IMT Institute for Advanced Studies, Lucca

Lucca, Italy

Stochastic Approximations in Model Checking: A New Scalable Approach to Collective Systems Verification

PhD Program in Computer Science (CDSS\CS)

XXVIII Cycle

By

Roberta Lanciani

2017

Program Coordinator: Prof. Rocco De Nicola, IMT Insitute for Advanced Studies Lucca

Supervisor: Prof. Rocco De Nicola, IMT Insitute for Advanced Studies Lucca

Supervisor: Prof. Luca Bortolussi, Università degli Studi di Trieste

IMT Institute for Advanced Studies, Lucca

2017

To my family

Contents

List of Figures

List of Tables

Acknowledgements

I would like to kindly thank Prof. Luca Bortolussi for his patient guidance and precious encouragement. Without his support and enthusiasm, especially during the last difficult months, all this work would not have been achievable.

I thank Prof. Rocco De Nicola for trusting and supporting me during all this years, and I express my sincere gratitude to all the members of the Quanticol Project for letting me be part of the good and exciting side of research.

I gratefully acknowledge Prof. Jane Hillston and Prof. Enrico Vicario for their kind reviews and their valuable comments.

I would like to deeply thank my wonderful colleague and kind friend Laura Nenzi, who is my special twin-star, and without whom the life-changing experience of the PhD would have been all the more difficult.

Finally, I would like to thank and dedicate this work to my family, whose love and encouragement are at the basis of all my achievements, including this work. To Alessandro, my precious island; to my parents, the most wonderful example of love, trust and strength; and to Carlo, Ilaria, Leonardo e Matilde, who always surround me with affection, smiles and inspiration...grazie!

Most part of this thesis has been published. In particular: most of Chapter 3, and Chapters 4 and 5 are based on (BL13a; BL14; BL15), respectively, all coauthored by Luca Bortolussi, University of Trieste. Part of Chapter 3 is based on a paper in preparation (BLN17), a joint work with Luca Bortolussi and Laura Nenzi, IMT Lucca.

Vita

March 18, 1987	Born, Macerata, Italy
September 2007 - October 2010	Bachelor's Degree in Mathematics Università degli Studi di Camerino, Italy Final grade: 110/110 cum laudae
August 2012 - October 2012	Visiting Student School of Informatics, University of Edinburgh, UK Granted by the mobility program of UNITS
October 2010 - December 2012	Master's Degree in Mathematics Università degli Studi di Trieste, Trieste, Italy Final grade: 110/110
February 2012 - Present	PhD Candidate in Computer Science IMT Lucca, Italy
April 2012 - Present	Active member of EU-FET project Quanticol (nr. 600708)
October 2014 - May 2015	Visiting Student University of Saarland, Germany Granted by the Erasmus Program

Publications

Conference papers:

1. L. Bortolussi, R. Lanciani, "Model Checking Markov Population Models by Central Limit Approximation," in Proc. of *10th International Conference on Quantitative Evaluation of SysTems (QEST)*, Buenos Aires, Argentina, 2013.

2. L. Bortolussi, R. Lanciani, "Stochastic Approximation of Global Reachability Probabilities of Markov Population Models," in Proc. of *11th European Workshop on Performance Engineering (EPEW)*, Florence, Italy, 2014.

3. L. Bortolussi, R. Lanciani, "Fluid Model Checking of Timed Properties," in Proc. of *13th International Conference on Formal Modeling and Analysis of Timed Systems (FORMATS)*, Madrid, Spain, 2015.

Journal papers:

4. L. Bortolussi, L. Nenzi, R. Lanciani, "Model Checking Markov Population Models by Stochastic Approximation," In Preparation.

Abstract

A *collective system* is a complex model comprised of a large number of individual entities, whose interaction gives rise to non-trivial emergent behaviours. The automatic verification of the intrinsically noisy dynamics of this type of models, by means of *Stochastic Model Checking* techniques, is severely hampered by the large size of their state space. In this project, we consider a new scalable and effective technique to validate the performance of these systems, based on *Stochastic Approximations* of the dynamics of the model. In this context, this works merges and extends the few preliminary results available in the literature at the beginning of this project and defines some interesting contributions leading the investigation in two major directions. First, we consider various types of Stochastic Approximations to accurately capture the probabilistic noise that characterises the evolution of collective systems when the number of individuals in the population is limited (*mesoscopic collective systems*). Second, we extend the set of properties that can be validated exploiting the efficiency of Stochastic Approximations. In particular, we consider requirements on the behaviour of the individuals (*local properties*), of the population (*global properties*) and of the individuals in the global context (*local-to-global properties*). Moreover, we develop procedures to verify the dynamics of *time-critical* systems. Finally, we prove the *theoretical results* that guarantee the quality of the developed model checking procedures, showing the asymptotic convergence of the results and the exactness in the limit of an infinite population size.

Chapter 1

Introduction

1.1 Motivation

Many real-life examples of large complex systems, ranging from (natural) biological mixtures to (artificial) computer networks, exhibit *collective behaviours*. These global dynamics are the result of intricate interactions between the large number of individual entities that comprise the populations of these systems. Understanding, predicting and controlling these emergent behaviours is becoming an increasingly important challenge for the scientists of the modern era. In particular, the development of an efficient and well-founded mathematical and computational modelling framework is essential to master the analysis of these complex collective systems.

In the Formal Methods community, powerful automatic verification techniques have been developed to validate the performance of a model of a system. In such procedures, called *model checkers* (BK08), the model and a property of interest are given in input to an algorithm which verifies whether or not the requirement is satisfied by the representation of the system. In the standard model checking techniques, the model is specified as a variant of a transition system, while the properties are usually instances of temporal logics.

In this context, it is essential to consider that the dynamics of a collec-

tive system is intrinsically subject to noisy behaviours, especially when the population is not very large and the stochastic evolution of the single individuals become relevant. Hence, the formal analysis and verification of a collective system have to rely on appropriate probabilistic extensions of the standard model checking techniques. In *Stochastic Model Checking* (BBHK00) the representation of the system is validated taking into account its stochastic dynamic behaviour, considering transition systems enriched with probabilities and extending the specification languages of the temporal logics to deal with stochastic constraints.

In the last decade, very powerful and successful verification algorithms have been developed in a stochastic framework, but unfortunately they all suffer from the well known curse of the *state space explosion*: when the number of interacting agents in the population increases, the variety of possible behaviours exhibited by a collective model can hamper the efficiency and applicability of the standard model checking procedures. Indeed, these verification techniques are based on an exhaustive exploration of the state space the model, which can be simply too large in the case of collective systems. To deal with this problem, some of the most successful applications of Stochastic Model Checking to large population models are based on numerical integration and statistical analysis (KNP11; JCL+09; BMS16), which however to date remain costly from a computational point of view.

1.2 Approach

In this work, we efficiently tackle the problem of *model checking collective systems* by designing fast and efficient Stochastic Model Checking procedures in which we exploit a powerful class of methods to accurately approximate the dynamics of the individuals and the population: the Stochastic Approximations.

Stochastic Approximations (BHLM13) have been successfully used in recent years in the Biology community (Gri10; Van92) to approximate the noisy behaviour of collective systems with a stochastic process, whose dynamics is encoded in a (numerically integrable) set of *Differential Equa-*

tions (DEs). Hence, when we make use of the right formal framework to describe the collective system, Stochastic Approximations represent a fast and easy way of obtaining an estimation of the dynamics of the model. Moreover, for almost all the techniques that we are going to consider in this work, the quality of the estimations improves as the number of agents in the system increases, keeping constant the computational cost and reaching exactness in the limit of an infinite population. In this way, these approximation methods actually take advantage of the large sizes of the collective systems, making them a fast, accurate and reliable approach to deal with the curse of the state space explosion. Among the many types of Stochastic Approximations present in the literature, we are going to exploit the *Fluid Approximation* (FA) (BH12b), the *Central Limit Approximation* (CLA) (Van92; EK05), and the *System Size Expansion* (SSE) (SSG16), and in some cases, we are going to compare them with the results obtained by the *Moment Closure* (MC) combined with a distribution reconstruction based on the *Maximum Entropy Principle* (AMW15a).

This thesis extends the few preliminary works that, before the beginning of this project, had already applied Stochastic Approximation procedures to the verification of collective systems. In particular, we extend the work done in (BH12b) and in (HSB12; HBC13). In the former, the authors exploit Fluid Approximation to verify properties of a single individual in a collective system. This was done by approximating the dynamics of the stochastic model with a deterministic trajectory, encoded in a system of Ordinary Differential Equations (ODEs). A similar approach was taken in (HSB12; HBC13), where also Moment Closure techniques where taken into account. In this framework, this work defines some interesting contributions in the verification of collective systems by merging and extending the work of (BH12b; HSB12; HBC13) in two major directions:

- *Accurate estimation of the stochastic noise in a collective system* - To go beyond the *deterministic* approximation of the dynamics of the collective system given by the Fluid Approximation, we consider higher order estimations like the Central Limit Approximation, the System Size Expansion, and the Moment Closure. These latter meth-

ods, indeed, can be used to define a description of the stochastic fluctuations of the probability distribution that represents the state of a collective system at a given time instant. As we have said, the behaviour of populations is intrinsically stochastic, especially when the number of agents is not very large (in these cases, we refer to the system as a *mesoscopic collective system*). Hence, model checking procedures that involve higher order approximations prove to be the perfect tool for a fast and accurate verification of mesoscopic collective systems.

- *Extension of the set of properties that can be validated* - Together with the possibility of verifying requirements that characterise the behaviour of the single individual (*local properties*) as in the work of (BH12b), we want to extend the set of properties that can be validated to comprise requirements on the behaviour of the entire population (*global properties*) and on the actions of the individuals in the global context (*local-to-global properties*). Moreover, we want to be able to model-check *time-critical* collective systems, meaning systems that show behaviours that are subject to time constraints.

1.3 Contributions

In this work, we efficiently tackle the problem of model checking collective systems by designing fast and efficient Stochastic Model Checking procedures to validate a wide set of properties exploiting different types of Stochastic Approximations.

We define a procedure to verify the behaviour of individual agents in the global context (*Local-to-Global properties*) exploiting *Central Limit Approximations*, *System Size Expansion* and *Moment Closure* (BL13a; BLN17). The requirements that we investigate in this context are stochastic properties of the type: "it is almost sure that 95% of the population will satisfy a specific behaviour within a time instant T". Hence, in this type of properties, we are interested in a specific behaviour that has to be met by the single agents in the population (a local property), and we look at the frac-

tion of individuals that satisfy it at the global level (local-to-global property), assigning a probability measure to the fact that this group of agents reaches at least 95% of the population within time T. The local property is expressed as a *Deterministic Time Automata* (DTA) (AD94), which, as in (BBHK00), is synchronised with an appropriate representation of the collective system, i.e. a *(Markov) Population Model*, in order to keep track of the behaviour of the single agents in the population. The dynamics of the synchronised model is then approximated using the Central Limit Approximation, the System Size Expansion and the Moment Closure to compute the probability measure of the Local-to-Global requirement. As it was already said in the previous sections, in this work we shall see how this type of model checking procedure becomes extremely powerful in the validation of *mesoscopic systems*, meaning systems with populations comprised of a limited number of agents and where the dynamics is intrinsically stochastic.

We also tackle the problem of validating *Global Reachability Properties* (BL14), meaning that we are interested in the fast computation of an accurate approximation of the probability that, within a given time horizon T, the collective system reaches a specific region of the state space, defined by a linear inequality of the counting variables that identify the state of the population. The stochastic model checking procedure that we develop in this case, transforms this type of reachability problems into a form of *Hitting-Time Problem*. Indeed, based on (EK05), we validate the Global Reachability Property by actually approximating the probability distribution of the time $t_{\mathcal{R}}$ in which the system enters the target region $\mathcal{R}$. To estimate $t_{\mathcal{R}}$, we again exploit the accuracy and efficiency of the *Central Limit Approximation* and the *System Size Expansion*.

Finally, we extend the procedure of Fluid Model Checking of (BH12b) to deal with timed local properties of single agents, specified by a *Deterministic Timed Automata* (DTA) endowed with a single clock that is allowed to be *reset*. Hence, building on the theory of the Fluid Approximation, to estimate the statisfaction probability of the local properties, we actually compute the probability that a properly defined subclass of *Time-Inhomogeneous Markov Renewal Processes* (CHKM11a), with expo-

nentially and deterministically-timed transitions, reaches an absorbing region of its states space. In doing this, we reduce the computation of the satisfaction probability to the numerical integration of a set of *Delay Differential Equations* (DDE).

In the following, we are going to show that all the Stochastic Model Checking techniques that we develop are reliable, fast and accurate. Moreover, we are going to prove the *theoretical results* that guarantee their quality, by showing their asymptotic convergence and their exactness in the limit of an infinite population size.

1.4 Structure of the thesis

In Chapter 2, we review the theoretical background that is at the basis of the model checking procedures that we develop. In particular, we look at the definition of a collective system as a Markov Population Model whose evolution is encoded in a Markov Chain. Afterwards, we review the theory of Stochastic Approximations, describing in detail the Fluid Approximation, the Central Limit Approximation, the System Size Expansion and the Moment Closure combined with the Maximum Entropy Principle. Finally, we move to the description of the properties and of the specification languages that we consider in this work.

In Chapter 3, we review the first results, defining a model checking procedure to validate local-to-global properties of collective systems applying Central Limit Approximation, Sistem Size Expansions and Moment Closure.

In Chapter 4, we illustrate the model checking algorithm that efficiently and accurately validates Global Reachability Properties computing the probability distribution of a Hitting-Time Problem.

In Chapter 5, we review the validation techniques for time-critical systems, approximating the dynamics of the system with a Time Inhomogeneous Markov Renewal Process, and integrating the set of DDEs that defines the its evolution in time.

Finally, in Chapter 6, we summarise the major results of this project, we draw some conclusions and we discuss new possible lines of research.

Chapter 2

Background

2.1 Overview

In this chapter, we review the theoretical background behind the model checking procedures that we build in this work. In particular, we illustrate the definition of the formal models that describe the collective systems we are interested in, the *Markov Population Models* and the mathematical specification that enables the approximation of their behaviour, the *Markov Chains* (Sections 2.2 and 2.3); we introduce the theory of *Stochastic Approximations* (Section 2.4), focusing on the Fluid Approximation, the Central Limit Approximation, the System Size Expansion and the Moment Closure; and finally, in Section 2.5 we review the *properties* and the *specification languages* that we exploit in our model checking procedures.

2.2 Markov Chains

Markovian processes (MP) are stochastic processes (i.e. dynamical processes taking a random value at any time instant (Dur10)) that enjoy the *Markov property*: the evolution of the probability distribution depends only on the current state of the process and is not altered by additional knowledge concerning its past behaviour. This means that this sort of

processes retain *no memory*. Moreover, MPs can be classified according to the cardinality of their state space or whether they depend on discrete or continuous time. The processes that have finite or countably infinite state spaces are called *Markov chains* (Nor97; Dur10; EK05).

Discrete Time Markov Chains

A discrete time stochastic process $\{X(n) \in S \mid n \in \mathbb{N}\}$ with finite or countably infinite state space S is called a *Discrete Time Markov Chain* (DTMC) if, for every $n \in \mathbb{N}$ and $j, i, i_1, \ldots, i_n \in S$,

$$\mathbb{P}\left(X(n+1) = j \mid X(n) = i, X(n-1) = i_1, \ldots, X(0) = i_n\right) =$$

$$= \mathbb{P}\left(X(n+1) = j \mid X(n) = i\right). \tag{2.1}$$

Condition (2.1), known as the *Markov property*, states that the conditional probability on $X(n+1) = j$ depends only on the previous state $X(n) = i$ and not on the entire history of the stochastic process $X(n-1) = i_1, \ldots, X(0) = i_n$.

A DTMC is *(temporally) homogeneous* if the conditional probability (2.1) is independent of $n \in \mathbb{N}$, i.e. there exists $p_{ij} \in [0, 1]$ such that

$$\mathbb{P}(X(n+1) = j \mid X(n) = i) = p_{ij}, \qquad \forall n \in \mathbb{N}.$$

The element p_{ij} is called *transition probability from i to j* and a state $i \in S$ is *absorbing* if

$$p_{ij} = 0, \qquad \forall j \in S,\ j \neq i.$$

The matrix $\mathbf{P} = (p_{ij}) \in [0, 1]^S \times [0, 1]^S$ is the *transition matrix* of the DTMC and is a *stochastic matrix*, meaning that every row in P is a probability distribution, i.e.

$$p_{ij} \in [0, 1],\ \ \forall i, j \in S, \qquad \text{and} \qquad \sum_i p_{ij} = 1.$$

The first formulation of DTMCs goes back to Markov in 1906, and since then it has been applied and studied in different research fields: from systems biology and social sciences, to electrical engineering and

information theory. For a full review good textbooks are (Nor97; Dur10; EK05).

The power of this mathematical model relies totally in the memoryless property (2.1). This principle not only turns out to be adequate to model a great variety of interesting real-life examples of random phenomena, but also makes the mathematical formulation simple and intuitive from a computational point of view. Indeed, given a *(finite) path* σ of the form

$$\sigma := i_0 \longrightarrow i_1 \longrightarrow i_2 \longrightarrow \ldots \longrightarrow i_{T-1} \longrightarrow i_T, \qquad T \in \mathbb{N}, \tag{2.2}$$

where $i_j \in S$ and $p_{i_j i_{j+1}} > 0$, $\forall\ j \in \{0, 1 \ldots, T-1\}$, thanks to the Markov property (2.1), the probability for the evolution of the DTMC $\{X(n) \mid n \in \mathbb{N}\}$ to coincide with σ is simply the initial probability of being in i_0 multiplied by the product of the transition probabilities within the states, i.e.

$$\mathbb{P}(\sigma) := \mathbb{P}(X(0) = i_0, X(1) = i_1, \ldots, X(T) = i_T) = p_{0,i_0} \prod_{j=1}^{T} p_{i_j i_{j+1}},$$

where $p_{0,i_0} = \mathbb{P}(X(0) = i_0)$. Analogously, if we consider the *transient-state probability* $\pi(s_0, s, T)$ *of being in state* s *at time* T *starting at* s_0,

$$\pi(s_0, s, T) = \sum_{\sigma_\alpha} \mathbb{P}(\sigma_\alpha),$$

where the sum is made over all the paths $\sigma_\alpha \in Path(\mathcal{X})$ such that $\sigma_\alpha(0) = s_0$ and $\sigma_\alpha(T) = s$, if we apply property (2.1), it is not difficult to prove that

$$\pi(s_0, s, T) = p_{s_0 s}^{(T)},$$

where $p_{s_0 s}^{(T)}$ is the (s_0, s)-th entry of the T-th power of the transition matrix $\mathbf{P}$.

Furthermore, there is a one-to-one correspondence between transition matrices and *labelled graphs* (Nor97) (an example is illustrated in Figure 1). For this reason, DTMCs can be intuitively seen as variants of transition systems enriched with probabilities. Indeed, labelled graphs

$$\mathbf{P} = \begin{pmatrix} 0 & 1 & 0 \\ 0 & \frac{1}{2} & \frac{1}{2} \\ \frac{1}{3} & \frac{2}{3} & 0 \end{pmatrix}$$

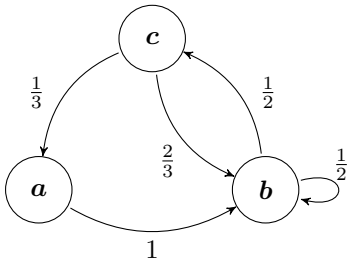

Figure 1: Example of the graph representing a DTMC with state space $S = \{a, b, c\}$ and transition matrix $\mathbf{P}$.

are a simple class of models that admit *probabilistic choice*, meaning that one can specify the probability of making a transition form one state (or vertex) to the other, and they are a simple and intuitive descriptions of DTMCs. As we shall see in the following, the same is true for Continuous Timed Markov Chains (CTMCs) and what distinguishes the two types of probabilistic models is the interpretation of *time*: the underlying time domain of DTMCs is discrete and each transition is assumed to take a single time unit.

Continuous Time Markov Chains

Continuous Time Markov Chains (CTMCs) are well known probabilistic models which admits *continuous time* and *no memory*. We shall see that a CTMC is completely specified by the definition of its initial state and of the *rates* of taking a transition form one state to the other.

Formally, a continuous time stochastic process $\left\{X(t) \in S \mid t \in \mathbb{R}^{\geq 0}\right\}$ with finite or countably infinite state space S is a *Continuous Time Markov Chain* (CTMC) if it enjoys the *Markov property*, which in the continuous domain takes the following form: for every $t, s \in \mathbb{R}^{\geq 0}$ and $i, j \in S$,

$$\mathbb{P}(X(t+s) = j \mid X(s) = i, X(r) = i_r, i_r \in S, r \in \mathbb{R}^{\geq 0}, r < s) =$$

$$= \mathbb{P}(X(t+s) = j \mid X(s) = i). \tag{2.3}$$

A CTMC is *(temporally) homogeneous* if its probability distribution is constant with respect to the translation in time, i.e. for every $t, s \in \mathbb{R}^{\geq 0}$

and $i, j \in S$,

$$\mathbb{P}(X(t+s) = j \mid X(s) = i) = \mathbb{P}(X(t) = j \mid X(0) = i) = p_{ij}(t).$$

The function $\mathbf{P} : \mathbb{R}^{\geq 0} \longrightarrow [0,1]^S \times [0,1]^S$, $\mathbf{P}(t) = (p_{ij}(t))$, is called the *transition probability function* and we require for its elements $p_{ij} : \mathbb{R}^{\geq 0} \longrightarrow [0,1]$ to be *differentiable* in t for every $i, j \in S$.

Given a CTMC with finite state space S and transition probability function $\mathbf{P} : \mathbb{R}^{\geq 0} \longrightarrow [0,1]^S \times [0,1]^S$, $\mathbf{P}(t) = (p_{ij}(t))$, we define the *infinitesimal generator* or *rate matrix* $\mathbf{Q} = (q_{ij})$ of the CTMC to be the $\mathbb{R}^S \times \mathbb{R}^S$-matrix whose elements are such that, for small intervals of time $\triangle t$,

$$p_{ij}(\triangle t) = \delta_{ij} + q_{ij}\triangle t + o(\triangle t),$$

where δ_{ij} is the *Kronecker delta*, i.e. the matrix that has $\delta_{ij} = 0$, for all $i \neq j$, and $\delta_{ii} = 1$. The elements q_{ij} of $\mathbf{Q}$, instead, are called *rate functions* and, since $p_{ij}(t)$ is a probability, we have

$$0 \leq -q_{ii} < \infty \; \forall i, \quad q_{ij} \geq 0 \text{ for } i \neq j, \quad \sum_j q_{ij} = 0 \quad \text{and} \quad q_{ii} = -\sum_{j \neq i} q_{ij}.$$

A CTMC with finite state space S is well-defined when we specify its state space S, its infinitesimal generator $\mathbf{Q} \in \mathbb{R}^S \times \mathbb{R}^S$, and the initial distribution $\mathbf{p}_0 = (p_j)_{j \in S}$. Indeed, given the tuple $(S, \mathbf{Q}, \mathbf{p}_0)$, the transition probability function $\mathbf{P} : \mathbb{R}^{\geq 0} \longrightarrow [0,1]^S \times [0,1]^S$, $\mathbf{P}(t) = (p_{ij}(t))$, can be computed according to the following Theorem (Nor97).

Theorem 2.1 (Kolmogorov's forward and backward equations) *Let* $\mathbf{Q}$ *be the infinitesimal generator of a CTMC* $\{X(t) \in S \mid t \in \mathbb{R}^{\geq 0}\}$ *with finite state space* S*. Then,*

- *the transition probability function* $\mathbf{P} : \mathbb{R}^{\geq 0} \longrightarrow [0,1]^S \times [0,1]^S$ *of* $\{X(t)\}$ *is the minimal non-negative solution of the matrix differential equation*

$$\frac{\partial \mathbf{P}}{\partial t} = \mathbf{Q}\mathbf{P}(t), \qquad \mathbf{P}(0) = I; \tag{2.4}$$

- $\mathbf{P}$ *is also the minimal non-negative solution of*

$$\frac{\partial \mathbf{P}}{\partial t} = \mathbf{P}(t)\mathbf{Q}, \qquad \mathbf{P}(0) = I; \tag{2.5}$$

$$\mathbf{Q} = \begin{pmatrix} -1 & \frac{1}{2} & \frac{1}{4} & \frac{1}{4} \\ 0 & -\frac{1}{2} & 0 & \frac{1}{2} \\ \frac{1}{3} & 0 & -\frac{2}{3} & \frac{1}{3} \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

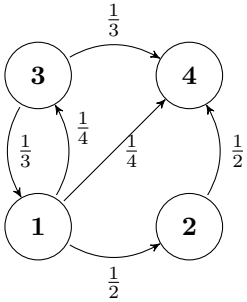

Figure 2: Example of the graph representing the CTMC given by $(S = \{1, 2, 3, 4\}, \mathbf{Q}, \mathbf{p}_0)$. Only arrows with positive rates are drawn.

- $\mathbf{P}$ *can be expressed in terms of the matrix exponential as*

$$\mathbf{P}(t) = e^{\mathbf{Q}t} = \sum_{n=0}^{\infty} \frac{\mathbf{Q}^n t^n}{n!};$$

- *the semigroup or memoryless property holds true:*

$$\mathbf{P}(s+t) = \mathbf{P}(s)\mathbf{P}(t) \quad \forall s, t \in \mathbb{R}^{\geq 0}.$$

In analogy with the discrete time case, there is a one-to-one correspondence between rate matrices $\mathbf{Q}$ and labelled graphs (see Figure 2), and CTMCs can be seen as another way to add probabilities to transition systems. As in the case of DTMCs, the Markov property makes the CTMCs a very powerful mathematical tool to model a great variety of natural probabilistic phenomena and the continuous real-time framework widens the area of applicability even more. For this reason, CTMCs are among the most popular operational models in performance evaluation and are the basis of this work.

2.3 Markov Population Models

In this section, we start to tackle the problem of the design of a collective system and we introduce the modelling framework that enables the

stochastic approximation of its behaviour. In particular, we define an automata-based formalism to specify *Markov Population Models* consisting of large collections of interacting components, or *agents*. Each component is a finite transition system, instance of an *Agent Class* $\mathbb{A}$ that defines its (finite) state space and its (finite) set of *local* transitions. The material is mainly based on (BL13b).

Definition 2.1 (Agent Class) *An* Agent Class $\mathbb{A}$ *is a pair*

$$\mathbb{A} = (S, E)$$

where $S = \{1, \ldots, n\}$ *is the state space of the agent and* $E = \{\epsilon_1, \ldots, \epsilon_m\}$ *is the set of local transitions of the form*

$$\epsilon_i = s_i \xrightarrow{\alpha_i} s_i', \qquad i \in \{1, \ldots, m\},$$

where α_i *is the transition label, taken from the label set* $\mathscr{L}$*.*

An agent belonging to class $\mathbb{A} = (S, E)$ defines a continuous time random variable $Y(t) \in S$, which denotes the state of the agent at time t. Moreover, let $Y(0) \in S$ be its initial state.

In the following, we consider populations of N agents $Y_k^{(N)}$, $k \in \{1, \ldots, N\}$, all belonging to the same class $\mathbb{A} = (S, E)$ with $S = \{1, \ldots, n\}$ We further make the classical assumption that agents in the same state are *indistinguishable*, hence the state of the population model can be described by *collective* or *counting variables*

$$\mathbf{X}^{(N)} = (X_1^{(N)}, \ldots, X_n^{(N)}), \qquad X_j^{(N)} \in \{0, \ldots, N\},$$

defined by

$$X_j^{(N)} = \sum_{k=1}^{N} \mathbb{1}\{Y_k^{(N)} = j\}.$$

The initial state $\mathbf{x}_0^{(N)}$ is given by $\mathbf{x}_0^{(N)} = \mathbf{X}^{(N)}(0)$, and the counting variables satisfy the conservation relation $\sum_{j \in S} X_j^{(N)} = N$. To complete the definition of a Markov Population Model, we need to specify its *global* transitions, describing all possible events that can change the state of the system.

Definition 2.2 (Markov Population model) *A Markov Population Model $\mathcal{X}^{(N)}$ of size N is a tuple*

$$\mathcal{X}^{(N)} = (\mathbb{A}, \mathcal{T}^{(N)}, \mathbf{x}_0^{(N)}),$$

where:

- $\mathbb{A}$ *is an Agent Class, as in Definition 2.1;*
- $\mathcal{T}^{(N)} = \{\tau_1, \ldots, \tau_\ell\}$ *is the set of global transitions of the form*

$$\tau_i = (\mathbb{S}i, f_i^{(N)}),$$

  where:

  - $\mathbb{S}i = \{s_1 \xrightarrow{\alpha_1} s_1', \ldots, s_p \xrightarrow{\alpha_p} s_p'\}$ *is the (finite) set of local transitions synchronized by τ_i;*
  - $f_i^{(N)} : \mathbb{R}^n \longrightarrow \mathbb{R}_{\geq 0}$ *is the (Lipschitz continuous) global rate function.*

- $\mathbf{x}_0^{(N)}$ *is the initial state.*

The rate $f_i^{(N)}$ gives the expected frequency of transition τ_i as a function of the state of the system. We assume $f_i^{(N)}$ equal to zero if there are not enough agents available to perform the transition. The synchronization set $\mathbb{S}i$, instead, specifies how many agents are involved in the transition τ_i and how they change state: when τ_i occurs, we see the local transitions $s_1 \xrightarrow{\alpha_1} s_1', \ldots, s_p \xrightarrow{\alpha_p} s_p'$ fire at the (local) level of the p agents involved in τ_i.

Remark 2.1 *The population models we consider in this thesis assume that the size N of the population is constant. This limitation eases the presentation (and the notation) of the model checking procedures that we define and that exploit Stochastic Approximations, however the assumption can be removed, as it is not a necessary condition. In doing so, though, extra care has to be taken in treating local properties, as discussed in (BH12b).*

Given a Markov Population Model $\mathcal{X}^{(N)} = (\mathbb{A}, \mathcal{T}^{(N)}, \mathbf{x}_0^{(N)})$ and a global transition $\tau = (\mathbb{S}\tau, f_\tau^{(N)}) \in \mathcal{T}^{(N)}$ with $\mathbb{S}\tau = \{s_1 \xrightarrow{\alpha_1} s_1', \ldots, s_p \xrightarrow{\alpha_p} s_p'\}$, we encode the net change in $\mathbf{X}^{(N)}$ due to τ in the *update vector*

$$\mathbf{v}_\tau = \sum_{i=1}^{p} (\mathbf{e}_{s_i} - \mathbf{e}_{s_i'}),$$

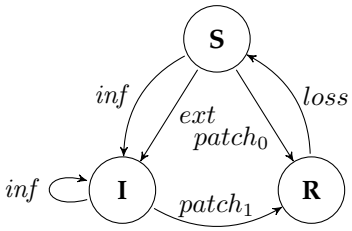

Figure 3: The automaton representation of a network node.

where $\mathbf{e}_{s_i}$ is the vector that is equal to 1 in position s_i and zero elsewhere.

We define the *CTMC* $\mathbf{X}^{(N)}(t)$ *associated with* $\mathcal{X}^{(N)}$ as the continuous time stochastic process that has state space

$$\mathcal{S}^{(N)} = \{(z_1, \ldots, z_n) \in \mathbb{N}^n \mid \sum_{i=1}^{n} z_i = N\},$$

initial probability distribution concentrated on $\mathbf{x}_0^{(N)}$, and infinitesimal generator matrix $\mathbf{Q}$ defined for $\mathbf{x}, \mathbf{x}' \in \mathcal{S}^{(N)}$, $\mathbf{x} \neq \mathbf{x}'$, by

$$q_{\mathbf{x},\mathbf{x}'} = \sum_{\tau \in \mathcal{T} | \mathbf{v}_\tau = \mathbf{x}' - \mathbf{x}} f_\tau(\mathbf{x}).$$

Example. To illustrate the modelling technique, we consider a simple example of a worm epidemic in a peer-to-peer network composed of N nodes. Each node is modelled by the simple agent shown in Figure 3, which has three states: susceptible to infection (S), infected (I), and patched/immune to infection (R). The contagion of a susceptible node can occur due to an event external to the network (ext), like the reception of an infected email, or by file sharing with an infected node within the network (inf). Nodes can also be patched, at different rates, depending if they are infected ($patch_1$) or not ($patch_0$). A patched node remains immune from the worm for some time, until immunity is lost ($loss$), modelling for instance the appearance of a new version of the worm.

The Agent Class of the network node has the form

$$\mathbb{A}_{node} = (S_{node}, E_{node})$$

and can be easily reconstructed form the automaton representation in Figure 3. The Markov Population Model

$$\mathcal{X}_{net}^{(N)} = (\mathbb{A}_{node}, \mathcal{T}^{(N)}, \mathbf{x}_0^{(N)})$$

with population variables

$$\mathbf{X} = (X_S, X_I, X_R)$$

is obtained by specifying transitions and initial conditions. We start the model with a simple network of susceptible nodes, hence the initial conditions are

$$\mathbf{x}_0^{(N)} = (N, 0, 0).$$

The transition set of the Markov Population Model instead is given by five global transitions: τ_{ext}, τ_{loss}, τ_{patch_0}, τ_{patch_1}, $\tau_{inf} \in \mathcal{T}^{(N)}$. For example, the external infection is defined by

$$\tau_{ext} = (\{S \xrightarrow{ext} I\}, f_{ext}),$$

where the synchronisation set specifies that only one susceptible node is involved and changes state from S to I at a rate given by

$$f_{ext}(\mathbf{X}) = \kappa_{ext} X_S,$$

corresponding to a rate of infection κ_{ext} per node. The transitions τ_{loss}, τ_{patch_0}, τ_{patch_1} have a similar format, while the internal infection is described by

$$\tau_{inf} = (\{I \xrightarrow{inf} I, S \xrightarrow{inf} I\}, f_{inf}),$$

and involves one S-node and one I-node. Furthermore, in this case of τ_{inf}, we assume that an infected node sends infectious messages at rate κ_{inf} to a random node, giving a classical density dependent rate function (AB00)

$$f_{inf}(\mathbf{X}) = \frac{1}{N} \kappa_{inf} X_S X_I.$$

2.4 Stochastic Approximations

The Stochastic Approximation techniques that we are about to describe have been studied and applied in recent years to tackle the *state space explosion problem* in the analysis of collective systems. Indeed, when the size of the system increases, and thus the set of reachable states enlarges, the numerical integration of the CTMCs that describe the evolution of the marginal probability distribution of a collective system over its the states, becomes infeasible. Hence, various types of *Stochastic Approximations* have been studied in recent years, in order to provide a *fast* and *accurate* estimation of the probability distributions for collective systems like the ones we are interested in.

Among the possible approximation methods, in this work, we focus on five particular techniques: the *Fluid Approximation*, the *Fast Simulation*, the *Central Limit Approximation*, the *System Size Expansion* and the *Moment Closure*. The first four methods represent a direct analytical approximation of the probability that describes the evolution of the collective system, and are all based on *Van Kampen's system size expansion* (EK05). The latter method, the *Moment Closure*, instead, *reconstructs* an approximation of the probability distribution that describes the evolution in time of the population model, starting from an estimate of the *moments* of the real distribution. Moreover, since a finite set of moments defines a set of probability distributions, the Moment Closure relies on the *Maximum Entropy Principle* to choose as the best approximation of the probability distribution of the collective system, the one that maximises the entropy (i.e. the *probabilistic noise*) of the system.

As we shall see in the following, all the Stochastic Approximations that we present in this work are particularly fruitful in cases like the ones modelled in Section 2.3, where we consider large populations comprised of big clusters, or classes, of identical interacting agents. Indeed, the complexity of these methods is *independent of the population size*. Actually, for the methods based on the system size expansion, we are going to prove that the accuracy of the approximation *increases* with the population size and is exact in the limit of an infinite population.

2.4.1 Fluid Approximation and Fast Simulation

Fluid Approximation, also known as *Mean Field Approximation*, has been successfully applied in a great variety of research fields including biology (BP09; Car08), game theory (BW03), computer viruses (BGH08), gossip protocols (BCFH09), crowd models (MLBH10; MLBH11) and performance evaluation of computer networks (TG11), Petri Nets (SR04) and computational grids (BCGH05). In Fluid Approximation the discrete, stochastic behaviour of the system is approximated by that of a continuous, *deterministic* model. This is done by treating the global transition rates as *flows*, thus obtaining a set of ODEs that estimate the behaviour of the underlying CTMC of the system.

Formally, we initially consider the Markov Population Model $\mathcal{X}^{(N)} = (\mathbb{A}, \mathcal{T}^{(N)}, \mathbf{x}_0^{(N)})$ of the system, built as described in Section 2.3. We then construct an infinite sequence $(\mathcal{X}^{(N)})_{N\in\mathbb{N}}$ of population models, all sharing the same structure, for increasing population size $N \in \mathbb{N}$ (e.g. the network models $(\mathcal{X}_{net}^{(N)})_{N\in\mathbb{N}}$ of the example in Section 2.3 with an increasing number of network nodes). The Fluid Approximation technique works by approximating the stochastic dynamics of $\mathcal{X}^{(N)}$ by the behaviour of the sequence $(\mathcal{X}^{(N)})_{N\in\mathbb{N}}$ in the limit $N \longrightarrow \infty$.

To compare the dynamics of the models in the sequence $(\mathcal{X}^{(N)})_{N\in\mathbb{N}}$, we consider the *normalised counting variables* $\hat{\mathbf{X}} = \frac{1}{N}\mathbf{X}$ (known also as *population densities* or *occupancy measures*, see (BHLM13) for further details) and we define the *normalized (Markov) Population Models* $\hat{\mathcal{X}}^{(N)} = (\mathbb{A}, \hat{\mathcal{T}}^{(N)}, \hat{\mathbf{x}}_0^{(N)})$, obtained from $\mathcal{X}^{(N)}$ by making the rate functions depend on the normalised variables and rescaling the initial conditions.

For simplicity, we assume that the rate function of each transition $\tau \in \hat{\mathcal{T}}^{(N)}$ satisfies the *density dependent condition*

$$\frac{1}{N} f_\tau^{(N)}(\hat{\mathbf{X}}) = f_\tau(\hat{\mathbf{X}})$$

for some Lipschitz function $f_\tau : \mathbb{R}^n \longrightarrow \mathbb{R}_{\geq 0}$, i.e. rates on normalised variables are independent of N.

To define the limit ODEs, we introduce the *drift* of $\mathcal{X}^{(N)}$, which is the

mean instantaneous change of the normalized models and is given by

$$\mathbf{F}(\hat{\mathbf{X}}) = \sum_{\tau \in \hat{\mathcal{T}}^{(N)}} \mathbf{v}_\tau f_\tau(\hat{\mathbf{X}}).$$

The *Fluid Approximation* of the CTMC $\hat{\mathbf{X}}^{(N)}(t)$ associated with $\hat{\mathcal{X}}^{(N)}$ is the unique solution of the ODE system

$$\begin{cases} \frac{d\mathbf{\Phi}(t)}{dt} = \mathbf{F}(\mathbf{\Phi}(t)); \\ \mathbf{\Phi}(0) = \hat{\mathbf{x}}_0^{(N)}. \end{cases} \tag{2.6}$$

The existence and uniqueness of $\hat{\mathbf{X}}^{(N)}(t)$ is guaranteed by the fact that all f_τ are Lipschitz continuous, thus $\mathbf{F}$ is.

The correctness of the estimate of the Fluid Approximation in the limit of an infinite population is guaranteed by the following Theorem. The proof of this result involves the theory of Markov Processes and the interested reader can refer to (EK05).

Theorem 2.2 *Let $\hat{\mathbf{X}}^{(N)}(t)$ and $\mathbf{\Phi}(t)$ be defined as before. Assume that there exists $\mathbf{x_0} \in S$ such that $\lim_{N\to\infty} \hat{\mathbf{X}}^{(N)}(0) = \mathbf{x_0}$. For any finite time horizon $T < \infty$, it holds that:*

$$\lim_{\Omega\to\infty} \sup_{t\in[0,T]} \left\| \hat{\mathbf{X}}^{(N)}(t) - \mathbf{\Phi}(t) \right\| = 0 \qquad \textit{almost surely}.$$

Fast Simulation

In this work, we are also interested in the behaviour of a (random) *single agent* inside a population. As we have just seen, the dynamics of a large population can be accurately described by a *deterministic* limit, the Fluid Approximation. But when we focus on one single agent in a collective system, we need to keep in mind that its behaviour in time will always remain a *stochastic* process, even in large populations. Nevertheless, the *Fast Simulation Theorem* (DN08; BLB08; GB10) guarantees that in the limit of an infinite population size, the stochastic process of the single agent senses only the *mean* behaviour of the rest of the agents (i.e. there is no need to keep track of all the states of all the other entities in the population). This means that, when the population size is large enough, to

analyse the dynamics the single agent, we can define the Fluid Approximation of the population model, and then use its state (i.e. the mean state of the rest of the agents) to compute the rates of a *Time-Inhomogeneous CTMC* (ICTMC) (BH15) that describes the behaviour of the single agent.

Formally, let $Y^{(N)}(t)$ be the stochastic process that describes the state of the single agent in the population model $\mathcal{X}^{(N)} = (\mathbb{A}, \mathcal{T}^{(N)}, \mathbf{x}_0^{(N)})$ with state vector $\mathbf{X}^{(N)}(t)$. By definition, $Y^{(N)}(t)$ *is not independent of* $\mathbf{X}^{(N)}(t)$. Now consider the normalised model $\widehat{\mathcal{X}}^{(N)}$ described by $\widehat{\mathbf{X}}^{(N)}(t)$, and let $\mathbf{\Phi}(t)$ be the Fluid Approximation of $\mathcal{X}^{(N)}$. Define the *generator matrix* $\boldsymbol{Q}^{(N)}(\boldsymbol{x}) = (q_{ij}^{(N)}(\boldsymbol{x}))$ of $Y^{(N)}(t)$ as a function of the normalised counting variables, i.e. $\forall\, q_{ij}^{(N)}(\boldsymbol{x})$,

$$Prob\left\{Y^{(N)}(t+dt) = j \mid Y^{(N)}(t) = i, \widehat{\mathbf{X}}^{(N)}(t) = \boldsymbol{x}\right\} = q_{ij}^{(N)}(\boldsymbol{x})dt.$$

Notice that $\boldsymbol{Q}^{(N)}(\boldsymbol{x})$ can be computed from the rates in $\mathcal{X}^{(N)}$. Indeed, for $i \neq j$,

$$q_{ij}^{(N)}(\boldsymbol{x}) = \sum_{\tau \in \mathcal{T}} \left[\frac{|\{\!\{i \to j \in \mathbb{S}_\tau\}\!\}|}{X_i} \frac{\widehat{f}_\tau^{(N)}(\widehat{\mathbf{X}})}{N} \right],$$

where $|\{\!\{i \to j \in \mathbb{S}_\tau\}\!\}|$ is the multiplicity of $i \to j$ in the transition set $\mathbb{S}_\tau$ of τ, i.e. the number of agents that take this transition according to τ. Furthermore, as customary, let $q_{ii}^{(N)}(\boldsymbol{x}) = -\sum_{j \neq i} q_{ij}^{(N)}(\boldsymbol{x})$. Then, since $\widehat{f}_i^{(N)}(\widehat{\mathbf{X}})/N \xrightarrow{N \to +\infty} f_i(\widehat{\mathbf{X}})$, we have that $\boldsymbol{Q}^{(N)}(\boldsymbol{x}) \to \boldsymbol{Q}(\boldsymbol{x})$, where $\boldsymbol{Q}(\boldsymbol{x})$ is computed in terms of the Lipschitz limits $f_i(\widehat{\mathbf{X}})$. Now, define the stochastic processes:

1. $Z^{(N)}(t)$, that describes the state of the process $Y^{(N)}(t)$ for the single agent in class $\mathbb{A}$, when $Y^{(N)}(t)$ is marginalised from $\mathbf{X}^{(N)}(t)$;

2. $Z(t)$, that is the ICTMC, defined on the same state space of $Z^{(N)}(t)$, with *time-dependent* generator matrix $\boldsymbol{Q}(\mathbf{\Phi}(t))$, i.e. the generator matrix $\boldsymbol{Q}(t)$, where the Lipschitz limits $f_i(t)$ are computed over the components of $\mathbf{\Phi}(t)$.

Then, the following theorem can be proved (DN08).

Theorem 2.3 (Fast Simulation) *For any time horizon $T < +\infty$ and $\epsilon > 0$,*

$$Prob\left\{ \sup_{0\leq t\leq T} ||Z^{(N)}(t) - Z(t)|| > \epsilon \right\} \xrightarrow{N\to+\infty} 0.$$

2.4.2 Central Limit Approximation

While the Fluid Approximation correctly describes the transient collective behaviour for very large populations, it is less accurate when one has to deal with a *mesoscopic* system, meaning a system with a population in the order of hundreds of individuals and whose dynamics turn out to be intrinsically probabilistic. Indeed, the (stochastic) behaviour of single agents becomes increasingly relevant as the size of the population decreases. The technique of *Central Limit Approximation* (EK05), also known as *Linear Noise Approximation* (Van92), provides an alternative and more accurate estimation of the stochastic dynamics of mesoscopic systems. In particular, in this technique, the probabilistic fluctuations about the average deterministic behaviour (described by the fluid limit) are approximated by a *Gaussian process*.

Consider the setting (notions and notations) of the Fluid Approximation described in Section 2.4.1. Define the stochastic process given by

$$\mathbf{Z}^{(N)}(t) \;:=\; N^{\frac{1}{2}} \left(\hat{\mathbf{X}}^{(N)}(t) - \mathbf{\Phi}(t) \right),$$

which captures the (rescaled) fluctuations of the CTMC $\hat{\mathbf{X}}^{(N)}(t)$ associated with $\hat{\mathcal{X}}^{(N)}$ around the fluid limit $\mathbf{\Phi}(t)$ given by (2.6). Then, by relying on the theory of Markov Processes, one can prove (Van92; EK05) that, in systems with large population sizes N, $\mathbf{Z}^{(N)}(t)$ can be approximated by a Gaussian process $\{\mathbf{Z}(t) \in \mathbb{R}^n \mid t \in \mathbb{R}\}$ (*independent of* N), whose mean $\mathbf{E}[t]$ and covariance $\mathbf{C}[t]$ are the unique solutions of the following ODE systems,

$$\begin{cases} \frac{\partial \mathbf{E}[t]}{\partial t} = \mathbf{J_F}(\mathbf{\Phi}(t))\mathbf{E}[t] \\ \mathbf{E}[0] = 0 \end{cases} \tag{2.7}$$

and

$$\begin{cases} \frac{\partial \mathbf{C}[t]}{\partial t} = \mathbf{J_F}(\mathbf{\Phi}(t))\mathbf{C}[t] + \mathbf{C}[t]\mathbf{J}_\mathbf{F}^T(\mathbf{\Phi}(t)) + \mathbf{G}(\mathbf{\Phi}(t)) \\ \mathbf{C}[0] = 0, \end{cases} \tag{2.8}$$

where $\mathbf{J_F}(\mathbf{\Phi}(t))$ denotes the Jacobian of the limit drift $\mathbf{F}$ calculated along the deterministic fluid limit $\mathbf{\Phi} : \mathbb{R}_{\geq 0} \longrightarrow \mathbb{R}^n$, and

$$\mathbf{G}(\hat{\mathbf{X}}) = \sum_{\tau \in \hat{\mathcal{T}}^{(N)}} \mathbf{v}_\tau \mathbf{v}_\tau^T f_\tau(\hat{\mathbf{X}})$$

is called the *diffusion* term. Formally, the following Theorem holds true (EK05).

Theorem 2.4 *Let $\mathbf{Z}^{(N)}(t)$ be the random variable given by*

$$\mathbf{Z}^{(N)}(t) \ := \ N^{\frac{1}{2}} \left(\hat{\mathbf{X}}^{(N)}(t) - \mathbf{\Phi}(t) \right),$$

and $\mathbf{Z}(t)$ be the Gaussian process with mean (2.7) and covariance (2.8). Assume that $\lim_{N\to\infty} \mathbf{Z}^{(N)}(0) = \mathbf{Z}(0)$. Then, $\mathbf{Z}^{(N)}(t)$ converges in distribution to $\mathbf{Z}(t)$ ($\mathbf{Z}^{(N)}(t) \Rightarrow \mathbf{Z}(t)$).

In conclusion, the *Central Limit Approximation* of the normalized CTMC

$$\hat{\mathbf{X}}^{(N)}(t) = \mathbf{\Phi}(t) + N^{-\frac{1}{2}} \mathbf{Z}^{(N)}(t)$$

associated with $\hat{\mathcal{X}}^{(N)}$ is the stochastic process

$$\mathbf{\Phi}(t) + N^{-\frac{1}{2}} \mathbf{Z}(t), \tag{2.9}$$

and Theorem 2.4 guarantees that the approximation is correct in the limit of an infinite population.

2.4.3 System Size Expansion

As discussed in the previous section, the Central Limit Approximation (CLA) is an estimation of the behaviour of a population model which is exact in the limit of an infinite population size, but can be efficiently applied even when considering mesoscopic systems. Indeed, the CLA provides a Gaussian estimation of the stochastic fluctuations of the dynamics of the population model around the (deterministic) average behaviour described by the fluid limit. Again, this Gaussian approximation is asymptotically correct, but in the case of mesoscopic populations it can

happen that even the CLA fails to properly describe the dynamics of the population model: the fluid limit itself may indeed fail to accurately describe the average behaviour of the system and/or the stochastic fluctuations around the fluid estimation could be not normally distributed. In these cases, to tackle the error in the estimation of the CLA, *higher-order approximations* of the system behaviour that are referred to as the *System Size Expansion* (SSE) (ABG+15) or *Inverse Omega Square,* (IOS) (Gri10), have been proposed and applied in the literature.

Let us introduce the higher-order approximation of the CTMC $\mathbf{X}^{(N)}(t)$ of a Markov Population Model $\mathcal{X}^{(N)}$. To ease the presentation, we will describe just the fundamental steps of the definition, leaving aside most of the mathematical details (the interested reader can refer to (Gri10) and (Van92)). As in the case of the CLA, we are interested in estimating the (normalised) process $\mathbf{Z}^{(N)}(t) := N^{\frac{1}{2}}(\hat{\mathbf{X}}^{(N)}(t) - \mathbf{\Phi}(t))$, capturing the noise of $\hat{\mathbf{X}}^{(N)}(t)$ around the deterministic Fluid Approximation $\mathbf{\Phi}(t)$ given by (2.6). To achieve this, we write $\hat{\mathbf{X}}^{(N)}(t) = \mathbf{\Phi}(t) + N^{-\frac{1}{2}}\mathbf{Z}^{(N)}(t)$ and we substitute this formula in an expansion in powers of N of the Master Equation associated with $\hat{\mathbf{X}}^{(N)}(t)$ (Nor97), describing the evolution in time of the transient probability $\mathbf{P}(\hat{\mathbf{X}}^{(N)}(t))$ of being in state $\hat{\mathbf{X}}^{(N)}(t) = \mathbf{\Phi}(t) + N^{-\frac{1}{2}}\mathbf{Z}^{(N)}(t)$ at time t. By dropping high order terms in N in the so-obtained form of the Master Equation, we can control the level of accuracy and define different higher-order approximations of $\hat{\mathbf{X}}^{(N)}(t)$. The simplest correction to the CLA, known also as IOS (Gri10), defines a stochastic process $\mathbf{Z}^*(t)$ whose first and second moments are given by

$$\frac{\partial \mathbf{E}^*(t)}{\partial t} = \mathbf{J}(\mathbf{\Phi}(t))\mathbf{E}^*(t) + N^{-1/2}\mathbf{\Delta}(\mathbf{C}^*(t)) + O(N^{-1}), \quad \mathbf{E}^*(0) = 0, \tag{2.10}$$

and

$$\frac{\partial \mathbf{C}^*(t)}{\partial t} = \mathbf{J}(\mathbf{\Phi}(t))\mathbf{C}^*(t) + \mathbf{C}^*(t)\mathbf{J}^T(\mathbf{\Phi}(t)) + \mathbf{G}(\mathbf{\Phi}(t)) + O(N^{-\frac{1}{2}}),\ \mathbf{C}^*(0) = 0, \tag{2.11}$$

where: $\mathbf{F}$ and $\mathbf{G}$ are the drift and diffusion terms, respectively; $\mathbf{J}$ is the Jacobian of $\mathbf{F}$; and $\mathbf{\Delta}(\mathbf{C}^*(t))$ is the vector whose i-th component is given

by

$$\Delta_i(\mathbf{C}^*(t)) = -\frac{1}{2}\left(\sum_{j,k}\frac{\partial^2}{\partial\Phi_j\partial\Phi_k}F_i(\mathbf{\Phi})C^*_{ij} - \sum_j \Phi_j\frac{\partial^2}{\partial{\Phi_j}^2}F_i(\mathbf{\Phi})\right). \quad (2.12)$$

Notice that $\mathbf{Z}^*(t)$ depends on N, and moreover, if in Equation (2.10) we drop the term that is $O(N^{-1/2})$, Equations (2.10) and (2.11) describe the mean and covariance of the Central Limit Approximation (i.e. they correspond to the ODE systems (2.7) and (2.8)). For this reason, we can indeed say that the System Size Expansion is a higher order estimation of the Central Limit Approximation. Furthermore, due to this relation between the CLA and the IOS, we can also state that the quality of the estimation of the IOS is guaranteed by its definition and Theorem 2.4. Hence, we have the following result.

Theorem 2.5 *Let $\boldsymbol{Z}^{(N)}(t)$ be the random variable given by*

$$\boldsymbol{Z}^{(N)}(t) \;:=\; N^{\frac{1}{2}}\left(\hat{\mathbf{X}}^{(N)}(t) - \mathbf{\Phi}(t)\right),$$

and $\boldsymbol{Z}^*(t)$ be the Gaussian process with mean (2.10) and covariance (2.11). Assume that $\lim_{N\to\infty}\boldsymbol{Z}^{(N)}(0) = \boldsymbol{Z}^*(0)$. Then, $\boldsymbol{Z}^{(N)}(t)$ converges in distribution to $\boldsymbol{Z}^*(t)$ ($\boldsymbol{Z}^{(N)}(t) \Rightarrow \boldsymbol{Z}^*(t)$).

2.4.4 Moment Closure and Maximum Entropy Principle

The last approximation technique that we consider in this work is the Moment Closure or Method of Moments combined with the Maximum Entropy Principle (AMW15a; SSG16; ABG+15). This technique gives an estimation of the probability distribution of the CTMC that describes the evolution of a population model, starting from an approximation of its moments. With this aim, let $\mathbf{P} : \mathbb{R}^{\geq 0} \to [0,1]^S \times [0,1]^S$, $\mathbf{P}(\mathbf{X}(t))$, be the transient probability distribution of the CTMC describing the evolution of a Markov Population Model $\mathcal{X}^{(N)} = (\mathbb{A}, \mathcal{T}^{(N)}, \mathbf{x}_0^{(N)})$, given a time instant $T \in \mathbb{R}^{\geq 0}$. The approximation technique we are going to describe follows of two steps:

(1) *Estimation of the moments of* $\mathbf{P}(\mathbf{X}(T))$ - Based on the Dynkin Formula, an estimation of moments of $\mathbf{P}(\mathbf{X}(T))$ of order lower than $K \geq 0$ is computed as the solution of a finite set of ODEs.

(2) *Definition of the approximation of* $\mathbf{P}(\mathbf{X}(T))$ - Relying on the Maximum Entropy Principle, the approximation of $\mathbf{P}(\mathbf{X}(T))$ is given by the probability distribution that has the first K moments defined in Step (1) and that maximizes the Shannon entropy of the system.

In the following, we review some of the mathematical background that defines the two steps of this approximation technique. We are going to omit the details regarding the functional analysis results that guarantee some of the results exploited in the following. For a full review see (AMW15a; SSG16; ABG+15).

Estimation of the moments of $\mathbb{P}[\mathbf{X}(T)]$. Let $\tau \in \mathcal{T}^{(N)}$ be the transitions of the Markov Population Model $\mathcal{X}^{(N)}$, each with rate $f_\tau^{(N)}(\mathbf{X}(t))$ and update vector $\mathbf{v}_\tau$. According to the Dynkin Formula (Kal06; AKS13), the moments of $\mathbf{P}(\mathbf{X}(t))$ satisfy the following relation:

$$\frac{d}{dt}\mathbf{E}\left[h(\mathbf{X}(t))\right] = \sum_\tau \mathbf{E}\left[\left(h(\mathbf{X}(t)+\mathbf{v}_\tau) - h(\mathbf{X}(t))\right) f_\tau^{(N)}(\mathbf{X}(t))\right] \quad (2.13)$$

where $h : \mathbb{R}^{(N)} \to \mathbb{R}^{(N)}$ is a suitable sufficiently smooth function. Starting form Equation (2.13), we can easily obtain the exact ODEs for the moments of $\mathbf{P}(\mathbf{X}(t))$ by choosing the right polynomial form for the function h. For example, given the identity $h(\mathbf{X}(t)) = \mathbf{X}(t)$, we obtain an ODE system that is satisfied by the *mean* of $\mathbf{P}(\mathbf{X}(t))$, indeed we have:

$$\frac{d}{dt}\mathbf{E}\left[\mathbf{X}(t)\right] = \sum_\tau \mathbf{v}_\tau \mathbf{E}\left[f_\tau^{(N)}(\mathbf{X}(t))\right]. \quad (2.14)$$

Notice that system (2.14) is closed and fully integrable when the rate functions $f_\tau^{(N)}(\mathbf{X}(t))$ are linear. If, instead, we consider transition functions that are more than linear (like the infection f_{inf} in the running Example (3)), higher order moments appear in system (2.13), that is no more closed. Differential equations for the covariances $\mathbf{C}\left[\mathbf{X}(t)\right]$ can be

obtained from (2.13) setting $h(\mathbf{X}(t)) = \mathbf{X}^2(t)$ and letting $\mathbf{C}[\mathbf{X}(t)] = \mathbf{E}\left[\mathbf{X}^2(t)\right] - \mathbf{E}[\mathbf{X}(t)]^2$, but again, the system of the Dynkin Formula is not closed when the transition functions are more than linear, as it comprises moments of order greater than 2.

In general, the Dynkin Formula (2.13) is an infinite hierarchy of ODEs that couples lower and higher order moments. *Moment Closure Approximations* truncate (and close) this infinite set of equations at a certain order K in an appropriate way in order to be able to integrate the system at time $T \in \mathbb{R}^{\geq 0}$ and find an approximation of the first K moments of $\mathbf{P}(\mathbf{X}(T))$. Many different strategies of closure have been proposed in the literature, but in this work, as in the experimental analysis we are going to exploit in the tools STAR (LMW11) and CERENA (KFR$^+$16), we are going to rely on the *Low Dispersion Moment Closure*, where all the central moments of $\mathbf{P}(\mathbf{X}(T))$ above order K are set to 0.

Definition of the approximation of $\mathbf{P}(\mathbf{X}(T))$. Given the estimation of the first K moments of $\mathbf{P}(\mathbf{X}(T))$ obtained in the previous step, we can now reconstruct an approximation of $\mathbf{P}(\mathbf{X}(T))$. To do this, we exploit the *Maximum Entropy Principle* (A$^+$10; AMW15a; AMW15b) and among all the probability distributions that share the moments of order K computed in Step (1), we choose as best approximation of $\mathbf{P}(\mathbf{X}(T))$ the distribution $\mathbf{P}^*(\mathbf{X}(T))$ that maximizes the Shannon entropy of the system.

Formally, to ease the notation consider a one dimensional Markov Population Model with one counting variable $X : \mathbb{R}^{\geq 0} \to \mathbb{R}^{\geq 0}$ and let $\mu^{(k)}$, $k = 0, 1, \ldots, K$, be the moments of $\mathbf{P}(X(T)) = \mathbf{P}(x)$ of order lower than K computed in Step (1). Let $\mathcal{G}$ be the set of non-negative distribution function g that share the non-central moments $\mu^{(k)}$, i.e. such that

$$\mathbf{E}_g[x^k] = \int x^k g(x)dx = \mu^{(k)}, \qquad k = 0, 1, \ldots, K, \quad \mu_0 = 1. \tag{2.15}$$

In accordance with the Maximum Entropy Principle, we choose as best approximation $\mathbf{P}^*(x)$ of $\mathbf{P}(x)$, the probability distribution that maximises the *Shannon Entropy* $H(g)$, i.e.

$$\mathbf{P}^*(x) = \arg\max_{g \in \mathcal{G}} H(g) = \arg\max_{g \in \mathcal{G}} \left(- \int g(x) \ln g(x) dx \right). \tag{2.16}$$

Remark 2.2 *The Shannon Entropy $H(g)$ is related to the amplitude of the stochastic fluctuations (or noise) of the stochastic process whose probability distribution is $\mathbf{P}(x)$. In particular, by choosing $\mathbf{P}^*(x)$ as the probability distribution that maximizes the Shannon Entropy, we actually choose, among all the stochastic processes described by $\mathbf{P}(x)$, the one that has the maximum noise, hence we put the minimum number of assumptions on the behaviour of the process described by $\mathbf{P}^*(x)$.*

In order to compute the approximation $\mathbf{P}^*(x)$ defined by Equation 2.16, we need to solve a non-linear constrained optimisation problem, and to achieve this, we leverage the functional theory of the *Lagrangian*. It is beyond the scope of this thesis to go in deep detail of the procedure that gives as a result the right formula for $\mathbf{P}^*(x)$, but in the following we try to give a general idea of the theory that leads us to the definition of $\mathbf{P}^*(x)$. The interested reader may refer to (A^+10; AMW15a) for a full review.

The Kuhn-Tucker Theorem states that, if a solution $\mathbf{P}^*(x)$ of the non-linear constrained optimisation problem (2.16) given (2.15) exists, then this solution is a stable point for the *Lagrangian Function* given by

$$\mathcal{L}(g, \lambda) = H(g) - \sum_{k=0}^{K} \lambda_k \left(\int x^k g(x) dx - \mu^{(k)} \right). \tag{2.17}$$

In the above function, the vector $\lambda = (\lambda_1, \ldots, \lambda_K)$ is called the vector of the *Lagrangian Multipliers*. Setting to zero all the first derivatives of $\mathcal{L}(g, \lambda)$ with respect to g, we obtain that the stable points of $\mathcal{L}(g, \lambda)$ have the form

$$q(x, \lambda) = \exp\left(-1 - \sum_{k=0}^{K} \lambda_k x^k \right). \tag{2.18}$$

Among all the stable points that satisfy (2.18), we need to identify the function $\mathbf{P}^*(x)$, i.e. we need to find the right vector of Lagrangian multipliers $\lambda^* = (\lambda_1^*, \ldots, \lambda_K^*)$, such that $\mathbf{P}^*(x) = q(x, \lambda^*)$ and $\mathbf{P}^*(x)$ is indeed the maximum of the constrained problem (2.16) given (2.15). To do this, we exploit a corollary of the Kuhn-Tucker Theorem, that states that the vector of Lagrangian multipliers $\lambda^* = (\lambda_1^*, \ldots, \lambda_K^*)$ is the solution of the

dual problem given by

$$\lambda^* = \arg\min \Psi(\lambda), \tag{2.19}$$

where $\Psi(\lambda)$ is the dual function of $H(q)$ given the form of $q(x, \lambda)$ obtained in (2.18). The dual problem (2.19) is an unconstrained convex minimization problem and can be numerically solved by considering that its solution $\lambda^* = (\lambda_1^*, \ldots, \lambda_K^*)$ is the point where the first derivatives of $\Psi(\lambda)$ are zero and its Hessian is positive definite. The final value of λ^* defines the approximation $\mathbf{P}^*(x)$ of the distribution $\mathbf{P}(T)$ as the following function

$$\mathbf{P}^*(x) = q(x, \lambda^*) = \exp\left(-1 - \sum_{k=0}^{K} \lambda_k^* x^k\right).$$

Remark 2.3 *Theoretically, the greater the order K chosen in Step (1), the more accurate should be the estimate of the lower moments, like the mean. In reality, this is not always the case as the accuracy of the Moment Closure actually depends (in complex ways) on the structure of the population model under consideration. Indeed, among the possible drawbacks of this approach there is the fact that the ODEs for higher order moments defined in Step (1) tend to be stiff and difficult to integrate. Moreover the multidimensional optimisation problem of Step (2) could be also unstable for complex systems. For a more detailed discussion in this sense, see (ABG+15).*

2.5 Formalization of Behavioural Properties

In the model checking procedure, the model of the system under consideration has to be accompanied with a specification of the property of interest that has to be verified. The properties of interest in this work can be classified under the following general paradigms.

Stochastic properties

In this project we focus on stochastic model checking, considering stochastic models of collective systems, and thus we consider *stochastic properties*

to specify their behaviour. In particular, instead of focusing on the absolute correctness of a requirement ("the system will not fail"), we are interested in the likelihood (probability) of occurrence of the requirement ("with 95% chance the system will not fail").

Individual, global and local-to-global properties

When we consider a population model the analysis of its behaviour can follow three different approaches: we can specify the evolution of a single agent (*individual properties*); we can look at the behaviour of the entire population (*global properties*) or we can choose to characterise the actions of the individuals in the global context (looking at the fraction of population that satisfies a given individual property). We refer to the last class of requirements as the *local-to-global properties*. While the collective behaviour of a large system is almost deterministic, the evolution of a single agent is always intrinsically stochastic. For this reason, different approaches and methodologies have to be investigated and developed for each of these class of properties.

Timed and time bounded properties

We are interested in the evaluation of the evolution of population models in time. A system is said to be *time-critical* if its behaviour is subject to timing constraints (e.g. requirements over the residence time in a state or the possibility of taking a transition within a particular time interval). In this context, the *timed properties* that are used to specify the behaviour of the system can assume or enforce the existence of a *time horizon* within which something should happen (think about processes subject to deadlines or timeouts). These constraints fall into the class of *time bounded properties*. This classification is particularly relevant to this work, since, as shown in Section 2.4, in most cases the results of the limit theorems for the Fluid, the Central Limit and the System Size Expansion approximations hold true only for a finite time horizon $T > 0$.

Safety, invariant, liveness and reachability properties

These are among the most classical and well-known categories of properties. Safety and liveness properties can be used to specify the behaviour of reactive systems and their definition goes back to the seventies (Lam77). *Safety properties* state that "nothing bad should happen" (a typical example is the mutual exclusion problem: always at most one process is in its critical section). *Invariants* are particular safety properties that require that a particular condition holds for all reachable states (if we consider the well known problem of the Dining Philosophers the constraint "at least one philosopher is not waiting to pick the chopsticks" is an example of an invariant). *Liveness propeties*, instead, specify that "something good will eventually happen" (an example in concurrency requires that a process will enter its critical section infinitely often). For a survey see (Kin94). Finally, we should remark that in the stochastic validation framework, both the safety and liveness properties, and many other interesting requirements, can be grouped under the fundamental and most discussed class of the *stochastic reachability properties*. In these requirements, we measure the probability associated with the system computations that "visit" (i.e. reach) a certain target subset of the state space, while avoiding some other specific regions. In this sense, a liveness property, in which we state that the system will eventually reach a good configuration, is a reachability requirement. And the same is true for safety properties, where the state of the system should be kept in a good region of the space (avoiding the bad configurations). For a good review on some of the methodologies related to this subject see (Buj12).

2.5.1 Specification of Local Properties: the DTA

In Chapters 3 and 5, we are interested in properties specifying how the agents in a population model behave in *time*. In order to monitor these requirements, we assign to them a *personal* or a *global clock*, which start at time 0 and can be either reset whenever the agent undergoes specific transitions or not. The personal clock is used to keep track of the actions of a single agent, thus it can be reset as it is influenced by the actions

taken by the single individual; the global clock, instead, monitors the behaviour of the agents within the population, recording the time at global level, and thus it cannot be influenced (or reset) by the actions taken by the single agents. Hence, in this work, we exploit *personal clocks* to validate *local properties* characterising the behaviour of single agents, while we consider *global clocks* to monitor *local-to-global properties*, that describes the actions of the single agents in the global context of the population.

In this section, we introduce the specification language that we use to express and monitor the local (timed) properties considered in Chapter 5: the *single-clock Deterministic Timed Automata* (DTA)(AD94; CHKM11a). We start from this formal setting, because, as we shall see in Chapter 3, the specification of the local-to-global requirements is very similar and in some way, it can be seen as an instance of the DTAs introduced in this section, where no reset is allowed.

The *single-clock Deterministic Timed Automata* (DTA) exploited in Chapter 5 to specify local properties keeps track of the behaviour of the single agent with respect to a personal clock, that, as we have just said, starts at time 0 and can be reset whenever the agent undergoes specific transitions. Moreover, since we want to exploit the Stochastic Approximations illustrated in Section 2.4, we restrict ourselves to *time bounded* properties and, hence, we assign to the DTA a finite *time horizon* $T < +\infty$, within which the requirement must be true. Considering the formal settings of the Markov Population Models of Section 2.3, we introduce the following definition.

Definition 2.3 (Timed Properties) *A timed property for a single agent in Agent Class $\mathbb{A}$ is specified as a single-clock DTA of the form*

$$\mathbb{D} = \mathbb{D}(T) = (T, \mathscr{L}, c, \mathcal{CC}, Q, q_0, F, \rightarrow),$$

where $T < +\infty$ is the finite time horizon; $\mathscr{L}$ is the label set of $\mathbb{A}$; c is the personal clock; $\mathcal{CC}$ is the set of clock constraints, which are conjunctions of atoms of the form $c < \lambda$, $c \leq \lambda$, $c \geq \lambda$ or $c > \lambda$ for $\lambda \in \mathbb{Q}$; Q is the (finite) set of states; $q_0 \in Q$ is the initial state; $F \subseteq Q$ is the set of final (or accepting) states; and $\rightarrow \subseteq Q \times \mathscr{L} \times \mathcal{CC} \times \{\varnothing, \{c\}\} \times Q$ is the edge relation. Moreover, $\mathbb{D}$ has to satisfy:

- *(determinism) for each initial state $q \in Q$, label $\alpha \in \mathscr{L}$, clock constraint $c_{\bowtie} \in \mathcal{CC}$, and clock valuation $\eta(c) \in \mathbb{R}_{\geq 0}$, there exists exactly one edge $q \xrightarrow{\alpha, c_{\bowtie}, r} q'$ such that $\eta(c) \models_{\mathcal{CC}} c_{\bowtie}$[1];*
- *(absorption) the final states are all absorbing.*

The timed property $\mathbb{D}$ is assessed over the time-bounded paths (of total duration T) of the Agent Class $\mathbb{A}$ sampled from the stochastic processes $Y^{(N)}(t)$ that describes the state of the agent belonging to $\mathbb{A}$. The labels of the transitions of $\mathbb{A}$ act as inputs for the DTA $\mathbb{D}$, and the latter is defined in such a way that it *accepts* a time-bounded path σ if and only if the behaviour of the single agent encoded in σ satisfies the property represented by $\mathbb{D}$. Formally, a time-bounded path $\sigma = s_0 \xrightarrow{\alpha_0, t_0} s_1 \xrightarrow{\alpha_1, t_1} \dots \xrightarrow{\alpha_n, t_n} s_{n+1}$ of $\mathbb{A}$ sampled from $Z^{(N)}(t)$, with $\sum_{j=0}^{n} t_j \leq T$, is *accepted* by $\mathbb{D}$ if and only if there exists a path $q_0 \xrightarrow{\alpha_0} q^{(1)} \xrightarrow{\alpha_1} q^{(2)} \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} q^{(n+1)}$ of $\mathbb{D}$ such that $q^{(n+1)} \in F$. In the path of $\mathbb{D}$, $q^{(i+1)} \in Q$ denotes the (unique) state that can be reached from $q^{(i)} \in Q$ taking the action $q^{(i)} \xrightarrow{\alpha_i, c_{\bowtie}, r} q^{(i+1)}$ whose clock constraint $c_{\bowtie}$ is satisfied by the clock valuation $\eta(c)$ updated according to time t_i. In the following, we will denote by $\Sigma_{\mathbb{A},\mathbb{D},T}$ the *set of time-bounded paths of $\mathbb{A}$ accepted by $\mathbb{D}$*.

[1]The notation $\eta(c) \models_{\mathcal{CC}} c_{\bowtie}$ stands for the fact that the value of the valuation $\eta(c)$ of c satisfies the clock constraint $c_{\bowtie}$.

Chapter 3

Stochastic Approximations for Local-to-Global Properties

3.1 Overview

The first Stochastic Model Checking procedures based on Fluid Approximations (FA) of the behaviour of the system appeared in the literature only a few years ago (BH12b; HSB12; HBC13). In (BH12b), the authors exploit FA to construct an approximate model of a single individual agent in a (large) population, and check efficiently Continuous Stochastic Logic (CSL) properties for that individual. A similar approach is taken in (HBC13), restricting to path properties specified by Deterministic Finite Automata (DFA, (BK08)). In (HSB12; HBC13), the authors consider also global properties concerned with the fraction of agents satisfying local specifications, using moment closure techniques to find approximate bounds on the associated probabilities.

In this chapter, we continue along this direction, focussing on the lifting of local specifications to the global level, but using different Stochastic Approximations to provide a more accurate estimate of the satisfaction probabilities: the *Central Limit Approximation* (CLA (EK05)), also

known as Linear Noise Approximation (Van92), the *System Size Expansion* (Gri10), and the *Moment Closure* technique combined with the *Maximum Entropy Principle* (AMW15b). In this respect, our approach complements that of (HSB12; HBC13). We also consider a richer class of path properties, expressed by Deterministic Timed Automata (DTA) with 1 global clock, i.e. a clock referring to the global time of the model. Hence, this work goes in the direction of merging the approaches of (BH12b) and (HSB12; HBC13) in the light of the logics asCSL (BCH$^+$07) or CSL-TA (DHS09), in which until path properties of CSL are replaced by DFA or DTA specifications. The link between local and global properties, with exclusive focus on average collective properties estimated using the fluid limit, has also been discussed in a logical setting in (KRdH13).

The chapter is organised as follows. In Section 3.2, we discuss the DTA specification of local properties and their lifting to the global level. In Section 3.3, we discuss how to combine a population model and a DTA specification into a larger sequence of population models, which is the key step of the algorithm of Section 3.4.1, based on the Central Limit Approximation. In the Section 3.4.2, we also discuss higher order approximations like the System Size Expansion and the Moment Closure combined with the Maximum Entropy distribution reconstruction. Finally, in Section 3.5 we discuss the quality of the approximations, and, in Section 3.6, we discuss the result and the future prospectives.

Part of the content of this chapter has beed published in (BL13a).

3.2 Local-to-Global Properties

We start the presentation of our model checking procedure by considering the theoretical framework and notation of Section 2.3. Hence, we consider a Markov Population Model $\mathcal{X}^{(N)} = (\mathbb{A}, \mathcal{T}^{(N)}, \mathbf{x}_0^{(N)})$ of size N, and in this section, we introduce the class of properties and the specification languages that we want to exploit.

As it was described in Section 2.5, we distinguish two levels of properties: *local properties*, describing the behaviour of individual agents, and *global properties*, describing the collective behaviour of agents with re-

spect to a local property of interest. In this classification, our approach is similar to (KRdH13; HSB12). Moreover, we focus *time-bounded* local properties specified by Deterministic Timed Automata (DTA). The restriction to finite time horizons is justified because the analysis of steady state properties is always problematic in the context of Stochastic Approximations (see (BH12b; BHLM13; HSB12) for further discussion on this point).

The global property layer, instead, allows us to specify queries about the fraction of agents that satisfies a given local specification. In particular, given a (local and time-bounded) path property ϕ, we want to compute the probability that the fraction of agents that satisfies ϕ at time T is smaller or larger than a threshold α. In the following, these requirements are captured by an appropriate operator, that can then be combined to specify more complex global queries, as in (KRdH13).

Formally, consider the theoretical setting (and the notation) of Section 2.3, and let us fix a Markov Population Model composed of N agents belonging to a class $\mathbb{A} = (S, E)$. We consider local path properties specified by *1-global-clock Deterministic Timed Automata* (1gDTA), which are DTAs similar to those introduced in Section 2.5.1, but are endowed with one single clock variable $x \in \mathbb{R}_{\geq 0}$, called *global clock*, that is never reset. As in Definition 2.3, we call $\mathcal{V}$ the set of *valuations of* x, i.e. functions $\eta : \{x\} \longrightarrow \mathbb{R}^{\geq 0}$ that assign a nonnegative real-value to the global clock x, and $\mathcal{CC}$ the set of *clock constraints*, which are positive boolean combinations of basic clock constraints of the form $x \leq a$ or $x \geq a$, where $a \in \mathbb{Q}^{\geq 0}$. We write $\eta(x) \models_{\mathcal{CC}} c$ if and only if $c \in \mathcal{CC}$ is satisfied when the clock variable takes the value $\eta(x)$. In addition to actions and clock constraints, we also label the edges of 1gDTA by a boolean formula, interpreted on the states $s \in S$ of agent $\mathbb{A}$, similarly to asCSL (BCH$^+$07) and CSL-TA (DHS09). Let $\Gamma_{\mathcal{S}}$ be the set of these *(atomic) state propositions over* S, and $\mathcal{B}(\Gamma_{\mathcal{S}})$ the set of boolean combinations over $\Gamma_{\mathcal{S}}$. We use the letter ϕ to range over formulae in $\mathcal{B}(\Gamma_{\mathcal{S}})$ and we denote by $\models_{\Gamma_{\mathcal{S}}}$ the satisfaction relation over $\mathcal{B}(\Gamma_{\mathcal{S}})$-formulae. In this way, a local transition $s \xrightarrow{\alpha_\tau} s'$ matches an edge with label α, c, ϕ in the 1gDTA if and only if the action name is the same, the clock constraint c is satisfied and the

$\mathcal{B}(\Gamma_{\mathcal{S}})$-formulae holds on the initial state s, i.e. $\alpha_\tau = \alpha$, $\eta(x) \models_{\mathcal{CC}} c$, and $s \models_{\Gamma_{\mathcal{S}}} \phi$. We obtain the following definition.

Definition 3.1 (1-global-clock DTA) *A 1-global-clock Deterministic Timed Automaton (1gDTA) is specified by the tuple*

$$\mathbb{D} = (\mathscr{L}, \Gamma_{\mathcal{S}}, Q, q_0, F, \rightarrow)$$

where:

- $\mathscr{L}$ *is the label set of* $\mathbb{A}$*;*
- $\Gamma_{\mathcal{S}}$ *is the set of atomic state propositions;*
- Q *is the (finite) set of states of the DTA, with initial state* $q_0 \in Q$*;*
- $F \subseteq Q$ *is the set of final (or accepting) states;*
- $\rightarrow \subseteq Q \times \mathscr{L} \times \mathcal{B}(\Gamma_{\mathcal{S}}) \times \mathcal{CC} \times Q$ *is the edge relation, where* $(q, \alpha, \phi, c, q') \in \rightarrow$ *is usually denoted by* $q \xrightarrow{\alpha,\phi,c} q'$*.*

Moreover, $\mathbb{D}$ *satisfies:*

- *(determinism) for each* $q \in Q$*,* $\alpha \in \mathscr{L}$*,* $s \in S$ *and clock valuation* $\eta(x) \in \mathbb{R}_{\geq 0}$*, there is exactly one edge* $q \xrightarrow{\alpha,\phi,c} q'$ *such that* $s \models_{\Gamma_{\mathcal{S}}} \phi$ *and* $\eta(x) \models_{\mathcal{CC}} c$*;*
- *(absorption) the final states* F *are all absorbing, i.e. they only have looping transitions out of them.*

When we write a 1gDTA, we stick to the convention that all non-specified edges are self-loops on the automata states. Hence, given α, s, and $\eta(x)$, if there is no specified edge from state q with label α, with formula satisfied by s and clock constraint satisfied by $\eta(x)$, then we assume the existence of an edge looping on q and satisfying all conditions.

Example 3.1 *As a running example for this chapter, we consider the Agent Class and the Markov Population Model of the network epidemic model of Figure 3 of Section 2.3, and the 1gDTA specification at the top of Figure 4 (a), where the formula* at_S *is true in local state* S *and false in states* I *and* R*. The property is satisfied when a susceptible node is infected by an internal infection after the first* τ *units of time. The sink state* q_b *is used to discard agents infected before* τ

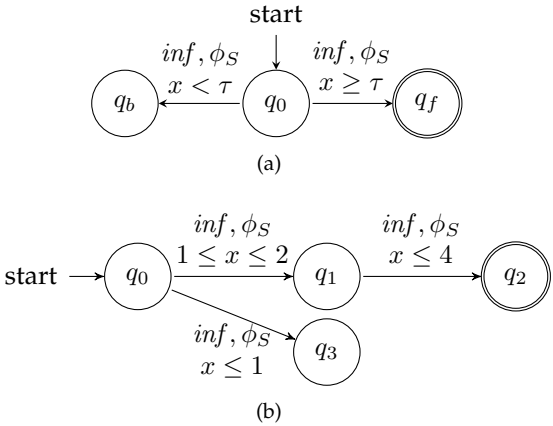

Figure 4: The 1gDTA specification discussed in the running example.

units of time. The use of the state formula ϕ_S allows us to focus only on agents that get infected, rather than also on agents that spread the contagion.

Exploiting the same specification language, we can of course define more complex requirements like the automaton of Figure 4 (b). This local timed property states that an agent is infected by internal contact twice, the first infection happening between time 1 and 2, and the second infection happening before time 4. The sink state q_3 is used again to discard agents being infected for the first time before time 1.

A *run* ρ of a 1gDTA $\mathbb{D}$ is a sequence

$$q_0 \xrightarrow{\alpha_0,t_0} q_1 \xrightarrow{\alpha_1,t_1} \dots q_n,$$

where $q_0, q_1, q_2 \in Q$ are states, α_0, α_1 are actions and t_0, t_1 correspond to the times taken by α_0, α_1. Moreover, we require for t_0, t_1 to satisfy the clock constraints. Finally, a run is said to be *accepting* if $q_n \in F$.

Consider now a Markov Population Model $\mathcal{X}^{(N)}$, and focus on a single individual agent of class $\mathbb{A}$ in the population. A path σ of length n for the agent is a sequence of the form

$$s_0 \xrightarrow{\alpha_0,t_0} s_1 \xrightarrow{\alpha_1,t_1} s_2 \xrightarrow{\alpha_2,t_2} \dots s_n,$$

where $s_i \in S$, $t_i \in \mathbb{R}_{\geq} 0$ is the time spent in the local state s_i, and α_i is the action taken at step i. The set of those paths will be denoted by $Path^n[\mathbb{A}]$ and the set of paths of finite length will be indicated by $Path^*[\mathbb{A}]$. Given σ, we let $\tau[\sigma] = \sum_{i=0}^{|\sigma|-1} t_i$ be the total time taken to go from state s_0 to state s_n, and with $\tau_i[\sigma]$ the time taken to reach state s_i. The set of paths of total duration equal to $T \in \mathbb{R}_{\geq 0}$ is denoted by $Path^T[\mathbb{A}]$. Given a path σ of length n, we define the run ρ_σ of a 1gDTA $\mathbb{D}$ induced by σ to be the sequence

$$q_0 \xrightarrow{\alpha_0, t_0} q_1 \xrightarrow{\alpha_1, t_1} \dots q_n,$$

where state q_{i+1} is determined by the unique transition

$$q_i \xrightarrow{\alpha_i, \phi, c} q_{i+1},$$

such that

$$s_i \models_{\Gamma_\mathcal{S}} \phi \qquad \text{and} \qquad T_{i+1}[\sigma] \models \mathcal{CC}.$$

If ρ_σ is accepting, we write $\sigma \models \mathbb{D}$.

Given a 1gDTA $\mathbb{D}$, we denote it by $\mathbb{D}[\phi_1, \dots, \phi_k]$, when we want to explicitly list all the atomic propositions $\Gamma_\mathcal{S}$ used to build the state propositions $\mathcal{B}(\Gamma_\mathcal{S})$.

Definition 3.2 (CSL-TA) *A CSL-TA formula Φ on a Agent Class $\mathbb{A}$ is defined recursively as*

$$\texttt{true} \mid a \mid \neg\Phi \mid \Phi_1 \wedge \Phi_2 \mid \mathbf{P}_{\bowtie p}^{\leq T}\left(\mathbb{D}[\Phi_1, \dots, \Phi_k]\right),$$

where

- *a is an atomic proposition interpreted on S;*
- *$T \in \mathbb{R}_{\geq 0}$ is the time horizon;*
- *$\bowtie p$ is the bound on the probability, with $p \in [0, 1]$ and $\bowtie \in \{<, \leq, \geq, >\}$;*
- *$\mathbb{D}[\Phi_1, \dots, \Phi_k]$ is a 1gDTA with atomic formulae taken to be CLS-TA formulae $\Phi_1, \dots, \Phi_k$.*

This definition is similar to (DHS09), with the only difference being the use of a restricted class of DTA, and the time bound on the probability

operator T. The satisfaction relation is defined relatively to state $s \in S$ of an individual agent $Y(t)$ in $\mathcal{X}^{(N)}$ of class $\mathbb{A}$ and an initial time t_0. The only interesting case is the one involving 1gDTA specifications, for which the relation is

$$s, t_0 \models \mathbf{P}^{\leq T}_{\bowtie p}\left(\mathbb{D}[\Phi_1, \ldots, \Phi_k]\right) \text{ iff } \mathbb{P}\{\sigma \in Path^T[\mathbb{A}] \mid \sigma \models \mathbb{D}[\Phi_1, \ldots, \Phi_k]\} \bowtie p$$

An individual agent in a population model satisfies the local property specified by a 1gDTA $\mathbb{D}$ at time T if, feeding to $\mathbb{D}$ the agent trajectory up to time T, we reach a final state. This can be formalised in a standard way, see for instance (CHKM11a; DHS09). In order to lift these local specifications to the collective level, we count the number of agents that satisfy the 1gDTA $\mathbb{D}$ at time T. More specifically, we check if the fraction of agents satisfying $\mathbb{D}$ is included in the interval $[a, b]$, which we write as $\mathbb{D}(T) \in [a, b]$, where the bounds a, b are specified in terms of the fraction of agents or population density (the number of agents divided by the total population size). To verify the random event $\mathbb{D}(T) \in [a, b]$, we compute its probability, which is then compared with a given threshold. The atomic global properties can be combined together by boolean operators, as in (KRdH13), to define more expressive queries.

Definition 3.3 (Syntax of global properties) *Given a Markov Population Model $\mathcal{X}^{(N)}$, a collective/global property on $\mathcal{X}^{(N)}$ is given by the following syntax:*

$$\Psi = \texttt{true} \mid \mathbb{P}_{\bowtie p}(\mathbb{D}(T) \in [a, b]) \mid \neg\Psi \mid \Psi_1 \wedge \Psi_2,$$

where $\mathbb{P}_{\bowtie p}(\mathbb{D}(T) \in [a, b])$ is true if and only if $q \bowtie p$, for $\bowtie \in \{<, \leq, \geq, >\}$, with q being the probability that at time T the number of agents that satisfies the local path property $\mathbb{D}$ is contained in the interval $[a, b]$.

As an example, consider again the 1gDTA property $\mathbb{D}$ of Figure 4 (a). The atomic global property $\mathbb{P}_{\geq 0.8}(\mathbb{D}(4) \leq \frac{1}{3})$ specifies that, with probability at least 0.8, less than one third of network nodes will be infected after 4 time units by an internal contact.

Remark 3.1 *In addition to path properties specified by 1gDTA, we could have considered state properties in the style of CSL-TA (DHS09). This can be done at the price of dealing with nesting of path and state properties, which for local specifications raises issues of time-dependency of truth values similar to those discussed in (BH12b). We leave this for future work.*

Remark 3.2 *The fact that final states are absorbing implies that we are looking for properties in which an accepting state of the 1gDTA must be reached at a time instant within* $[0, T]$*. Punctual properties, looking at satisfaction exactly at time* T*, can be obtained by dropping the absorbing condition in Definition 3.1.*

Remark 3.3 *In Definition 3.2, we allow the arbitrary nesting of CSL properties within 1gDTA. By the discussion of (BH12b), this operation requires some care. The problem is that individual agents are non-Markov processes (in fact, they are projections of Markov processes, the global model), for which the satisfaction of a CSL-TA formula involving the probability quantifier depends on the initial time at which the formula is evaluated. Hence, the satisfaction of a CSL-TA formula is a time-sdependent function, while 1gDTA require time independent state formulae. This discrepancy can be reconciled by encoding this time dependency in the 1gDTA using clock constraints. Hence, a state formula that is true in* s *up to time 5 and false afterwards, will give rise to two edges in the 1gDTA, the first considering a state formula in which* s *is true, and with an additional clock constraint* $x \leq 5$*, while the second corresponding to an edge with* s *false in its state formula, and additional clock constraint* $x > 5$*.*

However, for simplicity, in this thesis we will consider only non-nested CSL-TA formulae, leaving the formal definition of the so-modified 1gDTA for future work.

3.3 Model-Property Synchronization

In this section, we present the model checking procedure for the verification of global atomic properties. We aim at approximating these probabilities by means of central limit results (EK05; Van92). The first step is to synchronize the agent and the property, constructing an extended Markov Population Model in which the state space of each agent is combined with the specific path property we are observing. The Central Limit Approximation is then applied to the so-obtained model.

The main difficulty in this procedure is the presence of time constraints in the path property specification. However, thanks to the restriction to a single global clock, we can partition the time interval of interest into a finite set of subintervals, within which no clock constraint changes status. Thus, in each subinterval, we can remove the clock con-

straints, deleting all the edges that cannot fire since their clock constraint false. In this way, we generate a sequence of Deterministic Finite Automata (DFA), that are then combined with the local model $\mathbb{A}$ by a standard product of automata. Then, we construct the Markov Population Models associated with the local model (paying attention to the rates) and we obtain a *sequence* of population CTMC models to which we apply the Central Limit Approximation.

Let $\mathbb{A} = (S, E)$ be an Agent Class, $\mathbb{D} = (\mathscr{L}, \Gamma_{\mathcal{S}}, Q, q_0, F, \rightarrow)$ be a local path property, and $T > 0$ be the time horizon.

First step: uniqueness of transition labels. We define a new Agent Class $\bar{\mathbb{A}} = (S, \bar{E})$ by renaming the local transitions in E to make their labels unique. This allows us to remove edge formulae in $\mathbb{D}$, simplifying the product construction. In particular, if there exist

$$s_1 \xrightarrow{\alpha} s_1', \ldots, s_m \xrightarrow{\alpha} s_m' \in E$$

having the same label α, we rename them by $\alpha_{s_1}, \ldots, \alpha_{s_m}$, obtaining

$$s_1 \xrightarrow{\alpha_{s_1}} s_1', \ldots, s_m \xrightarrow{\alpha_{s_m}} s_m' \in \bar{E}.$$

The 1gDTA $\mathbb{D}$ is updated accordingly, by substituting each edge $q \xrightarrow{\alpha,\phi,c} q'$ with the set of edges $q \xrightarrow{\alpha_{s_i},\phi,c} q'$, for $i = 1, \ldots, m$. We call $\bar{\mathscr{L}}$ the label set of $\bar{\mathbb{A}}$.

Second step: removal of state conditions. We remove from the edge relation of $\mathbb{D}$ all the edges $q \xrightarrow{\alpha_{s_i},\phi,c} q'$ such that $s_i \not\models_{\Gamma_{\mathcal{S}}} \phi$, where s_i is the source state of the (now unique) transition of $\bar{\mathbb{A}}$ labeled by α_{s_i}. At this point, the information carried by state propositions becomes redundant, thus we drop them, writing $q \xrightarrow{\alpha_{s_i},c} q'$ in place of $q \xrightarrow{\alpha_{s_i},\phi,c} q'$.

Third step: removal of clock constraints. Let $t_1, \ldots, t_k$ be the ordered sequence of constants (smaller than T) appearing in the clock constraints of the edges of $\mathbb{D}$. We extend this sequence by letting $t_0 = 0$ and $t_{k+1} = T$. Let $I_j = [t_{j-1}, t_j]$, $j = 1, \ldots, k+1$, be the j-th sub-interval of $[0, T]$ identified by the given sequence. For each I_j, we define a DFA, $\mathbb{D}j = (\mathscr{L}, Q, q_0, F, \rightarrow_j)$, whose edge relation $\rightarrow_j$ is obtained from that of $\mathbb{D}$ by

selecting only the edges for which the clock constraints are satisfied in I_j, and dropping the clock constraint. Hence, from $q \xrightarrow{\alpha_{s_i}, c} q'$ such that $\eta(x) \models_{\mathcal{CC}} c$ whenever $\eta(x) \in (t_{j-1}, t_j)$, we obtain the DFA edge $(q, \alpha_{s_i}, q') \in \rightarrow_j$, denoted also by $q \xrightarrow{\alpha_{s_i}}_j q'$.

Fourth step: synchronization. To keep track of the behaviour of the agents with respect to the property specified by $\mathbb{D}$, we synchronize the Agent Class $\bar{\mathbb{A}} = (S, \bar{E})$ with each DFA $\mathbb{D}j$ through the standard product of automata. The sequence of deterministic automata obtained in this procedure is called the *Agent Class associated with the local property* $\mathbb{D}$.

Definition 3.4 (Agent Class associated with the local property $\mathbb{D}$) *The Agent Class $\mathscr{P}$ associated with the local property $\mathbb{D}$ is the sequence*

$$\mathscr{P} = (\mathscr{P}_{I_1}, \ldots, \mathscr{P}_{I_{k+1}})$$

of deterministic automata

$$\mathscr{P}_{I_j} = (\hat{S}, \hat{E}_j), \qquad j = 1, \ldots, k+1,$$

where $\hat{S} = S \times Q$ is the state space and $\hat{E}_j$ is the set of local transitions $\epsilon_i^j = (s, q) \xrightarrow{\alpha_s} (s', q')$, such that $s \xrightarrow{\alpha_s} s'$ is a local transition in $\bar{\mathbb{A}}$ and $q \xrightarrow{\alpha_s} q'$ is an edge in $\mathbb{D}j$.

Synchronisation of global properties

The Markov Population Model $\mathcal{X}^{(N)} = (\mathbb{A}, \mathcal{T}^{(N)}, \mathbf{x}_0^{(N)})$ has to be updated to follow the new specifications at the local level. We do this by defining the Markov Population Model associated with the local property $\mathbb{D}$ as a sequence $\boldsymbol{\mathcal{X}}^{(N)} = (\mathcal{X}_{I_1}^{(N)}, \ldots, \mathcal{X}_{I_k}^{(N)})$ of Markov Population Models. Since the agent states are synchronized with the property automaton, each transition in the population model needs to be replicated many times to account for all possible combinations of the extended local state space. Furthermore, we also need to take care of rate functions in order not to change the global rate. Fix the j-th element $\mathscr{P}_{I_j}$ in the Agent Class $\mathscr{P}$ associated with the property $\mathbb{D}$. The state space of $\mathscr{P}_{I_j}$ is $S \times Q$, hence to construct the global model we need nm counting variables ($n = |S|$, $m = |Q|$), where $X_{s,q}$ counts how many agents are in the

local state (s, q). Let $\tau = (\mathbb{S}\tau, f^{(N)}) \in \mathcal{T}^{(N)}$ be a global transition, apply the relabeling of action labels, according to step 1 above, and focus on the synchronisation set

$$\mathbb{S}\tau = \{s_1 \xrightarrow{\alpha_{s_1}} s'_1, \ldots, s_k \xrightarrow{\alpha_{s_k}} s'_k\}.$$

We need to consider all possible ways of associating states of Q with the different states $s_1, \ldots, s_k$ in $\mathbb{S}\tau$. Indeed, each choice $(q_1, \ldots, q_k) \in Q^k$ generates a different transition in $\mathcal{X}_{I_j}^{(N)}$, with synchronization set

$$\mathbb{S}\tau, r = \{(s_1, q_1) \xrightarrow{\alpha_{s_1}} (s'_1, q'_1), \ldots, (s_k, q_k) \xrightarrow{\alpha_{s_k}} (s'_k, q'_k)\},$$

where q'_i is the unique state of Q such that

$$q_i \xrightarrow{\alpha_{s_i}} q'_i.$$

The rate function $f_r^{(N)}$ associated with this instance of τ is a fraction of the total rate function $f^{(N)}$ of τ. Moreover, for all $s_i \xrightarrow{\alpha_{s_i}} s'_i \in \mathbb{S}\tau$, $f_r^{(N)}$ is proportional to the fraction of agents that before the synchronisation were in s_i and are now in state (s_i, q_i), i.e. X_{s_i,q_i} divided by $X_{s_i} = \sum_{q \in Q} X_{s_i,q}$. Formally,

$$f_r^{(N)}(\mathbf{X}) = \prod_{s_i \xrightarrow{\alpha_{s_i}} s'_i \in \mathbb{S}\tau} \left(\frac{X_{s_i,q_i}}{\sum_{q \in Q} X_{s_i,q}} \right) f^{(N)}(\widetilde{\mathbf{X}}), \tag{3.1}$$

where $\widetilde{\mathbf{X}} = (X_1, \ldots, X_n)$ with $X_s = \sum_{r=1}^{m} X_{s,r}$. Due to the restrictions enforced in Definition 2.2, summing up the rates $f_r^{(N)}(\mathbf{X})$ for all possible choices of $(q_1, \ldots, q_k) \in Q^k$, we obtain $f^{(N)}(\widetilde{\mathbf{X}})$.

Definition 3.5 (Markov Population Model associated with a local property) *The Markov Population Model associated with the local property $\mathbb{D}$ is the sequence*

$$\boldsymbol{\mathcal{X}}^{(N)} = (\mathcal{X}_{I_1}^{(N)}, \ldots, \mathcal{X}_{I_k}^{(N)}).$$

The elements $\mathcal{X}_{I_j}^{(N)} = (\mathscr{P}_{I_j}, \mathcal{T}_j^{(N)})$ are such that $\mathscr{P}_{I_j}$ is the j-th element of the Agent Class associated with $\mathbb{D}$ and $\mathcal{T}_j^{(N)}$ is the set of global transitions of the form $\tau_i^j = (\mathbb{S}i^j, f_{j,i}^{(N)})$, as defined above.[1]

[1]Initial conditions of the population models in $\boldsymbol{\mathcal{X}}^{(N)}$ are dropped, as they are not re-

3.4 Theoretical Results

In this section, we introduce the final step of our model checking procedure in which we estimate the probability of global properties by exploiting Central Limit Approximation. We then discuss how to possibly improve this estimate by using the System Size Expansion and the Moment Closure techniques combined with the distribution reconstruction based on the Maximum Entropy Principle.

3.4.1 Model Checking by Central Limit Approximation

Consider a Markov Population Model $\mathcal{X}^{(N)}$, for a fixed population size N, and a global property $\mathbb{P}_{\bowtie p}(\mathbb{D}(T) \in [a,b])$. In order to verify the latter, we need to compute the probability $\mathbb{P}(\mathbb{D}(T) \in [a,b])$ that, at time T, the fraction of agents satisfying the local specification $\mathbb{D}$ is contained in $[a,b]$. This probability can be computed exploiting the construction of Section 3.3, according to which we obtain a sequence of population models $\boldsymbol{\mathcal{X}}^{(N)} = (\mathcal{X}_{I_1}^{(N)}, \ldots, \mathcal{X}_{I_k}^{(N)})$, synchronising local agents with the sequence of deterministic automata associated with $\mathbb{D}$. In this construction we identified a sequence of times $0 = t_0, t_1, \ldots, t_k = T$ and in each interval $I_j = [t_{j-1}, t_j]$ the satisfaction of clock constraints does not change.

Therefore, in order to compute $\mathbb{P}(\mathbb{D}(T) \in [a,b])$, we can rely on *transient analysis* algorithms for CTMCs (BBHK00): first we compute the probability distribution at time t_1 for the first population model $\mathcal{X}_{I_1}^{(N)}$; then we use this result as the initial distribution for the CTMC associated with the population model $\mathcal{X}_{I_2}^{(N)}$ and we compute its probability distribution at time t_2; and so on, until we obtain the probability distribution for $\mathcal{X}_{I_k}^{(N)}$ at time $t_k = T$. Once we have this result, we can find the desired probability by summing the probability of all those states $\mathbf{X} \in \mathcal{S}^{(N)}$ such that $\sum_{s \in S, q \in F} \hat{X}_{s,q} \in [a,b]$.

Unfortunately, this approach suffers from state space explosion, which is severe even for a population size of a few hundreds of individuals.

quired in the following. The initial condition at time zero is obtained from that of $\mathcal{X}^{(N)}$ by letting $(x_0)_{s,q_0} = (x_0)_s$, where q_0 the initial state of $\mathbb{D}$ and $s \in S$.

Furthermore, for these population levels we cannot rely on the Fluid Approximation, as it would only give us an estimate of the average of the counting variables, while we need information about their distribution. It is here that the Central Limit Approximation enters the picture.

The idea is simply to compute the average and covariance matrix of the approximating Gaussian Process by solving the ODEs shown at the end of the previous section. In doing this, we have to take proper care of the different population models associated with the time intervals I_j. Then, we integrate the Gaussian density of the approximating distribution at time T to estimate of the probability $\mathbb{P}(\mathbb{D}(T) \in [a, b])$. The justification of this approach is in Theorem 2.4, which guarantees that the estimated probability is asymptotically correct, but in practice, we can obtain good approximations also for relatively small populations, in the order of hundreds of individuals.

Verification algorithm

The *inputs* of the verification algorithm are:

- an Agent Class $\mathbb{A} = (S, E)$ and a Markov Population Model $\mathcal{X}^{(N)} = (\mathbb{A}, \mathcal{T}^{(N)}, \mathbf{x}_0^{(N)})$;
- a local property specified by a 1gDTA $\mathbb{D} = (\mathscr{L}, \Gamma_{\mathcal{S}}, Q, q_0, F, \rightarrow)$;
- a global property $\mathbb{P}_{\bowtie p}(\mathbb{D}(T) \in [a, b])$ with time horizon $T > 0$.

The *steps* of the algorithm are:

1. **Construction of the Markov Population Model associated with $\mathbb{D}$.** Construct the *normalised* Markov Population Model

$$\hat{\mathcal{X}}^{(N)} = (\hat{\mathcal{X}}_{I_1}^{(N)}, \ldots, \hat{\mathcal{X}}_{I_k}^{(N)})$$

   associated with $\mathbb{D}$ according to the procedure of Section 3.3. Then modify it by adding to its vector of counting variables $\hat{\mathbf{X}}^{(N)}$ a new variable $\hat{X}_{Final}$ that keeps track of the fraction of agents entering one of the final states (s, q), $q \in F$.

2. **Integration of the Central Limit equations**. For each $j = 1, \ldots, k$, generate and solve numerically the systems of ODEs of the form similar to (2.7) and (2.8), that describe the fluid limit $\mathbf{\Phi}_j(t)$ and the Gaussian covariance $\mathbf{C}_j[\mathbf{Z}(t)]$ for the population model $\mathcal{X}_{I_j}^{(N)}$ in the interval $I_j = [t_{j-1}, t_j]$, with initial conditions

$$\forall j > 1, \quad \begin{cases} \mathbf{\Phi}_j(t_{j-1}) = \mathbf{\Phi}_{j-1}(t_{j-1}), \\ \mathbf{C}_j[\mathbf{Z}(t_{j-1})] = \mathbf{C}_{j-1}[\mathbf{Z}(t_{j-1})], \end{cases}$$

and

$$\begin{cases} \mathbf{\Phi}_1(0) = \mathbf{x}_0, \\ \mathbf{C}_1[\mathbf{Z}(0)] = Id. \end{cases}$$

For $t \in I_j$, define the population mean as

$$\mathbf{E}^{(N)}[\mathbf{X}(t)] = N\mathbf{\Phi}_j(t)$$

and the population covariance as

$$\mathbf{C}^{(N)}[\mathbf{X}(t)] = N\mathbf{C}_j[\mathbf{Z}(t)].$$

Finally, identify the component $E_{Final}^{(N)}[\mathbf{X}(t)]$ and the diagonal entry $C_{Final}^{(N)}[\mathbf{X}(t)]$ corresponding to X_{Final}.

3. **Computation of the probability**. Let $g(x \mid \mu, \sigma^2)$ be the probability density of a Gaussian distribution with mean μ and variance σ^2. Then, approximate $\mathbb{P}(\mathbb{D}(T) \in [a, b])$ by

$$\tilde{P}_{\mathbb{D}}^{(N)}(T) = \int_{Na}^{Nb} g(x \mid E_{Final}^{(N)}[\mathbf{X}(t)], C_{Final}^{(N)}[\mathbf{X}(t)])\mathrm{d}x,$$

and compare the result with the probability bound $\bowtie p$.

The asymptotic correctness of this procedure is captured in the next theorem, whose proof is a consequence of Theorem 2.4. We denote by $P_{\mathbb{D}}^{(N)}(T)$ the exact value of $\mathbb{P}(\mathbb{D}(T) \in [a, b])$ and by $\tilde{P}_{\mathbb{D}}^{(N)}(T)$ the approximate value computed by the Central Limit Approximation.

Theorem 3.1 *Under the hypothesis of Theorem 2.4, it holds that*

$$\lim_{N\to\infty} \|P_{\mathbb{D}}^{(N)}(T) - \tilde{P}_{\mathbb{D}}^{(N)}(T)\| = 0. \qquad \blacksquare$$

Proof 3.1 *Let us start by assuming that the sequence $\boldsymbol{\mathcal{X}}^{(N)}$ of the Markov Population Model associated with the property $\mathbb{D}$ (Definition 3.5) is composed of a single model. It is easy to verify that all rate functions of the modified population model of Step 1 in the verification algorithm are Lipschitz continuous, and that the conditions of Theorem 2.4 are satisfied. In particular, the initial conditions for $\boldsymbol{Z}^{(N)}(t)$ and $\boldsymbol{Z}(t)$ converge by definition. Moreover, as we are interested in the value of $\boldsymbol{Z}^{(N)}(t)$ and $\boldsymbol{Z}(t)$ at a fixed time $T > 0$, let $\boldsymbol{Z}^{(N)} = \boldsymbol{Z}^{(N)}(T)$ and $\boldsymbol{Z} = \boldsymbol{Z}(T)$. Theorem 2.4 implies that*

$$\boldsymbol{Z}^{(N)} \Rightarrow \boldsymbol{Z} \qquad \textit{(weak convergence)}.$$

To prove the convergence of $P_{\mathbb{D}}^{(N)}(T)$ to $\tilde{P}_{\mathbb{D}}^{(N)}(T)$, let us consider the N-dependent interval $[a^{(N)}, b^{(N)}]$ where

$$a^{(N)} = N^{\frac{1}{2}} \left(a - \boldsymbol{\Phi}_{Final}(T)\right)$$

and

$$b^{(N)} = N^{\frac{1}{2}} \left(b - \boldsymbol{\Phi}_{Final}(T)\right).$$

In terms of $[a^{(N)}, b^{(N)}]$, we can write

$$P_{\mathbb{D}}^{(N)}(T) = \mathbb{P}\{Z_{Final}^{(N)} \in [a^{(N)}, b^{(N)}]\}$$

and

$$\tilde{P}_{\mathbb{D}}^{(N)}(T) = \mathbb{P}\{Z_{Final} \in [a^{(N)}, b^{(N)}]\},$$

where $Z_{Final}^{(N)}$ and Z_{Final} are the marginal distributions of $\boldsymbol{Z}^{(N)}$ and $\boldsymbol{Z}$ on the coordinate corresponding to X_{Final}.

By the triangular inequality, we have

$$\begin{aligned}\|\mathbb{P}\{Z_{Final}^{(N)} \in [a^{(N)}, b^{(N)}]\} - \mathbb{P}\{Z_{Final} \in [a^{(N)}, b^{(N)}]\}\| \leq \\ \underbrace{\|\mathbb{P}\{Z_{Final}^{(N)} \in [a^{(N)}, b^{(N)}]\} - \mathbb{P}\{Z_{Final}^{(N)} \in [a^{\infty}, b^{\infty}]\}\|}_{(a)} + \\ \underbrace{\|\mathbb{P}\{Z_{Final} \in [a^{\infty}, b^{\infty}]\} - \mathbb{P}\{Z_{Final} \in [a^{(N)}, b^{(N)}]\}\|}_{(b)}\end{aligned}$$

where $[a^\infty, b^\infty]$ is the the limit set to which $[a^{(N)}, b^{(N)}]$ converges as N goes to infinity. Hence, by definition, term (b) in the inequality above goes to zero for $N \to \infty$.

To deal with term (a), we consider that $[a^\infty, b^\infty]$ can assume only one of the following four forms, depending on the relative value of a and b with respect to $\mathbf{\Phi}_{Final}(T)$:

1. *if $a, b > \mathbf{\Phi}_{Final}(T)$ or $a, b < \mathbf{\Phi}_{Final}(T)$, then $[a^\infty, b^\infty] = \emptyset$;*
2. *if $a < \mathbf{\Phi}_{Final}(T)$ and $b > \mathbf{\Phi}_{Final}(T)$, then $[a^\infty, b^\infty] = [-\infty, +\infty] = \mathbb{R}$;*
3. *if $a = \mathbf{\Phi}_{Final}(T)$ and $b > \mathbf{\Phi}_{Final}(T)$, then $[a^\infty, b^\infty] = [0, +\infty]$;*
4. *if $a < \mathbf{\Phi}_{Final}(T)$ and $b = \mathbf{\Phi}_{Final}(T)$, then $[a^\infty, b^\infty] = [-\infty, 0]$;*

Moreover, we can exploit the fact that, as $Z^{(N)}_{Final} \Rightarrow Z_{Final}$ and $\mathbb{R}$ is a Polish space, by the Prohorov theorem, $Z^{(N)}_{Final}$ is uniformly tight. This means that, for each $\epsilon > 0$, there exists $k_\epsilon > 0$ such that, for all N,

$$\mathbb{P}\{Z^{(N)}_{Final} \in [-k_\epsilon, k_\epsilon]\} > 1 - \epsilon.$$

Looking at each of the four forms of $[a^\infty, b^\infty]$ separately, we have the following results.

1. *Fix $\epsilon > 0$ and let N_0 be such that, for $N \geq N_0$,*

   $$[a^{(N)}, b^{(N)}] \cap [-k_\epsilon, k_\epsilon] = \emptyset.$$

   It follows that $\mathbb{P}\{Z_{Final} \in [a^{(N)}, b^{(N)}]\} < \epsilon$. Moreover, due to the definition of $[a^{(N)}, b^{(N)}]$, we have also that $\mathbb{P}\{Z^{(N)}_{Final} \in [a^\infty, b^\infty]\} = 0$, and thus the value of the whole term (a) is less than ϵ, which implies that (a) goes to zero for N going to infinity.

2. *Fix $\epsilon > 0$ and let N_0 be such that, for $N \geq N_0$,*

   $$[a^{(N)}, b^{(N)}] \cap [-k_\epsilon, k_\epsilon] = [-k_\epsilon, k_\epsilon].$$

   As $\mathbb{P}\{Z^{(N)}_{Final} \in [a^\infty, b^\infty]\} = 1$, it follows that (a) is smaller than ϵ, hence it has limit 0.

3. *Fix $\epsilon > 0$ and let N_0 be such that, for $N \geq N_0$,*

$$[a^{(N)}, b^{(N)}] \cap [-k_\epsilon, k_\epsilon] = [0, k_\epsilon].$$

By the monotonicity of the probability distributions, term (a) is smaller than $\mathbb{P}\{Z_{Final} > k_\epsilon\}$, which is itself smaller than ϵ. Also in this case, it follows that term (a) has limit 0.

4. *This case is similar to case 3.*

Hence, as desired, we have proven that

$$\lim_{N\to\infty} \|\mathbb{P}\{Z^{(N)}_{Final} \in [a^{(N)}, b^{(N)}]\} - \mathbb{P}\{Z_{Final} \in [a^{(N)}, b^{(N)}]\}\| = 0$$

To deal with the case in which the Markov Population Model associated with the property $\mathbb{D}$ is a sequence of $k > 1$ models, we can rely on the fact that the time constants defining intervals I_j are fixed, hence Theorem 2.4 holds inductively for each model of the sequence. Indeed, the initial conditions of model $\hat{\mathcal{X}}_{I_j}$ are given by the final state of model $\hat{\mathcal{X}}_{I_{j-1}}$, which converges by inductive hypothesis. Therefore, to prove the convergence of the probability $P^{(N)}_{\mathbb{D}}(T)$ to $\tilde{P}^{(N)}_{\mathbb{D}}(T)$, we just need to apply the argument discussed above to the final model of the sequence. ∎

Remark 3.4 *The introduction of the counting variable X_{Final} is needed to correctly capture the variance in entering one of the final states of the property. Indeed, it holds that $X_{Final} = \sum_{s\in S, q\in F} X_{s,q}$, and in principle we could have applied the Central Limit Approximation to the model without X_{Final}, by exploiting the fact that the sum of Gaussian variables is Gaussian (with mean and variance given by the sum of means and variances of the addends). In doing this, though, we overestimate the variance of X_{Final}, because we implicitly take into account the dynamics within the final components. The introduction of X_{Final}, instead, avoids this problem, as its variance depends only on the events that allow the agents to enter one of the final states.*

3.4.2 Model Checking by System Size Expansion and Moment Closure

The method just presented relies on the Central Limit Approximation (CLA), hence its accuracy depends on the quality of the estimation of the probability distribution given by the CLA at time T. In particular, in

some specific cases, we found the accuracy of the approach to be hampered: for example when if the probability $X^{(N)}_{final}(T)$ is not symmetric, or when its distribution deviates from a Gaussian, or when the Fluid Approximation fails in giving a good estimation of the mean of $X^{(N)}_{final}(T)$. One way of tackling these and other problems is to exploit different types of Stochastic Approximations, like the System Size Expansion or the Moment Closure combined with the Moment Reconstruction Principle.

As it was discussed in Section 2.4.3, the System Size Expansion (SSE), or Inverse Omega Square (IOS), is an approximation which estimates the dynamics of a Markov Population Model by means of a Gaussian probability distribution with mean $\mathbf{E}^*[\mathbf{X}(t)]$ and covariance $\mathbf{C}^*[\mathbf{X}(t)]$ given by (2.10) and (2.11), respectively. Moreover, when we drop the terms of order $O(N^{1/2})$ in the definitions of $\mathbf{E}^*[\mathbf{X}(t)]$ and $\mathbf{C}^*[\mathbf{X}(t)]$, we obtain equations (2.7) and (2.8) for the mean and covariance of the Central Limit Approximation. Hence, the System Size Expansion is indeed a higher order correction of the CLA, that can be efficiently exploited when the we want to improve the accuracy of the mean (i.e. the Fluid Approximation) and of the covariance of the CLA.

Due to the fact that the System Size Expansion (IOS) is indeed a higher order correction of the Central Limit Approximation, and hence their formal definitions are very similar, we can easily exploit the IOS in our model checking procedure starting from the the verification algorithm of Section 3.4.1. Indeed, to adjust the procedure to the IOS, we just need to substitute in Step 2, the integration of the equations for $\mathbf{E}[\mathbf{X}(t)]$ and $\mathbf{C}[\mathbf{X}(t)]$ with those of $\mathbf{E}^*[\mathbf{X}(t)]$ and $\mathbf{C}^*[\mathbf{X}(t)]$. Moreover, as for the CLA, the quality of the model checking procedure that is based on the System Size Expansion is guaranteed by a limit result like Theorem 3.1, whose proof is a straightforward adjustment of Proof 3.1 to exploit Theorem 2.5 (instead of Theorem 2.4).

When the knowledge of the typology of the probability distribution $X^{(N)}_{final}(T)$ is limited, or more in general, when a Gaussian estimation fails to accurately describe $X^{(N)}_{final}(T)$, then the IOS faces the same drop in the accuracy as the CLA, but we can still rely on the Moment Closure and the Maximum Entropy principle. Indeed, in the approximation

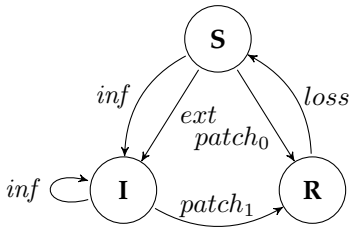

Figure 5: The automaton representation of the Agent Class of the running example.

technique introduced in Section 2.4.4 does not rely on any assumption about the typology of the distribution $X^{(N)}_{final}(T)$ (such as Gaussian for the CLA and the IOS), hence we can exploit it on less regular cases. To apply this Stochastic Approximation to our model checking procedure, in the algorithm of Section 3.4.1, we need to substitute: Step 2 with the integration of the K moments obtained by the Moment Reconstruction; and Step 3 with the computation of the estimate of the transient probability $\mathbb{P}(\mathbb{D}(T) \in [a, b]) = P^*(T)$ following the distribution reconstruction based on the Maximum Entropy Principle described in Section 2.4.4. In this case, we cannot provide a limit theorem to ensure the quality of the approximation, but, as we shall see in the following section, the experimental results are promising.

3.5 Experimental Analysis

3.5.1 Results of Central Limit Approximation

We discuss now the quality of the Central Limit Approximation for mesoscopic populations from an experimental perspective. We present a detailed investigation of the behaviour of the example describing a network epidemics introduced in Section 2.3 (whose automaton representation is depicted again in Figure 5 to ease the readability of this Section).

We consider the two local properties expressed as 1gDTAs shown in Figure 6. The first property $\mathbb{D}_1$ has no clock constraints on the edges of

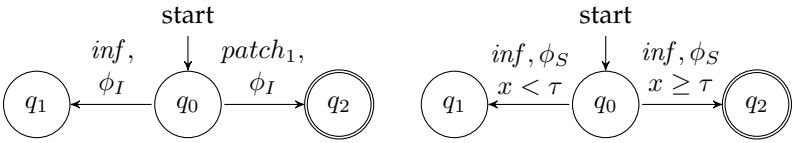

Figure 6: The 1gDTA specifications experimentally analysed in Section 3.5.

the automaton, therefore the 1gDTA reduces to a DFA. The property is satisfied if an infected node is patched before being able to infect other nodes in the network, thus checking the effectiveness of the antivirus deployment strategy. The second property $\mathbb{D}_2$, instead, is properly timed. It is satisfied when a susceptible node is infected by an internal infection after the first τ units of time. The corresponding global properties that we consider are $\mathbb{P}(\mathbb{D}_1(T) \geq \alpha_1)$ and $\mathbb{P}(\mathbb{D}_2(T) \geq \alpha_2)$.

In Figure 7, we show the probability of the two global properties as a function of the time horizon T, for different values of N and a specific configuration of the rates of the transitions and of the threshold of the properties ($\kappa_{inf} = 0.05$, $\kappa_{patch_1} = 0.02$, $\kappa_{loss} = 0.01$, $\kappa_{ext} = 0.05$, $\kappa_{patch_0} = 0.001$, $\alpha_1 = 0.5$, $\alpha_2 = 0.2$). The CLA is compared with a statistical estimate, obtained from 10000 simulation runs. As we can see, the accuracy in the transient phase increases rapidly with N, and the estimate is very good for both properties already for $N = 100$. The same parameter configuration was used to compute the computational costs (in Seconds), showed in Table 1. As we have seen, by definition the Central Limit Approximation is independent of the population size N and its computational costs is hundreds of times less than that of the statistical estimate (the Gillespie Algorithm) for both the first and the second properties.

Furthermore, in order to check more extensively the quality of the approximation also as a function of the system parameters, we ran the following experiment. We considered five different values of N ($N = 20, 50, 100, 200, 500$). For each of these values, we randomly chose 20 different combinations of parameter values, sampling uniformly from:

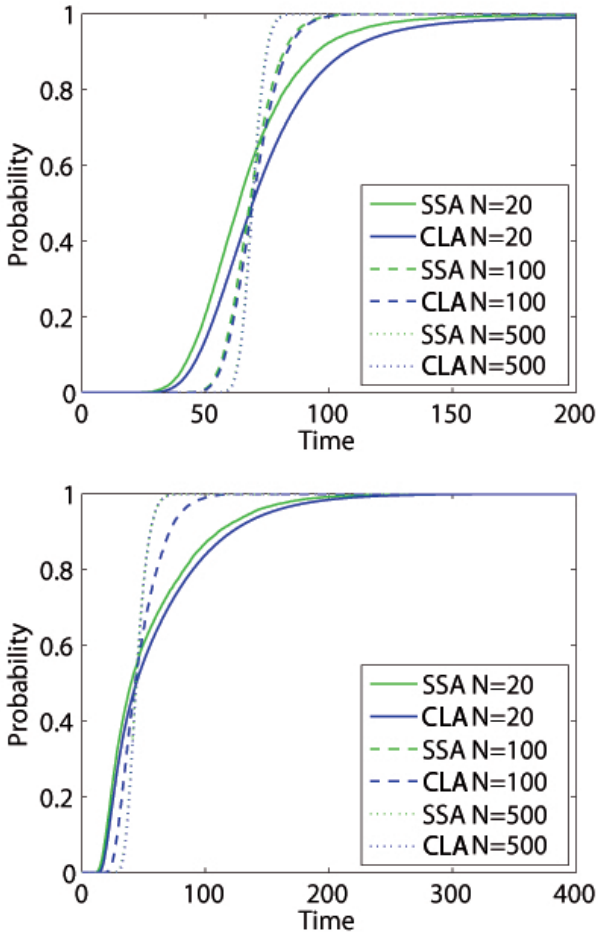

Figure 7: Comparison of Central Limit Approximation (CLA) and a statistical estimate (using the Gillespie algorithm, SSA) of the path probabilities of the 1gDTA properties of Figure 6 computed on the network epidemic model for different values of the population size N.

$\kappa_{inf} \in [0.05, 5]$, $\kappa_{patch_1} \in [0.02, 2]$, $\kappa_{loss} \in [0.01, 1]$, $\kappa_{ext} \in [0.05, 5]$, $\kappa_{patch_0} \in [0.001, 0.1]$, $\alpha_1 \in [0.1, 0.95]$, $\alpha_2 \in [0.1, 0.3]$. For each parameter set, we

First Property

N	SSAcost	CLAcost	Speedup
20	22.4114	0.0618	362.6440
50	23.3467	0.0618	377.7783
100	24.2689	0.0618	392.7006
200	26.1074	0.0618	442.4498
500	28.8754	0.0618	467.2395

Second Property

N	SSAcost	CLAcost	Speedup
20	32.0598	0.3035	105.6336
50	29.0915	0.3035	95.8534
100	28.8651	0.3035	95.1074
200	33.9825	0.3035	111.9687
500	43.4737	0.3035	143.2412

Table 1: Average computational costs (in Seconds) of the Gillespie Algorithm (SSAcost) and the Central Limit Approximation (CLAcost), and the relative SpeedUp (CLAcost/SSAcost). The data are shown as a function of the population size N (by definition the CLA is independent of N).

compared the CLA of the probability of each global property with a statistical estimate (from 5000 runs), measuring the error in a grid of 1000 equi-spaced time points. We then computed the maximum error and the average error. In Table 2, we report the mean and maximum values of these quantities over the 20 runs, for each considered value of N. We also report the error at the final time of the simulation, when the probability has stabilised to its limit value.[2] It can be seen that both the average and the maximum errors decrease with N, as expected, and are already quite small for $N = 100$ (for the first property, the maximum difference in the path probability for all runs is of the order of 0.06, while the average error is 0.003). For $N = 500$, the CLA is practically indistinguishable from the (estimated) true probability. For the second property, the errors are slightly worse, but still reasonably small.

Finally, we considered the problem of understanding what are the

[2]For this model, we can extend the analysis to steady state, as the fluid limit has a unique, globally attracting steady state. This is not possible in general, cf. (BHLM13).

First Property

N	MaxEr	$\mathbb{E}$[MaxEr]	Max$\mathbb{E}$[Er]	$\mathbb{E}[\mathbb{E}[\text{Er}]]$	MaxEr(T)	$\mathbb{E}$[Er(T)]
20	0.1336	0.0420	0.0491	0.0094	0.0442	0.0037
50	0.0866	0.0366	0.0631	0.0067	0.0128	0.0018
100	0.0611	0.0266	0.0249	0.0030	0.0307	0.0017
200	0.0504	0.0191	0.0055	0.0003	0.0033	0.0002
500	0.0336	0.0120	0.0024	0.0003	0.0002	9.5e-6

Second Property

N	MaxEr	$\mathbb{E}$[MaxEr]	Max$\mathbb{E}$[Er]	$\mathbb{E}[\mathbb{E}[\text{Er}]]$	MaxEr(T)	$\mathbb{E}$[Er(T)]
20	0.2478	0.1173	0.1552	0.0450	0.1662	0.0448
50	0.2216	0.0767	0.1233	0.0340	0.1337	0.0361
100	0.1380	0.0620	0.0887	0.0216	0.0979	0.0208
200	0.1365	0.0538	0.0716	0.0053	0.0779	0.0162
500	0.1187	0.0398	0.0585	0.0100	0.0725	0.0108

Table 2: Errors obtained by the Central Limit Approximation in the validation of Local-to-Global Properties. Maximum and mean of the maximum error (MaxEr, $\mathbb{E}$[MaxEr]) for each parameter configuration; maximum and mean of the average error with respect to time (Max$\mathbb{E}$[Er]), $\mathbb{E}[\mathbb{E}[\text{Er}]]$) for each parameter configuration; maximum and average error at the final time horizon T (MaxEr(T), $\mathbb{E}$[Er(T)]) for each parameter configuration. Data is shown as a function of the network size N.

most important aspects that determine the error. To this end, we regressed the observed error against the following features: estimated probability value by CLA, error in the predicted average and variance of X_{Final} (between the CLA and the statistical estimates), and statistical estimates of the mean, variance, skewness and kurtosis of X_{Final}. We used Gaussian Process regression with Adaptive Relevance Detection (GP-ADR, (RW06)), which performs a regularised regression searching the best fit on an infinite dimensional subspace of continuous functions, and permitted us to identify the most relevant features by learning the hyperparameters of the kernel function. We used both a squared exponential kernel, a quadratic kernel, and a combination of the two, with a training set of 500 points, selected randomly from the experiments performed. The mean prediction error on a test set of other 500 points (independently of N) is around 0.015 for all the considered kernels. Fur-

thermore, GP-ADR selected as most relevant the quadratic kernel, and in particular the following two features: the estimated probability and the error in the mean of X_{Final}. This suggests that moment closure techniques improving the prediction of the average can possibly reduce the error of the method.

Finite-Size Threshold Correction

Results obtained by CLA can be further improved for small values of N by introducing a correction on the thresholds a and b of a property $\mathbb{P}_{\bowtie p}(\mathbb{D}(T) \in [a, b])$, taking into account the discrepancy between the discrete nature of population counts and its continuous approximation. To better understand the correction, let us consider a property of the form $\mathbb{P}_{\bowtie p}(\mathbb{D}(T) \geq \alpha)$. In the algorithm presented in Section 3.4.1, the CLA works by integrating the Gaussian approximation of the variable X_{final} from αN to infinity. However, in this way, for small N, we neglect the discrete nature of the state space. Suppose we would like to compute the probability of $X_{final} = i$. Using the Gaussian approximation, we would always obtain zero, unless we integrate in a region around i. The obvious candidate is $[i - \frac{1}{2}, i + \frac{1}{2}]$, which correspond to a partition of the interval $[0, N]$ into subintervals of the form $[i - \frac{1}{2}, i + \frac{1}{2}]$[3]. Following this line of reasoning, instead of integrating the Gaussian approximation for X_{final} from αN, we should start from $j - \frac{1}{2}$, where j is the smallest integer greater than or equal to αN, i.e. $j = \lceil \alpha N \rceil$. Note that j is the smallest value that X_{final} can take to satisfy the property, when verifying it in the discrete stochastic model. Similarly, when dealing with properties of the form $\mathbb{P}_{\bowtie p}(\mathbb{D}(T) \leq \alpha)$, we would need to integrate up to $\lfloor \alpha N \rfloor + \frac{1}{2}$, combining the two corrections with dealing with threshold intervals $[a, b]$. In several experimental tests, we observed that this simple correction improves considerably the approximation, becoming less significant for large N.

Example 3.2 *In Figure 8 we see the correction at work for $N = 20$ and the first property of Figure 6, in which $\alpha = 0.5$, hence $\lceil \alpha N \rceil = 10$. We can see*

[3]The extremes 0 and N has to be treated in a special way: $(-\infty, \frac{1}{2}]$ for 0 and $[N - \frac{1}{2}, \infty)$ for N

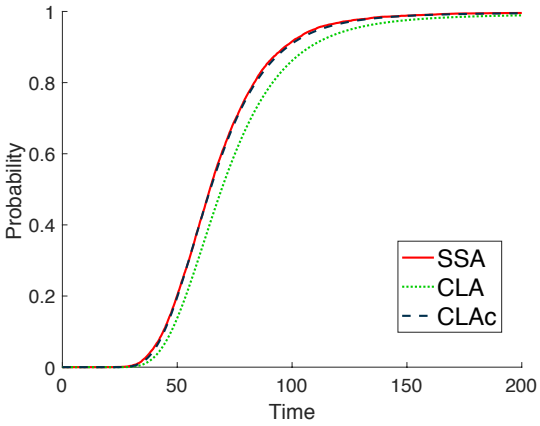

Figure 8: Comparison of a statistical estimate (using the Gillespie algorithm, SSA), the Central Limit Approximation (CLA), and the CLA with the finite-size threshold correction (CLAc) for the first property of Figure 6, with $N = 20$ and $\alpha = 0.5$.

what happens if we integrate from 9.5 instead of 10. Integrating from 10, some probability mass is lost, and the CLA under-approximates the true solution. The correction allows us to recover some of this lost mass, improving considerably the quality of the approximation.

3.5.2 Results of System Size Expansion and Moment Closure

At the moment of the completion of this thesis, a full and extensive experimental analysis of the stochastic model checking techniques based on the System Size Expansion (IOS) and the Moment Closure (MC) with distribution reconstruction is still an on going work which will be presented in the paper (BLN17). In the following, we illustrate the first (and very promising) results for these model checking techniques, leaving the detailed error analysis to (BLN17).

To analyse the model checking procedure for the IOS and the Moment Closure we have considered the untimed property of Figure 6 left. The

set of parameters considered in the experiment are the following: $\kappa_{inf} = 0.04$, $\kappa_{patch_1} = 0.02$, $\kappa_{loss} = 0.01$, $\kappa_{ext} = 0.05$, $\kappa_{patch_0} = 0.001$, $\alpha_1 = 0.6$. For the Moment Closure, we have considered a Low Dispersion of grade 3 technique, hence we have set to zero all the moments of grade greater or equal to 4. In Figure 9, we are able to compare the results obtained for the CLA and the Gillespie's statistical estimates (with 10000 runs) (SSA), with the probabilities estimated by the IOS and the MC, for two values of the population size: $N = 20$ (left) and $N = 100$ (right). As we can immediately see, on the case of this simple property, the performance of the three types of approximation (CLA, IOS and MC) is comparable (almost the same). Moreover, we can also notice how the IOS and the MC probabilities converge to the statistical estimates quite fast, and already obtain almost exact results at $N = 100$. In our next work (BLN17), we plan to stress the three model checking procedure on more complex and timed property, to understand better the different performances and the quality of the estimations.

3.6 Discussion

In this chapter, we have considered Markov Population Models and timed properties of individual agents specified by DTAs endowed with a one global clock. We introduced stochastic model checking methods based on Central Limit Approximation, System Size Expansion, and Moment Closure. The Stochastic Approximations are exploited to accurately estimate the collective probability with which a given fraction of agents satisfies the local specification. The correctness of our method is guaranteed by a convergence result and validated experimentally on a network epidemics model.

The results of this chapter could be extended in several directions. One possibility is to consider more complex DTA properties, for instance allowing clock resets. This line of research is going to be further discussed in Chapter 5, but only for individual local properties. In that chapter, we also discuss the difficulties of passing from local to collective properties when multiple clocks are involved.

On the theoretical side, it would be interesting to compute the speed of convergence of the approximations in Theorems 3.1 and 2.5 in order to possibly compare it with the results obtained for the Fluid Approximation (see e.g. (BHLM13)). Experimentally, we are currently expanding and improving the analysis of Section 3.5.2, considering also more complex requirements, to better capture the performances of the different approximation techniques. Moreover, a possible challenging area of investigation looks into the topology of the phase space of the ODEs of the Stochastic Approximations to extensively understand the error and the performance of the estimations.

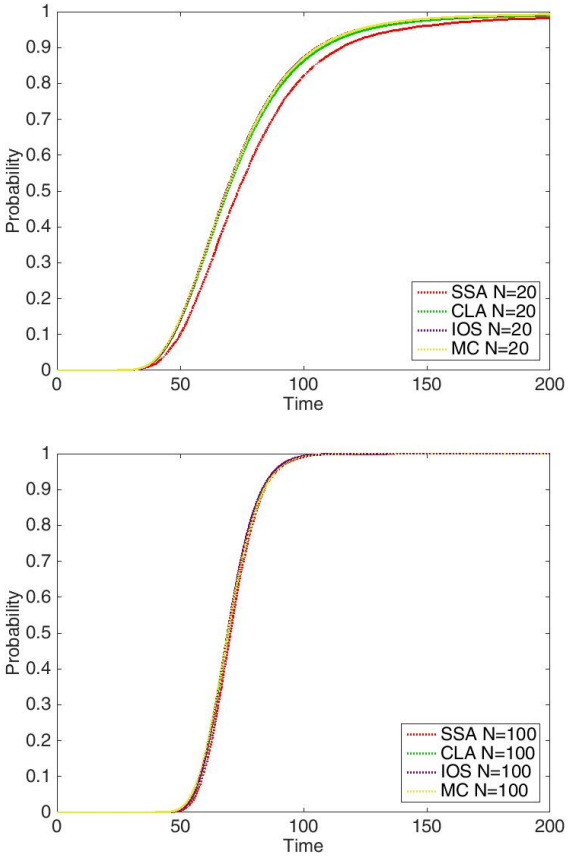

Figure 9: Comparison of the results obtained by the CLA, the statistical estimate (SSA), the System Size Expansion (IOS), and the Moment Closure (MC) for $N = 20$ and $N = 100$.

Chapter 4

Hitting Time Approximation for Global Reachability Properties

4.1 Overview

In this chapter, we apply Stochastic Approximations to validate a class of global properties of the Markov Population Models: the *global reachability properties*. In particular, we are interested in the fast and accurate estimation of the probability that a fraction of the population in the system reaches, within a given time horizon $T < \infty$, a certain region of the state space, the *target region*, defined by a non-linear inequality of the counting variables that identity the state of the population. An example of these properties is the probability that a large fraction of users in a peer-to-peer network downloads an updated piece of information, or the probability that a given fraction of computers in a LAN becomes infected by a virus. Reachability queries are important in many respects: safety properties belong to this class of properties, and moreover, these requirements constitute the core subroutine to check time-bounded CSL properties (BK08).

The main idea of our approach to validate global reachability prop-

erties is to transform the reachability problem into a *hitting time problem*, and thus approximate the reachability probabilities by computing an estimation of the *hitting time probabilities*, exploiting Central Limit Approximation and System Size Expansion (EK05).

Consider the Fluid Approximation, which as we have seen in Section 2.4 of Chapter 2 is a deterministic process $\mathbf{\Phi}(t)$ described by a set of ODEs, and assume that $\mathbf{\Phi}(t)$ enters the target region $\mathcal{R}$ at a given time $t_{\mathcal{R}}$. Then, $t_{\mathcal{R}}$ is an accurate estimate of the true time instant in which the Markov Population Model enters $\mathcal{R}$ when the population size is large. However, for populations of the order of hundreds of individuals, a typical size of *mesoscopic models*, this approximation loses its accuracy, as the stochastic noise generated at the local level of the single agents cannot be neglected. Hence, our idea is to exploit the *Central Limit Approximation* and the *System Size Expansion*, as in Chapter 3, and estimate the hitting time probability (and thus the global reachability probability) by a Gaussian process. Throughout the chapter, the feasibility and accuracy of the approach is going to be discussed on an example of software update process in a peer to peer network, inspired by (Hay12).

The chapter is organised as follows. In Section 4.2, we introduce the running example. In Section 4.3, we discuss the global reachability problem and the formalism we consider. Section 4.4 contains the main theoretical results: the Gaussian approximation of the hitting time, its use for estimating the reachabilty probability, and how to exploit System Size Expansion to improve accuracy. Section 4.5 shows the method in practice on the peer-to-peer example. A discussion can be found in Section 4.6.

The results of this paper have been published in (BL14).

4.2 Running Example

To illustrate the model checking procedure, we consider a simple variant of the peer-to-peer software update process introduced in (Hay12). In the modelled network, a node can be *old*, meaning that it has an old version of the software, or *updated*, when it has been able to receive the update.

In both cases, the node can be switched *ON* and *OFF*, and an old node can update only when it is on. The search for the update in the network lasts until a certain timeout is reached, after which the old node gives up and reaches a *oldOUT* state from which it can be eventually switched off. Finally, we mimic also the possibility that an *oldOUT* node obtains the update from an external source (ext_O) and that the license of the updated version of the software eventually expires or a new version is released (exp_U).

The Agent Class $\mathbb{A}_{node} = (S_{node}, E_{node})$ of the network nodes can be easily derived from the automaton representation depicted in Figure 10. The population model $\mathcal{X}_{network} = (\mathbb{A}_{node}, \mathcal{T}, \mathbf{x}_0)$ is described by the vector of counting variables

$$\mathbf{X} = (X_{oldOFF}, X_{oldON}, X_{oldOUT}, X_{updatedOFF}, X_{updatedON})$$

and the set of global transition is given by

$$\mathcal{T} = \{\tau_{\text{on}_O}, \tau_{\text{off}_O}, \tau_{\text{out}_O}, \tau_{\text{off}_T}, \tau_{\text{ext}_O}, \tau_{\text{off}_U}, \tau_{\text{on}_U}, \tau_{\text{update}}, \tau_{\text{exp}_U}\}.$$

For example, the switching on of an old node is described by

$$\tau_{\text{on}_O} = \{\{oldOFF \xrightarrow{on_O} oldON\}, f_{\text{on}_O}\},$$

where the synchronisation set specifies that only one (old) node is involved and changes state from *oldOFF* to *oldON* at an expected rate given by the function

$$f_{\text{on}_O}(\mathbf{X}) = \lambda_{\text{on}_O} X_{oldOFF},$$

in which λ_{on_O} is the constant indicating the rate of switching on of old nodes per single unit. The other global transitions τ_{off_O}, τ_{out_O}, τ_{off_T}, τ_{ext_O},τ_{off_U}, τ_{on_U}, τ_{exp_U} have a similar form (with λ_{ext_O} and λ_{exp_U} having low values to implement the fact that ext_O and exp_U happen on a much lower time-scale than the others). The global transition τ_{update}, instead, synchronises two local transitions. In particular, τ_{update} involves an *oldON*-node and an *updatedON*-node and we have

$$\tau_{\text{update}} = \{\{oldON \xrightarrow{update_O} updatedON, \\ updatedON \xrightarrow{update_U} updatedON\}, f_{\text{update}}\}.$$

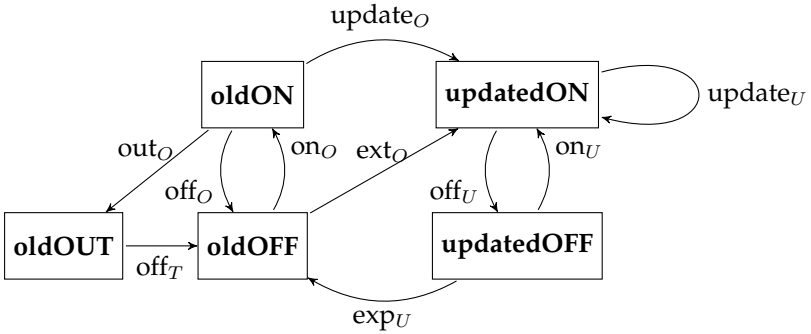

Figure 10: The automaton representation of the peer-to-peer software update process of Section 4.2.

In this case, we assume that an *updatedON*-node sends the update to an *oldON*-one at an instantaneous rate given by λ_{update} and the rate function has the classical mass action form

$$f_{\text{update}}(\mathbf{X}) = \lambda_{\text{update}} X_{oldON} X_{updatedON},$$

depending on the number of pairs of nodes that are ready to communicate (AB00).

4.3 Global Reachabilty Properties

In this chapter, we introduce a model checking procedure for the validation of *reachability properties*. In particular, we consider instances of *global* reachability properties, describing the dynamics of the system at the population level, i.e. characterising the collective behaviour of all agents. In order to verify these requirements, we compute the probability of reaching, within a given time horizon T, a specific *target region* $\mathcal{R} \subset \mathcal{S}$ of the state space $\mathcal{S}$ of the population model, starting from the initial state $\mathbf{x}_0$:

$$\mathbb{P}_{\mathcal{R}}(T) = \mathbb{P}\{\mathbf{X}(t) \in \mathcal{R} \mid t \in [0, T]\}. \tag{4.1}$$

Reachability is a fundamental notion in the analysis and verification of complex systems and it has been widely studied in many disciplines, in-

cluding physics, biology and computer science. In the latter community, the investigation of reachability has been usually motivated by the *safety verification problem*, that checks the performance of a model by computing the probability associated with its failure, i.e. with those trajectories that end up in a dead-lock or error state. This type of analysis is indeed fundamental for a sound and reliable verification of software and hardware systems, and in recent years great variety of stochastic model checking techniques have been developed in order to efficiently tackle the problem (Buj12).

The standard *stochastic model checking* procedures address the reachability problem (4.1) by making *absorbing* the states in the target region $\mathcal{R}$ and computing the transient probability of being in those states at time T. However, all these methods severely suffer by the *state space explosion* of population models, which hampers the computability of transient and steady-state probabilities. In this chapter, we introduce a model checking procedure, which tackles the problem of the state space explosion by considering scalable approximations of the population dynamics.

In the following, we reformulate the reachability problem (4.1) as a *hitting time problem*. In particular, instead of computing the probability of reaching the target region $\mathcal{R}$ before time T, we consider the *hitting time* $t_{\mathcal{R}}$, the instant in which the trajectory of the population model enters $\mathcal{R}$, and we compute the probability that $t_{\mathcal{R}} < T$:

$$\mathbb{P}^{hit}_{\mathcal{R}}(T) = \mathbb{P}\{t_{\mathcal{R}} \leq T\} \qquad \text{with} \qquad t_{\mathcal{R}} = \inf\{t > 0 \mid \mathbf{X}(t) \in \mathcal{R}\}. \quad (4.2)$$

It is straightforward to prove that $\mathbb{P}_{\mathcal{R}}(T) = \mathbb{P}^{hit}_{\mathcal{R}}(T)$.

To compute the reachability probability $\mathbb{P}_{\mathcal{R}}(T)$, we define a Gaussian estimation of the hitting time $t_{\mathcal{R}}$, starting from the Central Limit Approximation of the dynamics of the population model $\mathcal{X}$. Hence, following the procedure illustrated in Section 2.4 of Chapter 2, we define the sequence of Markov Population Models $(\mathcal{X}^{(N)})_{N\in\mathbb{N}}$ and, normalising with respect to the population size N, we consider the following reachability probability:

$$\mathbb{P}^{(N)}_{\mathcal{R}}(T) = \mathbb{P}\{t^{(N)}_{\mathcal{R}} \leq T\} \qquad \text{with} \qquad t^{(N)}_{\mathcal{R}} = \inf\{t > 0 \mid \hat{\mathbf{X}}^{(N)}(t) \in \mathcal{R}\}.$$

Moreover, we assume that the (normalised) target region $\mathcal{R}$ is defined by an inequality on population variables. Formally, we introduce a suitable *target function* $\rho : \mathcal{D} \to \mathbb{R}$, such that $\mathcal{R}$ is the subset of $\mathcal{D}$ where ρ is negative. The function ρ is defined on the compact set $\mathcal{D} \subseteq [0,1]^n$ such that $\bigcup_N \hat{\mathcal{S}}^{(N)} \subseteq \mathcal{D}$ and comes in the form of a nonlinear differentiable function of the normalised counting variables $\hat{X}_1(t), \ldots, \hat{X}_n(t)$. Hence, the final form of the *reachability problem* we want to solve is given by

$$\mathbb{P}^{(N)}_{\mathcal{R}}(T) = \mathbb{P}\{t^{(N)}_{\mathcal{R}} \leq T\} \qquad \text{with} \qquad t^{(N)}_{\mathcal{R}} = \inf\{t > 0 \mid \rho(\hat{\mathbf{X}}^{(N)}(t)) < 0\}. \tag{4.3}$$

Example 4.1 *As an example, consider the peer-to-peer software update process described in Section 4.2 in a network with 100 nodes, i.e. the population size is N=100. We can validate the performance of the model by considering the simple reachability property which controls the time in which* 95% *of the nodes have been updated. In this case, the target region* $\mathcal{R}$ *is that in which the number of agents that have received the update, i.e.* $X_{updatedOFF} + X_{updatedON}$*, is greater or equal to* 95% *of the population, i.e.* $0.95 \cdot 100$*. Hence,*

$$\mathcal{R} := \{\mathbf{X}(t) \in \mathcal{S} \mid X_{updatedOFF}(t) + X_{updatedON}(t) \geq 0.95 \cdot 100\},$$

and the target function $\rho : \mathcal{D} \to \mathbb{R}$ *is given by*

$$\rho(\hat{\mathbf{X}}^{(N)}(t)) := 0.95 - \hat{X}_{updatedOFF}(t) - \hat{X}_{updatedON}(t).$$

4.4 Theoretical Results

In this section, we present the theory behind our model checking procedure. First, we define the Central Limit Approximation of the hitting time $t_{\mathcal{R}}$, then we build an efficient algorithm to estimate the reachability probability, exploiting the duality between hitting times and reachability. The algorithm relies on an hypothesis, namely that the fluid trajectories enter the target region. We will discuss more this restriction at the end of the chapter. At the end of the section, we also discuss how to improve the approximation by exploiting the System Size Expansion and the Moment Closure.

4.4.1 Central Limit Approximation of the Hitting Time Distribution

To compute the cumulative probability distribution associated with the reachability problem (4.3), the model checking procedure that we are presenting exploits a Corollary of Theorem 2.4, which provides a Gaussian estimation of $t_{\mathcal{R}}^{(N)}$ based on the Fluid and Central Limit Approximations of $\hat{\mathbf{X}}^{(N)}(t)$. In this section, we review how to define this estimation.

The Fluid Approximation of $\hat{\mathbf{X}}^{(N)}(t)$ provides a deterministic approximation $t_{\mathcal{R}}$ (*independent of* N) of the hitting time $t_{\mathcal{R}}^{(N)}$ of the reachability problem (4.3), namely

$$t_{\mathcal{R}} = \inf\{t > 0 \mid \mathbf{\Phi}(t) \in \mathcal{R}\}. \tag{4.4}$$

As a direct consequence of Kurtz's Theorem on the fluid limit, this approximation is exact in the limit of an infinite population. Consider now the (normalised) stochastic variable $\varepsilon^{(N)} := \sqrt{N}(t_{\mathcal{R}}^{(N)} - t_{\mathcal{R}})$, which captures the noise of $t_{\mathcal{R}}^{(N)}$ around the deterministic fluid estimation $t_{\mathcal{R}}$. Let $\{\mathbf{Z}(t) \in \mathbb{R}^n \mid t \in \mathbb{R}\}$ be the Gaussian noise of the Central Limit Approximation with mean and covariance given by (2.7) and (2.8) respectively, and let ε be the random variable given by

$$\varepsilon := -\frac{\nabla\rho(\mathbf{\Phi}(t_{\mathcal{R}})) \cdot \mathbf{Z}(t_{\mathcal{R}})}{\nabla\rho(\mathbf{\Phi}(t_{\mathcal{R}})) \cdot \mathbf{F}(\mathbf{\Phi}(t_{\mathcal{R}}))},$$

where: $\rho : \mathcal{D} \to \mathbb{R}$ is the target function identifying the (normalised) target region $\hat{\mathcal{R}}$ in (4.3); ∇ is the gradient; and $\cdot$ is the Euclidean scalar product. Then, the following result holds true ((EK05), Ch 11, Theorem 4.1).

Theorem 4.1 *Assume that* $\lim_{N\to\infty} \boldsymbol{Z}^{(N)}(0) = \boldsymbol{Z}(0)$ *as in Theorem 2.4. If* $t_{\mathcal{R}} < \infty$ *and* $\nabla\rho(\mathbf{\Phi}(t_{\mathcal{R}}))\cdot\boldsymbol{F}(\mathbf{\Phi}(t_{\mathcal{R}})) < 0$, *then* $\varepsilon^{(N)}(t)$ *converges in distribution to* $\varepsilon(t)$.

In conclusion, the Gaussian approximation of the hitting time $t_{\mathcal{R}}^{(N)}$ that we consider in our model checking procedure is given by

$$t_{\mathcal{R}} - \frac{1}{\sqrt{N}} \frac{\nabla\rho(\mathbf{\Phi}(t_{\mathcal{R}})) \cdot \mathbf{Z}(t_{\mathcal{R}})}{\nabla\rho(\mathbf{\Phi}(t_{\mathcal{R}})) \cdot \mathbf{F}(\mathbf{\Phi}(t_{\mathcal{R}}))} \tag{4.5}$$

and Theorem 4.1 guarantees that the estimation is exact in the limit of an infinite population size.

4.4.2 Verification Algorithm

The algorithm of our model checking procedure for the verification of reachability properties over Markov Population Models has the following form.

Input:

- an Agent Class $\mathbb{A} = (S, E)$ as described in Definition 2.1;
- a Markov Population Model $\mathcal{X} = (\mathbb{A}, \mathcal{T}, \mathbf{x}_0)$ as described in Definition 2.2;
- a global reachability property with target region $\mathcal{R}$ identified by a target function $\rho : \mathcal{D} \to \mathbb{R}$.

Steps:

1. *Integration of the Fluid and Central Limit differential equations.* Numerically solve the ODE systems (2.6) for the Fluid Approximation $\mathbf{\Phi}(t)$, and (2.7) and (2.8) for the mean $\mathbf{E}(t)$ and covariance $\mathbf{C}(t)$ of the Gaussian noise of the Central limit Approximation;

2. *Computation of the fluid estimation $t_{\mathcal{R}}$.* Compute the fluid estimation $t_{\mathcal{R}}$ of the hitting time by solving $t_{\mathcal{R}} = \inf\{t > 0 \mid \rho(\mathbf{\Phi}(t)) < 0\}$;

3. *Computation of the mean and covariance of the Gaussian approximation.* Identify the mean μ_{hit} and variance σ^2_{hit} of the Gaussian approximation of the hitting time defined in (4.5) by solving

$$\mu_{hit} = t_{\mathcal{R}} - \frac{1}{\sqrt{N}} \frac{\nabla\rho(\mathbf{\Phi}(t_{\mathcal{R}})) \cdot \mathbf{E}(t_{\mathcal{R}})}{\nabla\rho(\mathbf{\Phi}(t_{\mathcal{R}})) \cdot \mathbf{F}(\mathbf{\Phi}(t_{\mathcal{R}}))}$$

and

$$\sigma^2_{hit} = -\frac{1}{\sqrt{N}} \frac{\nabla\rho(\mathbf{\Phi}(t_{\mathcal{R}})) \cdot \mathrm{diag}(\mathbf{C}(t_{\mathcal{R}}))}{\nabla\rho(\mathbf{\Phi}(t_{\mathcal{R}})) \cdot \mathbf{F}(\mathbf{\Phi}(t_{\mathcal{R}}))},$$

where $\mathrm{diag}(\mathbf{C}(t_{\mathcal{R}}))$ is the vector of diagonal elements of $\mathbf{C}(t_{\mathcal{R}})$.

4. *Computation of the reachability probability.* Let $f(t \mid \mu, \sigma^2)$ be the probability density function of a Gaussian distribution in t with mean μ and variance σ^2. Approximate the global reachability probability $\mathbb{P}_{\mathcal{R}}^{(N)}(T)$ by

$$\mathbb{P}_{\mathcal{R}}^{(N)}(T) \sim \int_{-\infty}^{T} f(t \mid \mu_{hit}, \sigma_{hit}^2) dt. \tag{4.6}$$

The asymptotic correctness of the approximation of the reachability probability $\mathbb{P}_{\mathcal{R}}(T)$ is guaranteed by the following result, which is a straightforward corollary of Theorem 4.1.

Theorem 4.2 *Let $\mathbb{P}_{\mathcal{R}}^{(N)}(T)$ be the exact value of the global reachability probability defined in (4.3), and let $\widetilde{\mathbb{P}}_{\mathcal{R}}^{(N)}(T) = \int_0^T f(t \mid \boldsymbol{\mu}_{hit}, \boldsymbol{\sigma}_{hit}^{\mathbf{2}}) dt$ be the Gaussian approximation computed in (4.6). Then, under the assumptions of Theorem 4.1, it holds that $\lim_{N\to\infty} || \mathbb{P}_{\mathcal{R}}^{(N)}(T) - \widetilde{\mathbb{P}}_{\mathcal{R}}^{(N)}(T) || = 0$.*

4.4.3 System Size Expansion

In the same setting of the algorithm defined in Section 4.4.2, let $\boldsymbol{\Phi}(t)$ be the Fluid Approximation of a Markov Population Model described by a stochastic process $\hat{\mathbf{X}}^{(N)}(t)$, and let $\mathbf{Z}^*(t)$ be the System Size Expansion of Section 2.4.3 with mean $\mathbf{E}^*(t)$ and covariance $\mathbf{C}^*(t)$ given by 2.10 and 2.11, respectively. The approximation $\mathbf{Z}^*(t)$ of the noise around the fluid limit $\boldsymbol{\Phi}(t)$ of $\hat{\mathbf{X}}^{(N)}(t)$ can be used to define the *System Size Expansion or Higher Order Approximation $\widetilde{t}_{\mathcal{R}}^{(N)}$ of the hitting-time $t_{\mathcal{R}}^{(N)}$* given by

$$\widetilde{t}_{\mathcal{R}}^{(N)} = t_{\mathcal{R}} - \frac{1}{\sqrt{N}} \frac{\nabla\rho(\boldsymbol{\Phi}(t_{\mathcal{R}})) \cdot \mathbf{Z}^*(t_{\mathcal{R}})}{\nabla\rho(\boldsymbol{\Phi}(t_{\mathcal{R}})) \cdot \mathbf{F}(\boldsymbol{\Phi}(t_{\mathcal{R}}))}. \tag{4.7}$$

While in (4.5) the CLA guarantees that $\mathbf{Z}(t)$ is a Gaussian process (EK05), in (4.7) there is no limit result that characterises the nature of the distribution of the higher-order approximation $\mathbf{Z}^*(t)$ (and, thus, of the stochastic variable $\widetilde{t}_{\mathcal{R}}^{(N)}$ defined in (4.7)). To tackle this problem and construct a plausible probability density function for $\widetilde{t}_{\mathcal{R}}^{(N)}$, we leverage the same *moment reconstruction* technique based on the *maximum entropy principle* of Section 2.4.4. Hence, to improve the estimation of the hitting

time $t_{\mathcal{R}}$ given by (4.5), we consider the first and second moments of $\widetilde{t}_{\mathcal{R}}^{(N)}$, which are given by

$$\mathbf{E}\left[\widetilde{t}_{\mathcal{R}}^{(N)}\right] = t_{\mathcal{R}} - \frac{1}{\sqrt{N}} \frac{\nabla\rho(\mathbf{\Phi}(t_{\mathcal{R}})) \cdot \mathbf{E}^*(t_{\mathcal{R}})}{\nabla\rho(\mathbf{\Phi}(t_{\mathcal{R}})) \cdot \mathbf{F}(\mathbf{\Phi}(t_{\mathcal{R}}))} \tag{4.8}$$

and

$$\mathbf{C}\left[\widetilde{t}_{\mathcal{R}}^{(N)}\right] = t_{\mathcal{R}} - \frac{1}{\sqrt{N}} \frac{\nabla\rho(\mathbf{\Phi}(t_{\mathcal{R}})) \cdot \mathrm{diag}(\mathbf{C}^*(t_{\mathcal{R}}))}{\nabla\rho(\mathbf{\Phi}(t_{\mathcal{R}})) \cdot \mathbf{F}(\mathbf{\Phi}(t_{\mathcal{R}}))}. \tag{4.9}$$

Then, in this case, the maximum entropy principle states that the best approximation (in terms of the Shannon entropy) for a one-dimensional probability distribution given its first and second moments, μ and σ^2 respectively, is the Gaussian distribution $\mathcal{N}(\mu, \sigma^2)$. Hence, we conclude that $\widetilde{t}_{\mathcal{R}}^{(N)} \sim \mathcal{N}(\mu^*_{hit}, \sigma^{2*}_{hit})$, where $\mu^*_{hit} = \mathbf{E}[\widetilde{t}_{\mathcal{R}}^{(N)}]$ and $\sigma^{2*}_{hit} = \mathbf{C}[\widetilde{t}_{\mathcal{R}}^{(N)}]$ are the mean (4.8) and the variance (4.9), respectively.

In conclusion, if we consider the higher-order approximation $\widetilde{t}_{\mathcal{R}}^{(N)}$ of the hitting time $t_{\mathcal{R}}^{(N)}$ defined in (4.7), the algorithm of our model checking procedure keeps the form described in Section 4.4.2, substituting the occurrences of the mean $\mathbf{E}(t)$ and covariance $\mathbf{C}(t)$ of the Central Limit Approximation with the mean $\mathbf{E}^*(t)$ and covariance $\mathbf{C}^*(t)$ of the higher-order approximation given by (2.11) and (2.10).

Remark 4.1 *Since by definition $\mathbf{C}(t) = \mathbf{C}^*(t)$, in practice the higher-order approximation (4.7) improves the estimation given by the Central Limit Approximation (4.5) by subtracting an N-dependent correction term from its mean, which was actually equal to the Fluid estimation $t_{\mathcal{R}}$ (indeed, by integration of (2.7), we have that $\mathbf{E}(t) \equiv 0$).*

4.5 Expertimental Results

In the following, we describe the experimental results obtained on the peer-to-peer software update process introduced in Section 4.2 where we set $\lambda_{\mathrm{on}_O} = \lambda_{\mathrm{off}_O} = \lambda_{\mathrm{off}_T} = \lambda_{\mathrm{on}_U} = \lambda_{\mathrm{off}_U} = 0.4$, $\lambda_{\mathrm{out}_O} = 0.9$, $\lambda_{\mathrm{ext}_O} = 0.008$, $\lambda_{\mathrm{exp}_U} = 0.001$ and $\lambda_{\mathrm{update}} = 1$.

First, we considered reachability properties characterised by *linear* target functions. In particular, we chose $\rho(\hat{\mathbf{X}}) = 0.95 - \hat{X}_{oldON} - \hat{X}_{upON}$,

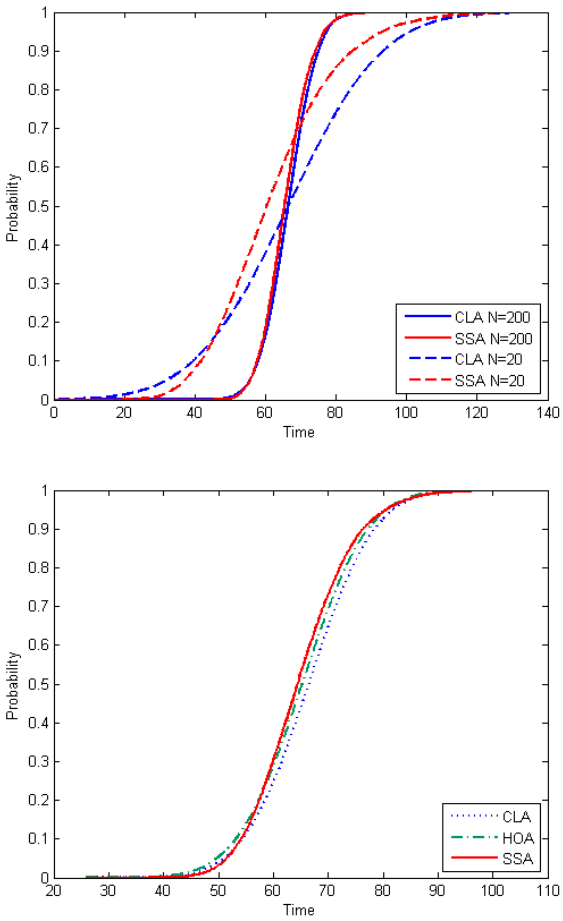

Figure 11: Results of the experimental analysis of the running example with $\rho(\hat{\mathbf{X}}) = 0.95 - \hat{X}_{oldON} - \hat{X}_{upON}$ and $\hat{\mathbf{x}}_0 = (0.9, 0, 0, 0.1, 0)$. Top: Comparison of reachability probabilities obtained by Central Limit Approximation (CLA) and Gillespie's statistical algorithm (SSA) $N = 20$ and $N = 200$. Bottom: Comparison of reachability probabilities obtained by CLA, SSA and the System Size Expansion of Section 4.4.3 (HOA) for $N = 100$.

which is used to check whenever the number of updated nodes reaches 95% of the network size (see Example 4.1). Figure 11 shows the reachability probabilities $\mathbb{P}^{(N)}_{\mathcal{R}}(T) = \mathbb{P}\{t^{(N)}_{\mathcal{R}} \leq T\}$ for three population sizes: $N = 20, 100, 200$. On the left of Figure 11, the Gaussian estimation (4.5) obtained by the Central Limit Approximation (CLA) is compared with a statistical estimation (the Gillespie's Stochastic Simulation Algorithm (SSA)) computed over 10000 simulation runs. As expected the accuracy of the estimation increases with N and is already very good for $N = 200$. On the right of Figure 11, for the case $N = 100$, we show also the System Size Expansion defined in Section 4.4.3, which corrects the mean obtained by the CLA, improving the quality of the approximation. In Table 3 Top, we report the maximum and mean absolute errors obtained by the CLA and the higher-order approximation. Again as expected, the results improve with N and the quality is already quite good for $N = 100$. Moreover, when $N = 100$ and $N = 200$, the higher-order approximation reduces the errors of more than 50%. In Table 3 Bottom, we show also the gain in terms of computational cost achieved by the CLA and the Higher Order Approximation. Indeed, the approximation methods illustrated in this chapter are up to 16 times faster than the statistical estimate even for small population sizes N.

In the second set of experiments, instead, we considered reachability properties identified by *non-linear* target functions. We chose to verify the efficiency of the communication across the network by checking whenever the throughput $X_{oldON} * X_{upON}$ gets below a certain N-dependent threshold. In particular, we set $\rho(\hat{\mathbf{X}}) = \hat{X}_{oldON} * \hat{X}_{upON} - 0.006$. In this case, the experimental results showed that, to reach a good level of accuracy in the approximation, much larger population sizes have to be considered (probably due to the fact that the error in the estimation of gets amplified by the product $\hat{X}_{oldON} * \hat{X}_{upON}$, hence larger N are needed for the method to converge). Indeed, Figure 12 compares the results obtained by the CLA, the SSA and the higher-order approximation for two population sizes: $N = 1000$ and $N = 10000$. When $N = 1000$, the CLA performs poorly in estimating the reachability probability and the quality of the approximation slightly improves when we compute

N	max(erCLA)	$\mathbb{E}$[erCLA]	max(erHOA)	$\mathbb{E}$[erHOA]
20	0.1453	0.0443	0.0908	0.0313
100	0.0931	0.0128	0.0415	0.0066
200	0.0623	0.0063	0.0171	0.0023

N	costSSA	costCLA	costHOA
20	6.1828	1.0017	0.8678
100	10.9279	1.0017	0.8678
200	16.6240	1.0017	0.8678

Table 3: Top: Maximum and mean absolute error on the reachability probability estimations obtained by the Central Limit Approximation (max(erCLA), $\mathbb{E}$[erCLA]) and the System Size Expansion of Section 4.4.3 (max(erHOA), $\mathbb{E}$[erHOA]) in the experiments of Figure 11. Bottom: Average computational costs in Seconds of the Gillespie Simulation (costSSA), the Central Limit Approximation (costCLA) and the System Size Expansion (costHOA) in the case of the Experiment of Figure 11.

the System Size Expansion, which however still fails to capture the slope of the cumulative distribution function obtained by the SSA, due to an inaccuracy in the prediction of the variance. Only when we consider population sizes of the order of 10000, the method starts to converge and the higher-order approximation is finally able to efficiently predict the (estimated) true reachability probability. Including higher-order terms in the method of Section 4.4.3 may improve this scenario, too.

4.6 Discussion

In this chapter, we have presented a model checking procedure for global reachability properties of Markov Population Models based on stochastic approximations of the system behaviour. In particular, we exploited the Central Limit Approximation (CLA) and the System Size Expansion (IOS) as in Chapter 3, but for a different class of requirements: the global reachability properties. Indeed, in Chapter 3, we considered queries based on counting how many agents satisfy a local specification, obtaining a reachability problem in which the target region $\mathcal{R}$ is guaranteed

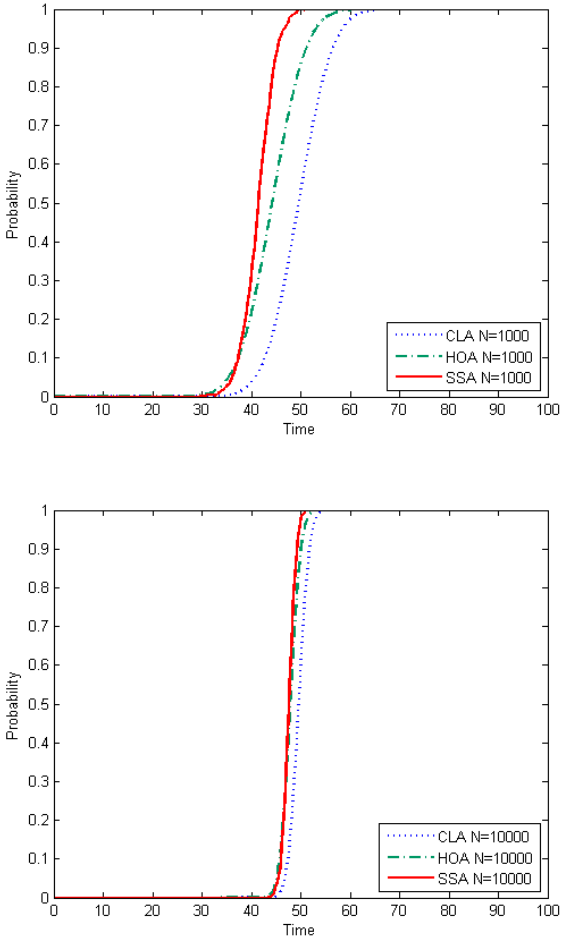

Figure 12: Results of the experimental analysis of the running example with $\rho(\hat{\mathbf{X}}) = \hat{X}_{oldON} * \hat{X}_{upON} - 0.006$ and $\hat{\mathbf{x}}_0 = (0, 0.9, 0, 0, 0.1)$. Comparison of reachability probabilities obtained by Central Limit Approximation (CLA), Gillespie's simulation algorithm (SSA) and the higher-order approximation of Section 4.4.3 (HOA) for $N = 1000$ (top) and $N = 10000$ (bottom).

to be absorbing. Here, instead, we consider arbitrary regions $\mathcal{R}$, defined by differentiable functions on collective variables, which cannot be made absorbing in a consistent way with the CLA or the IOS. Hence, we relied on a different mathematical machinery, based on a Gaussian approximation of the time instant in which the trajectory of the population model enters $\mathcal{R}$. Moreover, we improved the accuracy of the estimation considering the IOS as an higher-order approximation of the (first two) moments of the reachability probability distribution. The method was experimentally validated on a peer-to-peer software update process.

The main limitation of our methodology is that it requires the fluid limit trajectory to enter the target region $\mathcal{R}$ associated with the reachability constraint. And even when this happens, the quality of our approximation is correlated with the unimodality of the hitting time distribution: if the true distribution is multimodal, then the accuracy of our method will be hampered (BGH12). This can happen if the fluid trajectory passes close to the boundary of $\mathcal{R}$ without crossing it. We are currently investigating possible ways of overcoming these limitations.

Other directions for future work are the release of an implementation, the investigation and characterisation of the effect of higher-order approximations on the estimate of the reachability probability, and the application of the framework on larger case studies.

Chapter 5

Mean-Field Approximation for Timed Properties

5.1 Overview

In this chapter, we extend (BH15) to more complex *time-bounded properties* specified by *Deterministic Timed Automata* endowed with a single clock (AD94; BK08; DHS09). As in (BH15; HBC13; BL13a; CHKM11a), we combine the agent and the DTA specification with a product construction, obtaining a *Time-Inhomogeneous Markov Renewal Process* (Cin13). We then exploit results (BH12a; Hay12), defining the Fluid Approximation of this type of models as the solution of a system of *Delay Differential Equations* (DDE) (DN08). Other works dealing with the verification of DTA properties are (Fu13; BCH$^+$11; CDKM13; CHKM11b).

We introduce a new fast and efficient Fluid Model Checking procedure to accurately approximate the probability that a single agent satisfies a single-clock DTA specification up to time T. Similarly to (BH15), the technique is based of the Fast Simulation Theorem, and couples the Fluid Approximation of the collective system with a set of Delay Differential Equations for the transient probability of the Time-Inhomogeneous Markov Renewal Process obtained by the product construction between the single agent and the DTA specification.

In the chapter, we discuss the *theoretical aspects* of our approach, proving the *convergence* of the estimated probability to the true one in the limit of an infinite population. We also show the procedure at work on a running example of a simple epidemic process, emphasising the quality of the approximation and the gain in terms of computational time. Finally, by exploiting the construction of (BL13a; Hay12), we also show how to define a set of DDEs approximating the mean number of agents satisfying a single-clock DTA specification up to time T.

The chapter is organized as follows: In Section 5.2, we introduce the main example, recalling also the modelling notation. The DTA specification for the timed properties is discussed in Section 5.3. In Section 5.4, we present our FMC procedure, defining the DDEs for the probability that the single agent satisfies the timed property. We also show how to adapt our verification technique to compute the mean number of agents that meet the DTA requirement. In Section 5.5, we discuss the quality of the approximation on the epidemic example. Finally, in Section 5.6, we discuss the results of the chapter and possible extensions.

The content of this chapter was published in (BL15).

5.2 Running Example

The running example that we consider is a simple *SIS model*, describing the spreading of a disease inside a population. All agents belong to the same Agent Class $\mathbb{A}$, depicted in Figure 13, and can be either *susceptible* (S) or *infected* (I). When they are *susceptible*, they can be infected (*inf*), and when they are *infected*, they can either pass the infection (*pass*) or recover (*rec*). Hence, the state $\mathbf{X}^{(N)}(t)$ of the Markov Population Model is

$$\mathbf{X}^{(N)}(t) = (X_S^{(N)}(t), X_I^{(N)}(t)),$$

and we define 2 global transitions:

$$\tau_r = (\{I \xrightarrow{rec} S\}, f_r^{(N)})$$

and

$$\tau_i = (\{S \xrightarrow{inf} I, I \xrightarrow{pass} I\}, f_i^{(N)}).$$

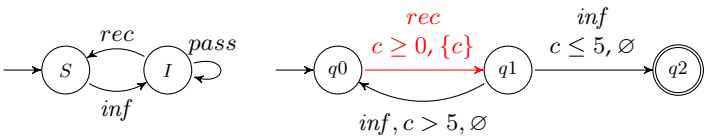

Figure 13: The agent class $\mathbb{A}$ (left) and property $\mathbb{D}$ (right) of the running example illustrated in Section 5.2.

The former, τ_r, mimics the recovery of one entity inside the population, while τ_i synchronises two local actions, namely $S \xrightarrow{inf} I$ and $I \xrightarrow{pass} I$, and models the transmission of the virus from an infected agent to a susceptible one. Finally, the rate functions depend on the number of agents involved in the transitions and follow the classical *rule of mass action* (AB00):

$$f_r^{(N)}(t) = k_r X_I^{(N)}(t)$$

and

$$f_i^{(N)}(t) = \frac{1}{N} k_i X_S^{(N)}(t) X_I^{(N)}(t),$$

where $k_r, k_i \in \mathbb{R}_{\geq 0}$.

5.3 Local Timed Properties

We are interested in properties specifying how a single agent behaves in *time*. In order to monitor such requirements, we assign to it a unique *personal clock*, which starts at time 0 and can be reset whenever the agent undergoes specific transitions. In this way, the properties that we consider can be specified by a *single-clock Deterministic Timed Automata* (DTA) (AD94; CHKM11a), which keeps track of the behaviour of the single agent with respect to its personal clock. Moreover, since we want to exploit the Fast Simulation Theorem 2.3, we restrict ourselves to *time bounded* properties and, hence, we assign to the DTA a finite *time horizon* $T < +\infty$, within which the requirement must be true.

Definition 5.1 (Timed Properties) *A timed property for a single agent in Agent Class $\mathbb{A}$ is specified as a single-clock DTA of the form*

$$\mathbb{D} = \mathbb{D}(T) = (T, \mathscr{L}, c, \mathcal{CC}, Q, q_0, F, \rightarrow),$$

where:

- *$T < +\infty$ is the finite time horizon;*
- *$\mathscr{L}$ is the label set of the Agent Class $\mathbb{A}$;*
- *c is the personal clock;*
- *$\mathcal{CC}$ is the set of constraints on the clock c, and are thus conjunctions of atoms of the form $c < \lambda$, $c \leq \lambda$, $c \geq \lambda$ or $c > \lambda$ for $\lambda \in \mathbb{Q}$;*
- *Q is the (finite) set of states;*
- *$q_0 \in Q$ is the initial state;*
- *$F \subseteq Q$ is the set of final (or accepting) states;*
- *$\rightarrow \subseteq Q \times \mathscr{L} \times \mathcal{CC} \times \{\varnothing, \{c\}\} \times Q$ is the edge relation.*

Moreover, the single-clock DTA $\mathbb{D}$ has to satisfy:

- *(determinism) for each initial state $q \in Q$, label $\alpha \in \mathscr{L}$, clock constraint $c_{\bowtie} \in \mathcal{CC}$, and clock valuation $\eta(c) \in \mathbb{R}_{\geq 0}$, there exists exactly one edge $q \xrightarrow{\alpha, c_{\bowtie}, r} q'$ such that $\eta(c) \models_{\mathcal{CC}} c_{\bowtie}$[1];*
- *(absorption) the final states are all absorbing.*

A timed property $\mathbb{D}$ is assessed over the time-bounded paths (of total duration T) of the Agent Class $\mathbb{A}$ sampled from the stochastic processes $Z^{(N)}(t)$ and $Z(t)$ defined for the Fast Simulation in Section 2.4.1. The labels of the transitions of $\mathbb{A}$ act as inputs for the DTA $\mathbb{D}$, and the latter is defined in such a way that it *accepts* a time-bounded path σ if and only if the behaviour of the single agent encoded in σ satisfies the property represented by $\mathbb{D}$. Formally, a time-bounded path of $\mathbb{A}$ sampled from $Z^{(N)}(t)$ (resp. $Z(t)$) of the form

$$\sigma = s_0 \xrightarrow{\alpha_0, t_0} s_1 \xrightarrow{\alpha_1, t_1} \dots \xrightarrow{\alpha_n, t_n} s_{n+1}$$

[1]The notation $\eta(c) \models_{\mathcal{CC}} c_{\bowtie}$ stands for the fact that the value of the valuation $\eta(c)$ of c satisfies the clock constraint $c_{\bowtie}$.

such that

$$\sum_{j=0}^{n} t_j \leq T,$$

is *accepted* by $\mathbb{D}$ if and only if there exists a path of $\mathbb{D}$ of the form

$$q_0 \xrightarrow{\alpha_0} q^{(1)} \xrightarrow{\alpha_1} q^{(2)} \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} q^{(n+1)}$$

such that

$$q^{(n+1)} \in F.$$

In the path of $\mathbb{D}$, $q^{(i+1)} \in Q$ denotes the (unique) state that can be reached form $q^{(i)} \in Q$ taking the action $q^{(i)} \xrightarrow{\alpha_i, c_{\bowtie}, r} q^{(i+1)}$ whose clock constraint $c_{\bowtie}$ is satisfied by the clock valuation $\eta(c)$ updated according to time t_i. In the following, we will denote by $\Sigma_{\mathbb{A},\mathbb{D},T}$ the *set of time-bounded paths of* $\mathbb{A}$ *accepted by* $\mathbb{D}$.

Example. We consider the following property for the running example: *within time* T*, the agent gets infected at least once during the* $\Delta = 5$ *time units that follow a recovery*. To verify this requirement, we use the DTA $\mathbb{D} = \mathbb{D}(T)$ represented in Figure 13. If we record the actions of the single agent on $\mathbb{D}$, i.e. we synchronise $\mathbb{A}$ and $\mathbb{D}$, when the agent recovers (rec), $\mathbb{D}$ passes from state $q0$ to $q1$, resetting the personal clock c. After that, if the agent gets infected (inf) within 5 time units, the property is satisfied, and $\mathbb{D}$ passes from state $q1$ to $q2$, which is accepting. If instead the agent is infected (inf) after 5 units of time, $\mathbb{D}$ moves back to state $q0$, and we start monitoring the behaviour of the agent again. In red we highlight the transition that resets the personal clock c in $\mathbb{D}$.

5.4 Theoretical Results

Consider a single agent of class $\mathbb{A} = (S, E)$ in a Markov Population Model

$$\mathcal{X}^{(N)} = (\mathbb{A}, \mathcal{T}^{(N)}, \mathbf{x}_0^{(N)}),$$

and a timed property

$$\mathbb{D} = \mathbb{D}(T) = (T, \mathscr{L}, \Gamma_{\mathcal{S}}, \mathcal{CC}, Q, q_0, F, \rightarrow).$$

Let $\Sigma_{\mathbb{A},\mathbb{D},T}$ be the set of time-bounded paths of $\mathbb{A}$ accepted by $\mathbb{D}$. Moreover, let $Z^{(N)}(t)$ and $Z(t)$ be the two stochastic processes defined for the Fast Simulation in Section 2.4.1. The following result holds true.

Proposition 5.1 *The set $\Sigma_{\mathbb{A},\mathbb{D},T}$ of time-bounded paths of $\mathbb{A}$ accepted by $\mathbb{D}$ is measurable for the probability measures $Prob_{Z^{(N)}}$ and $Prob_Z$ defined over the paths of $Z^{(N)}(t)$ and $Z(t)$, respectively.* ∎

In this chapter, we are interested in the *satisfaction probability* given by

$$P^{(N)}(T) = Prob_{Z^{(N)}}\{\Sigma_{\mathbb{A},\mathbb{D},T}\},$$

that corresponds to the probability that the single agent satisfies property $\mathbb{D}$ within time T in $\mathcal{X}^{(N)}$. Then, the main result that we exploit in our Fluid Model Checking procedure is that, when the population is large enough (i.e N is large enough), $P^{(N)}(T)$ can be accurately approximated by

$$P(T) = Prob_Z\{\Sigma_{\mathbb{A},\mathbb{D},T}\},$$

which is computed over the ICTMC $Z(t)$, whose rates are defined in terms of the Fluid Approximation $\mathbf{\Phi}(t)$ of $\mathcal{X}^{(N)}$. The correctness of the approximation relies on the Fast Simulation Theorem and is guaranteed by the following result.

Theorem 5.1 *For any $T < +\infty$, it holds true that*

$$\lim_{N\to\infty} P^{(N)}(T) = P(T). \quad ∎$$

Moreover, to compute $P(T)$, we define a suitable product construction $\mathbb{A}_\mathbb{D} = \mathbb{A}\otimes\mathbb{D}$, whose state is described by a *Time-Inhomogeneous Markov Renewal Process* (IMRP) (Cin13) that we denote by $Z_{\mathbb{A}_\mathbb{D}}(t)$. In the rest of this section, we define $\mathbb{A}_\mathbb{D}$ and $Z_{\mathbb{A}_\mathbb{D}}(t)$, and we show how to compute the *satisfaction probability* $P(T)$ in terms of the *transient probability* $\boldsymbol{P}(T)$ of $Z_{\mathbb{A}_\mathbb{D}}(t)$.

The Product $\mathbb{A}_\mathbb{D}$

We now introduce the product $\mathbb{A}_\mathbb{D}$ between $\mathbb{A}$ and $\mathbb{D}$, whose state is described by a *Time-Inhomogeneous Markov Renewal Process* (IMRP) $Z_{\mathbb{A}_\mathbb{D}}(t)$ that has rates computed over the Fluid Approximation $\mathbf{\Phi}(t)$ of $\mathcal{X}^{(N)}$.

A *Markov Renewal Process* (MRP) (Cin13) is a jump-process, where the sojourn times in the states can have a general probability distribution. In particular, in the MRP $Z_{\mathbb{A}_\mathbb{D}}(t)$, we allow both *exponentially* and *deterministically-timed* transitions, and in the following, we refer to them as the *Markovian* and *deterministic transitions*, respectively. Since the transition rates of $Z_{\mathbb{A}_\mathbb{D}}(t)$ are time-dependent, $Z_{\mathbb{A}_\mathbb{D}}(t)$ is indeed a *Time Inhomogeneous* MRP.

The product $\mathbb{A}_\mathbb{D}$ has the form

$$\mathbb{A}_\mathbb{D} = (\mathbb{A}, S_\mathbb{D}, \{\mathcal{M}, \mathcal{E}\}, s_{0,\mathbb{D}}, F_\mathbb{D}),$$

and to define it let

$$\delta_1 < \ldots < \delta_k$$

be the (ordered) constants that appear in the clock constraints of $\mathbb{D}$, and extend the sequence with $\delta_0 = 0$ and $\delta_{k+1} = T$. The *state space* $S_\mathbb{D}$ of $\mathbb{A}_\mathbb{D}$ is given by $\{1, \ldots, k+1\} \times S \times Q$. The first element of $S_\mathbb{D}$ identifies a time region of the clock c, and we refer to

$$S_{\mathbb{D}_i} = \{(i, s, q) \mid s \in S, q \in Q\}$$

as the *i-th Time Region* of $S_\mathbb{D}$. The rest of $\mathbb{A}_\mathbb{D}$ is going to be defined in such a way that the agent is in $S_{\mathbb{D}_i}$ if and only if the time of the personal clock c is between δ_{i-1} and δ_i, i.e. c satisfies

$$\delta_{i-1} \leq \eta(c) \leq \delta_i,$$

where η is the valuation of c.

The *set $\mathcal{M}$ of Markovian transitions* of the product $\mathbb{A}_\mathbb{D}$ is the smallest relation such that

$$\forall\, i \in 1, \ldots, k+1, \quad \frac{s \xrightarrow{\alpha} s' \in E \ \wedge \ q \xrightarrow{\alpha, c_{\bowtie}, \varnothing} q' \in \rightarrow \ \wedge \ [\delta_{i-1}, \delta_i] \models c_{\bowtie}}{(i, s, q) \xrightarrow{\alpha} (i, s', q') \in \mathcal{M}}, \tag{5.1}$$

$$\forall\, i \in 1, \ldots, k+1, \quad \frac{s \xrightarrow{\alpha} s' \in E \ \wedge \ q \xrightarrow{\alpha, c_{\bowtie}, \{c\}} q' \in \rightarrow \ \wedge \ [\delta_{i-1}, \delta_i] \models c_{\bowtie}}{(i, s, q) \xrightarrow{\alpha} (1, s', q') \in \mathcal{M}}. \tag{5.2}$$

Intuitively, rule (5.1) synchronises the local transitions

$$s \xrightarrow{\alpha} s' \in E$$

of the Agent Class $\mathbb{A} = (S, E)$ with the transition

$$q \xrightarrow{\alpha, c_{\bowtie}, \varnothing} q' \in \rightarrow$$

of property $\mathbb{D}$ that has the same label α, obtaining a local transition of the form

$$(i, s, q) \xrightarrow{\alpha} (i, s', q') \in \mathcal{M}$$

in $\mathbb{A}_{\mathbb{D}}$ for each time region i whose time interval $[\delta_{i-1}, \delta_i] \subseteq [0, T]$ satisfies the clock constraint $c_{\bowtie}$, i.e. such that $\forall t \in [\delta_{i-1}, \delta_i],\ t \models c_{\bowtie}$.

Rule (5.2), instead, defines the *reset transitions* of the form

$$(i, s, q) \xrightarrow{\alpha} (1, s', q') \in \mathcal{M},$$

that reset the personal clock c either within the 1^{st} Time Region (when $i = 1$), or by bringing the agent *back to* the 1^{st} Time Region. In the following, we denote by $\mathcal{R} \subset \mathcal{M}$ the *set of the reset transitions*.

To describe the deterministic transitions of $\mathbb{A}_{\mathbb{D}}$, instead, we define a set $\mathcal{E}$ of *clock events*. Each clock event has the form

$$e = (\mathcal{A}_e, \Delta, p_e),$$

where

- $\mathcal{A}_e \subset S_{\mathbb{D}_i}$ is the *active set*,
- Δ is the *duration*,
- $p_e : \mathcal{A}_e \times S_{\mathbb{D}} \longrightarrow [0, 1]$ is the *probability distribution*.

If the agent enters $\mathcal{A}_e$, that is the sets of states in which e can be active, a countdown starts from Δ. When this elapses, e_i is deactivated and the agent is immediately moved to a new state sampled from the distribution

$$p_e((i, s, q), \cdot) : S_{\mathbb{D}} \longrightarrow [0, 1],$$

where $(i, s, q) \in \mathcal{A}_e$ is the state in which the agent is when the countdown hits 0. Moreover, e_i is deactivated also when the agent takes a reset transition. In $\mathbb{A}_{\mathbb{D}}$, we define:

- one clock event $e_i \in \mathcal{E}$ for each time region $S_{\mathbb{D}i}$, $i = 2, \ldots, k$;
- $\ell + 1$ clock events $e_1^0, e_1^1, \ldots, e_1^\ell \in \mathcal{E}$ for the 1^{st} Time Region, where ℓ is the number of reset events $(1, s, q) \xrightarrow{\alpha} (1, s', q') \in \mathcal{R}$ defined by (5.2) with $i = 1$.

For $i > 1$, we have $\mathcal{A}_i = S_{\mathbb{D}i}$, $\Delta_i = \delta_i - \delta_{i-1}$, and the probability distribution given by

$$p_i((i, s, q), (i', s', q')) = \begin{cases} 1 & \text{if } i' = i + 1, s' = s, q' = q, \\ 0 & \text{otherwise.} \end{cases} \tag{5.3}$$

By definition, each clock event e_i with $i > 1$ connects each state $(i, s, q) \in \mathcal{A}_i$ with $(i + 1, s, q) \in S_{\mathbb{D}i+1}$, hence, when the duration Δ_i of e_i elapses, the clock event moves the agent from its state to the equivalent one in the next time region. When $i = 1$, instead, the duration and the probability distribution of each clock event e_1^j, $j = 1, \ldots, \ell$, are defined in the same way as before (i.e. $\Delta_1^j = \delta_1 - \delta_0 = \delta_1$, and p_1^j is given by (5.3)), but the activation sets are now subsets of $S_{\mathbb{D}_1}$. Indeed, since each reset transition

$$(1, s, q) \xrightarrow{\alpha_j} (1, s', q') \in \mathcal{R}$$

initiates the clock, for each of them, we define a clock event e_1^j, whose activation set $\mathcal{A}_i^j$ is the set of states in $S_{\mathbb{D}_1}$ that can be reached by the agent *after* it has taken the reset transition $(1, s, q) \xrightarrow{\alpha_j} (1, s', q')$. Furthermore, we have to define an extra clock event e_1^0, with $\mathcal{A}_1^0 = S_{\mathbb{D}_1}$, $\Delta_1^0 = \delta_1$, and p_1^0 given by (5.3), that is the only clock event initiated at time $t = 0$ (and not by the agent entering $\mathcal{A}_1^0$). Indeed, we require for the *initial state* $s_{0,\mathbb{D}}$ of $\mathbb{A}_\mathbb{D}$ to be one of the states of the form $(1, s, q_0)$, where $s \in S$ and q_0 is the initial state of $\mathbb{D}$ (hence, $s_{0,\mathbb{D}}$ belongs to $\mathcal{A}_1^0$). Finally, since the probability distributions p_1^j, $\forall j$, are all defined as in (5.3), also the clock events of the 1^{st} Time Region move the agent from a state to the equivalent one in the next time region (the 2^{nd}), when the countdown from $\Delta_1^j = \delta_1$ elapses. In the following, we denote by

$$(i, s, q) \dashrightarrow_e (i + 1, s, q)$$

the deterministic transition from $(i, s, q) \in S_{\mathbb{D}i}$ to $(i+1, s, q) \in S_{\mathbb{D}i+1}$ encoded by $e \in \mathcal{E}$, and by

$$\boldsymbol{\nu}_{e,s,q} = \mathbf{1}_{(i+1,s,q)} - \mathbf{1}_{(i,s,q)}$$

its update vector. The last component of $\mathbb{A}_{\mathbb{D}}$ that we define is the *set of final states* $F_{\mathbb{D}}$, which is given by

$$F_{\mathbb{D}} = \{(i, s, q) \in S_{\mathbb{D}} \mid q \in F\}.$$

Example. Figure 14 represents the product $\mathbb{A}_{\mathbb{D}}$ between the Agent Class $\mathbb{A}$ and the property $\mathbb{D}$ of the running example described in Section 5.2 and depicted in Figure 13. The state $(1, I, q1)$ that cannot be reached by the single agent is omitted. The black transitions are the Markovian transitions without reset; the red transitions are the Markovian transitions that reset the clock; finally, we define 2 clock events, e_1^0 and e_1^1, with duration $\Delta = 5$ for the 1^{st} Time Region, and the dashed green (resp. blue) transitions are the deterministic transitions encoded by e_1^0 (resp. e_1^1). In blue, we also highlight the states that belong to the activation set $\mathcal{A}_{e_1^1}$ (while $\mathcal{A}_{e_1^0}$ is the whole 1^{st} Time Region). Intuitively, the agent can be found in one of the states belonging to the 1^{st} Time Region whenever its personal clock c is between 0 and 5, i.e. less that 5 time units have passed since $t = 0$ or since a recovery rec. In a similar way, the agent is in the 2^{nd} Time Region when the valuation of c is above 5. Moreover, when the the duration of the clock events elapses (i.e. the countdown from 5 hits 0), the agent is moved from the 1^{st} Time Region to the 2^{nd} Time Region by the deterministic green and blue transitions, that indeed have duration $\Delta = 5$ and are initiated at $t = 0$ or by the reset actions rec, respectively. At the end, the agent is in one of the final states $((1, S, q2),\ (1, I, q2),\ (2, S, q2)$ or $(2, I, q2))$ at time T, if it meets property $\mathbb{D}$ within time T, i.e. within T, the agent has been infected during the 5 time units that follow a recovery. Hence, to verify $\mathbb{D}$, we will compute the probability of being in one of the final states of $\mathbb{A}_{\mathbb{D}}$ at time T.

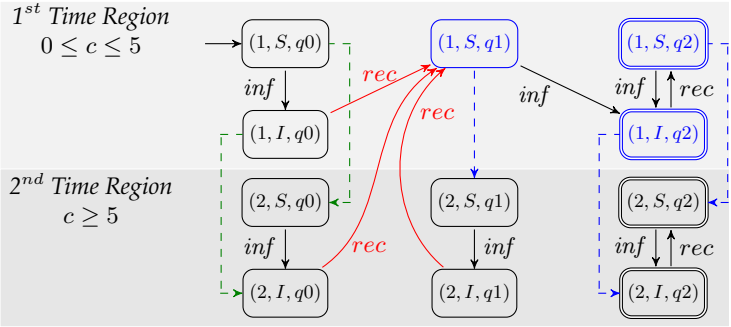

Figure 14: The Agent Class $\mathbb{A}_{\mathbb{D}}$ associated with the DTA $\mathbb{D}$ of the running example.

The IMRP $Z_{\mathbb{A}_{\mathbb{D}}}(t)$ and the Satisfaction Probability $P(T)$

Now we show how to formally define the IMRP $Z_{\mathbb{A}_{\mathbb{D}}}(t)$ that describes the state of the product $\mathbb{A}_{\mathbb{D}}$ in exploiting the Fluid Approximation and the Fast Simulation Theorem. In particular, we derive the *Delay Differential Equations* (DDE) (Cin13) for the *transient probability* $\boldsymbol{P}(t)$ of $Z_{\mathbb{A}_{\mathbb{D}}}(t)$, in terms of which we compute the *satisfaction probability* $P(T)$.

Let $\boldsymbol{\Phi}(t)$ be the Fluid Approximation of the Markov Population Model $\mathcal{X}^{(N)}$. To define the transient probability $\boldsymbol{P}(t)$ of $Z_{\mathbb{A}_{\mathbb{D}}}(t)$, we exploit the fact that, in the case of an IMRP, we have (cf. (Cin13))

$$\frac{d\boldsymbol{P}}{dt}(t) = \boldsymbol{M}\left(\boldsymbol{\Phi}(t)\right)\boldsymbol{P}(t) + \boldsymbol{D}\left(\boldsymbol{\Phi}(t), \boldsymbol{P}(t)\right).$$

In this equation, $\boldsymbol{M}(\boldsymbol{\Phi}(t))$ is the *generator matrix* for the Markovian transitions, and $\boldsymbol{D}(\boldsymbol{\Phi}(t), \boldsymbol{P}(t))$ accounts for the deterministic events. The elements of $\boldsymbol{M}(\boldsymbol{\Phi}(t))$ are computed following the same procedure that was described in Section 2.4.1, where the multiplicity of each transition $(i, s, q) \xrightarrow{\alpha} (i, s', q') \in \mathcal{M}$ in $\mathbb{A}_{\mathbb{D}}$ is always equal to 1 (one single agent) and the Lipschitz limit $f_\alpha(\boldsymbol{\Phi}(t))$ of α is that of the rate of the transition $s \xrightarrow{\alpha} s'$

in $\mathcal{X}^{(N)}$ from which α was derived (by rules (5.1) or (5.2)).

To define the components of $\boldsymbol{D}(\boldsymbol{\Phi}(t), \boldsymbol{P}(t))$, instead, consider any clock event $e = (\mathcal{A}_i, \Delta_i, p_i) \in \mathcal{E}$, *except* e_1^0 (whose contribution will be computed later on[2]). Choose one of the deterministic transitions

$$(i, s, q) \dashrightarrow_{e_i} (i+1, s, q) \tag{5.4}$$

encoded by e_i. The agent takes this transition at time t when:

1. it entered $\mathcal{A}_i \subseteq S_{\mathbb{D}_i}$ at time $t - \Delta_i$ (initiating its personal clock),
2. it is in state $(i, s, q) \in \mathcal{A}_i$ at time t (when the duration of e_i elapses)

Hence, to compute the term that corresponds to transition (5.4) in the deterministic term $\boldsymbol{D}(\boldsymbol{\Phi}(t), \boldsymbol{P}(t))$, we need to:

1. record the flux of probability that entered $\mathcal{A}_i$ at time $t - \Delta_i$,
2. multiply the flux of Step (1) by the probability that the agent reaches $(i, s, q) \in \mathcal{A}_i$ at time t, conditional on the state at which it entered $\mathcal{A}_i$ at $t - \Delta_i$.

To compute the probability of Step (2), we need to keep track of the dynamics of the agent while the clock event e_i is active. For this purpose, let $\bar{\mathcal{A}}_i$ be the activation set $\mathcal{A}_i$ of e_i extended to contain an extra state $s_{out} = (i, s_{out}, q_{out})$. Moreover, define $\bar{\mathcal{M}}$ to be the new set of Markovian transitions obtained from $\mathcal{M} \in \mathbb{A}_{\mathbb{D}}$ in the following way:

- make the reset transitions that start in $\mathcal{A}_i$ finish in s_{out}, i.e. for every

  $$(i, s, q) \xrightarrow{\alpha} (i', s', q') \in \mathcal{R} \subset \mathcal{M},$$

  we define a transition

  $$(i, s, q) \xrightarrow{\alpha} s_{out} \in \bar{\mathcal{M}}$$

- make s_{out} absorbing[3].

[2] If e is one of the events of the 1^{st} Time Region, i.e. $e = e_1^j$, for some $j = 1, \ldots, \ell$, in this section, we drop the index j to ease the notation, i.e. we write $e_1^j = e_1 = (\mathcal{A}_1, \Delta_1, p_1)$.

[3] The absorbing state s_{out} is needed for the probability $\boldsymbol{Y}_i(t)$ of step (2) to be well defined. Indeed, the agent can deactivate e_i by taking a reset transition.

Let $\mathbf{G}_i(\mathbf{\Phi}(t)) \in \text{Matr}(|\bar{\mathcal{A}}_i| \times |\bar{\mathcal{A}}_i|)$ be the time-dependent matrix s.t.

$$(\mathbf{G}_i(\mathbf{\Phi}(t)))_{(i,s,q),(i,s',q')} = \sum_{(i,s,q) \xrightarrow{\alpha} (i,s',q') \in \bar{\mathcal{M}}} \left[\frac{1}{\Phi_s(t)} f_\alpha(\mathbf{\Phi}(t)) \right], \quad (5.5)$$

where again the Lipschitz limit $f_\alpha(t)$ of each $\alpha \in \bar{\mathcal{M}}$ is that of the transition $s \xrightarrow{\alpha} s'$ in $\mathcal{X}^{(N)}$ from which its copy $\alpha \in \mathcal{M}$ was derived (by (5.1) and (5.2)). Moreover, let the diagonal elements of $\mathbf{G}_i(\mathbf{\Phi}(t))$ to be defined so that the rows sum up to zero. Then, we introduce the *probability matrix* $\boldsymbol{Y}_i(t)$, which is computed in terms of the *generator* $\mathbf{G}_i(\mathbf{\Phi}(t))$ according to the following ODEs (see also (BH15)):

$$\begin{cases} \frac{d\boldsymbol{Y}_i}{dt}(t) = \boldsymbol{Y}_i(t)\mathbf{G}_i(\mathbf{\Phi}(t)) - \mathbf{G}(\mathbf{\Phi}(t-\Delta_i))\boldsymbol{Y}_i(t), & \Delta_i \leq t \leq T, \\ \frac{d\boldsymbol{Y}_i}{dt}(t) = \boldsymbol{Y}_i(t)\mathbf{G}_i(\mathbf{\Phi}(t)), & 0 \leq t \leq \Delta_i, \end{cases} \quad (5.6)$$

with $\boldsymbol{Y}_i(0) = \mathbf{I}$. By definition, $(\boldsymbol{Y}_i(t))_{(i,s',q'),(i,s,q)}$ is the Fluid Approximation of the probability of Step (2), i.e. the probability that the agent, which has entered $\mathcal{A}_i$ in state (i,s',q') at time $t-\Delta_i$, moves (in a Markovian way) within $\mathcal{A}_i$ for Δ_i units of time, and reaches $(i,s,q) \in \mathcal{A}_i$ at time t (exactly when e_i elapses).

In terms of the probability matrix $\boldsymbol{Y}_i(t)$, we can now define the component of $\boldsymbol{D}(\mathbf{\Phi}(t), \boldsymbol{P}(t))$ that corresponds to the deterministic transition $(i,s,q) \dashrightarrow_{e_i} (i+1,s,q)$ of the clock event $e_i \in \mathcal{E}$. This component is the element in position $((i,s,q),(i+1,s,q))$ in $\boldsymbol{D}(\mathbf{\Phi}(t), \boldsymbol{P}(t))$, we call it $D_{e_i,s,q}(\mathbf{\Phi}(t), \boldsymbol{P}(t))$, and is given by

$$\begin{aligned} &D_{e_i,s,q}(\mathbf{\Phi}(t), \boldsymbol{P}(t)) = \\ &= \sum_{(i,\bar{s},\bar{q}) \in \mathcal{A}_i} (\boldsymbol{Y}_i(t))_{(i,\bar{s},\bar{q}),(i,s,q)} \Bigg[\mathbf{1}_{\{i>1\}} D_{e_{i-1},\bar{s},\bar{q}}(\mathbf{\Phi}(t-\Delta_i), \boldsymbol{P}(t-\Delta_i)) + \\ &\quad + \mathbf{1}_{\{i=1\}} \sum_{(i',s',q') \xrightarrow{\alpha} (1,\bar{s},\bar{q}) \in \mathcal{R}} \frac{1}{\Phi_{s'}(t)} f_\alpha(\mathbf{\Phi}(t-\Delta_1))(\boldsymbol{P}(t-\Delta_1))_{(i',s',q')} \Bigg], \end{aligned} \quad (5.7)$$

where $(\boldsymbol{P}(t-\Delta_1))_{(i',s',q')}$ is the component in position $(i',s',q') \in S_{\mathbb{D}_{i'}}$ in the vector of the transient probability $\boldsymbol{P}(t-\Delta_1)$ of $Z_{\mathbb{A}_\mathbb{D}}$ at time $t-\Delta_1$. In (5.7), for each state $(i,\bar{s},\bar{q})$ in the activation set $\mathcal{A}_i$, the quantity inside the squared brackets is the probability flux that entered $(i,\bar{s},\bar{q})$ at time $t-\Delta_i$. In particular, when $i>1$, $D_{e_{i-1},\bar{s},\bar{q}}(\boldsymbol{\Phi}(t-\Delta_i),\boldsymbol{P}(t-\Delta_i))$ accounts for the termination of clock event e_{i-1} (i.e. the deterministic transition $(i-1,\bar{s},\bar{q}) \dashrightarrow_{e_i} (i,\bar{s},\bar{q})$ fired at time $t-\Delta_i$). When we consider the 1^{st} Time Region, i.e. $i=1$, instead, each term in the sum over the reset transitions is the flux of probability entering $(1,\bar{s},\bar{q})$ at time $t-\Delta_1$ due to a clock reset. Finally, $(\boldsymbol{Y}_i(t))_{(i,\bar{s},\bar{q}),(i,s,q)}$ is again the probability of reaching $(i,s,q) \in \mathcal{A}_i$ from $(i,\bar{s},\bar{q}) \in \mathcal{A}_i$ in Δ_i units of time.

All the other off-diagonal elements of $\boldsymbol{D}(\boldsymbol{\Phi}(t),\boldsymbol{P}(t))$ can be computed in a similar way, while the diagonal ones are defined so that the rows sum up to zero. Moreover, since at the end $\boldsymbol{D}(\boldsymbol{\Phi}(t),\boldsymbol{P}(t))$ depends on the state of the system at times $t-\Delta_1,\dots,t-\Delta_k$ (through the probabilities $\boldsymbol{Y}_i(t),\ i=1,\dots,k$), we write $\boldsymbol{D}(\boldsymbol{\Phi}(t)) = \boldsymbol{D}(\boldsymbol{\Phi},\boldsymbol{P},\Delta_1,\dots,\Delta_k,t)$. Then, we define the *transient probability* $\boldsymbol{P}(t)$ of the IMRP $Z_{\mathbb{A}_\mathbb{D}}(t)$ as the solution of the following system of DDEs:

$$\begin{aligned}\boldsymbol{P}(t) =& \int_0^t \boldsymbol{M}(s)\boldsymbol{P}(t)ds + \int_0^t \boldsymbol{D}(\boldsymbol{\Phi},\boldsymbol{P},\Delta_1,\dots,\Delta_k,s)ds + \\ &+ \mathbf{1}_{t\geq\Delta_1} \sum_{(s,q)\in S\times Q} y_{e_1^0}\boldsymbol{\nu}_{e_1^0,s,q}. \end{aligned} \tag{5.8}$$

In (5.8), the third term is a deterministic jump in the value of $\boldsymbol{P}(t)$ at time $t=\Delta_1$, and represents the contribution of the clock event e_1^0. In such term, the vectors $\boldsymbol{\nu}_{e_1^0,s,q}$ are the update vectors of the deterministic transitions encoded by e_0^1 (hence, the sum is computed over all the transitions), and the probability $y_{e_1^0}$ is the value at time $t=\Delta_1$ of the component in position $(s_{0,\mathbb{D}},(1,s,q))$ (where $s_{0,\mathbb{D}}$ is the initial state of $\mathbb{A}_\mathbb{D}$) in the matrix $\boldsymbol{Y}_{e_1^0}(t)$ defined by:

$$\frac{d\boldsymbol{Y}_{e_1^0}}{dt}(t) = \boldsymbol{Y}_{e_1^0}(t)\mathbf{G}_1(\boldsymbol{\Phi}(t)), \quad 0\leq t\leq\Delta_1,$$

with $\mathbf{G}_1(\mathbf{\Phi}(t))$ defined as in (5.5), and $\boldsymbol{Y}_{e_1^0}(0) = \boldsymbol{I}$. Hence,

$$y_{e_1^0} = (\boldsymbol{Y}_{e_1^0}(\Delta_1))_{s_{0,\mathbb{D}},(1,s,q)}$$

is the probability that, starting from $s_{0,\mathbb{D}}$, the agents reaches $(1, s, q) \in S_{\mathbb{D}_1}$ at time $t = \Delta_1$ (exactly when the deterministic event

$$(1, s, q) \dashrightarrow_{e_1^0} (2, s, q)$$

fires).

Given the product $\mathbb{A}_{\mathbb{D}}$, the IMRP $Z_{\mathbb{A}_{\mathbb{D}}}(t)$, and its transient probability $\boldsymbol{P}(t)$, the following result holds true.

Proposition 5.2 *There is a one-to-one correspondence between the set $\Sigma_{\mathbb{A},\mathbb{D},T}$ of of $\mathbb{A}$ of duration T accepted by the property $\mathbb{D}$ and the set $AccPath(\mathbb{A}_{\mathbb{D}}, T)$ of accepted paths of duration T of $\mathbb{A}_{\mathbb{D}}$. Hence,*

$$P(T) = Prob_Z\{\Sigma_{\mathbb{A},\mathbb{D},T}\} = Prob_{Z_{\mathbb{A}_{\mathbb{D}}}}\{AccPath(\mathbb{A}_{\mathbb{D}}, T)\} = P_{F_{\mathbb{D}}}(T),$$

where $Prob_{Z_{\mathbb{A}_{\mathbb{D}}}}$ is the probability measure defined by the IMRP $Z_{\mathbb{A}_{\mathbb{D}}}$, and $P_{F_{\mathbb{D}}}(T)$ is the sum of the components of $\boldsymbol{P}(T)$ corresponding to the final states $F_{\mathbb{D}}$ of $\mathbb{A}_{\mathbb{D}}$. ∎

In other words, according to Proposition 5.2, when the population of $\mathcal{X}^{(N)}$ is large enough, $P_{F_{\mathbb{D}}}(T)$ is an accurate approximation of the probability that a (random) single agent in $\mathcal{X}^{(N)}$ satisfies property $\mathbb{D}$ within time T.

Example. For the product $\mathbb{A}_{\mathbb{D}}$ in Figure 14, the non-zero off-diagonal entries of the generator matrix $\boldsymbol{G}_{e_1^1}(\mathbf{\Phi}(t))$ of the clock event e_1^1 are

$$G_{(S,q1)(I,q2)}(t) = k_i \Phi_I(t);$$

$$G_{(S,q2)(I,q2)}(t) = k_i \Phi_I(t);$$

and

$$G_{(I,q2)(S,q2)}(t) = k_r.$$

In terms of $\boldsymbol{G}_{e_1^1}(\mathbf{\Phi}(t))$, we can define $\boldsymbol{Y}_{e_1^1}(t)$, as in (5.6), that is then used in the DDEs (5.8) for the probability $\boldsymbol{P}(t)$. In this latter set of 9 DDEs (one for each state of $\mathbb{A}_{\mathbb{D}}$), we have:

$$\boldsymbol{P}_{(1,S,q1)}(t) = \int_0^t k_r \boldsymbol{P}_{(1,S,q1)}(s)ds - \int_0^t k_i \Phi_I(s)\boldsymbol{P}_{(1,S,q1)}(s)ds + \\ - \int_0^t k_r \boldsymbol{Y}_{(1,S,q1),(1,S,q1)}(s-5,s)\boldsymbol{P}_{(1,S,q1)}(s)ds.$$

Remark 5.1 *The presence of only one clock in $\mathbb{D}$ enables us to define $\mathbb{A}_\mathbb{D}$ in such a way that $Z_{\mathbb{A}_\mathbb{D}}(t)$ is an IMRP. This cannot be done when we consider multiple clocks in $\mathbb{D}$. Indeed, in the latter case, the definition of the stochastic process which describes the state of the product $\mathbb{A}_\mathbb{D}$ is much more complicated, since, when a reset event occurs, we still need to keep track of the valuations of all the other clocks in the model (hence, the dynamics between the time regions of $\mathbb{A}_\mathbb{D}$ is not as simple as in the case of one single clock). In the future, we plan to investigate possible extensions of our model checking procedure to timed properties with multiple clocks, also taking into account the results of (Fu13) and (BCH+11).*

5.4.1 Mean Behaviour of Markov Population Models

The Fluid Model Checking procedure described in the previous sections can be modified to compute the *mean* number of agents that satisfy the timed property $\mathbb{D}$. This can be done by assigning a personal clock to each agent, and monitoring all of them using as Agent Class the product $\mathbb{A}_\mathbb{D}$ defined in Section 5.4. In terms of $\mathbb{A}_\mathbb{D}$, we can build the population model $\mathcal{X}_\mathbb{D}$, with $\mathbb{A}_\mathbb{D}$ as the only Agent Class, and the sum $P_{F_\mathbb{D}}(T)$ of the components corresponding to the final states of $\mathbb{A}_\mathbb{D}$ in the Fluid Approximation $\boldsymbol{\Phi}(t)$ of $\mathcal{X}_\mathbb{D}$ computed at $t = T$ is indeed the mean number of agents satisfying property $\mathbb{D}$ within time T. The construction of $\mathcal{X}_\mathbb{D}$ is not difficult: it follows the procedure described in Chapter 3, where a little extra care has to be taken just in the definition of the global transitions of $\mathcal{X}_\mathbb{D}$. Indeed, if we build for instance the Markov Population Model $\mathcal{X}_\mathbb{D}$ of the running example, we need to consider that the infected individual that passes the virus to an agent in state $(1, S, q0)$ can be now in one of *five* infected states:

$$(1, I, q0), \quad (1, I, q2), \quad (2, I, q0), \quad (2, I, q1), \quad (2, I, q2). \tag{5.9}$$

For this reason, we have to define *five* Markovian global transition in $\mathcal{X}_{\mathbb{D}}$, each of which moves an agent from $(1, S, q0)$ to $(1, I, q0)$ at a rate that is influenced by the number of individuals that are in the infected states (5.9) of $\mathbb{A}_{\mathbb{D}}$. Hence, the Markovian global transitions depend on the counting variables

$$X_{(1,I,q0)}(t), \quad X_{(1,I,q2)}(t), \quad X_{(2,I,q0)}(t), \quad X_{(2,I,q1)}(t), \quad X_{(2,I,q2)}(t).$$

The same reasoning has to be followed for the definition of the infections of the agents in states

$$(1, S, q1), \quad (1, S, q2), \quad (2, S, q0), \quad (2, S, q1), \quad (2, S, q2).$$

At the end, as for the single agent, due to the deterministic events, the Fluid Approximation $\mathbf{\Phi}(t)$ of $\mathcal{X}_{\mathbb{D}}$ is the solution of a system of DDEs similar to (5.8). The definition of these approximating equations for a population model with exponential and deterministic transitions is not new (Hay12), but, even if the results are promising (see Section 5.5), to our knowledge, nobody has yet proven the convergence of the estimation in the limit $N \to +\infty$. We save the investigation of this result for future work.

5.5 Experimental Results

To validate the procedures of Section 5.4, we performed a set of experiments on the running example, where we fixed: $k_i = 1.2$, $k_r = 1$, $\Delta = 5$, and an initial state of the population model with a susceptible-infected ratio of 9:1. As in Figure 14, we let the single agent start in the susceptible state, and we considered three different values of the population size: $N = 250, 500, 1000$. For each N, we compared our procedures with a statistical estimate from 10000 runs, obtained by a dedicated Java implementation of a Discrete Event Simulator (DES).

The errors and the execution times obtained by our FMC procedure (top) and the Fluid Approximation of the mean behaviour (bottom) are reported in Table 4. Regarding the errors, we would like to remark that

Fluid Model Checking

N	MeanRelErr	MaxRelErr	RelErr(T)
250	0.0927	6.4512	0.1043
500	0.0204	1.7191	0.0048
1000	0.0118	0.7846	0.0003

N	TimeDES	TimeFMC	Speedup
250	58.2960	0.4236	137.6204
500	44.0631	0.4236	104.0205
1000	170.9154	0.4236	403.4830

Fluid Approximation of the mean behaviour

N	MeanRelErr	MaxRelErr	RelErr(T)
250	0.1127	0.2316	0.0921
500	0.0289	0.3177	0.0289
1000	0.0117	0.2216	0.0117

N	TimeDES	TimeFluid	Speedup
250	105.5647	0.4294	245.8423
500	415.0635	0.4294	966.6127
1000	1547.0340	0.4294	3602.7806

Table 4: Mean Relative Error (MeanRelErr), Maximum Relative Error (MaxRelErr), and Relative Error at final time (RelErr(T)) of the FMC (top) and the Fluid Approximation of the mean behaviour (bottom) for different values of N. We also enlist the execution times (in seconds) of the DES (TimeDES) and the approximations (TimeFMC and TimeFluid), and the speedups (TimeDES divided by the other times).

the Relative Errors (RE) of both the FMC and the Fluid Approximation reach their maximum in the very first instants of time, when the true satisfaction probability (i.e. the denominator of the REs) is indeed really small, but then they decay really fast as the values of $P_{F_{\mathbb{D}}}(t)$ increase (this can be easily deduced from the values of the mean REs and the REs at final time). As expected, the accuracy of the approximations increases with the population size, and is already reasonably good for $N = 500$. Moreover, the resolution of the DDEs is computationally independent of N, and also much faster (approximatively 3 orders of magnitude in the case of the Fluid for $N = 1000$) than the simulation based method.

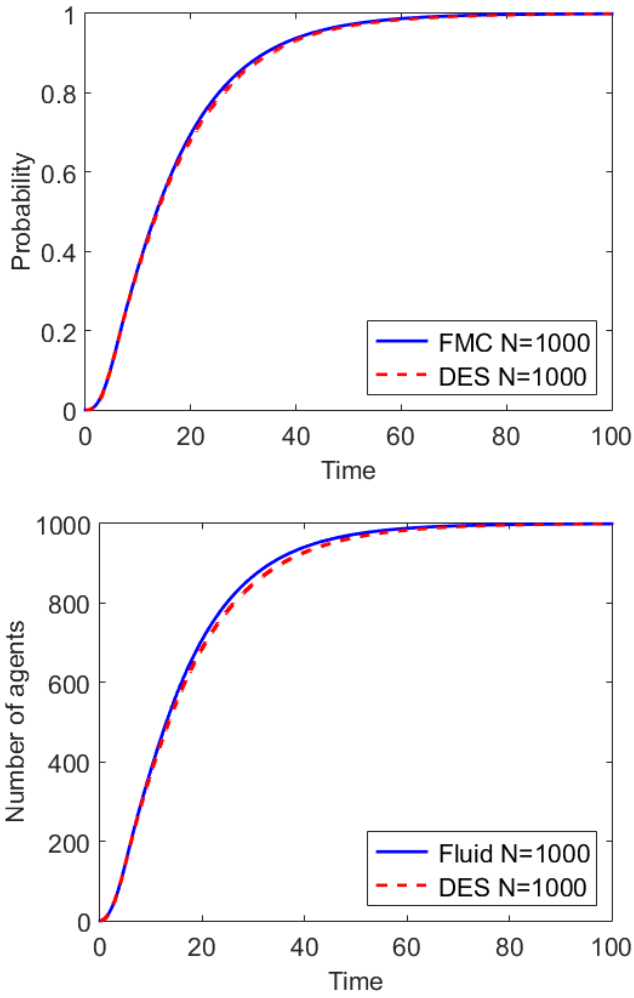

Figure 15: The satisfaction probability $P(T) = P_{F_{\mathbb{D}}}(T)$ obtained by the Fluid Model Checking (top) and the Fluid Approximation of the mean behaviour (bottom) in the case $N = 1000$. The results are compared with those obtained by the DES.

Moreover, notice also that the computational costs are comparable for the FMC and the Fluid Approximation, while the times of the DES are much higher in the second case. This is due to the fact that in the FMC, we have to keep track of the behaviour of just one single agent, while for the mean behaviour all the single entities in the population have to be tagged and thus the size of the model (and the computational cost of the DES) is much bigger. Figure 15 shows the results of the FMC and the Fluid Approximation in the case N=1000.

5.6 Discussion

We defined a fast and efficient Fluid Model Checking ((BH12b)) procedure that accurately estimates the probability that a single agent inside a large collective system satisfies a time-bounded property specified by a single-clock DTA. The method requires the integration of a system of DDEs for the transient probability of an IMRP, and the exactness of the estimation is guaranteed in the limit of an infinite population.

During the experimental analysis, we realised that, on certain models and properties, the DDEs (5.6) can be *stiff*, and their numerical integration in MATLAB is unstable (see also (BH15)). This is an issue that should be addressed by considering alternative integration methods (GH01), investigating also numerical techniques for MRP with time-dependent rates(ZFGH00).

The procedure of this chapter can be extended and improved in several directions: proving the convergence of the Fluid Approximation of Section 5.4.1, investigating higher-order estimates as in (BL13a; BL14), extend the FMC procedure to validate requirements specified in the logic CSL^{TA} (DHS09) and DTA properties endowed with multiple clocks (possibly considering the approximation techniques defined in (Fu13) and (BCH^{+}11)).

Chapter 6

Conclusions

6.1 Summary and Discussion

We have developed new, fast and reliable Stochastic Model Checking techniques for the analysis and verification of the behaviour of collective systems. In these procedures, we have exploited the efficiency and the accuracy of different types of Stochastic Approximations, including Fluid Approximation, Central Limit Approximation, System Size Expansion and Moment Closure combined with a distribution reconstruction based on the Maximum Entropy Principle. These estimation techniques turn out to be extremely useful in the validation of collective systems as they can be used to efficiently tackle the problem of the state space explosion. Indeed, the computational cost of the Stochastic Approximations that we have considered is independent of the population size, and the accuracy of the estimation actually increases with the number of agents comprised in the collective system.

In this work, we have merged and extended the very few model checking techniques that, at the beginning of this project, had already applied the Fluid Approximation and the Moment Closure to the validation of collective systems (BH12b; HSB12; HBC13), defining new interesting contributions in two directions: (1) we have considered numerous types of Stochastic Approximations, better capturing the probabilistic noise

that is intrinsic in the dynamics of mesoscopic collective systems; (2) we have widened the type of requirements on the behaviour of a collective system that can be validated applying Stochastic Approximations in model checking procedures. In particular, we have considered instances of local, global and local-to-global properties, and we have validated timed-critical properties in which the behaviour of the single individuals in the population is monitored by a single clock that can be reset.

Moreover, we have proven the theoretical results that guarantee the quality and correctness of our model checking procedures. In particular, we have proven the asymptotic convergence of the results and the correctness of the estimation in the limit of an infinite population.

6.2 Perspectives

The model checking procedures that we have developed in this work can be extended in several directions.

The main and most interesting line of research that easily arise form the presentation of this work goes in the direction of giving a solid and uniform theoretical structure to all the three promising model checking procedures of this thesis. This means that we should give a uniform classification to the stochastic approximations and the properties that can be exploited and validated in the procedures suggested in this work. Indeed, as we have seen, higher order approximations like the System Size Expansion (SSE) and the Methods of Moments (MM) provide promising results in the analysis of the Local-to-Global requirements. Further effort should be put in the investigation of the applicability of MM to the validation of global properties, and of both SSE and MM to the investigation of local requirements with clock resets. While in the first case, we expect to obtain quite good results (since the model checking procedure is based on the solution of a set of ordinary differential equations for the hitting time), in the validation of local requirements the applicability of SSE and especially of the MM to differential equations involving time delays, seems a greater and interesting challenge since, to the authors knowledge, there are no related works in the literature at the

moment of completion of this thesis. Moreover, the possibility of validation of properties with clock resets should be investigated for the model checking procedures that consider local-to-global and global requirements. Finally, also the coexistence of multiple clocks in the population model should be studied in all three cases. As it was shown in (Fu13) and (BCH$^+$11), the presence of more than one clock ends up in the definition of Partial Differential Equations (PPE) for the evolution of the state of the system. The analysis and solution of PPEs would be a challenging and most interesting novelty for all the model checking procedures illustrated in this thesis.

The unification of the theoretical background of the model checking procedures of this work will of course lead the way to a more accurate classification of the properties that can be validated exploiting stochastic approximation techniques. Indeed, as it was discussed along this thesis, the properties considered in this work are indeed very simple and quite far from the requirements of interest in the case of population models representing real life examples of epidemics or computer networks. But, indeed, the choice of the properties considered in this thesis was mainly focused on obtaining a easy and readable presentation of all the aspects of the model checking procedures, especially of the theory behind the synchronisation between the population model and the property, and of the differential equations to be defined and solved in the process. The investigation of more complex and expressive requirements is left for future (and less theoretical) works.

From a theoretical point of view, other very challenging lines of research include the definition of the speed of convergences of all stochastic procedures, in order for them to be theoretically comparable (in this respect, the work of (BHLM13) would be an interesting starting point). Also the investigation of the error bounds should be an interesting challenge (for primary results involving Fluid Approximation see (BH13)). Moreover, the investigation of the phase space of the differential equations could help the classification of the properties that can be considered (for example distinguishing between cyclic and acyclic automaton representations of the requirements). Finally, in the validation of the Global

Reachability Properties, our verification procedure relies on the hypothesis that the Fluid Approximation actually enters the target region and that the probability distribution of the hitting time is unimodal. This assumption limits the number of collective models and properties that we can validate, hence it would be (tough but) extremely interesting to find a way of lifting this hypothesis.

Finally, form an experimental point of view, the classification of the properties and the unification of the stochastic approximation techniques applicable in this framework of model checking techniques will also ease the design and implementation of a tool taking care of the automation of the steps required by the validation procedures. In this context, we need to extensively investigate new powerful integration methods to numerically solve the differential equations defined by the stochastic approximations. Indeed, as discussed in Chapter 5, the stiffness of the DDEs, that describes the dynamics of the IMRP in the validation of timed properties with clock reset, can deeply affect the computations. In this context, we found the numerical integration in MATLAB to be unstable and alternatives solvers must be consider if we want to extend the validation procedures to more complex and expressive models and properties.

References

[A+10] Rafail V Abramov et al. The multidimensional maximum entropy moment problem: A review of numerical methods. *Communications in Mathematical Sciences*, 2010. 26, 27

[AB00] H. Andersson and T. Britton. *Stochastic Epidemic Models and their Statistical Analysis*. Springer New York, 2000. 16, 64, 78

[ABG+15] Alexander Andreychenko, Luca Bortolussi, Ramon Grima, Philipp Thomas, and Verena Wolf. Distribution approximations for the chemical master equation: comparison of the method of moments and the system size expansion. *arXiv preprint arXiv:1509.09104*, 2015. 23, 24, 25, 28

[AD94] R. Alur and D. L. Dill. A Theory of Timed Automata. *Theor. Comput. Sci.*, 1994. 5, 31, 76, 78

[AKS13] Angelique Ale, Paul Kirk, and Michael PH Stumpf. A general moment expansion method for stochastic kinetic models. *The Journal of chemical physics*, 2013. 25

[AMW15a] Alexander Andreychenko, Linar Mikeev, and Verena Wolf. Model reconstruction for moment-based stochastic chemical kinetics. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 2015. 3, 24, 25, 26, 27

[AMW15b] Alexander Andreychenko, Linar Mikeev, and Verena Wolf. Reconstruction of multimodal distributions for hybrid moment-based chemical kinetics. *Journal of Coupled Systems and Multiscale Dynamics*, 2015. 26, 34

[BBHK00] C. Baier, H. Boudewijn, H. Hermanns, and J.P. Katoen. Model checking continuous-time Markov chains by transient analysis. In *Computer Aided Verification*. Springer, 2000. 2, 5, 44

[BCFH09] R. Bakhshi, L. Cloth, W. Fokkink, and B. Haverkort. Mean-field analysis for the evaluation of gossip protocols. In *Quantitative Evaluation of Systems, 2009. QEST'09. Sixth International Conference on the*, pages 247–256. IEEE, 2009. 18

[BCGH05] A. Benoit, M. Cole, S. Gilmore, and J. Hillston. Enhancing the effective utilisation of grid clusters by exploiting on-line performability analysis. In *Proceedings of the Fifth IEEE International Symposium on Cluster Computing and the Grid - Volume 01*, CCGRID '05, pages 317–324, Washington, DC, USA, 2005. IEEE Computer Society. 18

[BCH+07] C. Baier, L. Cloth, B.R. Haverkort, M. Kuntz, and M. Siegle. Model checking markov chains with actions and state labels. *IEEE Trans. Software Eng.*, 33(4):209–224, 2007. 34, 35

[BCH+11] B. Barbot, T. Chen, T. Han, J.-P. Katoen, and A. Mereacre. Efficient CTMC model checking of linear real-time objectives. In *Tools and Algorithms for the Construction and Analysis of Systems*. 2011. 76, 91, 95, 98

[BGH08] J. T. Bradley, S. T. Gilmore, and J. Hillston. Analysing distributed internet worm attacks using continuous state-space approximation of process algebra models. *Journal of Computer and System Sciences*, 74(6):1013–1032, 2008. 18

[BGH12] Luca Bortolussi, Vashti Galpin, and Jane Hillston. Hybrid performance modelling of opportunistic networks. In Herbert Wiklicky and Mieke Massink, editors, *QAPL 2012*. Open Publishing Association, 2012. 75

[BH12a] L. Bortolussi and J. Hillston. Fluid Approximation of CTMC with Deterministic Delays. In *Proceedings of QEST*, 2012. 76

[BH12b] L. Bortolussi and J. Hillston. Fluid model checking. In *CONCUR 2012–Concurrency Theory*, pages 333–347. Springer, 2012. 3, 4, 5, 14, 33, 34, 35, 39, 40, 95, 96

[BH13] L. Bortolussi and R. Hayden. Bounds on the deviation of discrete-time Markov chains from their mean-field model. *Performance Evaluation*, 2013. 98

[BH15] L. Bortolussi and J. Hillston. Model Checking Single Agent Behaviours by Fluid Approximation. *Inform. Comput.*, 2015. 20, 76, 88, 95

[BHLM13] L. Bortolussi, J. Hillston, D. Latella, and M. Massink. Continuous approximation of collective systems behaviour: a tutorial. *Perf. Eval.*, 2013. 2, 18, 35, 54, 59, 98

[BK08] C. Baier and J.P. Katoen. *Principles of Model Checking*. MIT press, 2008. 1, 33, 61, 76

[BL13a] L. Bortolussi and R. Lanciani. Model Checking Markov Population Models by Central Limit Approximation. In *Proceedings of QEST*, 2013. xiii, 4, 34, 76, 77, 95

[BL13b] L. Bortolussi and R. Lanciani. Model Checking Markov Population Models by Central Limit Approximation. In *10th International Conference on Quantitative Evaluation of SysTems, QEST 2013*, Buenos Aires, Argentina, 2013. Springer Verlag, Springer Verlag. 13

[BL14] L. Bortolussi and R. Lanciani. Stochastic Approximation of Global Reachability Probabilities of Markov Population Models. In *Proceedings of EPEW*, 2014. xiii, 5, 62, 95

[BL15] Luca Bortolussi and Roberta Lanciani. Fluid Model Checking of Timed Properties. In *International Conference on Formal Modeling and Analysis of Timed Systems*, pages 172–188. Springer, 2015. xiii, 77

[BLB08] M. Benaim and J.-Y. Le Boudec. A Class of Mean Field Interaction Models for Computer and Communication Systems. *Perform. Evaluation*, 2008. 19

[BLN17] L. Bortolussi, R. Lanciani, and L. Nenzi. Model Checking Markov Population Models by Stochastic Approximation, 2017. xiii, 4, 57, 58

[BMS16] Luca Bortolussi, Dimitrios Milios, and Guido Sanguinetti. Smoothed model checking for uncertain continuous-time Markov chains. *Information and Computation*, 2016. 2

[BP09] L. Bortolussi and A. Policriti. Dynamical systems and stochastic programming: To ordinary differential equations and back. In *Transactions on Computational Systems Biology XI*, pages 216–267. Springer, 2009. 18

[Buj12] L. M. Bujorianu. *Stochastic reachability analysis of hybrid systems*. Springer, 2012. 30, 65

[BW03] M. Benaïm and J. W. Weibull. Deterministic approximation of stochastic evolution in games. *Econometrica*, 71(3):873–903, 2003. 18

[Car08] L. Cardelli. On process rate semantics. *Theoretical Computer Science*, 391(3):190–215, 2008. 18

[CDKM13] T. Chen, M. Diciolla, M. Kwiatkowska, and A. Mereacre. Verification of linear duration properties over CTMCs. *Proceedings of TOCL*, 2013. 76

[CHKM11a] T. Chen, T. Han, J.-P. Katoen, and A. Mereacre. Model Checking of Continuous-Time Markov Chains Against Timed Automata Specifications. *Logical Methods in Computer Science*, 2011. 5, 31, 39, 76, 78

[CHKM11b] T. Chen, T. Han, J.-P. Katoen, and A. Mereacre. Observing continuous-time MDPs by 1-clock timed automata. In *Reachability Problems*. Springer, 2011. 76

[Cin13] E. Cinlar. *Introduction to Stochastic Processes*. Courier Corporation, 2013. 76, 81, 82, 86

[DHS09] S. Donatelli, S. Haddad, and J. Sproston. Model checking timed and stochastic properties with CSL^{TA}. *IEEE Trans. Software Eng.*, 35(2):224–240, 2009. 34, 35, 38, 39, 76, 95

[DN08] R. Darling and J. Norris. Differential Equation Approximations for Markov Chains. *Probability Surveys*, 2008. 19, 20, 76

[Dur10] R. Durrett. *Essentials of Stochastic Processes*. Springer Texts in Statistics. Springer, 2010. 7, 8, 9

[EK05] S. N. Ethier and T. G. Kurtz. *Markov Processes: Characterization and Convergence*. Wiley, 2005. 3, 5, 8, 9, 17, 19, 21, 22, 33, 40, 62, 67, 69

[Fu13] H. Fu. Approximating acceptance probabilities of ctmc-paths on multi-clock deterministic timed automata. In *Proceedings of HSCC*, 2013. 76, 91, 95, 98

[GB10] N. Gast and G. Bruno. A Mean Field Model of Work Stealing in Large-Scale Systems. *ACM SIGMETRICS Performance Evaluation Review*, 2010. 19

[GH01] N. Guglielmi and E. Hairer. Implementing Radau IIA Methods for Stiff Delay Differential Equations. *Computing*, 2001. 95

[Gri10] R. Grima. An effective rate equation approach to reaction kinetics in small volumes: Theory and application to biochemical reactions in nonequilibrium steady-state conditions. *The Journal of Chemical Physics*, 2010. 2, 23, 34

[Hay12] R. A. Hayden. Mean Field for Performance Models with Deterministically-Timed Transitions. In *Proceedings of QEST*, 2012. 62, 76, 77, 92

[HBC13] R. A. Hayden, J. T. Bradley, and A. Clark. Performance specification and evaluation with unified stochastic probes and fluid analysis. *IEEE Trans. Software Eng.*, 39(1):97–118, 2013. 3, 33, 34, 76, 96

[HSB12] R.A. Hayden, A. Stefanek, and J.T. Bradley. Fluid computation of passage-time distributions in large Markov models. *Theor. Comput. Sci.*, 413(1):106–141, 2012. 3, 33, 34, 35, 96

[JCL+09] Sumit Kumar Jha, Edmund M Clarke, Christopher J Langmead, Axel Legay, Andre Platzer, and Paolo Zuliani. Statistical model checking for complex stochastic models in systems biology. 2009. 2

[Kal06] Olav Kallenberg. *Foundations of modern probability*. Springer Science & Business Media, 2006. 25

[KFR+16] Atefeh Kazeroonian, Fabian Fröhlich, Andreas Raue, Fabian J Theis, and Jan Hasenauer. Cerena: Chemical reaction network analyzera toolbox for the simulation and analysis of stochastic chemical kinetics. *PloS one*, 2016. 26

[Kin94] E. Kindler. Safety and liveness properties: A survey. *Bulletin of the European Association for Theoretical Computer Science*, 53:268–272, 1994. 30

[KNP11] M. Kwiatkowska, G. Norman, and D. Parker. PRISM 4.0: Verification of Probabilistic Real-Time Systems. In *Computer Aided Verification*. Springer, 2011. 2

[KRdH13] A. Kolesnichenko, A. Remke, P.T. de Boer, and B.R. Haverkort. A logic for model-checking of mean-field models. In *Proc. of DSN*, 2013. 34, 35, 39

[Lam77] L. Lamport. Proving the correctness of multiprocess programs. *Software Engineering, IEEE Transactions on*, (2):125–143, 1977. 30

[LMW11] Maksim Lapin, Linar Mikeev, and Verena Wolf. Shave: stochastic hybrid analysis of markov population models. In *Proceedings of the 14th international conference on Hybrid systems: computation and control*. ACM, 2011. 26

[MLBH10] M. Massink, D. Latella, A. Bracciali, and M. D. Harrison. A scalable fluid flow process algebraic approach to emergency egress analysis. In *Software Engineering and Formal Methods (SEFM), 2010 8th IEEE International Conference on*, pages 169–180. IEEE, 2010. 18

[MLBH11] M. Massink, D. Latella, A. Bracciali, and J. Hillston. Modelling non-linear crowd dynamics in bio-pepa. In *Fundamental Approaches to Software Engineering*. Springer Berlin Heidelberg, 2011. 18

[Nor97] J. R. Norris. *Markov Chains*. Cambridge University Press, 1997. 8, 9, 11, 23

[RW06] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006. 55

[SR04] M. Silva and M. Recalde. On fluidification of Petri Nets: from discrete to hybrid and continuous models. *Annual Reviews in Control*, 28(2):253 – 266, 2004. 18

[SSG16] David Schnoerr, Guido Sanguinetti, and Ramon Grima. Approximation and inference methods for stochastic biochemical kinetics-a tutorial review. *Journal of Physics A: Mathematical and Theoretical*, 2016. 3, 24, 25

[TG11] M. Tribastone and S. Gilmore. Scaling performance analysis using fluid-flow approximation. In *Rigorous software engineering for service-oriented systems*, pages 486–505. Springer, 2011. 18

[Van92] N. G. Van Kampen. *Stochastic Processes in Physics and Chemistry*. Elsevier, 1992. 2, 3, 21, 23, 34, 40

[ZFGH00] A. Zimmermann, J. Freiheit, R. German, and G. Hommel. Petri Net Modelling and Performability Evaluation with TimeNET 3.0. In *Computer Performance Evaluation*. 2000. 95