

A Weakest Pre–Expectation Semantics for Mixed–Sign Expectations

Benjamin Lucien Kaminski and Joost-Pieter Katoen
 Software Modeling and Verification Group
 RWTH Aachen University, Aachen, Germany

I. INTRODUCTION

We consider probabilistic programs of the form

$$C \longrightarrow x := \text{expr} \mid C; C \mid \text{if}(\xi)\{C\}\text{else}\{C\} \\ \mid \text{while}(\xi)\{C\},$$

where ξ is a *probabilistic guard* which behaves as follows: Let Σ denote the set of program states, i.e. mappings from program variables to valuations. If the computation is currently in a state $\sigma \in \Sigma$ then ξ evaluates to `true` with probability $\llbracket \xi \rrbracket(\sigma)$ and to `false` with probability $1 - \llbracket \xi \rrbracket(\sigma) = \llbracket \neg \xi \rrbracket(\sigma)$. For example, the probabilistic guard

$$[x \text{ is even}] \cdot (2/3\langle \text{true} \rangle + 1/3\langle \text{false} \rangle) + [x \text{ is odd}] \cdot \langle \text{false} \rangle$$

evaluates to `false` with probability $1/3$ if in the current program state x is even and likewise with probability 1 if x is odd.

Given a program C , a random variable f mapping program states to reals, and an initial state σ , we are now interested in the following question: *What is the expected value of f after termination of C on input σ ?* This expected value is referred to as the *pre–expectation* of C with respect to *post–expectation* f . For example, what is the pre–expectation of

$$C_{geo} \triangleright \text{while}(1/2\langle \text{true} \rangle + 1/2\langle \text{false} \rangle)\{x := x + 1\}$$

with respect to post–expectation $f = x$? The answer is $x + 1$, since C_{geo} terminates with probability 1 and the execution of C_{geo} increases x on average by 1 .

If we stay in the realm of *non–negative* expectations, i.e. random variables $f \in \mathbb{E}_{\geq 0}^{\infty} := \{f \mid f: \Sigma \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}\}$, the situation is fairly well understood, e.g. by means of the weakest pre–expectation calculus [1], [2], [3], [?]. It gives a meaning to each program C by means of a transformer $\text{wp}[C]: \mathbb{E}_{\geq 0}^{\infty} \rightarrow \mathbb{E}_{\geq 0}^{\infty}$ that maps any post–expectation $f \in \mathbb{E}_{\geq 0}^{\infty}$ to a pre–expectation $\text{wp}[C](f) \in \mathbb{E}_{\geq 0}^{\infty}$, such that

$$\forall \sigma \in \Sigma: \text{wp}[C](f)(\sigma) = \mathbf{E}_{\llbracket C \rrbracket(\sigma)}(f),$$

where $\mathbf{E}_{\mu}(h)$ denotes the expected value of random variable h with respect to distribution μ and $\llbracket C \rrbracket(\sigma)$ denotes the distribution obtained by executing program C on input σ . Hence $\text{wp}[C](f)(\sigma)$ is the expected value of f after termination of program C executed on input σ .

What happens, however, if we drop the requirement of f being non–negative? Suppose we want to reason about C_{geo} with respect to post–expectation $f = (-2)^x/x$. Is

$\text{wp}[x := 1; C_{geo}](f)$ well–defined in that case? Intuitively, the pre–expectation is given by the series

$$-1 + \frac{1}{2} - \frac{1}{3} + \frac{1}{4} - \frac{1}{5} + \frac{1}{6} - \frac{1}{7} + \dots = -\ln(2).$$

The problem with this series is that—though it converges—it does not converge absolutely and thus—as a consequence of the well–known Riemann Series Theorem [4]—by reordering of the summands we can make the above series converge to *any real value* or even make it *tend to $+\infty$ or $-\infty$* . For representing expected values, however, the ordering of the summands should not matter as there is no natural ordering in which the mass of an expected value should be accumulated. For that reason, the expected value of a mixed–sign random variable f should exist if and only if f is integrable, i.e. if and only if the expected value of $|f|$ is finite. We therefore propose a weakest pre–expectation semantics that internally keeps track of the integrability of f , while still being well–defined for any program with respect to any mixed–sign post–expectation f .

II. OUR PROPOSAL

A. Integrability Witnessing Expectations

To keep track of the integrability of expectations, we accompany each mixed–sign expectation $f \in \mathbb{E}^* := \{f \mid f: \Sigma \rightarrow \mathbb{R}\}$ with a *non–negative* expectation $g \in \mathbb{E}_{\geq 0}^{\infty}$, such that $|f| \leq g$. We call such a pair (f, g) an *integrability witnessing pair*.

The intuition behind an integrability witnessing pair (f, g) is the following: If $\text{wp}[C](g)(\sigma) < \infty$, then the pre–expectation of C in σ with respect to $f \in \mathbb{E}^*$ should exist as f is integrable with respect to distribution $\llbracket C \rrbracket(\sigma)$. If, however, $\text{wp}[C](g)(\sigma) = \infty$ then we should not care about the pre–expectation of C in σ with respect to f since it should not be defined because in that case f is not integrable. This intuition leads to a quasi–order \lesssim on integrability witnessing expectations given by $(f, g) \lesssim (f', g')$ iff for all $\sigma \in \Sigma$,

$$g'(\sigma) \neq \infty \text{ implies } f(\sigma) \leq f'(\sigma) \text{ and } g(\sigma) \leq g'(\sigma).$$

Notice that, indeed, \lesssim is not a partial order as it is *not antisymmetric*: we can have two integrability witnessing pairs $(f, g) \neq (f', g')$ with $(f, g) \lesssim (f', g')$ and $(f, g) \gtrsim (f', g')$. This is the case if for some state $\sigma \in \Sigma$ we have $g(\sigma) = \infty = g'(\sigma)$, but $f(\sigma) \neq f'(\sigma)$.

On the other hand, two integrability witnessing pairs (f, g) and (f', g') , for which $f(\sigma) \neq f'(\sigma)$ holds only for those states in which $g(\sigma) = \infty = g'(\sigma)$, should really be

considered *equivalent*, even though they are not equal. This is because for states σ in which $g(\sigma) = \infty = g'(\sigma)$, the evaluations of $f(\sigma)$ and $f'(\sigma)$ should be irrelevant since integrability is not ensured. Consequently, we need a notion of equivalence of integrability witnessing pairs, given by $\approx = \lesssim \cap \gtrsim$. We denote the \approx -equivalence class of an integrability witnessing pair (f, g) by $\langle f, g \rangle$ and call such an equivalence class an *integrability witnessing expectation*. We denote by \mathbb{IE} the set of integrability witnessing expectations.

There is a canonical [?] partial order \sqsubseteq on \mathbb{IE} given by $\langle f_1, g_1 \rangle \sqsubseteq \langle f_2, g_2 \rangle$ iff $(f_1, g_1) \lesssim (f_2, g_2)$. As for an intuitive interpretation of this partial order, we note that if $\langle f_1, g_1 \rangle \sqsubseteq \langle f_2, g_2 \rangle$ holds, then we have $f'_1(\sigma) = f_1(\sigma) \leq f_2(\sigma) = f'_2(\sigma)$ for all $(f'_1, g'_1) \in \langle f_1, g_1 \rangle$, $(f'_2, g'_2) \in \langle f_2, g_2 \rangle$, and all states in which $g_2(\sigma) \neq \infty$ holds. Thus if integrability in σ is ensured, the first components compare in σ , which is the comparison we are mainly interested in.

The partial order \sqsubseteq on \mathbb{IE} expectations is complete in the sense that every *non-empty* subset has a supremum. However, \mathbb{IE} has *no least element*; in particular $\langle \mathbf{0}, \mathbf{0} \rangle$ is not a least element of \mathbb{IE} . This fact prevents us from applying the Kleene Fixed Point Theorem to ensure existence of least fixed points for defining a \mathbf{wp} -calculus acting on \mathbb{IE} .

B. Mixed-Sign Weakest Pre-Expectations

We now propose a weakest pre-expectation transformer $\widetilde{\mathbf{wp}}[C]: \mathbb{IE} \rightarrow \mathbb{IE}$ defined compositionally by induction on the structure of C , see **Table I**. Let us shortly go over these definitions: $\widetilde{\mathbf{wp}}[x := E] \langle f, g \rangle$ takes a representative $(f, g) \in \langle f, g \rangle$, performs the assignment $x := E$ on both components to obtain $(f[x/E], g[x/E])$ and then returns the according equivalence class $\langle f[x/E], g[x/E] \rangle$. Notice that assignments preserve \approx -equivalence, so doing the update on the representative is a sound and sufficient course of action.

$\widetilde{\mathbf{wp}}[C_1; C_2] \langle f, g \rangle$ obtains a pre-expectation for the program $C_1; C_2$ by applying $\widetilde{\mathbf{wp}}[C_1]$ to the intermediate integrability witnessing expectation obtained from $\widetilde{\mathbf{wp}}[C_2] \langle f, g \rangle$.

TABLE I

DEFINITIONS FOR $\widetilde{\mathbf{wp}}$. $f[x/expr]$ DENOTES $\lambda\sigma \bullet f(\sigma[x \mapsto \sigma(expr)])$, WHERE $\sigma[x \mapsto \sigma(expr)]$ IS THE PROGRAM STATE OBTAINED BY UPDATING IN σ THE VALUE OF x TO $\sigma(expr)$. FOR THE DEFINITION OF \mathbf{while} , $C'^{\xi}F_{\langle f, g \rangle}^n$ DENOTES APPLYING $C'^{\xi}F_{\langle f, g \rangle}$ n -FOLD TO ITS ARGUMENT.

C	$\widetilde{\mathbf{wp}}[C] \langle f, g \rangle$
$x := expr$	$\langle f[x/expr], g[x/expr] \rangle$
$C_1; C_2$	$\widetilde{\mathbf{wp}}[C_1](\widetilde{\mathbf{wp}}[C_2] \langle f, g \rangle)$
$\mathbf{if}(\xi) \{C_1\} \mathbf{else} \{C_2\}$	$\llbracket \xi \rrbracket \cdot \widetilde{\mathbf{wp}}[C_1] \langle f, g \rangle$ $+ \llbracket \neg\xi \rrbracket \cdot \widetilde{\mathbf{wp}}[C_2] \langle f, g \rangle$
$\mathbf{while}(\xi) \{C'\}$	$\lim_{n \rightarrow \omega} C'^{\xi}F_{\langle f, g \rangle}^n \langle \mathbf{0}, \mathbf{0} \rangle$
<hr/>	
	$C'^{\xi}F_{\langle f, g \rangle} \langle X, Y \rangle = \llbracket \neg\xi \rrbracket \cdot \langle f, g \rangle + \llbracket \xi \rrbracket \cdot \widetilde{\mathbf{wp}}[C] \langle X, Y \rangle$

$\widetilde{\mathbf{wp}}[\mathbf{if}(\xi) \{C_1\} \mathbf{else} \{C_2\}] \langle f, g \rangle$ weights $\widetilde{\mathbf{wp}}[C_1] \langle f, g \rangle$ and $\widetilde{\mathbf{wp}}[C_2] \langle f, g \rangle$ according to the probability of the guard ξ evaluating to true and false. Here, $g' \cdot \langle f, g \rangle := \langle g' \cdot f, g' \cdot g \rangle$ and $\langle f', g' \rangle + \langle f, g \rangle := \langle f' + f, g' + g \rangle$.

$\widetilde{\mathbf{wp}}$ of $\mathbf{while}(\xi) \{C'\}$ is defined as the limit of iteratively applying the characteristic functional

$$C'^{\xi}F_{\langle f, g \rangle} \langle X, Y \rangle = \llbracket \neg\xi \rrbracket \cdot \langle f, g \rangle + \llbracket \xi \rrbracket \cdot \widetilde{\mathbf{wp}}[C'] \langle X, Y \rangle,$$

to $\langle \mathbf{0}, \mathbf{0} \rangle$. This limit can be shown to be existent and unique. Recall that $\langle \mathbf{0}, \mathbf{0} \rangle$ is not the least element of \mathbb{IE} and thus the Kleene Fixed Point Theorem *cannot be applied* to deduct the existence of this limit. Instead, the core idea for proving this fact is adopted from a well-known proof proving that every absolutely convergent series is also convergent.

C. Soundness, Monotonicity, and Loop Invariants

Our mixed-sign weakest pre-expectation transformer is sound, meaning that if we can establish $\widetilde{\mathbf{wp}}[C] \langle f, |f| \rangle = \langle f', g' \rangle$ with $g'(\sigma) < \infty$, then $f'(\sigma)$ is in fact the expected value of f after termination of C on initial state σ .

Furthermore, our transformer is monotonic, meaning that

$$\langle f, g \rangle \sqsubseteq \langle f', g' \rangle \text{ implies } \widetilde{\mathbf{wp}}[C] \langle f, g \rangle \sqsubseteq \widetilde{\mathbf{wp}}[C] \langle f', g' \rangle.$$

Since monotonicity of a weakest pre-expectation transformer is as vital to reasoning about programs as the consequence rule is to Hoare logic, this is a key feature.

Lastly, we note that we can give a proof rule for loops based on loop invariants. Going into the details of this rule would go beyond the scope of this short abstract but suffice it to say that using our proof rule we can for instance prove that the pre-expectation of

$$C_{Op} \triangleright \Phi := \Phi + 1; \mathbf{while} \ (1/2) \ \{ \Phi := \Phi - 3 \}$$

with respect to post-expectation Φ is *at most* $\Phi - 2$. In other words: C_{Op} *decreases* Φ at least by 2.

Notice that such an analysis would neither be possible using the deduction rules of PPDL [1] nor the invariant-based approach of McIver & Morgan's \mathbf{wp} -calculus [2] off-the-shelf. Instead, a tailor made argument would be needed for reasoning about the mixed-sign post-expectation Φ .

ACKNOWLEDGMENT

This work was supported by the Excellence Initiative of the German federal and state government and by the CDZ project CAP (GZ 1023).

REFERENCES

- [1] D. Kozen, "A probabilistic PDL," *J. Comput. Syst. Sci.*, vol. 30, no. 2, pp. 162–178, 1985.
- [2] A. McIver and C. Morgan, *Abstraction, Refinement and Proof for Probabilistic Systems*. Springer, 2004.
- [3] F. Gretz, J. Katoen, and A. McIver, "Operational versus Weakest Pre-Expectation Semantics for the Probabilistic Guarded Command Language," *Performance Evaluation*, vol. 73, pp. 110–132, 2014.
- [4] B. Riemann, *Ueber die Darstellbarkeit einer Function durch eine trigonometrische Reihe*. Königliche Gesellschaft der Wissenschaften zu Göttingen, 1867.