

# On Probabilistic Parallel Programs with Process Creation and Synchronisation <sup>\*</sup>

Stefan Kiefer and Dominik Wojtczak

Oxford University Computing Laboratory, UK

{stefan.kiefer,dominik.wojtczak}@comlab.ox.ac.uk

**Abstract.** We initiate the study of probabilistic parallel programs with dynamic process creation and synchronisation. To this end, we introduce *probabilistic split-join systems (pSJSs)*, a model for parallel programs, generalising both probabilistic pushdown systems (a model for sequential probabilistic procedural programs which is equivalent to recursive Markov chains) and stochastic branching processes (a classical mathematical model with applications in various areas such as biology, physics, and language processing). Our pSJS model allows for a possibly recursive spawning of parallel processes; the spawned processes can synchronise and return values. We study the basic performance measures of pSJSs, especially the distribution and expectation of space, work and time. Our results extend and improve previously known results on the subsumed models. We also show how to do performance analysis in practice, and present two case studies illustrating the modelling power of pSJSs.

## 1 Introduction

The verification of probabilistic programs with possibly recursive procedures has been intensely studied in the last years. The Markov chains or Markov Decision Processes underlying these systems may have infinitely many states. Despite this fact, which prevents the direct application of the rich theory of finite Markov chains, many positive results have been obtained. Model-checking algorithms have been proposed for both linear and branching temporal logics [12, 16, 26], algorithms deciding properties of several kinds of games have been described (see e.g. [15]), and distributions and expectations of performance measures such as run-time and memory consumption have been investigated [13, 5, 6].

In all these papers programs are modelled as *probabilistic pushdown systems (pPDSs)* or, equivalently [10], as recursive Markov chains. Loosely speaking, a pPDS is a pushdown automaton whose transitions carry probabilities. The *configurations* of a pPDS are pairs containing the current control state and the current stack content. In each *step*, a new configuration is obtained from its predecessor by applying a transition rule, which may modify the control state and the top of the stack.

The programs modelled by pPDSs are necessarily sequential: at each point in time, only the procedure represented by the topmost stack symbol is active. Recursion, however, is a useful language feature also for multithreaded and other parallel programming

---

<sup>\*</sup> The first author is supported by a postdoctoral fellowship of the German Academic Exchange Service (DAAD). The second author is supported by EPSRC grant EP/G050112/1.

languages, such as Cilk and JCilk, which allow, e.g., for a natural parallelisation of divide-and-conquer algorithms [8, 9]. To model parallel programs in probabilistic scenarios, one may be tempted to use *stochastic multitype branching processes*, a classical mathematical model with applications in numerous fields including biology, physics and natural language processing [18, 2]. In this model, each process has a type, and each type is associated with a probability distribution on transition rules. For instance, a branching process with the transition rules  $X \xrightarrow{2/3} \{\}$ ,  $X \xrightarrow{1/3} \{X, Y\}$ ,  $Y \xrightarrow{1} \{X\}$  can be thought of describing a parallel program with two types of processes,  $X$  and  $Y$ . A process of type  $X$  terminates with probability  $2/3$ , and with probability  $1/3$  stays active and spawns a new process of type  $Y$ . A process of type  $Y$  changes its type to  $X$ . A *configuration* of a branching process consists of a pool of currently active processes. In each *step*, all active processes develop in parallel, each one according to a rule which is chosen probabilistically. For instance, a step transforms the configuration  $\langle XY \rangle$  into  $\langle XYX \rangle$  with probability  $\frac{1}{3} \cdot 1$ , by applying the second  $X$ -rule to the  $X$ -process and, in parallel, the  $Y$ -rule to the  $Y$ -process.

Branching processes do not satisfactorily model parallel programs, because they lack two key features: synchronisation and returning values. In this paper we introduce *probabilistic split-join systems (pSJSs)*, a model which offers these features. Parallel spawns are modelled by rules of the form  $X \hookrightarrow \langle YZ \rangle$ . The spawned processes  $Y$  and  $Z$  develop independently; e.g., a rule  $Y \hookrightarrow Y'$  may be applied to the  $Y$ -process, replacing  $Y$  by  $Y'$ . When terminating, a process enters a *synchronisation state*, e.g. with rules  $Y' \hookrightarrow q$  and  $Z \hookrightarrow r$  (where  $q$  and  $r$  are synchronisation states). Once a process terminates in a synchronisation state, it waits for its *sibling* to terminate in a synchronisation state as well. In the above example, the spawned processes wait for each other, until they terminate in  $q$  and  $r$ . At that point, they may *join* to form a single process, e.g. with a rule  $\langle qr \rangle \hookrightarrow W$ . So, synchronisation is achieved by the siblings waiting for each other to terminate. All rules could be probabilistic. Notice that synchronisation states can be used to return values; e.g., if the  $Y$ -process returns  $q'$  instead of  $q$ , this can be recorded by the existence of a rule  $\langle q'r \rangle \hookrightarrow W'$ , so that the resulting process (i.e.,  $W$  or  $W'$ ) depends on the values computed by the joined processes. For the notion of siblings to make sense, a *configuration* of a pSJS is not a set, but a binary tree whose leaves are *process symbols* (such as  $X, Y, Z$ ) or synchronisation states (such as  $q, r$ ). A *step* transforms the leaves of the binary tree in parallel by applying rules; if a leaf is not a process symbol but a synchronisation state, it remains unchanged unless its sibling is also a synchronisation state and a joining rule (such as  $\langle qr \rangle \hookrightarrow W$ ) exists, which removes the siblings and replaces their parent node with the right hand side.

*Related work.* The probabilistic models closest to ours are pPDSs, recursive Markov chains, and stochastic branching processes, as described above. The non-probabilistic (i.e., nondeterministic) version of pSJSs (SJSs, say) can be regarded as a special case of *ground tree rewriting systems*, see [20] and the references therein. A configuration of a ground tree rewriting system is a node-labelled tree, and a rewrite rule replaces a subtree. The *process rewrite system* (PRS) hierarchy of [22] features sequential and parallel process composition. Due to its syntactic differences, it is not obvious whether SJSs are in that hierarchy. They would be above pushdown systems (which is the sequential fragment of PRSs), because SJSs subsume pushdown systems, as we show in

Section 3.1 for the probabilistic models. *Dynamic pushdown networks* (DPNs) [4] are a parallel extension of pushdown systems. A configuration of a DPN is a list of configurations of pushdown systems running in parallel. DPNs feature the spawning of parallel threads, and an extension of DPNs, called *constrained DPNs*, can also model joins via regular expressions on spawned children. The DPN model is more powerful and more complicated than SJSs. All those models are non-probabilistic.

*Organisation of the paper.* In Section 2 we formally define our model and provide further preliminaries. Section 3 contains our main results: we study the relationship between pSJSs and pPDSs (Section 3.1), we show how to compute the probabilities for termination and finite space, respectively (Sections 3.2 and 3.3), and investigate the distribution and expectation of work and time (Section 3.4). In Section 4 we present two case studies illustrating the modelling power of pSJSs. We conclude in Section 5. All proofs are provided in the appendix.

## 2 Preliminaries

For a finite or infinite word  $w$ , we write  $w(0), w(1), \dots$  to refer to its individual letters. We assume throughout the paper that  $\mathcal{B}$  is a fixed infinite set of *basic process symbols*. We use the symbols ‘ $\langle$ ’ and ‘ $\rangle$ ’ as special letters not contained in  $\mathcal{B}$ . For an alphabet  $\Sigma$ , we write  $\langle \Sigma \Sigma \rangle$  to denote the language  $\{\langle \sigma_1 \sigma_2 \rangle \mid \sigma_1, \sigma_2 \in \Sigma\}$  and  $\Sigma^{1,2}$  to denote  $\Sigma \cup \langle \Sigma \Sigma \rangle$ . To a set  $\Sigma$  we associate a set  $T(\Sigma)$  of binary trees whose leaves are labelled with elements of  $\Sigma$ . Formally,  $T(\Sigma)$  is the smallest language that contains  $\Sigma$  and  $\langle T(\Sigma)T(\Sigma) \rangle$ . For instance,  $\langle \langle \sigma \sigma \rangle \sigma \rangle \in T(\{\sigma\})$ .

**Definition 1 (pSJS).** *Let  $Q$  be a finite set of synchronisation states disjoint from  $\mathcal{B}$  and not containing ‘ $\langle$ ’ or ‘ $\rangle$ ’. Let  $\Gamma$  be a finite set of process symbols, such that  $\Gamma \subset \mathcal{B} \cup \langle QQ \rangle$ . Define the alphabet  $\Sigma := \Gamma \cup Q$ . Let  $\delta \subseteq \Gamma \times \Sigma^{1,2}$  be a transition relation. Let  $Prob : \delta \rightarrow (0, 1]$  be a function so that for all  $a \in \Gamma$  we have  $\sum_{a \hookrightarrow \alpha \in \delta} Prob(a \hookrightarrow \alpha) = 1$ . Then the tuple  $S = (\Gamma, Q, \delta, Prob)$  is a probabilistic split-join system (pSJS). A pSJS with  $\Gamma \cap \langle QQ \rangle = \emptyset$  is called branching process.*

We usually write  $a \xrightarrow{p} \alpha$  instead of  $Prob(a \hookrightarrow \alpha) = p$ . For technical reasons we allow branching processes of “degree 3”, i.e., branching processes where  $\Sigma^{1,2}$  may be extended to  $\Sigma^{1,2,3} := \Sigma^{1,2} \cup \{\langle \sigma_1 \sigma_2 \sigma_3 \rangle \mid \sigma_1, \sigma_2, \sigma_3 \in \Sigma\}$ . In branching processes, it is usually sufficient to have  $|Q| = 1$ .


A *Markov chain* is a stochastic process that can be described by a triple  $M = (D, \rightarrow, Prob)$  where  $D$  is a finite or countably infinite set of *states*,  $\rightarrow \subseteq D \times D$  is a *transition relation*, and  $Prob$  is a function which to each transition  $s \rightarrow t$  of  $M$  assigns its probability  $Prob(s \rightarrow t) > 0$  so that for every  $s \in D$  we have  $\sum_{s \rightarrow t} Prob(s \rightarrow t) = 1$  (as usual, we write  $s \xrightarrow{x} t$  instead of  $Prob(s \rightarrow t) = x$ ). A *path* (or *run*) in  $M$  is a finite (or infinite, resp.) word  $u \in D^+ \cup D^\omega$ , such that  $u(i-1) \rightarrow u(i)$  for every  $1 \leq i < |u|$ . The set of all runs that start with a given path  $u$  is denoted by  $Run[M](u)$  (or  $Run(u)$ , if  $M$  is understood). To every  $s \in D$  we associate the probability space  $(Run(s), \mathcal{F}, \mathcal{P})$  where  $\mathcal{F}$  is the  $\sigma$ -field generated by all *basic cylinders*  $Run(u)$  where  $u$  is a path starting with  $s$ , and  $\mathcal{P} : \mathcal{F} \rightarrow [0, 1]$  is the

unique probability measure such that  $\mathcal{P}(Run(u)) = \prod_{i=1}^{|u|-1} x_i$  where  $u(i-1) \xrightarrow{x_i} u(i)$  for every  $1 \leq i < |u|$ . Only certain subsets of  $Run(s)$  are  $\mathcal{P}$ -measurable, but in this paper we only deal with “safe” subsets that are guaranteed to be in  $\mathcal{F}$ . If  $\mathbf{X}_s$  is a random variable over  $Run(s)$ , we write  $\mathbb{E}[\mathbf{X}_s]$  for its expectation. For  $s, t \in D$ , we define  $Run(s \downarrow t) := \{w \in Run(s) \mid \exists i \geq 0 : w(i) = t\}$  and  $[s \downarrow t] := \mathcal{P}(Run(s \downarrow t))$ .

To a pSJS  $S = (\Gamma, Q, \delta, Prob)$  with alphabet  $\Sigma = \Gamma \cup Q$  we associate a Markov chain  $M_S$  with  $T(\Sigma)$  as set of states. For  $t \in T(\Sigma)$ , we define  $Front(t) = a_1, \dots, a_k$  as the unique finite sequence of subwords of  $t$  (read from left to right) with  $a_i \in \Gamma$  for all  $1 \leq i \leq k$ . We write  $|Front(t)| = k$ . If  $k = 0$ , then  $t$  is called *terminal*. The Markov chain  $M_S$  has a transition  $t \xrightarrow{p} t'$ , if:  $Front(t) = a_1, \dots, a_k$ ;  $a_i \xrightarrow{p_i} \alpha_i$  are transitions in  $S$  for all  $i$ ;  $t'$  is obtained from  $t$  by replacing  $a_i$  with  $\alpha_i$  for all  $i$ ; and  $p = \prod_{i=1}^k p_i$ . Note that  $t \xrightarrow{1} t$ , if  $t$  is terminal. For branching processes of degree 3, the set  $T(\Sigma)$  is extended in the obvious way to trees whose nodes may have two or three children.

Denote by  $\mathbf{T}_\sigma$  a random variable over  $Run(\sigma)$  where  $\mathbf{T}_\sigma(w)$  is either the least  $i \in \mathbb{N}$  such that  $w(i)$  is terminal, or  $\infty$ , if no such  $i$  exists. Intuitively,  $\mathbf{T}_\sigma(w)$  is the number of steps in which  $w$  terminates, i.e., the *termination time*. Denote by  $\mathbf{W}_\sigma$  a random variable over  $Run(\sigma)$  where  $\mathbf{W}_\sigma(w) := \sum_{i=0}^{\infty} |Front(w(i))|$ . Intuitively,  $\mathbf{W}_\sigma(w)$  is the total *work* in  $w$ . Denote by  $\mathbf{S}_\sigma$  a random variable over  $Run(\sigma)$  where  $\mathbf{S}_\sigma(w) := \sup_{i=0}^{\infty} |w(i)|$ , and  $|w(i)|$  is the length of  $w(i)$  not counting the symbols ‘ $\langle$ ’ and ‘ $\rangle$ ’. Intuitively,  $\mathbf{S}_\sigma(w)$  is the maximal number of processes during the computation, or, short, the *space* of  $w$ .

*Example 2.* Consider the pSJS with  $\Gamma = \{X, \langle qr \rangle\}$  and  $Q = \{q, r\}$  and the transitions  $X \xrightarrow{0.5} \langle XX \rangle$ ,  $X \xrightarrow{0.3} q$ ,  $X \xrightarrow{0.2} r$ ,  $\langle qr \rangle \xrightarrow{1} X$ . Let  $u = X \langle XX \rangle \langle qr \rangle X q q$ . Then  $u$  is a path, because we have  $X \xrightarrow{0.5} \langle XX \rangle \xrightarrow{0.06} \langle qr \rangle \xrightarrow{1} X \xrightarrow{0.3} q \xrightarrow{1} q$ . Note that  $q$  is terminal. The set  $Run(u)$  contains only one run, namely  $w := u(0)u(1)u(2)u(3)u(4)u(4) \dots$ . We have  $\mathcal{P}(Run(u)) = 0.5 \cdot 0.06 \cdot 0.3$ , and  $\mathbf{T}_X(w) = 4$ ,  $\mathbf{W}_X(w) = 5$ , and  $\mathbf{S}_X(w) = 2$ . The dags in Figure 1 graphically represent this run (on the left), and another example run (on the right) with  $\mathbf{T}_X = 3$ ,  $\mathbf{W}_X = 5$ , and  $\mathbf{S}_X = 3$ .



**Fig. 1.** Two terminating runs

*Remark 3.* Our definition of pSJSs may be more general than needed from a modelling perspective: e.g., our rules allow for both synchronisation and splitting in a single step. We choose this definition for technical convenience and to allow for easy comparisons with pPDSs (Section 3.1).

The complexity-theoretic statements in this paper are with respect to the *size* of the given pSJS  $S = (\Gamma, Q, \delta, Prob)$ , which is defined as  $|\Gamma| + |Q| + |\delta| + |Prob|$ , where  $|Prob|$  equals the sum of the sizes of the binary representations of the values of  $Prob$ . A formula of  $ExTh(\mathbb{R})$ , the existential fragment of the first-order theory of the reals, is of the form  $\exists x_1 \dots \exists x_m R(x_1, \dots, x_n)$ , where  $R(x_1, \dots, x_n)$  is a boolean combination of comparisons of the form  $p(x_1, \dots, x_n) \sim 0$ , where  $p(x_1, \dots, x_n)$  is a multivariate polynomial and  $\sim \in \{<, >, \leq, \geq, =, \neq\}$ . The validity of closed formulas ( $m = n$ ) is decidable in PSPACE [7, 23]. We say that one can *efficiently express* a value  $c \in \mathbb{R}$  associated with a pSJS, if one can, in polynomial space, construct a formula  $\phi(x)$  in  $ExTh(\mathbb{R})$  of polynomial length such that  $x$  is the only free variable in  $\phi(x)$ , and  $\phi(x)$  is true if and only if  $x = c$ . Notice that if  $c$  is efficiently expressible, then  $c \sim \tau$  for  $\tau \in \mathbb{Q}$  is decidable in PSPACE for  $\sim \in \{<, >, \leq, \geq, =, \neq\}$ .

For some lower bounds, we prove hardness (with respect to P-time many-one reductions) in terms of the PosSLP decision problem. The PosSLP (Positive Straight-Line Program) problem asks whether a given straight-line program or, equivalently, arithmetic circuit with operations  $+$ ,  $-$ ,  $\cdot$ , and inputs 0 and 1, and a designated output gate, outputs a positive integer or not. PosSLP is in PSPACE. More precisely, it is known to be on the 4th level of the Counting Hierarchy [1]; it is not known to be in NP. The PosSLP problem is a fundamental problem for numerical computation; it is complete for the class of decision problems that can be solved in polynomial time on models with unit-cost exact rational arithmetic, see [1, 16] for more details.

### 3 Results

#### 3.1 Relationship with probabilistic pushdown systems (pPDSs)

We show that pSJSs subsume pPDSs. A *probabilistic pushdown system (pPDS)* [12, 13, 5, 6] is a tuple  $S = (\Gamma, Q, \delta, Prob)$ , where  $\Gamma$  is a finite *stack alphabet*,  $Q$  is a finite set of *control states*,  $\delta \subseteq Q \times \Gamma \times Q \times \Gamma^{\leq 2}$  (where  $\Gamma^{\leq 2} = \{\alpha \in \Gamma^*, |\alpha| \leq 2\}$ ) is a *transition relation*, and  $Prob : \delta \rightarrow (0, 1]$  is a function so that for all  $q \in Q$  and  $a \in \Gamma$  we have  $\sum_{qa \hookrightarrow r\alpha} Prob(qa \hookrightarrow r\alpha) = 1$ . One usually writes  $qa \xrightarrow{p} r\alpha$  instead of  $Prob(qa \rightarrow r\alpha) = p$ . To a pPDS  $S = (\Gamma, Q, \delta, Prob)$  one associates a Markov chain  $M_S$  with  $Q \times \Gamma^*$  as set of states, and transitions  $q \xrightarrow{1} q$  for all  $q \in Q$ , and  $qa\beta \xrightarrow{p} r\alpha\beta$  for all  $qa \xrightarrow{p} r\alpha$  and all  $\beta \in \Gamma^*$ .

A pPDS  $S_P$  with  $\Gamma_P$  as stack alphabet,  $Q_P$  as set of control states, and transitions  $\xrightarrow{p}_P$  can be transformed to an equivalent pSJS  $S$ : Take  $Q := Q_P \cup \Gamma_P$  as synchronisation states;  $\Gamma := \{\langle qa \rangle \mid q \in Q_P, a \in \Gamma_P\}$  as process symbols; and transitions  $\langle qa \rangle \xrightarrow{p} \langle \langle rb \rangle c \rangle$  for all  $qa \xrightarrow{p}_P rbc$ ,  $\langle qa \rangle \xrightarrow{p} \langle rb \rangle$  for all  $qa \xrightarrow{p}_P rb$ , and  $\langle qa \rangle \xrightarrow{p} r$  for all  $qa \xrightarrow{p}_P r$ . The Markov chains  $M_{S_P}$  and  $M_S$  are isomorphic. Therefore, we

occasionally say that a pSJS is a pPDS, if it can be obtained from a pPDS by this transformation. Observe that in pPDSs, we have  $\mathbf{T} = \mathbf{W}$ , because there is no parallelism.

Conversely, a pSJS  $S$  with alphabet  $\Sigma = \Gamma \cup Q$  can be transformed into a pPDS  $S_P$  by “serialising”  $S$ : Take  $Q_P := \{\square\} \cup \{\bar{q} \mid q \in Q\}$  as control states;  $\Gamma_P := \Gamma \cup Q \cup \{\tilde{q} \mid q \in Q\}$  as stack alphabet; and transitions  $\square a \xrightarrow{p}_P \square \sigma_1 \sigma_2$  for all  $a \xrightarrow{p} \langle \sigma_1 \sigma_2 \rangle$ ,  $\square a \xrightarrow{p}_P \square \sigma$  for all  $a \xrightarrow{p} \sigma$  with  $\sigma \in \Sigma \setminus \langle QQ \rangle$ , and  $\square q \xrightarrow{1}_P \bar{q}$  for all  $q \in Q$ , and  $\bar{q} \sigma \xrightarrow{1}_P \square \sigma \bar{q}$  for all  $q \in Q$  and  $\sigma \in \Sigma$ , and  $\bar{r} \bar{q} \xrightarrow{1}_P \square \langle qr \rangle$  for all  $q, r \in Q$ . The Markov chains  $M_S$  and  $M_{S_P}$  are *not* isomorphic. However, we have:

**Proposition 4.** *There is a probability-preserving bijection between the runs  $Run(\sigma \downarrow q)$  in  $M_S$  and the runs  $Run(\square \sigma \downarrow \bar{q})$  in  $M_{S_P}$ . In particular, we have  $[\sigma \downarrow q] = [\square \sigma \downarrow \bar{q}]$ .*

For example, the pSJS run on the left side of Figure 1 corresponds to the pPDS run  $\square X \xrightarrow{0.5} \square X X \xrightarrow{0.3} \square q X \xrightarrow{1} \bar{q} X \xrightarrow{1} \square X \tilde{q} \xrightarrow{0.2} \square r \tilde{q} \xrightarrow{1} \bar{r} \tilde{q} \xrightarrow{1} \square \langle qr \rangle \xrightarrow{1} \square X \xrightarrow{0.3} \square q \xrightarrow{1} \bar{q} \xrightarrow{1} \bar{q} \xrightarrow{1} \dots$

### 3.2 Probability of Termination

We call a run *terminating*, if it reaches a terminal tree. Such a tree can be a single synchronisation state (e.g.,  $q$  on the left of Figure 1), or another terminal tree (e.g.,  $\langle q(rq) \rangle$  on the right of Figure 1). For any  $\sigma \in \Sigma$ , we denote by  $[\sigma \downarrow]$  the termination probability when starting in  $\sigma$ ; i.e.,  $[\sigma \downarrow] = \sum_{t \text{ is terminal}} [\sigma \downarrow t]$ . One can transform any pSJS  $S$  into a pSJS  $S'$  such that whenever a run in  $S$  terminates, then a corresponding run in  $S'$  terminates in a synchronisation state. This transformation is by adding a fresh state  $\tilde{q}$ , and transitions  $\langle rs \rangle \xrightarrow{1} \tilde{q}$  for all  $r, s \in Q$  with  $\langle rs \rangle \notin \Gamma$ , and  $\langle \tilde{q} r \rangle \xrightarrow{1} \tilde{q}$  and  $\langle r \tilde{q} \rangle \xrightarrow{1} \tilde{q}$  for all  $r \in Q$ . It is easy to see that this keeps the probability of termination unchanged, and modifies the random variables  $\mathbf{T}_\sigma$  and  $\mathbf{W}_\sigma$  by at most a factor 2. Notice that the transformation can be performed in polynomial time. After the transformation we have  $[\sigma \downarrow] = \sum_{q \in Q} [\sigma \downarrow q]$ . A pSJS which satisfies this equality will be called *normalised* in the following. From a modelling point of view, pSJSs may be expected to be normalised in the first place: a terminating program should terminate all its processes.

We set up an equation system for the probabilities  $[\sigma \downarrow q]$ . For each  $\sigma \in \Sigma$  and  $q \in Q$ , the equation system has a variable of the form  $\llbracket \sigma \downarrow q \rrbracket$  and an equation of the form  $\llbracket \sigma \downarrow q \rrbracket = f_{\llbracket \sigma \downarrow q \rrbracket}$ , where  $f_{\llbracket \sigma \downarrow q \rrbracket}$  is a multivariate polynomial with nonnegative coefficients. More concretely: If  $q \in Q$ , then we set  $\llbracket q \downarrow q \rrbracket = 1$ ; if  $r \in Q \setminus \{q\}$ , then we set  $\llbracket r \downarrow q \rrbracket = 0$ ; if  $a \in \Gamma$ , then we set

$$\llbracket a \downarrow q \rrbracket = \sum_{\substack{a \xrightarrow{p} \langle \sigma_1 \sigma_2 \rangle \\ \langle q_1 q_2 \rangle \in \Gamma \cap \langle QQ \rangle}} p \cdot \llbracket \sigma_1 \downarrow q_1 \rrbracket \cdot \llbracket \sigma_2 \downarrow q_2 \rrbracket \cdot \llbracket \langle q_1 q_2 \rangle \downarrow q \rrbracket + \sum_{\substack{a \xrightarrow{p} \sigma' \\ \sigma' \in \Sigma \setminus \langle QQ \rangle}} p \cdot \llbracket \sigma' \downarrow q \rrbracket.$$

**Proposition 5.** *Let  $\sigma \in \Sigma$  and  $q \in Q$ . Then  $[\sigma \downarrow q]$  is the value for  $\llbracket \sigma \downarrow q \rrbracket$  in the least (w.r.t. componentwise ordering) nonnegative solution of the above equation system.*

One can efficiently approximate  $[\sigma \downarrow q]$  by applying Newton’s method to the fixed-point equation system from Proposition 5, cf. [16]. The convergence speed of Newton’s method for such equation systems was recently studied in detail [11]. The simpler

“Kleene” method (sometimes called “fixed-point iteration”) often suffices, but can be much slower. In the case studies of Section 4, using Kleene for computing the termination probabilities up to machine accuracy was not a bottleneck. The following theorem essentially follows from similar results for pPDSs:

**Theorem 6 (cf. [14, 16]).** *Consider a pSJS with alphabet  $\Sigma = \Gamma \cup Q$ . Let  $\sigma \in \Sigma$  and  $q \in Q$ . Then (1) one can efficiently express (in the sense defined in Section 2) the value of  $[\sigma \downarrow q]$ , (2) deciding whether  $[\sigma \downarrow q] = 0$  is in P, and (3) deciding whether  $[\sigma \downarrow q] < 1$  is PosSLP-hard even for pPDSs.*

### 3.3 Probability of Finite Space

A run  $w \in \text{Run}(\sigma)$  is either (i) terminating, or (ii) nonterminating with  $\mathbf{S}_\sigma < \infty$ , or (iii) nonterminating with  $\mathbf{S}_\sigma = \infty$ . From a modelling point of view, some programs may be considered incorrect, if they do not terminate with probability 1. As is well-known, this does not apply to programs like operating systems, network servers, system daemons, etc., where nontermination may be tolerated or desirable. Such programs may be expected not to need an infinite amount of space; i.e.,  $\mathbf{S}_\sigma$  should be finite.

Given a pSJS  $S$  with alphabet  $\Sigma = \Gamma \cup Q$ , we show how to construct, in polynomial time, a normalised pSJS  $\bar{S}$  with alphabet  $\bar{\Sigma} = \bar{\Gamma} \cup \bar{Q} \supseteq \Sigma$  where  $\bar{Q} = Q \cup \{\bar{q}\}$  for a fresh synchronisation state  $\bar{q}$ , and  $\mathcal{P}(\mathbf{S}_a < \infty = \mathbf{T}_a \mid \text{Run}(a)) = [a \downarrow \bar{q}]$  for all  $a \in \Gamma$ . Having done that, we can compute this probability according to Section 3.2.

For the construction, we can assume w.l.o.g. that  $S$  has been normalised using the procedure of Section 3.2. Let  $U := \{a \in \Gamma \mid \forall n \in \mathbb{N} : \mathcal{P}(\mathbf{S}_a > n) > 0\}$ .

**Lemma 7.** *The set  $U$  can be computed in polynomial time.*

Let  $B := \{a \in \Gamma \setminus U \mid \forall q \in Q : [a \downarrow q] = 0\}$ , so  $B$  is the set of process symbols  $a$  that are both “bounded above” (because  $a \notin U$ ) and “bounded below” (because  $a$  cannot terminate). By Theorem 6 (2) and Lemma 7 we can compute  $B$  in polynomial time. Now we construct  $\bar{S}$  by modifying  $S$  as follows: we set  $\bar{Q} := Q \cup \{\bar{q}\}$  for a fresh synchronisation state  $\bar{q}$ ; we remove all transitions with symbols  $b \in B$  on the left hand side and replace them with a new transition  $b \xrightarrow{1} \bar{q}$ ; we add transitions  $\langle q_1 q_2 \rangle \xrightarrow{1} \bar{q}$  for all  $q_1, q_2 \in \bar{Q}$  with  $\bar{q} \in \{q_1, q_2\}$ . We have the following proposition.

**Proposition 8.** *(1) The pSJS  $\bar{S}$  is normalised; (2) the value  $[a \downarrow q]$  for  $a \in \Gamma$  and  $q \in Q$  is the same in  $S$  and  $\bar{S}$ ; (3) we have  $\mathcal{P}(\mathbf{S}_a < \infty = \mathbf{T}_a \mid \text{Run}(a)) = [a \downarrow \bar{q}]$  for all  $a \in \Gamma$ .*

Proposition 8 allows for the following theorem.

**Theorem 9.** *Consider a pSJS with alphabet  $\Sigma = \Gamma \cup Q$  and  $a \in \Gamma$ . Let  $s := \mathcal{P}(\mathbf{S}_a < \infty)$ . Then (1) one can efficiently express  $s$ , (2) deciding whether  $s = 0$  is in P, and (3) deciding whether  $s < 1$  is PosSLP-hard even for pPDSs.*

Theorem 9, applied to pPDSs, improves Corollary 6.3 of [13]. There it is shown for pPDSs that comparing  $\mathcal{P}(\mathbf{S}_a < \infty)$  with  $\tau \in \mathbb{Q}$  is in EXPTIME, and in PSPACE if  $\tau \in \{0, 1\}$ . With Theorem 9 we get PSPACE for  $\tau \in \mathbb{Q}$ , and P for  $\tau = 0$ .

### 3.4 Work and Time

We show how to compute the distribution and expectation of work and time of a given pSJS  $S$  with alphabet  $\Sigma = \Gamma \cup Q$ .

*Distribution.* For  $\sigma \in \Sigma$  and  $q \in Q$ , let  $T_{\sigma \downarrow q}(k) := \mathcal{P}(\text{Run}(\sigma \downarrow q), \mathbf{T}_\sigma = k \mid \text{Run}(\sigma))$ . It is easy to see that, for  $k \geq 1$  and  $a \in \Gamma$  and  $q \in Q$ , we have

$$T_{a \downarrow q}(k) = \sum_{\substack{a \xrightarrow{p} \langle \sigma_1 \sigma_2 \rangle \\ \langle q_1 q_2 \rangle \in \Gamma \cap \langle QQ \rangle}} p \cdot \sum_{\substack{\ell_1, \ell_2, \ell_3 \geq 0 \\ \max\{\ell_1, \ell_2\} + \ell_3 = k-1}} T_{\sigma_1 \downarrow q_1}(\ell_1) \cdot T_{\sigma_2 \downarrow q_2}(\ell_2) \cdot T_{\langle q_1 q_2 \rangle \downarrow q}(\ell_3) + \sum_{\substack{a \xrightarrow{p} \sigma' \\ \sigma' \in \Sigma \setminus \langle QQ \rangle}} p \cdot T_{\sigma' \downarrow q}(k-1).$$

This allows to compute the distribution of time (and, similarly, work) using dynamic programming. In particular, for any  $k$ , one can compute  $\vec{T}_{\sigma \downarrow q}(k) := \mathcal{P}(\mathbf{T}_\sigma > k \mid \text{Run}(\sigma \downarrow q)) = 1 - \frac{1}{[\sigma \downarrow q]} \sum_{i=0}^k T_{\sigma \downarrow q}(i)$ .

*Expectation.* For any random variable  $Z$  taking positive integers as value, it holds  $\mathbb{E}Z = \sum_{k=0}^{\infty} \mathcal{P}(Z > k)$ . Hence, one can approximate  $\mathbb{E}[\mathbf{T}_\sigma \mid \text{Run}(\sigma \downarrow q)] = \sum_{k=0}^{\infty} \vec{T}_{\sigma \downarrow q}(k)$  by computing  $\sum_{k=0}^{\ell} \vec{T}_{\sigma \downarrow q}(k)$  for large  $\ell$ . In the rest of the section we show how to decide on the finiteness of expected work and time. It follows from Propositions 10 and 11 below that the expected work  $\mathbb{E}[\mathbf{W}_\sigma \mid \text{Run}(\sigma \downarrow q)]$  is easier to compute: it is the solution of a linear equation system.

We construct a branching process  $\bar{S}$  with process symbols  $\bar{T} = \{\langle aq \rangle \mid a \in \Gamma, q \in Q, [a \downarrow q] > 0\}$ , synchronisation states  $\bar{Q} = \{\perp\}$  and transitions as follows. For notational convenience, we identify  $\perp$  and  $\langle qq \rangle$  for all  $q \in Q$ . For  $\langle aq \rangle \in \bar{T}$ , we set

$$\begin{aligned} & - \langle aq \rangle \xrightarrow{y/[a \downarrow q]} \langle \langle \sigma_1 q_1 \rangle \langle \sigma_2 q_2 \rangle \langle \langle q_1 q_2 \rangle q \rangle \rangle \text{ for all } a \xrightarrow{p} \langle \sigma_1 \sigma_2 \rangle \text{ and } \langle q_1 q_2 \rangle \in \Gamma \cap \langle QQ \rangle, \\ & \text{where } y := p \cdot [\sigma_1 \downarrow q_1] \cdot [\sigma_2 \downarrow q_2] \cdot [\langle q_1 q_2 \rangle \downarrow q] > 0; \\ & - \langle aq \rangle \xrightarrow{y/[a \downarrow q]} \langle \sigma' q \rangle \text{ for all } a \xrightarrow{p} \sigma' \text{ with } \sigma' \in \Sigma \setminus \langle QQ \rangle, \text{ where } y := p \cdot [\sigma' \downarrow q] > 0. \end{aligned}$$

The following proposition (inspired by a statement on pPDSs [6]) links the distributions of  $\mathbf{W}_\sigma$  and  $\mathbf{T}_\sigma$  conditioned under termination in  $q$  with the distributions of  $\mathbf{W}_{\langle \sigma q \rangle}$  and  $\mathbf{T}_{\langle \sigma q \rangle}$ .

**Proposition 10.** *Let  $\sigma \in \Sigma$  and  $q \in Q$  with  $[\sigma \downarrow q] > 0$ . Then*

$$\begin{aligned} \mathcal{P}(\mathbf{W}_\sigma = n \mid \text{Run}(\sigma \downarrow q)) &= \mathcal{P}(\mathbf{W}_{\langle \sigma q \rangle} = n \mid \text{Run}(\langle \sigma q \rangle)) \quad \text{for all } n \geq 0 \quad \text{and} \\ \mathcal{P}(\mathbf{T}_\sigma \leq n \mid \text{Run}(\sigma \downarrow q)) &\leq \mathcal{P}(\mathbf{T}_{\langle \sigma q \rangle} \leq n \mid \text{Run}(\langle \sigma q \rangle)) \quad \text{for all } n \geq 0. \end{aligned}$$

In particular, we have  $[\langle \sigma q \rangle \downarrow] = 1$ .

Proposition 10 allows us to focus on branching processes. For  $X \in \Gamma$  and a finite sequence  $\sigma_1, \dots, \sigma_k$  with  $\sigma_i \in \Sigma$ , define  $|\sigma_1, \dots, \sigma_k|_X := |\{i \mid 1 \leq i \leq k, \sigma_i = X\}|$ ,



i.e., the number of  $X$ -symbols in the sequence. We define the *characteristic matrix*  $A \in \mathbb{R}^{\Gamma \times \Gamma}$  of a branching process by setting

$$A_{X,Y} := \sum_{X \xrightarrow{p} \langle \sigma_1 \sigma_2 \sigma_3 \rangle} p \cdot |\sigma_1, \sigma_2, \sigma_3|_Y + \sum_{X \xrightarrow{p} \langle \sigma_1 \sigma_2 \rangle} p \cdot |\sigma_1, \sigma_2|_Y + \sum_{X \xrightarrow{p} \sigma_1} p \cdot |\sigma_1|_Y.$$

It is easy to see that the  $(X, Y)$ -entry of  $A$  is the expected number of  $Y$ -processes after the first step, if starting in a single  $X$ -process. If  $S$  is a branching process and  $X_0 \in \Gamma$ , we call the pair  $(S, X_0)$  a *reduced branching process*, if for all  $X \in \Gamma$  there is  $i \in \mathbb{N}$  such that  $(A^i)_{X_0, X} > 0$ . Intuitively,  $(S, X_0)$  is reduced, if, starting in  $X_0$ , all process symbols can be reached with positive probability. If  $(S, X_0)$  is not reduced, it is easy to reduce it in polynomial time by eliminating all non-reachable process symbols.

The following proposition characterises the finiteness of both expected work and expected time in terms of the spectral radius  $\rho(A)$  of  $A$ . (Recall that  $\rho(A)$  is the largest absolute value of the eigenvalues of  $A$ .)

**Proposition 11.** *Let  $(S, X_0)$  be a reduced branching process. Let  $A$  be the associated characteristic matrix. Then the following statements are equivalent:*

- (1)  $\mathbb{E}\mathbf{W}_{X_0}$  is finite;      (2)  $\mathbb{E}\mathbf{T}_{X_0}$  is finite;      (3)  $\rho(A) < 1$ .

Further, if  $\mathbb{E}\mathbf{W}_{X_0}$  is finite, then it equals the  $X_0$ -component of  $(I - A)^{-1} \cdot \mathbf{1}$ , where  $I$  is the identity matrix, and  $\mathbf{1}$  is the column vector with all ones.

Statements similar to Proposition 11 do appear in the standard branching process literature [18, 2], however, not explicitly enough to cite directly or with stronger assumptions<sup>1</sup>. Our proof adapts a technique which was developed in [5] for a different purpose. It uses only basic tools and Perron-Frobenius theory, the spectral theory of nonnegative matrices. Proposition 11 has the following consequence:

**Corollary 12.** *Consider a branching process with process symbols  $\Gamma$  and  $X_0 \in \Gamma$ . Then  $\mathbb{E}\mathbf{W}_{X_0}$  and  $\mathbb{E}\mathbf{T}_{X_0}$  are both finite or both infinite. Distinguishing between those cases is in P.*

By combining the previous results we obtain the following theorem.

**Theorem 13.** *Consider a pSJS  $S$  with alphabet  $\Sigma = \Gamma \cup Q$ . Let  $a \in \Gamma$ . Then  $\mathbb{E}\mathbf{W}_a$  and  $\mathbb{E}\mathbf{T}_a$  are both finite or both infinite. Distinguishing between those cases is in PSPACE, and PosSLP-hard even for pPDSs. Further, if  $S$  is normalised and  $\mathbb{E}\mathbf{W}_a$  is finite, one can efficiently express  $\mathbb{E}\mathbf{W}_a$ .*

Theorem 13 can be interpreted as saying that, although the pSJS model does not impose a bound on the number of active processes at a time, its parallelism *cannot* be used to do an infinite expected amount of work in a finite expected time. However, the “speedup”  $\mathbb{E}[\mathbf{W}] / \mathbb{E}[\mathbf{T}]$  may be unbounded:

**Proposition 14.** *Consider the family of branching processes with transitions  $X \xrightarrow{p} \langle XX \rangle$  and  $X \xrightarrow{1-p} \perp$ , where  $0 < p < 1/2$ . Then the ratio  $\mathbb{E}[\mathbf{W}_X] / \mathbb{E}[\mathbf{T}_X]$  is unbounded for  $p \rightarrow 1/2$ .*

<sup>1</sup> For example, [2] assumes that there is  $n \in \mathbb{N}$  such that  $A^n$  is positive in all entries, a restriction which is not natural for our setting.

## 4 Case Studies

We have implemented a prototype tool in the form of a Maple worksheet, which allows to compute some of the quantities from the previous section: the termination probabilities, and distributions and expectations of work and time. In this section, we use our tool for two case studies<sup>2</sup>, which also illustrate how probabilistic parallel programs can be modelled with pSJSs. We only deal with normalised pSJSs in this section.

### 4.1 Divide and Conquer

The pSJS model lends itself to analyse parallel divide-and-conquer programs. For simplicity, we assume that the problem is already given as a binary tree, and solving it means traversing the tree and combining the results of the children. Figure 2 shows generic parallel code for such a problem.

```

function divCon(node)
  if node.leaf() then return node.val()
  else parallel < val1 := divCon(node.c1), val2 := divCon(node.c2) >
    return combine(val1, val2)

```

**Fig. 2.** A generic parallel divide-and-conquer program.

For an example, think of a routine for numerically approximating an integral  $\int_0^1 f(x) dx$ . Given the integrand  $f$  and a subinterval  $I \subseteq [0, 1]$ , we assume that there is a function which computes  $osc_f(I) \in \mathbb{N}$ , the “oscillation” of  $f$  in the interval  $I$ , a measure for the need for further refinement. If  $osc_f(I) = 0$ , then the integration routine returns the approximation  $1 \cdot f(1/2)$ , otherwise it returns  $I_1 + I_2$ , where  $I_1$  and  $I_2$  are recursive approximations of  $\int_0^{1/2} f(x) dx$  and  $\int_{1/2}^1 f(x) dx$ , respectively.<sup>3</sup>

We analyse such a routine using probabilistic assumptions on the integrand: Let  $n, n_1, n_2$  be nonnegative integers such that  $0 \leq n_1 + n_2 \leq n$ . If  $osc_f([a, b]) = n$ , then  $osc_f([a, (a+b)/2]) = n_1$  and  $osc_f([(a+b)/2, b]) = n_2$  with probability  $x(n, n_1, n_2) := \binom{n}{n_1} \cdot \binom{n-n_1}{n_2} \cdot \left(\frac{p}{2}\right)^{n_1} \cdot \left(\frac{p}{2}\right)^{n_2} \cdot (1-p)^{n-n_1-n_2}$ , where  $0 < p < 1$  is some parameter.<sup>4</sup> Of course, other distributions could be used as well. The integration routine can then be modelled by the pSJS with  $Q = \{q\}$  and  $\Gamma = \{\langle qq \rangle, 0, \dots, n_{\max}\}$  and the following rules:

$$0 \xrightarrow{1} q \quad \text{and} \quad \langle qq \rangle \xrightarrow{1} q \quad \text{and} \quad n \xrightarrow{x(n, n_1, n_2)} \langle n_1 n_2 \rangle \quad \text{for all } 1 \leq n \leq n_{\max} ,$$

<sup>2</sup> Available at <http://www.comlab.ox.ac.uk/people/Stefan.Kiefer/case-studies.mws>.

<sup>3</sup> Such an adaptive approximation scheme is called “local” in [21].

<sup>4</sup> That means, the oscillation  $n$  in the interval  $[a, b]$  can be thought of as distributed between  $[a, (a+b)/2]$  and  $[(a+b)/2, b]$  according to a ball-and-urn experiment, where each of the  $n$  balls is placed in the  $[a, (a+b)/2]$ -urn and the  $[(a+b)/2, b]$ -urn with probability  $p/2$ , respectively, and in a trash urn with probability  $1-p$ .

where  $0 \leq n_1 + n_2 \leq n$ . (Since we are merely interested in the performance of the algorithm, we can identify all return values with a single synchronisation state  $q$ .)

Using our prototype, we computed  $\mathbb{E}[\mathbf{W}_n]$  and  $\mathbb{E}[\mathbf{T}_n]$  for  $p = 0.8$  and  $n = 0, 1, \dots, 10$ . Figure 3 shows that  $\mathbb{E}[\mathbf{W}_n]$  increases faster with  $n$  than  $\mathbb{E}[\mathbf{T}_n]$ ; i.e., the parallelism increases.




Fig. 3. Expectations of time and work.

## 4.2 Evaluation of Game Trees

The evaluation of game trees is a central task of programs that are equipped with “artificial intelligence” to play games such as chess. These game trees are min-max trees (see Figure 4): each node corresponds to a position of the game, and each edge from a




Fig. 4. A game tree with value 3.

parent to a child corresponds to a move that transforms the position represented by the parent to a child position. Since the players have opposing objectives, the nodes alternate between max-nodes and min-nodes (denoted  $\vee$  and  $\wedge$ , respectively). A leaf of a game tree corresponds either to a final game position or to a position which is evaluated heuristically by the game-playing program; in both cases, the leaf is assigned a number.

Given such a leaf labelling, a number can be assigned to each node in the tree in the straightforward way; in particular, *evaluating a tree* means computing the root value.

In the following, we assume for simplicity that each node is either a leaf or has exactly three children. Figure 5 shows a straightforward recursive parallel procedure for evaluating a max-node of a game tree. (Of course, there is a symmetrical procedure for min-nodes.)

```

function parMax(node)
  if node.leaf() then return node.val()
  else parallel ⟨ val1 := parMin(node.c1), val2 := parMin(node.c2), val3 := parMin(node.c3) ⟩
    return max{val1, val2, val3}

```

**Fig. 5.** A simple parallel program for evaluating a game tree.

Notice that in Figure 4 the value of the root is 3, independent of some missing leaf values. Game-playing programs aim at evaluating a tree as fast as possible, possibly by not evaluating nodes which are irrelevant for the root value. The classic technique is called *alpha-beta pruning*: it maintains an interval  $[\alpha, \beta]$  in which the value of the current node is to be determined exactly. If the value turns out to be below  $\alpha$  or above  $\beta$ , it is safe to return  $\alpha$  or  $\beta$ , respectively. This may be the case even before all children have been evaluated (a so-called *cut-off*). Figure 6 shows a sequential program for alpha-beta pruning, initially to be called “seqMax(root,  $-\infty$ ,  $+\infty$ )”. Applying seqMax to the tree from Figure 4 results in several cut-offs: all non-labelled leaves are pruned.

```

function seqMax(node,  $\alpha$ ,  $\beta$ )
  if node.leaf() then if node.val()  $\leq \alpha$  then return  $\alpha$ 
    elseif node.val()  $\geq \beta$  then return  $\beta$ 
    else return node.val()
  else val1 := seqMin(node.c1,  $\alpha$ ,  $\beta$ )
    if val1 =  $\beta$  then return  $\beta$ 
    else val2 := seqMin(node.c2, val1,  $\beta$ )
      if val2 =  $\beta$  then return  $\beta$ 
      else return seqMin(node.c3, val2,  $\beta$ )

```

**Fig. 6.** A sequential program for evaluating a game tree using alpha-beta pruning.

Although alpha-beta pruning may seem inherently sequential, parallel versions have been developed, often involving the *Young Brothers Wait (YBW)* strategy [17]. It relies on a good ordering heuristic, i.e., a method that sorts the children of a max-node (resp. min-node) in increasing (resp. decreasing) order, without actually evaluating the children. Such an ordering heuristic is often available, but usually not perfect. The tree in Figure 4 is ordered in this way. If alpha-beta pruning is performed on such an ordered tree, then either all children of a node are evaluated or only the first one. The YBW

method first evaluates the first child only and hopes that this creates a cut-off or, at least, decreases the interval  $[\alpha, \beta]$ . If the first child fails to cause a cut-off, YBW *speculates* that both “younger brothers“ need to be evaluated, which can be done in parallel without wasting work. A wrong speculation may affect the performance, but not the correctness. Figure 7 shows a YBW-based program. Similar code is given in [8] using *Cilk*, a C-based parallel programming language.

```

function YBWMax(node,  $\alpha$ ,  $\beta$ )
  if node.leaf() then if node.val()  $\leq \alpha$  then return  $\alpha$ 
    elseif node.val()  $\geq \beta$  then return  $\beta$ 
    else return node.val()
  else val1 := YBWMin(node.c1,  $\alpha$ ,  $\beta$ )
    if val1 =  $\beta$  then return  $\beta$ 
    else parallel  $\langle$  val2 := YBWMin(node.c2, val1,  $\beta$ ), val3 := YBWMin(node.c3, val1,  $\beta$ )  $\rangle$ 
      return max{val2, val3}

```

**Fig. 7.** A parallel program based on YBW for evaluating a game tree.


We evaluate the performance of these three (deterministic) programs using probabilistic assumptions about the game trees. More precisely, we assume the following: Each node has exactly three children with probability  $p$ , and is a leaf with probability  $1-p$ . A leaf (and hence any node) takes as value a number from  $\mathbb{N}_4 := \{0, 1, 2, 3, 4\}$ , according to a distribution described below. In order to model an ordering heuristic on the children, each node carries a parameter  $e \in \mathbb{N}_4$  which intuitively corresponds to its expected value. If a max-node with parameter  $e$  has children, then they are min-nodes with parameters  $e, e\ominus 1, e\ominus 2$ , respectively, where  $a\ominus b := \max\{a-b, 0\}$ ; similarly, the children of a min-node with parameter  $e$  are max-nodes with parameters  $e, e\oplus 1, e\oplus 2$ , where  $a\oplus b := \min\{a+b, 4\}$ . A leaf-node with parameter  $e$  takes value  $k$  with probability  $\binom{4}{k} \cdot (e/4)^k \cdot (1-e/4)^{4-k}$ ; i.e., a leaf value is binomially distributed with expectation  $e$ . One could think of a game tree as the terminal tree of a branching process with  $\Gamma = \{Max(e), Min(e) \mid e \in \{0, \dots, 4\}\}$  and  $Q = \mathbb{N}_4$  and the rules  $Max(e) \xrightarrow{p} \langle Min(e) Min(e\ominus 1) Min(e\ominus 2) \rangle$  and  $Max(e) \xrightarrow{x(k)} k$ , with  $x(k) := (1-p) \cdot \binom{4}{k} \cdot (e/4)^k \cdot (1-e/4)^{4-k}$  for all  $e, k \in \mathbb{N}_4$ , and similar rules for  $Min(e)$ .

We model the YBW-program from Figure 7 running on such random game trees by the pSJS with  $Q = \{0, 1, 2, 3, 4, q(\vee), q(\wedge)\} \cup \{q(\alpha, \beta, \vee, e), q(\alpha, \beta, \wedge, e) \mid 0 \leq \alpha < \beta \leq 4, 0 \leq e \leq 4\} \cup \{q(a, b) \mid 0 \leq a, b \leq 4\}$  and the following rules:

$$\begin{aligned}
Max(\alpha, \beta, e) &\xrightarrow{x(0)+\dots+x(\alpha)} \alpha, & Max(\alpha, \beta, e) &\xrightarrow{x(\beta)+\dots+x(4)} \beta, & Max(\alpha, \beta, e) &\xrightarrow{x(k)} k \\
Max(\alpha, \beta, e) &\xrightarrow{p} \langle Min(\alpha, \beta, e) \ q(\alpha, \beta, \vee, e\ominus 1) \rangle \\
\langle \beta \ q(\alpha, \beta, \vee, e) \rangle &\xrightarrow{1} \beta, & \langle \gamma \ q(\alpha, \beta, \vee, e) \rangle &\xrightarrow{1} \langle Max2(\gamma, \beta, e) \ q(\vee) \rangle \\
Max2(\alpha, \beta, e) &\xrightarrow{1} \langle Min(\alpha, \beta, e) \ Min(\alpha, \beta, e\ominus 1) \rangle \\
\langle a \ b \rangle &\xrightarrow{1} q(a, b), & \langle q(a, b) \ q(\vee) \rangle &\xrightarrow{1} \max\{a, b\},
\end{aligned}$$

where  $0 \leq \alpha \leq \gamma < \beta \leq 4$  and  $\alpha < k < \beta$  and  $0 \leq e \leq 4$  and  $0 \leq a, b \leq 4$ . There are analogous rules with *Min* and *Max* exchanged. Notice that the rules closely follow the program from Figure 7. The programs *parMax* and *seqMax* from Figures 5 and 6 can be modelled similarly.

Let  $T(\text{YBW}, p) := \mathbb{E} [\mathbf{T}_{\text{Max}(0,4,2)} \mid \text{Run}(\text{Max}(0,4,2) \downarrow 2)]$ ; i.e.,  $T(\text{YBW}, p)$  is the expected time of the YBW-program called with a tree with value 2 and whose root is a max-node with parameter 2. (Recall that  $p$  is the probability that a node has children.) Let  $W(\text{YBW}, p)$  defined similarly for the expected work, and define these numbers also for *par* and *seq* instead of YBW, i.e., for the programs from Figures 5 and 6. Using our prototype we computed  $W(\text{seq}, p) = 1.00, 1.43, 1.96, 2.63, 3.50, 4.68, 6.33$  for  $p = 0.00, 0.05, 0.10, 0.15, 0.20, 0.25, 0.30$ . Since the program *seq* is sequential, we have the same sequence for  $T(\text{seq}, p)$ . To assess the speed of the parallel programs *par* and YBW, we also computed the percentaged increase of their runtime relative to *seq*, i.e.,  $100 \cdot (T(\text{par}, p)/T(\text{seq}, p) - 1)$ , and similarly for YBW. Figure 8 shows the results. One



**Fig. 8.** Percentaged runtime increase of *par* and YBW relative to *seq*.

can observe that for small values of  $p$  (i.e., small trees), the program *par* is slightly faster than *seq* because of its parallelism. For larger values of  $p$ , *par* still evaluates all nodes in the tree, whereas *seq* increasingly benefits from cut-offs of potentially deep branches. Using Proposition 11, one can prove  $W(\text{par}, \frac{1}{3}) = T(\text{par}, \frac{1}{3}) = \infty > W(\text{seq}, \frac{1}{3})$ .<sup>5</sup> The figure also shows that the YBW-program is faster than *seq*: the advantage of YBW increases with  $p$  up to about 10%.

We also compared the *work* of YBW with *seq*, and found that the percentaged increase ranges from 0 to about +0.4% for  $p$  between 0 and 0.3. This means that YBW wastes almost no work; in other words, a sequential version of YBW would be almost as fast as *seq*. An interpretation is that the second child rarely causes large cut-offs. Of

<sup>5</sup> In fact,  $W(\text{seq}, p)$  is finite even for values of  $p$  which are slightly larger than  $\frac{1}{3}$ ; in other words, *seq* cuts off infinite branches.

course, all of these findings could depend on the exact probabilistic assumptions on the game trees.

## 5 Conclusions and Future Work

We have introduced pSJSs, a model for probabilistic parallel programs with process spawning and synchronisation. We have studied the basic performance measures of termination probability, space, work, and time. In our results the upper complexity bounds coincide with the best ones known for pPDSs, and the lower bounds also hold for pPDSs. This suggests that analysing pSJSs is no more expensive than analysing pPDSs. The pSJS model is amenable to a practical performance analysis. Our two case studies have demonstrated the modelling power of pSJSs: one can use pSJSs to model, analyse, and compare the performance of parallel programs under probabilistic assumptions.

We intend to develop model-checking algorithms for pSJSs. It seems to us that a meaningful functional analysis should not only model-check the Markov chain induced by the pSJS, but rather take the individual process “histories” into account.

*Acknowledgements.* We thank Javier Esparza, Alastair Donaldson, Markus Müller-Olm, Luke Ong and Thomas Wahl for helpful discussions on the non-probabilistic version of pSJSs. We also thank the anonymous referees for valuable comments.

## References

1. E. Allender, P. Bürgisser, J. Kjeldgaard-Pedersen, and P. B. Miltersen. On the complexity of numerical analysis. In *IEEE Conf. on Computational Complexity*, pages 331–339, 2006.
2. K.B. Athreya and P.E. Ney. *Branching Processes*. Springer, 1972.
3. A. Berman and R.J. Plemmons. *Nonnegative matrices in the mathematical sciences*. Academic Press, 1979.
4. A. Bouajjani, M. Müller-Olm, and T. Touili. Regular symbolic analysis of dynamic networks of pushdown systems. In *Proceedings of CONCUR’05*, pages 473–487. Springer, 2005.
5. T. Brázdil, J. Esparza, and S. Kiefer. On the memory consumption of probabilistic pushdown automata. In *Proceedings of FSTTCS*, pages 49–60, 2009.
6. T. Brázdil, S. Kiefer, A. Kučera, and I.H. Vařeková. Runtime analysis of probabilistic programs with unbounded recursion. 2010. Submitted for publication. Available at <http://arxiv.org/abs/1007.1710>.
7. J. Canny. Some algebraic and geometric computations in PSPACE. In *STOC’88*, pages 460–467, 1988.
8. D. Dailey and C.E. Leiserson. Using Cilk to write multiprocessor chess programs. *The Journal of the International Computer Chess Association*, 2002.
9. J.S. Danaher, I.A. Lee, and C.E. Leiserson. Programming with exceptions in JCilk. *Science of Computer Programming (SCP)*, 63(2):147–171, 2006.
10. J. Esparza and K. Etessami. Verifying probabilistic procedural programs. In *FSTTCS 2004*, pages 16–31, 2004.
11. J. Esparza, S. Kiefer, and M. Luttenberger. Computing the least fixed point of positive polynomial systems. *SIAM Journal on Computing*, 39(6):2282–2335, 2010.
12. J. Esparza, A. Kučera, and R. Mayr. Model checking probabilistic pushdown automata. In *LICS’04*, pages 12–21. IEEE, 2004.

13. J. Esparza, A. Kučera, and R. Mayr. Quantitative analysis of probabilistic pushdown automata: Expectations and variances. In *LICS'05*, pages 117–126. IEEE, 2005.
14. K. Etessami and M. Yannakakis. Algorithmic verification of recursive probabilistic state machines. In *TACAS 2005*, pages 253–270, 2005.
15. K. Etessami and M. Yannakakis. Recursive concurrent stochastic games. *Logical Methods in Computer Science*, 4(4), 2008.
16. K. Etessami and M. Yannakakis. Recursive markov chains, stochastic grammars, and monotone systems of nonlinear equations. *Journal of the ACM*, 56(1):1–66, 2009.
17. R. Feldmann, B. Monien, P. Mysliwicz, and O. Vornberger. Distributed game-tree search. *ICCA Journal*, 12(2):65–73, 1989.
18. T.E. Harris. *The Theory of Branching Processes*. Springer, 1963.
19. R.A. Horn and C.A. Johnson. *Matrix Analysis*. Cambridge University Press, 1985.
20. C. Löding. Reachability problems on regular ground tree rewriting graphs. *Theory of Computing Systems*, 39:347–383, 2006.
21. M.A. Malcolm and R.B. Simpson. Local versus global strategies for adaptive quadrature. *ACM Transactions on Mathematical Software*, 1(2):129–146, 1975.
22. R. Mayr. Process rewrite systems. *Information and Computation*, 156(1-2):264–286, 2000.
23. J. Renegar. On the computational complexity and geometry of the first-order theory of the reals. Parts I–III. *Journal of Symbolic Computation*, 13(3):255–352, 1992.
24. E. Seneta. On asymptotic properties of sub-critical branching processes. *Journal of the Australian Mathematical Society*, 8:671–682, 1968.
25. F. Spitzer. *Principles of Random Walk*. Springer, 1976.
26. M. Yannakakis and K. Etessami. Checking LTL properties of recursive Markov chains. In *QEST 2005*, pages 155–165, 2005.



## A Proofs of Section 3.1

Here is a restatement of Proposition 4.

**Proposition 4.** *There is a probability-preserving bijection between the runs  $Run(\sigma \downarrow q)$  in  $M_S$  and the runs  $Run(\Box \sigma \downarrow \bar{q})$  in  $M_{S_P}$ . In particular, we have  $[\sigma \downarrow q] = [\Box \sigma \downarrow \bar{q}]$ .*

*Proof.* We define a bijection  $b : Run[M_S](\sigma \downarrow q) \rightarrow Run[M_{S_P}](\Box \sigma \downarrow \bar{q})$ . Since these sets only contain runs that reach a terminal state (namely,  $q$  and  $\bar{q}$ , respectively), we identify in this proof a run with the (finite) path that leads to the terminal state. For a path  $w$  in  $M_{S_P}$  with length  $n$  we write  $w \upharpoonright \sigma$  for the path  $z$  of length  $n$  with  $z(i) = w(i)\sigma$  for all  $0 \leq i \leq n-1$ .

Let  $w \in Run[M_S](\sigma \downarrow q)$ . We define  $b(w)$  inductively by order of the length  $n$  of  $w$ . The run  $w$  has one of the following forms:

- Let  $w = w(0)$ , where  $w(0) = \sigma = q$ . Then we set  $b(w) := \Box q, \bar{q}$ .
- Let  $w = a, \langle \sigma_1 \sigma_2 \rangle, \dots, w(k), \dots, q$ , where  $a \in \Gamma$  and  $w(k) = \langle q_1 q_2 \rangle$  for some  $k$  with  $1 \leq k \leq n-2$ , such that  $w(i) \notin \Gamma$  for  $i \in \{1, \dots, k-1\}$ . Consider the left and the right subtree of the root nodes in the run between  $w(1) = \langle \sigma_1 \sigma_2 \rangle$  and  $w(k) = \langle q_1 q_2 \rangle$ . By the semantics of pSJSs, there are corresponding runs  $u_1 = \sigma_1, \dots, q_1$  and  $u_2 = \sigma_2, \dots, q_2$  such that  $u_1 \in Run[M_S](\sigma_1 \downarrow q_1)$  and  $u_2 \in Run[M_S](\sigma_2 \downarrow q_2)$  and  $\max\{|u_1|, |u_2|\} = k$ . Let  $z_1 := b(u_1)$  and  $z_2 := b(u_2)$  and  $z_3 := b(w(k), \dots, w(n-1))$ . Then we set  $b(w) := \Box a, z_1 \upharpoonright \sigma_2, z_2 \upharpoonright \bar{q}_1, z_3$ .
- Let  $w = a, \sigma_1, \dots, q$ , where  $a \in \Gamma$  and  $\sigma_1 \in \Sigma \setminus \langle QQ \rangle$ . Then we set  $b(w) := \Box a, b(\sigma_1, \dots, q)$ .

Notice that in all cases  $b(w) \in Run[M_{S_P}](\Box \sigma \downarrow \bar{q})$ . It is easy to check that  $b$  is a bijection. It is also easy to see that  $\mathcal{P}(\{w\}) = \mathcal{P}(\{b(w)\})$ , because in both cases the probability is the product of the probabilities of the applied transition rules (of course, taking multiplicities into account).  $\square$

## B Proofs of Section 3.2

### B.1 Proof of Proposition 5

Here is a restatement of Proposition 5.

**Proposition 5.** *Let  $\sigma \in \Sigma$  and  $q \in Q$ . Then  $[\sigma \downarrow q]$  is the value for  $\llbracket \sigma \downarrow q \rrbracket$  in the least (w.r.t. componentwise ordering) nonnegative solution of the above equation system.*

*Proof.* The proposition can be proved by adapting the corresponding proofs for pPDSs from [12] or [16]. Alternatively, we can use the “serialisation” procedure from Section 3.1 and the equality  $[\sigma \downarrow q] = [\Box \sigma \downarrow \bar{q}]$ , reducing the problem from pSJSs to pPDSs. Then we can take the equation system from [12, 16] for  $[\Box \sigma \downarrow \bar{q}]$  and compress it by substituting variables with the right-hand side of their equation. This gives the same equation system as the one above.  $\square$

## B.2 Proof of Theorem 6

Here is a restatement of Theorem 6.

**Theorem 6.** *Consider a pSJS with alphabet  $\Sigma = \Gamma \cup Q$ . Let  $\sigma \in \Sigma$  and  $q \in Q$ . Then (1) one can efficiently express (in the sense defined in Section 2) the value of  $[\sigma \downarrow q]$ , (2) deciding whether  $[\sigma \downarrow q] = 0$  is in P, and (3) deciding whether  $[\sigma \downarrow q] < 1$  is PosSLP-hard even for pPDSs.*

*Proof.* The respective claims are shown for pPDSs in [14, 16]. For statements (1) and (2), we use Proposition 4 to reduce the problem to pPDSs. For statement (3), recall from Section 3.1 that pPDSs can be encoded as pSJSs.  $\square$

## C Proofs of Section 3.3

### C.1 Proof of Lemma 7

Here is a restatement of Lemma 7.

**Lemma 7.** *The set  $U$  can be computed in polynomial time.*

*Proof.* We write  $t_1 \rightarrow^* t_2$  if  $t_2$  can be reached from  $t_1$  in the Markov chain  $M_S$  induced by the pSJS  $S$ ; i.e.,  $\rightarrow^*$  is the reflexive and transitive closure of  $\rightarrow$ . Define  $\Rightarrow := \rightarrow^* \cap (\Sigma \times \Sigma)$ ; i.e., we have  $\sigma_1 \Rightarrow \sigma_2$  if and only if  $\sigma_2 \in \Sigma$  can be reached from  $\sigma_1 \in \Sigma$ . The relation  $\Rightarrow$  can be computed in polynomial time using the fact that it is the smallest subset of  $\Sigma \times \Sigma$  that satisfies:

- $\sigma \Rightarrow \sigma$  for all  $\sigma \in \Sigma$ ;
- $\sigma_1 \hookrightarrow \sigma_2 \Rightarrow \sigma_3$  implies  $\sigma_1 \Rightarrow \sigma_3$ ;
- $\sigma_1 \hookrightarrow \langle \sigma_2 \sigma_3 \rangle$  and  $[\sigma_2 \downarrow q_2] > 0$  and  $[\sigma_3 \downarrow q_3] > 0$  and  $\langle q_2 q_3 \rangle \Rightarrow \sigma_4$  imply  $\sigma_1 \Rightarrow \sigma_4$ .

For a tree  $t \in T(\Sigma)$ , let  $h(t)$  denote its *height*, i.e., the maximal distance of a leaf to the root. Moreover, we define for each  $k \in \mathbb{N}$ :

$$U_k := \{\sigma \in \Sigma \mid \text{there is a tree } t \in T(\Sigma) \text{ with } \sigma \rightarrow^* t \text{ and } h(t) \geq k\}.$$

Notice that  $\Sigma = U_0 \supseteq U_1 \supseteq U_2 \dots$ . It is easy to see that we have for all  $k \in \mathbb{N}$ :

$$U_{k+1} = \{a \in \Gamma \mid \exists \text{ transition } b \hookrightarrow \langle \sigma_1 \sigma_2 \rangle : a \Rightarrow b \text{ and } \{\sigma_1, \sigma_2\} \cap U_k \neq \emptyset\}.$$

It follows from this characterisation that the sequence  $U_0 \supseteq U_1 \supseteq U_2 \dots$  stabilises after at most  $|\Sigma|$  steps; i.e., we have  $U_{|\Sigma|} = U_{|\Sigma|+1} = \dots$ . In other words, we have:

$$U_{|\Sigma|} = \{a \in \Gamma \mid \forall n \in \mathbb{N} : \exists t \in T(\Sigma) : a \rightarrow^* t \text{ and } h(t) > n\}.$$

For each tree  $t \in T(\Sigma)$  we have  $h(t) \leq |t| \leq 2^{h(t)}$  (recall that  $|t|$  is the length of  $t$  not counting the symbols ' $\langle$ ' and ' $\rangle$ '). Therefore, we have  $U = U_{|\Sigma|}$ , so it suffices to compute  $U_{|\Sigma|}$ , which can be done in polynomial time with the above characterisation.  $\square$

## C.2 Proof of Proposition 8

Here is a restatement of Proposition 8.

**Proposition 8.** (1) The pSJS  $\overline{S}$  is normalised; (2) the value  $[a \downarrow q]$  for  $a \in \Gamma$  and  $q \in Q$  is the same in  $S$  and  $\overline{S}$ ; (3) we have  $\mathcal{P}(\mathbf{S}_a < \infty = \mathbf{T}_a \mid \text{Run}(a)) = [a \downarrow \overline{q}]$  for all  $a \in \Gamma$ .

*Proof.* Statement (1) follows from the fact that  $\overline{S}$  has, by construction, for all  $q_1, q_2 \in \overline{Q}$  a transition with  $\langle q_1, q_2 \rangle$  on the left hand side. (In other words,  $\overline{T} \supseteq \overline{Q} \times \overline{Q}$ .)

For statement (2) observe that  $\overline{S}$  is obtained from  $S$  via two steps: a normalisation step, and a step where transitions with process symbols  $b \in B$  on the left hand side are replaced. We argue that neither of those steps modifies the value  $[a \downarrow q]$  for  $a \in \Gamma$  and  $q \in Q$ . The first step does not modify  $[a \downarrow q]$ , because the normalisation affects only those runs that would not terminate in a synchronisation state without normalisation. With normalisation, those runs may terminate in the new synchronisation state introduced by the normalisation, but not in  $q$ . The second step does not modify  $[a \downarrow q]$ , because it affects only those runs that would otherwise not terminate. With the modification in the second step, those runs may terminate in  $\overline{q}$ , but not in  $q$ .

For statement (3), it is convenient to consider a version of  $S$  “in between” the first and the second modification step. More precisely, let  $S^{\text{norm}}$  denote the pSJS after the normalisation step, and let  $S^{\text{stat}}$  denote the pSJS obtained from  $S^{\text{norm}}$  by removing all transitions with symbols  $b \in B$  on the left hand side and replacing them with a transition  $b \xrightarrow{1} b$ . Notice that  $\mathcal{P}(\mathbf{S}_a < \infty = \mathbf{T}_a)$  is the same in  $S$  and  $S^{\text{norm}}$  and  $S^{\text{stat}}$  for all  $a \in \Gamma$ . Denote by  $\text{Good}(a)$  the set of those runs  $w \in \text{Run}[M_{S^{\text{stat}}}] (a)$  that reach a bottom strongly connected component (BSCC) of  $M_{S^{\text{stat}}}$ . For any  $n \in \mathbb{N}$  there are only finitely many trees  $t$  of  $M_{S^{\text{stat}}}$  such that  $|t| \leq n$ . Hence it follows using standard arguments on finite Markov chains that in  $S^{\text{stat}}$  we have for all  $n \in \mathbb{N}$

$$\begin{aligned} \mathcal{P}(\mathbf{S}_a \leq n) &= \mathcal{P}(\mathbf{S}_a \leq n \text{ and } \text{Good}(a)) && , \text{ and so} \\ \mathcal{P}(\mathbf{S}_a \leq n \text{ and } \mathbf{T}_a = \infty) &= \mathcal{P}(\mathbf{S}_a \leq n \text{ and } \mathbf{T}_a = \infty \text{ and } \text{Good}(a)) && , \text{ and so} \\ \mathcal{P}(\mathbf{S}_a < \infty = \mathbf{T}_a) &= \mathcal{P}(\mathbf{S}_a < \infty = \mathbf{T}_a \text{ and } \text{Good}(a)) . && (1) \end{aligned}$$

Observe that each BSCC of  $M_{S^{\text{stat}}}$  consists of exactly one tree  $t$  such that  $\text{Front}(t)$  is either empty or consists only of elements of  $B$ . So there is a natural 1-to-1 correspondence between the runs of  $S^{\text{stat}}$  that satisfy “ $\mathbf{S}_a < \infty = \mathbf{T}_a$  and  $\text{Good}(a)$ ” and the runs of  $\overline{S}$  that are in  $\text{Run}(a \downarrow \overline{q})$ . With (1) we get  $\mathcal{P}(\mathbf{S}_a < \infty = \mathbf{T}_a) = [a \downarrow \overline{q}]$ .  $\square$

## C.3 Proof of Theorem 9

Here is a restatement of Theorem 9.

**Theorem 9.** Consider a pSJS with alphabet  $\Sigma = \Gamma \cup Q$  and  $a \in \Gamma$ . Let  $s := \mathcal{P}(\mathbf{S}_a < \infty)$ . Then (1) one can efficiently express  $s$ , (2) deciding whether  $s = 0$  is in  $\mathbf{P}$ , and (3) deciding whether  $s < 1$  is PosSLP-hard even for pPDSs.

*Proof.* It follows from Proposition 8 that  $s = \sum_{q \in \overline{Q}} [a \downarrow q]$ , hence  $s$  is expressible by Theorem 6 (1). Moreover, it follows using Theorem 6 (2) that deciding whether  $s = 0$  is in P. It remains to show statement (3).

We draw from a reduction in [16], where it is shown that deciding if a pPDS terminates with probability 1 is PosSLP-hard. As a gadget for that, Etesami and Yannakakis [16] compute, given a PosSLP instance, a pPDS  $S$  with the following properties. (More precisely, they construct an equivalent recursive Markov chain.) The starting configuration is  $qa$ , and after having left the initial configuration,  $S$  reaches a configuration of the form  $q\alpha$  with  $\alpha \in \Gamma^*$  again with probability 1. At that time, the configuration is  $qaa$  with some probability  $p$ , and  $q$  with probability  $1 - p$ . Moreover, the time (and hence space) needed to reach either of those configurations is essentially bounded by the size of the given PosSLP instance, so it is finite. Furthermore, the given PosSLP instance is a “yes instance” if and only if  $p > \frac{1}{2}$ . It is easy to see that  $\mathcal{P}(\mathbf{S}_{qa} = \infty) > 0$  in  $S$  if and only if  $\mathcal{P}(\mathbf{S}_{qa} = \infty) > 0$  in the pPDS  $\overline{S}$  that consists only of the transitions  $qa \xrightarrow{p} qaa$  and  $qa \xrightarrow{1-p} q$ . The Markov chain induced by  $\overline{S}$ , in turn, is isomorphic to the simple random walk  $X_0, X_1, \dots$  on  $\mathbb{N}$  with  $X_0 = 1$  and  $\mathcal{P}(X_{i+1} = 0 \mid X_i = 0) = 1$  and

$$\begin{aligned} \mathcal{P}(X_{i+1} = n + 1 \mid X_i = n > 0) &= p & \text{and} \\ \mathcal{P}(X_{i+1} = n - 1 \mid X_i = n > 0) &= 1 - p. \end{aligned}$$

It is well-known (see e.g. [25]) that for this random walk we have  $\mathcal{P}(\sup_{i \in \mathbb{N}} X_i = \infty) > 0$  if and only if  $p > \frac{1}{2}$ . It follows that we have  $\mathcal{P}(\mathbf{S}_{qa} < \infty) < 1$  in  $S$  if and only if  $p > \frac{1}{2}$ .  $\square$

## D Proofs of Section 3.4

### D.1 Proof of Proposition 10

Here is a restatement of Proposition 10.

**Proposition 10.** *Let  $\sigma \in \Sigma$  and  $q \in Q$  with  $[\sigma \downarrow q] > 0$ . Then*

$$\begin{aligned} \mathcal{P}(\mathbf{W}_\sigma = n \mid \text{Run}(\sigma \downarrow q)) &= \mathcal{P}(\mathbf{W}_{\langle \sigma q \rangle} = n \mid \text{Run}(\langle \sigma q \rangle)) & \text{for all } n \geq 0 & \text{ and} \\ \mathcal{P}(\mathbf{T}_\sigma \leq n \mid \text{Run}(\sigma \downarrow q)) &\leq \mathcal{P}(\mathbf{T}_{\langle \sigma q \rangle} \leq n \mid \text{Run}(\langle \sigma q \rangle)) & \text{for all } n \geq 0. \end{aligned}$$

*In particular, we have  $[\langle \sigma q \rangle \downarrow] = 1$ .*

*Proof.* Define, for  $q \in Q$  and  $\sigma \in \Sigma$ ,

$$\begin{aligned} D_{\sigma q}(n) &:= \mathcal{P}(\text{Run}(\sigma \downarrow q), \mathbf{W}_\sigma = n \mid \text{Run}(\sigma)) & \text{and} \\ D_{\langle \sigma q \rangle}(n) &:= \mathcal{P}(\mathbf{W}_{\langle \sigma q \rangle} = n \mid \text{Run}(\langle \sigma q \rangle)) . \end{aligned}$$

Notice that  $\langle \sigma q \rangle$  and thus  $D_{\langle \sigma q \rangle}(n)$  are undefined, if and only if  $[\sigma \downarrow q] = 0$ . For the arithmetical expressions in the rest of this proof, we assume  $0 \cdot \text{undefined} = 0$ . For the

statement on **W** in the proposition, it suffices to show  $D_{\sigma q}(n) = [\sigma \downarrow q] \cdot D_{\langle \sigma q \rangle}(n)$  for  $n \geq 0$ . We proceed by induction on  $n$ .

Let  $n = 0$ . If  $[\sigma \downarrow q] = 0$ , then  $D_{\sigma q}(0) = 0 = 0 \cdot \text{undefined} = [\sigma \downarrow q] \cdot D_{\langle \sigma q \rangle}(0)$ . If  $[\sigma \downarrow q] > 0$  and  $\sigma = q$ , then  $D_{\sigma q}(0) = 1 = 1 \cdot 1 = [\sigma \downarrow q] \cdot D_{\perp}(0) = [\sigma \downarrow q] \cdot D_{\langle \sigma q \rangle}(0)$ . If  $[\sigma \downarrow q] > 0$  and  $\sigma \in \Gamma$ , then  $D_{\sigma q}(0) = 0 = [\sigma \downarrow q] \cdot 0 = [\sigma \downarrow q] \cdot D_{\langle \sigma q \rangle}(0)$ .

Let  $n > 0$ . If  $[\sigma \downarrow q] = 0$ , then  $D_{\sigma q}(n) = 0 = 0 \cdot \text{undefined} = [\sigma \downarrow q] \cdot D_{\langle \sigma q \rangle}(n)$ . If  $[\sigma \downarrow q] > 0$  and  $\sigma = q$ , then  $D_{\sigma q}(n) = 0 = 1 \cdot 0 = [\sigma \downarrow q] \cdot D_{\langle \sigma q \rangle}(n)$ . If  $[\sigma \downarrow q] > 0$  and  $\sigma = a \in \Gamma$ , then

$$\begin{aligned}
D_{a q}(n) &= \sum_{\substack{a \xrightarrow{p} \langle \sigma_1 \sigma_2 \rangle \\ \langle q_1 q_2 \rangle \in \Gamma \cap \langle Q Q \rangle}} \sum_{\substack{i, j, k \in \mathbb{N} \\ 1+i+j+k=n}} p \cdot D_{\sigma_1 q_1}(i) \cdot D_{\sigma_2 q_2}(j) \cdot D_{\langle q_1 q_2 \rangle q}(k) \\
&\quad + \sum_{\substack{a \xrightarrow{p} \sigma' \\ \sigma' \in \Sigma \setminus \langle Q Q \rangle}} p \cdot D_{\sigma' q}(n-1) \\
&= \sum_{\substack{a \xrightarrow{p} \langle \sigma_1 \sigma_2 \rangle \\ \langle q_1 q_2 \rangle \in \Gamma \cap \langle Q Q \rangle}} \sum_{\substack{i, j, k \in \mathbb{N} \\ 1+i+j+k=n}} p \cdot [\sigma_1 \downarrow q_1] \cdot [\sigma_2 \downarrow q_2] \cdot [\langle q_1 q_2 \rangle \downarrow q] \cdot \\
&\quad \cdot D_{\langle \sigma_1 q_1 \rangle}(i) \cdot D_{\langle \sigma_2 q_2 \rangle}(j) \cdot D_{\langle \langle q_1 q_2 \rangle q \rangle}(k) \\
&\quad + \sum_{\substack{a \xrightarrow{p} \sigma' \\ \sigma' \in \Sigma \setminus \langle Q Q \rangle}} p \cdot [\sigma' \downarrow q] \cdot D_{\langle \sigma' q \rangle}(n-1) \\
&= [a \downarrow q] \cdot \left( \sum_{\langle a q \rangle \xrightarrow{p'} \langle \langle \sigma_1 q_1 \rangle \langle \sigma_2 q_2 \rangle \langle \langle q_1 q_2 \rangle q \rangle \rangle} \sum_{\substack{i, j, k \in \mathbb{N} \\ 1+i+j+k=n}} p' \cdot D_{\langle \sigma_1 q_1 \rangle}(i) \cdot D_{\langle \sigma_2 q_2 \rangle}(j) \cdot D_{\langle \langle q_1 q_2 \rangle q \rangle}(k) \right. \\
&\quad \left. + \sum_{\langle a q \rangle \xrightarrow{p'} \langle \sigma' q \rangle} p' \cdot D_{\langle \sigma' q \rangle}(n-1) \right) \\
&= [a \downarrow q] \cdot D_{\langle a q \rangle}(n),
\end{aligned}$$

where the first and the last equality are by the definition of **W** and the semantics of pSJSs, the second equality is by the induction hypothesis, and the third equality is by the definition of  $\bar{\Gamma}$ . This proves the statement on **W** in the proposition.

The proof of the statement on **T** is similar. Define, for  $q \in Q$  and  $\sigma \in \Sigma$ ,

$$\begin{aligned}
E_{\sigma q}(n) &:= \mathcal{P}(\text{Run}(\sigma \downarrow q), \mathbf{T}_{\sigma} \leq n \mid \text{Run}(\sigma)) \quad \text{and} \\
E_{\sigma q}^{\bar{\cdot}}(n) &:= \mathcal{P}(\text{Run}(\sigma \downarrow q), \mathbf{T}_{\sigma} = n \mid \text{Run}(\sigma)) \quad \text{and} \\
E_{\langle \sigma q \rangle}(n) &:= \mathcal{P}(\mathbf{T}_{\langle \sigma q \rangle} \leq n \mid \text{Run}(\langle \sigma q \rangle)) .
\end{aligned}$$

Notice that  $\langle \sigma q \rangle$  and thus  $E_{\langle \sigma q \rangle}(n)$  are undefined, if and only if  $[\sigma \downarrow q] = 0$ . It suffices to show  $E_{\sigma q}(n) \leq [\sigma \downarrow q] \cdot E_{\langle \sigma q \rangle}(n)$  for  $n \geq 0$ . We proceed by induction on  $n$ .

Let  $n = 0$ . If  $[\sigma \downarrow q] = 0$ , then  $E_{\sigma q}(0) = 0 = 0 \cdot \text{undefined} = [\sigma \downarrow q] \cdot E_{\langle \sigma q \rangle}(0)$ . If  $[\sigma \downarrow q] > 0$  and  $\sigma = q$ , then  $E_{\sigma q}(0) = 1 = 1 \cdot 1 = [\sigma \downarrow q] \cdot E_{\perp}(0) = [\sigma \downarrow q] \cdot E_{\langle \sigma q \rangle}(0)$ . If  $[\sigma \downarrow q] > 0$  and  $\sigma \in \Gamma$ , then  $E_{\sigma q}(0) = 0 = [\sigma \downarrow q] \cdot 0 = [\sigma \downarrow q] \cdot E_{\langle \sigma q \rangle}(0)$ .

Let  $n > 0$ . If  $[\sigma \downarrow q] = 0$ , then  $E_{\sigma q}(n) = 0 = 0 \cdot \text{undefined} = [\sigma \downarrow q] \cdot E_{\langle \sigma q \rangle}(n)$ . If  $[\sigma \downarrow q] > 0$  and  $\sigma = q$ , then  $E_{\sigma q}(n) = 1 = 1 \cdot 1 = [\sigma \downarrow q] \cdot E_{\langle \sigma q \rangle}(n)$ . If  $[\sigma \downarrow q] > 0$  and  $\sigma = a \in \Gamma$ , then

$$\begin{aligned}
E_{aq}(n) &= \sum_{\substack{a \xrightarrow{p} \langle \sigma_1 \sigma_2 \rangle \\ \langle q_1 q_2 \rangle \in \Gamma \cap \langle QQ \rangle}} \sum_{\substack{i, j, k \in \mathbb{N} \\ 1 + \max\{i, j\} + k \leq n}} p \cdot E_{\sigma_1 q_1}^{\overline{=}}(i) \cdot E_{\sigma_2 q_2}^{\overline{=}}(j) \cdot E_{\langle q_1 q_2 \rangle q}^{\overline{=}}(k) \\
&\quad + \sum_{\substack{a \xrightarrow{p} \sigma' \\ \sigma' \in \Sigma \setminus \langle QQ \rangle}} p \cdot E_{\sigma' q}(n-1) \\
&\leq \sum_{\substack{a \xrightarrow{p} \langle \sigma_1 \sigma_2 \rangle \\ \langle q_1 q_2 \rangle \in \Gamma \cap \langle QQ \rangle}} \sum_{\substack{i, j, k \in \mathbb{N} \\ 1 + \max\{i, j, k\} \leq n}} p \cdot E_{\sigma_1 q_1}^{\overline{=}}(i) \cdot E_{\sigma_2 q_2}^{\overline{=}}(j) \cdot E_{\langle q_1 q_2 \rangle q}^{\overline{=}}(k) \\
&\quad + \sum_{\substack{a \xrightarrow{p} \sigma' \\ \sigma' \in \Sigma \setminus \langle QQ \rangle}} p \cdot E_{\sigma' q}(n-1) \\
&= \sum_{\substack{a \xrightarrow{p} \langle \sigma_1 \sigma_2 \rangle \\ \langle q_1 q_2 \rangle \in \Gamma \cap \langle QQ \rangle}} p \cdot E_{\sigma_1 q_1}(n-1) \cdot E_{\sigma_2 q_2}(n-1) \cdot E_{\langle q_1 q_2 \rangle q}(n-1) \\
&\quad + \sum_{\substack{a \xrightarrow{p} \sigma' \\ \sigma' \in \Sigma \setminus \langle QQ \rangle}} p \cdot E_{\sigma' q}(n-1) \\
&\leq \sum_{\substack{a \xrightarrow{p} \langle \sigma_1 \sigma_2 \rangle \\ \langle q_1 q_2 \rangle \in \Gamma \cap \langle QQ \rangle}} p \cdot [\sigma_1 \downarrow q_1] \cdot [\sigma_2 \downarrow q_2] \cdot [\langle q_1 q_2 \rangle \downarrow q] \cdot \\
&\quad \cdot E_{\langle \sigma_1 q_1 \rangle}(n-1) \cdot E_{\langle \sigma_2 q_2 \rangle}(n-1) \cdot E_{\langle \langle q_1 q_2 \rangle q \rangle}(n-1) \\
&\quad + \sum_{\substack{a \xrightarrow{p} \sigma' \\ \sigma' \in \Sigma \setminus \langle QQ \rangle}} p \cdot [\sigma' \downarrow q] \cdot E_{\langle \sigma' q \rangle}(n-1) \\
&= [a \downarrow q] \cdot \left( \sum_{\langle aq \rangle \xrightarrow{p'} \langle \langle \sigma_1 q_1 \rangle \langle \sigma_2 q_2 \rangle \langle \langle q_1 q_2 \rangle q \rangle \rangle} p' \cdot E_{\langle \sigma_1 q_1 \rangle}(n-1) \cdot E_{\langle \sigma_2 q_2 \rangle}(n-1) \cdot E_{\langle \langle q_1 q_2 \rangle q \rangle}(n-1) \right. \\
&\quad \left. + \sum_{\langle aq \rangle \xrightarrow{p'} \langle \sigma' q \rangle} p' \cdot E_{\langle \sigma' q \rangle}(n-1) \right) \\
&= [a \downarrow q] \cdot E_{\langle aq \rangle}(n),
\end{aligned}$$

where the first and the last equality are by the definition of  $\mathbf{T}$  and the semantics of pSJSs, the first inequality and the second equality are trivial, the second inequality is by the induction hypothesis, and the third equality is by the definition of  $\overline{T}$ .

For the final statement, observe that

$$[(\sigma q)\downarrow] = \lim_{n \rightarrow \infty} \mathcal{P}(\mathbf{T}_{(\sigma q)} \leq n \mid \text{Run}((\sigma q))) \geq \lim_{n \rightarrow \infty} \mathcal{P}(\mathbf{T}_\sigma \leq n \mid \text{Run}(\sigma \downarrow q)) = 1.$$

□

## D.2 Proof of Proposition 11

Here is a restatement of Proposition 11.

**Proposition 11.** *Let  $(S, X_0)$  be a reduced branching process. Let  $A$  be the associated characteristic matrix. Then the following statements are equivalent:*

$$(1) \mathbb{E}\mathbf{W}_{X_0} \text{ is finite;} \quad (2) \mathbb{E}\mathbf{T}_{X_0} \text{ is finite;} \quad (3) \rho(A) < 1.$$

Further, if  $\mathbb{E}\mathbf{W}_{X_0}$  is finite, then it equals the  $X_0$ -component of  $(I - A)^{-1} \cdot \mathbf{1}$ , where  $I$  is the identity matrix, and  $\mathbf{1}$  is the column vector with all ones.

*Proof.* Clearly, statement (1) implies statement (2), because, by definition,  $\mathbf{W}_{X_0} \geq \mathbf{T}_{X_0}$ . Next, we prove that statement (3) implies statement (1). For each  $X \in \Gamma$  and  $i \in \mathbb{N}$ , we define a random variable  $z_X^{(i)}$  over  $\text{Run}(X_0)$  by setting  $z_X^{(i)}(w) := |\text{Front}(w(i))|_X$ ; i.e.,  $z_X^{(i)}$  is the number of active  $X$ -processes at time  $i$ . We assemble the  $z_X^{(i)}$  in a row vector  $\mathbf{z}^{(i)}$ . Note that  $|\text{Front}(w(i))| = \|\mathbf{z}^{(i)}(w)\|_1$ . It is easy to see (by induction, see also [2, p. 184]) that  $\mathbb{E}[\mathbf{z}^{(i)}] = e^{(X_0)} \cdot A^i$ , where by  $e^{(X_0)} \in \mathbb{N}^\Gamma$  we mean the row vector whose only nonzero component is the  $X_0$ -component, which is 1. Consequently, we have

$$\mathbb{E}\mathbf{W}_{X_0} = \sum_{i=0}^{\infty} \mathbb{E}[|\text{Front}(w(i))| \mid w \in \text{Run}(X_0)] = \sum_{i=0}^{\infty} \left\| \mathbb{E}[\mathbf{z}^{(i)}] \right\|_1 = \left\| e^{(X_0)} \cdot A^* \right\|_1,$$

where  $A^* := \sum_{i=0}^{\infty} A^i$ . It is known [19] that the matrix series  $A^*$  converges if and only if  $\rho(A) < 1$ . It follows that statement (3) implies statement (1). Further, it is known [19] that if  $\rho(A) < 1$ , then  $A^* = (I - A)^{-1}$ , so then  $\mathbb{E}\mathbf{W}_{X_0} = \left\| e^{(X_0)} \cdot (I - A)^{-1} \right\|_1$ , which equals the  $X_0$ -component of  $(I - A)^{-1} \cdot \mathbf{1}$ .

It remains to show that statement (2) implies statement (3). For this part, we rely on Perron-Frobenius theory. Let  $\rho(A) \geq 1$ . Call a matrix  $B \in \mathbb{R}^{n \times n}$  *strongly connected*, if for all  $1 \leq i, j \leq n$  there is  $k > 0$  such that  $(B^k)_{ij} \neq 0$ . Since  $A$  is nonnegative, Corollary 2.1.6 of [3] asserts that there exists a strongly connected principal submatrix  $A'$  of  $A$  such that  $\rho(A') \geq 1$ ; i.e., there is  $\Gamma' \subseteq \Gamma$  such that the matrix  $A' \in \mathbb{R}^{\Gamma' \times \Gamma'}$  obtained from  $A$  by deleting all rows and columns not indexed with elements of  $\Gamma'$  is strongly connected. We will show that  $\mathbb{E}\mathbf{T}_X = \infty$  for all  $X \in \Gamma'$ . Since  $(S, X_0)$  is reduced, this implies that  $\mathbb{E}\mathbf{T}_{X_0}$  is infinite. Therefore, to simplify the notation, we

assume in the following w.l.o.g. that  $A = A'$ , i.e.,  $A$  is strongly connected. We will show  $\mathbb{E}\mathbf{T}_X = \infty$  for all  $X \in \Gamma$ .

Define, for each  $i \in \{0, 1, 2, 3\}$ , a function  $L^{(i)} : (\mathbb{R}^\Gamma)^i \rightarrow \mathbb{R}^\Gamma$  as follows, where  $\mathbf{y}, \mathbf{z}, \mathbf{w} \in \mathbb{R}^\Gamma$  are column vectors (we use subscripts to refer to indices):

$$\begin{aligned} L_X^{(0)} &:= \sum_{X \xrightarrow{p} \perp} p; & L^{(2)}(\mathbf{y}, \mathbf{z})_X &:= \sum_{X \xrightarrow{p} \langle YZ \rangle} p \cdot \mathbf{y}_Y \cdot \mathbf{z}_Z; \\ L^{(1)}(\mathbf{y})_X &:= \sum_{X \xrightarrow{p} Y} p \cdot \mathbf{y}_Y; & L^{(3)}(\mathbf{y}, \mathbf{z}, \mathbf{w})_X &:= \sum_{X \xrightarrow{p} \langle YZW \rangle} p \cdot \mathbf{y}_Y \cdot \mathbf{z}_Z \cdot \mathbf{w}_W. \end{aligned}$$

Notice that  $L^{(0)}$  is a vector of constants, and  $L^{(1)}, L^{(2)}, L^{(3)}$  are linear, bilinear, trilinear vector functions, respectively. We write  $L^{(2)}(\mathbf{y}, \cdot)$  and  $L^{(3)}(\cdot, \mathbf{z}, \mathbf{w})$  etc. to mean the matrices  $U, V \in \mathbb{R}^{\Gamma \times \Gamma}$  such that  $U \cdot \mathbf{x} = L^{(2)}(\mathbf{y}, \mathbf{x})$  and  $V \cdot \mathbf{x} = L^{(3)}(\mathbf{x}, \mathbf{z}, \mathbf{w})$  etc. for all  $\mathbf{x} \in \mathbb{R}^\Gamma$ . It is straightforward to verify that

$$A = L^{(1)}(\cdot) + L^{(2)}(\mathbf{1}, \cdot) + L^{(2)}(\cdot, \mathbf{1}) + L^{(3)}(\mathbf{1}, \mathbf{1}, \cdot) + L^{(3)}(\mathbf{1}, \cdot, \mathbf{1}) + L^{(3)}(\cdot, \mathbf{1}, \mathbf{1}),$$

where  $\mathbf{1}$  denotes the vector with all ones. Let  $\mathbf{f} : \mathbb{R}^\Gamma \rightarrow \mathbb{R}^\Gamma$  be the function with

$$\mathbf{f}(\mathbf{x}) := L^{(0)} + L^{(1)}(\mathbf{x}) + L^{(2)}(\mathbf{x}, \mathbf{x}) + L^{(3)}(\mathbf{x}, \mathbf{x}, \mathbf{x}).$$

This ‘‘generating function’’  $\mathbf{f}$  plays a central role in the branching process literature. Notice that  $\mathbf{f}$  characterises the branching process  $S$  up to the order of the children. Define  $\mathbf{q}^{(0)} := \mathbf{0}$  (i.e., the vector with all zeros) and  $\mathbf{q}^{(i+1)} := \mathbf{f}(\mathbf{q}^{(i)})$  for all  $i \in \mathbb{N}$ . It is well-known [18] (and straightforward to show by induction on  $i$ ) that  $\mathcal{P}(\mathbf{T}_X \leq i) = \mathbf{q}_X^{(i)}$ . Define  $\mathbf{r}^{(i)} := \mathbf{1} - \mathbf{q}^{(i)}$ . Using either the definition of  $\mathbf{r}^{(i)}$  or the fact that  $\mathbf{r}_X^{(i)} = \mathcal{P}(\mathbf{T}_X > i)$  one can easily check

$$\begin{aligned} \mathbf{r}^{(i+1)} &= L^{(1)}(\mathbf{r}^{(i)}) + L^{(2)}(\mathbf{1}, \mathbf{r}^{(i)}) + L^{(2)}(\mathbf{r}^{(i)}, \mathbf{1} - \mathbf{r}^{(i)}) \\ &\quad + L^{(3)}(\mathbf{1}, \mathbf{1}, \mathbf{r}^{(i)}) + L^{(3)}(\mathbf{1}, \mathbf{r}^{(i)}, \mathbf{1} - \mathbf{r}^{(i)}) + L^{(3)}(\mathbf{r}^{(i)}, \mathbf{1} - \mathbf{r}^{(i)}, \mathbf{1} - \mathbf{r}^{(i)}) \\ &= A \cdot \mathbf{r}^{(i)} - L^{(2)}(\mathbf{r}^{(i)}, \mathbf{r}^{(i)}) - L^{(3)}(\mathbf{1}, \mathbf{r}^{(i)}, \mathbf{r}^{(i)}) - L^{(3)}(\mathbf{r}^{(i)}, \mathbf{1}, \mathbf{r}^{(i)}) \\ &\quad - L^{(3)}(\mathbf{r}^{(i)}, \mathbf{r}^{(i)}, \mathbf{1}) + L^{(3)}(\mathbf{r}^{(i)}, \mathbf{r}^{(i)}, \mathbf{r}^{(i)}). \end{aligned}$$

By defining

$$B(\mathbf{x}) := L^{(2)}(\mathbf{x}, \cdot) + L^{(3)}(\mathbf{1}, \mathbf{x}, \cdot) + L^{(3)}(\mathbf{x}, \mathbf{1}, \cdot) + L^{(3)}(\mathbf{x}, \cdot, \mathbf{1})$$

we get

$$\mathbf{r}^{(i+1)} = (A - B(\mathbf{r}^{(i)})) \cdot \mathbf{r}^{(i)} + L^{(3)}(\mathbf{r}^{(i)}, \mathbf{r}^{(i)}, \mathbf{r}^{(i)}). \quad (2)$$

Note that  $A - B(\mathbf{x})$  and  $B(\mathbf{x})$  are nonnegative for  $\mathbf{x} \in [0, 1]^\Gamma$  and that  $B(\varepsilon \cdot \mathbf{x}) = \varepsilon \cdot B(\mathbf{x})$  for all  $\varepsilon \in \mathbb{R}$ .

In the following, for a vector  $\mathbf{x}$ , we write  $\mathbf{x}_{min}$  and  $\mathbf{x}_{max}$  for the minimal and maximal entry of  $\mathbf{x}$ . It is easy to see that there is  $s \in (0, 1]$  such that for all  $X, Y \in \Gamma$  we have  $\mathcal{P}(\mathbf{T}_X > n) \geq s \cdot \mathcal{P}(\mathbf{T}_Y > n)$  for all  $n \in \mathbb{N}$ . (For instance, take for  $s$  the



probability to reach, starting in  $X$ , a tree with  $Y$  in at most  $|\Gamma|$  steps. This probability is positive, because  $A$  is strongly connected.) It follows that  $\mathbf{r}_{\min}^{(n)} \geq s \cdot \mathbf{r}_{\max}^{(n)}$ .

As  $\rho(A) \geq 1$  and  $A$  is strongly connected, Perron-Frobenius theory (see [3]) asserts that there is a vector  $\mathbf{u} \in (0, 1]^\Gamma$ , strictly positive in all components, such that  $A \cdot \mathbf{u} \geq \mathbf{u}$ . (For instance, one can take the dominant eigenvector of  $A$ .) W.l.o.g. we can take  $\mathbf{u}_{\max} = s$ . Choose  $d \in \mathbb{R}_+$  such that

$$B(\mathbf{1}) \cdot \mathbf{u} \leq d \cdot \mathbf{u}. \quad (3)$$

Define a sequence  $(\varepsilon_n)_{n \in \mathbb{N}}$  by setting  $\varepsilon_n := \mathbf{r}_{\max}^{(n)}$ . As  $\mathbf{r}_{\min}^{(n)} \geq s \cdot \varepsilon_n = \varepsilon_n \cdot \mathbf{u}_{\max}$ , we have

$$\varepsilon_n \cdot \mathbf{1} \geq \mathbf{r}^{(n)} \geq \varepsilon_n \cdot \mathbf{u}. \quad (4)$$

Observe that, since  $A$  is strongly connected, we have  $\mathcal{P}(\mathbf{T}_X > n) > 0$  for all  $X$  and all  $n$ , and hence  $\varepsilon_n \geq \varepsilon_{n+1} > 0$ . If  $(\varepsilon_n)_n$  does not converge to 0, then there is  $c > 0$  and  $X \in \Gamma$  such that  $\mathcal{P}(\mathbf{T}_X \geq n) \geq c$  for all  $n \in \mathbb{N}$ , already implying that  $\mathbb{E}\mathbf{T}_X = \infty$ . So we assume in the following that  $\lim_{n \rightarrow \infty} \varepsilon_n = 0$ . In particular, there is  $n_\perp \in \mathbb{N}$  such that  $\varepsilon_n d < 1$  for all  $n \geq n_\perp$ . Now we show, for all  $n \geq n_\perp$  and all  $i \in \mathbb{N}$ , that

$$\mathbf{r}^{(n+i)} \geq (1 - \varepsilon_n d)^i \varepsilon_n \mathbf{u}. \quad (5)$$

We proceed by induction on  $i$ . The induction base ( $i = 0$ ) follows from (4). Let  $i \geq 0$ . We have:

$$\begin{aligned} \mathbf{r}^{(n+i+1)} &\geq (A - B(\mathbf{r}^{(n+i)})) \cdot \mathbf{r}^{(n+i)} && \text{(by (2))} \\ &\geq (A - B(\mathbf{r}^{(n+i)})) \cdot (1 - \varepsilon_n d)^i \varepsilon_n \cdot \mathbf{u} && \text{(by induction hypothesis)} \\ &\geq (1 - \varepsilon_n d)^i \varepsilon_n \cdot (\mathbf{u} - B(\mathbf{r}^{(n+i)}) \cdot \mathbf{u}) && \text{(as } A\mathbf{u} \geq \mathbf{u}\text{)} \\ &\geq (1 - \varepsilon_n d)^i \varepsilon_n \cdot (\mathbf{u} - B(\varepsilon_n \cdot \mathbf{1}) \cdot \mathbf{u}) && \text{(by (4))} \\ &\geq (1 - \varepsilon_n d)^i \varepsilon_n \cdot (\mathbf{u} - \varepsilon_n d \cdot \mathbf{u}) && \text{(by (3))} \\ &= (1 - \varepsilon_n d)^{i+1} \varepsilon_n \cdot \mathbf{u} \end{aligned}$$

This proves (5). Now we have for all  $n \geq n_\perp$ :

$$\begin{aligned} \sum_{i=0}^k \mathbf{r}^{(n+i)} &\geq \sum_{i=0}^k (1 - \varepsilon_n d)^i \varepsilon_n \mathbf{u} && \text{(by (5))} \\ &= \frac{1 - (1 - \varepsilon_n d)^{k+1}}{1 - (1 - \varepsilon_n d)} \varepsilon_n \mathbf{u} \\ &= \frac{1 - (1 - \varepsilon_n d)^{k+1}}{d} \cdot \mathbf{u}, \end{aligned}$$

so, for every  $n \geq n_\perp$  there exists some  $k(n) \in \mathbb{N}$  such that  $\sum_{i=n}^{k(n)} \mathbf{r}^{(i)} \geq \frac{1}{2d} \cdot \mathbf{u}$ . Hence, for any  $X \in \Gamma$  we have

$$\begin{aligned} \mathbb{E}\mathbf{T}_X &= \sum_{i=0}^{\infty} \mathcal{P}(\mathbf{T}_X > i) = \sum_{i=0}^{\infty} \mathbf{r}_X^{(i)} \geq \sum_{i=n_\perp}^{\infty} \mathbf{r}_X^{(i)} = \sum_{i=n_\perp}^{k(n_\perp)} \mathbf{r}_X^{(i)} + \sum_{i=k(n_\perp)+1}^{k(k(n_\perp)+1)} \mathbf{r}_X^{(i)} + \cdots \\ &\geq \frac{1}{2d} \mathbf{u}_X + \frac{1}{2d} \mathbf{u}_X + \cdots = \infty, \end{aligned}$$

because  $\mathbf{u}_X > 0$ . This completes the proof.  $\square$

### D.3 Proof of Corollary 12

Here is a restatement of Corollary 12.

**Corollary 12.** *Consider a branching process with process symbols  $\Gamma$  and  $X_0 \in \Gamma$ . Then  $\mathbb{E}\mathbf{W}_{X_0}$  and  $\mathbb{E}\mathbf{T}_{X_0}$  are both finite or both infinite. Distinguishing between those cases is in P.*

*Proof.* By Proposition 11, the expectations  $\mathbb{E}\mathbf{W}_{X_0}$  and  $\mathbb{E}\mathbf{T}_{X_0}$  are both finite or both infinite. To distinguish between those cases, compute in polynomial time the matrix  $A$ . As  $A$  is nonnegative, we have  $\rho(A) \geq 1$  if and only if there is a nonnegative vector  $\mathbf{x}$  which is nonzero in at least one component and  $A\mathbf{x} \geq \mathbf{x}$  holds (i.e., “ $\geq$ ” holds in all components), see [3]. Therefore,  $\rho(A) \geq 1$  if and only if the linear programming (LP) problem “ $A\mathbf{x} \geq \mathbf{x} \geq \mathbf{0}$  and  $\|\mathbf{x}\|_1 = 1$ ” is feasible. This can be decided in P. We remark that a similar method was used in [16].  $\square$

### D.4 Proof of Theorem 13

Here is a restatement of Theorem 13.

**Theorem 13.** *Consider a pSJS  $S$  with alphabet  $\Sigma = \Gamma \cup Q$ . Let  $a \in \Gamma$ . Then  $\mathbb{E}\mathbf{W}_a$  and  $\mathbb{E}\mathbf{T}_a$  are both finite or both infinite. Distinguishing between those cases is in PSPACE, and PosSLP-hard even for pPDSs. Further, if  $S$  is normalised and  $\mathbb{E}\mathbf{W}_a$  is finite, one can efficiently express  $\mathbb{E}\mathbf{W}_a$ .*

*Proof.* If  $S$  is not normalised, we normalise it in polynomial time (Section 3.2). This does not change the finiteness of  $\mathbb{E}\mathbf{W}_a$  or  $\mathbb{E}\mathbf{T}_a$ . Then we compute in polynomial time (Theorem 6) the set  $Q' := \{q \in Q \mid [a \downarrow q] > 0\}$ . The values  $[a \downarrow q]$  for  $q \in Q'$  can be efficiently expressed (Theorem 6). Therefore, we can also efficiently express  $[a \downarrow] = \sum_{q \in Q'} [a \downarrow q]$  and decide in PSPACE if  $[a \downarrow] = 1$  or  $[a \downarrow] < 1$ . If  $[a \downarrow] < 1$ , then nonterminating runs have a positive probability and hence  $\mathbb{E}\mathbf{W}_a = \mathbb{E}\mathbf{T}_a = \infty$ . Otherwise, we have, with Proposition 10,

$$\mathbb{E}\mathbf{W}_a = \sum_{q \in Q'} [a \downarrow q] \cdot \mathbb{E}[\mathbf{W}_a \mid \text{Run}(a \downarrow q)] = \sum_{q \in Q'} [a \downarrow q] \cdot \mathbb{E}\mathbf{W}_{(aq)} =: \widetilde{W} \quad (6)$$

$$\text{and } \mathbb{E}\mathbf{T}_a = \sum_{q \in Q'} [a \downarrow q] \cdot \mathbb{E}[\mathbf{T}_a \mid \text{Run}(a \downarrow q)] \geq \sum_{q \in Q'} [a \downarrow q] \cdot \mathbb{E}\mathbf{T}_{(aq)} =: \widetilde{T}. \quad (7)$$

If  $\mathbb{E}\mathbf{W}_a$  is finite, then  $\mathbb{E}\mathbf{T}_a$  is finite, because  $\mathbf{W}_a \geq \mathbf{T}_a$ . If  $\mathbb{E}\mathbf{T}_a$  is finite, then  $\widetilde{T}$  is finite by (7), hence  $\widetilde{W}$  is finite by Corollary 12, hence  $\mathbb{E}\mathbf{W}_a$  is finite by (6). Therefore,  $\mathbb{E}\mathbf{W}_a$  is finite if and only if  $\mathbb{E}\mathbf{T}_a$  is finite.

In order to decide if  $\mathbb{E}\mathbf{W}_a$  is finite, by (6) it suffices to decide if  $\mathbb{E}\mathbf{W}_{(aq)}$  is finite for all  $q \in Q'$ . We cannot use Corollary 12 directly, because the coefficients of the branching process from Proposition 10 are not explicitly given. However, by Theorem 6

they are efficiently expressible and, therefore, so is the matrix  $A$  from Proposition 11. So we can decide in PSPACE whether  $\mathbb{E}\mathbf{W}_{(aq)}$  is finite by deciding whether the formula “ $\exists \mathbf{x} : A\mathbf{x} \geq \mathbf{x} \geq \mathbf{0}$  and  $\|\mathbf{x}\|_1 = 1$ ” in  $ExTh(\mathbb{R})$  is true (cf. the proof of Corollary 12).

If  $\mathbb{E}\mathbf{W}_a$  is finite, we wish to efficiently express it. We use again  $\widetilde{W}$  from (6), so it suffices to show that  $\mathbb{E}\mathbf{W}_{(aq)}$  is efficiently expressible for all  $q \in Q'$ . For that, notice again that the matrix  $A$  from Proposition 11 is efficiently expressible, and consider the formula “ $(I - A)\mathbf{x} = \mathbf{1}$ ” in  $ExTh(\mathbb{R})$ . Existentially quantify all variables in  $\mathbf{x}$  except the  $(aq)$ -component. By Proposition 11, the resulting formula efficiently expresses  $\mathbb{E}\mathbf{W}_{(aq)}$ .

It remains to show PosSLP-hardness for pPDSs. We draw from a reduction in [16], where it is shown that deciding if a pPDS terminates with probability 1 is PosSLP-hard. As a gadget for that, Etessami and Yannakakis [16] compute, given a PosSLP instance, a pPDS  $S$  with the following properties. (More precisely, they construct an equivalent recursive Markov chain.) The starting configuration is  $qa$ , and after having left the initial configuration,  $S$  reaches a configuration of the form  $q\alpha$  with  $\alpha \in \Gamma^*$  again with probability 1. At that time, the configuration is  $q$  with some probability  $p$ , and  $qaa$  with probability  $1 - p$ . Moreover, the time (and hence work) needed to reach either of those configurations is essentially bounded by the size of the given PosSLP instance, so it is finite. Furthermore, the given PosSLP instance is a “yes instance” if and only if  $p > \frac{1}{2}$ . It is easy to see that  $\mathbb{E}\mathbf{W}_{qa}$  is finite in  $S$  if and only if  $\mathbb{E}\mathbf{W}_X$  is finite in the branching process  $\overline{S}$  that consists only of the transitions  $X \xrightarrow{p} \perp$  and  $X \xrightarrow{1-p} \langle XX \rangle$ . The  $1 \times 1$ -matrix  $A$  for  $\overline{S}$  from Proposition 11 consists of a single entry  $2 \cdot (1 - p)$ . Consequently,  $\mathbb{E}\mathbf{W}_{qa}$  is finite in  $S$  if and only if  $2 \cdot (1 - p) < 1$ , which is equivalent to  $p > \frac{1}{2}$ . This completes the reduction. We remark that the reduction does not show that deciding if  $\mathbb{E}\mathbf{W}$  is finite is PosSLP-hard for branching processes, because  $S$  has more control states than just  $q$ . (Recall that the problem for branching processes is in P by Corollary 12.)  $\square$

## D.5 Proof of Proposition 14

Here is a restatement of Proposition 14.

**Proposition 14.** *Consider the family of branching processes with transitions  $X \xrightarrow{p} \langle XX \rangle$  and  $X \xrightarrow{1-p} \perp$ , where  $0 < p < 1/2$ . Then the ratio  $\mathbb{E}[\mathbf{W}_X]/\mathbb{E}[\mathbf{T}_X]$  is unbounded for  $p \rightarrow 1/2$ .*

*Proof.* By Proposition 11, we have  $\mathbb{E}\mathbf{W}_X = 1/(1 - 2p)$ . It is shown in [24] that  $\lim_{p \rightarrow \frac{1}{2}} \frac{\mathbb{E}\mathbf{T}_X}{-2 \ln(1 - 2p)} = 1$ . Consequently, we have  $\lim_{p \rightarrow \frac{1}{2}} \mathbb{E}[\mathbf{W}_X]/\mathbb{E}[\mathbf{T}_X] = \infty$ .  $\square$