

A Bayesian Approach to Model Checking Biological Systems*

Sumit K. Jha¹, Edmund M. Clarke¹, Christopher J. Langmead^{1,2},
Axel Legay³, André Platzer¹, and Paolo Zuliani¹

¹ Computer Science Department, Carnegie Mellon University, USA

² Lane Center for Computational Biology, Carnegie Mellon University, USA

³ Institut d'Informatique INRIA, Rennes, France

Abstract. Recently, there has been considerable interest in the use of Model Checking for Systems Biology. Unfortunately, the state space of stochastic biological models is often too large for classical Model Checking techniques. For these models, a statistical approach to Model Checking has been shown to be an effective alternative. Extending our earlier work, we present the first algorithm for performing statistical Model Checking using Bayesian Sequential Hypothesis Testing. We show that our Bayesian approach outperforms current statistical Model Checking techniques, which rely on tests from Classical (aka Frequentist) statistics, by requiring fewer system simulations. Another advantage of our approach is the ability to incorporate prior Biological knowledge about the model being verified. We demonstrate our algorithm on a variety of models from the Systems Biology literature and show that it enables faster verification than state-of-the-art techniques, even when no prior knowledge is available.

1 Introduction

Computational models are increasingly used in the field of Systems Biology to examine the dynamics of biological processes (e.g., [8,10,20,30,34,37]). By ‘computational’, we mean discrete-variable and continuous or discrete-time models [4], where the components of the system interact and evolve by obeying a set of instructions or rules. In contrast to differential equation-based models, which are also widely used in Systems Biology, computational models can provide insights into the role of stochastic effects over discrete-populations of molecules or cells. Recently, there has been considerable interest in the application of Model

* This research was sponsored by the GSRC (University of California) under contract no. SA423679952, National Science Foundation under contracts no. CCF0429120, no. CNS0411152, and no. CCF0541245, Semiconductor Research Corporation under contract no. 2005TJ1366, Air Force (University of Vanderbilt) under contract no. 18727S3, International Collaboration for Advanced Security Technology of the National Science Council, Taiwan, under contract no. 1010717, the U.S. Department of Energy Career Award (DE-FG02-05ER25696), and a Pittsburgh Life-Sciences Greenhouse Young Pioneer Award.

Checking [15] as a powerful tool for *formally* reasoning about the dynamic properties of such models (e.g., [1,6,9,11,14,18,24,38]). This paper presents a new Model Checking algorithm that is well-suited for verifying properties of very large stochastic models, such as those created and used in Systems Biology.

The stochastic nature of most computational models from Systems Biology gives rise to an instance of the *Probabilistic Model Checking* (PMC) problem [13,15,31]. Suppose \mathcal{M} is a stochastic model over a set of states S , s_0 is a starting state, ϕ is a dynamic property expressed as a formula in temporal logic, and $\theta \in [0, 1]$ is a probability threshold. The PMC problem is: given the 4-tuple $(\mathcal{M}, s_0, \phi, \theta)$, to decide algorithmically whether $\mathcal{M}, s_0 \models P_{\geq\theta}(\phi)$. In this paper, property ϕ is expressed in BLTL - Bounded Linear Temporal Logic [36,35,19]. Given these, PMC algorithms decide whether the model satisfies the property with at least probability θ .

Existing algorithms for solving the PMC problem fall into one of two categories. The first category comprises numerical methods (e.g. [2,3,12,16,31]) which can compute the probability with which the property holds with high precision. Numerical methods are generally only suitable for small systems ($\approx 10^6$ to 10^7 states). In a Biological System, the number of states can easily exceed this limit, which motivates the need for algorithms for solving the PMC problem in an approximate fashion. Approximate methods (e.g., [23,26,39,46]) work by sampling a set of *traces* from the model. Each trace is then evaluated to determine whether it satisfies the property. The number of satisfying traces is used to (approximately) decide whether $\mathcal{M}, s_0 \models P_{\geq\theta}(\phi)$.

Approximate PMC methods can be further divided into two sub-categories: (i) those that seek to *estimate* the probability that the property holds and then compare that estimate to θ (e.g., [26,39]), and (ii) those that reduce the PMC problem to a *hypothesis testing* problem (e.g., [46,47]). That is, deciding between two hypotheses — $H_0 : P_{\geq\theta}(\phi)$ versus $H_1 : P_{<\theta}(\phi)$. Hypothesis-testing based methods are more efficient than those based on estimation when θ (which is specified by the user) is significantly different than the true probability that the property holds (which is determined by \mathcal{M} and s_0) [45].

Existing PMC methods based on hypothesis testing rely on *Classical* (aka *Frequentist*) statistical procedures, like Wald’s Sequential Probability Ratio Test (SPRT) [42], to answer the decision problem. Our algorithm performs hypothesis testing, but uses *Bayesian* statistical procedures. This distinction is not trivial, as Bayesian and Classical statistics are two very different fields. We will show that in practice, our Bayesian approach requires fewer samples than Wald’s SPRT. Finally, we note that because we adopt a Bayesian approach, our algorithm can incorporate prior knowledge, in the form of a probability distribution, $P(\theta)$, when available. This is relevant because in a Biological setting, it is often the case that prior knowledge is available.

The contributions of this paper are as follows:

- The first application of Bayesian Sequential Hypothesis Testing to statistical Model Checking,

- The first hypothesis-testing based statistical Model Checking algorithm designed for composite hypotheses, which can in particular include prior knowledge via a mixture of prior distributions,
- A theorem proving that our algorithm terminates with probability 1,
- Error bounds for our algorithm, and
- A series of case studies using Systems Biology models demonstrating that our method is empirically more efficient than existing algorithms for statistical Model Checking.

2 Background and Related Work

Our algorithm can be applied to any stochastic model \mathcal{M} with a well-defined probability space over traces. Several well-studied stochastic models like (discrete and continuous) Markov Chains satisfy this property [47]. We assume that each execution of the system can be represented by a sequence of states and the time spent in these states. The sequence $\sigma = (s_0, t_0), (s_1, t_1), \dots$ denotes an execution of the system along states s_0, s_1, \dots with durations $t_0, t_1, \dots \in \mathbb{R}$. The system stays in state s_i for duration t_i and makes a transition to s_{i+1} . We require that the sum $\sum_i^\infty t_i$ must diverge, that is, the system can not make infinitely many state switches in finite time.

2.1 Specifying Properties in Temporal Logic

Our algorithm verifies properties of \mathcal{M} expressed as formulas in *Probabilistic Bounded Linear Temporal Logic* (PBLTL). We first define the syntax and semantics of *Bounded Linear Temporal Logic* (BLTL) [36,35,19] and then extend that logic to PBLTL.

For a stochastic model \mathcal{M} , let the set of state variables SV be a finite set of real-valued variables. A Boolean predicate over SV is a constraint of the form $x \sim v$, where $x \in SV$, $\sim \in \{\geq, \leq, =\}$, and $v \in \mathbb{R}$. A BLTL property is built on a finite set of Boolean predicates over SV using Boolean connectives and temporal operators. The syntax of the logic is given by the following grammar:

$$\phi ::= x \sim v \mid (\phi_1 \vee \phi_2) \mid (\phi_1 \wedge \phi_2) \mid \neg \phi_1 \mid (\phi_1 \mathbf{U}^t \phi_2),$$

where $\sim \in \{\geq, \leq, =\}$, $x \in SV$, $v \in \mathbb{Q}$, and $t \in \mathbb{Q}_{\geq 0}$. We can define additional temporal operators such as $\mathbf{F}^t \psi = \mathbf{True} \mathbf{U}^t \psi$, or $\mathbf{G}^t \psi = \neg \mathbf{F}^t \neg \psi$ in terms of the bounded until \mathbf{U}^t .

We define the semantics of BLTL with respect to executions of \mathcal{M} . The fact that an execution σ satisfies property ϕ is denoted by $\sigma \models \phi$. Let $\sigma = (s_0, t_0), (s_1, t_1), \dots$ be an execution of the model along states s_0, s_1, \dots with durations $t_0, t_1, \dots \in \mathbb{R}$. We denote the execution trace starting at state i by σ^i (in particular, σ^0 denotes the original execution σ). The value of the state variable x in σ at the state i is denoted by $V(\sigma, i, x)$. The semantics of BLTL for a trace σ^k starting at the k^{th} state ($k \in \mathbb{N}$) is defined as follows:

- $\sigma^k \models x \sim v$ if and only if $V(\sigma, k, x) \sim v$;
- $\sigma^k \models \phi_1 \vee \phi_2$ if and only if $\sigma^k \models \phi_1$ or $\sigma^k \models \phi_2$;

- $\sigma^k \models \phi_1 \wedge \phi_2$ if and only if $\sigma^k \models \phi_1$ and $\sigma^k \models \phi_2$;
- $\sigma^k \models \neg\phi_1$ if and only if $\sigma^k \models \phi_1$ does not hold (written $\sigma^k \not\models \phi_1$);
- $\sigma^k \models \phi_1 \mathbf{U}^t \phi_2$ if and only if there exists $i \in \mathbb{N}$ such that (a) $\sum_{0 \leq l < i} t_{k+l} \leq t$, (b) $\sigma^{k+i} \models \phi_2$ and (c) for each $0 \leq j < i$, $\sigma^{k+j} \models \phi_1$.

Statistical Model Checking is based on evaluating whether $\sigma \models \phi$ holds on sample simulations σ of the system. In practice, sample simulations only have a finite duration. The question is how long these simulations have to be for the formula ϕ to have a well-defined semantics such that $\sigma \models \phi$ can be checked. If σ is too short, say of duration 2, the semantics of $\phi_1 \mathbf{U}^5 \phi_2$ may be unclear. But at what duration of the simulation can we stop because we know that the truth-value for $\sigma \models \phi$ will never change by continuing the simulation? In [29], we prove that finite simulations of bounded duration are always sufficient for Model Checking BLTL on traces.

We can now define Probabilistic Bounded Linear Temporal Logic.

Definition 1. A *Probabilistic Bounded LTL (PBLTL) formula* is a formula of the form $P_{\geq \theta}(\phi)$, where ϕ is a BLTL formula and $\theta \in (0, 1)$.

We say that \mathcal{M} satisfies PBLTL property $P_{\geq \theta}(\phi)$, denoted by $\mathcal{M} \models P_{\geq \theta}(\phi)$, if and only if the probability that an execution of \mathcal{M} satisfies BLTL property ϕ is greater than or equal to θ . The problem is well-defined [47] since one can always assign a unique probability measure to the set of executions of \mathcal{M} that satisfy a formula in BLTL. Note that counterexamples to the BLTL property ϕ are *not* counterexamples to the PBLTL property $P_{\geq \theta}(\phi)$, because the truth of $P_{\geq \theta}(\phi)$ depends on the likelihood of all counterexamples to ϕ . This makes PMC more difficult than standard Model Checking, because one counterexample to ϕ is not enough to answer $P_{\geq \theta}(\phi)$.

2.2 Existing Statistical Probabilistic Model Checking Algorithms

As outlined in the introduction, Probabilistic Model Checking algorithms can either be exact (e.g. [2,3,12,16,31]), or statistical in nature. In practice, statistical methods (e.g., [23,26,32,39,46]), which iteratively draw sample traces from the model, are generally better suited to Model Checking Biological systems because they scale better. Our method is statistical, and so we will compare and contrast our method to existing statistical methods in this section.

Existing PMC methods based on hypothesis testing rely on *Classical* (aka *Frequentist*) statistical procedures, like Wald's Sequential Probability Ratio Test (SPRT) [42], to answer the decision problem. Younes and Simmons introduced the first algorithm for statistical Model Checking [45,46,47] for verifying probabilistic temporal properties of stochastic systems. Their work uses the SPRT, which is designed for *simple* hypothesis testing¹. Specifically, the SPRT decides between

¹ A simple hypothesis completely specifies a distribution. For example, a Bernoulli distribution of parameter p is fully specified by the hypothesis $p = 0.5$ (or some other fixed value). A composite hypothesis has instead free parameters, e.g. the hypothesis $p < 0.3$, for a Bernoulli distribution.

the simple null hypothesis $H'_0 : \mathcal{M}, s_0 \models P_{=\theta_0}(\phi)$ against the simple alternate hypothesis $H'_1 : \mathcal{M}, s_0 \models P_{=\theta_1}(\phi)$, where $\theta_0 < \theta_1$. It can be shown that the SPRT is optimal for simple hypothesis testing, in the sense that it minimizes the expected number of samples among all the tests satisfying the same Type I and II errors [43], when either H'_0 or H'_1 is true. The PMC problem is instead a choice between two *composite* hypotheses $H_0 : \mathcal{M}, s_0 \models P_{\geq\theta}[\phi]$ versus $H_1 : \mathcal{M}, s_0 \models P_{<\theta}[\phi]$. The SPRT is not defined unless $\theta_0 \neq \theta_1$, so Younes and Simmons overcome this problem by separating the two hypotheses by an *indifference region* $(\theta - \delta, \theta + \delta)$, where $0 < \delta < 1$ is a user-specified parameter. It can be shown that the SPRT with indifference region can be used for testing composite hypotheses, while respecting the same Type I and II errors of a standard SPRT [21, Section 3.4]. However, in this case the test is no longer optimal, and the maximum expected sample size may be much bigger than the optimal fixed sample size sampling test - see [7] and [21, Section 3.6]. We note that our algorithm solves the composite hypothesis testing problem, but does so using Bayesian statistics, and thus requires no indifference region.

The method of [26] uses a fixed number of samples and estimates the probability the property holds as the number of satisfying traces divided by the number of sampled traces. Their algorithm guarantees the accuracy of the results using Chernoff-Hoeffding bounds. In particular, their algorithm can guarantee that the difference in the estimated and the true probability is less than ϵ , with probability ρ , where $\rho < 1$ and $\epsilon > 0$ are user-specified parameters. Grosu and Smolka use a similar technique for verifying formulas in LTL [23]. Their algorithm randomly samples lassos from a Büchi automaton in an on-the-fly fashion. The method of [32] is also Bayesian, like the algorithm in this paper, but estimates the probability the property holds and does not invoke hypothesis testing. Unlike the algorithm in this paper, [32] is fully Bayesian in the sense that it explicitly considers the prior distributions over the initial state and parameters of the model, in addition to the prior over the property.

Finally, Sen *et al.* [39,40] used the *p-value* for the null hypothesis as a statistic for hypothesis testing. The *p-value* is defined as the probability of obtaining observations at least as extreme as the one that was actually seen, given that the null hypothesis is true. It is important to realize that a *p-value* is *not* the probability that the null hypothesis is true. Sen *et al.*'s method does not have a way to control the Type I and II errors.

3 Bayesian Statistical Model Checking

In this section, we first review some important concepts from statistical Model Checking, and then introduce theory and terminology from Bayesian statistics. We then present our algorithm in Sec. 3.2.

Recall that the PMC problem is to decide whether $\mathcal{M} \models P_{\geq\theta}(\phi)$, where $\theta \in (0, 1)$ and ϕ is a BLTL formula. Let p be the (unknown but fixed) probability of the model satisfying ϕ : thus, the PMC problem can now be stated as deciding between two hypotheses:

$$H_0 : p \geq \theta \quad H_1 : p < \theta.$$

For any trace σ_i of the system, we can deterministically decide whether σ_i satisfies ϕ . Therefore, we can define a Bernoulli random variable X_i denoting the outcome of $\sigma_i \models \phi$. The probability mass function associated with X_i is thus:

$$f(x_i|u) = p^{x_i}(1-p)^{1-x_i}$$

where $x_i = 1$ iff $\sigma_i \models \phi$, otherwise $x_i = 0$. Note that the X_i are independent and identically distributed, as each trace is given by an independent execution of the model. Since p is unknown, we assume that it is given by a random variable, whose density $g(\cdot)$ is called the *prior* density. The prior is usually based on our previous experiences and beliefs about the system. A complete lack of information about the probability of the system satisfying the formula is usually summarized by a *non-informative* or *objective* prior probability.

3.1 Bayesian Statistics

Suppose we have a sequence of random variables X_1, \dots, X_n defined as above, and let $d = (x_1, \dots, x_n)$ denote a sample of those variables. Then Bayes' theorem states that the *posterior odds* are

$$P(H_0|d) = \frac{P(d|H_0)P(H_0)}{P(d)} \quad P(H_1|d) = \frac{P(d|H_1)P(H_1)}{P(d)}$$

where $P(d) = P(d|H_0)P(H_0) + P(d|H_1)P(H_1)$, which in our case is always non-zero. The ratio of the posterior odds for hypotheses H_0 and H_1 given data d is

$$\frac{P(H_0|d)}{P(H_1|d)} = \frac{P(d|H_0)}{P(d|H_1)} \frac{P(H_0)}{P(H_1)}. \quad (1)$$

Definition 2. The Bayes factor \mathcal{B} of sample d and hypotheses H_0 and H_1 is

$$\mathcal{B} = \frac{P(d|H_0)}{P(d|H_1)}.$$

For fixed priors in a given example, the Bayes factor is directly proportional to the posterior odds ratio by Equation (1). Thus, it may be used as a measure of relative confidence in H_0 vs. H_1 , as proposed by Jeffreys [28]. In particular, he suggested that a value of the Bayes factor greater than 100 provides decisive evidence in favor of H_0 . To test H_0 vs. H_1 we compute the Bayes factor \mathcal{B} of the available data and then compare it against a fixed threshold $T > 1$: we shall accept H_0 iff $\mathcal{B} > T$. Jeffreys interprets the value of the Bayes factor as a measure of the evidence in favor of H_0 (dually, $\frac{1}{\mathcal{B}}$ is the evidence in favor of H_1).

We now show how to compute the Bayes factor. According to Definition 2, we have to calculate the probability of the observed sample $d = (x_1, \dots, x_n)$ given H_0 and H_1 . They are given by integrating the joint density $h(d|\cdot)$ with respect to the prior $g(\cdot)$, and since we assume that the sample is drawn from iid variables, we have that $h(d|\cdot) = f(x_1|\cdot) \cdots f(x_n|\cdot)$. Therefore, the Bayes factor is the ratio:

$$\mathcal{B} = \frac{P(x_1, \dots, x_n | H_0)}{P(x_1, \dots, x_n | H_1)} = \frac{\int_{\theta}^1 f(x_1|u) \cdots f(x_n|u) \cdot g(u) \, du}{\int_0^{\theta} f(x_1|u) \cdots f(x_n|u) \cdot g(u) \, du} . \tag{2}$$

We observe that the Bayes factor depends on the data d and on the prior g , so it may be considered a measure of confidence in H_0 vs. H_1 provided by the data x_1, \dots, x_n , and “weighted” by the prior g . Hence, the choice of the threshold Bayes Factor (T) in Sec. 3.2 also indicates an objective degree of confidence in the accepted hypothesis when the Bayesian Statistical Model Checking algorithm stops.

3.2 Algorithm

Our algorithm is essentially a sequential version of Jeffreys’ test. Remember we want to establish whether $\mathcal{M} \models P_{\geq \theta}(\phi)$, where $\theta \in (0, 1)$ and ϕ is a BLTL formula. Like all statistical Model Checking algorithms, we assume that it is possible to generate unbiased samples from the model. The algorithm iteratively draws independent and identically distributed sample traces $\sigma_1, \sigma_2, \dots$, and checks whether they satisfy ϕ . As explained above, we can model this procedure as independent sampling from a Bernoulli distribution X of unknown parameter p - the actual probability of the model satisfying ϕ . At stage n the algorithm has drawn samples x_1, \dots, x_n iid like X . It then computes the Bayes factor \mathcal{B}_n according to (2), and it stops iff $(\mathcal{B}_n > T \vee \mathcal{B}_n < \frac{1}{T})$. When this occurs, it will accept H_0 iff $\mathcal{B}_n > T$, and will accept H_1 iff $\mathcal{B}_n < \frac{1}{T}$. The algorithm is shown below.

From (2) we see that the algorithm can incorporate prior knowledge through g , when computing the Bayes factor. Our examples focus on Beta priors which are defined over the $(0, 1)$ interval by the following probability density (for real parameters $\alpha, \beta > 0$):

$$\forall u \in (0, 1) \quad g(u, \alpha, \beta) \hat{=} \frac{1}{B(\alpha, \beta)} u^{\alpha-1} (1-u)^{\beta-1} \tag{3}$$

where the Beta function $B(\alpha, \beta)$ is defined as:

$$B(\alpha, \beta) \hat{=} \int_0^1 t^{\alpha-1} (1-t)^{\beta-1} dt . \tag{4}$$

By varying the parameters α and β , one can approximate other smooth unimodal densities on $(0, 1)$ by a Beta density (e.g., the uniform density over $(0, 1)$ is a Beta with $\alpha = \beta = 1$). We also define the Beta distribution function $F_{(\alpha, \beta)}(u)$:

$$\forall u \in (0, 1) \quad F_{(\alpha, \beta)}(u) \hat{=} \int_0^u g(t, \alpha, \beta) dt = \frac{1}{B(\alpha, \beta)} \int_0^u t^{\alpha-1} (1-t)^{\beta-1} dt \tag{5}$$

which is just the usual distribution function for a Beta random variable of parameters α, β (i.e., the probability that it takes values less than or equal to u).

Algorithm 1. Bayesian Statistical Model Checking

Require: PBLTL Property $P_{\geq\theta}(\phi)$, Threshold $T > 1$, Prior density g for unknown parameter p

```

 $n := 0$       {number of traces drawn so far}
 $x := 0$       {number of traces satisfying  $\phi$  so far}
repeat
   $\sigma :=$  draw a sample trace of the system (iid)
   $n := n + 1$ 
  if  $\sigma \models \phi$  then
     $x := x + 1$ 
  end if
   $\mathcal{B}_n :=$  BayesFactor( $n, x$ )    {compute according to Equation (2)}
until ( $\mathcal{B}_n > T \vee \mathcal{B}_n < \frac{1}{T}$ )
if ( $\mathcal{B}_n > T$ ) then
  return  $H_0$  accepted
else
  return  $H_1$  accepted
end if

```

The choice of the Beta density is not arbitrary. It is well-known that the Beta distribution is the *conjugate prior* to the Bernoulli distribution². This relationship gives rise to closed-form solutions to the *posterior* density over θ (i.e., $P(\theta|d)$), thus avoiding numerical integration when calculating the Bayes factor. Our data (x_1, \dots, x_n) are assumed to be iid samples drawn from a Bernoulli distribution of unknown parameter p . We write $x = \sum_{i=1}^n x_i$ for the number of successes in (x_1, \dots, x_n) . The prior density $g(\cdot)$ is assumed to be a Beta density with fixed parameters $\alpha, \beta > 0$. In [29], we show that the Bayes factor \mathcal{B}_n at stage n can be computed in terms of the Beta distribution function:

$$\mathcal{B}_n = \frac{1}{F_{(x+\alpha, n-x+\beta)}(\theta)} - 1 .$$

The Beta distribution function can be computed with high accuracy by standard mathematical libraries (e.g. the GNU Scientific Library) or software (e.g. Matlab). Hence, the Beta distribution is the appropriate choice for summarizing the prior probability distribution in Statistical Model Checking.

We present the following two Theorems:

Theorem 1 (Termination). *The Bayesian Statistical Model Checking algorithm terminates with probability one, for Beta priors and Bernoulli samples. (See [29] for a proof.)*

² A distribution $P(\theta)$ is said to be a conjugate prior for a likelihood function, $P(d|\theta)$, if the posterior, $P(\theta|d)$ is in the same family of distributions.

Theorem 2. *If the Bayesian Model Checking algorithm terminates after observing n sample traces, an upper bound on the probability of the Type I error is*

$$\sum_{x=0}^n I_{\{\mathcal{B}(n,x) < 1/T\}}(x) \binom{n}{x} t_{max}^x (1 - t_{max})^{n-x}$$

where t_{max} is the value of t that maximizes the expression $t^i(1-t)^{n-i}$ defined on $[\theta, 1]$, T is the Bayes Factor threshold used in the Bayesian Model Checking algorithm, and I is the indicator function. (See [29] for a proof.)

3.3 Verification over General Priors

The use of conjugate priors does not pose restrictions, in practice. It is known that any prior distribution (with or without a density) can be well approximated by a *finite* mixture of conjugate priors [17]. Thus, we can approximate an arbitrary prior over $(0, 1)$ by constructing a density $G(\cdot)$ of the form:

$$G(u) \hat{=} \sum_{i=1}^N r_i \cdot g_i(u, \alpha_i, \beta_i)$$

where N is a positive integer which depends on the level of accuracy required, the g_i 's are Beta densities (of possibly different parameters α_i, β_i), and the r_i 's are positive reals summing up to 1 - this ensures that G is a proper density.

For such priors, the computation of the Bayes factor is slightly more complicated. In [29], we show that the Bayes factor at stage n is given by:

$$\mathcal{B}_n = \frac{\sum_{i=1}^N r'_i \cdot B(x + \alpha_i, n - x + \beta_i)}{\sum_{i=1}^N r'_i \cdot B(x + \alpha_i, n - x + \beta_i) \cdot F_{(x+\alpha_i, n-x+\beta_i)}(\theta)} - 1$$

where $r'_i = \frac{r_i}{B(\alpha_i, \beta_i)}$. Again, we see that the Bayes factor can be computed by means of standard, well-known numerical methods, thereby simplifying the implementation of the algorithm. Theorem 1 can be extended to handle this case, too [29].

4 Benchmarks

In this section, we analyze the performance of our algorithm on five benchmark models from the Systems Biology literature. Three of the models are written in the PRISM Model Checking tool's specification language [27,31], and the remaining two are written in SBML and were obtained from the Matlab Systems Biology Toolbox. The PRISM Model Checker tool is capable of both symbolic (i.e., exact) Probabilistic Model Checking, and statistical Probabilistic Model Checking. PRISM's statistical Probabilistic Model Checking Algorithm implements the algorithm of [26] which uses a fixed sized sampling approach and estimates the true probability as the number of satisfying traces over the number of sampled

traces. We note that for each of the benchmark sets, we consider models that are too large for symbolic model checking.

Our experiments demonstrate two important properties of our algorithm: (i) we show that our algorithm requires fewer traces than either the algorithm of [26] implemented in PRISM or Wald's SPRT algorithm - while retaining the same bounds on the frequentist Type-I and Type-II error probabilities. (ii) The performance of both the Wald's algorithm [42] and our Bayesian Model Checking algorithm degrades as the threshold probability (i.e., θ) in the PBLTL temporal logic formula gets close to the actual probability of the model satisfying the BLTL formula. However, the Bayesian algorithm shows a more graceful degradation compared to Wald's SPRT approach.

4.1 PRISM Benchmarks

We studied three large PRISM benchmarks which are not well suited for numerical approaches to Probabilistic Model Checking. In our experiments, the Bayesian Model Checking algorithm used uniform priors, and accepted a hypothesis when it was 10000 times more likely than the other hypothesis (Bayes Factor threshold $T = 10000$). Our experiments with Wald's SPRT used Type I and II error bounds of 0.01. We chose an indifference region δ so as to make the Type I and Type II errors for both the Wald's Test and the Bayes Factor test equal. The statistical estimation engine of the PRISM model checker always needed 92042 samples to estimate the probability of the BLTL formulae being true.

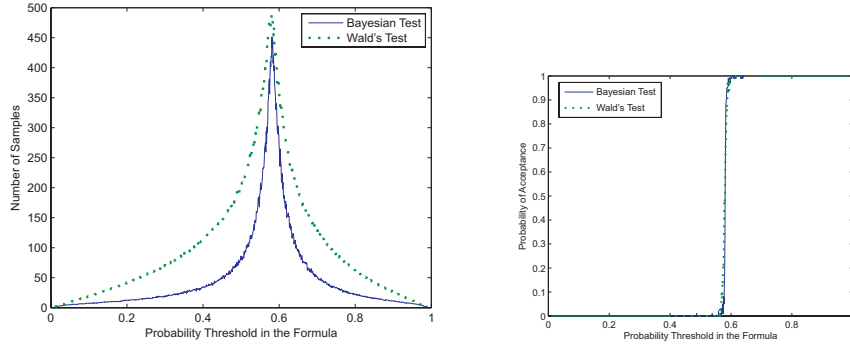
The results of experiments with the Fibroblast Growth Factor Signaling Model (see [24], [25] for details) are presented. We checked the property whether the probability that Grb2 binds to FRS2 within 20 time units exceeds θ (for several values of θ):

$$H_0 : \mathcal{M} \models P_{\geq \theta} [\mathbf{F}^{20} (FRS2_GRB > 0)]$$

The power curves and the number of samples for this benchmark are plotted in Fig. 2(a) and Fig. 2(b) respectively. A power curve indicates the probability of accepting the null hypothesis for various values of the threshold probability θ in the PBLTL formula. We chose the Wald's Test so that its power curve matched that of the Bayesian Test at the 0.01 and 0.99 acceptance probability. The goal is to make sure that the two tests have equal statistical power. From Figure 2(b), it is clear that both the power curves are almost on top of each other and hence, both the tests have indeed been calibrated to be equally powerful. The Bayesian algorithm needs fewer samples than Wald's SPRT test for this benchmark. This shows that the Bayesian Statistical Model Checking performs better than an approach based on Wald's SPRT.

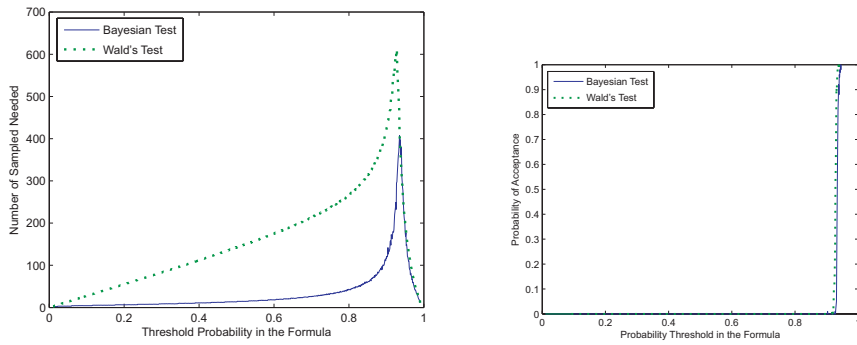
We also studied the continuous time Markov Chain model [5,41] for circadian rhythm. We checked the property that the probability of the number of activated messenger RNAs exceeding 5 units within 0.25 time units is more than θ (for various values of θ):

$$H_0 : \mathcal{M} \models P_{\geq \theta} [\mathbf{F}^{0.25} (ma > 5)]$$



(a) Number of Samples for various probability thresholds in the formula (b) Power Curve of the Bayesian and Wald's approach

Fig. 1. Fibroblast Growth Factor Signaling Model: The system satisfies the formula with probability 0.58. (Bayes Factor=10000)

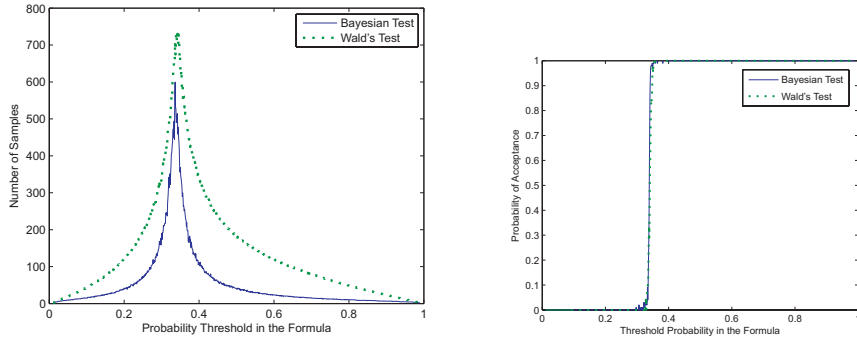


(a) Number of Samples for various probability thresholds in the formula (b) Power Curve of the Bayesian and Wald's approach

Fig. 2. Circadian Rhythm: The system satisfies the formula with probability 0.93. (Bayes Factor=10000)

The power curves and the number of samples for this benchmark are plotted in Fig. 2(b) and Fig. 2(a) respectively. We calibrated Wald's test so that its power curve closely matched that of the Bayesian Test so as to make a fair comparison. From the figure, we observe that the Bayesian algorithm always needs fewer samples than the Wald's SPRT test for this benchmark.

We also analyzed the model on Cell cycle control [33] and studied the probability that Cyclin gets bound within the first 0.5 time units. We check the property that the probability of the number of bound Cyclin molecules exceeds 3 units within 0.5 time units exceeds θ (for various values of θ):



(a) Number of Samples for various probability thresholds in the formula (b) Power Curve of the Bayesian and Wald's approach

Fig. 3. Cell Cycle Control: The system satisfies the formula with probability 0.34. (Bayes Factor=10000)

$$H_0 : \mathcal{M} \models P_{\geq \theta} [\mathbf{F}^{0.5} (\text{cyclin_bound} > 3)]$$

The results of our experiment are presented in Fig. 3(a). The Bayesian Statistical Model Checking algorithm usually required fewer samples than the approach based on Wald's SPRT.

4.2 SBML Experiments

We also studied SBML models using the implementation of Gillespie's Stochastic Simulation Algorithm in Matlab's Systems Biology Toolbox. We analyzed two large models with over 10^8 and 10^{17} species. We used monitors written in Matlab to verify the BLTL properties on traces. Our analysis of the experiments in this section is purely Bayesian, i.e., we have studied the performance of the algorithm over only one run (using uniform priors). In the previous sections, we had compared the performance of our algorithm with Wald's SPRT by running the algorithm several times on the same model - a frequentist approach.

We analyzed the Yeast Heterotrimeric G Protein Cycle benchmark [44]. We analyzed the property that the G protein stays above the threshold of 6000 units for 2 time units and falls below 6000 before 20 time units.

$$H_0 : \mathcal{M} \models P_{\geq \theta} [\mathbf{G}^2(GProtein > 6000) \text{ and } \mathbf{F}^{20}(GProtein < 6000)] .$$

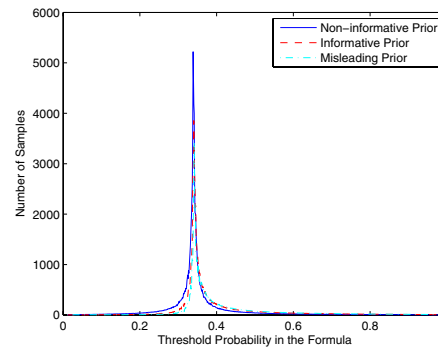
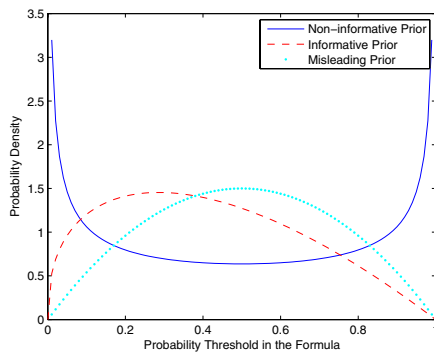
We also ran experiments using the Lotka model [22] and verified the property that the number of copies of the x species rises to a threshold level within 0.01 time units.

$$H_0 : \mathcal{M} \models P_{\geq \theta} [\mathbf{F}^{0.01}(x > 1.4 * 10^7)]$$

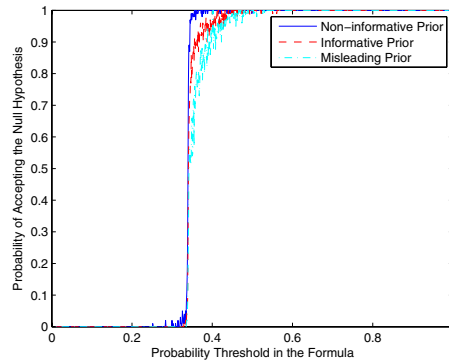
Table 1. Performance on the G Protein (left) and Lotka Benchmark (right)

Probability	# Samples Needed
0.2	3
0.6	8
0.8	14
0.9	23
0.9999	99

Probability	# Samples Needed
0.1	2
0.5	6
0.7	10
0.9	23
0.99	69



(a) Shape of the Priors used in our Experiments (b) Number of Samples with Different Classes of Priors



(c) Power curve of the tests of the Algorithm

Fig. 4. Different Classes of Priors

The results of our experiments are shown in Table 1: both hypotheses are always accepted, although the number of samples increases with the probability threshold of the temporal formula.

4.3 Experiment with Different Classes of Priors

We investigated the effect of priors on the performance of the Bayesian Model Checking algorithm. We used three different priors - non-informative prior, an informative prior and a misleading prior. The priors, the number of samples needed by the Bayesian algorithm for these priors, and the power curve for each of these priors is also plotted in Fig. 4(a), Fig. 4(b) and Fig. 4(c) respectively. The priors used are Beta distributions with different shape parameters: (i) $\alpha = 1/2, \beta = 1/2$: non-informative prior, (ii) $\alpha = 1.4, \beta = 2$: informative prior with a peak around 0.34 (iii) $\alpha = 2, \beta = 2$: a misleading prior with peak around 0.5.

Fig. 4(b) shows that the number of samples needed by the Bayesian algorithm becomes smaller when the prior probability distribution is informative and supports the true hypothesis. Also, the power curve (see Fig. 4(c)) becomes sharper when the Bayesian algorithm is given a correct and informative prior probability distribution. A completely non-informative prior also performs well both in the number of samples and the power of the test. Strongly misleading priors make the power curve less steep. However, the algorithm still performs quite well when the actual probability of the system is away from the threshold probability in the formula.

5 Conclusions and Future Work

We have introduced the first algorithm for Probabilistic Model Checking based on Bayesian Sequential Hypothesis Testing. Our algorithm terminates with probability 1, and provides bounds on the probability of returning an incorrect answer. Empirically, we have shown that our algorithm requires fewer traces to terminate than techniques based on Classical Statistics. This is not surprising as the Bayesian method comparing composite hypotheses whereas techniques like Wald's SPRT are comparing simple hypotheses. This advantage in efficiency is important in the context of Systems Biology as the cost of generating traces is not necessarily negligible. Bayesian methods also afford a convenient means for incorporating domain knowledge through the prior distributions.

Our algorithm is presently limited to incorporating prior information on the probability that the property is true. A more fully Bayesian approach would incorporate prior information on not just the property, but also the starting state and parameters of the model. We are presently extending our method to address this limitation.

Acknowledgments

The authors would like to thank H.L.S. Younes for comments on a draft of this paper.

References

1. Antoniotti, M., Policriti, A., Ugel, N., Mishra, B.: Model building and model checking for biochemical processes. *Cell Biochem. Biophys.* 38(3), 271–286 (2003)
2. Baier, C., Clarke, E.M., Hartonas-Garmhausen, V., Kwiatkowska, M.Z., Ryan, M.: Symbolic model checking for probabilistic processes. In: Degano, P., Gorrieri, R., Marchetti-Spaccamela, A. (eds.) *ICALP 1997*. LNCS, vol. 1256, pp. 430–440. Springer, Heidelberg (1997)
3. Baier, C., Haverkort, B.R., Hermanns, H., Katoen, J.-P.: Model-checking algorithms for continuous-time markov chains. *IEEE Trans. Software Eng.* 29(6), 524–541 (2003)
4. Bailey, N.: *The Elements of Stochastic Processes with Applications to the Natural Sciences*. Wiley-IEEE (1990)
5. Barkai, N., Leibler, S.: Biological rhythms: Circadian clocks limited by noise. *Nature* 403, 267–268 (2000)
6. Batt, G., Ropers, D., de Jong, H., Geiselman, J., Mateescu, R., Page, M., Schneider, D.: Validation of qualitative models of genetic regulatory networks by model checking: analysis of the nutritional stress response in *Escherichia coli*. *Bioinformatics* 25(1), i19–i28 (2005)
7. Bechhofer, R.: A note on the limiting relative efficiency of the Wald sequential probability ratio test. *J. Amer. Statist. Assoc.* 55, 660–663 (1960)
8. Calder, M., Gilmore, S., Hillston, J.: Modelling the influence of RKIP on the ERK signalling pathway using the stochastic process algebra PEPA. *Transactions on Computational Systems Biology* (2006) (in press)
9. Calder, M., Vysheirsky, V., Gilbert, D., Orton, R.: Analysis of signalling pathways using the PRISM model checker. In: *Proc. Computational Methods in Systems Biology (CMSB 2005)*, pp. 179–190 (2005)
10. Cardelli, L.: Abstract machines of systems biology. In: Priami, C., Merelli, E., Gonzalez, P., Omicini, A. (eds.) *Transactions on Computational Systems Biology III*. LNCS (LNBI), vol. 3737, pp. 145–168. Springer, Heidelberg (2005)
11. Chabrier, N., Fages, F.: Symbolic Model Checking of Biochemical Networks. In: Priami, C. (ed.) *CMSB 2003*. LNCS, vol. 2602, pp. 149–162. Springer, Heidelberg (2003)
12. Ciesinski, F., Größer, M.: On probabilistic computation tree logic. In: Baier, C., Haverkort, B.R., Hermanns, H., Katoen, J.-P., Siegle, M. (eds.) *Validation of Stochastic Systems*. LNCS, vol. 2925, pp. 147–188. Springer, Heidelberg (2004)
13. Clarke, E.M., Emerson, E.A.: Design and synthesis of synchronization skeletons using branching-time temporal logic. In: *Logic of Programs, Workshop*, London, UK, pp. 52–71. Springer, Heidelberg (1982)
14. Clarke, E.M., Faeder, J.R., Langmead, C.J., Harris, L.A., Jha, S.K., Legay, A.: Statistical model checking in biolab: Applications to the automated analysis of t-cell receptor signaling pathway. In: Heiner, M., Uhrmacher, A.M. (eds.) *CMSB 2008*. LNCS (LNBI), vol. 5307, pp. 231–250. Springer, Heidelberg (2008)
15. Clarke, E.M., Grumberg, O., Peled, D.A.: *Model Checking*. MIT Press, Cambridge (1999)
16. Courcoubetis, C., Yannakakis, M.: The complexity of probabilistic verification. *Journal of the ACM* 42(4), 857–907 (1995)
17. Diaconis, P., Ylvisaker, D.: Quantifying prior opinion. In: Bernardo, J.M., De Groot, M.H., Lindley, D.B., Smith, A.F.M. (eds.) *Bayesian Statistics 2: Proceedings of the 2nd Valencia International Meeting*. Elsevier Science Publisher, Amsterdam (1985)

18. Fages, F.: Temporal logic constraints in the biochemical abstract machine biocham. In: Hill, P.M. (ed.) LOPSTR 2005. LNCS, vol. 3901, pp. 1–5. Springer, Heidelberg (2006)
19. Finkbeiner, B., Sipma, H.: Checking finite traces using alternating automata. In: Proceedings of Runtime Verification (RV 2001), pp. 44–60 (2001)
20. Fisher, J., Piterman, N., Hubbard, E.J., Stern, M.J., Harel, D.: Computational insights into *caenorhabditis elegans* vulval development. *Proc. Natl. Acad. Sci. U S A* 102(6), 1951–1956 (2005)
21. Ghosh, B., Sen, P. (eds.): Handbook of sequential analysis. Dekker, New York (1991)
22. Gillespie, D.T.: Exact stochastic simulation of coupled chemical reactions. *The Journal of Physical Chemistry* 81(25), 2340–2361 (1977)
23. Grosu, R., Smolka, S.: Monte Carlo Model Checking. In: CAV, pp. 271–286 (2005)
24. Heath, J., Kwiatkowska, M., Norman, G., Parker, D., Tymchyshyn, O.: Probabilistic model checking of complex biological pathways. In: Priami, C. (ed.) CMSB 2006. LNCS (LNBI), vol. 4210, pp. 32–47. Springer, Heidelberg (2006)
25. Heath, J., Kwiatkowska, M., Norman, G., Parker, D., Tymchyshyn, O.: Probabilistic model checking of complex biological pathways. *Theoretical Computer Science* 319(3), 239–257 (2008)
26. Hérault, T., Lassaigne, R., Magniette, F., Peyronnet, S.: Approximate probabilistic model checking. In: Steffen, B., Levi, G. (eds.) VMCAI 2004. LNCS, vol. 2937, pp. 73–84. Springer, Heidelberg (2004)
27. Hinton, A., Kwiatkowska, M., Norman, G., Parker, D.: PRISM: A tool for automatic verification of probabilistic systems. In: Hermanns, H., Palsberg, J. (eds.) TACAS 2006. LNCS, vol. 3920, pp. 441–444. Springer, Heidelberg (2006)
28. Jeffreys, H.: *Theory of Probability*. Clarendon Press, Oxford (1961)
29. Jha, S.K., Clarke, E.M., Langmead, C.J., Legay, A., Platzer, A., Zuliani, P.: A bayesian approach to model checking biological systems. Technical Report CMU-CS-09-110, Computer Science Department, Carnegie Mellon University (2009)
30. Kam, N., Harel, D., Cohen, I.R.: Modeling biological reactivity: Statecharts vs. boolean logic. In: Proceedings of the Second International Conference on Systems Biology (2001)
31. Kwiatkowska, M.Z., Norman, G., Parker, D.: Prism 2.0: A tool for probabilistic model checking. In: QEST, pp. 322–323. IEEE, Los Alamitos (2004)
32. Langmead, C.J.: Generalized Queries and Bayesian Statistical Model Checking in Dynamic Bayesian Networks: Application to Personalized Medicine. In: Proc. 8th Ann. Intl Conf. on Comput. Sys. Bioinf. (CSB), pp. 201–212 (2009)
33. Lecca, P., Priami, C.: Cell cycle control in eukaryotes: A BioSpi model. In: Proc. Workshop on Concurrent Models in Molecular Biology (BioConcur 2003). ENTCS (2003)
34. McAdams, H., Shapiro, L.: Circuit simulation of genetic networks. *Science* 269, 650–656 (1995)
35. Owicki, S.S., Lamport, L.: Proving liveness properties of concurrent programs. *ACM Trans. Program. Lang. Syst.* 4(3), 455–495 (1982)
36. Pnueli, A.: The temporal logic of programs. In: FOCS, pp. 46–57. IEEE, Los Alamitos (1977)
37. Priami, C., Regev, A., Shapiro, E., Silverman, W.: Application of a stochastic name-passing calculus to representation and simulation of molecular processes. *Inf. Process. Lett.* 80(1), 25–31 (2001)

38. Sadot, A., Fisher, J., Barak, D., Admanit, Y., Stern, M.J., Hubbard, E.J.A., Harel, D.: Toward verified biological models. *IEEE/ACM Transactions on Computational Biology and Bioinformatics* 5(2), 223–234 (2008)
39. Sen, K., Viswanathan, M., Agha, G.: Statistical model checking of black-box probabilistic systems. In: Alur, R., Peled, D.A. (eds.) *CAV 2004*. LNCS, vol. 3114, pp. 202–215. Springer, Heidelberg (2004)
40. Sen, K., Viswanathan, M., Agha, G.: On statistical model checking of stochastic systems. In: Etesami, K., Rajamani, S.K. (eds.) *CAV 2005*. LNCS, vol. 3576, pp. 266–280. Springer, Heidelberg (2005)
41. Vilar, J., Kueh, H.-Y., Barkai, N., Leibler, S.: Mechanisms of noise-resistance in genetic oscillators. *Proc. Nat. Acad. Sci. USA* 99(9), 5988–5992 (2002)
42. Wald, A.: Sequential tests of statistical hypotheses. *Annals of Mathematical Statistics* 16(2), 117–186 (1945)
43. Wald, A.: *Sequential Analysis*. Dover Publications (June 2004)
44. Yi, T.M., Kitano, H., Simon, M.I.: A quantitative characterization of the yeast heterotrimeric g protein cycle. *Proc. Natl. Acad. Sci. USA* 100(19), 10764–10769 (2003)
45. Younes, H.L.S., Kwiatkowska, M.Z., Norman, G., Parker, D.: Numerical vs. statistical probabilistic model checking. *STTT* 8(3), 216–228 (2006)
46. Younes, H.L.S., Simmons, R.G.: Probabilistic verification of discrete event systems using acceptance sampling. In: Brinksma, E., Larsen, K.G. (eds.) *CAV 2002*. LNCS, vol. 2404, pp. 223–235. Springer, Heidelberg (2002)
47. Younes, H.L.S., Simmons, R.G.: Statistical probabilistic model checking with a focus on time-bounded properties. *Information and Computation* 204(9), 1368–1409 (2006)