A note on the attractor-property of infinite-state Markov chains

Christel Baier [a], Nathalie Bertrand [b], Philippe Schnoebelen [b]

[a] *Universität Bonn, Institut für Informatik I, Germany*

[b] *Lab. Specification & Verification, CNRS UMR 8643 & ENS de Cachan, France*

Abstract

In the past 5 years, a series of verification algorithms has been proposed for infinite Markov chains that have a finite attractor, i.e., a set that will be visited infinitely often almost surely starting from any state. In this paper, we establish a sufficient criterion for the existence of an attractor. We show that if the states of a Markov chain can be given levels (positive integers) such that the expected next level for states at some level $n > 0$ is less than $n - \Delta$ for some positive Δ, then the states at level 0 constitute an attractor for the chain. As an application, we obtain a direct proof that some probabilistic channel systems combining message losses with duplication and insertion errors have a finite attractor.

Key words: Theory of computation; Attractors in Markov chains; Verification of probabilistic systems; Lossy channel systems.

1 Introduction

In the past two decades, several methods for the automatic verification of systems modeled by finite Markov chains have been proposed, see, e.g., [22,15,10,23,9], and have been implemented in model checkers like PRISM [17]. A striking feature of these methods is that, for checking *qualitative* properties, [1] they are very similar to well-known methods used for classical model checking of nondeterministic systems modeled by finite transition systems. In particular, these methods are mainly concerned with what connected components are reachable from where, and the actual values of the probabilities appearing in the Markov chain are not relevant.

Email addresses: `baier@cs.uni-bonn.de` (Christel Baier), `bertrand@lsv.ens-cachan.fr` (Nathalie Bertrand), `phs@lsv.ens-cachan.fr` (Philippe Schnoebelen).

[1] I.e., properties holding with probability 1, "*almost surely*", or, dually, with probability 0.

The situation is not so easy with *infinite* Markov chains. Indeed, among the numerous infinite-state nondeterministic models investigated in the model checking literature, research on verification algorithms for infinite-state Markov chains is comparatively rare. Beside model checking algorithms for probabilistic timed automata and related models [4,18], we are only aware of two examples where researchers considered extending infinite-state models with probabilistic aspects: probabilistic pushdown systems and probabilistic lossy channel systems. While the methods for analyzing *probabilistic pushdown systems* usually abstract the set of configurations into finitely many classes that behave uniformly w.r.t. probabilistic aspects [12,11,13,14], the verification algorithms for *channel systems with probabilistic message losses* are based on the existence of a *finite attractor*, a set of configurations that will almost surely be visited infinitely often. In [5], the finite attractor property has been used for solving the qualitative LTL$\backslash$X model checking problem by a reachability analysis between the configurations of the attractor. The finite attractor approach of [5] was streamlined in [3,6] where a variant model for probabilistic message losses was introduced, and where several other kinds of unreliability in message transfers (spurious message duplications or insertions) were considered. For the same model, [20] considered the verification of qualitative properties and here again the positive results crucially rely on the existence of a finite attractor (see [21]). Recently Abdulla *et al.* presented a general framework for verifying qualitative and quantitative aspects in infinite Markov chains with a finite attractor [1].
Proving that a given set of configurations is an attractor is a required step in all the aforementioned works on probabilistic lossy channel systems. However only [2] gives a real proof (for channel systems with probabilistic losses). The proof is tedious and is not extended to other kinds of corruption, such as duplications or insertions, for which the attractor is stated without proof.

In the aforementioned examples, the underlying reason behind the existence of a finite attractor is the same: when enough messages are present in the channels, the system is more likely to lose messages than to create new ones (by writing, by spurious duplications, etc.). Hence, from "large" configurations, i.e., configurations with many messages, the systems tend to drift towards "small" configurations, with few messages.

Contribution of this paper. In the rest of this note we prove a general result that validates the above informal reasoning. We consider Markov chains where the set S of configurations is partitioned in "levels" $S_0, S_1, S_2, \ldots$ and show that if the system tends to go to smaller levels in the sense that, from any state in some level $k \neq 0$, the *expected next level* is less than k, then the lowest level is an attractor. This result can be seen as a variant of Foster's Theorem from martingales theory [8, Ch. 5], where we do not require strong connectivity, and for which we provide an elementary proof. As an application, we consider the channel systems with probabilistic message losses and duplications from [3,2] and prove the existence of a finite attractor under general conditions, thus providing the proofs that are omitted in all the aforementioned works.

2 A sufficient condition for the existence of an attractor

Markov chains. A Markov chain is a tuple $\mathcal{M} = (S, \mathbf{P})$ where S is a countable *state space* and $\mathbf{P} : S \times S \to [0,1]$ is the *transition probability matrix* where we require that $\sum_{t \in S} \mathbf{P}(s,t) = 1$ for any state $s \in S$. The intuitive operational behavior is that if the current state is s then the next state is chosen according to a probabilistic choice which selects state t with probability $\mathbf{P}(s,t)$. For $n \in \mathbb{N}$, we write $\mathbf{P}^n(s,t)$ for the probability to be in state t after exactly n steps when starting in state s. Formally, $\mathbf{P}^{n+1}(s,t) \stackrel{\text{def}}{=} \sum_{u \in S} \mathbf{P}^n(s,u) \cdot \mathbf{P}(u,t)$, starting from $\mathbf{P}^0(s,t) \stackrel{\text{def}}{=} 1$ if $s = t$, $\mathbf{P}^0(s,t) \stackrel{\text{def}}{=} 0$ otherwise. For $T \subseteq S$ we let $\mathbf{P}^n(s,T) \stackrel{\text{def}}{=} \sum_{t \in T} \mathbf{P}^n(s,t)$ and $\mathbf{P}(s,T) \stackrel{\text{def}}{=} \mathbf{P}^1(s,T)$. We assume here the standard sigma-field and probability measure (denoted $\Pr(s_0, \cdot)$) on the infinite paths starting in a given starting state s_0, see, e.g., [16,19]. If $T \subseteq S$ then $\Diamond T$ denotes the set of infinite paths in $\mathcal{M}$ that eventually visit T. $\Pr(s, \Diamond T)$ denotes the probability to reach T from state s. T is called an *attractor* for $\mathcal{M}$ iff $\Pr(s, \Diamond T) = 1$ for all $s \in S$. It then follows that, for any starting state s, almost surely the attractor T is visited infinitely often. Observe that S itself is a (trivial) attractor.

Left-oriented Markov chains. We deal here with a special type of markov chains where the state space S is partitioned into infinitely many levels labeled with non-negative integers. Formally, we assume a partition $S = \bigcup_{i \in \mathbb{N}} S_i$ with pairwise disjoint (possibly empty) subsets S_i of S. We refer to S_i as the i-th *level* in $\mathcal{M}$ and think of S_i as standing on the right of level S_{i-1} and on the left of level S_{i+1}. (However, there is no topological requirement that justifies the notions "left" or "right": transitions may go from any level S_i to any level S_j.) If $s \in S$ then $\mathit{level}(s)$ denotes the unique index $i \in \mathbb{N}$ such that $s \in S_i$. Then, $\mathbb{E}(s) \stackrel{\text{def}}{=} \sum_{j=0}^{\infty} \mathbf{P}(s, S_j) \cdot j$ denotes the *expected next level* for state s.[2] Assuming a given partition, $\mathcal{M}$ is called *left-oriented* iff there exists a positive constant $\Delta > 0$ such that $\mathbb{E}(s) \le \mathit{level}(s) - \Delta$ for all states $s \in S \setminus S_0$, i.e., all s with $\mathit{level}(s) \ge 1$.

Theorem 2.1 *For any left-oriented Markov chain $\mathcal{M}$, the leftmost level S_0 is an attractor.*

PROOF. We must show that $\Pr(s, \Diamond S_0) = 1$ for all states s. To simplify the following calculations, we assume that S_0 is a sink, i.e., once S_0 has been entered, it can never be left. Formally, $\mathbf{P}(s, S_0) = 1$ for all states $s \in S_0$. This is no loss of generality since, given an arbitrary $\mathcal{M}$, changing the outgoing transitions of the states in S_0 does not affect the probabilities to reach S_0 and hence does not influence whether

[2] This infinite series needs not converge. Hence, $\mathbb{E}(s) = +\infty$ is possible. However, we will only consider Markov chains where $\mathbb{E}(s)$ is finite for all $s \in S$.

S_0 is an attractor or not.
This assumption yields that $\mathbf{P}^n(s, S_0)$ is the probability to reach S_0 from s within n or less steps. Hence, $\Pr(s, \Diamond S_0) = \lim_{n\to\infty} \mathbf{P}^n(s, S_0)$. We now show by induction on n that for the expected level after n steps from state $s \in S \setminus S_0$ the following inequality holds:

$$\sum_{j=0}^{\infty} \mathbf{P}^n(s, S_j) \cdot j \;\leq\; \mathit{level}(s) \;-\; n\Delta \;+\; \Delta \sum_{\ell=1}^{n-1} \mathbf{P}^{\ell}(s, S_0) \qquad (*)$$

Here, Δ is the positive constant such that $\mathbb{E}(s) \leq n - \Delta$ for all states $s \in S \setminus S_0$.
For $n = 1$, (*) coincides with the statement that $\mathcal{M}$ is left-oriented. We now assume that $n \geq 2$ and that the induction hypothesis holds for $n-1$. Then, for all states $s \in S \setminus S_0$:

$$\begin{aligned}
& \sum_{j=0}^{\infty} \mathbf{P}^n(s, S_j) \cdot j \;=\; \sum_{j=1}^{\infty} \mathbf{P}^n(s, S_j) \cdot j \;=\; \sum_{j=1}^{\infty} \sum_{k=0}^{\infty} \sum_{t \in S_k} \mathbf{P}^{n-1}(s,t) \cdot \mathbf{P}(t, S_j) \cdot j \\
= \quad & \sum_{j=1}^{\infty} \sum_{k=1}^{\infty} \sum_{t \in S_k} \mathbf{P}^{n-1}(s,t) \cdot \mathbf{P}(t, S_j) \cdot j \quad \text{(since } \mathbf{P}(t, S_j) = 0 \text{ if } t \in S_0 \text{ and } j \geq 1) \\
= \quad & \sum_{k=1}^{\infty} \sum_{t \in S_k} \mathbf{P}^{n-1}(s,t) \cdot \underbrace{\sum_{j=1}^{\infty} \mathbf{P}(t, S_j) \cdot j}_{\leq \mathit{level}(t) - \Delta = k - \Delta} \\
\leq \quad & \sum_{k=1}^{\infty} \sum_{t \in S_k} \mathbf{P}^{n-1}(s,t) \cdot (k - \Delta) \\
= \quad & \sum_{k=1}^{\infty} \sum_{t \in S_k} \mathbf{P}^{n-1}(s,t) \cdot k \;-\; \Delta \cdot \underbrace{\sum_{k=1}^{\infty} \sum_{t \in S_k} \mathbf{P}^{n-1}(s,t)}_{= 1 - \mathbf{P}^{n-1}(s, S_0)} \\
\overset{\text{ind. hypo.}}{\leq} \quad & \mathit{level}(s) \;-\; (n-1)\Delta \;+\; \Delta \sum_{\ell=1}^{n-2} \mathbf{P}^{\ell}(s, S_0) \;-\; \Delta \cdot (1 - \mathbf{P}^{n-1}(s, S_0)) \\
= \quad & \mathit{level}(s) \;-\; n\Delta \;+\; \Delta \sum_{\ell=1}^{n-1} \mathbf{P}^{\ell}(s, S_0).
\end{aligned}$$

Let us now assume that $\Pr(s, \Diamond S_0) < 1$ for some state s. Then, $s \notin S_0$. Let n be a natural number with $n > \mathit{level}(s)/\Delta(1 - \Pr(s, \Diamond S_0))$. As $\mathbf{P}^{\ell}(s, S_0) \leq \Pr(s, \Diamond S_0)$ we get:

$$\begin{aligned}
& -n\Delta + \Delta \textstyle\sum_{\ell=1}^{n-1} \mathbf{P}^{\ell}(s, S_0) \;\leq\; -n\Delta + (n-1)\Delta \Pr(s, \Diamond S_0) \\
& = \; -\underbrace{n\Delta\big(1 - \Pr(s, \Diamond S_0)\big)}_{> \mathit{level}(s)} - \underbrace{\Delta \Pr(s, \Diamond S_0)}_{\geq 0} \;<\; -\mathit{level}(s)
\end{aligned}$$

Thus, inequality (*) for the expected level after n steps from s yields:

$$0 \;\leq\; \sum_{j=0}^{\infty} \mathbf{P}^n(s, S_j) \cdot j \;\leq\; \mathit{level}(s) - n\Delta \;+\; \Delta \sum_{\ell=1}^{n-1} \mathbf{P}^{\ell}(s, S_0) \;<\; \mathit{level}(s) - \mathit{level}(s) \;=\; 0$$

This is a contradiction! Hence, $\Pr(s, \Diamond S_0) = 1$ for all $s \in S$. $\square$

Almost left-oriented chains. In some situations the requirement that $\mathbb{E}(s) \le level(s) - \Delta$ for all states s such that $level(s) > 0$, is too restrictive. In fact, the requirement can be relaxed as follows. We say $\mathcal{M}$ is *almost left-oriented* if there is some positive constant Δ and some $n_0 \in \mathbb{N}$ such that $\mathbb{E}(s) \le level(s) - \Delta$ for all states s where $level(s) > n_0$.

Theorem 2.2 *If $\mathcal{M}$ is almost left-oriented (for a given n_0), the levels S_i, $i \le n_0$, are finite and S_0 is reachable from all states in $S_1 \cup \cdots \cup S_{n_0}$, then S_0 is an attractor.*

PROOF. We consider the partitioning $S'_0 = \bigcup_{i \le n_0} S_i$, $S'_i = \emptyset$ for $1 \le i \le n_0$ and $S'_i = S_i$ for all $i > n_0$. For this new partitioning, $\mathcal{M}$ is left-oriented. Thus, we may apply Theorem 2.1 to obtain that S'_0 is an attractor. Hence, independent on the starting state, S'_0 is visited infinitely often with probability 1. Now, since S'_0 is finite and S_0 is reachable from each state in S'_0, visiting S'_0 infinitely often entails visiting S_0 almost surely. $\square$

Chains with no orientation. The requirement that Δ is strictly positive may seem too restrictive. However, if we consider partitions such that $\mathbb{E}(s) \le level(s)$ for all $s \notin S_0$, no general statement can be made. For example, the classic 1-dimensional random walk with a barrier at 0 [3] has $\mathbb{E}(i) = i$ for all $i > 0$ and it admits $\{0\}$ as an attractor. On the other hand, the Markov chain $\mathcal{M} = (\mathbb{N}, \mathbf{P})$ with $\mathbf{P}(i,i) = 1$ for all $i \in \mathbb{N}$ has $\mathbb{E}(i) = i$ but it has no nontrivial attractors.

Right-oriented chains. One might expect that if the Markov chain tends to the right, in the sense that $\mathbb{E}(s) > level(s)$ for all states s, then no nontrivial attractors exist. However, no such general statement is possible. There is even *no* function $f : \mathbb{N} \to \mathbb{N}$ such that $\mathbb{E}(s) > f(level(s))$ for all states s can ensure that the given Markov chain has no nontrivial attractors. To see why, consider a Markov chain with state space $S = \mathbb{N}$ and the partition $S_i = \{i\}$, $i = 0, 1, \ldots$ For the transition probabilities we assume that $\mathbf{P}(i,0) = \frac{1}{2}$ and $\mathbf{P}(i, 4f(i)) \ge \frac{1}{4}$. The remaining probabilities for the successors of i are arbitrary. E.g., we may assume that i has an edge to all states $j \in S$, or we simply may deal with $\mathbf{P}(i, 4f(i)) = \frac{1}{2}$. We then have $\mathbb{E}(i) \ge f(i)$ for all states i, but $S_0 = \{0\}$ is an attractor for $\mathcal{M}$ as S_0 is reachable from any state with probability $\frac{1}{2}$ by a single transition.

[3] Formally, we consider the chain $(\mathbb{N}, \mathbf{P})$ where $\mathbf{P}(i, i+1) = \mathbf{P}(i, i-1) = \frac{1}{2}$ for $i \ge 1$ and $\mathbf{P}(0,1) = 1$. For this, it is known that the probability to visit eventually state 0 is 1 for any starting state i.

3 Probabilistic lossy channel systems

Channel systems [7] are a natural model for asynchronous systems that communicate by messages sent along FIFO links. In this section, we will only give an informal description of them and refer to the survey paper [21] (and the references therein) for motivations and formal definitions.

A channel system is made up of some number of finite-state component that communicate through some number of channels. Fig. 1 displays a schematic example with $n = 2$ components and $m = 2$ channels. The operational semantics of a

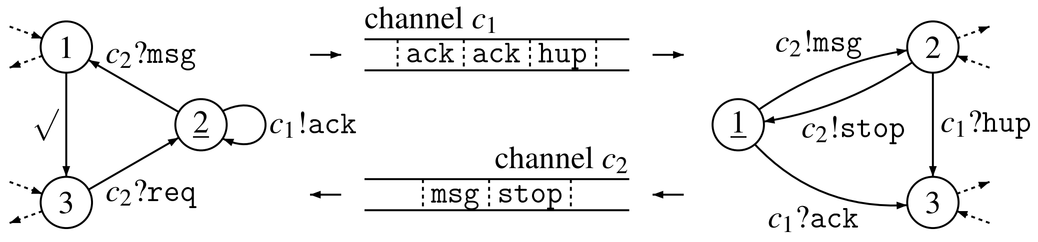

Fig. 1. A channel system

channel system is given by a transition system where a *configuration* is a tuple $s = \langle q_1, \ldots, q_n, w_1, \ldots, w_m \rangle$ of *n local control states* and *m channel contents*. Here w_i is a word over the alphabet of messages, describing what messages are currently in transit in the *i*-th channel. For example, the current configuration in Fig. 1 is $\langle 2,\ 1,\ \texttt{hup} \cdot \texttt{ack} \cdot \texttt{ack},\ \texttt{msg} \cdot \texttt{stop} \rangle$.
We often use vector notation $s = \langle \mathbf{q}, \mathbf{w} \rangle$ for configurations. The *size* $|\langle \mathbf{q}, \mathbf{w} \rangle|$ of a configuration is the number of messages currently in $\mathbf{w}$, i.e., $\sum_{i=1}^{m} |w_i|$. We write S for the set of all configurations (of some system), and S_k, where $k \in \mathbb{N}$, for the subset of configurations having size k. Transitions $s \to s'$ are defined in the obvious way: components asynchronously change their local states by following the edges of their control graph and performing the action labeling the edge. Send actions $c!m$ enqueue message m in channel c, read actions $c?m$ consume m from the head of c (this is only allowed if indeed m is the first message in c, hence send actions also act as guards), and null actions $\surd$ do not test or modify the channels. For simplicity, we only consider deadlock-free channel systems, i.e., systems where every configuration has at least one possible step allowed.

Probabilistic lossy channels systems. The model we consider here is a probabilistic extension where the choice of the next step (and of the performing component) is made probabilistically. Additionally, transmission errors can occur: messages can be lost from the channels, they can be duplicated spuriously, all this according to some probabilistic laws. More formally, a *probabilistic lossy channel system* (a PLCS) is a channel system equipped with a probability distribution θ for the choice of the next step, with a *message loss probability* $\tau \in [0,1]$, and with a *message duplication probability* $\lambda \in [0,1]$. Usually θ is given by assigning positive weights to the rules of the components (the edges in their control graph): such weights translate into a mapping $\theta : S \to \mathsf{Dist}(S)$ in the standard way (recall that $\mathsf{Dist}(S)$ is the set of *probability distributions* over S). After a next step is chosen probabilistically, message losses and message duplications may occur spuriously. More precisely, each message in $\mathbf{w}$ is lost with probability τ, and each remaining message is duplicated with probability λ.

The operational semantics of a PLCS L is given by a Markov chain $\mathcal{M}_{\mathsf{L}} = (S, \mathbf{P})$ where the states are the configurations of L. Defining the probabilities $\mathbf{P}(s, s')$ is tedious because one has to combine steps of the components, losses, and duplications, and because there usually exist several different way of reaching a same s' by one step with losses and duplications. We do not recall the definition here (see [3,2]). However, some properties of $\mathbf{P}$ can be explained without a full definition: given some s with $|s| = n$, the probability $\mathbf{Q}(s, S_\ell)$ that one round of losses and duplications transforms s into some s' with $|s'| = \ell$, written $\mathbf{Q}(n, \ell)$, is given by

$$\mathbf{Q}(n,\ell) \;=\; \sum_{i=n-\ell}^{\lfloor n-\frac{\ell}{2}\rfloor} \tau^i \cdot (1-\tau)^{n-i} \cdot \binom{n}{i} \cdot \lambda^{\ell-(n-i)} \cdot (1-\lambda)^{2(n-i)-\ell} \cdot \binom{n-i}{\ell-(n-i)} \quad (+)$$

The summation in (+) is nonempty if $\ell \leq 2n$ and we have $\mathbf{Q}(n, \ell) = 0$ if $\ell > 2n$. (+) can be explained as follows: it considers that first exactly i messages are lost, and then exactly $\ell - (n - i)$ of the remaining $n - i$ messages are duplicated, which gives a total amount of

$$2\underbrace{(\ell-(n-i))}_{\text{duplicated}} + \underbrace{((n-i)-(\ell-(n-i)))}_{\text{not duplicated}} \;=\; 2\ell - 2(n-i) + 2(n-i) - \ell \;=\; \ell$$

messages. Here, i has to fulfill the constraint $n - i \leq \ell \leq 2(n - i)$. One derives the expected size of the channels contents after losses and duplications:

$$\sum_{j=0}^{\infty} \mathbf{Q}(n,j) \cdot j \;=\; n - n\tau + n(1-\tau)\lambda \;=\; n(1-\tau)(1+\lambda). \qquad (3.1)$$

For a configuration s of L, let us write $p_!(s)$, $p_?(s)$ and $p_{\surd}(s)$ for the probabilities that the next step will be, respectively, a send action, a read action, or an internal action. These probabilities only depend on θ and, for all s, $p_!(s) + p_?(s) + p_{\surd}(s) = 1$. Furthermore $p_?(s) = 0$ when $|s| = 0$. Now, the transition probability matrix $\mathbf{P}$ in $\mathcal{M}_{\mathsf{L}}$ satisfies for $|s| = n$:

$$\mathbf{P}(s, S_\ell) = p_!(s)\mathbf{Q}(n+1,\ell) + p_?(s)\mathbf{Q}(n-1,\ell) + p_{\surd}(s)\mathbf{Q}(n,\ell) \qquad (3.2)$$

Theorem 3.1 *For any PLCS* L *with* $(1-\tau)(1+\lambda) < 1$, S_0*, the set of configurations where all channels are empty, is an attractor of* $\mathcal{M}_{\mathsf{L}}$.

PROOF. We apply Theorem 2.2 for the partition $(S_i)_{i\in\mathbb{N}}$ induced by channels contents size, and we show that it makes $\mathcal{M}_{\mathsf{L}}$ almost left-oriented. Consider a configuration s with $n = |s| \geq 1$. Combining (3.1) and (3.2) yields

$$\mathbb{E}(s) \;=\; \sum_{j=0}^{\infty} \mathbf{P}(s, S_j) \cdot j \;=\; (n + p_!(s) - p_?(s))(1-\tau)(1+\lambda) \;\leq\; (n+1)(1-\tau)(1+\lambda).$$

If $(1-\tau)(1+\lambda) < 1$ then there exists $n_0 \in \mathbb{N}$ such that $(n+1)(1-\tau)(1+\lambda) \;\leq\; n - \frac{1}{2}$ for all $n \geq n_0$. Hence $\mathbb{E}(s) \leq \mathit{level}(s) - \frac{1}{2}$ for all s with $\mathit{level}(s) \geq n_0$, $\mathcal{M}_{\mathsf{L}}$ is almost left-oriented, and S_0 is an attractor. $\square$

Theorem 3.1 proves that S_0 is an attractor in $\mathcal{M}_{\mathsf{L}}$ when $\lambda = 0$ and $\tau > 0$, i.e., for PLCS's without duplication errors as used in [6]. One also sees that if $\tau > 0$ and $\lambda \leq \tau$, i.e., if duplication errors are not more likely than message losses, then again S_0 is an attractor in $\mathcal{M}_{\mathsf{L}}$. This result was given in [3, lemma 10] without an explicit proof. But Theorem 3.1 shows that S_0 is an attractor under even weaker conditions: it is enough that $\lambda < \frac{\tau}{1-\tau}$.

PLCS with insertion errors. As in [2], we may also consider PLCS where *insertion errors* may occur in any step after losing and duplicating certain messages. In the approach of [2], there is a fixed distribution that specifies the probabilities for an insertion of k messages at any configuration (and hence independently of said configuration). Let K be the expected number of inserted messages. Then, the expected next level for any state s at level n is $\mathbb{E}(s) \leq (n+1)(1-\tau)(1+\lambda) + K$. Again, if $(1-\tau)(1+\lambda) < 1$ then $\mathbb{E}(s) \leq \mathit{level}(s) - \frac{1}{2}$ for all configurations of level higher than some n_0 large enough. Thus, $\mathcal{M}_{\mathsf{L}}$ is almost left-oriented and S_0 is an attractor.

References

[1] P. Abdulla, N. Ben Henda, and R. Mayr. Verifying infinite Markov chains with a finite attractor or the global coarseness property. In *Proc. Logic in Computer Science (LICS 2005)*, pages 127–136, 2005.

[2] P. A. Abdulla, N. Bertrand, A. Rabinovich, and Ph. Schnoebelen. Verification of probabilistic systems with faulty communication. *Inf. & Comp.*, 202(2):141–165, 2005.

[3] P. A. Abdulla and A. Rabinovich. Verification of probabilistic systems with faulty communication. In *Proc. Foundations of Software Science and Computation Structures (FOSSACS 2003), LNCS 2620*, pages 39–53, 2003.

[4] R. Alur, C. Courcoubetis, and D. Dill. Model-checking for probabilistic real-time systems. In *Proc. Automata, Languages, and Programming (ICALP '91), LNCS 510*, pages 115–126, 1991.

[5] C. Baier and B. Engelen. Establishing qualitative properties for probabilistic lossy channel systems: An algorithmic approach. In *Proc. Formal Methods for Real-Time and Probabilistic Systems (ARTS '99), LNCS 1601*, pages 34–52, 1999.

[6] N. Bertrand and Ph. Schnoebelen. Model checking lossy channels systems is probably decidable. In *Proc. Foundations of Software Science and Computation Structures (FOSSACS 2003), LNCS 2620*, pages 120–135, 2003.

[7] D. Brand and P. Zafiropulo. On communicating finite-state machines. *J. ACM*, 30(2):323–342, 1983.

[8] P. Brémaud. *Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues*, volume 31 of *Texts in Applied Mathematics*. Springer, 1999.

[9] D. Bustan, S. Rubin, and M. Vardi. Verifying ω-regular properties of Markov chains. In *Proc. Computer Aided Verification (CAV 2004), LNCS 3114*, pages 189–201, 2004.

[10] C. Courcoubetis and M. Yannakakis. The complexity of probabilistic verification. *J. ACM*, 42(4):857–907, 1995.

[11] J. Esparza and K. Etessami. Verifying probabilistic procedural programs. In *Proc. Found. of Software Technology and Theor. Comp. Sci. (FST&TCS 2004), LNCS 3328*, pages 16–31, 2004.

[12] J. Esparza, A. Kučera, and R. Mayr. Model checking probabilistic pushdown automata. In *Proc. Logic in Computer Science (LICS 2004)*, pages 12–21, 2004.

[13] K. Etessami and M. Yannakakis. Algorithmic verification of recursive probabilistic state machines. In *Proc. Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2005), LNCS 3440*, pages 253–270, 2005.

[14] K. Etessami and M. Yannakakis. Recursive Markov decision processes and recursive stochastic games. In *Proc. Automata, Languages, and Programming (ICALP 2005), LNCS 3580*, pages 891–903, 2005.

[15] H. Hansson and B. Jonsson. A logic for reasoning about time and reliability. *Formal Aspects of Computing*, 6(5):512–535, 1994.

[16] J. G. Kemeny, J. L. Snell, and A. W. Knapp. *Denumerable Markov Chains*. D. Van Nostrand Co., Princeton, NJ, USA, 1966.

[17] M. Z. Kwiatkowska, G. Norman, and D. Parker. PRISM 2.0: A tool for probabilistic model checking. In *Proc. Quantitative Evaluation of Systems (QEST 2004)*, pages 322–323, 2004.

[18] M. Z. Kwiatkowska, G. Norman, R. Segala, and J. Sproston. Automatic verification of real-time systems with discrete probability distributions. *Theor. Comp. Sci.*, 282(1):101–150, 2002.

[19] P. Panangaden. Measure and probability for concurrency theorists. *Theor. Comp. Sci.*, 253(2):287–309, 2001.

[20] A. Rabinovich. Quantitative analysis of probabilistic lossy channel systems. In *Proc. Automata, Languages, and Programming (ICALP 2003), LNCS 2719*, pages 1008–1021, 2003.

[21] Ph. Schnoebelen. The verification of probabilistic lossy channel systems. In *Validation of Stochastic Systems – A Guide to Current Research, LNCS 2925*, pages 445–465. 2004.

[22] M. Vardi. Automatic verification of probabilistic concurrent finite-state programs. In *Proc. Foundations of Computer Science (FOCS '85)*, pages 327–338, 1985.

[23] M. Vardi. Probabilistic linear-time model checking: An overview of the automata-theoretic approach. In *Proc. Formal Methods for Real-Time and Probabilistic Systems (ARTS '99), LNCS 1601*, pages 265–276, 1999.