

Quantitative Refinement for Weighted Modal Transition Systems

Claus Thrane, Sebastian S. Bauer, Uli Fahrenberg, Kim Guldstrand Larsen,
Line Juhl, Axel Legay

► **To cite this version:**

Claus Thrane, Sebastian S. Bauer, Uli Fahrenberg, Kim Guldstrand Larsen, Line Juhl, et al.. Quantitative Refinement for Weighted Modal Transition Systems. MFCS, Aug 2011, Warszawa, Poland. pp.60 - 71, 2011, <10.1007/978-3-642-22993-0_9>. <hal-01088046>

HAL Id: hal-01088046

<https://hal.inria.fr/hal-01088046>

Submitted on 27 Nov 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Quantitative Refinement for Weighted Modal Transition Systems

Sebastian S. Bauer¹, Uli Fahrenberg², Line Juhl³,
Kim G. Larsen³, Axel Legay², and Claus Thrane³

¹ Ludwig-Maximilians-Universität München, Germany

² Irisa/INRIA Rennes, France

³ Aalborg University, Denmark

Abstract. Specification theories as a tool in the development process of component-based software systems have recently attracted a considerable attention. Current specification theories are however qualitative in nature and hence fragile and unsuited for modern software systems. We propose the first specification theory which allows to capture quantitative aspects during the refinement and implementation process.

Keywords: reducing complexity of design, modal specifications, quantitative reasoning

1 Introduction

Rigorous design of modern computer systems faces the major challenge that the systems are too complex to reason about [16]. Hence it is necessary to reason at the level of specification rather than at the one of implementations. Such specifications, which act as finite and concise abstractions for possibly infinite sets of implementations, allow not only to decrease the complexity of the design, but also permit to reason on subsystems independently.

Any reasonable specification theory is equipped with a *satisfaction relation* to decide whether an implementation matches the requirements of a specification, and a *refinement relation* that allows to compare specifications (hence sets of implementations). Moreover, the theory needs a notion of *logical composition* which allows to infer larger specifications as logical combinations of smaller ones. Another important ingredient is a notion of *structural composition* that allows to build overall specifications from subspecifications, mimicking at the implementation level *e.g.* the interaction of components in a distributed system. A partial inverse of this operation is given by the notion of *quotient* which allows to synthesize a subspecification from an overall specification and an implementation which realizes a part of the overall specification.

Over the years, there have been a series of advances on specification theories [2, 14, 5, 13]. The predominant approaches are based on modal logics and process algebras but have the drawback that they cannot naturally embed both logical and structural composition within the same formalism. Moreover, such

formalisms do not permit to reason from specification to implementation through stepwise refinement.

In order to leverage those problems, the concept of modal transition systems was introduced [12]. In short, modal transition systems are labeled transition systems equipped with two types of transitions: *must* transitions which are mandatory for any implementation, and *may* transitions which are optional for implementations. It is well admitted that modal transition systems match all the requirements of a reasonable specification theory (see *e.g.* [15] for motivations). Also, practical experience shows that the formalism is expressive enough to handle complex industrial problems [6, 17].

In a series of recent work [3, 10], the modal transition system framework has been extended in order to reason on *quantitative* aspects, hence providing a new specification theory for more elaborated structures, with the objective to better meet practical needs. In this quantitative setting however, the standard Boolean satisfaction and refinement relations are too fragile. Indeed, either an implementation satisfies a specification or it does not. This means that minor and major modifications in the implementation cannot be distinguished, as both of them may reverse the Boolean answer. As observed by de Alvaro *et al.* for the logical framework of CTL [1], this view is obsolete; engineers need quantitative notions on how modified implementations differ.

The main contribution of this paper is to mitigate the above problem by lifting the satisfaction and refinement relations into the quantitative framework, hence completing the quantitative approach to reason on modal transition systems. More precisely, and similarly to what has been proposed in the logical framework, we introduce a notion of *distance* between both specifications and implementations, which permits quantitative comparison. Given two implementations that do not necessarily satisfy a specification, we can decide through quantitative reasoning which one is the better match for the specification's requirements.

To facilitate this reasoning, we develop a notion of *modal* distance between specifications, which approximates the distances between their implementations. This preserves the relation between modal refinement and satisfaction checking in the Boolean setting. We show that computing distances between implementation sets is EXPTIME-hard, whereas modal distances are computable in $NP \cap CO-NP$ (which is higher than for Boolean modal refinement). Akin to *discounted games* [19] we can reason on behaviors in a discounted manner, giving more importance to differences that happen in the near future, while accumulating the amount by which the specifications fail to be compatible at each step. As for the games, the semantics of the outcome is considered application specific.

Modifying the semantic outcome of satisfaction has strong impact on operations between specifications. As a second contribution of this paper, we propose quantitative versions of structural composition and quotient which inherit the good properties from the Boolean setting. We also propose a new notion of *relaxation*, which is inherent to the quantitative framework and allows *e.g.* to calibrate the quotient operator.

However, there is no free lunch, and working with distances has a price: some of the properties of logical conjunction and determinization are not preserved in the quantitative setting. More precisely, conjunction is not the greatest lower bound with respect to refinement distance as it is in the Boolean setting, and deterministic overapproximation is too coarse. In fact we show that this is a fundamental limitation of *any* reasonable quantitative specification formalism.

Structure of the paper. We start out by introducing our quantitative formalism which has weighted transition systems as implementations and weighted modal transition systems as specifications. In Section 3 we introduce the distances we use for quantitative comparison of both implementations and specifications. Section 4 is devoted to a formalization of the notion of relaxation which is of great use in quantitative design. In the next section we see some inherent limitations of the quantitative approach, and Section 6 finishes the paper by showing that structural composition works as expected in the quantitative framework and links relaxation to quotients.

Acknowledgment. The authors wish to thank Jiří Srba for fruitful discussions during the preparation of this work.

2 Weighted Modal Transition Systems

In this section we present the formalism we use for implementations and specifications. As implementations we choose the model of *weighted transition systems*, *i.e.* labeled transition systems with integer weights at transitions. Specifications both have a *modal* dimension, specifying discrete behavior which *must* be implemented and behavior which *may* be present in implementations, and a *quantitative* dimension, specifying intervals of weights on each transition which an implementation must choose from.

Let $\mathbb{I} = \{[x, y] \mid x \in \mathbb{Z} \cup \{-\infty\}, y \in \mathbb{Z} \cup \{\infty\}, x \leq y\}$ be the set of closed extended-integer intervals and let Σ be a finite set of actions. Our set of *specification labels* is $\mathbf{Spec} = (\Sigma \times \mathbb{I}) \cup \{\perp\}$, where the special symbol \perp models *inconsistency*. The set of *implementation labels* is defined as $\mathbf{Imp} = \Sigma \times \{[x, x] \mid x \in \mathbb{Z}\} \approx \Sigma \times \mathbb{Z}$. Hence a specification imposes labels and integer *intervals* which constrain the possible weights of an implementation.

We define a partial order on \mathbb{I} (representing inclusion of intervals) by $[x, y] \sqsubseteq [x', y']$ if $x' \leq x$ and $y \leq y'$, and we extend this order to specification labels by $(a, I) \sqsubseteq (a', I')$ if $a = a'$ and $I \sqsubseteq I'$, and $\perp \sqsubseteq (a, I)$ for all $(a, I) \in \mathbf{Spec}$. The partial order on \mathbf{Spec} is hence a *refinement* order; if $k_1 \sqsubseteq k_2$, then no more implementation labels are contained in k_1 than in k_2 .

Specifications and implementations are defined as follows:

Definition 1. A *weighted modal transition system* (WMTS) is a four-tuple $(S, s^0, \dashrightarrow, \longrightarrow)$ consisting of a set of states S with an initial state $s^0 \in S$ and *must* and *may* transition relations $\longrightarrow \subseteq \dashrightarrow \subseteq S \times \mathbf{Spec} \times S$. A WMTS is an *implementation* if $\longrightarrow = \dashrightarrow \subseteq S \times \mathbf{Imp} \times S$.

A WMTS is *finite* if S and \dashrightarrow (and hence also \rightarrow) are finite sets, and it is *deterministic* if it holds that for any $s \in S$ and $a \in \Sigma$, $(s, (a, I_1), t_1), (s, (a, I_2), t_2) \in \dashrightarrow$ imply $I_1 = I_2$ and $t_1 = t_2$. Hence a deterministic specification allows at most one transition under each discrete action from every state. In the rest of the paper we will write $s \xrightarrow{k} s'$ for $(s, k, s') \in \dashrightarrow$ and similarly for \rightarrow , and we will always write $S = (S, s^0, \dashrightarrow, \rightarrow)$ or $S_i = (S_i, s_i^0, \dashrightarrow_i, \rightarrow_i)$ for WMTS and $I = (I, i^0, \rightarrow)$ for implementations. Note that an implementation is just a usual integer-weighted transition system.

The implementation semantics of a specification is given through modal refinement, as follows: A *modal refinement* of WMTS S_1, S_2 is a relation $R \subseteq S_1 \times S_2$ such that for any $(s_1, s_2) \in R$ and any *may* transition $s_1 \xrightarrow{k_1} t_1$ in S_1 , there exists $s_2 \xrightarrow{k_2} t_2$ in S_2 for which $k_1 \sqsubseteq k_2$ and $(t_1, t_2) \in R$, and for any *must* transition $s_2 \xrightarrow{k_2} t_2$ in S_2 , there exists $s_1 \xrightarrow{k_1} t_1$ in S_1 for which $k_1 \sqsubseteq k_2$ and $(t_1, t_2) \in R$. Hence in such a modal refinement, behavior which is required in S_2 is also required in S_1 , no more behavior is allowed in S_1 than in S_2 , and the quantitative requirements in S_1 are refinements of the ones in S_2 . We write $S_1 \leq_m S_2$ if there is a modal refinement relation R for which $(s_1^0, s_2^0) \in R$. The implementation semantics of a specification can then be defined as the set of all implementations which are also refinements:

Definition 2. The *implementation semantics* of a WMTS S is the set $\llbracket S \rrbracket = \{I \mid I \leq_m S, I \text{ implementation}\}$.

We say that a WMTS S is *consistent* if it has an implementation, *i.e.* if $\llbracket S \rrbracket \neq \emptyset$. A useful over-approximation of consistency is *local consistency*: a WMTS S is said to be locally consistent if $s \xrightarrow{k} t$ implies $k \neq \perp$, *i.e.* if no \perp -labeled *must* transitions appear in S . Local consistency implies consistency, but the inverse is not true; *e.g.* the WMTS $s_0 \xrightarrow{a_2} s_1 \xrightarrow{a_1} s_2 \xrightarrow{\perp} s_3$ has an implementation $i_0 \xrightarrow{a_2} i_1$. Local inconsistencies may be removed recursively as follows:

Definition 3. For a WMTS S , let $\text{pre} : 2^S \rightarrow 2^S$ be given by $\text{pre}(B) = \{s \in S \mid s \xrightarrow{k} t \in B \text{ for some } k\}$, and let $S^\perp = \{s \in S \mid s \xrightarrow{\perp} t \text{ for some } t \in S\}$. If $s^0 \notin \text{pre}^*(S^\perp)$, then the *pruning* $\rho(S) = (S_\rho, s^0, \dashrightarrow_\rho, \rightarrow_\rho)$ is defined by $S_\rho = S \setminus \text{pre}^*(S^\perp)$, $\dashrightarrow_\rho = \dashrightarrow \cap (S_\rho \times (\text{Spec} \setminus \{\perp\}) \times S_\rho)$ and $\rightarrow_\rho = \rightarrow \cap (S_\rho \times (\text{Spec} \setminus \{\perp\}) \times S_\rho)$.

Note that if $\rho(S)$ exists, then it is locally consistent, and if $\rho(S)$ does not exist ($s^0 \in \text{pre}^*(S^\perp)$), then S is inconsistent. Also, $\rho(S) \leq_m S$ and $\llbracket \rho(S) \rrbracket = \llbracket S \rrbracket$.

3 Thorough and Modal Refinement Distances

For the quantitative specification formalism we have introduced in the last section, the standard Boolean notions of satisfaction and refinement are too fragile. To be able to reason not only whether a given quantitative implementation satisfies a given quantitative specification, but also *to what extent*, we introduce a notion of *distance* between both implementations and specifications.




Fig. 1. Two weighted transition systems with branching distance $d(I_1, I_2) = 18$.

We first define the distance between *implementations*; for this we introduce a distance on implementation labels by

$$d_{\text{imp}}((a_1, x_1), (a_2, x_2)) = \begin{cases} \infty & \text{if } a_1 \neq a_2, \\ |x_1 - x_2| & \text{if } a_1 = a_2. \end{cases} \quad (1)$$

In the rest of the paper, let $\lambda \in \mathbb{R}$ with $0 < \lambda < 1$ be a *discounting factor*.

Definition 4. Let I_1, I_2 be implementations (weighted transition systems). The *implementation distance* $d : I_1 \times I_2 \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ between the states of I_1 and I_2 is the least fixed point of the equations

$$d(i_1, i_2) = \max \begin{cases} \sup_{i_1 \xrightarrow{k_1} j_1} \inf_{i_2 \xrightarrow{k_2} j_2} d_{\text{imp}}(k_1, k_2) + \lambda d(j_1, j_2), \\ \sup_{i_2 \xrightarrow{k_2} j_2} \inf_{i_1 \xrightarrow{k_1} j_1} d_{\text{imp}}(k_1, k_2) + \lambda d(j_1, j_2). \end{cases}$$

We define $d(I_1, I_2) = d(i_1^0, i_2^0)$.

Except for the symmetrizing max operation, this is precisely the *accumulating branching distance* which is introduced in [18]; see also [8, 9] for a thorough introduction to linear and branching distances as we use them here. As the equations in the definition define a *contraction*, they have indeed a unique least fixed point; note that $d(i_1, i_2) = \infty$ is also a fixed point, *cf.* [11].

We remark that besides this accumulating distance, other interesting system distances may be defined depending on the application at hand, but we concentrate here on this distance and leave a generalization to other distances for future work.

Example 1. Consider the two implementations I_1 and I_2 in Figure 1 with a single action (elided for simplicity) and with discounting factor $\lambda = .9$. The equations in the illustration have already been simplified by removing all expressions that evaluate to ∞ . What remains to be done is to compute the least fixed point of the equation $d(k_1, i_2) = \max\{2 + .9 d(k_1, i_2), 0\}$ which is $d(k_1, i_2) = 20$. Hence $d(i_1, i_2) = \max\{3, .9 \cdot 20\} = 18$.

To lift implementation distance to specifications, we need first to consider the distance between *sets* of implementations. Given implementation sets $\mathcal{I}_1, \mathcal{I}_2$, we define

$$d(\mathcal{I}_1, \mathcal{I}_2) = \sup_{I_1 \in \mathcal{I}_1} \inf_{I_2 \in \mathcal{I}_2} d(I_1, I_2)$$

Note that in case \mathcal{I}_2 is finite, we have that for all $\varepsilon \geq 0$, $d(\mathcal{I}_1, \mathcal{I}_2) \leq \varepsilon$ if and only if for each implementation $I_1 \in \mathcal{I}_1$ there exists $I_2 \in \mathcal{I}_2$ for which $d(I_1, I_2) \leq \varepsilon$, hence this is a natural notion of distance. Especially, $d(\mathcal{I}_1, \mathcal{I}_2) = 0$ if and only if \mathcal{I}_1 is a subset of \mathcal{I}_2 up to bisimilarity. For infinite \mathcal{I}_2 , we have the slightly more complicated property that $d(\mathcal{I}_1, \mathcal{I}_2) \leq \varepsilon$ if and only if for all $\delta > 0$ and any $I_1 \in \mathcal{I}_1$, there is $I_2 \in \mathcal{I}_2$ for which $d(I_1, I_2) \leq \varepsilon + \delta$.

Note that in general, our distance on sets of implementations is *asymmetric*; we may well have $d(\mathcal{I}_1, \mathcal{I}_2) \neq d(\mathcal{I}_2, \mathcal{I}_1)$. We lift this distance to specifications as follows:

Definition 5. The *thorough refinement distance* between WMTS S_1 and S_2 is defined as $d_t(S_1, S_2) = d(\llbracket S_1 \rrbracket, \llbracket S_2 \rrbracket)$. We write $S_1 \leq_t^\varepsilon S_2$ if $d_t(S_1, S_2) \leq \varepsilon$.

Indeed this permits us to measure incompatibility of specifications; intuitively, if two specifications have thorough distance ε , then any implementation of the first specification can be matched by an implementation of the second up to ε . Also observe the special case where $S_1 = I_1$ is an implementation: then $d_t(I_1, S_2) = \inf_{I_2 \in \llbracket S_2 \rrbracket} d(I_1, I_2)$, which measures how close I_1 is to satisfy the specification S_2 .

To facilitate computation and comparison of refinement distance, we introduce modal refinement distance as an overapproximation. We will show in Theorem 2 below that similarly to the Boolean setting [4], computation of thorough refinement distance is EXPTIME-hard, whereas modal refinement distance is computable in $\text{NP} \cap \text{CO-NP}$. First we generalize the distance on implementation labels from Equation (1) to specification labels so that for $k, \ell \in \text{Spec}$ we define

$$d_{\text{Spec}}(k, \ell) = \sup_{k' \sqsubseteq k, k' \in \text{Imp}} \inf_{\ell' \sqsubseteq \ell, \ell' \in \text{Imp}} d_{\text{Imp}}(k', \ell').$$

Note that d_{Spec} is asymmetric, and that $d_{\text{Spec}}(k, \ell) = 0$ if and only if $k \sqsubseteq \ell$. Also, $d_{\text{Spec}}(k, \ell) = d_{\text{Imp}}(k, \ell)$ for all $k, \ell \in \text{Imp}$. Using the $\dot{-}$ operation defined on integers by $x_1 \dot{-} x_2 = \max(x_1 - x_2, 0)$, we can express d_{Spec} as follows:

$$\begin{aligned} d_{\text{Spec}}((a_1, I_1), (a_2, I_2)) &= \infty \quad \text{if } a_1 \neq a_2 \\ d_{\text{Spec}}((a, [x_1, y_1]), (a, [x_2, y_2])) &= \max(x_2 \dot{-} x_1, y_1 \dot{-} y_2) \\ d_{\text{Spec}}(\perp, (a, I_2)) &= 0 \qquad \qquad \qquad d_{\text{Spec}}((a, I_1), \perp) = \infty \end{aligned}$$

Definition 6. Let S_1, S_2 be WMTS. The *modal refinement distance* $d_m : S_1 \times S_2 \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ from states of S_1 to states of S_2 is the least fixed point of the equations

$$d_m(s_1, s_2) = \max \begin{cases} \sup_{s_1 \xrightarrow{k_1} t_1} \inf_{s_2 \xrightarrow{k_2} t_2} d_{\text{Spec}}(k_1, k_2) + \lambda d_m(t_1, t_2), \\ \sup_{s_2 \xrightarrow{k_2} t_2} \inf_{s_1 \xrightarrow{k_1} t_1} d_{\text{Spec}}(k_1, k_2) + \lambda d_m(t_1, t_2). \end{cases}$$

We define $d_m(S_1, S_2) = d_m(s_1^0, s_2^0)$, and we write $S_1 \leq_m^\varepsilon S_2$ if $d_m(S_1, S_2) \leq \varepsilon$.




Fig. 2. Incompleteness of modal refinement distance.

The argument for existence and uniqueness of the least fixed point is exactly the same as for implementation distance in Definition 4. Like thorough refinement distance, modal refinement distance may be asymmetric.

The next theorem shows that modal refinement distance indeed overapproximates thorough refinement distance, and that it is exact for deterministic WMTS. Note that nothing general can be said about the precision of the overapproximation in the nondeterministic case; as an example observe the two specifications in Figure 2 for which $d_t(S_1, S_2) = 0$ but $d_m(S_1, S_2) = \infty$.

The fact that modal refinement only equals thorough refinement for deterministic specifications is well-known from the theory of modal transition systems [12], and the special case of S_1 locally consistent and S_2 deterministic is important, as it can be argued [12] that indeed, deterministic specifications are sufficient for applications.

Theorem 1. *For WMTS S_1, S_2 we have $d_t(S_1, S_2) \leq d_m(S_1, S_2)$. If S_1 is locally consistent and S_2 is deterministic, then $d_t(S_1, S_2) = d_m(S_1, S_2)$.*

The complexity results in the next theorem show that modal refinement distance can serve as a useful approximation of thorough refinement distance.

Theorem 2. *For finite WMTS S_1, S_2 and $\varepsilon \geq 0$, it is EXPTIME-hard to decide whether $S_1 \leq_t^\varepsilon S_2$. The problem whether $S_1 \leq_m^\varepsilon S_2$ is decidable in $\text{NP} \cap \text{co-NP}$.*

4 Relaxation

We introduce here a notion of *relaxation* which is specific to the quantitative setting. Intuitively, relaxing a specification means to weaken the quantitative constraints, while the discrete demands on which transitions may or must be present in implementations are kept. A similar notion of strengthening may be defined, but we do not use this here.

Definition 7. For WMTS S, S' and $\varepsilon \geq 0$, S' is an ε -relaxation of S if $S \leq_m S'$ and $S' \leq_m^\varepsilon S$.

Hence the quantitative constraints in S' may be more permissive than the ones in S , but no new discrete behavior may be introduced. Also note that any implementation of S is also an implementation of S' , and no implementation of S' is further than ε away from an implementation of S . The following proposition relates specifications to relaxed specifications:

Proposition 1. *If S'_1 and S'_2 are ε -relaxations of S_1 and S_2 , respectively, then $d_m(S_1, S_2) - \varepsilon \leq d_m(S_1, S'_2) \leq d_m(S_1, S_2)$ and $d_m(S_1, S_2) \leq d_m(S'_1, S_2) \leq d_m(S_1, S_2) + \varepsilon$.*

On the syntactic level, we can introduce the following *widening* operator which relaxes all quantitative constraints in a systematic manner. We write $I \pm \delta = [x - \delta, y + \delta]$ for an interval $I = [x, y]$ and $\delta \in \mathbb{N}$.

Definition 8. Given $\delta \in \mathbb{N}$, the δ -widening of a WMTS S is the WMTS $S^{+\delta}$ with transitions $s \xrightarrow{a, I \pm \delta} t$ in $S^{+\delta}$ for all $s \xrightarrow{a, I} t$ in S , and $s \xrightarrow{a, I \pm \delta} t$ in $S^{+\delta}$ for all $s \xrightarrow{a, I} t$ in S .

Widening and relaxation are related as follows; note also that as widening is a global operation whereas relaxation may be achieved entirely locally, not all relaxations may be obtained as widenings.

Proposition 2. *The δ -widening of any WMTS S is a $(1 - \lambda)^{-1}\delta$ -relaxation.*

There is also an implementation-level notion which corresponds to relaxation:

Definition 9. The ε -extended implementation semantics, for $\varepsilon \geq 0$, of a WMTS S is $\llbracket S \rrbracket^{+\varepsilon} = \{I \mid I \leq_m^\varepsilon S, I \text{ implementation}\}$.

Proposition 3. *If S' is an ε -relaxation of S , then $\llbracket S' \rrbracket \subseteq \llbracket S \rrbracket^{+\varepsilon}$.*

It can be shown that there are WMTS S, S' such that S' is an ε -relaxation of S but the inclusion $\llbracket S' \rrbracket \subseteq \llbracket S \rrbracket^{+\varepsilon}$ is strict.

5 Limitations of the Quantitative Approach

In this section we turn our attention towards some of the standard operators for specification theories; determinization and logical conjunction. Quite surprisingly, we show that in the quantitative setting, there are problems with these notions which do not appear in the Boolean theory. More specifically, we show that there is no determinization operator which always yields a *smallest* deterministic overapproximation, and there is no conjunction operator which acts as a greatest lower bound.

Theorem 3. *There is no unary operator \mathcal{D} on WMTS for which it holds that*

- (3.1) $\mathcal{D}(S)$ is deterministic for any WMTS S ,
- (3.2) $S \leq_m \mathcal{D}(S)$ for any WMTS S ,
- (3.3) $S \leq_m^\varepsilon D$ implies $\mathcal{D}(S) \leq_m^\varepsilon D$ for any WMTS S , any deterministic WMTS D , and any $\varepsilon \geq 0$.

In the standard Boolean setting, there is indeed a determinization operator which satisfies properties similar to the above, and which is useful because it enables checking thorough refinement, *cf.* Theorem 1. Likewise, the greatest-lower-bound property of logical conjunction in the Boolean setting ensures that the set of implementations of a conjunction of specifications is precisely the intersection of the implementation sets of the two specifications.

Theorem 4. *There is no partial binary operator \wedge on WMTS for which it holds that*

- (4.1) $S_1 \wedge S_2 \leq_m S_1$ and $S_1 \wedge S_2 \leq_m S_2$ for all locally consistent WMTS S_1, S_2 for which $S_1 \wedge S_2$ is defined,
- (4.2) for any locally consistent WMTS S and all deterministic and locally consistent WMTS S_1, S_2 such that $S \leq_m S_1$ and $S \leq_m S_2$, $S_1 \wedge S_2$ is defined and $S \leq_m S_1 \wedge S_2$,
- (4.3) for any $\varepsilon \geq 0$, there exist $\varepsilon_1 \geq 0$ and $\varepsilon_2 \geq 0$ such that for any locally consistent WMTS S and all deterministic and locally consistent WMTS S_1, S_2 for which $S_1 \wedge S_2$ is defined, $S \leq_m^{\varepsilon_1} S_1$ and $S \leq_m^{\varepsilon_2} S_2$ imply $S \leq_m^\varepsilon S_1 \wedge S_2$.

The counterexamples used in the proofs of Theorems 3 and 4 are quite general and apply to a large class of distances, rather than only to the accumulating distance discussed in this paper. Hence it can be argued that what we have exposed here is a fundamental limitation of any quantitative approach to modal specifications.

6 Structural Composition and Quotient

In this section we show that in our quantitative setting, notions of structural composition and quotient can be defined which obey the properties expected of such operations. In particular, structural composition satisfies independent implementability [2], hence the refinement distance between structural composites can be bounded by the distances between their respective components.

First we define partial synchronization operators \oplus and \ominus on specification labels which will be used for synchronizing transitions. We let $(a_1, I_1) \oplus (a_2, I_2)$ and $(a_1, I_1) \ominus (a_2, I_2)$ be undefined if $a_1 \neq a_2$, and otherwise

$$\begin{aligned} (a, [x_1, y_1]) \oplus (a, [x_2, y_2]) &= (a, [x_1 + x_2, y_1 + y_2]), \\ (a, I_1) \oplus \perp &= \perp \oplus (a, I_2) = \perp; \\ (a, [x_1, y_1]) \ominus (a, [x_2, y_2]) &= \begin{cases} \perp & \text{if } x_1 - x_2 > y_1 - y_2, \\ (a, [x_1 - x_2, y_1 - y_2]) & \text{if } x_1 - x_2 \leq y_1 - y_2, \end{cases} \\ (a, I_1) \ominus \perp &= \perp \ominus (a, I_2) = \perp. \end{aligned}$$

Note that we use CSP-style synchronization, but other types of synchronization can easily be defined. Also, defining \oplus to add intervals (and \ominus to subtract them) is only one particular choice; depending on the application, one can also *e.g.* let \oplus be intersection of intervals or some other operation. It is not difficult to see that these alternative synchronization operators would lead to properties similar to those we show here.

Definition 10. Let S_1 and S_2 be WMTS. The *structural composition* of S_1 and S_2 is $S_1 \parallel S_2 = (S_1 \times S_2, (s_1^0, s_2^0), \text{Spec}, \dashrightarrow, \longrightarrow)$ with transitions given as follows:

$$\frac{s_1 \xrightarrow{k_1} t_1 \quad s_2 \xrightarrow{k_2} t_2 \quad k_1 \oplus k_2 \text{ defined}}{(s_1, s_2) \xrightarrow{k_1 \oplus k_2} (t_1, t_2)} \quad \frac{s_1 \xrightarrow{k_1} t_1 \quad s_2 \xrightarrow{k_2} t_2 \quad k_1 \oplus k_2 \text{ defined}}{(s_1, s_2) \xrightarrow{k_1 \oplus k_2} (t_1, t_2)}$$




Fig. 3. WMTS for which $d_m(S_3, S_1 \parallel S_2) \neq d_m(S_2 \parallel S_3, S_1)$.

The *quotient* of S_1 by S_2 is $S_1 \parallel S_2 = \rho(S_1 \times S_2 \cup \{u\}, (s_1^0, s_2^0), \text{Spec}, \dashrightarrow, \longrightarrow)$ with transitions given as follows:

$$\begin{array}{c}
 \frac{s_1 \xrightarrow{k_1} t_1 \quad s_2 \xrightarrow{k_2} t_2 \quad k_1 \ominus k_2 \text{ defined}}{(s_1, s_2) \xrightarrow{k_1 \ominus k_2} (t_1, t_2)} \quad \frac{s_1 \xrightarrow{k_1} t_1 \quad s_2 \xrightarrow{k_2} t_2 \quad k_1 \oplus k_2 \text{ defined}}{(s_1, s_2) \xrightarrow{k_1 \oplus k_2} (t_1, t_2)} \\
 \frac{s_1 \xrightarrow{k_1} t_1 \quad \forall s_2 \xrightarrow{k_2} t_2 : k_1 \ominus k_2 \text{ undefined}}{(s_1, s_2) \xrightarrow{\perp} (s_1, s_2)} \\
 \frac{k \in \text{Spec} \quad \forall s_2 \xrightarrow{k_2} t_2 : k \oplus k_2 \text{ undefined}}{(s_1, s_2) \xrightarrow{-k} u} \quad \frac{k \in \text{Spec}}{u \xrightarrow{-k} u}
 \end{array}$$

Note that we ensure that the quotient $S_1 \parallel S_2$ is locally consistent by recursively removing \perp -labeled *must* transitions using pruning, see Definition 3. The following theorem shows that structural composition is well-behaved with respect to modal refinement distance in the sense that the distance between the composed systems is bounded by the distances of the individual systems. Note also the special case in the theorem of $S_1 \leq_m S_2$ and $S_3 \leq_m S_4$ implying $S_1 \parallel S_3 \leq_m S_2 \parallel S_4$.

Theorem 5 (Independent implementability). *For WMTS S_1, S_2, S_3, S_4 we have $d_m(S_1 \parallel S_3, S_2 \parallel S_4) \leq d_m(S_1, S_2) + d_m(S_3, S_4)$.*

The following theorem expresses the fact that quotient is a partial inverse to structural composition. Intuitively, the theorem shows that the quotient $S_1 \parallel S_2$ is maximal among all WMTS S_3 with respect to any distance $S_2 \parallel S_3 \leq_m^\varepsilon S_1$; note the special case of $S_3 \leq_m S_1 \parallel S_2$ if and only if $S_2 \parallel S_3 \leq_m S_1$.

Theorem 6 (Soundness and maximality of quotient). *Let S_1, S_2 and S_3 be locally consistent WMTS such that S_2 is deterministic and $S_1 \parallel S_2$ is defined. If $d_m(S_3, S_1 \parallel S_2) < \infty$, then $d_m(S_3, S_1 \parallel S_2) = d_m(S_2 \parallel S_3, S_1)$.*

The example depicted in Figure 3 shows that the condition $d_m(S_3, S_1 \parallel S_2) < \infty$ in Theorem 6 is necessary. Here $d_m(S_2 \parallel S_3, S_1) = 1$, but $d_m(S_3, S_1 \parallel S_2) = \infty$ because of inconsistency between the transitions $s_1 \xrightarrow{a,[0,0]} t_1$ and $s_2 \xrightarrow{a,[0,1]} t_2$ for which $k_1 \ominus k_2$ is defined.

As a practical application, we notice that *relaxation* as defined in Section 4 can be useful when computing quotients. The quotient construction in Definition 10 introduces local inconsistencies (which afterwards are pruned) whenever

there is a pair of transitions $s_1 \xrightarrow{k_1} t_1$, $s_2 \xrightarrow{k_2} t_2$ (or $s_1 \xrightarrow{k_1} t_1$, $s_2 \xrightarrow{k_2} t_2$) for which $k_1 \ominus k_2 = \perp$. Looking at the definition of \ominus , we see that this is the case if $k_1 = (a, [x_1, y_1])$ and $k_2 = (a, [x_2, y_2])$ are such that $x_1 - x_2 > y_1 - y_2$; hence these local inconsistencies can be avoided by *enlarging* k_1 .

Enlarging quantitative constraints is exactly the intuition of relaxation, thus in practical cases where we get a quotient $S_1 \parallel S_2$ which is “too inconsistent”, we may be able to solve this problem by constructing a suitable ε -relaxation S'_1 of S_1 . Theorems 5 and 6 can then be used to ensure that also $S'_1 \parallel S_2$ is a relaxation of $S_1 \parallel S_2$.

7 Conclusion and Further Work

We have shown in this paper that within the quantitative specification framework of weighted modal transition systems, refinement and implementation distances provide a useful tool for robust compositional reasoning. Note that these distances permit us not only to reason about differences between implementations and from implementations to specifications, but they also provide a means by which we can compare specifications directly at the abstract level.

We have shown that for some of the ingredients of our specification theory, namely structural composition and quotient, our formalism is a conservative extension of the standard Boolean notions. We have also noted however, that for determinization and logical conjunction, the properties of the Boolean notions are not preserved, and that this is a fundamental limitation of any reasonable quantitative specification theory. The precise practical implications of this for the applicability of our quantitative specification framework are subject to future work.

References

1. Luca de Alfaro, Marco Faella, and Mariëlle Stoelinga. Linear and branching system metrics. *IEEE Trans. Software Eng.*, 35(2):258–273, 2009.
2. Luca de Alfaro and Thomas Henzinger. Interface-based design. In Manfred Broy, Johannes Grünbauer, David Harel, and Tony Hoare, editors, *Engineering Theories of Software Intensive Systems*, volume 195 of *NATO Science Series II: Mathematics, Physics and Chemistry*, pages 83–104. Springer-Verlag, 2005.
3. Sebastian S. Bauer, Line Juhl, Kim G. Larsen, Axel Legay, and Jiří Srba. Extending modal transition systems with structured labels. 2011. Submitted.
4. Nikola Beneš, Jan Křetínský, Kim G. Larsen, and Jiří Srba. Checking thorough refinement on modal transition systems is EXPTIME-complete. In Martin Leucker and Carroll Morgan, editors, *ICTAC*, volume 5684 of *Lecture Notes in Computer Science*, pages 112–126. Springer-Verlag, 2009.
5. Arindam Chakrabarti, Luca de Alfaro, Thomas A. Henzinger, and Freddy Y. C. Mang. Synchronous and bidirectional component interfaces. In *Proc. 14th Int. Conference on Computer Aided Verification (CAV)*, volume 2404 of *Lecture Notes in Computer Science*, pages 414–427, 2002.

6. STREP COMBEST (COMponent-Based Embedded Systems design Techniques). <http://www.combest.eu/home/>.
7. Anne Condon. The complexity of stochastic games. *Information and Computation*, 96(2):203–224, 1992.
8. Uli Fahrenberg, Kim G. Larsen, and Claus Thrane. A quantitative characterization of weighted Kripke structures in temporal logic. *Computing and Informatics*, 29(6+):1311–1324, 2010.
9. Uli Fahrenberg, Claus Thrane, and Kim G. Larsen. Distances for weighted transition systems: Games and properties. In *Proc. QAPL'11*, Electronic Proceedings in Theoretical Computer Science, 2011. To be published.
10. Line Juhl, Kim G. Larsen, and Jiří Srba. Modal transition systems with weight intervals. *Journal of Logic and Algebraic Programming*, 2011. To be published.
11. Kim G. Larsen, Uli Fahrenberg, and Claus Thrane. Metrics for weighted transition systems: Axiomatization and complexity. *Theoretical Computer Science*, 2011. 10.1016/j.tcs.2011.04.003.
12. Kim Guldstrand Larsen. Modal specifications. In *Automatic Verification Methods for Finite State Systems*, volume 407 of *Lecture Notes in Computer Science*, pages 232–246. Springer, 1989.
13. Nancy Lynch and Mark R. Tuttle. An introduction to input/output automata. *CWI-Quarterly*, 2(3), 1989.
14. Ulrik Nyman. *Modal Transition Systems as the Basis for Interface Theories and Product Lines*. PhD thesis, Aalborg University, Department of Computer Science, September 2008.
15. Jean-Baptiste Raclet. Residual for component specifications. *Electr. Notes in Theor. Comput. Sci.*, 215:93–110, 2008.
16. Joseph Sifakis. A vision for computer science – the system perspective. *Central European Journal of Computer Science*, 1(1):108–116, 2011.
17. SPEEDS (SPEculative and Exploratory Design in Systems Engineering). <http://www.speeds.eu.com/>.
18. Claus Thrane, Uli Fahrenberg, and Kim G. Larsen. Quantitative simulations of weighted transition systems. *Journal of Logic and Algebraic Programming*, 79(7):689–703, 2010.
19. Uri Zwick and Mike Paterson. The complexity of mean payoff games. In Ding-Zhu Du and Ming Li, editors, *COCOON*, volume 959 of *Lecture Notes in Computer Science*, pages 1–10. Springer-Verlag, 1995.