

Fluid Model Checking of Timed Properties^{*}

Luca Bortolussi¹ and Roberta Lanciani²

¹ Modelling and Simulation Group, Saarland University, Germany
DMG, University of Trieste, Italy

CNR/ISTI, Pisa, Italy
luca@dmi.units.it

² IMT Lucca, Italy

roberta.lanciani@imtlucca.it

Abstract. We address the problem of verifying timed properties of Markovian models of large populations of interacting agents, modelled as finite state automata. In particular, we focus on time-bounded properties of (random) individual agents specified by Deterministic Timed Automata (DTA) endowed with a single clock. Exploiting ideas from fluid approximation, we estimate the satisfaction probability of the DTA properties by reducing it to the computation of the transient probability of a subclass of Time-Inhomogeneous Markov Renewal Processes with exponentially and deterministically-timed transitions, and a small state space. For this subclass of models, we show how to derive a set of Delay Differential Equations (DDE), whose numerical solution provides a fast and accurate estimate of the satisfaction probability. In the paper, we also prove the asymptotic convergence of the approach, and exemplify the method on a simple epidemic spreading model. Finally, we also show how to construct a system of DDEs to efficiently approximate the average number of agents that satisfy the DTA specification.

Keywords: Stochastic Model Checking, Fluid Model Checking, Deterministic Timed Automata, Time-Inhomogeneous Markov Renewal Processes, Fluid Approximation, Delay Differential Equations.

1 Introduction

One of the major technological challenges in computer science and engineering is the design and analysis of large-scale distributed systems, where many autonomous components interact in an open environment. Examples include the public and shared transportation in smart cities, the power distribution in smart grids, and the robust communication protocols of online multimedia services. In this context, the mathematical and computational modelling plays a crucial role in the management of such *Collective Adaptive Systems* (CAS), due to the

^{*} This research has been partially funded by the EU-FET project QUANTICOL (nr. 600708) and by the German Research Council (DFG) as part of the Cluster of Excellence on Multimodal Computing and Interaction at Saarland University.

need of understanding and control of their emergent behaviours in open working conditions. The intrinsic uncertainty of CAS can be properly captured by *stochastic models*, but the large number of interacting entities always results in a severe *state space explosion*, introducing exceptional computational challenges. In particular, the scalability of the models and of their analysis techniques is a major issue in the development of *stochastic model checking* procedures for the verification of formal properties. In this context, up to now, the numerical approaches [24] are deeply hampered by the state space explosion of the large stochastic models, and the statistical methods based on simulation require a large computational effort.

A recent line of work tries to address the issue of scalability by exploiting stochastic approximation techniques [10, 11], like the *Fluid Approximation* [8, 9, 18]. In this method, a stochastic discrete model is replaced by a simpler continuous one, whose dynamics is described by a set of differential equations. In [8], the authors exploit this limit construction to verify properties that assess the behaviour of a single individual in a collective system, and define a procedure called the *Fluid Model Checking* (FMC) [7, 25]. This technique is based on the *Fast Simulation Theorem* [16], which ensures that in a large population, a single entity is influenced only by the mean behaviour of the rest of the agents.

In this work, we extend [8] to more complex *time-bounded properties* specified by *Deterministic Timed Automata* endowed with a single clock [1, 3, 17]. As in [8, 10, 13, 23], we combine the agent and the DTA specification with a product construction, obtaining a *Time-Inhomogeneous Markov Renewal Process* [15]. We then exploit results [6, 22], defining the Fluid Approximation of this type of models as the solution of a system of *Delay Differential Equations* (DDE) [16]. Other works dealing with the verification of DTA properties are [4, 12, 14, 19].

Main Result. 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 [8], the technique is based on 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 paper, 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 [10, 22], 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 .

Paper structure. In Sec. 2, we introduce the modelling language, the Fluid Approximation, the Fast Simulation Theorem, and the DTA specification for the timed properties. In Sec. 3, we present our FMC procedure, defining the DDEs for the probability that the single agent satisfies the timed property. In Sec. 3.1, we adapt our verification technique to compute the mean number of agents that

meet the DTA requirement. In Sec. 4, we discuss the quality of the approximation on the epidemic example. Finally, in Sec. 5, we draw the final conclusions. The proofs of the theoretical results are omitted and can be found in the extended version available at <http://arxiv.org/abs/1506.05909>.

2 Background and Modelling Language

Agent Classes and Population Models. A collective system is comprised of a large number of interacting *agents*. To describe its dynamics, we define a *population model* [10,11] in which the agents are divided into classes, called *agent classes*, according to their behaviour.

Definition 1 (Agent Class). *An agent class \mathcal{A} is a pair (S, E) in which $S = \{1, \dots, m\}$ is the (finite) state space and $E = \{\epsilon_1, \dots, \epsilon_n\} \subseteq S \times \mathcal{L} \times S$ is the (finite) set of local transitions of the form $\epsilon_i = s_i \xrightarrow{\alpha_i} s'_i$, where $s_i, s'_i \in S$ are the initial and arrival states, and $\alpha_i \in \mathcal{L}$ is the unique label of ϵ_i , i.e. $\alpha_i \neq \alpha_j$ for $i \neq j$ ³.*

Intuitively, an agent in class $\mathcal{A} = (S, E)$ is a finite state automaton that can change state by performing the actions in E . Then, assuming that agents in the same state are indistinguishable, to define the population model, we rely on the *counting abstraction*, counting how many agents are in each state at time t . Hence, for each agent class, we define the *collective* or *counting variables* $X_1^{(N)}(t), \dots, X_m^{(N)}(t)$ given by $X_j^{(N)}(t) = \sum_k \mathbb{1}_{\{Y_k^{(N)}(t)=j\}}$, where $Y_k^{(N)}(t) \in \{1, \dots, m\}$ is the random variable denoting the state of agent k at time t , and $N = \sum_{\mathcal{A}} \sum_j X_j^{(N)}$ is the *population size*, that we assume constant in time (cf. also [8]). Then, given $n = \sum_{\mathcal{A}} |S|$, the state of the population model is given by the vector $\mathbf{X}^{(N)}(t) \in (\mathbb{R}_{\geq 0})^n$ that enlists the counting variables of the agent classes.

Definition 2 (Population Model). *A population model $\mathcal{X}^{(N)}$ is a tuple $\mathcal{X}^{(N)} = (\mathbb{A}, \mathcal{T}^{(N)}, \mathbf{x}_0^{(N)})$, where $\mathbb{A} = \{\mathcal{A}_1, \dots, \mathcal{A}_\nu\}$ is the set of agent classes, as in Definition 1; $\mathbf{x}_0^{(N)} = \mathbf{X}^{(N)}(0)$ is the initial state; and $\mathcal{T}^{(N)} = \{\tau_1, \dots, \tau_\ell\}$ is the set of global transitions of the form $\tau_i = (\mathbb{S}_i, f_i^{(N)}, \mathbf{v}_i^{(N)})$, where:*

- $\mathbb{S}_i = \{s_1 \xrightarrow{\alpha_1} s'_1, \dots, s_p \xrightarrow{\alpha_p} s'_p\}$ is the (finite) multi-set of local transitions synchronized by τ_i ;
- $f_i^{(N)} : (\mathbb{R}_{\geq 0})^n \rightarrow \mathbb{R}_{\geq 0}$ is the (Lipschitz continuous) global rate function;
- $\mathbf{v}_i = \sum_{\alpha_j \in \mathbb{S}_i} |\{\alpha_j\}| (\mathbb{1}_{s_j} - \mathbb{1}_{s'_j})$ is the update vector, where $|\{\alpha_j\}|$ is the multiplicity of α_i in \mathbb{S}_i , and $\mathbb{1}_{s_j}$ is the vector equal to 1 on s_j and 0 elsewhere.

³ The restriction on the uniqueness of the labels can be dropped (as in [10]) at the price of heavier notation and combinatorics in the definitions of the rest of the paper.




Fig. 1. The agent class \mathcal{A} (left) and property \mathbb{D} (right) of the running example.

When a global transition $\tau_i = (\mathbb{S}_i, f_i^{(N)}, \mathbf{v}_i)$ is taken, the transitions in \mathbb{S}_i fire at the local level, meaning that, for each $s \xrightarrow{\alpha} s'$ in \mathbb{S}_i , an agent moves from s to s' . Hence, the update vector \mathbf{v}_i encodes the net change in the state $\mathbf{X}^{(N)}(t)$ of $\mathcal{X}^{(N)}$ due to transition τ_i . Moreover, for the model to be meaningful, whenever at time t it is not possible to execute τ_i , because there are not enough agents available, i.e. $(\mathbf{X}^{(N)}(t) - \mathbf{v}_i)_j < 0$ for some $j \in \{1, \dots, n\}$ with $n = |\mathbf{X}^{(N)}(t)|$, we require the rate function to be zero, i.e. $f_i^{(N)}(\mathbf{X}^{(N)}(t)) = 0$.

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 \mathcal{A} , depicted in Fig. 1, 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 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 $\tau_i = (\{S \xrightarrow{inf} I, I \xrightarrow{pass} I\}, f_i^{(N)})$. 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* [2]: $f_r^{(N)}(t) = k_r X_I^{(N)}(t)$ and $f_i^{(N)}(t) = \frac{1}{N} k_i X_S^{(N)}(t) X_I^{(N)}(t)$, where $k_r, k_i \in \mathbb{R}_{\geq 0}$.

Fluid Approximation. The *Fluid Approximation* [8, 9, 18] of a population model $\mathcal{X}^{(N)} = (\mathbb{A}, \mathcal{T}^{(N)}, \mathbf{x}_0^{(N)})$ is an estimate of the *mean* behaviour of its agents. To compute this approximation, we first *normalise* $\mathcal{X}^{(N)}$ by dividing the state vector $\mathbf{X}^{(N)}(t)$ and the initial state $\mathbf{x}_0^{(N)}$ by the population size N , i.e. we define $\widehat{\mathbf{X}}^{(N)}(t) = \mathbf{X}^{(N)}(t)/N$ and $\widehat{\mathbf{x}}_0^{(N)} = \mathbf{x}_0^{(N)}/N$. Then, for all transition $\tau_i \in \mathcal{T}^{(N)}$, we let $\widehat{f}_i^{(N)}(\widehat{\mathbf{X}})$ be the rate function, where we substitute the counting variables of $\mathbf{X}^{(N)}(t)$ with the new normalised counting variables of $\widehat{\mathbf{X}}(t)$. Moreover, we assume that for each $\widehat{f}_i^{(N)}(\widehat{\mathbf{X}})$, there exist a *Lipschitz function* $f_i(\widehat{\mathbf{X}})$ such that $\widehat{f}_i^{(N)}(\widehat{\mathbf{X}})/N \xrightarrow{N \rightarrow +\infty} f_i(\widehat{\mathbf{X}})$ uniformly. Finally, in terms of $f_i(\widehat{\mathbf{X}})$, we define the *drift* $\mathbf{F}(\widehat{\mathbf{X}})$ given by $\mathbf{F}(\widehat{\mathbf{X}}) = \sum_{\tau_i} \mathbf{v}_i f_i(\widehat{\mathbf{X}})$, whose components represent the instantaneous net flux of agents in each state of the model. Then, given a *time*

horizon $T < +\infty$, the *Fluid Approximation* $\Phi(t)$ of $\mathcal{X}^{(N)}$ is the unique⁴ solution of the system of *Ordinary Differential Equations* (ODEs) given by

$$\frac{d\Phi}{dt}(t) = \mathbf{F}(\Phi(t)), \quad \text{for } 0 \leq t \leq T,$$

with $\Phi(0) = \mathbf{x}_0$. The accuracy of the approximation *improves* the larger is the ensemble of agents that we consider, i.e. *the larger is N* , and is exact in the limit of an infinite population. Indeed, the following theorem holds true [18].

Theorem 1 (Fluid Approximation). *For any $T < +\infty$ and $\epsilon > 0$,*

$$\text{Prob} \left\{ \sup_{0 \leq t \leq T} \|\widehat{\mathbf{X}}^{(N)}(t) - \Phi(t)\| > \epsilon \right\} \xrightarrow{N \rightarrow +\infty} 0.$$

Fast Simulation. In this paper, we are 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* [5, 16, 20] 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) [8] 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 $\Phi(t)$ be the Fluid Approximation of $\mathcal{X}^{(N)}$. Define the *generator matrix* $\mathbf{Q}^{(N)}(\mathbf{x}) = (q_{ij}^{(N)}(\mathbf{x}))$ of $Y^{(N)}(t)$ as a function of the normalised counting variables, i.e. $\forall q_{ij}^{(N)}(\mathbf{x})$,

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

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

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

where $|\{i \rightarrow j \in \mathbb{S}_\tau\}|$ is the multiplicity of $i \rightarrow j$ in the transition set \mathbb{S}_τ of τ , i.e. the number of agents that take such transition according to τ . Furthermore,

⁴ Existence and uniqueness of $\Phi(t)$ are guaranteed by the Lipschitzianity of the $f_i(\widehat{\mathbf{X}})$.

as customary, let $q_{ii}^{(N)}(\mathbf{x}) = -\sum_{j \neq i} q_{ij}^{(N)}(\mathbf{x})$. Then, since $\widehat{f}_i^{(N)}(\widehat{\mathbf{X}})/N \xrightarrow{N \rightarrow +\infty} f_i(\widehat{\mathbf{X}})$, we have that $\mathbf{Q}^{(N)}(\mathbf{x}) \rightarrow \mathbf{Q}(\mathbf{x})$, where $\mathbf{Q}(\mathbf{x})$ 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 \mathcal{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 $\mathbf{Q}(\Phi(t))$, i.e. the generator matrix $\mathbf{Q}(t)$, where the Lipschitz limits $f_i(t)$ are computed over the components of $\Phi(t)$.

Then, the following theorem can be proved [16].

Theorem 2 (Fast Simulation). *For any time horizon $T < +\infty$ and $\epsilon > 0$,*

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

Example. For the running example, if we consider a population of 1000 agents, i.e. $N = 1000$, and an initial state $\mathbf{x}_0^{(N)} = (900, 100)$, then the Fluid Approximation $\Phi(t)$ of the population model is the unique solution of the following ODEs:

$$\begin{cases} \frac{d\Phi_S}{dt}(t) = -k_i\Phi_I(t)\Phi_S(t) + k_r\Phi_I(t); \\ \frac{d\Phi_I}{dt}(t) = +k_i\Phi_I(t)\Phi_S(t) - k_r\Phi_I(t); \end{cases} \quad \text{with} \quad \begin{cases} \Phi_S(0) = 0.9; \\ \Phi_I(0) = 0.1. \end{cases} \quad (1)$$

The generator $\mathbf{Q}(\Phi(t))$ of the ICTMC $Z(t)$ for the single agent, instead, is:

$$q_{S,S}(t) = -q_{S,I}(t); \quad q_{S,I}(t) = k_i\Phi_I(t); \quad q_{I,S}(t) = k_r; \quad q_{I,I}(t) = -q_{I,S}(t). \quad (2)$$

2.1 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) [1, 13], 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, 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 3 (Timed Properties). *A timed property for a single agent in agent class \mathcal{A} is specified as a single-clock DTA of the form $\mathbb{D} = \mathbb{D}(T) = (T, \mathcal{L}, c, \mathcal{CC}, Q, q_0, F, \rightarrow)$, where $T < +\infty$ is the finite time horizon; \mathcal{L} is the label set of \mathcal{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 $\rightarrow \subseteq Q \times \mathcal{L} \times \mathcal{CC} \times \{\emptyset, \{c\}\} \times Q$ is the edge relation. Moreover, \mathbb{D} has to satisfy:*

- (determinism) for each initial state $q \in Q$, label $\alpha \in \mathcal{L}$, 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 $\eta(c) \models_{\mathcal{CC}} c_{\bowtie}$ ⁵;
- (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 \mathcal{A} sampled from the stochastic processes $Z^{(N)}(t)$ and $Z(t)$ defined for the Fast Simulation in Sec. 2. The labels of the transitions of \mathcal{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 \mathcal{A} sampled from $Z^{(N)}(t)$ (resp. $Z(t)$), 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 $\Sigma_{\mathcal{A}, \mathbb{D}, T}$ the set of time-bounded paths of \mathcal{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 such requirement, we use the DTA $\mathbb{D} = \mathbb{D}(T)$ represented in Fig. 1. If we record the actions of the single agent on \mathbb{D} , i.e. we synchronise \mathcal{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} .

3 Fluid Model Checking of Timed Properties

Consider a single agent of class $\mathcal{A} = (S, E)$ in a population model $\mathcal{X}^{(N)} = (\mathcal{A}, \mathcal{T}^{(N)}, \mathbf{x}_0^{(N)})$, and a timed property $\mathbb{D} = \mathbb{D}(T) = (T, \mathcal{L}, \Gamma_S, \mathcal{CC}, Q, q_0, F, \rightarrow)$. Let $\Sigma_{\mathcal{A}, \mathbb{D}, T}$ be the set of time-bounded paths of \mathcal{A} accepted by \mathbb{D} . Moreover, let $Z^{(N)}(t)$ and $Z(t)$ be the two stochastic processes defined for the Fast Simulation in Sec. 2. The following result holds true.

Proposition 1. *The set $\Sigma_{\mathcal{A}, \mathbb{D}, T}$ is measurable for the probability measures $Prob_{Z^{(N)}}$ and $Prob_Z$ defined over the paths of $Z^{(N)}(t)$ and $Z(t)$, respectively. \square*

Let $P^{(N)}(T) = Prob_{Z^{(N)}}\{\Sigma_{\mathcal{A}, \mathbb{D}, T}\}$ and $P(T) = Prob_Z\{\Sigma_{\mathcal{A}, \mathbb{D}, T}\}$. In this paper, we are interested in the *satisfaction probability* $P^{(N)}(T)$, i.e. the probability

⁵ 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} .

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

Theorem 3. *For any $T < +\infty$, $\lim_{N \rightarrow \infty} P^{(N)}(T) = P(T)$. \square*

Moreover, to compute $P(T)$, we consider a suitable product construction $\mathcal{A}_{\mathbb{D}} = \mathcal{A} \otimes \mathbb{D}$, whose state is described by a *Time-Inhomogeneous Markov Renewal Process* (IMRP) [15] that we denote by $Z_{\mathcal{A}_{\mathbb{D}}}(t)$. In the rest of this section, we define $\mathcal{A}_{\mathbb{D}}$ and $Z_{\mathcal{A}_{\mathbb{D}}}(t)$, and we show how to compute the *satisfaction probability* $P(T)$ in terms of the *transient probability* $\mathbf{P}(T)$ of $Z_{\mathcal{A}_{\mathbb{D}}}(t)$.

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

A *Markov Renewal Process* (MRP) [15] is a jump-process, where the sojourn times in the states can have a general probability distribution. In particular, in the MRP $Z_{\mathcal{A}_{\mathbb{D}}}(t)$, we will allow both *exponentially* and *deterministically-timed* transitions, and in the following, we will refer to them as the *Markovian* and *deterministic transitions*, respectively. Since the transition rates of $Z_{\mathcal{A}_{\mathbb{D}}}(t)$ will be time-dependent, $Z_{\mathcal{A}_{\mathbb{D}}}(t)$ will be a *Time-Inhomogeneous* MRP.

To define the product $\mathcal{A}_{\mathbb{D}} = (\mathcal{A}, S_{\mathbb{D}}, \{\mathcal{M}, \mathcal{E}\}, s_{0, \mathbb{D}}, F_{\mathbb{D}})$, let $\delta_1 < \dots < \delta_k$ 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 $\mathcal{A}_{\mathbb{D}}$ is given by $\{1, \dots, k+1\} \times S \times Q$. 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 $\mathcal{A}_{\mathbb{D}}$ will be defined in such a way that the agent is in $S_{\mathbb{D}_i}$ if and only if 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 $\mathcal{A}_{\mathbb{D}}$ is the smallest relation such that

$$\forall i \in 1, \dots, k+1, \quad \frac{s \xrightarrow{\alpha} s' \in E \wedge q \xrightarrow{\alpha, c_{\bowtie}, \emptyset} q' \in \rightarrow \wedge [\delta_{i-1}, \delta_i] \models c_{\bowtie}}{(i, s, q) \xrightarrow{\alpha} (i, s', q') \in \mathcal{M}}, \quad (3)$$

$$\forall i \in 1, \dots, k+1, \quad \frac{s \xrightarrow{\alpha} s' \in E \wedge q \xrightarrow{\alpha, c_{\bowtie}, \{c\}} q' \in \rightarrow \wedge [\delta_{i-1}, \delta_i] \models c_{\bowtie}}{(i, s, q) \xrightarrow{\alpha} (1, s', q') \in \mathcal{M}}. \quad (4)$$

Intuitively, rule (3) synchronises the local transitions $s \xrightarrow{\alpha} s' \in E$ of the agent class $\mathcal{A} = (S, E)$ with the transition $q \xrightarrow{\alpha, c_{\bowtie}, \emptyset} q' \in \rightarrow$ that has the same label in \mathbb{D} , obtaining a local transition $(i, s, q) \xrightarrow{\alpha} (i, s', q') \in \mathcal{M}$ in $\mathcal{A}_{\mathbb{D}}$ for each time region i whose time interval $[\delta_{i-1}, \delta_i] \subseteq [0, T]$ satisfies the clock constraint c_{\bowtie} , meaning that $\forall t \in [\delta_{i-1}, \delta_i], t \models c_{\bowtie}$. Rule (4), instead, defines the *reset transitions* $(i, s, q) \xrightarrow{\alpha} (1, s', q') \in \mathcal{M}$ that reset the personal clock c either within

the 1st Time Region (when $i = 1$), or by bringing the agent *back to the 1st 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 $\mathcal{A}_{\mathbb{D}}$, 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*, and $p_e : \mathcal{A}_e \times S_{\mathbb{D}} \rightarrow [0, 1]$ 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 $p_e((i, s, q), \cdot) : S_{\mathbb{D}} \rightarrow [0, 1]$, where $(i, s, q) \in \mathcal{A}_e$ is the state in which the agent is when the countdown hits zero. Moreover, e_i is deactivated also when the agent takes a reset transition. In $\mathcal{A}_{\mathbb{D}}$, we define:

- one clock event $e_i \in \mathcal{E}$ for each time region $S_{\mathbb{D}_i}$, $i = 2, \dots, k$;
- $\ell + 1$ clock events $e_1^0, e_1^1, \dots, e_1^\ell \in \mathcal{E}$ for the 1st Time Region, where ℓ is the number of reset events $(1, s, q) \xrightarrow{\alpha} (1, s', q') \in \mathcal{R}$ defined by (4) with $i = 1$.

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

$$p_i((i, s, q), (i', s', q')) = \begin{cases} 1 & \text{if } i' = i + 1, s' = s, q' = q, \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

As it is defined, 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 , $j = 1, \dots, \ell$, 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)), 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 need to define a clock event e_1^j , whose activation set \mathcal{A}_1^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), 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 $\mathcal{A}_{\mathbb{D}}$ 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), also the clock events of the 1st Time Region move the agent from a state to the equivalent one in the next time region (the 2nd), when the countdown from $\Delta_1^j = \delta_1$ elapses. In the following, we denote by $(i, s, q) \xrightarrow{e} (i + 1, s, q)$ 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 $\nu_{e, s, q} = \mathbb{1}_{(i+1, s, q)} - \mathbb{1}_{(i, s, q)}$ its update vector. The last component of $\mathcal{A}_{\mathbb{D}}$ 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. Fig. 2 represents the product $\mathcal{A}_{\mathbb{D}}$ between the agent class \mathcal{A} and the property \mathbb{D} of the running example (Fig. 1). The state $(1, I, q1)$ that cannot be

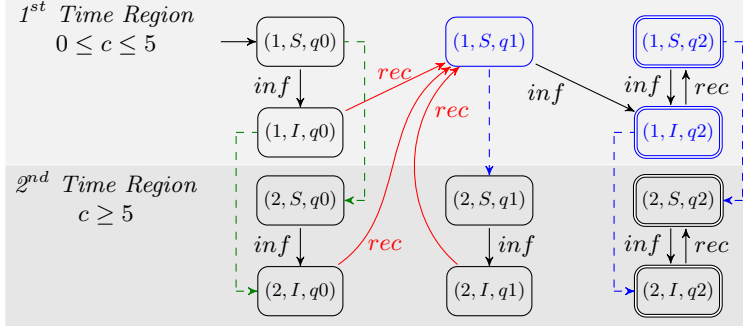


Fig. 2. The agent class $\mathcal{A}_{\mathbb{D}}$ associated with the DTA \mathbb{D} of the running example.

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 1st 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 1st Time Region). Intuitively, the agent can be found in one of the states belonging to the 1st Time Region whenever its personal clock c is between 0 and 5, i.e. less than 5 time units have passed since $t = 0$ or since a recovery *rec*. In a similar way, the agent is in the 2nd Time Region when the valuation of c is above 5. Moreover, when the duration of the clock events elapses (i.e. the countdown from 5 hits 0), the agent is moved from the 1st Time Region to the 2nd 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 $\mathcal{A}_{\mathbb{D}}$ at time T .

The IMRP $Z_{\mathcal{A}_{\mathbb{D}}}(t)$ and the Satisfaction Probability $P(T)$. Now we show how to formally define the IMRP $Z_{\mathcal{A}_{\mathbb{D}}}(t)$ that describes the state of the product $\mathcal{A}_{\mathbb{D}}$ in the mean field regime. In particular, when deriving the *Delay Differential Equations* (DDE) [15] for the *transient probability* $\mathbf{P}(t)$ of $Z_{\mathcal{A}_{\mathbb{D}}}(t)$, in terms of which we compute the *satisfaction probability* $P(T)$.

Let $\Phi(t)$ be the Fluid Approximation of the population model $\mathcal{X}^{(N)}$. To define the transient probability $\mathbf{P}(t)$ of $Z_{\mathcal{A}_{\mathbb{D}}}(t)$, we exploit the fact that, in the case of an IMRP, we have: $\frac{d\mathbf{P}}{dt}(t) = \mathbf{M}(\Phi(t))\mathbf{P}(t) + \mathbf{D}(\Phi(t), \mathbf{P}(t))$ (cf. [15]). In this equation, $\mathbf{M}(\Phi(t))$ is the *generator matrix* for the Markovian transitions, and

$\mathbf{D}(\Phi(t), \mathbf{P}(t))$ accounts for the deterministic events. The elements of $\mathbf{M}(\Phi(t))$ are computed following the same procedure that was described in Sec. 2, where the multiplicity of each transition $(i, s, q) \xrightarrow{\alpha} (i, s', q') \in \mathcal{M}$ in $\mathcal{A}_{\mathbb{D}}$ is always equal to 1 (one single agent) and the Lipschitz limit $f_{\alpha}(\Phi(t))$ 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 (3) or (4)).

To define the components of $\mathbf{D}(\Phi(t), \mathbf{P}(t))$, instead, consider any clock event $e = (\mathcal{A}_i, \Delta_i, p_i) \in \mathcal{E}$, *except* e_1^0 , whose contribute will be computed later on⁶. Choose one of the deterministic transitions $(i, s, q) \dashrightarrow_{e_i} (i+1, s, q)$ 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), and (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 this transition in $\mathbf{D}(\Phi(t), \mathbf{P}(t))$, we need to: (1) record the flux of probability that entered \mathcal{A}_i at time $t - \Delta_i$, and (2) multiply it 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 $s_{out} = (i, s_{out}, q_{out})$, and let $\bar{\mathcal{M}}$ be the set \mathcal{M} of Markovian transitions in $\mathcal{A}_{\mathbb{D}}$ modified in order to make the reset transitions that start in \mathcal{A}_i finish into s_{out} (i.e. for every $(i, s, q) \xrightarrow{\alpha} (i', s', q') \in \mathcal{R} \subset \mathcal{M}$, we define $(i, s, q) \xrightarrow{\alpha} s_{out} \in \bar{\mathcal{M}}$), and to have s_{out} absorbing⁷. Let $\mathbf{G}_i(\Phi(t)) \in \text{Matr}(|\bar{\mathcal{A}}_i| \times |\bar{\mathcal{A}}_i|)$ be the time-dependent matrix s.t.

$$(\mathbf{G}_i(\Phi(t)))_{(i,s,q),(i,s',q')} = \sum_{(i,s,q) \xrightarrow{\alpha} (i,s',q') \in \bar{\mathcal{M}}} \left[\frac{1}{\Phi_s(t)} f_{\alpha}(\Phi(t)) \right], \quad (6)$$

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 (3) and (4)). Moreover, let the diagonal elements of $\mathbf{G}_i(\Phi(t))$ to be defined so that the rows sum up to zero. Then, we introduce the *probability matrix* $\mathbf{Y}_i(t)$, which is computed in terms of the *generator* $\mathbf{G}_i(\Phi(t))$ according to the following ODEs (see also [8]):

$$\begin{cases} \frac{d\mathbf{Y}_i}{dt}(t) = \mathbf{Y}_i(t)\mathbf{G}_i(\Phi(t)) - \mathbf{G}_i(\Phi(t - \Delta_i))\mathbf{Y}_i(t), & \Delta_i \leq t \leq T, \\ \frac{d\mathbf{Y}_i}{dt}(t) = \mathbf{Y}_i(t)\mathbf{G}_i(\Phi(t)), & 0 \leq t \leq \Delta_i, \end{cases} \quad (7)$$

with $\mathbf{Y}_i(0) = \mathbf{I}$. By definition, $(\mathbf{Y}_i(t))_{(i,s',q'),(i,s,q)}$ 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 (Markovianly) 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 $\mathbf{Y}_i(t)$, we can now define the component of $\mathbf{D}(\Phi(t), \mathbf{P}(t))$ that corresponds to the deterministic transition $(i, s, q) \dashrightarrow_{e_i}$

⁶ If e is one of events of the 1st Time Region, i.e. $e = e_1^j$, for some $j = 1, \dots, \ell$, in this section, we drop the index j to ease the notation, i.e. we write $e_1^j = e_1 = (\mathcal{A}_1, \Delta_1, p_1)$.

⁷ The absorbing state s_{out} is needed for the probability $\mathbf{Y}_i(t)$ of step (2) to be well defined. Indeed, the agent can deactivate e_i by taking a reset transition.

$(i + 1, s, q)$ of the clock event $e_i \in \mathcal{E}$. This component is the element in position $((i, s, q), (i + 1, s, q))$ in $\mathbf{D}(\Phi(t), \mathbf{P}(t))$, we call it $D_{e_i, s, q}(\Phi(t), \mathbf{P}(t))$, and is given by

$$D_{e_i, s, q}(\Phi(t), \mathbf{P}(t)) = \sum_{(i, \bar{s}, \bar{q}) \in \mathcal{A}_i} \left[\mathbf{1}_{\{i > 1\}} D_{e_{i-1}, \bar{s}, \bar{q}}(\Phi(t - \Delta_i), \mathbf{P}(t - \Delta_i)) + \mathbf{1}_{\{i=1\}} \times \right. \\ \left. \times \sum_{(i', s', q') \xrightarrow{\alpha} (1, \bar{s}, \bar{q}) \in \mathcal{R}} \frac{1}{\Phi_{s'}(t)} f_{\alpha}(\Phi(t - \Delta_1))(\mathbf{P}(t - \Delta_1))_{(i', s', q')} \right] (\mathbf{Y}_i(t))_{(i, \bar{s}, \bar{q}), (i, s, q)}, \quad (8)$$

where $(\mathbf{P}(t - \Delta_1))_{(i', s', q')}$ is the component in position $(i', s', q') \in S_{\mathbb{D}_{i'}}$ in the vector of the transient probability $\mathbf{P}(t - \Delta_1)$ of $Z_{\mathcal{A}_{\mathbb{D}}}$ at time $t - \Delta_1$. In (8), 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 $i > 1$, $D_{e_{i-1}, \bar{s}, \bar{q}}(\Phi(t - \Delta_i), \mathbf{P}(t - \Delta_i))$ 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 1st 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, $(\mathbf{Y}_i(t))_{(i, \bar{s}, \bar{q}), (i, s, q)}$ 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 $\mathbf{D}(\Phi(t), \mathbf{P}(t))$ 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 $\mathbf{D}(\Phi(t), \mathbf{P}(t))$ depends on the state of the system at times $t - \Delta_1, \dots, t - \Delta_k$ (through the probabilities $\mathbf{Y}_i(t)$, $i = 1, \dots, k$), we write $\mathbf{D}(\Phi(t)) = \mathbf{D}(\Phi, \mathbf{P}, \Delta_1, \dots, \Delta_k, t)$. Then, we define the *transient probability* $\mathbf{P}(t)$ of the IMRP $Z_{\mathcal{A}_{\mathbb{D}}}(t)$ as the solution of the following system of DDEs:

$$\mathbf{P}(t) = \int_0^t \mathbf{M}(s) \mathbf{P}(s) ds + \int_0^t \mathbf{D}(\Phi, \mathbf{P}, \Delta_1, \dots, \Delta_k, s) ds + \mathbf{1}_{t \geq \Delta_1} \sum_{(s, q) \in S \times Q} y_{e_1^0} \nu_{e_1^0, s, q}. \quad (9)$$

In (9), the third term is a deterministic jump in the value of $\mathbf{P}(t)$ at time $t = \Delta_1$, and represents the contribute of the clock event e_1^0 . In such term, the vectors $\nu_{e_1^0, s, q}$ are the update vectors of the deterministic transitions encoded by e_1^0 (hence, the sum is computed over all such 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 $\mathcal{A}_{\mathbb{D}}$) in the matrix $\mathbf{Y}_{e_1^0}(t)$ defined by:

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

with $\mathbf{G}_1(\Phi(t))$ defined as in (6), and $\mathbf{Y}_{e_1^0}(0) = \mathbf{I}$. Hence, $y_{e_1^0} = (\mathbf{Y}_{e_1^0}(\Delta_1))_{s_{0, \mathbb{D}}, (1, s, q)}$ 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 $\mathcal{A}_{\mathbb{D}}$, the IMRP $Z_{\mathcal{A}_{\mathbb{D}}}(t)$, and its transient probability $\mathbf{P}(t)$, the following result holds true.

Proposition 2. *There is a 1:1 correspondence between $\Sigma_{\mathcal{A},\mathbb{D},T}$ and the set $\text{AccPath}(\mathcal{A}_{\mathbb{D}},T)$ of accepted paths of duration T of $\mathcal{A}_{\mathbb{D}}$. Hence,*

$$P(T) = \text{Prob}_Z\{\Sigma_{\mathcal{A},\mathbb{D},T}\} = \text{Prob}_{Z_{\mathcal{A}_{\mathbb{D}}}}\{\text{AccPath}(\mathcal{A}_{\mathbb{D}},T)\} = P_{F_{\mathbb{D}}}(T),$$

where $\text{Prob}_{Z_{\mathcal{A}_{\mathbb{D}}}}$ is the probability measure defined by $Z_{\mathcal{A}_{\mathbb{D}}}$, and $P_{F_{\mathbb{D}}}(T)$ is the sum of the components of $\mathbf{P}(T)$ corresponding to the final states $F_{\mathbb{D}}$ of $\mathcal{A}_{\mathbb{D}}$. \square

In other words, according to Proposition 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 $\mathcal{A}_{\mathbb{D}}$ in Fig. 2, the non-zero off-diagonal entries of the generator matrix $\mathbf{G}_{e_1^1}(\Phi(t))$ 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 $\mathbf{G}_{e_1^1}(\Phi(t))$, we can define $\mathbf{Y}_{e_1^1}(t)$, as in (7), that is then used in the DDEs (9) for the probability $\mathbf{P}(t)$. In this latter set of 9 DDEs (one for each state of $\mathcal{A}_{\mathbb{D}}$), we have:

$$\begin{aligned} P_{(1,S,q1)}(t) = & \int_0^t k_r \mathbf{P}_{(1,S,q1)}(s) ds - \int_0^t k_i \Phi_I(s) \mathbf{P}_{(1,S,q1)}(s) ds + \\ & - \int_0^t k_r \mathbf{Y}_{(1,S,q1),(1,S,q1)}(s-5,s) \mathbf{P}_{(1,S,q1)}(s) ds. \end{aligned}$$

Remark. The presence of *only one clock* in \mathbb{D} enables us to define $\mathcal{A}_{\mathbb{D}}$ in such a way that $Z_{\mathcal{A}_{\mathbb{D}}}(t)$ 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 $\mathcal{A}_{\mathbb{D}}$ 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 $\mathcal{A}_{\mathbb{D}}$ 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 [19] and [4].

3.1 The Mean Behaviour of the Population Model

It is possible to modify our FMC procedure in order to compute the *mean* number of agents that satisfy \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 $\mathcal{A}_{\mathbb{D}}$ defined in Sec. 3. In terms of $\mathcal{A}_{\mathbb{D}}$, we can build the population model $\mathcal{X}_{\mathbb{D}}$, with $\mathcal{A}_{\mathbb{D}}$ as the only agent class, and the sum $P_{F_{\mathbb{D}}}(T)$ of the components corresponding to the final states of $\mathcal{A}_{\mathbb{D}}$ in the Fluid Approximation $\Phi(t)$ 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 of [10], 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 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* states: $(1, I, q0)$, $(1, I, q2)$, $(2, I, q0)$, $(2, I, q1)$ or $(2, I, q2)$. For this reason, we have to

Fluid Model Checking

N	MeanRelErr	MaxRelErr	RelErr(T)	TimeDES	TimeFMC	Speedup
250	0.0927	6.4512	0.1043	11.0273	0.4731	23.3086
500	0.0204	1.7191	0.0048	44.0631	0.3980	110.7113
1000	0.0118	0.7846	0.0003	170.9154	0.3998	427.5022

Fluid Approximation of the mean behaviour

N	MeanRelErr	MaxRelErr	RelErr(T)	TimeDES	TimeFluid	Speedup
250	0.1127	0.2316	0.0921	105.5647	0.4432	339.7217
500	0.0289	0.3177	0.0289	415.0635	0.4237	979.6165
1000	0.0117	0.2216	0.0117	1547.0340	0.4213	3672.0484

Table 1. 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 . The table enlists also the execution times (in seconds) of the DES (TimeDES) and the approximations (TimeFMC and TimeFluid), and the speedups (TimeDES divided by the other times).

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 of $\mathcal{A}_{\mathbb{D}}$, recorded in the counting variables $X_{(1,I,q0)}(t)$, $X_{(1,I,q2)}(t)$, $X_{(2,I,q0)}(t)$, $X_{(2,I,q1)}(t)$ or $X_{(2,I,q2)}(t)$. The same reasoning has to be followed for the definition of the infections of the agents in states $(1, S, q1)$, $(1, S, q2)$, $(2, S, q0)$, $(2, S, q1)$ and $(2, S, q2)$. At the end, as for the single agent, due to the deterministic events, the Fluid Approximation $\Phi(t)$ of $\mathcal{X}_{\mathbb{D}}$ is the solution of a system of DDEs similar to (9). The definition of such approximating equations for a population model with exponential and deterministic transitions is not new [22], but, even if the results are promising (see Sec. 4), to our knowledge, nobody has yet proven the convergence of the estimation in the limit $N \rightarrow +\infty$. We save the investigation of this result for future work.

4 Experimental Results

To validate the procedures of Sec. 3, 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 Fig. 2, 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 Tab. 1. Regarding the errors, we would like to remark that 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




Fig. 3. The satisfaction probability $P(T) = P_{F_D}(T)$ obtained by the Fluid Model Checking (left) and the Fluid Approximation of the mean behaviour (right) in the case $N = 1000$. The results are compared with those obtained by the DES.

of $P_{F_D}(t)$ 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 (approximately 3 orders of magnitude in the case of the Fluid for $N = 1000$) than the simulation based method. Fig. 3 shows the results of the FMC and the Fluid Approximation in the case $N=1000$.

5 Conclusions

We defined a fast and efficient FMC 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.

Future Work. During the experimental analysis, we realised that, on certain models and properties, the DDEs (7) can be *stiff*, and their numerical integration in MATLAB is unstable (see also [8]). In the future, we want to address this issue by considering alternative integration methods [21], investigating also numerical techniques for MRP with time-dependent rates [26]. Furthermore, we plan to prove the convergence of the Fluid Approximation of Sec. 3.1, and to investigate higher-order estimates as in [10, 11]. Finally, we want to extend the FMC procedure of this paper to validate requirements specified in the logic CSL^{TA} [17] and DTA properties endowed with multiple clocks (possibly considering the approximation techniques defined in [19] and [4]).

References

1. R. Alur and D. L. Dill. A Theory of Timed Automata. *Theor. Comput. Sci.*, 1994.
2. H. Andersson and T. Britton. *Stochastic Epidemic Models and their Statistical Analysis*. Springer New York, 2000.

3. C. Baier and J.P. Katoen. *Principles of Model Checking*. MIT press, 2008.
4. 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.
5. M. Benaim and J.-Y. Le Boudec. A Class of Mean Field Interaction Models for Computer and Communication Systems. *Perform. Evaluation*, 2008.
6. L. Bortolussi and J. Hillston. Fluid Approximation of CTMC with Deterministic Delays. In *Proceedings of QEST*, 2012.
7. L. Bortolussi and J. Hillston. Efficient Checking of Individual Rewards Properties in Markov Population Models. In *Proceedings of QAPL*, 2015.
8. L. Bortolussi and J. Hillston. Model Checking Single Agent Behaviours by Fluid Approximation. *Inform. Comput.*, 2015.
9. L. Bortolussi, J. Hillston, D. Latella, and M. Massink. Continuous Approximation of Collective Systems Behaviour: a Tutorial. *Perform. Evaluation*, 2013.
10. L. Bortolussi and R. Lanciani. Model Checking Markov Population Models by Central Limit Approximation. In *Proceedings of QEST*, 2013.
11. L. Bortolussi and R. Lanciani. Stochastic Approximation of Global Reachability Probabilities of Markov Population Models. In *Proceedings of EPEW*, 2014.
12. T. Chen, M. Diciolla, M. Kwiatkowska, and A. Mereacre. Verification of linear duration properties over CTMCs. *Proceedings of TOCL*, 2013.
13. 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.
14. T. Chen, T. Han, J.-P. Katoen, and A. Mereacre. Observing continuous-time MDPs by 1-clock timed automata. In *Reachability Problems*. Springer, 2011.
15. E. Çinlar. *Introduction to Stochastic Processes*. Courier Corporation, 2013.
16. R. Darling, J. Norris, et al. Differential Equation Approximations for Markov Chains. *Probability Surveys*, 2008.
17. S. Donatelli, S. Haddad, and J. Sproston. Model Checking Timed and Stochastic Properties with CSL^{TA}. *IEEE Trans. Software Eng.*, 2009.
18. S. N. Ethier and T. G. Kurtz. *Markov Processes: Characterization and Convergence*. Wiley, 2005.
19. H. Fu. Approximating acceptance probabilities of ctmc-paths on multi-clock deterministic timed automata. In *Proceedings of HSCC*, 2013.
20. N. Gast and G. Bruno. A Mean Field Model of Work Stealing in Large-Scale Systems. *ACM SIGMETRICS Performance Evaluation Review*, 2010.
21. N. Guglielmi and E. Hairer. Implementing Radau IIA Methods for Stiff Delay Differential Equations. *Computing*, 2001.
22. R. A. Hayden. Mean Field for Performance Models with Deterministically-Timed Transitions. In *Proceedings of QEST*, 2012.
23. R.A. Hayden, J.T. Bradley, and A. Clark. Performance Specification and Evaluation with Unified Stochastic Probes and Fluid Analysis. *IEEE Trans. Software Eng.*, 2013.
24. M. Kwiatkowska, G. Norman, and D. Parker. PRISM 4.0: Verification of Probabilistic Real-Time Systems. In *Computer Aided Verification*. Springer, 2011.
25. D. Latella, M. Loreti, and M. Massink. On-the-Fly Fast Mean-Field Model-Checking. In *Proceedings of TGC*. 2014.
26. A. Zimmermann, J. Freiheit, R. German, and G. Hommel. Petri Net Modelling and Performability Evaluation with TimeNET 3.0. In *Computer Performance Evaluation*. 2000.