

Waiting for locks: How long does it usually take? *

Christel Baier¹, Marcus Daum¹, Benjamin Engel², Hermann Härtig², Joachim Klein¹, Sascha Klüppelholz¹, Steffen Märcker¹, Hendrik Tews², Marcus Völpl²

¹Institute for Theoretical Computer Science and ²Operating-Systems Group
Technische Universität Dresden, Germany

Abstract. Reliability of low-level operating-system (OS) code is an indispensable requirement. This includes functional properties from the safety-liveness spectrum, but also quantitative properties stating, e.g., that the average waiting time on locks is sufficiently small or that the energy requirement of a certain system call is below a given threshold with a high probability. This paper reports on our experiences made in a running project where the goal is to apply probabilistic model checking techniques and to align the results of the model checker with measurements to predict quantitative properties of low-level OS code.

1 Introduction

For safety-critical systems such as space, flight, and automotive control systems one wants correctness guarantees for the software not only for the functional behaviour of components but also, e.g., for their timing behaviour. Worst-case execution-time analyses (see e.g. [3, 11, 22]) are able to provide these guarantees, but only in the form of upper bounds on the execution times of all involved components, which hold even in the most extreme situations. Many computer systems are however either not safety critical or they include fail-safe mechanisms that prevent damage in highly exceptional situations. Quantitative analyses can provide detailed information on the probabilities of certain events or on the average behaviour. First, the requirement that certain requirements hold for all possible execution sequences is a very strong condition. E.g., for uncontended locks the property that a process will find a lock free without waiting might not hold universally, but with some high probability. Second, probabilistic features are crucial for the evaluation of complex architectures such as x86 that are optimised according to their average-performance, for systems that rely on imprecise real-time computing techniques to deal with transient overload [18, 6], or for systems that may fail in extreme cases. Third, the probabilistic analysis may guide OS level optimisation justifying, for instance, the use of a simple test-and-test-and-set (TTS) lock implementation over more complex ticket [1] or queue-based locks [15].

* This work was in part funded by the German Research Council (DFG) through the QuaOS project and the collaborative research center 912 Highly-Adaptive Energy-Efficient Computing (HAEC).

In this paper, we report on a running project of the operating-system and the formal-methods group at TU Dresden whose aim it is to establish quantitative properties of low-level operating-system (OS) code using probabilistic model checking techniques. By low-level OS code, we mean drivers, the kernel of monolithic operating systems, microkernels or microhypervisors and similar code that directly interacts with hardware devices and that is therefore often optimised to fully exploit the intrinsic behaviour of modern processor architectures. Although the applicability of probabilistic model checking techniques is expected to work in principle, several non-trivial problems have to be addressed.

Modelling. The first challenge is to find a reasonable abstraction level for the formal model based on the probabilistic analysis that will be carried out. E.g., there are several details on the realisation of hardware primitives (such as caches, busses and controllers of the memory subsystem) that have impact on the timing behaviour of low-level OS code. The model must cover all features that dominate the quantitative behaviour, while still being compact enough for the algorithmic analysis. The latter requires to abstract away from details that have negligible impact on the quantitative behaviour or that would render the model unmanageable. The abstraction of many of these details is indispensable because only little information on the hardware realisation is available, and even if these details are known, too fine-grained hardware models make the state-explosion problem unscalable and lead to quantitative results that are too hardware specific. Instead, we incorporate hardware timing effects in the distributions for the execution times and use measurement-based simulation techniques to obtain empirical evidence for the models and the model checking results.

Measurement-based simulation. Generating this evidence from measurement data is the second major challenge because the quantities of interest for many relevant OS-level properties are in a range where measurements significantly disturb the normal system behaviour and where instrumentation-induced noise blurs the results. For instance, the update rate and resolution of CPU-internal energy sensors necessitate a statistical analysis over a multitude of measurements to extract an energy profile for a single system call [5]. To counteract these effects and to obtain empirical evidence for the models and model checking results, we construct microbenchmarks that place the to-be-measured code into a manageable environment and that mimic the formalisation as close as possible.

Quantitative properties. A third major step is the identification of the types of quantitative properties that are relevant for low-level OS code. At a first glance, it seems that constrained reachability conditions such as “what is the probability for threads to find the requested resource locked for longer than 1 microsecond?” can be expressed as probabilistic queries of the form $\mathbb{P}_{=?}(\varphi)$ using comparably simple patterns of path formulas φ in standard temporal logics, such as PCTL. (The notation $\mathbb{P}_{=?}(\varphi)$ refers to the probability for the event specified by φ .) However, the main interest is in deducing probabilities of this kind for the *long run* rather than for fixed initial distributions. Typically, the long-run behaviour of programs (e.g., the time programs hold a certain resource) shows fundamentally different characteristics when compared to the initialisation-phase

behaviour (when resources are requested for the first time). These differences are caused by the fact that the system had time to learn and adjust to the program characteristics, e.g., by warming up the disk or processor caches or by adjusting the scheduling parameters of the program to meet its responsiveness and interactivity demands. Queries for questions of the above type must therefore be able to ask for long-run probabilities of path formulas *under the condition* that the system is in a certain set of states. The above question translates into a condition about the states in the long run (e.g., the probability of finding a resource held in the long run) and a temporal formula on the paths starting in these states (e.g., “what is the likelihood that such a held resource will be released and granted to the requesting thread within 1 microsecond?”). We refer to these type of queries as *conditional steady-state queries*. A second class of important queries for low-level OS code asks for the value of a quantity that is not exceeded in the majority of all cases: the *quantile*. Two examples of important quantile-based queries are “how long does a thread wait for a resource in 99.9% of all cases?” and “what is the energy that must remain in the battery to guarantee the complete playback of a certain video with a probability greater than 95%?”. To our surprise, neither conditional steady-state queries nor quantile-based queries are supported by state-of-the-art probabilistic model checkers.

Outline. In this paper, we report on our experiences with a simple TTS spinlock as a starting point and initial experiment for a more elaborate investigation of the general feasibility of probabilistic model checking techniques and the limitations of existing tool support. Sec. 2 presents the TTS spinlock and a set of relevant quantitative properties. Sec. 3 explains our discrete-time Markov chain model. Sec. 4 presents the results of the quantitative analysis that we have carried out with the model checker PRISM [13]. We explain how we dealt with conditional long-run probabilities and quantile-based properties and report on the results of the measurement-based simulation and the lessons learned. Sec. 5 concludes this paper. Due to space limitations, we present the detailed model checking statistics in an extended version¹ and provide here only a brief summary.

2 A Test-And-Test-And-Set Lock

Fig. 1 shows the C/C++ code of a simple TTS lock. To acquire the lock, the requesting process executes the atomic swap operation in Line 3 to atomically read the value of the shared variable `occupied` and to then set it to `true`. The loop exits if the process was first to perform this swap after another process has released the lock by setting `occupied` to `false`. For as long as in Line 4 `occupied` is `true`, the process only reads this variable to avoid unnecessary contention on the core-to-core interconnect.

```

1 volatile bool occupied = false;
2 void lock(){
3     while(atomic_swap(occupied,true)){
4         while(occupied){}
5     } }
6 void unlock(){
7     occupied = false
8 }
```

Fig. 1. Simple TTS spinlock

¹ <http://wwwwtcs.inf.tu-dresden.de/ALGI/spinlock-FMICS2012.pdf>

Properties. For our case study, we investigate four questions as representatives for complex conditional long-run and quantile-based properties:

- (A1) probability that a process finds a free lock when it seeks to acquire this lock;
- (A2) probability of re-acquiring a previously held lock (without spinning) without other processes having acquired the lock in the meantime;
- (A3) expected amount of time a process waits for a lock;
- (A4) the 95% quantile of the time processes wait for a lock.

Properties such as (A1)–(A4), which characterise the quantitative behaviour of locks, are highly relevant to guide design decisions and optimisation of low-level OS code. For instance, high probabilities in (A1) and (A2) justify the use of less complex lock implementations, respectively of simpler execution-time analyses. An analysis which assumes low fixed costs for acquiring and releasing a lock is justified by a high probability of (A2) because cache eviction of the *occupied* variable is unlikely. The expected waiting time (A3) is important to judge whether a lock implementation is suitable for the common cases of a given scenario. And the quantile in (A4) replaces the worst-case lock acquisition time in imprecise real-time systems [18, 6] and in systems with a fail safe override in case of late results. It returns an upper bound t for the lock-acquisition time that will be met with the specified probability (here 95%).

We investigate these measures in a scenario where a fixed number of processes repeatedly acquire a single shared lock, execute a critical section while holding the lock, and then wait for some time after they have released the lock before they attempt to re-acquire it. We call the time between release and the attempt to re-acquire the *interim time* and refer to the code that is executed during this time as the *interim section*. We assume here, that the length of the critical section is more or less constant, while the interim time varies. Further, critical sections are typically very short in comparison with the times when no lock is held. For the distribution of the length of the interim section we draw inspiration from a video decoding example. For video decoding, the different frame types (I and P-frames) lead to clusters of the interim time in certain small intervals. Our approach used for both model checking and the measure-based techniques relies on discretisation of these distributions using finitely many sampling points.

The modelled lock-acquisition pattern allows for the derivation of results about the common-case behaviour of applications when they use a certain operating-system functionality, but gives also rise to extreme-case analyses where one assumes malicious or erroneous applications to attack the operating system. For example, by setting the interim time to the execution time of a system call minus its critical sections it is possible to deduce the contention of locks that protect these sections under the assumption that malicious applications invoke this system call as fast as they can.

3 Markov chain model for the spinlock protocol

To model the spinlock protocol, we have chosen a *discrete-time Markov chain* (DTMC) model for the following reasons. The clear demand for probabilis-

tic guarantees about the *long-run* behaviour requires a model where steady-state probabilities are mathematically well defined and supported by model checking tools. This, for instance, rules out probabilistic timed automata and other stochastic models with nondeterminism (e.g., Markov decision processes). Continuous-time Markov chains are not adequate given that the distribution specifying the duration of the critical and interim sections are not exponential. Approximations with phase-type distributions lead to large and unmanageable state spaces. From a mathematical point of view, we could use semi-Markovian models with continuous-time uniform distributions, but we are not aware of tools that provide engines for all queries (A1)–(A4).

Preliminaries: discrete-time Markov chains. We briefly explain our notations used for DTMCs and refer to standard textbooks for further details, see e.g. [12, 7]. A (probabilistic) *distribution* on a countable set X is a function $\mu : X \rightarrow [0, 1]$ such that $\sum_{x \in X} \mu(x) = 1$. The support $\text{supp}(\mu)$ of μ consists of all elements $x \in X$ with $\mu(x) > 0$. μ is called a Dirac distribution if its support is a singleton. If $C \subseteq X$ then $\mu(C) = \sum_{x \in C} \mu(x)$.

In our approach, a DTMC is a tuple $\mathcal{M} = (S, \mathbf{P}, s_{init}, \text{rew})$ where S is a finite state space, $s_{init} \in S$ the initial state, $\mathbf{P} : S \times S \rightarrow [0, 1]$ the transition probability matrix and $\text{rew} : S \rightarrow \mathbb{N}$ the reward function. We require $\sum_{u \in S} \mathbf{P}(s, u) = 1$ for all $s \in S$ and refer to $\mathbf{P}(s, u)$ as the probability to move from s to u within one (time) step. A path in \mathcal{M} is a finite or infinite sequence $\pi = s_0 s_1 s_2 \dots$ of states with $\mathbf{P}(s_i, s_{i+1}) > 0$ for all i . Let $\pi(k) = s_k$ be the $(k+1)$ -st state and $\pi \downarrow k = s_0 s_1 \dots s_k$ the prefix consisting of the first $k+1$ states. $\text{Paths}(s)$ denotes the set of infinite paths starting in state s . The accumulated reward for a finite path $\pi = s_0 s_1 \dots s_k$ is $\text{Rew}(\pi) = \text{rew}(s_0) + \dots + \text{rew}(s_{k-1})$. If π is a finite path then the cylinder set $\text{Cyl}(\pi)$ spanned by π consists of all infinite paths π' where π is a prefix of π' . Using well-known concepts of measure theory, given some probabilistic distribution $\mu : S \rightarrow [0, 1]$, there exists a unique *probability measure* $\text{Pr}_\mu^{\mathcal{M}}$ on the σ -algebra generated by the cylinder sets of finite paths such that $\text{Pr}_\mu^{\mathcal{M}}(\text{Cyl}(s_0 s_1 \dots s_k)) = \mu(s_0) \cdot \prod_{0 \leq i < k} \mathbf{P}(s_i, s_{i+1})$. If μ is a Dirac distribution with $\text{supp}(\mu) = \{s\}$ we simply write Pr_s or $\text{Pr}_s^{\mathcal{M}}$ for $\text{Pr}_\mu^{\mathcal{M}}$.

For specifying measurable path events (i.e., sets of paths that belong to the σ -algebra generated by the cylinder sets of finite paths), we use the standard LTL-like notations with the symbols \bigcirc (next), \mathcal{U} (until) and \diamond (eventually) and time-bounded variants thereof. For $X, Y \subseteq S$, $\bigcirc X$ stands for the set of infinite paths π with $\pi(1) \in X$. $X\mathcal{U}^=k Y$ denotes the set of infinite paths π such that $\pi(n) \in X \setminus Y$ for $0 \leq n < k$ and $\pi(k) \in Y$. We write $X\mathcal{U}^{\leq K} Y$ for the union of the sets $X\mathcal{U}^=k Y$ where k ranges over the elements in $\{0, 1, \dots, K\}$ and $X\mathcal{U} Y$ for the union of the sets $X\mathcal{U}^=k Y$ where $k \in \mathbb{N}$. $\diamond Y$, $\diamond^=k Y$ and $\diamond^{\leq K} Y$ are short forms of $S\mathcal{U} Y$, $S\mathcal{U}^=k Y$ and $S\mathcal{U}^{\leq K} Y$, respectively. To deal with query (A2), we will also use LTL-like formulas with cascades of until-operators and suppose here the standard LTL-semantics for paths. For further details see [20, 4, 21].

For $m \in \mathbb{N}$, $\theta_m : S \rightarrow [0, 1]$ denotes the state distribution for \mathcal{M} after m steps. Formally, θ_0 is the Dirac distribution with $\text{supp}(\theta_0) = \{s_{init}\}$ and $\theta_{m+1} = \mathbf{P} \cdot \theta_m$ for $m \geq 0$. The function $\theta : S \rightarrow [0, 1]$,

$$\theta(s) = \lim_{k \rightarrow \infty} \frac{1}{k+1} \cdot \sum_{m=0}^k \theta_m(s),$$

is called the *steady-state distribution* for \mathcal{M} . Then, $\theta(s) > 0$ iff s belongs to a bottom strongly connected component (BSCC) that is accessible from s_{init} . If C is a set of states with $\theta(C) > 0$ and Π a measurable set of paths then the *conditional long-run probability* for Π (under condition C) is defined by:

$$\mathbb{P}^{\mathcal{M}}(\Pi | C) \stackrel{\text{def}}{=} \sum_{s \in C} \theta(s)/\theta(C) \cdot \text{Pr}_s^{\mathcal{M}}(\Pi)$$

Here, $\theta(s)/\theta(C)$ is the conditional steady-state probability for state s , again under condition C . The intuitive meaning of $\theta(s)/\theta(C)$ is the portion of time spent in state s on long runs relative to the total time spent in states of C . With the factor $\text{Pr}_s^{\mathcal{M}}(\Pi)$, the above weighted sum represents the long-run probability for the event specified by Π under condition C . Analogously, conditional steady-state average values of random variables can be defined as weighted sums. (A3) will be formalised as an instance of *conditional long-run accumulated reward* for reaching a goal set Y defined by:

$$\mathbb{R}^{\mathcal{M}}(\diamond Y | C) \stackrel{\text{def}}{=} \sum_{s \in C} \theta(s)/\theta(C) \cdot \text{ExpAccRew}_s^{\mathcal{M}}(\diamond Y)$$

where $\text{ExpAccRew}_s^{\mathcal{M}}(\diamond Y)$ denotes the expected accumulated reward for reaching Y from state s . It is defined by:

$$\sum_{r=0}^{\infty} r \cdot \text{Pr}_s^{\mathcal{M}}\{\pi \in \text{Paths} : \exists k \in \mathbb{N} \text{ s.t. } \pi \in \diamond^{=k} Y \wedge \text{Rew}(\pi \downarrow k) = r\}$$

Markov chain model for the spinlock protocol. To analyse the quantitative behaviour of the spinlock protocol for n processes P_1, \dots, P_n , we use a DTMC that results as the synchronous parallel composition of one module representing the spinlock in Fig. 1 (see Fig. 3) and one module for each of the processes (see Fig. 2). The t_i 's are integer variables that serve as timers for the critical and noncritical section. Distribution ν models the *interim time*, i.e., the time required for the (noncritical) activities of the processes between two critical sections, including the request to acquire the lock, but without the spinning time. Distributions γ_0 and γ_1 serve to model the total length of the critical

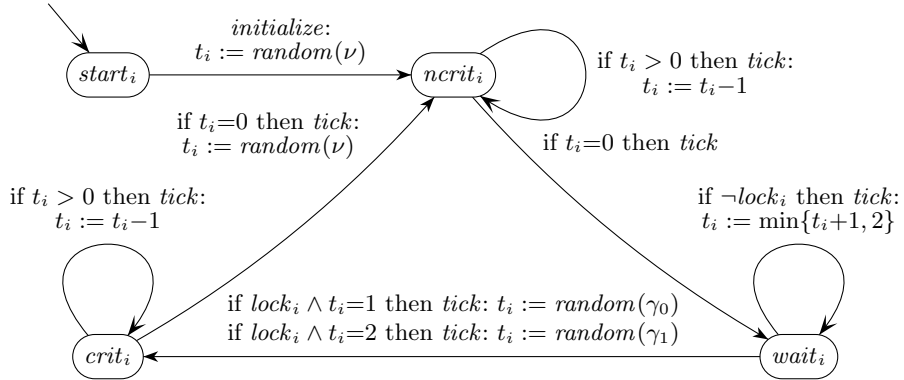


Fig. 2. Control flow graph of process P_i




Fig. 3. Control flow graph of the spinlock

section (including lock acquisition and release). In order to account for cache effects that lead to different time behaviour for the lock acquisition, depending on whether the lock is taken with and without spinning, we use two distributions for the critical section length. Distribution γ_0 is used when the lock was obtained without spinning, while distribution γ_1 is used if some spinning was necessary. To obtain a DTMC model, ν and γ_0, γ_1 are discrete distributions with finitely many sampling points in the relevant time intervals. E.g., the distribution $\nu(8) = \nu(10) = \nu(12) = \frac{1}{3}$ indicates that with equal probability the duration of the interim time is 8, 10 or 12 time units. The assignment $t_i := \text{random}(\nu)$ means that a sample is drawn according to distribution ν and assigned to t_i .

For each process P_i , we distinguish four locations. Location $start_i$ serves to set the timers for the first noncritical phase. The other three locations have the obvious meaning. Location $wait_i$ signals that P_i is trying to acquire the lock. If the lock is granted to P_i , the lock process switches to location $lock_i$, which in turn enables the transition from $wait_i$ to $crit_i$. In location $wait_i$, process P_i waits until the lock has been made accessible for it. In location $wait_i$, variable t_i does not serve as a timer. Instead, t_i indicates whether process P_i has just entered the waiting location ($t_i \in \{0, 1\}$) or P_i is spinning as some other process is holding the lock ($t_i=2$). The control flow graph of the lock process contains for each P_i one location $lock_i$ (indicating that P_i may take or holds the lock) and one location $unlock$ (the lock is free). For the synchronisation, we followed the approach of PRISM's input language with synchronisation over shared actions. Action *initialize* has to be synchronised by all processes, while *tick* indicates one time step and must be executed synchronously by all processes and the lock.

The states of the DTMC \mathcal{M} for the composite model have the form $s = \langle \ell_1, \dots, \ell_n, m, t_1=b_1, \dots, t_n=b_n \rangle$ where ℓ_i is the current location of process P_i , m the current location of the lock and b_i the current value of variable t_i . Then, P_i is spinning in state s iff $s.m \neq lock_i$, $s.\ell_i = wait_i$ and $s.b_i = 2$, where we use the dot notation to refer to parts of a state s . If process P_i performs its last critical action and moves from location $crit_i$ to $ncrit_i$ then either the lock returns to

its initial location *unlock* (if no process P_k is spinning) or there is a uniform probabilistic choice for the lock to move to one of the locations $lock_k$ where process P_k is spinning. To compute the long-run average spinning time (query (A3)), we deal with the reward function $rew_spin_i(s) = 1$ for each state s where process P_i is spinning. For all other states s , we have $rew_spin_i(s) = 0$.

Formalisation of queries (A1)–(A4). In the sequel, we use propositional formulas over the locations and conditions on the values of t_1, \dots, t_n to characterise sets of states. For instance, $crit_i$ is identified with the set $\{s \in S : s.l_i = crit_i\}$, where S denotes the state space of the DTMC \mathcal{M} for the composite system. Condition $request_i = wait_i \wedge t_i=0$ characterises the set of states s in \mathcal{M} where process P_i has just requested the lock. The set of states where process P_i spins is specified by $spin_i = wait_i \wedge t_i=2 \wedge \neg lock_i$. Finally, $release_i = crit_i \wedge t_i=0$ characterises the states s where process P_i is just performing its last critical actions and the lock is to be released next. The relevant quantitative measures (A1)–(A4) of Section 2 now correspond to the following values.

- (A1) $\mathbb{P}^{\mathcal{M}}(\varphi_1 \mid request_i)$ where $\varphi_1 = \bigcirc lock_i$
- (A2) $\mathbb{P}^{\mathcal{M}}(\varphi_2 \mid release_i)$ where $\varphi_2 = \bigcirc (unlock \mathcal{U} lock_i)$
- (A3) $\mathbb{R}^{\mathcal{M}}(\diamond lock_i \mid request_i)$
- (A4) $\min\{t \in \mathbb{N} : \mathbb{P}^{\mathcal{M}}(\diamond^{\leq t+1} lock_i \mid request_i) \geq 0.95\}$

(A1), (A3) and (A4) refer to the conditional steady-state distribution for the condition that process P_i has just performed its first request operation. (A1) corresponds to the long-run probability under the condition $request_i$ for the path event $\varphi_1 = \bigcirc lock_i$ stating that in the next time step process P_i will win the race between the waiting processes, i.e., it will enter its critical section without spinning. (A3) stands for the long-run average spinning time from states where P_i has just requested its critical section. Replacing the condition $request_i$ in (A3) with $wait_i \wedge t_1=1 \wedge \neg lock_i$, we obtain the long-run average spinning time, provided that the first attempt to acquire the lock was not successful. The quantile in (A4) corresponds to the minimal number t of time steps minus one such that the long-run probability from the $request_i$ -states for the path event $\diamond^{\leq t+1} lock_i$ stating that the lock will be granted for process P_i in $t+1$ or fewer steps is at least 0.95. For the long-run probability of acquiring the lock, again without interference by other processes (see (A2)), there are several reasonable formalisations. For the constraint that no process other than P_i requested the lock in the mean time, we can deal with the path event $\varphi_2 = \bigcirc (unlock \mathcal{U} lock_i)$ under the condition that process P_i will release the lock in the next step (condition $release_i$). The variant of (A2) where other processes may have acquired and released the lock in between the critical sections of P_i and where P_i has not to spin and hence experiences a low overhead lock can be formalised as $\mathbb{P}^{\mathcal{M}}(\varphi'_2 \mid request_i)$ for the event specified by the LTL-formula $\varphi'_2 = \bigcirc (lock_i \mathcal{U} (\neg crit_i \wedge (\neg spin_i \mathcal{U} crit_i)))$. The treatment of (A2) using φ_2 or φ'_2 is rather complex since the standard treatment of LTL-queries relies on a probabilistic reachability analysis of a product construction of a deterministic ω -automaton and the DTMC (see e.g. [2]). To avoid

this automaton-based approach, one can replace φ_2 and φ'_2 with a simpler reachability condition when refining the control flow graph of P_i by duplicating the control loop $ncrit_i \text{ wait}_i \text{ crit}_i$. This can be realised by introducing a Boolean variable b that flips its value after leaving the critical section, see the extended version of this paper, and justified using an appropriate notion of bisimulation. The control flow graphs for the lock and the other processes remain unchanged. Then, instead of φ_2 and φ'_2 we can then deal with $\psi_2 = (\text{lock}_i \vee \text{unlock}) \mathcal{U}(\text{crit}_i \wedge b)$ and $\psi'_2 = \neg \text{spin}_i \mathcal{U}(\text{crit}_i \wedge b)$ under the condition $\text{release}_i \wedge \neg b$ and $\text{request}_i \wedge \neg b$, respectively. Indeed, our experiments show that the analysis of \mathcal{M}' with the modified queries is more efficient than the analysis of \mathcal{M} .

4 Quantitative analysis of the TTS spinlock

Our general approach to the quantitative analysis of low-level operating-system code proceeds in four steps: (1) The targeted operating-system code is formalised at a suitable level of abstraction in the input language of a probabilistic model checker. For our TTS lock case study, we used the prominent model checker PRISM [13]. (2) The parameters of the model are determined with the help of measurements in the targeted setting. If this is not possible due to unduly high interference with the instrumentation code, we extract the relevant code and measure it in the form of a microbenchmark in a controlled environment. (3) The queries of interest are evaluated both by the model checker and with a microbenchmark that again executes the code in a manageable environment. This step is necessary to determine how well the model corresponds to the targeted setting. (4) The queries are evaluated on the parameters obtained for the real workload and if possible cross-checked against measurements in this setting. There are two situations in which such a comparison is infeasible, namely when the interference between the instrumentation code and the targeted operating-system code fundamentally changes the behaviour of the latter or when the analysis is performed with parameters of a system that does not yet exist.

The measurements done in the microbenchmarks do as well interfere with the to-be-analysed code. However, these measurements only extract the required model parameters (possibly only one at a time). The challenge is to construct a setting where this interference is limited only to those parameters that are not currently extracted. We now report on our experiences in adjusting the model with the help of a microbenchmark and measurement-based simulation.

Measurement-based simulation. For the extraction of the two distributions γ_0 and γ_1 for the critical section and for the distribution ν for the interim section, we can resort to random sampling and similar techniques [14] to deduce the characteristics of the workload we seek to investigate. Very short critical sections and the acquire and release costs of the lock must however be measured in an environment where it is possible to control side effects on the quantity of interest. For the TTS lock, we construct such an environment by mimicking the behaviour of the formal model in the critical and interim section. When a process enters a critical section, it selects pseudorandomly a sampling point of the distribution γ_0 if the lock was free and of γ_1 if it had to spin. A counter in the

spinning loop (in Line 4 in Fig. 1) reveals whether or not the process was able to obtain the lock without spinning. We pass this counter in a processor register not to disturb the timing of the TTS code. The actual instrumentation consists of three reads of the per core time stamp counter, before and after acquiring the invocation of the `lock()` function in Line 2 and after the `unlock()` function returns. Both the critical and interim section consist of a loop of an integer instruction: `rep; dec %eax`. We have confirmed that the execution time of this loop is very regular. The loop executes for a time that is proportional to the sampling point selected for the respective critical or interim section.

Quantitative analysis using PRISM. For the quantitative analysis of the DTMC of the TTS spinlock we used the probabilistic model checker PRISM [13]. We mainly concentrated on (A1)–(A4), but also considered functional and a few more quantitative queries. To obtain empirical evidence in the model and in the model checking results we compare the model checking results with measurements of the model mimicking a microbenchmark. Unfortunately, PRISM has no direct support for computing conditional long-run probabilities or quantile-based queries. We extended the PRISM code by operators that compute conditional long-run probability $\mathbb{P}^{\mathcal{M}}(\varphi \mid C)$ and conditional long-run accumulated rewards $\mathbb{R}^{\mathcal{M}}(\diamond Y \mid C)$ where φ is a PCTL path formula and Y, C sets of states. Although there is also no direct support for (A4) in PRISM, quantiles that refer to the amount of time until some event occurs can be calculated with the same iterative bottom-up computation scheme as for bounded reachability properties. In (A4), we are interested in the minimal $t \in \mathbb{N}$ such that the conditional long-run probability is greater than some fixed probability value. For finding the minimal value $t \in \mathbb{N}$ with the above property efficiently, we modified the implementation of the bounded until operator in PRISM to store the intermediate probabilities for all $0 \leq j \leq t$. Instead of a more direct evaluation at the MTBDD level of PRISM, this storing of intermediate results allows for an external script to check the gradually increasing values of the bounded until formula without restarting the model checking for each such check.

Lessons learned. During the analysis of the queries (A1)–(A4) and in the course of performing and evaluating the measurements for the simulation, we encountered several difficulties that we would like to share.

Cascade effects. We first evaluated (A1)–(A4) using a model where the length of the critical section and the interim time are deterministic, i.e., where $\gamma_0 = \gamma_1$ and ν are Dirac distributions with values T_{crit} and T_{int} and where $T_{crit} \ll T_{int}$. A simulation run of the model revealed a probabilistic choice of the order in which processes acquire the lock for the first time. For all subsequent turns, the processes received the lock in this same order and without having to spin. Small variations in the lock acquisition times and in the points in time when the processes start prevented the measurement-based simulation from entering a similar cascade. These cascade effects, however, do not appear for more realistic models where at least ν is a distribution with $|supp(\nu)| \geq 2$. Assuming that the values for ν are much larger than those for γ_0 and γ_1 , such DTMC models only

have one bottom strongly component, which justifies to apply the measurements for just a few simulation runs.

Short critical sections. One of the first workloads we evaluated, was a system call of the Nova microhypervisor [19] in which very short TTS-lock protected critical sections alternate with relatively long sequences of interim activities. The measure-based approach encountered situations where the lock-acquisition times exceeded the critical section length. To avoid a too fine granular (discrete) time domain for the DTMC model that would rule out the feasibility of model checking techniques, we used time domains of different granularity for the measurements and the DTMC model and relate them via a scaling factor sf . E.g., for the scaling factor $sf = 1000$, one time unit in the DTMC model corresponds to 1000 cycles (approx. 362 ns) on the target system.

Varying acquisition times. In an early version of the microbenchmark, we adjusted the `rep; dec %eax` loops, which together with the pseudo random choice of a sampling point mimic the distributions of the critical and interim sections, to the same constant value for all processes. More precisely, for $\gamma(x) = 1$, we adjusted the loop in critical to consume as close as possible to 1000 cycles minus the time required to execute the instrumentation code. However, the average costs for acquiring and releasing a lock increase with the number of processes that require this lock. One explanation for this behaviour could be that the costs for invalidating a cache line when acquiring or releasing the lock vary with varying number of cores. This is because the on-chip networks in modern multicore processors connect cores in a point-to-point fashion while maintaining information about the locations of copies of cache lines.

Controlled measurement environment. We realised early variants of our microbenchmark as a Linux user-level application, disabled all obvious sources of interference and raised the priority of this application into the otherwise empty real-time priority band. From the results, we classified spikes as interference and ignored these points in our comparison with the formal model. Still, we experienced high fluctuations of the measurement results, which could not easily be explained. In a repeated measurement on top of a small kernel binary, which we just used to bootstrap our microbenchmark and to communicate the results after the measurement part completed, these variations did not reappear. We therefore take this effect as an indicator and as a warning that the interference of large operating-system kernels on short executing microbenchmarks should not be underestimated and best be avoided whenever this is possible.

Need for two critical distributions. After having performed the above adjustments, we observed a discrepancy between the model checking results and the measurements of approx. 20% (see the model I results in Fig. 4(a)–4(c)). Following an in-depth search for possible causes both in the model and in the microbenchmark, we identified small variations between the costs of acquiring a lock with and without spinning. In the course of this search, we encountered several other possible causes such as the quantisation due to scaling, which showed similar small variations in the measurements. To confirm these factors, we adjusted a copy of the measurement data to mimic the effect we suspected to

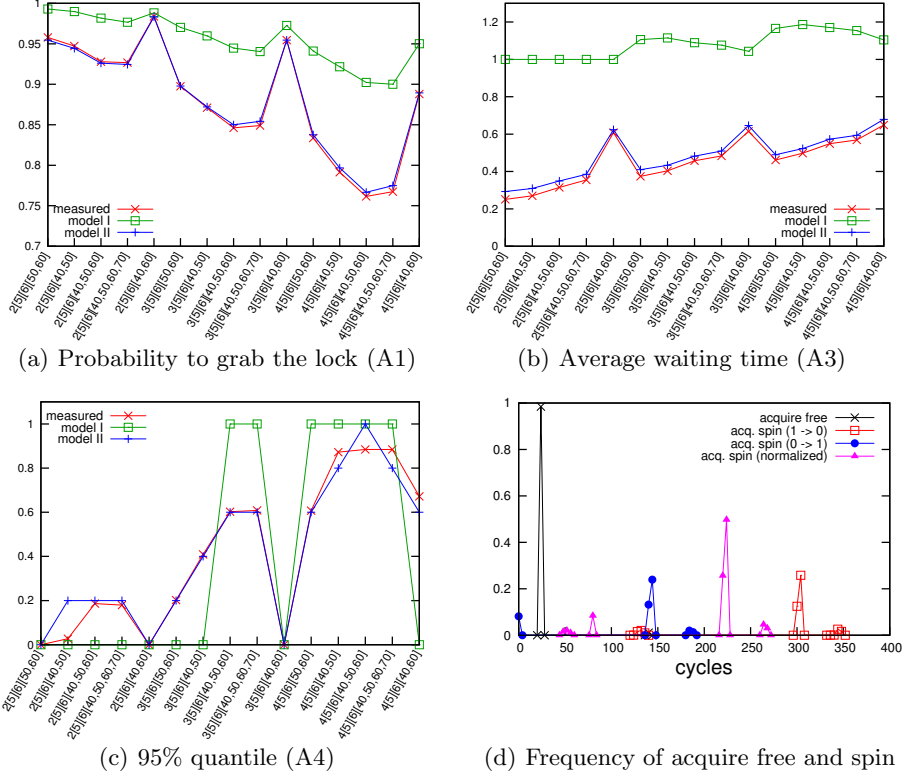


Fig. 4. (a)-(c) results for model I and II and (d) the histogram for acquire time

determine the magnitude of the impact this effect could have. If this impact was in a range where it could have explained a significant part of this discrepancy, we adjusted the formal model accordingly. For the small variations between lock acquisition times, we changed the model from a single Dirac distribution for the critical section length to the two Dirac distributions γ_0 and γ_1 in Fig. 2. Unfortunately, the value of the singleton sampling point of γ_1 could not be directly measured because it would require the inclusion of `rdtsc` in the spinning loop (Line 4 in Fig. 1) to read the core cycle counter, which would significantly change the timing behaviour of the lock. We therefore calculate the average costs for acquiring a lock under the condition that a process had to spin by comparing the time stamps of the releasing cores with the time stamps of the lock acquiring cores. However, although the time stamp counter for cores on the same die are derived from the same clock, there is an offset caused by the barrier on which these cores synchronise to start the measurement at approx. the same point in time. The two dashed lines in Fig. 4(d) show the acquire spin costs per core. The figure also shows the acquire spin costs after normalising the offset between the two clocks. We choose the spike as the sampling point for γ_1 . The scaling factor $sf = 1000$ that we used to deal with the single-distribution model (where $\gamma(1) = 1$) was no longer adequate for integrating the resulting

distributions ($\gamma_0(5)=\gamma_1(6)=1$) into the DTMC model. We therefore decreased the scaling factor to $sf=200$.

Evaluation of the results. Fig. 4(a)–4(c) show the results of our quantitative analysis of the queries (A1)–(A4). We omit the plot for (A2) because, for selected distributions, the chance to re-acquire the lock without interference by other processes is close to zero. The plots compare our measurements with the results from two models. Model I is our earlier, simpler model which uses only one distribution for the critical section ($\gamma_0=\gamma_1$). Model II is more precise because it uses different distributions γ_0 and γ_1 as explained before. The x-axis displays the number of processes and the distributions used for model II and, with scaling factor $sf=200$, for the measurement. The label 3[5][6][40, 50, 60] stands for $n=3$ processes and the distributions $\gamma_0(5)=\gamma_1(6)=1$ and $\nu(40)=\nu(50)=\nu(60)=\frac{1}{3}$. The distribution for model I, which uses scaling factor $sf=1000$, is obtained by dividing all values by 5 and using the value of γ_0 for γ_1 too. Thus, for model I, the same label stands for the model with 3 processes, $\gamma_0(1)=\gamma_1(1)=1$ and $\nu(8)=\nu(10)=\nu(12)=\frac{1}{3}$. We performed the measurements on an Intel i7 920 quadcore machine at 2.67 GHz. The benchmark and the sampling area for storing the measurement results fitted completely in the on-die caches. Fig. 4(a) shows that the introduction of γ_1 reduced the error between the measured and analysed results from 20% for model I to below 1% for model II. The discrepancy between model II and the measurements in Fig. 4(c) is due to the quantisation of the model. That is, the model considers changes of the waiting time only in steps of $sf=200$ cycles, which corresponds to one fifth (0.2) of the critical section length. The differences between the results obtained for the models with parameters $n[5][6][40, 50, 60]$ and $n[5][6][40, 60]$ (where $n \in \{2, 3, 4\}$) illustrate that not only the mean value, but also the variance of distribution ν for the interim time has non-negligible impact on (A1)–(A4).

For model I ($sf=1000$ and $\gamma_0=\gamma_1$), we used PRISM to compute (A1)–(A4) for the DTMC with up to six processes. For model II ($sf=200$, $\gamma_0 \neq \gamma_1$) the analysis has been carried out with up to four processes. Because of the reduced scaling factor, model II is far more complex. E.g., for $n=4$ processes and the distribution $\nu(8) = \nu(14) = \frac{1}{2}$ the DTMC for model I has about 10^4 states, while for model II with the corresponding distribution $\nu(40) = \nu(70) = \frac{1}{2}$ the DTMC has about $2.5 \cdot 10^6$ states. In a nutshell, for the DTMC with $n=6$ and $sf=1000$ PRISM needs a few minutes for all queries, while for $n=4$ and $sf=200$ the computation can take a few hours. In most cases, the computation of the steady-state probabilities is the most time consuming part. For more information on the PRISM statistics (MTBDD sizes, time for the model construction and the quantitative analysis) we refer to the extended version.

5 Conclusions

The paper presents a first step towards the application of probabilistic model checking techniques for the quantitative analysis of low-level operating-system code. We reported on the difficulties we encountered when analysing a simple test-and-test-and-set (TTS) spinlock with the model checker PRISM and on our

solutions to address them. A major challenge was to find an appropriate level of abstraction that allows to capture all relevant behaviour and allows to abstract from the precise timing behaviour of the cache and other CPU parts. We performed extensive measurement-based simulations of real spinlocks to demonstrate that our abstract model does indeed reproduce important aspects of the timing behaviour. We considered a representative list of properties that are of high interest to the system designer when he has to choose the right lock implementation. These properties involve conditional steady-state probabilities and quantiles. To overcome the lack of direct tool support for both types of queries, we added the relevant features in the PRISM code.

Related work. Many researchers performed case studies with probabilistic model checkers for mutual exclusion protocols and other coordination algorithms for distributed systems (see e.g. [8, 16, 17] or the PRISM web pages [13]). While some of these case studies address the analysis of randomised protocols, we deal with non-randomised operating-system primitives (the TTS spinlock). Our models rely on stochastic assumptions on the execution times of the critical section and interim sections of competing processes). Unlike the wide range of case studies with continuous-time models, where rates of exponential distributions specify, e.g., the frequency of the arrival of requests or the average duration of events (see e.g. [7]), we deal with a DTMC model and discretisations of non-exponential distributions. Of course, there have been plenty of case studies with the DTMC engine of PRISM, but we are not aware of experiments that have been carried out where the modelling and evaluation process was accompanied by measure-based techniques. To the best of our knowledge, none of these case studies considers conditional steady-state probabilities or quantile-based queries. The majority of work on model checking low-level operating-system code concentrates on properties in the safety-liveness domain. E.g., [23] used model checking to find serious file system bugs. Formal quantitative analyses often consider only worst-case execution times as measure. For the special case of probabilistic worst-case execution times (i.e. queries similar to Query (A3)), [3] presents a timing-schema based on independent or only pairwise dependent (i.e., joint) execution profiles.

Future work. It would be very interesting to scale the analysis of basic and more advanced spinlocks for CPUs with more than 100 cores. Such results could justify the use of more simple locks with less overhead for certain tasks. Given the exponential growth of the system model with the number of cores, this is an extremely challenging task and requires clever encoding, abstraction and reduction techniques such as symmetry reduction. Other promising candidates are bisimulation quotienting techniques as supported by the model checker MRMC [9, 10] and the use of sophisticated (MT)BDD-based techniques to increase the efficiency of PRISM's symbolic engine.

References

1. T. E. Anderson. The performance of spin lock alternatives for shared-memory multiprocessors. *IEEE Trans. Parallel Distrib. Syst.*, 1(1):6–16, January 1990.
2. C. Baier and J.-P. Katoen. *Principles of Model Checking*. MIT Press, 2008.

3. G. Bernat, A. Colin, and S. Petters. WCET analysis of probabilistic hard real-time systems. In *RTSS'02*, pages 279–288. IEEE, 2002.
4. C. Courcoubetis and M. Yannakakis. The complexity of probabilistic verification. *Journal of the ACM*, 42(4):857–907, 1995.
5. M. Hähnel. Energy-utility functions. Diploma thesis, TU Dresden, Germany, 2012.
6. C.-J. Hamann, J. Löser, L. Reuther, S. Schönberg, J. Wolter, and H. Härtig. Quality-assuring scheduling - using stochastic behavior to improve resource utilization. In *RTSS'01*, pages 119–128. IEEE, 2001.
7. B. Haverkort. *Performance of Computer Communication Systems: A Model-Based Approach*. Wiley, 1998.
8. S. Irani, G. Singh, S. K. Shukla, and R. Gupta. An overview of the competitive and adversarial approaches to designing dynamic power management strategies. *IEEE Trans. VLSI Syst.*, 13(12):1349–1361, 2005.
9. J.-P. Katoen, T. Kemna, I. Zapreev, and D. Jansen. Bisimulation minimisation mostly speeds up probabilistic model checking. In *TACAS'07*, volume 4424 of *LNCS*, pages 87–101, 2007.
10. J.-P. Katoen, I. Zapreev, E. Hahn, H. Hermanns, and D. Jansen. The ins and outs of the probabilistic model checker MRMC. *Perform. Eval.*, 68(2):90–104, 2011.
11. S. Knapp and W. Paul. Realistic worst-case execution time analysis in the context of pervasive system verification. In *Program Analysis and Compilation*, volume 4444 of *LNCS*, pages 53–81, 2007.
12. V. Kulkarni. *Modeling and Analysis of Stochastic Systems*. Chapman & Hall, 1995.
13. M. Kwiatkowska, G. Norman, and D. Parker. Probabilistic symbolic model checking with PRISM: A hybrid approach. *STTT*, 6(2):128–142, 2004.
14. J. Liedtke, N. Islam, T. Jaeger, V. Panteleenko, and Y. Park. Irreproducible benchmarks might be sometimes helpful. In *ACM SIGOPS European Workshop*, pages 242–246. ACM, 1998.
15. J. Mellor-Crummey and M. Scott. Scalable reader-writer synchronization for shared-memory multiprocessors. In *PPOPP'91*, pages 106–113. ACM, April 1991.
16. G. Norman. Analysing randomized distributed algorithms. In *Validation of Stochastic Systems: A Guide to Current Research*, volume 2925 of *LNCS*, pages 384–418, 2004.
17. G. Norman, D. Parker, M. Kwiatkowska, S. Shukla, and R. Gupta. Using probabilistic model checking for dynamic power management. *Formal Aspects of Computing*, 17(2):160–176, 2005.
18. W. K. Shih, J. W.-S. Liu, and J.-Y. Chung. Algorithms for scheduling imprecise computations with timing constraints. *SIAM J. Comput.*, 20(3):537–552, 1991.
19. U. Steinberg and B. Kauer. NOVA: a microhypervisor-based secure virtualization architecture. In *EuroSys'10*, pages 209–222. ACM, 2010.
20. M. Vardi. Automatic verification of probabilistic concurrent finite-state programs. In *FOCS'85*, pages 327–338. IEEE, 1985.
21. M. Vardi. Probabilistic linear-time model checking: An overview of the automata-theoretic approach. In *ARTS'99*, volume 1601 of *LNCS*, pages 265–276, 1999.
22. R. Wilhelm, J. Engblom, A. Ermedahl, N. Holsti, S. Thesing, D. Whalley, G. Bernat, C. Ferdinand, R. Heckmann, T. Mitra, F. Mueller, I. Puaut, P. Puschner, J. Staschulat, and P. Stenström. The worst-case execution-time problem - overview of methods and survey of tools. *Trans. Embedded Comput. Syst.*, 7(3):1–53, 2008.
23. J. Yang, P. Twohey, D. Engler, and M. Musuvathi. Using model checking to find serious file system errors. *ACM Trans. Comput. Syst.*, 24(4):393–423, 2006.

A Appendix

A.1 Simplifications of model I for scaling factor 1000

In Section 4, we described model I to use only one distribution $\gamma \stackrel{\text{def}}{=} \gamma_0 = \gamma_1$ for the length of the critical section. In addition to the larger scaling factor for model I (i.e., $sf=1000$ in contrast to $sf=200$ for model II), there are several possibilities for simplifying the control flow graph.

The following simplifications have been used for the evaluation of model I. In the following, we consider $i = 1$ and analyze the quantitative behavior of process P_1 against queries (A1)–(A4).

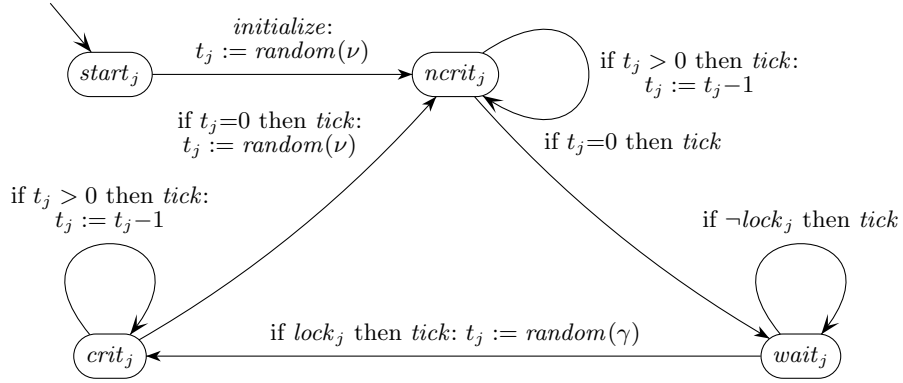


Fig. 5. Simplified control flow graph of P_j , $2 \leq j \leq n$, for Model I

For all considered queries (A1), (A2), (A2'), (A3), (A3') and (A4), the control flow of all other processes P_j ($2 \leq j \leq n$) can be simplified by dropping the assignments for t_j in location $wait_j$. The simplified control flow graphs for P_2, \dots, P_n are shown in Fig. 5.

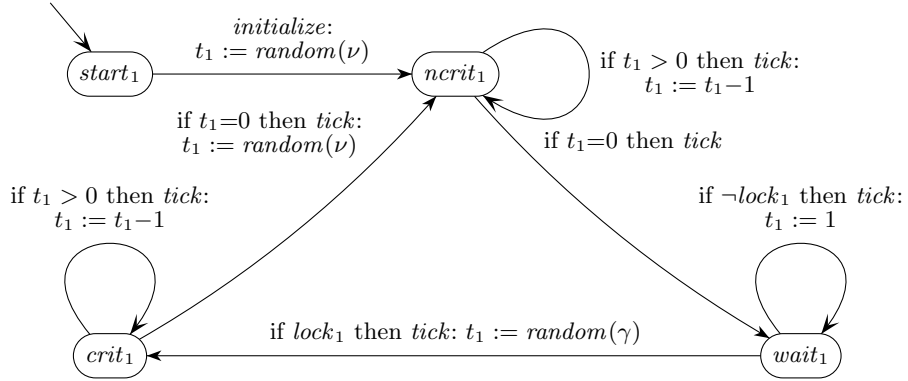


Fig. 6. Simplified control flow graph of process P_1 for model I and (A1), (A3), (A4)

Let us now turn to the simplifications of the control graph for process P_1 in Model I. Queries (A1), (A3) and (A4) reason about the waiting times for the chosen process P_1 . Thus, we cannot use the same simplified control flow graphs as for P_2, \dots, P_n . Nevertheless, in (A1), (A3) and (A4) there is no need to distinguish between states in which $wait_1 \wedge t_1=1$ (i.e., P_1 just requested the lock) and $wait_1 \wedge t_1=2$ (i.e., P_1 is spinning). This leads to the simplified control flow graph for process P_1 as shown in Fig. 6.

In Section 3, we presented a variant of the conditional average waiting time that differs from (A3) in the condition.

$$(A3') \quad \mathbb{R}^{\mathcal{M}}(\diamond lock_1 \mid wait_1 \wedge t_1=1 \wedge \neg lock_1)$$

For (A3'), it is crucial to distinguish those states where $wait_1 \wedge t_1=1$ holds and those states where $wait_1 \wedge t_1=2$ holds. Thus, the simplification of Fig. 6 is not adequate for dealing with (A3'). To evaluate (A3') in model I, we used the control flow graph of P_1 as shown in Fig. 7.

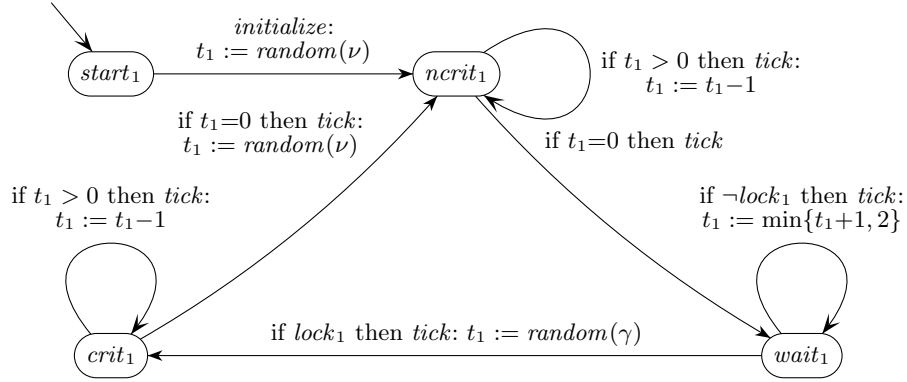


Fig. 7. Simplified control flow graph of process P_1 for model I and (A3')

A.2 Variant of model I and II for the treatment of (A2)

For the treatment of (A2), we mentioned in Section 3 the possibility to deal with an until-property (ψ_2) instead of the LTL-formula φ_2 , when duplicating the control loop $ncrit_1 wait_1 crit_1$ for process P_1 (where we again assume $i = 1$).

Let \mathcal{M} be the DTMC obtained by the control graphs shown in Section 3 (see Fig. 2) for the two distribution variant of model II (i.e., $\gamma_0 \neq \gamma_1$ and scaling factor $sf = 200$). For model I (i.e., $\gamma = \gamma_0 \neq \gamma_1$ and $sf = 1000$) the DTMC \mathcal{M} is obtained by using

- the control flow graphs shown in Fig. 5 for processes P_2, \dots, P_n , and
- the control flow graph shown in Section 3 for process P_1 (see Fig. 2).

In both cases, the control graph for the spinlock is as described in Section 3 (see Fig. 3).

To avoid the costly LTL-queries for the variants (A2) and (A2'), the control flow graph for the process P_1 can be extended with a Boolean variable b , which unrolls the control loop $ncrit_1 wait_1 crit_1$. Let \mathcal{M}' denote the DTMC for this extended control graph and for process P_1 as shown in Fig. 8 (for model I) and in Fig. 9 (for model II). For all other processes and for the lock, the control graphs remain unchanged.

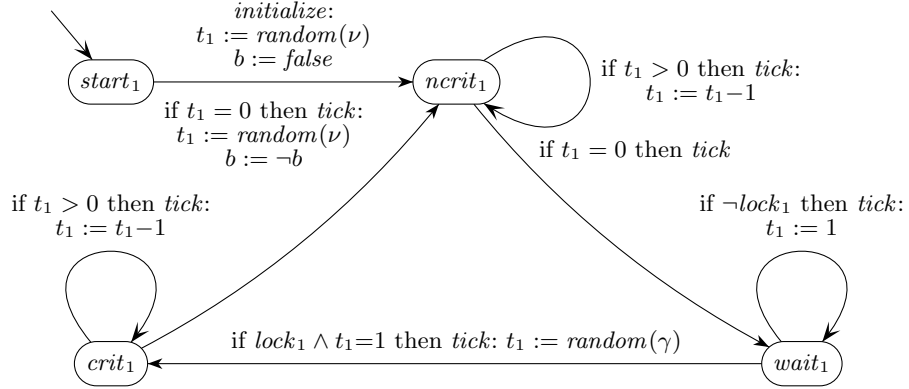


Fig. 8. Extended control flow graph of process P_1 for model I and (A2), (A2')

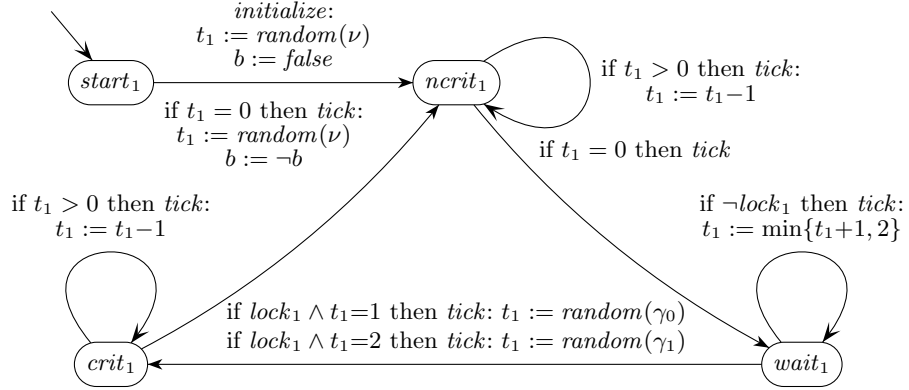


Fig. 9. Extended control flow graph of process P_1 for model II and (A2), (A2')

The variants (A2) and (A2') refer in \mathcal{M} to the steady-state probabilities under the condition $release_i = crit_i \wedge t_i = 0$ and $request_i = wait_i \wedge t_i = 0$, respectively. In \mathcal{M}' these conditions correspond to $release'_i = crit_i \wedge t_i = 0 \wedge \neg b$ and

$request'_i = wait_i \wedge t_i=0 \wedge \neg b$, respectively. In the original DTMC \mathcal{M} , we deal with the LTL-formulas φ_2 and φ'_2 . The switch to \mathcal{M}' however permits us to use the PCTL-formulas ψ_2 and ψ'_2 :

$$(A2) \quad \mathbb{P}^{\mathcal{M}}(\varphi_2 \mid release_i) = \mathbb{P}^{\mathcal{M}'}(\psi_2 \mid release'_i)$$

$$(A2') \quad \mathbb{P}^{\mathcal{M}}(\varphi'_2 \mid request_i) = \mathbb{P}^{\mathcal{M}'}(\psi'_2 \mid request'_i)$$

where

$$\varphi_2 = \bigcirc(\text{unlock } \mathcal{U} \text{ lock}_i)$$

$$\psi_2 = (\text{lock}_i \vee \text{unlock}) \mathcal{U} (\text{crit}_i \wedge b)$$

$$\varphi'_2 = \bigcirc(\text{lock}_i \mathcal{U} (\neg \text{crit}_i \wedge (\neg \text{spin}_i \mathcal{U} \text{crit}_i)))$$

$$\psi'_2 = \neg \text{spin}_i \mathcal{U} (\text{crit}_i \wedge b)$$

A.3 Results and Statistics of the PRISM calculations

In this section, we present results and statistics of the calculations carried out for queries (A1), (A2) and its variant (A2'), (A3) and its variant (A3') and (A4) using our extended version of PRISM. All the queries were evaluated for model I with the simplifications as described in Section A.1 and scaling factor $sf = 1000$ and for model II with scaling factor $sf = 200$ as described in Section 3, with the appropriate modifications for (A2) and (A2') as described in Section A.2. For model I, we present results for two to six processes, for model II, we present results for two to four processes.

We used the sparse engine of PRISM, applying the backward Gauß-Seidel iteration method. The steady-state distribution which is calculated to determine the conditional long-run probabilities and rewards is cached to allow re-use for subsequent queries carried out in the same run. The symmetry reduction implemented in PRISM was not applicable to our models. We performed all calculations on a dual-socket Intel Xeon L5630 (Quad-Core) system at 2.13 GHz with 32 GB total amount of RAM.

We will now provide tables listing the results and statistics of the various PRISM calculations. All times are in seconds. The first columns list the number of processes and the distributions considered in each run of PRISM using the notation introduced on page 13. The next two columns provide statistics about the system size, given by the number of states and the number of MTBDD nodes used for the transition matrix. The next column lists the time spent by PRISM for constructing the internal representation of the model and determining reachability. The column “time steady” then shows the time spent calculating the steady-state distribution $\theta(s)$ used for the calculation of conditional queries. The remaining columns then provide the results and time spent for the evaluation of the queries.

Queries (A1), (A3) and (A3')

For model I and scaling factor $sf = 1000$, the statistics for the computation of queries (A1) and (A3) are shown in Table 1 and in Table 2 for query (A3'). For model II and scaling factor $sf = 200$, Table 3 shows the statistics for all three queries (A1), (A3) and (A3').

The calculated value for (A1) corresponds to the probability to acquire the lock at the first attempt. The values calculated for (A3) are the expected time a process waits to grab the lock and the values calculated for (A3') correspond to the expected time a process is spinning after not acquiring the lock at the first attempt. The graphs in Fig. 4(b) show the results normalized against the base critical section length, i.e. 5 for model II, $sf = 200$ and 1000 for the measurements.

number procs.	distributions		system size		time	time	(A1)		(A3)	
	γ_0	ν	states	MTBDD	build	steady	value	time	value	time
2	[1]	[8,10]	243	244	0.02 s	0.02 s	0.990	0.00 s	0.010	0.00 s
2	[1]	[8,12]	339	239	0.03 s	0.03 s	0.988	0.00 s	0.012	0.01 s
2	[1]	[10,12]	339	235	0.03 s	0.04 s	0.993	0.00 s	0.007	0.01 s
2	[1]	[8,10,12]	339	247	0.02 s	0.02 s	0.982	0.00 s	0.018	0.00 s
2	[1]	[8,10,12,14]	451	254	0.02 s	0.02 s	0.977	0.00 s	0.023	0.01 s
3	[1]	[8,10]	2,575	1,052	0.08 s	0.08 s	0.960	0.00 s	0.045	0.02 s
3	[1]	[8,12]	4,155	1,693	0.09 s	0.14 s	0.973	0.01 s	0.029	0.02 s
3	[1]	[10,12]	4,287	1,039	0.13 s	0.12 s	0.970	0.01 s	0.033	0.02 s
3	[1]	[8,10,12]	4,299	1,091	0.06 s	0.07 s	0.945	0.00 s	0.060	0.02 s
3	[1]	[8,10,12,14]	6,647	1,132	0.06 s	0.08 s	0.940	0.01 s	0.064	0.03 s
4	[1]	[8,10]	26,070	7,207	0.30 s	0.49 s	0.922	0.02 s	0.093	0.13 s
4	[1]	[8,12]	48,600	11,210	0.40 s	1.12 s	0.950	0.03 s	0.055	0.26 s
4	[1]	[10,12]	52,554	8,572	0.54 s	0.96 s	0.941	0.02 s	0.069	0.21 s
4	[1]	[8,10,12]	53,466	4,358	0.22 s	0.53 s	0.902	0.02 s	0.115	0.13 s
4	[1]	[8,10,12,14]	96,550	4,536	0.19 s	0.74 s	0.900	0.03 s	0.115	0.21 s
5	[1]	[8,10]	249,966	33,500	1.21 s	2.99 s	0.880	0.08 s	0.152	0.88 s
5	[1]	[8,12]	543,366	52,890	2.40 s	7.67 s	0.913	0.17 s	0.104	1.80 s
5	[1]	[10,12]	621,302	45,617	2.77 s	7.47 s	0.909	0.15 s	0.112	1.89 s
5	[1]	[8,10,12]	650,782	19,685	1.10 s	4.23 s	0.857	0.12 s	0.180	1.31 s
5	[1]	[8,10,12,14]	1,383,077	16,255	0.93 s	7.21 s	0.857	0.23 s	0.177	2.50 s
6	[1]	[8,10]	2,259,830	131,803	7.37 s	21.49 s	0.833	0.51 s	0.223	6.19 s
6	[1]	[8,12]	5,737,572	218,922	17.14 s	51.77 s	0.865	1.17 s	0.178	14.92 s
6	[1]	[10,12]	7,055,432	205,199	19.92 s	69.37 s	0.873	1.34 s	0.163	16.78 s
6	[1]	[8,10,12]	7,704,497	89,681	5.94 s	39.53 s	0.810	1.20 s	0.257	14.54 s
6	[1]	[8,10,12,14]	19,478,654	61,027	5.94 s	85.94 s	0.813	2.70 s	0.250	34.88 s

Table 1. (A1) and (A3) for $sf = 1000$

number procs.	distributions		system size		time	time	(A3')	
	γ_0	ν	states	MTBDD	build	steady	value	time
2	[1]	[8,10]	245	269	0.02 s	0.02 s	1.000	0.01 s
2	[1]	[8,12]	341	264	0.03 s	0.03 s	1.000	0.01 s
2	[1]	[10,12]	341	260	0.03 s	0.04 s	1.000	0.01 s
2	[1]	[8,10,12]	342	275	0.02 s	0.03 s	1.000	0.01 s
2	[1]	[8,10,12,14]	455	283	0.03 s	0.02 s	1.000	0.01 s
3	[1]	[8,10]	2,618	1,156	0.08 s	0.08 s	1.114	0.02 s
3	[1]	[8,12]	4,206	1,824	0.09 s	0.14 s	1.043	0.04 s
3	[1]	[10,12]	4,338	1,141	0.13 s	0.13 s	1.106	0.02 s
3	[1]	[8,10,12]	4,373	1,198	0.06 s	0.08 s	1.089	0.02 s
3	[1]	[8,10,12,14]	6,758	1,248	0.06 s	0.09 s	1.076	0.03 s
4	[1]	[8,10]	26,748	7,538	0.31 s	0.50 s	1.186	0.16 s
4	[1]	[8,12]	49,548	11,746	0.42 s	1.12 s	1.105	0.28 s
4	[1]	[10,12]	53,514	8,965	0.56 s	0.99 s	1.166	0.22 s
4	[1]	[8,10,12]	54,817	4,734	0.24 s	0.55 s	1.171	0.16 s
4	[1]	[8,10,12,14]	98,844	4,930	0.20 s	0.76 s	1.154	0.25 s
5	[1]	[8,10]	258,990	34,675	1.24 s	3.00 s	1.260	1.04 s
5	[1]	[8,12]	558,258	54,803	2.50 s	7.86 s	1.203	2.11 s
5	[1]	[10,12]	636,782	47,135	2.80 s	7.56 s	1.225	2.04 s
5	[1]	[8,10,12]	672,378	20,604	1.13 s	4.42 s	1.259	1.44 s
5	[1]	[8,10,12,14]	1,424,878	17,286	0.98 s	7.41 s	1.240	2.64 s
6	[1]	[8,10]	2,366,890	134,861	7.42 s	21.08 s	1.334	7.59 s
6	[1]	[8,12]	5,946,062	223,652	17.81 s	52.39 s	1.318	16.34 s
6	[1]	[10,12]	7,281,652	209,554	19.36 s	70.54 s	1.283	17.42 s
6	[1]	[8,10,12]	8,022,062	92,093	6.37 s	40.72 s	1.354	15.93 s
6	[1]	[8,10,12,14]	20,185,714	63,320	6.25 s	89.70 s	1.336	36.10 s

Table 2. (A3') for $sf = 1000$

number procs.	distributions			system size		time	time	(A1)		(A3)		(A3')	
	γ_0	γ_1	ν	states	MTBDD	build	steady	value	time	value	time	value	time
2	[5]	[6]	[40,50]	1,342	1,445	0.12 s	0.29 s	0.944	0.00 s	0.086	0.02 s	1.545	0.01 s
2	[5]	[6]	[40,60]	2,202	1,687	0.20 s	0.51 s	0.983	0.00 s	0.052	0.03 s	3.118	0.02 s
2	[5]	[6]	[50,60]	1,852	1,628	0.18 s	0.43 s	0.955	0.00 s	0.066	0.02 s	1.462	0.02 s
2	[5]	[6]	[40,50,60]	2,506	1,029	0.12 s	0.25 s	0.926	0.00 s	0.129	0.02 s	1.745	0.02 s
2	[5]	[6]	[40,50,60,70]	3,350	1,227	0.15 s	0.31 s	0.925	0.00 s	0.145	0.03 s	1.926	0.02 s
3	[5]	[6]	[40,50]	67,001	32,772	4.05 s	3.45 s	0.872	0.03 s	0.277	0.57 s	2.166	0.45 s
3	[5]	[6]	[40,60]	142,575	40,645	11.57 s	9.24 s	0.955	0.05 s	0.147	0.99 s	3.229	0.72 s
3	[5]	[6]	[50,60]	110,013	42,545	7.77 s	5.94 s	0.898	0.05 s	0.209	0.98 s	2.047	0.87 s
3	[5]	[6]	[40,50,60]	195,849	10,892	3.56 s	6.38 s	0.850	0.04 s	0.362	0.67 s	2.410	0.54 s
3	[5]	[6]	[40,50,60,70]	306,697	12,344	4.06 s	8.77 s	0.854	0.06 s	0.371	1.06 s	2.547	0.85 s
4	[5]	[6]	[40,50]	4,082,808	569,046	316.31 s	149.10 s	0.797	1.12 s	0.531	25.34 s	2.612	22.13 s
4	[5]	[6]	[40,60]	8,776,938	822,186	819.24 s	385.80 s	0.889	2.11 s	0.376	48.96 s	3.395	41.39 s
4	[5]	[6]	[50,60]	8,325,516	878,109	759.38 s	313.11 s	0.838	2.06 s	0.397	48.57 s	2.446	38.44 s
4	[5]	[6]	[40,50,60]	13,580,130	117,898	127.21 s	150.41 s	0.767	2.06 s	0.669	35.78 s	2.865	33.06 s
4	[5]	[6]	[40,50,60,70]	25,406,006	113,008	144.22 s	299.64 s	0.775	3.71 s	0.668	70.44 s	2.967	65.21 s

Table 3. (A1), (A3) and (A3') for $sf = 200$

Queries (A2) and (A2')

For queries (A2) and (A2') the model variants and PCTL queries as described in Sec. A.2 where used. For model I and scaling factor $sf = 1000$, the results and statistics are shown in Table 4, while the results and statistics for model II and scaling factor $sf = 200$ are shown in Table 5. The calculated value for (A2) corresponds to the probability to re-acquire a previously held lock without other processes having acquired the lock in the meantime. The values for (A2') are the probability to acquire the lock without spinning twice in a row.

number procs.	distributions		system size		time	time	(A2)		(A2')	
	γ_0	ν	states	MTBDD	build	steady	value	time	value	time
2	[1]	[8,10]	482	350	0.03 s	0.04 s	0.051	0.01 s	0.980	0.01 s
2	[1]	[8,12]	670	416	0.04 s	0.05 s	0.094	0.02 s	0.976	0.02 s
2	[1]	[10,12]	674	339	0.04 s	0.06 s	0.042	0.01 s	0.986	0.02 s
2	[1]	[8,10,12]	674	354	0.03 s	0.04 s	0.081	0.01 s	0.963	0.02 s
2	[1]	[8,10,12,14]	898	364	0.03 s	0.03 s	0.105	0.02 s	0.953	0.02 s
3	[1]	[8,10]	5,103	1,547	0.12 s	0.13 s	0.003	0.04 s	0.923	0.12 s
3	[1]	[8,12]	8,235	2,190	0.14 s	0.21 s	0.017	0.06 s	0.946	0.16 s
3	[1]	[10,12]	8,523	1,650	0.20 s	0.21 s	0.002	0.04 s	0.943	0.15 s
3	[1]	[8,10,12]	8,554	1,356	0.09 s	0.12 s	0.011	0.05 s	0.893	0.15 s
3	[1]	[8,10,12,14]	13,238	1,406	0.11 s	0.14 s	0.019	0.06 s	0.884	0.19 s
4	[1]	[8,10]	51,817	8,154	0.51 s	0.81 s	0.000	0.18 s	0.854	0.81 s
4	[1]	[8,12]	96,841	12,150	0.71 s	1.77 s	0.003	0.30 s	0.903	1.47 s
4	[1]	[10,12]	104,725	9,695	0.93 s	1.56 s	0.000	0.27 s	0.889	1.42 s
4	[1]	[8,10,12]	106,576	5,160	0.48 s	0.96 s	0.001	0.20 s	0.817	1.13 s
4	[1]	[8,10,12,14]	192,560	5,214	0.44 s	1.32 s	0.004	0.29 s	0.811	1.61 s
5	[1]	[8,10]	499,170	34,427	2.21 s	5.20 s	0.000	0.74 s	0.781	5.26 s
5	[1]	[8,12]	1,085,326	54,670	4.35 s	13.09 s	0.000	1.66 s	0.835	10.59 s
5	[1]	[10,12]	1,241,718	46,530	5.20 s	13.72 s	0.000	1.69 s	0.831	10.83 s
5	[1]	[8,10,12]	1,299,121	22,551	2.49 s	7.68 s	0.000	1.24 s	0.739	7.96 s
5	[1]	[8,10,12,14]	2,761,581	17,411	2.41 s	13.41 s	0.001	2.32 s	0.737	13.15 s
6	[1]	[8,10]	4,516,160	133,779	11.53 s	39.53 s	0.000	3.92 s	0.703	25.45 s
6	[1]	[8,12]	11,465,152	223,238	25.80 s	89.18 s	0.000	10.17 s	0.751	59.16 s
6	[1]	[10,12]	14,107,112	207,169	30.52 s	130.06 s	0.000	11.51 s	0.770	65.64 s
6	[1]	[8,10,12]	15,390,508	97,733	13.10 s	70.81 s	0.000	10.70 s	0.662	55.19 s
6	[1]	[8,10,12,14]	38,925,105	67,238	14.68 s	158.19 s	0.000	25.76 s	0.664	113.91 s

Table 4. (A2) and (A2') for $sf = 1000$

number procs.	distributions			system size		time	time	(A2)		(A2')	
	γ_0	γ_1	ν	states	MTBDD	build	steady	value	time	value	time
2	[5]	[6]	[40,50]	2,674	1,554	0.17 s	0.38 s	0.056	0.06 s	0.889	0.14 s
2	[5]	[6]	[40,60]	4,384	1,799	0.28 s	0.77 s	0.100	0.09 s	0.967	0.20 s
2	[5]	[6]	[50,60]	3,694	1,739	0.26 s	0.59 s	0.045	0.08 s	0.909	0.20 s
2	[5]	[6]	[40,50,60]	5,002	1,036	0.19 s	0.38 s	0.085	0.08 s	0.852	0.23 s
2	[5]	[6]	[40,50,60,70]	6,690	1,220	0.21 s	0.46 s	0.107	0.11 s	0.849	0.31 s
3	[5]	[6]	[40,50]	133,604	33,875	5.80 s	5.11 s	0.000	1.05 s	0.760	6.59 s
3	[5]	[6]	[40,60]	284,682	42,058	15.53 s	14.11 s	0.020	2.12 s	0.910	12.11 s
3	[5]	[6]	[50,60]	219,558	45,011	11.32 s	9.05 s	0.000	1.85 s	0.806	10.88 s
3	[5]	[6]	[40,50,60]	391,588	10,523	6.50 s	10.84 s	0.008	1.25 s	0.720	12.66 s
3	[5]	[6]	[40,50,60,70]	613,254	11,610	7.84 s	14.96 s	0.017	1.99 s	0.726	21.21 s
4	[5]	[6]	[40,50]	8,164,362	652,002	394.71 s	254.47 s	0.000	33.16 s	0.639	247.55 s
4	[5]	[6]	[40,60]	17,552,628	939,739	1,063.92 s	669.62 s	0.003	68.19 s	0.785	513.10 s
4	[5]	[6]	[50,60]	16,649,628	1,037,676	1,134.80 s	529.04 s	0.000	68.11 s	0.706	487.28 s
4	[5]	[6]	[40,50,60]	27,155,404	126,006	218.60 s	263.96 s	0.001	36.76 s	0.590	453.80 s
4	[5]	[6]	[40,50,60,70]	50,810,662	115,527	278.36 s	528.41 s	0.003	74.33 s	0.600	878.63 s

Table 5. (A2) and (A2') for $sf = 200$

Query (A4)

The statistics for the computation of the quantile in (A4) are shown in Table 6 for model I and scaling factor $sf = 1000$ and in Table 7 for model II and scaling factor $sf = 200$. The threshold is 0.95. Column “steps” provides the resulting bound t necessary for surpassing the threshold probability, with column “prob.” providing the probability for the calculated bound t . The last column reports the time needed by the external script for the evaluation of the calculated values in each iteration of the bounded until calculation to determine at which step the threshold is surpassed, with PRISM calculating the necessary data for all bounds $t \leq 5$. To facilitate a comparison, we normalized the values t obtained from model II, $sf = 200$ and the measurements for Fig. 4(c), against the base critical section length, i.e. 5 and 1000, respectively.

number procs.	distributions		system size		time	time	(A4)	time	time	
	γ_0	ν	states	MTBDD	build	steady				steps
2	[1]	[8,10]	243	244	0.02 s	0.03 s	0	0.990	0.01 s	0.03 s
2	[1]	[8,12]	339	239	0.03 s	0.03 s	0	0.988	0.01 s	0.03 s
2	[1]	[10,12]	339	235	0.03 s	0.04 s	0	0.993	0.01 s	0.03 s
2	[1]	[8,10,12]	339	247	0.02 s	0.02 s	0	0.982	0.01 s	0.03 s
2	[1]	[8,10,12,14]	451	254	0.02 s	0.03 s	0	0.977	0.01 s	0.04 s
3	[1]	[8,10]	2,575	1,052	0.08 s	0.08 s	0	0.960	0.03 s	0.25 s
3	[1]	[8,12]	4,155	1,693	0.09 s	0.15 s	0	0.973	0.04 s	0.37 s
3	[1]	[10,12]	4,287	1,039	0.13 s	0.13 s	0	0.970	0.04 s	0.37 s
3	[1]	[8,10,12]	4,299	1,091	0.06 s	0.08 s	1	0.995	0.04 s	0.38 s
3	[1]	[8,10,12,14]	6,647	1,132	0.06 s	0.09 s	1	0.995	0.05 s	0.51 s
4	[1]	[8,10]	26,070	7,207	0.30 s	0.49 s	1	0.986	0.39 s	2.65 s
4	[1]	[8,12]	48,600	11,210	0.41 s	1.12 s	0	0.950	0.47 s	4.18 s
4	[1]	[10,12]	52,554	8,572	0.54 s	0.96 s	1	0.991	0.50 s	4.50 s
4	[1]	[8,10,12]	53,466	4,358	0.22 s	0.53 s	1	0.984	0.38 s	4.58 s
4	[1]	[8,10,12,14]	96,550	4,536	0.19 s	0.74 s	1	0.986	0.57 s	7.21 s
5	[1]	[8,10]	249,966	33,500	1.21 s	2.98 s	1	0.973	3.50 s	25.03 s
5	[1]	[8,12]	543,366	52,890	2.40 s	7.52 s	1	0.985	7.02 s	48.43 s
5	[1]	[10,12]	621,302	45,617	2.78 s	7.51 s	1	0.981	8.96 s	53.40 s
5	[1]	[8,10,12]	650,782	19,685	1.10 s	4.25 s	1	0.968	7.46 s	55.68 s
5	[1]	[8,10,12,14]	1,383,077	16,255	0.93 s	7.21 s	1	0.971	16.81 s	104.21 s
6	[1]	[8,10]	2,259,830	131,803	7.35 s	21.56 s	1	0.954	33.74 s	235.32 s
6	[1]	[8,12]	5,737,572	218,922	17.22 s	51.95 s	1	0.967	101.20 s	525.70 s
6	[1]	[10,12]	7,055,432	205,199	20.03 s	69.37 s	1	0.969	120.14 s	632.48 s
6	[1]	[8,10,12]	7,704,497	89,681	5.91 s	39.32 s	2	0.988	126.60 s	683.72 s
6	[1]	[8,10,12,14]	19,478,654	61,027	5.96 s	86.34 s	1	0.951	291.18 s	1,576.28 s

Table 6. (A4) for $sf = 1000$ and threshold 0.95

number procs.	distributions			system size		time	time	steps	(A4)		time script
	γ_0	γ_1	ν	states	MTBDD	build	steady		prob.	time	
2	[5]	[6]	[40,50]	1,342	1,445	0.12 s	0.29 s	1	0.990	0.02 s	0.04 s
2	[5]	[6]	[40,60]	2,202	1,687	0.21 s	0.52 s	0	0.983	0.03 s	0.06 s
2	[5]	[6]	[50,60]	1,852	1,628	0.18 s	0.45 s	0	0.955	0.02 s	0.05 s
2	[5]	[6]	[40,50,60]	2,506	1,029	0.12 s	0.25 s	1	0.982	0.02 s	0.07 s
2	[5]	[6]	[40,50,60,70]	3,350	1,227	0.15 s	0.32 s	1	0.977	0.03 s	0.09 s
3	[5]	[6]	[40,50]	67,001	32,772	4.06 s	3.45 s	2	0.955	0.59 s	2.40 s
3	[5]	[6]	[40,60]	142,575	40,645	15.25 s	9.73 s	0	0.955	1.02 s	4.63 s
3	[5]	[6]	[50,60]	110,013	42,545	7.72 s	5.85 s	1	0.961	0.95 s	3.54 s
3	[5]	[6]	[40,50,60]	195,849	10,892	3.59 s	6.38 s	3	0.960	0.72 s	6.08 s
3	[5]	[6]	[40,50,60,70]	306,697	12,344	4.04 s	8.72 s	3	0.955	1.05 s	8.72 s
4	[5]	[6]	[40,50]	4,082,808	569,046	305.10 s	145.49 s	4	0.954	35.48 s	141.65 s
4	[5]	[6]	[40,60]	8,776,938	822,186	825.78 s	394.35 s	3	0.950	60.60 s	290.65 s
4	[5]	[6]	[50,60]	8,325,516	878,109	754.58 s	309.86 s	3	0.957	65.04 s	259.34 s
4	[5]	[6]	[40,50,60]	13,580,130	117,898	131.74 s	158.47 s	5	0.986	69.82 s	422.28 s
4	[5]	[6]	[40,50,60,70]	25,406,006	113,008	149.04 s	313.46 s	4	0.953	111.51 s	729.37 s

Table 7. (A4) for $sf = 200$ and threshold 0.95