Quantitative Analysis and Model Checking

Michael Huth
Department of Computing and Information Sciences
Kansas State University, Manhattan, KS66506, USA
huth@cis.ksu.edu

Marta Kwiatkowska
School of Computer Science
University of Birmingham, Edgbaston, B15 2TT, United Kingdom
mzk@cs.bham.ac.uk

Abstract

Many notions of models in computer science provide quantitative information, or uncertainties, which necessitate a quantitative model checking paradigm. We present such a framework for reactive and generative systems based on a non-standard interpretation of the modal mu-calculus, where $\mu x.\phi$/$\nu x.\phi$ are interpreted as least/greatest fixed points over the infinite lattice of maps from states to the unit interval. By letting formulas denote lower bounds of probabilistic evidence of properties, the values computed by our quantitative model checker can serve as satisfactory correctness guarantees in cases where conventional qualitative model checking fails. Since fixed point iteration in this infinite domain is computationally unfeasible, we establish that the computation of fixed points may be restated as a conventional, and on average efficient, optimization problem in linear programming; this holds for a fragment of the modal mu-calculus which subsumes `CTL`*. Our semantics induces a state equivalence which is strictly in between probabilistic bisimulation and probabilistic ready bisimulation.*

1. Introduction

1.1. Motivation

Model checking [7, 17] has been established as an extremely successful mechanical tool of automated verification, especially in hardware design. Model checking as such is *qualitative* as the answers to checks of satisfiability are either "yes" or "no". Our work is motivated by the observation that there is a wide range of models and applications in computer science requiring a *quantitative* answer to such checks. The interaction geometry of modern distributed systems and network protocols calls for probabilistic, or more generally, quantitative estimates of e.g. performance and cost measures. Increasingly, randomized algorithms are utilized to achieve high performance at a cost of obtaining correct answers only with some high probability. In cases such as this conventional model checking, based on e.g. `CTL`, fails to deliver, as many relevant properties are simply not true in the model, and one is forced to consider suitable extensions of logics with quantitative information, in particular probabilistic logics.

The prevailing work on probabilistic temporal logics e.g. [11, 21] (and also modal logics e.g. `PML` [16]) uses *probability thresholds for truth* in the syntax of an underlying language of logic such as `CTL`, in the sense that the probabilistic formula $(\phi)_\rho$ is satisfied if the `CTL` formula ϕ holds with probability *at least* ρ for some real number ρ in the unit interval. While this permits qualitative checks of properties such as 'the probability of ϕ being satisfied within t steps is at least ρ' through a mix of conventional model checking algorithms and linear algebra techniques (see e.g. [11, 10]), it inevitably leads to potentially severe problems due to the continuous nature of the probabilities. Quite apart from the inherent difficulties of 'guessing' which threshold to use in a given system, the decision as to the formula being satisfied or not is reduced to deciding whether the probability of the formula being true in a state is above the threshold ρ. Although the inequality $p < r$ is decidable for rational numbers, and this has been advocated in [2], this approach does not lead to efficient automated procedures. Furthermore, thresholds do not permit continuous dependence of the outcome of checks on the probability distribution and the handling of infinite state systems, and may lead to discontinuity and even inconsistencies when adding fixed points to the logic.

While *state-based* quantitative semantics is a firmly established paradigm for Markovian models, with the latter often used as models for probabilistic extensions of temporal logics, *action-based* quantitative semantics has not yet been explored from a foundational point of view. Such action-based semantics is necessary for probabilistic and stochastic extensions of process algebras (see e.g. [9]), of which we mention PEPA [12], a CCS-like calculus whose transitions are enriched with rates of activity interpreted as parameters of exponential distribution. In models such as PEPA both the state- and action-based semantics arise: the former is used in the derivation of transient or steady state probability distributions based on the underlying continuous time Markov chain process, while the latter lends itself to a range of performance analyses, e.g. the throughput of a process or the utilization of resources, which are obtained by forming the scalar product of the steady state probability distribution with a vector of *rewards*.

This paper presents a conceptual and mathematical foundation for a viable alternative to the model checking based on the threshold approach, which avoids the difficulties mentioned previously by allowing the outcome of model 'checks' to range over the full unit interval, instead of just the truth values. Due to space restrictions, we confine this study to *reactive* probabilistic transition systems; in such systems, as a result of the interaction of action-based and probabilistic choice, the best probabilistic estimates that we can calculate are greatest lower bounds. However, this work does adapt to the probabilistically simpler *generative* settings and to models enriched with probabilistic information given in terms of rates of activity as in PEPA.

The natural choice for an action-based specification logic is Kozen's modal mu-calculus $\mu\mathtt{M}$ [15]. Since a host of temporal logics, such as CTL, have efficient translations into $\mu\mathtt{M}$, our results will apply to those logics as well.

1.2. Main Contribution

Our main technical contribution is threefold. First, we take a probabilistic transition system and let the formulas denote the (bounds on the) *probabilistic evidence* of properties through giving the modal mu-calculus ($\mu\mathtt{M}$) a non-standard interpretation over the infinite lattice of maps from states to the unit interval. This semantics is achieved by defining a monotone Σ-algebra and interpreting $\mu x.\phi$/$\nu x.\phi$ as least/greatest fixed points of monotone maps over that domain. By letting the formulas denote all intermediate reals in the unit interval, as well as 0 and 1, the values computed by our quantitative model checker can serve as satisfactory correctness guarantees (i.e. the property is shown to be satisfied with some high probability) in cases where conventional qualitative model checking fails. In addition, the framework overcomes the difficulties of finding the correct threshold, allows continuous dependence on distributions and handles infinite-state systems, while at the same time avoiding the discontinuity problems and inconsistencies.

Second, for a sublanguage of the modal mu-calculus, which subsumes the translation of CTL into $\mu\mathtt{M}$, we reduce the generally infinite fixed point iterations to an, on average efficient, standard optimization problem in linear programming. This is the basis for a global model checking algorithm for reactive and generative systems.

Third, we realize the induced notion of process equivalence for reactive systems as being strictly contained in the union of all probabilistic bisimulations and containing probabilistic ready bisimulation, thus allowing state reduction techniques based on probabilistic bisimulation.

1.3. Technical Overview

Probabilistic transition systems are labeled transition systems (Q, Δ) such that all action types a and states q with $q \ \to^a$ have a probability distribution on the possible a-moves at state q [16]. The semantic type for qualitative model checking, which makes no use of those distributions, is $\mathbf{2}^Q$, where $0 \in \mathbf{2}$ stands for "no" and $1 \in \mathbf{2}$ for "yes". The domain for *quantitative* model checking we consider here is $[0,1]^Q$. Our framework should allow for adaptations to other quantitative domains such as the reals.

The choice of Kozen's modal mu-calculus $\mu\mathtt{M}$ [15] as a specification logic allows us to show results in a fairly uniform way for a host of temporal logics, such as CTL, which have efficient translations into $\mu\mathtt{M}$. The meaning of constants can now be any function of type $[0,1]^Q$, but `true` and `false` are modeled via the constant $\mathbf{1}$, respectively, $\mathbf{0}$ function. The meaning of the modalities $\langle a\rangle\phi$ at state q is defined as a *weighted average* of the meaning of ϕ at the states reached via a-moves. Conjunction and disjunction get a *fuzzy* interpretation as

$$[\![\phi \wedge \psi]\!] = \max(\mathbf{0}, [\![\phi]\!] + [\![\psi]\!] - \mathbf{1}), \tag{1}$$

making $[0,1]^Q$ into a modal algebra. One may think of these values as greatest lower bounds for actual probabilities [23], since not all properties ϕ correspond to measurable sets via our semantics. This is similar to taking *inner measures* in measure theory. Note that

$$\max(0, [\![\phi]\!]q + [\![\psi]\!]q - 1) \leq [\![\phi]\!]q \cdot [\![\psi]\!]q,$$

where the latter would be a semantics for independent events. A motivation for equation (1) can be found in [23] (see also [19]) for a semantics of probabilistic predicate transformers. We simply adopt the argument in [23] to our different setting:

Imagine that $\prec$ is a preorder on $\mu\mathtt{M}$ such that its quotient $\mu\mathtt{M}/\equiv_\prec$ is a lattice, where conjunction and disjunction become the lattice operations up to equivalence. Let us write $(L; \leq, \wedge, \vee)$ for this lattice. We identify any $\phi \in \mu\mathtt{M}$ with its equivalence class in L and assume ϕ to be closed in order to keep this discussion simple. We also temporarily assume that $\mu\mathtt{M}$ allows for negation such that $\phi \wedge \neg\phi = \mathtt{false}$ and $\phi \vee \neg\phi = \mathtt{true}$ in L; otherwise a comparison to measure-theoretic notions will not be possible since such concepts require set-complements. For example, Kozen's standard interpretation of $\mu\mathtt{M}$ gives rise to such a lattice, where the equivalence classes are predicates: $[\phi]_{\equiv_\prec}$ is the set of all states q satisfying ϕ. Incidentally, the results of this paper extend to a modal mu-calculus *with* negation where the bodies of fixed points are restricted to formally monotone ones.

Given any state $q \in Q$ suppose that there were a way of assigning quantitative evidence $\mu_q(\phi) \in [0,1]$ to each $\phi \in L$ such that

$$\mu_q : L \to [0,1] \tag{2}$$

satisfies $\mu_q(\mathtt{false}) = 0$ and the usual modular law of measure theory:

$$\mu_q(\phi \wedge \psi) + \mu_q(\phi \vee \psi) = \mu_q(\phi) + \mu_q(\psi) \tag{3}$$

for all $\phi, \psi \in L$. If we now want to design a semantics $[\![\cdot]\!]^{D_\mathcal{P}}$ which provides lower bounds for the evidence given via μ_q, then our task at hand it to make $[\![\phi \wedge \psi]\!]^{D_\mathcal{P}} q$ a lower bound of $\mu_q(\phi \wedge \psi)$ which is also the *maximal* choice given our lack of knowledge of μ_q. So let us define $[\![\phi \wedge \psi]\!]^{D_\mathcal{P}} q$ according to equation (1). This is clearly a lower bound of $\mu_q(\phi \wedge \psi)$ if $[\![\phi \wedge \psi]\!]^{D_\mathcal{P}} q$ equals 0. Otherwise, if $[\![\phi]\!]^{D_\mathcal{P}} q \leq \mu_q(\phi)$ and $[\![\psi]\!]^{D_\mathcal{P}} q \leq \mu_q(\psi)$, we compute

$$\begin{array}{rcl} [\![\phi \wedge \psi]\!]^{D_\mathcal{P}} q & = & [\![\phi]\!]^{D_\mathcal{P}} q + [\![\psi]\!]^{D_\mathcal{P}} q - 1 \quad \text{(by equa. (1))} \\ & \leq & [\![\phi]\!]^{D_\mathcal{P}} q + [\![\psi]\!]^{D_\mathcal{P}} q - \mu_q(\phi \vee \psi) \\ & \leq & \mu_q(\phi) + \mu_q(\psi) - \mu_q(\phi \vee \psi) \\ & = & \mu_q(\phi \wedge \psi) \quad \text{(by equa. (3))}. \end{array}$$

Hence, equation (1) gives us a *compositional* lower estimate. Moreover, this estimate is as good as it could possibly be since we do not know anything about the concrete nature of μ_q. For example, let L be such that $\mathtt{false} \neq \phi \wedge \psi$ and $\phi \neq \psi$. Then we may define a strict function $\mu_q : L \to [0,1]$ by

$$\begin{array}{rcl} \mu_q(\phi \wedge \neg\psi) & = & [\![\phi]\!]^{D_\mathcal{P}} q \wedge (1 - [\![\psi]\!]^{D_\mathcal{P}} q) \\ \mu_q(\psi \wedge \neg\phi) & = & [\![\psi]\!]^{D_\mathcal{P}} q \wedge (1 - [\![\phi]\!]^{D_\mathcal{P}} q) \\ \mu_q(\phi \wedge \psi) & = & [\![\phi \wedge \psi]\!]^{D_\mathcal{P}} q. \end{array}$$

The reader may verify that $\mu_q : L \to [0,1]$ is modular with

$$\mu_q(\phi \wedge \psi) = [\![\phi \wedge \psi]\!]^{D_\mathcal{P}} q. \tag{4}$$

This concludes our motivation and justification for the semantics of conjunction.

Finally, we may interpret least and greatest fixed points semantically using the Knaster-Tarski Theorem for the complete lattice $[0,1]^Q$.

Some examples of this semantics reveal its relationship to probabilistic testing preorders, results in discrete Markov chain processes, and queueing network theory. We briefly sketch which fragments of the modal mu-calculus transform our semantics into the standard qualitative one by taking an optimistic/pessimistic stance towards quantitative evidence.

Our semantics makes $[0,1]^Q$ into a bi-continuous Σ-algebra, i.e. continuous on $[0,1]^Q$ and on its order dual; but mixed fixed-point computations might require more than ω iterations since the least fixed point operator

$$\mathtt{fix}_{[0,1]^Q} : [[0,1]^Q \to [0,1]^Q] \to [0,1]^Q$$

is not *bi*-continuous. It is an open problem whether this phenomenon can occur for a finite set of states Q. However, we show how these fixed point computations can be captured as standard optimization problems in linear programming. This yields an efficient alternative to fixed point iteration, which is likely to be amenable to further efficiency gains via techniques such as MTBDDs [8, 4].

The substitution principle for equality makes it possible to reduce the state space of a model, underlying a process term, provided that the notion of equivalence being considered is a congruence. Probabilistic bisimulations, which have their variants in generative systems or in Markov chain processes (the lumpability condition in [14]), are a robust notion of process equivalence for probabilistic transition systems [16] with polynomial-time decision procedures [3]. We compare this notion of equivalence to an equivalence relation $\equiv_{\mathtt{HML}}$ for reactive systems defined by our semantics via all formulas of Hennessy-Milner logic. For $=^{\mathtt{r}}$, an adaptation of probabilistic ready bisimulation [13] from generative to reactive systems, we show that $=^{\mathtt{r}}$ is strictly included in $\equiv_{\mathtt{HML}}$, and the latter is a proper subset of the union of all probabilistic bisimulations.

2. Quantitative Semantics

Given sets Act, C and Var such that Var is countable we obtain Kozen's positive modal mu-calculus $\mu\mathtt{M}$ by considering the following grammar:

$$\phi ::= c \mid \phi \wedge \phi \mid \phi \vee \phi \mid [a]\phi \mid \langle a \rangle \phi \mid x \mid \mu x.\phi \mid \nu x.\phi$$

where $a \in \mathsf{Act}$, $c \in \mathsf{C}$ and $x \in \mathsf{Var}$. The signature

$$\Sigma = \Sigma_0 \cup \Sigma_1 \cup \Sigma_2$$

of this language has

$$\Sigma_0 = \mathsf{C}$$

as the set of nullary operations,

$$\Sigma_1 = \{[a], \langle a\rangle \mid a \in \mathsf{Act}\}$$

contains all unary and

$$\Sigma_2 = \{\wedge, \vee\}$$

all binary operations.

We use the probabilistic transition systems of [16], which consist of a set of states Q, a set of labeled transitions

$$\Delta \subseteq Q \times \mathsf{Act} \times Q,$$

and a probability distribution $(\pi(q,a,q'))_{(q,a,q')\in\Delta}$ for each $q \in Q$ and $a \in \mathsf{Act}$ with $q \rightarrow^a$:

$$\sum_{q\rightarrow^a q'} \pi(q,a,q') = 1 \qquad (\pi(q,a,q') > 0).$$

We write $\mathcal{P}$ for such triples (Q, Δ, π) and make $[0,1]^Q$ into three different monotone Σ-algebras $D_{\mathcal{P}}$ by setting:

$$\begin{array}{rcl} \mathtt{true}^{D_{\mathcal{P}}} & = & \lambda q.1 \\ \mathtt{false}^{D_{\mathcal{P}}} & = & \lambda q.0 \\ f \wedge^{D_{\mathcal{P}}} g & = & \lambda q.f(q) * g(q) \\ \neg^{D_{\mathcal{P}}} f & = & \lambda q.1 - f(q) \\ f \vee^{D_{\mathcal{P}}} g & = & \neg^{D_{\mathcal{P}}}(\neg^{D_{\mathcal{P}}} f \wedge^{D_{\mathcal{P}}} \neg^{D_{\mathcal{P}}} g) \\ \langle a\rangle^{D_{\mathcal{P}}} f\, q & = & \displaystyle\sum_{(q,a,q')\in\Delta} \pi(q,a,q')\cdot f(q') \\ [a]^{D_{\mathcal{P}}} & = & \neg^{D_{\mathcal{P}}} \circ \langle a\rangle^{D_{\mathcal{P}}} \circ \neg^{D_{\mathcal{P}}}. \end{array}$$

Constants other than the truth values may be assigned arbitrary functions $f\colon Q \rightarrow [0,1]$ as meaning. For $f(q) * g(q)$ we consider *minimum*, *multiplication* or

$$\max(0, f(q) + g(q) - 1)$$

as in equation (1). The reader should check that the de Morgan dual of this fuzzy interpretation renders

$$(f \vee^{D_{\mathcal{P}}} g)\, q = \min(1, f(q) + g(q)) \qquad (5)$$

for all $q \in Q$. We already motivated that fuzzy choice in Subsection (1.3); in the presence of action-based choice between probabilistic distributions we cannot calculate exact probabilities; instead, the last version of $\wedge^{D_{\mathcal{P}}}$ computes lower probability bounds similar to taking inner measures [23]. Further, only the last version of $\wedge^{D_{\mathcal{P}}}$ makes $[a]^{D_{\mathcal{P}}}$ preserve finite conjunctions in $D_{\mathcal{P}}$ (modal algebra), with the other two operations being appropriate in the context of cost and/or time measures. If we state results about $D_{\mathcal{P}}$ they implicitly hold for all three versions of $\wedge^{D_{\mathcal{P}}}$ above, *unless stated otherwise*. Although we decided on a positive form of $\mu\mathtt{M}$, we have negation in the semantics and use it to define the semantics of $[a]$ and $\vee$ via the classical de Morgan dualities. Incidentally, our semantics is close to that of a linear mu-calculus in the sense that the meanings of $\langle a\rangle\psi$ and $[a]\psi$ coincide at all states q for which an a-move is possible (basic arithmetic, noting that the sum of all $\pi(q,a,q')$ with $(q,a,q') \in \Delta$ equals 1). However, if there is no such a-move, the former returns 0 (sum over the empty index set), whereas the second one computes 1. This is in accord with the standard interpretation on $\mathbf{2}^Q$.

The interpretations of conjunction and disjunction are *truly* fuzzy and abandon the realm of classical logic. For example, if we allowed negation in the syntax then the meaning of $\mathtt{true}$ and $\phi \vee \neg\phi$ would be different in general for the first two versions of $\vee$ above (for the third version consider $\mathtt{false}$ and $\phi \wedge \neg\phi$). This is to be expected since we only calculate lower bounds for probabilities.

Proposition 1 *The completely distributive lattice $D_{\mathcal{P}}$ is a monotone Σ-algebra with respect to the operations $\mathsf{op}^{D_{\mathcal{P}}}$, $\mathsf{op} \in \Sigma$, defined above. In particular, the least/greatest fixed point semantics*

$$[\![\,]\!]^{D_{\mathcal{P}}}\colon \mu\mathtt{M} \rightarrow \mathsf{Env}_{D_{\mathcal{P}}} \rightarrow D_{\mathcal{P}}$$

is well-defined.

Note that $[0,1]^Q$ is a complete lattice in the pointwise order. Thus, the semantics for least/greatest fixed points is the usual one using the Knaster-Tarski fixed point theorem. We point out that the semantics above can equally well be applied to *generative* systems, where

$$\sum_{a\in\mathsf{Act},\ q'\in Q} \pi(q,a,q') = 1$$

for all $q \rightarrow$; but then the semantics of $[a]\psi$ is quite different from the one for $\langle a\rangle\psi$.

3. Quantitative Model Checks

The first example checks a formula without any fixed points. It illustrates that our semantics is implicitly based on a probabilistic testing order. Such operational notions have explicitly been studied in [6]. Given the probabilistic transition system $\mathcal{P}$ below we check the formula

$$\phi \equiv [a]\langle b\rangle\mathtt{true}.$$

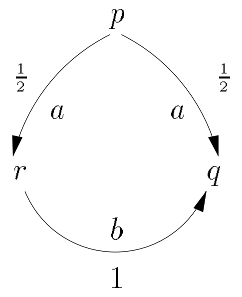

In the standard qualitative interpretation this yields the vector $(0, 1, 1)$ assuming the coordinates correspond to values at p, r, and q in that order. In our quantitative analysis, however, we compute $(0.5, 1, 1)$. Our semantics simulates drawing from a distribution whenever a possible a-move has to be made. That draw determines a path and the property in question will be satisfied by that path (success), or not (failure). Our semantics of $\langle a \rangle$ as a *weighted average* then reflects the expected value of the ratio of success and failure if we were to repeat that test and those draws infinitely often. This links our semantics to established testing paradigms and induces a testing preorder $q \sqsubseteq q'$ on Q defined as $[\![\phi]\!]^{D_{\mathcal{P}}} q \leq [\![\phi]\!]^{D_{\mathcal{P}}} q'$ for all formulas ϕ of Hennessy-Milner logic.

Note that combinations of box and diamond modalities render a non-classical result above. Combining diamonds and greatest fixed points, and dually boxes and least fixed points, is another source of non-classical qualitative outcome. In the next section we show that these are the only possible deviations if we re-interpret quantities as qualities. For an example of that second kind consider

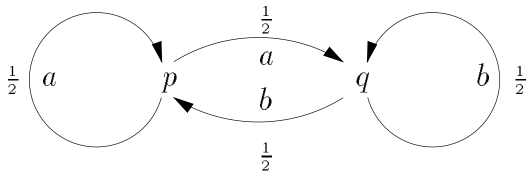

Intuitively, this model tosses a fair coin at each state deciding whether to change state or not. We evaluate the formula

$$\phi \equiv \nu x.\langle a \rangle x$$

for the states p and q: observe that $\langle a \rangle^{D_{\mathcal{P}}} f$ sends p to

$$\frac{1}{2} \cdot f(p) + \frac{1}{2} \cdot f(q)$$

and q to 0, and so $(\langle a \rangle^{D_{\mathcal{P}}})^{k+1} f$ sends p to

$$(\frac{1}{2})^{k+1} \cdot f(p)$$

and q to 0. Clearly, the ω limit of these iterations (the pointwise infimum) is $\lambda s.0$, the denotation of `false`. Our result matches the usual state-based probabilistic analysis of the event $\{a^\omega\}$. Note that we would have obtained the denotation of `true` with a qualitative model check.

To see the fuzzy semantics of disjunction at work, let us check

$$\psi \equiv \nu x.((\mu y.[a]y) \wedge (\mu y.[b]y)) \wedge (\langle a \rangle x \vee \langle b \rangle x),$$

which classically singles out those states allowing *only fair* infinite sequence of interleavings of a's and b's; i.e. the classical meaning is $\lambda s.0$. First, $F(f) = [a]^{D_{\mathcal{P}}} f$ maps p to $\frac{1}{2} \cdot f(p) + \frac{1}{2} \cdot f(q)$ and q to 1 so the *only* fixed point of F is $\lambda s.1$. By symmetry, $G(f) = [b]^{D_{\mathcal{P}}} f$ has the same least fixed point $\lambda s.1$. Since $\max(0, 1 + 1 - 1) = 1$ we obtain

$$[\![(\mu y.[a]y) \wedge (\mu y.[b]y)]\!]^{D_{\mathcal{P}}} = \lambda s.1.$$

The greatest fixed point of $H \colon D_{\mathcal{P}} \to D_{\mathcal{P}}$ with

$$H(f) = \lambda s.1 \wedge^{D_{\mathcal{P}}} (\langle a \rangle^{D_{\mathcal{P}}} f \vee^{D_{\mathcal{P}}} \langle b \rangle^{D_{\mathcal{P}}} f)$$

is $\lambda s.1$ as well, for $H(f)$ maps p to

$$\max(0, 1 + \min(1, \frac{1}{2} \cdot f(p) + \frac{1}{2} \cdot f(q) + 0) - 1),$$

which is

$$\min(1, \frac{1}{2} \cdot f(p) + \frac{1}{2} \cdot f(q) + 0),$$

and it maps q to $\min(1, \frac{1}{2} \cdot f(p) + \frac{1}{2} \cdot f(q) + 0)$. Thus,

$$[\![\psi]\!]^{D_{\mathcal{P}}} = \lambda s.1$$

shows that quantitative checking, unlike the qualitative one, semantically captures a *fair merge*. In particular, we reached the fixed point after only two iteration steps.

The next example illustrates the connections to queueing theory and linear algebra. It also makes the point that some questions can be asked in action-based *and* state-based semantics at the same time. We reason about the existence of deadlock in a simple system, a pipe shown below. The formula we want to check is

$$\phi = \mu x.[a]x,$$

which classically means that there are no infinite paths of a-actions. For our pipe $\mathcal{P}$ this just means that we can reach a state which deadlocks. One may ask that same question in a state-based semantics, by wanting to know the probability of the limit distribution at state q_n if we view $\mathcal{P}$ as a discrete random process. Ordering the states from q_0 up to q_n we may represent any $f \in [0, 1]^Q$ as an $(n + 1)$-dimensional vector $\mathbf{x}$ and the meaning of $[a]$ as an almost conventional $(n + 1) \times (n + 1)$ matrix $\mathbf{A}$. Yet, recall that the meaning of $[a]$ equals that of $\langle a \rangle$ for all states but q_n since $q_n \not\rightarrow^a$; for the latter state $[a]$ evaluates to 1. Therefore, if we want to encode in $\mathbf{A}$ all the necessary information about $[a]$ we compute $\mathbf{A}$ in two stages.

1. First, we let (a_{ij}) be the $(n+1) \times (n+1)$ matrix corresponding to the linear map induced by $\langle a \rangle$ on $D_{\mathcal{P}}$.

2. Whenever $q_i \not\rightarrow^a$ we set $x_i = 1$ and we may drop the i-th row vector from (a_{ij}). Thus, we arrive at a $(n + 1 - s) \times (n+1)$ matrix $\mathbf{B} = (b_{ij})$, where s is the size of the set

$$S = \{q_i \in Q \mid q_i \not\rightarrow^a\}. \tag{6}$$

   Then, we solve $\mathbf{Bx} = \mathbf{x}$ with the understanding that $x_i = 1$ for all $q_i \in S$. Of course, that can be summarized as

$$\mathbf{Cy} = \mathbf{y} \tag{7}$$

   where $\mathbf{y}$ is an $(n+1-s)$-dimensional vector and $\mathbf{C}$ a $(n+1-s) \times (n+1-s)$ matrix.

In our example this is all very simple since $S = \{p_n\}$, $\mathbf{C}$ equals

$$\begin{pmatrix} l_0 & 1-l_0 & 0 & \dots & \dots \\ 0 & l_1 & 1-l_1 & 0 & \dots \\ \vdots & & \ddots & & \dots \\ 0 & \dots & \dots & 0 & l_{n-1} \end{pmatrix}$$

and $\mathbf{y}$ is an $(n+1-1)$-dimensional vector. The existence of a minimal solution to $\mathbf{Cy} = \mathbf{y}$ is guaranteed by our denotational semantics. Assuming $0 < l_i < 1$ for all i, the constant one vector $\mathbf{1}$ solves $\mathbf{Cy} = \mathbf{y}$. Let $\mathbf{I}$ be the $n \times n$ unit matrix. Then

$$\mathbf{det}(\mathbf{C} - \mathbf{I}) = (1-l_0) \cdot (1-l_1) \cdot \ldots \cdot (1-l_n) \tag{8}$$

is non-zero, so the n-dimensional solution $\mathbf{1}$ of (7) is unique. Since x_{n+1} has to be 1 we see that the $(n+1)$-dimensional vector $\mathbf{1}$ equals the least fixed point computed by $[\![\phi]\!]^{D_{\mathcal{P}}}$.

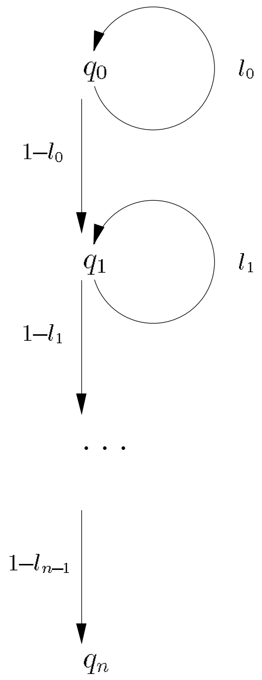

Probabilistic modal logic (`PML`) [16] is a modification of Hennessy-Milner logic without $[a]$, where the modality $\langle a \rangle$ is enriched with quantitative information $\langle a \rangle_\rho$, for $\rho \in [0,1]$. In this way one gets a semantics of `PML` of type 2^Q such that ρ acts as a threshold for truth in the meaning of $\langle a \rangle_\rho$. This allows one to characterize probabilistic bisimulation logically for certain finitely branching probabilistic transition systems [16]. The sole purpose of the next example is to demonstrate that the attempt to extend `PML` with fixed points *through relying on least/greatest fixed point semantics of type* 2^Q must fail (assume both transitions are labeled with a):

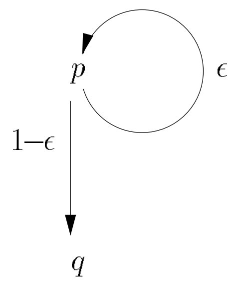

Checking $\nu x.\langle a \rangle x$ in our semantics yields $(0,0)$. However, if we check

$$\nu x.\langle a \rangle_\rho x$$

as a greatest fixed point of the `PML` semantics in 2^Q we get $(1,0)$ for $0 < \rho \leq \epsilon$, so adjusting the threshold parameter ρ may lead to flawed results. The next example also shows that it is impossible to fix this shortcoming by choosing the right threshold; the latter has to vary with its state. Another reason for preferring a numerical semantics is that it is, unlike the one for PML, *continuously dependent* on the actual numbers in

$$\{\pi(q,a,q') \mid (q,a,q') \in \Delta\}.$$

We conclude our series of examples with a non-trivial *infinite state* system:

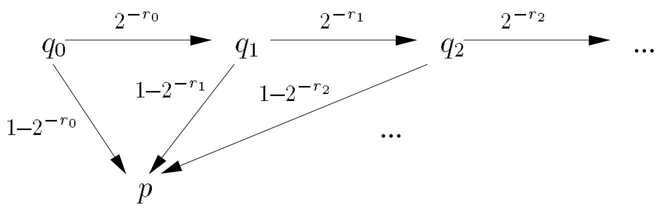

where all r_i are positive real numbers such that the infinite sum $\sum_{i=0}^{\infty} r_i$ converges. We check

$$\nu x.\langle a\rangle x$$

by calculating the greatest fixed point starting from $f = \lambda s.1$. Clearly,

$$(\langle a\rangle^{D_\mathcal{P}})^{\alpha}\,\lambda s.1\,p = 0$$

for all $0 \leq \alpha \leq \omega$. For any $i \geq 0$ we get

$$(\langle a\rangle^{D_\mathcal{P}})^{1}\,\lambda s.1\,q_i = (1-2^{-r_i})\cdot 1 + 2^{-r_i}\cdot 1 = 1$$

and therefore

$$(\langle a\rangle^{D_\mathcal{P}})^{2}\,\lambda s.1\,q_i = (1-2^{-r_i})\cdot 0 + 2^{-r_i}\cdot 1 = 2^{-r_i}.$$

By induction, we obtain

$$(\langle a\rangle^{D_\mathcal{P}})^{k+1}\,\lambda s.1\,q_i = 2^{-(\Sigma_{j=i}^{k+i} r_j)} \qquad (i \geq 0)$$

for all $0 \leq k \leq \omega$ resulting in the limit (coordinatewise infimum)

$$(\langle a\rangle^{D_\mathcal{P}})^{\omega}\,\lambda s.1\,q_i = 2^{-(\Sigma_{j=i}^{\infty} r_j)} \qquad (i \geq 0).$$

The latter infinite vector (with initial 0 for p) is clearly a fixed point of $\langle a\rangle^{D_\mathcal{P}}$ and therefore equal to $[\![\nu x.\langle a\rangle x]\!]^{D_\mathcal{P}}$. These results are in accord with the limiting probability transitions from state q_i to p. Note that filtered infima and directed suprema of fixed point computations in $D_\mathcal{P}$ are the topological limits of their finite unwindings in the Euclidean topology.

4. On the Convergence of Fixed Points

One of the center pieces of Denotational Semantics is the fact that, given a continuous map $f \in [D \to D]$ on a dcpo D with bottom $\bot$, we may compute the least fixed point $\mathtt{fix}_D(f)$ of f as the directed supremum of the finite iterations $f^i(\bot)$ [1]. This models incremental and computationally meaningful gain of information. Moreover, the function

$$\mathtt{fix}_D(f)\colon [D \to D] \to D$$

is continuous [1]. However, in the presence of two kinds of fixed points we do not get a perfectly symmetric transfer of these results.

Restricting to *finitely branching* models, the Σ-algebra $D_\mathcal{P}$ is not only continuous on $[0,1]^Q$ but also on its order dual (= bi-continuous). So fixed point iterations of formulas which contain only one kind of fixed point reach their meaning at level ω. However, this is no longer true if we nest least and greatest fixed points in specifications. For example, there are finitely branching labeled transition systems, where for each q and each a there is at most one a-move such that the computation of Kozen's standard fixed point semantics requires $\omega+1$ steps for a particular ϕ [5]. We may simply enrich such a model by assigning probability 1 to each possible transition, and in this special scenario our semantics computes just the qualitative result. At first glance something seems wrong with this. But the reason why this central fact of Denotational Semantics does not hold in this mixed setting is that $\mathtt{fix}_{D_\mathcal{P}}\colon [D_\mathcal{P} \to D_\mathcal{P}] \to D_\mathcal{P}$ is not *bi*-continuous!

The examples alluded to above all involve infinitely many states. So it is an interesting theoretical problem whether fixed point iterations need not go beyond ω for our semantics in *finite state* models. Notice that this has a trivial affirmative answer in the qualitative case as then $\mathbf{2}^Q$ is a finite lattice. No matter what the answer to this problem will be, we later sketch an approach to computing fixed points based on *linear programming*, which does not rely on fixed point iterations at all. Nonetheless, our (numerical) fixed point iterations do reveal the nature of the gain of information as our model check of the infinite system above reveals.

5. Re-covering Quality in Quantity

How may we compare quantitative and qualitative results? This requires a function

$$t\colon [0,1] \to \mathbf{2},$$

where $t^{-1}(1)$ are all the quantities we deem to result in "yes" answers. Obviously, there are many choices, even if we commit ourselves to monotone ones. We just state two of those, the extreme optimist, $t_\mathtt{o}$, and the extreme pessimist, $t_\mathtt{p}$, and match them to corresponding sublanguages of $\mu\mathtt{M}$. The extreme optimist satisfies

$$t_\mathtt{o}^{-1}(1) = (0,1].$$

So it is only willing to reply "no" to a check if the computed quantity equals 0. Dually,

$$t_{\mathtt{p}}^{-1}(1) = \{1\}$$

makes our pessimist reply "yes" only if the quantitative evidence is 1.

The domain for the qualitative semantics is $D_{\mathcal{L}}$, the lattice $\mathbf{2}^Q$ with Kozen's standard interpretation of $\mu\mathtt{M}$. Thus, we have an action

$$t^*\colon D_{\mathcal{P}} \to D_{\mathcal{L}}$$

via composition

$$t^*(f) = t \circ f.$$

One readily sees that $t^*_{\circ}\colon D_{\mathcal{P}} \to D_{\mathcal{L}}$ is a strict continuous map. This allows us to use the principle of uniform fixed points [1] to relate least fixed points of $D_{\mathcal{P}}$ to those of $D_{\mathcal{L}}$. Routine calculations show that $t^*_{\circ}$ is a homomorphism of Σ-algebras for all operations *but* (semantic) negation and $[a]$. Putting these pieces together we obtain:

Theorem 1 *Let ϕ be any formula in the following sublanguage of $\mu\mathtt{M}$ (positive existential fragment):*

$$\phi ::= \mathtt{true} \mid \mathtt{false} \mid \phi \wedge \phi \mid \phi \vee \phi \mid \langle a \rangle \phi \mid x \mid \mu x.\phi$$

where $a \in \mathsf{Act}$ and $x \in \mathsf{Var}$. Given $\mathcal{P} = (Q, \Delta, \pi)$ and $\mathcal{L} = (Q, \Delta)$ we have

$$t_{\circ} \circ [\![\phi]\!]^{D_{\mathcal{P}}} \rho = [\![\phi]\!]^{D_{\mathcal{L}}} (t_{\circ} \circ \rho)$$

for all environments ρ, where $[\![\cdot]\!]^{D_{\mathcal{L}}}$ is the semantics of qualitative model checking, provided that the semantics of $\wedge$ is multiplication or minimum.

By duality, we get a result involving $t_{\mathtt{p}}$ by replacing $\langle a \rangle$ with $[a]$ and μx with νx. We should point out that the combinations excluded in these results are those where a deviation from the usual qualitative results is desirable in capturing probabilistic finite or infinite behavior.

6. Semantic Equivalence of Processes

Probabilistic bisimulations have been introduced in [16] as certain equivalence relations $\equiv$ such that for all states $p \equiv q$ and all equivalence classes S the sum of the probabilities of a-moves $p \to^a p'$ into S equals to the sum of the probabilities of a-moves $q \to^a q'$ into S. Such relations play a similar role as strong bisimulations [18, 20], but for the class of probabilistic transition systems.

Assuming that we have some notion of a probabilistic process algebra such that the models of terms are probabilistic transition systems, one would like a semantics of a specification language not to distinguish probabilistically bisimilar states. The latter has polynomial time decision procedures [3] and can be used in simplifying the process term without constructing the entire state space in the first place. Thus, a soundness result is mandatory for using such state-aggregation techniques.

Theorem 2 *Given that p and q are probabilistically bisimilar, we have $[\![\phi]\!]^{D_{\mathcal{P}}} p = [\![\phi]\!]^{D_{\mathcal{P}}} q$ for all formulas ϕ of Hennessy-Milner logic.*

Proof: The proof proceeds by induction on the modal depth of ϕ. We focus on the case where ϕ equals $\langle a \rangle \psi \in \mathtt{HML}_{k+1}$. The set $\mathtt{HML}_n$ denotes all formulas of Hennessy-Milner logic of modal depth n. So let $\equiv$ be a probabilistic bisimulation such that $p \equiv q$. Given

$$\begin{aligned} F &= \{p' \in Q \mid (p, a, p') \in \Delta\} \\ G &= \{q' \in Q \mid (q, a, q') \in \Delta\}, \end{aligned}$$

the equivalence relation $\equiv$ induces a partition

$$(M_1, M_2, \ldots, M_n)$$

on the set $F \cup G$. Since $\equiv$ is a probabilistic bisimulation, we have

$$\sum_{p' \in M_i \cap F} \pi(p, a, p') = \sum_{q' \in M_i \cap G} \pi(q, a, q') \tag{9}$$

for all $1 \leq i \leq n$. By our induction hypothesis, we know that $[\![\psi]\!]^{D_{\mathcal{P}}} p' = [\![\psi]\!]^{D_{\mathcal{P}}} q'$ for all $p' \equiv q'$. Thus equation (9) entails

$$\begin{aligned} [\![\phi]\!]^{D_{\mathcal{P}}} p &= \sum_{p' \in F} \pi(p, a, p') \cdot [\![\psi]\!]^{D_{\mathcal{P}}} p' \\ &= \sum_{i=1}^{n} \sum_{p' \in M_i \cap F} \pi(p, a, p') \cdot [\![\psi]\!]^{D_{\mathcal{P}}} p' \\ &= \sum_{i=1}^{n} \sum_{q' \in M_i \cap G} \pi(q, a, q') \cdot [\![\psi]\!]^{D_{\mathcal{P}}} q' \\ &= [\![\phi]\!]^{D_{\mathcal{P}}} q. \end{aligned}$$

End of Proof.

Since our semantics captures a notion of probabilistic testing, we cannot expect a completeness result for probabilistic bisimulation itself. As a counter-example consider the two trees in the figure below. The roots are not probabilistically bisimilar, but one can prove that they compute the same numbers for every formula of Hennessy-Milner logic.

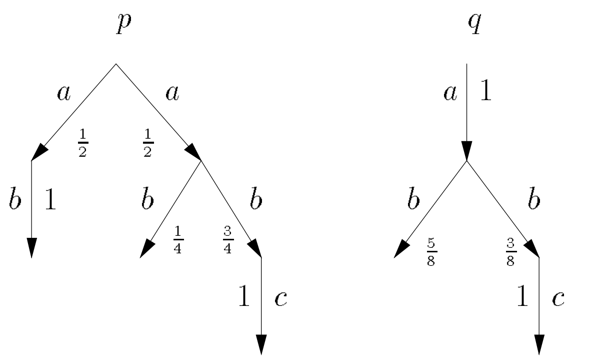

This example also shows that the pseudo-metric

$$d(p,q) = \sum\{\frac{|\,[\![\phi]\!]^{D_\mathcal{P}}p - [\![\phi]\!]^{D_\mathcal{P}}q\,|}{2^n} \;\mid\; n \geq 1, \mathsf{ht}(\phi) \leq n\}$$

on Q induced by our denotational semantics is not a metric. We assume sort-finiteness of $\mathcal{P}$ here to ensure that the pseudo-metric d is well-defined. Of course, $\mathsf{ht}(\phi)$ denotes the height of the parse tree of ϕ.

Probabilistic ready bisimulation has been defined in [13] for generative systems; it has a re-definition for reactive systems.

Definition 1 *We define a function*

$$R\colon Q \to \mathsf{Act}^* \to 2^{\mathsf{Act}} \to [0,1]$$

by induction on the length of $w \in \mathsf{Act}^*$. Let $\mathsf{init}(p)$ be the set of initial actions of p in $\mathcal{P}$. Then

$$R\;p\;\epsilon\;S = \begin{cases} 1 & \textit{if } S = \mathsf{init}(p) \\ 0 & \textit{otherwise} \end{cases}$$

defines R on the empty word ϵ and

$$R\;p\;aw\;S = \sum_{(p,a,p')\in\Delta} \pi(p,a,p') \cdot (R\;p\;w\;S)$$

specifies the inductive case. We say that p and q are probabilistically ready bisimilar *if $R\;p = R\;q$ in the space*

$$\mathsf{Act}^* \to 2^{\mathsf{Act}} \to [0,1].$$

Let us assume that $\mathcal{P}$ is sort-finite. Then

$$\phi_S \equiv \bigwedge_{a\in S} \langle a \rangle \mathtt{true}$$

is a formula of Hennessy-Milner logic for any $S \subseteq \mathsf{Act}$. Clearly, $[\![\phi_S]\!]^{D_\mathcal{P}}p$ returns 1 if $S = \mathsf{init}(p)$ and 0 otherwise. Since the inductive case of the definition of R is just like our semantics of $\langle a \rangle$, one can show

$$[\![\langle a_1\rangle\langle a_2\rangle \ldots \langle a_n\rangle \phi_S]\!]^{D_\mathcal{P}} p = R\;p\;a_1a_2\ldots a_n\;S$$

for all $S \subseteq \mathsf{Act}$, $p \in Q$, and $a_1a_2\ldots a_n \in \mathsf{Act}^*$ by induction on $n \geq 0$.

Theorem 3 *Let $\mathcal{P}$ be sort-finite. If $[\![\phi]\!]^{D_\mathcal{P}}p = [\![\phi]\!]^{D_\mathcal{P}}q$ for all formulas ϕ of Hennessy-Milner logic, then p and q are probabilistically ready bisimilar.*

The converse of the above is not true. The roots of the two trees below are probabilistically ready bisimilar, but we can distinguish them using the HML formula

$$\phi \equiv \langle a \rangle(\langle c\rangle\mathtt{true} \vee \langle b\rangle\langle d\rangle\mathtt{true}),$$

for

$$\begin{aligned} [\![\phi]\!]^{D_\mathcal{P}}p &= \frac{1}{2}\cdot\min(1,1+1) + \frac{1}{2}\cdot\min(1,0+0) = \frac{1}{2} \\ [\![\phi]\!]^{D_\mathcal{P}}q &= \frac{1}{2}\cdot\min(1,1+0) + \frac{1}{2}\cdot\min(1,0+1) = 1. \end{aligned}$$

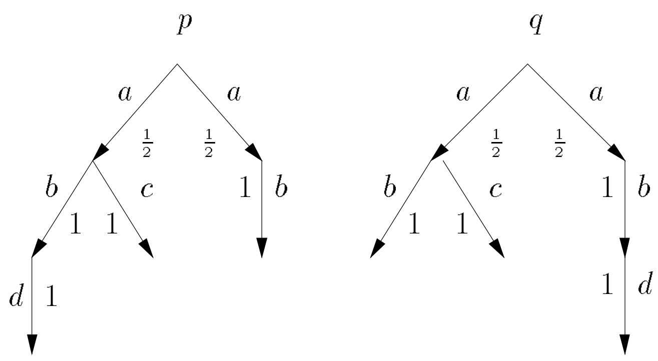

7. Global Model Checking Algorithm

Let Q be $\{q_1, q_2, \ldots, q_n\}$. We think of elements of $D_\mathcal{P}$ as real vectors $\mathbf{x} = (x_1, x_2, \ldots, x_n)$ in $\mathbf{R}^n$. We sketch a model checking algorithm for all closed formulas ϕ which contain only *unnested* fixed points [5]: if $\sigma x.\psi$ is a subformula of ϕ with $\sigma \in \{\mu,\nu\}$ then ψ has at most x as a free variable. Note that this includes the translation of all CTL formulas into μM [17]. We write $[\![\phi]\!]$ for the output vector of this algorithm. The vector $[\![\phi]\!]$ is computed by structural induction on ϕ. Before we sketch that algorithm we need to prove a technical fact that ensures that this algorithm works for fixed point formulas.

Proposition 2 *Let $\psi \in \mu$M be a closed formula which contains only* unnested *fixed points and has only x as a free variable ($\sigma \in \{\mu,\nu\}$). If the algorithm for computing $[\![\,]\!]$ is already well-defined for all subformulas of ψ then there exists a real number r_i^ψ and an n-dimensional vector $\mathbf{b}_i^\psi$ such that*

$$[\![\psi]\!]_i = \,<\mathbf{b}_i^\psi, \mathbf{x}>\, + r_i^\psi \quad (1 \leq i \leq n) \qquad (10)$$

where

$$<\mathbf{x}, \mathbf{y}> \,= \sum_{i=1}^n x_i \cdot y_i \qquad (11)$$

is the scalar product in $\mathbf{R}^n$.

Proof: We prove this by structural induction on ψ.

1. If $\psi \equiv x$ we set $r_i^\psi = 0$ and $\mathbf{b}_i^\psi = \mathbf{e}_i$, the i-th unit vector. This definition is in accord with the stipulation in (15).

2. If $\psi \equiv \psi_1 \wedge \psi_2$ then we choose $r_i^\psi = 0$ and $\mathbf{b}_i^\psi = \mathbf{0}$ if $[\![\psi_1 \wedge \psi_2]\!]_i = 0$. Otherwise, we use induction on ψ_1 and ψ_2 and get

$$\begin{aligned} [\![\psi_1 \wedge \psi_2]\!]_i &= [\![\psi_1]\!]_i + [\![\psi_2]\!]_i - 1 \\ &= (<\mathbf{b}_i^{\psi_1}, \mathbf{x}> + r_i^{\psi_1}) \\ &\quad + (<\mathbf{b}_i^{\psi_2}, \mathbf{x}> + r_i^{\psi_2}) - 1 \\ &= <\mathbf{b}_i^{\psi_1} + \mathbf{b}_i^{\psi_2}, \mathbf{x}> \\ &\quad + (r_i^{\psi_1} + r_i^{\psi_2} - 1) \end{aligned}$$

which suggests defining

$$\begin{aligned} \mathbf{b}_i^\psi &= \mathbf{b}_i^{\psi_1} + \mathbf{b}_i^{\psi_2} \\ r_i^\psi &= r_i^{\psi_1} + r_i^{\psi_2} - 1. \end{aligned}$$

3. The case $\psi \equiv \psi_1 \vee \psi_2$ is similar.

4. The cases $\psi \equiv \sigma' z.\psi_1$ and $\psi \equiv y$ with $y \not\equiv x$ cannot occur since x is free in ψ and ψ contains only unnested fixed points.

5. If $\psi \equiv \langle a \rangle \psi_1$ then $[\![\psi]\!]_i$ is defined as

$$\sum_{q_i \to^a q_k} \pi(q_i, a, q_k) \cdot [\![\psi_1]\!]_k. \tag{12}$$

By induction we have

$$[\![\psi_1]\!]_k = <\mathbf{b}_k^{\psi_1}, \mathbf{x}> + r_k^{\psi_1} \tag{13}$$

for all $1 \leq k \leq n$. Using equation (12), we compute

$$\begin{aligned} [\![\psi]\!]_i &= \sum_{q_i \to^a q_k} \pi(q_i, a, q_k) \cdot [\![\psi_1]\!]_k \\ &= \sum_{q_i \to^a q_k} \pi(q_i, a, q_k) \cdot (<\mathbf{b}_k^{\psi_1}, \mathbf{x}> + r_k^{\psi_1}) \\ &= (\sum_{q_i \to^a q_k} \pi(q_i, a, q_k) \cdot <\mathbf{b}_k^{\psi_1}, \mathbf{x}>) \\ &\quad + (\sum_{q_i \to^a q_k} \pi(q_i, a, q_k) \cdot r_k^{\psi_1}) \\ &= <\sum_{q_i \to^a q_k} \pi(q_i, a, q_k) \cdot \mathbf{b}_k^{\psi_1}, \mathbf{x}> \\ &\quad + (\sum_{q_i \to^a q_k} \pi(q_i, a, q_k) \cdot r_k^{\psi_1}) \end{aligned}$$

which suggests defining

$$\begin{aligned} \mathbf{b}_i^\psi &= \sum_{q_i \to^a q_k} \pi(q_i, a, q_k) \cdot \mathbf{b}_k^{\psi_1} \\ r_i^\psi &= \sum_{q_i \to^a q_k} \pi(q_i, a, q_k) \cdot r_k^{\psi_1}. \end{aligned}$$

6. Finally, if $\psi \equiv [a]\psi_1$ we have

$$[\![\psi]\!]_i = 1 - \sum_{q_i \to^a q_k} \pi(q_i, a, q_k) \cdot (1 - [\![\psi_1]\!]_k). \tag{14}$$

In particular, if $q_i \not\to^a$ then $[\![\psi]\!]_i = 1$ and we choose $\mathbf{b}_i^\psi = \mathbf{0}$ and $r_i^\psi = 1$. Otherwise, $[\![\psi]\!]_i$ equals $[\![\langle a \rangle \psi_1]\!]_i$ and we covered that case already.

End of Proof.

Now we are in position to specify our global model checking algorithm:

Definition 2 Global Model Checking Algorithm.

1. *$\phi \equiv \mathsf{c}$; then $[\![\mathsf{c}]\!]$ is the interpretation of the constant* c.

2. *$\phi \equiv \phi_1 \wedge \phi_2$; then $[\![\phi]\!]_i = \max(0, [\![\phi_1]\!]_i + [\![\phi_2]\!]_i - 1)$ for $1 \leq i \leq n$.*

3. *$\phi \equiv \phi_1 \vee \phi_2$; then $[\![\phi]\!]_i = \min(1, [\![\phi_1]\!]_i + [\![\phi_2]\!]_i)$.*

4. *$\phi \equiv [a]\psi$; then $[\![\phi]\!]_i$ is defined as in (14).*

5. *$\phi \equiv \langle a \rangle \psi$; then $[\![\phi]\!]_i$ is defined as in (12).*

6. *$\phi \equiv \mu x.\psi$; if x is not free in ψ, we return $[\![\psi]\!]$; otherwise, x is the* only *free variable in ψ, for ϕ is an unnested fixed point. The latter also ensures that free occurrences of x are never under the scope of some σy ($\sigma \in \{\mu, \nu\}$) in ψ. We stipulate that*

$$[\![x]\!]_i = x_i \tag{15}$$

for all occurrences x which are subformulas of ψ. It then remains to minimize $\mathbf{x}$ in $[0, 1]^Q$ for all solutions $\mathbf{x}$ of

$$[\![\psi]\!]_i = x_i \quad (1 \leq i \leq n) \tag{16}$$

subject to $\mathbf{0} \leq \mathbf{x} \leq \mathbf{1}$. By Proposition 2 we have $\mathbf{b}_i^\psi$ and r_i^ψ such that equation (10) holds. From this it follows that me may rephrase the equations in (16) as

$$<\mathbf{b}_i^\psi, \mathbf{x}> + r_i^\psi = x_i \quad (1 \leq i \leq n). \tag{17}$$

Clearly, we may write the system of equations in (17) in the form

$$\mathbf{B}_\psi \mathbf{x} = \mathbf{r}_\psi \tag{18}$$

for a $n \times n$-matrix $\mathbf{B}_\psi$ and an n-dimensional vector $\mathbf{r}_\psi$ by putting real numbers to the right hand and scalar multiples of components of $\mathbf{x}$ to the left hand side of the equations in (17). These rewrites do not affect the space of solutions $\mathbf{x}$. Since our denotational semantics guarantees a unique minimal fixed point for $\mu x.\psi$, we *may compute this vector as follows: minimize the performance measure*

$$P(\mathbf{x}) = \sum_{i=1}^{n} x_i \qquad (19)$$

for all solutions of (18) subject to the constraints

$$\mathbf{0} \leq \mathbf{x} \leq \mathbf{1}. \qquad (20)$$

The latter may be reduced to the standard constraints $\mathbf{0} \leq \mathbf{x}$ *using the Upper-Bounding Algorithm for the Simplex method [22]. However, the special nature of our semantics (*unique *optimal vertex) might be exploited to derive specialized optimization algorithms.*

7. $\phi \equiv \nu x.\psi$*; this is dual to the case before. We work with the linear system as in (18), the performance measure* P *as in (19), and the boundary conditions as in (20), but now need to* maximize $\mathbf{x}$ *with respect to* P.

Note that this algorithm can be seen as a *local* model checking algorithm if we omit fixed points altogether, i.e. if we restrict μM to HML.

Postscript: Since the system of linear equations in (18) is guaranteed to have a solution by our denotational semantics, one may solve it using Gaussian elimination and then substituting 1 for the free choices of x_i if the maximal solution within $\mathbf{0} \leq \mathbf{x} \leq \mathbf{1}$ is sought; the minimal solution is obtained by setting $x_i = 0$ for all x_i that we may choose freely. This greatly reduces the complexity of our global model checking algorithm for μM since it circumvents the linear programming problem entirely.

8. Outlook

We presented a non-standard semantics of the modal mu-calculus as a mathematical foundations for quantitative model checking of reactive and generative systems. The key technical insight was that infinite fixed point iterations can be overcome by restating these fixed point equations as equivalent problems in linear programming. The key conceptual insight was that quantitative model checking allows us to calculate lower bounds for probabilistic evidence of properties, which can provide satisfactory guarantees in cases where conventional qualitative model checking fails. The obvious extensions of this work will be such foundations for *generative* systems which use pairs of actions and rates as action types as in PEPA. Since our reactive semantics does not replace non-determinism with probabilistic choice, the outcomes of our model checks are probabilistic estimates in much the same way as the values obtained by *inner* measures in the Dempster-Shafer theory of evidence for dealing with uncertainty in AI. It would be interesting to link our results with such belief measures. A generative semantics, however, does resolve non-determinism probabilistically, so there is hope that generative checks have a precise probabilistic underpinning as can be glimpsed from some of our examples. The derivation of more efficient methods for our global model checking algorithm would be a worthwhile goal; for example, MTBDDs [8], which are similar to the OBDD representation of the transition relation [7, 17], allow for a compact representation of matrices that performs no worse than sparse matrices and can be applied to solve linear equations efficiently. MTBDDs have been used to specify and develop a practical model checker for probabilistic CTL with very encouraging results [4], and we anticipate that it should not be too difficult to apply them in the context of the modal mu-calculus. They are unlikely, though, to handle infinite state systems, for which suitable methods would be desirable, unless combined with induction. Finally, we remark that our approach is robust enough to deal with probabilistic transition systems enriched with durations of actions and quantitative checks of properties being realizable within time t, thus generalizing the work of [11, 10].

9. Acknowledgements

We would like to thank Samson Abramsky, Christel Baier, Julian Bradfield, Ed Clarke, Jane Hillston, Achim Jung, Matthias Kegelmann, Carroll Morgan, Gethin Norman, Edmund Robinson, Jeff Sanders, David Schmidt, and the anonymous referees for their stimulating and insightful comments about this work. This work has been supported in part by EPSRC grant GR/K42028.

References

[1] S. Abramsky and A. Jung. Domain theory. In S. Abramsky, D. M. Gabbay, and T. S. E. Maibaum, editors, *Handbook of Logic in Computer Science*, volume 3, pages 1–168. Clarendon Press, 1994.

[2] A. Aziz, K. Sanwal, V. Singhal and R. Brayton. Verifying Continuous Time Markov Chains. In *Proceedings of CAV'96*, volume 1102 of *Lecture Notes in Computer Science*. Springer Verlag, 1996.

[3] C. Baier. Polynomial Time Algorithms for Testing Probabilistic Bisimulation and Simulation. In *Proceedings of CAV'96*, volume 1102 of *Lecture Notes in Computer Science*, pages 38–49. Springer Verlag, 1996.

[4] C. Baier, M. Kwiatkowska, M. Ryan, E. Clarke, and V. Hartonas-Garmhausen. Symbolic Model Checking for Probabilistic Processes. In *Proceedings of* ICALP'97, 1997. To appear.

[5] J. C. Bradfield. *Verifying Temporal Properties Of Systems*. Birkhaeuser, Boston, Mass., 1991.

[6] I. Christoff. *Testing Equivalences for Probabilistic Processes*. PhD thesis, Department of Computer Systems, Uppsala University, Sweden, 1990.

[7] E. M. Clarke J. R. Burch, D. L. Dill K. L. McMillan, and J. Hwang. Symbolic model checking: 10^{20} states and beyond. *Information and Computation*, 98(2):142–170, 1992.

[8] E. M. Clarke, M. Fujita, X. Zhao. Multi-Terminal Binary Decision Diagrams and Hybrid Decision Diagrams, Representations of Discrete Functions, Edited by Tsutomu Sasao and Masahiro Fujita, Kluwer Academic Publishers, 1996, pages 93-108.

[9] S. Gilmore and J. Hillston (guest editors). *Process Algebra and Performance Modelling*, volume 38 of *the Computer Journal*. Oxford University Press (The British Computer Society), 1995.

[10] H. Hansson. *Time and Probability in Formal Design of Distributed Systems*. Elsevier, 1994.

[11] H. Hansson and B. Jonsson: A Logic for Reasoning about Time and Probability, *Formal Aspects of Computing* 6:512-535, 1994.

[12] J. Hillston. *A Compositional Approach to Performance Modelling*. Cambridge University Press. Distinguished Dissertation Series, 1996.

[13] C. Jou and S. Smolka. Equivalences, congruences and complete axiomatizations for probabilistic processes. In *Proceedings CONCUR'90*, volume 458 of *Lecture Notes in Computer Science*. Springer Verlag, pages 367-383, 1990.

[14] J.G. Kemeny and J.L. Snell. *Finite Markov Chains*. Van Nostrand, 1960.

[15] D. Kozen. Results on the propositional mu-calculus. *Theoretical Computer Science*, 27:333–354.

[16] K. G. Larsen and A. Skou. Bisimulation through Probabilistic Testing. *Information and Computation*, 94(1):1–28, September 1991.

[17] K. L. McMillan. *Symbolic Model Checking*. Kluwer Academic Publishers, 1993.

[18] R. Milner. *Communication and Concurrency*. Series in Computer Science. Prentice-Hall International, 1989.

[19] C. Morgan, A. McIver and K. Seidel. Probabilistic Predicate Transformers. *ACM Transactions on Programming Languages and Systems*, 8(1):1–30, January 1999.

[20] D. M. Park. On the Semantics of Fair Parallelism. In D. Bjorner, editor, *Abstract Software Specifications*, volume 86 of *Lecture Notes in Computer Science*, pages 504–526. Springer Verlag, 1980.

[21] R. Segala and N. Lynch: Probabilistic Simulations for Probabilistic Processes. In *Proceedings CONCUR 94*, volume 836 of *Lecture Notes in Computer Science*, Springer Verlag, pages 481-496, 1994.

[22] D. A. Pierre. *Optimization Theory with Applications*. Dover Publications, Inc., 2nd edition edition, 1986.

[23] K. Seidel C. Morgan, A. McIver and J.W. Sanders. Probabilistic Predicate Transformers. Technical Report PRG-TR-4-95, Oxford University Computing Laboratory, 1995.