

Quantitative Analysis of Probabilistic Pushdown Automata: Expectations and Variances (Extended Abstract)

Javier Esparza

Institute for Formal Methods in Computer Science,
University of Stuttgart,
Universität str. 38, 70569 Stuttgart, Germany.
esparza@informatik.uni-stuttgart.de

Antonín Kučera*

Faculty of Informatics, Masaryk University,
Botanická 68a, 60200 Brno,
Czech Republic.
tony@fi.muni.cz

Richard Mayr

Department of Computer Science, Campus Box 8207,
North Carolina State University, Raleigh, NC 27695-8207, USA
mayr@csc.ncsu.edu

Abstract

Probabilistic pushdown automata (pPDA) have been identified as a natural model for probabilistic programs with recursive procedure calls. Previous works considered the decidability and complexity of the model-checking problem for pPDA and various probabilistic temporal logics. In this paper we concentrate on computing the expected values and variances of various random variables defined over runs of a given probabilistic pushdown automaton. In particular, we show how to compute the expected accumulated reward and the expected gain for certain classes of reward functions. Using these results, we show how to analyze various quantitative properties of pPDA that are not expressible in conventional probabilistic temporal logics.

1. Introduction

Pushdown automata (or recursive state machines) are a natural model for sequential programs with recursive procedure calls [5, 2, 10, 4]. Recent papers are, e.g., [1, 8, 3]. Recently, probabilistic aspects of such programs have been taken into account, and several papers have studied in detail the decidability and complexity of model-checking both linear and branching-time probabilistic temporal logics for probabilistic pushdown automata (pPDA) [6, 11, 13, 12]. Using these results it is possible to decide if, say, the proba-

bility that a pPDA terminates is at least 0.98, even though its associated Markov chain may have infinitely many states.

However, we are often interested not only in the probability of termination, but also in the expected termination time, defined over the runs of the program that terminate. Moreover, we may wish to obtain some more information about the probability distribution of the termination time, in particular its variance.

In this paper we address this problem. We obtain some generic results about computing expectations and variance of reward functions, and apply them to three fundamental problems: termination time, renewal time (time between two visits to a given control state), and stack length. All three are obviously relevant for the design and performance analysis of probabilistic programs.

The semantic of a pPDA is a possibly infinite Markov chain whose states are *configurations*, i.e., pairs consisting of a control state and a stack content. A trajectory in this chain is called a *run*. We consider reward functions that assign a reward to every configuration of the pPDA. When the pPDA enters a configuration, it collects its associated reward. Given a reward function, our goal is to compute (a) the expectation of the reward accumulated during a finite run. For infinite runs this reward is usually infinite, and so our goal is to compute (b) the expectation of the *gain*, defined as the average reward earned per transition. (Formally, the gain is the random variable that assigns to an infinite run the limit, as n approaches infinity, of the reward accumulated during the first n transitions divided by n .) We are also interested in the variances of the corresponding random variables.

* Supported by the Alexander von Humboldt Foundation and by the research centre Institute for Theoretical Computer Science (ITI), project No. 1M0021620808.

The first part of the paper deals with reward functions that only depend on the control state, which we call *simple* functions. In a first step (Section 3) we show how to compute the expected accumulated reward and its variance for finite runs of the pPDA starting at a configuration of the form pX (control state p and stack content X , of length 1), and ending at a configuration $q\varepsilon$ (control state q and empty stack). The termination time corresponds to the simple reward function that assigns 1 to each control state.

In a second and more involved step (Section 4), we use the results of Section 3 to compute the expected value of the gain for simple reward functions. We apply the technique developed in [11] (see also [12]) to solve the LTL model-checking problem. This technique consists of computing a finite Markov chain which records the *minima* of an infinite run of the pPDA, defined as the configurations at which the run reaches for the first time a certain length, but in such a way that in the future it never goes below that length. A transition of this finite chain, which in the rest of the introduction we call a ‘jump’, corresponds to the (finite) sequence of transitions carried out by the pPDA in order to move from one minimum to the next. We define a modification of this chain, and show how to compute the expected accumulated reward for each jump. Then, we show how to use this information to obtain the expected gain. In fact, the technique has a limitation: It only works if the expected reward for each jump is finite. Fortunately, the case of infinite expected reward per jump is rather pathological, and should correspond to a design fault in most cases. The renewal time for a set P of control states, i.e., the expected time between two visits to configurations whose control states belong to P , corresponds to the simple function which assigns 1 to the control states of P , and 0 to the others. More precisely, the renewal time is the inverse of the gain of this function.

In the second part of the paper, we extend the techniques of the first part to reward functions that depend on the control state and (linearly) on the length of the stack, which we call *linear* reward functions. Again, this part is divided into two steps. The first step (Section 5) extends the results of Section 3, and is rather straightforward. The second step (Section 6) requires some more care. The reason is that, in general, the reward obtained when executing a jump now may depend on the past, i.e., on all the previous jumps. Fortunately, we are able to show that when this is the case the expected reward is infinite. So we only need to worry about the memoryless case, and we are thus able to apply the results of Section 4. The expected average stack length corresponds to the linear function that assigns to a configuration the length of its stack content.

While the expectation of the average stack length provides some useful information, there is a more interesting parameter, namely the *maximal* stack length, i.e., the random variable that assigns to a run the maximal stack length

reached along it. In Section 6.1 we provide some partial results on this variable. More precisely, we show how to compute the probability that the maximal stack length remains bounded, and the probability, for each value n , that the maximal stack length is at most n .

2. Preliminaries

Let \mathbb{N} , \mathbb{N}_0 , \mathbb{Q} , \mathbb{Q}^+ , \mathbb{R} , and \mathbb{R}^+ denote the sets of positive integers, non-negative integers, rational numbers, non-negative rational numbers, real numbers, and non-negative real numbers, respectively. We also use \mathbb{Q}_ω^+ and \mathbb{R}_ω^+ to denote the sets $\mathbb{Q}^+ \cup \{\omega\}$ and $\mathbb{R}^+ \cup \{\omega\}$, respectively, where $\omega \notin \mathbb{Q} \cup \mathbb{R}$ is a special symbol. We stipulate that $c + \omega = c + c = \omega + \omega = \omega$ and $c/\omega = 0$ for each $c \in \mathbb{R}^+$, and $c \cdot \omega = \omega$ for each positive $c \in \mathbb{R}^+$.

For a given alphabet Σ , the symbol Σ^* denotes the set of all finite words over Σ . The length of a given $w \in \Sigma^*$ is denoted by $|w|$, and the empty word is denoted by ε .

Definition 2.1. A (fully) probabilistic transition system is a triple $\mathcal{T} = (S, \rightarrow, Prob)$ where S is a finite or countably infinite set of states, $\rightarrow \subseteq S \times S$ is a transition relation, and $Prob$ is a function which to each transition $s \rightarrow t$ of \mathcal{T} assigns its probability $Prob(s \rightarrow t) \in (0, 1]$ so that for every $s \in S$ we have $\sum_{s \rightarrow t} Prob(s \rightarrow t) \in \{0, 1\}$. (The sum above can be 0 if s does not have any outgoing transitions.)

In the rest of this paper we also write $s \xrightarrow{x} t$ instead of $Prob(s \rightarrow t) = x$. A *path* in \mathcal{T} is a finite or infinite sequence $w = s_0, s_1, \dots$ of states such that $s_i \rightarrow s_{i+1}$ for every i . We also use $w(i)$ to denote the state s_i of w (by writing $w(i) = s$ we implicitly impose the condition that the length of w is at least $i + 1$). The prefix s_0, s_1, \dots, s_i of w is denoted by w^i . A *run* is a maximal path, i.e., a path which cannot be prolonged. The sets of all finite paths and all runs of \mathcal{T} are denoted $FPath$ and Run , respectively. Similarly, the sets of all finite paths and runs that start in a given $s \in S$ are denoted $FPath(s)$ and $Run(s)$, respectively.

A strongly connected component (SCC) of \mathcal{T} is a subset $C \subseteq S$ where for all $x, y \in C$ we have that $x \rightarrow^* y$. A SCC is a bottom SCC (BSCC) if no other SCC is reachable from it. It follows that all BSCC are maximal SCC.

A *reward function* is a function $f: S \rightarrow \mathbb{Q}^+$ that assigns to a state s a reward $f(s)$. We assume that when a process enters a configuration s it collects the reward $f(s)$. Given a reward function f , we extend it to a function $F: FPath \rightarrow \mathbb{Q}^+$ by $F(s_0, \dots, s_n) = \sum_{i=1}^n f(s_i)$. Hence, F assigns to each path its *accumulated reward*. Note that $f(s_0)$ is not included in the sum.

Each $w \in FPath$ determines a *basic cylinder* $Run(w)$ which consists of all runs that start with w . To every $s \in S$ we associate the probabilistic space $(Run(s), \mathcal{F}, \mathcal{P})$ where

\mathcal{F} is the σ -field generated by all basic cylinders $Run(w)$ where w starts with s , and $\mathcal{P} : \mathcal{F} \rightarrow [0, 1]$ is the unique probability function such that $\mathcal{P}(Run(w)) = \prod_{i=0}^{m-1} x_i$ where $w = s_0, \dots, s_m$ and $s_i \xrightarrow{x_i} s_{i+1}$ for every $0 \leq i < m$ (if $m = 0$, we put $\mathcal{P}(Run(w)) = 1$).

Definition 2.2. A probabilistic pushdown automaton (pPDA) is a tuple $\Delta = (Q, \Gamma, \delta, Prob)$ where Q is a finite set of control states, Γ is a finite stack alphabet, $\delta \subseteq Q \times \Gamma \times Q \times \Gamma^*$ is a transition relation such that whenever $(p, X, q, \alpha) \in \delta$, then $|\alpha| \leq 2$, and $Prob$ is a function which to each transition $pX \rightarrow q\alpha$ assigns a rational probability $Prob(pX \rightarrow q\alpha) \in (0, 1]$ so that for all $p \in Q$ and $X \in \Gamma$ we have that $\sum_{pX \rightarrow q\alpha} Prob(pX \rightarrow q\alpha) \in \{0, 1\}$.

In the rest of this paper we adopt a more intuitive notation, writing $pX \rightarrow q\alpha$ instead of $(p, X, q, \alpha) \in \delta$ and $pX \xrightarrow{x} q\alpha$ instead of $Prob(pX \rightarrow q\alpha) = x$. The set $Q \times \Gamma^*$ of all configurations of Δ is denoted by $\mathcal{C}(\Delta)$. Given a configuration $pX\alpha$ of Δ , we call pX the *head* and α the *tail* of $pX\alpha$.

To Δ we associate the probabilistic transition system \mathcal{T}_Δ where $\mathcal{C}(\Delta)$ is the set of states and the probabilistic transition relation is determined by $pX\beta \xrightarrow{x} q\alpha\beta$ iff $pX \xrightarrow{x} q\alpha$ is a transition of Δ .

Now we recall some known results which will be used in the following sections. Let us fix a pPDA $\Delta = (Q, \Gamma, \delta, Prob)$. For all $p, q \in Q$ and $X \in \Gamma$, the symbol $[pXq]$ denotes the probability that a run initiated in pX hits $q\varepsilon$. More precisely, $[pXq] = \mathcal{P}(w \in Run(pX) \mid \exists i \in \mathbb{N} : w(i) = q\varepsilon)$. It has been shown in [11] (see also [13]) that there effectively exists a finite system of recursive quadratic equations with variables of the form $\langle pXq \rangle$ (i.e., there is a variable $\langle pXq \rangle$ for all $p, q \in Q$ and $X \in \Gamma$) such that the family of all $[pXq]$ probabilities forms the least solution of this system of equations with respect to component-wise ordering. Since all terms in these equations are built using just summation, multiplication, and rational constants, each of the $[pXq]$ probabilities is effectively expressible in $(\mathbb{R}, +, *, \leq)$ in the following sense: there effectively exists a formula Φ of $(\mathbb{R}, +, *, \leq)$ with one free variable x such that $\Phi[x/c]$ holds iff $c = [pXq]$. Hence, for all $\sim \in \{<, >, \leq, \geq, =\}$ and all rational constants ϱ one can decide if $[pXq] \sim \varrho$ simply by checking if the formula $\exists x(\Phi \wedge x \sim \varrho)$ is valid or not. One can also compute the value $[pXq]$ up to an arbitrarily small non-zero error (for example, by a simple binary search). Since the value of $[pXq]$ can be irrational [13], it cannot be computed precisely in general.

The decidability of $(\mathbb{R}, +, *, \leq)$ is due to Tarski [15], and some fragments of $(\mathbb{R}, +, *, \leq)$ are known to have a relatively reasonable complexity. For example, the existential (and hence also the universal) fragment is in **PSPACE** [7], and each fragment with a bounded alternation depth of

quantifiers is in **EXPTIME** [14]. We use these results to estimate the complexity of our algorithms.

In this paper we show that a number of quantitative features of pPDA are effectively expressible in $(\mathbb{R}, +, *, \leq)$ (we just say that a given random variable, or its expected value, or its variance, is “expressible”). This also applies to features which can be infinite. For example, in the next section we consider the expected length of a terminating computation (i.e., we express the average length of the subset of all finite runs) which can be infinite even if a given pPDA configuration terminates with probability one. In that case, the associated formula Φ does not hold for any $c \in \mathbb{R}_0^+$. So, the problem whether the expected time is finite or infinite can be decided by checking whether the formula $\exists x \geq 0. \Phi$ is valid or invalid, respectively. If the time is finite, we can use the formula Φ in the ways described above.

Let us note that once a certain feature is shown expressible, it can be used to define other features which then become expressible as well. In some cases, the structure of the resulting formula is quite complicated. Our complexity results are based on evaluating the size and structure of these formulae. The size remains typically polynomial in the size of the original pPDA. Sometimes we obtain a formula of the existential fragment, and sometimes we need to nest the quantifiers to some fixed depth. Therefore, typical upper bounds presented in this paper are **PSPACE** and **EXPTIME**.

In the following sections we also use $[qY\uparrow]$ to denote the probability that a run initiated in qY is infinite. Observe that since $[qY\uparrow] = 1 - \sum_{p \in Q} [qYp]$, this probability is expressible. Finally, let us note that since the variance of a given random variable Y is equal to $E(Y^2) - (E(Y))^2$, the variance of Y is expressible if both $E(Y)$ and $E(Y^2)$ are expressible (and this is what we usually prove).

3. Simple reward functions: Expected accumulated reward

In this section, let us fix a pPDA $\Delta = (Q, \Gamma, \delta, Prob)$ such that for each transition $pX \xrightarrow{x} q\alpha$ we have that $|\alpha| \in \{0, 2\}$. This assumption is not restrictive (for every PDA there is an equivalent one in this form, up to transition graph isomorphism), and becomes particularly useful in this section (otherwise, the systems of equations considered in Theorem 3.1 and Theorem 3.4 would take even less readable form).

Since the probabilities $[pXq]$ are known to be expressible (see Section 2), we assume that they are already “known” and can safely be used in expressions for other variables (see the discussion in Section 2).

In this section we consider *simple* reward functions. A reward function f is simple if $f(p\alpha)$ only depends on p . For the rest of the section we fix a simple reward function f , and write $f(p)$ instead of $f(p\alpha)$.

For all $p, q \in Q$ and $X \in \Gamma$, we compute the conditional expectation of the reward accumulated by the pPDA along a path, under the condition that the path starts at pX and ends at $q\varepsilon$. We also compute the corresponding conditional variance.

Consider the probabilistic space $(Run(pX), \mathcal{F}, \mathcal{P})$. For each $q \in Q$, let $Run(pXq)$ be the set of all $w \in Run(pX)$ such that $w(i) = q\varepsilon$ for some $i \in \mathbb{N}$. For a given pX we consider only those $q \in Q$ such that $[pXq] > 0$. The other control states are irrelevant (and the notions introduced below do not make sense if $[pXq] = 0$).

We define a random variable R_{pXq} over $Run(pX)$ in the following way:

$$R_{pXq}(w) = \begin{cases} 0 & \text{if } w \notin Run(pXq) \\ F(w^\ell) & \text{if } w(\ell) = q\varepsilon \end{cases}$$

Then $E(R_{pXq} \mid Run(pXq))$ is the conditional expected accumulated reward from pX to $q\varepsilon$, under the condition that $q\varepsilon$ is reached. From now on we write $[E(pXq)]$ instead of $E(R_{pXq} \mid Run(pXq))$.

We show that $[E(pXq)]$ can be computed as the minimal solution of a system of linear equations over \mathbb{R}_ω^+ . Let $\mathcal{V} := \{\langle E(pXq) \rangle \mid p, q \in Q, X \in \Gamma, [pXq] > 0\}$ be a set of variables over \mathbb{R}_ω . That is, for every $[E(pXq)]$ there is the associated variable $\langle E(pXq) \rangle$. Consider the following system of recursive equations: $\langle E(pXq) \rangle = 0$ for $[pXq] = 0$. Otherwise,

$$\langle E(pXq) \rangle = \frac{1}{[pXq]} \left(\sum_{pX \xrightarrow{x} q\varepsilon} x \cdot f(q) + \sum_{pX \xrightarrow{x} rYZ} x \cdot K_{pX, rYZ} \right)$$

where the term $K_{pX, rYZ}$ is given by

$$\sum_{s \in Q} [rYs][sZq](f(r) + \langle E(rYs) \rangle + \langle E(sZq) \rangle)$$

If $[rYs]$ or $[sZq]$ is zero, then the corresponding summand of $K_{pX, rYZ}$ is removed (thus we avoid problems with undefined expressions like $0 \cdot \omega$; note that, e.g., $[E(rYs)]$ can be ω).

Theorem 3.1. *The tuple of all $[E(pXq)]$ values is exactly the least solution of the above system of equations in \mathbb{R}_ω^+ with respect to component-wise ordering.*

Proof sketch. The system of equations determines a unique operator $\mathcal{F} : (\mathbb{R}_\omega^+)^{|\mathcal{V}|} \rightarrow (\mathbb{R}_\omega^+)^{|\mathcal{V}|}$ where $\mathcal{F}(t)$ is the tuple of values obtained by evaluating the right-hand sides of the equations where each variable of \mathcal{V} is substituted with its associated value in t . Since \mathcal{F} is monotonic and continuous, \mathcal{F} has the least fixed-point μ . We show that μ is exactly the tuple of all $[E(pXq)]$ values.

We first prove that μ is smaller than or equal to the tuple of $[E(pXq)]$ values. We show that the equations hold if each $\langle E(pXq) \rangle$ is substituted with $[E(pXq)]$ (which means

that the tuple of all $[E(pXq)]$ values is a fixed-point of \mathcal{F} , and hence this tuple can only be larger than μ). A run that starts with a transition $pX \xrightarrow{x} rYZ$ and ends at $q\varepsilon$ must go through a first configuration of stack length 1. Moreover, this configuration must have Z as stack content. Fix a state s , and consider the accumulated reward under the assumption that this intermediate state of the run is s . The run can be split into three parts as follows: $pX \xrightarrow{x} rYZ \rightarrow^* sZ \rightarrow^* q\varepsilon$. The expectation of the reward accumulated during the path is the sum of the expectations of the accumulated rewards, and so equal to $f(rYZ) + [E(rYZ, sZ)] + [E(sZq)]$, where $[E(rYZ, sZ)]$ denotes the conditional expected accumulated reward between the configurations rYZ and sZ . Since f is simple, we have $f(rYZ) = f(r)$, and, moreover, the rewards accumulated during the path $rYZ \rightarrow^* sZ$ and the path $rY \rightarrow^* s\varepsilon$ obtained by removing Z from all configurations of $rYZ \rightarrow^* sZ$ coincide. So $[E(rYZ, sZ)] = [E(rYs)]$, and so the expected value for paths of the form $pX \xrightarrow{x} rYZ \rightarrow^* sZ \rightarrow^* q\varepsilon$ is $f(r) + [E(rYs)] + [E(sZq)]$. The conditional expectation under the assumption that the run starts with $pX \xrightarrow{x} rYZ$ and has s as intermediate state is then given by $[rYs][sZq](f(r) + \langle E(rYs) \rangle + \langle E(sZq) \rangle)$.

The other inequality is proven inductively—for each $i \in \mathbb{N}$ we define a random variable $R_{pXq,i}$ over $Run(pX)$, which returns $F(w^\ell)$ if $w(\ell) = q\varepsilon$ for some $\ell \leq i$, and zero otherwise. Then we can also approximate $[E(pXq)]$ by a family of conditional expectations $[E(pXq)]_i = E(R_{pXq,i} \mid Run(pXq))$. Obviously, $[E(pXq)] = \lim_{i \rightarrow \infty} [E(pXq)]_i$. Then, it is inductively shown that for each $i \in \mathbb{N}$, the tuple of all $[E(pXq)]_i$ values can only be less than μ . Hence, the same holds for the limit $[E(pXq)]$. \square

Corollary 3.2. *$[E(pXq)]$ is expressible. Moreover, the problem whether $[E(pXq)] \sim \varrho$, where $\varrho \in \mathbb{Q}_\omega$ and $\sim \in \{<, >, \leq, \geq, =\}$, is in **PSPACE**.*

In Corollary 3.2, some extra care is needed to compute all $[E(pXq)]$ that take the value ω .

Example 3.3 (Termination time). *If we assume that every transition of the pPDA takes one time unit, the expected termination time of the system when started in the configuration p_0X_0 can be computed as follows. Let f be the simple reward function that assigns 1 to each control state. Then, the expected termination time under the condition that the pPDA terminates (i.e., reaches a configuration with empty stack) is given by*

$$\frac{\sum_{q \in Q} [E(p_0X_0q)] \cdot [p_0X_0q]}{1 - [p_0X_0\uparrow]}$$

Hence, the conditional expected termination time is expressible (of course, the fraction only makes sense if the probability of termination is non-zero).

Let us note that the problem of computing the expected accumulated reward between a given pair of (arbitrary) configurations $p\alpha$ and $q\beta$ is easily reducible to the problem of computing $[E(pXq)]$. Hence, we can also solve this problem.

Now we show how to compute the conditional variance of the accumulated reward of a path under the condition that it starts at pX and ends at $q\varepsilon$. Since we already know how to compute $E(R_{pXq} \mid \text{Run}(pXq))$, it only remains to compute the conditional second moment $E(R_{pXq}^2 \mid \text{Run}(pXq))$.

Similarly as before, consider the probabilistic space $(\text{Run}(pX), \mathcal{F}, \mathcal{P})$. For every $q \in Q$ such that $[pXq] > 0$ we define a random variable Q_{pXq} over $\text{Run}(pX)$ as follows:

$$Q_{pXq}(w) = \begin{cases} 0 & \text{if } w \notin \text{Run}(pXq) \\ (F(w^\ell))^2 & \text{if } w(\ell) = q\varepsilon \end{cases}$$

$E(Q_{pXq} \mid \text{Run}(pXq))$ is the conditional expected square of the accumulated reward from pX to $q\varepsilon$, under the condition that $q\varepsilon$ is reached. From now on we write $[Q(pXq)]$ instead of $E(Q_{pXq} \mid \text{Run}(pXq))$.

Analogously as for $[E(pXq)]$ we now show that the tuple of all $[Q(pXq)]$ values forms the least solution of an effectively constructible system of linear equations. Since the values of $[E(pXq)]$ are expressible, they can be used as coefficients in the system. So, let $\mathcal{V} = \{\langle Q(pXq) \rangle \mid p, q \in Q, X \in \Gamma, [pXq] > 0\}$ be a set of variables over \mathbb{R}_ω^+ . Consider the following system of linear equations (linear in the $\langle Q(aBc) \rangle$ variables): $\langle Q(pXq) \rangle = 0$ for $[pXq] = 0$, else

$$\langle Q(pXq) \rangle = \frac{1}{[pXq]} \left(\sum_{pX \xrightarrow{x} q\varepsilon} x \cdot f(q)^2 + \sum_{pX \xrightarrow{x} rYZ} x \cdot \sum_{s \in Q} [rYs][sZq] K_{pX, rYZ, s} \right)$$

where the expression $K_{pX, rYZ, s}$ stands for

$$\langle Q(rYs) \rangle + \langle Q(sZq) \rangle + f(r)^2 + 2[E(rYs)][E(sZq)] + 2f(r)[E(rYs)] + 2f(r)[E(sZq)]$$

If $[rYs]$, $[sZq]$, $[E(rYs)]$, or $[E(sZq)]$ is zero, then the corresponding summands are eliminated. Now we derive analogous results as in the case of $[E(pXq)]$:

Theorem 3.4. *The tuple of all $[Q(pXq)]$ is exactly the least solution of the above system of equations in \mathbb{R}_ω^+ .*

Proof sketch. The structure of the proof is the same as for Theorem 3.1. Let us just explain the key points for obtaining $K_{pX, rYZ, s}$. Recall that if A and B are independent random variables, then $E((A + B)^2) = E(A^2) + E(B^2) + 2E(A) \cdot E(B)$. In general, if $X = \sum_{i=1}^n X_i$ and the X_i are pairwise independent, then $E(X^2)$ is expressible as a polynomial in $E(X_i)$ and $E(X_i^2)$.

We are interested in the conditional second moment of the accumulated reward of a path under the assumption that the path has the form $pX \xrightarrow{x} rYZ \rightarrow^* sZ \rightarrow^* q\varepsilon$ for fixed pX , rYZ , sZ , and $q\varepsilon$, as in Theorem 3.1. Now we make two observations. First, for simple reward functions the expectation of the reward accumulated along the path is given by $f(r) + [E(rYs)] + [E(sZq)]$ (see the proof of Theorem 3.1). Second, for arbitrary reward functions the random variables that assign to each part of the run its accumulated reward are pairwise independent (follows from the fact that we have fixed the configurations rYZ , sZ , $q\varepsilon$, and so the initial configuration of a part of the path does not depend on the previous parts). From these two observations follows that the square of the accumulated reward can be expressed in terms of $(f(r))^2$, $[Q(rYs)]$, $[Q(sZq)]$, $f(r)$, $[E(rYs)]$, $[E(sZq)]$. The term $K_{pX, rYZ, s}$ is now obtained by a routine calculation. \square

Corollary 3.5. *$[Q(pXq)]$ is expressible. Moreover, the problem whether $[Q(pXq)] \sim \varrho$, where $\varrho \in \mathbb{Q}_\omega$ and $\sim \in \{<, >, \leq, \geq, =\}$, is in PSPACE.*

Example 3.6. *Consider the pPDA model of the ‘‘gambler’s ruin’’ problem. We have the rules $pC \xrightarrow{x} pCC$, $pC \xrightarrow{1-x} p\varepsilon$, and the initial configuration pC . The minimal solution of the recursive equation system for $[pCp]$ yields that $[pCp] = 1$ if $x \leq 1/2$ and $[pCp] = (1-x)/x$ if $x > 1/2$.*

By solving the recursive equations of Theorem 3.1 and Theorem 3.4, one obtains the following results. The conditional expectation $[E(pCp)]$ of the distance from pC to $p\varepsilon$ is 0 for $x = 1$, ω for $x = 1/2$, and otherwise

$$[E(pCp)] = \frac{1-x+x[pCp]^2}{[pCp]-2x[pCp]^2}$$

The conditional expectation $[Q(pCp)]$ of the square of the distance is 0 for $x = 1$, ω for $x = 1/2$, and otherwise

$$[Q(pCp)] = \frac{(1-x)+x[pCp]^2(2[E(pCp)]^2+4[E(pCp)]+1)}{[pCp]-2x[pCp]^2}$$

For $x = 3/4$ one obtains $[pCp] = 1/3$, $[E(pCp)] = 2$, and $[Q(pCp)] = 10$. Hence, the conditional variance is 6.

For $x = 1/2$ one obtains the well-known result $[pCp] = 1$ and $[E(pCp)] = \omega$, i.e., although $p\varepsilon$ is reached almost surely, the expected number of steps to get there is infinite.

Finally, let us note that the approach of this section can be used to compute the conditional k^{th} moment $E(R_{pXq}^k \mid \text{Run}(pXq))$ for every $k \in \mathbb{N}$, which can be useful for a deeper analysis of $\text{Run}(pXq)$.

4. Simple reward functions: Expected gain

Similarly as in Section 3, let us fix a simple reward function f and its associated function F which assigns an ac-

accumulated reward to each finite computation path. We also fix a pPDA $\Delta = (Q, \Gamma, \delta, Prob)$. To simplify our presentation, we assume that there is a special *initial configuration* $q_0 Z_0$, where $q_0 \in Q$ and $Z_0 \in \Gamma$, such that the symbol Z_0 can never be removed from the stack.

We define a function $G_f: Run(q_0 Z_0) \rightarrow \mathbb{R}^+$ as follows.

$$G_f(w) = \begin{cases} \lim_{n \rightarrow \infty} \frac{F(w^n)}{n} & \text{if the limit exists;} \\ \perp & \text{otherwise.} \end{cases}$$

If $\mathcal{P}(G_f = \perp) = 0$ and $\mathcal{P}(G_f \leq x)$ exists for each $x \in \mathbb{R}^+$, then G_f is a random variable where $G_f(w)$ corresponds to the average reward earned per transition during the execution of w , which we call the *gain* of w . Our aim is to compute the expected value of G_f , which is given by $E(G_f) = \int_{Run(q_0 Z_0)} G_f(w) d\mathcal{P}$, assuming that the integral exists. Generally, computing $E(G_f)$ appears to be a difficult problem. Nevertheless, it becomes solvable under a relatively mild assumption. In order to formulate the assumption, we need to recall a definition of [11]. (In fact, this definition is slightly extended to fit our present needs.)

Definition 4.1. Let $w = p_1 \alpha_1, p_2 \alpha_2 \dots$ be an (infinite) run in $Run(q_0 Z_0)$. For each $i \in \mathbb{N}$ we define the i^{th} minimum of w , denoted $\min_i(w)$, inductively as follows. The i^{th} minimum can be either increasing or non-increasing.

- $\min_1(w) = p_1 \alpha_1$ (i.e., $\min_1(w)$ is the starting configuration $q_0 Z_0$ of w). We stipulate that $\min_1(w)$ is non-increasing.
- Let $\min_i(w) = p_\ell \alpha_\ell$. Then $\min_{i+1}(w) = p_k \alpha_k$ where k is the least number such that $k > \ell$ and $|\alpha_{k'}| \geq |\alpha_k|$ for each $k' \geq k$. Observe that $|\alpha_k| - |\alpha_\ell|$ equals either 1 or 0. In the first case, $\min_{i+1}(w)$ is increasing. Otherwise, $\min_{i+1}(w)$ is non-increasing.

Our assumption is the following: *the expected accumulated reward between any two consecutive minima is finite*. A precise formulation of this condition is given below; as we shall see, the condition can be effectively checked in polynomial space. From a practical point of view, the introduced restriction is not strong. In the context of programs with procedures, one sufficient condition which implies that our assumption is satisfied is that the expected termination time of each procedure is finite. One can argue that if the expected termination time for some procedure is infinite, there is a design error in the system.

The Markov chain \mathbf{X} . For each $i \in \mathbb{N}$ we define a random variable X_i over $Run(q_0 Z_0)$ as follows: $X_i(w) = (qY, m)$, where qY is the head of $\min_i(w)$, and m is either + or 0 depending on whether $\min_i(w)$ is increasing or non-increasing, respectively. The next lemma reveals that the sequence $\mathbf{X} = X_1, X_2, \dots$ is a homogeneous Markov chain.

Lemma 4.2. For all $n \geq 2$ and $(q_1 Y_1, m_1), \dots, (q_n Y_n, m_n)$ where $\mathcal{P}(\bigwedge_{i=1}^{n-1} X_i = (q_i Y_i, m_i)) > 0$ we have that

$$\mathcal{P}(X_n = (q_n Y_n, m_n) \mid \bigwedge_{i=1}^{n-1} X_i = (q_i Y_i, m_i)) = \frac{K}{[q_{n-1} Y_{n-1} \uparrow]}$$

where K is equal either to $\sum_{q_{n-1} Y_{n-1} \xrightarrow{x} q_n Y_n Z} x \cdot [q_n Y_n \uparrow]$ or to $\sum_{q_{n-1} Y_{n-1} \xrightarrow{x} r Z Y_n} x \cdot [r Z q_n] \cdot [q_n Y_n \uparrow] + \sum_{q_{n-1} Y_{n-1} \xrightarrow{x} q_n Y_n} x \cdot [q_n Y_n \uparrow]$

depending on whether m_n is equal to + or 0, respectively.

Proof sketch. Let us assume that the current minimum is $pX\alpha$. This assumption means that we only consider those runs from $pX\alpha$ which never access α (and hence α is completely irrelevant). Under this assumption, the probability that the next minimum will be increasing and of the form $qYZ\alpha$ is equal to the probability that we execute a transition of the form $pX \xrightarrow{x} qYZ$ from $pX\alpha$ and the stack is never decreased to $Z\alpha$ in the future. Hence, the considered conditional probability is equal to $\sum_{pX \xrightarrow{x} qYZ} x \cdot \frac{[qY \uparrow]}{[pX \uparrow]}$. Similarly, the conditional probability that the next minimum will be non-increasing and of the form $qY\alpha$ is also evaluated by considering the first transition of $pX\alpha$. This transition is either of the form $pX \xrightarrow{x} qY$, in which case the considered conditional probability equals $x \cdot \frac{[qY \uparrow]}{[pX \uparrow]}$, or of the form $pX \xrightarrow{x} rZY$, where we have to get rid of the symbol Z by a sequence of transitions of the form $rZ \rightarrow^* q\epsilon$. Hence, in the second case the conditional probability equals $x[rZq] \frac{[qY \uparrow]}{[pX \uparrow]}$. \square

Observe that the expression given in Lemma 4.2 depends just on the values of X_n and X_{n-1} . Since all probabilities which appear in this expression are expressible, the transition probabilities of \mathbf{X} are expressible as well.

A *trajectory* in \mathbf{X} is an infinite sequence s_0, s_1, \dots of states of \mathbf{X} such that $s_0 = (q_0 Z_0, 0)$ and the probability of $s_i \rightarrow s_{i+1}$ is non-zero for each $i \in \mathbb{N}_0$.

To each $w \in Run(q_0 Z_0)$ we can associate its *footprint* $X_1(w), X_2(w), \dots$. Note that there can be runs whose footprints are *not* trajectories in \mathbf{X} . Let C_1, \dots, C_k be the BSCC of \mathbf{X} . To each C_i we associate the set $Run(q_0 Z_0, C_i)$ consisting of all $w \in Run(q_0 Z_0)$ such that the footprint of w is a trajectory in \mathbf{X} which hits the component C_i . Note that since \mathbf{X} has finitely many states, $\mathcal{P}(Run(q_0 Z_0, C_i))$ is computable by standard methods for finite-state Markov chains. Moreover, it can easily be shown that

$$\sum_{i=1}^k \mathcal{P}(Run(q_0 Z_0, C_i)) = 1 \quad (1)$$

In the following we show that various quantitative properties of $Run(q_0Z_0)$ can be analyzed by considering the property for each $Run(q_0Z_0, C_i)$ separately and combining the obtained results. For this we use generic results which are described next.

The random variable M^f . For each $i \in \mathbb{N}$ we define a random variable M_i^f over $Run(q_0Z_0)$ as follows: Let $w = s_1, s_2, \dots$ be a run of $Run(q_0Z_0)$. Then $M_i^f(w) = F(s_k, \dots, s_\ell)$, where $s_k = \min_i(w)$ and $s_\ell = \min_{i+1}(w)$. In other words, $M_i^f(w)$ is the reward accumulated between $\min_i(w)$ and $\min_{i+1}(w)$.

Lemma 4.3. *Let $i \in \mathbb{N}$, and let $(pX, m), (qY, n)$ be two states of \mathbf{X} (not necessarily different) such that $\mathcal{P}(X_i=(pX, m) \wedge X_{i+1}=(qY, n)) > 0$. The conditional expectation*

$$E(M_i^f \mid X_i=(pX, m) \wedge X_{i+1}=(qY, n))$$

is equal either to $f(q)$ or to

$$\frac{\sum_{pX \xrightarrow{r} qY} x[rZq][qY\uparrow](f(r) + [E(rZq)]) + \sum_{pX \xrightarrow{r} qY} x[qY\uparrow]f(q)}{\text{Prob}((pX, m) \rightarrow (qY, n))}$$

depending on whether n is equal to $+$ or 0 , respectively.

Proof sketch. We use a similar approach as in Lemma 4.2, only that now the associated analysis of possible runs between two consecutive minima must be carried out rather carefully. \square

Since $E(M_i^f \mid X_i=(pX, m) \wedge X_{i+1}=(qY, n))$ is independent of i as long as $\mathcal{P}(X_i=(pX, m) \wedge X_{i+1}=(qY, n)) > 0$, this conditional expectation can be associated directly with the edge $(pX, m) \rightarrow (qY, n)$ in \mathbf{X} , and will be denoted by $E^f((pX, m) \rightarrow (qY, n))$ in the rest of this paper. Observe that $E^f((pX, m) \rightarrow (qY, n))$ is expressible due to the results of Section 3.

Similarly, we can also express the conditional second moment $E((M_i^f)^2 \mid X_i=(pX, m) \wedge X_{i+1}=(qY, n))$ by employing the results of Section 3. The conditional second moment (and hence also the conditional variance) are thus expressible.

Now we introduce another random variable M^f over $Run(q_0Z_0)$ which corresponds to the average accumulated reward between two consecutive minima. Formally, for each $w \in Run(q_0Z_0)$ we define

$$M^f(w) = \begin{cases} \lim_{n \rightarrow \infty} \frac{M_1^f(w) + \dots + M_n^f(w)}{n} & \text{if the limit exists;} \\ \perp & \text{otherwise.} \end{cases}$$

Theorem 4.4. *Let C be a BSCC of \mathbf{X} , and let μ_C be the invariant probability distribution for C (here we view C as an irreducible finite-state Markov chain). Let $k_C = \sum_{s \in C} \mu_C(s) \cdot \sum_{s \xrightarrow{x} t} x \cdot E^f(s \rightarrow t)$. Then $\mathcal{P}(M^f = k_C \mid Run(q_0Z_0, C)) = 1$.*

Proof sketch. Realize that if the transitions of C were assigned fixed values, we could apply standard results for finite-state Markov chains to compute the average reward of a transition. The resulting expression would be the one given in our theorem. Since we deal with expected rewards between two states, a full proof is somewhat technical. \square

According to Theorem 4.4, the variable M^f takes the same value k_C for almost all runs of $Run(q_0Z_0, C)$. Together with Equation (1), this implies that $\mathcal{P}(M^f = \perp) = 0$. Moreover, we have the following corollary:

Corollary 4.5. *Let C be a BSCC of \mathbf{X} . Then*

$$E(M^f \mid Run(q_0Z_0, C)) = \sum_{s \in C} \mu_C(s) \cdot \sum_{s \xrightarrow{x} t} x \cdot E^f(s \rightarrow t)$$

and thus

$$E(M^f) = \sum_{C \in \mathcal{C}} \mathcal{P}(Run(q_0Z_0, C)) \cdot \sum_{s \in C} \mu_C(s) \cdot \sum_{s \xrightarrow{x} t} x \cdot E^f(s \rightarrow t)$$

Hence, $E(M^f)$ is expressible.

Consider the reward function $\mathbf{1}$ that assigns 1 to each control state. Let C be a BSCC of \mathbf{X} . The next lemma tells how to compute $G_f(w)$ for the runs of $Run(q_0Z_0, C)$ when $E^f(s \rightarrow t)$ is finite for all states s, t of C .

Lemma 4.6. *Let C be a BSCC of \mathbf{X} , and let us assume that for all states $s, t \in C$ we have that $E^f(s \rightarrow t)$ is finite. Then for almost all $w \in Run(q_0Z_0, C)$ we have that*

$$G_f(w) = \frac{\sum_{s \in C} \mu_C(s) \cdot \sum_{s \xrightarrow{x} t} x \cdot E^f(s \rightarrow t)}{\sum_{s \in C} \mu_C(s) \cdot \sum_{s \xrightarrow{x} t} x \cdot E^1(s \rightarrow t)}$$

Hence, $G_f(w)$ is the same for almost all $w \in Run(q_0Z_0, C)$, and thus we finally obtain:

Theorem 4.7. *Let \mathcal{C} be the set of all BSCC of \mathbf{X} . Let us assume that for each $C \in \mathcal{C}$ and all states $s, t \in C$ we have that $E^f(s \rightarrow t)$ is finite. Then $E(G_f) = \sum_{C \in \mathcal{C}} \mathcal{P}(Run(q_0Z_0, C)) \cdot G_f(C)$, where*

$$G_f(C) = \frac{\sum_{s \in C} \mu_C(s) \cdot \sum_{s \xrightarrow{x} t} x \cdot E^f(s \rightarrow t)}{\sum_{s \in C} \mu_C(s) \cdot \sum_{s \xrightarrow{x} t} x \cdot E^1(s \rightarrow t)}$$

Remark 4.8. *Note that Lemma 4.6 actually says that the variable G_f only takes one of the finitely many given values for almost all runs of $Run(q_0Z_0)$. Since these values and the associated probabilities are expressible, we have a bit more detailed information about the behaviour of q_0Z_0 , which is not reflected in the average $E(G_f)$.*

The complexity bounds associated to Theorem 4.7 are given in the following corollary:

Corollary 4.9. *The problem whether the assumption of Theorem 4.7 is satisfied for a given pPDA Δ is in **PSPACE**. If the assumption is satisfied, then the problem whether $E(G_f) \sim \varrho$, where $\sim \in \{<, >, \leq, \geq, =\}$ and $\varrho \in \mathbb{Q}$, is in **EXPTIME**.*

Finally, let us note that $E(G_f)$ can exist even if the assumption of Theorem 4.7 is not satisfied. This is demonstrated in the following example:

Example 4.10. *Let us consider a pPDA given by the rules*

$$\begin{array}{cccc} qZ \xrightarrow{1/2} pIZ & qZ \xrightarrow{1/2} pZ & qI \xrightarrow{1/2} pII & qI \xrightarrow{1/2} pE \\ pZ \xrightarrow{1/2} qIZ & pZ \xrightarrow{1/2} qZ & pI \xrightarrow{1/2} qII & pI \xrightarrow{1/2} qE \end{array}$$

where the initial configuration is qZ . Then the chain \mathbf{X} looks as follows (we omit the states that are not reachable from the state $(qZ, 0)$):

$$1/2 \begin{array}{c} \curvearrowright \\ \curvearrowleft \end{array} (pZ, 0) \xrightleftharpoons[1/2]{1/2} (qZ, 0) \begin{array}{c} \curvearrowright \\ \curvearrowleft \end{array} 1/2$$

Let $f(p) = 1$ and $f(q) = 0$. Since $E^f((pZ, 0) \rightarrow (qZ, 0))$ and $E^f((qZ, 0) \rightarrow (pZ, 0))$ are infinite, Theorem 4.7 cannot be applied. Nevertheless, $E(G_f) = 1/2$ because the control states p and q regularly alternate in each $w \in \text{Run}(qZ)$.

Renewal times. Given a set $P \subseteq Q$ of control states, let f_P be the reward function given by $f_P(p) = 1$ if $p \in P$ and $f_P(p) = 0$ otherwise. The variable G_{f_P} assigns to an infinite run the average number of visits to states of P per transition. Therefore, if $E(G_{f_P}) \neq 0$, then $1/E(G_{f_P})$ gives the average number of transitions between any two visits to P , i.e., the average renewal time associated to P (if $E(G_{f_P}) = 0$ then we can say that the average renewal time is infinite).

5. Linear reward functions: Expected accumulated reward

Let us fix a pPDA $\Delta = (Q, \Gamma, \delta, \text{Prob})$ as in Section 3. A reward function $f: \mathcal{C}(\Delta) \rightarrow \mathbb{R}^+$ is *linear* if there are functions $g: Q \rightarrow \mathbb{R}^+$ and $c: \Gamma \rightarrow \mathbb{R}^+$ such that for every $p\alpha \in \mathcal{C}(\Delta)$ we have that $f(p\alpha) = g(p) + \sum_{Y \in \Gamma} c(Y) \cdot \#_Y(\alpha)$, where $\#_Y(\alpha)$ denotes the number of occurrences of Y in α . Notice that the simple reward functions correspond to the special case when $c(Y) = 0$ for every $Y \in \Gamma$.

For the rest of this section we fix a linear reward function ℓ given by the functions g and c . The associated reward function for finite paths is denoted L .

We use the notation introduced in Section 3. Observe that the definition of conditional expectation $[E(pXq)]$ makes sense for an arbitrary reward function f . We write $[E(pXq), f]$ to denote this conditional expectation for a

given f . In particular, we shall consider the function ℓ introduced above, and the function $\mathbf{1}$ of the previous section which assigns 1 to all configurations. Since the function $\mathbf{1}$ is simple, $[E(pXq), \mathbf{1}]$ is expressible (see Corollary 3.2).

Let $\langle E(pXq), \ell \rangle$ be a variable for all $p, q \in Q$ and $X \in \Gamma$ such that $[pXq] > 0$. Now consider the system of recursive equations, where each variable $\langle E(pXq), \ell \rangle$ is equal to

$$\frac{1}{[pXq]} \left(\sum_{pX \xrightarrow{x} q\epsilon} x \cdot \ell(q\epsilon) + \sum_{pX \xrightarrow{x} rYZ} x \cdot \sum_{s \in Q} [rYs][sZq] K'_{pX, rYZ} \right)$$

where the term $K'_{pX, rYZ}$ is given by

$$\langle E(rYs), \ell \rangle + \langle E(sZq), \ell \rangle + \ell(rYZ) + c(Z)[E(rYs), \mathbf{1}]$$

Observe that in the case when $c(Y) = 0$ for every $Y \in \Gamma$ we recover the system of Section 3. Also note that since $[E(pXq), \mathbf{1}]$ appears in the above equation, we would still have to handle simple reward functions separately if we started directly with linear reward functions in Section 3.

Theorem 5.1. *The tuple of all $[E(pXq), \ell]$ values is exactly the least solution of the above system of equations in \mathbb{R}_ω^+ with respect to component-wise ordering.*

Proof sketch. The proof is very similar to that of Theorem 3.1. The only difference is the following. As in Theorem 3.1, let $rY \rightarrow^* s\epsilon$ be the path obtained by removing Z from all configurations of $rYZ \rightarrow^* sZ$. In the case of simple reward functions, the rewards accumulated during $rYZ \rightarrow^* sZ$ and $rY \rightarrow^* s\epsilon$ coincide. In the case of linear functions, the reward accumulated during $rYZ \rightarrow^* sZ$ is equal to the reward accumulated during $rY \rightarrow^* s\epsilon$ plus $c(Z)$ times the length of $rY \rightarrow^* s\epsilon$. So, in average this reward equals $[E(rYs), \ell] + c(Z)[E(rYs), \mathbf{1}]$, because $[E(rYs), \mathbf{1}]$ is the expected length of the path $rYZ \rightarrow^* sZ$. This leads to the term $K'_{pX, rYZ}$. \square

For the conditional second moment, we adopt a similar notation as above. Consider the system of equations

$$\langle Q(pXq), \ell \rangle = \frac{1}{[pXq]} \left(\sum_{pX \xrightarrow{x} q\epsilon} x \cdot \ell(q\epsilon)^2 + \sum_{pX \xrightarrow{x} rYZ} x \cdot \sum_{s \in Q} [rYs][sZq] K'_{pX, rYZ, s} \right)$$

where the expression $K'_{pX, rYZ, s}$ stands for

$$\begin{aligned} & \langle Q(rYs), \ell \rangle + \langle Q(sZq), \ell \rangle + \ell(rYZ)^2 + 2[E(rYs), \ell][E(sZq), \ell] \\ & + 2\ell(rYZ)[E(rYs), \ell] + 2\ell(rYZ)[E(sZq), \ell] \\ & + 2c(Z)([E(rYs), \ell] + \ell(rYZ) + [E(sZq), \ell])[E(rYs), \mathbf{1}]) + \\ & c(Z)^2[E(rYs), \mathbf{1}] \end{aligned}$$

Again, taking $c(Y) = 0$ for every $Y \in \Gamma$ we recover the system of Section 3. We have the following result, which is proved by combining the observations presented in proofs of Theorem 3.4 and Theorem 5.1:

Theorem 5.2. *The tuple of all $[Q(pXq, \ell)]$ values is exactly the least solution of the above system of equations in \mathbb{R}_ω^+ with respect to component-wise ordering.*

6. Linear reward functions: Expected gain

Let us fix a pPDA $\Delta = (Q, \Gamma, \delta, Prob)$ and its initial configuration q_0Z_0 as in Section 4. We also fix a linear reward function ℓ given by the functions g and c . In this section we show how to compute $E(G_\ell)$.

We say that a transition $(qY, m) \rightarrow (rZ, n)$ of the Markov chain \mathbf{X} is *bounded* if either $n = +$ and for all transitions of the form $qY \xrightarrow{x} rZT$ we have that $c(T) = 0$, or $n = 0$ and for all transitions of the form $qY \xrightarrow{x} tTZ$ we have that $[E(tTr), \ell]$ is finite. Note that we can effectively check if a given transition is bounded by using the results of Section 5. A transition which is not bounded is *unbounded*.

Lemma 6.1. *Let C be a BSCC of \mathbf{X} which contains an unbounded transition. Then for almost all $w \in Run(q_0Z_0)$ we have that $G_\ell(w)$ is infinite.*

Now assume that C is a BSCC of \mathbf{X} where all transitions are bounded. For simplicity, consider first the case when \mathbf{X} is strongly connected. As in Section 4, define a random variable M_i^ℓ which for every $w \in Run(q_0Z_0)$ returns the accumulated reward between $\min_i(w)$ and $\min_{i+1}(w)$, and a random variable M^ℓ as the average reward collected when moving from one minimum to the next.

In the case of a simple function f that depends only on the control state, the variable M_i^f depends only on X_i and X_{i+1} . This is no longer the case for a linear function ℓ . The reason is that the variable X_i records only the head of the i^{th} minimum, but not the stack content which is needed to compute the reward. The stack content of the i^{th} minimum depends on the values of all of X_1, \dots, X_i . Whenever one of these variables is of the form $(pX, +)$, the stack length of the i^{th} minimum increases by 1. Fortunately, since we assume that all transitions are bounded, increasing minima are no longer a problem, because the c -value of symbols that are pushed is zero. Then M_i^ℓ depends only on X_i and X_{i+1} , and we can reuse all the results of Section 4. In particular, Lemma 4.3 still holds after some straightforward modifications (these modifications are based on the same idea which was used in Theorem 5.1 to modify the equations of Theorem 3.1). This shows how to compute $E^\ell((pX, m) \rightarrow (qY, n))$. Corollary 4.5 shows how to compute $E(M^\ell)$ after replacing f by ℓ , and Theorem 4.7 shows how to compute $E(G^\ell)$ after replacing f by ℓ .

If \mathbf{X} is not strongly connected, then the problem is slightly more complicated, and we only sketch the argument. Let B be the random variable which for every $w \in Run(q_0Z_0)$ returns either \perp if the footprint of w does not

hit a BSCC of \mathbf{X} , or $c(\beta)$ where β is the tail of the first minimal configuration which hits a BSCC of \mathbf{X} . For each BSCC C , we express $E(B_C)$, the conditional expected value of B under the assumption that the BSCC reached by the run is C . Now we express $E(G_{\ell, C})$, the conditional expected value of G_ℓ under the condition that the run starts in C (this can be done by the method described in the previous paragraph). Thus, for each BSCC C of Δ we get $E(G_\ell | Run(q_0Z_0, C)) = E(B_C) + E(G_{\ell, C})$, hence

$$E(G_\ell) = \sum_C \mathcal{P}(Run(q_0Z_0, C)) \cdot (E(B_C) + E(G_{\ell, C}))$$

Average stack length. The average stack length corresponds to the linear function $\ell(p\alpha) = |\alpha|$, and so its expectation can be computed using the results of this section.

6.1. Maximal stack length

For many applications, the maximal stack length of a run is perhaps more interesting than the average stack length which can be computed by applying the results of the previous section.

Formally, let us define the random variable ML over $Run(q_0Z_0)$ as follows: $ML(w)$ is the least $\rho \in \mathbb{R}_\omega^+$ such that $\ell(w(i)) \leq \rho$ for all $i \in \mathbb{N}$. If $ML(w) < \omega$, then w is called *bounded*. Observe that in the special case when $g(p) = 0$ and $c(Y) = 1$ for all $p \in Q$ and $Y \in \Gamma$ we have that $ML(w)$ is the *maximal stack length* in w . We are interested in the probability $\mathcal{P}(ML=\omega)$ of unbounded runs. The next theorem says how to compute this probability, but we need a preliminary definition.

We say that a transition $(qY, m) \rightarrow (rZ, n)$ of \mathbf{X} is *limited* if either $n = +$ and for all transitions of the form $qY \xrightarrow{x} rZT$ we have that $c(T) = 0$, or $n = 0$ and there is $\rho \in \mathbb{R}^+$ such that for every path w from qY to rZ we have that $\ell(w(i)) \leq \rho$ for every state $w(i)$ of w . Observe that the exact values of transition probabilities in Δ do not matter here. Hence, one can rely on standard results for non-probabilistic PDA and conclude that the problem whether a given transition is limited is decidable in polynomial time. (The problem whether $(qY, m) \rightarrow (rZ, 0)$ is limited can be decided, e.g., using the results of [10]: One can compute the set $post^*(qY)$ of all successor configurations of qY , the set $pre^*(rZ)$ of all predecessor configurations of rZ . Since these sets are regular, their homomorphic images obtained by replacing all $Y \in \Gamma$ such that $c(Y) = 0$ with ε are also regular. Obviously, the considered transition is limited iff the intersection of these two images is finite. The whole procedure can be implemented in polynomial time.)

Theorem 6.2. *Let \mathcal{C} be the set of all BSCC of \mathbf{X} which contain at least one non-limited transition. Then $\mathcal{P}(ML=\omega)$ is equal to $\sum_{C \in \mathcal{C}} \mathcal{P}(Run(q_0Z_0, C))$.*

Proof sketch. Let C be a BSCC of \mathbf{X} . We show that (1) if $C \in \mathcal{C}$ then almost all runs of $Run(q_0Z_0, C)$ are unbounded, and (2) if $C \notin \mathcal{C}$ then almost all runs of $Run(q_0Z_0, C)$ are bounded. For (1) we can distinguish two cases:

(a) If C contains a non-limited transition of the form $(qY, m) \rightarrow (rZ, +)$, one can argue that almost all runs of $Run(q_0Z_0, C)$ contain infinitely many pairs of consecutive configurations of the form $qY\alpha, rZT\alpha$, where $c(T) > 0$, which are both minimal (realize that if $w(i)$ is an increasing minimum of a run w , then $w(i-1)$ is also a minimum). Hence, the T is pushed infinitely many times when entering a minimal configuration, and hence almost all runs of $Run(q_0Z_0, C)$ are unbounded.

(b) If C contains a non-limited transition of the form $(qY, m) \rightarrow (rZ, 0)$, one has to realize that almost all runs of $Run(q_0Z_0, C)$ have infinitely many pairs of consecutive minima whose \mathbf{X} -values are (qY, m) and $(rZ, 0)$. Since each finite path between qY and rZ has a non-zero fixed probability, almost all runs of $Run(q_0Z_0, C)$ execute each of these finite paths infinitely many times. Hence, almost all runs are again unbounded.

Finally, (2) follows by observing that almost all runs of $Run(q_0Z_0, C)$ have only finitely many “properly increasing” minima, i.e., those increasing minima where the incoming transition pushes a symbol T such that $c(T) > 0$. Hence, the value of ℓ remains bounded if we restrict ourselves to the minimal configurations. However, from the definition of limited transitions it follows that the value of ℓ is bounded also between the minimal configurations by a global constant. \square

Corollary 6.3. *The problem whether $\mathcal{P}(ML=\omega) \sim \varrho$, where $\sim \in \{<, >, \leq, \geq, =\}$ and $\varrho \in \mathbb{Q}$, is in **EXPTIME**. In the special case when $\varrho \in \{0, 1\}$ the problem belongs to **PSPACE**.*

Theorem 6.2 shows that the probability $\mathcal{P}(ML=\omega)$ is expressible. We can easily show that also $\mathcal{P}(ML=\varrho)$ and $\mathcal{P}(ML\leq\varrho)$ are expressible for every $\varrho \in \mathbb{Q}$ by applying results about the quantitative model-checking problem for LTL properties [11, 6, 12]. Computing the expectation $E(ML)$ seems to be a harder problem which is left for future work.

7. Conclusions and future work

The results about expected gain for simple reward functions indicate that our proof techniques might also be used for analysis of long-run average behavior of probabilistic systems in the style of [9]. In certain situations, properties of individual runs are more relevant than ensemble averages computed over all runs. For example, one can ask what is the probability of all runs where the average reward per

transition stays within certain bounds. In fact, using our results we *can* answer even this question, at least for simple reward functions (see Remark 4.8). Hence, an interesting open problem is whether one can extend our results to answer more complicated quantitative questions of this kind.

Acknowledgments We thank Olivier Serre for drawing our attention to the stack length problem, and Tomáš Brázdil for reading a preliminary draft of this paper.

References

- [1] R. Alur, K. Etessami, and P. Madhusudan. A temporal logic of nested calls and returns. In *Proc. TACAS 2004, LNCS 2988*, p. 467–481. Springer, 2004.
- [2] R. Alur, K. Etessami, and M. Yannakakis. Analysis of recursive state machines. In *Proc. CAV 2001, LNCS 2102*, p. 207–220. Springer, 2001.
- [3] R. Alur and P. Madhusudan. Visibly pushdown languages. In *Proc. STOC 2004*, p. 202–211. ACM Press, 2004.
- [4] M. Benedikt, P. Godefroid, and T.W. Reps. Model checking of unrestricted hierarchical state machines. In *Proc. ICALP’2001, LNCS 2076*, p. 652–666. Springer, 2001.
- [5] A. Bouajjani, J. Esparza, and O. Maler. Reachability analysis of pushdown automata: application to model checking. In *Proc. CONCUR’97, LNCS 1243*, 135–150. Springer, 1997.
- [6] T. Brázdil, A. Kučera, and O. Stražovský. On the decidability of temporal properties of probabilistic pushdown automata. In *Proc. STACS’2005, LNCS 3404*, p. 145–157. Springer, 2005.
- [7] J. Canny. Some algebraic and geometric computations in PSPACE. In *Proc. STOC’88*, p. 460–467. ACM Press, 1988.
- [8] S. Chaudhuri, R. Alur, K. Etessami, and P. Madhusudan. On-the-fly reachability and cycle detection for recursive state machines. In *Proc. TACAS 2005, LNCS 3440*. Springer, 2005.
- [9] L. de Alfaro. How to specify and verify the long-run average behavior of probabilistic systems. In *Proc. LICS’98*, p. 454–465. IEEE, 1998.
- [10] J. Esparza, D. Hansel, P. Rossmanith, and S. Schwoon. Efficient algorithms for model checking pushdown systems. In *Proc. CAV 2000, LNCS 1855*, p. 232–247. Springer, 2000.
- [11] J. Esparza, A. Kučera, and R. Mayr. Model-checking probabilistic pushdown automata. In *Proc. LICS 2004*, p. 12–21. IEEE, 2004.
- [12] K. Etessami and M. Yannakakis. Algorithmic verification of recursive probabilistic systems. In *Proc. TACAS 2005, LNCS 3440*, p. 253–270. Springer, 2005.
- [13] K. Etessami and M. Yannakakis. Recursive Markov chains, stochastic grammars, and monotone systems of non-linear equations. In *Proc. STACS’2005, LNCS 3404*, p. 340–352. Springer, 2005.
- [14] D. Grigoriev. Complexity of deciding Tarski algebra. *Journal of Symbolic Computation*, 5(1–2):65–108, 1988.
- [15] A. Tarski. *A Decision Method for Elementary Algebra and Geometry*. Univ. of California Press, Berkeley, 1951.