

Central Limit Model Checking*

LUCA BORTOLUSSI, Department of Mathematics and Geosciences, University of Trieste

LUCA CARDELLI, Microsoft Research & University of Oxford

MARTA KWIATKOWSKA, University of Oxford

LUCA LAURENTI[†], University of Oxford

We consider probabilistic model checking for continuous-time Markov chains (CTMCs) induced from Stochastic Reaction Networks (SRNs) against a fragment of Continuous Stochastic Logic (CSL) extended with reward operators. Classical numerical algorithms for CSL model checking based on uniformisation are limited to finite CTMCs and suffer from exponential growth of the state space with respect to the number of species. On the other hand, approximate techniques such as mean-field approximations and simulations combined with statistical inference are more scalable, but can be time consuming and do not support the full expressiveness of CSL. In this paper we employ a continuous-space approximation of the CTMC in terms of a Gaussian process based on the Central Limit Approximation (CLA), also known as the Linear Noise Approximation (LNA), whose solution requires solving a number of differential equations that is quadratic in the number of species and independent of the population size. We then develop efficient and scalable approximate model checking algorithms on the resulting Gaussian process, where we restrict the target regions for probabilistic reachability to convex polytopes. This allows us to derive an abstraction in terms of a time-inhomogeneous discrete-time Markov chain (DTMC), whose dimension is independent of the number of species, on which model checking is performed. Using results from probability theory, we prove the convergence in distribution of our algorithms to the corresponding measures on the original CTMC. We implement the techniques and, on a set of examples, demonstrate that they allow us to overcome the state space explosion problem, while still correctly characterizing the stochastic behaviour of the system. Our methods can be used for formal analysis of a wide range of distributed stochastic systems, including biochemical systems, sensor networks and population protocols.

CCS Concepts: • **Mathematics of computing** → **Probability and statistics**; • **Theory of computation** → **Logic and verification**;

ACM Reference Format:

Luca Bortolussi, Luca Cardelli, Marta Kwiatkowska, and Luca Laurenti. 2019. Central Limit Model Checking. 1, 1 (May 2019), 33 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Distributed systems with Markovian interactions can be modelled as continuous-time Markov chains [29]. Examples include randomised population protocols [5], genetic regulatory networks [53] and biochemical systems evolving in a spatially homogeneous environment, at constant volume and temperature [29, 32]. For

*This research is supported by a Royal Society Research Professorship. LB is supported by EU-FET project QUANTICOL (nr 600708).

[†]This is the corresponding author

Authors' addresses: Luca Bortolussi, Department of Mathematics and Geosciences, University of Trieste, lbortolussi@units.it; Luca Cardelli, Microsoft Research & University of Oxford, luca@microsoft.com; Marta Kwiatkowska, University of Oxford, marta.kwiatkowska@cs.ox.ac.uk; Luca Laurenti, University of Oxford, luca.laurenti@cs.ox.ac.uk.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Association for Computing Machinery.

XXXX-XXXX/2019/5-ART \$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

such systems, stochastic modelling is necessary to describe stochastic fluctuations for low/medium population counts that deterministic fluid techniques cannot capture [29].

A versatile programming language for modelling the behaviour of Markovian distributed systems is that of *Stochastic Reaction Networks (SRNs)*, which induce CTMCs under certain mild restrictions. Computing the probability distributions of the species of a SRN over time is achieved by solving the Kolmogorov Equation, also known in the biochemical literature as the Chemical Master Equation (CME) [54]. Unfortunately, classical numerical solution methods for computing transient probability based on uniformisation [9] are often infeasible because of the state space explosion problem, that is, the number of states of the resulting Markov chain grows exponentially with respect to the number of species and may be infinite. A more scalable transient analysis can be achieved by employing simulations combined with statistical inference [31], but to obtain good accuracy large numbers of simulations are needed, which for some systems can be very time consuming.

A promising approach, which we explore in this paper, is to instead approximate the CTMC induced by a Stochastic Reaction Network as a *continuous-space* stochastic process by means of the *Central Limit Approximation (CLA)* [29], also known in statistical physics as the *Linear Noise Approximation (LNA)*. That is, a Gaussian process is derived to approximate the original CTMC [54]. As the marginals of a Gaussian process are fully determined by its expectation and covariances, its solution requires solving a number of differential equations that is quadratic in the number of species and independent of the population size. As a consequence, the CLA is generally much more scalable than a discrete-state stochastic representation and has been successfully used for analysis of large Stochastic Reaction Networks [18, 21, 23, 43]. However, none of these works enables the computation of complex temporal properties such as global *probabilistic reachability* properties, which quantify the probability of reaching a particular region of the state space in a particular time interval. This property is fundamental for verification of more complex temporal logic properties, for example *probabilistic until* properties, where the probability of reaching a certain region within a certain time bound while remaining in another region is quantified. Such properties can be expressed in Continuous Stochastic Logic (CSL) [6] or Linear Temporal Logic (LTL) [49], whose formulae are verified by reduction to the computation of the reachability properties [10].

1.0.1 Contributions. We derive fast and scalable approximate probabilistic model checking algorithms for CTMCs induced by Stochastic Reaction Networks against a time-bounded fragment of CSL extended with reward operators. Our model checking algorithms are numerical and explore a continuous-space approximation of the CTMC in terms of a Gaussian process. One of our key results is a novel scalable algorithm for computing probabilistic reachability for Gaussian processes over target regions of the state space that are assumed to be convex polytopes, i.e. intersections of a finite set of linear inequalities. More specifically, for a CTMC approximated as a Gaussian process, the resulting algorithm computes the probability that the system falls in the target region within a specified time interval. Given a set of k linear inequalities, and relying on the fact that a linear combination of the components of a Gaussian distribution is still Gaussian, we discretize time and space for the k -dimensional stochastic process defined by the particular linear combinations. This allows us to derive an abstraction in terms of a time-inhomogeneous *discrete-time Markov chain (DTMC)*, whose dimension is independent of the number of species, since a linear combination is a uni-dimensional entity. The method ensures scalability, as in general we are interested in a small number, i.e., one or at most two, of linear inequalities. This abstraction is then used to perform model checking of time-bounded CSL properties [9, 40]. To compute such an abstraction, the most delicate aspect is to derive equations for the transition kernel of the resulting DTMC. This is formulated as the conditional probability at the next discrete time step given the system in a particular state. Reachability probabilities are then computed by making the target set absorbing. We then extend CSL with the reward operators as in [40]. We derive approximate reward measures for such operators using the CLA, and prove the convergence in distribution of our algorithms to the original measures when the size of the system (number of molecules) tends to infinity.

We show the effectiveness of our approach on a set of case studies taken from the biological literature, also in cases where existing numerical model checking techniques are infeasible.

A preliminary version of this work has appeared in [14]. This paper extends [14] in several aspects. While in [14] we only consider probabilistic reachability, here we generalise our algorithms to the time-bounded fragment of CSL, which we also extend with reward operators. Furthermore, we prove weak convergence of our algorithms and significantly extend the experimental evaluation.

1.0.2 Related work. Algorithms for model checking CSL properties for continuous-time Markov chains have been introduced and then improved with techniques based on uniformization [8] (essentially a discretisation of the original CTMC) and reward computation [40]. The analysis typically involves computing the transient probability of the system residing in a state at a given time, or, for a model annotated with rewards, the expected reward that can be obtained. Despite improvements such as symmetry reduction [35], sliding window [56] and fast adaptive uniformisation [28], their practical use for Stochastic Reaction Networks is severely hindered by state space explosion [35], which in a SRN grows exponentially with the number of molecules when finite, and may be infinite, in which case finite projection methods have to be used [47]. As a consequence, approximate but faster algorithms are appealing. The mainstream solution is to rely on simulations combined with statistical inference to obtain estimates [20, 41]. These methods, however, are still computationally expensive. A recent trend of works explored as an alternative whether estimates could be obtained by relying on approximations of the stochastic process based on mean-field [15] or linear noise [18, 19, 23]. However, CSL and some classes of reward properties, like those considered here, are very challenging. In fact, most approaches consider either local properties of individual molecules [15], or properties obtained by observing the behaviour of individual molecules and restricting the target region to an absorbing subspace of the (modified) model [18]. The only approach dealing with more general subsets, [19], imposes restrictions on the behaviour of the mean-field approximation, whose trajectory has to enter the reachability region in a finite time. Another interesting approach has been developed in [46, 51], where model checking of time-bounded properties for CTMCs is expressed as a Bayesian inference problem, and approximate model checking algorithms are derived. However, no guarantees on the convergence of the resulting algorithms are given. Recent works also considered approximations of the CTMC induced by a SRN in terms of a *stochastic hybrid system* [13, 24, 34, 36]. The idea is to approximate the species in high population as a continuous-space process, while keeping the subset of species in low counts as a discrete process. Although this approach can be effective and capture multimodal dynamics, it has convergence guarantees only in terms of stochastic hybrid systems with deterministic continuous dynamics [13], and there is no convergence guarantee when the continuous dynamics is expressed by Gaussian processes [23]. Moreover, methods based on moment closure [34] do not have any convergence or error guarantees at all. In addition, model checking of stochastic hybrid systems against CSL properties requires discretization of the continuous state space, and are thus constrained by state space explosion [42].


Our approach differs in that it is based on the CLA and considers regions defined by polytopes, which encompasses most properties of practical interest. The simplest idea would be to consider the CLA and compute reachability probabilities for this stochastic process, invoking convergence theorems for the CLA to prove the asymptotic correctness. Unfortunately, there is no straightforward way to do this, since dealing with a continuous space and continuous time diffusion process, e.g., Gaussian, is computationally hard, and computing reachability is challenging (see [1]). As a consequence, discrete abstractions are appealing.

2 BACKGROUND


Stochastic Reaction Networks. Stochastic Reaction Networks are a versatile formalism used to model the stochastic evolution of populations of indistinguishable agents, where the species represent the states of the agents. SRNs have found applications in areas ranging from biological systems to sensor networks as a model of

distributed systems with asynchronous interactions [5, 25], and are a stochastic continuous-time generalization of Petri nets [48], Vector Addition Systems (VAS) [38] and distributed population protocols [5]. A *Stochastic Reaction Network (SRN)* $C = (\Lambda, R)$ is a pair of finite sets, where Λ is a set of *species* or *agents*, $|\Lambda|$ denotes its size, and R is a set of reactions. Species $\lambda \in \Lambda$ interact according to the reactions in R . A *reaction* $\tau \in R$ is a triple $\tau = (r_\tau, p_\tau, \alpha_\tau)$, where $r_\tau \in \mathbb{N}^{|\Lambda|}$ is the *reactant complex*, $p_\tau \in \mathbb{N}^{|\Lambda|}$ is the *product complex* and $\alpha_\tau : \mathbb{R}_{\geq 0}^{|\Lambda|} \rightarrow \mathbb{R}_{\geq 0}$ is the *reaction rate* associated to τ . r_τ and p_τ represent the stoichiometry of reactants and products. Given a reaction $\tau_1 = ([1, 1, 0]^T, [0, 0, 2]^T, \alpha_1)$, where \cdot^T is the transpose of a vector, we often refer to it as $\tau_1 : \lambda_1 + \lambda_2 \xrightarrow{\alpha_1} 2\lambda_3$. The *state change* associated to a reaction τ is defined by $v_\tau = p_\tau - r_\tau$. For example, for τ_1 as above, we have $v_{\tau_1} = [-1, -1, 2]^T$. A *configuration* or *state* $x \in \mathbb{N}^{|\Lambda|}$ of the system is given by a vector of the number of molecules of each species. Given a configuration x then x_{λ_i} represents the number of molecules of λ_i in the configuration and $\hat{x}_{\lambda_i} = \frac{x_{\lambda_i}}{N}$ is a real number describing the *concentration* or *density* of λ_i in the same configuration, where N is the *system size*, i.e. roughly a measure that describes how large the populations of interacting agents are. The notion of system size is specific to the system under consideration: for molecular systems, N typically represents the volume of the solution, while in many population models with constant population N is the total population size. In this latter case, \hat{x}_{λ_i} is typically known as an *occupancy measure*. In this paper, however, we will refer to \hat{x}_{λ_i} as a concentration, as is common in the literature on reaction networks [29].

Example 2.1. As a running example we consider the following simple model of gene expression [52], where the mRNA is produced by an always active promoter, and then catalyzes the production of the protein. We have $\Lambda = \{mRNA, Pro\}$ and the following set of reactions R :



Example 2.2. As a further illustrative example we consider the *Approximate Majority algorithm*, which provides the asymptotically fastest way to reach a common decision by all members of a population between two possible outcomes, where the decision approximately matches the initial relative majority [5, 22]. Consider two species (agents) A and B with initial population A_0 and B_0 , then the Approximate Majority algorithm can be represented by the following reactions:



where J is an auxiliary species. An intuitive description of the process is that agents in state J are undecided, while agents in states A and B are attempting to convert each other. However, at first they can just convert an opposing agent to an undecided state J . Under the above reactions, any initial state with zero agents in state J converges rapidly to a uniform population of size $A_0 + B_0$ corresponding with high probability to the initial majority of A_0 vs. B_0 (or to a random uniform population if $A_0 = B_0$) [5].

2.1 Stochastic Semantics of Stochastic Reaction Networks

Under the well-mixed assumption [3], a Stochastic Reaction Network $C = (\Lambda, R)$ induces a *discrete-state* Markov process. For a reaction τ , α_τ is also called the *propensity rate* of reaction τ and is a function of the current configuration x of the system, such that $\alpha_\tau(x)dt$ is the probability that a reaction event occurs in the next time interval dt . For instance, in case of mass action kinetics, $\alpha_\tau(x) = k_\tau \frac{\prod_{i=1}^{|\Lambda|} r_{i,\tau}!}{N^{|r_\tau|-1}} \prod_{i=1}^{|\Lambda|} \binom{x_{\lambda_i}}{r_{i,\tau}}$, where $r_{i,\tau}!$ is the factorial of $r_{i,\tau}$, $|r_\tau| = \sum_{i=1}^{|\Lambda|} r_{i,\tau}$, and x_{λ_i} is the component of vector x relative to species λ_i [4]. In this paper we assume $\alpha_\tau : \mathbb{R}_{\geq 0}^{|\Lambda|} \rightarrow \mathbb{R}_{\geq 0}$ is a real analytic function [15], that is, a function that locally coincides with its Taylor expansion. This is not restrictive, as it includes all the more commonly used kinetics such as mass action or Hill. We also

require that the SRN satisfies the *density dependent rate* condition¹, that is, for any α_τ , there exists a function $\beta_\tau : \mathbb{R}_{\geq 0}^{|\Lambda|} \rightarrow \mathbb{R}_{\geq 0}$ such that for $x \in \mathbb{R}_{\geq 0}^{|\Lambda|}$ it holds that $\alpha_\tau(x) = N\beta_\tau(\hat{x})$, where $\hat{x} = \frac{x}{N}$ represents the concentration of the species in Λ in configuration x . Consequently, a SRN $C = (\Lambda, R)$ is modelled in terms of a *time-homogeneous continuous-time Markov chain* (CTMC) [29] $(X^N(t), t \in \mathbb{R}_{\geq 0})$ with state space S given by the set of possible configurations of the system, where in X^N we made explicit the dependence on the system size N . Thus, $X^N(t)$ is a random vector describing the population count of each species at time t . Given X^N , we denote by $\hat{X}^N = \frac{X^N}{N}$ the CTMC describing the evolution of the species in Λ in terms of concentrations. The transient evolution of X^N , and consequently also of the concentrations \hat{X}^N , is described by the Kolmogorov equations, also called the Chemical Master Equation (CME), namely, a set of differential equations describing the transient evolution of the reachable states x .

Definition 2.3. (Kolmogorov Equations) Let $x_0 \in \mathbb{N}^{|\Lambda|}$ be the initial configuration of X^N . For $x \in S$, we define $P(x, t|x_0) = \text{Probability}(X^N(t) = x | X^N(0) = x_0)$. $P(x, t|x_0)$ describes the transient evolution of X^N , and is the solution of the following system of ordinary differential equations (ODEs):

$$\frac{d}{dt} (P(x, t|x_0)) = \sum_{\tau \in R} \{ \alpha_\tau(x - v_\tau) P(x - v_\tau, t|x_0) - \alpha_\tau(x) P(x, t|x_0) \}. \quad (1)$$

Solving Eqn (1) requires computing the solution of a differential equation for each reachable state. The size of the reachable state space is exponential in the number of the species, and may be infinite. As a consequence, solving the CME is generally feasible only for SRNs with very few species and small molecular populations. This is the so-called state space explosion problem, which strongly limits the applicability of the CME in practice. Finite projection methods have been developed to numerically solve Eqn (1) when the state space is not finite [47]. However, they still suffer from the state space explosion problem and are limited to SRNs with few species and moderate population counts.

Often, Eqn (1) is approximated with a deterministic model using fluid techniques [15], where the concentrations of the species are approximated over time as the solution $\Phi(t)$ of the following set of ODEs, the so-called *rate equations*:

$$\frac{d\Phi(t)}{dt} = F(\Phi(t)) = \sum_{\tau \in R} v_\tau \cdot \beta_\tau(\Phi(t)), \quad (2)$$

where in case of mass action kinetics we have $\beta_\tau(\Phi(t)) = (k_\tau \prod_{i=1}^{|\Lambda|} \Phi_i^{r_{i,\tau}}(t))$, for $\Phi_i^{r_{i,\tau}}(t)$ the i -th component of vector $\Phi(t)$ raised to the power of $r_{i,\tau}$, i -th component of vector r_τ . The initial condition is $\Phi(0) = \frac{x_0}{N} = \hat{x}_0$. Eqn (2) converges to $\hat{X}^N(t)$, $t \in \mathbb{R}_{\geq 0}$ when N , the system size, tends to infinity [29]. However, Eqn (2) completely neglects the stochastic fluctuations, which may be essential to understand the behaviour of the system being modelled [23].

Example 2.4. Consider the SRN introduced in Example 2.1. Then, for $t \in \mathbb{R}_{\geq 0}$, we have that $X^N(t) = [X_{mRNA}^N(t), X_{Pro}^N(t)]$ is a random variable describing the number of molecules in the system at time t . Given an initial condition $x_0 \in \mathbb{N}_{\geq 0}^2$, S , the state space of X^N is given by the set of states reachable from x_0 . That is, for any $x \in S$ there is a sequence of reactions $\tau_1, \dots, \tau_n \in R$ such that $x = x_0 + v_{\tau_1} + \dots + v_{\tau_n}$. Note that the presence of the reaction $\rightarrow^{0.5} mRNA$ implies that, in this example, S is not finite. Thus, most of the techniques commonly used for model checking CTMCs would not be directly applicable in this case [40]. $\hat{X}^N(t) = [\hat{X}_{mRNA}^N(t), \hat{X}_{Pro}^N(t)] = [\frac{X_{mRNA}^N(t)}{N}, \frac{X_{Pro}^N(t)}{N}]$ describes the evolution of mRNA and Pro in terms of concentrations.

¹Note that this condition is not strictly necessary for our results, but guarantees a simpler form for equations [29].

2.2 Central Limit Approximation

The *Central Limit Approximation (CLA)*, also called the *Linear Noise Approximation (LNA)*, is a *continuous-space* approximation of the CTMC in terms of a Gaussian process based on the Central Limit theorem [29, 54].

The CLA at time t approximates the distribution of $X^N(t)$ with the distribution of the random vector $Y^N(t)$ such that:

$$X^N(t) \approx Y^N(t) = N\Phi(t) + N^{\frac{1}{2}}G(t) \quad (3)$$

where $G(t) = (G_1(t), G_2(t), \dots, G_{|\Lambda|}(t))$ is a random vector, independent of the system size N , representing the stochastic fluctuations at time t around $\Phi(t)$, the solution of Eqn (2). The probability distribution of $G(t)$ is given by the solution of a linear Fokker-Planck equation [55]. As a consequence, for any time instant t , $G(t)$ has a multivariate normal distribution whose expected value $\mathbb{E}[G(t)]$ and covariance matrix $\text{cov}(G(t))$ are the solution of the following differential equations:

$$\frac{d\mathbb{E}[G(t)]}{dt} = J_F(\Phi(t))\mathbb{E}[G(t)] \quad (4)$$

$$\frac{d\text{cov}(G(t))}{dt} = J_F(\Phi(t))\text{cov}(G(t)) + \text{cov}(G(t))J_F^T(\Phi(t)) + W(\Phi(t)) \quad (5)$$

where $J_F(\Phi(t))$ is the Jacobian of $F(\Phi(t))$, $J_F^T(\Phi(t))$ its transpose, $W(\Phi(t)) = \sum_{\tau \in R} v_\tau v_\tau^T \alpha_{c,\tau}(\Phi(t))$ and $F_j(\Phi(t))$ the j th component of $F(\Phi(t))$. We assume $X^N(0) = x_0$ with probability 1; as a consequence $\mathbb{E}[G(0)] = 0$ and $\text{C}[G(0)] = 0$, which implies $\mathbb{E}[G(t)] = 0$ for every t . The following theorem illustrates the nature of the approximation using the CLA.

THEOREM 2.5 ([29]). *Let $C = (\Lambda, R)$ be a SRN, X^N the discrete state space Markov process induced by C and $\hat{X}^N = \frac{X^N}{N}$. Let $\Phi(t)$ be the solution of Eqn (2) with initial condition $\Phi(0) = \hat{x}$ and G be the Gaussian process with expected value and variance given by Eqns (4) and (5). Then, for any $t \in \mathbb{R}_{\geq 0}$ we have:*

$$N^{\frac{1}{2}} |\hat{X}^N(t) - \Phi(t)| \Rightarrow_{N \rightarrow \infty} G(t). \quad (6)$$

In the above, $\Rightarrow_{N \rightarrow \infty}$ indicates convergence in distribution as the system size parameter N increases [12]. The CLA is exact in the limit of high populations, but has also been successfully used in many different scenarios showing surprisingly good results [33, 55]. To compute the CLA it is necessary to solve $O(|\Lambda|^2)$ first order differential equations, and the complexity is independent of the initial number of molecules of each species. Therefore, one can avoid the exploration of the state space that methods based on uniformization rely upon, taking an important step towards scalable stochastic analysis of reaction systems.

By Eqn (3), we have that the distribution of $Y^N(t)$ is Gaussian with expected value and covariance matrix given by:

$$\mathbb{E}[Y^N(t)] = N\Phi(t)$$

$$\text{cov}(Y^N(t)) = N^{\frac{1}{2}} \text{cov}(G(t)) N^{\frac{1}{2}} = N \text{cov}(G(t)).$$

Then, the following standard proposition guarantees that a set of linear combinations of the components of Y^N is still Gaussian.

PROPOSITION 2.6 ([2]). *Let $B \in \mathbb{Z}^{m \times |\Lambda|}$ be a matrix and Y^N a $|\Lambda|$ -dimensional Gaussian process. Then, $Z^N = B \cdot Y^N$ is a m -dimensional Gaussian process. For any $t \in \mathbb{R}_{\geq 0}$, we have that $Z^N(t)$ is characterized by the following mean and covariance:*

$$\mathbb{E}[Z^N(t)] = B\mathbb{E}[Y^N(t)] \quad (7)$$

$$\text{cov}(Z^N(t)) = B \text{cov}(Y^N(t)) B^T. \quad (8)$$

Example 2.7. Consider the SRN introduced in example 2.1. According to Theorem 2.5 we can associate to C a Gaussian process $Y^N(t)$ with values in \mathbb{R}^2 . Suppose we want to know the distribution of $Z_{mRNA+Pro}^N(t) = Y_{mRNA}^N(t) + Y_{Pro}^N(t)$, where Y_{mRNA}^N and Y_{Pro}^N are the components of Y^N relative to *mRNA* and *Pro*. Then, we have that $Z_{mRNA+Pro}^N(t)$ is still Gaussian and with mean and variance given by

$$E[Z_{mRNA+Pro}^N(t)] = E[Y_{mRNA}^N(t)] + E[Y_{Pro}^N(t)] \quad cov(Z_{mRNA+Pro}^N(t)) = [1, 1]cov(Y^N(t))[1, 1]^T.$$

Thus, Z^N represents the time evolution of m linear combinations of the population counts of the species defined by B over time. Importantly, Z^N is still a Gaussian process, and hence completely characterized by its mean and covariance matrix. Note also that the distribution of $\hat{Z}^N = \frac{Z^N}{N}$ (concentrations) depends on Y^N *only via its mean and covariance*, which are obtained by solving ODEs in Eqns (4) and (5). This is a key feature that we will use to obtain an effective dimensionality reduction in our model checking algorithms.

3 CONTINUOUS STOCHASTIC LOGIC (CSL)

Temporal properties of continuous time Markov chains can be expressed using *Continuous Time Stochastic Logic (CSL)* [7], which can thus be used for the CTMC X^N induced from a SRN $C = (\Lambda, R)$. We will develop approximate model checking algorithms for CSL based on the CLA. Since the CLA is correct in the limit of diverging system size N , we will define CSL for the *normalized* process $\hat{X}^N = \frac{X^N}{N}$, as introduced in the previous section. Therefore, we will be working in terms of concentrations instead of population counts. This is not a limitation: if we are interested in a fixed value of N , population counts can always be rescaled to population densities, and vice versa, by dividing or multiplying them by N . In the following, we will thus refer to states and concentrations interchangeably without loss of generality.

Given a SRN $C = (\Lambda, R)$, a *path* of the induced CTMC \hat{X}^N is defined as $\omega = \hat{x}_0 t_0 \hat{x}_1 t_1 \dots$ where $\hat{x}_k \in \mathbb{R}_{\geq 0}^{|\Lambda|}$, $t_k \in \mathbb{R}_{\geq 0}$ and for all $k \geq 0$ there exists $\tau \in R$ such that $\beta_\tau(\hat{x}_k) > 0$ and $\hat{x}_k + \frac{\nu_\tau}{N} = \hat{x}_{k+1}$, where β_τ is the density dependent rate. That is, ω is an alternating sequence of states (equivalently, concentrations) and residence times in those states. Let Ω be the set of all paths of \hat{X}^N and $\Omega_{\hat{x}_0}$ the set of all paths of \hat{X}^N starting from \hat{x}_0 . Call $\omega(t)$ the state of the path at time t , i.e. $\omega(t) = \hat{x}_n$ where $\sum_{k=0}^n t_k \leq t \leq \sum_{k=0}^{n+1} t_k$. Then, a probability measure, Prob , for \hat{X}^N can be defined using cylinder sets of paths [40]. For further details on the measure-theoretic properties we refer to [9].

Since \hat{X}^N takes values in $\mathbb{R}_{\geq 0}^{|\Lambda|}$, we will work with predicates over concentrations, similarly to how real-time signals are verified in *Signal Temporal Logic (STL)* [45], instead of the conventional atomic propositions defined in states of the Markov chain [40].

Definition 3.1. (Convex Predicate). Let $\eta : \mathbb{R}^{|\Lambda|} \rightarrow \{\text{true}, \text{false}\}$ be a predicate. We call η a *convex predicate* if there exist $B_1, \dots, B_m \in \mathbb{Z}^{|\Lambda|}$, $l_1, \dots, l_m \in \mathbb{R}$, $m > 0$, such that for $\hat{x} \in \mathbb{R}^{|\Lambda|}$ it holds that:

$$\eta(\hat{x}) = (B_1 \cdot \hat{x} \leq l_1) \wedge \dots \wedge (B_m \cdot \hat{x} \leq l_m)$$

Hence, convex predicates are true for concentration \hat{x} belonging to a, not necessarily bounded, convex polytope. We denote by Θ the set of all convex predicates with domain in $\mathbb{R}_{\geq 0}^{|\Lambda|}$.

We now define the time-bounded fragment of CSL for SRNs as follows. We do not consider time-unbounded properties because of the nature of the convergence of CLA, which is guaranteed just for finite time. In Section 7 we extend this fragment with reward operators.

Definition 3.2. (CSL Syntax) Given a SRN $C = (\Lambda, R)$, and the induced CTMC \hat{X}^N , we define the syntax of CSL as:

$$\Psi ::= \neg \Psi \mid \Psi_1 \wedge \Psi_2 \mid P_{\sim p}(F^{[t_1, t_2]} \eta) \mid P_{\sim p}(\eta_1 U^{[t_1, t_2]} \eta_2)$$

where $\eta, \eta_1, \eta_2 \in \Theta$, $t_1, t_2 \in \mathbb{R}_{\geq 0}$, $\in [0, 1]$ and $\sim \in \{<, >\}$.

The above definition slightly differs from the usual definition of CSL in that the reachability ($F^{[t_1, t_2]}$) and until ($U^{[t_1, t_2]}$) operators work directly with predicates over concentrations, rather state labels. Note also that, in Definition 3.1, we do not allow nesting of CSL properties, and we restrict predicates to sets that are convex polytopes. This latter point does not limit the expressivity of the logic. However, it is a fundamental requirement for our model checking algorithms, which allows us to obtain an exponential speed up compared to existing algorithms.

Example 3.3. Given the SRN C of Example 2.1 for $N = 100$, the property "is the probability that the concentration of Pro remains below 0.1 until there is a concentration of mRNA of at least 0.3 in the first 50 time units greater than 0.6?" can be expressed as:

$$P_{>0.6}[(\hat{Pro} < 0.1) U^{[0, 50]} (\hat{mRNA} > 0.3)],$$

where with an abuse of notation we call \hat{Pro} and \hat{mRNA} the components of vector \hat{X}^N relative to species Pro and $mRNA$. Obviously, this property is equivalent to the following one, but expressed on the rescaled process X^N :

$$P_{>0.6}[(Pro < 10) U^{[0, 50]} (mRNA > 30)].$$

Definition 3.4. (Semantics of CSL) Let \hat{X}^N be the CTMC induced by SRN C . Given $\hat{x} \in \mathbb{R}_{\geq 0}^{|\Lambda|}$, the semantics of CSL is defined as follows:

$$\begin{aligned} \hat{X}^N, \hat{x} \models \neg \Psi &\leftrightarrow \hat{X}^N, \hat{x} \not\models \Psi \\ \hat{X}^N, \hat{x} \models \Psi_1 \wedge \Psi_2 &\leftrightarrow \hat{X}^N \models \Psi_1 \wedge \hat{X}^N \models \Psi_2 \\ \hat{X}^N, \hat{x} \models P_{\sim p}(F^{[t_1, t_2]} \eta) &\leftrightarrow \text{Prob}(\exists t \in [t_1, t_2] \text{ s.t. } \eta(\omega(t)) \mid \omega \in \Omega_{\hat{x}}) \sim p \\ \hat{X}^N, \hat{x} \models P_{\sim p}(\eta_1 U^{[t_1, t_2]} \eta_2) &\leftrightarrow \text{Prob}(\exists t \in [t_1, t_2] \text{ s.t. } \eta_2(\omega(t)) \wedge \forall t' \in [0, t] \eta_1(\omega(t')) \mid \omega \in \Omega_{\hat{x}}) \sim p \end{aligned}$$

Note that the reachability operator can be expressed with the until. For example, $P_{>0.9}[F^{[0, 1]} mRNA > 0]$ is equivalent to $P_{>0.9}[mRNA \geq 0 U^{[0, 1]} mRNA > 0]$. Similarly to classical CSL, \sim can be replaced with $=?$, in the style of quantitative model checking, indicating the probability of satisfaction [37].

Model checking procedures for CTMCs against CSL specifications are well known [10, 40]. They reduce to computing the probability of reaching a given set, and hence to solving Eqn (1), albeit resulting in the well known state space explosion problem. Here, we explore the usage of the CLA to derive approximate model checking procedures that converge to the original CTMC values but do not suffer from the state space explosion problem, therefore enabling fast stochastic characterization of the system.

4 THE REACHABILITY OPERATOR

In this section we define our CLA-based algorithm to verify the probabilistic reachability operator $P_{\sim p}(F^{[t_1, t_2]} \eta)$, which is the key procedure for model checking of more complex CSL properties. As η is a convex predicate, in order to check this property, for a convex polytope A defined as $A = \{x \in \mathbb{R}_{\geq 0}^{|\Lambda|} \text{ s.t. } \forall i \in \{1, \dots, m\} (Bx)_i \leq b_i\}$ where $B \in \mathbb{Z}^{m \times |\Lambda|}$, $b \in \mathbb{R}^m$, we need to compute:

$$P_{reach}^A(\hat{x}_0, t_1, t_2) = \text{Prob}(\exists t \in [t_1, t_2] \text{ s.t. } \omega(t) \in A \mid \omega \in \Omega_{\hat{x}_0}),$$

where $\Omega_{\hat{x}_0}$ is the set of paths of \hat{X}^N starting from \hat{x}_0 as defined in Section 3. We will compute such a probability for $\hat{Y}^N = \frac{\hat{X}^N}{N}$, the CLA of X^N expressed in terms of concentrations, and then show how the computed measure converges to the original process \hat{X}^N , but guaranteeing much greater scalability. Computing the reachability probability for \hat{Y}^N is not straightforward, because the system evolves in continuous time and analytic solutions cannot be derived in general. As a consequence, we need to devise numerical algorithms and prove their correctness. Here, we derive a scalable numerical algorithm based on time and space discretization of linear

projections of \hat{Y}^N , and, using properties of Gaussian processes, we then prove the convergence of the algorithm to the original measure.

In order to exploit the CLA, we first discretize time for the Gaussian process given by the CLA, with a fixed (or adaptive) step size h , which we can do effectively owing to the Markov property and the knowledge of its mean and covariance. As a result, we obtain a *discrete-time, continuous-space*, Markov process with a Gaussian transition kernel. Then, by resorting to state space discretization with parameter $\Delta z > 0$, we compute the reachability probability on this new process, obtaining an approximation in terms of time-inhomogeneous discrete-time Markov chain (DTMC) converging to the CLA approximation uniformly, when h and Δz go to 0. At first sight, there seems to be little gain, as we now have to deal with a $|\Lambda|$ -dimensional continuous state space. Indeed, for general regions this can be the case. However, if we restrict to regions defined by intersections of linear inequalities (i.e. polytopes), we can exploit properties of Gaussian distributions (i.e. their closure with respect to linear combinations), reducing the dimension of the continuous space to the number of different linear combinations used in the definition of the linear inequalities (in fact, the same hyperplane can be used to fix both an upper and a lower bound). As we are generally interested only in one or few projections, the complexity will then be dramatically reduced.

4.1 Time Discretization Scheme

Given \hat{Y}^N , the CLA of \hat{X}^N expressed in terms of concentrations, and matrix $B \in \mathbb{Z}^{m \times |\Lambda|}$, we introduce an exact time discretization scheme for $\hat{Z}^N = B\hat{Y}^N$. For simplicity we assume $m = 1$, but all the results extend to $m > 1$. Fix a small time step $h > 0$. By sampling \hat{Y}^N at step h and invoking the Markov property,² we obtain a *discrete-time Markov process* (DTMP) $\hat{Y}^{h,N}(k) = \hat{Y}^N(kh)$ on continuous space. Applying the linear projection mapping \hat{Z}^N to $\hat{Y}^N(k)$, and leveraging its Gaussian nature, we obtain a process $\hat{Z}^{h,N}(k) = \hat{Z}^N(kh)$ which is also a DTMP, though with a kernel depending on time through the mean and variance of Y^N .

Definition 4.1. A (time-inhomogeneous) *discrete-time Markov process* (DTMP) $(\hat{Z}^{h,N}(k), k \in [0, I] \subseteq \mathbb{N})$ is uniquely defined by a triple $(S, \mathcal{B}(S), \mathcal{T})$, where $(S, \mathcal{B}(S))$ is a measurable space and $\mathcal{T} : \mathcal{B}(S) \times S \times \mathbb{N} \rightarrow [0, 1]$ is a transition kernel such that, for any $z \in S$, $A \in \mathcal{B}(S)$ and $k \in \mathbb{N}$, $\mathcal{T}(A, z, k)$ is the probability that $\hat{Z}^{h,N}(k+1) \in A$ conditioned on $\hat{Z}^{h,N}(k) = z$.

From Definition 4.1, it follows that, for $[0, I] \subseteq \mathbb{N}$, $\hat{Z}^{h,N}$ is a discrete-time stochastic process defined on the sample space given by the product space $\Omega = S^{I+1}$, endowed with the sigma-algebra, $\mathcal{B}(\Omega)$, generated by the product topology, and with a probability measure $Prob^h$, which is uniquely defined by the transition kernel \mathcal{T} and the initial condition [11].

Thus, in order to characterize $\hat{Z}^{h,N}$, we need to compute its transition kernel, \mathcal{T} . This is equivalent to computing $f_{\hat{Z}^N(t+h)|\hat{Z}^N(t)=\bar{z}}(z)$, i.e. the density function of $\hat{Z}^N(t+h)$ given the event $\hat{Z}^N(t) = \bar{z}$.

Consider the joint distribution $(\hat{Y}^N(t), \hat{Y}^N(t+h))$, which is Gaussian. Its projected counterpart $(\hat{Z}^N(t), \hat{Z}^N(t+h))$ is thus also Gaussian, with covariance function:

$$cov(\hat{Z}^N(t), \hat{Z}^N(t+h)) = B cov(\hat{Y}^N(t), \hat{Y}^N(t+h)) B^T = \frac{1}{N} B cov(Y^N(t), Y^N(t+h)) B^T,$$

where $cov(Y^N(t), Y^N(t+h))$ is the covariance function of Y^N at times t and $t+h$. It follows by the closure properties of Gaussian processes that $(\hat{Z}^N(t+h)|\hat{Z}^N(t) = \bar{z})$ is Gaussian too, and thus fully characterized by its mean and variance. Hence, we need to derive $cov(Y^N(t), Y^N(t+h))$. From now on, we denote $cov(Y^N(t+h), Y^N(t)) = C_{Y^N}(t+h, t)$ and $cov(\hat{Z}^N(t+h), \hat{Z}^N(t)) = C_{\hat{Z}^N}(t+h, t)$. Following [29], we introduce the following matrix

²The Gaussian process obtained by the Linear Noise Approximation is Markovian, as it is the solution of a linear Fokker-Planck equation (stochastic differential equation) [54].

differential equation:

$$\frac{dY(t, s)}{dt} = J_F(\Phi(t))Y(t, s) \quad (9)$$

with $t \geq s$ and initial condition $Y(s, s) = Id$, where Id is the identity matrix of dimension $|\Lambda|$. Then, as illustrated in [29], we have:

$$C_{YN}(t, t+h) = \int_0^t Y(t, s)W(\Phi(s))[Y(t+h, s)]^T ds, \quad (10)$$

where W is the matrix introduced in Eqn (5). This is an integral equation, which has to be computed numerically. To simplify this task, we derive an equivalent representation in terms of differential equations. This is given by the following lemma.

LEMMA 4.2. *Solution of Eqn (10) is given by the solution of the following differential equations:*

$$\frac{dC_{YN}(t, t+h)}{dt} = W(\Phi(t))\Psi^T(t+h, t) + J_F(\Phi(t))C_{YN}(t, t+h) + C_{YN}(t, t+h)J_F^T(\Phi(t+h)) \quad (11)$$

with initial condition $C_{YN}(0, h)$ computed as the solution of:

$$\frac{dC_{YN}(0, s)}{ds} = C_{YN}(0, 0+s)J_F^T(\Phi(s)).$$

PROOF. Applying the general form of the Fundamental Theorem of Calculus to Eqn (10) with respect to t we get:

$$\begin{aligned} \frac{dC_{YN}(t, t+h)}{dt} &= Y(t, t)W(\Phi(t))Y(t+h, t)^T + \int_0^t \frac{d}{dt}(Y(s, t)W(\Phi(s))Y(t+h, s)^T) ds \\ &= Id \cdot W(\Phi(t))Y(t+h, t)^T + \int_0^t \frac{dY(s, t)}{dt}W(\Phi(s))Y(t+h, s)^T ds + \int_0^t Y(s, t)W(\Phi(s))\frac{dY(t+h, s)}{dt} ds. \end{aligned}$$

As $\frac{dY(t, s)}{dt} = J_F(\Phi(t))Y(t, s)$, we get

$$\begin{aligned} \frac{dC_{YN}(t, t+h)}{dt} &= \\ &W(\Phi(t))Y(t+h, t)^T + J_F(\Phi(t)) \int_0^t Y(t, s)W(\Phi(s))Y(t+h, s)^T ds + \int_0^t Y(s, t)W(\Phi(s))Y(t+h, s)^T ds J_F^T(\Phi(t+h))^T. \end{aligned}$$

By substituting Eqn (10) we have the result. Similarly, to derive the initial condition $C_{YN}(0, h)$ we can apply the Fundamental Theorem of Calculus to Eqn (10), but with respect to h . \square

$Y(t+h, t)$ can be computed by solving Eqn (9). Knowledge of $C_{YN}(t, t+h)$ allows us to directly compute:

$$C_{\hat{Z}^N}(t, t+h) = \frac{1}{N}B C_{YN}(t, t+h) B^T.$$

Then, by using the law for conditional expectation of a Gaussian distribution, we finally have:

$$\mathbb{E}[\hat{Z}^N(t+h)|\hat{Z}^N(t) = \bar{z}] = \mathbb{E}[\hat{Z}^N(t+h)] + C_{\hat{Z}^N}(\hat{Z}^N(t+h), Z(t))C[\hat{Z}^N(t)]^{-1}(\bar{z} - \mathbb{E}[\hat{Z}^N(t)]) \quad (12)$$

$$C[\hat{Z}^N(t+h)|\hat{Z}^N(t) = \bar{z}] = C[\hat{Z}^N(t+h)] - C_{\hat{Z}^N}(t, t+h)C_{\hat{Z}^N}(t, t)^{-1}C_{\hat{Z}^N}(t, t+h). \quad (13)$$

As the kernel is Gaussian, it is completely determined by its expectation and covariance matrix over time. Note that the resulting kernel is time-inhomogeneous. The dependence on time is via the mean and covariance of Y^N , which are functions of time and define completely the distribution of Y^N . The following result, which is a corollary of Theorem 3 in [42], guarantees the correctness of the approximation.

THEOREM 4.3. Given vector $B \in \mathbb{Z}^{|\Lambda|}$, $b \in \mathbb{R}$, measurable set $A = \{x \in \mathbb{R}^{|\Lambda|} \mid Bx \leq b\}$ and process $\hat{Z}^N = B\hat{Y}^N$ with initial condition $z_0 = B\hat{x}_0 \in \mathbb{R}$, call

$$P_{reach}^{\hat{Y}^N, A}(\hat{x}_0, t_1, t_2) = \text{Prob}^{\hat{Y}^N}(\exists t \in [t_1, t_2] \text{ s.t. } \hat{Y}^N(t) \in A \mid \hat{Y}^N(0) = \hat{x}_0),$$

where $\text{Prob}^{\hat{Y}^N}$ is the Gaussian probability measure of the process \hat{Y}^N . Further, let $\hat{Z}^{h, N}$ be the DTMP obtained by discretizing \hat{Z}^N at time step $h > 0$. Then, for $t_1, t_2 \in \mathbb{R}_{\geq 0}$, it holds that

$$|P_{reach}^{\hat{Y}^N, A}(\hat{x}_0, t_1, t_2) - \text{Prob}^h(\exists k \in [\lfloor \frac{t_1}{h} \rfloor, \lceil \frac{t_2}{h} \rceil] \text{ s.t. } \hat{Z}^{h, N}(k) \leq b)| \rightarrow_{h \rightarrow 0} 0,$$

uniformly.

4.2 Space Discretization

In order to compute the reachability probability for the DTMP $\hat{Z}^{h, N}$, we discretize its continuous state space into a countable set of non-overlapping cells (regions) of constant size $\Delta z > 0$ (except for at most regions of measure 0, i.e. the boundaries of the cells), obtaining an abstraction in terms of a discrete-time Markov chain $\hat{Z}^{\Delta z, h, N}$ with state space $S^{\Delta z}$. Specifically, given S , the state space of $\hat{Z}^{h, N}$, $A = \{x \in \mathbb{R}^{|\Lambda|} \mid Bx \leq b\}$ the target set for $B \in \mathbb{R}^{|\Lambda|}$, $b \in \mathbb{R}$, we divide $S \setminus A$ into a grid of cells of length $2\Delta z$, where Δz defines how fine our space discretization is. For each of the resulting regions we consider a representative point, given by the median of the set. We call the set of representative points $\hat{S}^{\Delta z}$. Then, we have $S^{\Delta z} = \hat{S}^{\Delta z} \cup \{z_d^A\}$, where z_d^A is the state representing the target set. Theorem 4.4 guarantees that for $\Delta z \rightarrow 0$ the error introduced by the space discretization tends to zero. However, for a fixed N , a possible choice of Δz is $\Delta z = \frac{0.5}{N}$, which means that the rescaled process $N\hat{Z}^{\Delta z, h, N}$ takes values in \mathbb{Z} . Nevertheless, when the population is of the order of hundreds or thousands, it can be beneficial to consider $\Delta z > \frac{0.5}{N}$, since a coarser state space aggregation is reasonable.

Similarly to the previous section (see Definition 4.1), as $\hat{Z}^{\Delta z, h, N}$ is a discrete-time stochastic process, given $[0, I] \subseteq \mathbb{N}$ we can associate to $\hat{Z}^{\Delta z, h, N}$ a probability space with sample space given by the product space $(S^{\Delta z})^{I+1}$, and with a probability measure $\text{Prob}^{\Delta z, h}$ uniquely defined by $\mathcal{T}^{\Delta z}$, the transition kernel of $\hat{Z}^{\Delta z, h, N}$, which is defined as follows. For $z'_d, z_d \in \hat{S}^{\Delta z}$, $\mathcal{T}^{\Delta z}(z'_d, z_d, k)$ is defined as:

$$\mathcal{T}^{\Delta z}(z'_d, z_d, k) = \int_{z'_d - \Delta z}^{z'_d + \Delta z} f_{\hat{Z}^N(hk+h) | \hat{Z}^N(hk)=z_d}(x) dx, \quad (14)$$

where h is the discrete time step, assumed to be fixed to simplify the notation. For $z_d \in \hat{S}^{\Delta z}$, we have:

$$\mathcal{T}^{\Delta z}(z_d^A, z_d, k) = \int_A f_{\hat{Z}^N(hk+h) | \hat{Z}^N(hk)=z_d}(x) dz, \quad (15)$$

and for the last case, we have:

$$\mathcal{T}^{\Delta z}(z_d, z_d^A, k) = \begin{cases} 1 & \text{if } z_d = z_d^A \\ 0 & \text{otherwise} \end{cases}.$$

That is, z_d^A is made absorbing. Finally, we define:

$$P_{reach}^{\Delta z, h, A}(z_d, t_1, t_2) = \text{Prob}^{\Delta z, h}(\exists k \in [\lfloor \frac{t_1}{h} \rfloor, \lceil \frac{t_2}{h} \rceil] \text{ s.t. } \hat{Z}^{\Delta z, h, N}(k) \in z_d^A \mid \hat{Z}^{\Delta z, h, N}(0) = z_d).$$

The following theorem, which is a corollary of Theorem 2 in [1], guarantees that the error introduced by the state space approximation tends to zero, decreasing Δz .

THEOREM 4.4. Let $\hat{Z}^{h,N}$ be a DTMP, and $\hat{Z}^{\Delta z, h, N}$ the DTMC obtained by space discretization of $\hat{Z}^{h,N}$ with space discretization step $\Delta z > 0$. Call z_0 the initial state of $\hat{Z}^{h,N}$ and $z_{d,0} \in S^{\Delta z}$ the discrete state representing the region containing z_0 . Then, for $t_1, t_2 \in \mathbb{R}_{\geq 0}$, and measurable set $A \subseteq \mathbb{R}$,

$$|\text{Prob}^h(\exists k \in [\lfloor \frac{t_1}{h} \rfloor, \lceil \frac{t_2}{h} \rceil] \text{ s.t. } \hat{Z}^{h,N}(k) \in A | \hat{Z}^{h,N}(0) = z_0) - P_{reach}^{\Delta z, h, A}(z_{d,0}, t_1, t_2)| \rightarrow_{\Delta z} 0,$$

uniformly.

4.3 Correctness

To prove the correctness of our numerical algorithm we need to show that, for any measurable set, the reachability measure computed on \hat{X}^N converges to that computed on \hat{Y}^N . This is guaranteed by the following theorem.

THEOREM 4.5. Let $C = (\Lambda, R)$ be a SRN with induced CTMC \hat{X}^N and $\hat{Z}^{\Delta z, h, N}$ be the DTMC obtained by space and time discretization of $B\hat{Y}^N$. Assume $\hat{X}^N(0) = \hat{x}_0$ and the corresponding initial state for $\hat{Z}^{\Delta z, h, N}$ is $z_{d,0}$. Then, for $t_1, t_2 \in \mathbb{R}_{\geq 0}$, $B \in \mathbb{R}^{m \times |\Lambda|}$ and $b \in \mathbb{R}^m$ and $A = \{x \in \mathbb{R}_{\geq 0}^{|\Lambda|} \text{ s.t. } \forall i \in \{1, \dots, m\} (Bx)_i \leq b_i\}$, it holds that:

$$\lim_{N \rightarrow \infty} \lim_{h \rightarrow 0} \lim_{\Delta z \rightarrow 0} |P_{reach}^A(\hat{x}_0, t_1, t_2) - P_{reach}^{\Delta z, h, A}(z_{d,0}, t_1, t_2)| = 0.$$

The proof of Theorem 4.5 is detailed in the Appendix. The main idea is to use Theorems 4.4 and 4.3 to show that the numerical model checking algorithms on the Gaussian process \hat{Y}^N are sound. Then, we employ Theorem 2.5 and the theory of weak convergence to show the convergence in distribution of the reachability measure on \hat{X}^N to that on \hat{Y}^N . The proof is complicated by the fact that both \hat{Y}^N and \hat{X}^N depend on N .

4.4 Computation of Reachability Probabilities

Our approach for computing reachability probabilities is summarized in Algorithm 1.

In Line 1, we initialize the system at time 0. In the context of the algorithm, \mathcal{S} is a set containing the states at a particular time with probability mass greater than the threshold Th . In our implementation we partition \mathbb{R} with a grid of cells of constant size $\Delta z > 0$. Then, for each cell we select a representative point, and denote the set of representative points $\mathcal{P}_{\Delta z}$. \mathcal{S} , for any time $t > 0$, will be a subset of this set. Th equals 10^{-14} in all our experiments. The use of a threshold guarantees that the algorithm always terminates in finite time. This introduces a truncation error, which can be easily estimated as shown in [56]. Initially, we have that \mathcal{S} contains only one state $B \cdot \hat{x}_0$. Then, in Lines 3 – 7, we propagate the probability for any discrete step while $t < t_1$, according to classical algorithms for DTMCs [40]. For generality, we assume that the time step h is chosen adaptively, according to the system dynamics. Propagating probability is possible, as for any $z'_d \in \mathcal{S}$, $\text{Prob}^{\Delta z, h}(\hat{Z}^{\Delta z, h, N}(k+1) = z'_d | \hat{Z}^{\Delta z, h, N}(k) = z_d) = \mathcal{T}^{\Delta z}(z'_d, z_d, k)$. From Line 8 to 15, we compute the probabilistic reachability, $P_{reach}^A(\hat{x}_0, t_1, t_2)$, by propagating the probability only for states that are not in A . In fact, states in A are made absorbing. When we reach $t \geq t_2$, we have that $P_{reach}^A(\hat{x}_0, t_1, t_2) \approx \sum_{z \in \mathcal{S} \cap I} \text{Prob}^{\Delta z, h}(\hat{Z}^{\Delta z, h, N}(\lceil \frac{t_2}{h} \rceil) = z | \hat{Z}^{\Delta z, h, N}(0) = z_{d,0})$.

Example 4.6. We return to the SRN introduced in Example 2.1, and, for $N = 100$, we consider the following reachability property:

$$P_{=?}(F^{[0, T]} \hat{mRNA} > \hat{Pro} + 0.2), T \in [0, 100]$$

where $=?$, in the style of PRISM [41] or PEPA [27], represents the quantitative value of a property. The above formula asks for the probability that, during the first 100 seconds, the system reaches a state where the mRNA concentration exceeds the protein concentration by more than 0.2. In Figure 1 we compare the probability value computed by Algorithm 1 with the same property computed on the CTMC \hat{X}^N using PRISM for different values of h . We assume $\Delta z = \frac{0.5}{N}$, which is justified by the fact that the number of molecules is an integer.

ALGORITHM 1: Compute Time-Bounded Probabilistic Reachability

Input: SRN $C = (\Lambda, R)$ with initial concentration \hat{x}_0 , $B \in \mathbb{Z}^{|\Lambda|}$, $I \subseteq \mathbb{R}$, a finite time interval $[t_1, t_2]$, a partition of the real numbers with the set of representative points $\mathbf{P}_{\Delta z}$, a target set $A = \{x \in \mathbb{R} \text{ s.t. } Bx \in I\}$ and a threshold \mathbf{Th} .

Output: $P_{reach}^A(\hat{x}_0, t_1, t_2)$.

Set $t = 0, k = 0, \mathcal{S} = \{B \cdot \hat{x}_0\}$ with $Prob^{\Delta z, h}(\hat{Z}^{\Delta z, h, N}(0) = B \cdot \hat{x}_0) = 1$;

while $t < t_1$ **do**

Compute time step h ;

$\mathcal{S} \leftarrow \{z_d \in \mathbf{P}_{\Delta z} \text{ s.t. } Prob^{\Delta z, h}(\hat{Z}^{\Delta z, h, N}(t+h) = z_d) \geq \mathbf{Th} \text{ where}$

$Prob^{\Delta z, h}(\hat{Z}^{\Delta z, h, N}(k+1) = z_d) = \sum_{z'_d \in \mathcal{S}} Prob^{\Delta z, h}(\hat{Z}^{\Delta z, h, N}(k+1) = z_d | \hat{Z}^{\Delta z, h, N}(k) = z'_d) Prob^{\Delta z, h}(\hat{Z}^{\Delta z, h, N}(k) = z'_d)\}$

$t \leftarrow t + h$;

$k \leftarrow k + 1$;

end

while $t < t_2$ **do**

Compute time step h ;

$\mathcal{S}' \leftarrow \mathcal{S} \setminus I$;

$\mathcal{S}_1 \leftarrow \{z_d \in \mathbf{P}_{\Delta z} \setminus I \text{ s.t. } Prob^{\Delta z, h}(\hat{Z}^{\Delta z, h, N}(k+1) = z_d) \geq \mathbf{Th}, \text{ where}$

$Prob^{\Delta z, h}(\hat{Z}^{\Delta z, h, N}(k+1) = z_d) = \sum_{z'_d \in \mathcal{S}'} Prob^{\Delta z, h}(\hat{Z}^{\Delta z, h, N}(k+1) = z_d | \hat{Z}^{\Delta z, h, N}(k) = z'_d) Prob^{\Delta z, h}(\hat{Z}^{\Delta z, h, N}(k) = z'_d)\}$

$\mathcal{S}_2 \leftarrow \{z_d \in \mathbf{P}_{\Delta z} \cap I \text{ s.t. } Prob^{\Delta z, h}(\hat{Z}^{\Delta z, h, N}(k+1) = z_d) \geq \mathbf{Th}, \text{ where}$

$Prob^{\Delta z, h}(\hat{Z}^{\Delta z, h, N}(k+1) = z_d) = Prob^{\Delta z, h}(\hat{Z}^{\Delta z, h, N}(k) = z_d) +$

$\sum_{z'_d \in \mathcal{S}'} Prob^{\Delta z, h}(\hat{Z}^{\Delta z, h, N}(k+1) = z_d | \hat{Z}^{\Delta z, h, N}(k) = z'_d) Prob^{\Delta z, h}(\hat{Z}^{\Delta z, h, N}(k) = z'_d)\}$

$\mathcal{S} \leftarrow \mathcal{S}_1 \cup \mathcal{S}_2$;

$t \leftarrow t + h$;

$k \leftarrow k + 1$;

end

return $P_{reach}^A(\hat{x}_0, t_1, t_2) = \sum_{z_d \in \mathcal{S} \cap I} Prob^{\Delta z, h}(\hat{Z}^{\Delta z, h, N}(k) = z_d)$;

5 UNTIL OPERATOR

We show how to generalize the computation of the reachability probabilities of the previous section to the until operator. For $\hat{x} \in \mathbb{R}_{\geq 0}^{|\Lambda|}$, let $\eta_1(\hat{x}) = B_1 \hat{x} \leq l_1$ and $\eta_2(\hat{x}) = B_2 \hat{x} \leq l_2$, then, by definition we have:

$$\hat{X}^N, \hat{x} \models P_{\sim p}(\eta_1 U^{[t_1, t_2]} \eta_2) \iff Prob(\exists t \in [t_1, t_2] \text{ s.t. } \eta_2(\omega(t)) \wedge \forall t' \in [0, t), \eta_1(\omega(t')) | \omega \in \Omega_{\hat{x}}),$$

where $\Omega_{\hat{x}}$ is the set of paths of \hat{X}^N starting in \hat{x} . To solve this problem we can construct the following stochastic process:

$$\hat{Z}^N = B \hat{Y}^N$$

where $B = (B_1, B_2)^T$, and \hat{Y}^N is the CLA of \hat{X}^N . By the properties of multivariate Gaussian distribution, \hat{Z}^N is still a Gaussian process with mean and covariance matrix given by

$$\mathbb{E}[\hat{Z}^N(t)] = B \mathbb{E}[Y^N(t)] \quad C_{\hat{Z}^N}(t) = \frac{1}{N} B C_{Y^N}(t) B^T, \quad t \in \mathbb{R}_{\geq 0}.$$




Fig. 1. Comparison of the evaluation of $P_{=?}(F^{[0,T]}mRNA > Pro + 20), T \in [0, 100]$, on the CTMC as estimated by PRISM [41], and on the CLA approximation for a fixed Δz and four different values of h .

Note that \hat{Z}^N is again a time-inhomogeneous Markov process, as its kernel depends on the statistics of Y^N . Following the approach of the previous section, we can discretize time and space for \hat{Z}^N , and thus obtain a DTMC $\hat{Z}^{\Delta z, h, N}$. At this point, the problem reduces to computing the probability for until on the DTMC. Algorithms for computing the resulting measure on a time-inhomogeneous DTMC exist and are well studied [26]. In fact, to compute $P_{-p}(\eta_1 U^{[t_1, t_2]} \eta_2)$, we can simply make the regions that do not satisfy η_1 and those for which η_2 holds absorbing, and then compute the probability of reaching a region for which η_2 is satisfied. This can be computed by resorting on Algorithm 1, as presented in the previous section. Theorem 4.5 then guarantees the following proposition.

PROPOSITION 5.1. *Let $\eta_1(\hat{x}) = B_1 \hat{x} \sim l_1$, $\eta_2(\hat{x}) = B_2 \hat{x} \sim l_2$, and $B = \begin{bmatrix} B_1 \\ B_2 \end{bmatrix}$. For $\hat{x}_0 \in \mathbb{R}_{\geq 0}^{|\Lambda|}$, let $z_{d,0}$ be the state in the state space of $Z^{\Delta z, h, N}$ corresponding to the region containing $B\hat{x}_0$. Call*

$$P_{until}((\hat{x}_0, \eta_1, \eta_2, \hat{X}_1^N, t_1, t_2)) = \text{Prob}(\exists t \in [t_1, t_2] \text{ s.t. } \eta_2(\omega(t)) \wedge \forall t' \in [0, t), \eta_1(\omega(t')) \mid \omega \in \Omega_{\hat{x}_0}),$$

$$P_{until}^{\Delta z, h}((z_{d,0}, \eta_1, \eta_2, \hat{Z}^{\Delta z, h, N}, t_1, t_2)) =$$

$$\text{Prob}^{\Delta z, h}(\exists k \in [\lfloor \frac{t_1}{h} \rfloor, \lfloor \frac{t_2}{h} \rfloor] \text{ s.t. } \eta_2(Z^{\Delta z, h, N}(k)) \wedge \forall k' \in [0, k-1], \eta_1(Z^{\Delta z, h, N}(k')) \mid Z^{\Delta z, h, N}(0) = z_{d,0}).$$

Then, it holds that

$$\lim_{N \rightarrow \infty} \lim_{h \rightarrow 0} \lim_{\Delta z \rightarrow 0} |P_{until}((\hat{x}_0, \eta_1, \eta_2, \hat{X}_1^N, [t_1, t_2])) - P_{until}^{\Delta z, h}((z_{d,0}, \eta_1, \eta_2, \hat{Z}^{\Delta z, h, N}, [t_1, t_2]))| = 0.$$




Fig. 2. Comparison of the evaluation of $P_{[0,T]}[(Pro < 10) U^{[0,T]}(mRNA > 30)]$ on a CTMC as estimated by PRISM [41], and on the CLA approximation for $\Delta z = \frac{0.5}{N}$ and three different values of h .

Example 5.2. Let us return to the SRN introduced in Example 2.1. We consider the following quantitative property:

$$P_{=?}[(Pro < 10) U^{[0,T]}(mRNA > 30)], T \in [0, 100],$$

which is satisfied for those paths in which the mRNA population becomes greater than 30 before the protein population hits 10 molecules. In Figure 2 we evaluate the property for different values of h and fixed $N = 100$. Already for $h = 5$ the property is surprisingly close to the same measure computed on the original CTMC using uniformization techniques as implemented in PRISM [41]. Note that here the property is expressed in terms of number of molecules. In fact, as we explained in Section 3 for the CSL properties considered in here the two representations are equivalent.

6 CORRECTNESS

The method we present is approximate. In particular, errors are introduced in two ways: by resorting to the CLA and by discretisation of time and space of the CLA. The quality of these two approximations is controlled by three parameters: (a) N , the system size, which influences the accuracy of CLA, (b) h , the time step size, and (c) Δz , the space discretization step, which influences the quality of the approximation of the reachability probability of the CLA. Then, we have the following result.

THEOREM 6.1. *Let Ψ be a CSL formula as defined in Definition 3.2, $\hat{x} \in \mathbb{R}_{\geq 0}^{|\Lambda|}$ and $z_{0,d}$ be the state in $Z^{\Delta z, h, N}$ corresponding to the region containing \hat{x}_0 . Then, for $N \rightarrow \infty, h \rightarrow 0, \Delta z \rightarrow 0$, it holds that:*

$$\hat{X}^N, \hat{x} \models \Psi \leftrightarrow \hat{Z}^{\Delta z, h, N}, z_{d_0} \models \Psi,$$

except for a set of thresholds of Lebesgue measure 0.

PROOF. The proof is by induction on the terms in Definition 3.2. The interesting cases are $\Psi = P_{\sim p}(F^{[t_1, t_2]}\eta)$ and $\Psi = P_{\sim p}(\eta_1 U^{[t_1, t_2]}\eta_2)$. Theorem 4.5 guarantees that, for $N \rightarrow \infty, h \rightarrow 0, \Delta z \rightarrow 0$, the difference between the probability of the above properties computed on \hat{Y}^N , the CLA of \hat{X}^N , and on \hat{X}^N is equal to $\epsilon \rightarrow 0$. Assume $Prob(\exists t \in [t_1, t_2] \text{ s.t. } \eta(\omega(t)) | \omega \in \Omega_{\hat{x}}) = q$, and consider the reachability property $P_{\sim q}(F^{[t_1, t_2]}\eta)$. In this case, no

approximation algorithm can guarantee to give the right answer, as the threshold is exactly the value of the reachability property. However, the point q is a set of Lebesgue measure zero with respect to the set of all possible thresholds, which is a subset of the reals. Same reasoning can be applied to the until case. \square

The convergence stated in Theorem 6.1 means that, since N is fixed for a given SRN, then, even if we have control over h and Δz , the quality of the approximation depends on how well the CLA approximates the SRN. Error bounds would be a viable companion to estimate the committed error, and although these could be estimated for time and space discretization following the approaches in [1, 42], we are not aware of any explicit formulation of them for the convergence of the CLA. However, experimental results in Section 8 show that the error committed is generally limited also for moderately small N and for h relatively large compared to the time evolution of the particular CRN.

6.1 Complexity

Complexity of the method depends on the following: (a) the equations we need to solve, (b) the time step size h , and (c) the space discretization step Δz . Algorithm 1 requires solving Eqns (11) and (5), that is, a set of differential equations quadratic in the number of species. In fact, solving these equations requires computing J_F , Jacobian of F . However, the number of equations we need to solve is independent of the number of molecules in the system. This guarantees the scalability of our approach. An important point is that Eqn (11) requires solving Eqn (10) once for each sampling point of the numerical solution of Eqn (11). A possible way to avoid this is to consider approximate solutions to Eqn (10), which are accurate in the limit of $h \rightarrow 0$. However, to keep this approximation under control, h has to be chosen really small, slowing down the computation. Moreover, for any sample point, Eqn (10) is solved only for a small time interval (between t and $t + h$). As a consequence, in practice, the computational cost introduced in solving Eqn (10) is under control.

A smaller value of h implies that, for a given time interval, we have a greater number of discrete time steps, which can slow down the computation in some cases. The value of Δz determines the number of states of the resulting DTMC. However, we stress that we discretize $\hat{Z}^N(t)$, a uni-dimensional distribution (or m -dimensional in the case we have $m > 1$ linear inequalities). As a consequence, the number of reachable states with significant probability mass is generally limited and under control. Obviously, if the number of molecules is large and Δz extremely small, then this is detrimental to performance.

7 REWARDS

We extend CSL properties with reward operators as in [40]. As for probabilistic reachability, we will define the reward structure on the normalised process \hat{X}^N . Formally, we define the *state reward* function $\rho : \mathbb{R}^{|\Lambda|} \rightarrow \mathbb{R}$, which associates a real-valued reward with any point of the normalised state space of $\hat{X}^N(t)$, $t \in \mathbb{R}_{\geq 0}$. In this paper, we make a few assumptions about the regularity of ρ :

- ρ is bounded, i.e. there exists a constant $C > 0$ such that $\rho(\hat{x}) \leq C$ for each $\hat{x} \in \mathbb{R}^{|\Lambda|}$;
- ρ is Lipschitz continuous on $\mathbb{R}^{|\Lambda|}$, i.e. there is a constant L_ρ such that, for each $\hat{x}, \hat{x}' \in \mathbb{R}^{|\Lambda|}$, $|\rho(\hat{x}) - \rho(\hat{x}')| \leq L_\rho |\hat{x} - \hat{x}'|$.

These requirements are important to show the convergence of rewards computed on \hat{X}^N with the same measure but computed on the normalised CLA $\hat{Y}^N = \frac{Y^N}{N}$. Moreover, they do not limit the expressiveness of our framework: for a fixed N , we are interested only in the value of ρ at a finite number of points. Such a function can always be extended to a Lipschitz continuous one. Boundedness is also not problematic, as we can always assume an upper bound on a physically meaningful population size, meaning that we can restrict ourselves to a compact set and define ρ to be constant outside such a set.

Given a reward structure ρ , we consider the following three kinds of rewards.

- **Instantaneous Rewards** up to finite time T . $\rho_I(\hat{x}_0, \hat{X}^N, T)$ is the expectation of $\rho(\hat{X}^N(T))$, i.e., the expectation of the reward structure at time T over all the trajectories of \hat{X}^N that start from state \hat{x}_0 . More precisely, for $\Omega_{\hat{x}_0}$, the set of paths of \hat{X}^N starting from \hat{x}_0 :

$$\rho_I(\hat{x}_0, \hat{X}^N, T) = \sum_{x \in \mathbb{R}^{|\Lambda|}} \rho(\hat{x}) \text{Prob}(\omega(T) = \hat{x} | \omega \in \Omega_{\hat{x}_0}). \quad (16)$$

- **Cumulative Rewards** up to a finite time T . Given $\omega : \mathbb{R}_{\geq 0} \rightarrow \mathbb{N}^{|\Lambda|}$, a path of \hat{X}^N , the cumulative reward for ω is defined as:

$$\rho_C(\omega, T) = \int_0^T \rho(\omega(t)) dt = \sum_{i=1}^{|\text{jumps}(\omega)|} \rho(\omega(t_{i-1}))(t_i - t_{i-1}) + \quad (17)$$

$$\rho(\omega(T))(T - t_{|\text{jumps}(\omega, t)|}) \quad (18)$$

where $\text{jumps}(\omega, t) \subset \mathbb{R}_{\geq 0}$ is the set of time instants at which ω changes state between 0 and T . Then, we define:

$$\rho_C(\hat{x}_0, \hat{X}^N, T) = \mathbb{E}[\rho_C(\omega, T) | \omega(0) = \Omega_{\hat{x}_0}],$$

where the expectation is intended over the trajectories of \hat{X}^N starting from state \hat{x}_0

- **Bounded Reachability Rewards.** Given a target set $A \in \mathbb{R}^{|\Lambda|}$, for the normalized process \hat{X}^N , define $\rho_{\text{reach}}(\hat{x}_0, \hat{X}^N, A, T)$, the cumulative reward until we enter the target set A within time T . Formally, we can define $\rho_{\text{reach}}(\hat{x}_0, \hat{X}^N, A, T)$ as the cumulative reward until time T for the modified process \tilde{X}^N where all states in A are made absorbing, and where we consider the modified state rewards:

$$\bar{\rho}(\hat{x}) = \begin{cases} 0 & \text{if } \hat{x} \in A \\ \rho(\hat{x}) & \text{otherwise} \end{cases}$$

REMARK 1. Note that here ρ is a state reward, that is, a function that associates a real value with any given state of the process. An alternative reward structure could be based on the transition reward function [16], which can be used for checking how many times a given reaction fires up to a certain time. However, in the context of SRNs, such a quantity can be easily expressed with an additional species counting how many times a subset of the transitions fire. Then, instantaneous rewards can be used to "read" its value. For instance, in Example 2.1, if we want to count the number of times a mRNA molecule is produced, we can consider an additional species C and modify the SRN such that:



Then, for $\hat{x} \in \mathbb{R}_{\geq 0}^{|\Lambda|}$ we have $\rho(x) = \hat{x}_C$, where \hat{x}_C is the component of vector \hat{x} relative to species C , and $N\rho_I(\hat{x}_0, \hat{X}^N, T)$ will give the desired measure at time T for \hat{x}_0 , initial concentration of the species.³

7.1 Extending CSL with Rewards

In order to incorporate rewards in our framework, given a SRN $C = (\Lambda, R)$ and the induced CTMC X^N , we extend CSL with the following formulae, whose semantics will depend on the particular reward structure ρ :

$$R_{\sim r}[C_{\rho}^{[\leq T]}] \mid R_{\sim r}[I_{\rho}^{=T}] \mid R_{\sim r}[F_{\rho}^{\leq T} \eta]$$

where $\eta : \mathbb{R}^{|\Lambda|} \rightarrow \{\text{true}, \text{false}\}$ is a convex predicate over \hat{X}^N , $T \in \mathbb{R}_{\geq 0}$, $r \in \mathbb{R}_{\geq 0}$, and $\sim \in \{>, <\}$. For $\hat{x} \in \mathbb{R}_{\geq 0}^{|\Lambda|}$, the semantics of such formulae is as follows:

$$\hat{X}^N, \hat{x} \models R_{\sim r}[C_{\rho}^{[\leq T]}] \quad \text{iff} \quad \rho_C(\hat{x}, \hat{X}^N, T) \sim r$$

³The multiplication of ρ_I by N is needed to convert from the normalized process back to the integer population count.

$$\begin{aligned}\hat{X}^N, \hat{x} \models R_{\sim r}[I_{\rho}^{\leq T}] & \text{ iff } \rho_I(\hat{x}, \hat{X}^N, T) \sim r \\ \hat{X}^N, \hat{x} \models R_{\sim r}[F_{\rho}^{\leq T} \eta] & \text{ iff } \rho_{reach}(\hat{x}, \hat{X}^N, A, T) \sim r,\end{aligned}$$

where $A = \{\hat{x}' \in \mathbb{R}^{|\Lambda|} \text{ s.t. } \eta(\hat{x}')\}$.

7.2 Computing Expectation and Variance Using Reward Operators

Two of the most common statistics needed when studying stochastic processes are expectation and variance (or covariance) of some set of variables. Suppose we have a CTMC X^N with values in $\mathbb{R}^{|\Lambda|}$, and we want to compute expectation and variance of one of its components X_i^N at time t . Then, for $\hat{x} \in \mathbb{R}^{|\Lambda|}$, we define the following reward structures on the normalised process:

$$\rho^{size}(\hat{x}) = \begin{cases} \hat{x}_i & \text{if } \hat{x}_i < K \\ K & \text{if } \hat{x}_i \geq K \end{cases} \quad \rho^{size^2}(\hat{x}) = \begin{cases} \hat{x}_i^2 & \text{if } \hat{x}_i < K \\ K^2 & \text{if } \hat{x}_i \geq K, \end{cases}$$

where $K \in \mathbb{R}$ can be any real number, typically an upper bound on the physically admissible population size. For instance, for biochemical processes, we can choose K to be of the order of 10^{80} , the estimated number of atoms of the universe. Then, we have

$$\begin{aligned}\mathbb{E}[X_i^N(t)] &= N R_{=?}[I_{\rho^{size}}^{\leq t}] \\ cov[X_i^N(t)] &= N(R_{=?}[I_{\rho^{size^2}}^{\leq t}] - (R_{=?}[I_{\rho^{size}}^{\leq t}])^2).\end{aligned}$$

Since T is finite and K non-negative, the above equality holds for any SRN whose species count remains finite at least for a finite time interval. Note that, as rewards are defined for the normalised process, we need to rescale them back to population counts to compute variance and average of the non-normalised process.

7.3 Computing Rewards through Central Limit Approximation

Computing reward properties over \hat{X}^N is generally not possible because of the state space explosion problem. As a consequence, we compute such properties using \hat{Y}^N , the CLA of \hat{X}^N . We show that the values computed on \hat{Y}^N converge (weakly) to those computed on \hat{X}^N . We stress again how working in terms of the normalised processes is not a limitation. For instance, consider the reward for expectation. If we are interested in the expectation of population counts for a fixed N , we can either define the reward for \hat{x} in the normalised space as $\rho(x) = N\hat{x}$, for N fixed, or rescale the computed reward as done in the previous section.

7.3.1 Instantaneous Rewards. Given a reward structure ρ , instantaneous rewards can be computed on $\hat{Y}^N = \frac{Y^N}{N}$ as:

$$\rho_I^{CLA}(\hat{x}, \hat{Y}^N, t) \approx \mathbb{E}[\rho(\hat{Y}^N(t))] = \int_K \rho(x) \mathcal{N}(x | \mathbb{E}[\hat{Y}^N(t)], cov[\hat{Y}^N(t)]) dx,$$

where $\mathcal{N}(x | \mathbb{E}[\hat{Y}^N(t)], cov[\hat{Y}^N(t)])$ is the normal distribution with mean and covariance matrix respectively, $\mathbb{E}[\hat{Y}^N(t)], cov[\hat{Y}^N(t)]$ for $\hat{Y}^N(0) = \hat{x}$. Furthermore, $K \subseteq \mathbb{R}^{|\Lambda|}$ is a compact set in which we restrict integration for numerical purposes. The choice of K is such that the error we incur is bounded by a chosen tolerance level. The following proposition guarantees that $\rho_I^{CLA}(\hat{x}, \hat{Y}^N, t)$ converges to $\rho_I(\hat{x}, \hat{X}^N, T)$.

PROPOSITION 7.1. *Let $T \in \mathbb{R}_{\geq 0}$, then it holds that:*

$$\lim_{N \rightarrow \infty} \rho_I(\hat{x}, \hat{X}^N, T) = \lim_{N \rightarrow \infty} \rho_I^{CLA}(\hat{x}, \hat{Y}^N, t)$$

PROOF. We want to prove that, for fixed $T > 0$, $\mathbb{E}[\rho(\hat{X}^N)]$ converges to $\mathbb{E}[\rho(\hat{Y}^N(T))]$ as N tends to infinity. Using the triangular inequality, it holds that:

$$\begin{aligned} |\mathbb{E}[\rho(\hat{X}^N(T))] - \mathbb{E}[\rho(\hat{Y}^N(T))]| &\leq \\ |\mathbb{E}[\rho(\hat{X}^N)] - \mathbb{E}[\rho(\Phi(T))]| &+ |\mathbb{E}[\rho(\Phi(T))] - \mathbb{E}[\rho(\hat{Y}^N(T))]|, \end{aligned}$$

where $\rho(\Phi(T))$ is the reward ρ evaluated on the fluid limit $\Phi(T)$. Invoking the fluid approximation theorem [17], it holds that $\hat{X}^N(T) \Rightarrow_{N \rightarrow \infty} \Phi(T)$ (as convergence in probability implies weak convergence). Furthermore, $\hat{Y}^N(T) = \frac{G(T)}{\sqrt{N}} + \Phi(T) \Rightarrow_{N \rightarrow \infty} \Phi(T)$, as G is independent of N and it has a bounded covariance matrix for each T (which implies convergence in probability). Therefore, recalling that ρ is bounded and continuous by assumption, both terms on the right hand side of the triangular inequality converge to zero by virtue of the Portmanteau theorem [12] stating that, for any continuous and bounded functional f on \mathcal{D} , the space of $\mathbb{R}^{|\Lambda|}$ -valued Cadlag functions (i.e. right continuous functions with left limit) [12], it holds that $\mathbb{E}[f(X^N)] \rightarrow_{N \rightarrow \infty} \mathbb{E}[f(X)]$ whenever $X^N \Rightarrow X$. Thus, we can conclude:

$$\rho_I(\hat{x}, \hat{X}^N, T) \rightarrow_{N \rightarrow \infty} \rho_I^{CLA}(\hat{x}, \hat{Y}^N, T).$$

□

Example 7.2. We consider the SRN introduced in Example 2.1. We are interested in knowing the expectation and variance of $mRNA - P$ over time. This can be computed using the following reward structures:

$$\begin{aligned} \rho^{mRNA-P}(\hat{x}) &= \begin{cases} \hat{x}(mRNA) - \hat{x}(P) & \text{if } \hat{x}(mRNA) - \hat{x}(P) < 10^{80} \\ 10^{80} & \text{otherwise} \end{cases} \\ \rho^{(mRNA-P)^2}(\hat{x}) &= \begin{cases} (\hat{x}(mRNA) - \hat{x}(P))^2 & \text{if } (\hat{x}(mRNA) - \hat{x}(P))^2 < 10^{80} \\ 10^{80} & \text{otherwise} \end{cases}. \end{aligned}$$

Then, we have:

$$\begin{aligned} \mathbb{E}[X_{mRNA}^N(t) - X_P^N(t)] &= NR_{=?}[I_{\rho^{mRNA-P}}^{-t}], \quad t \in [0, 100], \\ Cov(X_{mRNA}^N(t) - X_P^N(t)) &= N(R_{=?}[I_{\rho^{(mRNA-P)^2}}^{-t}] - (R_{=?}[I_{\rho_{size}}^{-t}])^2), \quad t \in [0, 100], \end{aligned}$$

where the rewards are computed on the normalised process \hat{X}^N . The resulting expectation and variance is plotted in Figure 3. Note that, in this case, the resulting variance and expectation, as estimated by the CLA, are guaranteed to be exact for any N . This is because the SRN is linear [33]. It is easy to observe that our algorithms are exponentially faster than the computation of the same measures on the original CTMC, because of the continuous nature of the CLA.

7.3.2 Cumulative Rewards. Cumulative rewards can also be computed exploiting \hat{Y}^N , the CLA approximation of \hat{X}^N , as shown in the following proposition

PROPOSITION 7.3. *Let $T \in \mathbb{R}_{\geq 0}$. Then, $\rho_C^{CLA}(\hat{x}, \hat{Y}^N, T)$, the cumulative reward for \hat{Y}^N starting from $\hat{Y}^N = \hat{x}$, can be computed as follows*

$$\rho_C^{CLA}(\hat{x}, \hat{Y}^N, T) = \int_0^T \rho_I^{CLA}(\hat{x}, \hat{Y}^N, s) ds,$$

PROOF. Let $\omega : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^{|\Lambda|}$ be a path of \hat{Y}^N . Then, we have that

$$\rho_C^{CLA}(\hat{x}, \hat{Y}^N, T) = \mathbb{E}[\rho_C(\omega, T) | \omega(0) = \hat{x}] = \mathbb{E}\left[\int_0^T \rho(\omega(t)) dt \mid \omega(0) = \hat{x}\right].$$




Fig. 3. $N\rho_I(\hat{X}^N, T), T \in [0, 100]$ for reward structure $\rho^{\text{mRNA-P}}$.

Now, in order to conclude, we can apply Fubini's theorem [50], which allows us to compute a double integral using iterated integrals. Thus, switching the order of integration. Being both a probability measure and the Lebesgue measure over the reals σ -finite measures, a sufficient condition for application of Fubini's theorem is that $\mathbb{E}[\int_0^T |\rho(\omega(t))|dt]$ is finite. Owing to boundedness of ρ , there is an $M > 0$ such that, for all $x \in \mathbb{R}^{|\Lambda|}$, we have that $|\rho(x)| \leq M$. Thus, we have,

$$\mathbb{E} \left[\int_0^T |\rho(\omega(t))|dt \right] \leq \mathbb{E} \left[\int_0^T Mdt \right] = MT,$$

which is finite as T and M are both finite. \square

The following proposition, for $\hat{x} \in \mathbb{R}^{|\Lambda|}$, guarantees that $\rho_C^{CLA}(\hat{x}, \hat{Y}^N, T)$ converges to $\rho_C(\hat{x}, \hat{X}^N, T)$

PROPOSITION 7.4. *Let $T \in \mathbb{R}_{\geq 0}$, then it holds that*

$$\lim_{N \rightarrow \infty} \rho_C(\hat{x}, \hat{X}^N, T) = \lim_{N \rightarrow \infty} \rho_C^{CLA}(\hat{x}, \hat{Y}^N, T)$$

PROOF. For a path $\omega : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^{|\Lambda|}$, define the following functional $\mathcal{R}_C(T, \omega) = \int_0^T \rho(\omega(t))dt$, which is defined on \mathcal{D} , the space of $\mathbb{R}^{|\Lambda|}$ -valued Cadlag functions (i.e. right continuous functions with left limit) [12]. $\rho_C(\hat{x}, \hat{X}^N, T) = \mathbb{E}[\mathcal{R}_C(T, \omega)]$, where the expectation is taken over $\Omega_{\hat{x}}$, the paths of \hat{X}^N starting from \hat{x} . As T and ρ are bounded, for each ω , $\mathcal{R}_C(T, \omega)$ is bounded. It is also continuous, due to the continuity of ρ . Thus, we can apply same reasoning as in the proof of Proposition 7.1, applying Portmanteau theorem to conclude. \square

Example 7.5. We consider the SRN introduced in Example 2.1. We are interested in knowing the expected cumulative reward of *mRNA* – *P* to understand if during the time interval there have been, on average, more




Fig. 4. $N\rho_C(\hat{X}^N, T)$, $T \in [0, 500]$, for reward structure $\rho^{\text{mRNA-P}}$.

mRNA or more *P* molecules in the system. This can be computed using the reward structure $\rho^{\text{mRNA-P}}$ introduced in Example 7.2, and the following cumulative reward:

$$NR_{=?}[C_{\rho^{\text{mRNA-P}}}^{\leq T}], T \in [0, 500],$$

where $R_{=?}[C_{\rho^{\text{mRNA-P}}}^{\leq T}]$ is intended to be computed on \hat{X}^N . The resulting expectation and variance are plotted in Figure 4. We stress again how in this case, since the SRN is linear, the measure estimated by the CLA is exact for any N .

7.3.3 Bounded Reachability Rewards. Bounded reachability rewards can be computed efficiently on the CLA under a further assumption on the reward structure ρ . Specifically, for $\hat{x} \in \mathbb{R}_{\geq 0}^{|\Lambda|}$, consider the predicate $\eta(\hat{x}) = B\hat{x} < b$, $b \in (\mathbb{R} \cup \{\infty\})^m$, $m > 0$. We assume that the reward structure is defined on the projection of \hat{X}^N spanned by the columns of matrix defining the η predicate, namely $\rho : \mathbb{R}^m \rightarrow \mathbb{R}$ assigns a reward to each state of $B\hat{X}^N$. Consider the CSL reward property $R_{\sim r}[F_{\rho}^{\leq T}\eta]$. That is, $R_{\sim r}[F_{\rho}^{\leq T}\eta]$ is the bounded reachability reward until reaching a state in $A = \{\hat{x} \in \mathbb{R}^{|\Lambda|} \text{ s.t. } \forall i \in \{1, \dots, m\}, (B\hat{x})_i \geq b_i\}$ during the time interval $[0, T]$. Such a reward can be computed by exploring the approximation of the CLA in terms of the DTMC $\hat{Z}^{\Delta z, h, N}$, which is obtained by time and space discretization of the process $\hat{Z}^N = B\hat{X}^N$. We call $\rho_{\text{reach}}(\hat{x}_0, \hat{Z}^{\Delta z, h, N}, \lfloor \frac{T}{h} \rfloor, A)$ the bounded reachability reward computed on $\hat{Z}^{\Delta z, h, N}$ for a number of discrete steps $\lfloor \frac{T}{h} \rfloor$, where $h > 0$. Then $\rho_{\text{reach}}(\hat{x}_0, \hat{Z}^{\Delta z, h, N}, \lfloor \frac{T}{h} \rfloor, A)$ can be computed by considering the modified process $\hat{Z}^{\Delta z, h, N}$ where the target states are made absorbing, and modifying the reward structure ρ to $\bar{\rho}$ such that $\bar{\rho}(\hat{x}) = 0$ for all absorbing states. Then, for $\hat{x}_0 \in \mathbb{R}^{|\Lambda|}$ and $z_{d,0}$, the state in the state space if $\hat{Z}^{\Delta z, h, N}$ corresponding to the region containing \hat{x}_0 , cumulative rewards can be




Fig. 5. $\rho_{reach}(x_0, T)$, $T \in [0, 35]$, for reward structure ρ^{mRNA-P} estimated using 10000 simulations compared with the CLA approximation for different values of sampling time h . $\delta z = 0.5$ for all experiments.

computed with the following algorithm for $n > 0$:

$$\rho_{reach}(\hat{x}_0, \hat{Z}^{\Delta z, h, N}, n, A) = \sum_{z \in S^{\Delta z}} \bar{\rho}(z) \text{Prob}(\hat{Z}^{\Delta z, h, N}(n-1) = z | \hat{Z}^{\Delta z, h, N}(0) = z_{d,0}) h + \rho_{reach}(\hat{Z}^{\Delta z, h, N}, \hat{x}_0, n-1, A). \quad (19)$$

and such that $\rho_{reach}(\hat{Z}^{\Delta z, h, N}, \hat{x}_0, 0, A) = 0$ for $\hat{x}_0 \notin A$. The proof of the following proposition can be found in the Appendix.

PROPOSITION 7.6. *For $T \in \mathbb{R}_{\geq 0}$ and $B \in \mathbb{R}^{|\Lambda| \times k}$ let A be the set defined as $A = \{x \in \mathbb{R}^{|\Lambda|} \text{ s.t. } \forall i \in \{1, \dots, k\}, (Bx)_i \geq b_i\}$. Then, for $\hat{x}_0 \in \mathbb{R}^{|\Lambda|}$ and $z_{d,0}$, the state in the state space of $\hat{Z}^{\Delta z, h, N}$ corresponding to the region containing \hat{x}_0 , it holds that:*

$$\lim_{N \rightarrow \infty} \lim_{h \rightarrow 0} \lim_{\Delta z \rightarrow 0} |\rho_{reach}(\hat{x}_0, \hat{X}^N, T, A) - \rho_{reach}(z_{d,0}, \hat{Z}^{\Delta z, h, N}, \lfloor \frac{T}{h} \rfloor, A)| = 0.$$

Example 7.7. We consider the SRN introduced in Example 2.1. We are interested in knowing the expected cumulative reward of $mRNA - P$ for all those paths for which the $mRNA$ does not reach 30 individuals within $[0, T]$ for $T \in [0, 50]$. We consider the reward structure $\rho^{mRNA-P}(x)$ introduced in Example 7.2, and the following cumulative reward $\rho_{reach}(X^N, T, mRNA < 30) = N \rho_{reach}(\hat{X}^N, T, mRNA < \frac{30}{N}), T \in [0, 50]$. To compute such a reward we explore the CLA approximation of X^N . We consider $B_1 = [1, 0]$, $B_2 = [-1, 1]$ and $B = [B_1, B_2]$, where we assume the first component of X^N represents the number of protein molecules in that state. Then, we consider \hat{Z}^N , the projection of the CLA of \hat{X}^N over B , namely, $\hat{Z}^N = B\hat{Y}^N$. At this point we discretize \hat{Z}^N with sampling time $h > 0$ and a grid of cells of size $\delta z > 0$, and compute the above rewards using Eqn (19). The solution to Eqn

(19) is approximate, and errors are introduced by two factors: firstly, by the usage of the CLA approximation of X^N , and, secondly, by the discretization of the resulting Gaussian process. Thus, we compare our reward value with the value computed on X^N using 10000 simulations for each time point. The resulting values are plotted in Figure 5. To perform a further comparison we employ the following metrics, ϵ_{max}^{rel} and ϵ_{avg}^{rel} , defined as follows:

$$\epsilon_{max}^{rel} = \max_{T \in \Sigma_h} \frac{|Rew_T^{CLA} - Rew_T|}{|Rew|}, \quad \epsilon_{avg}^{rel} = \sum_{T \in \Sigma_h} \frac{|Rew_T^{CLA} - Rew_T|}{|Rew|} \frac{1}{|\Sigma_h|}$$

where Σ_h is the set of sampling points for sampling time h , $Rew_T^{CLA} = \rho_{reach}^{mRNA-P}(\hat{Z}^{\Delta z, h, N}, \hat{x}_0, \lfloor \frac{T}{h} \rfloor, A)$ and $Rew_T = \rho_{reach}^{mRNA-P}(\hat{X}, \hat{x}_0, T, A)$. The calculated metrics are summarised in the table below for three different values of h .

h	ϵ_{avg}^{rel}	ϵ_{max}^{rel}
5	1.5468	7.96
3	0.196	0.88
1.5	0.041	0.24

It is possible to observe how the two measures converge very fast. In fact, already for $h = 1.5$, which corresponds to 25 sampling times, the two measures have an average relative error of less than 0.041.

8 EXPERIMENTAL RESULTS

We implemented our algorithms in Matlab and evaluated them on two case studies. All the experiments were run on an Intel Dual Core i7 machine with 8 GB of RAM. The first case study is a Gene Expression Network as introduced in Example 2.1. We use this example to demonstrate that our approach is more powerful than existing approximate techniques. Specifically, we show how our CLA approach, based on a Gaussian process approximation, is able to correctly evaluate properties that methods based on Fluid Limit Approximation [15] cannot, while still guaranteeing comparable scalability. The second example is a Phosphorelay Network with 7 species. We use this example to show the trade-off between the different parameters and the molecular population. More precisely, we show that the accuracy of our approach increases as the number of molecules grows, but can still give fast and accurate results when the molecular population is relatively small. We validate our results by comparing our method with statistical model checking (SMC) as implemented in PRISM [41]. In fact, for both examples, exact numerical computation of the reachability probabilities using uniformisation techniques on the induced CTMC is not feasible because of state space explosion.

8.1 Gene Expression

We consider the following gene expression model, as introduced in Example 2.1 with initial counts of all the species equal to 0. We consider the property $P_{=?}(F^{[0, Time]}(mRNA \geq 175))$, which quantifies the probability that at least 175 *mRNA* molecules are produced during the first *Time* seconds, for $Time \in [0, 1000]$. This is a particularly difficult property because the trajectory of the mean-field of the model, and so the expected value of the CLA, does not enter the target region. As a consequence, approximate approaches introduced in [29] and [19], which are based on the hitting times of the mean-field model, fail and evaluate the probability as always equal to 0.

Conversely, our approach is able to correctly evaluate such a property. Figure 6 compares the value computed by our method with statistical model checking of the same property as implemented in PRISM over 30000 simulations for each time point and confidence interval 0.01. In Figure 6 we consider $h = 1.8$ and $\Delta z = 0.5$ and demonstrate that our approach is able to correctly estimate such a difficult property. Note that, as the mean-field does not enter the target region, for each time point the probability to enter the target region depends on a portion of the tail of the Gaussian given by the CLA. As a consequence, the accuracy of our results strictly depends on how well the CLA approximates the original CTMC, much more than for properties where the mean-field enters





Fig. 6. The figure plots $F_{=?}[mRNA \geq 174]_{[0, Time]}$ for $h = 1.85$ and $\Delta z = 0.5$. The x-axis represents the value of *Time* and the y-axis the quantitative value of the formula for that value of *Time*.

the target region. In the following table, we evaluate our results for two different values of h and $\Delta z = 0.5$. In order to compare the accuracy we consider the following metrics as defined in Example 7.7, ϵ_{avg}^{rel} and ϵ_{max}^{rel} . The comparison is shown in the table below.

Property	Ex. Time	h	ϵ_{avg}^{rel}	ϵ_{max}^{rel}
$P_{=?}(F_{[0, Time]}^{[0, Time]}(mRNA \geq 174)), Time \in [0, 100]$	298 sec	1.85	0.0075	0.022
$P_{=?}(F_{[0, Time]}^{[0, Time]}(mRNA \geq 174)), Time \in [0, 100]$	152 sec	5	0.0147	0.13

8.2 Phosphorelay Network

We now consider a three-layer phosphorelay network consisting of 7 species given by the following reactions:



There are 3 layers, $(L1, L2, L3)$, which can be found in phosphorylate form $(L1p, L2p, L3p)$, and the ligand B . We consider the initial condition $x_0 \in \mathbb{N}^7$ such that $x_0(L1) = x_0(L2) = x_0(L3) = 0.25N$, $x_0(L1p) = x_0(L2p) = x_0(L3p) = 0$ and $x_0(B) = 0.15N$. In Figure 7, we compare the estimates obtained by our approach for two different initial conditions ($N = 400$ and $N = 800$) with statistical model checking as implemented in PRISM [41], with 30000 simulations and confidence interval equal to 0.01. In both experiments we set $\Delta z = 0.5$.

In Figure 7a we can see that, if we increase the time interval of interest, the error tends to increase. This is because, for $N = 400$, the CLA and CME do not have perfect convergence. As a consequence, at every step of the discretized DTMC, a small error is introduced. This source of error is present until we enter the target region with probability 1. If we increase N the error disappears, and the inaccuracies are due to the finiteness of h and Δz . However, already for $h = 0.1$ and $N = 800$, the CLA produces a fast and reasonably accurate approximation. In the following table we compare our approach and PRISM evaluations for different values of N and h and $\Delta z = \frac{0.5}{N}$ in the normalised space, which implies the resulting discrete state process takes values in \mathbb{Z} .




Fig. 7. Comparison of the evaluation of $F_{[0,Time]}[L3p > 80]$ (a) with $N = 400$ and $F_{[0,Time]}[L3p > 180]$ (b) with $N = 800$ using statistical model checking as implemented in PRISM and our approach. In both figures we considered $h = 0.1$, $\Delta z = 0.5$.

Property	Ex. Time	h	N	ϵ_{avg}^{rel}	ϵ_{avg}^{max}
$P_{=?}(F_{[0,Time]}(L3p > 80)), Time \in [0, 10]$	97 sec	0.1	400	0.0088	0.11
$P_{=?}(F_{[0,Time]}(L3p > 180)), Time \in [0, 10]$	130 sec	0.1	800	0.0015	0.0217
$P_{=?}(F_{[0,Time]}(L3p > 80)), Time \in [0, 10]$	28 sec	0.5	400	0.0381	0.24
$P_{=?}(F_{[0,Time]}(L3p > 180)), Time \in [0, 10]$	39 sec	0.5	800	0.0289	0.14

The results show that the best accuracy is obtained for $h = 0.1$ and $N = 800$, where $h = 0.1$ induces a finer time discretization, whereas the worst is for $h = 0.5$ and $N = 400$. We comment that the numerical solution of the CME using PRISM is not feasible for this model, and our method is several orders of magnitude faster than statistical model checking with PRISM (30000 simulations for each time point), which takes several hours for each property.

9 CONCLUSION

We presented a framework for approximate model checking of a time-bounded fragment of CSL extended with rewards for CTMCs that are induced from Stochastic Reaction Networks. Our approach employs an abstraction based on the Central Limit Approximation to approximate the CTMC as a Gaussian process. Then, numerical procedures for model checking CSL formulae on the resulting Gaussian process are derived. We do not consider time-unbounded properties because of the nature of the convergence of CLA, which is guaranteed just for finite time. Since the CLA requires solving a number of differential equations that is quadratic in the number of species and independent of the population size, our methods enable formal analysis of possibly infinite-state CTMCs that cannot be analysed using classical methods based on uniformization [28, 56].

Deriving model checking algorithms was challenging because the CLA yields a continuous time stochastic process with an uncountable state space. As a consequence, existing methods that rely on finite state spaces cannot be used directly and discretizing the uncountable state space induced by the CLA naturally leads to state space explosion. In order to overcome these issues, we considered reachability regions defined as polytopes. Using the fact that the CLA is a Gaussian Markov process, for a given linear combination of the species of a SRN we are able to project the original, multi-dimensional Gaussian process onto a uni-dimensional stochastic process. We then derived an abstraction in terms of a time-inhomogeneous DTMC, whose state space is independent of the

number of the species of a SRN, as it is derived by discretizing linear combinations of the species. This ensures scalability. On different case studies, we showed that our approach outperforms existing methods and permits fast and scalable probabilistic analysis of SRNs. The accuracy depends on parameters controlling space and time discretization, as well as on the accuracy of the CLA. Using the theory of convergence in distribution we proved the convergence of our algorithms in the limit of high populations. As future work we plan to release a tool for scalable model checking of SRNs. Moreover, we wish to investigate the speed of convergence of our methods.

A PROOFS

Theorem 4.5 Let $C = (\Lambda, R)$ be a SRN with induced CTMC \hat{X}^N and $\hat{Z}^{\Delta z, h, N}$ be the DTMC obtained by space and time discretization of $B\hat{Y}^N$. Assume $\hat{X}^N(0) = \hat{x}_0$ and the corresponding initial state for $\hat{Z}^{\Delta z, h, N}$ is $z_{d,0}$. Then, for $t_1, t_2 \in \mathbb{R}_{\geq 0}$, and $A = \{x \in \mathbb{R}^{|\Lambda|} \text{ s.t. } \forall i \in \{1, \dots, m\} (Bx)_i \leq b_i\}$, for $B \in \mathbb{R}^{m \times |\Lambda|}$ and $b \in \mathbb{R}^m$, it holds that:

$$\lim_{N \rightarrow \infty} \lim_{h \rightarrow 0} \lim_{\Delta z \rightarrow 0} |P_{reach}^A(\hat{x}_0, t_1, t_2) - P_{reach}^{\Delta z, h, A}(z_{d,0}, t_1, t_2)| = 0.$$

Proof. Without any loss of generality, we assume $t_1 = 0, t_2 = T$. Call

$$P_{reach}^{h, A}(\hat{x}_0, 0, T) = Prob^h(\exists t \in [0, \lceil \frac{T}{h} \rceil] \text{ s.t. } \hat{Z}^{h, N}(k) \leq b \mid \hat{Z}^{h, N}(0) = B\hat{x}_0),$$

and

$$P_{reach}^{\hat{Y}^N, A}(\hat{x}_0, 0, T) = Prob^{\hat{Y}^N}(\exists t \in [0, T] \text{ s.t. } \hat{Y}^N(t) \in A \mid \hat{Y}^N(0) = \hat{x}_0),$$

where $Prob^{\hat{Y}^N}$ is the Gaussian probability measure of \hat{Y}^N .

By application of the triangular inequality we have that:

$$\begin{aligned} |P_{reach}^A(\hat{x}_0, 0, T) - P_{reach}^{\Delta z, h, A}(z_{d,0}, 0, T)| &\leq \\ &|P_{reach}^A(\hat{x}_0, 0, T) - P_{reach}^{\hat{Y}^N, A}(\hat{x}_0, 0, T)| + \\ &|P_{reach}^{\hat{Y}^N, A}(\hat{x}_0, 0, T) - P_{reach}^{h, A}(\hat{x}_0, 0, T)| + \\ &|P_{reach}^{h, A}(\hat{x}_0, 0, T) - P_{reach}^{\Delta z, h, A}(z_{d,0}, 0, T)|. \end{aligned}$$

The convergence of the third and second components is a consequence of Theorems 4.4 and 4.3. We need to show that:

$$\lim_{N \rightarrow \infty} |P_{reach}^A(\hat{x}_0, 0, T) - P_{reach}^{\hat{Y}^N, A}(\hat{x}_0, 0, T)| = 0.$$

Note that we removed the limits for Δz and h , as this term is independent of time and space discretization. In what follows we assume that BX^N is a uni-dimensional process. Generalization for $m > 1$ follows from this case. Intuitively, this holds due to the convergence of X^N to its CLA Y^N . A formal proof requires a more involved machinery. In fact, Theorem 6 states that:

$$\sqrt{N} \left(\hat{X}^N(t) - \Phi(t) \right) \Rightarrow G(t),$$

hence, to rely on it, we need to reason on the modified stochastic model:

$$G^N(t) = \sqrt{N} \left(\hat{X}_N(t) - \Phi(t) \right),$$

rather than on the original CTMC $\hat{X}^N(t)$. Now, consider the reachability problem $B\hat{X}^N \leq b$; rephrasing it in terms of G^N we get:

$$B\hat{X}^N(t) \leq b \text{ iff } BG^N(t) \leq \sqrt{N}(b - B\Phi(t)) = b^N(t).$$

As we can see, the reachability problem for G^N has a different nature: the threshold b becomes both N dependent and time dependent! In addition, we see that for the CLA, $BY^N(t) \leq b$ iff $BG(t) \leq b^N(t)$. Let's look at this reachability problem from the point of view of the trajectory space, i.e. the space of cadlag function $\omega : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$. Both G^N and G can be seen as probability measures over this space. The reachable set in the trajectory space depends on N , precisely being $R_N = \{\omega \mid \exists t \in [0, T] : \omega(t) \leq b^N(t)\}$. We also consider the complement of this set, $R_N^c = \{\omega \mid \forall t \in [0, T] : \omega(t) > b^N(t)\}$, and the boundary of the set $\partial R_N = \{\omega \mid \forall t \in [0, T] : \omega(t) \geq b^N(t) \wedge \exists t \in [0, T] : \omega(t) = b^N(t)\}$.

Before proceeding further, we need to understand how the set R_N changes as N goes to infinity. Consider the threshold $b^N(t) = \sqrt{N}(b - B\Phi(t))$. There are three cases:

- (1) if $b > B\Phi(t)$, then $b^N(t) \rightarrow +\infty$;
- (2) if $b < B\Phi(t)$, then $b^N(t) \rightarrow -\infty$;
- (3) if $b = B\Phi(t)$, then $b^N(t) = 0$.

In the first case, the reachable set at time t converges to \mathbb{R} , in the second case to the empty set, in the third case to $(-\infty, 0]$. Therefore, the limit reachable set R in the trajectory space will be the union for each t of one of these three kind of sets.

By the assumption that rate functions are real analytic, it follows that $\Phi(t)$ is also a real analytic function, and therefore $B\Phi(t)$ will equal b only in a finite number of points of $[0, T]$, or in the whole interval (a degenerate case which is easily tractable) [39]. It then follows that $b(t) = \lim_{N \rightarrow \infty} b^N(t)$ changes value a finite number of times, say at times t_1, \dots, t_n , where it equals zero. Outside these points, it is either plus or minus infinity. The reachable set R , in the limit of infinite N , is thus a finite union of sets of the form $t_i \times (-\infty, 0]$ at times t_i and either \emptyset or \mathbb{R} for each t in between t_{i-1} and t_i .

Now, if in such a union the set $(t_{i-1}, t_i) \times \mathbb{R}$ is present at least once, then the reachability probability in the limit equals exactly one. This is because any trajectory ω will enter the set R in that subregion. In this case, convergence is easily shown. In fact, being the Skorokhod space a Polish space, any converging sequence $G^N \Rightarrow G$ of random variables in that space is uniformly tight, meaning that for each ϵ there is a compact space K_ϵ such that, outside it, all random variables and the limit have probability less than ϵ . In particular, a compact set of trajectories is bounded in $[0, T]$ with respect to the sup norm [2], meaning that for each ϵ there is a $k_\epsilon > 0$ such that the probability that a trajectory ω has modulus $|\omega(t)| \leq k$ uniformly in $[0, T]$ is more than $1 - \epsilon$ for all N . Now, consider the time interval (t_{i-1}, t_i) where the reachable set converges to $(t_{i-1}, t_i) \times \mathbb{R}$ in the limit. As the threshold $b^N(t)$ is an analytic function of t , removing a region of length Δ near t_{i-1} and t_i (i.e. restricting to $[t_{i-1} + \Delta, t_i - \Delta]$), we can find an N_0 such that, for $N > N_0$, $b^N(t)$ is greater than k_ϵ uniformly in $[t_{i-1} + \Delta, t_i - \Delta]$. Then the limit of the reachability for G^N is greater than $1 - \epsilon$ for any epsilon, that is, it equals one. The case in which the limit region R is the empty set for every t is easily proved along the same lines.

The interesting case is the one in which there are some t_i 's where $b^N(t_i) = 0$ for all N , but it is always negative outside them, implying the reachable region R converges to the empty set everywhere but in the t_i 's, where it equals $(-\infty, 0]$. This corresponds to the scenario in which the fluid limit $\Phi(t)$ is tangent to the reachable set, but never enters it, a scenario known to cause trouble in the use of mean field to estimate hitting times [13].

To deal with this last case, let us denote with P^N the probability in the trajectory space for BG^N , and with P the probability for BG .

As before, denote with R_N the reachability set for G^N and with R the limit set, taking the threshold b^N to infinity. We now introduce a set which over-approximates R_N for N large. This set is defined as follows: invoking uniform tightness, we fix a large value k_ϵ as before, so that trajectories of BG^N and of BG are contained in

$[-k_\epsilon, k_\epsilon]$ with probability $1 - \epsilon$, uniformly for $t \in [0, T]$. Furthermore, we consider points t_i where $b^N(t_i)$ is zero, and take a small neighborhood B_i^Δ of width Δ around them. Define the set R_ϵ in the trajectory space as:

$$R_\epsilon = \{\omega(t) | \omega(t) \leq 0, \text{ for } t \in B_i^\Delta, \omega(t) \leq -k_\epsilon \text{ elsewhere in } [0, T]\}.$$

By relying on the continuity of the set R for G , we can choose Δ small enough so as to enforce that $|P(R_\epsilon) - P(R)| \leq \epsilon$. The continuity of R for G follows from the fact that $\omega \in R$ if and only if $\omega(t_i) \leq 0$ for $i = 1, \dots, n$, i.e. R is a finite dimensional projection on t_i 's. Therefore, its boundary is a set of topological dimension less than n in \mathbb{R}^n , which has probability zero under the finite dimensional projection of G on t_i 's (which is Gaussian). Now, using triangular inequality, we get:

$$\begin{aligned} |P^N(R_N) - P(R)| &\leq |P^N(R_N) - P^N(R)| + |P^N(R) - P(R)| \\ &\leq |P^N(R_\epsilon) - P^N(R)| + |P^N(R) - P(R)| \\ &\leq |P^N(R_\epsilon) - P(R_\epsilon)| \\ &\quad + |P(R) - P(R_\epsilon)| + 2|P^N(R) - P(R)|. \end{aligned}$$

The second inequality above follows from the monotonic behaviour of probability distributions, as for each Δ and k_ϵ there is an N_0 such that, for all $N \geq N_0$, $R \subset R_N \subset R_\epsilon$, hence $|P^N(R_N) - P^N(R)| \leq |P^N(R_\epsilon) - P^N(R)|$.

Furthermore, $|P^N(R) - P(R)| \rightarrow 0$, by the continuity of the set Y . In R , by virtue of Lemmas 1 and 2 below, it also follows that $|P^N(R_\epsilon) - P(R_\epsilon)| \rightarrow 0$, and hence:

$$\limsup_{N \rightarrow \infty} |P^N(R_N) - P(R)| \leq \epsilon,$$

which holds for any $\epsilon > 0$, allowing us to conclude that:

$$\lim_{N \rightarrow \infty} |P^N(R_N) - P(R)| = 0,$$

as desired.

Lemma 1. Let $b \in \mathbb{R}$, and consider the reachable set $R = \{\omega | \exists t : \omega(t) \leq b\}$. Then $P^N(R) \rightarrow P(R)$, with P^N, P as above.

Proof. The boundary of the reachable set R is the set of trajectories ω such that $\inf_{t \in [0, T]} \omega(t) = b$. In order to conclude, we need to show that this set has measure 0. As G is a Gaussian process, assuming the covariance function is non-zero, we have that the distribution of the infimum (or equivalently the supremum) is absolutely continuous [44], which implies that the set of trajectories for which $\inf_{t \in [0, T]} \omega(t) = b$ has measure 0. Hence, R is a continuity set for G , which prove the thesis due to the Portmanteau theorem.

Lemma 2. Consider a reachable set R defined by a piecewise constant threshold. Hence, fix $0 = t_1, \dots, t_{n+1} = T \in [0, T]$, and $b_i \in \mathbb{R}$, for $i = 1, \dots, n$, and let $R = \{\omega | \exists i \in \{1, \dots, n\}, \exists t \in [t_i, t_{i+1}] : \omega(t) \leq b_i\}$. Then $P^N(R) \rightarrow P(R)$, with P^N, P as above.

Proof. We proceed by induction on j , showing that R is a continuity set for G . The case for $j = 1$ follows from Lemma 1 above. Suppose we proved it up to $j - 1$. Then, conditioned on an initial trajectory ω from time zero to t_j , with $\omega(t_j) = y$, G restricted in $[t_j, t_{j+1}]$ is a Gaussian process, and we can apply Lemma 1 to show that the probability of ∂R , restricted in this time span, is zero. Now, the probability of ∂R restricted to $[0, t_{j+1}]$ can be bounded by the sum of two terms. The first is the probability of ∂R in $[0, t_j]$, which is zero, the second is probability of $\partial R \cup R^c$ up to time t_j times the probability of ∂R in $[t_j, t_{j+1}]$, conditional on being in $\partial R \cup R^c$ up to time t_j . Also this second term is zero, as the conditional probability is zero for any initial trajectory ω . The bound on the probability of ∂R follows because any trajectory in ∂R up to time t_{j+1} is either touching b_j between

$[t_j, t_{j+1}]$ (second term), or before t_j (first term). The second case overlaps with the first for all trajectories that touch the threshold both before t_j and between $[t_j, t_{j+1}]$.

Proposition 7.6. For $T \in \mathbb{R}_{\geq 0}$ and $B \in \mathbb{R}^{|\Lambda| \times k}$ let A be the set defined as $A = \{x \in \mathbb{R}^{|\Lambda|} \text{ s.t. } \forall i \in \{1, \dots, m\}, (Bx)_i \leq b_i\}$. Then, for $\hat{x}_0 \in \mathbb{R}^{|\Lambda|}$ and $z_{d,0}$, the state in the state space of $\hat{Z}^{\Delta z, h, N}$ corresponding to the region containing \hat{x}_0 , it holds that:

$$\lim_{N \rightarrow \infty} \lim_{h \rightarrow 0} \lim_{\Delta z \rightarrow 0} |\rho_{reach}(\hat{x}_0, \hat{X}^N, T, A) - \rho_{reach}(z_{d,0}, \hat{Z}^{\Delta z, h, N}, \lfloor \frac{T}{h} \rfloor, A)| = 0.$$

Proof. In order to prove the convergence, we start by introducing some notation. First of all, $B\hat{X}^N$ and $B\hat{Y}^N$ are the CTMC and its CLA projected on the inequalities defining the region A . Additionally we denote by $\hat{Z}^{h, N}$ the DTMP obtained by time discretization of $B\hat{Y}^N$, and $\hat{Z}^{\Delta z, h, N}$ is the space discretization of $\hat{Z}^{h, N}$.

We now introduce the following stopping times, which are random variables on $\mathbb{R}_{\geq 0}$ denoting the random time in which a certain event happens. In particular, we are interested in the stopping times corresponding to the event of entering into the region A , usually known as hitting times, for the different processes we consider:

- τ^N is the hitting time for \hat{X}^N ;
- $\bar{\tau}^N$ is the hitting time for \hat{Y}^N ;
- $\tau^{N, h}$ is the hitting time for $\hat{Z}^{h, N}$;
- $\tau^{N, h, \Delta z}$ is the hitting time for $\hat{Z}^{\Delta z, h, N}$.

Hitting times are strictly related to the reachability probability. For instance, $Prob\{\exists t \leq T : \hat{X}^N(t) \in A\} = Prob\{\tau^N \leq T\}$. Furthermore, we introduce also the stopping time τ^G , which is the hitting time for the Gaussian process $G(t)$ to enter the rescaled region A^∞ , which is the limiting region, similarly to what we do in the proof of the Theorem 4.5. We have the following weak convergence relationships for such hitting times:

- $\tau^{N, h, \Delta z} \Rightarrow \tau^{N, h}$ as $\Delta z \rightarrow 0$;
- $\tau^{N, h} \Rightarrow \bar{\tau}^N$ as $h \rightarrow 0$;
- $\bar{\tau}^N \Rightarrow \tau^G$ as $N \rightarrow \infty$;
- $\tau^N \Rightarrow \tau^G$ as $N \rightarrow \infty$;

To show these relationships, one just has to use the correspondence of hitting times with the reachability probability, and the convergence of the latter by virtue of the proof of Theorem 4. For instance $Prob\{\tau^N \leq T\} = Prob\{\exists t \leq T : \hat{X}^N(t) \in A\} \rightarrow_{N \rightarrow \infty} Prob\{\exists t \leq T : G(t) \in A^\infty\} = Prob\{\tau^G \leq T\}$. The pointwise convergence of the cumulative distributions function of τ^N to that of τ implies weak convergence by the Portmanteau theorem [12].

In order to prove the convergence of rewards, given a reward structure ρ on \mathbb{R}^m and a path $\omega : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^m, m > 0$, we define the functional $\mathcal{R}(\omega, T) = \int_0^T \rho(\omega(s)) ds$. In order to evaluate the desired reward, we need to stop the integration as soon as the process enters the target region A , hence $\rho_{reach}(\hat{X}^N, T, A) = \mathbb{E}[\mathcal{R}(\hat{X}^N, \tau^N)]$, where the expectation is taken with respect to both X^N and τ^N . Then, by triangular inequality, we have:

$$\begin{aligned} |\mathbb{E}[\mathcal{R}(B\hat{X}^N, \tau^N)] - \mathbb{E}[\mathcal{R}(\hat{Z}^{\Delta z, h, N}, \tau^{N, h, \Delta z})]| &\leq \\ |\mathbb{E}[\mathcal{R}(B\hat{X}^N, \tau^N)] - \mathbb{E}[\mathcal{R}(B\hat{Y}^N, \bar{\tau}^N)]| &+ \\ |\mathbb{E}[\mathcal{R}(B\hat{Y}^N, \bar{\tau}^N)] - \mathbb{E}[\mathcal{R}(\hat{Z}^{h, N}, \tau^{N, h})]| &+ \\ |\mathbb{E}[\mathcal{R}(\hat{Z}^{h, N}, \tau^{N, h})] - \mathbb{E}[\mathcal{R}(\hat{Z}^{\Delta z, h, N}, \tau^{N, h, \Delta z})]| &. \end{aligned}$$

We will prove the proposition by showing that all three terms on the right hand side of the above inequality converge to zero. In particular, the third term can be sent to zero for only $\Delta z \rightarrow 0$, and the second term by sending only $h \rightarrow 0$, as both are related to the discretization of $B\hat{Y}^N$. Instead, the first term depends only on N .

We will start with the second term. First, results in [42] imply that $\hat{Z}^{h,N} \rightarrow B\hat{Y}^N$ in probability as $h \rightarrow 0$. Furthermore, Theorem 4.3 gives us weak convergence of the hitting times: $\tau^{N,h} \Rightarrow \bar{\tau}^N$. The challenge in the second term lies in the fact that it depends on two random variables, so we need to rely again on triangular inequality to separate them:

$$\begin{aligned} |\mathbb{E}[\mathcal{R}(B\hat{Y}^N, \bar{\tau}^N)] - \mathbb{E}[\mathcal{R}(\hat{Z}^{h,N}, \tau^{N,h})]| &\leq \\ |\mathbb{E}[\mathcal{R}(B\hat{Y}^N, \bar{\tau}^N)] - \mathbb{E}[\mathcal{R}(\hat{Z}^{h,N}, \bar{\tau}^N)]| &+ \\ |\mathbb{E}[\mathcal{R}(\hat{Z}^{h,N}, \bar{\tau}^N)] - \mathbb{E}[\mathcal{R}(\hat{Z}^{h,N}, \tau^{N,h})]| & \end{aligned}$$

Consider now a term appearing in the right hand side, e.g. $\mathbb{E}[\mathcal{R}(B\hat{Y}^N, \bar{\tau}^N)]$. As the expectation is taken with respect to both $B\hat{Y}^N$ and $\bar{\tau}^N$, we can rely on the following conditional expectation decomposition:

$$\mathbb{E}_{B\hat{Y}^N, \bar{\tau}^N}[\mathcal{R}(B\hat{Y}^N, \bar{\tau}^N)] = \mathbb{E}_{\bar{\tau}^N}[\mathbb{E}_{B\hat{Y}^N}[\mathcal{R}(B\hat{Y}^N, t) \mid \bar{\tau}^N = t]].$$

Furthermore, recall that:

$$\mathbb{E}_{B\hat{Y}^N}[\mathcal{R}(B\hat{Y}^N, t)] = \int_0^t \mathbb{E}[\rho(B\hat{Y}^N(s))ds].$$

Now, consider the term $|\mathbb{E}[\mathcal{R}(B\hat{Y}^N, \bar{\tau}^N)] - \mathbb{E}[\mathcal{R}(\hat{Z}^{h,N}, \bar{\tau}^N)]|$. Applying the previous decomposition, we can upper bound it by:

$$\mathbb{E}_{\bar{\tau}^N} \left[\int_0^t |\mathbb{E}[\rho(B\hat{Y}^N(s) \mid \bar{\tau}^N = t)] - \mathbb{E}[\rho(\hat{Z}^{h,N}(s) \mid \bar{\tau}^N = t)]| ds \right],$$

where we assume that $\hat{Z}^{h,N}(s)$ is a piecewise constant function in between each step at distance h , to write its cumulative reward as an integral.

We have that $\sup_{s \leq t} |B\hat{Y}^N(s) - \hat{Z}^{h,N}(s)|$ converges to zero in probability as $h \rightarrow 0$ [42]. From this, we can deduce that $\mathbb{E}[\sup_{s \leq t} |B\hat{Y}^N(s) - \hat{Z}^{h,N}(s)|]$ converges to zero. This proof presented in [42] is consequence of the *Borell-TIS inequality* [2], which guarantees that the supremum of a Gaussian process is still normally distributed. Hence: $|\mathbb{E}[\rho(B\hat{Y}^N(s) \mid \bar{\tau}^N = t)] - \mathbb{E}[\rho(\hat{Z}^{h,N}(s) \mid \bar{\tau}^N = t)]| \leq \mathbb{E}[|\rho(B\hat{Y}^N(s)) - \rho(\hat{Z}^{h,N}(s))| \mid \bar{\tau}^N = t] \leq L_\rho \mathbb{E}[|B\hat{Y}^N(s) - \hat{Z}^{h,N}(s)| \mid \bar{\tau}^N = t] \leq L_\rho \mathbb{E}[\sup_{s \leq T} |B\hat{Y}^N(s) - \hat{Z}^{h,N}(s)| \mid \bar{\tau}^N = t] = L_\rho \Delta_h$, which converges to zero by the discussion above. Recall that in the above L_ρ is the Lipschitz constant of reward ρ . Hence we can bound the first term by $\mathbb{E}[\int_0^t \Delta_h ds] \leq \Delta_h T$, which goes to zero as $h \rightarrow 0$.

Consider now the term $|\mathbb{E}[\mathcal{R}(\hat{Z}^{h,N}, \bar{\tau}^N)] - \mathbb{E}[\mathcal{R}(\hat{Z}^{h,N}, \tau^{N,h})]|$: it tends to zero by application of the Portmanteau theorem, owing to the weak convergence of $\tau^{N,h}$ to $\bar{\tau}^N$, and the fact that $\mathcal{R}(\hat{Z}^{h,N}, t)$ is a bounded and continuous function of t (being the cumulative reward up to time t of a bounded function ρ).

The third term in the main inequality, $|\mathbb{E}[\mathcal{R}(\hat{Z}^{h,N}, \tau^{N,h})] - \mathbb{E}[\mathcal{R}(\hat{Z}^{\Delta z, h, N}, \tau^{N, h, \Delta z})]|$, can be shown to converge to zero using a similar approach, owing to the convergence of the space discretization to the DTMP $\hat{Z}^{h,N}$, and the convergence of the hitting times.

What is left is the first term of the main inequality of the theorem, namely $|\mathbb{E}[\mathcal{R}(B\hat{X}^N, \tau^N)] - \mathbb{E}[\mathcal{R}(B\hat{Y}^N, \bar{\tau}^N)]|$, which has to converge to zero as N diverges.

To simplify the notation below, let us define:

- $g^N(t) = \mathbb{E}[\mathcal{R}(B\hat{X}^N, t)]$ is the cumulative reward for $B\hat{X}^N$ up to time t
- $\gamma^N(t) = \mathbb{E}[\mathcal{R}(B\hat{Y}^N, t)]$ is the cumulative reward for $B\hat{Y}^N$ up to time t .

Then the first term can be bounded by:

$$\begin{aligned} \|\mathbb{E}[g^N(\tau^N)] - \mathbb{E}[Y^N(\bar{\tau}^N)]\| &\leq \|\mathbb{E}[g^N(\tau^N)] - \mathbb{E}[g^\infty(\tau^N)]\| \\ &+ \|\mathbb{E}[Y^\infty(\tau^N)] - \mathbb{E}[Y^\infty(\tau)]\| \\ &+ \|\mathbb{E}[Y^\infty(\tau)] - \mathbb{E}[Y^\infty(\bar{\tau}^N)]\| \\ &+ \|\mathbb{E}[Y^\infty(\bar{\tau}^N)] - \mathbb{E}[Y^N(\bar{\tau}^N)]\| \end{aligned}$$

where $g^\infty = Y^\infty$ is the cumulative reward for the fluid limit $B\hat{X}^\infty = B\Phi$.

Consider the first term in the above inequality:

$$\begin{aligned} \|\mathbb{E}[g^N(\tau^N)] - \mathbb{E}[g^\infty(\tau^N)]\| &\leq \mathbb{E}_{t \sim \tau^N} \left[\mathbb{E} \left[\int_0^t \|\rho(\hat{X}^N(s)) - \rho(\Phi(s))\| ds \right] \right] \\ &\leq \mathbb{E}_{t \sim \tau^N} \left[\int_0^t L_\rho \mathbb{E}[\|X^N(s) - \Phi(s)\|] \right] \\ &\leq \mathbb{E}_{t \sim \tau^N} \left[\int_0^t L_\rho \sup_{s \leq T} \mathbb{E}[\|X^N(s) - \Phi(s)\|] \right]. \end{aligned}$$

Now $\sup_{s \leq T} \mathbb{E}[\|X^N(s) - \Phi(s)\|]$ converges to zero by virtue of a corollary of the fluid approximation theorem on the rate of convergence of expectations [30], meaning that there is N_1 such that, for $N \geq N_1$, it is less than $\epsilon/(4T)$. For all such N , it follows that $\|\mathbb{E}[g^N(\tau^N)] - \mathbb{E}[g^\infty(\tau^N)]\| \leq \epsilon/4$.

Let us deal with the fourth term:

$$\|\mathbb{E}[Y^\infty(\bar{\tau}^N)] - \mathbb{E}[Y^N(\bar{\tau}^N)]\| \leq \mathbb{E}_{t \sim \bar{\tau}^N} [\|Y^\infty(t) - Y^N(t)\|].$$

For a fixed t , we have that $\|Y^\infty(t) - Y^N(t)\| \leq \int_0^t \mathbb{E}[\|\rho(\Phi(s) + G(s)/\sqrt{N}) - \rho(\Phi(s))\|] \leq \int_0^t \mathbb{E}[L_\rho \|G(s)/\sqrt{N}\|] = L_\rho \int_0^t \mathbb{E}[\sup_{s \leq T} |G(s)|/\sqrt{N}]$. Now, as $G(t)$ has bounded covariance matrix in $[0, T]$, $\mathbb{E}[\sup_{s \leq T} |G(s)|]$ is finite, say equal to M_G , hence $\|Y^\infty(t) - Y^N(t)\| \leq L_\rho M_G t/\sqrt{N}$, and so $\|\mathbb{E}[Y^\infty(\bar{\tau}^N)] - \mathbb{E}[Y^N(\bar{\tau}^N)]\| \leq L_\rho M_G T/\sqrt{N}$ which is less than $\epsilon/4$ for $N \geq N_4$, for some $N_4 > 0$.

Terms two and three in the inequality above, instead, converge by virtue of the Portmanteau theorem and of the weak convergence of τ^N or $\bar{\tau}^N$ to τ^G , hence there is N_2 such that they are less than $\epsilon/4$ for $N \geq N_2$. It then follows that:

$$\limsup_{N \rightarrow \infty} \|\mathbb{E}[g^N(\tau^N)] - \mathbb{E}[Y^N(\bar{\tau}^N)]\| < \epsilon$$

for an arbitrary ϵ , implying:

$$\lim_{N \rightarrow \infty} \|\mathbb{E}[g^N(\tau^N)] - \mathbb{E}[Y^N(\bar{\tau}^N)]\| = 0.$$

Thus, we showed that $\|\mathbb{E}[\mathcal{R}(B\hat{X}^N, \tau^N)] - \mathbb{E}[\mathcal{R}(\hat{Z}^{\Delta z, h, N}, \tau^{N, h, \Delta z})]\|$ converges to zero for $\Delta z, h$ tending to zero and N diverging, as so do all the three terms bounding it.

REFERENCES

- [1] Alessandro Abate, Joost-Pieter Katoen, John Lygeros, and Maria Prandini. 2010. Approximate model checking of stochastic hybrid systems. *European Journal of Control* 16, 6 (2010), 624–641.
- [2] Robert J Adler. 2010. *The geometry of random fields*. SIAM.
- [3] David F Anderson and Thomas G Kurtz. [n. d.]. *Stochastic analysis of biochemical systems*. Springer.
- [4] David F Anderson and Thomas G Kurtz. 2011. Continuous time Markov chain models for chemical reaction networks. In *Design and analysis of biomolecular circuits*. Springer, 3–42.
- [5] Dana Angluin, James Aspnes, and David Eisenstat. 2008. A simple population protocol for fast robust approximate majority. *Distributed Computing* 21, 2 (2008), 87–102.

- [6] Adnan Aziz, Kumud Sanwal, Vigyan Singhal, and Robert Brayton. 1996. Verifying continuous time Markov chains. In *Computer Aided Verification*. Springer, 269–276.
- [7] Adnan Aziz, Kumud Sanwal, Vigyan Singhal, and Robert Brayton. 2000. Model-checking continuous-time Markov chains. *ACM Transactions on Computational Logic (TOCL)* 1, 1 (2000), 162–170.
- [8] Christel Baier, Boudewijn Haverkort, Holger Hermanns, and Joost-Pieter Katoen. 2000. Model checking continuous-time Markov chains by transient analysis. In *International Conference on Computer Aided Verification*. Springer, 358–372.
- [9] Christel Baier, Boudewijn Haverkort, Holger Hermanns, and Joost-Pieter Katoen. 2003. Model-checking algorithms for continuous-time Markov chains. *Software Engineering, IEEE Transactions on* 29, 6 (2003), 524–541.
- [10] Christel Baier, Joost-Pieter Katoen, et al. 2008. *Principles of model checking*. Vol. 26202649. MIT press Cambridge.
- [11] Dimitir P Bertsekas and Steven Shreve. 2004. *Stochastic optimal control: the discrete-time case*.
- [12] Patrick Billingsley. 2013. *Convergence of probability measures*. John Wiley & Sons.
- [13] Luca Bortolussi. 2016. Hybrid behaviour of Markov population models. *Information and Computation* 247 (2016), 37–86.
- [14] Luca Bortolussi, Luca Cardelli, Marta Kwiatkowska, and Luca Laurenti. 2016. Approximation of probabilistic reachability for chemical reaction networks using the linear noise approximation. In *International Conference on Quantitative Evaluation of Systems*. Springer, 72–88.
- [15] Luca Bortolussi and Jane Hillston. 2012. Fluid model checking. In *CONCUR 2012—Concurrency Theory*. Springer, 333–347.
- [16] Luca Bortolussi and Jane Hillston. 2015. Efficient checking of individual rewards properties in Markov population models. *arXiv preprint arXiv:1509.08561* (2015).
- [17] Luca Bortolussi, Jane Hillston, Diego Latella, and Mieke Massink. 2013. Continuous approximation of collective system behaviour: A tutorial. *Performance Evaluation* 70, 5 (2013), 317–349.
- [18] Luca Bortolussi and Roberta Lanciani. 2013. Model checking Markov population models by central limit approximation. In *Quantitative Evaluation of Systems*. Springer, 123–138.
- [19] Luca Bortolussi and Roberta Lanciani. 2014. Stochastic Approximation of Global Reachability Probabilities of Markov population models. In *Computer Performance Engineering*. Springer, 224–239.
- [20] Luca Bortolussi, Dimitrios Milios, and Guido Sanguinetti. 2016. Smoothed model checking for uncertain Continuous-Time Markov Chains. *Information and Computation* (2016).
- [21] Luca Cardelli, Milan Češka, Martin Fränzle, Marta Kwiatkowska, Luca Laurenti, Nicola Paoletti, and Max Whitby. 2017. Syntax-Guided Optimal Synthesis for Chemical Reaction Networks. In *International Conference on Computer Aided Verification*. Springer, 375–395.
- [22] Luca Cardelli and Attila Csikász-Nagy. 2012. The cell cycle switch computes approximate majority. *Scientific reports* 2 (2012), 656.
- [23] Luca Cardelli, Marta Kwiatkowska, and Luca Laurenti. 2016. Stochastic analysis of chemical reaction networks using linear noise approximation. *Biosystems* 149 (2016), 26–33.
- [24] Luca Cardelli, Marta Kwiatkowska, and Luca Laurenti. 2016. A stochastic hybrid approximation for chemical kinetics based on the linear noise approximation. In *International Conference on Computational Methods in Systems Biology*. Springer, 147–167.
- [25] Luca Cardelli, Marta Kwiatkowska, and Luca Laurenti. 2018. Programming discrete distributions with chemical reaction networks. *Natural computing* 17, 1 (2018), 131–145.
- [26] Taolue Chen, Tingting Han, Joost-Pieter Katoen, and Alexandru Mereacre. 2009. LTL model checking of time-inhomogeneous Markov chains. *Automated Technology for Verification and Analysis* (2009), 104–119.
- [27] Federica Ciocchetta and Jane Hillston. 2009. Bio-PEPA: A framework for the modelling and analysis of biological systems. *Theoretical Computer Science* 410, 33-34 (2009), 3065–3084.
- [28] Frits Dannenberg, Ernst Moritz Hahn, and Marta Kwiatkowska. 2015. Computing Cumulative Rewards Using Fast Adaptive Uniformization. *ACM Trans. Model. Comput. Simul.* 25, 2, Article 9 (Feb. 2015), 23 pages. <https://doi.org/10.1145/2688907>
- [29] Stewart N Ethier and Thomas G Kurtz. 2009. *Markov processes: characterization and convergence*. Vol. 282. John Wiley & Sons.
- [30] N. Gast. 2017. Expected Values Estimated via Mean-Field Approximation Are 1/N-Accurate. *Proc. ACM Meas. Anal. Comput. Syst.* 1, 1 (2017), 17:1–17:26. <https://doi.org/10.1145/3084454>
- [31] Daniel T Gillespie. 1977. Exact stochastic simulation of coupled chemical reactions. *The journal of physical chemistry* 81, 25 (1977), 2340–2361.
- [32] Daniel T Gillespie. 1992. A rigorous derivation of the chemical master equation. *Physica A: Statistical Mechanics and its Applications* 188, 1 (1992), 404–425.
- [33] Ramon Grima. 2015. Linear-noise approximation and the chemical master equation agree up to second-order moments for a class of chemical systems. *Physical Review E* 92, 4 (2015), 042124.
- [34] J Hasenauer, V Wolf, A Kazerooni, and FJ Theis. 2014. Method of conditional moments (MCM) for the chemical master equation. *Journal of mathematical biology* 69, 3 (2014), 687–735.
- [35] John Heath, Marta Kwiatkowska, Gethin Norman, David Parker, and Oksana Tymchyshyn. 2008. Probabilistic model checking of complex biological pathways. *Theoretical Computer Science* 391, 3 (2008), 239–257.

- [36] Thomas A Henzinger, Linar Mikeev, Maria Mateescu, and Verena Wolf. 2010. Hybrid numerical solution of the chemical master equation. In *Proceedings of the 8th International Conference on Computational Methods in Systems Biology*. ACM, 55–65.
- [37] Jane Hillston. 2005. *A compositional approach to performance modelling*. Vol. 12. Cambridge University Press.
- [38] Richard M Karp and Raymond E Miller. 1969. Parallel program schemata. *Journal of Computer and system Sciences* 3, 2 (1969), 147–195.
- [39] S. Krantz and P.R. Harold. 2002. *A Primer of Real Analytic Functions (Second ed.)*. Birkhäuser.
- [40] Marta Kwiatkowska, Gethin Norman, and David Parker. 2007. Stochastic model checking. In *Formal methods for performance evaluation*. Springer, 220–270.
- [41] Marta Kwiatkowska, Gethin Norman, and David Parker. 2011. PRISM 4.0: Verification of probabilistic real-time systems. In *Computer aided verification*. Springer, 585–591.
- [42] Luca Laurenti, Alessandro Abate, Luca Bortolussi, Luca Cardelli, Milan Ceska, and Marta Kwiatkowska. 2017. Reachability Computation for Switching Diffusions: Finite Abstractions with Certifiable and Tuneable Precision. In *Proceedings of the 20th International Conference on Hybrid Systems: Computation and Control*. ACM, 55–64.
- [43] Luca Laurenti, Attila Csikasz-Nagy, Marta Kwiatkowska, and Luca Cardelli. 2018. Molecular Filters for Noise Reduction. *Biophysical Journal* 114, 12 (2018), 3000–3011.
- [44] MA Lifshits. 1984. Absolute continuity of functionals of \sup type for Gaussian processes. *Journal of Mathematical Sciences* 27, 5 (1984), 3103–3112.
- [45] Oded Maler and Dejan Nickovic. 2004. Monitoring temporal properties of continuous signals. In *Formal Techniques, Modelling and Analysis of Timed and Fault-Tolerant Systems*. Springer, 152–166.
- [46] Dimitrios Milios, Guido Sanguinetti, and David Schnoerr. 2017. Probabilistic Model Checking for Continuous Time Markov Chains via Sequential Bayesian Inference. *arXiv preprint arXiv:1711.01863* (2017).
- [47] Brian Munsky and Mustafa Khammash. 2006. The finite state projection algorithm for the solution of the chemical master equation. *The Journal of chemical physics* 124, 4 (2006), 044104.
- [48] Tadao Murata. 1989. Petri nets: Properties, analysis and applications. *Proc. IEEE* 77, 4 (1989), 541–580.
- [49] Amir Pnueli. 1977. The temporal logic of programs. In *Foundations of Computer Science, 1977., 18th Annual Symposium on*. IEEE, 46–57.
- [50] René L Schilling. 2017. *Measures, integrals and martingales*. Cambridge University Press.
- [51] David Schnoerr, Botond Cseke, Ramon Grima, and Guido Sanguinetti. 2017. Efficient low-order approximation of first-passage time distributions. *Physical review letters* 119, 21 (2017), 210601.
- [52] Vahid Shahrezaei and Peter S Swain. 2008. Analytical distributions for stochastic gene expression. *Proceedings of the National Academy of Sciences* 105, 45 (2008), 17256–17261.
- [53] Mukund Thattai and Alexander Van Oudenaarden. 2001. Intrinsic noise in gene regulatory networks. *Proceedings of the National Academy of Sciences* 98, 15 (2001), 8614–8619.
- [54] Nicolaas Godfried Van Kampen. 1992. *Stochastic processes in physics and chemistry*. Vol. 1. Elsevier.
- [55] EWJ Wallace, DT Gillespie, KR Sanft, and LR Petzold. 2012. Linear noise approximation is valid over limited times for any chemical system that is sufficiently large. *IET systems biology* 6, 4 (2012), 102–115.
- [56] Verena Wolf, Rushil Goel, Maria Mateescu, and Thomas A Henzinger. 2010. Solving the chemical master equation using sliding windows. *BMC systems biology* 4, 1 (2010), 1.