0](r)) \leq d([\sigma_0](r), [\rho_0](r)) \leq pd(\sigma_0, \rho_0) \leq td_\alpha(\sigma, \rho).$$

If $\varphi = \Diamond_\alpha \psi$, by inductive hypothesis we have that, for all $i \in \mathbb{N}$, $[\![\psi]\!](\rho^i) - [\![\psi]\!](\sigma^i) \leq td_\alpha(\rho^i, \sigma^i)$ and thus $\alpha^i \cdot [\![\psi]\!](\rho^i) - \alpha^i \cdot [\![\psi]\!](\sigma^i) \leq \alpha^i \cdot td_\alpha(\rho^i, \sigma^i) \leq td_\alpha(\rho, \sigma)$. Then, by Lemma 1,

$$[\![\varphi]\!](\rho) - [\![\varphi]\!](\sigma) = \sup_{i \in \mathbb{N}} \alpha^i \cdot [\![\psi]\!](\rho^i) - \sup_{j \in \mathbb{N}} \alpha^j \cdot [\![\psi]\!](\sigma^j) \leq td_\alpha(\rho, \sigma).$$

Similar observations hold for the remaining cases.

The second assertion can be proved in symmetrical fashion. The third assertion can be easily proved along similar lines. □

The first result of the previous lemma is tight in two respects: both replacing $\text{QLTL}_\alpha \setminus \{D(r,c)\}$ with QLTL_α and replacing $[\![\varphi]\!](\rho) - [\![\varphi]\!](\sigma)$ with $|[\![\varphi]\!](\rho) - [\![\varphi]\!](\sigma)|$ render the result false. The second assertion is also tight in a similar sense. The following theorem uses the linear distances to provide the desired bounds for QLTL.

Theorem 4 *For all MTSs $(S, \tau, \Sigma, [\cdot])$, all $\alpha \in (0, 1]$ and $s, t \in S$, we have: For all $\varphi \in \text{QLTL}_\alpha \setminus \{D(r,c)\}$:*

$$ld^{\mathrm{a}}_\alpha(s,t) \geq [\![\varphi]\!](t) - [\![\varphi]\!](s) \text{ and } \overline{ld}^{\mathrm{a}}_\alpha(s,t) \geq |[\![\varphi]\!](t) - [\![\varphi]\!](s)|;$$

For all $\varphi \in \text{QLTL}_\alpha$:

$$ld^{\mathrm{s}}_\alpha(s,t) \geq [\![\varphi]\!](t) - [\![\varphi]\!](s) \text{ and } \overline{ld}^{\mathrm{s}}_\alpha(s,t) \geq |[\![\varphi]\!](t) - [\![\varphi]\!](s)|.$$

Proof. We first prove that $ld^{\mathrm{a}}_\alpha(s,t) \geq [\![\varphi]\!](t) - [\![\varphi]\!](s)$.

$$\begin{aligned} ld^{\mathrm{a}}_\alpha(s,t) &= \sup_{\sigma \in Tr(s)} \inf_{\rho \in Tr(t)} td_\alpha(\sigma, \rho) \\ &\geq \sup_{\sigma \in Tr(s)} \inf_{\rho \in Tr(t)} ([\![\varphi]\!](\rho) - [\![\varphi]\!](\sigma)) \qquad \text{by Lemma 2,} \\ &= \inf_{\rho \in Tr(t)} [\![\varphi]\!](\rho) - \inf_{\sigma \in Tr(s)} [\![\varphi]\!](\sigma) \\ &= [\![\varphi]\!](t) - [\![\varphi]\!](s). \end{aligned}$$

The result for $\overline{ld}^{\mathrm{a}}_\alpha$ is an immediate consequence. The statements concerning ld^{s}_α and $\overline{ld}^{\mathrm{s}}_\alpha$ follow in a similar way from Lemma 2. □

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

We next show that, for finitely branching systems, QLTL provides a full logical characterization of the linear distances, meaning that the distinguishing power of the logic is exactly the same as the one of the distances. We start with a technical lemma. Given two traces σ and ρ, an integer m and a discount factor α, let the *bounded distance* between σ and ρ be defined as $btd^m_\alpha(\sigma, \rho) = \max_{0 \leq i \leq m} \alpha^i pd(\sigma_i, \rho_i)$. Clearly, $td_\alpha(\sigma, \rho) = \lim_m btd^m_\alpha(\sigma, \rho)$.

Lemma 3 *If the MTS M is finitely branching, then for all traces σ, discount factors $\alpha \in (0,1]$ and $t \in S$, we have*

$$\sup_{m\in\mathbb{N}} \inf_{\rho\in Tr(t)} btd^m_\alpha(\sigma,\rho) = \inf_{\rho\in Tr(t)} \sup_{m\in\mathbb{N}} btd^m_\alpha(\sigma,\rho).$$

Proof. Since the l.h.s. is trivially smaller than or equal to the r.h.s., we are left to prove that $(l.h.s.) \geq (r.h.s.)$. Specifically, we prove that, for all $\epsilon > 0$, $(r.h.s.) \leq (l.h.s.) + \epsilon$. Fix $\epsilon > 0$. For all $m > 0$, there exists $\rho_m \in Tr(t)$ such that

$$btd^m_\alpha(\sigma,\rho_m) \leq \inf_{\rho\in Tr(t)} btd^m_\alpha(\sigma,\rho) + \epsilon.$$

For all $m \geq 0$, let γ_m be the prefix of ρ_m up to the $m+1$-th valuation. The set $\{\gamma_m \mid m \geq 0\}$ can be arranged into a tree that is a subtree of the unrolling of t. Since this tree contains infinitely many nodes and is finitely branching, by König's lemma it must contain an infinite trace $\rho^* \in Tr(t)$. The trace ρ^* has infinitely many prefixes in $\{\gamma_m \mid m \geq 0\}$. Therefore, there is an increasing sequence $(i_m)_{m>0}$ such that, for all $m \geq 0$, γ_{i_m} is a prefix of ρ^*. It follows that

$$\begin{aligned}
(r.h.s.) \leq td_\alpha(\sigma,\rho^*) &= \lim_m btd^m_\alpha(\sigma,\rho^*)\\
&= \lim_m btd^{i_m}_\alpha(\sigma,\rho^*)\\
&\leq \lim_m btd^{i_m}_\alpha(\sigma,\gamma_{i_m})\\
&= \lim_m btd^{i_m}_\alpha(\sigma,\rho_{i_m})\\
&\leq \lim_m \inf_{\rho\in Tr(t)} btd^{i_m}_\alpha(\sigma,\rho) + \epsilon = (l.h.s.) + \epsilon. \quad \square
\end{aligned}$$

The following theorem states which fragment of the logic is necessary to characterize each linear distance. In particular, the operator $\Diamond$ is never needed. Together with Theorem 4, this result constitutes a full characterization of linear distances in terms of QLTL.

Theorem 5 *If an MTS $M = (S, \tau, \Sigma, [\cdot])$ is finitely branching, then for all $\alpha \in (0,1]$ and $s, t \in S$,*

$$\begin{aligned}
ld^{\mathrm{a}}_\alpha(s,t) &= \sup_{\varphi\in \mathrm{QLTL}_\alpha\setminus\{D(r,c),\Diamond\}} [\![\varphi]\!](t) - [\![\varphi]\!](s)\\
\overline{ld}^{\mathrm{a}}_\alpha(s,t) &= \sup_{\varphi\in \mathrm{QLTL}_\alpha\setminus\{D(r,c),\Diamond\}} |[\![\varphi]\!](t) - [\![\varphi]\!](s)|\\
ld^{\mathrm{s}}_\alpha(s,t) &= \sup_{\varphi\in \mathrm{QLTL}_\alpha\setminus\{\Diamond\}} [\![\varphi]\!](t) - [\![\varphi]\!](s)\\
\overline{ld}^{\mathrm{s}}_\alpha(s,t) &= \sup_{\varphi\in \mathrm{QLTL}_\alpha\setminus\{\Diamond\}} |[\![\varphi]\!](t) - [\![\varphi]\!](s)|
\end{aligned}$$

Proof. By Theorem 4, we only need to prove the "$\leq$" part of the equalities. We first prove the statement involving ld^{a}_α. For sake of simplicity, assume $\Sigma = \{r\}$. Let $ld^{\mathrm{a}}_\alpha(s,t) = x$, we show that for all $\epsilon > 0$ there is a formula φ such that

$[\![\varphi]\!](t) - [\![\varphi]\!](s) > x - \epsilon$. Let $\sigma^* \in Tr(s)$ be a path such that $\inf_{\rho \in Tr(t)} td_\alpha(\sigma^*, \rho) > x - \epsilon$. For all $m \geq 0$, we set

$$\varphi_m = \bigvee_{0 \leq i \leq m} \bigcirc_\alpha^i D([\sigma_i^*](r), r),$$

where $\bigcirc_\alpha^i$ stands for i repetitions of the operator $\bigcirc_\alpha$. Intuitively, when formula φ_m is evaluated on a trace σ', it measures the asymmetric distance between σ' and σ^*, up to the m-th step. Obviously, it is $[\![\varphi_m]\!](s) = 0$ for all $m \geq 0$. Then, the value of φ_m on a state s' measures the distance between σ^* and the trace in $Tr(s')$ which is closest to it. For all $t \in S$, it holds that

$$\begin{aligned}
\sup_m [\![\varphi_m]\!](t) = \lim_m [\![\varphi_m]\!](t) &= \lim_m \inf_{\rho \in Tr(t)} \max_{0 \leq i \leq m} \alpha^i D([\sigma_i^*](r), [\rho_i](r)) \\
&= \lim_m \inf_{\rho \in Tr(t)} btd_\alpha^m(\sigma^*, \rho) \\
&= \inf_{\rho \in Tr(t)} td_\alpha(\sigma^*, \rho) \qquad \text{by Lemma 3} \\
&> x - \epsilon.
\end{aligned}$$

Consequently,

$$\begin{aligned}
\sup_{\varphi \in \mathrm{QLTL}_\alpha \setminus \{D(r,c)\}} [\![\varphi]\!](t) - [\![\varphi]\!](s) &\geq \sup_{m \in \mathbb{N}} [\![\varphi_m]\!](t) - [\![\varphi_m]\!](s) \\
&= \sup_{m \in \mathbb{N}} [\![\varphi_m]\!](t) - 0 \\
&> x - \epsilon.
\end{aligned}$$

The statement about $\overline{ld}_\alpha^{\mathrm{a}}$ is an easy consequence: Assume first that $\overline{ld}_\alpha^{\mathrm{a}}(s,t) = ld_\alpha^{\mathrm{a}}(s,t)$. Then,

$$\overline{ld}_\alpha^{\mathrm{a}}(s,t) = \sup_{\varphi \in \mathrm{QLTL}_\alpha \setminus \{D(r,c)\}} [\![\varphi]\!](s) - [\![\varphi]\!](t) \leq \sup_{\varphi \in \mathrm{QLTL}_\alpha \setminus \{D(r,c)\}} |[\![\varphi]\!](s) - [\![\varphi]\!](t)|.$$

If instead $\overline{ld}_\alpha^{\mathrm{a}}(s,t) = ld_\alpha^{\mathrm{a}}(t,s)$, we have

$$\overline{ld}_\alpha^{\mathrm{a}}(s,t) = \sup_{\varphi \in \mathrm{QLTL}_\alpha \setminus \{D(r,c)\}} [\![\varphi]\!](t) - [\![\varphi]\!](s) \leq \sup_{\varphi \in \mathrm{QLTL}_\alpha \setminus \{D(r,c)\}} |[\![\varphi]\!](s) - [\![\varphi]\!](t)|.$$

We now consider the statement about ld_α^{s}. The proof proceeds similarly to the one involving ld_α^{a}, using as distinguishing formula the following.

$$\varphi_m = \bigvee_{0 \leq i \leq m} \bigcirc_\alpha^i D([\sigma_i^*](r), r) \vee \bigcirc_\alpha^i D(r, [\sigma_i^*](r)).$$

Finally, the statement involving $\overline{ld}_\alpha^{\mathrm{s}}$ can be easily obtained from the proof that $ld_\alpha^{\mathrm{s}}(s,t) = \sup_{\varphi \in \mathrm{QLTL}_\alpha} [\![\varphi]\!](t) - [\![\varphi]\!](s)$ and the fact that $\overline{ld}_\alpha^{\mathrm{s}}(s,t) = ld_\alpha^{\mathrm{s}}(s,t) \sqcup ld_\alpha^{\mathrm{s}}(t,s)$. $\square$

The next example shows that finitely branching is necessary for Theorem 5 to hold.

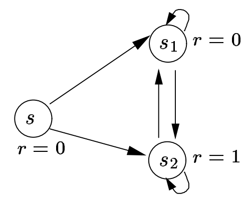

Fig. 3. An MTS exhibiting the language $0\{0,1\}^\omega$; the single predicate is evaluated in the metric space $\mathbf{X}_\mathbb{B}$.

Theorem 6 *There is an infinitely branching MTS such that*

$$ld^{\mathrm{s}}_{\alpha}(s,t) > \sup_{\varphi \in \mathrm{QLTL}_\alpha} [\![\varphi]\!](s) - [\![\varphi]\!](t).$$

Proof. Consider the system in Figure 3, where $\Sigma = \{r\}$. Informally, $Tr(s) = 0\{0,1\}^\omega$. Let σ be a trace such that $\{\sigma\}$ is not a regular language over the alphabet $\{0,1\}$ (it would be sufficient for σ to be not star-free regular). For instance, let $\sigma = 01\,001\,0001\ldots$. Consider a second system, containing a state t such that $Tr(t) = Tr(s) \setminus \{\sigma\}$. Notice that, in order to have such a set of traces, t must be infinitely branching, since if a finitely branching tree contains all prefixes of an infinite path, it must also contain the path itself. We have $ld^{\mathrm{s}}_1(s,t) = 1$. We know that ordinary LTL cannot distinguish s from t, otherwise there would be a formula $\psi \in$ LTL such that $L(\psi) = \{\sigma\}$. We argue that QLTL is also unable to distinguish s from t. To prove it, we have to show that discounting does not give any advantage. □

3.4 Logical characterization via logics with countably many symbols

Above, we have provided a logical characterization for the linear distances in terms of a logic that contains a potentially uncountable set of constants: in general, we need one constant for each element of a metric space corresponding to a predicate. Here, we show how, for separable MTSs, we can provide a characterization in terms of logics with countably many symbols. First, we state a useful result, namely, that the logic is robust with respect to changes in the constants occurring in the formulas: a small change in the constants causes a small change in the value of the formulas.

Theorem 7 *Consider a formula φ of* QLTL *containing the constants $c_1, \ldots, c_n$, belonging respectively to the metric spaces $(q_1, d_1), \ldots, (q_n, d_n)$. Let $\psi = \varphi[c'_1, \ldots, c'_n / c_1, \ldots, c_n]$ be the result of replacing each c_i with c'_i, for $1 \le i \le n$, and let $\delta = \max_{i=1}^{n}(d_i(c_i, c'_i) \sqcup d_i(c'_i, c_i))$ be the maximal distance between the new and old value of each constant. Then, for all $s \in S$, we have $|[\![\varphi]\!](s) - [\![\psi]\!](s)| \le \delta$.*

Proof. The result follows by a straightforward structural induction. The only interesting case is the one for $D(r, c_i)$, for some $1 \le i \le n$; in this case, using the triangular inequality we have

$$|[\![D(r, c_i)]\!](s) - [\![D(r, c_i')]\!](s)| = |d_i([s](r), c) - d_i([s](r), c')| \leq d_i(c', c);$$

the case for $D(c_i, r)$ is similar. □

From the robustness of the logic with respect to the constants, it follows that if an MTS is separable, we can obtain a logical characterization of the linear distances in terms of logics that consist only of countably many symbols. The idea, essentially, is to replace each constant with a nearby element of a countable base in the formulas used to characterize the distances.

Theorem 8 *If an MTS $M = (S, \tau, \Sigma, [\cdot])$ is both finitely branching and separable, then the characterizations provided by Theorem 5 hold also when we restrict the formulas of* QLTL *to contain only constants from the countable set $\bigcup_{r \in \Sigma} \mathcal{B}_r$, where $\mathcal{B}_r$ is a countable basis for the metric space (X_r, d_r), for each $r \in \Sigma$.*

Proof. The result follows immediately from the observation that by Theorem 7 the value of a formula, at every state, can be approximated arbitrarily closely by the value of a formula containing only constants that belong to the countable bases of the metric spaces. □

3.5 A note on algorithmic complexity

The following section describes an algorithm that takes as input a finite MTS M over a directed metric space (X, d), and computes the value of a linear distance between all pairs of states. To discuss its complexity, we need to fix a finite representation for the input data. Considering that all the linear distances have as starting point the propositional distance pd, it is sufficient to provide as input the $|S| \times |S|$ matrix $A = (a_{s,t})_{s,t \in S}$, where $a_{s,t} = pd(s, t)$.

We assume that the values $pd(s, t)$ are rational numbers encoded in fixed-precision binary representation; we denote by $|x|_b$ the number of bits in the encoding of the rational number x. We define the size of a finite MTS $M = (S, \tau, \Sigma, [\cdot])$ by $|M| = \sum_{s,t \in S} |pd(s, t)|_b$. The size of an MTS is thus quadratic in $|S|$. We further assume that arithmetic operations can be carried out in constant time.

3.6 Computing the linear distance

Given as inputs a finite MTS $M = (S, \tau, \Sigma, [\cdot])$, a discount factor $\alpha \in (0, 1]$ (the case $\alpha = 0$ is trivial), and $x \in \{\mathrm{a}, \mathrm{s}\}$, we wish to compute $ld^x_\alpha(s_0, t_0)$, for all $s_0, t_0 \in S$.

We describe the computation of ld^{a}_α, as the computation of ld^{s}_α is analogous. We can read the definition of ld^{a}_α 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[5] of $t_0 \cdots t_n$. Such a game can be solved via a variation of the subset construction [16]. 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 $v(s_0 \cdots s_n, t_0 \cdots t_n) = \max_{0 \le k \le n} \alpha^{k-n} pd(s_k, t_k)$.

From M, we construct another MTS $M' = (S', \tau', \{r\}, [\cdot]')$, 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 $|\mathbb{D}| \le |S|^2$. For $\alpha \in (0,1)$, we take $\mathbb{D} = \{pd(s,t)/\alpha^k \mid s,t \in S \wedge k \in \mathbb{N} \wedge pd(s,t) \le \alpha^k\} \cup \{1\}$, so that $|\mathbb{D}| \le |S|^2 \cdot \lceil \log_\alpha \min\{pd(s,t) \mid s,t \in S \wedge pd(s,t) > 0\} \rceil + 1$. 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 M' with the highest possible (discounted) value or r. Thus, for all $s, t \in S$, we have $ld_\alpha^{\mathrm{a}}(s,t) = [\![\exists \diamond_\alpha r]\!]_{M'}(\langle s, \{\langle t, pd(s,t) \rangle\} \rangle)$, where the right-hand side is to be computed on M'. This expression can be evaluated by a depth-first traversal of the state space of M', noting that no state of M' needs to be visited twice, as subsequent visits do not increase the value of $\diamond_\alpha r$. This leads to the following complexity result.

Theorem 9 *For all $x \in \{\mathrm{a}, \mathrm{s}\}$, the following assertions hold:*

1. *Computing ld_α^x for $\alpha \in (0,1]$ and MTS M is PSPACE-complete in $|M| + |\alpha|_b$.*
2. *Computing ld_α^x for $\alpha \in (0,1]$ and deterministic MTS M is PSPACE-complete in $|M| + |\alpha|_b$.*
3. *Computing ld_α^x for $\alpha \in (0,1]$ and boolean, deterministic MTS M is in time $O(|M|^4)$.*

Proof. For Part 1, the upper complexity bound comes from the above algorithm, noticing that the subset construction can be done on the fly; the lower bound comes from a reduction from the corresponding result for trace inclusion [17].

Part 2 states that, unlike in the boolean case, the problem remains PSPACE-complete even for deterministic MTSs. This result is proved by an nlogspace reduction from the problem of computing the distance between nondeterministic systems to the one of computing it between deterministic ones. More precisely, let M be a nondeterministic MTS and let m be the number of bits needed to represent each quantity in M. Assume that α is also encoded as a fixed-precision number of m bits. Then, from an analysis of the algorithm, we see that the minimum difference between two possible answers returned by the algorithm is a number with $(n+1)m$ bits, where $n = |S|$. This is essentially α^n times the least difference of value among two non-equal valuations. We then build a

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

deterministic MTS M', by copying every valuation and padding it to $(n+1)m+1$ bits, thus using $\log_2 |S|$ additional bits to uniquely identify each state of S. Once the algorithm returns an answer for the deterministic system, the answer for the original nondeterministic one can be recovered by rounding to $(n+1)m$ bits of precision.

Part 3 is a consequence of Theorems 17 and 18. □

3.7 Discussion

In Definition 10, we could have defined the propositional distance between two states using the L_2 norm, via $pd(u,v) = \left(\sum_{r \in \Sigma} d(u(r), v(r))^2\right)^{1/2}$ (or in general using the L_n norm, for $n > 0$). The reason why in Definition 10 we chose the L_∞ norm is that this definition leads to a logical characterization of the distances, since the max in the L_∞ norm corresponds to the $\vee$ of the logics. It is easy to see that, aside from the logical characterizations, the results of the paper would hold if we replaced in Definition 10 the L_∞ norm with L_n, for any $n > 0$.

4 Branching Distances and Logics

4.1 Branching distances

Definition 13 (branching distances) For $\alpha \in (0,1]$ and $x \in \{\text{Aa}, \text{As}, \text{Sa}, \text{Ss}\}$, consider the four operators $H_\alpha^x : (S^2 \to \widehat{\mathbb{R}}_+) \to (S^2 \to \widehat{\mathbb{R}}_+)$ defined as follows, for $d : S^2 \to \widehat{\mathbb{R}}_+$:

$$H_\alpha^{\text{Aa}}(d)(s,t) = pd(s,t) \sqcup \alpha \cdot \sup_{s' \in \tau(s)} \inf_{t' \in \tau(t)} d(s',t')$$

$$H_\alpha^{\text{As}}(d)(s,t) = \overline{pd}(s,t) \sqcup \alpha \cdot \sup_{s' \in \tau(s)} \inf_{t' \in \tau(t)} d(s',t')$$

$$H_\alpha^{\text{Sa}}(d)(s,t) = pd(s,t) \sqcup \alpha \cdot \sup_{s' \in \tau(s)} \inf_{t' \in \tau(t)} d(s',t') \sqcup \alpha \cdot \sup_{t' \in \tau(t)} \inf_{s' \in \tau(s)} d(s',t')$$

$$H_\alpha^{\text{Ss}}(d)(s,t) = \overline{pd}(s,t) \sqcup \alpha \cdot \sup_{s' \in \tau(s)} \inf_{t' \in \tau(t)} d(s',t') \sqcup \alpha \cdot \sup_{t' \in \tau(t)} \inf_{s' \in \tau(s)} d(s',t')$$

For $x \in \{\text{Aa}, \text{As}, \text{Sa}, \text{Ss}\}$, we define the branching distance bd_α^x as the least fixpoints of the operators H^x. □

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}, $\overline{bd}_\alpha^{\text{Aa}}$, $\overline{bd}_\alpha^{\text{As}}$, and $\overline{bd}_\alpha^{\text{Sa}}$ are undirected metrics.

The distance bd_α^{Ss} is a quantitative generalization of bisimulation, and it essentially coincides with the metrics of [7, 18, 4]; as it is already symmetrical, we have $\overline{bd}_\alpha^{\text{Ss}} = bd_\alpha^{\text{Ss}}$. Similarly, the distance bd_α^{As} generalizes simulation, and $\overline{bd}_\alpha^{\text{As}}$ generalizes mutual simulation.

Theorem 10 *For all MTSs $(S, \tau, \Sigma, [\cdot])$ where d_r is a directed distance for all $r \in \Sigma$, and for all $\alpha \in (0,1]$, we have $\preceq_{sim} = \text{Zero}(bd_\alpha^{\text{As}})$ and $\approx_{bis} = \text{Zero}(bd_\alpha^{\text{Ss}})$.*

The distances bd_α^{Aa} and bd_α^{Sa} correspond to quantitative notions of simulation and bisimulation with respect to the asymmetrical propositional distance pd; these distances are not symmetrical, and we indicate their symmetrical versions by $\overline{bd}_\alpha^{\mathrm{Aa}}$ and $\overline{bd}_\alpha^{\mathrm{Sa}}$. Just as in the boolean case mutual similarity is not equivalent to bisimulation, so in our quantitative setting $\overline{bd}_\alpha^{\mathrm{As}}$ can be strictly smaller than bd_α^{Ss}, and $\overline{bd}_\alpha^{\mathrm{Aa}}$ can be strictly smaller than bd_α^{Sa}.

Theorem 11 *The relations in Figure 5(b) hold for all MTS and for all* $\alpha \in (0,1]$. *For* $\alpha \in (0,1]$, *no other inequalities hold on all MTSs.*

Proof. The inequalities $bd_\alpha^{\mathrm{Aa}} \leq bd_\alpha^{\mathrm{Sa}} \leq bd_\alpha^{\mathrm{Ss}}$ and $bd_\alpha^{\mathrm{Aa}} \leq bd_\alpha^{\mathrm{As}} \leq bd_\alpha^{\mathrm{Ss}}$ shown in the figure are immediate. Let $\alpha \in (0,1]$ and consider the MTS in Figure 5(a) again. In this MTS, we have $ld_\alpha^{\mathrm{a}} = bd_\alpha^{\mathrm{Aa}}$, $ld_\alpha^{\mathrm{s}} = \overline{bd}_\alpha^{\mathrm{As}}$, $\overline{ld}_\alpha^{\mathrm{a}} = \overline{bd}_\alpha^{\mathrm{Sa}}$, $\overline{ld}_\alpha^{\mathrm{s}} = bd_\alpha^{\mathrm{Ss}}$ Hence, the results for the linear distances (see Theorem 2) show that $bd_\alpha^{\mathrm{Aa}} \neq \overline{bd}_\alpha^{\mathrm{As}}$, $bd_\alpha^{\mathrm{Aa}} \neq \overline{bd}_\alpha^{\mathrm{Sa}}$, $bd_\alpha^{\mathrm{As}} \neq bd_\alpha^{\mathrm{Ss}}$, $bd_\alpha^{\mathrm{Sa}} \neq bd_\alpha^{\mathrm{Ss}}$, and neither $bd_\alpha^{\mathrm{As}} \leq bd_\alpha^{\mathrm{Sa}}$ nor $bd_\alpha^{\mathrm{As}} \geq bd_\alpha^{\mathrm{Sa}}$. □

The branching distances, like the linear ones, are robust with respect to perturbations in the state valuations: small changes in the proposition valuations cause small changes in the distances. To state the theorem, given a state valuation $f : S \to \mathcal{U}[\Sigma]$, $x \in \{\mathrm{Aa}, \mathrm{As}, \mathrm{Sa}, \mathrm{Ss}\}$, and $\alpha \in (0,1]$, we write $bd_{f,\alpha}^x$ for the distances defined as in Definition 13, using f as the state valuation.

Theorem 12 (branching distance robustness) *For all* $\alpha \in (0,1]$, *all* $x \in \{\mathrm{As}, \mathrm{Sa}, \mathrm{Ss}\}$, *all predicate valuations* $[\cdot]_1, [\cdot]_2$, *and all* $s, t \in S$, *we have*

$$bd_{[\cdot]_1,\alpha}^{Aa}(s,t) - bd_{[\cdot]_2,\alpha}^{Aa}(s,t) \leq d([\cdot]_1,[\cdot]_2) + d([\cdot]_2,[\cdot]_1)$$
$$|bd_{[\cdot]_1,\alpha}^{x}(s,t) - bd_{[\cdot]_2,\alpha}^{x}(s,t)| \leq 2 \cdot \overline{d}([\cdot]_1,[\cdot]_2).$$

4.2 Quantitative μ-calculus

We define quantitative μ-calculus after [5, 4]. Given a set of variables X and a set of propositions Σ, the formulas of the *quantitative* μ*-calculus* are generated by the grammar:

$$\varphi \ ::= \ D(r,c) \mid D(c,r) \mid x \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \exists \bigcirc_\alpha \varphi \mid \forall \bigcirc_\alpha \varphi \mid \mu x \,.\, \varphi \mid \nu x \,.\, \varphi$$

for propositions $r \in \Sigma$, variables $x \in X$, constants $c \in \bigcup_{r \in AP} X_r$, and discount factors $\alpha \in (0,1]$. We assume that, in a term of the form $D(r,c)$ or $D(c,r)$, we have $c \in X_r$. Denoting by $\mathcal{F} = (S \to \widehat{\mathbb{R}}_+)$, a (variable) interpretation is a function $\mathcal{E} : X \to \mathcal{F}$. Given an interpretation $\mathcal{E}$, a variable $x \in X$ and a function $f \in \mathcal{F}$, we denote by $\mathcal{E}[x := f]$ the interpretation $\mathcal{E}'$ such that $\mathcal{E}'(x) = f$ and, for all $y \neq x$, $\mathcal{E}'(y) = \mathcal{E}(y)$. Given an MTS and an interpretation $\mathcal{E}$, every formula

φ of the quantitative μ-calculus defines a valuation $[\![\varphi]\!]_{\mathcal{E}} : S \to \widehat{\mathbb{R}}_+$:

$$\begin{array}{llll} [\![D(r,c)]\!]_{\mathcal{E}}(s) &= d([s](r),c) & [\![\exists\bigcirc_\alpha\varphi]\!]_{\mathcal{E}}(s) &= \alpha \cdot \sup_{s'\in\tau(s)}[\![\varphi]\!]_{\mathcal{E}}(s') \\ [\![D(c,r)]\!]_{\mathcal{E}}(s) &= d(c,[s](r)) & [\![\forall\bigcirc_\alpha\varphi]\!]_{\mathcal{E}}(s) &= \alpha \cdot \inf_{s'\in\tau(s)}[\![\varphi]\!]_{\mathcal{E}}(s') \\ [\![x]\!]_{\mathcal{E}} &= \mathcal{E}(x) & [\![\mu x \,.\, \varphi]\!]_{\mathcal{E}} &= \inf\{f \in \mathcal{F} \mid f = [\![\varphi]\!]_{\mathcal{E}[x:=f]}\} \\ [\![\varphi_1 \wedge \varphi_2]\!]_{\mathcal{E}} &= [\![\varphi_1]\!]_{\mathcal{E}} \sqcap [\![\varphi_2]\!]_{\mathcal{E}} & [\![\nu x \,.\, \varphi]\!]_{\mathcal{E}} &= \sup\{f \in \mathcal{F} \mid f = [\![\varphi]\!]_{\mathcal{E}[x:=f]}\}. \\ [\![\varphi_1 \vee \varphi_2]\!]_{\mathcal{E}} &= [\![\varphi_1]\!]_{\mathcal{E}} \sqcup [\![\varphi_2]\!]_{\mathcal{E}} & & \end{array}$$

The existence of the required fixpoints is guaranteed by the monotonicity and continuity of all operators. A variable x is *bound* in φ if it is in the scope of a quantifier μx or νx; otherwise, it is called *free*. A formula is *closed* if all variables are bound. If φ is closed, we write $[\![\varphi]\!]$ for $[\![\varphi]\!]_{\mathcal{E}}$. For all $\alpha \in (0,1]$, we call QMU_α the set of quantitative μ-calculus formulas where all discount factors are smaller than or equal to α. We denote by CLQMU_α the subset of QMU_α containing only closed formulas. For $ops \subseteq \{D(c,r), D(r,c), \exists, \forall, \mu, \nu\}$, we denote by $\text{QMU}_\alpha \setminus ops$ and $\text{CLQMU}_\alpha \setminus ops$ the respective subsets of formulas that do not employ operators in ops. Notice that, if 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.

4.3 Logical characterizations of branching distances

In the following theorem, we write $\varphi(x_1, \ldots, x_n)$ to signify that the free variables in φ are among $x_1, \ldots, x_n$.

Lemma 4 *Given an MTS $(S, \tau, \Sigma, [\cdot])$ and a discount factor $\alpha \in (0,1]$, the following holds.*

1. *For all $\varphi(x_1, \ldots, x_n) \in \text{QMU}_\alpha \setminus \{\exists, D(r,c)\}$, for all variable environments $\mathcal{E}$, and for all $f_1, \ldots, f_n \in \mathcal{F}$, if for all $s,t \in S$ and all $i = 1, \ldots, n$, $f_i(t) - f_i(s) \leq bd^{\text{Aa}}_\alpha(s,t)$, then, for all $s,t \in S$,*
$$[\![\varphi]\!]_{\mathcal{E}[x_i:=f_i]}(t) - [\![\varphi]\!]_{\mathcal{E}[x_i:=f_i]}(s) \leq bd^{\text{Aa}}_\alpha(s,t).$$
2. *For all $\varphi(x_1, \ldots, x_n) \in \text{QMU}_\alpha \setminus \{\exists\}$, and for all $f_1, \ldots, f_n \in \mathcal{F}$, if for all $s,t \in S$ and all $i = 1, \ldots, n$, $f_i(t) - f_i(s) \leq bd^{\text{As}}_\alpha(s,t)$, then, for all $s,t \in S$,*
$$[\![\varphi]\!]_{\mathcal{E}[x_i:=f_i]}(t) - [\![\varphi]\!]_{\mathcal{E}[x_i:=f_i]}(s) \leq bd^{\text{As}}_\alpha(s,t).$$
3. *For all $\varphi(x_1, \ldots, x_n) \in \text{QMU}_\alpha \setminus \{D(r,c)\}$, and for all $f_1, \ldots, f_n \in \mathcal{F}$, if for all $s,t \in S$ and all $i = 1, \ldots, n$, $f_i(t) - f_i(s) \leq bd^{\text{Sa}}_\alpha(s,t)$, then, for all $s,t \in S$,*
$$[\![\varphi]\!]_{\mathcal{E}[x_i:=f_i]}(t) - [\![\varphi]\!]_{\mathcal{E}[x_i:=f_i]}(s) \leq bd^{\text{Sa}}_\alpha(s,t).$$
4. *For all $\varphi(x_1, \ldots, x_n) \in \text{QMU}_\alpha$, and for all $f_1, \ldots, f_n \in \mathcal{F}$, if for all $s,t \in S$ and all $i = 1, \ldots, n$, $|f_i(t) - f_i(s)| \leq bd^{\text{Ss}}_\alpha(s,t)$, then, for all $s,t \in S$,*
$$|[\![\varphi]\!]_{\mathcal{E}[x_i:=f_i]}(t) - [\![\varphi]\!]_{\mathcal{E}[x_i:=f_i]}(s)| \leq bd^{\text{Ss}}_\alpha(s,t).$$

Proof. We prove statements 1 and 3; the other two statements can be proved in similar fashion.

Statement 1. We prove the result concerning bd_α^{Aa} by structural induction on the formula. For $\varphi = D(c,r)$, we obtain by triangle inequality $[\![\varphi]\!](t) - [\![\varphi]\!](s) = d(c,[t](r)) - d(c,[s](r)) \le d([s](r),[t](r)) \le pd(s,t) \le bd_\alpha^{\mathrm{Aa}}(s,t)$.

The cases $\varphi = x$, $\varphi = \varphi_1 \wedge \varphi_2$ and $\varphi = \varphi_1 \vee \varphi_2$ are also trivial.

Consider the case $\varphi = \forall\bigcirc_\beta \psi$, for some $\beta \le \alpha$: we prove that, for all states $s,t \in S$ and all $\epsilon > 0$, $[\![\varphi]\!](t) - [\![\varphi]\!](s) \le bd_\alpha^{\mathrm{Aa}}(s,t) + \epsilon$. For ease of notation, in this part of the proof we write $[\![\cdot]\!]$ for $[\![\cdot]\!]_{\mathcal{E}[x_i := f_i]}$, as the variable interpretation is not the issue here. Recall that, for all $t \in S$, we have by definition $[\![\varphi]\!](t) = \beta \inf_{t' \in \tau(t)} [\![\psi]\!](t')$, By inductive hypothesis, for all $s', t' \in S$, $[\![\psi]\!](t') - [\![\psi]\!](s') \le bd_\alpha^{\mathrm{Aa}}(s',t')$. For all $s^* \in \tau(s)$ and $\delta > 0$, we define $closer(t, s^*, \delta)$ to contain all states $t^* \in \tau(t)$ such that $bd^{\mathrm{Aa}}(s^*,t^*) \le \delta + \inf_{t' \in \tau(t)} bd_\alpha^{\mathrm{Aa}}(s^*,t')$. Intuitively, $closer(t,s^*,\delta)$ contains those successors of t that are closer than δ to the best match for s^*. For all $s^* \in \tau(s)$ and $t^* \in closer(t,s^*,\delta)$, we have that

$$\begin{aligned}
\alpha \cdot \big([\![\psi]\!](t^*) - [\![\psi]\!](s^*)\big) &\le \alpha \cdot bd_\alpha^{\mathrm{Aa}}(s^*,t^*) \\
&\le \alpha \cdot \Big(\inf_{t' \in \tau(t)} bd_\alpha^{\mathrm{Aa}}(s^*,t') + \delta\Big) \\
&\le \alpha \cdot \Big(\sup_{s' \in \tau(s)} \inf_{t' \in \tau(t)} bd_\alpha^{\mathrm{Aa}}(s',t') + \delta\Big) \\
&\le \alpha\delta + bd_\alpha^{\mathrm{Aa}}(s,t). \qquad (\S)
\end{aligned}$$

Finally, let $s^* \in \tau(s)$ be such that $[\![\psi]\!](s^*) \le \inf_{s' \in \tau(s)} [\![\psi]\!](s') + \frac{\epsilon}{2\alpha}$ and $t^* \in closer(t, s^*, \frac{\epsilon}{2\alpha})$, we have

$$\begin{aligned}
[\![\varphi]\!](t) - [\![\varphi]\!](s) &= \beta \inf_{t' \in \tau(t)} [\![\psi]\!](t') - \beta \inf_{s' \in \tau(s)} [\![\psi]\!](s') \\
&\le \alpha\Big([\![\psi]\!](t^*) - [\![\psi]\!](s^*) + \frac{\epsilon}{2\alpha}\Big) \qquad (\dagger) \\
&\le \frac{\epsilon}{2} + \alpha([\![\psi]\!](t^*) - [\![\psi]\!](s^*)) \\
&\le \frac{\epsilon}{2} + \frac{\epsilon}{2} + bd_\alpha^{\mathrm{Aa}}(s,t). \qquad (\ddagger)
\end{aligned}$$

To obtain (†), we have used $[\![\psi]\!](t^*) \ge \inf_{t' \in \tau(t)} [\![\psi]\!](t')$ and our choice of s^*; to obtain (‡), we have used $t^* \in closer(t, s^*, \frac{\epsilon}{2\alpha})$, along with the previous result (§). This concludes this case.

If $\varphi = \mu y \,.\, \psi$, then $[\![\varphi]\!] = \lim_n g_n$, where $g_0(s) = 0$ for all $s \in S$, and $g_{n+1} = [\![\psi]\!]_{\mathcal{E}[y := g_n]}$. This is a consequence of the fact that, when the MTS is finitely branching, all operators of the μ-calculus are continuous: that is, for each operator $F \in \{\wedge, \vee, \exists\bigcirc, \forall\bigcirc\}$ and each sequence $g_{n\, n \ge 0}$ of functions $S^2 \to \widehat{\mathbb{R}}_+$, we have $F(\lim_n g_n) = \lim_n F(g_n)$. Since $g_0(t) - g_0(s) = 0 \le bd_\alpha^{\mathrm{Aa}}(s,t)$, by inductive hypothesis we obtain that, for all $n \in \mathbb{N}$, $g_n(t) - g_n(s) \le bd_\alpha^{\mathrm{Aa}}(s,t)$, and thus the thesis. By taking $g_0(s) = \infty$ for all $s \in S$, we obtain the argument for $\varphi = \nu y \,.\, \psi$.

Statement 3. The cases $\varphi = r$, $\varphi = x$, $\varphi = \psi_1 \wedge \psi_2$ and $\varphi = \psi_1 \vee \psi_2$ are trivial, while the proofs for $\varphi = \forall\bigcirc_\beta \psi$, $\varphi = \mu y \,.\, \psi$ and $\varphi = \nu y \,.\, \psi$ are similar to the ones of Part 1.

Let $\varphi = \exists \bigcirc_\beta \psi$, for some $\beta \leq \alpha$. We prove that, for all states $s, t \in S$ and all $\epsilon > 0$, $[\![\varphi]\!](t) - [\![\varphi]\!](s) \leq bd_\alpha^{\mathrm{Sa}}(s,t) + \epsilon$. For ease of notation, we again write $[\![\cdot]\!]$ for $[\![\cdot]\!]_{\mathcal{E}[x_i:=f_i]}$. By inductive hypothesis, for all $s', t' \in S$, $[\![\psi]\!](t') - [\![\psi]\!](s') \leq bd_\alpha^{\mathrm{Sa}}(s', t')$.

For all $s^* \in \tau(s)$ and $\delta > 0$, we define $closer(s, t^*, \delta)$ to contain all states $s^* \in \tau(t)$ such that $bd^{\mathrm{Sa}}(s^*, t^*) \leq \delta + \inf_{s' \in \tau(s)} bd_\alpha^{\mathrm{Sa}}(s', t^*)$. Again, $closer(s, t^*, \delta)$ contains those successors of s that are closer than δ to the best match for t^*. For all $t^* \in \tau(t)$ and $s^* \in closer(s, t^*, \delta)$, we have that $\alpha \cdot bd_\alpha^{\mathrm{Sa}}(s^*, t^*) \leq \alpha\delta + bd_\alpha^{\mathrm{Sa}}(s,t)$, and thus

$$\begin{aligned} \alpha \cdot \big([\![\psi]\!](t^*) - [\![\psi]\!](s^*)\big) &\leq \alpha \cdot bd_\alpha^{\mathrm{Sa}}(s^*, t^*) \\ &\leq \alpha\delta + bd_\alpha^{\mathrm{Sa}}(s,t). \qquad (\S\S) \end{aligned}$$

There are now three cases.

1. If $[\![\varphi]\!](t) = \beta \sup_{t' \in \tau(t)} [\![\psi]\!](t') < \infty$, then let $t^* \in \tau(t)$ be such that $[\![\psi]\!](t^*) \geq \sup_{t' \in \tau(t)} [\![\psi]\!](t') - \frac{\epsilon}{2\alpha}$ and $s^* \in closer(s, t^*, \frac{\epsilon}{2\alpha})$. We have

$$\begin{aligned} [\![\varphi]\!](t) - [\![\varphi]\!](s) &= \beta \sup_{t' \in \tau(t)} [\![\psi]\!](t') - \beta \sup_{s' \in \tau(s)} [\![\psi]\!](s') \\ &\leq \alpha\big([\![\psi]\!](t^*) + \frac{\epsilon}{2\alpha} - [\![\psi]\!](s^*)\big) \\ &\leq \frac{\epsilon}{2} + \alpha([\![\psi]\!](t^*) - [\![\psi]\!](s^*)) \\ &\leq \frac{\epsilon}{2} + \frac{\epsilon}{2} + bd_\alpha^{\mathrm{Sa}}(s,t), \end{aligned}$$

   leading to the desired result.
2. If $[\![\varphi]\!](t) = \infty$ and $bd_\alpha^{\mathrm{Sa}}(s,t) = \infty$, then we are done.
3. If $[\![\varphi]\!](t) = \infty$ and $bd_\alpha^{\mathrm{Sa}}(s,t) < \infty$, then for every $c \in \mathbb{R}$, we can find $t^* \in \tau(t)$ such that $[\![\psi]\!](t^*) \geq c$. From (§§), we can thus find $s^* \in closer(s, t^*, 1)$ such that

$$\alpha(c-1) - bd_\alpha^{\mathrm{Sa}}(s,t) \leq \alpha[\![\psi]\!](t^*) - \alpha - bd_\alpha^{\mathrm{Sa}}(s,t) \leq [\![\psi]\!](s^*).$$

   From $[\![\varphi]\!](s) = \beta \sup_{s' \in \tau(s)} [\![\psi]\!](s') \geq [\![\psi]\!](s^*)$, since $bd_\alpha^{\mathrm{Sa}}(s,t) < \infty$ and since c is arbitrary, we obtain $[\![\varphi]\!](s) = \infty = [\![\varphi]\!](t)$, concluding the proof. □

From the preceding lemma, we immediately obtain a theorem stating that the branching distances provide bounds for the corresponding fragments of the μ-calculus. The statement for bd_α^{Ss} is very similar to a result in [7].

Theorem 13 *For all MTSs $(S, \tau, \Sigma, [\cdot])$, states $s, t \in S$, and $\alpha \in (0, 1]$, we have*

for all $\varphi \in \mathrm{ClQmu}_\alpha \setminus \{\exists, D(r,c)\}$	$bd_\alpha^{\mathrm{Aa}}(s,t) \geq [\![\varphi]\!](t) - [\![\varphi]\!](s)$
for all $\varphi \in \mathrm{ClQmu}_\alpha \setminus \{\exists\}$	$bd_\alpha^{\mathrm{As}}(s,t) \geq [\![\varphi]\!](t) - [\![\varphi]\!](s)$
for all $\varphi \in \mathrm{ClQmu}_\alpha \setminus \{D(r,c)\}$	$bd_\alpha^{\mathrm{Sa}}(s,t) \geq [\![\varphi]\!](t) - [\![\varphi]\!](s)$
for all $\varphi \in \mathrm{ClQmu}_\alpha$	$bd_\alpha^{\mathrm{Ss}}(s,t) \geq \lvert [\![\varphi]\!](t) - [\![\varphi]\!](s) \rvert$

As noted before, each bound of the form $d(s,t) \geq [\![\varphi]\!](t) - [\![\varphi]\!](s)$ trivially leads to a bound of the form $\overline{d}(s,t) \geq |[\![\varphi]\!](t) - [\![\varphi]\!](s)|$. The bounds are tight for finitely branching systems, and the following theorem identifies which fragments of quantitative μ-calculus suffice for characterizing each branching distance. The formula scheme used to characterize bd_α^{Ss} is reminiscent of the one used in [1] for bisimulation.

Theorem 14 *For all finitely branching MTSs $(S, \tau, \Sigma, [\cdot])$, states $s, t \in S$, and $\alpha \in (0,1]$, we have*

$$\begin{aligned} bd_\alpha^{\mathrm{Aa}}(s,t) &= \sup\nolimits_{\varphi \in \mathrm{CLQMU}_\alpha \setminus \{\exists, D(r,c), \mu, \nu\}} [\![\varphi]\!](t) - [\![\varphi]\!](s), \\ bd_\alpha^{\mathrm{As}}(s,t) &= \sup\nolimits_{\varphi \in \mathrm{CLQMU}_\alpha \setminus \{\exists, \mu, \nu\}} [\![\varphi]\!](t) - [\![\varphi]\!](s), \\ bd_\alpha^{\mathrm{Sa}}(s,t) &= \sup\nolimits_{\varphi \in \mathrm{CLQMU}_\alpha \setminus \{D(r,c), \mu, \nu\}} [\![\varphi]\!](t) - [\![\varphi]\!](s), \\ bd_\alpha^{\mathrm{Ss}}(s,t) &= \sup\nolimits_{\varphi \in \mathrm{CLQMU}_\alpha \setminus \{\mu, \nu\}} [\![\varphi]\!](t) - [\![\varphi]\!](s). \end{aligned}$$

Proof.

Part 1. Consider the statement about bd_α^{Aa}. For all $s \in S$, we define the sequence of formulas $(\varphi_s^k)_{k \geq 0}$ as follows.

$$\begin{aligned} \varphi_s^0 &= \bigvee_{r \in \Sigma} D([s](r), r), \\ \varphi_s^{k+1} &= \varphi_s^0 \vee \bigvee_{s' \in \tau(s)} \forall \bigcirc_\alpha \varphi_{s'}^k. \end{aligned}$$

First, one can easily prove by induction that, for all $k \in \mathbb{N}$ and $s \in S$, $[\![\varphi_s^k]\!](s) = 0$. The distance bd_α^{Aa} is defined as the least fixpoint of H_α^{Aa}. Denoting by $(H_\alpha^{\mathrm{Aa}})^k$ a sequence of k applications of H_α^{Aa}, since the MTS is finitely branching, we have that $bd_\alpha^{\mathrm{Aa}} = \lim_k (H_\alpha^{\mathrm{Aa}})^k(pd)$. We prove by induction on k that, for all $s, t \in S$, $[\![\varphi_s^k]\!](t) = (H_\alpha^{\mathrm{Aa}})^k(pd)(s,t)$.

$$\begin{aligned} [\![\varphi_s^0]\!](t) &= \max_{r \in \Sigma} d([s](r), [t](r)) \\ &= pd(s,t) = (H_\alpha^{\mathrm{Aa}})^0(pd)(s,t); \end{aligned}$$

$$\begin{aligned} [\![\varphi_s^{k+1}]\!](t) &= [\![\varphi_s^0]\!](t) \sqcup \max_{s' \in \tau(s)} \min_{t' \in \tau(t)} \alpha [\![\varphi_{s'}^k]\!](t') \\ &= pd(s,t) \sqcup \max_{s' \in \tau(s)} \min_{t' \in \tau(t)} \alpha \cdot (H_\alpha^{\mathrm{Aa}})^k(pd)(s',t') \\ &= (H_\alpha^{\mathrm{Aa}})^{k+1}(pd)(s,t). \end{aligned}$$

It follows that

$$\begin{aligned} \sup_{\varphi \in \mathrm{CLQMU}_\alpha \setminus \{\exists, D(r,c), \mu, \nu\}} [\![\varphi]\!](t) - [\![\varphi]\!](s) &\geq \sup_{k \in \mathbb{N}} [\![\varphi_s^k]\!](t) - [\![\varphi_s^k]\!](s) \\ &= \sup_{k \in \mathbb{N}} (H_\alpha^{\mathrm{Aa}})^k(pd)(s,t) - 0 \\ &= bd_\alpha^{\mathrm{Aa}}(s,t). \end{aligned}$$

Part 2. To prove the statement concerning $bd_\alpha^{\mathrm{As}}(s,t)$, we define the following sequence of formulas $(\varphi_s^k)_{k\in\mathbb{N}}$.

$$\varphi_s^0 = \bigvee_{r\in\Sigma} D([s](r),r) \vee D(r,[s](r)),$$
$$\varphi_s^{k+1} = \varphi_s^0 \vee \bigvee_{s'\in\tau(s)} \forall\bigcirc_\alpha \varphi_{s'}^k.$$

We then proceed similarly to the previous part.

Part 3. To prove the bound on $bd_\alpha^{\mathrm{Sa}}(s,t)$, we use the formulas:

$$\varphi_s^0 = \bigvee_{r\in\Sigma} D([s](r),r)$$
$$\varphi_s^{k+1} = \varphi_s^0 \vee \bigvee_{s'\in\tau(s)} \forall\bigcirc_\alpha \varphi_{s'}^k \vee \exists\bigcirc_\alpha\Big(\bigwedge_{s'\in\tau(s)} \varphi_{s'}^k\Big).$$

Once again, one can easily prove by induction that, for all $k\in\mathbb{N}$ and $s\in S$, $[\![\varphi_s^k]\!](s) = 0$. The distance bd_α^{Sa} is defined as the least fixpoint of H_α^{Sa}. In particular, denoting by $(H_\alpha^{\mathrm{Sa}})^k$ a sequence of k applications of H_α^{Sa}, again due to the fact that the MTS is finitely branching we have $bd_\alpha^{\mathrm{Sa}} = \lim_k (H_\alpha^{\mathrm{Sa}})^k(pd)$. We prove by induction on k that, for all $s,t\in S$, $[\![\varphi_s^k]\!](t) = (H_\alpha^{\mathrm{Sa}})^k(pd)(s,t)$.

$$\begin{aligned}[\![\varphi_s^0]\!](t) &= \max_{r\in\Sigma}\big(d([s](r),[t](r)) \sqcup d([t](r),[s](r))\big)\\ &= pd(s,t) = (H_\alpha^{\mathrm{Sa}})^0(pd)(s,t);\end{aligned}$$

$$\begin{aligned}[\![\varphi_s^{k+1}]\!](t) &= [\![\varphi_s^0]\!](t) \sqcup \max_{s'\in\tau(s)}\min_{t'\in\tau(t)} \alpha[\![\varphi_{s'}^k]\!](t') \sqcup \max_{t'\in\tau(t)}\min_{s'\in\tau(s)} \alpha[\![\varphi_{s'}^k]\!](t')\\ &= pd(s,t) \sqcup \alpha \max_{s'\in\tau(s)}\min_{t'\in\tau(t)} (H_\alpha^{\mathrm{Sa}})^k(pd)(s',t')\\ &\qquad \sqcup \alpha \max_{t'\in\tau(t)}\min_{s'\in\tau(s)} (H_\alpha^{\mathrm{Sa}})^k(pd)(s',t')\\ &= (H_\alpha^{\mathrm{Sa}})^{k+1}(pd)(s,t).\end{aligned}$$

It follows that

$$\begin{aligned}\sup_{\varphi\in\mathrm{CLQMU}_\alpha\setminus\{D(r,c),\mu,\nu\}} [\![\varphi]\!](t) - [\![\varphi]\!](s) &\geq \sup_{k\in\mathbb{N}} [\![\varphi_s^k]\!](t) - [\![\varphi_s^k]\!](s)\\ &= \sup_{k\in\mathbb{N}} (H_\alpha^{\mathrm{Sa}})^k(pd)(s,t) - 0\\ &= bd_\alpha^{\mathrm{Sa}}(s,t).\end{aligned}$$

Part 4. To prove the bound on $bd_\alpha^{\mathrm{Ss}}(s,t)$, we use the formulas:

$$\varphi_s^0 = \bigvee_{r\in\Sigma} D([s](r),r) \vee D(r,[s](r)),$$
$$\varphi_s^{k+1} = \varphi_s^0 \vee \bigvee_{s'\in\tau(s)} \forall\bigcirc_\alpha \varphi_{s'}^k \vee \exists\bigcirc_\alpha\Big(\bigwedge_{s'\in\tau(s)} \varphi_{s'}^k\Big).$$

We then proceed similarly to the previous parts. □

4.4 Logical characterization via logics with countably many symbols

Again, the logical characterization above is in terms of formulas defined over a potentially uncountable set of constants: in general, we need one constant for each element of a metric space corresponding to a predicate. As in the linear case, we show that if the MTS is separable, then it suffices to consider formulas defined over the countable set of constants corresponding to the countable bases of the metric spaces for the various predicates. We start once more with a result that expresses the robustness of the calculus with respect to changes in the valuation of the constants.

Theorem 15 *Consider a formula φ of the quantitative μ-calculus containing the constants $c_1, \ldots, c_n$, belonging respectively to the metric spaces $(q_1, d_1), \ldots, (q_n, d_n)$. Let $\psi = \varphi[c'_1, \ldots, c'_n / c_1, \ldots, c_n]$ be the result of replacing each c_i with c'_i, for $1 \le i \le n$, and let $\delta = \max_{i=1}^{n}(d_i(c_i, c'_i) \sqcup d_i(c'_i, c_i))$ be the maximal distance between the new and old value of each constant. Then, for all $s \in S$ and all variable environments $\mathcal{E}$, we have $|[\![\varphi]\!]_{\mathcal{E}}(s) - [\![\psi]\!]_{\mathcal{E}}(s)| \le \delta$.*

Proof. The result is obtained by a straightforward induction on the structure of the formula; the only interesting case is the base case for D, which is proved as in the proof of Theorem 7. □

Again, for separable MTSs this result leads to logical characterizations based on languages with countable sets of constants, corresponding to the bases of the metric spaces.

Theorem 16 *If an MTS $M = (S, \tau, \Sigma, [\cdot])$ is both finitely branching and separable, then the characterizations provided by Theorem 14 hold also when we restrict the formulas of quantitative μ-calculus to contain only constants from the countable set $\bigcup_{r \in \Sigma} \mathcal{B}_r$, where $\mathcal{B}_r$ is a countable basis for the metric space (X_r, d_r), for each $r \in \Sigma$.*

Proof. Similarly to the linear case, the result follows from the observation that by Theorem 15 the value of a formula, at every state, can be approximated arbitrarily closely by the value of a formula containing only constants that belong to the countable bases of the metric spaces. □

4.5 Computing the branching distances

Given a finite MTS $M = (S, \tau, \Sigma, [\cdot])$ 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 13. 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.

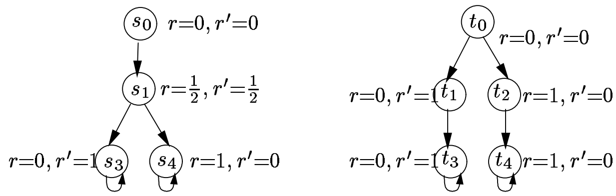

Fig. 4. Linear versus branching distances on a deterministic MTS.

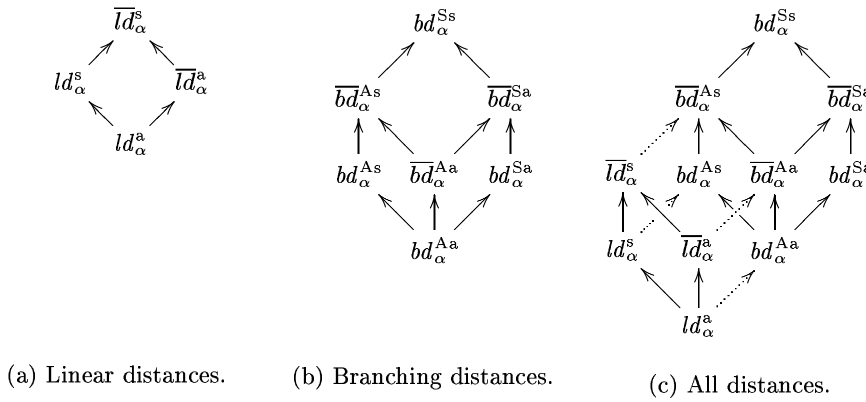

Fig. 5. Relations between distances, where $f \to g$ means $f \leq g$. In (c), the dotted arrows collapse to equality for boolean, deterministic MTSs.

Theorem 17 *Computing* bd_α^x *for* $x \in \{\mathrm{Ss}, \mathrm{Sa}, \mathrm{As}, \mathrm{Aa}\}$, $\alpha \in (0, 1]$ *and an MTS* M *can be done in time* $O(|M|^4)$.

5 Comparing the Linear and Branching Distances

Last, we provide a comparison between linear and branching distances. Just as similarity implies trace inclusion, we have both $ld_\alpha^{\mathrm{a}} \leq bd_\alpha^{\mathrm{Aa}}$ and $ld_\alpha^{\mathrm{s}} \leq \overline{bd}_\alpha^{\mathrm{As}}$; just as bisimilarity implies trace equivalence, we have $\overline{ld}_\alpha^{\mathrm{s}} \leq bd_\alpha^{\mathrm{Ss}}$ and $\overline{ld}_\alpha^{\mathrm{a}} \leq \overline{bd}_\alpha^{\mathrm{Sa}}$. Moreover, in the non-quantitative setting, trace inclusion (resp. trace equivalence) coincides with (bi-)similarity on deterministic systems. This result generalizes to distances over MTSs that are both deterministic and boolean, but not to distances over MTSs that are just deterministic. To formalize these results, we say that an MTS is *boolean* if all its predicates are evaluated in the metric space $\mathbf{X}_{\mathbb{B}}$.

Theorem 18 *The following properties hold.*

1. *For all MTSs and all* $\alpha \in (0,1]$*, we have*

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

Moreover, for $\alpha \in (0,1]$*, the inequalities cannot be replaced by equalities.*

2. *For all boolean, deterministic MTSs and for 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 MTSs.

The relations of Part 1 are illustrated in Figure 5(c).

In order to prove this theorem, we proceed in steps. First, we provide a relation between the fixpoints of the operators used to define linear and branching distances. For $\alpha \in (0,1]$ and $x \in \{\mathrm{a},\mathrm{s}\}$, we define the operators $F^x_\alpha, \overline{F}^x_\alpha : (S^2 \to \widehat{\mathbb{R}}_+) \to (S^2 \to \widehat{\mathbb{R}}_+)$ as follows, for $d : S^2 \to \widehat{\mathbb{R}}_+$:

$$\begin{aligned}
F^{\mathrm{a}}_\alpha(d)(s,t) &= pd(s,t) \sqcup \sup_{\sigma \in Paths(s)} \inf_{\rho \in Paths(t)} \sup_{i \in \mathbb{N}} \alpha^i d(\sigma_i, \rho_i) \\
F^{\mathrm{s}}_\alpha(d)(s,t) &= \overline{pd}(s,t) \sqcup \sup_{\sigma \in Paths(s)} \inf_{\rho \in Paths(t)} \sup_{i \in \mathbb{N}} \alpha^i d(\sigma_i, \rho_i) \\
\overline{F}^{\mathrm{a}}_\alpha(d)(s,t) &= pd(s,t) \sqcup \sup_{\sigma \in Paths(s)} \inf_{\rho \in Paths(t)} \sup_{i \in \mathbb{N}} \alpha^i d(\sigma_i, \rho_i) \\
&\quad \sqcup \sup_{\rho \in Paths(t)} \inf_{\sigma \in Paths(s)} \sup_{i \in \mathbb{N}} \alpha^i d(\rho_i, \sigma_i) \\
\overline{F}^{\mathrm{s}}_\alpha(d)(s,t) &= \overline{pd}(s,t) \sqcup \sup_{\sigma \in Paths(s)} \inf_{\rho \in Paths(t)} \sup_{i \in \mathbb{N}} \alpha^i d(\sigma_i, \rho_i) \\
&\quad \sqcup \sup_{\rho \in Paths(t)} \inf_{\sigma \in Paths(s)} \sup_{i \in \mathbb{N}} \alpha^i d(\rho_i, \sigma_i).
\end{aligned}$$

These operators should be compared with the fixpoint operators used in Definition 13 to define the branching distances. Essentially, the operators F^x_α above share the same structure of the operators H^x_α, except that F^x_α looks at the infinite paths originating from states, whereas H^x_α looks just at the successor states. The following lemma follows immediately from the definitions.

Lemma 5 *Denoting by* $\mathbf{0} : \lambda(s,t).0$ *the zero function* $S^2 \to \widehat{\mathbb{R}}_+$*. For* $\alpha \in (0,1]$ *and* $x \in \{\mathrm{a},\mathrm{s}\}$*, we have:*

$$\begin{aligned}
ld^{\mathrm{a}}_\alpha &= F^{\mathrm{a}}_\alpha(F^{\mathrm{a}}_\alpha(\mathbf{0})) \\
ld^{\mathrm{s}}_\alpha &= F^{\mathrm{s}}_\alpha(F^{\mathrm{s}}_\alpha(\mathbf{0})) \\
\overline{ld}^{\mathrm{a}}_\alpha &= \overline{F}^{\mathrm{a}}_\alpha(\overline{F}^{\mathrm{a}}_\alpha(\mathbf{0})) \\
\overline{ld}^{\mathrm{s}}_\alpha &= \overline{F}^{\mathrm{s}}_\alpha(\overline{F}^{\mathrm{s}}_\alpha(\mathbf{0})).
\end{aligned}$$

For $\alpha \in (0,1]$ denote the least fixpoints of these operators by:

$$\begin{aligned}
fd^{\mathrm{Aa}}_\alpha &= \inf\{d : S^2 \to \widehat{\mathbb{R}}_+ \mid d = F^{\mathrm{a}}_\alpha(d)\} \\
fd^{\mathrm{As}}_\alpha &= \inf\{d : S^2 \to \widehat{\mathbb{R}}_+ \mid d = F^{\mathrm{s}}_\alpha(d)\} \\
fd^{\mathrm{Sa}}_\alpha &= \inf\{d : S^2 \to \widehat{\mathbb{R}}_+ \mid d = \overline{F}^{\mathrm{a}}_\alpha(d)\} \\
fd^{\mathrm{Ss}}_\alpha &= \inf\{d : S^2 \to \widehat{\mathbb{R}}_+ \mid d = \overline{F}^{\mathrm{s}}_\alpha(d)\}
\end{aligned}$$

(where we have preferred to avoid the μ-notation for least fixpoints not to generate confusion with μ-calculus over MTSs). The following lemma states that these fixpoints are branching distances.

Lemma 6 *For all $\alpha \in (0,1]$, we have that*

$$fd_\alpha^{\mathrm{Aa}} = bd_\alpha^{\mathrm{Aa}}$$
$$fd_\alpha^{\mathrm{As}} = bd_\alpha^{\mathrm{As}}$$
$$fd_\alpha^{\mathrm{Sa}} = fd_\alpha^{\mathrm{Ss}} = bd_\alpha^{\mathrm{Ss}}.$$

Proof. Let $\alpha \in (0,1]$. We show that $fd_\alpha^{\mathrm{Aa}} = bd_\alpha^{\mathrm{Aa}}$; the other cases are similar. First, note that the operator H_α^{Aa} used in Definition 13 to define the branching distances can be equivalently replaced by the following operator $G : (S^2 \to \widehat{\mathbb{R}}_+) \to (S^2 \to \widehat{\mathbb{R}}_+)$ by

$$G(d)(s,t) = pd(s,t) \sqcup d(s,t) \sqcup \sup_{s' \in \tau(s)} \inf_{t' \in \tau(t)} \alpha \cdot d(s',t').$$

For convenience, let also $F = F_\alpha^{\mathrm{a}}$. Then bd_α^{Aa} is the least fixpoint of G and fd_α^{Aa} is the least fixpoint of F. Since $G(d) \leq F(d)$ for all $d : S^2 \to \widehat{\mathbb{R}}_+$, we get by monotonicity of G and F that $bd_\alpha^{\mathrm{Aa}} \leq fd_\alpha^{\mathrm{Aa}}$. To prove that $fd_\alpha^{\mathrm{Aa}} \leq bd_\alpha^{\mathrm{Aa}}$, we define for each $k \in \mathbb{N}$

$$F_k(d)(s,t) = pd(s,t) \sqcup \sup_{\sigma \in Paths(s)} \inf_{\rho \in Paths(t)} \sup_{0 \leq i \leq k} \alpha^i d(\sigma_i, \rho_i).$$

We denote by G^k the operator G iterated k times, i.e. $G^0(d) = d$ and $G^{k+1}(d) = G(G^k(d))$. We show by induction that $F_k \leq G^k$ for all $k \geq 1$. For $k = 1$, we have $F_1(d) = pd \sqcup d \leq G^1(d)$. For $k+1$, we have:

$$\begin{aligned}
&F_{k+1}(d)(s,t) \\
&= pd(s,t) \sqcup \sup_{\sigma \in Paths(s)} \inf_{\rho \in Paths(t)} \sup_{0 \leq i \leq k+1} \alpha^i d(\sigma_i, \rho_i) \\
&= pd(s,t) \sqcup \sup_{s' \in \tau(s)} \sup_{\sigma' \in Paths(s')} \inf_{t' \in \tau(t)} \inf_{\rho' \in Paths(t')} \sup_{0 \leq i \leq k} (d(s,t) \sqcup \alpha^{i+1} d(\sigma'_i, \rho'_i)) \\
&\leq pd(s,t) \sqcup d(s,t) \sqcup \sup_{s' \in \tau(s)} \inf_{t' \in \tau(t)} \sup_{\sigma' \in Paths(s')} \inf_{\rho' \in Paths(t')} \sup_{0 \leq i \leq k} \alpha^{i+1} d(\sigma'_i, \rho'_i) \\
&= pd(s,t) \sqcup d(s,t) \sqcup \sup_{s' \in \tau(s)} \inf_{t' \in \tau(t)} F_k(d)(s',t') \\
&\leq pd(s,t) \sqcup d(s,t) \sqcup \sup_{s' \in \tau(s)} \inf_{t' \in \tau(t)} \alpha G^k(d)(s',t') \\
&= G^{k+1}(d)(s,t).
\end{aligned}$$

Then,

$$F(bd_\alpha^{\mathrm{Aa}}) = \lim_k F_k(bd_\alpha^{\mathrm{Aa}}) \leq \lim_k G^k(bd_\alpha^{\mathrm{Aa}}) = bd_\alpha^{\mathrm{Aa}}.$$

Together with $F(d) \geq d$ for all d, this shows $F(bd_\alpha^{\mathrm{Aa}}) = bd_\alpha^{\mathrm{Aa}}$, i.e. bd_α^{Aa} is a fixpoint of F. Hence, $bd_\alpha^{\mathrm{Aa}} \geq fd_\alpha^{\mathrm{Aa}}$, since fd_α^{Aa} is the least fixpoint of F. □

With this result, we can finally prove Theorem 18.

Proof of Theorem 18.

1. The inequalities follow from Lemmas 5 and 6, and from the monotonicity of the F^x_α, $\overline{F}^x_\alpha$ operators for $\alpha \in (0,1]$ and $x \in \{\mathrm{a},\mathrm{s}\}$. To see that on deterministic, non-boolean MTSs, the linear distances between states can be strictly smaller than the corresponding branching ones, consider the MTS in Figure 4. 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 $bd^{\mathrm{Aa}}_\alpha(s,t) = bd^{\mathrm{As}}_\alpha(s,t) = \overline{bd}^{\mathrm{Aa}}_\alpha(s,t) = \overline{bd}^{\mathrm{As}}_\alpha(s,t) = \alpha^2$.
2. Let $M = (S,\tau,\Sigma,[\cdot])$ be a boolean, deterministic MTS, let $\alpha \in (0,1]$ and $s,t \in S$. We show that $ld^{\mathrm{a}}_\alpha = bd^{\mathrm{Aa}}_\alpha$. The other cases are similar. By part 1 of this theorem, we know that $ld^{\mathrm{a}}_\alpha \leq bd^{\mathrm{Aa}}_\alpha$. To prove that $ld^{\mathrm{a}}_\alpha \geq bd^{\mathrm{Aa}}_\alpha$, we show that $H^{\mathrm{Aa}}(ld^{\mathrm{a}}_\alpha) = ld^{\mathrm{a}}_\alpha$, i.e. that ld^{a}_α is a fixpoint of H^{Aa}. As bd^{Aa}_α is the least fixpoint of H^{Aa}, we obtain $ld^{\mathrm{a}}_\alpha \geq bd^{\mathrm{Aa}}_\alpha$. First, we observe that

$$\begin{aligned}
H^{\mathrm{Aa}}(ld^{\mathrm{a}}_\alpha)(s,t) &= pd(s,t) \sqcup \alpha \sup_{s' \in \tau(s)} \inf_{t' \in \tau(t)} ld^{\mathrm{a}}_\alpha(s',t') \\
&= pd(s,t) \sqcup \alpha \sup_{s' \in \tau(s)} \inf_{t' \in \tau(t)} \sup_{\sigma' \in Paths(s')} \inf_{\rho' \in Paths(t')} td_\alpha(\sigma',\rho') \\
&\geq pd(s,t) \sqcup \alpha \sup_{s' \in \tau(s)} \sup_{\sigma' \in Paths(s')} \inf_{t' \in \tau(t)} \inf_{\rho' \in Paths(t')} td_\alpha(\sigma',\rho') \\
&= \sup_{\sigma \in Paths(s)} \inf_{\rho \in Paths(t)} td_\alpha(\sigma,\rho) \\
&= ld^{\mathrm{a}}_\alpha(s,t).
\end{aligned}$$

So $H^{\mathrm{Aa}}(ld^{\mathrm{a}}_\alpha)(s,t) \geq ld^{\mathrm{a}}_\alpha(s,t)$. We show that also $H^{\mathrm{Aa}}(ld^{\mathrm{a}}_\alpha)(s,t) \leq ld^{\mathrm{a}}_\alpha(s,t)$. If $pd(s,t) = 1$, then $H^{\mathrm{Aa}}(ld^{\mathrm{a}}_\alpha)(s,t) = ld^{\mathrm{a}}_\alpha(s,t) = 1$. Hence, assume $pd(s,t) = 0$. We distinguish two cases.

Case 1: $\sup_{s' \in \tau(s)} \inf_{t' \in \tau(t)} pd_\alpha(s',t') = 1$. Then one easily shows that

$$H^{\mathrm{Aa}}(ld^{\mathrm{a}}_\alpha)(s,t) = \alpha = ld^{\mathrm{a}}_\alpha(s,t).$$

Case 2: $\sup_{s' \in \tau(s)} \inf_{t' \in \tau(t)} pd_\alpha(s',t') = 0$.
Since M is deterministic and boolean, we know that for all $s' \in \tau(s)$, there is a $t_{s'} \in \tau(t)$ such that $pd_\alpha(s',t_{s'}) = 0$ and $pd_\alpha(s',t') = 1$ for $t' \neq t_{s'}$. Then, we have for all $s' \in \tau(s), t' \in \tau(t), t' \neq t_{s'}$, $\sigma' \in Paths(s')$, $\rho' \in Paths(t')$, and $\rho_{s'} \in Paths(t_{s'})$ that

$$td_\alpha(\sigma',\rho_{t_{s'}}) \leq \alpha \quad \text{and} \quad td_\alpha(\sigma',\rho') = 1$$

and therefore

$$\inf_{\rho' \in Paths(t_{s'})} td_\alpha(\sigma',\rho') \leq \inf_{\rho' \in Paths(t')} td_\alpha(\sigma',\rho')$$

so

$$\inf_{\rho' \in Paths(t_{s'})} td_\alpha(\sigma',\rho') \leq \inf_{t' \in \tau(t)} \inf_{\rho' \in Paths(t')} td_\alpha(\sigma',\rho'). \qquad (*)$$

Recalling that $pd(s,t)=0$, we get

$$\begin{aligned}
H^{\mathrm{Aa}}(ld_\alpha^{\mathrm{a}})(s,t) &= \alpha \sup_{s'\in\tau(s)} \inf_{t'\in\tau(t)} \sup_{\sigma'\in Paths(s')} \inf_{\rho'\in Paths(t')} td_\alpha(\sigma',\rho') \\
&\le \alpha \sup_{s'\in\tau(s)} \sup_{\sigma'\in Paths(s')} \inf_{\rho'\in Paths(t_{s'})} td_\alpha(\sigma',\rho') && \text{by (*)} \\
&\le \alpha \sup_{s'\in\tau(s)} \sup_{\sigma'\in Paths(s')} \inf_{t'\in\tau(t)} \inf_{\rho'\in Paths(t')} td_\alpha(\sigma',\rho') \\
&= \sup_{\sigma\in Paths(s)} \inf_{\rho\in Paths(t)} td_\alpha(\sigma,\rho) \\
&= ld_\alpha^{\mathrm{a}}(s,t).
\end{aligned}$$

To see that the equalities cannot be strengthened to equalities, consider $\alpha \in (0,1]$. We give the proof for $\alpha > \frac{1}{2}$; a similar example works if $\alpha \le \frac{1}{2}$. Consider the MTS in Figure 4. Then $ld_\alpha^x(s,t) = \frac{1}{2}\alpha$, while $bd_\alpha^x(s,t) = \alpha^2$. □

6 Conclusions

In this paper, we have provided metric extensions of the classical linear and branching relations: trace inclusion, trace equivalence, simulation, and bisimulation. We remark that, while metric analogous of bisimulation had been known for some time [7, 18], this is not the case for the other notions, which had escaped attention thus far. We hope that the introduction of these quantitative asymmetrical and symmetrical distances constitutes a useful step toward achieving a *quantitative theory of systems,* in which the classical boolean setting of specification and verification is replaced by a setting in which properties have (real-valued, or general) values, and verification can yield not only yes/no answers, but also measures of quality, adequacy, and cost.

We have provided three main classes of characterizations for linear and branching distances:

1. *Distances as upper bounds for logic valuations.* Results in this class state that the distances provide an upper bound for the difference in value of formulas of linear (QLTL) and branching (QMU) logics. Results of this type are Theorems 4 and 13.
2. *Logics as full characterizations of distances.* Results in this class state that the distances are equal to the supremum of the difference in value of all linear, or branching formulas. Results of this type are Theorems 5 and 14.
3. *Relations among distances.* Results in this class compare the value of linear and branching distances; results of this type are Theorems 2, 11, and 18.

Results in classes 1 and 3 hold for general MTSs, and are thus particularly satisfying. In contrast, as we have seen, results in class 2 hold only for finitely branching MTSs. Many MTSs of interest are not finite branching: for instance, in a hybrid system, there can be uncountably many successors of a state, corresponding to the real-valued length of time steps possible from the state. It is an interesting open problem to investigate classes of MTSs that are more general than finitely branching MTSs, and for which results of class 2 still hold.

References

1. M.C. Browne, E.M. Clarke and O. Grumberg. Characterizing finite Kripke structures in propositional temporal logic. In *Theoretical Computer Science*, vol. 59, Springer 1988.
2. P. Caspi and A. Benveniste. Toward an approximation theory for computerized control. In *Proceedings of EMSOFT*, volume 2491 of *LNCS*, pages 294–304, 2002.
3. L. de Alfaro, M. Faella, and M.I.A. Stoelinga. Linear and Branching Metrics for Quantitative transition systems. In *Proceedings of ICALP*, volume 3142 of *LNCS*, pages 97–109 Springer, 2004.
4. L. de Alfaro, T.A. Henzinger, and R. Majumdar. Discounting the future in systems theory. In *Proceedings of ICALP*, volume 2719 of *LNCS*, pages 1022–1037. Springer, 2003.
5. L. de Alfaro and R. Majumdar. Quantitative solution of omega-regular games. In *Proceedings STOC*, pages 675–683. ACM Press, 2001.
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 *Proceedings of CONCUR*, volume 1664 of *LNCS*, pages 258–273, 1999.
8. J. Filar and K. Vrieze. *Competitive Markov Decision Processes.* Springer-Verlag, 1997.
9. P. Fletcher and W.F. Lindgren. *Quasi-uniform spaces*, volume 77 of *Lecture Notes in Pure and Applied Mathematics.* Marcel Dekker Inc., New York, 1982.
10. M. Huth and M.Z. Kwiatkowska. Quantitative analysis and model checking. In *Proceedings of LICS*, pages 111–122, 1997.
11. D. Kozen. A probabilistic PDL. In *Proc. 15th ACM Symp. Theory of Comp.*, pages 291–297, 1983.
12. 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.
13. R. Majumdar. *Symbolic algorithms for verification and control.* PhD thesis, University of California, Berkeley, 2003.
14. Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems: Specification.* Springer-Verlag, New York, 1991.
15. A. McIver and C. Morgan. Games, probability, and the quantitative μ-calculus q$M\mu$. In *Proceedings of LPAR*, pages 292–310, 2002.
16. J.H. Reif. Universal games of incomplete information. In *11th Annual ACM Symposium on Theory of Computing*, pages 288–308, April, Atlanta, Georgia 1979.
17. 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.
18. F. van Breugel and J. Worrel. An algorithm for quantitative verification of probabilistic transition systems. In *Proceedings of CONCUR*, volume 2154 of *LNCS*, pages 336–350, 2001.