Contextual Behavioural Metrics

Ugo Dal Lago
University of Bologna, Italy
INRIA Sophia Antipolis, France

Maurizio Murgia
Gran Sasso Science Institute, L'Aquila, Italy

Abstract

We introduce contextual behavioural metrics (CBMs) as a novel way of measuring the discrepancy in behaviour between processes, taking into account both quantitative aspects and contextual information. This way, process distances by construction take the environment into account: two (non-equivalent) processes may still exhibit very similar behaviour in some contexts, e.g., when certain actions are never performed. We first show how CBMs capture many well-known notions of equivalence and metric, including Larsen's environmental parametrized bisimulation. We then study compositional properties of CBMs with respect to some common process algebraic operators, namely prefixing, restriction, non-deterministic sum, parallel composition and replication.

2012 ACM Subject Classification Theory of computation → Program reasoning; Theory of computation → Process calculi

Keywords and phrases Behavioural metrics, Labelled Transition Systems, Differential Semantics

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2023.38

Related Version *Full Version*: `https://arxiv.org/abs/2307.07400`

Funding The first author is supported by the ERC CoG "DIAPASoN", GA 818616, and by PNRR – M4C2 – I1.4 Project CN00000013 "National Centre for HPC, Big Data and Quantum Computing". The second author is supported by MUR project PON REACT EU DM 1062/21.

Acknowledgements We thank the anonymous reviewers for their feedback.

1 Introduction

Simulation and bisimulation relations are often the methodology of choice for reasoning relationally about the behaviour of systems specified in the form of LTSs. On the one hand, most of them can be proved to be congruences, therefore enabling modular equivalence proofs. On the other hand, not being based on any universal quantification (e.g. on tests or on traces), they enable simpler relational arguments, especially when combined with enhancements such as the so-called up-to techniques [29].

The outcome of relational reasoning as supported by (bi)simulation relations is inherently binary: two programs or systems are *either* (bi)similar *or not so*. As an example, all pairs of non-equivalent elements have the same status, i.e. the bisimulation game gives no information on the degree of dissimilarity between non-equivalent states. This can be a problem in those contexts, such as that of probabilistic systems, in which non-equivalent states can give rise to completely different but also extremely similar behaviours.

This led to the introduction of a generalization of bisimulation relations, i.e. the so-called bisimulation *metrics* [7], which rather than being binary relations on the underlying set of states S, are binary *maps* from S to a quantale (most often of real numbers) satisfying the axioms of (pseudo)metrics. In that context, the bisimulation game becomes inherently quantitative: the defender aims at proving that the two states at hand are *close* to each other, while the attacker tries to prove that they are *far apart*. The outcome of this game is a quantity representing a bound not only on any discrepancy about the *immediate* behaviour of the two involved states, (e.g. the fact that some action is available in s but not in t), but also providing some information about differences which will only show up in the *future*, all this *regardless* of the actions chosen by the attacker. In this sense, therefore, bisimulation metrics condense a great deal of information in just one number.

Notions of bisimulation metrics have indeed be defined for various sequential and concurrent calculi (see, e.g., [4, 9, 11, 13, 14, 33]), allowing a form of metric reasoning on program behaviour. But when could any of such techniques be said to be compositional? This amounts to be able to *derive* an upper bound on the distance $\delta(C[t], C[s])$ between two programs in the form $C[t]$ and $C[s]$ *from* the distance $\delta(s, t)$ between s and t. Typically, the latter is required to be itself an upper bound on the former, giving rise to *non-expansiveness* as a possible generalization of the notion of a congruence. This, however, significantly *restricts* the class of environments C to which the aforementioned analysis can be applied, since being able to amplify differences is a very natural property of processes. Indeed, an inherent tension exists between expressiveness and compositionality in metric reasoning [14].

But there is another reason why behavioural metrics can be seen as less informative than they could be. As already mentioned, any number measuring the distance between two states s and t implicitly accounts for all the possible ways of comparing s and t, i.e. any context. Often, however, only contexts that act in a certain very specific way could highlight large differences between s and t, while others might simply see s and t as very similar, or even equivalent. This further dimension is abstracted away in compositional metric analysis: if the distance between s and t is very high, but C does not "take advantage" of such large differences, $C[s]$ and $C[t]$ should be *close* to each other, but are dubbed being *far away* from each other, due to the aforementioned abstraction step. It is thus natural to wonder whether metric analysis can be made contextual. In the realm of process equivalences, this is known to be possible through, e.g. Larsen's environmental parametrized bisimulation [23], but not much is known about contextual enhancements of bisimulation *metrics*. Other notions of program equivalence, like logical relations or denotational semantics, have been shown to have metric analogues [6, 30], which in some cases can be made contextual [17, 22].

In this paper, we introduce the novel notion of *contextual behavioural metric* (CBM in the following) through which it is possible to fine-tune the abstraction step mentioned above and which thus represents a refinement over behavioural metrics. In CBMs, the distance between two states s, t of an LTS is measured by an object d having a richer structure than that of a number. Specifically, d is taken to be an element of a metric transition system, in which the contextual and temporal dimensions of the differences can be taken into account. In addition to the mere introduction of this new notion of distance, our contributions are threefold:

- On the one hand, we show that metric labelled transition systems (MLTSs in the following), namely the kind of structures meant to model differences, are indeed quantales, this way allowing us to prove that CBMs are generalized metrics. This is in Section 3.
- On the other hand, we prove that some well-known methodologies for qualitative and quantitative relational reasoning on processes, namely (strong) bisimulation relations and metrics, and environmental parametrized bisimulations [23], can all be seen as CBMs where the underlying MLTS corresponds to the original quantale. This is in Section 4.
- Finally, we prove that CBMs have some interesting compositional properties, and that this allows one to derive approximations to the distance between processes following their syntactic structure. This is in Section 5.

Many of the aforementioned works about behavioural metrics are concerned with probabilistic forms of LTSs. In this work, instead, we have deliberately chosen to focus on usual nondeterministic transition systems. On the one hand, the quantitative aspects can be handled through the so-called *immediate* distance between states, see below. On the other hand, it is well known that probabilistic transition systems can be seen as (non)deterministic systems whose underlying reduction relation is defined between state distributions. Focusing on ordinary LTSs has the advantage of allowing us to concentrate our attention on those aspects related to metrics, allowing for a separation of concerns. This being said, we are confident that most of the results described here could hold for probabilistic LTSs, too.

2 Why the Environment Matters

The purpose of this section is to explain why purely numerical quantales do *not* precisely capture differences between states of an LTS and how a more structured approach to distances can be helpful to tackle this problem. We will do this through an example drawn from the realm of higher-order programs, the latter seen as states of the LTS induced by Abramsky's applicative bisimilarity [1].

Let us start with a pair of programs written in a typed λ-calculus, both of them having type $(\mathtt{Nat} \to \mathtt{Nat}) \to \mathtt{Nat}$, namely M_2 and M_4, where $M_n \triangleq \lambda x.xn$. These terms can indeed be seen as states of an LTS, whose relevant fragment is the following one:

$$M_2 \xrightarrow{\quad V \quad} V\ 2 \xrightarrow{\quad eval \quad} \mathsf{E}(V\ 2)$$

$$M_4 \xrightarrow{\quad V \quad} V\ 4 \xrightarrow{\quad eval \quad} \mathsf{E}(V\ 4)$$

Labelled transitions correspond to either parameter passing (each actual parameter being captured by a distinct label V) or evaluation. It is indeed convenient to see the underlying LTS as a bipartite structure whose states are either computations or values. The two states $\mathsf{E}(V\ 2)$ and $\mathsf{E}(V\ 4)$ are the natural number values to which $V\ 2$ and $V\ 4$ evaluate, respectively. Clearly, the latter are not to be considered equivalent whenever different, and this can be captured, e.g., by either exposing the underlying numerical value through a labelled self-transition or by stipulating that base type values, contrary to higher-order values, can be explicitly observed, thus being equivalent precisely when equal. If one plays the bisimulation game on top of this LTS, the resulting notion of equivalence turns out to be precisely Abramsky's applicative bisimilarity. For very good reasons, M_2 and M_4 are dubbed as *not* equivalent: they can be separated by feeding, e.g. $V = \lambda x.x$ to them.

But now, *how far apart* should M_2 and M_4 be? The answer provided by behavioural metrics consists in saying that M_2 and M_4 are at distance *at most* $x \in \mathbb{R}^{\infty}_{+}$ iff x is an upper bound on the differences any adversary observes while interacting with them, *independently* on how the adversary behaves. As a consequence, if the underlying λ-calculus provides a primitive for multiplication, then it is indeed possible to define values of the form $V_n \triangleq \lambda x.x \times n$ for every n, allowing the environment to observe arbitrarily large differences of the form

$$|\ \mathsf{E}(V_n\ 2) - \mathsf{E}(V_n\ 4)\ | = |\ 2n - 4n\ | = 2n.$$

In other words, the distance between M_2 and M_4 is $+\infty$. The possibility of arbitrarily amplifying distances is well-known, and can be tackled, e.g., by switching to a calculus in which *all* functions are *non-expansive*, ruling out terms such as V_n where $n > 1$. In other words, the distance between M_2 and M_4 is indeed 2, because no input term V can "stretch" the distance between 2 to 4 to anything more than 2. This is what happens, e.g., in Fuzz [30].

But is this the end of the story? Are we somehow losing too much information by stipulating that M_2 and M_4 are, say, at distance 2? Actually, the only moment in which the environment observes the state with which it is interacting is at the end of the dialogue, namely after feeding it with a function $V : \texttt{Nat} \to \texttt{Nat}$. If, for example, the environment picks $V_q \triangleq \lambda x.(x-3)^2+2$, then the observed difference is 0, while if it picks $V_l \triangleq \lambda x.x+2$ then the observed distance is maximal, i.e. 2. In other words, the observed distance strictly depends on how the environment behaves and should arguably be parametrised on it. This is indeed the main idea behind Larsen's environmental parametrised bisimulation, but also behind our contextual behavioural metrics. In the latter, differences can be faithfully captured by the states of *another* labelled transition system, called a metric labelled transition system, in which observed distances are associated to states. In our example, the difference between M_2 and M_4 *is* the state s of a metric labeled transition system whose relevant fragment is:

$$\begin{tikzcd}
 & t_q \arrow[r, "eval"] & u_q \\
s \arrow[ur, "V_q"] \arrow[dr, "V_l"'] & & \\
 & t_l \arrow[r, "eval"] & u_l
\end{tikzcd}$$

Crucially, while s, t_s, t_l, u_l are all mapped to the null observable difference, u_q is associated to 2. This allows to discriminate between those environments which are able to see large differences from those which are not. This is achieved by allowing differences to be modelled by the states of a transition system *themselves*. Using a categorical jargon, it looks potentially useful, but also very tempting, to impose the structure of a coalgebra to the underlying space of distances rather than taking it as a monolithical, numeric, quantale. The rest of this paper can be seen as an attempt to make this idea formal.

3 Contextual Behavioural Metrics, Formally

This section is devoted to introducing contextual behavioural metrics, namely the concept we aim at studying in this paper. We start with the definition of quantale [31], the canonical codomain of generalized metrics [24]. The notion of quantale used in this paper is that of unital integral commutative quantale:

▶ **Definition 1** (Quantale). *A* quantale *is a structure* $\mathbb{Q} = (Q, \bigwedge, \bigvee, \perp, \top, +)$ *such that* $\bigwedge, \bigvee : 2^Q \to Q$*, the two objects* $\perp, \top$ *are in* Q*, and* $+$ *is a binary operation on* Q*, where:*
- $(Q, \bigwedge, \bigvee, \perp, \top)$ *is a complete lattice;*
- $(Q, +, \perp)$ *is a commutative monoid;*
- *for every* $e \in Q$ *and every* $A \subseteq Q$ *it holds that* $e + \bigwedge A = \bigwedge\{e + f \mid f \in A\}$*.*

We write $e \leq f$ *when* $e = \bigwedge\{e, f\}$*.*

Generalized metrics are maps which associate an element of a given quantale to each pair of elements. As customary in behavioural metrics, we work with *pseudo*metrics, in which distinct elements may be at minimal distance:

▶ **Definition 2** (Metrics). *A* pseudometric *over a set* A *with values in a quantale* $\mathbb{Q}$ *is a map* $m : A \times A \to \mathbb{Q}$ *satisfying:*
- *for all* $a \in A : m(a, a) = \perp$*;*
- *for all* $a, b \in A : m(a, b) = m(b, a)$*;*
- *for all* $a, b, c \in A : m(a, c) \leq m(a, b) + m(b, c)$*.*

In the rest of this paper, we refer to pseudometrics simply as metrics.

It is now time to introduce our notion of a *process*, namely of the computational objects we want to compare. We do not fix a syntax, and work with abstract labelled transition systems (LTSs in the following). In order to enable (possibly quantitative) metric reasoning, we equip states of our LTS with an immediate metric D, namely a metric measuring the observable distance between two states.

▶ **Definition 3** (Process LTS). *We define a $\mathbb{Q}$-LTS as a quadruple $(P, \mathcal{L}, \rightarrow, D)$ where:*
- *P is the set of processes;*
- *$\mathcal{L}$ is the set of labels;*
- *$\rightarrow \subseteq P \times \mathcal{L} \times P$ is the transition relation;*
- *$D : P \times P \rightarrow \mathbb{Q}$ is a metric.*

▶ **Example 4.** The example LTS from Section 2 should be helpful in understanding why the metric D is needed: terms and values of distinct types are at maximal immediate distance, while terms and values of the same type are at minimal distance, except when the type is `Nat`, whereas the immediate distance is just the absolute value between the two numbers.

We now need to introduce *another* notion of transition system, this time meant to model *differences* between computations. This kind of structure can be interpreted as a quantale, and will form the codomain of Contextual Bisimulation Metrics. Intuitively, a Metric LTS is an LTS endowed with a function from states to a quantale $\mathbb{Q}$. This allows to keep track of immediate distance changes. Let us start with the notion of a *pre*-metric LTS:

▶ **Definition 5** (Pre-metric LTS). *A* pre-metric $\mathbb{Q}$-LTS *is a quadruple $\mathbb{V} = (S, \mathcal{L}, \rightarrow, \Downarrow)$ where:*
- *S is the set of states;*
- *$\mathcal{L}$ is the set of labels;*
- *$\rightarrow \subseteq S \times \mathcal{L} \times S$ is the transition relation;*
- *$\Downarrow : S \rightarrow \mathbb{Q}$ is a function which assigns values in $\mathbb{Q}$ to states in S.*

A pre-metric LTS does not necessarily form a quantale, because S does not necessarily have, e.g. the structure of a monoid or a lattice. In order to be proper codomains for metrics, pre-metric LTSs need to be endowed with some additional structure, which will be proved to be enough to form a quantale.

▶ **Definition 6** (Metric LTS). *A metric $\mathbb{Q}$-LTS $\mathbb{V} = (S, \mathcal{L}, \rightarrow, \Downarrow)$ is a pre-metric $\mathbb{Q}$-LTS endowed with two elements $\bot_{\mathbb{V}}, \top_{\mathbb{V}} \in S$, and three operators $\bigwedge_{\mathbb{V}}, \bigvee_{\mathbb{V}} : 2^S \rightarrow S$ and $+_{\mathbb{V}} : S \times S \rightarrow S$, where the conditions hold for all possible values of the involved metavariables:*

$$\bot_{\mathbb{V}} \xrightarrow{\ell} s \iff s = \bot_{\mathbb{V}} \qquad\qquad \Downarrow \bot_{\mathbb{V}} = \bot_{\mathbb{Q}}$$

$$\forall \ell \in \mathcal{L} : \top_{\mathbb{V}} \not\xrightarrow{\ell} \qquad\qquad \Downarrow \top_{\mathbb{V}} = \top_{\mathbb{Q}}$$

$$\textstyle\bigwedge_{\mathbb{V}} S' \xrightarrow{\ell} s \iff \exists s' \in S' : s' \xrightarrow{\ell} s \qquad\qquad \Downarrow \bigwedge_{\mathbb{V}} S' = \bigwedge_{\mathbb{Q}} \{\Downarrow s \mid s \in S'\}$$

$$\begin{array}{c}\textstyle\bigvee_{\mathbb{V}} S' \xrightarrow{\ell} s \iff \exists S'' : s = \bigvee_{\mathbb{V}} S'' \textit{ and} \\ \exists \textit{ surjective } f : S' \rightarrow S'' : \forall s' \in S' : s' \xrightarrow{\ell} f(s')\end{array} \qquad\qquad \Downarrow \textstyle\bigvee_{\mathbb{V}} S' = \bigvee_{\mathbb{Q}} \{\Downarrow s \mid s \in S'\}$$

$$\begin{array}{c}s_1 +_{\mathbb{V}} s_2 \xrightarrow{\ell} s' \iff s' = s_1' +_{\mathbb{V}} s_2' \textit{ for some } s_1', s_2' \\ \textit{such that: } s_1 \xrightarrow{\ell} s_1' \textit{ and } s_2 \xrightarrow{\ell} s_2'\end{array} \qquad\qquad \Downarrow (s_1 +_{\mathbb{V}} s_2) = \Downarrow s_1 +_{\mathbb{Q}} \Downarrow s_2$$

Axioms ensures that $\bot_{\mathbb{V}}$ allows every possible behaviour (somehow capturing every context), and dually $\top_{\mathbb{V}}$ disallows every behaviour. $\bigwedge_{\mathbb{V}} S'$ allows all and only the behaviours in S' (union of contexts), while $\bigvee_{\mathbb{V}} S'$ enables all and only the behaviours allowed by *every* element in S' (intersection of contexts). The sum $+_{\mathbb{V}}$ has a behaviour similar to $\bigvee_{\mathbb{V}}$, but it is binary and differs on the value returned by $\Downarrow$.

▶ Remark 7 (On The Existance Of Non-Trivial MLTSs). Due to the requirements about joins and meets over potentially infinite sets, MLTSs are not easy to define directly. We argue, however, that an MLTS can be defined as the closure of a pre-MLTS. If the underlying quantale $\mathbb{Q}$ is boolean, one can get the desired structure by considering 2^{2^X}, where X is the carrier of the given pre-MLTS: it suffices to take subsets in "conjunctive" normal form. For the general case, the class $\cup_{n\in\mathbb{N}} \overbrace{2^{\cdot^{\cdot^{\cdot^{2^X}}}}}^{n \text{ times}}$, which is indeed a set in ZFC, suffices.

The axiomatics above is still not sufficient to give the status of a quantale to $\mathbb{Q}$-MLTSs. The reason behind all this is that there could be equivalent but distinct states in S. We then define a preorder $\leq_{\mathbb{V}}$ on the states of any MLTS $\mathbb{V}$:

▶ **Definition 8.** *A relation $\mathcal{R} \subseteq S \times S$ is a $\leq_{\mathbb{Q}}$-*preserving simulation[1] *if, whenever $s_1 \mathcal{R} s_2$, it holds that:*

1. $\Downarrow s_1 \ \leq_{\mathbb{Q}} \Downarrow s_2$*;*
2. $\forall \ell \in \mathcal{L} : s_2 \xrightarrow{\ell} s_2' \implies \exists s_1' : s_1 \xrightarrow{\ell} s_1'$ *and* $s_1'\mathcal{R}s_2'$*.*

We define $\leq_{\mathbb{V}} \subseteq S \times S$ as the largest $\leq_{\mathbb{Q}}$-preserving simulation. We use the notation $\leq\geq_{\mathbb{V}}$ for mutual $\leq_{\mathbb{Q}}$-preserving simulation, that is $\leq\geq_{\mathbb{V}} = \leq_{\mathbb{V}} \cap \geq_{\mathbb{V}}$. We say that s is a lower (resp. upper) bound of $S' \subseteq S$ if $s \leq_{\mathbb{V}} s'$ (resp. $s' \leq_{\mathbb{V}} s$) for all $s' \in S'$.

The forthcoming result states that, in general, MLTSs *almost* form quantales. We can recover a proper quantale by quotienting S modulo $\leq\geq_{\mathbb{V}}$.

▶ **Proposition 9** (Properties of MLTSs)**.** *Let $\mathbb{V} = (S, \mathcal{L}, \rightarrow, \Downarrow)$ be a MLTS. Then:*

1. *$\leq_{\mathbb{V}}$ is a preorder relation;*
2. *For all s: $\perp_{\mathbb{V}} \leq_{\mathbb{V}} s$ and $s \leq_{\mathbb{V}} \top_{\mathbb{V}}$;*
3. *For all $S' \subseteq S$: $\bigwedge_{\mathbb{V}} S'$ is a lower bound of S', and if s' is a lower bound of S' then $s' \leq_{\mathbb{V}} \bigwedge_{\mathbb{V}} S'$.*
4. *For all $S' \subseteq S$: $\bigvee_{\mathbb{V}} S'$ is an upper bound of S', and if s' is an upper bound of S' then $\bigwedge_{\mathbb{V}} S' \leq_{\mathbb{V}} s'$.*
5. *For all $s \in S, S' \subseteq S : s +_{\mathbb{V}} \bigwedge_{\mathbb{V}} S' \leq\geq_{\mathbb{V}} \bigwedge_{\mathbb{V}} \{s + s' \mid s' \in S'\}$.*
6. *For all $s \in S : s +_{\mathbb{V}} \perp \leq\geq_{\mathbb{V}} s$.*
7. *For all $s, s' \in S : s +_{\mathbb{V}} s' \leq\geq_{\mathbb{V}} s' + s$.*
8. *For all $s, s', s'' \in S : (s +_{\mathbb{V}} s') +_{\mathbb{V}} s'' \leq\geq_{\mathbb{V}} s' +_{\mathbb{V}} (s +_{\mathbb{V}} s'')$.*
9. *If $\leq_{\mathbb{V}}$ is a partial order relation, then $\mathbb{V}$ is a quantale.*

Unless stated otherwise, we assume that every MLTS $\mathbb{V}$ we work with is a quantale.

▶ **Definition 10** (Contextual Behavioural Metrics)**.** *Let $(P, \mathcal{L}, \rightarrow, D)$ and $\mathbb{V} = (S, \mathcal{L}, \rightarrow, \Downarrow)$ be, respectively, a $\mathbb{Q}$-LTS and a $\mathbb{Q}$-MLTS. Then, a map $m : P \times P \rightarrow S$ is a* contextual bisimulation map *if:*

1. $D(p, q) \ \leq_{\mathbb{Q}} \ \Downarrow m(p, q)$*;*
2. *if $m(p, q) \xrightarrow{\ell} s'$, then the following holds:*
   a. $p \xrightarrow{\ell} p' \implies \exists q' : q \xrightarrow{\ell} q'$ *and* $m(p', q') \leq_{\mathbb{V}} s'$*;*
   b. $q \xrightarrow{\ell} q' \implies \exists p' : p \xrightarrow{\ell} p'$ *and* $m(p', q') \leq_{\mathbb{V}} s'$*.*

We say that m is a contextual bisimulation metric *(CBM) if m is both a contextual bisimulation map and a metric. We define the contextual bisimilarity map δ as follows:*

$$\delta(p, q) = \bigwedge_{\mathbb{V}} \{m(p, q) \mid m \text{ is a contextual bisimulation map}\}$$

[1] Technically, it is a reverse simulation. We call it simulation for brevity.

The following result states that the contextual bisimilarity map is well behaved, being a contextual bisimulation map upper bounding any other such map:

▶ **Lemma 11.** *δ is a contextual bisimulation map. Moreover, for all contextual bisimulation maps m, and processes p, q, it holds that $\delta(p, q) \leq_{\mathbb{V}} m(p, q)$.*

We still do not know whether δ is a *metric*. We need a handy characterization of δ for that.

A Useful Characterization of CBMs

Larsen's environment parametrized bisimulations [23] is a variation on ordinary bisimulation in which the compared states are tested against environments of a specific kind, this way giving rise to a ternary relation. We here show that CBMs can be captured along the same lines. A formal comparison between CBMs and Larsen's approach is deferred to Section 4.3.

▶ **Definition 12** (Parametrized Bisimulation). *Let $(P, \mathcal{L}, \to, D)$ and $(S, \mathcal{L}, \to, \Downarrow)$ be, respectively, a $\mathbb{Q}$-LTS and a $\mathbb{Q}$-MLTS. An S-indexed family of relations $\{\mathcal{R}_s\}$ such that $\mathcal{R}_s \subseteq P \times P$ is said to be a* parametrized bisimulation *iff, whenever $p \mathcal{R}_s q$, it holds that $D(p, q) \leq_{\mathbb{Q}} \Downarrow s$, and $s \xrightarrow{\ell} s'$ implies:*

- $p \xrightarrow{\ell} p' \implies \exists q' : q \xrightarrow{\ell} q'$ *and* $p' \mathcal{R}_{s'} q'$*;*
- $q \xrightarrow{\ell} q' \implies \exists p' : p \xrightarrow{\ell} p'$ *and* $p' \mathcal{R}_{s'} q'$*.*

Parametrized bisimilarity is the largest parametrized bisimulation, namely the largest family $\{\sim_s\}$ such that $p \sim_s q$ if $p \mathcal{R}_s q$ for some parametrized bisimulation $\{\mathcal{R}_s\}$.

The fact that $\{\sim_s\}$ is indeed a parametrized bisimulation holds because parametrized bisimulations are closed under unions (defined point-wise), something which can be proved with a simple generalisation of standard techniques [26, 27]. Parametrized bisimilarity turns out to be strongly related to δ, this way providing a simple proof technique that will be heavily used in the rest of the paper.

▶ **Proposition 13.** *For all p, q, s, it holds that $\delta(p, q) \leq_{\mathbb{V}} s \iff p \sim_s q$.*

We are finally ready to state that δ satisfies the axioms of a metric.

▶ **Theorem 14.** *The contextual bisimulation map δ is a metric.*

4 Some Relevant Examples

This section is devoted to showing how well-known and heterogeneous notions of equivalence and distance can be recovered as CBMs for appropriate quantales and MLTSs.

4.1 Strong Bisimilarity as a CBM

We start recalling that strong bisimilarity [26, 27] is the largest strong bisimulation relation, that is a relation $\mathcal{R} \subseteq P \times P$ on the states of a plain LTS $(P, \mathcal{L}, \to)$ such that $p \mathcal{R} q$ implies:

- $p \xrightarrow{\ell} p' \implies \exists q' : q \xrightarrow{\ell} q'$ and $p' \mathcal{R} q'$;
- $q \xrightarrow{\ell} q' \implies \exists p' : p \xrightarrow{\ell} p'$ and $p' \mathcal{R} q'$.

The first thing we have to do to turn strong bisimilarity into a CBM is to define, given such an LTS $(P, \mathcal{L}, \to)$, a canonical immediate distance D on the boolean quantale $\mathbb{B}$, which we call the *canonical* distance:

$$D(p, q) = \begin{cases} \bot \text{ if } \forall \ell : p \xrightarrow{\ell} \iff q \xrightarrow{\ell} \\ \top \text{ otherwise} \end{cases}$$

That is, the immediate distance is $\bot$ precisely when the processes expose the same labels. Notice that immediate distance is not affected by possible future behavioural differences. Any LTS like this is said to be a *boolean* LTS. The boolean quantale can be turned very naturally into a MLTS: let $\mathbb{V}$ be $(\{\bot_{\mathbb{V}}, \top_{\mathbb{V}}\}, \mathcal{L}, \to, \Downarrow)$ where the transitions are self loops $\bot_{\mathbb{V}} \xrightarrow{\ell} \bot_{\mathbb{V}}$ for every $\ell \in \mathcal{L}$, and $\Downarrow$ just associates $\bot_{\mathbb{Q}}$ to $\bot_{\mathbb{V}}$ and $\top_{\mathbb{Q}}$ to $\top_{\mathbb{V}}$.

▶ **Proposition 15.** *Given any boolean LTS, δ is the characteristic function of bisimilarity, i.e. $\delta(p, q) = \bot_{\mathbb{V}} \iff p \sim q$.*

4.2 Behavioural CBMs

Most behavioural metrics from the literature are defined on *probabilistic* transition systems [8, 10, 34], differently from CBMs. Some probabilistic behavioural metrics can still be captured in our framework by using as states of the process LTS (sub)distributions of states of the original PLTS, e.g. the distribution based metric in [12]. Non-probabilistic behavioural metrics exist, e.g., the so-called "branching metrics" [5], which are indeed instances of behavioural metrics as defined below. Notice that our definition has a generic quantale $\mathbb{Q}$ as its codomain, while usually behavioural metrics take values in the interval $\mathbb{R}_{[0,1]}$.

Let us first recall what we mean by a behavioural metric here. A metric $M : P \times P \to \mathbb{Q}$ is said to be a *behavioural metric* if, for all pairs of states p, q, it holds that $D(p, q) \leq_{\mathbb{Q}} M(p, q)$ and, whenever $M(p, q) <_{\mathbb{Q}} \top_{\mathbb{Q}}$, we have that:

- $p \xrightarrow{\ell} p' \implies \exists q' : q \xrightarrow{\ell} q'$ and $M(p, q) \geq_{\mathbb{Q}} M(p', q')$;
- $q \xrightarrow{\ell} q' \implies \exists p' : p \xrightarrow{\ell} p'$ and $M(p, q) \geq_{\mathbb{Q}} M(p', q')$.

Intuitively, behavioural metrics can be seen as quantitative variations on the theme of a bisimulation: they associate a value from a quantale to each pair of processes (rather than a boolean), they are coinductive in nature. Moreover, they are based on the bisimulation game, i.e., any move of one of the two processes needs to be matched by some move of the other, at least when their distance is not maximal. Our definition is similar to the one in [12]. However, many behavioural metrics in literature deal with non-determinism through the Hausdorff lifting, that is by stipulating that $D(p, q) \leq_{\mathbb{Q}} M(p, q)$ and for all ℓ:

$$M(p, q) \geq_{\mathbb{Q}} \bigvee_{p \xrightarrow{\ell} p'} \bigwedge_{q \xrightarrow{\ell} q'} M(p', q') \qquad \text{and} \qquad M(p, q) \geq_{\mathbb{Q}} \bigvee_{q \xrightarrow{\ell} q'} \bigwedge_{p \xrightarrow{\ell} p'} M(p', q')$$

The two notions are equivalent if the process LTS is image-finite and $\mathbb{Q}$ is totally ordered, both conditions are often assumed to be true in the literature.

We now show how to interpret $\mathbb{Q}$ as a MLTS. Morally, we just fix $\mathbb{Q}$ as the set of states, the identity as $\Downarrow$, and self loops as transitions. This however violates the requirement that the top element has no outgoing transitions. We therefore add the element $\top_{\mathbb{V}}$. Notice that we still need $\top_{\mathbb{Q}}$, as it ensures that $\mathbb{V}$ is closed under $+_{\mathbb{V}}$. Let $\mathbb{V} = (S, \mathcal{L}, \to, \Downarrow)$ where $S = \mathbb{Q} \uplus \{\top_{\mathbb{V}}\}$, $\mathcal{L}$ is as in the underlying process LTS, transitions are the self loops of the form $s \xrightarrow{\ell} s$ for every $\ell \in \mathcal{L}$, and $s \in \mathbb{Q}$, $\Downarrow$ is the identity on $\mathbb{Q}$, and $\Downarrow (\top_{\mathbb{V}}) = \top_{\mathbb{Q}}$. Notice that, when $\Downarrow s <_{\mathbb{Q}} \top_{\mathbb{Q}}$, we have that:

$$\Downarrow s \leq_{\mathbb{Q}} \Downarrow s' \iff s \leq_{\mathbb{V}} s'. \tag{1}$$

We also have that for every behavioural metric there is a CBM that "agrees" on the quantitative distance between processes. This intuition is formalized as follows:

▶ **Proposition 16.** *Let M be a behavioural metric, and let m_M be defined as:*

$$m_M(p,q) = \begin{cases} M(p,q) & \text{if } M(p,q) <_{\mathbb{Q}} \top_{\mathbb{Q}} \\ \top_{\mathbb{V}} & \text{otherwise} \end{cases}$$

Then, m_M is a CBM and for every p,q it holds that $\Downarrow m_M(p,q) = M(p,q)$.

The agreement of M and m_M holds by definition. The fact that m_M is a CBM, instead is a consequence of the fact that transitions preserve M distances (by definition of behavioural metric), and that behavioural metrics are metrics, indeed.

4.3 On Environment Parametrised Bisimulation and CBMs

As already mentioned, the concept of a CBM is inspired by Larsen's environment parametrized bisimulation [23]. It should then come with no surprise that there is a relationship between the two, which is the topic of this section.

First, let us recall what an environment parametrized bisimulation is. Let $(P, \mathcal{L}, \to)$ and $(E, \mathcal{L}, \to)$ be LTSs. Elements of P are called processes, while elements of E are called environments. A E-indexed family of relations $\{\mathcal{R}_e\}$, where $\mathcal{R}_e \subseteq P \times P$ is a *environment parametrized bisimulation* (EPB in the following) if, whenever $p\ \mathcal{R}_e\ q$ and $e \xrightarrow{\ell} e'$:

- $p \xrightarrow{\ell} p' \implies \exists q' : q \xrightarrow{\ell} q'$ and $p' \mathcal{R}_{e'} q'$;
- $q \xrightarrow{\ell} q' \implies \exists p' : p \xrightarrow{\ell} p'$ and $p' \mathcal{R}_{e'} q'$.

Environment parametrized bisimilarity, denoted as $\sim_e$, is defined as $p \sim_e q$ iff $p\ \mathcal{R}_e\ q$ for some EPB $\mathcal{R}$. It turns out that $\sim_e$ is the largest EPB [23].

EPBs can be embedded into the CBMs framework as follows:

- fix $\mathbb{Q}$ as the boolean quantale $\mathbb{B}$, and define $D(p,q) = \begin{cases} \bot_{\mathbb{B}} & \text{if } \exists \ell : p \xrightarrow{\ell} \text{ and } q \xrightarrow{\ell} \\ \top_{\mathbb{B}} & \text{otherwise} \end{cases}$
- let $\mathbb{V}_E = (S, \mathcal{L}, \to, \Downarrow)$ be any MLTS such that for all $s \in S$ it holds that $\Downarrow s = \bot_{\mathbb{B}} \iff s \neq \top_{\mathbb{V}}$, and for all $e \in E$ there is $s_e \in S$ such that $e \precsim\succsim s_e$. Here $\precsim\succsim$ is strong mutual similarity on the disjoint union of $\mathbb{V}$ (forgetting $\Downarrow$) and E. When such conditions hold, we say that E is embedded into $\mathbb{V}_E$.

We remark that, for every E, there is an MLTS $\mathbb{V}_E$ enjoying the properties above, obtained by augmenting E with the immediate metric defined above (this gives rise to a pre-metric LTS, Definition 5) and by closing it with respect to the operations and constants $\bigvee, \bigwedge, \top, \bot$ of Definition 6. The intuition is that:

- Two processes should have minimal immediate distance if there is a non-empty context in which their immediate behaviour is equivalent. This is ensured by the fact that they exhibit at least a common label from their current state.
- $\mathbb{V}_E$ needs to precisely simulate the behaviours in E. We therefore require that every element of E has a corresponding element in s, with "equivalent behaviour". In this setting, mutual simulation turns out to be the appropriate notion of behavioural equivalence.

The link between environment parametrized bisimulations and CBMs is made formal by the following proposition.

▶ **Proposition 17.** *Let E be an environment LTS embedded into an MLTS $\mathbb{V}_E$. For every p, q and e, it holds that $p \sim_e q \implies \delta(p,q) \leq_{\mathbb{V}_E} s_e$.*

The proposition above ultimately follows from the fact that $p \sim_e q \iff p \sim_{s_e} q$ (where $\sim_{s_e}$ is parametrized bisimilarity Definition 12) together with Proposition 13.

5 About the Compositionality of CBMs

One of the greatest advantages of the bisimulation proof method is its modularity, which comes from the fact that, under reasonable assumptions, bisimilarity is a congruence. In a metric setting, one strives to obtain similar properties [14, 16], which take the form of non-expansiveness, or variations thereof.

In this section we study the compositionality properties of CBMs with respect to some standard process algebraic operators. We are interested in properties that generalise the concept of a congruence. Following the lines of [22, 28], our treatment will be contextual, meaning that the environment in which processes are deployed can indeed contribute to altering their distance, although in a controlled way.

In order to keep our theory syntax independent, we model operators f as functions $f : P^n \to P$ (where n is the arity of the operator). In particular, for each process operator f of arity n we define the function $\hat{f} : P^n \times S^n \to S$ as follows:

$$\hat{f}(p_1, \dots, p_n, s_1, \dots, s_n) = \bigvee_{\mathbb{V}} \{\delta(f(p_1, \dots, p_n), f(p'_1, \dots, p'_n)) \mid \forall 1 \leq i \leq n : \delta(p_i, p'_i) \leq_{\mathbb{V}} s_i\}$$

Intuitively, $\hat{f}(\vec{p}, \vec{s})$ bounds $\delta(f(\vec{p}), f(\vec{q}))$ whenever $\vec{q}$ is such that $\delta(p_i, q_i) \leq_{\mathbb{V}} s_i$ for every i. Moreover, $\hat{f}(\vec{p}, \vec{s})$ is the lowest among such bounds.

Of course, our compositionality results rely on some assumptions on the compositionality of the immediate metric D. Formally, we require that, for all operators f (with arity n), the following holds for every $p_1, \dots, p_n, q_1, \dots, q_n$:

$$D(f(p_1, \dots, p_n), f(q_1, \dots, q_n)) \leq_{\mathbb{Q}} D(p_1, q_1) +_{\mathbb{Q}} \dots +_{\mathbb{Q}} D(p_n, q_n). \tag{2}$$

Below, we will give results about when and under which condition the value of the operator $\hat{f}$ can be upper-bounded by a function on its parameters. We remark that our compositionality results apply to each operator independently.

For the sake of concreteness, we give some examples of processes and their metric analysis. To this purpose, let $\mathcal{L} = \{a, b\}$, fix $\mathbb{Q}$ as the boolean quantale $\mathbb{B}$ and let D be defined exactly as we did in Section 4.3 (i.e., D returns $\bot$ if the processes can fire some common action, $\top$ otherwise). Distances will take values from a MLTS $\mathbb{V}_0$ over $\mathbb{B}$. Similarly to Section 4.3, we require $\mathbb{V}_0$ to be such that for every $s \in S$ it holds that $\Downarrow s = \bot_{\mathbb{B}} \iff s \neq \top_{\mathbb{V}_0}$. Moreover, we assume that $\mathbb{V}_0$ is able to represent at least Milner's synchronisation trees [25]. For simplicity, we omit self loops of $\bot_{\mathbb{V}_0}$ from all the graphical representations of our MLTS. Of course, these assumptions hold only in the examples, while our results hold for general MLTSs.

5.1 Restriction

We assume restriction to be modelled by a $\mathcal{L}$-indexed family of unary operators ν_ℓ, and that P is closed under these operators. Their semantics is standard.

▶ **Example 18.** Let p_0 and q_0 be as in the following figure. We have that p_0 and q_0 have the exact same behaviour on the b branch, while we can observe differences on the a branch (q_1 can perform an action, p_1 is terminated). State s_0 captures exactly the similarities between p_0 and q_0: after a b move it reduces to $\bot$; after an a move, it reduces to s_1. We argue that s_1 captures the similarities between p_1 and q_1: since neither of the two can perform the action a, s_1 reduces to $\bot$ with label a, while it does not perform b actions because p_1 and q_1 "disagree" on such label. So $\delta(p_0, q_0) = s_0$. Processes $(\nu a)p_0$ and $(\nu a)q_0$ exhibit equivalent behaviour instead. In fact, operator (νa) filters out the problematic a branch. It is therefore the case that $\delta((\nu a)p_0, (\nu a)q_0) = \bot$.

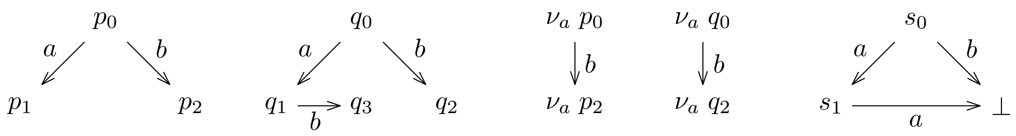

◀

The restriction operator does not add new behaviours to the original process, as it can only restrict it. We can then expect that the differences between any two processes do not increase if such processes are placed in a restriction context. Proposition below indeed shows that $\hat{\nu_\ell}$ enjoys a property similar to non-expansiveness, that is the distance between any two processes p and q bounds the distance between $\nu_\ell\ p$ and $\nu_\ell\ q$.

▶ **Proposition 19.** $\hat{\nu_\ell}(p, s) \leq_{\mathbb{V}} s$.

5.2 Prefixing

We assume that P is closed under operator $. : \mathcal{L} \times P \to P$, whose semantics is standard. We proceed similarly to the case of ν: we treat the prefix operator . as an $\mathcal{L}$-indexed family of unary operators $._\ell$.

▶ **Example 20.** Let p_0 and q_0 be as in Example 18. Since $b.p_0$ and $b.q_0$ can only reduce with a b move to, respectively, p_0 and q_0, their distance $\delta(b.p_0, b.q_0)$ should reduce to $\delta(p_0, q_0) = s_0$. Moreover, after performing an a action, $\delta(b.p_0, b.q_0)$ should reduce to $\bot$.

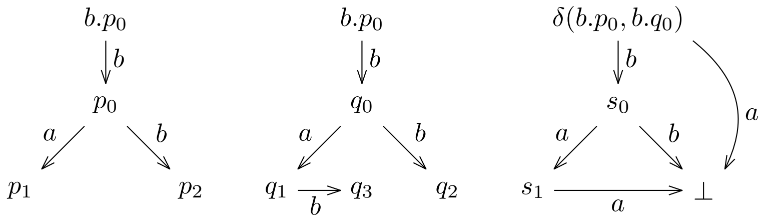

◀

In our contextual setting, prefixing of processes can change the distance, and the new distance may be incomparable to the original one. Therefore properties like non-expansiveness do not hold in general for $\hat{._\ell}$. Among the compositionality properties appeared in literature, uniform continuity [15] seems appropriate for prefixing. Uniform continuity holds when for all $s_\epsilon >_{\mathbb{V}} \bot_{\mathbb{V}}$ there is $s_\delta >_{\mathbb{V}} \bot_{\mathbb{V}}$ such that $\hat{._\ell}(p, s_\delta) \leq_{\mathbb{V}} s_\epsilon$. Such condition is too strong: for instance if $s_\epsilon \xrightarrow{\ell} \bot_{\mathbb{V}}$ the only option is to take $s_\delta = \bot_{\mathbb{V}}$, hence $s_\delta \not>_{\mathbb{V}} \bot_{\mathbb{V}}$.

For this reason, we need a stronger property for s_ϵ, namely that the meet of the set of ℓ reducts of s_ϵ is strictly greater than $\bot_{\mathbb{V}}$ and its immediate value is lower than that of s_ϵ.

▶ **Proposition 21.** *For all $s_\epsilon >_{\mathbb{V}} \bot_{\mathbb{V}}$ such that $s_\ell = \bigwedge_{\mathbb{V}}\{s \mid s_\epsilon \xrightarrow{\ell} s\} > \bot$ and $\Downarrow s_\ell \leq_{\mathbb{Q}} \Downarrow s_\epsilon$, there is $s_\delta >_{\mathbb{V}} \bot_{\mathbb{V}}$ such that $\hat{._\ell}(p, s_\delta) \leq_{\mathbb{V}} s_\epsilon$.*

5.3 Non-deterministic Sum

We assume that P is closed under binary operator $+$, whose semantics is again standard.

▶ **Example 22.** Let p_0, q_0 and s_0 be as in Example 18, and r_0 as in the picture below. We have that $\delta(p_0 + r_0, q_0 + r_0) = s_0$: it reduces to $\bot$ after a b move (both processes indeed terminate after a b action). An a action instead leads to a state that can only perform a a action towards $\bot$. This is because $q_0 + r_0$ can reduce to q_1 with a a move, while $p_0 + r_0$ cannot match that action exactly: it can reduce to p_1 or r_1, that are not bisimilar to q_1.

$$\begin{tikzcd}
r_0 \arrow[d, "a"] & & p_0 + r_0 \arrow[dl, "a"'] \arrow[d, "b"] \arrow[dr, "a"] & & & & q_0 + r_0 \arrow[dl, "a"'] \arrow[d, "b"] \arrow[dr, "a"] & \\
r_1 & p_1 & p_2 & r_1 & q_3 & q_1 \arrow[l, "b"'] & q_2 & r_1
\end{tikzcd}$$

┘

Intuitively, the non-deterministic sum of two precesses can behave as the former process or as the latter (but not as both). Therefore we can expect that the distance between two sums is bounded by the join of the distances of the components. This is however not always the case, as the immediate distance is not necessarily non-expansive. The sum operator $+_{\mathbb{V}}$ from Definition 6, instead, turns out to be sufficient for our purposes. Proposition below indeed shows that $\hat{+}$ is non-extensive.

▶ **Proposition 23.** *For every p_1, p_2, s_1, s_2 : it holds that $\hat{+}(p_1, p_2, s_1, s_2) \leq_{\mathbb{V}} s_1 +_{\mathbb{V}} s_2$.*

5.4 Parallel Composition

We assume P to be closed under the binary operator $|$, whose semantics is defined below:

$$\frac{p \xrightarrow{\ell} p'}{p|q \xrightarrow{\ell} p'|q} \qquad \frac{p \xrightarrow{\ell} p' \quad q \xrightarrow{\ell} q'}{p|q \xrightarrow{\ell} p'|q'} \qquad \frac{q \xrightarrow{\ell} q'}{p|q \xrightarrow{\ell} p|q'}$$

The notion of synchronisation considered in this paper is the one pioneered in CSP [19, 35]. This choice is motivated by the fact that, in comparison with CCS-like communication [25] (which requires dual actions to synchronise resulting in an invisible τ-action), CSP notion does not change the label: this simplifies the technical development and enables stronger compositionality properties. Most of the works on compositionality of metrics for parallel composition we are aware of use CSP synchronisation, e.g. [2, 14, 15].

▶ **Example 24.** Let p_0 and q_0 be as in Example 18, and r_0 as in Example 22. We have that $\delta(p_0|r_0, q_0|r_0)$ is as the figure below. Indeed, $p_0|r_0$ and $q_0|r_0$ necessarily reduce to bisimilar states after a b action: therefore their distance b-reduces to $\bot$. The situation for a actions is more involved, due the the presence of several a-reducts for both processes. So, consider the transition $p_0|r_0 \xrightarrow{a} p_1|r_0$. We need to find the matching move of $q_0|r_0$ that minimises the distance between the reducts. So, consider the transition $q_0|r_0 \xrightarrow{a} q_1|r_1$. Since $p_1|r_0$ can only perform a actions while $q_1|r_1$ only b ones, we have that $\delta(p_1|r_0, q_1|r_1) = \top$. If we instead consider transition $q_0|r_0 \xrightarrow{a} q_1|r_0$, we have that $\delta(p_1|r_0, q_1|r_0) = s_1'$. Indeed, $q_1|r_0 \xrightarrow{b}$ while $p_1|r_0$ does not: hence $s_1' \not\xrightarrow{b}$. Moreover, $s_1' \xrightarrow{a} s_2'$. The only a-reducts of $p_1|r_0$ and $q_1|r_0$ are, respectively, $p_1|r_1$ and $q_1|r_1$. It is easy to verify that $\delta(p_1|r_1, q_1|r_1) = s_2'$. The last possible matching choice is $q_0|r_0 \xrightarrow{a} q_0|r_1$, for which we have that $\delta(p_1|r_0, q_0|r_1) = s_1'$: the argument is similar to the previous case. All the other starting a-moves of $p_0|r_0$, and those of $q_0|r_0$, have matching moves leading to distances greater or equal than s_1'.

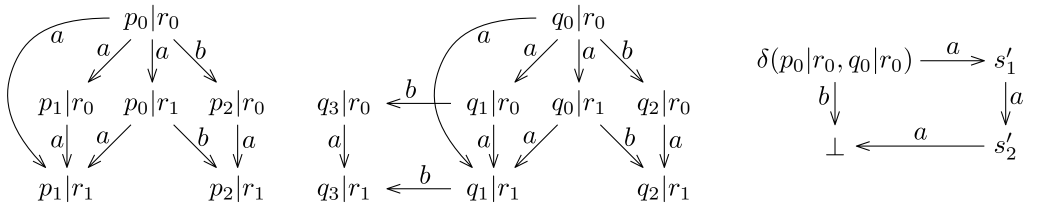

Parallel composition does not enjoy strong compositionality properties. Indeed in general $\hat{|}(p_1, p_2, s_1, s_2)$ is related neither to s_1 nor to s_2, and even $\hat{|}(p_1, p_2, s_1, \bot_{\mathbb{V}})$ is not related to s_1. Consider for instance the case where p_2 "consumes" a s_1 move.

However, our metric domain $\mathbb{V}$ contains "contextual" information. We exploit this fact to show that a nice compositionality property, similar to non-extensivity [15], holds when the context and the distance are "compatible". A formal definition of compatibility follows.

▶ **Definition 25.** *A relation $\mathcal{R} \subseteq S \times P$ is a compatibility relation if, whenever $s \;\mathcal{R}\; p$:*

1. $s \xrightarrow{\ell} s' \implies s' \;\mathcal{R}\; p$*;*
2. $s \xrightarrow{\ell} s'$ *and* $p \xrightarrow{\ell} p' \implies s' \;\mathcal{R}\; p'$ *and* $s \leq_{\mathbb{V}} s'$*.*

We say that s is p-compatible iff $s \;\mathcal{R}\; p$ for some compatibility relation $\mathcal{R}$.

▶ **Example 26.** Consider again p_0, s_0, s_1 from Example 18 and $\delta(p_0|r_0, q_0|r_0)$ of Example 24. We have that s_0 is not p_0-compatible as Condition 2 from Definition 25 is violated: $s_0 \xrightarrow{a} s_1$ and $p_0 \xrightarrow{a} p_1$ but $s_0 \not\leq_{\mathbb{V}_0} s_1$. Instead, s''_0 below is p_0-compatible: it follows from the facts that s''_0 necessarily reduces to a greater or equal state, p_0 reduces to terminated states, which are vacuously compatible with every distance. Note that $\delta(p_0, q_0) = s_0 \leq_{\mathbb{V}_0} s''_0$ and $\delta(p_0|r_0, q_0|r_0) \leq_{\mathbb{V}_0} s''_0$. The second inclusion follows from the first by Proposition 27.

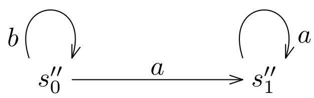

▶ **Proposition 27.** *If s_1 is p_2-compatible and s_2 is p_1-compatible, then $\hat{|}(p_1, p_2, s_1, s_2) \leq_{\mathbb{V}} s_1 +_{\mathbb{V}} s_2$.*

5.5 Replication

We assume that P is closed both under operator $|$ (as defined in Section 5.4) and under $! : P \to P$, whose semantics is standard. In general, replication has bad compositionality properties: since it allows infinite behaviour, even a small distance in the parameter can get amplified to a much larger value. However, we show that $\hat{!}$ is not expansive under the assumption that the parameter s always reduces to a larger or equal value and $+_{\mathbb{Q}}$ is idempotent. Such condition is of course quite strong, but it holds for instance when interpreting bisimilarity as a contextual bisimulation metric (see Section 4.1).

▶ **Example 28.** We have that $!p_0$ and $!q_0$ can both fire a a or b action and reduce to a process with the same behaviour (the simplest state with this property is drawn in the figure). Therefore, the distance $\delta(!p_0, !q_0) = \bot_{\mathbb{V}}$. In general, however, the distance among processes is not preserved by replication, as shown below:

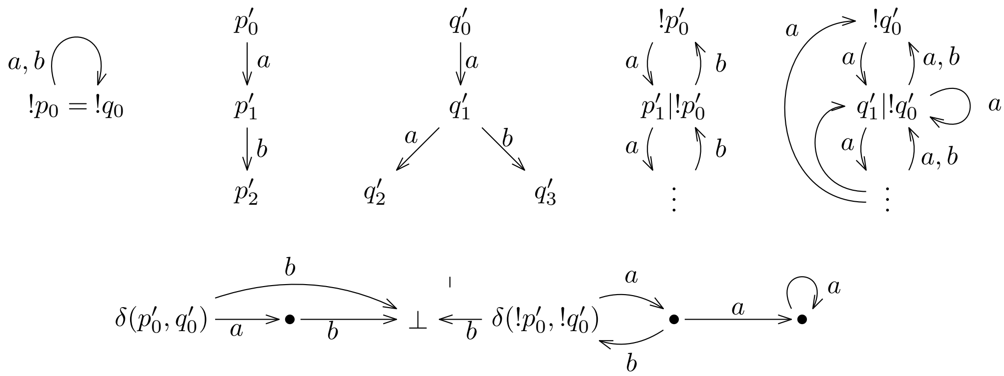

⌟

▶ **Definition 29.** *We define* **Inc***, the set of increasing states, as the largest set $S' \subseteq S$ such that, whenever $s \in S'$ and $s \xrightarrow{\ell} s' : s \leq_{\mathbb{V}} s'$ and $s' \in S'$.*

▶ **Proposition 30.** *If s is increasing and $+_{\mathbb{Q}}$ is idempotent, then $\hat{!}(p, s) \leq_{\mathbb{V}} s$.*

6 Related Work & Conclusion

Quite a few works in the literature study context dependent relations. The closest to our work is the already mentioned study about environment paremetrized bisimilarity [23]. Our definition of CBM is similar to theirs, where the main differences are that we also consider quantitative aspects and that we explicitly work with a metric. The same work also provides an interesting logical characterisation of their relation in terms of Hennessy-Milner logic, but does not study compositionality. Since environment parametrized bisimilarity can be embedded into our framework, our compositionality results also hold for [23]. A closely related line of research [3, 20, 21] (non-exhaustive list) studies conditional bisimulations in an abstract categorical framework, where conditions are used to make assumptions on the environment. In particular, [21] introduces a notion of conditional bisimilarity for reactive systems and shows that conditional bisimilarity is a congruence. In [18], an early and a late notion of symbolic bisimilarity for value passing processes are introduced, where actual values are symbolically represented with boolean expressions with free variables. Symbolic bisimilarities are parametric w.r.t. a predicate that, in a sense, allows to make assumptions on the values that the context can send. Our notion of contextuality instead restricts the choices of the environment, and we do not consider explicit value passing.

Compositionality of behavioural metrics has been studied in the probabilistic setting [4, 8]. In [2], it has been shown that parallel composition is non-extensive. We remark that our notion of parallel composition is slightly more general than the one considered in [2], as in there processes necessarily synchronise on common actions. The work [14] studies compositionality for quite a few process algebraic operators, showing e.g. that non-deterministic sum is non-expansive, while parallel composition is non-extensive. The bang operator is shown Lipschitz continuous for the discounted metric, while not even uniformly continuous w.r.t. the non-discounted one. [16] introduces structural operational semantics formats that guarantee compositionality of operators. Basically, compositionality depends on how many parameters of the operator are copied from the source to the destination of the rules, weighted by probabilities and the discount factor.

Concluding Remarks. This paper introduces a new form of metric on the states of a LTS, called contextual behavioural metric, which enables contextual and quantitative reasoning. We study compositional properties of CBMs w.r.t. some operators, showing that, under the assumption that the immediate metric is non-extensive, the following hold: restriction is non-expansive, non-deterministic sum is non-extensive, prefixing enjoys a property slightly weaker than uniform continuity, parallel composition is non-extensive when the distance between components is compatible with the context and replication enjoys non-expansiveness under some (rather strong) assumptions on the underling quantale $\mathbb{Q}$.

Due to the generality of CBMs, our compositionality results extend to behavioural metrics as defined in Section 4.2. For instance, since the compatibility relation of Definition 25 holds trivially for the MLTS of behavioural metrics, we have that compositionality of parallel composition only depends on the compositionality of the immediate metric.

Our work is still preliminary, and indeed we are yet in the quest for an appropriate general notion of compositionality: here we tried to adapt concepts from the probabilistic setting [14, 16], where uniform continuity is considered as the most general notion of compositionality. In our setting not even prefixing enjoys uniform continuity, which should not come as a surprise, as quantales are not totally ordered in general. Our compositionality results have heterogeneous side conditions. Spelling out all the compositionality results in a uniform way would come with a high price: operators for which compositionality holds without any side condition, such as restriction, would have to be treated as those for which compositionality holds only modulo appropriate (and strong) hypotheses, such as replication. An interesting future work would be to infer the side conditions directly from SOS rules, or studying more operators or rule formats as in [16].

Another direction of future research would be to consider calculi with value and/or channel passing like the π-calculus: since strong bisimilarity is not a congruence in such settings, a promising approach could be a "contextualisation" of open-bisimilarity [32].

References

1. S. Abramsky. The lazy λ-calculus. In D. Turner, editor, *Research Topics in Functional Programming*, pages 65–117. Addison Wesley, 1990.
2. Giorgio Bacci, Giovanni Bacci, Kim G. Larsen, and Radu Mardare. Computing behavioral distances, compositionally. In *Proc. of MFCS*, volume 8087 of *Lecture Notes in Computer Science*, pages 74–85. Springer, 2013.
3. Harsh Beohar, Barbara König, Sebastian Küpper, and Alexandra Silva. Conditional transition systems with upgrades. *Sci. Comput. Program.*, 186, 2020.
4. Konstantinos Chatzikokolakis, Daniel Gebler, Catuscia Palamidessi, and Lili Xu. Generalized bisimulation metrics. In *Proc. of CONCUR*, volume 8704 of *Lecture Notes in Computer Science*, pages 32–46. Springer, 2014.
5. Luca de Alfaro, Marco Faella, and Mariëlle Stoelinga. Linear and branching metrics for quantitative transition systems. In *Proc. of ICALP*, volume 3142 of *Lecture Notes in Computer Science*, pages 97–109. Springer, 2004.
6. Arthur Azevedo de Amorim, Marco Gaboardi, Justin Hsu, Shin-ya Katsumata, and Ikram Cherigui. A semantic account of metric preservation. In *Proc. of POPL*, pages 545–556. ACM, 2017.
7. Josée Desharnais, Vineet Gupta, Radha Jagadeesan, and Prakash Panangaden. Metrics for labeled markov systems. In *Proc. of CONCUR*, volume 1664 of *LNCS*, pages 258–273. Springer, 1999.
8. Josée Desharnais, Vineet Gupta, Radha Jagadeesan, and Prakash Panangaden. Metrics for labelled markov processes. *Theor. Comput. Sci.*, 318(3):323–354, 2004.
9. Josée Desharnais, Radha Jagadeesan, Vineet Gupta, and Prakash Panangaden. The metric analogue of weak bisimulation for probabilistic processes. In *Proc. of LICS*, pages 413–422. IEEE Computer Society, 2002.

10 Josée Desharnais, François Laviolette, and Mathieu Tracol. Approximate analysis of probabilistic processes: Logic, simulation and games. In *Proc. of QEST*, pages 264–273. IEEE Computer Society, 2008.

11 Wenjie Du, Yuxin Deng, and Daniel Gebler. Behavioural pseudometrics for nondeterministic probabilistic systems. In *Proc. of SETTA*, volume 9984 of *Lecture Notes in Computer Science*, pages 67–84, 2016.

12 Wenjie Du, Yuxin Deng, and Daniel Gebler. Behavioural pseudometrics for nondeterministic probabilistic systems. *Sci. Ann. Comput. Sci.*, 32(2):211–254, 2022.

13 Norm Ferns, Prakash Panangaden, and Doina Precup. Metrics for finite markov decision processes. In *Proc. of AAAI*, pages 950–951. AAAI Press / The MIT Press, 2004.

14 Daniel Gebler, Kim G. Larsen, and Simone Tini. Compositional bisimulation metric reasoning with probabilistic process calculi. *Log. Methods Comput. Sci.*, 12(4), 2016.

15 Daniel Gebler, Kim Guldstrand Larsen, and Simone Tini. Compositional metric reasoning with probabilistic process calculi. In *Proc. of FoSSaCS*, volume 9034 of *Lecture Notes in Computer Science*, pages 230–245. Springer, 2015.

16 Daniel Gebler and Simone Tini. SOS specifications for uniformly continuous operators. *J. Comput. Syst. Sci.*, 92:113–151, 2018.

17 Guillaume Geoffroy and Paolo Pistone. A partial metric semantics of higher-order types and approximate program transformations. In *Proc. of CSL*, volume 183 of *LIPIcs*, pages 23:1–23:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.

18 Matthew Hennessy and Huimin Lin. Symbolic bisimulations. *Theor. Comput. Sci.*, 138(2):353–389, 1995.

19 C. A. R. Hoare. Communicating sequential processes. *Commun. ACM*, 21(8):666–677, 1978.

20 Mathias Hülsbusch and Barbara König. Deriving bisimulation congruences for conditional reactive systems. In *Proc. of FoSSaCS*, volume 7213 of *Lecture Notes in Computer Science*, pages 361–375. Springer, 2012.

21 Mathias Hülsbusch, Barbara König, Sebastian Küpper, and Lars Stoltenow. Conditional bisimilarity for reactive systems. In *Proc. of FSCD*, volume 167 of *LIPIcs*, pages 10:1–10:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.

22 Ugo Dal Lago, Francesco Gavazzo, and Akira Yoshimizu. Differential logical relations, part I: the simply-typed case. In *Proc. of ICALP*, volume 132 of *LIPIcs*, pages 111:1–111:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

23 Kim Guldstrand Larsen. A context dependent equivalence between processes. *Theor. Comput. Sci.*, 49:184–215, 1987.

24 F. William Lawvere. Metric spaces, generalized logic, and closed categories. In *Rend. Sem. Mat. Fis. Milano*, volume 43, pages 135–166, 1973.

25 Robin Milner. *A Calculus of Communicating Systems*, volume 92 of *Lecture Notes in Computer Science*. Springer, 1980.

26 Robin Milner. *Communication and concurrency.* PHI Series in computer science. Prentice Hall, 1989.

27 David Park. Concurrency and automata on infinite sequences. In *Theoretical Computer Science*, pages 167–183, 1981.

28 Paolo Pistone. On generalized metric spaces for the simply typed lambda-calculus. In *Proc. of LICS*, pages 1–14. IEEE, 2021.

29 Damien Pous and Davide Sangiorgi. *Enhancements of the bisimulation proof method*, pages 233–289. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2011. doi:10.1017/CBO9780511792588.007.

30 Jason Reed and Benjamin C. Pierce. Distance makes the types grow stronger: a calculus for differential privacy. In *Proc. of ICFP*, pages 157–168. ACM, 2010.

31 Kimmo I. Rosenthal. *Quantales and their applications / Kimmo I. Rosenthal.* Pitman research notes in mathematics series ; 234. Longman Scientific & Technical, Essex, England, 1990.

32 Davide Sangiorgi. A theory of bisimulation for the pi-calculus. In *Proc. of CONCUR*, volume 715 of *Lecture Notes in Computer Science*, pages 127–142. Springer, 1993.

33 Franck van Breugel, Claudio Hermida, Michael Makkai, and James Worrell. An accessible approach to behavioural pseudometrics. In *Proc. of ICALP*, volume 3580 of *LNCS*, pages 1018–1030. Springer, 2005.

34 Franck van Breugel and James Worrell. Approximating and computing behavioural distances in probabilistic transition systems. *Theor. Comput. Sci.*, 360(1-3):373–385, 2006.

35 Rob J. van Glabbeek. Notes on the methodology of CCS and CSP. *Theor. Comput. Sci.*, 177(2):329–349, 1997.