A Proof System for Reasoning about Probabilistic Concurrent Processes

Matteo Mio

LFCS, School of Informatics, University of Edinburgh
M.Mio@sms.ed.ac.uk

The need of formally specify and verify the behavior of increasingly complex systems is posing many new challenges, one of them being the difficulty of proving and reasoning over properties of systems that exhibit some kind of probabilistic behavior. Following [1], we consider probabilistic concurrent processes formalized as probabilistic transition systems given as a pair $\langle P, \{\xrightarrow{a}\}_{a\in L}\rangle$, where L is some set of labels, $\mathcal{D}(P)$ is the set of probability distributions on the set of states P, and $\xrightarrow{a} \subseteq P \times \mathcal{D}(P)$ is the a-labeled transition relation. This definition generalizes that of standard labeled transition system.

The modal μ-calculus has proven to be a very expressive and yet tractable logic for expressing properties of labeled transition systems. In [2, 3] an extension of this logic to the context of probabilistic transition systems, called quantitative modal μ-calculus or just $qM\mu$, is studied and its semantics (assigning real values in $[0, 1]$ to pairs of states and formulae) is given denotationally by means of fixed points. However understanding μ-calculus (and qMμ) formulae denotationally can be very hard. In [4] the authors propose a different semantics based on Two Player Stochastic games for qMμ very similar to the, nowadays standard, Two Player game semantics for the modal μ-calculus introduced in [5]: the main difference being that in the game configuration $(p, \langle a\rangle F)$ (respectively $(p, [a]\, F)$), Player 1 (respectively Player 2) chooses an a-successor of p (i.e. a distribution); then the next game configuration is chosen probabilistically according to the distribution. This allows one to understand operationally the meaning of a qMμ formula as a limit probability, called the *value*, of winning the game according to the usual rules for winning μ-calculus games [5]; for instance the value associated to the pair $(p, \mu X.\, [a]\, X)$ is the greatest lower bound for the probability of making a sequence of a-actions and eventually reaching a state from which no a-action is possible. This semantics is arguably easier to understand than the denotational semantics and therefore offers a tractable way to deal with the meaning of qMμ formulae.

In the past 15 years, several proof systems have been proposed for reasoning over μ-calculus properties. In [6, 7] the author introduces a sound and complete sequent based proof system for proving Hennessy-Milner properties for processes described by a class of well behaved process calculi. At the same time a sequent based proof system for CCS processes and general μ-calculus properties is introduced in [8]. These papers provided evidence for the many advantages offered by the use of sequents as basic judgments in the proposed proof systems (see [7] for a detailed overview). More recent work on sequent calculi for the modal μ-calculus include [9] and [10]. One of the most interesting aspects of these proof systems is that derivations are in general infinite (finite branching) trees, where the (left and right) rules for fixed points formulae are just unfoldings. Proofs are therefore not just finite derivations with axioms at the leafs, but infinite derivations satisfying some *proof condition*. The general study of this kind of sequent based infinitary proof systems is developed in [11] in the context of first-order logic extended with inductive definitions. The proof condition in these systems is always expressed as a ω-regular property that every infinite path in the derivation must satisfy, which in turn is based

$$\frac{x \xrightarrow{a} \alpha}{x|y \xrightarrow{a} \alpha|y} \mid \mathcal{R} \qquad \frac{y \xrightarrow{a} \alpha}{x|y \xrightarrow{a} x|\alpha} \mid \mathcal{L}$$

Figure 1: Rules for the non-communicating asynchronous parallel operator $\_|\_$

on properties of the sequents forming the path.

This abstract describes the ongoing research aimed to the definition of a sequent based infinitary proof system for proving qMμ properties of processed described by a class of well behaved process calculi ([12]). The first notion that needs to be revisited is the meaning of sequents. As the game-based semantics of the logic qMμ assigns a value $[\![p{:}F]\!] \in [0,1]$ to a pair of process p and formula F, the standard classical (conjunction and disjunction) interpretation for the commas in the sequents needs to be changed. We found that interpreting left commas as multiplications, many properties of interest can be expressed by sequents. Therefore fixed an interpretation ρ for the process variables we define the meaning of the sequent

$$p_1{:}F_1, ..., p_n{:}F_n \vdash q_1{:}G_1, ..., q_m{:}G_m$$

as

$$[\![p_1{:}F_1]\!]_\rho \cdot ... \cdot [\![p_n{:}F_n]\!]_\rho \leq [\![q_1{:}G_1]\!]_\rho \odot ... \odot [\![q_m{:}G_m]\!]_\rho$$

where $\odot$ is the De Morgan dual of $\cdot$ under the involution $\neg x = 1 - x$, i.e. $x \odot y = \neg(\neg x \cdot \neg y)$. A sequent is valid if the inequality holds for every interpretation. Observe that adopting this semantics, the familiar left and right contraction rules become unsound.

The following is a simple but surprisingly not trivial example of valid sequent only involving the non-communicating asynchronous parallel ($\_|\_$) term constructor (whose probabilistic operational semantics is given, as in [12], by the rules of Figure 1) and the qMμ formula $\mu X.\,[a]\,X$ whose meaning has been discussed above:

$$x{:}\mu X.\,[a]\,X, y{:}\mu X.\,[a]\,X \vdash x|y{:}\mu X.\,[a]\,X$$

This is one of many examples that motivated the choice of multiplication (and its dual) as interpretation of the commas. It seems that this choice fits quite well with the operational semantics associated with any process operator given in the general SOS format for probabilistic process operators of [12]. As a brief remark note that the quantitative entailment expressed by the above valid sequent, implies a qualitative one: if the value of all assertions on the left part of the sequent is 1 then the assertion on the right have value 1; in other words when reasoning qualitatively, the meaning of the commas in the sequents collapses to the classical one. This seems to support the choice of working with a quantitative interpretation for the sequents, as it can express more refined properties.

Our sequent calculus proof system includes a number of interesting features due to the quantitative interpretation of sequents. As already observed, contraction is unsound, so our sequent calculus implements an *affine* logic. More significantly, the inclusion of probabilistic process operators means that it is necessary to include proof rules for dealing with probability distributions over processes. Let us consider for example the following right rule that operates with the probability distribution $p +_{\frac{1}{3}} q$ which assigns probability $\frac{1}{3}$ to p and $\frac{2}{3}$ to q:

$$\frac{\Gamma \vdash \Delta, p{:}F \qquad \Gamma \vdash \Delta, q{:}F}{\Gamma \vdash \Delta, p +_{\frac{1}{3}} q{:}F} +\mathcal{R}$$

Roughly speaking, this rule must be understood bottom-up, as corresponding to the probabilistic choice associated with the distribution $p +_{\frac{1}{3}} q$. This amounts to consider the two premises of the end sequent as having different probabilistic weights. Therefore the end sequent of the rule $+_{\mathcal{R}}$ can be considered as a probabilistic node in a derivation tree and the whole derivation tree, under this interpretation, is a Markov Process, i.e. some of its nodes are probabilistic, and some other nodes offer genuine choices. As in [9, 10, 11, 8] our proof system is based on infinite (non-well-founded) derivations, and so requires a global proof condition to impose on proofs to ensure soundness. To formulate the proof condition we first define the notion of *Markov path* through a proof. This is like a path through the derivation tree, but it builds in probabilistic choices at probabilistic proof rules, giving rise to a Markov chain of paths through the derivation. The *proof condition* is then the following: for every Markov path through the proof, the set of paths in the Markov path having a left μ-trace or a right ν-trace has measure 1. The notion of (μ, ν)-trace is fairly standard, as it coincides with analogous definitions appearing in [11, 8, 9, 10]. Our main result is the soundness of this proof system. This is nontrivial and the proof involves novel constructions on Two Player Stochastic games, in particular novel notions of product and coproduct game related to the quantitative interpretation of sequents.

In addition to the soundness result, we have examples of derivations of nontrivial properties in the system; for instance, the example presented above.

As a final remark, observe that in any practical use of the proof system, some finite description for the infinite derivation must be provided. Among the classes of infinite derivations finitely describable by different structures, the simplest is that of regular derivations, those representable by cyclic graphs. An interesting non-trivial problem is then associated with the decidability of the proof condition on derivations represented as cyclic graphs.

References

[1] R. Segala. *Modeling and Verification of Randomized Distributed Real-Time Systems*, PhD thesis, Laboratory for Computer Science, Massachusetts Institute of Technology, 1995.

[2] M. Huth, M. Kwiatkowska. *Quantitative Analysis and Model Checking*, Proceedings of the 12th Annual IEEE Symposium on Logic in Computer Science, 1997.

[3] C. Morgan, A. McIver. *A Probabilistic Temporal Calculus Based on Expectations*, In Lindsay Groves and Steve Reeves, editors, Proc. Formal Methods, 1997.

[4] C. Morgan, A. McIver. *Results on the quantitative μ-calculus $qM\mu$*, ACM Trans. Comput. Logic, Volume 8, Issue 1, 2007.

[5] C. Stirling. *Modal and Temporal Properties of Processes* , Springer Texts in Computer Science, 2001.

[6] A. Simpson. *Compositionality via cut-elimination: Hennessy-Milner logic for an arbitrary GSOS*, Proceedings of the 10th Annual IEEE Symposium on Logic in Computer Science, 1995.

[7] A. Simpson. *Sequent Calculi for Process Verification: Hennessy-Milner Logic for an Arbitrary GSOS*, Journal of Logic and Algebraic Programming, 60-61:287-322, 2004.

[8] M. Dam, D. Gurov. *Compositional Verification of CCS Processes*, Proceedings of the Third international Conference on Perspectives of System informatics, 1999.

[9] C. Dax, M. Hofmann, M. Lange. *A Proof System for the Linear Time μ-Calculus*, Lecture Notes in Computer Science, Volume 4337, 2006.

[10] T. Studer. *On the proof theory of the Modal μ-Calculus*, Studia Logica, Volume 89, Number 3, 2007.

[11] J. Brotherston, A. Simpson. *Complete Sequent Calculi for Induction and Infinite Descent*, Proceedings of the 22nd Annual IEEE Symposium on Logic in Computer Science, 2007.

[12] F. Bartels. *GSOS for Probabilistic Transition Systems*, Electronic Notes in Theoretical Computer Science, Volume 65, Issue 1, 2002.