Fair Termination for Parameterized Probabilistic Concurrent Systems

Ondřej Lengál[1], Anthony W. Lin[2], Rupak Majumdar[3], and Philipp Rümmer[4]

[1] FIT, Brno University of Technology, Czech Republic
[2] Department of Computer Science, University of Oxford, UK
[3] MPI-SWS Kaiserslautern, Germany
[4] Uppsala University, Sweden

Abstract. We consider the problem of automatically verifying that a parameterized family of probabilistic concurrent systems terminates with probability one for all instances against adversarial schedulers. A parameterized family defines an infinite-state system: for each number n, the family consists of an instance with n finite-state processes. In contrast to safety, the parameterized verification of liveness is currently still considered extremely challenging especially in the presence of probabilities in the model. One major challenge is to provide a sufficiently powerful symbolic framework. One well-known symbolic framework for the parameterized verification of non-probabilistic concurrent systems is *regular model checking*. Although the framework was recently extended to probabilistic systems, incorporating fairness in the framework — often crucial for verifying termination — has been especially difficult due to the presence of an infinite number of fairness constraints (one for each process). Our main contribution is a systematic, regularity-preserving, encoding of *finitary fairness* (a realistic notion of fairness proposed by Alur & Henzinger) in the framework of regular model checking for probabilistic parameterized systems. Our encoding reduces termination with finitary fairness to verifying parameterized termination *without fairness* over probabilistic systems in regular model checking (for which a verification framework already exists). We show that our algorithm could verify termination for many interesting examples from distributed algorithms (Herman's protocol) and evolutionary biology (Moran process, cell cycle switch), which do not hold under the standard notion of fairness. To the best of our knowledge, our algorithm is the first fully-automatic method that can prove termination for these examples.

1 Introduction

In parameterized probabilistic concurrent systems, a population of *agents*, each typically modeled as a finite-state probabilistic program, run concurrently in discrete time and update their states based on probabilistic transition rules. The interaction is governed by an underlying *topology*, which determines which agents can interact in one step, and a *scheduler*, which picks the specific agents involved in the interaction. Concurrent probabilistic systems arise as models of distributed algorithms [38,29,31,34,25], where each agent is a processor, the interaction between processors is determined by a communication topology, and the processor can update its internal state based on the communication as well as randomization. In each step, the scheduler adversarially chooses a processor to run. Concurrent probabilistic populations also arise in agent-based population models in biology [35], wherein an agent can represent an allele, a

cell, or a species, and the interaction between agents describes how these entities evolve over time. For a population of a fixed size, there is a rich theory of probabilistic verification [20,7,1,48] based on finite-state Markov decision processes (MDPs). Verification questions for population models, however, ask if a property holds for populations of *all* sizes: even if each agent is finite-state, the family of all processes (for each population size) is an infinite-state MDP. Indeed, for many simple population models, one can show that the verification question is undecidable, even for reachability or safety properties in the non-probabilistic setting [6,11,22]. Consequently, the verification question for populations requires techniques beyond finite-state probabilistic verification, and requires symbolic techniques to represent potentially infinite sets of states.

One well-known symbolic framework for verifying parameterized non-probabilistic concurrent systems is *regular model checking* [3,42,4,13,41,47], where states of a population are modeled using words over a suitable alphabet, sets of states are represented as regular languages, and the transition relation is defined as a regular transducer. From parameterized verification of non-probabilistic processes, it is known that regular languages provide a robust symbolic representation of infinite sets, and automata-theoretic algorithms provide the basis of checking safety or termination properties.

In this paper, we consider the problem of verifying that a given parameterized family of probabilistic concurrent systems *almost surely terminates*, i.e., reaches certain final states with probability 1 from each initial state regardless of the behaviour of the schedulers. Termination is a fundamental property when verifying parameterized probabilistic systems. Since termination typically, however, fails without imposing certain *fairness* conditions on the scheduler, it is crucial to be able to incorporate fairness assumptions into a termination analysis. Therefore, although the framework of regular model checking has recently been extended for proving termination (without fairness) over parameterized probabilistic concurrent systems [36], it still cannot be used to prove termination for many interesting parameterized probabilistic concurrent systems.

What notion of fairness should we consider for proving termination for parameterized probabilistic concurrent systems? To answer this question, one would naturally start by looking at standard notions of fairness in probabilistic model checking [7], which asserts that every process must be chosen infinitely often. However, this notion seems to be too weak to prove termination for many of our examples, notably Herman's self-stabilizing protocol [29] in an asynchronous setting, and population models from biology (e.g. Moran's process [35]). The standard notion of fairness gives rise to a rather unintuitive and unrealistic strategy for the scheduler, which could delay an enabled process for as long as it desires while still being fair (see [15, Example 8] and the Herman's protocol example in Section 3). For this reason, we propose to consider Alur & Henzinger's [5] *finitary fairness* — a stronger notion of fairness that allows the scheduler to delaying executing an enabled process in an infinite run for at most k steps, for some unknown but fixed bound $k \in \mathbb{N}$. Alur & Henzinger argued that this fairness notion is more realistic in practice, but it is not as restrictive as the notion of k*-fairness*, which fixes the bound k a priori. In addition, it should be noted that finitary fairness is strictly weaker than probabilistic fairness (scheduler chooses processes randomly) for almost-sure termination over finite MDPs and parameterized probabilistic systems (an infinite family of finite MDPs). We will show in this paper that there are many interest-

ing examples of parameterized probabilistic concurrent systems for which termination is satisfied under finitary fairness, but *not* under the most general notion of fairness.

Contributions. Our main contribution is a systematic, regularity-preserving, encoding of finitary fairness in the framework of regular model checking for parameterized probabilistic concurrent systems. More precisely, our encoding reduces the problem of verifying almost sure termination under finitary fairness to almost sure termination *without fairness* in regular model checking, for which a verification framework exists [36].

In general, the difficulty with finding an encoding of fairness is how to deal with an infinite number of fairness requirements (one for each process) in a systematic and regularity-preserving manner. There are known encodings of general notions of fairness in regular model checking, e.g., by using a token that is passed to the next process (with respect to some ordering of the processes) when the current process is executed, and ensuring that the first process holds the token and passes it to the right infinitely many times (e.g. see [4,42]). However, these encodings do not work in in our case for several reasons. Firstly, they do not take into account the unknown upper bound (from finitary fairness) within which time a process has to be executed. Adapting these encodings to finitary fairness would require *the use of unbounded counters*, which do not preserve regularity. Secondly, such encodings would yield the problem of verifying an almost-sure Rabin property (of the form $\Box\Diamond A \wedge \Diamond B$ in LTL notation, where A and B are regular sets). Although we could reduce this to an almost-sure termination property by means of product automata construction (i.e. by first converting the formula to deterministic Rabin automaton), the target set B in the resulting termination property $\Diamond B$ (consisting of configurations in strongly connected components satisfying some properties) is *not* necessarily regular.

Instead, we revisit the well-known *abstract program transformation* in the setting of non-probabilistic concurrent systems [26] encoding fairness into the program by associating to each process an unbounded counter that acts as an "alarm clock", which will "set off" if an enabled process has not been chosen by the scheduler for "too long." This abstract program transformation has been adapted by Alur & Henzinger [5] in the case of finitary fairness by additionally incorporating an extra counter n that stores the unknown upper bound and resetting the value of a counter belonging to a chosen process to the "default value" n. Our contributions are as follows:

1. We show how Alur & Henzinger's program transformation could be adapted to the setting of probabilistic parameterized concurrent systems (infinite family of finite MDPs). This involves constructing a new parameterization of the system (using the idea of weakly finite systems) and a proof that the transformation preserves reachability probabilities.
2. We show how the resulting abstract program transformation could be made concrete in the setting of regular model checking *without using automata models beyond regular automata*.
3. We have implemented this transformation in FAIRYTAIL. Combined with the existing algorithm [36] for verifying almost sure termination (without fairness) in regular model checking, we have successfully verified a number of models obtained from distributed algorithms and biological systems including Herman's protocol [29], Moran processes in a linear array [40,35], and the cell cycle switch

model [17] on ring and line topologies. To the best of our knowledge, our algorithm is the first fully-automatic method that can prove termination for these examples.

Related work. There are few techniques for automatic verification of liveness properties of parameterized probabilistic programs. Almost sure verification of probabilistic finite-state programs goes back to Pnueli and co-workers [28,45]. Esparza et al. [23] generalize the reasoning to weakly finite programs, and describe a heuristic to guess a *terminating pattern* by constructing a nondeterministic program from a given probabilistic program and a terminating pattern candidate. This allows them to exploit model checkers and termination provers for nondeterministic programs. More recently, Lin and Rümmer [36] consider unconditional termination for parameterized probabilistic programs. While our work builds on these techniques, our main contribution is the incorporation of fairness in regular model checking of probabilistic programs, which was not considered before.

Fairness for concurrent probabilistic systems was considered by Vardi [48] and by Hart, Sharir, and Pnueli [28], and generalized later [45,21,8]. The focus was, however, on a fixed number of processes. The notion of fairness through explicit scheduling was developed by Olderog and Apt [43]. More recently, notions of fairness for infinitary control (i.e., where an infinite number of processes can be created) was considered by Hoenicke, Olderog, and Podelski [44,30].

Martingale techniques have been used to prove termination of sequential, infinite-state, probabilistic programs [18,39,24,19,32]. These results are not comparable to our results, as they do not consider unbounded families of fairness constraints nor communication topologies.

2 Preliminaries

General notations: For any two given real numbers $i \leq j$, we use a standard notation (with an extra subscript) to denote real intervals, e.g., $[i,j]_{\mathbb{R}} = \{k \in \mathbb{R} : i \leq k \leq j\}$ and $(i,j]_{\mathbb{R}} = \{k \in \mathbb{R} : i < k \leq j\}$. We will denote intervals over integers by removing the subscript, i.e., $[i,j] = [i,j]_{\mathbb{R}} \cap \mathbb{Z}$. Given a set S, we use S^* to denote the set of all finite sequences of elements from S. The set S^* always includes the empty sequence, which we denote by ϵ. We use S^+ to denote the set $S^* \setminus \{\epsilon\}$. Given two sets of words S_1, S_2, we use $S_1 \cdot S_2$ to denote the set $\{v \cdot w : v \in S_1, w \in S_2\}$ of words formed by concatenating words from S_1 with words from S_2. Given two relations $R_1, R_2 \subseteq S \times S$, we define their composition as $R_1 \circ R_2 = \{(s_1, s_3) : \exists s_2((s_1, s_2) \in R_1 \wedge (s_2, s_3) \in R_2)\}$.

Transition systems: We fix the (countably infinite) set AP of *atomic propositions*. Let ACT be a finite set of *action symbols*. A *transition system* over ACT is a tuple $\mathfrak{S} = \langle S; \{\rightarrow_a\}_{a \in \mathsf{ACT}}, \ell \rangle$, where S is a set of *configurations*, $\rightarrow_a \subseteq S \times S$ is a binary relation over S, and $\ell : \mathsf{AP} \rightarrow 2^S$ maps atomic propositions to sets of configurations (we omit ℓ if it is not important). We use $\rightarrow$ to denote the relation $\left(\bigcup_{a \in \mathsf{ACT}} \rightarrow_a\right)$. The notation $\rightarrow^+$ (resp. $\rightarrow^*$) is used to denote the transitive (resp. transitive-reflexive) closure of $\rightarrow$. We say that a sequence $s_1 \rightarrow \cdots \rightarrow s_n$ is a *path* (or *run*) in $\mathfrak{S}$ (or in $\rightarrow$). Given two paths $\pi_1 : s_1 \rightarrow^* s_2$ and $\pi_2 : s_2 \rightarrow^* s_3$ in $\rightarrow$, we may concatenate them to obtain $\pi_1 \odot \pi_2$ (by gluing together s_2). We call π_1 a *prefix* of $\pi_1 \odot \pi_2$. For each $S' \subseteq S$, we use the

notations $pre_{\rightarrow}(S')$ and $post_{\rightarrow}(S')$ to denote the pre/post image of S' under $\rightarrow$. That is, $pre_{\rightarrow}(S') = \{p \in S : \exists q \in S'(p \rightarrow q)\}$ and $post_{\rightarrow}(S') = \{q \in S : \exists p \in S'(p \rightarrow q)\}$.

Words and automata: We assume basic familiarity with finite word automata. Fix a finite alphabet Σ. For each finite word $w = w_1 \dots w_n \in \Sigma^*$, we write $w[i, j]$, where $1 \leq i \leq j \leq n$, to denote the segment $w_i \dots w_j$. Given an automaton $\mathcal{A} = (\Sigma, Q, \delta, q_0, F)$, a run of $\mathcal{A}$ on w is a function $\rho : \{0, \dots, n\} \rightarrow Q$ with $\rho(0) = q_0$ that obeys the transition relation δ. We may also denote the run ρ by the word $\rho(0) \cdots \rho(n)$ over the alphabet Q. The run ρ is said to be *accepting* if $\rho(n) \in F$, in which case we say that w is *accepted* by $\mathcal{A}$. The language $L(\mathcal{A})$ of $\mathcal{A}$ is the set of words in Σ^* accepted by $\mathcal{A}$.

Reachability games: We recall some basic concepts on 2-player reachability games (see e.g. [27, Chapter 2] on games with 1-accepting conditions). An *arena* is a transition system $\mathfrak{S} = \langle S = V_1 \cup V_2; \rightarrow_1, \rightarrow_2 \rangle$, where S (i.e. the set of "game configurations") is partitioned into two disjoint sets V_1 and V_2 such that $pre_{\rightarrow_i}(S) \subseteq V_i$ for each $i \in \{1, 2\}$. The transition relation $\rightarrow_i$ denotes the actions of Player i. Similarly, for each $i \in \{1, 2\}$, the configurations V_i are controlled by Player i. In the following, Player 1 will also be called "Scheduler," and Player 2 "Process". Given a set $I_0 \subseteq S$ of initial configurations and a set $F \subseteq S$ of final (a.k.a. target) configurations, the goal of Player 2 is to reach F from I_0, while the goal of Player 1 is to avoid it. More formally, a *strategy* for Player i is a partial function $f : S^* V_i \rightarrow S$ such that, for each $v \in S^*$ and $p \in V_i$, if vp is a path in $\mathfrak{S}$ and p is not a dead end (i.e., $p \rightarrow_i q$ for some q), then $f(vp)$ is defined in such a way that $p \rightarrow_i f(vp)$. Given a strategy f_i for Player $i \in \{1, 2\}$ and an initial configuration $s_0 \in S$, we can define a unique (finite or infinite) path in $\mathfrak{S}$ such that $\pi : s_0 \rightarrow_{j_1} s_1 \rightarrow_{j_2} \cdots$ where $s_{j_{k+1}} = f_i(s_0 s_1 \dots s_{j_k})$ for $i \in \{1, 2\}$ is the (unique) configuration s.t. $s_{j_k} \in V_i$. Player 2 *wins* iff some configuration in F appears in π, or if the path is finite and the last configuration belongs to Player 1. Player 1 *wins* iff Player 2 does not win; we say Player 2 *loses*. A strategy f for Player i is *winning* from I_0 if for each strategy g of Player $3 - i$, the unique path in $\mathfrak{S}$ from each $s_0 \in I_0$ witnesses a win for Player i. Such games (a.k.a. *reachability games*) are *determined* (see e.g. [27, Proposition 2.21]): either Player 1 has a winning strategy or Player 2 has a winning strategy.

Convention. *For notational simplicity, w.l.o.g., we make the following assumptions on our reachability games. They suffice for the purpose of proving liveness for parameterised systems.*

- **(A0)** *Arenas are strictly alternating, i.e., a move made by a player does not take the game back to her configuration ($post_{\rightarrow_i}(S) \cap V_i = \emptyset$, for each $i \in \{1, 2\}$).*
- **(A1)** *Initial and final configurations belong to Player 1, i.e., $I_0, F \subseteq V_1$*
- **(A2)** *Non-final configurations are not dead ends: $\forall x \in S \setminus F, \exists y : x \rightarrow_1 y \vee x \rightarrow_2 y$.*

Markov chains: A (discrete-time) *Markov chain* (a.k.a. *DTMC*) is a structure of the form $\mathfrak{S} = \langle S; \delta, \ell \rangle$ where S is a set of configurations, δ is a function that associates a configuration $s \in S$ with a probability distribution over a sample space $D \subseteq S$ (i.e. the probability of going to a certain configuration from s), and $\ell : \mathsf{AP} \rightarrow 2^S$ maps atomic propositions to subsets of S. In what follows, we will assume that each $\delta(s)$ is a discrete probability distribution with a finite sample space. This assumption allows us

to simplify our notation: a DTMC $\langle S; \delta, \ell \rangle$ can be seen as a transition system $\langle S; \rightarrow, \ell \rangle$ with a transition probability function δ mapping a transition $t = (s, s') \in \rightarrow$ to a value $\delta(t) \in (0, 1]$ such that $\sum_{s' \in post(s)} \delta((s, s')) = 1$. That is, transitions with zero probabilities are removed from $\rightarrow$. We will write $s \xrightarrow{p} s'$ to denote $s \rightarrow s'$ and that $\delta((s, s')) = p$. The *underlying transition graph* of a DTMC $\langle S; \delta, \ell \rangle$ is the transition system $\langle S; \rightarrow, \ell \rangle$ with δ omitted. Given a finite path $\pi = s_0 \rightarrow \cdots \rightarrow s_n$ from the initial configuration $s_0 \in S$, let Run_π be the set of all finite/infinite paths with π as a prefix, i.e., of the form $\pi \odot \pi'$ for some finite/infinite path π'. Given a set $F \subseteq S$ of target configurations, the probability $\text{Prob}_{\mathfrak{S}}(s_0 \models \Diamond F)$ (the subscript $\mathfrak{S}$ may be omitted when understood) of reaching F from s_0 in $\mathfrak{S}$ can be defined using a standard cylinder construction (see e.g [33]) as follows. For each finite path $\pi = s_0 \rightarrow \cdots \rightarrow s_n$ in $\mathfrak{S}$ from s_0, we set Run_π to be a basic cylinder, to which we associate the probability $\text{Prob}(Run_\pi) = \prod_{i=0}^{n-1} \delta((s_i, s_{i+1}))$. This gives rise to a unique probability measure for the σ-algebra over the set of all runs from s_0. The probability $\text{Prob}(s_0 \models \Diamond F)$ is then the probability of the event F containing all paths in $\mathfrak{S}$ with some "accepting" finite path as a prefix, i.e., a finite path from s_0 ending in some configuration in F. In general, given an LTL formula φ over AP, the event containing all paths from s_0 in $\mathfrak{S}$ satisfying φ is measurable [48] and its probability value $\text{Prob}(s_0 \models \varphi)$ is well-defined.
Notation: Whenever understood, we will omit mention of ℓ from $\langle S; \delta, \ell \rangle$.

3 Abstract Models of Probabilistic Concurrent Programs

In this section, we recall the notion of Markov Decision Processes (MDPs) and fair MDPs [7]. These serve as our abstract models of probabilistic concurrent programs. We then define the notion of finitary fairness [5] and discuss its basic properties in the setting of MDPs.

3.1 Markov Decision Processes

A *Markov decision process* (*MDP*) is a strictly alternating arena $\mathfrak{S} = \langle S = V_1 \cup V_2; \rightarrow_1, \rightarrow_2 \rangle$ such that $\langle S; \rightarrow_2 \rangle$ is a DTMC, i.e., $\rightarrow_2$ is associated with some transition probability function, and that the atomic propositions are not important. Intuitively, the transition relation $\rightarrow_1$ is nondeterministic (controlled by a "demonic" scheduler), whereas the transition relation $\rightarrow_2$ is probabilistic. By definition of arenas, the configurations of the MDPs are partitioned into the set V_1 of *nondeterministic states* (controlled by Scheduler) and the set V_2 of *probabilistic states*. Formally, $pre_{\rightarrow_1}(S) \cap pre_{\rightarrow_2}(S) = \emptyset$. Each Scheduler's strategy[5] $f : S^* V_1 \rightarrow S$ gives rise to an infinite-state DTMC with the underlying transition system $\mathfrak{S}_f = \langle S'; \rightarrow_3, \ell \rangle$ and the transition probability function δ' defined as follows. Here, S' is the set of all finite/infinite paths π from s_0. For each state $s' \in S$ and each path π from s_0 ending in some state $s \in S$, we define $\pi \rightarrow_3 \pi s'$ iff: (1) if $s \in V_1$ is a nondeterministic state, then $f(\pi) = s'$, and (2) if $s \in V_2$ is a probabilistic state, then $s \rightarrow_2 s'$. Intuitively, $\mathfrak{S}_f$ is an unfolding of the game arena $\mathfrak{S}$ (i.e. a disjoint union of trees) where branching only occurs on probabilistic states. Transitions $\pi \rightarrow_3 \pi s'$ satisfying Case (1) have the probability $\delta'((\pi, \pi s')) = 1$;

[5] Also called "scheduler" or "adversary" for short.

otherwise, its probability is $\delta'((\pi, \pi s')) = \delta((s, s'))$. We let ℓ be a function mapping each subset $X \subseteq S$ (used as an atomic proposition) to the set of all finite paths in $\mathfrak{S}_f$ from s_0 to X. Since $\mathfrak{S}_f$ is a DTMC, given an LTL formula φ over subsets of S as atomic propositions, the probability $\text{Prob}_{\mathfrak{S}_f}(s_0 \models \varphi)$ of satisfying φ in $\mathfrak{S}$ from s_0 under the scheduler f is well-defined. In particular, $\text{Prob}_{\mathfrak{S}_f}(s_0 \models \Diamond F)$ is the probability of reaching F from s_0 in $\mathfrak{S}$ under the scheduler f. The probability $\text{Prob}_{\mathfrak{S},\mathcal{C}}(s_0 \models \varphi)$ of satisfying φ from s_0 in the MDP $\mathfrak{S}$ under a class $\mathcal{C}$ of schedulers is defined to be the infimum of the set of all probabilities $\text{Prob}_{\mathfrak{S}_f}(s_0 \models \varphi)$ over all $f \in \mathcal{C}$. We will omit mention of $\mathcal{C}$ when it denotes the class of all schedulers.

An MDP is *weakly-finite* [23] if from each configuration, the set of all configurations that are reachable from it (in the underlying transition system of the MDP) is finite. Note that the state space of weakly-finite MDPs can be infinite. The restriction of weak finiteness is another way of defining the notion of *parameterized systems*, which are an infinite family of finite-state systems. Weakly-finite MDPs capture many interesting probabilistic concurrent systems in which each process is finite-state; this is the case for many probabilistic distributed protocols.

3.2 Fair Markov Decision Processes

A *fair Markov decision process (FMDP)* is a structure of the form $\mathfrak{S} = \langle S = V_1 \cup V_2; \rightarrow_1, \rightarrow_2, \mathfrak{C}, \mathfrak{J} \rangle$, where $\langle S = V_1 \cup V_2; \rightarrow_1, \rightarrow_2 \rangle$ is an MDP, $\mathfrak{J}$ is a weak fairness (a.k.a. *justice*) requirement, and $\mathfrak{C}$ is a strong fairness (a.k.a. *compassion*) requirement. More precisely, a *weak fairness requirement* is a set (at most countably infinite) of *atomic weak fairness requirements* of the form $\Diamond\Box A \Rightarrow \Box\Diamond B$, for some $A, B \subseteq S$. Here, the $\Box$ and $\Diamond$ modalities are the standard "always" and "eventually" LTL operators. The set A (resp. B) will be called the *premise* (resp. *consequence*). Intuitively, if A is interpreted as "Process 1 is waiting to move" and B as "Process 1 is chosen", then this fairness requirement may be read as: at no point can Process 1 be continuously waiting to move without being chosen. In addition, a *strong fairness requirement* is a set (again, at most countably infinite) of *atomic strong fairness requirements* of the form $\Box\Diamond A \Rightarrow \Box\Diamond B$, for some $A, B \subseteq S$. Using the above example, a strong fairness requirement reads: if Process 1 is waiting to move infinitely often, then it is chosen infinitely often. As before, the set A (resp. B) will be called the *premise* (resp. *consequence*). In the following, when it is clear whether a fairness requirement is a justice or a compassion, we will denote it by the pair (A, B) of premise and consequence.

Given an FMDP $\mathfrak{S} = \langle S = V_1 \cup V_2; \rightarrow_1, \rightarrow_2, \mathfrak{C}, \mathfrak{J} \rangle$, a configuration $s_0 \in S$, and a scheduler f, since each atomic fairness requirement is an LTL formula and there are at most countably many atomic fairness requirements, the set of paths from s_0 in the DTMC $\mathfrak{S}_f$ induced by f satisfying $\mathfrak{C}$ and $\mathfrak{J}$ is measurable. We say that a scheduler f is $\mathfrak{S}$-*fair* if $\text{Prob}_{\mathfrak{S}_f}(s_0 \models \mathfrak{C} \wedge \mathfrak{J}) = 1$ for every initial configuration s_0. The fairness conditions $(\mathfrak{C}, \mathfrak{J})$ are *realizable* in $\mathfrak{S}$ if there exists at least one $\mathfrak{S}$-fair scheduler.

A natural fairness notion we consider in this paper is *process fairness*, which asserts that each process is chosen infinitely often. For this notion of fairness, we can assume that the consequence B of each atomic fairness requirement asserts that a particular process is chosen. We make one simplifying assumption: *each process is always enabled* (i.e., can always be chosen by the scheduler). This assumption is reasonable since we

can always introduce an idle transition for each process. Under this assumption, we have that *from each $v_1 \in V_1$, there exists a transition $v_1 \rightarrow_1 v_2$ for some $v_2 \in B$.* This implies that our fairness conditions are always realizable and that the probability $\text{Prob}_{\mathfrak{S},\mathcal{C}}(E)$ of event E over the set of all $\mathfrak{S}$-fair schedulers is well-defined.

3.3 Finitary Fairness

Given an FMDP $\mathfrak{S} = \langle S = V_1 \cup V_2; \rightarrow_1, \rightarrow_2, \mathfrak{C}, \mathfrak{J} \rangle$, a configuration $s_0 \in S$, and a number $k \in \mathbb{N}$, we say that a scheduler f is $\mathfrak{S}$-*k-fair* (or *k-fair* whenever $\mathfrak{S}$ is understood) if for each atomic fairness requirement (A, B):

1. if (A, B) is justice, then (the underlying graph of) $\mathfrak{S}_f$ contains no path π of length k satisfying the LTL formula $\Box(A \wedge \neg B)$.
2. if (A, B) is compassion, then $\mathfrak{S}_f$ contains no path π satisfying the LTL formula $\psi_k \wedge \Box\neg B$, where $\psi_0 := \mathit{true}$ and $\psi_i := \Diamond(A \wedge \psi_{i-1})$ for each $i > 0$.

In other words, a premise in a justice requirement cannot be satisfied for k *consecutive* steps without satisfying a consequence, while a premise in a compassion requirement cannot be satisfied for k (not necessarily consecutive) steps without satisfying a consequence. A scheduler is said to be *finitary fair (fin-fair)* if it is k-fair for some k. The fairness conditions $(\mathfrak{C}, \mathfrak{J})$ are said to be *finitary-realizable (fin-realizable)* in $\mathfrak{S}$ if there exists at least one fin-fair scheduler. Under this assumption, the probability $\text{Prob}_{\mathfrak{S},\mathcal{C}}(E)$ of an event E over the set $\mathcal{C}$ of all fin-fair schedulers is well-defined. In what follows, for an FMDP $\mathfrak{S}$, we will simply denote $\text{Prob}_{\mathfrak{S},\mathcal{C}}(E)$ as $\text{Prob}_{\mathfrak{S}}(E)$. In this paper, we propose to study *termination of probabilistic concurrent programs under finitary fairness*, i.e., to determine whether $\text{Prob}_{\mathfrak{S},\mathcal{C}}(s_0 \models \Diamond F) = 1$, where $\mathcal{C}$ is the class of all fin-fair schedulers.

The following proposition states one special property of weakly-finite MDPs.

Proposition 1. *Let $\mathfrak{S}$ and $\mathfrak{S}'$ be two weakly-finite fair MDPs with identical underlying transition systems (but possibly different probability values). For each set F of final states, and each initial configuration s_0, it is the case that $\text{Prob}_{\mathfrak{S}}(s_0 \models \Diamond F) = 1$ iff $\text{Prob}_{\mathfrak{S}'}(s_0 \models \Diamond F) = 1$.*

By Proposition 1, when dealing with almost-sure finitary-fair termination of weakly-finite MDPs, we only care whether a transition has a zero or a non-zero probability, i.e., if it is non-zero, then the exact value is irrelevant. Incidentally, the same also holds for other properties including almost-sure termination without fairness and qualitative temporal specifications [28,45,36]. *For this reason, we may simply omit these probability values from our symbolic representation of weakly-finite MDPs, which we will do from the next section onwards.*

3.4 Herman's Protocol

Herman's protocol [29] is a distributed self-stabilization algorithm for a population of processes organized in a ring. The *correct* configurations are those where exactly one process holds a token. If, through some error, the ring enters an *erroneous* configuration

(in which multiple processes hold tokens), Herman's protocol ensures that the system will *self-stabilize*: it will almost surely go back to a configuration with only one token.

Let us discuss how the protocol works in more detail. Fix $N \geq 3$ processors organized in a ring. If a chosen process does not hold a token, then it can perform an idle transition (i.e. do nothing). If a chosen process holds a token, then it can keep holding the token with probability $\frac{1}{2}$ or pass it on to its clockwise neighbor (the process $(i+1) \bmod N$, for processes numbered $0, \ldots, N-1$) with probability $\frac{1}{2}$. If a process currently holds a token and receives another token from its (counter-clockwise) neighbor, then the two tokens are merged[6] into one, leaving the process with one token.

Formally, Hermann's protocol can be modeled as a weakly-finite Markov decision process whose states are vectors in $\{\bot, \top\}^*$. For each N, the state of the protocol is described by a bitvector of N bits, with the i-th bit being one iff the i-th process holds a token. From a state $\mathbf{v}$, the scheduler picks a process $i \in \{0, \ldots, N-1\}$. Given a chosen process i, the new state remains $\mathbf{v}$ if the chosen process i did not hold a token ($\mathbf{v}(i) = \bot$). If $\mathbf{v}(i) = \top$, the new state is $\mathbf{v}$ with probability $\frac{1}{2}$ and $\mathbf{v} \oplus e_i \oplus e_{(i+1) \bmod N}$ with probability $\frac{1}{2}$. Here, e_i denotes a vector with $\top$ in the i-th position and $\bot$ everywhere else, and $\oplus$ is the XOR operation. We want to ensure that, starting from an arbitrary initial assignment of tokens, any population self-stabilizes with probability one.

Process fairness for Herman's protocol is a set of N atomic fairness requirements, each asserting that the process i is executed infinitely often, for each $i \in \{1, \ldots, N\}$. Unfortunately, Herman's protocol does *not* terminate with probability one against some fair schedulers. To see this, consider the start state $s_0 = (\top, \bot, \top)$. Let us call the token held by Process 0 "the first token", and the token held by Process 2 "the second token". Define a *round* as the following sequence of moves by the scheduler: keep choosing the process that holds the first token until it passes the token to the right, and do the same to the same to the second token. For example, the two configurations obtained after completing the first and second rounds from s_0 are, respectively, $(\top, \top, \bot)$ and $(\bot, \top, \top)$. To see that the scheduler is fair, for each integer $i > 0$, the probability that the i-th round is not completed is 0 since the probability that one of the tokens will be kept at the same process for an infinite amount of time is 0. Therefore, the probability that some round is not completed is also 0. Completing two rounds ensure that all the processes are picked. Therefore, every process will be chosen with probability 1. On the other hand, observe that correct configurations are not seen in the induced DTMC, showing that self-stabilization holds with probability 0 under this scheduler.

Herman's protocol can be shown to self-stabilize with probability one under all fin-fair schedulers, which can be proved by our fully-automatic verification algorithm (presented later in the paper).

4 Regular Model Checking: A Symbolic Framework

In this section, we recall *regular model checking* (see e.g. [3,42,46]), a symbolic framework for specifying infinite-state systems based on finite automata and regular transducers and developing automatic verification (semi-)algorithms.

A transition system $\mathfrak{S} = \langle S = V_1 \cup V_2; \rightarrow_1, \rightarrow_2 \rangle$ is specified in the framework as a regular language S (e.g. as a regular expression over some alphabet Σ), and two

[6] Herman [29] describes a more general protocol in which tokens can be merged/destroyed with some probability. We consider this restriction for simplicity of presentation.

"regular relations" $\rightarrow_1, \rightarrow_2 \subseteq \Sigma^* \times \Sigma^*$. For simplicity, in the following we will assume that $S = \Sigma^*$. How do we specify regular relations? One standard way is to restrict to length-preserving relations (i.e. the relation may only contain a pair of words of the same length) and specify such relations as regular languages over the alphabet $\Sigma \times \Sigma$. There is, then, a simple one-to-one correspondence between the set of words over $\Sigma \times \Sigma$ and the set of all pairs of words over Σ of the same length. This can be achieved by mapping a pair (v, w) of words Σ with $|v| = |w| = n$ to a word $v \otimes w$, defined as $(v_1, w_1)(v_2, w_2) \cdots (v_n, w_n)$ whenever $v = v_1 \cdots v_n$ and $w = w_1 \cdots w_n$.

Proving that a property φ holds over a transition system $\mathfrak{S}$ is done "in a regular way,", by finding a "regular proof" for the property. For example, if φ asserts that the set Bad of bad states can never be reached, then a regular proof amounts to finding an inductive invariant Inv in the form of a regular language [3,42] that does not intersect with Bad, i.e., $Bad \cap Inv = \emptyset$, $S_0 \subseteq Inv$ (S_0 is a regular set of initial states), and $post_{\rightarrow}(Inv) \subseteq Inv$, where $\rightarrow = \rightarrow_1 \cup \rightarrow_2$. Since regular languages are effectively closed under boolean operations and taking pre/post images w.r.t. regular transducers, an algorithm for verifying the correctness of a given regular proof can be obtained by using language inclusion algorithms for regular automata, e.g., [14,2]. The framework of regular proofs is incomplete in general since it could happen that there is a proof, but no regular proof. The pathological cases when only non-regular proofs exist do not, however, seem to frequently occur in practice, e.g., see [12,16,41,3,47,9,10,42,37].

The framework of regular proofs has been extended to deal with almost-sure termination for weakly-finite probabilistic concurrent programs in [36]. We briefly summarise the main idea, since we reduce the fair termination problem to their setting. By Proposition 1, the actual probability values do not matter in proving almost-sure termination. For this reason, we may specify a weakly-finite MDP $\mathfrak{S} = \langle S = V_1 \cup V_2; \rightarrow_1, \rightarrow_2 \rangle$ as a regular specification in the same way as we specify a non-probabilistic transition system in our regular specification language. Given an MDP $\mathfrak{S} = \langle S = V_1 \cup V_2; \rightarrow_1, \rightarrow_2 \rangle$, a set $I_0 \subseteq V_1$ of initial configurations, and a set $F \subseteq V_1$ of final configurations, a regular proof for $\text{Prob}(s_0 \models F) = 1$ for each $s_0 \in I_0$ is a pair $\langle Inv, \prec \rangle$ consisting of a regular inductive invariant $Inv \subseteq S$ and a regular relation $\prec \subseteq S \times S$ such that:

1. $I_0 \subseteq Inv$ and $post_{\rightarrow}(Inv) \subseteq Inv$.
2. $\prec$ is a strict preorder on S, i.e., it is irreflexive ($\forall s \in S : s \not\prec s$) and transitive ($\forall s, s', s'' \in S : s \prec s' \wedge s' \prec s'' \rightarrow s \prec s''$).
3. irrespective of the nondeterministic transitions from any configuration in Inv, there is a probabilistic transition to a configuration in Inv that decreases its rank with respect to $\prec$:

$$\forall x \in Inv \setminus F, y \in S \setminus F : \quad \big((x \rightarrow_1 y) \Rightarrow (\exists z \in Inv : (y \rightarrow_2 z) \wedge x \succ z)\big) .$$

An automata-theoretic algorithm can then be devised for checking the above verification conditions with respect to a given regular proof [36].

Example 1. **[Herman's protocol, continued]** We provide a regular encoding of Herman's protocol. The configurations are words over the alphabet $\{\top, \bot, \overline{\top}, \overline{\bot}\}$, where $\top$ (resp. $\bot$) signifies that a process holds (resp. does not hold) a token, while overlining

the character signifies that the process is chosen by the scheduler. We set $\Sigma = \{\top, \bot\}$. The set S_0 of initial configurations is $\Sigma^* \top \Sigma^*$, i.e., at least one process holds a token. The set of final configurations is $\bot^* \top \bot^*$, i.e., there is only a single token in the system. The actions of the scheduler is to choose a process; this can be expressed as the regular expression $I^*((\top, \overline{\top}) + (\bot, \overline{\bot}))I^*$, where I denotes the regular language $(\top, \top) + (\bot, \bot)$. The probabilistic actions can be expressed as a union of the following three regular expressions:

$$I^*((\overline{\top}, \top) + (\overline{\bot}, \bot))I^* \qquad \textbf{(idle)}$$

$$I^*(\overline{\top}, \bot)((\bot, \top)) + (\top, \top))I^*, \quad ((\bot, \top) + (\top, \top))I^*(\overline{\top}, \bot)) \qquad \textbf{(pass token right)}$$

5 Handling Fairness Requirements

We now describe the main result of the paper: a general method for embedding finitary fairness into regular model checking for probabilistic concurrent systems.

5.1 Regular Specifications of Fairness

When a complex system or a distributed protocol is being modelled in regular model checking, it is often necessary to add an *infinite* number of fairness requirements. This is because such a system admits a finite but arbitrary number of agents or processes, each with its own fairness requirement (e.g. that the process should be executed infinitely often). For this reason, it is not enough to simply express the fairness requirements as a finite set of pairs of regular languages (one for the premise, and one for the consequence). We describe a regular way of specifying infinitely many fairness constraints. Our presentation is a generalisation of the regular specification of fairness from [4,42].

The general idea is to define a "regular function" $\mathcal{T}$ that maps a configuration $s = s_1 \cdots s_n \in S$ to a word $w = w_1 \cdots w_n$ such that w_i contains: (1) a bit b_i indicating whether s is in the premise of the i-th fairness requirement, (2) a bit b'_i indicating whether s is in the consequence of the i-th fairness requirement, and (3) a bit t indicating whether the i-th fairness requirement is justice or compassion. Such a regular specification of fairness allows an infinite number of fairness constraints since S is potentially infinite (i.e., containing words of unbounded lengths), though only the first $|s|$ fairness requirements matter for a word $s \in S$. This is sufficient for weakly-finite MDPs since the set of reachable configurations from any given configuration s is finite and so, among the infinite number of fairness constraints, only finitely many are distinguishable. The regular function can be defined by a letter-to-letter tranducer with input alphabet Σ and output alphabet $\Gamma := \{0,1\} \times \{0,1\} \times \{0,1\}$. Without loss of generality, we assume that the i-th letter in the output of every input word of $\mathcal{T}$ agree on the third bit (i.e., whether the fairness requirement is justice or compassion is well-defined): for every $s, s' \in S$ and $i \in \mathbb{N}$, if $\mathcal{T}(s)[i] = (a, b, c)$ and $\mathcal{T}(s')[i] = (a', b', c')$, then $c = c'$. Observe this condition on $\mathcal{T}$ can be algorithmically checked by using a simple automata-theoretic method: find two accepted words in which in some position their third bits differ.

In this case, $\mathcal{T}$ gives rise to compassion requirements $\mathfrak{C}$ and justice requirements $\mathfrak{J}$ by associating the i-th position in all output words by a unique fairness constraint.

More precisely, let $A_i = \{s : \mathcal{T}(s)[i] = (1, j, t), \text{ for some } j, t \in \{0, 1\}\}$ and $B_i = \{s : \mathcal{T}(s)[i] = (j, 1, t), \text{ for some } j, t \in \{0, 1\}\}$. Define: (i) $\mathfrak{J} = \{\Diamond\Box A_i \Rightarrow \Box\Diamond B_i : \mathcal{T}(s)[i] = (i, j, 0), \text{ for some } s \in S, \text{ for some } j \in \{0, 1\}\}$, (ii) $\mathfrak{C} = \{\Box\Diamond A_i \Rightarrow \Box\Diamond B_i : \mathcal{T}(s)[i] = (i, j, 1), \text{ for some } s \in S, \text{ for some } j \in \{0, 1\}\}$. Therefore, by Proposition 1, our regular fairness specification allows us to define weakly-finite fair MDPs $\langle S = V_1 \cup V_2; \rightarrow_1, \rightarrow_2, \mathfrak{C}, \mathfrak{J}\rangle$. In the following, we shall call such fair MDPs *regular*.

Our main theorem is a regularity-preserving reduction from proving almost sure termination for regular FMDPs (under finitary fairness) to proving almost sure termination for regular MDPs (without fairness).

Theorem 1. *Let $\mathfrak{S} = \langle S = V_1 \cup V_2; \rightarrow_1, \rightarrow_2, \mathfrak{C}, \mathfrak{J}\rangle$ be a regular representation of an FMDP, $I_0 \subseteq V_1$ be a regular set of initial configurations, and $F \subseteq V_1$ be a regular set of final configurations. Then one can compute a regular presentation of MDP $\mathfrak{S}' = \langle S = V_1' \cup V_2'; \rightsquigarrow_1, \rightsquigarrow_2\rangle$ and two regular sets $I_0', F' \subseteq V_1'$ such that it holds that if $\mathfrak{C}$ and $\mathfrak{J}$ are realizable, then $Prob_{\mathfrak{S}'}(I_0' \models \Diamond F') = 1$ iff $Prob_{\mathfrak{S}}(I_0 \models \Diamond F) = 1$.*

5.2 Abstract Program Transformation

Before proving Theorem 1, let us first recall an abstract program transformation *à la* Alur & Henzinger [5], which encodes finitary fairness into a program using integer counter variables. Intuitively, we reserve one variable for each atomic fairness condition as an "alarm clock" that will set off if its corresponding process has not been executed for a long time, and one global variable n that acts as a *default* value to reset a clock to as soon as the corresponding process is executed. Although Alur & Henzinger [5] did not discuss about probabilistic programs, their transformation can be easily adapted to the setting of MDPs, though correctness still has to be proven.

We now elaborate on the details of the transformation. Given an FMDP $\mathfrak{S} = \langle S = V_1 \cup V_2; \rightarrow_1, \rightarrow_2, \mathfrak{C}, \mathfrak{J}\rangle$ with a probability distribution δ, the transformation will produce an MDP $\mathfrak{S}' = \langle S = V_1' \cup V_2'; \rightsquigarrow_1, \rightsquigarrow_2\rangle$ with a probability distribution δ' as follows. Introduce a set $\mathcal{V}$ of "counter" variables that range over natural numbers: x_j (for each $j \in \mathfrak{J}$), y_c (for each $c \in \mathfrak{C}$), and n. Let $\mathfrak{F}$ be the set of all valuations f mapping each variable in $\mathcal{V}$ to a natural number such that $f(x_j), f(y_c) \leq f(n)$ for each $j \in \mathfrak{J}$ and $c \in \mathfrak{C}$. We define $V_1' = V_1 \times \mathfrak{F}$ and $V_2' = V_2 \times \mathfrak{F}$. We now define the transition relation $\rightsquigarrow_i$ such that $(s, f) \rightsquigarrow_i (s', f')$ if $s \rightarrow_i s'$ and

- for each $z \in \mathcal{V}$, $f(z) > 0$,
- $f'(n) := f(n)$,
- x_j (for $j = (A, B) \in \mathfrak{J}$) and y_c (for $c = (A, B) \in \mathfrak{C}$) change as follows:

$$f'(x_j) = \begin{cases} f(x_j) - 1 & \text{if } s \in A \cap \overline{B} \\ f(n) & \text{if } s \in \overline{A} \cup B \end{cases} \qquad f'(y_c) = \begin{cases} f(n) & \text{if } s \in \overline{A} \cap \overline{B} \\ f(y_c) - 1 & \text{if } s \in A \cap \overline{B} \\ f(n) & \text{if } s \in B \end{cases}$$

($\overline{A}$ denotes the set-complement of A). Finally, we define the probability distribution δ' underlying $\rightsquigarrow_2$ as $\delta'((s, f), (s', f')) = \delta(s, s')$ whenever $s \in V_2$.

Lemma 1. *If $\mathfrak{S}$ is a weakly-finite FMDP, then $\mathfrak{S}'$ is weakly-finite.*

Intuitively, the variables x_j's and y_c's keep track of how long the scheduler has delayed choosing an enabled process, while the variable n (unchanged once the initial configuration of the MDP is fixed) aims to ensure that the scheduler is n-fair. Since n is a variable (not a constant), the resulting MDP $\mathfrak{S}'$ captures precisely the behaviour of $\mathfrak{S}$ under fin-fair schedulers.

We next state a correctness lemma for the transformation (proof in the full version). To this end, given a set $S_0 \subseteq S$ of initial configurations in $\mathfrak{S}$, we define:

- $S'_0 := S_0 \times \mathfrak{F}_=$, where $\mathfrak{F}_=$ contains functions $f \in \mathfrak{F}$ such that $f(x_j) = f(y_c) = f(n)$ for each $j \in \mathfrak{J}$ and $c \in \mathfrak{C}$.
- $F' = (F \times \mathfrak{F}_{>0}) \cup (S \times \mathfrak{F}_0)$, where $\mathfrak{F}_0$ contains all $f \in \mathfrak{F}$ such that $f(x_j) = 0$ for some $j \in \mathfrak{J}$ or $f(y_c) = 0$ for some $c \in \mathfrak{C}$ (i.e. one of the alarms has been triggered), and $\mathfrak{F}_{>0} := \mathfrak{F} \setminus \mathfrak{F}_0$.

Lemma 2 (Correctness). *For weakly-finite fair MDPs $\mathfrak{S}$, it is the case that* $\mathit{Prob}_{\mathfrak{S}}(S_0 \models \Diamond F) = \mathit{Prob}_{\mathfrak{S}'}(S'_0 \models \Diamond F')$.

These two lemmas immediately imply Theorem 1.

5.3 Finitary Fairness in Regular Model Checking

We now show how to implement the aforementioned abstract program transformation in our regular model checking framework. Fix a regular presentation of an FMDP $\mathfrak{S} = \langle S = V_1 \cup V_2; \rightarrow_1, \rightarrow_2, \mathfrak{C}, \mathfrak{J}\rangle$, which includes two automata over the alphabet $\Sigma \times \Sigma$ representing $\rightarrow_1$ and $\rightarrow_2$, and an automaton over the alphabet $\Sigma \times \Gamma$ representing the regular specification of the fairness conditions $\mathfrak{C}$ and $\mathfrak{J}$. [Recall that $\Gamma := \{0,1\} \times \{0,1\} \times \{0,1\}$.] We describe the construction of $\rightsquigarrow_1$ (the construction for $\rightsquigarrow_2$ is similar). Let $\mathcal{A} = (\Sigma \times \Sigma, Q, \Delta, q_0, F)$ be an automaton representing $\rightarrow_1$ and $\mathcal{A}^f = (\Sigma \times \Gamma, Q^f, \Delta^f, q^f_0, F^f)$ be an automaton representing the regular specification of fairness. The construction of the automaton for $\rightsquigarrow_1$ has two stages.

Stage 1: compute an intermediate automaton. The intermediate automaton $\mathcal{B}$ will have the alphabet $\Sigma' := (\Sigma \times \Sigma) \cup \Gamma$ and recognize a subset of $((\Sigma \times \Sigma)\Gamma)^*$. Intuitively, on input $(a,b) \in \Sigma \times \Sigma$, the automaton $\mathcal{B}$ simultaneously runs both $\mathcal{A}$ and $\mathcal{A}^f$. Here, the automaton $\mathcal{A}^f$ will nondeterministically guess a letter $c \in \Gamma$ and make a transition on the letter (a,c). The automaton $\mathcal{B}$, then, immediately consumes the letter c. This process is repeated until both $\mathcal{A}$ and $\mathcal{A}^f$ accept. More precisely, the automaton is defined as $\mathcal{B} := (\Sigma', Q^B, \Delta^B, q^B_0, F^B)$ where:

- $Q^B = Q \times Q^f \times (\Gamma \cup \{?\})$, $q^B_0 = (q_0, q^f_0, ?)$, and $F^B = F \times F^f \times \{?\}$
- Δ_B has the following transitions:
  - $((p_1, q^f_1, ?), (a,b), (p_2, q^f_2, c))$ if $(p_1, (a,b), p_2) \in \Delta$ and $(q^f_1, (a,c), q^f_2) \in \Delta^f$.
  - $((p, q^f, c), c, (p, q^f, ?))$ for each $c \in \Gamma$.

Stage 2: regular substitution of letters in Γ. Define the following regular languages

- (Identity) $\texttt{ID} := (1,1)^+(?,?)^*$,
- (Decrement) $\texttt{DEC} := (1,1)^*(1,?)(?,?)^*$, and
- (Reset) $\texttt{RES} := (1,1)^+(?,1)^*$.

Define the regular substitution σ mapping letters in Γ to regular languages:

- if (x, y, z) is $(i, 1, j)$ or $(0, i, 0)$ (for $i, j \in \{0, 1\}$), then $\sigma((x, y, z)) = \texttt{RES}$.
- if (x, y, z) is of the form $(1, 0, i)$ (for some $i \in \{0, 1\}$), then $\sigma((x, y, z)) = \texttt{DEC}$.
- define $\sigma((0, 0, 1)) = \texttt{ID}$.

We then apply the regular substitution σ to the letters Γ appearing in our intermediate automaton $\mathcal{B}$. The resulting automaton implements the desired automaton for $\rightsquigarrow_1$.

Finishing off the rest of the construction. Computing S_0', F' is easy. Define S_0' to be the set of all words $a_1 w_1 a_2 w_2 \cdots a_m w_m$ — where $a_i \in \Sigma$ and $w_i \in 1^+$ for each $i \in \{1, \ldots, m\}$ — such that $a_1 \cdots a_m \in S_0$. Similarly, define F' to be the set of all words $a_1 w_1 a_2 w_2 \cdots a_m w_m$ — where $a_i \in \Sigma$ and $w_i \in 1^+?^* \cup ?^+$ for each $i \in \{1, \ldots, m\}$ — such that $a_1 \cdots a_m \in F$ or $w_i \in ?^+$ for some $i \in \{1, \ldots, m\}$. Regular automata for these sets could be easily constructed given automata for S_0 and F.

Example 2. **[Herman's protocol]** We encode process fairness in the following way. The counters use the unary encoding, their values represented as the lengths of sequences of 1's padded on the right by the symbol ? (crucial to keep the transducers length-preserving). For example, the number 3 is represented by any word of the form $111?^*$. Define $\mathcal{X} = 1^*?^*$, i.e., the set of all valid counters. The set of initial configurations can be expressed using the regular expression $(\Sigma \cdot \mathcal{X})^*(\top \cdot \mathcal{X})(\Sigma \cdot \mathcal{X})^*$, i.e., counters for all processes are initialized to an arbitrary value. The set of final configurations is now $(\bot \cdot \mathcal{X})^*(\top \cdot \mathcal{X})(\bot \cdot \mathcal{X})^* \cup (\Sigma \cdot \mathcal{X})^*(\Sigma \cdot ?^*)(\Sigma \cdot \mathcal{X})^*$, i.e., either there is exactly one token in the system, or (at least) one counter has reached 0. Scheduler now also performs operations on the counters for processes: for a chosen process, the counter is reset, for other processes, the counter is decremented. This can be expressed as the language $(I \cdot \texttt{DEC})^*\big(((\bot, \overline{\bot}) + (\top, \overline{\top})) \cdot \texttt{RES}\big)(I \cdot \texttt{DEC})^*$. Actions of the protocol are the same as in the original encoding and the values of counters are left unmodified:

$$(I \cdot \texttt{ID})^*\big(((\overline{\bot}, \bot) + (\overline{\top}, \top)) \cdot \texttt{ID}\big)(I \cdot \texttt{ID})^* \quad \textbf{(idle)}$$

$$(I \cdot \texttt{ID})^*\big((\overline{\top}, \bot) \cdot \texttt{ID}\big)\big(((\bot, \top) + (\top, \top)) \cdot \texttt{ID}\big)(I \cdot \texttt{ID})^* \quad \textbf{(pass token right}_1\textbf{)}$$

$$\big(((\bot, \top) + (\top, \top)) \cdot \texttt{ID}\big)(I \cdot \texttt{ID})^*\big((\overline{\top}, \bot) \cdot \texttt{ID}\big) \quad \textbf{(pass token right}_2\textbf{)}$$

At this point, we can use existing tools for checking termination (without fairness constraints), e.g. [36]. Indeed, we can automatically check that the system after reduction terminates with probability one, thus proving that Herman's protocol fairly terminates with probability one (under finitary process-fair schedulers).

6 Implementation and Experiments

The approach presented in this paper has been implemented in the tool FAIRYTAIL.[7] For evaluation, we extracted models of a number of probabilistic parameterized systems. The tool receives a system with fairness conditions and transforms it into a system without fairness conditions, where fairness of the original system is encoded using counters. For solving liveness in the output transformed system, we use SLRP [36] (in the *incremental liveness proofs* setting) as the underlying liveness checker for parameterized systems.

[7] https://github.com/uuverifiers/autosat/tree/master/Fairness

Table 6 shows the results of our experiments. The times given are the wall clock times for the individual benchmarks on a PC with 4 Quad-Core AMD Opteron 8389 processors with Java heap memory limited to 64 GiB. The time included translation of the system into a system without fairness (always less than 1s) and the runtime of SLRP.

Table 1. Times of analyses of probabilistic paremeterised systems. The timeout was set to 10 hours (timeout is denoted as T/O).

Case study	Time
Herman's protocol (merge, line)	3.64 s
Herman's protocol (annih., line)	4.33 s
Herman's protocol (merge, ring)	4.31 s
Herman's protocol (annih., ring)	4.61 s
Moran process (2 types, line)	2 m 48 s
Moran process (3 types, line)	56 m 14 s
Cell cycle switch (1 types, line)	43.94 s
Cell cycle switch (2 types, line)	9 h 46 m
Clustering (2 types, line)	10 m 30 s
Clustering (3 types, line)	T/O
Coin game ($k = 3$, clique)	1 m 0 s

We consider two versions of *Herman's protocol* and two topologies. *Moran process*, is a model of genetic drift [40] with individuals of $N \geq 2$ types. When an individual is chosen by the scheduler, it can either idle or infect a neighbor. The model of *cell cycle switch* is a simplification of the model of [17]. Individuals can be committed to a decision from $N \geq 1$ types. An individual neighboring another one not sharing the same decision can make him *undecided*, or persuade an undecided individual to commit to his decision. *Clustering* considers a population model of alleles of $N \geq 2$ (resp. 3) types on a line, that can change position with their neighbours of a different type. *Coin game* is a population protocol where every agent has one of two types of coins. If an agent is chosen by the scheduler, it can change its coin to the one held by the majority of k other randomly selected agents.

In our experiments, we verify that a given property holds under every finitary process-fair scheduler with probability one. For *clustering*, the property is that the system eventually reaches a configuration with N clusters of the same type, while for the other population protocols, the property is that the system reaches a stable configuration.

The experiments show that our encoding of fairness into systems is viable and can be used for verification of parameterized systems with fairness by their reduction to systems without fairness. On the other hand, when the size of the regular proof is large, we observe that the problem for the underlying solver gets significantly more difficult (as can be seen on the example of *clustering* for three types of alleles). We conjecture that the performance can be improved significantly by making the solver take into account the (not arbitrary) structure of the problem, which we leave for future work.

Future work. We leave the reader with several research challenges. A natural question is how to deal with non-finitary fairness for parameterized probabilistic concurrent systems in general and in the framework of regular model checking. Secondly, since there are numerous examples of population models over more complex topologies (e.g. grids), how do you deal with termination and fair termination over such models in the parameterized setting?

Acknowledgement. We thank anonymous reviewers and Dave Parker for their helpful feedback. This work was supported by the Czech Science Foundation (project 16-24707Y), the BUT FIT project FIT-S-17-4014, the IT4IXS: IT4Innovations Excellence in Science project (LQ1602), Yale-NUS Starting Grant, the European Research Council under ERC Grant Agreement no. 610150, and Swedish Research Council (2014-5484).

References

1. PRISM website (referred in July 2015). `http://http://www.prismmodelchecker.org/`.
2. P. A. Abdulla, Y.-F. Chen, L. Holík, R. Mayr, and T. Vojnar. When simulation meets antichains. In *TACAS*, pages 158–174, 2010.
3. Parosh Aziz Abdulla. Regular model checking. *STTT*, 14(2):109–118, 2012.
4. Parosh Aziz Abdulla, Bengt Jonsson, Marcus Nilsson, Julien d'Orso, and Mayank Saksena. Regular model checking for LTL(MSO). *STTT*, 14(2):223–241, 2012.
5. Rajeev Alur and Thomas A. Henzinger. Finitary fairness. *ACM Trans. Program. Lang. Syst.*, 20(6):1171–1194, 1998.
6. Krzysztof R. Apt and Dexter Kozen. Limits for automatic verification of finite-state concurrent systems. *Inf. Process. Lett.*, 22(6):307–309, 1986.
7. Christel Baier and Joost-Pieter Katoen. *Principles of Model Checking*. MIT Press, 2008.
8. Christel Baier and Marta Z. Kwiatkowska. On the verification of qualitative properties of probabilistic processes under fairness constraints. *Inf. Process. Lett.*, 66(2):71–79, 1998.
9. S. Bardin, A. Finkel, J. Leroux, and L. Petrucci. FAST: acceleration from theory to practice. *STTT*, 10(5):401–424, 2008.
10. S. Bardin, A. Finkel, J. Leroux, and Ph. Schnoebelen. Flat acceleration in symbolic model checking. In *ATVA*, pages 474–488, 2005.
11. Nathalie Bertrand and Paulin Fournier. Parameterized verification of many identical probabilistic timed processes. In *FSTTCS'13*, volume 24 of *LIPIcs*, pages 501–513. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2013.
12. B. Boigelot and F. Herbreteau. The power of hybrid acceleration. In *CAV*, pages 438–451, 2006.
13. B. Boigelot, A. Legay, and P. Wolper. Iterating transducers in the large (extended abstract). In *CAV*, pages 223–235, 2003.
14. Filippo Bonchi and Damien Pous. Checking NFA equivalence with bisimulations up to congruence. In *POPL'13*, pages 457–468. ACM, 2013.
15. Rémi Bonnet, Stefan Kiefer, and Anthony Widjaja Lin. Analysis of probabilistic basic parallel processes. In *FOSSACS*, pages 43–57, 2014.
16. Ahmed Bouajjani, Peter Habermehl, Adam Rogalewicz, and Tomás Vojnar. Abstract regular (tree) model checking. *STTT*, 14(2):167–191, 2012.
17. Luca Cardelli and Attila Csikász-Nagy. The cell cycle switch computes approximate majority. *Scientific Reports*, 2(656), 2012.
18. Aleksandar Chakarov and Sriram Sankaranarayanan. Probabilistic program analysis with martingales. In *CAV*, pages 511–526, 2013.
19. Aleksandar Chakarov, Yuen-Lam Voronin, and Sriram Sankaranarayanan. Deductive proofs of almost sure persistence and recurrence properties. In *TACAS*, pages 260–279, 2016.
20. C. Courcoubetis and M. Yannakakis. Minimum and maximum delay problems in real-time systems. *Formal Methods in System Design*, 1(4):385–415, 1992.
21. Luca de Alfaro. Temporal logics for the specification of performance and reliability. In *STACS 97, Symposium on Theoretical Aspects of Computer Science*, volume 1200 of *Lecture Notes in Computer Science*, pages 165–176. Springer, 1997.
22. Javier Esparza. Parameterized verification of crowds of anonymous processes. *Dependable Software Systems Engineering*, 45:59–71, 2016.
23. Javier Esparza, Andreas Gaiser, and Stefan Kiefer. Proving termination of probabilistic programs using patterns. In *CAV*, pages 123–138, 2012.
24. Luis María Ferrer Fioriti and Holger Hermanns. Probabilistic termination: Soundness, completeness, and compositionality. In *POPL'15*, pages 489–501. ACM, 2015.

25. Wan Fokkink. *Distributed Algorithms*. MIT Press, 2013.
26. Nissim Francez. *Fairness*. Springer-Verlag New York, Inc., New York, NY, USA, 1986.
27. Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research [outcome of a Dagstuhl seminar, February 2001]*, volume 2500 of *Lecture Notes in Computer Science*. Springer, 2002.
28. Sergiu Hart, Micha Sharir, and Amir Pnueli. Termination of probabilistic concurrent program. *ACM Trans. Program. Lang. Syst.*, 5(3):356–380, 1983.
29. Ted Herman. Probabilistic self-stabilization. *Inf. Process. Lett.*, 35(2):63–67, 1990.
30. Jochen Hoenicke, Ernst-Rüdiger Olderog, and Andreas Podelski. Fairness for dynamic control. In *TACAS'10*, volume 6015 of *LNCS*, pages 251–265. Springer, 2010.
31. Amos Israeli and Marc Jalfon. Token management schemes and random walks yield self-stabilizing mutual exclusion. In *PODC*, pages 119–131, 1990.
32. Benjamin Lucien Kaminski, Joost-Pieter Katoen, Christoph Matheja, and Federico Olmedo. Weakest precondition reasoning for expected run-times of probabilistic programs. In *ESOP'16*, volume 9632 of *LNCS*, pages 364–389. Springer, 2016.
33. Marta Z. Kwiatkowska. Model checking for probability and time: from theory to practice. In *LICS*, page 351, 2003.
34. D. Lehmann and M. Rabin. On the advantage of free choice: A symmetric and fully distributed solution to the dining philosophers problem (extended abstract). In *POPL*, pages 133–138, 1981.
35. E Lieberman, C Hauert, and M A Nowak. Evolutionary dynamics on graphs. *Nature*, 433(7023):312–316, January 2005.
36. Anthony W. Lin and Philipp Rümmer. Liveness of randomised parameterised systems under arbitrary schedulers. In *CAV'16 (2)*, volume 9779 of *LNCS*, pages 112–133. Springer, 2016.
37. Anthony Widjaja Lin. Accelerating tree-automatic relations. In *FSTTCS*, pages 313–324, 2012.
38. Nancy A. Lynch, Isaac Saias, and Roberto Segala. Proving time bounds for randomized distributed algorithms. In *PODC*, pages 314–323, 1994.
39. David Monniaux. An abstract analysis of the probabilistic termination of programs. In *SAS*, pages 111–126. Springer, 2001.
40. Patrick A.P. Moran. Random processes in genetics. *Mathematical Proceedings of the Cambridge Philosophical Society*, 54(1):60–71, Jan 1958.
41. Daniel Neider and Nils Jansen. Regular model checking using solver technologies and automata learning. In *NFM*, pages 16–31, 2013.
42. M. Nilsson. *Regular Model Checking*. PhD thesis, Uppsala Universitet, 2005.
43. Ernst-Rüdiger Olderog and Krzysztof R. Apt. Fairness in parallel programs: The transformational approach. *ACM Trans. Program. Lang. Syst.*, 10(3):420–455, 1988.
44. Ernst-Rüdiger Olderog and Andreas Podelski. Explicit fair scheduling for dynamic control. In *Concurrency, Compositionality, and Correctness, Essays in Honor of Willem-Paul de Roever*, volume 5930 of *Lecture Notes in Computer Science*, pages 96–117. Springer, 2010.
45. Amir Pnueli and Lenore D. Zuck. Verification of multiprocess probabilistic protocols. *Distributed Computing*, 1(1):53–72, 1986.
46. A. W. To. *Model Checking Infinite-State Systems: Generic and Specific Approaches*. PhD thesis, LFCS, School of Informatics, University of Edinburgh, 2010.
47. Anthony Widjaja To and Leonid Libkin. Algorithmic metatheorems for decidable LTL model checking over infinite systems. In *FoSSaCS*, pages 221–236, 2010.
48. Moshe Y. Vardi. Automatic verification of probabilistic concurrent finite-state programs. In *FOCS*, pages 327–338, 1985.