

Probabilistic Verification of Herman’s Self-Stabilisation Algorithm

Marta Kwiatkowska¹, Gethin Norman² and David Parker¹

¹Department of Computer Science, University of Oxford, Oxford, UK

²School of Computing Science, University of Glasgow, Glasgow, UK

Abstract. Herman’s self-stabilisation algorithm provides a simple randomised solution to the problem of recovering from faults in an N -process token ring. However, a precise analysis of the algorithm’s maximum execution time proves to be surprisingly difficult. McIver & Morgan have conjectured that the worst-case behaviour results from a ring configuration of 3 evenly spaced tokens, giving an expected time of approximately $0.15N^2$. However, the tightest upper bound proved to date is $0.64N^2$. We apply probabilistic verification techniques, using the probabilistic model checker PRISM, to analyse the conjecture, showing it to be correct for all sizes of the ring that can be exhaustively analysed. We furthermore demonstrate that the worst-case execution time of the algorithm can be reduced by using a biased coin.

Keywords: Self-stabilisation; Herman’s algorithm; Probabilistic model checking

1. Introduction

Self-stabilisation [Dij74] is an approach for designing fault-tolerant systems in a distributed setting. A self-stabilisation algorithm for a network of processes is a protocol that, starting from an arbitrary initial configuration, returns to a “stable” (legal) configuration within a finite number of steps and without any outside intervention. It is designed to allow the system to recover from transient faults, such as those caused by the states of individual processes becoming corrupted.

In this paper, we study the self-stabilisation algorithm of Herman [Her90], designed for a token ring of N synchronous processes. The algorithm is randomised and guarantees that the system eventually self-stabilises with probability 1. Another important property of the algorithm is the *expected time* required for stabilisation and, in particular, the *worst-case* expected time from an arbitrary system configuration. Despite the simplicity of Herman’s algorithm, a precise analysis of this property turns out to be surprisingly difficult. An upper bound of $O(N^2)$ has been proven by several people [MM05, Nak05, FMP06] and recent work in [KMO⁺11] established the tightest known upper bound of approximately $0.64N^2$. McIver and Morgan [MM05] also provided an exact expression for the expected stabilisation time from any configuration containing 3 tokens and posed an interesting conjecture that the case of 3 tokens spaced evenly around the ring always yields the worst-case behaviour of the algorithm. However, a proof of this conjecture remains elusive.

We analyse Herman’s algorithm using *probabilistic model checking*, which provides a way to formally specify and verify quantitative system properties, based on the construction and analysis of a finite-state probabilistic model. We use the probabilistic model checker PRISM [KNP11] to perform an exhaustive analysis of Herman’s algorithm on all system configurations up to a ring size of $N = 21$. We show that, in all these cases, the conjecture of McIver and Morgan does indeed hold. Furthermore, we then use PRISM to study the performance of the algorithm when using a biased coin, rather than a fair coin. Firstly, we discover that, in this case, 3 evenly spaced tokens does *not* always result in worst-case behaviour. More interestingly, we also find that, by selecting an appropriate coin bias, we can actually improve the performance of the algorithm, by reducing the worst-case expected stabilisation time.

Related Work. The problem of designing self-stabilising systems has received considerable attention; see for example [Dol00] for a detailed overview. Other self-stabilisation algorithms include those of Israeli & Jalfon [IJ90] and Beauquier, Gradinariu & Johnen [BGJ99]; however, both of these are designed for rings of asynchronous processes, rather than synchronous ones. A general approach for studying dependability properties of self-stabilisation algorithms, based on metrics that can be evaluated using probabilistic model checking was presented in [DTC⁺09].

As stated earlier, Herman’s algorithm has been studied by several people. McIver & Morgan [MM05], in addition to making the conjecture described above, presented an elementary proof of the $O(N^2)$ upper bound and related this to more general notions of abstraction for probabilistic programs [MM04]. Fribourg et al. [FMP06] proved the same upper bound, but using the notion of coupling from probability theory. Nakata, on the other hand, derived a more precise bound (of approximately $0.93N^2$) based on techniques for the analysis of coalescing random walks. Most recently, Kiefer et al. [KMO⁺11] applied techniques from annihilating particle systems to show an upper bound of approximately $0.64N^2$. They also show that stabilisation is significantly faster for specific classes of faults and study an asynchronous version of the algorithm. Finally, probabilistic model checking of Herman’s algorithm is used as an example in [KNP05], but only for a subset of the properties considered here, and without studying the effect of biased coins.

2. Herman’s Algorithm

Herman’s self-stabilisation algorithm assumes a network of N identical processes in a ring, indexed 1 to N and ordered anticlockwise. The protocol operates synchronously, the ring is oriented and N must be odd. Processes can possess *tokens*, which are passed unidirectionally around the ring. At every step of the algorithm, each process with a token decides whether to keep it or pass it on to its left neighbour, with the decision being based on the outcome of a random coin toss. When two tokens meet (i.e. are held by the same process), they are both eliminated. Thus, if the number of tokens is initially odd, it always remains odd. A configuration of the ring is said to be *stable* if exactly one process has a token. Once a stable configuration is reached, the token should be passed around the ring forever in a fair manner.

Following the presentation of Herman [Her90], the algorithm can be implemented using a single bit for each process. More precisely, the state of process i is represented by a bit x_i and $i \ominus 1$ denotes the index of process i ’s left (clockwise) neighbour. Process i is said to have a token if its bit x_i equals $x_{i \ominus 1}$. At each step of the algorithm, process i checks the values of x_i and $x_{i \ominus 1}$ for equality. If the values differ, then process i inverts its bit, making it the same as its neighbour ($x_i := x_{i \ominus 1}$); if the values are the same, then process i randomly sets x_i to either 0 (with probability p) or 1 (with probability $1-p$).

Figure 1 shows an example of two successive steps of the algorithm for $N = 9$ processes. We use x_i and \bar{x}_i to indicate that process i ’s bit is 1 or 0, respectively, and draw shaded grey circles to indicate processes that have tokens. Initially, processes 1, 4 and 6 have tokens; thus, each of the other six processes set their bit equal to their left neighbour’s, while processes 1, 4 and 6 set their bits randomly. In the example, they select values 0, 1 and 0, respectively. The resulting configuration is shown in the centre of Figure 1. Now, processes 1, 4 and 5 have tokens, i.e. process 6 has passed its token to process 5. In the second step, processes 1, 4 and 5 randomly select values 1, 0 and 1, respectively. This means that process 4 keeps its token, but process 5’s token is passed to the left, resulting in the elimination of both them. In the subsequent (stable) configuration, only process 1 has a token.




Fig. 1. Example successive configurations of Herman’s ring for $N = 9$; we write x_i and \bar{x}_i to denote the fact that bit x_i is 1 or 0, respectively; processes with a token are shaded grey.

3. Probabilistic Model Checking

Probabilistic model checking is an approach for formally verifying systems with stochastic behaviour. It works by constructing and analysing a probabilistic model of the system. In this paper, we assume that models are *discrete-time Markov chains* (DTMCs), which specify the probability of making a transition from one system state to another. Formally, we have the following definition.

Definition 1 (Discrete-time Markov chain). A *discrete-time Markov chain* (DTMC) is a tuple $D = (S, \bar{S}, \mathbf{P}, L)$ where S is a (finite) set of states, $\bar{S} \subseteq S$ is a set of possible initial states, $\mathbf{P} : S \times S \rightarrow [0, 1]$ is a transition probability matrix satisfying $\sum_{s' \in S} \mathbf{P}(s, s') = 1$ for all $s \in S$, and $L : S \rightarrow 2^{AP}$ is a labelling function mapping each state to a set of atomic propositions taken from a set AP .

An execution of a DTMC $D = (S, \bar{S}, \mathbf{P}, L)$ is represented by a *path*, which is an infinite sequence of states $\omega = s_0 s_1 s_2 \dots$ such that $\mathbf{P}(s_i, s_{i+1}) > 0$ for all $i \geq 0$. To reason about the likelihood of certain events occurring, we define, in standard fashion [KSK76], a probability measure Pr_s over the set of all paths starting from a given state s . Events of interest are then defined as measurable sets of paths.

We formally specify properties to be verified against DTMCs using probabilistic temporal logics such as PCTL [HJ94] or PCTL* [ASB⁺95]. A key element of these logics is the probabilistic operator \mathbf{P} . A formula $\mathbf{P}_{\bowtie p}[\psi]$, where $\bowtie \in \{\leq, <, \geq, >\}$, p is a probability and ψ is a path formula, asserts that the probability of ψ being true (starting from some state s) satisfies the bound $\bowtie p$. Typical examples are:

- $\mathbf{P}_{\geq 1}[\mathbf{F} \text{ “terminate”}]$ – “the algorithm eventually terminates with probability 1”;
- $\mathbf{P}_{\geq 0.98}[\mathbf{G}^{\leq 60} \neg \text{“fail”}]$ – “with probability at least 0.98, no failures occur within the first 60 steps”.

We also use *rewards* (or, equivalently, *costs*) to specify additional system properties, such as those concerning resource usage. A *reward structure* $r : S \times S \rightarrow \mathbb{R}_{\geq 0}$ augments the transitions of a DTMC with (non-negative) reward values. We then specify properties using an \mathbf{R} operator [KNP07], which takes a similar form to the \mathbf{P} operator from above. This refers to the *expected* amount of reward accumulated over a system’s execution. For example, using a simple reward structure *time* that assigns a reward of 1 to all transitions, we have:

- $\mathbf{R}_{< 20}^{\text{time}}[\mathbf{F} \text{ “terminate”}]$ – “the expected termination time of the algorithm is less than 20 seconds”;

For both the \mathbf{P} and \mathbf{R} operators, we also adopt a *quantitative* (numerical) form, which simply returns the actual probability or expected reward value, rather than comparing it to a specified threshold:

- $\mathbf{P}_{=?}[\mathbf{F} \text{ “terminate”}]$ – “what is the probability of the algorithm eventually terminating?”;
- $\mathbf{R}_{=?}^{\text{time}}[\mathbf{F} \text{ “terminate”}]$ – “what is the expected time taken for the algorithm to terminate?”.

Model checking for the properties given above reduces to a combination of graph-based algorithms (e.g., determining the set of states from which a target set can be reached) and numerical computation (typically solving linear equation systems). See, for example, [KNP07, BK08] for details.

4. Probabilistic Verification of Herman's Algorithm

We now describe how to model and analyse Herman's algorithm using probabilistic model checking. In particular, we will use the probabilistic model checker PRISM [KNP11]. We first show how the algorithm can be modelled and then describe the verification of various properties, focusing first on the correctness of the algorithm and then its performance, specifically its expected execution time. Finally, we use PRISM to evaluate the influence that a biased coin has on the performance of the algorithm.

4.1. Modelling

It is straightforward to model Herman's self-stabilisation algorithm as a DTMC. For N processes, and assuming that the algorithm works with a fair coin (i.e. $p = \frac{1}{2}$), we use a DTMC $D = (S, \bar{S}, \mathbf{P}, L)$ where:

- $S = \{0, 1\}^N$, i.e. each state is a vector $\underline{x} = (x_1, \dots, x_N)$, where the i th element corresponds to the value of process i 's bit;
- $\bar{S} = S$, i.e. the system can start in any of its possible states;
- $\mathbf{P} : S \times S \rightarrow [0, 1]$ is such that, for any states $\underline{x} = (x_1, \dots, x_N)$ and $\underline{y} = (y_1, \dots, y_N)$ we have $\mathbf{P}(\underline{x}, \underline{y}) = \prod_{i=1}^N \delta((x_{i \oplus 1}, x_i), y_i)$ where, for $x, x', y \in \{0, 1\}$:

$$\delta((x', x), y) = \begin{cases} 1 & \text{if } x \neq x' \text{ and } y = x' \\ \frac{1}{2} & \text{if } x = x' \\ 0 & \text{otherwise;} \end{cases}$$

- the set of atomic propositions is $AP = \{token_1, \dots, token_N, stable\}$ and, for any state $\underline{x} = (x_1, \dots, x_N)$, we have $token_i \in L(\underline{x})$ if $x_{i \oplus 1} = x_i$ and $stable \in L(\underline{x})$ if exactly one of $\{token_1, \dots, token_N\}$ is in $L(\underline{x})$.

Returning to the example steps of the algorithm shown in Figure 1 (where $N=9$), the corresponding path fragment of the DTMC is $(101100101)(010110010)(101010101)$ where:

$$\mathbf{P}((101100101), (010110010)) = \frac{1}{2} \cdot 1 \cdot 1 \cdot \frac{1}{2} \cdot 1 \cdot \frac{1}{2} \cdot 1 \cdot 1 \cdot 1 = \frac{1}{8}$$

$$\mathbf{P}((010110010), (101010101)) = \frac{1}{2} \cdot 1 \cdot 1 \cdot \frac{1}{2} \cdot 1 \cdot \frac{1}{2} \cdot 1 \cdot 1 \cdot 1 = \frac{1}{8}.$$

We next introduce a PRISM language description for this DTMC model of Herman's algorithm. Models in PRISM are described in a high-level language based on guarded commands. The basic components of the modelling language are *modules* and *variables*. A system model is defined by specifying a set of modules, each of whom's state is represented by a finite number of variables. The behaviour of each module is given by a set of guarded commands of the form:

$$[\langle \text{action} \rangle] \langle \text{guard} \rangle \rightarrow \langle \text{prob} \rangle : \langle \text{update} \rangle + \dots + \langle \text{prob} \rangle : \langle \text{update} \rangle;$$

The action label is used to force modules to synchronise (i.e. execute their commands simultaneously) and the guard is a predicate over all the variables of the model, indicating when the command is enabled. The updates describe the probabilistic transitions that the module can make when the command is executed, i.e. the changes made to its own variables; primed variables indicate the next values of variables.

A PRISM model of Herman's algorithm, for the case $N=5$, is shown in Figure 2 (left). We have one module for each process, and one variable for each module. Since the algorithm is synchronous, each command is labelled with the same action *step*, ensuring that all modules synchronise at each stage of the algorithm. The model description uses *module renaming* since the modules of processes 2, ..., 5 are identical in structure to the module for process 1. The fact that all possible system configurations can be initial states is expressed using the *init...endinit* construct. The model includes a reward structure *time*, which we will use to reason about the expected execution time of the algorithm. This is measured in terms of the number of steps performed, so it assigns a reward of 1 to every transition of the model. We also include a formula that calculates the number of tokens in a state, and then use this to define a label (an atomic proposition) for the stable configurations, i.e. states in which there is precisely one token. Figure 2 also shows the size of the resulting DTMC when constructed by PRISM, for varying N .

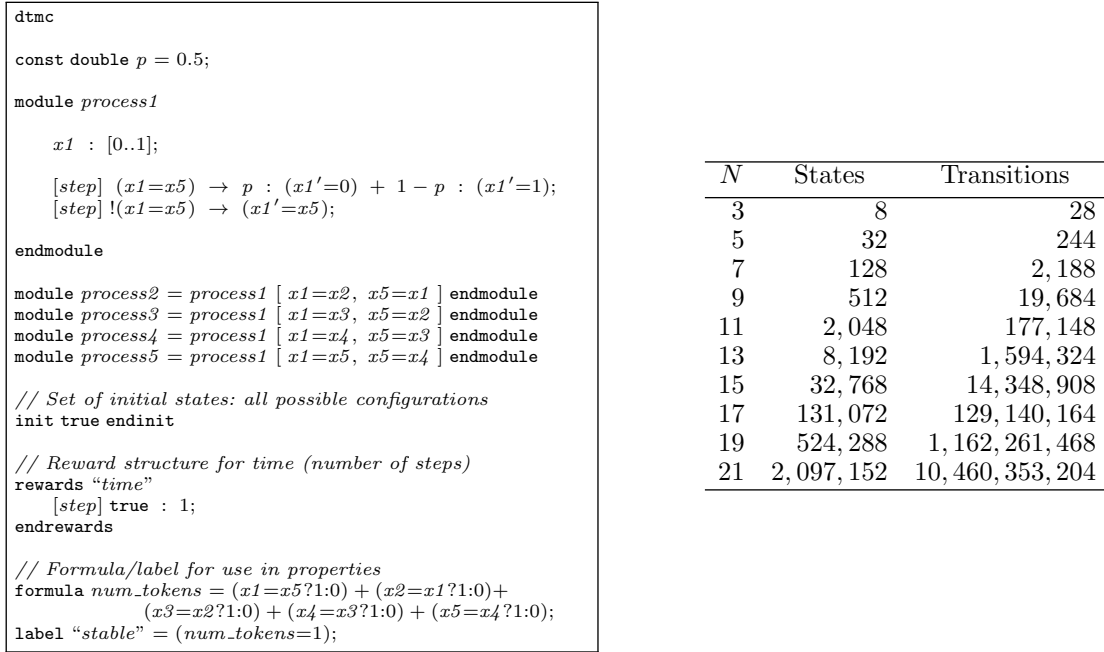


Fig. 2. Left: PRISM code for $N=5$; Right: Statistics showing model size as N varies.

4.2. Correctness

We begin by verifying that the protocol is correct, namely that, from any initial state, the ring eventually reaches a stable configuration with probability 1. In fact, it is relatively straightforward to prove that this property holds for any value of N (and any value of p) [Her90]. However, for completeness, we illustrate how this can be checked using PRISM (for fixed N). The required property is given by:

$$P_{\geq 1}[\mathbf{F} \text{ "stable"}]$$

where “stable”, as shown in Figure 2, identifies states of the model with one token. This property, along with the others that we will use in this section are also summarised in Table 1 for reference.


Since this a *qualitative* property (meaning that the probability bound is 0 or 1), model checking reduces to an analysis of the underlying graph of the DTMC, i.e., the exact probability values labelling its transitions are not important. This makes verification relatively efficient (compared to properties that also require numerical computation). Furthermore, PRISM’s *symbolic* implementation (using binary decision diagrams) allows the analysis to be performed for a large range of values of N . For example, in the case where $N = 31$, and the DTMC has $2^{31} = 2,147,483,648$ states, the model can be built and verified in under a second.

We also consider a second correctness property, namely that, after a stable configuration is reached, the token is passed around the ring in a fair manner. Since we already know that a stable configuration is reached with probability 1, to verify this property, it suffices to check that (from any possible initial state) each process gets a token infinitely often with probability 1. Formally this can be expressed by:

$$P_{\geq 1}[\wedge_{i=1}^N (\mathbf{G} \mathbf{F} \text{ "token}_i\text{"})]$$

where “token _{i} ” is true for states in which $x_i = x_{i \oplus 1}$.

As already mentioned, for qualitative properties such as the ones given here, the actual probability values of the DTMC are not used during the analysis. Therefore, although we have to specify a fixed value of the probability p to PRISM, it in fact follows from the definition of the model that checking for one value of $p \in (0, 1)$ implies that the above two properties hold for all values of $p \in (0, 1)$.



(a) Worst-case expected times from PRISM; with lower and upper bounds from [MM05] and [KMO⁺11], respectively.

(b) Worst-case expected times, over initial configurations containing k tokens

Fig. 3. Worst-case expected time for self-stabilisation.

4.3. Execution Time

Next, we consider a more interesting property of Herman’s algorithm: the expected amount of time required for its execution, i.e. the expected number of steps required until stabilisation occurs. From a fixed initial configuration of the ring, this corresponds to the PRISM query:

$$\mathbf{R}_{=?}^{time}[\mathbf{F} \text{ “stable”}]$$

where *time* denotes the reward structure shown in Figure 2 that assigns 1 to every transition of the model. Unlike the formulas used in the previous section, this is a *quantitative* property, whose verification requires numerical computation to be performed; here, this amounts to the solution of a linear equation system.

Of particular interest is the *worst-case* expected stabilisation time, from any possible initial configuration. We can check this with PRISM in a similar fashion. In fact, computationally, this is no more expensive than the previous property, since verification of that property proceeds by computing the expected time from all states of the model. In PRISM’s property specification language, the required query uses a *filter*:

$$\mathit{filter}(\mathit{max}, \mathbf{R}_{=?}^{time}[\mathbf{F} \text{ “stable”}], \text{“init”}).$$

A filter takes the form $\mathit{filter}(\mathit{op}, \mathit{prop}, \mathit{states})$, where *op* is a (commutative and associative) operator, *prop* is an (arbitrary) PRISM property and *states* is a Boolean-valued PRISM property. The filter applies operator *op* to values of *prop* over all states of the model that satisfy *states*.

A *lower* bound on the worst-case expected time for Herman’s algorithm can be established from the results of McIver and Morgan [MM05]. They show that, for an initial configuration containing 3 tokens, the expected stabilisation time is $4abc/N$, where a , b and c are the distances between the three tokens. The worst 3-token configuration, i.e. the one from which the execution time is the highest, is where the tokens are distributed evenly (or as evenly as possible) throughout the ring. In this case, since $a+b+c = N$, the execution time is $\frac{4}{27}N^2 \approx 0.15N^2$, which provides us with a lower bound on the worst-case expected time. McIver and Morgan also conjecture that this scenario actually constitutes the worst case over *all* possible initial configurations. However, a proof of this statement remains elusive.

The tightest *upper* bound on the worst-case time provided so far comes from recent work by Kiefer et al. [KMO⁺11], who prove a bound of $(\frac{\pi^2}{2} - \frac{116}{27})N^2 \approx 0.64N^2$. In Figure 3(a), we show results obtained from PRISM using the query given above, for values of N between 3 and 21, along with the lower and upper bounds from the results of [MM05] and [KMO⁺11], respectively. Notice that the worst-case times coincide precisely with the lower bound, corroborating the conjecture of [MM05].

The fact that worst-case expected times always result from configurations with 3 tokens (for $N=3, \dots, 21$) is more clearly illustrated using the following query:

$$\mathit{filter}(\mathit{max}, \mathbf{R}_{=?}^{time}[\mathbf{F} \text{ “stable”}], \mathit{num_tokens}=k)$$




Fig. 4. Worst-case expected times in the “random pass” model for various values of p ; individual points are obtained from PRISM and dotted lines show times for the worst 3-token configuration, from [KMO⁺11].

which gives the maximum expected time from k -token configurations. Figure 3(b) shows the results as the value of k varies.

Lastly, we note that, since PRISM performs an exhaustive search and analysis of the models it constructs, it can also report precisely which states result in the worst-case times. For all models analysed, these states are the ones proposed in [MM05], i.e. those in which the three tokens are evenly spaced around the ring.

4.4. Biased Coins

To complete our analysis of Herman’s algorithm, we consider how its performance varies for different values of the probability p , used to randomise the processes’ choices at each step. Herman’s original presentation of the algorithm includes p as a parameter, but the analysis of expected time assumes a fair coin (i.e. $p = 0.5$). We will show in this section that the worst-case performance can actually be improved using a biased coin.

Some previous analyses of Herman’s algorithm do consider variations in p , including [KMO⁺11]. However, there is a subtlety to be aware of in the way that biased coins are introduced. Our presentation of the algorithm, and our corresponding model in PRISM, follow the original description of Herman [Her90]: when process i has a token (i.e. $x_i = x_{i \oplus 1}$), he sets his bit x_i to 0 with probability p and to 1 with probability $1 - p$. If $p = 0.5$, as used in the analysis of [Her90], then an equivalent interpretation of the algorithm is as follows: when process i has a token, he passes it to his neighbour with probability p ; if this results in a process possessing two tokens, then both are removed. We will refer to these two interpretations of the algorithm as the “random-bit” model and “random-pass” model, respectively.

In the case where $p \neq 0.5$, however, these two interpretations do *not* coincide. When using the “random pass” model, in each state where a process has a token, the probability of him passing it to his neighbour is fixed and equals probability p . On the other hand, when using the “random-bit” model, his actions depend on both the relationship between his bit and his neighbour’s (i.e. whether he has a token) and the actual current value of his bit. For example, in the case where a process’s bit does equal that of his neighbour, the probability of him passing his token is:

$$\text{Prob}[\text{passes token}] = \begin{cases} (1-p)^2 + p^2 & \text{if left neighbour has a token} \\ p & \text{if left neighbour does not have a token and current bit value is 1} \\ 1-p & \text{if left neighbour does not have a token and current bit value is 0.} \end{cases}$$

From the above, it is clear that, for $p = 0.5$, the probability of passing a token is fixed at 0.5; however, for any other value of $p \in (0, 1)$, the three probability values in the expression are distinct. It is also worth noting that, for any $p \neq 0.5$, when its left neighbour has a token, the probability of passing on the token is strictly greater than 0.5.

Our analysis will focus mainly on the “random bit” model (which is the one shown in the PRISM code in Figure 2), however, we briefly also consider the “random pass” version, which is the one studied in [KMO⁺11]. Our PRISM model is easily adapted to this version by replacing the first command in module *process1* with:

$$[\text{step}] (x1 = x5) \rightarrow p : (x1' = 1 - x1) + 1 - p : (x1' = x1);$$

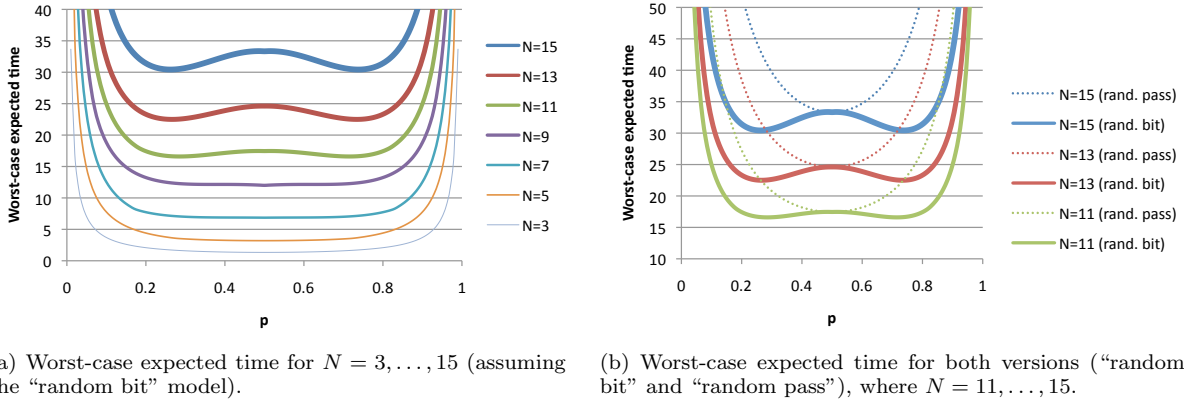

Fig. 5. Worst-case expected time for varying probability p .

Fig. 6. Magnified portions of the plots from Figure 5(a) for ring sizes 7, 9 and 11.

The analysis of [KMO⁺11] generalises the expression given in [MM05] for the expected time required from 3-token configurations to $abc/p(1-p)N$, where a , b and c are the distances between the tokens. In Figure 4, we plot the worst-case expected times over all configurations of the “random pass” model, computed by PRISM, against the worst-case expected time for a 3-token configuration, obtained from the expression above. It appears that, as for the case where $p = 0.5$, a 3-token configuration again always results in the worst-case behaviour.

Now, we consider the “random bit” model. Figure 5(a) shows the worst-case expected time over all configurations as p varies and for a range of values of N . Interestingly, we observe that, for $N \leq 9$, setting p to 0.5 results in the minimum worst-case expected time, but for larger values of N , this is not the case. In fact, for $N > 9$, we can actually improve the worst-case behaviour of the algorithm by using a biased coin with an appropriate value of p . By contrast, in the “random pass” version of the algorithm, the worst-case time is always minimised by choosing $p = 0.5$ and is higher than for the “random bit” model. Figure 5(b) illustrates this comparison for ring sizes 11, 13 and 15.

Also, whereas the worst-case time for the “random pass” model appears to be given by the simple expression $abc/p(1-p)N$ (i.e. assuming that the 3-token case is worst), the times for the “random bit” model follow a more complex pattern. This can be seen in Figure 6, which magnifies several of the plots from Figure 5(a) around the central point $p=0.5$.

Finally, we observe that, in the “random bit” model, 3-token configurations do *not* always result in the worst-case behaviour. The PRISM query:

$$\text{filter}(\text{range}, \text{num_tokens}, \text{filter}(\text{argmax}, \mathbf{R}_{=?}^{\text{time}}[\mathbf{F} \text{ “stable”}], \text{“init”}))$$

reveals the different possible numbers of the tokens that occur in the configurations yielding maximum expected times. For all values of N we considered, the worst-case number of tokens appears to be unique for each value of p , but varies between 3 and N for different values of p .

Table 1. Summary of the PRISM properties used in the paper.

Name	PRISM notation	Meaning
“ <i>terminate</i> ”	$P_{\geq 1}[F \text{ “stable”}]$	A stable configuration is eventually reached with probability 1.
“ <i>fairness</i> ”	$P_{\geq 1}[\bigwedge_{i=1}^N (G F \text{ “token}_i\text{”})]$	Each process gets a token infinitely often with probability 1.
“ <i>exp_time</i> ”	$R_{=?}^{time}[F \text{ “stable”}]$	Expected time (number of steps) until stabilisation.
“ <i>worst_time</i> ”	$filter(max, R_{=?}^{time}[F \text{ “stable”}], \text{“init”})$	Worst-case expected time until stabilisation, from any initial configuration.
“ <i>worst_time_k</i> ”	$filter(max, R_{=?}^{time}[F \text{ “stable”}], num_tokens=k)$	Worst-case expected time until stabilisation, from any initial configuration with k tokens.
“ <i>worst_states</i> ”	$filter(argmax, R_{=?}^{time}[F \text{ “stable”}], \text{“init”})$	True if a state results in the worst-case expected time until stabilisation.
“ <i>worst_tokens</i> ”	$filter(range, num_tokens, \text{“worst_states”})$	Range of values of the number of tokens in states satisfying property “ <i>worst_states</i> ”.

5. Conclusions

In this paper, we applied probabilistic model checking to analyse a simple randomised self-stabilisation algorithm due to Herman [Her90]. In particular, we investigated the worst-case expected time required for the algorithm to execute, providing experimental evidence to corroborate the unproven conjecture of McIver and Morgan [MM05] that 3 tokens, spaced evenly around the ring, always yields the maximum execution time. We also studied the effect of using a biased coin in the algorithm, discovering that this can actually improve the worst-case behaviour of the algorithm.

This example gives a good illustration of the benefits of probabilistic model checking. Firstly, it provides languages, techniques and tools that make it simple to model probabilistic systems and then analyse a variety of quantitative properties against them. Another key advantage of the approach is that it performs an *exhaustive* analysis: in this case, that amounts to computing values of a given property for all system configurations, then determining the maximum possible value and identifying the precise configurations that yield this value. By contrast, other analysis methods, such as those based on Monte Carlo simulation, are unable to perform an exhaustive analysis of this kind.

On the other hand, a disadvantage is that probabilistic model checking is typically restricted to analysing *finite-state* models. Furthermore, the size of such models that can be analysed is in practice always limited by the amount of time and space available. In this paper, the largest models for which we studied quantitative properties (that require numerical solution) contained about 2 million states and 10 billion transitions. In this instance, PRISM’s symbolic engines can handle these (and slightly larger) models without problems; the limiting factor is actually the time required to solve the models. But, more typically, it is the model size that represents the bottleneck in terms of scalability.

Progress has been made in recent years on the development of techniques to automatically construct finite abstractions of large or infinite-state probabilistic models [HWZ08, KKNP10]. However, for a comprehensive analysis of systems like Herman’s algorithm, we really require automated tools and techniques that can perform a *parametric* verification of the algorithm, i.e. for an arbitrary number of processes N . This represents a very challenging area of future work. On the other hand, promising steps have been taken [HHZ11] in the related problem of computing *symbolic* results expressed in terms of model parameters that affect probabilities (e.g. p in Herman’s algorithm), rather than the size of the model.

Acknowledgements

This work was part supported by ERC Advanced Grant VERIWARE.

References

- [ASB⁺95] A. Aziz, V. Singhal, F. Balarin, R. Brayton, and A. Sangiovanni-Vincentelli. It usually works: The temporal logic of stochastic systems. In P. Wolper, editor, *Proc. 7th Int. Conf. Computer Aided Verification (CAV'95)*, volume 939 of *LNCS*, pages 155–165. Springer, 1995.
- [BGJ99] J. Beauquier, M. Gradinariu, and C. Johnen. Memory space requirements for self-stabilizing leader election protocols. In *Proc. 18th ACM Symp. Principles of Distributed Computing (PODC'99)*, pages 199–208. ACM, 1999.
- [BK08] C. Baier and J.-P. Katoen. *Principles of Model Checking*. MIT Press, 2008.
- [Dij74] E. Dijkstra. Self-stabilizing systems in spite of distributed control. *Communications of the ACM*, 17(11):643–644, 1974.
- [Dol00] S. Dolev. *Self-Stabilization*. MIT Press, 2000.
- [DTC⁺09] A. Dhama, O. Theel, P. Crouzen, H. Hermanns, R. Wimmer, and B. Becker. Dependability engineering of silent self-stabilizing systems. In *Proc. 11th Int. Symp. Stabilization, Safety, and Security of Distributed Systems*, volume 5873 of *LNCS*, pages 238–253. Springer, 2009.
- [FMP06] L. Fribourg, S. Messika, and C. Picaronny. Coupling and self-stabilization. *Distributed Computing*, 18(3):221–232, 2006.
- [Her90] T. Herman. Probabilistic self-stabilization. *Information Processing Letters*, 35(2):63–67, 1990. An extended version is available at <ftp://ftp.math.uiowa.edu/pub/selfstab/H90.html>.
- [HHZ11] E. M. Hahn, H. Hermanns, and L. Zhang. Probabilistic reachability for parametric Markov models. *International Journal on Software Tools for Technology Transfer (STTT)*, 13(1):3–19, 2011.
- [HJ94] H. Hansson and B. Jonsson. A logic for reasoning about time and reliability. *Formal Aspects of Computing*, 6(5):512–535, 1994.
- [HWZ08] H. Hermanns, B. Wachter, and L. Zhang. Probabilistic CEGAR. In A. Gupta and S. Malik, editors, *Proc. 20th International Conference on Computer Aided Verification (CAV'08)*, volume 5123 of *LNCS*, pages 162–175. Springer, 2008.
- [IJ90] A. Israeli and M. Jalfon. Token management schemes and random walks yield self-stabilizing mutual exclusion. In *Proc. 9th ACM Symp. Principles of Distributed Computing (PODC'90)*, pages 119–131. ACM, 1990.
- [KKNP10] M. Kattenbelt, M. Kwiatkowska, G. Norman, and D. Parker. A game-based abstraction-refinement framework for Markov decision processes. *Formal Methods in System Design*, 36(3):246–280, 2010.
- [KMO⁺11] S. Kiefer, A. Murawski, J. Ouaknine, J. Worrell, and L. Zhang. On stabilization in Herman’s algorithm. In L. Aceto, M. Henzinger, and J. Sgall, editors, *Proc. 38th Int. Colloq. Automata, Languages and Programming (ICALP'11)*, volume 6756 of *LNCS*, pages 466–477. Springer, 2011.
- [KNP05] M. Kwiatkowska, G. Norman, and D. Parker. Quantitative analysis with the probabilistic model checker PRISM. *Electronic Notes in Theoretical Computer Science*, 153(2):5–31, 2005.
- [KNP07] M. Kwiatkowska, G. Norman, and D. Parker. Stochastic model checking. In M. Bernardo and J. Hillston, editors, *Formal Methods for the Design of Computer, Communication and Software Systems: Performance Evaluation (SFM'07)*, volume 4486 of *LNCS*, pages 220–270. Springer, 2007.
- [KNP11] M. Kwiatkowska, G. Norman, and D. Parker. PRISM 4.0: Verification of probabilistic real-time systems. In G. Gopalakrishnan and S. Qadeer, editors, *Proc. 23rd Int. Conf. Computer Aided Verification (CAV'11)*, volume 6806 of *LNCS*, pages 585–591. Springer, 2011.
- [KSK76] J. Kemeny, J. Snell, and A. Knapp. *Denumerable Markov Chains*. Springer, 2nd edition, 1976.
- [MM04] A. McIver and C. Morgan. *Abstraction, refinement and proof for probabilistic systems*. Springer, 2004.
- [MM05] A. McIver and C. Morgan. An elementary proof that Herman’s ring is $\Theta(N^2)$. *Information Processing Letters*, 94(2):79–84, 2005.
- [Nak05] T. Nakata. On the expected time for Herman’s probabilistic self-stabilizing algorithm. *Theoretical Computer Science*, 349(3):475–483, 2005.