QUANTICOL

A Quantitative Approach to Management and Design of Collective and Adaptive Behaviours

http://www.quanticol.eu

D1.2

A framework for hybrid limits under uncertainty

Revision: 1.0; September 29, 2015

Author(s): Luca Bortolussi (CNR), Nicolas Gast (INRIA), Jane Hillston (UEDIN), Mirco Tribastone (IMT)

Due date of deliverable: Month 30 (September 2015)

Actual submission date: Sept 30, 2015

Nature: R. Dissemination level: PU

Funding Scheme: Small or medium scale focused research project (STREP)
Topic: ICT-2011 9.10: FET-Proactive 'Fundamentals of Collective Adaptive Systems' (FOCAS)
Project number: 600708

Coordinator: Jane Hillston (UEDIN)
e-mail: Jane.Hillston@ed.ac.uk
Fax: +44 131 651 1426

Part. no.	Participant organisation name	Acronym	Country
1 (Coord.)	University of Edinburgh	UEDIN	UK
2	Consiglio Nazionale delle Ricerche – Istituto di Scienza e Tecnologie della Informazione "A. Faedo"	CNR	Italy
3	Ludwig-Maximilians-Universität München	LMU	Germany
4	Ecole Polytechnique Fédérale de Lausanne	EPFL	Switzerland
5	IMT Lucca	IMT	Italy
6	University of Southampton	SOTON	UK
7	Institut National de Recherche en Informatique et en Automatique	INRIA	France

COOPERATION

Executive Summary

This deliverable reports on the development of a theoretical framework to study the effect of multiple scales and imprecision in the emergent behaviour of collective adaptive systems (CAS). We show how to construct suitable mean-field approximations for such systems. It constitutes the main achievement of Task 1.1 and a first step towards Task 1.3 and the linking of language specification and mean-field techniques. This document is structured in two main sections, the first one presents results related to multiple scales, both in terms of time and of population levels. The second part focuses on mean field results in presence of uncertainty.

As for multiple scales, we discuss the following results in detail. We first present a general framework for mean field limits for systems with heterogeneous population size, following [Bor15]. This framework considers a very general class of population processes, allowing both immediate and stochastic transitions, guarded by Boolean predicates (to encode for example control actions), and obtaining limits in terms of stochastic hybrid systems, which are usually faster to simulate. This is discussed in detail in Section 2.2. Computing the transition rates of some immediate transitions requires the computation of stochastic hitting times, *i.e.*, the time for a stochastic system to hit a given domain. We show how to use a fluid approximation to compute this time in Section 2.3.

Next, in Section 2.4, we present a general framework to combine mean field limits with reduction of multiple time scales, with conditions providing guarantees on the correctness of exchanging these two operations [BP14]. This framework leads to a new simulation algorithm for Markov models with multiple time scales, leveraging powerful statistical abstraction tools [BMS15]. Finally, an integration of hybrid conditional moment techniques [Has+14] within the stochastic process algebra PEPA [Pou15] is discussed in Section 2.5.

The second part of the document is devoted to the analysis of CAS models in the presence of uncertainty. We distinguish an *uncertain* model – for which a parameter exists but is not known – and an *imprecise* model – for which some parameters may vary. We show how mean field limits can greatly simplify the study of uncertain and imprecise population models [BG15]. This setting encompasses the imprecise and the uncertain scenario, but also more classic models like Markov Decision Processes. It uses a differential inclusion to represent the limit. We discuss it in Section 3.1.

We develop some numerical methods to analyse the class of limit models for uncertain and imprecise population models. In particular, we discuss in Section 3.2 a method based on statistical emulation for the uncertain case [BS14], and two methods, one based on differential hulls [TT15] and one based on the Pontraygin optimal control principle [BG15].

Contents

1 Introduction

Collective adaptive systems (CAS) have many complex features that make the construction of mathematical models a challenging task. In this report, we focus mainly on a few crucial aspects:

1. Dynamical evolution happening at multiple temporal scales. An example is the different activation times of different electricity generation mechanisms in smart grids: switching off appliances to reduce the demand takes a few milliseconds while turning on thermo-generators requires several minutes.

2. The presence of different populations of interacting agents of heterogeneous sizes. For instance, in a smart grid scenario in which appliances can be controlled remotely, we have many appliances but only one or few central controllers.

3. The intrinsic uncertainty in the values of model parameters, due to lack of knowledge or due to intrinsic imprecision in their statistical estimate from available experimental data. We call this the uncertain scenario.

4. The necessity to consider open scenarios and interaction with an unpredictable environment, with the consequent need of under-specification of some model parameters. This will be referred to as the imprecise scenario.

Imprecision and uncertainty, in particular, are ubiquitous features of all CAS, and in general of all complex systems, due to our limited intellectual, experimental, and computational capabilities.

In this project, we investigate the use of mean field approximation as a viable tool to analyse large-scale stochastic models of CAS. The behaviour of mean field approximation in the presence of multiple scales in time and population size, or in the presence of parametric uncertainty and imprecision, is poorly understood. This deliverable collects results obtained within the project that aim at filling this gap. In particular, we present general mean field results for multiple population levels and in the presence of uncertainty and imprecision, paving the way to novel and efficient algorithms for the analysis of models of CAS exhibiting these features.

2 Hybrid Limits Arising in Mean-Field Interaction Systems

In this section, we present results on different kinds of hybridness that can emerge in mean field limits. We start from the results of [Bor15], a journal version of the technical report [Bor12], which has been discussed previously in Deliverable 1.1. As such, we will stress more the novel results on hybrid mean field limits in the presence of discontinuities, as these are based on similar techniques to those discussed in Section 3 about mean field under parameter uncertainty, hinting at how hybridness and uncertainty may be seen from a unified perspective, at least from a formal point of view. We also show how to compute rates of instantaneous transitions that are fired when the system hits a boundary [GG15].

In this section we will also report on work connecting mean field and temporal reduction of multi-scale systems [BP14; BMS15], on hybrid conditional moment techniques [Pou15]. Our starting point, however, will be a quick overview of Markov Population Models and standard mean field limits, following mainly [Bor+13], to fix notation.

2.1 Markov Population Models and Classic Mean-Field Limits

We recall here for the sake of completeness the definition of Markov population models, defining a (population) CTMC. We are interested in models in which populations of agents, possibly of different kinds, interact and evolve by changing state. We assume that the identity of individual agents does not influence the dynamics, hence we just need to count how many agents are in each state. To this end, we introduce integer-valued counting variables $\mathbf{X} = (X_1, \dots, X_n)$, with n the number of different agent states. These variables take values in a subset $S \subseteq \mathbb{N}^n$, which is the state space of the model. The dynamics of Markov population models is described, according to [Bor+13; Bor15], by a collection $\mathcal{T}$ of transition classes, each consisting in a tuple $\eta = (a, \phi(\mathbf{X}), \nu, f(\mathbf{X}))$. Specifically, a is the transition label, ν is the n-dimensional update vector, giving the net change in the number of agents for each state, $f(\mathbf{X})$ is the population dependent rate of the transition, encoding its average frequency, and $\phi(\mathbf{X})$ is a state dependent Boolean predicate such the the transition is active only when it evaluates to true.

A Markov population model is therefore defined by a vector of counting variables $\mathbf{X}$, a set of transition classes $\mathcal{T}$, and an initial state $\mathbf{x_0} \in S$. From this description, it is straightforward to derive the underlying CTMC, denoted by $\mathbf{X}(t)$, see [Bor+13] for further details. In the context of mean field limits, we are usually interested in the behaviour of the model when the size of the system N, usually the (initial) population of agents, tends to infinity. Classic mean field limits show that a Markov population model, with a proper scaling of rates and with variables divided by N, converges to a deterministic limit given by the solution of a differential equation. For this reason, the dependency on the total population N is made explicit in the notation: we write $\mathbf{X}^N$ in place of $\mathbf{X}$, and $f^N(N, \mathbf{X}^N)$ for the rate function. The rescaling of the model is done by defining normalised variables $\hat{\mathbf{X}} = \mathbf{X}/N$ and by changing variables from $\mathbf{X}$ to $\hat{\mathbf{X}}$ in $f^N(N, \mathbf{X}^N)$. For the classic mean field limit, we have further to require that $f^N(N, \hat{\mathbf{X}}^N)/N$ converges (uniformly) to a Lipschitz continuous function $f(\hat{\mathbf{X}})$, independent of N. This enables us to construct the limit drift $F(\hat{\mathbf{X}}) = \sum_{\eta_j \in \mathcal{T}} \nu_j f_j(\hat{\mathbf{X}})$, defining the mean field ODE

$$\frac{d}{dt}\mathbf{x}(t) = F(\mathbf{x}(t)). \tag{1}$$

to which the (normalised) sequence of CTMCs $\hat{\mathbf{X}}^N(t)$ converges (almost surely), when restricting to an arbitrary finite time interval $[0, T]$. This convergence can be extended to steady state, at the price of assuming ergodicity of the sequence of CTMCs and requiring that the ODE (1) has a unique globally attracting equilibrium (see [Bor+13] for more details).

2.2 Hybrid System as a Limit

The classic mean field limit works in scenarios in which all populations of agents are large, and can be, at least formally, taken to infinity. There are situations, however, in which this does not hold, as some populations are present in small numbers which do not grow with the system size N. This is the case when modelling genes explicitly in genetic networks [Bor15; BP13], or when modelling a centralised controller interacting with a distributed system, for instance, smart meters controlled centrally in a smart grid model.

In these cases, however, we can still rely on mean field results, at the price of constructing a more complex limit model, where discrete and continuous variables will coexist and coevolve. Formally, this can be described as a (stochastic) hybrid system. Here we will roughly sketch the main ideas behind these results, referring to [Bor15] for a detailed presentation and for the proofs of all the results.

The first step in the construction of a hybrid limit, starting from a Markov population model $(\mathbf{X}, \mathcal{T}, \mathbf{x_0})$, is to partition variables $\mathbf{X}$ into the two classes of discrete $\mathbf{Z}$ and continuous variables $\mathbf{Y}$, and identifying a notion of system or population size for continuous ones, thus writing $\mathbf{X}^N = (\mathbf{Z}, \mathbf{Y}^N)$. The idea now is to renormalise only continuous variables, hence changing coordinates from $(\mathbf{Z}, \mathbf{Y}^N)$ to $(\mathbf{Z}, \hat{\mathbf{Y}}^N) = (\mathbf{Z}, \mathbf{Y}^N/N)$. Such a change of variables is then applied to all rate functions of transition in $\mathcal{T}$.

The second step of the construction is concerned with a partition of transitions into continuous and discrete, depending on their effect on variables, namely if they modify only fast, only slow, or both kinds of variables, and on the dependency of the rate function on the system size N. For the moment, assume all guard predicates are tautologies, i.e. always evaluate to true. Then, in order to construct a hybrid limit, all transitions must belong to one of the three classes listed below.

- **Continuous transitions** modify only fast variables $\hat{\mathbf{Y}}^N$, and their rate $f^N(N, \mathbf{Z}, \hat{\mathbf{Y}}^N)$ is proportional to N, so that $f^N(N, \mathbf{Z}, \hat{\mathbf{Y}}^N)/N$ converges to a N-independent function $f(\mathbf{Z}, \hat{\mathbf{Y}})$.
- **Discrete stochastic transitions with no jumps** modify only discrete variables, and have a rate function $f^N(N, \mathbf{Z}, \hat{\mathbf{Y}}^N)$ which is independent of N: $f^N(N, \mathbf{Z}, \mathbf{Y}^N) = f(\mathbf{Z}, \hat{\mathbf{Y}})$. They can also modify continuous variables, but we assume the corresponding entry in the update vector is independent of N, so that the change in the normalised continuous variables in the limit is zero.
- **Discrete stochastic transitions with jumps** behave as the other class of discrete transitions in terms of rates, but the entry in the update vector for continuous variables can be proportional to N, so that in the limit each of these transitions induces a discontinuous jump in the continuous variables $\hat{\mathbf{Y}}$.

The hybrid limit. Once variables and transitions are partitioned into the aforementioned classes, we can easily construct a stochastic hybrid system, which can be shown to be the limit of the sequence $(\mathbf{Z}, \hat{\mathbf{Y}}^N)$ of population CTMCs. Formally, such a stochastic hybrid system can be defined in terms of a Piecewise Deterministic Markov Process, see [Dav93], but we refrain from providing formal details here. A more complete discussion can be found in [Bor15], where the convergence (in the weak sense) is formally stated and proved. Intuitively, the stochastic hybrid limit is defined as follows: its discrete skeleton is defined by all possible states that the discrete variables $\mathbf{Z}$ can take, while continuous variables $\hat{\mathbf{Y}}$ define its continuous state space. The dynamics of the system are given by a never-ending alternation of phases of continuous evolution and discrete jumps. In the continuous phase, continuous variables evolve following the solution of an ODE defined by a mode-specific vector field, obtained as in Equation (1) by restricting to continuous transitions only. Discrete jumps happen at exponentially distributed random times, according to the (continuous-state dependent) rate of discrete transitions. When a discrete transition fires, the discrete state of the automaton can change, as well as possibly the value of continuous variables, when a transition with jumps happens. After the jump, the system restarts its continuous dynamics from the new mode.

Guards and discontinuities. We turn now the focus on the effect of guards in the limit behaviour, distinguishing which kind of transitions are guarded.

Guards in continuous transitions induce discontinuities in the (mode-specific) vector field, as the corresponding entry in Equation (1) will be active only in a subset of the continuous state space (assuming guards involve continuous variables and not only discrete ones). This results in a mode-specific dynamics governed by a set of non-smooth ODEs, for which a mean field limit result similar to the classic one still holds, provided such discontinuous ODEs have a unique solution. This can be proved by working directly with non-smooth ODE [Bor11], or rephrasing the problem in the context of differential inclusions [GG12], which provides a link with mean field limits under uncertainty as will be discussed below.

Guards in discrete transitions, instead, result in discrete jumps active only in subsets of the (continuous) state space, but for the limit to be properly defined, extra conditions are required on the way the discontinuous surfaces defined by such guards interact with the vector field (essentially, it must be tangential to these surfaces). The corresponding conditions turn out to be difficult to check automatically, and randomisation of the continuous dynamics, for example in terms of stochastic differential equations, seems a possibility to circumvent such problems [Bor15].

Finally, in [Bor15] the author considers two additional classes of transitions: instantaneous ones, which fire as soon as their guard becomes true, and timed ones, which fire when the simulation time reaches a certain threshold. Both classes of transitions can be accommodated in the hybrid limit framework. In both cases they result in forced transitions, i.e. transitions of the stochastic hybrid system taken as soon as their guard becomes true.

We finally remark that we can also admit continuous transitions modifying discrete variables but with a rate growing linearly with N. In this case, however, a different construction applies [BLB08], as the dynamics of discrete variables induced in this way will happen on an infinitely faster time scale than the one caused by discrete transitions, leading to the necessity of computing the equilibrium distribution of discrete variables conditional on continuous ones and averaging away the effect of discrete variables from the rates of continuous transitions. In this sense, this phenomenon is closer to the time-scale separation scenario presented in the next subsection.

2.3 Computing the Jump Rate in Hybrid Systems

Some hybrid systems have transitions that fire as soon as the system hits a certain domain. In the hybrid framework described in Section 2.2 and in [Bor15], this corresponds to instantaneous transitions that are activated as soon as their guard becomes true. In this section, we propose a method based on a fluid approximation, to compute the time for a guard to become true, and hence to compute the jump rate of some hybrid systems.

Computing the time at which these transitions are fired is not always easy. For example, let us consider a population of individuals that can be either alive or dead and assume that an individual becomes dead after a random time that is exponentially distributed with mean 1. If N individuals are alive at time 0, then by using the mean-field framework of [BLB08], one can show that the proportion of living individual at time t is close to $\exp(-t)$ when N is large. Then, if an instantaneous transition is fired as soon as the proportion of living individuals becomes less than $1/2$, it is straightforward to show that this transition occurs after a time that tends to $\log 2$ as N becomes large [DN08]. This time is the solution of $\exp(-t) = 1/2$. However, if an instantaneous transition is fired only when the number of living individuals is zero, it can be shown that this instantaneous transition is fired at a time close to $\log N$. This time is not the solution of $\exp(-t) = 0$ but the solution of $\exp(-t) = 1/N$.

In [GG15], we develop an approach to compute the hitting time of a stochastic process whose fluid limit $m(t)$ tends to 0 as t goes to infinity. We establish two results. First, we show that the hitting time of the stochastic system is close to the time t_N for the fluid limit to reach $1/N$, *i.e.*, $m(t_N) = 1/N$. Then, we also compute an asymptotic development of t_N for a large class of stochastic systems and show a logarithmic trend in all cases: $t_N = c \log N + d \log\log N + O(1)$. For now, our method is limited to mean field models with few interactions between objects. We are now working on an extension of this method to more general stochastic systems. This will be helpful to understand the relationship between the extinction time of stochastic models and the time for its fluid approximation to get close to extinction.

2.4 Multi-Scale and Mean-Field

In this section we discuss a different kind of hybrid behaviour, emerging due to the presence of multiple time scales in a system. We stick here to the framework of quasi equilibrium (QE), which can be defined both in the stochastic and in the deterministic regime, following [BP14].

The starting point in this scenario is the separation of transitions of a given Markov population model into a fast and a slow class (more classes can be accommodated as well). In the QE framework, these are distinguished by the presence of a small parameter ϵ, defining the fast time scale of the order of $1/\epsilon$. Formally, one is interested in observing what happens in the limit $\epsilon \to 0$, giving an infinite separation of time scales. Identification of such a parameter is model specific, and often can be found among the parameters modulating rates of the model.

The second step in the construction of the QE reduction is to consider only the fast transitions (ignoring or freezing slow ones), and check if this modified model has additional conservation relations. These correspond to variables (possibly after a linear change of variables, see [BP14] for details), call them $\mathbf{Z}$, which are not modified by the fast transitions and thus identify, together with slow transitions, the *slow subsystem.* The remaining variables $\mathbf{Y}$, together with fast transitions, define the *fast subsystem.*

The main idea behind QE is that by taking the parameter ϵ to zero, the dynamics of the fast subsystem, relative to the slow one, will be faster and faster, to the point that one can safely assume the fast subsystem equilibrates instantaneously. The idea is then to compute the equilibrium distribution of the fast subsystem, conditional on a fixed value of the slow variables $\mathbf{Z}$, and remove the fast subsystem (variables and transitions) from the model, eliminating fast variables from the rates $f_j(\mathbf{Z}, \mathbf{Y})$ of slow transitions by taking their expectation $\mathbb{E}_{\mathbf{Y}|\mathbf{Z}}[f_j(\mathbf{Z}, \mathbf{Y})]$ with respect to the equilibrium distribution of $\mathbf{Y}$, conditional on the current value of $\mathbf{Z}$.

This operation reduces the dimensionality of the model, eliminating all fast variables. Moreover, the description above can be instantiated not only to a Population CTMC, but also to its mean field limit, with the difference that in this case the equilibrium distribution of $\mathbf{Y}$ will usually be given by an attracting equilibrium of the mean field ODE for the fast subsystem. This generality, however, poses the problem of understanding what happens in a population model in which we have a growing population N and an independent fast temporal scale $1/\epsilon$, when both are taken to the limit. In particular, in [BP14], we discussed formally when these two limits can be safely interchanged. It turns out that this is a non-trivial operation, requiring specific conditions on the mean field ODEs of the fast subsystem. More specifically, these ODEs must have a unique globally attracting equilibrium. These conditions are similar to those for extending classic mean field convergence to steady state, which is a consequence of the fact that the separation of times scales requires taking the fast subsystem to its equilibrium. Therefore, for limits on N and ϵ to be exchanged, mean field convergence must also hold at steady state for the fast subsystem. In [BP14], the authors provide also a counterexample violating such a condition on the fast subsystem, where limits cannot be exchanged. This provides a set of formal conditions to be checked, allowing time scale reduction at the level of the mean field ODEs to be safely carried out, which is the typical praxis, preserving mean field convergence.

Efficient simulation of stiff stochastic models. The time scale separation framework of [BP14] has been exploited in [BMS15] to define a novel simulation algorithm for population CTMCs with multiple time scales. The problem of simulating these systems in general is that the rate functions of the reduced models are seldom computable analytically, hence the fast subsystem has to be balanced by running a number of simulations of the fast subsystem after a slow transition occurs. The approach of [BMS15] obviates this problem by constructing an analytic approximation of slow rates $\mathbb{E}_{\mathbf{Y}|\mathbf{Z}} f_j(\mathbf{Z}, \mathbf{Y})$, as a function of $\mathbf{Z}$, sampling this function for a few values of the slow variables and then using machine learning tools to infer a statistical surrogate of the function $\mathbb{E}_{\mathbf{Y}|\mathbf{Z}} f_j(\mathbf{Z}, \mathbf{Y})$. In particular, the authors rely on Gaussian Process regression, which provides also an estimate of the error committed in the reconstruction. Learning the slow rates is a preliminary step that has to be done just once, resulting in considerable savings on the simulation time of the reduced system, compared with alternative slow scale simulation algorithms. Moreover, the authors show how to learn slow rates also as a function of some model parameters, at a negligible additional preprocessing cost, providing an efficient way to explore the parameter space and remove stiffness at the same time.

2.5 Hybrid Conditional Moments

In [Pou15], Pourranjbar (under the supervision of Hillston) has considered CAS which are multi-scale in the sense that they are comprised of some classes of agents with large populations interacting with other classes of agents which occur only in low copy numbers. Such models are readily identified when the model is developed in a formal modelling language, in this case PEPA [Hil95]. Direct application of fluid approximation [Hil05] or moment closure techniques to such models may lead to coarse approximation of the behaviour of the small population [PHB13]. Thus in [Pou15], the conditional moment closure technique of [Has+14], originally developed for modelling biological systems, has been adapted to the context described above. The small populations are represented explicitly and discretely, whilst the large populations are treated as continuous, thus a hybrid representation is generated. Using this approach, key characteristics of the probability distribution that represents the behaviour of the large populations, conditioned on the current state of the small populations, are approximated through conditional moments. The first-order moments are the expectations or averages related to the stochastic behaviour of the large populations given the stochastic evolution of the small populations. In contrast with fluid flow approximation, where the evolution of large populations is studied by a single expectation, here we calculate many expectations corresponding to the different configurations of the small populations. Working directly from the PEPA description, a set of differential algebraic equations (DAEs) are generated whose solution is the transient evolution of conditional expectations. A mode of operation is defined as a subset of states within the discrete state space which satisfy a property of interest (for instance, number of resources in failure mode). It is shown that the DAE solution can be used to derive coarse-grained conditional expectations for any mode of operation. In [Pou15] it is demonstrated that the analysis of conditional expectations can capture where significant probability masses are clustered.

The stochastic behaviour of large populations potentially makes deviations from the conditional expectations. Pourranjbar's work goes on to expand the analysis of conditional moments by including higher-order conditional moments. These include the conditional variances, conditional skewness, etc. given the different configurations of small populations. The higher-order moments enable us to obtain a richer representation of the conditional distributions. As for the conditional moments, a set of DAEs related to higher-order moments can be automatically derived from the model. This set of DAEs is larger than that constructed for conditional expectations; as the order increases, the set of equations is augmented with more equations and finding the solution becomes computationally more expensive. An analytical expression for the size of this set of equations is given and it is shown that, given the capabilities of current DAE solvers, the analysis for up to the third-order (conditional skewness) is practically possible for most models.

The usefulness of the technique of conditional moments has been demonstrated in the context of a simple scenario of a client-server system, and also for a more complex model of a two-tier wireless network, based on the femto-cell macro-cell architecture [CCU10].

3 Parameter Uncertainties: Estimation and Approximation

In this section, we briefly introduce a novel theoretical framework to describe population models in the presence of different kinds of uncertainty, and their mean field limits. We discus algorithmic issues to analyze the parameters of such models.

3.1 Uncertain Systems and Mean Field Limits as Differential Inclusions: Theory

As discussed in the introduction, uncertainty is a ubiquitous feature of models of complex systems, and collective adaptive systems are no exception to this rule. In the kind of models we are considering, uncertainty is mainly present in the values of model parameters. More specifically, we consider a population model according to the definition of Section 2.1, and assume one or more parameters ϑ are subject to some form of uncertainty. We mainly focus on two kinds of uncertainty:

1. Some parameters ϑ_t can depend on features of the environment external to the model. We may not know the precise value of ϑ_t, because we cannot measure these features. Furthermore, some environmental features like temperature, PH, atmospheric weather, light intensity, may be subject to variations during the time horizon T of interest, so that considering ϑ_t as having a fixed value may be an incorrect assumption that can lead to incorrect results. One way to capture such variability, without committing to assumptions on the form of dependency of ϑ_t on the external/environmental factors, is just to fix a set Θ of possible values for ϑ_t and assume

that ϑ_t depends on time t and can take any value of Θ at any time instant, i.e. that $\vartheta_t \in \Theta$. We call this the *imprecise* scenario.

2. In a simpler scenario, some parameters ϑ may be assumed fixed, but their values not known precisely. This may be the case, for instance, if ϑ has been estimated from experimental data. In this case, we just assume that $\vartheta \in \Theta$, where Θ is the possible set of values of ϑ, as above. This will be referred to as the *uncertain* scenario.

It is easy to see that the second case is a simpler subcase of the first, assuming the dependency of ϑ on time is constant. Hence, in the following we will deal mostly with imprecise models, though a dedicated discussion of the uncertain scenario is useful as analysing these models is simpler.

3.1.1 Imprecise Markov Population Models

We will start by the definition of imprecise CTMCs, introducing in the following the conditions to be satisfied by an uncertain population process. Consider a stochastic process $\mathbf{X} = (X_t)_{t\geq 0}$, adapted to a filtration[1] $\mathcal{F}$, that takes values in a state space $\mathbf{E} \subset \mathbb{R}^d$. The dynamics of the process depends on a parameter (or a vector of parameters) ϑ. We denote by Θ the set of possible parameter values of ϑ, and we consider a set of infinitesimal generators on $\mathbf{E}$, parametrised by $\vartheta \in \Theta$: for each $\vartheta \in \Theta$, Q^ϑ is a transition kernel, *i.e.* such that $Q^\vartheta_{xy} \geq 0$ for $x \neq y \in \mathbf{E}$ and $\sum_{y\in\mathbf{E}} Q^\vartheta_{xy} = 0$.

Definition 3.1. *An* imprecise continuous time Markov chain *is a stochastic process* $\mathbf{X}$ *together with a* $\mathcal{F}_t$*-adapted process* ϑ *such that for all* $t \geq 0$*:*

$$\lim_{h\to 0} \frac{1}{h}\mathbb{P}(X_{t+h} = y \mid \mathcal{F}_t \text{ and } X_t = x) = \begin{cases} Q^{\vartheta_t}_{xy} & \text{if } x \neq y \\ -\sum_{y\neq x} Q^{\vartheta_t}_{xy} & \text{otherwise} \end{cases}$$

The definition of an imprecise Markov chain makes no restriction on the set of processes to which the varying parameter ϑ belong to. In some cases, it can be interesting to focus on a subset of processes. In particular, if we assume that ϑ_t is deterministic and constant in time, we obtain the notion of *uncertain continuous-time Markov Chain*. The evolution in time of the probability mass of an imprecise CTMC is described by a generalised form of Kolmogorov equations, expressed in terms of a differential inclusion (see [BG15] and the next section for further details).

Imprecise population models are imprecise CTMC satisfying additional constraints. In particular, we consider sequences of population models depending on a scaling parameter N (typically, N is the population size of the considered model). Such a sequence is denoted $(X^N)_N$, and takes values on a sequence of subsets $\mathbf{E}^N \subseteq \mathbf{E} \subseteq \mathbb{R}^d$. The stochastic process $\mathbf{X}^N$ is an imprecise process of kernel $Q^{N,\vartheta}$.

Definition 3.2. *An imprecise (respectively uncertain) population process is a sequence of imprecise (respectively uncertain) Markov chains that satisfies the following assumptions:*

(i) The chains are uniformizable: i.e.*, for all* N*:* $\sup_{x\in\mathbf{E}^N,\vartheta\in\Theta} Q^{N,\vartheta}_{xx} < \infty$

(ii) The transitions become smaller as N *grows,* i.e.*, there exists* $\varepsilon > 0$ *such that*

$$\lim_{N\to\infty} \sup_{x\in\mathbf{E}^N,\vartheta\in\Theta} \sum_{y\in\mathbf{E}^N} Q^{N,\vartheta}_{xy} \|y - x\|^{1+\varepsilon} = 0$$

(iii) The drifts are well-defined and bounded:

$$\limsup_{N\to\infty} \sup_{x\in\mathbf{E}^N,\vartheta\in\Theta} \sum_{y\in\mathbf{E}^N} Q^{N,\vartheta}_{xy} \|y - x\| < \infty$$

[1] A filtration is a set of σ-algebra $(\mathcal{F}_t)_{t\geq 0}$ such that $\mathcal{F}_s \subset \mathcal{F}_t \subset \mathcal{F}$ for each $t \geq s \geq 0$. $\mathbf{X}$ is adapted to $\mathcal{F}$ means that X_t is $\mathcal{F}_t$ measurable for every t.

Typically, these models are derived from a description in terms of transition classes, as that presented in Section 2.1. The conditions for mean field described in that section, together with rate functions bounded uniformly on x and ϑ, guarantee the satisfaction of the previous definition. In particular, boundedness of update vectors with respect to N and the scaling condition of rate functions implies condition (iii) above.

Given an imprecise population model, we define its *imprecise drift* as

$$f^N(x, \vartheta) = \sum_{y \in \mathbf{E}} Q^{N,\vartheta}_{xy}(y - x),$$

which makes sense provided for all $x \in \mathbf{E}, \vartheta \in \Theta$ we have $\sum_{y \in \mathbf{E}} Q^{\vartheta}_{xy} \|y - x\| < \infty$. This quantity will play a central role in the definition of the mean field limits.

3.1.2 Differential Inclusions

Before stating the mean field theorem, we provide a very short introduction to differential inclusions (see [AC84] for further details).

Let F be a set-valued function on $\mathbf{E} \subset \mathbb{R}^d$ that assigns to each $x \in \mathbf{E}$ a set of vectors $F(x) \subset \mathbb{R}^d$. A solution to the differential inclusion $\dot{x} \in F(x)$ that starts in x_0 is a function $\mathbf{x} : [0, \infty) \to \mathbf{E}$ such that there exists a measurable function f satisfying for all $t \geq 0$ $f(t) \in F(x(t))$ and

$$\mathbf{x}(t) = x_0 + \int_0^t f(s)ds.$$

For an initial condition x_0, we denote by S_{F,x_0} the set of solutions of $\dot{x} \in F(x)$ that start in x_0. Note that the set S_{F,x_0} can be empty or be composed of multiple solutions, depending on the function F. When $\mathbf{E} = \mathbb{R}^d$, a sufficient condition for the existence of at least one solution is that (a) for all $x \in \mathbb{R}^d$, $F(x)$ is non-empty, convex and bounded (*i.e.*, $\sup_{x \in X, y \in F(x)} \|y\| < \infty$) and (b) F is upper-semi-continuous (*i.e.*, the graph of F, $\{(x, y), y \in F(x)\}$, is a closed set).

As time goes to infinity, solutions of a differential inclusion can show a variety of different behaviours, for instance, they can have a chaotic behaviour or they can oscillate around many points. What can be said in the generic case is that the set of points containing solutions of the differential inclusion in the limit of an infinite time are contained in the so-called Birkhoff centre, formally defined as the closure of the set of recurrent points of the differential inclusion. Intuitively, the Birkhoff centre contains all attractors, equilibria, limit cycles, and in general all points of the differential inclusion that are visited infinitely often.

3.1.3 Mean Field Limits as Differential Inclusions

The mean field limit of a sequence $\mathbf{X}^N$ of imprecise population processes is a differential inclusion, specified by the set-valued function constructed from the imprecise drift. More specifically, let $f^N(x, \vartheta)$ be the drift of the system for size N. We define the limit drift of the system as the convex closure of the set of the accumulation points of $f^N(x^N, \vartheta)$ as N goes to infinity, for all sequences x^N that converge to x:

$$F(x) = \overline{\lim_{N \to \infty}} \bigcup_{\vartheta \in \Theta} \left\{ f^N(x, \vartheta) \right\}. \tag{2}$$

Then the mean field differential inclusion is given by $\dot{x} \in F(x)$, and asymptotic convergence is proved in the following theorem (cf. [BG15] for a proof).

Theorem 3.1. *Let (X^N) be an imprecise population process. Then, if $\mathbf{X}^N(0)$ converges (in probability) to a point x, then for any finite time horizon T, the stochastic process $\mathbf{X}^N$ converges (in probability) to $S_{F,x}$, the set of solutions of the differential inclusion $\dot{x} \in F(x)$ starting in x.*

What happens as time tends to infinity? As usual in the mean field setting, only weaker results hold. What can be said in this case is that, despite the fact that an imprecise Markov population process does not necessarily have a stationary behaviour, it is possible to constrain its asymptotic regime, showing that with probability one it will be contained in the asymptotic reachable set of the differential inclusion A_F [BG15], defined by:

$$A_F = \bigcap_{T>0} \bigcup_{x,t\geq T,\mathbf{x}\in S_{F,x}} \{\mathbf{x}(t)\} \tag{3}$$

Note that the set A_F is included in the Birkhoff centre B_F and such an inclusion is in general strict.

3.2 Uncertain Systems and Mean Field Limits as Differential Inclusions: Numerical Methods

Mean field limits of imprecise population models are given by differential inclusions (of dimension much smaller than Kolmogorov differential inclusions). However, differential inclusions are much more complicated to analyse numerically than differential equations. Hence, to provide a practical value to the theoretical results of the previous section, we need to provide efficient numerical routines to compute the flow of a differential inclusion, or at least to over-approximate it.

In literature there has been much research, especially in the context of reachability computation for hybrid systems. Here we will discuss two methods for generic differential inclusions, one based on *differential hulls* [TT15] and one based on the Pontraygin principle [BG15]. We also discuss a more efficient method tailored on the simpler class of limit models for uncertain CTMC.

3.2.1 Differential Hull of a Differential Inclusion

The idea of this method is to construct rectangular bounds for the differential inclusion, *i.e.*, two functions $\underline{x}$ and $\overline{x}$ such that $\underline{x}(t) \leq x(t) \leq \overline{x}(t)$ for any solution $x(t)$ of the differential inclusion $\dot{x} \in F(x)$. The method presented in [TT15] is based on the idea of defining a simple set of differential equations that provide such upper and lower bounds, called the differential hull. These bounds are reasonably tight when the set of possible values that the parameters θ can take is small. However, the bounds provided by this approximation become too large when the possible values of ϑ increase (see also [BG15]).

It can be shown [TT15] that the tightest functions $\underline{f}$ and $\overline{f}$ that are a differential hull for a differential inclusion F are:

$$\underline{f}_i(\underline{x},\overline{x}) = \min_{x\in[\underline{x},\overline{x}]:x_i=\underline{x}_i(t)} F_i(x)$$
$$\overline{f}_i(\underline{x},\overline{x}) = \max_{x\in[\underline{x},\overline{x}]:x_i=\underline{x}_i(t)} F_i(x).$$

3.2.2 Reachability of a Differential Inclusion as an Optimal Control Problem

An alternative method discussed in [BG15] is based on Pontryagin's maximum principle to compute the exact minimal value $x^{\min}(t)$ and maximal value $x^{\max}(t)$ that can be reached by an imprecise fluid model at time t.

Let $T \geq 0$ be some fixed time and $i \in \{1 \dots d\}$ a coordinate. Let $x_i^{\min}(T) = \inf_{\mathbf{x}\in S_{F,x}} x_i(T)$ be the minimal value that the ith coordinate of the solution of a differential inclusion can take at time t. The quantity $x_i^{\min}(T)$ is the solution of the minimisation problem:

$$x_i^{\min}(T) := \min_{\theta} x_i(T) \text{ such that for all } t \in [0,T]\text{:} \begin{cases} x(t) = x + \int_0^t f(x(s),\theta(s))ds \\ \theta(t) \in [\vartheta_{\min},\vartheta_{\max}] \end{cases}$$

Pontryagin's maximum principle provides an algorithmic method to solve this optimisation problem. Following [Tod06, Section 3], if x is a trajectory that maximises $x_i(T)$, then there exists a *costate trajectory* $\mathbf{p}$ such that $p_i(T) = -1$, $p_j(T) = 0$ for $j \neq i$ and:

$$\dot{x}(t) = f(x(t), u(t)) \tag{4}$$

$$\vartheta(t) \in \arg\min_{\vartheta} f(x, \vartheta)^T p, \tag{5}$$

$$-\dot{p}(t) = \frac{\partial}{\partial x}\left(f(x, \vartheta)^T p\right) \tag{6}$$

where $f(x, \vartheta)^T p$ denotes the scalar product of $f(x, \vartheta)$ and p. This solution leads to iterative numerical methods that start from an initial costate p, then updates p by computing a trajectory x forward in time from x by using (4) and (5) and then computing a new p backward in time by solving the ODE (6).

3.2.3 Statistical-Based Computation of Reachable Sets for Parameter Uncertainty

In [BS14], the authors introduce a novel method to compute over-approximations of reachable sets for dynamical systems subject to parameter uncertainty, such as those obtained as mean field limits of uncertain population processes.

The authors assume uncertain models as defined in Section 3.1, meaning that parameters have fixed but unknown values, and belong to a certain real interval. The method presented in [BS14] works for computing the reachable set at time T, or up to time T, for a model specified by a set of ODEs, or for any quantity which is a deterministic function of time, like the expectation of a stochastic model. More specifically, letting $\mathbf{x}(t, \theta)$ be such a function, depending on time and on some uncertain parameters $\theta \in D$, the authors propose an approach constructing a polytope bounding the reachable set $R(T) = \{\mathbf{x}(T, \theta) \mid \theta \in D\}$. The idea is that the polytope $A\mathbf{x} \leq \mathbf{b}$ will be specified by fixing the matrix A (meaning fixing the direction of hyperplanes delimiting the polytope) and then by choosing the vector $\mathbf{b}$, of dimension k, as the (coordinate-wise) solution of k optimisation problems: $\mathbf{b} = \max_{\mathbf{x} \in R(T)} A\mathbf{x}$.

Bayesian optimisation is used to solve these problems. In practice, the function $\mathbf{x}(T, \theta)$ is sampled at few values $\theta_1, \ldots, \theta_N$ of the parameters, and then a statistical emulation of the function is obtained by means of Gaussian Process regression. Then an upper quantile of the so-obtained Gaussian Process is optimised, getting a candidate maximum at a certain θ^*. Next, the true function $\mathbf{x}(T, \theta^*)$ is computed in θ^*. These emulation and optimisation steps are carried out iteratively until no further improvement is possible. This algorithm is provably convergent (in probability) to the true optimum [Sri+12].

The advantage of this method is that it requires only few computations of the function $\mathbf{x}(t, \theta)$, a particularly convenient feature when such a function is expensive to compute, like the average or another statistic indicator of a stochastic process, obtained statistically from a number of simulation runs.

3.3 Representing Heterogeneous Systems with an Unknown Heterogeneity

The previous sections assume that a model of the system has a unique parameter that is unknown and that may vary in time. Yet many systems are composed of many entities that are structurally similar, yet behave differently. This can be represented by a mathematical model for which the equations describing the behaviour of all entities have the same shape but different parameters. In [TT15], we consider models defined as nonlinear ordinary differential equations (ODEs) and develop an approximate reduction of such models.

This method combines ideas from the theories of differential inequalities [RMC09] and lumpability of ODEs (for example [OM98]) to provide a method that can give bounds on the solution of a heterogeneous model by means of an ODE system which preserves the structure but is characterised by parameters that represent the extreme values found in the original model. As in Section 3.2.1, we use a differential hull to provide a lower and upper bound for every state of every distinct entity in the original system.

We then show that this differential hull can be automatically reduced, by employing model-order reduction techniques. We develop a generalisation of [TT12], called uniform lumpability, to partition the state space such that variables in the same partition block have equal solutions if they have equal initial conditions. This method is complementary to exact lumpability, which transforms of the original state space through a linear mapping [Tom+97]. In [TT15], we illustrate and compare these methods with two cases studies: a multiclass model of epidemic spread and a multiclass queuing network.

The differential hull approximation provides guarantees on upper and lower bounds on the behaviour of a heterogeneous system and are fast to compute. Hence, in practice they can be readily used in parameter design, to find values of parameters that provide some required dynamical behaviour. As shown in [BG15; TT15], these bounds are tight when the set of possible parameter values is small. However, when the set of parameter values is large, they are sometimes too loose to be exploitable. In such cases, the development of approximation techniques such as the one of Section 3.2.2 might prove to be useful.

4 Discussion

Summary. In this deliverable we report on the progress made on Task 1.1 since the first reporting period. We address two main questions that are important when modelling CAS. The first one is what types of models to use to describe a system that exhibits multiple scales. The second is how to deal with models for which some parameters are not known or are unpredictably varying in time.

We prove two different results related to multi-scale systems. When there are multiple organisational scales, we showed that it is reasonable to construct a model as a hybrid system, where part of the system is described by a fluid approximation while the rest is kept discrete and stochastic. We have shown convergence results in a general setting, including both instantaneous and guarded transitions. In addition, we investigated the use of hybrid conditional moment techniques in combination with models obtained from the stochastic process algebra PEPA. When multiple time scale are present, we provided conditions guaranteeing the correctness for the combined use of mean field and time-scale reduction techniques. Exploiting machine learning techniques, we also defined a novel fast simulation algorithm for stochastic systems with multiple time-scales.

Our approach to deal with the uncertainty inherent in CAS models is to consider some model parameters as fixed but unknown or time-varying. The complexity of the obtained model depends on how we allow these parameter to vary. We formulate mean-field convergence results, both for the transient and the steady state, and develop numerical techniques to compute the reachable states. We present in particular a method based on statistical evaluation for the uncertain case, and two methods, one based on differential hulls and one based on Pontraygin optimal control principle for the imprecise case. These methods are related to parameter estimation and prediction using qualitative observations of system behaviour, encoded as linear-time temporal logic properties.

Future work. At this stage of the project, we have developed solid theoretical foundations for scalable analysis of CAS using mean-field based techniques and proposed several algorithmic approaches for the analysis, which exploit our novel approximation results. The most important step now is to bring these methods to a practical stage, mostly in two directions:

- Connect the Carma language with the plethora of mean-field analysis methods, by developing appropriate primitives and semantics and by investigating ways to detect/ choose the best mean field approximation for a given system. This is discussed more in internal report 1.1.

- Provide efficient implementations of the algorithms for the project toolchain. In particular, we will provide implementations of a simulator for the hybrid mean field semantics and of the algorithms to study uncertainty and imprecision, possibly for models represented as Markov Population Processes, and link them to the Carma language by the appropriate semantic constructs that are being developed in Task 1.3.

In addition, we also need to extensively test our methods on the project case studies, to better understand the gains in using them and the margins of their applicability. In this direction, an important benchmark will be the prediction of availability in bike-sharing systems.

Connection with other work packages. We discuss now connections with other work packages and deliverables, also in terms of future work.

WP1 The discussion of Section 2 on multi-scale systems is a natural follow-up of Deliverable 1.1. In particular, here we present a finalised framework about hybrid mean field limits, particularly in the presence of guards and instantaneous transitions. We also discuss novel work with respect to multiple time-scales, particularly the relationship between mean field and time-scale reduction techniques. Furthermore, the outcomes of WP1 are collected both in this document and Deliverable 1.3, which focuses on adaptive policies. Deliverable 1.3 also contains results on heterogeneous and uncertain systems [GVH15] but in the context of the study of a particular adaptive policy.

WP2 Many techniques for the analysis of space developed in WP2 can be naturally combined with the framework for imprecision considered here. In particular, discrete representations of space can be dealt with in the classical mean-field framework, hence both the hybrid mean field and the mean field under imprecision can be applied to them. The way of dealing with uncertainty presented in this report, in particular, can be appealing to deal with spatial moment closures techniques (particularly spatial averages and higher order moments and pair approximation) in the presence of uncertainty in parameters and of spatial heterogeneity. We will pursue this direction, also from the point of view of tool integration. Theoretically, it would be interesting to investigate how uncertainty and imprecision can be included in spatial mean field limits resulting in partial differential equation limits.

WP3 The algorithmic techniques presented in this report, and in particular the Pontraygin principle, may be used as an alternative analysis approach in the framework of ODE aggregation under parameter variability, which results in a differential inclusion approximation of the original ODE system. This is a clear link with Task 3.2 in WP3. Furthermore, the algorithmic approach to deal with imprecision and uncertainty based on the Pontraygin principle can be lifted to compute the quantitative satisfaction score of spatio-temporal logic formulae, evaluated on the mean field limits.

WP4 An important aspect of the project is to lift the scalable analysis techniques developed in this work package to CARMA. In this respect, there are two lines of work. One direction consists of developing a hybrid semantics of CARMA models, approximating part of the system as continuous and keeping the rest discrete. In this respect, the environment in CARMA is a natural candidate to be modelled as a discrete component of the system, provided it changes at a speed independent of the total population. More generally, we need to develop appropriate static analysis routines to identify those system components to be kept discrete, for instance exploiting algorithms to detect conservation laws. Furthermore, in order to link CARMA with the imprecise mean field semantics, we need suitable language constructs to specify uncertainty in parameters and imprecision at the environment level. This is discussed in internal report 1.1.

WP5 In [Gas+15], we have used parameter estimation of an uncertain system for prediction of bike-sharing systems. This parameter estimation should be integrated in a tool-chain that would allow an automatic answer for queries about the quantitative quality of some predictors. Furthermore, as discussed above, by linking the hybrid and uncertain semantics with CARMA, we will integrate in the project toolchain the algorithms to simulate and analyse the class of mean field models considered in this report.

References (from the Quanticol project within the reporting period)

[BG15] L. Bortolussi and N. Gast. "Mean Field Approximation of Imprecise Population Processes". In: *QUANTICOL Technical Report TR-QC-07-2015* (2015).

[BMS15] L. Bortolussi, D. Milios, and G. Sanguinetti. "Efficient stochastic simulation of systems with multiple time scales via statistical abstraction". In: *Proceedings of Computational Methods in Systems Biology, CMSB 2015*. 2015.

[Bor15] L. Bortolussi. "Hybrid Behaviour of Markov Population Models". In: *Information and Computation* accepted (2015).

[BP14] L. Bortolussi and R. Paškauskas. "Mean-Field approximation and Quasi-Equilibrium reduction of Markov Population Models". In: *Proceedings of Eleventh International Conference on the Quantitative Evaluation of Systems, QEST 2014*. 2014.

[BS14] L. Bortolussi and G. Sanguinetti. "A statistical approach for computing reachability of non-linear and stochastic dynamical systems". In: *Proceedings of Quantitative Evaluation of Systems*. Springer, 2014, pp. 41–56.

[Gas+15] N. Gast, G. Massonnet, D. Reijsbergen, and M. Tribastone. "Probabilistic forecasts of bike-sharing systems for journey planning". In: *Proceeding of the 24th ACM International Conference on Information and Knowledge Management (CIKM'15)*. ACM, 2015, to appear.

[GG15] N. Gast and B. Gaujal. "Hitting time and fluid approximation: the Poisson approximation". In: *working paper* (2015).

[GVH15] N. Gast and B. Van Houdt. "Transient and Steady-state Regime of a Family of List-based Cache Replacement Algorithms". In: *ACM SIGMETRICS 2015*. 2015.

[PHB13] A. Pourranjbar, J. Hillston, and L. Bortolussi. "Don't Just Go with the Flow: Cautionary Tales of Fluid Flow Approximation". In: *EPEW*. Vol. 7587. LNCS. Springer, 2013, pp. 156–171. DOI: 10.1007/978-3-642-36781-6_11. URL: http://dx.doi.org/10.1007/978-3-642-36781-6_11.

[TT15] M. Tschaikowski and M. Tribastone. "Approximate Reduction of Heterogenous Nonlinear Models with Differential Hulls". In: *IEEE Transactions on Automatic Control* (2015).

References

[AC84] J. Aubin and A. Cellina. *Differential Inclusions*. Springer-Verlag, 1984.

[BLB08] M. Benaim and J.-Y. Le Boudec. "A class of mean field interaction models for computer and communication systems". In: *Performance Evaluation* 65.11 (2008), pp. 823–838.

[Bor+13] L. Bortolussi, J. Hillston, D. Latella, and M. Massink. "Continuous approximation of collective systems behaviour: a tutorial". In: *Performance Evaluation* (2013). URL: http://www.sciencedirect.com/science/article/pii/S0166531613000023.

[Bor11] L. Bortolussi. "Hybrid Limits of Continuous Time Markov Chains". In: *Proceedings of Eighth International Conference on the Quantitative Evaluation of Systems, QEST 2011*. IEEE Computer Society, 2011, pp. 3–12. ISBN: 978-0-7695-4491-5. DOI: 10.1109/QEST.2011.10.

[Bor12] L. Bortolussi. "Hybrid Behaviour of Markov Population Models". In: *CoRR* abs/1211.1643 (2012).

[BP13] L. Bortolussi and A. Policriti. "(Hybrid) automata and (stochastic) programs. The hybrid automata lattice of a stochastic program". In: *Journal of Logic and Computation* 23.4 (2013), pp. 761–798. ISSN: 0955-792X, 1465-363X. DOI: 10.1093/logcom/exr045. URL: http://logcom.oxfordjournals.org/cgi/doi/10.1093/logcom/exr045.

[CCU10] D. Calin, H. Claussen, and H. Uzunalioglu. "On femto deployment architectures and macrocell offloading benefits in joint macro-femto deployments". In: *IEEE Communications Magazine* 48.1 (2010), pp. 26–32. DOI: 10.1109/MCOM.2010.5394026. URL: http://dx.doi.org/10.1109/MCOM.2010.5394026.

[Dav93] M. H. A. Davis. *Markov Models and Optimization*. Chapman & Hall, 1993.

[DN08] R. Darling and J. R. Norris. "Differential equation approximations for Markov chains". In: *Probability surveys* 5 (2008), pp. 37–79.

[GG12] N. Gast and B. Gaujal. "Markov chains with discontinuous drifts have differential inclusion limits". In: *Performance Evaluation* 69.12 (2012), pp. 623–642.

[Has+14] J. Hasenauer, V. Wolf, A. Kazeroonian, and F. Theis. "Method of conditional moments (MCM) for the chemical master equation". In: *Journal of Mathematical Biology* 69.3 (2014), pp. 687–735.

[Hil05] J. Hillston. "Fluid Flow Approximation of PEPA Models". In: *2nd International Conference on the Quantitative Evaluation of Systems*. IEEE. 2005, pp. 33–42. DOI: 10.1109/QEST.2005.12.

[Hil95] J. Hillston. *A Compositional Approach to Performance Modelling*. CUP, 1995.

[OM98] M. S. Okino and M. L. Mavrovouniotis. "Simplification of mathematical models of chemical reaction systems". In: *Chemical reviews* 98.2 (1998), pp. 391–408.

[Pou15] A. Pourranjbar. "Performance Analysis of Large-Scale Resource-Bound Computer Systems". PhD thesis. School of Informatics, University of Edinburgh, 2015.

[RMC09] N. Ramdani, N. Meslem, and Y. Candau. "A hybrid bounding method for computing an over-approximation for the reachable set of uncertain nonlinear systems". In: *Automatic Control, IEEE Transactions on* 54.10 (2009), pp. 2352–2364.

[Sri+12] N. Srinivas, A. Krause, S. M. Kakade, and M. W. Seeger. "Information-Theoretic Regret Bounds for Gaussian Process Optimization in the Bandit Setting". In: *IEEE Transactions on Information Theory* 58.5 (2012), pp. 3250–3265. DOI: 10.1109/TIT.2011.2182033.

[Tod06] E. Todorov. "Optimal control theory". In: *Bayesian brain: probabilistic approaches to neural coding* (2006), pp. 269–298.

[Tom+97] A. S. Tomlin, G. Li, H. Rabitz, and J. Tóth. "The effect of lumping and expanding on kinetic differential equations". In: *SIAM Journal on Applied Mathematics* 57.6 (1997), pp. 1531–1556.

[TT12] M. Tschaikowski and M. Tribastone. "Exact fluid lumpability for Markovian process algebra". In: *CONCUR*. 2012, pp. 380–394.