



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Model Checking Single Agent Behaviours by Fluid Approximation

Citation for published version:

Bortolussi, L & Hillston, J 2015, 'Model Checking Single Agent Behaviours by Fluid Approximation', *Information and Computation*, vol. 242, pp. 183-226. <https://doi.org/10.1016/j.ic.2015.03.002>

Digital Object Identifier (DOI):

[10.1016/j.ic.2015.03.002](https://doi.org/10.1016/j.ic.2015.03.002)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Information and Computation

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Model Checking Single Agent Behaviours by Fluid Approximation

Luca Bortolussi^{a,b}, Jane Hillston^c

^a*Department of Mathematics and Geosciences
University of Trieste, Italy.*

^b*CNR/ISTI, Pisa, Italy.*

^c*Laboratory for the Foundations of Computer Science,
School of Informatics, University of Edinburgh, UK.*

Abstract

In this paper we investigate a potential use of fluid approximation techniques in the context of stochastic model checking of CSL formulae. We focus on properties describing the behaviour of an individual agent in a (large) population of agents, exploiting a limit result known as fast simulation. In particular, we will approximate the behaviour of a single agent with a time-inhomogeneous CTMC, which depends on the environment and on the other agents only through the solution of the fluid differential equation, and model check this process. In order to achieve this goal, we will present a novel procedure to model check time-inhomogeneous CTMC against CSL formulae, investigating also the decidability of this model checking problem. We will then prove the asymptotic correctness of our approach in terms of satisfiability of CSL formulae.

Keywords: Stochastic model checking, fluid approximation, mean field approximation, reachability probability, time-inhomogeneous Continuous Time Markov Chains

1. Introduction

In recent years, there has been a growing interest in fluid approximation techniques in the formal methods community [1, 2, 3, 4, 5, 6]. These

Email addresses: luca@dmf.units.it (Luca Bortolussi), jane.hillston@ed.ac.uk (Jane Hillston)

techniques, also known as mean field approximation, are useful for analysing quantitative models of systems based on continuous time Markov Chains (CTMC), possibly described in process algebraic terms. They work by approximating the discrete state space of the CTMC by a continuous one, and by approximating the stochastic dynamics of the process with a deterministic one, expressed by means of a set of differential equations. The asymptotic correctness of this approach is guaranteed by limit theorems [7, 8, 9], showing the convergence of the CTMC to the fluid ODE for systems of increasing size.

The notion of size can be different from domain to domain, yet in models of interacting agents usually considered in computer science, the size has the standard meaning of number of individuals in a population. All these fluid approaches, in particular, require a shift from an agent-based description to a population-based one, in which the system is represented by variables counting the number of agents in each possible state and so individual behaviours are abstracted. In fact, in large systems, the individual choices of single agents have a small impact, hence the whole system tends to evolve according to the average behaviour of agents. Therefore, the deterministic description of the fluid approximation is mainly related to the average behaviour of the model, and information about statistical properties is generally lost, although it can be partially recovered by introducing fluid equations of higher order moments of the stochastic process (moment closure techniques [10, 11, 12]).

When kept discrete, quantitative systems like those described by process algebras can be analysed using quantitative model checking. These techniques have a long tradition in computer science and are powerful ways of querying a model and extracting information about its behaviour. As far as stochastic model checking is considered, there are some established approaches based mainly on checking Continuous Stochastic Logic (CSL) formulae [13, 14, 15], which led to widely used software tools [16]. All these methods, however, suffer (in a more or less relevant way) from the curse of state space explosion, which severely hampers their practical applicability. In order to mitigate these combinatorial barriers, multiple techniques have been developed, many of them based on some notion of abstraction or approximation of the original process [17, 18].

In this paper, we will precisely target this problem, investigating the extent to which fluid approximation techniques can be used to speed up the model checking of CTMC. We will focus on a restricted subset of system

properties: We will consider population models in which many agents interact, and then focus on the behaviour of individual agents. In fact, even if large systems behave almost deterministically, the evolution of a single agent in a large population is always stochastic. Single agent properties are interesting in many application domains. For instance, in performance models of computer networks, like client-server interactions, one is often interested in the behaviour and quality-of-service metrics of a single client (or a single server), such as the waiting time of the client or the probability of a time-out.

Single agent properties may also be interesting in other contexts. In ecological models, one may be interested in the chances of survival or reproduction of an animal, or in its foraging patterns [19]. In biochemistry, there is some interest in the stochastic properties of single molecules in a mixture (single molecule enzyme kinetics [20, 21]). Other examples may include the time to reach a certain location in a traffic model of a city, or the chances to escape successfully from a building in case of emergency egress [22].

The use of fluid approximation in this restricted context is made possible by a corollary of the fluid convergence theorems, known by the name of *fast simulation* [23, 9], which provides a characterization of the behaviour of a single agent in terms of the solution of the fluid equation: the agent senses the rest of the population only through its “average” evolution, as given by the fluid equation. This characterization can be proved to be asymptotically correct.

The main idea of this paper is simply to use the CTMC for a single agent obtained from the fluid approximation instead of the full model with N interacting agents. In fact, extracting metrics from the description of the global system can be extremely expensive from a computational point of view. Fast simulation, instead, allows us to abstract the system and study the evolution of a single agent (or of a subset of agents) by decoupling its evolution from the evolution of its environment. This has the effect of drastically reducing the dimensionality of the state space by several orders of magnitude.

Of course, in applying the mean field limit, we are introducing an error which is difficult to control (there are error bounds but they depend on the final time and they are very loose [9]). However, this error decreases as the populations increase, and it is usually acceptable in practice, especially for systems with a large pool of agents, as certified by the widespread use of fluid approximation [24]. We stress that these are precisely the cases in which current tools suffer severely from state space explosion, and that can benefit most from a fluid approximation. However, we will see in the following

that in many cases the quality of the approximation is good also for small populations.

In the rest of the paper, we will basically focus on how to analyse single agent properties of three kinds:

- Next-state probabilities, i.e. the probability of jumping into a specific set of states, at a specific time.
- Reachability properties, i.e. the probability of reaching a set of states G , while avoiding unsafe states U .
- Branching temporal logic properties within a bounded amount of time, i.e. verifying time-bounded CSL formulae.

A central feature of the abstraction based on fluid approximation is that the limit of the model of a single agent has rates depending on time, via the solution of the fluid ODE. Hence, the limit models are time-inhomogeneous CTMC (ICTMC). This introduces some additional complexity in the approach, as model checking of an ICTMC is far more difficult than in the standard time-homogeneous case. To the best of the authors' knowledge, in fact, there is no known algorithm to solve this problem in general, although related work is presented in Section 2. We will discuss a general method in Section 5, based on the solution of variants of the Kolmogorov equations, which is expected to work for small state spaces and the controlled dynamics of the fluid approximation. The main difficulty with CSL model checking of ICTMC is that the truth of a formula can depend on the time at which the formula is evaluated. Hence, we need to impose some regularity on the dependency of rates on time to control the complexity of time-dependent truth. We will see that the requirement, piecewise analyticity of rate functions, is intimately connected not only with the decidability of the model checking for ICTMC, but also with the lifting of convergence results from CTMC to truth values of CSL formulae (Theorems 5.1 and 6.1).

Summarising, the main contributions of the paper are the following:

- Methodologically, we advocate the use of fluid approximation to efficiently verify properties of individual agents in large population models, dubbing this approach *fluid model checking*.
- We present a novel model checking algorithm for time-bounded CSL properties on time-inhomogeneous CTMCs.

- We prove that asymptotic correctness of fluid model checking for time-bounded CSL formulae.

The structure of the paper reflects the three main contributions listed above. We start by discussing related work in Section 2, and by introducing preliminary notions, fixing the class of models considered (Section 3.1), presenting fluid limit and fast simulation theorems (Sections 3.2 and 3.3), and introducing Continuous Stochastic Logic (Section 3.4) and time-inhomogeneous CTMCs (Section 3.5). In Section 4, we discuss the main idea of fluid model checking. In Section 5, instead, we present the CSL model checking algorithm for time-inhomogeneous CTMCs, discussing first how to compute next state (Section 5.1) and reachability probabilities (Section 5.3), and then how to combine these routines into an appropriate model checker (Section 5.4). Decidability and complexity are investigated in Section 5.5, under some constraints on rates discussed in Section 5.2. Readers interested only in CSL model checking for ICTMC can read this section independently from the rest of the paper, save for the required preliminary notions (Sections 3.4 and 3.5). In Sections 6 and 7, instead, we investigate the asymptotic correctness of fluid model checking. Finally, in Section 8, we discuss open issues and future work. All the proofs of propositions, lemmas, and theorems of the paper are presented in Appendix A. A preliminary version of this work has appeared in [25].

2. Related work

Our work is underpinned by the notion of fast simulation, which has previously been applied in a number of different contexts [9]. One recent case is a study of policies to balance the load between servers in large-scale clusters of heterogeneous processors [23]. A similar approach is adopted in [26], in the context of Markov games. These ideas also underlie the work of Hayden *et al.* in [27]. Here the authors extend the consideration of transient characteristics as captured by the fluid approximation, to approximation of first passage times, in the context of models generated from the stochastic process algebra PEPA. Their approach for passage times related to individual components is closely related to the fast simulation result and the work presented in this paper. Through fast simulation we are able to reduce the model checking problem on an extremely large CTMC to a model checking problem on a relatively small ICTMC.

Model checking (time homogeneous) Continuous Time Markov Chains (CTMC) against Continuous Stochastic Logic (CSL) specifications has a long tradition in computer science [13, 14, 15]. At the core of our approach to study time-bounded properties there are similarities to that developed in [13], because we consider a transient analysis of a Markov chain whose structure has been modified to reflect the formula under consideration. But the technical details of the transient analysis, and even the structural modification, differ to reflect the time-inhomogeneous nature of the process we are studying.

In contrast, the case of time-inhomogeneous CTMCs has received much less attention. To the best of the authors' knowledge, there has been no previous proposal of an algorithm to model check CSL formulae on a ICTMC. Nevertheless model checking of ICTMCs has been considered with respect to other logics. Specifically, previous work includes model checking of Hennessy-Milner Logic (HML) and Linear Time Logic (LTL) on ICTMC.

In [28], Katoen and Mereacre propose a model checking algorithm for HML on ICTMC. Their work is based on the assumption of piecewise constant rates (with a finite number of pieces) within the ICTMC. The model checking algorithm is based on the computation of integrals and the solution of algebraic equations with exponentials (for which a bound on the number of zeros can be found).

LTL model checking for ICTMC, instead, has been proposed by Chen *et al.* in [29]. The approach works for time-unbounded formulae by constructing the product of the CTMC with a generalized Büchi automaton constructed from the LTL formula, and then reducing the model checking problem to computation of reachability of bottom strongly connected components in this larger (pseudo)-CTMC. The authors also propose an algorithm for solving time bounded reachability similar to the one considered in this paper (for time-constant sets).

Another approach related to the work we present is the verification of CTMC against deterministic time automata (DTA) specifications [30], in which the verification works by taking the product of the CTMC with the DTA, which is converted into a Piecewise Deterministic Markov Process (PDMP, [31]), and then solving a reachability problem for the so-obtained PDMP. This extends earlier work by Baier *et al.* [32] and Donatelli *et al.* [33]. These approaches were limited to considering only a single clock. This means that they are able to avoid the consideration of ICTMC, in the case of [33], through the use of supplementary variables and subordinate CTMCs.

In [34], Chen *et al.* consider the verification of time-homogeneous CTMC against formulae in the metric temporal logic (MTL). This entails finding the probability of a set of timed paths that satisfy the formula over a fixed, bounded time interval. The approach taken is one of approximation, based on an estimate of the maximal number of discrete jumps that will be needed in the CTMC, N , and timed constraints over the residence time within states of a path with up to N steps. The probabilities are then determined by numerically computing a multidimensional integral.

3. Preliminaries

In this section, we will introduce some background material needed in the rest of the paper. First of all, we introduce a suitable notation to describe the population models we are interested in. This is done in Section 3.1. In particular, models will depend parametrically on the (initial) population size, so that we are in fact defining a sequence of models. Then, in Section 3.2, we present the classic fluid limit theorem, which proves convergence of a sequence of stochastic models to the solution of a differential equation. In Section 3.3, instead, we describe fast simulation, a consequence of the fluid limit theorem which connects the system view of the fluid limit to the single agent view, providing a description of single agent behaviour in the limit. In Section 3.4, we recall the basics of Continuous Stochastic Logic (CSL) model checking. Finally, in Section 3.5, we present time-inhomogeneous Continuous Time Markov Chains (ICTMC).

3.1. Modelling Language

In the following, we will describe a basic language for CTMC, in order to fix the notation. We have in mind population models, where a population of agents, possibly of different kinds, interact together through a finite set of possible actions. To avoid a notational overhead, we assume that the number of agents is constant during the simulation, and equal to N . Furthermore, we do not explicitly distinguish between different classes of agents in the notation.

In particular, let $Y_i^{(N)} \in S$ represent the state of agent i , where $S = \{1, 2, \dots, n\}$ is the state space of each agent. Multiple classes of agents can be represented in this way by suitably partitioning S into subsets, and allowing state changes only within a single class. Notice that we made explicit the dependence on N , the total population size.

A configuration of a system is thus represented by the tuple $(Y_1^{(N)}, \dots, Y_N^{(N)})$. When dealing with population models, it is customary to assume that single agents in the same internal state cannot be distinguished, hence we can move from the agent representation to the system representation by introducing variables counting how many agents are in each state. With this objective, define

$$X_j^{(N)} = \sum_{i=1}^N \mathbf{1}\{Y_i^{(N)} = j\}, \quad (1)$$

so that the system can be represented by the vector $\mathbf{X}^{(N)} = (X_1^{(N)}, \dots, X_n^{(N)})$, whose dimension is independent of N . The domain of each variable $X_j^{(N)}$ is obviously $\{1, \dots, N\}$.

We will describe the evolution of the system by a set of transition rules at this global level. This simplifies the description of synchronous interactions between agents. The evolution from the perspective of a single agent will be reconstructed from the system level dynamics. In particular, we assume that $\mathbf{X}^{(N)}$ is a CTMC (Continuous-Time Markov Chain), with a dynamics described by a fixed number of transitions, collected in the set $\mathcal{T}^{(N)}$. Each transition $\tau \in \mathcal{T}^{(N)}$ is defined by a *multi-set* of *update rules* R_τ and by a rate function $r_\tau^{(N)}$. The multi-set¹ R_τ contains update rules $\rho \in R_\tau$ of the form $i \rightarrow j$, where $i, j \in S$. Each rule specifies that an agent changes state from i to j . Let $m_{\tau, i \rightarrow j}$ denote the multiplicity of the rule $i \rightarrow j$ in R_τ . We assume that R_τ is independent of N , so that each transition involves a finite and fixed number of individuals. Given a multi-set of update rules R_τ , we can define the *update vector* \mathbf{v}_τ in the following way:

$$\mathbf{v}_\tau = \sum_{(i \rightarrow j) \in R_\tau} m_{\tau, i \rightarrow j} \mathbf{e}_j - \sum_{(i \rightarrow j) \in R_\tau} m_{\tau, i \rightarrow j} \mathbf{e}_i,$$

where \mathbf{e}_i is the vector equal to one in position i and zero elsewhere. The vector \mathbf{v}_τ gives the net change in the number of agents in each state i due to the happening of a τ transition, taking multiplicities into account.

Hence, each transition changes the state from $\mathbf{X}^{(N)}$ to $\mathbf{X}^{(N)} + \mathbf{v}_\tau$. The rate function $r_\tau^{(N)}(\mathbf{X})$ depends on the current state of the system, and specifies the speed of the corresponding transition. It is assumed to be equal to zero if

¹The fact that R_τ is a multi-set, allows us to model events in which agents in the same state synchronise.




Figure 1: Visual representation of the client server system of the running example.

there are not enough agents available to perform a τ transition. Furthermore, it is required to be *Lipschitz continuous*. We indicate such a model by $\mathcal{X}^{(N)} = (\mathbf{X}^{(N)}, \mathcal{T}^{(N)}, \mathbf{x}_0^{(N)})$, where $\mathbf{x}_0^{(N)}$ is the initial state of the model.

Given a model $\mathcal{X}^{(N)}$, it is straightforward to construct the CTMC associated with it, exhibiting its infinitesimal generator matrix. First, its state space is $\mathcal{D} = \{(x_1, \dots, x_n) \mid x_i \in \{1, \dots, N\}, \sum_i x_i = N\}$. The infinitesimal generator matrix Q , instead, is the $\mathcal{D} \times \mathcal{D}$ matrix defined by

$$q_{\mathbf{x}, \mathbf{x}'} = \sum \{r_{\tau}(\mathbf{x}) \mid \tau \in \mathcal{T}, \mathbf{x}' = \mathbf{x} + \mathbf{v}_{\tau}\}.$$

We will indicate the state of such a CTMC at time t by $\mathbf{X}(t)$.

Example. We introduce now the main running example of the paper: we will consider a model of a simple client-server system, in which a pool of clients submits queries to a group of servers, waiting for a reply. In particular, the client asks for information from a server and waits for it to reply. It can time-out if too much time passes. The server, instead, after receiving a request does some processing and then returns the answer. It can time-out while processing and while it is ready to reply. After an action, it always logs data.

The client and server agents are visually depicted in Figure 1. The global system is described by the following 8 variables:

- 4 variables for the client states: C_{rq} , C_w , C_{rc} , and C_t , with domain $[0, N_C]$, N_C the total number of clients.
- 4 variables for the server states: S_{rq} , S_p , S_{rp} , and S_l , with domain $[0, N_S]$, N_S the total number of servers.

Furthermore, there are 9 transitions in total, corresponding to all possible arrows of Figure 1. We list them below, stressing that synchronization between clients and servers has a rate computed using the minimum, in the PEPA style [35], reflecting the strict pairing and handshaking nature of the interaction.

- request: $R_{request} = \{C_{rq} \rightarrow C_w, S_{rq} \rightarrow S_p\}$, $r_{request} = k_r \cdot \min(C_{rq}, S_{rq})$;
- reply: $R_{reply} = \{C_w \rightarrow C_t, S_{rp} \rightarrow S_l\}$, $r_{reply} = \min(k_w C_w, k_{rp} S_{rp})$;
- timeout (client): $R_{timeout1} = \{C_w \rightarrow C_{rc}\}$, $r_{timeout1} = k_{to} C_w$;
- recover: $R_{recover} = \{C_{rc} \rightarrow C_{rq}\}$, $r_{recover} = k_{rec} C_{rc}$;
- think: $R_{think} = \{C_t \rightarrow C_{rq}\}$, $r_{think} = k_t C_t$;
- logging: $R_{logging} = \{S_l \rightarrow S_{rq}\}$, $r_{logging} = k_l S_l$;
- process: $R_{process} = \{S_p \rightarrow S_{rp}\}$, $r_{process} = k_p S_p$;
- timeout (server processing): $R_{timeout2} = \{S_p \rightarrow S_l\}$, $r_{timeout2} = k_{sto} S_p$;
- timeout (server replying): $R_{timeout3} = \{S_{rp} \rightarrow S_l\}$, $r_{timeout3} = k_{sto} S_{rp}$;

As an example of the construction of the state change vectors \mathbf{v}_τ , we report that for the first transition, which is equal to $\mathbf{v}_{request} = \mathbf{e}_{C_w} - \mathbf{e}_{C_{rq}} + \mathbf{e}_{S_p} - \mathbf{e}_{S_{rq}}$, i.e. 1 in the coordinates corresponding to variables C_w and S_p and -1 in the coordinates of C_{rq} and S_{rq} .

The system-level models we have defined depend on the total population $N = N_C + N_S$ and on the ratio N_S/N_C between server and clients, which is specified by the initial conditions. Increasing the total population N (keeping

fixed the client-server ratio), we obtain a sequence of models, and we are interested in their limit behaviour, for N going to infinity.

In order to compare the models of such a sequence, we will normalize them to the same scale, dividing each variable by N and thus introducing the normalized variables $\hat{\mathbf{X}}^{(N)} = \frac{\mathbf{X}^{(N)}}{N}$. In the case of a constant population, normalised variables are usually referred to as the *occupancy measure*, as they represent the fraction of agents in each state. Update vectors are scaled correspondingly, i.e. dividing them by N . Furthermore, we will also require an appropriate scaling (in the limit) of the rate functions of the normalized models. More precisely, let $\mathcal{X}^{(N)} = (\mathbf{X}^{(N)}, \mathcal{T}^{(N)}, \mathbf{X}_0^{(N)})$ be the N -th non-normalized model and $\hat{\mathcal{X}}^{(N)} = (\hat{\mathbf{X}}^{(N)}, \hat{\mathcal{T}}^{(N)}, \hat{\mathbf{X}}_0^{(N)})$ the corresponding normalized model. We require that:

- initial conditions scale appropriately: $\hat{\mathbf{X}}_0^{(N)} = \frac{\mathbf{X}_0^{(N)}}{N}$;
- for each transition $(\mathbf{v}_\tau, r_\tau^{(N)}(\mathbf{X}))$ of the non-normalized model, we let $\hat{r}_\tau^{(N)}(\hat{\mathbf{X}})$ be the rate function expressed in the normalized variables (i.e. after a change of variables). The corresponding transition in the normalized model is $(R_\tau, \hat{r}_\tau^{(N)}(\hat{\mathbf{X}}))$, with update vector equal to $\frac{1}{N}\mathbf{v}_\tau$. We assume that there exists a bounded and Lipschitz continuous function $f_\tau(\hat{\mathbf{X}}) : E \rightarrow \mathbb{R}^n$ on normalized variables (where E contains all domains of all $\hat{\mathcal{X}}^{(N)}$), independent of N , such that $\frac{\hat{r}_\tau^{(N)}(\mathbf{x})}{N} \rightarrow f_\tau(\mathbf{x})$ *uniformly* on E .

We will denote the state of the CTMC of the N -th non-normalized (resp. normalized) model at time t as $\mathbf{X}^{(N)}(t)$ (resp. $\hat{\mathbf{X}}^{(N)}(t)$).

Example. Consider again the running example. If we want to scale the model with respect to the scaling parameter N , we can increase the initial population of clients and servers by a factor k (hence keeping the client-server ratio constant), similarly to [36]. The condition on rates, in this case, automatically holds due to their (piecewise) linear nature.

Remark 3.1. The conditions discussed in this section are necessary for the fluid approximation theorem to hold. Fortunately, they are not very stringent. Lipschitz continuity holds for most practically used rate functions, including the minimum, the product and generalised mass action [24]. In any case, any differentiable function defined on a compact set will be Lipschitz continuous. The condition on convergence of rates rescaled by N ,

furthermore, holds also in most circumstances, being trivially true for linear functions, piecewise linear functions (e.g. defined by the minimum), multi-affine functions [24]. For general non-linear rate functions, the convergence of rates is usually enforced by appropriately scaling parameters with respect to the total population N . However, the meaningfulness of such a scaling has to be checked on a case by case basis.

3.2. Deterministic limit theorem

In order to present the “classic” deterministic limit theorem, we need to introduce a few more concepts needed to construct the limit ODE. Consider a sequence of normalized models $\hat{\mathcal{X}}^{(N)}$ and let \mathbf{v}_τ be the (non-normalised) update vectors. The drift $F^{(N)}(\hat{\mathbf{X}})$ of $\hat{\mathcal{X}}$ is defined as

$$F^{(N)}(\hat{\mathbf{X}}) = \sum_{\tau \in \hat{\mathcal{T}}} \frac{1}{N} \mathbf{v}_\tau \hat{r}_\tau^{(N)}(\hat{\mathbf{X}}) \quad (2)$$

Furthermore, let $f_\tau : E \rightarrow \mathbb{R}^n$, $\tau \in \hat{\mathcal{T}}$ be the limit rate functions of transitions of $\hat{\mathcal{X}}^{(N)}$. We define the *limit drift* of the model $\hat{\mathcal{X}}^{(N)}$ as

$$F(\hat{\mathbf{X}}) = \sum_{\tau \in \hat{\mathcal{T}}} \mathbf{v}_\tau f_\tau(\hat{\mathbf{X}}) \quad (3)$$

It is easily seen that the conditions of the previous subsection imply that $F^{(N)}(\mathbf{x}) \rightarrow F(\mathbf{x})$ uniformly. The limit drift F can be seen as a vector field in E , defining the limit ODE

$$\frac{d\mathbf{x}}{dt} = F(\mathbf{x}), \quad (4)$$

with initial conditions given by $\mathbf{x}(0) = \mathbf{x}_0 \in E$. Given that F is Lipschitz in E (as all f_τ are), the ODE has a unique solution $\mathbf{x}(t)$ in E starting from \mathbf{x}_0 . Then, the following theorem can be proved [7, 8]:

Theorem 3.1 (Deterministic approximation [7, 8]). *Let the sequence $\hat{\mathbf{X}}^{(N)}(t)$ of Markov processes and $\mathbf{x}(t)$ be defined as before, and assume that there is some point $\mathbf{x}_0 \in E$ such that $\hat{\mathbf{X}}^{(N)}(0) \rightarrow \mathbf{x}_0$ in probability. Then, for any finite time horizon $T < \infty$, it holds that:*

$$\mathbb{P} \left\{ \sup_{0 \leq t \leq T} \|\hat{\mathbf{X}}^{(N)}(t) - \mathbf{x}(t)\| > \varepsilon \right\} \rightarrow 0.$$

Notice that the theorem can be specialised to subsets $E' \subseteq E$, in which case it can also provide an estimate of exit times from set E' , see [8]. Furthermore, if the initial conditions converge almost surely, then it also holds that $\sup_{0 \leq t \leq T} \|\hat{\mathbf{X}}^{(N)}(t) - \mathbf{x}(t)\| \rightarrow 0$ almost surely [37].

3.3. Fast simulation

We now turn our attention back to a single individual in the population. Even if the system-level dynamics, in the limit of a large population, becomes deterministic, the dynamics of a single agent remains a stochastic process. However, the fluid limit theorem implies that the dynamics of a single agent, in the limit, becomes essentially dependent on the other agents only through the global system state. This asymptotic decoupling allows us to find a simpler Markov Chain for the evolution of the single agent. This result is often known in the literature [9] under the name of *fast simulation* [23].

To explain this point formally, let us focus on a single individual $Y_h^{(N)}$, which is a Markov process on the state space $S = \{1, \dots, n\}$, conditional on the global state of the population $\hat{\mathbf{X}}^{(N)}(t)$. Let $Q^{(N)}(\mathbf{x})$ be the infinitesimal generator matrix of $Y_h^{(N)}$, described as a function of the normalized state of the population $\hat{\mathbf{X}}^{(N)} = \mathbf{x}$, i.e.

$$\mathbb{P}\{Y_h^{(N)}(t + dt) = j \mid Y_h^{(N)}(t) = i, \hat{\mathbf{X}}^{(N)}(t) = \mathbf{x}\} = q_{i,j}^{(N)}(\mathbf{x})dt.$$

We stress that this is the exact Markov Chain for $Y_h^{(N)}$, conditional on $\hat{\mathbf{X}}^{(N)}(t)$, and that $Y_h^{(N)}(t)$ in general is *not independent* of $\hat{\mathbf{X}}^{(N)}(t)$.² In fact, without conditioning on $\hat{\mathbf{X}}^{(N)}$, $Y_h^{(N)}(t)$ is not a Markov process. This means that in order to capture its evolution in a Markovian setting, one has to consider the Markov chain $(Y_h^{(N)}(t), \hat{\mathbf{X}}^{(N)}(t))$.

Example. Consider the running example, and suppose we want to construct the CTMC for a single client. For this purpose, we have to extract from the specification of global transitions a set of local transitions for the client. The state space of a client will consist of four states, $S_c = \{rq, w, t, rc\}$.

Then, we need to define its rate matrix $Q^{(N)}$. In order to do this, we need to take into account all global transitions involving a client, and then extract the rate at which a specific client can perform such a transition. As

²Independence would imply that all agents evolve independently from one another, hence the model is in product form for any t .

a first example, consider the **think** transition, changing the state of a client from t to rq . Its global rate is $r_{think} = k_t C_t$. As we have C_t clients in state t , the rate at which a specific one will perform a think transition is $\frac{k_t C_t}{C_t} = k_t$. Hence, we just need to divide the global rate of observing a think transition by the total number of clients in state t . Notice that, as we are assuming that one specific client is in state t , then $C_t \geq 1$, hence we are not dividing by zero.

Consider now a **reply** transition. In this case, the transition involves a server and a client in state w . The global rate is $r_{reply} = \min(k_w C_w, k_{rp} S_{rp})$, and $C_w \geq 1$ (in the non-normalized model with total population N). Dividing this rate by C_w , we obtain $\min(k_w, k_{rp} \frac{S_{rp}}{C_w})$, which is defined for $C_w > 0$. If we switch to normalised variables, we obtain a similar expression: $\min(k_w, k_{rp} \frac{s_{rp}}{c_w})$, which is independent of N . However, in taking N to the limit we must be careful: even if in the non-normalized model C_w (and hence c_w) are always non-zero (if a specific agent is in state w), this may not be true in the limit: if only one client is in state w , then the limit fraction of clients in state w is zero (just take the limit of $\frac{1}{N}$). Hence, we need to take care of boundary conditions, guaranteeing that the single-agent rate is defined also in these circumstances. In this case, we can assume that the rate is zero if s_{rp} is zero (whatever the value of c_w), and that the rate is k_w if c_w is zero but $s_{rp} > 0$.

In order to treat the previous set of cases in a homogeneous way, we make the following assumption about rates:

Definition 3.1. Let $\tau \in \mathcal{T}$ be a transition such that its update rule set contains the rule $i \rightarrow j$, with multiplicity $m_{\tau, i \rightarrow j}$. The rate $r_{\tau}^{(N)}$ is *single-agent- i compatible* if there exists a Lipschitz continuous function $f_{\tau}^i(\mathbf{x})$ on normalized variables such that the limit rate on normalized variables $f_{\tau}(\mathbf{x})$ can be factorised as $f_{\tau}(\mathbf{x}) = x_i f_{\tau}^i(\mathbf{x})$. A transition τ is *single-agent compatible* if and only if it is single-agent- i compatible for any i appearing in the left-hand side of an update rule.

Hence, the limit rate of observing a transition from i to j for a specific agent in state i is $m_{\tau, i \rightarrow j} f_{\tau}^i(\mathbf{x})$, where the factor $m_{\tau, i \rightarrow j}$ comes from the fact that it is one out of $m_{\tau, i \rightarrow j}$ agents changing state from i to j due to τ .³

³The factor m stems from the following simple probabilistic argument: if we choose at random m agents out of X_i , then the probability to select a specific agent is $\frac{m}{X_i}$.

Then, assuming all transitions τ are single-agent compatible, we can define the rate $q_{i,j}^{(N)}$ as

$$q_{i,j}^{(N)}(\mathbf{x}) = \sum_{\tau \in \mathcal{T} \mid \{i \rightarrow j\} \subseteq R_\tau} m_{\tau, i \rightarrow j} \frac{r_\tau^{(N)}(\mathbf{x})}{x_i} = \sum_{\tau \in \hat{\mathcal{T}} \mid \{i \rightarrow j\} \subseteq R_\tau} m_{\tau, i \rightarrow j} \frac{\hat{r}_\tau^{(N)}(\hat{\mathbf{x}})}{\hat{x}_i} = q_{i,j}^{(N)}(\hat{\mathbf{x}}).$$

It is then easy to check that

$$q_{i,j}^{(N)}(\mathbf{x}) \rightarrow q_{i,j}(\mathbf{x}) = \sum_{\tau \in \mathcal{T} \mid \{i \rightarrow j\} \subseteq R_\tau} m_{\tau, i \rightarrow j} \frac{f_\tau(\mathbf{x})}{x_i} = \sum_{\tau \in \hat{\mathcal{T}} \mid \{i \rightarrow j\} \subseteq R_\tau} m_{\tau, i \rightarrow j} f_\tau^i(\mathbf{x}).$$

In the following, we fix an integer $k > 0$ and let $Z_k^{(N)} = (Y_1^{(N)}, \dots, Y_k^{(N)})$ be the process tracking the state of k selected agents among the population, with state space $\mathcal{S} = S^k$. Notice that k is fixed and independent of N , so that we will track k individuals embedded in a population that can be very large.

Let $\mathbf{x}(t)$ be the solution of the fluid ODE, and assume we are under the hypothesis of Theorem 3.1. Consider now $z_k^{(N)}(t)$ and $z_k(t)$, the *time-inhomogeneous* CTMCs on \mathcal{S} defined by the following infinitesimal generators (for any $h = 1, \dots, k$):

$$\mathbb{P}\{z_k^{(N)}(t+dt) = (z_1, \dots, j, \dots, z_k) \mid z_k^{(N)}(t) = (z_1, \dots, i, \dots, z_k)\} = q_{i,j}^{(N)}(\mathbf{x}(t))dt,$$

$$\mathbb{P}\{z_k(t+dt) = (z_1, \dots, j, \dots, z_k) \mid z_k(t) = (z_1, \dots, i, \dots, z_k)\} = q_{i,j}(\mathbf{x}(t))dt,$$

Notice that, while $Z_k^{(N)}$ describes exactly the evolution of k agents, $z_k^{(N)}$ and \mathbf{z}_k do not. In fact, they are CTMCs in which the k agents evolve independently, each one with the same infinitesimal generator, depending on the global state of the system via the fluid limit. We stress that the time-inhomogeneity comes from the dependence of the generator on the solution $\mathbf{x}(t)$ of the fluid ODE.

However, the following theorem can be proved [9]:

Theorem 3.2 (Fast simulation theorem). *For any $T < \infty$, $\mathbb{P}\{Z_k^{(N)}(t) \neq z_k^{(N)}(t), \text{ for some } t \leq T\} \rightarrow 0$, and $\mathbb{P}\{Z_k(t) \neq z_k(t), \text{ for some } t \leq T\} \rightarrow 0$, as $N \rightarrow \infty$.*

This theorem states that, in the limit of an infinite population, each fixed set of k agents will behave independently, sensing only the mean state of the

global system, described by the fluid limit $\mathbf{x}(t)$. Furthermore, those k agents will evolve independently, as if there was no synchronisation between them. This *asymptotic decoupling* of the system, holding for any set of k agents, is also known in the literature under the name of *propagation of chaos* [6]. In particular, this holds if we define the rate of the limit CTMC either by the single-agent rates for population N ($z_k^{(N)}$) or by the limit rates (z_k). Note that, when the CTMC has density dependent rates [37], then $z_k^{(N)}(t) = z_k(t)$, as their infinitesimal generators will be the same.

We stress once again that the process $Z_k^{(N)}(t)$ is not a Markov process. It becomes a Markov process when considered together with $\hat{\mathbf{X}}^{(N)}(t)$. This can be properly understood by observing that it is the projection of the Markov process $(Y_1^{(N)}(t), \dots, Y_N^{(N)}(t))$ on the first k coordinates, and recalling that a projection of a Markov process need not be Markov (intuitively, we can throw away some relevant information about the state of the process). However, being the projection of a Markov process, the probability of $Z_k^{(N)}(t)$ at each time t is perfectly defined, and it can be obtained by marginalising out the additional coordinates of $(Y_1^{(N)}(t), \dots, Y_N^{(N)}(t))$.⁴ Nevertheless, its non-Markovian nature has consequences for reachability probabilities and the satisfiability of CSL formulae.

Example. Consider again the client-server example, and focus on a single client. As said before, its state space is $S_c = \{rq, w, t, rc\}$, and the non-null rates of the infinitesimal generator Q for the process z_1 are:

- $q_{rq,w}(t) = k_r \min\{1, s_{rq}(t)/c_{rq}(t)\}$ (with appropriate boundary conditions);
- $q_{w,t}(t) = \min\{k_w, k_{rp}s_{rp}(t)/c_w(t)\}$;
- $q_{w,rc}(t) = k_{to}$;
- $q_{t,rq}(t) = k_t$;
- $q_{rc,rq}(t) = k_{rc}$.

In Figure 2, we show a comparison of the transient probabilities for the approximating chain for a single client and the true transient probabilities,

⁴Formally, $\mathbb{P}\{Z_k^{(N)}(t) = (s_1, \dots, s_k)\} = \sum_{\mathbf{s} \in S^{N-k}} \mathbb{P}\{(Y_1^{(N)}(t), \dots, Y_N^{(N)}(t)) = (s_1, \dots, s_k, \mathbf{s})\}$.

estimated by Monte Carlo sampling of the CTMC, for different population levels N . As we can see, the approximation is quite precise already for $N = 15$.

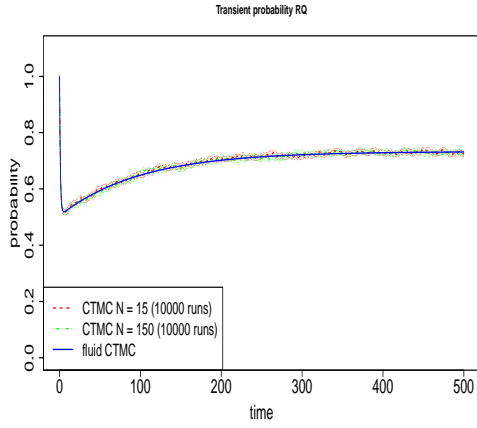
Remark 3.2. Single-agent consistency is not a very restrictive condition. However, there are cases in which it is not satisfied. One example is passive rates in PEPA [35]. In this case, in fact, the rate of the synchronization of $P = (\alpha, \top).P1$ and $Q = (\alpha, r).Q1$ is $rX_Q\mathbf{1}\{X_P > 0\}$. In particular, the rate is independent of the exact number of P agents. If we look at a single P -agent, its rate equals $r\frac{X_Q}{X_P}\mathbf{1}\{X_P > 0\}$. Normalising variables, we get the rate $r\frac{x_Q}{x_P}\mathbf{1}\{x_P > 0\}$, which approaches infinity as x_P goes to zero (for x_Q fixed). Hence, it cannot be extended to a Lipschitz continuous function. However, in the case $x_P = 0$ and $x_Q > 0$, if we look at a single agent, then the speed at which P changes state is in fact infinite. We can see this by letting $X_P = 1$ and $X_Q = Nq$, so that the rate of the transition from the point of view of P is $Nq \rightarrow \infty$. Thus, in the limit, the state $X_P = 0$ becomes vanishing.

Remark 3.3. The hypothesis of constant population, i.e. the absence of birth and death, can be relaxed. The fluid approximation continues to work also in the presence of birth and death events, and so does the fast simulation theorem. In our framework, birth and death can be easily introduced by allowing rules of the form $\emptyset \rightarrow i$ (for birth) and $i \rightarrow \emptyset$ (for death). In terms of a single agent, death can be dealt with by adding a single absorbing state to its local state space \mathcal{S} . Birth, instead, means that we can choose the time instant at which an agent enters the system (provided that there is a non-null rate for birth transitions at the chosen time).


Another solution would be to assume an infinite pool of agents, among which only finitely many can be alive, and the others are an infinite supply of “fresh souls”. Even if this is plausible from the point of view of a global model, it creates problems in terms of a single agent perspective (what is the rate of birth of a soul?). A solution can be to assume a large but finite pool of agents. But in this case birth becomes a passive action (and it introduces discontinuities in the model, even if in many cases one can guarantee to remain far away from the discontinuous boundary), hence we face the same issues discussed in Remark 3.2.

3.4. Continuous Stochastic Logic

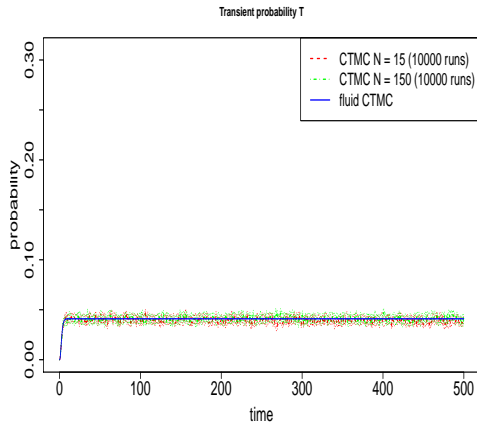
In this section we consider labelled stochastic processes. A labelled stochastic process is a random process $Z(t)$, with state space \mathcal{S} and a labelling




(a) $\mathbb{P}\{rq\}$



(b) $\mathbb{P}\{w\}$



(c) $\mathbb{P}\{t\}$



(d) $\mathbb{P}\{rc\}$

Figure 2: Comparison of the transient probability for all four states of the fluid model of the client-server system, computed solving the Kolmogorov forward equations, and the transient probability of CTMC models for $N = 15$ and $N = 150$ (2:1 client server ratio). Parameters are $k_r = 1$, $k_w = 100$, $k_{to} = 0.01$, $k_t = 1$, $k_{rc} = 100$, $k_l = 10$, $k_p = 0.1$, $k_{sto} = 0.005$, initial conditions of the full system are $C_{rq} = n$, $S_{rq} = m$, while the single client CTMC starts in state rq .

function $L : \mathcal{S} \rightarrow 2^{\mathcal{P}}$, associating with each state $s \in \mathcal{S}$ a subset of atomic propositions $L(s) \subset \mathcal{P} = \{a_1, \dots, a_k, \dots\}$ true in that state: each atomic proposition $a_i \in \mathcal{P}$ is true in s if and only if $a_i \in L(s)$. We require that all subsets of paths considered are measurable. This condition will be satisfied by all subsets considered in the paper. It is satisfied for a CTMC or the projection of a CTMC with constant rates, see [13]. It holds also for the subsets of paths of the time-inhomogeneous CTMC we will work with, as we will show in Section 5.

A path of $Z(t)$ starting in state s_0 at time t_0 is a sequence

$$\sigma = s_0 \xrightarrow{\delta_0} s_1 \xrightarrow{\delta_1} \dots,$$

such that the probability (density) of going from s_i to s_{i+1} at time $t_0 + t_\sigma[i]$, with $t_\sigma[i] = \sum_{j=0}^i \delta_j$, is greater than zero. For CTMCs, this condition is equivalent to $q_{s_i, s_{i+1}}(t_0 + t_\sigma[i]) > 0$. Denote with $\sigma@t$ the state of σ at time t , with $\sigma[i]$ the i -th state of σ , and with $t_\sigma[i]$ the time spent in the i -th state $\sigma[i]$ in σ .

A time-bounded CSL formula φ is defined by the following syntax:

$$\varphi = a \mid \varphi_1 \wedge \varphi_2 \mid \neg\varphi \mid P_{\bowtie p}(\mathbf{X}^{[T_1, T_2]}\varphi) \mid P_{\bowtie p}(\varphi_1 \mathbf{U}^{[T_1, T_2]}\varphi_2),$$

where $T_1, T_2 \in \mathbb{R}_{\geq 0}$, $T_1 \leq T_2 < \infty$. The satisfiability relation of φ with respect to a labelled stochastic process $Z(t)$ is given by the following rules:

- $s, t_0 \models a$ if and only if $a \in L(s)$;
- $s, t_0 \models \neg\varphi$ if and only if $s, t_0 \not\models \varphi$;
- $s, t_0 \models \varphi_1 \wedge \varphi_2$ if and only if $s, t_0 \models \varphi_1$ and $s, t_0 \models \varphi_2$;
- $s, t_0 \models P_{\bowtie p}(\mathbf{X}^{[T_1, T_2]}\varphi)$ if and only if $\mathbb{P}\{\sigma \mid \sigma, t_0 \models \mathbf{X}^{[T_1, T_2]}\varphi\} \bowtie p$.
- $s, t_0 \models P_{\bowtie p}(\varphi_1 \mathbf{U}^{[T_1, T_2]}\varphi_2)$ if and only if $\mathbb{P}\{\sigma \mid \sigma, t_0 \models \varphi_1 \mathbf{U}^{[T_1, T_2]}\varphi_2\} \bowtie p$.
- $\sigma, t_0 \models \mathbf{X}^{[T_1, T_2]}\varphi$ if and only if $t_\sigma[1] \in [T_1, T_2]$ and $\sigma[1], t_0 + t_\sigma[1] \models \varphi$.
- $\sigma, t_0 \models \varphi_1 \mathbf{U}^{[T_1, T_2]}\varphi_2$ if and only if $\exists \bar{t} \in [t_0 + T_1, t_0 + T_2]$ s.t. $\sigma@\bar{t}, \bar{t} \models \varphi_2$ and $\forall t_0 \leq t < \bar{t}$, $\sigma@t, t \models \varphi_1$.

Notice that we are considering a fragment of CSL without the steady state operator and allowing only time-bounded properties. This last restriction is connected with the nature of convergence theorems 3.1 and 3.2, which hold only on finite time horizons. However, see Remark 6.1 for possible relaxations of this restriction.

Model checking of a next CSL formula $P_{\bowtie p}(\mathbf{X}^{[T_1, T_2]}\varphi)$ is usually performed by computing the next-state probability via an integral, and then comparing the so-obtained value with the threshold p . Model checking of an until CSL formula $P_{\bowtie p}(\varphi_1 \mathbf{U}^{[T_1, T_2]}\varphi_2)$ in a *time-homogeneous CTMC* $Z(t)$, instead, can be reduced to the computation of two reachability problems, which themselves can be solved by transient analysis [13]. In particular, consider the sets of states $U = \llbracket \neg\varphi_1 \rrbracket = \{s \in \mathcal{S} \mid s \models \neg\varphi_1\}$ and $G = \llbracket \varphi_2 \rrbracket = \{s \in \mathcal{S} \mid s \models \varphi_2\}$ and compute the probability $\pi_{s_1, s_2}^1(T_1)$ of going from state $s_1 \notin U$ to a state $s_2 \notin U$ in T_1 time units, in the CTMC in which all U -states are made absorbing. Furthermore, consider the modified CTMC in which all U and G states are made absorbing, and denote by $\pi_{s_2, s_3}^2(T_2 - T_1)$ the probability of going from a state $s_2 \notin U$ to a state $s_3 \in G$ in $T_2 - T_1$ units of time in such a CTMC. Then the probability of the until formula in state s can be computed as $P_s(\varphi) = \sum_{s_3 \in G, s_2 \notin U} \pi_{s_1, s_2}^1(T_1) \pi_{s_2, s_3}^2(T_2 - T_1)$. The probabilities π^1 and π^2 can be computed using standard methods for transient analysis (e.g. by uniformisation [38] or by solving the Kolmogorov equations [39]). Then, to determine the truth value of the formula φ in state s , one has just to solve the inequality $P_s(\varphi) \bowtie p$. The truth value of a generic CSL formula can therefore be computed recursively on the structure of the formula.

3.5. Time-inhomogeneous Continuous Time Markov Chains

A time-inhomogeneous Continuous Time Markov Chain (ICTMC) is a generalisation of a CTMC with time-dependent rates for state transitions, so that the infinitesimal generator matrix $Q(t)$ is dependent on time [39]. It still satisfies the Markov property: the future behaviour depends only on the current state and the current time, not on the history of how the state was reached. We will denote a generic ICTMC on state space \mathcal{S} by $Z(t)$.

An ICTMC $Z(t)$ is fully specified by its *time-dependent infinitesimal generator matrix* $Q(t)$. From this, we can reconstruct the probability matrix $\Pi(t_1, t_2)$ of $Z(t)$, whose entry $\pi_{s_1, s_2}(t_1, t_2)$ gives the probability of being in state s_2 at time t_2 conditional on being in state s_1 at time t_1 , by solving the (generalised) *Kolmogorov forward and backward equations*. More specifically, the Kolmogorov forward equation describes the time evolution of $\Pi(t_1, t_2)$ as

a function of t_2 , and is written $\frac{\partial \Pi(t_1, t_2)}{\partial t_2} = \Pi(t_1, t_2)Q(t_2)$, while the backward equation expresses how $\Pi(t_1, t_2)$ varies with respect to t_1 , and it is $\frac{\partial \Pi(t_1, t_2)}{\partial t_1} = -Q(t_1)\Pi(t_1, t_2)$.

For model checking purposes, we need to compute probabilities of sets of trajectories of $Z(t)$, which must be measurable according to the sigma-algebra \mathcal{F} on the set *Paths* of paths of $Z(t)$. Such a sigma-algebra \mathcal{F} , as customary, is defined as the smallest sigma-algebra containing the cylinder sets \mathcal{C} . Let \mathcal{I} be the set of non-empty intervals of $\mathbb{R}_{\geq 0}$ with rational endpoints. A *cylinder set* $C_{t_0}(s_0, I_0, s_1, \dots, I_{n-1}, s_n)$, consists of all paths that are in s_0 at time t_0 and that jump to s_j , $1 \leq j \leq n$ at a time $t \in t_0 \oplus I_0 \oplus \dots \oplus I_{j-1}$, where \oplus is the Minkowsky sum of two sets, defined as $A \oplus B = \{a + b \mid a \in A, b \in B\}$. The (countable) collection of cylinder sets starting at time t_0 is called \mathcal{C}_{t_0} , while $\mathcal{C} = \bigcup_{t_0} \mathcal{C}_{t_0}$. The probability of a cylinder set $C_{t_0}(s_0, I_0, s_1, \dots, I_{n-1}, s_n)$ is defined recursively as

$$\mathbb{P}(C_{t_0}(s_0, I_0, \dots, s_n)) = \int_{t_0 \oplus I_0} q_{s_0, s_1}(t) e^{-\Lambda(t_0, t)[s_0]} \mathbb{P}(C_t(s_1, I_1, \dots, s_n)) dt,$$

where $\Lambda(t_0, t)[s] = \int_{t_0}^t -q_{s, s}(\tau) d\tau$ is the cumulative exit rate of state s from time t_0 to time t .

4. Fluid Model Checking

In this paper, we are mainly concerned with the following verification problem:

Given a population model $\mathcal{X}^{(N)}$ and a fixed subset of k agents, consider the process $Z_k^{(N)}(t)$ on the state space $\mathcal{S} = S^k$. We want to check whether $Z_k^{(N)}$ satisfies a time-bounded CSL property φ in a state $s \in \mathcal{S}$ at time t_0 .

Considering the joint process $(Z_k^{(N)}(t), \hat{\mathbf{X}}^{(N)}(t))$, which is a classic time-homogeneous CTMC with state space $\mathcal{S} \times \mathcal{D}$,⁵ we can try to solve the

⁵Of course, instead of $\mathbf{X}^{(N)}(t)$ we can consider variables counting the state of the remaining $N - k$ processes, but for large N and small k , the gain in terms of number of states is limited.

above mentioned verification problem with classic stochastic model checking tools. More specifically, suppose we wish to compute the probability of an until or a next path formula ψ . We can proceed by computing the probability $P_\psi(s, \mathbf{x}) = \mathbb{P}\{\sigma \models \psi \mid \sigma[0] = (s, \mathbf{x})\}$ for each state (s, \mathbf{x}) of $(Z_k^{(N)}, \hat{\mathbf{X}}^{(N)})$. As $(Z_k^{(N)}, \hat{\mathbf{X}}^{(N)})$ is a time-homogeneous CTMC, this probability is independent of the initial time, and can be computed with standard means [13]. Now, fix a state $s \in \mathcal{S}$ of $Z_k^{(N)}$, and consider the probability $P_{s,\mathbf{x}}(t \mid s) = \mathbb{P}\{(Z_k^{(N)}, \hat{\mathbf{X}}^{(N)})(t) = (s, \mathbf{x}) \mid Z_k^{(N)}(t) = s\}$ of being in (s, \mathbf{x}) at time t , conditional on being in s , i.e.

$$P_{s,\mathbf{x}}(t \mid s) = \mathbb{P}\{(Z_k^{(N)}, \hat{\mathbf{X}}^{(N)})(t) = (s, \mathbf{x})\} / \sum_{\mathbf{x}} \mathbb{P}\{(Z_k^{(N)}, \hat{\mathbf{X}}^{(N)})(t) = (s, \mathbf{x})\},$$

defined when the denominator is non-zero. Then, this is the initial distribution of $(Z_k^{(N)}, \hat{\mathbf{X}}^{(N)})$ that we have to take into account when computing the path probability $P_\psi(s, t_0)$ for $Z_k^{(N)}$, starting at time t_0 in state s . By marginalising $\hat{\mathbf{X}}^{(N)}$ ($\hat{\mathcal{D}}$ is the state space of $\hat{\mathbf{X}}$ in the equation below), it follows that

$$P_\psi(s, t_0) = \sum_{\mathbf{x} \in \hat{\mathcal{D}}} P_{s,\mathbf{x}}(t_0 \mid s) P_\psi(s, \mathbf{x}), \quad (5)$$

which depends on t_0 via $P_{s,\mathbf{x}}(t_0 \mid s)$. This has the consequence that, unlike for standard model checking questions in time-homogeneous CTMCs, the truth of $P_{\bowtie p}(\psi)$, $p \in [0, 1]$, for $Z_k^{(N)}$ depends on the initial time t at which the formula is evaluated. This immediately brings us to a choice, resulting in two different model checking algorithms. The first possibility is to apply standard model checking tools to the process $(Z_k^{(N)}, \hat{\mathbf{X}}^{(N)})$ up to the top path formulae (those not being subformulae of another path formula), and then apply the marginalisation of equation (5) to them for any initial time of interest. We will refer to this scheme as the $(Z_k^{(N)}, \hat{\mathbf{X}}^{(N)})$ model checking. Alternatively, we can marginalise the path probabilities for all path subformulae. This latter choice, however, changes the nature of the model checking problem for $(Z_k^{(N)}, \hat{\mathbf{X}}^{(N)})$, forcing us to work with a proper time dependent satisfaction relation while computing path probabilities of formulae containing nested path subformulae. This means that, in a formula like $\psi = \varphi_1 \mathbf{U}^{[T_1, t_2]} \varphi_2$, the states that satisfy φ_1 or φ_2 may depend on the time t at which such formulae are evaluated. This second approach will be called the $Z_k^{(N)}$ model checking. The details of the effect of having to work with time-dependent satisfaction relations are discussed in the next section.

It is clear that when the state space \mathcal{D} is large, as is often the case even for a moderate population size N , both approaches sketched above are cursed by state space explosion. However, for large populations, the fast simulation theorem (Theorem 3.2) tells us that $Z_k^{(N)}(t)$ and $z_k(t)$ are essentially indistinguishable for any finite time horizon. The idea of *fluid model checking* is simply that of approximating the model checking problem for $Z_k^{(N)}(t)$ by replacing $Z_k^{(N)}(t)$ with $z_k(t)$.

The advantage of this approach is that now we need to compute path probabilities on the state space \mathcal{S} , which is usually orders of magnitude smaller than $\mathcal{S} \times \mathcal{D}$. The disadvantage is that $z_k(t)$ is a time-inhomogeneous CTMC, and no CSL model checking algorithm is available off-the-shelf for this class of models. Indeed, the model checking problem for ICTMC is more difficult, especially when considering CSL formulae with nested path operators. In fact, due to time-inhomogeneity, the satisfaction probability of a path formula will depend on the initial time t_0 at which it is evaluated, as for $Z_k^{(N)}(t)$. When comparing this probability with a bound p , we obtain that the truth of a CSL formula in a given state can depend on the time at which it is evaluated. This time dependency creates problems when computing reachability probabilities for nested formulae, as the set to be reached can change with time. A CSL model checking algorithm for ICTMC will be discussed in Section 5.

Once we are able to check a CSL property in the limit model $z_k(t)$, the question becomes that of understanding how the computed results relate to the CSL model checking problem for $Z_k^{(N)}(t)$. More specifically, we would like to know whether the fact that a property φ is satisfied by $z_k(t)$ in state s at time t_0 tells us something about the satisfaction of the same property φ by $Z_k^{(N)}(t)$. A positive answer to this question will be provided in Section 6, where we will prove that the satisfaction will be asymptotically the same, i.e. that the fluid approximation will provide the correct answer for a sufficiently large population size. Practical considerations and experimental results will also be addressed.

Finally, Section 7 compares CSL model checking for $(Z_k^{(N)}(t), \mathbf{X}^{(N)}(t))$, for $Z_k^{(N)}(t)$, and for $z_k(t)$, showing again asymptotic consistency.

5. Model Checking ICTMC

In this section, we present the algorithmic procedures that underlie the CSL model checking algorithm for a generic ICTMC $Z(t)$ on state space \mathcal{S} .

We will build up to this incrementally, focusing first on next state probabilities in Section 5.1. In particular, we will show how to compute the probability that the next state to which the chain jumps belongs to a given set of goal states $G \subseteq \mathcal{S}$, constraining the jump to happen within time $[t_0 + T_1, t_0 + T_2]$, where t_0 is the current time. This is clearly at the basis of the computation of the probability of next path formulae. More specifically, we provide algorithms for ICTMC, focussing on two versions of the next state probability: the case in which the set G is constant, and the case in which the set G depends on time (i.e. a state may belong to G depending on time t). In this latter case, we will talk about *time-varying* or *time-dependent* sets, and we formally identify them with their (time-dependent) indicator function $G : \mathbb{R}_{\geq 0} \times \mathcal{S} \rightarrow \{0, 1\}$, required to be *measurable* with respect to the standard Borel sigma-algebra of $\mathbb{R}_{\geq 0}$. To ensure proper computability of these probabilities, we will mainly work with time-varying sets enjoying the *finite variability property*: $G(t, s)$ can change value only a finite number of time instants in any finite interval $[t_0, t_1]$. We will see how to enforce this condition for the time-varying sets generated during the model checking of a CSL formula by imposing an appropriate restriction on rate functions, namely that they should be piecewise analytic, discussed in Section 5.2.

In Section 5.3, we will focus on the computation of reachability probabilities. Essentially, we want to compute the probability of the set of traces reaching some goal state $G \subseteq \mathcal{S}$ within T units of time, starting at time t_0 and avoiding unsafe states in $U \subseteq \mathcal{S}$. Similarly to Section 5.1, we will consider two versions of the reachability problem: one for constant goal and unsafe sets, and one in which G and U depend on time.

Finally in Section 5.4, we will combine the previous methods into a CSL model checking algorithm for ICTMC. As already hinted, the major consequence of the time-inhomogeneity is that the truth value of φ in a state s depends on the time t at which we evaluate such a formula. In particular, φ may be true in state s at time t_1 , but false at a different time t_2 . Consequently, the set of states that satisfy a CSL formula φ can be time dependent, introducing an additional layer of complexity to the analysis. Indeed, this requires the computation of next-state and reachability probabilities for time-varying sets.

The problem of model checking CSL formulae on time-inhomogeneous CTMC is difficult, and to the authors' knowledge, there is no general method in the literature. An exception is [28], in which the authors show a model checking algorithm for the Hennessy-Milner logics, under the hypothesis of

piecewise constant rates. The method we put forward can cope with the difficulties, but in general may require a large computational effort (for until formulae, the systems of ODE to be solved is quadratic in the size of the state space of the ICTMC, and it depends on the number of discontinuity points of the sets U and G). However, we are interested in using this algorithm to check properties of the fluid approximation z_k , which is an abstract and approximate model of the behaviour of a single or few agents. Usually, a single agent has a very small state space; hence the given approaches to compute next-state probability and reachability of time-varying sets are feasible in practice for the application we have in mind.

5.1. Next-state Probability

We will start by defining the probabilities we want to compute.

Definition 5.1. Let $Z(t)$ be a ICTMC with state space \mathcal{S} and infinitesimal generator matrix $Q(t)$.

1. Let $G \subseteq \mathcal{S}$. The *constant-set next state probability* $P_{next}(Z, t_0, T_1, T_2, G)[s]$ is the probability of the set $Paths_{next}(Z, s, t_0, T_1, T_2, G)$ of trajectories of Z whose first jump happens within time $[t_0 + T_1, t_0 + T_2]$ and ends in a state in G , starting at time t_0 in state s . $P_{next}(Z, t_0, T_1, T_2, G)$ is the next state probability vector on \mathcal{S} .
2. Let $G : [t_0, t_1] \times \mathcal{S} \rightarrow \{0, 1\}$ be a time-dependent set of finite variability (identified with its indicator function, i.e. $G(t)$ is the goal set at time t). The *time-varying-set next state probability* $P_{next}(Z, t_0, T_1, T_2, G(t))[s]$ is the probability of the set $Paths_{next}(Z, s, t_0, T_1, T_2, G)$ of trajectories of Z whose first jump happens within time $[t_0 + T_1, t_0 + T_2]$ and ends in a state in G , starting at time t_0 in state s .

In order for the previous definition to make sense, we need to show that the sets of trajectories $Paths_{next}(Z, s, t_0, T_1, T_2, G)$ are measurable according to the sigma-algebra \mathcal{F} on paths defined for ICTMC (see Section 3.5). The proof of the following proposition involves simple measure-theoretic arguments and is reported in Appendix A.

Proposition 5.1. *Let $G : [t_0, t_1] \times \mathcal{S} \rightarrow \{0, 1\}$ be a time-dependent set and $Z(t)$ an ICTMC. Then $Paths_{next}(Z, s, t_0, T_1, T_2, G)$ is measurable. ■*

Consider a generic ICTMC $Z(t)$, and focus for the moment on a constant set $G \subseteq \mathcal{S}$. For any fixed t_0 , the probability $P_{next}(Z, t_0, T_1, T_2, G)[s]$ that $Z(t)$'s next jump happens at time $t \in [t_0 + T_1, t_0 + T_2]$ and ends in a state of G , given that $Z(t)$ is in state s at time t_0 , is given by the following integral [40, 28]

$$P_{next}(Z, t_0, T_1, T_2, G)[s] = \int_{t_0+T_1}^{t_0+T_2} q_{s,G}(t) \cdot e^{-\Lambda(t_0,t)[s]} dt, \quad (6)$$

where $\Lambda(t_0, t)[s] = \int_{t_0}^t -q_{s,s}(\tau) d\tau$, is the cumulative exit rate of state s from time t_0 to time t , and $q_{s,G}(t) = \sum_{s' \in G} q_{s,s'}(t)$ is the rate of jumping from s to a state $s' \in G$ at time t .

Equation 6 holds for the following reason. Let A_t be the event that the first jump is into a G -state at some time $\tau \in [t_0, t]$. Then $A_{t_1} \subseteq A_{t_2}$ for $t_1 \leq t_2$, and $\mathbb{P}\{A_t\} = \int_{t_0}^t q_{s,G}(\tau) \cdot e^{-\Lambda(t_0,\tau)[s]} d\tau$. We are interested in the probability of the event $A = A_{t_0+T_2} \setminus A_{t_0+T_1}$, which has probability

$$\mathbb{P}\{A\} = \mathbb{P}\{A_{t_0+T_2}\} - \mathbb{P}\{A_{t_0+T_1}\} = \int_{t_0+T_1}^{t_0+T_2} q_{s,G}(t) \cdot e^{-\Lambda(t_0,t)[s]} dt.$$

In order to compute $P_{next}(Z, t_0, T_1, T_2, G)[s]$ for a given t_0 , we can numerically compute the integral, or differentiate both sides obtaining an ODE, and integrate it with standard numerical methods. This simplifies the treatment of the nested integral $\Lambda(t_0, t)[s]$ involved in the computation of P_{next} . More specifically, we can introduce two functions, P and L , initialise $P(t_0+T_1) = 0$ and $L(t_0+T_1) = \Lambda(t_0, t_0+T_1)$, and then integrate the following two ODEs from time t_0+T_1 to time t_0+T_2 :

$$\begin{cases} \frac{d}{dt} P(t) = q_{s,G}(t) \cdot e^{-L(t)} \\ \frac{d}{dt} L(t) = -q_{s,s}(t) \end{cases} \quad (7)$$

However, for CSL model checking purposes, we need to compute $P_{next}(Z, t_0, T_1, T_2, G)[s]$ as a function of t_0 : $\bar{P}_s(t_0) = P_{next}(Z, t_0, T_1, T_2, G)[s]$. One way of doing this is to compute the integral (6) for any t_0 . A better approach is to use the differential formulation of the problem, and define a

```

function NEXT-STATE-PROBABILITY( $Z, G, T_1, T_2, t_0, t_1$ )
  for all  $s \in \mathcal{S}$  do
    Compute  $\bar{P}_s(t_0)$  by solving ODE (7)
    Compute  $\bar{P}_s(t)$  for  $t \in [t_0, t_1]$  by solving ODE (8)
  end for
  return  $\bar{P}(t), t \in [t_0, t_1]$ 
end function

```

Figure 3: Algorithm for the computation of next-state probability $\bar{P}(t)$, for any state s and $t \in [t_0, t_1]$. Other input parameters are as in the text.

set of ODEs with the initial time t_0 as independent variable. First, observe that the derivative of $\bar{P}_s(t_0)$ with respect to t_0 is

$$\begin{aligned}
\frac{d}{dt_0} \bar{P}_s(t_0) &= q_{s,G}(t_0 + T_2) \cdot e^{-\Lambda(t_0, t_0 + T_2)} - q_{s,G}(t_0 + T_1) \cdot e^{-\Lambda(t_0, t_0 + T_1)} \\
&\quad + \int_{t_0 + T_1}^{t_0 + T_2} \frac{\partial}{\partial t_0} q_{s,G}(t) \cdot e^{-\Lambda(t_0, t)} dt \\
&= q_{s,G}(t_0 + T_2) \cdot e^{-\Lambda(t_0, t_0 + T_2)} - q_{s,G}(t_0 + T_1) \cdot e^{-\Lambda(t_0, t_0 + T_1)} \\
&\quad - q_{s,s}(t_0) \bar{P}_s(t_0)
\end{aligned}$$

Consequently, we can compute the next-state probability as a function of t_0 by solving the following set of ODEs:

$$\left\{ \begin{array}{l} \frac{d}{dt} \bar{P}_s(t) = q_{s,G}(t + T_2) \cdot e^{-L_2(t)} - q_{s,G}(t + T_1) \cdot e^{-L_1(t)} - q_{s,s}(t) \bar{P}_s(t) \\ \frac{d}{dt} L_1(t) = -q_{s,s}(t) + q_{s,s}(t + T_1) \\ \frac{d}{dt} L_2(t) = -q_{s,s}(t) + q_{s,s}(t + T_2) \end{array} \right. \quad (8)$$

where $L_1(t) = \Lambda(t, t + T_1)$ and $L_2(t) = \Lambda(t, t + T_2)$.

Initial conditions are $P_s(t_0) = P_{next}(Z, t_0, T_1, T_2, G)[s]$, $L_1(t_0) = \Lambda(t_0, t_0 + T_1)$, and $L_2(t_0) = \Lambda(t_0, t_0 + T_2)$, and are computed solving the equations (7). The algorithm is sketched in Figure 3.

We turn now to discuss the case of a time-varying next-state set $G(t)$. In this case, the only difference with respect to the constant-set case is that the function $q_{\cdot,G(t)}$ is piecewise continuous, rather than continuous. In fact,

each time a state s' gains or loses membership of $G(t)$, the range of the sum defining $q_{\cdot, G(t)}$ changes, and a discontinuity can be introduced. However, as long as these discontinuities constitute a set of measure zero (for instance, they are finite in number), this is not a problem: the integral (6) is defined and absolutely continuous, and so is the solution of the set of ODEs (8) (because the functions involved are discontinuous with respect to time). It follows that the method for computing the next-state probability for constant sets works also for time-varying sets.

Now, if we want to use this algorithm in a model checking routine, we need to be able also to solve the equation $\bar{P}_s(t) = p$, for $p \in [0, 1]$ and each $s \in \mathcal{S}$. In particular, for obvious computability reasons, we want the number of solutions to this equation to be finite. This is unfortunately not true in general, as even a smooth function can be equal to zero on an uncountable and nowhere dense set of Lebesgue measure 0 (for instance, on the Cantor set [41]).

Consequently, we have to introduce some restrictions on the class of functions that we can use. In particular, we will require that the rate functions of $Z(t)$ are *piecewise real analytic functions*.

5.2. Piecewise Real Analytic Functions

A function $f : I \rightarrow \mathbb{R}$, I an open subset of \mathbb{R} , is said to be analytic [42] in I if and only if for each point t_0 of I there is an open neighbourhood of I in which f coincides with its Taylor series expansion around t_0 . Hence, f is locally a power series. For a piecewise analytic function, we intend a function from $I \rightarrow \mathbb{R}$, I an interval, such that there exists I_1, \dots, I_k disjoint open intervals, with $I = \bigcup_j \bar{I}_j$, such that f is analytic in each I_j . A similar definition holds for functions from \mathbb{R}^n to \mathbb{R} , considering their multi-dimensional Taylor expansion.

Analytic functions are a class of functions closed under addition, product, composition, division (for non-zero analytic functions), differentiation and integration. Piecewise analytic functions also satisfy these closure properties, by considering the intersections of their analytic sub-domains. Many functions are analytic: polynomials, the exponential, logarithm, sine, cosine. Using the previous closure properties, one can show that most of the functions we work with in practice are analytic.

Analytic functions have two additional properties that make them particularly suitable in this context:

1. The zeros of an analytic function f in I , different from the constant function zero, are isolated.⁶ In particular, if I is bounded, then the number of zeros is finite. This is true also for the derivatives of any order of the function f .
2. If f is analytic in a set E , then the solution $\mathbf{x}(t)$ of $d\mathbf{x}/dt = f(\mathbf{x})$ in E is also analytic (this is a consequence of the Cauchy-Kowalevski theorem [43]).

This second property, in particular, guarantees that if the rate functions of $Z(t)$ are piecewise analytic, then all the probability functions computed solving the differential equations, like those introduced above, are also piecewise analytic.

Example. If we consider our running example, then it is easy to check that the rate functions defining the infinitesimal generator matrices of interest are piecewise analytic. In fact, even if the vector field of the fluid ODE is not analytic, due to the minimum function, the two functions $g_1(\mathbf{x})$ and $g_2(\mathbf{x})$ of which we take the minimum are analytic. Piecewise analyticity follows from the fact that the solutions of the associated ODE cross the surface $g_1(\mathbf{x}) - g_2(\mathbf{x}) = 0$ (where the minimum is not analytic) only a finite number of times.

5.3. Reachability

In this section, we also consider two versions of the reachability problem: one for constant goal and unsafe sets, and one in which G and U depend on time. We will start by defining these problems for a generic ICTMC $Z(t)$ on state space \mathcal{S} :

Definition 5.2. Let $Z(t)$ be an ICTMC with state space \mathcal{S} and infinitesimal generator matrix $Q(t)$.

1. Let $U, G \subseteq \mathcal{S}$. The *constant-set reachability* $P_{reach}(Z, t_0, T, G, U)[s]$ is the probability of the set $Paths_{reach}(Z, s, t_0, T, G, U)$ of trajectories of Z reaching a state in G without passing through a state in U , within T time units, starting at time t_0 in state s . $P_{reach}(Z, t_0, T, G, U)$ is the reachability probability vector on \mathcal{S} .

⁶An isolated zero of a function $f : I \rightarrow \mathbb{R}$ is a point $z_0 \in I$ such that $f(z_0) = 0$ and there is a neighbourhood W of z_0 in I with $f(x) \neq 0$ for all $x \in W \setminus \{z_0\}$.

2. Let $U, G : [t_0, t_1] \times \mathcal{S} \rightarrow \{0, 1\}$ be time-dependent sets of finite variability. The *time-varying-set reachability* $P_{reach}(Z, t_0, T, G(t), U(t))[s]$ is the probability of the set $Paths_{reach}(Z, s, t_0, T, G(t), U(t))$ of trajectories of Z reaching a state in $G(t)$ at time $t \in [t_0, t_0 + T]$ without passing through a state in $U(t')$, for $t' \in [t_0, t]$, starting at time t_0 in state s .

The previous definition requires the measurability of the sets of paths considered, formally proved in Appendix A.

Proposition 5.2. *Let $G, U : [t_0, t_1] \times \mathcal{S} \rightarrow \{0, 1\}$ be time-dependent set of finite-variability and $Z(t)$ an ICTMC. Then $Paths_{reach}(Z, s, t_0, T, G, U)$, $s \in \mathcal{S}$, is measurable. ■*

5.3.1. Constant-set reachability

We now focus on constant-set reachability, according to Definition 5.2. As previously, let $Z(t)$ be an ICTMC on \mathcal{S} , with rate matrix $Q(t)$ and initial state $Z(0) = Z_0 \in \mathcal{S}$. We will solve the reachability problem in a standard way, by reducing it to the computation of transient probabilities in a modified ICTMC [13]. The solution is similar to the one proposed in [29].

Recall that $\Pi(t_1, t_2)$ is the probability matrix of $Z(t)$, in which entry $\pi_{s_1, s_2}(t_1, t_2)$ gives the probability of being in state s_2 at time t_2 , given that we were in state s_1 at time t_1 and the time evolution of $\Pi(t_1, t_2)$ is described by the Kolmogorov forward and backward equations. Specifically, the forward equation is $\frac{\partial \Pi(t_1, t_2)}{\partial t_2} = \Pi(t_1, t_2)Q(t_2)$, while the backward equation is $\frac{\partial \Pi(t_1, t_2)}{\partial t_1} = -Q(t_1)\Pi(t_1, t_2)$.

The constant-set reachability problem, for a given initial time t_0 , can be solved by integration of the forward Kolmogorov equation (with initial value given by the identity matrix) in the modified ICTMC $Z'(t)$, with infinitesimal generator matrix $Q'(t)$, in which all unsafe states and goal states are made absorbing [13] (i.e. $q'_{s_1, s_2}(t) = 0$, for each $s_1 \in G \cup U$). In particular, $P_{reach}(Z, t_0, T, G, U) = \Pi'(t_0, t_0 + T)\mathbf{e}_G$, where \mathbf{e}_G is an $n \times 1$ vector equal to 1 if $s \in G$ and 0 otherwise, and Π' is the probability matrix of the modified ICTMC Z' .⁷ We emphasise that, in order for the initial value problem

⁷Clearly, alternative ways of computing the transient probability, like uniformization for ICTMC [44], could also be used. However, we stick to the ODE formulation in order to deal with the dependency on the initial time t_0 .

defined by the Kolmogorov forward equation to be well posed, the infinitesimal generator matrix $Q(t)$ has to be sufficiently regular (e.g. bounded and integrable).

As already remarked, in contrast with time-homogeneous CTMC, the reachability probability for ICTMC can depend on the initial time t_0 at which we start the process. Consider now the problem of computing $P(t) = P_{reach}(Z, t, T, G, U)$ as a function of $t \in [t_0, t_1]$. To this end, we can solve the forward equation for t_0 and then use the chain rule to define a differential equation for $\Pi(t, t+T)$, solving it using $\Pi(t_0, t_0+T)$ as the initial condition, i.e.

$$\begin{aligned} \frac{d}{dt}\Pi(t, t+T) &= \frac{\partial}{\partial t}\Pi(t, t+T) + \frac{\partial}{\partial(t+T)}\Pi(t, t+T)\frac{d}{dt}(t+T) \\ &= -Q(t)\Pi(t, t+T) + \Pi(t, t+T)Q(t+T). \end{aligned} \quad (9)$$

Using a numerical solver for the ODE, this gives an effective algorithm (Figure 4) to compute the probability of interest (for any fixed error bound). Furthermore, if we can guarantee that the number of zeros of the equation $P(t) - p$ is finite, then we also have an effective procedure to compute the truth value of $P(t) \bowtie p$, as a function of time⁸, for $\bowtie \in \{<, \leq, \geq, >\}$ (provided we can find those zeros, as will be discussed in Section 5.4).

Example. We consider the time-inhomogeneous CTMC obtained by constructing the fluid limit of a single client in the client-server example of Section 3.1. We consider two reachability probabilities:

1. The probability of observing a time-out before being served for the first time within time T . This is a reachability problem with goal set $G = \{rc\}$ and unsafe set $U = \{rq, w\}$.
2. The probability of observing a timeout within time T . This is a reachability problem with goal set $G = \{rc\}$ and unsafe set $U = \emptyset$.⁹

In Figures 8(a) and 9(a), we can find these reachability probabilities computed with the method just presented, as a function of the time horizon T . In

⁸In fact, we can compute with arbitrary precision the times at which the truth value changes in any state. Therefore, we can compute a truth value function which will be arbitrarily close to the “true” one, in a functional sense, i.e. according to a metric synchronising nearby jumps, similar to the Skorokhod one [45].

⁹In fact, this is a first passage time problem.

function REACHABILITY-CONSTANT-SET(Z, T, G, U, t_0, t_1)

Construct the CTMC in which G and U states are absorbing, with rate matrix $Q'(t)$.

Compute $\Pi'(t_0, t_0 + T)$ by solving the forward Kolmogorov ODE for the modified CTMC.

Compute $\Pi'(t, t+T)$ for $t \in [t_0, t_1]$ by solving ODE (9) for the modified CTMC with initial conditions $\Pi'(t_0, t_0 + T)$.

return $P(t) = \Pi'(t_0, t_0 + T)\mathbf{e}_G, t \in [t_0, t_1]$.

end function

Figure 4: Algorithm for the computation of reachability probability $P(t)$ for $t \in [t_0, t_1]$ and constant goal and unsafe sets G and U . Other input parameters are as in the text.

Figures 8(c) and 9(c), instead, we plotted the reachability probability for both problems 1 and 2 for $T = 50$ as a function of the initial time $t_0 \in [0, 25]$ (blue solid line).

5.3.2. Time-varying set reachability

Now we turn our attention to the reachability problem for time-varying sets. The main difficulty in this case is that, at each time T_i in which the goal or the unsafe set changes, also the modified Markov chain that we need to consider to compute the reachability probability changes structure. This can have the effect of introducing a discontinuity in the probability matrix.

In particular, if at time T_i a state s becomes a goal state, then the probability $\pi_{s_1, s}(t, T_i)$ suddenly needs to be added to the reachability probability from state s_1 . Therefore, a change in the goal set at time T_i introduces a discontinuity in the reachability probability at time T_i . Similarly, if a state s was safe and then becomes unsafe, we have to discard the probability of trajectories that are in that state at time T_i , as those trajectories become suddenly unsafe.

As previously, let $G(t)$ and $U(t)$ be the goal and unsafe sets, satisfying the finite variability property, i.e. such that the set of time points at which G or U change value (at least in one state) is finite and equal to $T_1 \leq T_2 \dots \leq T_k$. This can be enforced by requiring that rate functions are piecewise analytic. Let $T_0 = t$ and $T_{k+1} = t + T$.

In order to compute the reachability probability, we can exploit the semi-group property of the Markov process, stating that $\Pi(T_0, T_{k+1}) = \prod_{i=0}^k \Pi(T_i, T_{i+1})$. Then, we also need to deal appropriately with the disconti-

nity effects at each time T_i , mentioned above. We proceed in the following way:

1. We double the state space, letting $\tilde{\mathcal{S}} = \mathcal{S} \cup \bar{\mathcal{S}}$, where a state $\bar{s} \in \bar{\mathcal{S}}$ represents state s when it is a goal state (cf. below for the definition of $\tilde{Q}(t)$). Hence, in the probability matrix $\tilde{\Pi}$, $\tilde{\pi}_{s_1, \bar{s}_2}$ is the probability of having reached s_2 avoiding unsafe states, while s_2 was a goal state.
2. Consider a discontinuity time T_i and let $t_1 \in [T_{i-1}, T_i)$ and $t_2 \in (T_i, T_{i+1}]$. Define $W(t) = S \setminus (G(t) \cup U(t))$. Then, for $s_1 \in W(t_1)$ and $s_2 \in W(t_2)$, the probability of being in s_2 at time t_2 , given that we were in s_1 at time t_1 and avoiding both unsafe and goal sets, can be written as $\tilde{\pi}_{s_1, s_2}(t_1, t_2) = \sum_{s \in W(t_1) \cap W(t_2)} \tilde{\pi}_{s_1, s}(t_1, T_i) \tilde{\pi}_{s, s_2}(T_i, t_2)$. Hence, we have to appropriately restrict the summation set at time T_i , to account for changes in W .
3. Consider again a discontinuity time T_i and let $t_1 \in [T_{i-1}, T_i)$ and $t_2 \in (T_i, T_{i+1}]$. Suppose $s_2 \in W(t_1)$ and $s_2 \in G(t_2)$. Then, the probability of reaching the goal state s_2 at time t_2 , given that at time t_1 we were in s_1 , can be written as

$$\tilde{\pi}_{s_1, s_2}(t_1, T_i) + \sum_{s \in W(t_1) \cap W(t_2)} \tilde{\pi}_{s_1, s}(t_1, T_i) \tilde{\pi}_{s, \bar{s}_2}(T_i, t_2).$$

The first term is needed because all safe trajectories that are in state s_2 at time T_i suddenly become trajectories satisfying the reachability problem, hence we have to add them to compute the reachability probability. The second term computes the probability of remaining in a safe path from time t_1 to T_i , being at time T_i in a state that remains safe even after the discontinuous change at time T_i , and then reaching s_2 via a safe path during time $[T_i, t_2]$. As the goal state s_2 cannot be reached in other ways during time $[t_1, t_2]$, the expression above computes the probability correctly.

All the previous remarks can be formally incorporated into the semi-group expansion of $\tilde{\Pi}(t, t + T)$ by multiplying on the right each term $\tilde{\Pi}(T_i, T_{i+1})$ by a suitable 0/1 matrix, depending only on the structural changes at time T_{i+1} . Let $|\mathcal{S}| = n$ and let $\zeta_W(T_i)$ be the $n \times n$ matrix equal to 1 only on the diagonal elements corresponding to states s_j belonging to both $W(T_i^-)$ and $W(T_i^+)$ (i.e. states that are safe and not goals both before and after

T_i), and equal to 0 elsewhere. Furthermore, let $\zeta_G(T_i)$ be the $n \times n$ matrix equal to 1 in the diagonal elements corresponding to states s_j belonging to $W(T_i^-) \cap G(T_i^+)$, and zero elsewhere. Finally, let $\zeta(T_i)$ be the $2n \times 2n$ matrix defined by:

$$\zeta(T_i) = \begin{pmatrix} \zeta_W(T_i) & \zeta_G(T_i) \\ 0 & I \end{pmatrix}.$$

Consider now the following ICTMC \tilde{Z} on $\tilde{\mathcal{S}}$, with rate matrix $\tilde{Q}(t)$, where

1. for $\bar{s}_1 \in \tilde{\mathcal{S}}$ and any $s_2 \in \tilde{\mathcal{S}}$, $\tilde{q}_{\bar{s}_1, s_2}(t) = 0$;
2. for $s_1 \notin W(t)$ and all $s_2 \in \tilde{\mathcal{S}}$, $\tilde{q}_{s_1, s_2}(t) = 0$
3. for $s_1 \in W(t)$ and $s_2 \in S \setminus G(t)$, $\tilde{q}_{s_1, s_2}(t) = q_{s_1, s_2}(t)$, while $\tilde{q}_{s_1, \bar{s}_2}(t) = 0$;
4. for $s_1 \in W(t)$ and $s_2 \in G(t)$, $\tilde{q}_{s_1, \bar{s}_2}(t) = q_{s_1, s_2}(t)$, while $\tilde{q}_{s_1, s_2}(t) = 0$.

In the previous chain, all unsafe and goal states are absorbing, while transitions leading from a safe state s to a goal state are readdressed to the copy \bar{s} of s . States in $\tilde{\mathcal{S}}$ are absorbing, too.

Now let $\tilde{\Pi}(t_1, t_2)$ be the probability matrix associated with the ICTMC $\tilde{Q}(t)$. Given the interval $I = [t, t + T]$, we indicate with T_1, \dots, T_{k_I} the ordered sequence of discontinuity points of goal and unsafe sets internal to I . Let

$$\Upsilon(t, t + T) = \tilde{\Pi}(t, T_1)\zeta(T_1)\tilde{\Pi}(T_1, T_2)\zeta(T_2) \cdots \zeta(T_{k_I})\tilde{\Pi}(T_{k_I}, t + T). \quad (10)$$

Then, we have that

$$P_s(t) = P_{reach}(Z, t, T, G, U)[s] = \sum_{\bar{s}_1 \in \tilde{\mathcal{S}}} \Upsilon_{s, \bar{s}_1}(t, t + T) + \mathbf{1}\{s \in G(t)\}, \quad (11)$$

where the first term takes into account the probability of reaching a goal state starting from a non-goal state, while the second term is needed to properly account for states $s \in G(t)$, for which $P_s(t)$ has to be equal to 1 (a formal proof can be given by induction on the number of discontinuity points). $\Upsilon(t, t + T)$ can be obtained by computing each $\tilde{\Pi}(T_i, T_{i+1})$ solving the associated forward Kolmogorov equation and then multiplying those matrices and the appropriate ζ ones, according to the definition of Υ .

If we want to compute $P(t)$ as a function of t , instead, we need a way to compute $\Upsilon(t, t + T)$ as a function of t . This can be done by observing that Υ depends on t only from the first and last factors in the multiplication. Defining $\Gamma(T_1, T_k) = \zeta(T_1)\tilde{\Pi}(T_1, T_2)\zeta(T_2) \cdots \tilde{\Pi}(T_{k-1}, T_k)\zeta(T_k)$, writing

$\Upsilon(t, t + T) = \tilde{\Pi}(t, T_1)\Gamma(T_1, T_k)\tilde{\Pi}(T_k, t + T)$, differentiating with respect to t and applying the forward or backward equation for $\tilde{\Pi}$, we find the following differential equation for Υ :

$$\frac{d\Upsilon(t, t + T)}{dt} = -\tilde{Q}(t)\Upsilon(t, t + T) + \Upsilon(t, t + T)\tilde{Q}(t + T). \quad (12)$$

This equation holds until either t or $t + T$ becomes equal to a discontinuity point. When this happens, the integration has to be stopped and restarted, recomputing Υ accordingly.

Practically, to solve this problem we can proceed as follows:

1. Given an interval $[t_0, t_1]$ of interest for the reachability, find all discontinuity points of the sets G and U contained in $[t_0, t_1 + T]$, and let them be $t_0 = T_0 < T_1 < \dots < T_k < T_{k+1} = t_1 + T$. Furthermore, let $T'_i = T_i + T$ for $i = 0, \dots, k$, let $pre(t)$ be the greatest T_j preceding t , and $post(t)$ the smallest T_j following t .
2. Compute $\tilde{\Pi}(T_i, T_{i+1})$ and $\tilde{\Pi}(pre(T'_i), T'_i)$ for $i \leq k$, using the forward Kolmogorov equations¹⁰. Compute also each $\zeta(T_i)$.
3. Compute $\Upsilon(t_0, t_0 + T)$ and integrate until time $t = \min\{T_1, T_{j+1} - T\}$, where $t_0 + T \in [T_j, T_{j+1}]$.
4. If $t + T = T_{j+1}$, multiply Υ on the right by $\zeta(T_{j+1})$ and continue the integration. If $t = T_1$, then recompute Υ as $\tilde{\Pi}(T_1, T_2)\Gamma(T_2, T_j)\tilde{\Pi}(T_j, T_1 + T)$, where $\tilde{\Pi}(T_j, T_1 + T) = \tilde{\Pi}(pre(T'_1), T'_1)$.
5. Integrate piecewise using the previous rules until time t_1 .

A more detailed algorithmic presentation of the procedure is shown in Figure 5. Notice that, if the infinitesimal generator matrix $Q(t)$ of Z is sufficiently well-behaved (for instance, Lipschitz continuous), then the function $P(t)$ will be at least piecewise continuous, with a finite number of discontinuity points at instants T_i and T'_i .

Remark 5.1. The precise behaviour of the G and U functions at their discontinuity points (i.e. if they are left-continuous or right-continuous) is irrelevant for the computation of Υ : the set of trajectories of Z differing in those time points has probability 0.

¹⁰Notice, that, if $T_j = pre(T'_i)$, then $\tilde{\Pi}(T_j, T'_i)$ and $\tilde{\Pi}(T_j, T_{j+1})$ can be computed during the same numerical integration of the forward equation.

function REACHABILITY(Z, T, G, U, t_0, t_1)

Construct the CTMC on the modified state space $\tilde{\mathcal{S}}$, according to the recipe in the text.

Let $t_0 = T_0, T_1, \dots, T_k, T_{k+1} = t_1 + T$ be the time instants at which G or U has a discontinuity.

for $i = 0$ to k **do**

 Compute $\tilde{\Pi}(T_i, T_{i+1})$ and $\tilde{\Pi}(pre(T_i + T), T_i + T)$ using the forward Kolmogorov equations

end for

 Compute $\Upsilon(t_0, t_0 + T)$ according to equation (10) and $P(t_0)$ according to equation (11)

$t \leftarrow t_0$

repeat

$T_a \leftarrow post(t)$

$T_b \leftarrow post(t + T)$

$\bar{t} \leftarrow \min\{T_a, T_b - T\}$

 Compute Υ and P from t to \bar{t} , according to ODE (12) and equation (11), with initial conditions $\Upsilon(t, t + T)$ (previously computed).

if $\bar{t} + T = T_b$ **then**

$\Upsilon(\bar{t}, \bar{t} + T) \leftarrow \Upsilon(\bar{t}, \bar{t} + T)\zeta(T_b)$

else if $\bar{t} = T_a$ **then**

$\Upsilon(\bar{t}, \bar{t} + T) \leftarrow \tilde{\Pi}(T_a, post(T_a))\Gamma(post(T_a), pre(T_b))\tilde{\Pi}(pre(T_b), T_b + T)$

end if

$t \leftarrow \bar{t}$

until $t \geq t_1$

return $[\Upsilon(t, t + T), P(t)], t \in [t_0, t_1]$.

end function

Figure 5: Algorithm for the computation of reachability probability $P(t)$ for $t \in [t_0, t_1]$ and time-varying goal and unsafe sets $G(t)$ and $U(t)$, with a finite number of discontinuities. Other input parameters are as in the text.

Remark 5.2. In the previous method, we need to repeatedly integrate a set $4n^2$ differential equations. However, most of these variables are redundant. In fact, we only need n^2 variables for the probability transition matrix Π on \mathcal{S} and an additional n variables to store the reachability probability vector. The method presented above can be easily reconfigured to this restricted set of variables.

5.4. The CSL Model Checking Algorithm

We now combine the previous numerical routines into a CSL model checking algorithm for ICTMC. Before doing this, we need to guarantee that the set of paths involved in the definition of the semantics of CSL formulae are measurable. This is proved in the following proposition (see Appendix A for technical details) by combining Propositions 5.1 and 5.2 with the fact that the piecewise analytic nature of rates preserves the finite variability property.

Proposition 5.3. *Let $Z(t)$ be a ICTMC with piecewise analytic time-dependent rate matrix $Q(t)$. Then*

1. *The time-dependent set of states $\llbracket \varphi \rrbracket = \llbracket \varphi \rrbracket(t)$ that satisfy a CSL formula φ has the finite variability property.*
2. *The set of paths $\text{Paths}(s, t_0, \psi)$ that satisfy a CSL path formula φ starting in state s at time t_0 is measurable. ■*

Consider now an until CSL formula $\varphi = \mathcal{P}_{\triangleright p}(\varphi_1 \mathbf{U}^{[0, T]} \varphi_2)$, where φ_1 and φ_2 are boolean combinations of atomic propositions. The major consequence of the time-inhomogeneity of $Z(t)$ is that the truth value of φ in a state s depends on the time t at which we evaluate such a formula. In particular, φ may be true in state s at time t_1 , but false at a different time t_2 . Consequently, the set of states that satisfy a CSL formula φ can be time dependent, hence nesting φ into a larger temporal formula requires the computation of next-state and reachability probabilities for time-varying sets. There is a similar problem with next formulae of the form $\varphi = \mathcal{P}_{\triangleright p}(\mathbf{X}^{[T_a, T_b]} \varphi_1)$, as the next-state probability also depends on the evaluation time t .

The computation of next-state probabilities for time-varying target sets can be done by the method presented in Section 5.1, in particular the algorithm in Figure 3.

The algorithm of Section 5.3.2 for computing reachability in the presence of piecewise constant goal and update sets, instead, is the core procedure to compute the probability of an until formula. In fact, consider the path formula $\varphi_1 \mathbf{U}^{[T_a, T_b]} \varphi_2$. To compute its probability for initial time in $[t_0, t_1]$,¹¹ we solve two reachability problems separately and then combine the results.

¹¹The appropriate values of t_0 and t_1 are to be deduced from φ_1 , φ_2 and the superformula of the until, in a standard way [46]

The first reachability problem is for unsafe set $U_1 = \llbracket \neg\varphi_1 \rrbracket$ and empty goal set $G(t + T_a) = \emptyset$. Let $\Upsilon^1(t, t + T_a)$ be the probability matrix of this reachability problem. In order for the computation of the until probability to work, we must then discard the probability of being in an unsafe state, essentially multiplying $\Upsilon^1(t, t + T_a)$ by $\zeta^1(t + T_a)$ on the right (see Section 5.3.2).¹²

The second reachability problem is for unsafe set $U = \llbracket \neg\varphi_1 \rrbracket$ and goal set $G = \llbracket \varphi_2 \rrbracket$, and is solved for initial time $t \in [t_0 + T_a, t_1 + T_a]$, and time horizon $T_b - T_a$. Let $\Upsilon^2(t + T_a, t + T_b)$ be the function computed by the algorithm in Section 5.3.2 for this second problem. Then, for each state s , safe at time t , we compute $P(t) = \Upsilon^1(t, t + T_a)\zeta^1(t + T_a)\Upsilon^2(t + T_a, t + T_b)\mathbf{e}_{\bar{S}}$, where $\mathbf{e}_{\bar{S}}$ is the vector equal to 1 for states $\bar{s} \in \bar{S}$ and zero elsewhere. $P_s(t)$ contains the probability of the until formula in state s . Then, we can determine if state s at time t satisfies $\mathcal{P}_{\bowtie p}(\varphi_1 \mathbf{U}^{[T_a, T_b]} \varphi_2)$ by solving the inequality $P_s(t) \bowtie p$.

This provides an algorithm to approximately solve the CSL model checking for ICTMC recursively on the structure of the formula, provided that the number of discontinuities of sets satisfying a formula is finite and that we are able to find all the zeros of the computed probability functions, to construct the proper time-dependent satisfiability sets (or approximations thereof). The full procedure is sketched in Figure 6.

Example. As an example of the functioning of the CSL model checking for ICTMC, we consider the client-server running example, more specifically the time-inhomogeneous CTMC obtained by constructing the fluid limit of a single client discussed at the end of Section 3.3. Consider the until path formula $trueU^{[0,50]} timeout$, where $timeout$ is true only in state rc . Its probability, as a function of the initial time, is shown in Figures 7(a), 7(b), and 7(c), for the states rq , w , and t , respectively. In the same figures, we also show the time-dependent truth of the CSL formula $\mathcal{P}_{<0.167}(trueU^{[0,50]} timeout)$, which is obtained by solving the inequality $P_s(t) < 0.167$, where $P_s(t)$ is a time-dependent probability function. In this case, we can observe that for time $t_0 \in [0, 100]$, there is only one solution, as the probability is monotone. This depends on the solution of the fluid equations. In this case, in fact, they

¹²In fact, this reachability problem can be solved in a simpler way: it just requires trajectories not to enter an unsafe state, and then collects the probability to be in a safe state at the time $t + T_a$. In particular, we can get rid of the copy \bar{S} of the state space, and define a simplified Υ function using ζ_W matrices instead of ζ ones.

function CSL_MC(Z, φ, t_0, t_1) ▷ Computes $V(t)$, where
 $V_{\uparrow s}(t) = \mathbf{I}\{s, t \models \varphi\}$ for $s \in \mathcal{S}$ and $t \in [t_0, t_1]$.
if $\varphi = p$ **then**
 $V(t)$ is such that $V_{\uparrow s}(t) \leftarrow \mathbf{I}\{p \in L(s)\}$, for $s \in \mathcal{S}$
else if $\varphi = \neg\varphi_1$ **then**
 $V_1 \leftarrow \text{CSL_MC}(Z, \varphi_1, t_0, t_1)$
 $V(t) \leftarrow 1 - V_1(t)$
else if $\varphi = \varphi_1 \wedge \varphi_2$ **then**
 $V_1 \leftarrow \text{CSL_MC}(Z, \varphi_1, t_0, t_1)$
 $V_2 \leftarrow \text{CSL_MC}(Z, \varphi_2, t_0, t_1)$
 $V(t) \leftarrow \min\{V_1(t), V_2(t)\}$
else if $\varphi = \mathcal{P}_{\bowtie p}(\mathbf{X}^{[T_a, T_b]}\varphi_1)$ **then**
 $V_1 \leftarrow \text{CSL_MC}(Z, \varphi_1, t_0, t_1)$
 $\bar{P} \leftarrow \text{NEXT-STATE-PROBABILITY}(Z, V_1, T_a, T_b, t_0, t_1)$
 $V(t) \leftarrow \mathbf{I}\{\bar{P}(t) \bowtie p\}$
else if $\varphi = \mathcal{P}_{\bowtie p}(\varphi_1 \mathbf{U}^{[T_a, T_b]}\varphi_2)$ **then**
 $V_1 \leftarrow \text{CSL_MC}(Z, \neg\varphi_1, t_0, t_1 + T_b)$
 $V_2 \leftarrow \text{CSL_MC}(Z, \varphi_2, t_0, t_1 + T_b)$
 $\Upsilon^1 \leftarrow \text{REACHABILITY}(Z, T_a, \emptyset, V_1, t_0, t_1 + T_a)[1]$ ▷ Returns Υ
component of REACHABILITY.
 $\Upsilon^2 \leftarrow \text{REACHABILITY}(Z, T_b - T_a, V_2, V_1, t_0 + T_a, t_1 + T_b)[1]$
 $P(t) = \Upsilon^1(t, t + T_a)\zeta^1(t + T_a)\Upsilon^2(t + T_a, t + T_b)\mathbf{e}_{\mathcal{S}}$
 $V(t) \leftarrow \mathbf{I}\{P(t) \bowtie p\}$
end if
return V
end function

Figure 6: Core algorithm for solving the CSL model checking problem, by computing the satisfiability of a CSL-formula φ as a function of the time $t \in [t_0, t_1]$ at which it is evaluated. The truth value of φ is then the value it has in t_0 , which is usually 0.

converge to a steady state, hence we do expect that also the time dependent truth value of CSL until formulae stabilises (when the fluid ODE are close to steady state, the rates of the ICTMC are practically constant). This suggests that in many practical cases, the number of changes of truth value of until formulae will be very small, as in the running example. Notice that in the case of the running example, if we had chosen a threshold bigger, say, than 0.25, then the time-dependent truth formulae would have been a constant

function.

In Figure 7(d), instead, we show the probability of the path formula

$$true \mathbf{U}^{[0,T]}(\mathcal{P}_{<0.167}(true U^{[0,50]} \text{ timeout})),$$

as a function of the time horizon T . In the plot, it is evident how this probability has discontinuities at those time instants when the truth value function of its until sub-formula changes. These discontinuities differentiate the model checking of ICTMC from that of time-homogeneous CTMC.

5.5. Decidability of the CSL Model Checking for ICTMC


We will now consider the model checking algorithm of the previous subsection in more detail, focussing particularly on correctness and termination. In this consideration we will make the following assumption about the numerical algorithms that it uses.

Assumption 1. There are interval arithmetic routines that can compute bounding sets for the rate functions of $Z(t)$, in such a way that the approximation error can be made arbitrary small. We call such functions *interval computable*.¹³


Notice that this assumption is not very restrictive. It applies to all the standard functions, and also to solutions of ODEs of functions which satisfy it, to derivatives of these functions and to their integrals [47, 48]. In particular, if the rate functions are interval computable, then so will be all the probabilities computed by solving reachability problems.

The approach presented above relies on, in addition to the solution of ODEs, also two other key numerical operations: given a computable real number p , determine if p is zero and given an analytic function f , find all the zeros of such a function (or better an interval approximation of these zeros of arbitrary accuracy). However, it is not clear if these two operations can be carried out effectively for any input that we can generate (see Remark 5.4 for further comments). Therefore, we need some further assumptions. Instead of restricting the class of functions (which seems a difficult problem since we have to consider the solution of differential equations), we will follow the


¹³More formally, for any interval computable function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ there is an effective algorithm that takes as input a point $\mathbf{x} \in \mathbb{R}^n$ and an $\varepsilon > 0$ and returns an interval $[a, b] \subseteq \mathbb{R}$ such that $b - a < \varepsilon$ and $f(\mathbf{x}) \in [a, b]$.




(a) $true \mathbf{U}^{[0,50]} \text{ timeout} - rq$



(b) $true \mathbf{U}^{[0,50]} \text{ timeout} - w$



(c) $true \mathbf{U}^{[0,50]} \text{ timeout} - t$



(d) $true \mathbf{U}^{[0,T]}(\mathcal{P}_{<0.167}(true \mathbf{U}^{[0,50]} \text{ timeout}))$

Figure 7: Figures 7(a), 7(b), and 7(c). Probability of the formula $true \mathbf{U}^{[0,50]} \text{ timeout}$, for varying initial time, and different initial states (rq , w , and t respectively). The dotted line shows the time varying truth function for the CSL formula $\mathcal{P}_{<0.167}(true \mathbf{U}^{[0,50]} \text{ timeout})$, which is obtained by finding the zeros of the initial-time dependent probability. Figure 7(d). Probability of the until path formula $true \mathbf{U}^{[0,T]}(\mathcal{P}_{<0.167}(true \mathbf{U}^{[0,50]} \text{ timeout}))$, as a function of time bound T . Vertical dotted lines show the discontinuity points of the time-dependent truth of the until sub-formula.

approach of [49], introducing a notion of *robust CSL formula* and proving decidability for this subset of formulae. This will not solve the decidability problem in theory, but makes it quasi-decidable [49], which may be enough in practice. As we will see, the set of CSL formulae which are not robust has measure zero (see Theorem 5.2).

5.5.1. Robust CSL formulae

In order to introduce the concept of robust CSL formula, we first need some ancillary notions, concerned with robustness of time-varying sets. The first definition introduces a robust time-varying set in terms of the notion of piecewise real analytic function. The second definition, instead, guarantees the preservation of robustness with respect to boolean operations on time-varying sets.

Definition 5.3. A time-dependent subset $V(t)$ of \mathcal{S} , $t \in I$, is *robust* if and only if there is a piecewise analytic function $h_V : \mathcal{S} \times I \rightarrow \mathbb{R}$ and an operator $\bowtie \in \{<, \leq, \geq, >\}$, such that for each $s \in \mathcal{S}$, the indicator function $V_{\uparrow s} : I \rightarrow \{0, 1\}$ of s is given by $\mathbf{1}\{h_V(s, t) \bowtie 0\}$, and it further satisfies:

1. the number of discontinuity points $Disc(V) = \{(s, \bar{t}) \mid V_{\uparrow s}(\bar{t}^-) \neq V_{\uparrow s}(\bar{t}^+)\}$ is finite;
2. if h_V is analytic in (s, t) and $h_V(s, t) = 0$, then $\frac{d}{dt}h_V(s, t) \neq 0$ (zeros of h_V are simple);
3. if h_V is *not* analytic in (s, t) , then $h_V(s, t^-) \neq 0$ and $h_V(s, t^+) \neq 0$.¹⁴

In the following, we will usually indicate with $V(t)$ both a time dependent set V and its indicator function (with values in $\{0, 1\}^m$, $m = |\mathcal{S}|$), and use h_V to denote the piecewise analytic function defining it.

Definition 5.4. Two time varying sets V^1 and V^2 are *compatible* if and only if they do not have a discontinuity at the same time for the same state s : $\forall s \in \mathcal{S}, Disc(V_{\uparrow s}^1) \cap Disc(V_{\uparrow s}^2) = \emptyset$.

Consider now a CSL formula φ and let p_1, \dots, p_k be the constants appearing in the $\mathcal{P}_{\bowtie p}$ path quantifiers of next and until sub-formulae of φ . We will treat $\varphi = \varphi(p_1, \dots, p_k)$ as a function of those p_1, \dots, p_k . Furthermore, we will call the next or until sub-formulae of φ *top next sub-formulae* or *top until sub-formulae* if they are not sub-formulae of other next or until formulae. The other next or until formulae will be called *dependent*.

¹⁴This condition states that, if h_V is continuous but not analytic in (s, t) , then it cannot be equal to zero in those points, implying that first order derivatives exist and are non-null in all continuity points in which h_V crosses zero. Moreover, if $h_V(s, t^-) \neq h_V(s, t^+)$, h_V can cross zero in (s, t) only if the jump contains zero, meaning that $\min\{h_V(s, t^-), h_V(s, t^+)\} < 0 < \max\{h_V(s, t^-), h_V(s, t^+)\}$.

Definition 5.5. A CSL formula $\varphi = \varphi(\mathbf{p})$, $\mathbf{p} \in [0, 1]^k$ is *robust* if and only if

1. There is an open neighbourhood W of \mathbf{p} in $[0, 1]^k$ such that for each $\mathbf{p}_1 \in W$,

$$s, 0 \models \varphi(\mathbf{p}) \Leftrightarrow s, 0 \models \varphi(\mathbf{p}_1).$$

2. The time-varying sets of any dependent next or until sub-formula of φ are *robust*.
3. The time-varying sets of sub-formulae of φ that are part of the same until formula or of a conjunction/disjunction are *compatible*.

Example. The CSL formula

$$\mathcal{P}_{>0.2}(\text{true } \mathbf{U}^{[0,T]}(\mathcal{P}_{<0.167}(\text{true } U^{[0,50]} \text{ timeout})))$$

is robust. In fact, as can be easily seen from visual inspection of Figure 7(d), the probability of the top until formula at time zero is equal to 0, which is different from 0.2. Furthermore, the time-varying set of its nested until formula is also robust, being defined by a piecewise analytic function (the time-dependent reachability probability for piecewise analytic rates is piecewise analytic), which crosses the threshold 0.167 only once per state, with non-null derivative (this can also be easily checked by visual inspection of Figure 7).

We now prove the following theorem, which states that the CSL model checking algorithm we put forward works at least for robust formulae:

Theorem 5.1. *The CSL model checking for ICTMC, for piecewise analytic interval computable rate functions, is decidable for a robust CSL formula $\varphi(p_1, \dots, p_k)$.*

Proof Sketch: The crucial point behind the proof is to show that the time-varying set of a robust CSL formula is robust and it can be effectively computed with arbitrary precision. This means that one can encapsulate the location of the finite number of discontinuity points with arbitrary small precision $\varepsilon > 0$, using interval arithmetic routines. This in turns guarantees that error introduced in the computation of the probability $P(0)$ at time zero of φ can be made arbitrarily small. By point 1 in the definition of robustness, this value will be different from the threshold p with which it is compared, meaning that for some $\varepsilon > 0$, the interval $[P(0) - \varepsilon, P(0) + \varepsilon]$ will be disjoint from p , allowing us to effectively determine the truth of φ . The full proof

can be found in Appendix A. ■

The following corollary is a straightforward consequences of the proof of the previous theorem:

Corollary 5.1. *The algorithm for CSL model checking presented in this section is correct for robust CSL formulae.* ■

5.5.2. Quasi-decidability

We turn now to characterise the set of robust formulae from a topological and measure-theoretic point on view. We have the following

Theorem 5.2. *Given a CSL formula $\varphi(\mathbf{p})$, with $\mathbf{p} \in [0, 1]^k$, then the set $\{\mathbf{p} \mid \varphi(\mathbf{p}) \text{ is robust}\}$ is relatively open¹⁵ in $[0, 1]^k$ and has Lebesgue measure 1.*

Proof Sketch: The theorem can be proved by structural induction using technical arguments of measure theory, the most important one being Fubini's Theorem [45]. Intuitively, the result follows because the set of thresholds \mathbf{p} on which robustness fails is (roughly speaking) a finite union of manifolds of topological dimension strictly less than k (hence of measure zero), so that it has measure zero. The properties of piecewise analytic functions play a crucial role in establishing this result. We refer the interested reader to Appendix A for further details. ■

The openness of the set of robust thresholds for a formula allows us to prove the following corollary about quasi-decidability. In this paper, we consider a notion of quasi-decidability, which is slightly different from the one defined in [49]. In fact, we take advantage of the fact that our input values belong to a compact subset $K \subseteq \mathbb{R}^n$, for which a standard notion of measure exists.

Definition 5.6. A problem with inputs in a compact subset $K \subseteq \mathbb{R}^n$ of Lebesgue measure $\mu_\ell(K) > 0$, is *quasi-decidable* if there is an algorithm that solves it correctly for an open subset $U \subset K$, with $\mu_\ell(U)/\mu_\ell(K) = 1$.

Combining Theorems 5.1 and 5.2, we obtain the following:

¹⁵A set $U \subset V$ is relatively open in $V \subset W$, where W is a topological space, if it is open in the subspace topology, i.e. if there exists an open subset $U_1 \subseteq W$ such that $U = V \cap U_1$.

Corollary 5.2. *The CSL model checking for ICTMC, for piecewise analytic interval computable rate functions is quasi-decidable for any formula φ . ■*

Remark 5.3. The notions of robustness and quasi-decidability have a practical side. First, the openness property of the set of robust thresholds for a formula $\varphi(\mathbf{p})$ guarantees that if we “perturb” a formula (by varying the set \mathbf{p} of threshold constants of the path probability operators), then the formula remains robust. Furthermore, by the definition of robustness, also its truth value remains the same (as the notion of quasi-decidability of [49] requires). This explains the use of the terminology “robust”.

Secondly, the characterisation of the set R of robust thresholds for a formula φ provided in Theorem 5.2, implies that if we choose thresholds at “random”, we are likely to select a robust formula. In fact, consider the grid of rational numbers with $\frac{1}{n}$ in $[0, 1]$, i.e. $GR_n = \{\frac{m}{n} \mid m < n, m, n \in \mathbb{N}\}$, and take the Cartesian product $GR_n^k \subset [0, 1]^k$. Let μ_n be the uniform distribution in GR_n^k , then $\mu_n \rightarrow \mu$, the uniform distribution on $[0, 1]^k$ (which coincides with the Lebesgue measure on Borel sets). Now, as R is open and has Lebesgue measure 1, then $\mu(R) = 1$ and $\mu(\partial R) = 0$, hence R is a continuity set for μ . Therefore, $\mu_n(R) \rightarrow \mu(R) = 1$ by the Portmanteau Theorem [45]. This means that, fixing $\varepsilon > 0$, if we choose the thresholds of the until subformulas from the set GR_n^k , for n large enough, the probability of choosing a bad set of thresholds, for which the formula is not robust and the CSL model checking algorithm may not terminate, will be less than ε .

Remark 5.4. The semi-decidability result presented here is in contrast with the decidability result of model checking for time-homogeneous CTMC. However, in that case the result follows because $P_s(0)$ has a special form allowing the application of the Lindeman-Weierstass Theorem for transcendental numbers (together with zero testing procedures for algebraic numbers) [50]. This, in turn, is a consequence of having constant (rational) rates. In our case, instead, rates are piecewise analytic functions, and we cannot rely on the method of [50] anymore. In fact, in the algorithm for computing the probability, there are two numerical operations that are potential sources of undecidability:

1. Given a number p , which is the analytic image of a rational, decide if it is zero. This is a classical problem whose decidability is not known, even restricting to expressions made up by polynomials and exponentials only [51, 52]. Indeed, its decidability is connected with the truth

of the Schanuel Conjecture [51, 52], which is in turn connected with decidability of the theory of reals extended by the exponential. However, even in case the Schanuel Conjecture holds, it is not clear if the zero problem will be decidable for any analytic function.

2. Detecting the zeros of an analytic function with arbitrary precision. In this case the problem is caused by non-simple zeros, i.e. points in which the function and some of its derivatives are zero. The method sketched in the proof of Theorem 5.1 does not work (see Appendix A), as it relies on the fact that we can bound the derivative away from zero on null points of the function. Furthermore, in the presence of non-simple zeros, detecting if a compact interval is bounded away from zero is semi-decidable (the decision procedure fails if the interval contains a non-simple zero). Whether there is a decidable algorithm for this problem is not known to the authors (even assuming the Schanuel Conjecture is true). It may be possible, however, to find algorithms for some subclass of analytic functions large enough for practical purposes. For instance, if we know a lower bound on the radius of convergence of power series in each analytic point, we can effectively extend the real analytic function to an open ball in the complex plane, and then use methods developed for complex analytic functions [53] which can effectively compute the number of zeros in any sufficiently simple open set, by integrating a function on its boundary with interval arithmetic routines [54, 53].

Our conjecture is that the model checking problem for time-inhomogeneous CTMC is not decidable in general, although it may be decidable for some restricted subclass of rate functions if the Schanuel Conjecture is true. Further investigations on this issue are required.

Remark 5.5 (Computational Complexity). Finding an upper bound on the complexity of the approximation algorithm, when it converges, requires us to find an upper bound on the number of zeros of the solution of a linear differential equation with piecewise analytic rates. This is a non trivial problem, for which some results are known for linear systems with bounded analytic rate functions [55], giving in some cases an upper bound Ψ on the number of zeros, expressible as an elementary function of the upper bound on coefficients of the ODE. However, these upper bounds are astronomically large, and we do not expect such huge complexity in practice. For this reason, we refrain from exploring further this direction in this paper, leaving as future

work the problem of finding tighter bounds under more restrictive hypothesis satisfied by many practical applications.

6. Convergence results for CSL Fluid Model Checking

In this section we reconsider CSL model checking for ICTMC in the light of fluid model checking. In particular, we will prove the asymptotic correctness of the approximation method of Section 4, when considering the sequence $Z_k^{(N)}$ and its fluid limit z_k . The goal is to prove that the truth value of a CSL formula φ computed in the limit model z_k will be the same as that of $Z_k^{(N)}$, for N large enough. Also in this case, we need to restrict our attention to robust CSL formulae. In order to leverage structural induction, as boolean operators pose no real problem, we need to concentrate on next formulae $\varphi = \mathbf{X}^{[T_a, T_b]} \varphi_2$ and on until formulae $\varphi = \mathcal{P}_{\bowtie p}(\varphi_1 \mathbf{U}^{[0, T]} \varphi_2)$, when the time-varying sets associated with the satisfaction of φ_1 and φ_2 are robust.

In particular, model checking these formulas can be reduced to the computation of the next-state probabilities $\bar{P}^{(N)}(t) = P_{next}(Z_k^{(N)}, t, T_a, T_b, G^{(N)})$ and $\bar{P}(t) = P_{next}(z_k, t, T_a, T_b, G)$ (for next formulae) or to reachability probabilities $P^{(N)}(t) = P_{reach}(Z_k^{(N)}, t, T, G^{(N)}, U^{(N)})$ and $P(t) = P_{reach}(z_k, t, T, G, U)$ (for until formulae), where $G^{(N)}(t)$ ($U^{(N)}(t)$) is the set of states satisfying φ_2 ($\neg\varphi_1$) at time t for $Z_k^{(N)}$, while G and U are defined similarly for z_k .¹⁶ We will prove convergence of $\bar{P}^{(N)}(t)$ to $\bar{P}(t)$ and of $P^{(N)}(t)$ to $P(t)$ in Lemmas 6.1 and 6.2.

However, in CSL model checking we are interested in truth values rather than in probabilities, and lifting the previous convergence to truth values is not so straightforward. Consider the path formula $\psi = \varphi_1 \mathbf{U}^{[0, T]} \varphi_2$, and the quantified state formula $\varphi = \mathcal{P}_{\bowtie p}(\psi)$. The problem is that we have to compute its probability $P(t)$ (depending on the initial time t) for z_k and then solve the algebraic equation $P_s(t) - p = 0$ for each state s , to identify for which time instants state s satisfies the formula. Now, the point is that, even in case $P^{(N)}(t) \rightarrow P(t)$ uniformly, we are not guaranteed that $P^{(N)}(t) \bowtie p \rightarrow P(t) \bowtie p$. For instance, if $P(t) = p$, and \bowtie is \leq , then if $P^{(N)}(t)$ converges to $P(t)$ from above, it never satisfies $P^{(N)}(t) \bowtie p$ for any N , hence convergence

¹⁶We can restrict our attention to until formulae with time between $[0, T]$, as intervals $[T_a, T_b]$ can be dealt with by essentially solving two reachability problems of this kind and combining their solution (or better, by computing two transient probabilities and then combining those probabilities, see [13]).

of $P^{(N)}(t) \bowtie p$ to $P(t) \bowtie p$ does not hold. However, things can go wrong only when $P(t) = p$, and the main point of the convergence theorem is to prove that this happens sufficiently “rarely” not to impact on the computation of probabilities of a next or of an until formula in which φ is a sub-formula.

6.1. Preliminary properties

Before establishing convergence results for next-state and reachability probability, we need some straightforward properties of piecewise analytic functions, and a notion of robust convergence of time-varying sets.

Proposition 6.1. *Let $f : I \rightarrow \mathbb{R}$ be a piecewise analytic function, with $I \subseteq \mathbb{R}$ a compact interval. Let $E_f = \{x \in \mathbb{R} \mid \mu_\ell(f^{-1}(\{x\})) = 0\}$ be the set of all values x such that f is not locally constantly equal to x , where μ_ℓ is the Lebesgue measure. Furthermore, let $Z_x = f^{-1}(\{x\})$ be the set of solutions of $f(t) = x$ and let $DZ_f = \{x \in \mathbb{R} \mid \forall t \in Z_x, f'(t) \neq 0\}$. Then*

1. $\forall x \in E_f, Z_x$ is finite.
2. $\mu_\ell(E_f \cap DZ_f) = 1$ ■

The notion of robust time-varying sets, introduced in Definition 5.3, has a counterpart in terms of convergence of time-varying sets:

Definition 6.1. A sequence of time-varying sets $V^{(N)}(t)$, $t \in I$ interval, converges *robustly* to a *robust* time-varying set $V(t)$, $t \in I$, if and only if, for each $s \in \mathcal{S}$ and each open neighbourhood U of $\text{Disc}(V_{\uparrow s})$ (i.e. the set of discontinuity points of $V_{\uparrow s}$), $V_{\uparrow s}^{(N)}(t) \rightarrow V_{\uparrow s}(t)$ *uniformly* in $I \setminus U$.¹⁷

Connecting the notions of robust set and robust convergence, we have the following proposition, proved in Appendix A:

Proposition 6.2. *Let $V^{(N)}(t)$ be a sequence of time varying sets converging robustly to a robust set $V(t)$, $t \in I$. Let $D_V^{(N)} = \{t \mid V^{(N)}(t) \neq V(t)\}$. Then $\mu_\ell(D_V^{(N)}) \rightarrow 0$, where μ_ℓ is the Lebesgue measure on \mathbb{R} . ■*

¹⁷This notion of robust convergence is weaker than convergence according to Skorokhod metric in the space of cadlag functions [45, 56] of the indicator functions of time-varying sets. The difference is in the fact that we do not require that the number of jumps of the sequence of time-varying sets to be definitively the same as in the limit set (as implied by the time-resynchronisation operation of the Skorokhod metric), just to be very close in time.

6.2. Convergence of next-state probability

We consider now the problem of relating the next-state probabilities for the limit single agent process $z_k(t)$ and the sequence of single agent processes $Z_k^{(N)}(t)$ in a population of size N . In particular, we want to show that the probability $\bar{P}_s^{(N)}(t) = P_{next}(Z_k^{(N)}, t_0, T_1, T_2, G)[s]$ converges to $\bar{P}_s(t) = P_{next}(z_k, t_0, T_1, T_2, G)[s]$ uniformly for $t \in [t_0, t_1]$, as N goes to infinity. We will prove this result in a general setting. More specifically, we will consider time-varying sets that can depend on N , and that converge to a robust limit time-varying set in the sense of Definition 6.1. This is needed because the time-varying sets we must consider are obtained by solving (for each $s \in \mathcal{S}$) equations of the form $\bar{P}_s^{(N)}(t) - p = 0$ or $\bar{P}_s(t) - p = 0$, which are generally different, but intuitively converge (as $\bar{P}_s^{(N)}(t)$ converges to $\bar{P}_s(t)$).

The following lemma will be one of the key ingredients to prove the inductive step in the convergence for truth of CSL formulae in Section 6.4.

Lemma 6.1. *Let $\mathcal{X}^{(N)}$ be a sequence of CTMC models, as defined in Section 3.1, and let $Z_k^{(N)}$ and z_k be defined from $\mathcal{X}^{(N)}$ as in Section 3.3, with piecewise real analytic rates, in a compact interval $[0, T']$, for $T' > t_1 + T_b$. Let $G(t)$, $t \in [t_0, t_1 + T_b]$ be a robust time-varying set, and let $G^{(N)}(t)$ be a sequence of time-varying sets converging robustly to G . Furthermore, let $\bar{P}(t) = P_{next}(z_k, t, T_a, T_b, G)$ and $\bar{P}^{(N)}(t) = P_{next}(Z_k^{(N)}, t, T_a, T_b, G^{(N)})$, $t \in [t_0, t_1]$. Finally, fix $p \in [0, 1]$, $\bowtie \in \{\leq, <, >, \geq\}$, and let $V_p(t) = \mathbf{1}\{\bar{P}(t) \bowtie p\}$, $V_p^{(N)}(t) = \mathbf{1}\{\bar{P}^{(N)}(t) \bowtie p\}$. Then*

1. $\bar{P}^{(N)}(t) \rightarrow \bar{P}(t)$, uniformly in $t \in [t_0, t_1]$.
2. For almost every $p \in [0, 1]$, V_p is robust and the sequence $V_p^{(N)}$ converges robustly to V_p .

Proof Sketch: The proof of point 1 combines two main arguments. The first is that, after coupling, convergence in probability of $Z_k^{(N)}$ to z_k implies that, for N large enough, at least a fraction $1 - \varepsilon$ of the trajectories of the two processes coincide for the first T units of time. Conditioning on these trajectories, the conclusive argument is that a sequence of time-varying sets that converges robustly differs from the limit set only in a neighbourhood W of the discontinuity points, which can be made as small as desired for large N (by Proposition 6.1), and that the probability of a process doing its first jump in a time $t \in W$ can also be made as small as desired. Point 2 follows easily after discarding all those values of p for which $\bar{P}(t)$ has non-simple

zeros or that are equal to the value of $\bar{P}(t)$ in a point of non-analyticity, and noting that the set of these thresholds is finite. A more formal treatment can be found in Appendix A. ■

6.3. Convergence of reachability probability

We turn now our attention to the convergence of reachability probabilities. We will first start with the simpler scenario in which goal and unsafe sets are constant, and then extend this result to time-varying sets.

6.3.1. Constant set reachability

Consider now the sequence of processes $Z_k^{(N)}$ defined in Section 3.3. We are interested in the asymptotic behaviour of $P_{reach}(Z_k^{(N)}, t, T, G, U)$ for constant sets G and U . The following result is an immediate consequence of Theorem 3.2:

Proposition 6.3. *Let $\mathcal{X}^{(N)}$ be a sequence of CTMC models, as defined in Section 3.1, and let $Z_k^{(N)}$ and z_k be defined from $\mathcal{X}^{(N)}$ as in Section 3.3. Assume that the infinitesimal generator matrix $Q(t)$ of z_k is bounded and integrable in every compact interval $[0, T]$. Then $P_{reach}(Z_k^{(N)}, t, T, G, U) \rightarrow P_{reach}(z_k, t, T, G, U)$ uniformly in $[t_0, t_1]$ as $N \rightarrow \infty$, i.e.*

$$\sup_{t \in [t_0, t_1]} \|P_{reach}(Z_k^{(N)}, t, T, G, U) - P_{reach}(z_k, t, T, G, U)\| \rightarrow 0.$$
■

The previous proposition shows that the reachability probability for $Z_k^{(N)}$ converges to the reachability probability for z_k , hence for large N we can approximate the former with the latter. Notice also that the hypotheses of the proposition are weaker than those of Lemma 6.1, due to the fact that we consider constant sets.

Example. We consider again the client-server example of Section 3.1 and the two reachability probabilities for a single client discussed in Section 5.3.1, which we report here for convenience:

1. The probability of observing a time-out before being served for the first time within time T .

2. The probability of observing a timeout within time T .

In Figures 8(a), 8(b), 9(a) and 9(b) we can observe a comparison between the values computed for the limit ICTMC z_k and the exact ICTMC $Z_k^{(N)}$, for $N = 15$ or $N = 150$ (with a client-server ratio of 2:1), as a function of the time horizon T . As can be seen, the probability for z_k is in very good agreement with that of $Z_k^{(N)}$ (computed using a statistical approach, from a sample of 10000 traces) even for N relatively small. As far as running time is concerned, the fluid model checking is 100 times faster for $N = 15$, and 1000 times faster for $N = 150$, than the stochastic simulation. What is even more important is that the complexity of the fluid approach is *independent of N* , hence its computational cost (on the order of 200 milliseconds for all cases considered here) can scale to much larger systems. Furthermore, another advantage of the fluid approach is that, by solving a set of differential equations, we are computing the reachability probability for each $t \in [0, T]$ (or better for any finite grid of points in $[0, T]$), while a method based on uniformisation (as in PRISM [16]) has to deal with each time point separately.


In Figures 8(c), 8(d), 9(c) and 9(d), instead, we focus on the reachability probability for both problem 1 and 2 for $T = 50$ as a function of the initial time $t_0 \in [0, 25]$. The value for the fluid model is compared with the probability of $Z_k^{(N)}$ obtained by simulating the full CTMC up to time t_0 and then focussing attention on a specific client in state *request* and starting the computation of the reachability probability.¹⁸ As we can see, the agreement is good also in this case.

Finally, in Figures 8(e), 8(f), 9(e) and 9(f), we compare the reachability probability for $T = 100$ (reachability problem 1) or $T = 250$ (reachability problem 2) of the ICTMC for different populations N and different proportions of clients (n) and servers (m), with the fluid limit. This data confirms that the agreement is good also for small populations for this model.

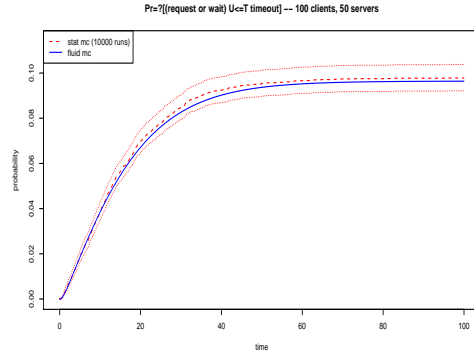
6.3.2. Time-varying set reachability

We consider now the limit behaviour of time-varying reachability probability for $Z_k^{(N)}$, proving that it converges (almost everywhere) to that of z_k . As in Section 5.1, we state this result in a more general form, assuming that

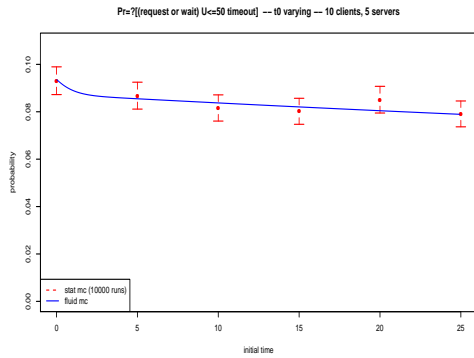
¹⁸This is done by using two indicator variables X_G and X_U that are set equal to one when a trajectory reaches a goal or an unsafe set, respectively. Then, we estimate the reachability probability by the sample mean of X_G at the desired time.



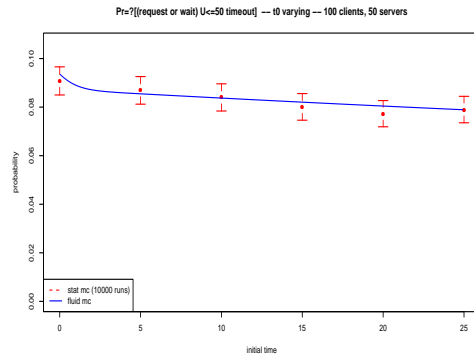
(a) T varying, $n = 10, m = 5$



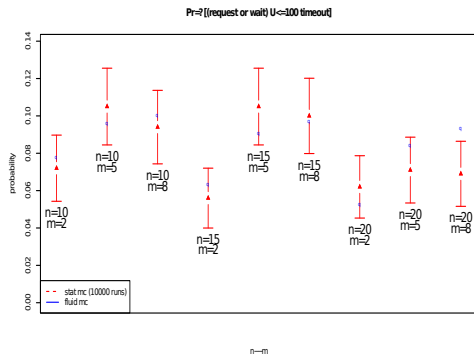
(b) T varying, $n = 100, m = 50$



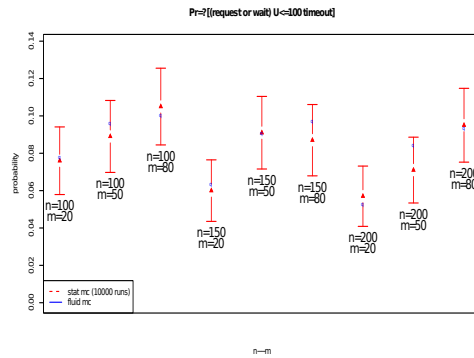
(c) t_0 varying, $n = 10, m = 5$



(d) t_0 varying, $n = 100, m = 50$




(e) n, m varying

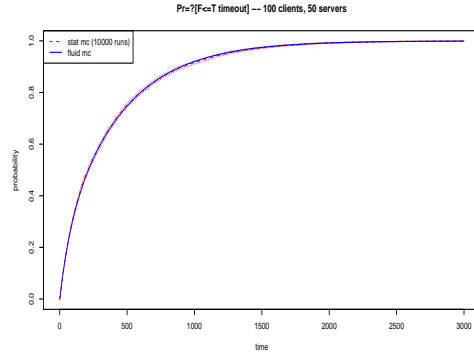


(f) n, m varying


Figure 8: Client-Server model of Section 3.1, single client CTMC. First line: comparison of time-out before being served probability (property 1) for fluid and CTMC models as a function of time horizon T . Second line: comparison of time-out before being served probability (property 1) for fixed time horizon $T = 50$ and variable initial time t_0 . Third line: time-out before being served probability (property 1) at time $T = 250$, and variable number of client and servers.




(a) T varying, $n = 10$, $m = 5$




(b) T varying, $n = 100$, $m = 50$




(c) t_0 varying, $n = 10$, $m = 5$



(d) t_0 varying, $n = 100$, $m = 50$



(e) n , m varying



(f) n , m varying

Figure 9: Client-Server model of Section 3.1, single client CTMC. First line: comparison of time-out probability (property 2) for fluid and CTMC models as a function of time horizon T . Second line: comparison of time-out probability (property 2) for fixed time horizon $T = 50$ and variable initial time t_0 . Third line: comparison of time-out probability (property 2) at time $T = 250$, and variable number of client and servers.

also the goal and unsafe sets depend on N , and converge robustly to some robust limit sets G and U . Furthermore, we require that G and U are *compatible* in the sense of Definition 5.4, i.e. that they do not have a discontinuity at the same time for the same state s : $\forall s \in \mathcal{S}, \text{Disc}(G_s) \cap \text{Disc}(U_s) = \emptyset$. The following lemma, which is also the basic inductive step to prove convergence for CSL model checking formulae, relies on the functions involved being piecewise analytic.

Lemma 6.2. *Let $\mathcal{X}^{(N)}$ be a sequence of CTMC models, as defined in Section 3.1, and let $Z_k^{(N)}$ and z_k be defined from $\mathcal{X}^{(N)}$ as in Section 3.3, with piecewise analytic rates, in a compact interval $[0, T']$, for T' sufficiently large.*

Let $G(t), U(t), t \in [t_0, t_1 + T]$ be compatible and robust time-varying sets, and let $G^{(N)}(t), U^{(N)}(t)$ be sequences of time-varying sets converging robustly to G and U , respectively.

Furthermore, let $P(t) = P_{\text{reach}}(z_k, t, T, G, U)$ and

$P^{(N)}(t) = P_{\text{reach}}(Z_k^{(N)}, t, T, G^{(N)}, U^{(N)}), t \in [t_0, t_1]$.

Finally, fix $p \in [0, 1]$, $\bowtie \in \{\leq, <, >, \geq\}$, and let $V_p(t) = \mathbf{1}\{P(t) \bowtie p\}$, $V_p^{(N)}(t) = \mathbf{1}\{P^{(N)}(t) \bowtie p\}$. Then

1. *For all but finitely many $t \in [t_0, t_1]$, $P^{(N)}(t) \rightarrow P(t)$, with uniform speed (i.e. independently of t).*
2. *For almost every $p \in [0, 1]$, V_p is robust and the sequence $V_p^{(N)}$ converges robustly to V_p .*

Proof Sketch: The proof, reported in Appendix A, is very similar to the one of Lemma 6.1. The only difference is that convergence to $P(t)$ can fail in all those time instants t in which the goal or the unsafe sets have a discontinuity, which are finite. This extra level of complexity is reflected in statement 1. ■

6.4. Convergence for CSL formulae

We are now ready to state a convergence result for CSL model checking. Also in this case, we will restrict our attention to robust CSL formulae. This is reasonable, as we want to use Lemmas 6.1 and 6.2, which require robustness of time-varying sets.

In particular, we can reduce this problem to the computation of the next-state probabilities $\bar{P}^{(N)}(t) = P_{\text{next}}(Z_k^{(N)}, t, T_a, T_b, G^{(N)})$ and $\bar{P}(t) = P_{\text{next}}(z_k, t, T_a, T_b, G)$ (for next formulae) or to reachability probabilities $P^{(N)}(t) =$

$P_{reach}(Z_k^{(N)}, t, T, G^{(N)}, U^{(N)})$ and $P(t) = P_{reach}(z_k, t, T, G, U)$ (for until formulae), where $G^{(N)}(t)$ ($U^{(N)}(t)$) is the set of states satisfying φ_2 ($\neg\varphi_1$) at time t for $Z_k^{(N)}$, while G and U are defined similarly for z_k .¹⁹ Then, we may resort to Lemmas 6.1 and 6.2 to prove convergence of $\bar{P}^{(N)}(t)$ to $\bar{P}(t)$ and of $P^{(N)}(t)$ to $P(t)$.

Theorem 6.1. *Let $\mathcal{X}^{(N)}$ be a sequence of CTMC models, as defined in Section 3.1, and let $Z_k^{(N)}$ and z_k be defined from $\mathcal{X}^{(N)}$ as in Section 3.3.*

Assume that $Z_k^{(N)}$, z_k have piecewise analytic infinitesimal generator matrices.

Let $\varphi(p_1, \dots, p_k)$ be a robust CSL formula. Then, there exists an N_0 such that, for $N \geq N_0$ and each $s \in \mathcal{S}$

$$s, 0 \models_{Z_k^{(N)}} \varphi \Leftrightarrow s, 0 \models_{z_k} \varphi.$$

Proof Sketch: The proof proceeds by structural induction, using Lemmas 6.1 and 6.2 to prove the inductive steps for path formulae, and leveraging the fact that the time dependent satisfaction of a subformula of a robust CSL formula is robust in the sense of Definition 5.3. The fact that at time 0 truth values converge is due to one of the properties of robust CSL formulae, requiring that the values of top path formulae at time zero are different from the thresholds they are compared to, hence a sequence converging to this value will definitively be greater or smaller than such a threshold. The full proof is in Appendix A. ■

Corollary 6.1. *Given a CSL formula $\varphi(\mathbf{p})$, with $\mathbf{p} \in [0, 1]^k$, then the subset of $[0, 1]^k$ in which convergence holds has Lebesgue measure 1 and is open in $[0, 1]^k$. ■*

The previous theorem shows that the results that we obtain abstracting a single agent in a population of size N with the fluid approximation is consistent. However, the theorem excludes the sets of constants \mathbf{p} for which the formula is not robust. Interestingly, this is the same condition required for decidability of the model checking problem for ICTMC, a fact that shows

¹⁹As previously remarked in footnote 16, we can restrict our attention to until formulae with time between $[0, T]$.

how these two aspects are intimately connected. Notice that, contrary to decidability, this limitation is unavoidable and is present also in the case of sequences of processes converging to a time-homogeneous CTMC. In this case, in fact, the next-state and reachability probabilities are constant with respect to the initial time, and their value p (in the limit model) can cause convergence of truth values to fail.

However, notice that the constants p appearing in a formula that can make convergence fail depend only on the limit CTMC z_k . Hence we can detect potentially dangerous situations while solving the CSL model checking for the limit process (in these cases the model checking algorithm may fail to provide an answer).

Remark 6.1. In this paper, we are considering only time bounded operators. This limitation is a consequence of the very nature of the approximation Theorem 3.2, which holds only for a finite time horizon. However, there are situations in which we can extend the validity of the theorem to the whole time domain, but this extension depends on properties of the phase space of the fluid ODE [57, 58, 59].

In those cases, we can prove convergence of the steady state behaviour of $Z_k^{(N)}$ to that of z_k , hence we can incorporate also operators dealing with steady state properties (see [60, 61] for a discussion of this issue).

In order to deal with time unbounded operators, instead, convergence to steady state is not enough. We also need to ensure that the equation $P(t) - p$ has a finite number of zeros on the whole positive time axis. The piecewise analytic property is not sufficient in this case (think about sine and cosine), and stronger conditions have to be required. However, for periodic functions, we may reason similarly to [29], if we can prove that periodicity of rate functions implies periodicity in the reachability probabilities as a function of initial time.

7. Comparison of CSL model checking for $Z_k^{(N)}$ and $(Z_k^{(N)}, \hat{\mathbf{X}}^{(N)})$

In this paper we have considered two possible descriptions of a single agent at a fixed population level N , i.e. $Z_k^{(N)}(t)$ and $(Z_k^{(N)}(t), \hat{\mathbf{X}}^{(N)}(t))$. From the discussion in Sections 3.3 and 4 we already know that, while $(Z_k^{(N)}(t), \hat{\mathbf{X}}^{(N)}(t))$ is a CTMC with finite (but extremely large) state space, $Z_k^{(N)}(t)$ has a much smaller state space but is not a Markov process. Furthermore, its behaviour is time dependent. The non-Markovian nature of $Z_k^{(N)}(t)$ has consequences

for its reachability probability (see Section 4), meaning that its value is dependent on the initial time at which we compute it. This implies that the satisfiability of a CSL formula (with the truth value of atomic propositions depending only on \mathcal{S}) for $Z_k^{(N)}(t)$ can depend on the time at which we evaluate it. Hence we need to consider time-dependent sets to compute the probabilities of next or until path formulae. But time-dependent sets can introduce discontinuities in such probabilities, as discussed in Section 5.3.2. On the other hand, $(Z_k^{(N)}(t), \hat{\mathbf{X}}^{(N)}(t))$ is a time-homogeneous CTMC, hence its next-state and reachability probabilities do not depend on time and no time-dependent notion of satisfaction has to be considered in this case. In particular, when considering $(Z_k^{(N)}(t), \hat{\mathbf{X}}^{(N)}(t))$, its reachability probability is always a continuous function. This implies that the truth value of a CSL formula containing nested next or until sub-formulae, can be different if we consider its satisfiability with respect to $Z_k^{(N)}(t)$ or $(Z_k^{(N)}(t), \hat{\mathbf{X}}^{(N)}(t))$.

However, despite this discrepancy for finite N , we will prove that the satisfiability for $Z_k^{(N)}(t)$ and $(Z_k^{(N)}(t), \hat{\mathbf{X}}^{(N)}(t))$ is asymptotically the same, at least if we restrict to robust CSL formulae. In order to show this, we will combine the convergence results of the previous sections with additional results relative to $(Z_k^{(N)}(t), \hat{\mathbf{X}}^{(N)}(t))$ and $(z(t), \mathbf{x}(t))$.

Example. If we observe Figures 9(c) and 9(d), we can easily convince ourselves that the reachability probability for $Z_k^{(N)}$ in the running example for the formula $\varphi_1 = \text{true } \mathbf{U}^{[0,50]} \text{ timeout}$ depends on the initial time. Hence it gives rise to a time-dependent set for the satisfiability of the formula $\varphi_2 = \mathcal{P}_{<0.167}(\varphi_1)$. This implies that for $Z_k^{(N)}$, the probability of the formula $\varphi = \text{true } \mathbf{U}^{[0,T]} \varphi_2$ will have discontinuities as a function of T , similarly to the case for z_k . However, if we compute the reachability probability for φ in $(Z_k^{(N)}, \hat{\mathbf{X}}^{(N)})$, in a state s, \mathbf{x}_0 , this will be a continuous function of T , hence the two probabilities are different.

We will now prove the convergence of the standard CSL model checking for $\mathbf{Y}^{(N)}(t) = (Z_k^{(N)}(t), \hat{\mathbf{X}}^{(N)}(t))$ in state s, \mathbf{x}_0 , to the CSL model checking procedure for $\mathbf{y}(t) = (z_k(t), \mathbf{x}(t))$, which is equivalent to the one for $z_k(t)$ alone. This procedure requires us to compute, given a next formula $\varphi = \mathbf{X}^{[T_a, T_b]} \varphi_1$ or an until formula $\varphi = \varphi_1 \mathbf{U}^{[T_a, T_b]} \varphi_2$, its probability $P(s, \mathbf{x})$ starting from time 0, in each point (s, \mathbf{x}) of the state space $\mathcal{S} \times E$ of $\mathbf{y}(t)$, and then solve the inequality $P(s, \mathbf{x}) \bowtie p$, to determine the truth of $\mathcal{P}_{\bowtie p}(\varphi)$ in (s, \mathbf{x}) . This defines a subset of $\mathcal{S} \times E$ where $\mathcal{P}_{\bowtie p}(\varphi)$ is true.

The intuition behind the proof is that the truth value of an until formula

in a state (s, \mathbf{x}_0) for $\mathbf{y}(t)$ does not depend on the whole state space $\mathcal{S} \times E$, but only on the points of E intersected by the solution of the fluid ODE starting in \mathbf{x}_0 , i.e. on $\mathcal{S} \times \Phi([0, T], \mathbf{x}_0)$, where $\Phi(t, \mathbf{x}_0)$ is the flow of the differential equation²⁰. Furthermore, the convergence of $\hat{\mathbf{X}}^{(N)}(t)$ to $\mathbf{x}(t)$ allows us to restrict attention to an arbitrary small neighbourhood of $\Phi([0, T], \mathbf{x}_0)$, in order to solve the model checking problem for $\mathbf{Y}^{(N)}(t)$, for N large enough.

In the following, we need some additional concepts and definitions.

Consider the domain $\hat{\mathcal{D}}^{(N)} \subset E$ of $\hat{\mathbf{X}}^{(N)}$. With each point $\mathbf{x} \in E$, we associate a point $\nu^{(N)}(\mathbf{x}) \in \hat{\mathcal{D}}^{(N)}$, such that $\|\mathbf{x} - \nu^{(N)}(\mathbf{x})\| < \frac{\epsilon}{N}$. The existence of such a point is guaranteed by the definition of E . Now, we further assume that, given a point $(s, \mathbf{x}) \in E$, the initial state $\mathbf{Y}^{(N)}(0)$ is $(s, \nu^{(N)}(\mathbf{x}))$, so that $\mathbf{Y}^{(N)}(0)$ converges to (s, \mathbf{x}) uniformly in space. This choice of $\mathbf{Y}^{(N)}(0)$ guarantees uniform bounds in space for Kurtz theorem and the fast simulation theorem, for convergence in probability.²¹

Now, consider the fluid limit differential equation, and let $\Phi(t, \mathbf{x}_0)$ be its flow. We assume that $\Phi(t, \mathbf{x}_0)$ is a piecewise analytic function with respect to t and \mathbf{x} . The T, ϵ -flow tube for \mathbf{x}_0 is the set $E_0 \subset E$, defined by $E_0 = \Phi([0, T], B_\epsilon(\mathbf{x}_0))$, i.e. the set of all trajectories up to time T starting in a ball of radius ϵ centred in \mathbf{x}_0 . Now, consider a T, ϵ -flow tube E_0 for \mathbf{x}_0 . For any $\mathbf{x} \in E_0$, let $T_{\mathbf{x}}^+ = T_{\mathbf{x}}^+(E_0) = \sup\{t \mid \Phi([0, t], \mathbf{x}) \in E_0\}$ be the time at which the trajectory starting in \mathbf{x} leaves E_0 . Furthermore, let $T_{\mathbf{x}}^- = T_{\mathbf{x}}^-(E_0) = \inf\{t \mid \Phi([t, 0], \mathbf{x}) \in E_0\}$ be the time at which the trajectory starting in \mathbf{x} enters E_0 .

A subset $D \subseteq \mathcal{S} \times E_0$ is a d -set for E_0 if and only if, (i) D is closed (in $\mathcal{S} \times E_0$), (ii) D is the union of a finite number of smooth manifolds²² of dimension $n - 1$ or less, and (iii) for each $\mathbf{x} \in E_0$, it holds that $\{s\} \times \Phi([T_{\mathbf{x}}^-(E_0), T_{\mathbf{x}}^+(E_0)], \mathbf{x}) \cap D$ contains at most k points in each state s . In other words, a d -set is a union of piecewise analytic manifolds that intersects each trajectory in at most k points. It can be easily checked that each d -set has (Lebesgue) measure zero.²³

²⁰The solution of the fluid ODE at time t starting in \mathbf{x}_0 at time 0.

²¹The speed of convergence to the fluid limit depends on the initial conditions only through $\|\hat{\mathbf{X}}^{(N)}(0) - \mathbf{x}(0)\|$; the choice of $\nu^{(N)}(\mathbf{x})$ guarantees the uniform convergence of this quantity with respect to \mathbf{x} .

²²A smooth manifold is the zero set of a sufficiently smooth function, in this paper at least having continuous first-order derivatives.

²³Any set of topological dimension $n - 1$ or less has Lebesgue measure zero in \mathbb{R}^n .

We also introduce a notion of *robust subset* of $\mathcal{S} \times E_0$, for a T, ε -flow tube E_0 in \mathbf{x}_0 . Consider a subset $V \subset \mathcal{S} \times E_0$. We say that V is robust in $\mathcal{S} \times E_0$ if and only if, (i) its boundary ∂V is a d-set in $\mathcal{S} \times E_0$, and (ii) for each $(s, \mathbf{x}) \in \mathcal{S} \times E_0$, the time-varying set $V_{\mathbf{x}}[s](t) = \mathbf{1}\{(s, \Phi(t, \mathbf{x})) \in V\}$, $T_{\mathbf{x}}^- < t < T_{\mathbf{x}}^+$, is robust in the sense of Definition 5.3 (notice that it contains at most $k < \infty$ discontinuity points, where k does not depend on \mathbf{x} , as ∂V is a d-set). We sometimes denote ∂V by $Disc(V)$. We also say that two robust subsets V_1 and V_2 of $\mathcal{S} \times E_0$ are *compatible* if $\partial V_1 \cap \partial V_2 = \emptyset$.

Similarly to Section 5.3.2, we say that a sequence of sets $V^{(N)} \subset \mathcal{S} \times E_0$ *converges robustly* to a robust set $V \subseteq \mathcal{S} \times E_0$, with E_0 a T, ε -flow tube in \mathbf{x}_0 , if and only if, for each open neighbourhood U of $Disc(V)$, there is $N_0 > 0$ such that, $\forall N \geq N_0$ and all $(s, \mathbf{x}) \in (\mathcal{S} \times E_0) \setminus U$, $(s, \mathbf{x}) \in V^{(N)}$ if and only if $(s, \mathbf{x}) \in V$.

We are now ready to state the following lemmas, which are space-versions of Lemmas 6.1 and 6.2 on time-varying sets, and are the key to the induction step of Lemma 7.4.

Lemma 7.1. *Let $E_0 \subset E$ be a T, ε_0 -flow tube for \mathbf{x}_0 . Let G be a robust subset of $\mathcal{S} \times E_0$, and $G^{(N)}$ be a sequence of subsets of $\mathcal{S} \times E_0$ that converge robustly to G .*

Let $\bar{P}(s, \mathbf{x}) = P_{next}(\mathbf{y}, s, \mathbf{x}, T_a, T_b, G)$ be the probability that the first jump of $\mathbf{y}(t)$ is into a state in G and happens at a time $t \in [T_a, T_b]$, given that \mathbf{y} started at time $t = 0$ in state $(s, \mathbf{x}) \in \mathcal{S} \times E_0$, and let $\bar{P}^{(N)}(s, \mathbf{x}) = P_{next}^{(N)}(\mathbf{Y}^{(N)}, s, \nu^{(N)}(\mathbf{x}), T_a, T_b, G^{(N)})$ be defined similarly, with G and \mathbf{x} replaced by $G^{(N)}$ and $\nu^{(N)}(\mathbf{x})$, respectively. Furthermore, define $V = \{(s, \mathbf{x}) \mid \bar{P}(s, \mathbf{x}) \triangleright p\}$ and $V^{(N)} = \{(s, \mathbf{x}) \mid \bar{P}^{(N)}(s, \mathbf{x}) \triangleright p\}$. Then there exists $\varepsilon_1 > 0$ such that, in E_1 , the $(T - T_b), \varepsilon_1$ -flow tube for \mathbf{x}_0 :

1. $\bar{P}^{(N)}(s, \mathbf{x}) \rightarrow \bar{P}(s, \mathbf{x})$ for all $\mathbf{x} \in E_1$, uniformly in (s, \mathbf{x}) .
2. If $V_{\mathbf{x}_0}(t)$, $t \in [T_{\mathbf{x}_0}^-(E_1), T_{\mathbf{x}_0}^+(E_1)]$, is a robust time-varying set, then V is robust in E_1 and $V^{(N)}$ converges robustly to V .

Proof Sketch: The proof is similar in spirit to that of Lemma 6.1, but with an extra level of complexity caused by the fact that now we need to take into account the spatial dimension in addition to the temporal one. Here we rely on the fact that, by robust convergence of $G^{(N)}$ to G , choosing a small neighbourhood W of the d-set $Disc(G)$, the time spent by $\hat{\mathbf{X}}^{(N)}$ or \mathbf{x} in W can be made arbitrarily small, so that with very high probability $\mathbf{Y}^{(N)}$ and \mathbf{y} will end up doing the first jump of the s -component outside it. More details

can be found in Appendix A. ■

Lemma 7.2. *Let $E_0 \subset E$ be a T, ε_0 -flow tube for \mathbf{x}_0 . Let U and G two robust and compatible subsets of $\mathcal{S} \times E_0$, and $U^{(N)}, G^{(N)}$ be sequences of subsets of $\mathcal{S} \times E_0$ that converge robustly to U and G , respectively.*

Let $P(s, \mathbf{x}) = P_{\text{reach}}(\mathbf{y}, s, \mathbf{x}, T_1, T_2, U, G)$ be the probability that $\mathbf{y}(t)$ reaches a state in G within time $[T_a, T_b]$, avoiding any unsafe state in U , given that \mathbf{y} started at time $t = 0$ in state $(s, \mathbf{x}) \in \mathcal{S} \times E_0$, and let $P^{(N)}(s, \mathbf{x}) = P_{\text{reach}}(\mathbf{Y}^{(N)}, s, \nu^{(N)}(\mathbf{x}), T_a, T_b, U^{(N)}, G^{(N)})$ be defined similarly, with G, U, \mathbf{x} replaced by $G^{(N)}, U^{(N)}$, and $\nu^{(N)}(\mathbf{x})$, respectively.

Furthermore, we define $V = \{(s, \mathbf{x}) \mid P(s, \mathbf{x}) \bowtie p\}$ and $V^{(N)} = \{(s, \mathbf{x}) \mid P^{(N)}(s, \mathbf{x}) \bowtie p\}$. Then there exists $\varepsilon_1 > 0$ such that, in E_1 , the $(T - T_b), \varepsilon_1$ -flow tube for \mathbf{x}_0 :

1. $P^{(N)}(s, \mathbf{x}) \rightarrow P(s, \mathbf{x})$ for all $\mathbf{x} \in E_1 \setminus D$, where D is a d -set, uniformly in (s, \mathbf{x}) .
2. If $V_{\mathbf{x}_0}(t)$, $t \in [T_{\mathbf{x}_0}^-(E_1), T_{\mathbf{x}_0}^+(E_1)]$, is a robust time-varying set, then V is robust in E_1 and $V^{(N)}$ converges robustly to V .

Proof Sketch: Similarly to Lemma 7.1, this is the spatial counterpart of Lemma 6.2, and this is reflected in a similar but more complex proof. In particular, also in this case convergence of $P^{(N)}(s, \mathbf{x})$ to $P(s, \mathbf{x})$ fails at points (s, \mathbf{x}) belonging to the boundary of the two robust sets U and G . A detailed proof can be found in Appendix A. ■

The previous lemmas are the key arguments used in the structural induction to prove the following result.

Lemma 7.3. *Let $\mathcal{X}^{(N)}$ be a sequence of CTMC models, as defined in Section 3.1, and let $Z_k^{(N)}$ and z_k be defined from $\mathcal{X}^{(N)}$ as in Section 3.3.*

Assume that there is a flow tube E_0 of \mathbf{x}_0 such that all trajectories in E_0 are piecewise analytic.

Let $\varphi = \varphi(\mathbf{p})$ be a robust CSL formula for the trajectory $\Phi(t, \mathbf{x}_0)$. Then, there is an N_0 such that, for all $N \geq N_0$,

$$s, \mathbf{x}_0 \models_{\mathbf{y}} \varphi \Leftrightarrow s, \nu^{(N)}(\mathbf{x}_0) \models_{\mathbf{Y}^{(N)}} \varphi.$$

Proof Sketch: The proof goes by structural induction on the CSL formula. The essential point is to show that, under the imposed restrictions, the satisfaction sets of subformulae converge robustly to the limit satisfaction set.

This property is propagated upwards in the formula tree by Lemmas 7.1 and 7.2 applied to path subformulae; see Appendix A. ■

We now turn to consider the relationship between the model checking problem of a CSL formula φ for $z_k(t)$ and the model checking problem for the same formula with respect to $\mathbf{y}(t)$. In this case, it is easy to see that a formula is true for $z_k(t)$ if and only if it is true for $\mathbf{y}(t)$. In fact, in this process the truth value of a formula in state (s, \mathbf{x}_0) depends only on the trajectory $\Phi(t, \mathbf{x}_0)$ starting in \mathbf{x}_0 . Furthermore, if we fix a time \bar{t} and consider the point $\mathbf{x}_{\bar{t}} = \Phi(\bar{t}, \mathbf{x}_0)$, then the process $\bar{z}_k(t)$, defined with respect to the trajectory $\Phi(t, \mathbf{x}_{\bar{t}})$ starting in point $\mathbf{x}_{\bar{t}}$ at time zero, equals the process $z_k(t + \bar{t})$, starting in \mathbf{x}_0 at time zero, due to the semi-group property of the flow $\Phi(\cdot, \cdot)$. Hence, any reachability probability for \mathbf{z}_k with respect to the initial time \bar{t} equals the reachability probability for \bar{z}_k at time 0: We can always turn a time-dependent reachability problem into a more classical space-dependent one. From the previous discussion, the following lemma follows:

Lemma 7.4. *Let $\mathcal{X}^{(N)}$ be a sequence of CTMC models, as defined in Section 3.1, and let $Z_k^{(N)}$ and z_k be defined from $\mathcal{X}^{(N)}$ as in Section 3.3. Let $\varphi = \varphi(\mathbf{p})$ be a robust CSL formula for the piecewise analytic trajectory $\Phi(t, \mathbf{x}_0)$, and let z_k be the ICTMC defined on \mathcal{S} with respect to trajectory $\Phi(t, \mathbf{x}_0)$. Then,*

$$s, \mathbf{x}_0 \models_{\mathbf{y}} \varphi \Leftrightarrow s \models_{z_k} \varphi.$$

■

Using the previous lemmas, we can show the following theorem.

Theorem 7.1. *Let $\mathcal{X}^{(N)}$ be a sequence of CTMC models, as defined in Section 3.1. Assume that there is a flow tube E_0 of \mathbf{x}_0 such that all trajectories in E_0 are piecewise analytic.*

Let $\varphi = \varphi(\mathbf{p})$ be a robust CSL formula for the trajectory $\Phi(t, \mathbf{x}_0)$, let $Z_k^{(N)}(t)$ and $z_k(t)$ be the stochastic processes on \mathcal{S} defined as in Section 3.3, and let $\mathbf{y}(t)$ and $\mathbf{Y}^{(N)}(t)$ be defined as in this section. Then, there is an N_0 such that, for all $N \geq N_0$,

$$s \models_{Z_k^{(N)}} \varphi \Leftrightarrow s, \nu^{(N)}(\mathbf{x}_0) \models_{\mathbf{Y}^{(N)}} \varphi.$$

Proof. There exists an N_0 , such that, for all $N \geq N_0$,

$$s \models_{Z_k^{(N)}} \varphi \Leftrightarrow s \models_{z_k} \varphi \Leftrightarrow s, \mathbf{x}_0 \models_{\mathbf{y}} \varphi \Leftrightarrow s, \nu^{(N)}(\mathbf{x}_0) \models_{\mathbf{Y}^{(N)}} \varphi,$$

where the first equivalence follows from Theorem 6.1, the second equivalence from Lemma 7.4, and the third equivalence from Lemma 7.3, while N_0 can be chosen as the largest one between that of Theorem 6.1 and that of Lemma 7.4. ■

Inspecting the proof of the previous theorem, the following corollary is straightforward.

Corollary 7.1. *Let $\varphi = \varphi(\mathbf{p})$ be a robust CSL formula for the trajectory $\Phi(t, \mathbf{x}_0)$. Then, there is an N_0 such that, for all $N \geq N_0$,*

$$s, \nu^{(N)}(\mathbf{x}_0) \models_{\mathbf{Y}^{(N)}} \varphi \Leftrightarrow s \models_{z_k} \varphi.$$
■

8. Conclusions

Summary. In this paper we exploited a corollary of fluid limit theorems to approximate properties of the behaviour of single agents in large population models. In particular, we focussed on reachability and stochastic model checking of CSL formulae. The method proposed requires us to model check a time-inhomogeneous CTMC of size equal to the number of internal states of the agent (which is usually rather small). Hence, it gives a large improvement in terms of computational efficiency. This is the main methodological contribution of the paper (Section 4).

The first theoretical result of this paper is a CSL model checking algorithm for ICTMC (Section 5). We first provided algorithms for the next state (Section 5.1) reachability problems (Section 5.3) for ICTMC, both in the case of time-constant and time-varying sets, and then combined them into a proper model checking algorithm for the time-bounded fragment of CSL (Section 5.4). We also gave a quasi-decidability result, showing that the algorithm works for all robust formulae, where the set of non-robust formulae has measure zero (Section 5.5).

The second theoretical contribution of the paper is a proof of correctness of the fluid approximation for CSL properties of individual agents (Section 6):

we first proved convergence of the next state (Section 6.2) and reachability probabilities (Section 6.3) computed for the single agent in a finite population of size N to those of the limit fluid CTMC, and then lifted this to the convergence of the truth value of CSL formulae (Section 6.4).

Practical considerations. The method presented in this paper is non-trivial, and relies on several assumptions. Here we will collect the various comments on the practical side, arguing that, despite the theoretical intricacies, the CSL model checking should work well at least for the application that brought us to investigate it, namely fluid model checking.

- One important requirement of the approach is the piecewise analyticity of rates. This is used to theoretically enforce that reachability or next state probabilities cross a threshold p only a finite number of times, i.e. that the solutions of $P(t) - p = 0$ are finite. Practically, most of the functions used in actual models (e.g. minimum, polynomials, rational functions, exponentials, logarithms) are piecewise analytic. Hence this is not very restrictive. However, our method should work, whenever the finiteness condition on time-varying satisfaction sets holds or for non-nested properties, under milder continuity hypothesis (i.e. Lipschitz continuity of rates, although piecewise smooth functions can be treated as well in some cases, see [62]).
- The complexity of ICTMC CSL model checking depends on several factors. One is the size n of the state space: We need to solve a non-autonomous linear system of differential equations quadratic in n . However, the algorithm we presented was designed with fluid model checking in mind. Here the state space is that of a single agent in a population model. Hence, in this context n will be really small, usually less than 10. Hence, solving such a system of ODEs not only is expected to be feasible, but also extremely fast. Another source of complexity, when nesting properties, depends on the number of solutions of the equation $P(t) - p = 0$ and on the levels of nesting. However, in most practical cases we do not expect to have to deal with many zeros of the previous equation. In all models we studied (see also [60]), this number was one or two per state. As for the nesting, one can safely argue that in practical applications nested properties are rare, and with at most one level of nesting being usual (see the illuminating discussion of [63]).

Summarising, the model checking algorithm presented in this paper is designed for the fluid approximation of individual agents in large population models. In these situations, we expect it not only to be computationally efficient in practice, but to be the only computationally feasible method to check such properties.

Future work. There are many issues that we wish to tackle in the future. First, we would like to better understand the quality of convergence. This can be accomplished by trying to derive theoretical error bounds (which may be too loose to be of practical interest) or by running many experiments to identify situations in which the approximation performs well (in terms of both classes of formulae and model structure). In addition, we would like to provide a working implementation of the model checking algorithm for ICTMC, studying its computational cost empirically (and exploring how easy it is in practice to find a non computable instance). Furthermore, we want to investigate the connections between single agent properties and system level properties. We believe this approach can become a powerful tool to investigate the relationship between microscopic and macroscopic characterisations of systems, and to understand their emergent behaviour.

As far as CSL model checking for ICTMC is concerned, we aim to extend it to include time unbounded and steady state operators, at least for those subsets of rate functions in which the algorithm can be shown to be decidable. We also need to consider rewards, at least for a finite time horizon (here we expect their inclusion to be relatively straightforward). Then, we would like to show convergence results also for this larger subset of CSL, under the hypothesis required for steady state convergence of the fluid approximation.

Another line of investigation would be to consider different temporal logics, such as MTL. For this logic, asymptotic correctness is relatively easy to prove, along the lines of Proposition 6.3. What is more difficult is to find an effective algorithm to model check MTL properties for ICTMC. One possibility may be to combine the approaches of [34, 29, 30], and exploit algorithms and techniques to compute reachability of PDMP [31].

References

- [1] R. Bakhshi, L. Cloth, W. Fokkink, B. Haverkort, Mean-field analysis for the evaluation of gossip protocols, in: Proceedings of Sixth International Conference on the Quantitative Evaluation of Systems, QEST'09, 2009, pp. 247–256.

- [2] R. Bakhshi, L. Cloth, W. Fokkink, B. R. Haverkort, Mean-field framework for performance evaluation of push-pull gossip protocols, *Performance Evaluation* 68 (2) (2011) 157–179.
- [3] A. Kolesnichenko, A. Remke, P. de Boer, B. Haverkort, Comparison of the mean-field approach and simulation in a peer-to-peer botnet case study, in: *Proceedings of 8th European Performance Engineering Workshop*, EPEW 2011, 2011, pp. 133–147.
- [4] J. Hillston, Fluid flow approximation of PEPA models, in: *Proceedings of the Second International Conference on the Quantitative Evaluation of SysTems*, QEST 2005, 2005, pp. 33 – 42. doi:10.1109/QEST.2005.12.
- [5] L. Bortolussi, A. Policriti, Dynamical systems and stochastic programming: To ordinary differential equations and back, in: *Transactions on Computational Systems Biology XI*, Vol. 5750 of *Lecture Notes in Computer Science*, Springer Berlin / Heidelberg, 2009, pp. 216–267, 10.1007/978-3-642-04186-0_11.
URL http://dx.doi.org/10.1007/978-3-642-04186-0_11
- [6] M. Benaïm, J. L. Boudec, A class of mean field interaction models for computer and communication systems, *Performance Evaluation* 65 (11-12) (2008) 823–838.
- [7] T. G. Kurtz, Solutions of ordinary differential equations as limits of pure jump Markov processes, *Journal of Applied Probability* 7 (1970) 49–58.
- [8] R. Darling, Fluid limits of pure jump Markov processes: A practical guide, *ArXiv e-prints arXiv:math/0210109*.
- [9] R. Darling, J. Norris, Differential equation approximations for Markov chains, *Probability Surveys* 5 (2008) 37–79.
- [10] R. A. Hayden, J. T. Bradley, A fluid analysis framework for a Markovian process algebra, *Theoretical Computer Science* 411 (22-24) (2010) 2260–2297.
- [11] A. Singh, J. Hespanha, Lognormal moment closures for biochemical reactions, in: *Proceedings of 45th IEEE Conference on Decision and Control*, 2006, pp. 2063–2068.

- [12] L. Bortolussi, On the approximation of stochastic concurrent constraint programming by master equation, in: Proceedings of the 6th International Workshop on Quantitative Aspects of Programming Languages, QAPL'08, Vol. 220, Electr. Notes Theor. Comput. Sci., 2008, pp. 163–180.
- [13] C. Baier, B. Haverkort, H. Hermanns, J. Katoen, Model checking continuous-time Markov chains by transient analysis, in: Proceedings of the 12th International Conference on Computer Aided Verification, CAV'00, 2000, pp. 358–372. doi:10.1007/10722167_28.
- [14] A. Aziz, V. Singhal, F. Balarin, R. Brayton, A. Sangiovanni-Vincentelli, Verifying continuous time Markov chains, in: Proceedings of the 8th International Conference on Computer Aided Verification, CAV'96, 1996, pp. 269–276.
- [15] J. Rutten, M. Kwiatkowska, G. Norman, D. Parker, Mathematical Techniques for Analyzing Concurrent and Probabilistic Systems, Vol. 23 of CRM Monograph Series, American Mathematical Society, 2004.
- [16] M. Kwiatkowska, G. Norman, D. Parker, Probabilistic symbolic model checking with PRISM: A hybrid approach, International Journal on Software Tools for Technology Transfer 6 (2) (2004) 128–142.
- [17] M. Kattenbelt, M. Kwiatkowska, G. Norman, D. Parker, Game-based probabilistic predicate abstraction in prism, Electr. Notes Theor. Comput. Sci. 220 (3) (2008) 5–21.
- [18] M. Kattenbelt, M. Kwiatkowska, G. Norman, D. Parker, Abstraction refinement for probabilistic software, in: Proceedings of the 10th International Conference on Verification, Model Checking, and Abstract Interpretation, VMCAI 2009, 2009, pp. 182–197.
- [19] D. Sumpter, From bee to society: An agent-based investigation of honey bee colonies, Ph.D. thesis, University of Manchester (2000).
- [20] H. Qian, E. Elson, Single-molecule enzymology: stochastic michaelis-menten kinetics, Biophysical Chemistry 101 (2002) 565–576.

- [21] K. Sanft, D. Gillespie, L. Petzold, Legitimacy of the stochastic Michaelis-Menten approximation, *IET Systems Biology* 5 (1) (2011) 58–69.
- [22] M. Massink, D. Latella, A. Bracciali, M. Harrison, J. Hillston, Scalable context-dependent analysis of emergency egress models, *Formal Aspects of Computing* 24 (12) (2012) 267–302. doi:10.1007/s00165-011-0188-1.
- [23] N. Gast, B. Gaujal, A mean field model of work stealing in large-scale systems, in: *Proceedings of the 2010 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, 2010, pp. 13–24.
- [24] L. Bortolussi, J. Hillston, D. Latella, M. Massink, Continuous approximation of collective system behaviour: A tutorial, *Performance Evaluation* 70 (5) (2013) 317–349. doi:10.1016/j.peva.2013.01.001.
- [25] L. Bortolussi, J. Hillston, Fluid model checking, in: *Proceedings of the 23rd International Conference on Concurrency Theory, CONCUR’12*, 2012, pp. 333–347.
- [26] H. Tembine, J. L. Boudec, R. El-Azouzi, E. Altman, Mean field asymptotics of markov decision evolutionary games and teams, in: *Proceedings of the First ICST International Conference on Game Theory for Networks, GameNet’09*, IEEE Press, 2009, pp. 140–150.
- [27] R. Hayden, A. Stefanek, J. Bradley, Fluid computation of passage-time distributions in large Markov models, *Theoretical Computer Science* 413 (1) (2012) 106–141.
- [28] J.-P. Katoen, A. Mereacre, Model checking hml on piecewise-constant inhomogeneous Markov chains, in: *Proceedings of the 6th International Conference on Formal Modeling and Analysis of Timed Systems, FORMATS’08*, 2008, pp. 203–217.
- [29] T. Chen, T. Han, J. Katoen, A. Mereacre, LTL model checking of time-inhomogeneous Markov chains, in: *Proceedings of the 7th International Symposium on Automated Technology for Verification and Analysis, ATVA’09*, 2009, pp. 104–119.

- [30] T. Chen, T. Han, J. Katoen, A. Mereacre, Model checking of continuous-time Markov chains against timed automata specifications, *Logical Methods in Computer Science* 7 (1).
- [31] M. Davis, *Markov Models and Optimization*, Chapman & Hall, 1993.
- [32] C. Baier, L. Cloth, B. Haverkort, M. Kurtz, M. Siegle, Model checking markov chains with actions and state labels, *IEEE Trans. on Software Engineering* 33 (4) (2007) 209–224.
- [33] S. Donatelli, S. Haddad, J. Sproston, Model checking timed and stochastic properties with csl^{TA} , *IEEE Trans. on Software Engineering* 35 (2) (2009) 224–240.
- [34] T. Chen, M. Diciolla, M. Kwiatkowska, A. Mereacre, Time-bounded verification of ctmc's against real-time specifications, in: *Proceedings of the 9th International Conference on Formal Modeling and Analysis of Timed Systems, FORMATS'11, 2011*, pp. 26–42.
- [35] J. Hillston, *A Compositional Approach to Performance Modelling*, Cambridge University Press, 1996.
- [36] M. Tribastone, S. Gilmore, J. Hillston, Scalable differential analysis of process algebra models, *IEEE Trans. Software Eng.* 38 (1) (2012) 205–219.
- [37] T. Kurtz, S. Ethier, *Markov Processes - Characterisation and Convergence*, Wiley, 1986.
- [38] A. Jensen, Markov chains as an aid in the study of Markov processes, *Skandinavisk Aktuarietidskrift* 36.
- [39] J. R. Norris, *Markov Chains*, Cambridge University Press, 1997.
- [40] A. P. A. van Moorsel, K. Wolter, Numerical solution of non-homogeneous Markov processes through uniformization, in: *Proceedings of the 12th European Simulation Multiconference - Simulation- Past, Present and Future, ESM'98, 1998*, pp. 710–717.
- [41] W. Rudin, *Principles of Mathematical Analysis*, McGraw-Hill, 1976.

- [42] S. Krantz, P. Harold, *A Primer of Real Analytic Functions* (Second ed.), Birkhäuser, 2002.
- [43] G. Folland, *Introduction to Partial Differential Equations*, Princeton University Press, 1995.
- [44] A. Andreychenko, P. Crouzen, V. Wolf, On-the-fly uniformization of time-inhomogeneous infinite Markov population models, in: *Proceedings Ninth Workshop on Quantitative Aspects of Programming Languages, QAPL 2011*, Vol. 57 of EPTCS, 2011, p. 1.
- [45] P. Billingsley, *Probability and Measure*, 3rd ed., John Wiley and Sons, 1995.
- [46] S. K. Jha, E. M. Clarke, C. J. Langmead, A. Legay, A. Platzer, P. Zuliani, A Bayesian approach to model checking biological systems, in: *Proceedings of the 7th International Conference on Computational Methods in Systems Biology, CMSB 2009*, 2009, pp. 218–234. doi:10.1007/978-3-642-03845-7_15.
- [47] A. Neumaier, *Interval Methods for Systems of Equations*, University Press, Cambridge, 1990.
- [48] G. Alefeld, G. Mayer, Interval analysis: theory and applications, *Journal of Computational and Applied Mathematics* 121 (2000) 421–464.
- [49] P. Franek, S. Ratschan, P. Zgliczynski, Satisfiability of systems of equations of real analytic functions is quasi-decidable, in: *Proceedings of the 36th international conference on Mathematical foundations of computer science, MFCS'11*, 2011, pp. 315–326.
- [50] A. Aziz, K. Sanwal, V. Singhal, R. Brayton, Model-checking continuous time Markov chains, *ACM Trans. Comp. Logic* 1 (2000) 162–170.
- [51] D. Richardson, Zero tests for constants in simple scientific computation, *Mathematics in Computer Science* 1 (1) (2007) 21–37.
- [52] D. Richardson, *Effective methods in algebraic geometry*, Birkhäuser, 1991, Ch. Finding roots of equations involving functions defined by first order differential equations.

- [53] T. Johnson, W. Tucker, Enclosing all zeros of an analytic function — a rigorous approach, *Journal of Computational and Applied Mathematics* 228 (1) (2009) 418–423. doi:10.1016/j.cam.2008.10.014.
- [54] L. Ahlfors, *Complex Analysis*, 1st ed., McGraw Hill, Cambridge, 1953.
- [55] D. Novikov, Systems of linear ordinary differential equations with bounded coefficients may have very oscillating solutions, *ArXiv Mathematics e-prints* arXiv:arXiv:math/0007110.
- [56] P. Billingsley, *Convergence of probability measures*, Wiley, 1999.
URL <http://site.ebrary.com/lib/alltitles/docDetail.action?docID=10344091>
- [57] M. Benaïm, J. Weibull, Deterministic approximation of stochastic evolution in games, *Econometrica* 71 (3) (2003) 873–903.
- [58] M. Benaïm, Recursive algorithms, urn processes and chaining number of chain recurrent sets, *Ergodic Theory and Dynamical Systems*.
- [59] M. Benaïm, J. L. Boudec, On mean field convergence and stationary regime, *CoRR* abs/1111.5710.
- [60] L. Bortolussi, J. Hillston, Checking Individual Agent Behaviours in Markov Population Models by Fluid Approximation, Vol. 7938 of LNCS, Springer, 2013, pp. 113–149. doi:10.1007/978-3-642-38874-3_4.
- [61] A. Kolesnichenko, P.-T. de Boer, A. Remke, B. R. Haverkort, A logic for model-checking mean-field models, in: *Proceedings of 43rd Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN'13*, 2013, pp. 1–12.
- [62] L. Bortolussi, Hybrid limits of continuous time markov chains, in: *Proceedings of the Eighth International Conference on Quantitative Evaluation of Systems, QEST'11*, IEEE Computer Society, Aachen, Germany, 2011, pp. 3–12. doi:10.1109/QEST.2011.10.
- [63] L. Grunske, Specification patterns for probabilistic quality properties, in: *Proceedings of the ACM/IEEE 30th International Conference on Software Engineering, ICSE'08*, 2008, pp. 31–40.

- [64] P. Taylor, A lambda calculus for real analysis, *Journal of Logic and Analysis* 2 (5) (2010) 1–115. doi:10.4115/jla.2010.2.5.

Appendix A. Proofs

In this appendix, we present the proofs of propositions, lemmas, and theorems of the paper. We will start by showing measurability of next state probabilities and reachability probabilities for ICTMC, turning then to prove their convergence for the fluid approximation. This will provide some key tools that are helpful to prove the quasi-decidability of CSL model checking for ICTMC. Convergence of CSL model checking will come next. The final part of the appendix will be devoted to showing the results of Section 7.

Appendix A.1. Measurability of Path Sets for ICTMC

Fix a ICTMC $Z(t)$ on state space \mathcal{S} with infinitesimal generator matrix $Q(t)$. We will start by recalling the definition of the sigma-algebra on the set *Paths* of paths of $Z(t)$. Let \mathcal{I} be the set of non-empty intervals of $\mathbb{R}_{\geq 0}$ with rational endpoints. For $I \in \mathcal{I}$, denote T_I^- and T_I^+ the upper and lower bounds of I . The sigma-algebra \mathcal{F} on *Paths* is the smallest sigma-algebra containing all the cylinder sets $C_{t_0}(s_0, I_0, s_1, \dots, I_{n-1}, s_n)$, consisting of all paths that are in s_0 at time t_0 and that jump to s_j , $1 \leq j \leq n$ at a time $t \in t_0 \oplus I_0 \oplus \dots \oplus I_{j-1}$, where \oplus is the Minkowsky sum. The collection of cylinder sets starting at time t_0 is called \mathcal{C}_{t_0} and it is countable. The collection of all cylinder sets is denoted by \mathcal{C} . In accordance to the notation of Section 5, the probability of a cylinder set $C_{t_0}(s_0, I_0, s_1, \dots, I_{n-1}, s_n)$ is defined recursively as

$$\mathbb{P}(C_{t_0}(s_0, I_0, \dots, s_n)) = \int_{t_0 \oplus I_0} q_{s_0, s_1}(t) e^{-\Lambda(t_0, t)[s_0]} \mathbb{P}(C_t(s_1, I_1, \dots, s_n)) dt$$

We now list some basic measurability properties for ICTMCs, which will be used in the following.

Proposition Appendix A.1. *Let $Z(t)$, *Paths*, and \mathcal{F} be defined as above. Call Σ^+ the set of sequences of states \mathcal{S} of finite length. The following facts hold:*

1. *The function $\tau_{t_0, n} : \text{Paths} \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ returning the time of the n -th jump after t_0 , relative to t_0 (and ∞ for Zeno paths exploding before t_0) is measurable.*
2. *For $W \subseteq \Sigma^+$, the set $\text{Paths}(t_0, W)$ of paths such that their sequence of states, starting in t_0 , belongs to W is measurable.*

3. For $W \subseteq \Sigma^+$, the set $Paths(t_0, t_1, W)$ of paths such that their sequence of states during time (t_0, t_1) belongs to W is measurable.

Proof. To prove point 1, note that, as the intervals \mathcal{I} generate the Borel sigma-algebra on $\mathbb{R}_{\geq 0}$, we just need to show that $\tau_{t_0, n}^{-1}(I) \in \mathcal{F}$ for $I \in \mathcal{I}$.²⁴ Given a cylinder set $C_{t_0}(s_0, I_0, \dots, s_n)$, denote by T_n^- the minimum time of the n -th jump, relative to t_0 (i.e. $T_n^- = \sum_{j < n} T_{I_j}^-$) and with T_n^+ the maximum time of the n -th jump. Let $\mathcal{C}_{t_0, n, I} \subset \mathcal{C}$ contain all cylinder sets $C_{t_0}(s_0, I_0, \dots, s_n)$ such that $T_n^- = T_I^-$ and $T_n^+ = T_I^+$, i.e. such that their n -th jump happens at time $t \in I$, relative to t_0 . Then $\tau_{t_0, n}^{-1}(I) = \bigcup_{C \in \mathcal{C}_{t_0, n, I}} C \in \mathcal{F}$ (as \mathcal{C}_{t_0} is countable).

Point 2 follows from the fact that W is countable, and that we can express $Paths(t_0, W)$ as the union of the cylinder sets of \mathcal{C}_{t_0} whose sequence of states is in W .

Point 3 follows similarly to point 2, expressing $Paths(t_0, t_1, W)$ as the union of cylinder sets $C_{t_0}(s_0, I_0, \dots, s_n, I_n, s_{n+1})$ of \mathcal{C}_{t_0} such that $s_0 \dots s_n \in W$ and $T_{n+1}^- \geq t_1 - t_0$.²⁵ ■

We are now ready to prove the propositions about measurability in the paper.

Proposition (5.1). *Let $G : [t_0, t_1] \times \mathcal{S} \rightarrow \{0, 1\}$ be a time-dependent set and $Z(t)$ an ICTMC. Then $Paths_{next}(Z, s_0, t_0, T_1, T_2, G)$ is measurable.*

Proof. Fix $s \in \mathcal{S}$ and let $W_{s_0, s} = \{s_0 s\} \subset \Sigma^+$ be the singleton set containing the sequence $s_0 s$. Furthermore, consider the measurable function $G(\cdot, s) : [t_0, t_1] \rightarrow \{0, 1\}$. Then, the set $A_s = G(\cdot, s)^{-1}([t_0 + T_1, t_0 + T_2]) \ominus t_0$, containing the time instants relative to t_0 such that s is in G , is measurable. We can express $Paths_{next}(Z, t_0, T_1, T_2, G)$ as

$$Paths_{next}(Z, s_0, t_0, T_1, T_2, G) = \bigcup_{s \in \mathcal{S}} (\tau_{t_0, 1}^{-1}(A_s) \cap Paths(t_0, W_{s_0, s})).$$

²⁴The set of paths exploding before t_0 is also measurable, so we avoid treating it explicitly for simplicity.

²⁵To be more precise, we need to treat separately the case of trajectories whose $n + 1$ -th jump after t_0 happens exactly at time t_1 . In fact, $t_1 - t_0$ can be irrational, hence necessarily $T_{n+1}^- > t_1 - t_0$. But such a set of trajectories is measurable due to point 1, hence this poses no problem.

The expression on the right is measurable, and it is the union for $s \in \mathcal{S}$ of the set of paths such that the first jump after t_0 ends in state s and happens at a time in which s is in G . \blacksquare

Proposition (5.2). *Let $G, U : [t_0, t_1] \times \mathcal{S} \rightarrow \{0, 1\}$ be time-dependent set of finite-variability and $Z(t)$ an ICTMC. Then $Paths_{reach}(Z, s, t_0, T, G, U)$, $s \in \mathcal{S}$, is measurable.*

Proof. Let $T_0 = t_0, T_1, T_2, \dots, T_n = T_0 + T$ be all the time instants in which G or U are discontinuous for some state $s \in \mathcal{S}$. Hence, in the interval (T_i, T_{i+1}) , the sets G and U are constant. Now we define the following subsets of paths: $Safe(I)$, of paths that visit only safe and non-goal states during the interval $I \subset [t_0, t_1]$, and $SafeGoal(I)$, the set of paths that, when restricting to times in I , visit safe sets and then reach a goal state. By proposition Appendix A.1, these sets are measurable for any $I \subseteq (T_i, T_{i+1})$ (just define the appropriate subsets W of sequences Σ^+ , i.e. made of all safe and non-goal sets, or of safe sets until reaching a goal one, and visiting arbitrary states afterwards). By measurability of finite dimensional projections, $Safe(I)$ and $SafeGoal(I)$ are measurable also for single time instants, $I = \{t\}$. Then also

$$Safe_i = Safe_{i-1} \cap Safe(\{T_{i-1}\}) \cap Safe((T_{i-1}, T_i)),$$

$0 < i \leq n$, is measurable for $0 < i \leq n$, with $Safe_0 = Safe(\{T_0\})$, and so are

$$Reach_i = Safe_i \cap ((Safe(\{T_i\}) \cap SafeGoal((T_i, T_{i+1}))) \cup SafeGoal(\{T_i\}))$$

for $0 \leq i < n$ and

$$Paths_{reach}(Z, s, t_0, T, G, U) = \left(\bigcup_{j < n} Reach_j \right) \cup (Safe_n \cap SafeGoal(\{T_n\})).$$

\blacksquare

Proposition (5.3). *Let $Z(t)$ be a ICTMC with piecewise analytic time-dependent rate matrix $Q(t)$. Then*

1. *The time-dependent set of states $\llbracket \varphi \rrbracket = \llbracket \varphi \rrbracket(t)$ that satisfy a CSL formula φ has the finite variability property.*

2. The set of paths $Paths(s, t_0, \psi)$ that satisfy a CSL path formula φ starting in state s at time t_0 is measurable.

Proof. The proof proceeds by structural induction on formulae. Point 1 for atomic boolean state formulae is trivial, while for quantified state formulae follows by the piecewise analyticity of the time dependent probabilities $P(t)$ for next and until path formulae, which guarantees that the inequality $P(s, t) - p \bowtie 0$ changes truth status a finite number of times in any bounded interval. Point 2 for next path formulae follows from Proposition 5.1, while for until path formulae follows from a straightforward modification of Proposition 5.2. ■

Appendix A.2. Convergence of Next-State Probability

Proposition (6.1). *Let $f : I \rightarrow \mathbb{R}$ be a piecewise analytic function, with $I \subseteq \mathbb{R}$ a compact interval. Let $E_f = \{x \in \mathbb{R} \mid \mu_\ell(f^{-1}(\{x\})) = 0\}$ be the set of all values x such that f is not locally constantly equal to x , where μ_ℓ is the Lebesgue measure. Furthermore, let $Z_x = f^{-1}(\{x\})$ be the set of solutions of $f(t) = x$ and let $DZ_f = \{x \in \mathbb{R} \mid \forall t \in Z_x, f'(t) \neq 0\}$. Then*

1. $\forall x \in E_f, Z_x$ is finite.
2. $\mu_\ell(E_f \cap DZ_f) = 1$

Proof. Point 1 follows from basic properties of the piecewise analytic function $(f - x)$: in any analytic piece, either the function is constantly equal to zero, or it has only a finite number of zeros. Point 2, instead, follows from the fact that the derivative $f'(t)$ of t is piecewise analytic, hence has only a finite number of zeros (in the analytic pieces in which f is not constant). ■

Proposition (6.2). *Let $V^{(N)}(t)$ be a sequence of time varying sets converging robustly to a robust set $V(t), t \in I$. Let $D_V^{(N)} = \{t \mid V^{(N)}(t) \neq V(t)\}$. Then $\mu_\ell(D_V^{(N)}) \rightarrow 0$, where μ_ℓ is the Lebesgue measure on \mathbb{R} .*

Proof. A straightforward consequence of the definition of robust convergence is that, for each open neighbourhood U of $Disc(V)$, there exists an N_0 such that, for all $N \geq N_0$, $V^{(N)}(t) = V(t)$ for $t \in I \setminus U$. Now, as V is robust, then $|Disc(V)| = m < \infty$. Fix $\varepsilon > 0$ and define $U_\varepsilon = \bigcup_{\bar{t} \in Disc(V)} B(\bar{t}, \varepsilon)$, where $B(\bar{t}, \varepsilon)$ is the open ball centred at \bar{t} of radius ε . Then $\mu_\ell(U_\varepsilon) \leq 2m\varepsilon$. Now, fix $\varepsilon_k \rightarrow 0$. For each k , there is an N_k such that, for all $N \geq N_k$, $V^{(N)}(t) = V(t)$

for $t \in I \setminus U_{\varepsilon_k}$, and therefore $D_V^{(N)} \subseteq U_{\varepsilon_k}$. ■

Lemma (6.1). *Let $\mathcal{X}^{(N)}$ be a sequence of CTMC models, as defined in Section 3.1, and let $Z_k^{(N)}$ and z_k be defined from $\mathcal{X}^{(N)}$ as in Section 3.3, with piecewise real analytic rates, in a compact interval $[0, T']$, for $T' > t_1 + T_b$.*

Let $G(t)$, $t \in [t_0, t_1 + T_b]$ be a robust time-varying set, and let $G^{(N)}(t)$ be a sequence of time-varying sets converging robustly to G .

Furthermore, let $\bar{P}(t) = P_{next}(z_k, t, T_a, T_b, G)$ and $\bar{P}^{(N)}(t) = P_{next}(Z_k^{(N)}, t, T_a, T_b, G)$, $t \in [t_0, t_1]$.

Finally, fix $p \in [0, 1]$, $\bowtie \in \{\leq, <, >, \geq\}$, and let $V_p(t) = \mathbf{1}\{\bar{P}(t) \bowtie p\}$, $V_p^{(N)}(t) = \mathbf{1}\{\bar{P}^{(N)}(t) \bowtie p\}$. Then

1. $\bar{P}^{(N)}(t) \rightarrow \bar{P}(t)$, uniformly in $t \in [t_0, t_1]$.
2. For almost every $p \in [0, 1]$, V_p is robust and the sequence $V_p^{(N)}$ converges robustly to V_p .

Proof. By a standard coupling argument, assume that z_k and $Z_k^{(N)}$ are defined on the same probability space Ω . Then, letting Y be either z_k or $Z_k^{(N)}$, for $\omega \in \Omega$, let $\chi(t, Y(\omega))$ be equal to one if trajectory $Y(\omega)$'s first jump, starting at time t , is into a state of G at time $t' \in [t + T_a, t + T_b]$, and zero otherwise. Similarly, let $\chi^{(N)}(t, Y(\omega))$ be 1 if $Y(\omega)$'s first jump, starting at time t , is into $G^{(N)}$ at time $t' \in [t + T_a, t + T_b]$, and zero otherwise. Then $\bar{P}(t) = \mathbb{E}[\chi(t, z_k)]$, and $\bar{P}^{(N)}(t) = \mathbb{E}[\chi^{(N)}(t, Z_k^{(N)})]$. It follows that

$$\begin{aligned} |\mathbb{E}[\chi(t, z_k)] - \mathbb{E}[\chi^{(N)}(t, Z_k^{(N)})]| &\leq \underbrace{\mathbb{E}[|\chi(t, z_k) - \chi^{(N)}(t, z_k)|]}_{(1)} \\ &+ \underbrace{\mathbb{E}[|\chi^{(N)}(t, z_k) - \chi^{(N)}(t, Z_k^{(N)})|]}_{(2)} \end{aligned}$$

Consider term (2) above. We can partition trajectories into two measurable subsets: $\Omega_1 = \{\omega \in \Omega \mid z_k(t, \omega) = Z_k^{(N)}(t, \omega), t \leq t_1 + T_b\}$ and $\Omega_0 = \Omega \setminus \Omega_1$. Let μ_Ω be the probability measure in Ω . Applying Theorem 3.2 up to time $t_1 + T_b$, we have that $\chi^{(N)}(t, Z_k^{(N)}(\omega)) = \chi^{(N)}(t, z_k(\omega))$ for $\omega \in \Omega_1$

and $\mathbb{P}(\Omega_0) \leq \varepsilon_N$. Hence, for any $t \in [t_0, t_1]$,

$$\begin{aligned} \mathbb{E}[|\chi^{(N)}(t, Z_k^{(N)}) - \chi^{(N)}(t, z_k)|] &= \int_{\Omega_1} |\chi^{(N)}(t, Z_k^{(N)}) - \chi^{(N)}(t, z_k)| d\mu_\Omega \\ &+ \int_{\Omega_0} |\chi^{(N)}(t, Z_k^{(N)}) - \chi^{(N)}(t, z_k)| d\mu_\Omega \\ &\leq \varepsilon_N \rightarrow 0. \end{aligned}$$

Notice that ε_N does not depend on t .

Let us focus now on term (1) in the inequality above. Let $T_1 < T_2 < \dots < T_h$ be all the points in $Disc(G)$ (which are finite in number as G is robust). Fix $t \in [t_0, t_1]$. As $G^{(N)}$ converges robustly to G , for $N \geq N_0$ they differ only in disjoint balls $B(T_i, \varepsilon)$, for ε small enough. Furthermore, if G has a discontinuity for state s in T_i , then the value of G on the left of $B(T_i, \varepsilon)$ is different from the value of G on the right of $B(T_i, \varepsilon)$.

It follows that the only trajectories of z_k for which $\chi(t, z_k) \neq \chi^{(N)}(t, z_k)$ are those jumping within the set $D_G^{(N)}$ (intersected with $[t, t + T_b]$).²⁶ As the rate functions of z_k are piecewise analytic, they are bounded by a constant Λ , thus the probability of a trajectory jumping in $D_G^{(N)}$ is bounded by $\int_{D_G^{(N)}} \Lambda e^{-\Lambda t} dt \leq \int_{D_G^{(N)}} \Lambda dt = \Lambda \mu_\ell(D_G^{(N)}) \rightarrow 0$ (independently of t). It follows that

$$|\mathbb{E}[\chi(t, z_k)] - \mathbb{E}[\chi^{(N)}(t, Z_k^{(N)})]| \leq \delta_N,$$

with $\delta_N = \varepsilon_N + \Lambda \mu_\ell(D_G^{(N)}) \rightarrow 0$ independently of t , which proves uniform convergence of $P^{(N)}(t)$ to $P(t)$.

Let us turn now to point 2 of the lemma.

Consider the set $H_{\bar{P}}$ of values $p \in [0, 1]$ for which either (i) $\bar{P}(t)$ is constantly equal to p in one analytic piece of \bar{P} , or (ii) $\bar{P}(t) = p$ and $\bar{P}'(t) = 0$ for some t , or (iii) $\bar{P}(t) = p$ and \bar{P} is not analytic in t . By Prop. 6.1 and the definition of piecewise analytic functions, the set $H_{\bar{P}}$ is finite. Fix a $p \notin H_{\bar{P}}$. For such a p , the function $\bar{P}(t) - p$ defines a robust time-varying set, as it has a finite number of simple zeros, all in analytic points of \bar{P} .

Call A the set of points in which V_p has a discontinuity, which is finite. Fix ε and define A_ε to be $\bigcup_{t \in A} B(t, \varepsilon)$, where $B_\varepsilon(t) = (t - \varepsilon, t + \varepsilon)$. Now, if W is a neighbourhood of A , then for a small $\varepsilon > 0$, $A_\varepsilon \subset W$. Let $f_p(t) = |\bar{P}(t) - p|$

²⁶Notice that robustness of G is not necessary for this proof, but we enforce it for uniformity with the convergence of reachability probabilities in Section 5.3.

and consider the set $I_\varepsilon = I \setminus A_\varepsilon$. Now, I_ε is compact and $f_p(t)$ is different from zero in I_ε , so that $\min\{f_p(t) \mid t \in I_\varepsilon\} = m_\varepsilon > 0$ (by the Weierstrass Theorem [41]). As $\bar{P}^{(N)}$ converges uniformly to \bar{P} , there is N_0 such that, for all $N \geq N_0$ and all $t \in I_\varepsilon$, $|\bar{P}^{(N)}(t) - \bar{P}(t)| \leq \frac{m_\varepsilon}{2}$, hence for all $N \geq N_0$ and all $t \in I_\varepsilon$, $V_p(t) = V_p^{(N)}(t)$. It follows that $V_p^{(N)}(t)$ converges robustly to $V_p(t)$. ■

Appendix A.3. Convergence of Reachability Probability

Proposition (6.3). *Let $\mathcal{X}^{(N)}$ be a sequence of CTMC models, as defined in Section 3.1, and let $Z_k^{(N)}$ and z_k be defined from $\mathcal{X}^{(N)}$ as in Section 3.3. Assume that the infinitesimal generator matrix $Q(t)$ of z_k is bounded and integrable in every compact interval $[0, T]$. Then*

$$P_{reach}(Z_k^{(N)}, t, T, G, U) \rightarrow P_{reach}(z_k, t, T, G, U), \text{ uniformly in } [t_0, t_1], \text{ as } N \rightarrow \infty$$

i.e. $\sup_{t \in [t_0, t_1]} \|P_{reach}(Z_k^{(N)}, t, T, G, U) - P_{reach}(z_k, t, T, G, U)\| \rightarrow 0$.

Proof. The proposition follows from a similar argument to that used in the first part of the proof of Lemma 6.1. By a standard coupling argument, we can assume that the processes $Z_k^{(N)}$ and z_k are defined on the same probability space Ω . Therefore, there exists a sequence $\varepsilon_N \in \mathbb{R}_+$, $\varepsilon_N \rightarrow 0$, such that $\mathbb{P}\{\omega \in \Omega \mid \forall t \leq T', Z_k^{(N)}(\omega, t) = z_k(\omega, t)\} \geq 1 - \varepsilon_N$. This means that with probability $1 - \varepsilon_N$, the trajectories of the two processes are the same up to time T' .

Now, we can define a (measurable) function $\chi = \chi_{t, T, G, U}$ on the trajectories of the CTMCs which is equal to 1 if they satisfy the reachability property, and 0 otherwise. Therefore, it holds that $P_{reach}(Z_k^{(N)}, t, T, G, U) = \mathbb{E}[\chi(Z_k^{(N)})]$, and similarly for z_k . With a similar notation as in Lemma 6.1, let $\Omega_1 = \{\omega \mid Z_k^{(N)}(t, \omega) = z_k(t, \omega), \forall t \leq t_0 + T\}$, $\Omega_0 = \{\omega \mid Z_k^{(N)}(t, \omega) \neq z_k(t, \omega)\}$, and μ_Ω be the probability measure in Ω (i.e. in the trajectory space). Observe that $\chi(Z_k^{(N)}) = \chi(z_k)$ on Ω_1 and $\mathbb{P}(\Omega_0) \leq \varepsilon_N$, hence

$$\begin{aligned} |\mathbb{E}[\chi(Z_k^{(N)})] - \mathbb{E}[\chi(z_k)]| &\leq \mathbb{E}[|\chi(Z_k^{(N)}) - \chi(z_k)|] \\ &= \int_{\Omega_1} |\chi(Z_k^{(N)}) - \chi(z_k)| d\mu_\Omega \\ &\quad + \int_{\Omega_0} |\chi(Z_k^{(N)}) - \chi(z_k)| d\mu_\Omega \\ &\leq \varepsilon_N \rightarrow 0. \end{aligned}$$

Uniform convergence follows from the fact that the sequence ε_N does not depend on the initial or the final time of the reachability property, if they are both less than $T + t_1$. \blacksquare

Lemma (6.2). *Let $\mathcal{X}^{(N)}$ be a sequence of CTMC models, as defined in Section 3.1, and let $Z_k^{(N)}$ and z_k be defined from $\mathcal{X}^{(N)}$ as in Section 3.3, with piecewise analytic rates, in a compact interval $[0, T']$, for T' sufficiently large.*

Let $G(t), U(t), t \in [t_0, t_1 + T]$ be compatible and robust time-varying sets, and let $G^{(N)}(t), U^{(N)}(t)$ be sequences of time-varying sets converging robustly to G and U , respectively.

Furthermore, let $P(t) = P_{reach}(z_k, t, T, G, U)$ and

$P^{(N)}(t) = P_{reach}(Z_k^{(N)}, t, T, G^{(N)}, U^{(N)})$, $t \in [t_0, t_1]$.

Finally, fix $p \in [0, 1]$, $\bowtie \in \{\leq, <, >, \geq\}$, and let $V_p(t) = \mathbf{1}\{P(t) \bowtie p\}$, $V_p^{(N)}(t) = \mathbf{1}\{P^{(N)}(t) \bowtie p\}$. Then

1. *For all but finitely many $t \in [t_0, t_1]$, $P^{(N)}(t) \rightarrow P(t)$, with uniform speed (i.e. independently of t).*
2. *For almost every $p \in [0, 1]$, V_p is robust and the sequence $V_p^{(N)}$ converges robustly to V_p .*

Proof. As in the proof of Lemma 6.1, by a standard coupling argument assume that z_k and $Z_k^{(N)}$ are defined on the same probability space Ω . Then, letting Y be either z_k or $Z_k^{(N)}$, for $\omega \in \Omega$, let $\chi(t, Y(\omega))$ be equal to 1 if trajectory $Y(\omega)$ satisfies the reachability problem with respect to G and U and zero otherwise, starting at time t and $\chi^{(N)}(t, Y(\omega))$ be 1 if $Y(\omega)$ satisfies the reachability problem for $G^{(N)}, U^{(N)}$, and zero otherwise, starting at time t . Then $P(t) = \mathbb{E}[\chi(t, z_k)]$, and $P^{(N)}(t) = \mathbb{E}[\chi^{(N)}(t, Z_k^{(N)})]$, and

$$\begin{aligned} |\mathbb{E}[\chi(t, z_k)] - \mathbb{E}[\chi^{(N)}(t, Z_k^{(N)})]| &\leq \underbrace{\mathbb{E}[|\chi(t, z_k) - \chi^{(N)}(t, z_k)|]}_{(1)} \\ &+ \underbrace{\mathbb{E}[|\chi^{(N)}(t, z_k) - \chi^{(N)}(t, Z_k^{(N)})|]}_{(2)} \end{aligned}$$

Term (2) above is bounded by $\varepsilon_N \rightarrow 0$, as in Lemma 6.1, by a straightforward application of Theorem 3.2. Term (1) is also treated similarly to Lemma 6.1, with an extra argument to deal with pointwise discontinuities in the reachability probability. Let $T_1 < T_2 < \dots < T_h$ be all the points

in $Disc(G) \cup Disc(U)$ (which is finite as G and U are robust). If we suppose neither t nor $t + T$ coincide with one of the previous points (i.e. all discontinuities are internal in the time domain), then by robust convergence of $G^{(N)}$ (resp. $U^{(N)}$) to G (resp. U), for $N \geq N_0$ they differ only in small disjoint balls $B(T_i, \varepsilon)$ internal to $[t, t + T]$. Reasoning as in Lemma 6.1, it follows that the only trajectories of z_k for which $\chi(t, z_k) \neq \chi^{(N)}(t, z_k)$ are those jumping within the set $D^{(N)} = D_G^{(N)} \cup D_U^{(N)}$.²⁷ As the rate functions of z_k are piecewise analytic, they are bounded by a constant Λ , thus the probability of a trajectory jumping in $D^{(N)}$ is bounded by $\int_{D^{(N)}} \Lambda e^{-\Lambda t} dt \leq \int_{D^{(N)}} \Lambda dt = \Lambda \mu_\ell(D^{(N)}) \rightarrow 0$ (independently of t). It follows that, if $t \notin T_d$, with $T_d = \{T_1, \dots, T_h, T_1 - T, \dots, T_h - T\}$, then

$$|\mathbb{E}[\chi(t, z_k)] - \mathbb{E}[\chi^{(N)}(t, Z_k^{(N)})]| \leq \delta_N,$$

with $\delta_N = \varepsilon_N + \Lambda \mu_\ell(D^{(N)}) \rightarrow 0$.

On the contrary, if $t \in T_d$, then a discontinuity of G or U happens exactly at the boundary of the time domain $[t, t + T]$ in which we have to verify the formula. In this case, the value of sets $G^{(N)}$ and G (or $U^{(N)}$ and U) may never be the same at this extreme point t^* , whatever small neighbourhood of t^* in $[t, t + T]$ one takes into account (e.g. if $t^* = T_i$ is the left extreme of the time domain, it may happen that all changes of $G^{(N)}$ occur before this point). Therefore, there can be a set of trajectories of measure > 0 that are accepted by $\chi^{(N)}$ and refused by χ (or vice versa). In particular, this can happen if and only if $P(t)$ has a discontinuity in one of those points (otherwise, convergence follows by continuity). Hence, in these time points, convergence may not hold. However, the set T_d is finite, hence point 1 of the Lemma is proved.

Let us turn now to point 2 of the lemma, which is similar to Lemma 6.1, with extra care for the discontinuities of P .

As in Lemma 6.1, construct the set $H_{\bar{P}}$ of values $p \in [0, 1]$ for which either (i) $\bar{P}(t)$ is constantly equal to p in one analytic piece of \bar{P} , or (ii) $\bar{P}(t) = p$ and $\bar{P}'(t) = 0$ for some t , or (iii) $\bar{P}(t) = p$ and \bar{P} is not analytic in t . This

²⁷If G or U are not robust, then even if they have a finite number of discontinuity points, the previous argument may not hold. In fact, they may have a discontinuity point T_i such that $G_s(T_i) = 1$ but $G_s(t) = 0$ in a neighbourhood $W \setminus \{T_i\}$ of T_i . In this case, it is possible that $G_s^{(N)}(t) = 0$ on all W , which implies that $\chi(t, z_k) \neq \chi^{(N)}(t, z_k)$ for all those trajectories that are in state s at time T_i .

set is finite, and for $p \notin H_{\bar{P}}$, the function $\bar{P}(t) - p$ is easily seen to define a robust time-varying set, as it has a finite number of simple zeros, all in analytic points of \bar{P} . Consider now the set A of discontinuity points of V_p . Fix ε and define A_ε to be $\bigcup_{t \in A} B(t, \varepsilon)$, where $B_\varepsilon(t) = (t - \varepsilon, t + \varepsilon)$. By reasoning as in the last part of the proof of Lemma 6.1, letting $I_\varepsilon = I \setminus A_\varepsilon$ and $\min\{|P(t) - p| \mid t \in I_\varepsilon\} = m_\varepsilon > 0$, as $P^{(N)}(t)$ converges in I_ε to $P(t)$ with uniform speed, there is N_0 such that, for all $N \geq N_0$ and all $t \in I_\varepsilon$, $|P^{(N)}(t) - P(t)| \leq \frac{m_\varepsilon}{2}$, hence for all $N \geq N_0$ and all $t \in I_\varepsilon$, $V_p(t) = V_p^{(N)}(t)$. It follows that $V_p^{(N)}(t)$ converges robustly to $V_p(t)$.

However, here we need extra care as the set I_ε may contain time instants \tilde{t} in which the convergence of $P^{(N)}$ to P does not hold, but that do not generate a discontinuity in V_p , because $P(\tilde{t}^+)$ and $P(\tilde{t}^-)$ are both greater or both less than p . These points do not create problems, essentially because the function $P^{(N)}$, for N large, remains close to P . In fact, convergence at \tilde{t} fails because the jumps in $G^{(N)}$ and $U^{(N)}$ happen at time instants converging to the ones of jumps in G and U , but not necessarily at \tilde{t} . This slightly puts out of synchronization the time at which the discontinuity happens, but the values of $P^{(N)}$ and P around such a discontinuity are close. This implies that, for N large, $P^{(N)}$ will remain below p if both $P(\tilde{t}^+)$ and $P(\tilde{t}^-)$ are below it, and similarly for the symmetric case.

A formalisation of this argument requires a more careful inspection of the behaviour of $G^{(N)}$ (respectively $U^{(N)}$) near a discontinuity of G (respectively U), and a clarification of the connection between discontinuities in G and U and discontinuities in P . For the former point, note that by the robust convergence property of $G^{(N)}$ to G , if G has a discontinuity for state s at time t , say from 0 to 1, then $G^{(N)}$ also has a discontinuity of the same kind near t . In fact, it can do more than one jump around t , but for sure, for any small $\varepsilon > 0$ and N large, it will equal 0 before $t - \varepsilon$ and 1 after $t + \varepsilon$. The point is that these additional jumps do not matter, as they happen so close to each other that almost no probability mass moves in between, hence they have a vanishing effect on $P^{(N)}$ (as N grows). As for the connection between discontinuities in G and U and the function P , observe that we can have a discontinuity in P at time t only if either G or U has a discontinuity at time t or at time $t + T$ (they cannot both have such a discontinuity, due to the compatibility condition). There are many cases to take into account (a change from goal to non-goal, or from non-goal to goal, and so on), but only a few of them induce a discontinuity, specifically a change from non-goal to

goal of a safe state s at time $t + T$ (inducing a discontinuity in any safe and non-goal state at time t), a change from goal to non-goal of a safe state s or from unsafe to safe of a non-goal state s at time t (inducing a discontinuity in s), and a change in the goal status of an unsafe state at time t . In the first case, we can have a discontinuous increase in P . In the second case, the value of P in s can drop from 1 to a value $p' < 1$. In the third case, the value of P can increase from 0 to a value $p' > 0$. In the fourth case, which is a rather strange case, the value of P changes from 0 to 1, or vice versa.

To understand the connection between P , $P^{(N)}$, G and $G^{(N)}$, consider a situation of the first kind, in which one or more safe states s change from non-goal to goal at time $t + T$. This creates a discontinuity in $P(t)$ for any safe state s' , such that there is a non null-probability of going from s' to s along a safe and non-goal path from t to $t + T$. This probability, in fact, is added to $P_{s'}(t)$, according to the discussion in Section 5.3. Suppose for simplicity that only a single state s changes status in $t + T$ from non-goal to goal. Then this happens close to $t + T$ also in $G^{(N)}$. In fact, s can change status many times in $G^{(N)}$, near $t + T$, but only the first one really matters for the discontinuity of $P^{(N)}$. This happens because the probability added to $P_{s'}^{(N)}$ for subsequent jumps of s from non-goal to goal state close to $t + T$ is only the probability of jumping into s from another safe state in the short time interval in which s is non-goal. To be more concrete, if s jumps from non-goal to goal at time $t_1^{(N)} + T$, then from goal to non-goal at time $t_2^{(N)} + T$ and back to goal at time $t_3^{(N)} + T$, then the probability added to $P_{s'}^{(N)}(t_3^{(N)})$ is bounded by the amount of probability mass that can flow into s in between times $t_2^{(N)} + T$ and $t_3^{(N)} + T$, which is of the order of $t_3^{(N)} - t_2^{(N)}$. Hence it vanishes as N grows (as $t_2^{(N)}$ and $t_3^{(N)}$ collapse to t). More precisely, if Λ is an upper bound for the exit rate of the single agent (uniform in N , which can be found as the exit rate of $Z_k^{(N)}$ converges to the exit rate of z_k , which is itself uniformly bounded), then the jump size at time $t_3^{(N)}$ is bounded by $2\Lambda(t_3^{(N)} - t_2^{(N)})$. Furthermore, the speed at which $P_{s'}^{(N)}$ can increase or decrease, excluding jumps, is also bounded by 2Λ , so that the value of $P_{s'}^{(N)}$ cannot vary too much after the first jump in a small time interval of size Δ_t around t . In fact, combining these two arguments, the total variation (after the first jump) is bounded by $2\Lambda\Delta_t$.²⁸ Note that, if more than one

²⁸The factor 2 comes from the fact that we are working with a combination of the backward and forward equation, both giving an upper bound of Λ on the rate of change

safe state changes goal status at time $t + T$ in G , then in $G^{(N)}$ these events can happen asynchronously, hence to see the full increase in $P_{s'}^{(N)}$ we need to wait until all those states have changed value in $G^{(N)}$. Yet the bound in terms of Λ and Δ_t remains valid. The other discontinuous jump types are treated analogously (with the exception of the jump from 0 to 1 or from 1 to 0, which however contains any threshold p in its interior). We can now give a formal argument that discontinuities in I_ε are not a problem. Assume that P has a discontinuity at \tilde{t} such that $\mu = \max\{P_{s'}(\tilde{t}^+), P_{s'}(\tilde{t}^-)\} < p$, and call $\varepsilon = p - \mu$. Now, choose δ such that $4\delta\Lambda < \varepsilon/4$, $\|P_{s'}(\tilde{t} - \delta) - P_{s'}(\tilde{t}^-)\| < \varepsilon/4$, and $\|P_{s'}(\tilde{t} + \delta) - P_{s'}(\tilde{t}^+)\| < \varepsilon/4$. Then, choose an N_0 such that, for $N \geq N_0$, all the jumps of $G^{(N)}$ are closer than δ to the jumps of G , and such that $\|P_{s'}^{(N)}(\tilde{t} \pm \delta) - P_{s'}(\tilde{t} \pm \delta)\| < \varepsilon/4$. Then, using the previous reasoning, we can see that $\sup_{t \in]\tilde{t} - \delta, \tilde{t} + \delta[} P_{s'}^{(N)}(t) < \max\{P_{s'}^{(N)}(\tilde{t} + \delta), P_{s'}^{(N)}(\tilde{t} - \delta)\} + 4\delta\Lambda \leq \mu + 3/4\varepsilon < p$. The first inequality holds because $\max\{P_{s'}^{(N)}(\tilde{t} + \delta), P_{s'}^{(N)}(\tilde{t} - \delta)\}$ is a value close to the value of $P^{(N)}$ after the first jump, and is combined with the bound $4\delta\Lambda$ on the variation. Hence $P_{s'}^{(N)}$ ultimately does not cross the line p around \tilde{t} . The case in which $\min\{P_{s'}(\tilde{t}^+), P_{s'}(\tilde{t}^-)\} > p$ is dealt with similarly. ■

Appendix A.4. CSL model checking

Theorem (5.1). *The CSL model checking for ICTMC, for piecewise analytic interval computable rate functions, is decidable for a robust CSL formula $\varphi(p_1, \dots, p_k)$.*

Proof. First of all, we prove that we can approximate the function $P(t)$ for any top next or until formula φ with arbitrary small precision. To start, notice that procedures for integrating ODEs and doing matrix multiplication (which are at the basis of the methods in Sections 5.1 and 5.3) can be computed with arbitrary precision, due to the assumptions of interval computability. Hence, let us focus on the set of zeros of $P(t) - p$, for a dependent next or until formula φ_1 . We want to prove that we can find those zeros with arbitrary precision, and that in doing this we will be able to compute the probability of any next or until formula which contains φ_1 as a subformula with arbitrary precision. If φ is robust, then the time-varying truth

of probability mass.

of formula φ_1 is robust. This means that $P(t) - p$ has a finite number of simple zeros (i.e. their derivatives are not zero). Hence, it is possible to effectively encapsulate them in disjoint intervals of size as small as desired [64]²⁹. Therefore, we can compute the time-varying truth value of the set of states satisfying the formula φ_1 with arbitrary precision, in the sense that for each $\varepsilon > 0$ small enough, we can provide intervals of size at most ε , each containing a single discontinuity point of the set in which one or more states change truth status. The condition on compatibility ensures that we can combine such approximation of time-varying sets and still obtain robust sets. (This may fail if we take the minimum (conjunction) of two truth sets which have a discontinuity for s in the same time point T : we can obtain a function which is neither left nor right continuous, a situation that cannot originate from a simple zero.) Furthermore we have the further property that we can always assume that there is a zero in every approximation interval³⁰. Consider now the problem of computing the probability of an until formula, having two approximations of time-varying truth as described above. Reasoning as in the proof of Lemma 6.2, we can see that if we choose an arbitrary point in each interval wrapping a discontinuity point in spite of the correct one, we commit an error in computing the probability of the until which is uniformly bounded by the total size of the approximation intervals. Hence, we can make such error as small as desired. A similar conclusion can be drawn for a next formula, invoking the line of reasoning of Lemma 6.1. Reasoning inductively, we can therefore compute with any arbitrary precision the probability $P(0)$ of any top until formula.

²⁹As the number of zeros is finite and their first-order derivative is non-zero, the function $f_p(t) = P(t) - p$ crosses zero in those points. Furthermore, notice that the absolute minimum value of the derivative in those zero points is > 0 . Hence, there is an ε such that each interval of size ε containing a zero point has different signs at the extremes and the derivative is provably different from zero. By iterated bisection, we can always find such intervals after a finite number of steps. Furthermore, all intervals J not containing a zero can be eventually discarded by bisection, computing an upper bound L on the absolute value of the derivative in such intervals and bisecting them until we can prove that they are disjoint from zero using the Lipschitz condition with Lipschitz constant L (compute f_p on a single point x in J of length δ , and discard J if $|f_p(x)| - L\delta > 0$).

³⁰If we take the minimum (conjunction) of two truth sets which have a discontinuity for s in the same time point T , then even if the conjunction is robust, when we have an approximation of the time-varying truth function, we can never know if both discontinuities happen in the same time point or in different ones.

Given this value, we then have to solve the inequality $P_s(0) < p_i$ (or $P_s(0) > p_i$) for any s and any top next or until formula φ_i . By the robustness of the CSL formula φ , it cannot be that $P_s(0) = p_i$, hence we can effectively solve that problem by computing $P_s(0)$ with precision $\varepsilon_i < |P_s(0) - p_i|$. As we are doing interval arithmetic computations, we can increase the precision until each p_i will be outside the approximation interval for $P_s(0)$. This proves that the algorithm presented is effective for robust formulae and eventually computes the exact answer. ■

Theorem (5.2). *Given a CSL formula $\varphi(\mathbf{p})$, with $\mathbf{p} \in [0, 1]^k$, then the set $\{\mathbf{p} \mid \varphi(\mathbf{p}) \text{ is robust}\}$ is relatively open³¹ in $[0, 1]^k$ and has Lebesgue measure 1.*

Proof. We will prove the theorem by structural induction on the formula φ . We first need some preliminary definitions. Consider an until formula $\varphi = \mathcal{P}_{\times p}(\varphi_1 \mathbf{U}^{[T_1, T_2]} \varphi_2)$ or a next formula $\varphi = \mathcal{P}_{\times p}(\mathbf{X}^{[T_a, T_b]} \varphi_1)$ and call \mathbf{q} a generic tuple of values for the thresholds on which φ_1 and (in the until case) φ_2 depend. Fix a \mathbf{q} such that the time-varying sets for φ_1 and φ_2 are robust. An open neighbourhood $U_{\mathbf{q}}$ of \mathbf{q} is *robust* if the time-varying sets for φ_1 and φ_2 are robust for each $\mathbf{q}' \in U_{\mathbf{q}}$. Observe that in $U_{\mathbf{q}}$ the number of discontinuities of time-varying truth sets of φ_1 and φ_2 does not change (φ_j is robust for each point in U , and a change in the number of discontinuities can happen only at a non-robust point) and the time-instants at which such discontinuities happen depend continuously on \mathbf{q} . Now, we define the set-valued function $b : U_{\mathbf{q}} \rightarrow 2^{[0, 1]}$ in the following way: Given \mathbf{q}' , $b(\mathbf{q}')$ is the set of values p which causes the time-varying truth set of φ to be non-robust. Therefore, $b(\mathbf{q}')$ contains the values of $P_s(t)$ for which $P'_s(t)$ is zero, the values $P_s(t^-)$ and $P_s(t^+)$ for each non-analytic point t of P , and the values of constant pieces of P_s , plus the value $P_s(0)$. Hence it is finite. By possibly restricting $U_{\mathbf{q}}$, we can also assume that the number of points in $b(\mathbf{q}')$ is bounded by $|b(\mathbf{q})|$ in $U_{\mathbf{q}}$ ³² and the value of such points depends

³¹A set $U \subset V$ is relatively open in $V \subset W$, where W is a topological space, if it is open in the subspace topology, i.e. if there exists an open subset $U_1 \subseteq W$ such that $U = V \cap U_1$.

³²We have to restrict U to avoid the appearance of further zeros of the derivatives away from current zeros. Note also that if a value $p \in b(\mathbf{q})$ corresponds to a non-simple zero at a time t_0 in which the derivative has a maximum or a minimum, a small perturbation of \mathbf{q} can split it in two, or make it disappear. Just think about raising or lowering a

continuously on \mathbf{q}' . Therefore, the set-valued map $b : U_{\mathbf{q}} \rightarrow 2^{[0,1]}$ is *upper-semicontinuous* in $U_{\mathbf{q}}$, i.e. for each neighbourhood $U_{b(\mathbf{x})}$ of $b(\mathbf{x})$ in $[0, 1]$, there is a neighbourhood $U_{\mathbf{x}}$ of \mathbf{x} in $U_{\mathbf{q}}$ such that $b(U_{\mathbf{x}}) \subseteq U_{b(\mathbf{x})}$.

Given a formula $\varphi = \varphi(\mathbf{p})$, $\mathbf{p} \in [0, 1]^k$, we define the set $R_{\varphi} \subset [0, 1]^k$ of all robust thresholds, i.e. $\mathbf{p}_0 \in R_{\varphi}$ if and only if $\varphi(\mathbf{p}_0)$ is robust for φ . Hence, our goal is to show that R_{φ} is open and has measure 1 for any formula φ .

We are now ready for the inductive argument.

Base case: The base case corresponds to (boolean combinations of) atomic formulae, which are robust for each $\mathbf{p} \in [0, 1]^k$.

Boolean combinations: The only non-trivial cases are the conjunctions or disjunctions of until or next formulae. In these cases, we have to enforce the compatibility condition by guaranteeing that the discontinuity times of truth-valued functions are disjoint for each pair of until or next formulae. Consider two until or next formulae φ_1 and φ_2 , and let $\mathbf{p} = (\mathbf{q}_1, p_1, \mathbf{q}_2, p_2)$, where p_j is the threshold for formula φ_j and \mathbf{q}_j is the set of constants which φ_j depends on. By inductive hypothesis, the robust sets $R_j = R_{\varphi_j}$ for φ_j are open and have measure 1. Now, let $P_j = P_j(t, \mathbf{q}_j)$ be the probability of the until or next path formula in φ_j , and fix a robust point $\mathbf{q} = (\mathbf{q}_1, p_1, \mathbf{q}_2)$. Let $g(\mathbf{q})$ be the set valued function $g(\mathbf{q}) = P_2(\{t \mid P_1(t, \mathbf{q}_1) = p_1\}, \mathbf{q}_2) \cup b_2(\mathbf{q}_2)$, where b_2 is the set of non-robust points for φ_2 , as defined above. The set $g(\mathbf{q})$ contains all thresholds for φ_2 for which φ_2 is non-robust and all thresholds that would make the boolean combination non-robust. By properties of the analytic functions, it follows that $g(\mathbf{q})$ is finite. Hence by arguments similar to the ones above, the function g is upper-semicontinuous³³ in a neighbourhood U of \mathbf{q} . Therefore, letting $p_2 \notin g(\mathbf{q})$ and $V \cap g(\mathbf{q}) = \emptyset$

curve having a local maximum or minimum with value zero. However, split zeros will be at points t_1 and t_2 arbitrarily close to t_0 , and therefore, by continuity of $P_s(t, \mathbf{q})$ with respect to \mathbf{q} , the values of $P_s(t_i, \mathbf{q}')$, for \mathbf{q}' close to \mathbf{q} will be close to $P_s(t_0, \mathbf{q})$. Zeros that disappear are not a problem for semicontinuity, as the empty set is contained in any set. Hence, we need to count those points twice in $|b(\mathbf{q})|$.

³³The number of solutions of $P_1(t, \mathbf{q}_1) = p_1$ in a sufficiently small neighbourhood U_1 of (p_1, \mathbf{q}_1) is constant and hence the set-valued function $g_1(p_1, \mathbf{q}_1) = \{t \mid P_1(t, \mathbf{q}_1) = p_1\}$ is upper semicontinuous. Furthermore, in a sufficiently small neighbourhood U_2 of \mathbf{q}_2 , the function $P_2(t, \mathbf{q}_2)$ is continuous in \mathbf{q}_2 for each continuity point t of $P_2(t, \mathbf{q}_2)$. Points for which $P_2(t, \mathbf{q}_2)$ is not continuous are covered by b , hence both $P_2(t^+, \mathbf{q}_2)$ and $P_2(t^-, \mathbf{q}_2)$ are in $g(\mathbf{q})$. Now, for each neighbourhood V of $g(\mathbf{q})$, by piecewise analyticity

a neighbourhood of p_2 in $[0, 1]$, we can find a neighbourhood U of \mathbf{q} such that $g(U) \cap V = \emptyset$, so that $W = U \times V$ is an open neighbourhood of $\mathbf{p} = (\mathbf{q}_1, p_1, \mathbf{q}_2, p_2)$ which contains only robust points, which proves that $R = R_\varphi$ is open.

Furthermore, R is *a fortiori* measurable. Now, let $h_R : [0, 1]^{k_1+k_2} \rightarrow \{0, 1\}$ be the indicator function of the set R in which the boolean combination φ of φ_1 and φ_2 is robust. Note that $R \subseteq R_1 \times R_2$, and call R'_2 the set of thresholds \mathbf{q}_2 for which sub-formulae of φ_2 are robust, which is open and has measure 1 in $[0, 1]^{k_2-1}$ by inductive hypothesis. By Fubini's Theorem:

$$\begin{aligned} \mu_\ell(R) &= \int_{[0,1]^{k_1+k_2}} h_R(\mathbf{q}_1, p_1, \mathbf{q}_2, p_2) \mu_\ell(d\mathbf{q}_1, dp_1, d\mathbf{q}_2, dp_2) \\ &= \int_{[0,1]^{k_1+k_2-1}} \int_{[0,1]} h_R(\mathbf{q}_1, p_1, \mathbf{q}_2, p_2) \mu_\ell(dp_2) \mu_\ell(d\mathbf{q}_1, dp_1, d\mathbf{q}_2) \\ &= \int_{R_1 \times R'_2} \mu_\ell(d\mathbf{q}_1, dp_1, d\mathbf{q}_2) = \int_{R_1} \mu_\ell(d\mathbf{q}_1, dp_1) \int_{R'_2} \mu_\ell(d\mathbf{q}_2) = 1, \end{aligned}$$

which proves that R has measure 1.

If we have a boolean combination of $j > 2$ until formulae, we simply reason pairwise and then take the intersection of the so-obtained robust sets, thus getting an open set of measure 1.

Until formulae: Let $\varphi = \mathcal{P}_{\succ p_k}(\varphi_1 \mathbf{U}^{[T_1, T_2]} \varphi_2)$. By inductive hypothesis, the set $R_{\varphi_1} \times R_{\varphi_2} \subset [0, 1]^{k-1}$ for which φ_1 and φ_2 are robust is open and has measure 1. By reasoning as in the boolean combination case (and considering all until and next conjunct/disjuncts of φ_1 and φ_2), we can immediately conclude that the set $R' \subseteq R_{\varphi_1} \times R_{\varphi_2}$ in which the time varying sets of φ_1 and φ_2 are robust and compatible is open and has measure 1.

Now, fix a point $\mathbf{q} \in R'$, let $U \subseteq R'$ be a robust neighbourhood of \mathbf{q} , and consider the set valued function $b : U \rightarrow 2^{[0,1]}$ as defined above. Now fix

and right/left continuity of P_2 , we can find a neighbourhood U_2 of \mathbf{q}_2 and a neighbourhood V_1 of $g_1(p_1, \mathbf{q}_1)$ such that $b_2(U_2) \subseteq V$ and both $\{p \mid p = P_2(t^+, \mathbf{q}_2), t \in V_1\} \subseteq V$ and $\{p \mid p = P_2(t^-, \mathbf{q}_2), t \in V_1\} \subseteq V$. Now, by upper-semicontinuity of g_1 , there is a neighbourhood U_1 of (\mathbf{q}_1, p_1) such that $g_1(U_1) \subseteq V_1$. It follows that $g(U_1 \times U_2) \subseteq V$, hence g is upper-semicontinuous in \mathbf{q} .

$p \notin b(\mathbf{q})$, and choose a neighbourhood V of p such that $V \cap b(\mathbf{q}) = \emptyset$. As b is upper-semicontinuous, there exists $W \subset U$ such that $b(W) \cap V = \emptyset$, hence φ is robust in $W \times V$. By the arbitrary choice of $\mathbf{p} = (\mathbf{q}, p)$, it follows that $R = R_\varphi$ is open, and hence measurable. Now, let $h_R : [0, 1]^k \rightarrow \{0, 1\}$ be the indicator function of the set R in which φ is robust. By Fubini's theorem, it follows that R has measure 1.

Next formulae: The argument for a next formula $\varphi = \mathcal{P}_{\triangleright p_k}(\mathbf{X}^{[T_a, T_b]}\varphi_1)$ is essentially the same as for until formulae, with the only difference that the inductive hypothesis is applied only to φ_1 and there is no need to ensure compatibility. \blacksquare

Theorem (6.1). *Let $\mathcal{X}^{(N)}$ be a sequence of CTMC models, as defined in Section 3.1, and let $Z_k^{(N)}$ and z_k be defined from $\mathcal{X}^{(N)}$ as in Section 3.3.*

Assume that $Z_k^{(N)}$, z_k have piecewise analytic infinitesimal generator matrices.

Let $\varphi(p_1, \dots, p_k)$ be a robust CSL formula. Then, there exists an N_0 such that, for $N \geq N_0$ and each $s \in \mathcal{S}$

$$s, 0 \models_{Z_k^{(N)}} \varphi \Leftrightarrow s, 0 \models_{z_k} \varphi.$$

Proof. We use structural induction to prove that, for each formula φ , the time-varying truth sets $V_\varphi^{(N)}$ of φ in $Z_k^{(N)}$ converge robustly to the robust time-varying truth set V_φ of φ in z_k .

Base case: The case for atomic propositions is trivial, as $V_\varphi^{(N)}$ and V_φ are constant and equal.

Negation: Let $\varphi = \neg\varphi_1$. The result follows because $V_\varphi(t) = 1 - V_{\varphi_1}$ and $V_\varphi^{(N)}(t) = 1 - V_{\varphi_1}^{(N)}$.

Conjunction/Disjunction: Let $\varphi = \varphi_1 \circ \varphi_2$, $\circ \in \{\wedge, \vee\}$. Due to the compatibility condition of robustness of φ with respect to z_k , the set $V_\varphi(t) = mm\{V_{\varphi_1}(t), V_{\varphi_2}(t)\}$, $mm \in \{\min, \max\}$ is robust, with $Disc(V_\varphi) \subseteq Disc(V_{\varphi_1}) \cup Disc(V_{\varphi_2})$. Using the inductive hypothesis, it easily follows that $V_\varphi^{(N)}(t) = mm\{V_{\varphi_1}^{(N)}(t), V_{\varphi_2}^{(N)}(t)\}$ converges robustly to $V_\varphi(t)$.

Next: Let $\varphi = \mathcal{P}_{\triangleright p}(\mathbf{X}^{[T_a, T_b]}\varphi_1)$. By inductive hypothesis, we can apply Lemma 6.1 and deduce that V_φ is robust and $V_\varphi^{(N)}$ converges robustly to V_φ .

Until: Let $\varphi = \mathcal{P}_{\triangleright p}(\varphi_1 \mathbf{U}^{[T_a, T_b]}\varphi_2)$. By inductive hypothesis (and the compatibility condition in the definition of robustness of φ), we can apply Lemma 6.2³⁴ and deduce that V_φ is robust and $V_\varphi^{(N)}$ converges robustly to V_φ .

The fact that $V_\varphi^{(N)}$ converges robustly to the robust set V_φ , combined with property 1 of robustness of φ , let us conclude that the truth value of φ at level N converges to the truth value of the limit ICTMC at time zero (if 0 was a point in which convergence of probability fails, then a small perturbation in \mathbf{p} could change the truth value of φ in the limit ICTMC, contradicting the robustness of φ ; furthermore, robustness of φ forbids that $P_s(0) = p$). ■

Appendix A.4.1. Comparison of CSL model checking for $Z_k^{(N)}$ and $(Z_k^{(N)}, \hat{\mathbf{X}}_k^{(N)})$

We will prove now the lemmas in Section 7 of the paper. We will start by an auxiliary result, which is needed to adapt the proof style of Lemmas 6.1 and 6.2 to the processes $(Z_k^{(N)}, \hat{\mathbf{X}}_k^{(N)})$ and $(z_k, \hat{\mathbf{x}}_k)$ discussed in this section. In particular, in Lemmas 6.1 and 6.2 we used the fact that processes jump with a small probability in a small temporal neighbourhood of the discontinuity points of time-varying sets. In the space-based setting, however, we need to consider neighbourhoods of the boundaries of goals and unsafe sets. Therefore, to use the same proof style, we need to bound the time trajectories spend in such a neighbourhood, in a uniform way in space. The key point is that, as the convergence results we are interested in depend only on a neighbourhood of the trajectory $\Phi([0, T], \mathbf{x}_0)$, we can always choose a small flow tube such that the velocity with which a trajectory crosses the boundary

³⁴We need to apply it twice for the two reachability problems involved in computing the probability of an until formula, noticing that the probability of the path formula within φ is an analytic combination of the two so-computed probabilities. Robustness of V_φ and robust convergence of $V_\varphi^{(N)}$ to V_φ follows from the same arguments of Lemma 6.2. Alternatively, one can modify Lemma 6.2 and tailor it to the reachability involved in the until case (which reduces the time window in which one can reach the goal set), by a straightforward modification of the definition of $\chi^{(N)}$ and χ and adaptation of the arguments for robust convergence.

of a goal or an unsafe set is close to that of $\Phi([0, T], \mathbf{x}_0)$, and so will be the time spent in a neighbourhood around such boundary. In the following, we will make this intuition formal.

First of all, observe that the notion of robust set V implies that when a trajectory crosses the boundary ∂V in a point \mathbf{x} , the function h defining the smooth manifold of the d-set ∂V around \mathbf{x} changes sign.

We will now prove an upper bound for the time spent by a trajectory in a neighbourhood of the d-set $D = \partial G$ of a robust set G in $\mathcal{S} \times E_0$, a T, ε_0 -flow tube of $\Phi([0, T], \mathbf{x}_0)$.

For each trajectory $\Phi([T_{\mathbf{x}}^-(E_0), T_{\mathbf{x}}^+(E_0)], \mathbf{x})$, $\mathbf{x} \in E_0$, consider the points $Disc(s, \mathbf{x}) = \{(s, \mathbf{x}_1) \in D \mid \mathbf{x}_1 \in \Phi([T_{\mathbf{x}}^-(E_0), T_{\mathbf{x}}^+(E_0)], \mathbf{x})\}$ in which it intersects D .

We define the ε -neighbourhood of D in $\mathcal{S} \times E_0$ as $D_\varepsilon = \bigcup_{(s, \mathbf{x}) \in D \cap (\mathcal{S} \times E_0)} B_\varepsilon(s, \mathbf{x})$, which is an open set. Note that $D_\varepsilon = \bigcup_{(s, \mathbf{x}) \in \mathcal{S} \times E_0} \bigcup_{(s, \mathbf{x}_1) \in Disc(s, \mathbf{x})} B_\varepsilon(s, \mathbf{x}_1)$, as $\mathcal{S} \times E_0$ is the union of a set of trajectories.

Now, by the robustness property of G in $\mathcal{S} \times E_0$, we have that $|Disc(s, \mathbf{x})| \leq k$. Furthermore, by the robustness of G , the trajectory $\Phi([0, T], \mathbf{x}_0)$ will cross D moving from the interior of G to the interior of its complement, or vice versa, for any point $(s, \mathbf{x}_i^s) \in Disc(s, \mathbf{x}_0)$ and any $s \in \mathcal{S}$. Consider a neighbourhood W of (s, \mathbf{x}_i^s) in which $D = \partial G$ is a smooth manifold. Therefore, there is a sufficiently smooth function h in W such that $D \cap W$ is the zero set of h . By the robustness property of G , the function $h(\Phi(t, \mathbf{x}_0))$ equals 0 at t_i (the time such that $\Phi(t, \mathbf{x}_0) = \mathbf{x}_i^s$), and changes sign around t_i (i.e. it is positive at $t_i - \delta$ and negative at $t_i + \delta$, for a $\delta > 0$). It follows that the derivative of $h(\Phi(t, \mathbf{x}_0))$ in t_i is non-null. As it equals $\nabla h(\mathbf{x}_i^s) \cdot F(\mathbf{x}_i^s)$, we have that $\nabla h(\mathbf{x}_i^s) \neq \mathbf{0}$, and hence $|\frac{\nabla h(\mathbf{x}_i^s)}{\|\nabla h(\mathbf{x}_i^s)\|} \cdot F(\mathbf{x}_i^s)| > 0$.

Now, by choosing a suitably small neighbourhood $W_1 \subset W$ of (s, \mathbf{x}_i^s) (we need to ensure that the manifold containing (s, \mathbf{x}_i^s) is the closest one in W_1 , among those constituting D), we obtain that the function $\rho(t, \mathbf{x}) = dist((s, \Phi(t, \mathbf{x})), \partial G) = \inf_{(s, \mathbf{y}) \in \partial G} \|\Phi(t, \mathbf{x}) - \mathbf{y}\|$ is differentiable (in t and \mathbf{x}), and its derivative in $(0, \mathbf{x}_i^s)$ is $\rho'(0, \mathbf{x}_i^s) = \frac{\partial \rho(0, \mathbf{x}_i^s)}{\partial t} = \frac{\nabla h(\mathbf{x}_i^s)}{\|\nabla h(\mathbf{x}_i^s)\|} \cdot F(\mathbf{x}_i^s)$, i.e. it is the projection of the vector field along the normal to the surface $\{h = 0\}$. Now, by continuity of ρ' , we find a neighbourhood $W_2 \subset W_1$ of (s, \mathbf{x}_i^s) such that $|\rho'(0, \mathbf{x})| \geq \rho_i^s/2$, where $\rho_i^s = |\rho'(0, \mathbf{x}_i^s)|$.

Now, choose $\varepsilon_1 < \varepsilon_0$ and $\bar{\varepsilon}$ that satisfies: (i) the number of intersections between D and a trajectory in $\mathcal{S} \times E_1$ is constant and equal to the number k of intersections of $\mathcal{S} \times \Phi([0, T], \mathbf{x}_0)$ with D (this is possible because the

flow tube $\mathcal{S} \times E_1$ is a small neighbourhood of $\mathcal{S} \times \Phi([0, T], \mathbf{x}_0)$, and (ii) $D_{\bar{\varepsilon}}$, the $\bar{\varepsilon}$ -neighbourhood of D in the T, ε_1 -flow tube $\mathcal{S} \times E_1$, is contained in the neighbourhood W_2 of (s, \mathbf{x}_i^s) identified above, for any $s \in \mathcal{S}$ and $i \leq k$.

By the choice of W_2 , it follows that the speed at which each trajectory of $\mathcal{S} \times E_1$ travels in $D_{\bar{\varepsilon}}$ (with respect to the distance from D) is bounded below by $\rho_0 = \min_{s,i} \rho_i^s/2$, and hence the total time $\tau_{\bar{\varepsilon}}$ a trajectory of $\mathcal{S} \times E_1$ spends in $D_{\bar{\varepsilon}}$ is bounded above by $\frac{2\bar{\varepsilon}k}{\rho_0}$, as it has to travel a total distance of $2\bar{\varepsilon}k$ with speed no less than ρ_0 . Notice that this bound is independent of the specific trajectory considered.

With the previous discussion, we have proved the following

Lemma Appendix A.1. *Let $E_0 \subset E$ be a T, ε_0 -flow tube for \mathbf{x}_0 . Let G be a robust subset of $\mathcal{S} \times E_0$. Then, there are positive constants $\varepsilon_1, \bar{\varepsilon}$, and ρ_0 such that, for any $\varepsilon' < \bar{\varepsilon}$, the total time $\tau_{\varepsilon'}$ a trajectory in $\mathcal{S} \times E_1$ (E_1 the T, ε_1 -flow tube for \mathbf{x}_0) spends in $D_{\varepsilon'}$, the ε' neighbourhood of $D = \partial G$, satisfies $\tau_{\varepsilon'} \leq \frac{2\varepsilon'k}{\rho_0}$, where k is the number of intersections of any trajectory with D .*

Equipped with this lemma, we can now prove the following one.

Lemma (7.1). *Let $E_0 \subset E$ be a T, ε_0 -flow tube for \mathbf{x}_0 . Let G be a robust subset of $\mathcal{S} \times E_0$, and $G^{(N)}$ be a sequence of subsets of $\mathcal{S} \times E_0$ that converge robustly to G .*

Let $\bar{P}(s, \mathbf{x}) = P_{next}(\mathbf{y}, s, \mathbf{x}, T_a, T_b, G)$ be the probability that the first jump of $\mathbf{y}(t)$ is into a state in G and happens at a time $t \in [T_1, T_2]$, given that \mathbf{y} started at time $t = 0$ in state $(s, \mathbf{x}) \in \mathcal{S} \times E_0$, and let $\bar{P}^{(N)}(s, \mathbf{x}) = P_{next}^{(N)}(\mathbf{Y}^{(N)}, s, \nu^{(N)}(\mathbf{x}), T_a, T_b, G^{(N)})$ be defined similarly, with G and \mathbf{x} replaced by $G^{(N)}$ and $\nu^{(N)}(\mathbf{x})$, respectively.

Furthermore, define $V = \{(s, \mathbf{x}) \mid \bar{P}(s, \mathbf{x}) \triangleright p\}$ and $V^{(N)} = \{(s, \mathbf{x}) \mid \bar{P}^{(N)}(s, \mathbf{x}) \triangleright p\}$. Then there exists $\varepsilon_1 > 0$ such that, in E_1 , the $(T - T_b), \varepsilon_1$ -flow tube for \mathbf{x}_0 :

1. $\bar{P}^{(N)}(s, \mathbf{x}) \rightarrow \bar{P}(s, \mathbf{x})$ for all $\mathbf{x} \in E_1$, uniformly in (s, \mathbf{x}) .
2. If $V_{\mathbf{x}_0}(t)$, $t \in [T_{\mathbf{x}_0}^-(E_1), T_{\mathbf{x}_0}^+(E_1)]$, is a robust time-varying set, then V is robust in E_1 and $V^{(N)}$ converges robustly to V .

Proof. First, notice that we can always restrict to an arbitrary small neighbourhood of $\Phi(t, \mathbf{x}_0)$, i.e. to a T, ε_1 -flow tube E_1 for \mathbf{x}_0 , with ε_1 as small as desired. This follows from the convergence in probability implied by Kurtz Theorem 3.1. Given $\delta > 0$, this guarantees that we can find an index N_0

such that, for any $N > N_0$, with probability at least $1 - \delta$, the trajectories of $\hat{\mathbf{X}}^{(N)}(t)$ are contained in E_1 . Furthermore, we can choose such an index N_0 independently of \mathbf{x} . Hence, given E_0 , if we consider a flow tube E_1 for \mathbf{x}_0 with radius $\varepsilon_1 < \varepsilon_0/2$, each flow tube of radius ε_1 wrapping a trajectory in E_1 will be contained in E_0 . This guarantees that the next-step probability for $\mathbf{Y}^{(N)}$ for any point in $\mathcal{S} \times E_1$ will ultimately depend only on the goal sets $G^{(N)}$ within $\mathcal{S} \times E_0$.

We first prove point 1 of the lemma in a $(T - T_b), \varepsilon_1$ -flow tube E_1 for x_0 , for an $\varepsilon_1 < \varepsilon_0/2$ to be fixed in the following. We will prove convergence of $\bar{P}^{(N)}$ to \bar{P} for each $(s, \mathbf{x}) \in \mathcal{S} \times E_1$.

We will now use an argument similar to the one of Lemma 6.1. Couple \mathbf{y} and $\mathbf{Y}^{(N)}$ on the same probability space Ω , and let χ (resp. $\chi^{(N)}$) be random variables defined on sample trajectories and equal to one if the trajectory's first jump of the s component is in G (resp. $G^{(N)}$). Then, as $P(s, \mathbf{x}) = \mathbb{E}[\chi(s, \mathbf{x}, \mathbf{y}(\omega))]$, where $\mathbf{y}(0) = (s, \mathbf{x})$, and similarly for $P^{(N)}(s, \mathbf{x})$, to show convergence we just need to prove that $|\mathbb{E}[\chi(s, \mathbf{x}, \mathbf{y})] - \mathbb{E}[\chi^{(N)}(s, \mathbf{x}, \mathbf{Y}^{(N)})]| \rightarrow 0$. It holds that:

$$\begin{aligned} |\mathbb{E}[\chi(s, \mathbf{x}, \mathbf{y})] - \mathbb{E}[\chi^{(N)}(s, \mathbf{x}, \mathbf{Y}^{(N)})]| &\leq \underbrace{\mathbb{E}[|\chi(s, \mathbf{x}, \mathbf{y}) - \chi^{(N)}(s, \mathbf{x}, \mathbf{y})|]}_{(1)} \\ &+ \underbrace{\mathbb{E}[|\chi^{(N)}(s, \mathbf{x}, \mathbf{y}) - \chi^{(N)}(s, \mathbf{x}, \mathbf{Y}^{(N)})|]}_{(2)}, \end{aligned}$$

To treat term (1), invoke Lemma Appendix A.1, assume ε_1 is smaller than the one required by the lemma, and let $\bar{\varepsilon}$ and ρ_0 the other two constants obtained from it. Now, as in Lemma 6.1, observe that for each $\varepsilon' < \bar{\varepsilon}$, the only trajectories of \mathbf{y} for which χ and $\chi^{(N)}$ can have a different value are those jumping at a time at which $\mathbf{y}(t)$ is in $D_{\varepsilon'}$. Now, the total amount of time \mathbf{y} spends in $D_{\varepsilon'}$ is uniformly bounded by $\tau_{\varepsilon'} \leq \frac{2\varepsilon'k}{\rho_0}$, where k is the number of intersections of a trajectory in $\mathcal{S} \times E_1$ with \bar{D} . It follows that term (1) can be bounded by $\frac{2\varepsilon'\Lambda k}{\rho_0}$, where Λ is an upper bound for the jump rate of z_k in $\mathcal{S} \times E_0$.

The bound on term (2), instead, follows from the convergence of $\mathbf{Y}^{(N)}$ to \mathbf{y} , but it requires a slightly different treatment than in Lemma 6.1, as now the time varying sets for $\mathbf{Y}^{(N)}$ depend on the sample trajectories of $\hat{\mathbf{X}}^{(N)}$, hence they are random quantities. Call $G_{\mathbf{x}, \hat{\mathbf{X}}^{(N)}}^{(N)}(t)$ the time-varying sets relative to $G^{(N)}$, but defined with respect to trajectories of $\mathbf{Y}^{(N)}(t)$. The time varying

sets for \mathbf{y} , with respect to $G^{(N)}$, are denoted by $G_{\mathbf{x}}^{(N)}(t)$, while that relative to G is $G_{\mathbf{x}}(t)$. We will need now to control two things: first, we will construct a neighbourhood of D in such a way that all the time varying sets are the same outside it, for N large enough. Then, we will bound the time taken by $\hat{\mathbf{X}}^{(N)}(t)$ to cross such a neighbourhood (again for N large enough).

Assume $\varepsilon' < \bar{\varepsilon}/2$, and consider the ε' -neighbourhood $D_{\varepsilon'}$ of D . Invoking robust convergence, choose N_0 such that for $N \geq N_0$, $G^{(N)}$ coincides with G outside $D_{\varepsilon'}$. We now want to find a neighbourhood $[\bar{t} - \tau', \bar{t} + \tau']$ of the time \bar{t} in which $\hat{\mathbf{x}}(\bar{t}) \in D$, such that we are guaranteed that if t falls outside this neighbourhood, both $\hat{\mathbf{x}}(t)$ and $\hat{\mathbf{X}}^{(N)}(t)$ are outside $D_{\varepsilon'}$. For any such time t , it clearly holds that $G_{\mathbf{x}, \hat{\mathbf{X}}^{(N)}}^{(N)}(t)$ and $G_{\mathbf{x}}^{(N)}(t)$ coincide. To find such a neighbourhood of \bar{t} , let N_1 be such that, for $N \geq N_1$, $\|\hat{\mathbf{X}}^{(N)}(t) - \hat{\mathbf{x}}(t)\|$ is less than ε' with probability $1 - \delta$ (δ to be fixed later). Call this event $\Omega_{\varepsilon'}$. Condition on it and consider $D_{2\varepsilon'}$. If $\hat{\mathbf{x}}(t) \notin D_{2\varepsilon'}$, then it follows that $\hat{\mathbf{X}}^{(N)}(t)$ will not belong to $D_{\varepsilon'}$. Hence, we just need to bound the time $\tau_{2\varepsilon'}$ that $\hat{\mathbf{x}}(t)$ spends in $D_{2\varepsilon'}$. By Lemma Appendix A.1, this time is no more than $\frac{4\varepsilon'k}{\rho_0}$.

Now, using Theorem 3.2, choose an N_2 such that, for $N \geq N_2$, $Z_k^{(N)}(t)$ coincides in $[0, T]$ with $z_k(t)$ with probability at least $1 - \delta$. To bound term 2, observe that if $Z_k^{(N)}$ and z_k are the same, and conditional on event $\Omega_{\varepsilon'}$, if both $Z_k^{(N)}$ and z_k jump at time instants in which $\hat{\mathbf{x}}(t)$ is outside $D_{2\varepsilon'}$, then $\chi^{(N)}(s, \mathbf{x}, \mathbf{y})$ and $\chi^{(N)}(s, \mathbf{x}, \mathbf{Y}^{(N)})$ will have the same value.

Therefore, we can bound term (2) by the probability of $\chi^{(N)}(s, \mathbf{x}, \mathbf{y}) \neq \chi^{(N)}(s, \mathbf{x}, \mathbf{Y}^{(N)})$, which is itself bounded by

$$\mathbb{P}\{z_k \text{ jumps in } D_{2\varepsilon'}\} + \mathbb{P}\{\Omega_{\varepsilon'}^c\} + \mathbb{P}\{Z_k^{(N)} \neq z_k\} \leq \frac{4\varepsilon'\Lambda k}{\rho_0} + 2\delta.$$

Now, fix $\varepsilon > 0$ and choose $\varepsilon' < \min\{\bar{\varepsilon}/4, \frac{\varepsilon\rho_0}{12\Lambda k}\}$, and $\delta < \varepsilon/4$. By combining the bounds on term (1) and term (2), we obtain that

$$\limsup_{N \rightarrow \infty} |\mathbb{E}[\chi(s, \mathbf{x}, \mathbf{y})] - \mathbb{E}[\chi^{(N)}(s, \mathbf{x}, \mathbf{Y}^{(N)})]| \leq \frac{6\varepsilon'\Lambda k}{\rho_0} + 2\delta \leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon,$$

which by the arbitrariness of ε implies that

$$\lim_{N \rightarrow \infty} |\mathbb{E}[\chi(s, \mathbf{x}, \mathbf{y})] - \mathbb{E}[\chi^{(N)}(s, \mathbf{x}, \mathbf{Y}^{(N)})]| = 0.$$

Therefore, we obtain that $\bar{P}^{(N)}(s, \mathbf{x}) \rightarrow \bar{P}(s, \mathbf{x})$, and this convergence is uniform with respect to $(s, \mathbf{x}) \in \mathcal{S} \times E_1$, as the bound derived above is independent of it.

As for point 2 of the lemma, observe that by the fact that the time-varying set $V_{\mathbf{x}_0}(t)$ associated with the fluid trajectory $\Phi(t, \mathbf{x}_0)$ is robust, and by piecewise analyticity of \bar{P} , we can choose an ε_1 sufficiently small not only to satisfy the constraints to derive convergence discussed above, but also such that all trajectories in the flow tube E_1 are robust, i.e. their time varying set with respect to \bar{P} are robust (just observe that the function $\bar{P}(s, \Phi(t, \mathbf{x}))$ is piecewise analytic in t and \mathbf{x} for each s). Furthermore, we have chosen ε_1 so that the number of intersections of $\Phi(t, \mathbf{x})$ with V in each state s , i.e. the number of times $\bar{P}(s, \Phi(t, \mathbf{x})) - p$ changes sign, is the same as that of $\Phi(t, \mathbf{x}_0)$. Now, consider the boundary ∂V in $\mathcal{S} \times E_1$, which is the zero set of the function $h(s, \mathbf{x}) = \bar{P}(s, \mathbf{x}) - p$. By continuity of \bar{P} and by the robustness property of $V_{\mathbf{x}_0}(t)$, we have that the trajectory $(s, \Phi([0, T], \mathbf{x}_0))$ intersects ∂V in points \mathbf{x}_i^s in which the function \bar{P} is analytic. Hence, \bar{P} will be analytic in a neighbourhood of \mathbf{x}_i^s , and, by a suitable choice of ε_1 , \bar{P} will be analytic in the whole component of ∂V containing \mathbf{x}_i^s . It follows that ∂V is the union of smooth manifolds (analytic in this case).

Thus, by choosing ε_1 suitably small, ∂V is a d-set and V is robust.

As for the robust convergence of $V^{(N)}$ to V , by the uniform convergence of $\bar{P}^{(N)}$ to \bar{P} outside an open neighbourhood of ∂V , we obtain the robust convergence of $V^{(N)}$ to V . \blacksquare

Lemma (7.2). *Let $E_0 \subset E$ be a T, ε_0 -flow tube for \mathbf{x}_0 . Let U and G two robust and compatible subsets of $\mathcal{S} \times E_0$, and $U^{(N)}, G^{(N)}$ be sequences of subsets of $\mathcal{S} \times E_0$ that converge robustly to U and G , respectively.*

Let $P(s, \mathbf{x}) = P_{\text{reach}}(\mathbf{y}, s, \mathbf{x}, T_a, T_b, U, G)$ be the probability that $\mathbf{y}(t)$ reaches a state in G within time $[T_a, T_b]$, avoiding any unsafe state in U , given that \mathbf{y} started at time $t = 0$ in state $(s, \mathbf{x}) \in \mathcal{S} \times E_0$, and let $P^{(N)}(s, \mathbf{x}) = P_{\text{reach}}(\mathbf{Y}^{(N)}, s, \nu^{(N)}(\mathbf{x}), T_a, T_b, U^{(N)}, G^{(N)})$ be defined similarly, with G, U, \mathbf{x} replaced by $G^{(N)}, U^{(N)}$, and $\nu^{(N)}(\mathbf{x})$, respectively. Furthermore, define $V = \{(s, \mathbf{x}) \mid P(s, \mathbf{x}) \bowtie p\}$ and $V^{(N)} = \{(s, \mathbf{x}) \mid P^{(N)}(s, \mathbf{x}) \bowtie p\}$. Then there exists $\varepsilon_1 > 0$ such that, in E_1 , the $(T - T_b), \varepsilon_1$ -flow tube for \mathbf{x}_0 :

1. $P^{(N)}(s, \mathbf{x}) \rightarrow P(s, \mathbf{x})$ for all $\mathbf{x} \in E_1 \setminus D$, where D is a d-set, uniformly in (s, \mathbf{x}) .
2. If $V_{\mathbf{x}_0}(t)$, $t \in [T_{\mathbf{x}_0}^-(E_1), T_{\mathbf{x}_0}^+(E_1)]$, is a robust time-varying set, then V is robust in E_1 and $V^{(N)}$ converges robustly to V .

Proof. First, notice that, as in Lemma 7.1, we can always restrict on an

arbitrary small neighbourhood of $\Phi(t, \mathbf{x}_0)$, i.e. on a T, ε_1 -flow tube E_1 for \mathbf{x}_0 , with ε_1 as small as desired, implying that the reachability problem for $\mathbf{Y}^{(N)}$ for any point in $\mathcal{S} \times E_1$ will eventually depend only on the goal sets $G^{(N)}$ and $U^{(N)}$ within $\mathcal{S} \times E_0$.

We first prove point 1 of the lemma in a $(T - T_b), \varepsilon_1$ -flow tube E_1 for \mathbf{x}_0 , for an $\varepsilon_1 < \varepsilon_0/2$ (ε_1 will be fixed in the following). Consider the set D in E_0 , $D = \text{Disc}(G) \cup \text{Disc}(U) \cup \Phi^{-1}(T_a, \text{Disc}(G)) \cup \Phi^{-1}(T_a, \text{Disc}(U)) \cup \Phi^{-1}(T_b, \text{Disc}(G)) \cup \Phi^{-1}(T_b, \text{Disc}(U))$, containing the discontinuity points of G and U and all points that are mapped by the flow to $\text{Disc}(G) \cup \text{Disc}(U)$ after T_a or T_b units of time. D is easily seen to be a d-set. In fact, it is closed and it intersects each trajectory a finite number of times, as $\text{Disc}(G)$ and $\text{Disc}(U)$ are d-sets. Furthermore, each smooth manifold of G or U , defined as the zero set of the function $h(\mathbf{x})$, will be mapped by $\Phi^{-1}(T_j, \cdot)$, $j = a, b$, into the smooth manifold defined by the function $h(\Phi(t_j, \mathbf{x}))$. This function is smooth as $\Phi(t, \mathbf{x})$ is piecewise analytic and it is at least of class \mathcal{C}^1 .

Differently from Lemma 7.1, we will prove convergence of $P^{(N)}$ to P for each $(s, \mathbf{x}) \in (\mathcal{S} \times E_1) \setminus D$.

Consider now the set $D_0 = \text{Disc}(G) \cup \text{Disc}(U)$, and define the ε -neighbourhood of D_0 , D_ε , as done in Lemma 7.1. It clearly holds that $D_\varepsilon \rightarrow D_0$, as $\varepsilon \rightarrow 0$.

We will now use an argument similar to the one of Lemma 6.2. Let χ (resp. $\chi^{(N)}$) be random variables defined on sample trajectories and equal to one if the trajectory satisfies the reachability problem of the Lemma with respect to G, U (resp. $G^{(N)}, U^{(N)}$). Then, as $P(s, \mathbf{x}) = \mathbb{E}[\chi(s, \mathbf{x}, \mathbf{y}(\omega))]$, where $\mathbf{y}(0) = (s, \mathbf{x})$, and similarly for $P^{(N)}(s, \mathbf{x})$, to show convergence we just need to prove that $|\mathbb{E}[\chi(s, \mathbf{x}, \mathbf{y})] - \mathbb{E}[\chi^{(N)}(s, \mathbf{x}, \mathbf{Y}^{(N)})]| \rightarrow 0$. It holds that:

$$\begin{aligned} |\mathbb{E}[\chi(s, \mathbf{x}, \mathbf{y})] - \mathbb{E}[\chi^{(N)}(s, \mathbf{x}, \mathbf{Y}^{(N)})]| &\leq \underbrace{\mathbb{E}[|\chi(s, \mathbf{x}, \mathbf{y}) - \chi^{(N)}(s, \mathbf{x}, \mathbf{y})|]}_{(1)} \\ &+ \underbrace{\mathbb{E}[|\chi^{(N)}(s, \mathbf{x}, \mathbf{y}) - \chi^{(N)}(s, \mathbf{x}, \mathbf{Y}^{(N)})|]}_{(2)}, \end{aligned}$$

From Lemma Appendix A.1, we obtain constants $\bar{\varepsilon}_1$ and $\bar{\varepsilon}$ that bound the size of the flow tube E_1 , and of the $D_{\varepsilon'}$ neighbourhood of D_0 . Under these constraints, we can reason exactly as in the proof of Lemma 7.1 to bound terms (1) and (2) by $\frac{6\varepsilon'\Lambda k}{\rho_0} + 2\delta$, where δ and $\varepsilon' < \bar{\varepsilon}/2$ can be chosen arbitrary small for N large enough, concluding that $|P^{(N)}(s, \mathbf{x}) - P(s, \mathbf{x})|$ converges to zero, uniformly in $(\mathcal{S} \times E_1) \setminus D$.

As for point 2 of the lemma, robustness of V follows by the same argument as Lemma 7.1. Notice that $Disc(V)$ is closed, as it is the union of the zero sets of continuous functions (the analytic pieces of P), plus the subset D_p of discontinuity points of P such that $\liminf P(s, \mathbf{x}) \leq p$ and $\limsup P(s, \mathbf{x}) \geq p$, which is also closed. Furthermore, by robustness of $V_{\mathbf{x}_0}(t)$, we can choose ε_1 such that all points in D_p satisfy $\liminf P(s, \mathbf{x}) < p$ and $\limsup P(s, \mathbf{x}) > p$ (we need this because $V_{\mathbf{x}}(t)$ has to be robust for all \mathbf{x} in $\mathcal{S} \times E_1$). For the robust convergence of $V^{(N)}$ to V , $Disc(V)$ is a d-set, hence we can use uniform convergence of $P^{(N)}$ to P outside an open neighbourhood of $Disc(V)$. Additionally, notice that, as in the proof of Lemma 6.2, the points in which we do not have convergence of $P^{(N)}$ to P and that are not in $Disc(V)$, do not create problems, as in a small neighbourhood of those points, P is always strictly above or below p , and the lim sup or the lim inf of $P^{(N)}$ in those points will uniformly satisfy the inequality defining $V^{(N)}$ (thanks to the compatibility condition of G and U). ■

Lemma (7.3). *Let $\mathcal{X}^{(N)}$ be a sequence of CTMC models, as defined in Section 3.1, and let $Z_k^{(N)}$ and z_k be defined from $\mathcal{X}^{(N)}$ as in Section 3.3.*

Assume that there is a flow tube E_0 of \mathbf{x}_0 such that all trajectories in E_0 are piecewise analytic.

Let $\varphi = \varphi(\mathbf{p})$ be a robust CSL formula for the trajectory $\Phi(t, \mathbf{x}_0)$. Then, there is an N_0 such that, for all $N \geq N_0$,

$$s, \mathbf{x}_0 \models_{\mathbf{y}} \varphi \iff s, \nu^{(N)}(\mathbf{x}_0) \models_{\mathbf{Y}^{(N)}} \varphi.$$

Proof. We will prove by structural induction on the formula φ , that there is a T, ε -flow tube E_φ of \mathbf{x}_0 such that $V_\varphi^{(N)}$ converges to V_φ robustly, where V_φ is the set of points $(s, \mathbf{x}) \in \mathcal{S} \times E_\varphi$ such that $s, \mathbf{x} \models_{\mathbf{y}} \varphi$ and $V_\varphi^{(N)}$ is the set of points $(s, \mathbf{x}) \in \mathcal{S} \times E_\varphi$ such that $s, \nu^{(N)}(\mathbf{x}_0) \models_{\mathbf{Y}^{(N)}} \varphi$.

Base case: the result for atomic formulae is trivial as their truth value depends only on s , hence we can choose $E_\varphi = E_0$ and $V_\varphi = V_\varphi^{(N)} = \mathcal{S}_\varphi \times E_0$, where $\mathcal{S}_\varphi = \{s \mid s \models \varphi\}$.

Negation: If $\varphi = \neg\varphi_1$, we can choose $E_\varphi = E_{\varphi_1}$, and simply observe that robust convergence of $V_{\varphi_1}^{(N)}$ to V_{φ_1} implies robust convergence of $V_\varphi^{(N)} = E_{\varphi_1} \setminus V_{\varphi_1}^{(N)}$ to $V_\varphi = E_{\varphi_1} \setminus V_{\varphi_1}$.

Conjunction and Disjunction: If $\varphi = \varphi_1 \circ \varphi_2$, $\circ \in \{\wedge, \vee\}$, consider E_{φ_i} , a T, ε_i -flow tube, and sets $V_{\varphi_i}^{(N)} \rightarrow V_{\varphi_i}$. As φ is robust for the trajectory starting in \mathbf{x}_0 and $\mathbf{x}_0 \in E_{\varphi_i}$, by the compatibility condition of φ there exists an ε such that the d-sets of V_{φ_1} and V_{φ_2} are disjoint in $\mathcal{S} \times E_{\varphi}$, for E_{φ} the T, ε -flow tube in \mathbf{x}_0 . It easily follows that $V_{\varphi} = V_{\varphi_1} \bullet V_{\varphi_2}$ is robust in $\mathcal{S} \times E_{\varphi}$, $\bullet \in \{\cap, \cup\}$, and $V_{\varphi_1}^{(N)} \bullet V_{\varphi_2}^{(N)} \rightarrow V_{\varphi_1} \bullet V_{\varphi_2}$ robustly.

Next: If $\varphi = \mathcal{P}_{\triangleright p}(\mathbf{X}^{[T_1, T_2]} \varphi_1)$, let E_{φ_1} be a T, ε -flow tube for φ_1 and let V_{φ_1} be a robust set, such that $V_{\varphi_1}^{(N)} \rightarrow V_{\varphi_1}$ robustly. By considering the ε, T -flow tube E_j for \mathbf{x}_0 , and using the robustness of φ , we satisfy the hypothesis of Lemma 7.1, hence there is an $(T - T_2), \varepsilon$ -flow tube E_{φ} for \mathbf{x}_0 such that V_{φ} is robust in $\mathcal{S} \times E_{\varphi}$ and $V_{\varphi}^{(N)} \rightarrow V_{\varphi}$ robustly.

Until: If $\varphi = \mathcal{P}_{\triangleright p}(\varphi_1 \mathbf{U}^{[T_1, T_2]} \varphi_2)$, let E_{φ_i} be T, ε_i -flow tubes for φ_i , $i = 1, 2$, and robust sets V_{φ_i} , such that $V_{\varphi_i}^{(N)} \rightarrow V_{\varphi_i}$ robustly. By letting $\varepsilon_0 < \min\{\varepsilon_1, \varepsilon_2\}$, such that $\partial V_{\varphi_1} \cap \partial V_{\varphi_2} = \emptyset$ (which can be found by the compatibility condition enforced by robustness of φ), considering the ε_0, T -flow tube E_j for \mathbf{x}_0 , and using the robustness of φ , we satisfy the hypothesis of Lemma 7.2, hence there is an $(T - T_2), \varepsilon$ -flow tube E_{φ} for \mathbf{x}_0 such that V_{φ} is robust in $\mathcal{S} \times E_{\varphi}$ and $V_{\varphi}^{(N)} \rightarrow V_{\varphi}$ robustly.

Given a formula φ , and the flow-tube E_{φ} for \mathbf{x}_0 , such that V_{φ} is robust in $\mathcal{S} \times E_{\varphi}$ and $V_{\varphi}^{(N)} \rightarrow V_{\varphi}$ robustly, then the lemma follows by observing that, due to robustness of φ , \mathbf{x}_0 does not belong to the d-set ∂V_{φ} , hence there is an N_0 such that, for all $N \geq N_0$, $(s, \mathbf{x}_0) \in V_{\varphi}^{(N)} \Leftrightarrow (s, \mathbf{x}_0) \in V_{\varphi}$. \blacksquare