QUANTICOL

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

http://www.quanticol.eu

D1.4

A mean field framework for evaluating and controlling CAS

Revision: 1.0; March 30, 2017

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

Due date of deliverable: Month 48 (March 2017)
Actual submission date: March 30, 2017
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

Contents

Executive Summary

This deliverable reports on the work that was accomplished within workpackage 1 to extend the applicability of mean-field approximation. Mean-field approximation refers to a collection of techniques that make the analysis of a large stochastic systems easier. Before the start of the QUANTICOL project, these approximations were known to be applicable to systems composed of a large population of homogeneous individuals. In the previous deliverables, we reported new techniques that allowed us to study multi-scale systems, hybrid behaviour and uncertain systems. We illustrated the use of these techniques with our case studies, and in particular bike-sharing and smart-grid. During this reporting period, we continued this work, by developing both the theory and its application. Several contributions are reported in this deliverable.

Our first contribution has been to analyse the accuracy of mean-field approximation. Classical mean-field approximation consists of replacing the study of a system with N objects by its limit as N goes to infinity. We obtained conditions under which the distance between these two quantities decreases as c/N, where c is a constant that depends on the system's parameters. This allowed us to defined a **refined approximation** that depends on the system size N. For small N (*e.g.*, $N = 10$ or $N = 20$), this new approximation is much more accurate that the original mean-field approximation.

The second reported contribution concerns the development of **numerical algorithms to quantity the effect of uncertainties** on the parameters of the systems. Within the QUANTICOL project, we construct and study stochastic models of collective adaptive systems. These models assume that the sequence of events that will occur in a system is unknown but they assume that the probability of the occurrence of an event is known. When these probabilities are not known, we talk of an uncertain system. In this work, we developed different numerical techniques that allowed us to study quantitatively the properties of such systems.

The above methods are descriptive. They allow one to estimate the performance of a given policy, which can in turn be used to compare between different heuristics and choose the most appropriate one. Our third contribution concerns the use of these approximations to **create efficient control policies**. We make contributions to the theory of control by comparing centralised and decentralised policies. In particular, we obtained a counter-intuitive result, that is that mean-field games are not the unique limit of stochastic N-players games.

Our last contribution concerns the design of **control algorithms for smart-grid systems**. These algorithms are motivated by the increasing share of renewable sources of energy – like photo-voltaic panels or wind turbines. This creates two problems: (1) these sources of energy are distributed and produce locally; and (2) they are volatile and intermittent. We develop two approaches to deal with this issue: a real-time control algorithm that we plan to deploy on a distribution network; a market framework that could be used to encourage people to smoothe their consumption.

1 Introduction

Collective adaptive systems (CAS) can be effectively described by stochastic population models. These systems are composed of a set of objects, agents, or entities interacting with each other. Each individual agent is typically described in a simple way, as a finite state machine with few states. An agent changes state spontaneously or by interacting with other agents in the system. All transitions happen probabilistically and take a random time to be completed. By choosing exponentially distributed times, the resulting stochastic process is a continuous-time Markov chain with a finite state space.

Markov chains suffer from the state-space explosion: the state-space needed to represent a system with N agents grows exponentially with the number of agents. As a result, one needs to resort to approximation or simulations. The starting point of WP1 is that mean-field approximation is a very effective technique to characterise the transient probability distribution or steady-state regime of such systems when the number of entities N grows very large. The idea of mean-field approximation is to replace a complex stochastic system – X^N – by a simpler deterministic dynamical system – x. This dynamical system is constructed by assuming that the objects are asymptotically independent. Each object is viewed as interacting with an average of the other objects (the mean-field). When each object has a finite or countable state-space, this dynamical system is usually a non-linear ordinary differential equation (ODE).

The overall objective of Workpackage 1 of the QUANTICOL project has been to extend and refine these techniques to adapt to specific properties of CAS. Some of these extensions, in particular related to multi-scale modelling and hybrid behaviour, are not reported here as they were already reported in Deliverable 1.1 and Deliverable 1.2. In this deliverable, we report on the most recent advances on mean-field techniques. We group in three categories:

1. **A refined mean-field approximation** – In many papers, the mean-field approximation is justified by showing that the original stochastic system X^N converges (in a sense made precise below) to a deterministic limit x. In our work, we studied different techniques to compute the exact rate of convergence. We then use these ideas to construct a refined approximation that, for a given time-budget, is more accurate than simulation. These results are reported in Section 2.

2. **Uncertain and heterogeneous systems** – The construction of a Markov model of a system requires setting the values of various parameters (*e.g.*, rate of transitions, transition probabilities, etc.) In many practical situations, an exact value of these parameters is hard to set because these parameters have to be estimated from experience or because the agents are all heterogeneous. We developed a few numerical methods to overcome this issue, some of them have been implemented in the QUANTICOL toolchain (see Deliverable 5.3). The results are reported in Section 3.

3. **Control algorithms with applications to smart grid** – The above methods are essentially descriptive: starting from a model of the system, they help to characterise its emergent behaviour. Using this emergent behaviour, we can then create control algorithms that ameliorate the performance of the system. In Section 4, we described some of these methods, with an emphasis on decentralised and centralised control. In Section 5, we describe our progress in the context of the optimal control of smart-grid systems. This section is a continuation of the work developed in Deliverable 1.3.

2 Classical and Refined Mean-Field Approximation

2.1 The Classical Mean-Field Approximation

Mean-field approximation consists of constructing a deterministic approximation of a stochastic population model. We recall here the definition and the class of models to which it applies. A more detailed introduction to this technique is provided in the book chapter [4].

A population process is given by a tuple $(X, \mathscr{T}, N)$ where:

- $X \in R^d$ is the population vector. $X_i(t)$ represents the number of agents that are in a state i at time t.

- $\mathscr{T}$ is the set of transitions, of the form $(\ell, \beta_\ell(\cdot))$. ℓ is an update vector and β_ℓ a rate function: at rate $\beta_\ell(x/N)$, the state x becomes $x+\ell$.
- N is a scaling factor (in general the total number of agents)

We call $X^{(N)}$ the normalised process : $X^{(N)}(t) = \frac{1}{N}X(t)$.

Drift and mean-field approximation. The main quantity required to define mean-field equations is the drift. The drift is the average direction of change of the population model, conditional on being in a certain state at some time t. For a state x, the drift is defined as

$$f(x) = \sum_{\ell \in \mathscr{T}} \ell \beta_\ell(x).$$

The mean-field approximation of the stochastic system is then given as the quantity:

$$\dot{x} = f(x).$$

It can be shown that under mild conditions (essentially if the drift f is Lipschitz-continuous f) then, as the population size N grows to infinity, the stochastic process $X^{(N)}$ converges to x as N goes to infinity [4, Theorem 1]. More precisely, if x denotes the solution starting in x_0 and if the initial conditions $X^{(N)}(0)$ converges to x_0 as N (in probability), then for all T: $\sup_{t \le T} \left\| X^{(N)}(t) - x(t) \right\|$ converges to 0 (in probability).

When one wants to compute values for a finite time-horizon T, the necessary conditions to apply this result can be done by a syntactic analysis of the model. This is more complicated for steady-state analysis [25], in which case special Lyapunov functions have to be found (see for example [10]).

This is illustrated in Figure 1(a) in which we consider the power-of-two-choice model studied in [34, 38, 11]. This model describes a load-balancing strategy in a cluster with N servers. We plot the proportion of servers that have 2 jobs to serve as a function of time. We plot one trajectory obtained by simulation of $X_2^{(N)}$ for each value of $N \in \{10, 100, 1000\}$. We also display the mean-field approximation. We observe that indeed, as N converges to $+\infty$, the trajectory converges to its mean-field approximation. For $N = 1000$, the behaviour is matched with high precision.

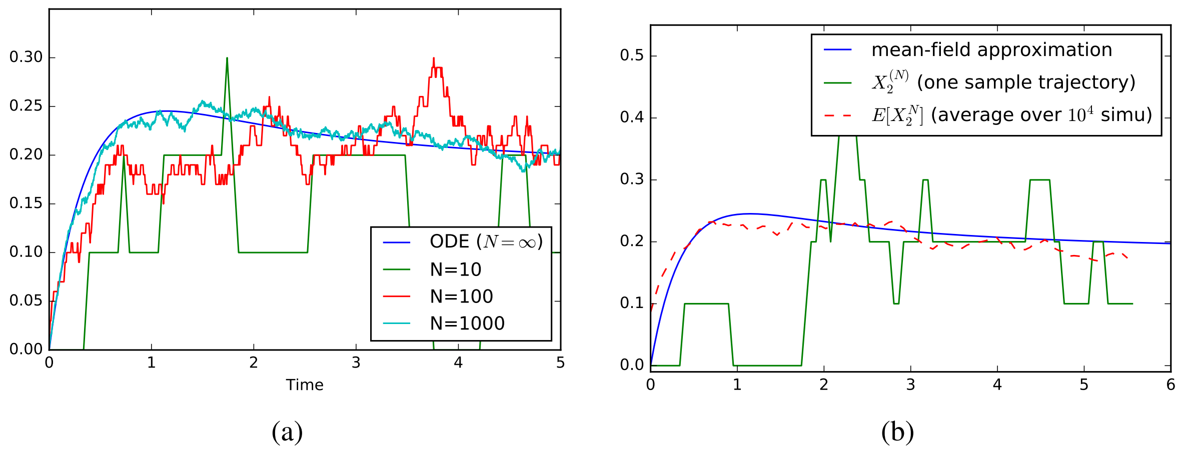

Figure 1: Two-choice model: Proportion of servers with 2 jobs as a function of time. (a) On the left panel we plot the result of one simulation for $N = 10$, $N = 100$ and $N = 1000$. (b) On the right panel, we plot one simulation for $N = 10$ as well as an average over 10^4 simulation.

2.2 The Classical Approach is $O(1/N)$-Accurate

The rate of convergence of $X_i^{(N)}$ to x_i has been studied in several papers, *e.g.* [22, 30, 40], that show that the expected distance between $X^{(N)}$ and x is of order $1/\sqrt{N}$:

$$\mathbb{E}\left[\left\|X^{(N)}-x\right\|\right] \approx \frac{C}{\sqrt{N}}. \tag{1}$$

This result is a central-limit-theorem for mean-field systems: $X^{(N)}(t)$ is equal to $x(t)$ plus $1/\sqrt{N}$ times a diffusion noise [30]. The Figure 1(a) illustrates that for $N = 100$ or $N = 1000$, the noise around the mean-field limit is similar to a Brownian motion that decreases at rate $O(1/\sqrt{N})$.

Yet, Equation (1) does not fully explain the accuracy of mean-field approximation. In [11], we show that when one wants to estimate an expected measure of $X^{(N)}$, the accuracy of mean-field approximation is $O(1/N)$ and not $O(1/\sqrt{N})$. More precisely, we show that if h is a twice-differentiable function, then the expectation of $h(X^{(N)})$ converges to

$$\mathbb{E}\left[h(X^{(N)})\right] = h(x) + O\left(\frac{1}{N}\right) \tag{2}$$

This result shows that an average value estimated via mean-field approximation is $1/N$-accurate. In [11], we show that this result holds for the classical model of [30], provided that the drift is twice-differentiable. It holds for the transient regime and can be extended to the stationary regime under the same conditions as [40].

This result has many applications. Here, we cite only two:

- For model-checking applications, it is often useful to compute the probability of having reached a given state s at a time t. This probability is exactly $\mathbb{E}\left[X_s^{(N)}\right]$ and our result shows that this probability is much more accurate than the $O(1/\sqrt{N})$ suggested. As an example, we show in Figure 1(b) the probability that a server has 2 servers at time t. In red we plot an estimate of this probability computed over 10^4 simulations. We observe that this probability is extremely close to its mean-field approximation. In terms of computation, the computation of the mean-field approximation takes less than $0.01s$ whereas the simulation time is several minutes.

- Another example is when one wants to compute expected performance metrics. For example, in a queuing network such as the two-choice model, the average queue length can be expressed as $\mathbb{E}[h(X^{(N)})]$. Equation (2) shows that the average queue length converges at rate $O(1/N)$ to its mean-field approximation. This is what is observed in Table 1, where we have $\mathbb{E}[h(X^{(N)})] \approx h(x) + 4/N$.

N	10	100	1000	$+\infty$
Average queue length (m^N)	2.81	2.39	2.36	2.35
Error ($m^N - m^\infty$)	0.45	0.039	0.004	0

Table 1: Average queue length in the two-choice model. The values for a finite number of servers N are obtained by simulation. The value for $N = +\infty$ is the mean-field approximation.

2.3 Corrected Approximation and Improved Rate of Convergence

Using the idea above, we show in [9] that under the same conditions, for any function h, there exists a constant C such that

$$\mathbb{E}\left[h(X^{(N)})\right] = h(x^*) + \frac{C}{N} + O\left(\frac{1}{N^2}\right).$$

The constant C can be expressed as a function of the Jacobian and the Hessian matrices of the drift. It can be computed by a numerical algorithm (that involves few matrix multiplications and inverse computation).

This result has two main implications:

- First, it guarantees that the error of a mean-field model is always in the same direction: if a quantity is over-estimated by mean-field approximation for $N = 100$, one can expect that mean-field will still over-estimate this quantity for $N = 1000$.
- It allows one to define a *refined approximation*, given by $h(x^*) + \frac{C}{N}$. For example, if we apply this to the two-choice model, the average queue length given by our new approximation is

	$N = 10$	$N = 20$	$N = 30$	$N = 50$
Average queue length (simulation)	2.8040	2.5665	2.4907	2.4344
Refined approximation	2.7513	2.5520	2.4855	2.4324
Classical mean-field approximation (independent of N)	2.3527	2.3527	2.3527	2.3527

For smaller values of N, it is much closer than the classical mean-field approximation, that is equal to 2.3527.

3 Uncertain and/or Heterogeneous Systems

3.1 Numerical Methods for Quantifying Uncertainties

In [3] (partially reported in Deliverable 1.2) we consider stochastic models in the presence of uncertainty, originating from lack of knowledge of parameters or by unpredictable effects of the environment. We set up a formal framework to study uncertain mean-field interaction models. We prove that the limiting behaviour of these systems as the population size goes to infinity can be described by a differential inclusion that can be constructed from the (imprecise) drift. In our more recent work [8], we focus on numerical methods.

Computing the reachable set of a differential inclusion is particularly challenging in the case of non-linear systems. As a result, several over-approximation [17, 23, 28, 31, 33, 35, 37] as well as under-approximation [24, 39] techniques have been developed, aiming to provide super- and subsets of the reachable set, respectively. Over-approximations can be used to verify whether an ODE system under study satisfies a certain property while under-approximations are useful to detect violations.

In [8], we present a tool – UTOPIC – that provides an under-approximation of the reachable-sets of differential inclusions – *i.e.* a subset of the reachable set – which can then be used to formally falsify properties of quantitative models. Our tool uses optimal control theory and Pontryagin's maximum principle and works for non-linear systems. We compare our tool with state-of-the-art tools Flow* and CORA on a few benchmarks. These benchmarks are classical benchmarks used to compare over-approximation tools (such as Brusselator, JetEngine, Helicopter, etc.) as well as the binding models used in chemical reaction networks. We show that UTOPIC provides tight under-approximations for all of these benchmarks. We also show that classical tools cannot scale to more than 20-30 dimensions while our tool scales to models with more than 100 dimensions. The tool is described in more detail in Deliverable 5.3.

3.2 Bayesian Learning

The approach described in the previous section was to consider that parameter values can vary in time in an uncontrolled manner. In some systems, the parameters of the models exists but are unknown to the modeler. A interesting approach in this case is to learn the parameters from observations while the system runs. In [14], we develop methods for learning the underlying parameters of stochastic population models from observed data. The methods follow the Bayesian paradigm of encoding uncertainty as probability distributions and can

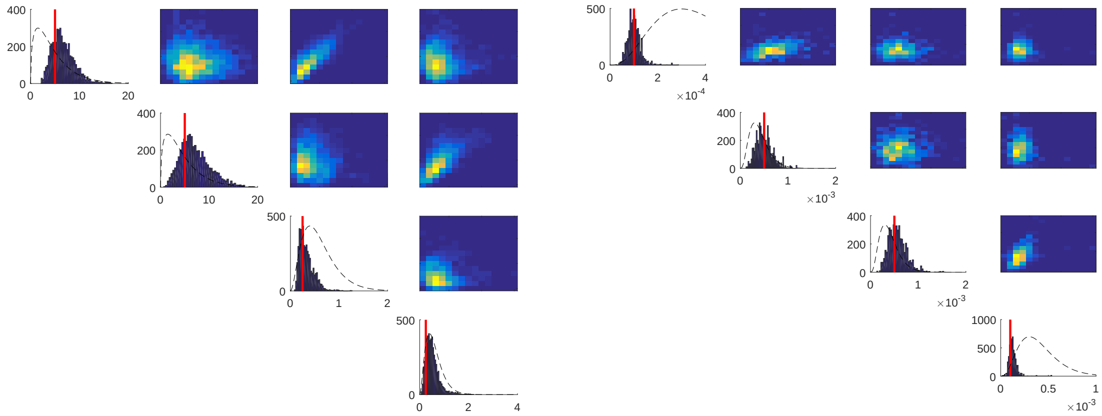

Figure 2: An example of the results generated by the estimation methods of [14]: histograms representing the posterior distribution for each parameter, and heat maps showing pairwise correlations.

take into account prior beliefs about the values of the parameters. By combining this prior knowledge with the measured behaviour, we obtain a new distribution, representing our updated belief about the parameter values. We use a Monte Carlo scheme: this posterior distribution is not obtained in an analytical form, but rather as a set of representative samples, which can be used to empirically reconstruct it or to estimate various quantities of interest (Figure 2).

A major novelty of the methodology presented in the paper is that it can be applied to systems with infinite state-spaces. Instead of employing a continuous approximation, we retain the discrete stochastic dynamics and use a random strategy to obtain a finite representation. The way in which this state-space truncation is performed ensures that the results obtained are statistically correct.

We have applied these methods to models of disease spread and predator-prey interactions, as well as a landmark synthetic genetic network. Analysis of our results suggests that our approach offers significant advantages in performance and accuracy over state-of-the-art methods in the machine learning and statistics literature.

3.3 Policy Learning for Continuous-Time Markov Decision Processes

Continuous-time Markov decision processes (CTMDPs) are a very powerful mathematical framework to solve control and dependability problems in real-time systems featuring both probabilistic and nondeterministic behaviours [18].

In [1], we focus on the core numerical tool for verification of temporal properties, namely time-bounded reachability – i.e. the probability of reaching a given set of states within a prescribed time bound. In our context, the answer to a reachability computation is not a probability value, but an interval of probabilities, whose extremes are the minimum and maximum achievable ones. An orthogonal problem to reachability computation is planning, meaning identifying policies maximising or minimising the satisfaction probability. We present a novel statistical approach to compute lower bounds on the maximum reachability probability of a CTMDP. Our method uses a basis-function regression approach to compactly encode schedulers and effectively search for an optimal one. We consider randomised time-dependent early schedulers, and focus on population models, which despite being so common, suffer severely from state space explosion, with the number of states growing exponentially with the number of variables. This reflects on the size of the schedulers: storing a function of time for each state of the CTMDP is infeasible.

The paper [1] contains two main novel insights. First, for population models we use the structure of the state space, embedding it as a discrete grid in real space and thus obtaining a continuous relaxation of the problem (similar to a mean-field approximation). This allows us to consider schedulers defined on such a continuous space, and to treat time and space uniformly, representing schedulers as continuous functions. This opens up the use of machine learning methods to represent continuous functions as combinations of basis functions, and allows us to define the optimisation problem as a search in such a continuous function space. The second main idea in [1] is to set up an efficient stochastic gradient descent search algorithm, which considerably speeds up the search in the space of functions. This is based on a novel algorithm using Gaussian Processes (GPs) and statistical model checking to sample in an unbiased manner the gradient of the functional associating a reachability probability with a randomised scheduler. This method allows us to effectively learn schedulers that minimise (locally) the reachability probability. A drawback of the approach is that it does not provide formal guarantees on the optimality, but it is feasible for models which are out of reach for classic numerical methods for CTMDPs. We tested this algorithm on an epidemic model in which there was a choice between administering or not a treatment with low probability of being lethal. We show that our algorithm is able to quickly retrieve effective policies.

3.4 Heterogeneity as a Form of Uncertainty

Some systems are composed of many objects that are similar in nature but still have different parameters. This occurs for example when one wants to model the cache of a video-streaming system: each video is an object that can be stored in the cache but each video has its own size and its own popularity. In this context, when one wants to consider the model of such a system, a mean-field approximation with uniform objects is not appropriate and one has to deal with heterogeneous objects. The definition and the numerical resolution of a mean-field system with a heterogeneous population is not obvious. A possible approach to this problem is to consider a worst-case approach. This leads to the differential hull approximation of [37], and that was reported in Deliverable 1.2. In [10, 12, 13], we work directly with the heterogeneous model and introduce a mean-field approximation of a family of cache replacement algorithms.

Computer system and network performance can be significantly improved by caching frequently used information. When the cache size is limited, the cache replacement algorithm has an important impact on the effectiveness of caching. Our approximation allows us to compute the cache hit probability of two recently introduced classes of cache replacement algorithms : LRU(m) and h-LRU [12] – both are variants of the “LRU” policy (least-recently used). The originality and the technical difficulty of our approach is that the studied models are composed of heterogeneous agents. These lead us to develop specific approximation techniques that are valid for these types of agents.

The approximation that we use is based on an equivalence between time-to-live (TTL) approximations and cache replacement polices. We provide both numerical and theoretical support for the claim that the proposed TTL approximations are asymptotically exact. In particular, we show that the transient hit probability converges to the solution of a set of ordinary differential equations (ODEs), where the fixed point of the set of ODEs corresponds to the TTL approximation. We further show that LRU(m) and h-LRU perform alike, while LRU(m) requires less work when a hit/miss occurs. This paves the way to design more efficient cache-replacement algorithms that are sensitive to the correlation between consecutive inter-request times.

4 Centralised and Decentralised Control of Large Systems

4.1 Decentralised Control and Mean-Field Games

The notion of mean-field games has been introduced by Lasry and Lions in [32] and has a growing popularity. Mean-field games model the rational behaviour of an infinite number of indistinguishable players in interaction. An important assumption of mean-field games is that, as the number of players is infinite, the decisions of an individual player do not affect the dynamics of the mass. Each player plays against the mass. A mean-field equilibrium corresponds to the case when the selfishly optimal decisions of a player coincide with the decisions of the mass. Many authors argue that mean-field games are a good approximation of symmetric stochastic games with a large number of players, the rationale behind this being that the impact of one player becomes negligible when the number of players goes to infinity.

In [5, 6], we question this assertion. We consider mean-field games with discrete state spaces (aka discrete mean-field games) and we analyse these games in continuous and discrete time, over finite as well as infinite time horizons. We provide the minimal assumptions under which the existence of a mean-field equilibrium can be guaranteed, namely continuity of the cost and of the drift, which mimics nicely the case of classical games. Besides, we also study the convergence of the equilibria of N-player games to mean-field equilibria. We show that, in general, this convergence does not hold. The "tit for tat" principle allows one to define many equilibria in repeated or stochastic games with N players. However, in mean-field games, the deviation of a single player is not visible to the population. The conclusion is that, even if N-player games have many equilibria with a good social cost, this may not be the case for the limit game.

In [7], we consider an instance of this general model. We study the dynamics of a system composed of an infinite number of homogeneous strategic players evolving in continuous time over a finite state space. Each player is driven by a simple SIR dynamics (Susceptible, Infected, Recovered) with an additional state (Vaccination), and each player can control its vaccination strategy. We describe a mean-field game model of the behaviour of the population and we show that for this type of game, the Nash equilibrium (that is guaranteed to exist here) can be seen as a fixed point of a Bellman optimality equation. Here, the population dynamics is simple enough that the structure of the unique symmetric Nash equilibrium can be described: The strategy of each player is to vaccinate at maximal rate until the proportion of susceptible players in the system becomes smaller than a certain threshold. At this point, the player does not vaccinate any longer. We also consider the associated centralised control problem where the goal is to find the vaccination strategy for the population that minimises the total cost of the system. We show that the solution of this problem is also a threshold strategy. We compare both strategies and we show that, to encourage selfish individuals to vaccinate optimally (i.e. the mean-field equilibrium reaches a social optimal), vaccination for individuals should be subsidised.

4.2 Symbolic Performance Adaptation via Mean-Field Representation

In this section we report on an approach to achieve centralised control of software performance based on a mean-field representation [15]. It considers queuing networks (QN) as performance models because they have been recognised as a well-assessed formalism in the area of software performance engineering [26, 20, 19, 36]. The popularity of QN is also due to the immediate mapping between its constituent elements and features of software systems: (i) a queueing *service center* may either model a software component or a hardware resource; (ii) each service centre can be composed of multiple *servers*, modelling software concurrent threads or hardware cores; (iii) *routing probabilities* between service stations model the operational profile. Previous work (performed within this project and reported in Deliverable 3.2) presented closed-form symbolic expressions for the estimation of typical QN steady-state performance indices (e.g., utilisation and response times) using a mean-field model [29]. This was achieved by allowing symbolic specifications for the number of servers, their service times, and the routing probabilities. The actual performance of a specific QN can then be obtained by simply plugging its concrete parameters into the symbolic expression. Notably, topological variations to a model (such as the removal/addition of a service centre) can be accommodated by allowing certain parameters to be zero.

The previous work [29] was mainly concerned with analysing many variants of a performance model effectively, but did not consider the problem of exploring these variants for control purposes, i.e., in order to find a configuration that satisfies certain QoS requirements. A naive strategy for doing this may be infeasible when the parameter space is very large, or even impossible if it is infinite — for instance when allowing rates and probabilities to vary within a range. The approach of [15], called *Symbolic Performance Adaptation* (SPA), tackles this issue by leveraging the symbolic QN solution as the underlying engine for control. Technically, the main novel contribution of [15] is to encode the QoS satisfaction problem itself symbolically as a satisfiability modulo theories (SMT) problem [21]. This enables the formulation of statements such as the existence of an assignment for the symbolic variables that satisfy a quality of service (QoS) constraint, specified in first-order logic. The peculiarity of this formulation is that, if the SMT problem is satisfiable, i.e., an assignment exists, then an SMT solver is able to produce a *witness*, i.e., one possible assignment. Instead, in the case the problem is not satisfiable then the modeller has the definite proof that no assignment can ensure the given QoS. This mandates a revision of the system and/or a relaxation of the QoS requirements. In other words, in doing so we exploit the symbolic variables as the run-time "knobs" for guiding the QoS-based control process.

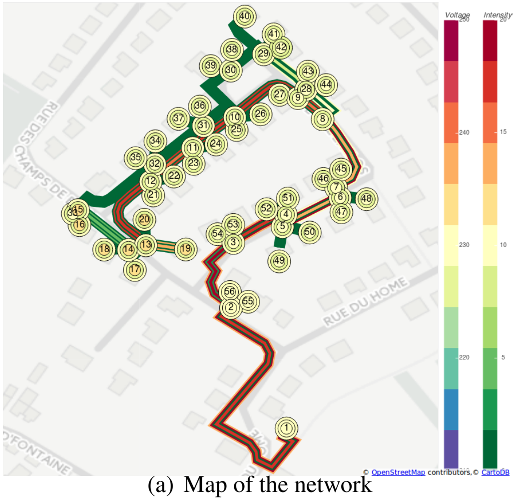

(a) Map of the network

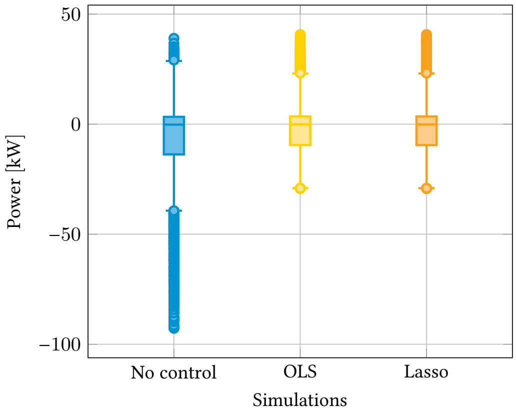

(b) Power flow at the transformer : the transformer should operate between -30 and +40 kW. Without control, this limit is often violated. This is no more the case with our control mechanism (OSL and Lasso correspond to two variants of our algorithm).

Figure 3: Map of the electrical distribution network and power flow at the transformer in the model of [16].

In practice, SPA enables analysis at two stages of software design: (i) at design-time one can provide assurances that the system can meet a desired QoS; (ii) at runtime, after detecting changing conditions, these can be encoded as further constraints on the model. These, in turn, trigger a new evaluation of the satisfiability problem, leading to a possible new configuration that meets the newly imposed constraints. Thus, it becomes crucial to evaluate the effectiveness and scalability of SPA in re-configuring the system at runtime. On a realistic QN model of a three-tiered distributed web system with service replications, SPA was reported to find re-configurations that satisfy the desired QoS in the face of random changes of workloads and injection of faults within a few hundred milliseconds.

5 Control of Energy Systems

The deployment of new renewable production capacities are creating two new problems that electrical network operators have to solve. First, renewable generation is hard to predict and is very volatile. Second, these new generators are distributed : most of the photo-voltaic panels are connected to distribution networks, that were not dimensioned for this and are less robust than transmission networks. This section describes two contributions that focus on how to control a smart electrical system in which renewable energy accounts for a large part of the production. The first one describes an algorithm to control in real time a distribution network. The second is to use an energy market to encourage people to smoothe their consumption.

5.1 Power-Based Control

In [16], we consider the problem of controlling an electrical distribution network, such as the one shown in Figure 3(a). This figure represents a small neighbourhood in which a transformer (at the bottom of the figure) feeds 44 private users that are both consuming and producing energy. A smart grid can be broadly defined as an electricity distribution system equipped with flexible generators and loads that are controlled in such a way that no congestion occurs in the power grid. To achieve this, a mathematical model that captures the physics of the electricity grid must be available. It is generally assumed in the literature that such a model can be derived from physical parameters such as cable type and length, resulting in the notorious “load flow problem”. In practice, however, this data is usually not fully or not accurately available to the distribution system operator. In [16], we show that it is possible to design an efficient control algorithm without knowing a physical model of the network. Our algorithm uses observation data to learn the sensitivity of the network to the consumption fluctuations. We show that standard learning techniques, such as Lasso, lead to models that are both simple and accurate. We develop a simulator of a real distribution network on which we study the optimal curtailment of electricity produced by photo-voltaic panels. By using our simulator, we show that our learning algorithm suppresses all over-voltage problems in the network and all constraints at the transformer. An example is shown in Figure 3(b) : in theory, the transformer should operate between -30kW and +40kW. Without control (blue curve) we see that the constraints are often violated while our learning algorithm (Yellow and Orange curve) suppresses all problems.

5.2 Market-Based Control

The second problem that we have studied is the impact of the volatility of renewable energy. The increased penetration of renewable energy sources in existing power systems has led to necessary developments in electricity market mechanisms. Renewable energy generation is recognised as being responsible for deviations between scheduled and actual energy generation. As a result, these are penalised by current electricity markets. However, there is currently no mechanism to enforce accountability for the additional costs induced by power fluctuations. These costs are socialised and eventually supported by electricity customers. In [2], we propose some metrics for assessing the contribution of all market participants to power regulation needs, as well as an attribution mechanism for fairly redistributing related power regulation costs. We discuss the effect of various metrics used by the attribution mechanisms, and we illustrate, in a game-theoretical framework, their consequences on the strategic behaviour of market participants. We also illustrate, by using the case of Western Denmark, how these mechanisms may affect revenues and the various market participants. These results suggest that simple incentives encourage people to adapt more efficient behaviour.

6 Conclusion

Summary – Before the start of the QUANTICOL project, the state-of-the-art in mean-field approximation was that it is a technique that can be applied to models composed of a large and homogeneous population. During this project, we extended the theory in multiple directions: multi-scale systems (Deliverable 1.1), hybrid limits (Deliverable 1.2) and distributed control (Deliverable 1.3).

In this deliverable, we report on the most recent progress that has been made during the last reporting period. One of the main achievements of this workpackage during this period was to show that it is possible to define an approximation that is more precise than mean-field. This will make mean-field applicable to systems of small or moderate population size (composed of a few tens of individuals for example) for which the mean-field approximation was not accurate enough before. This result opens many interesting questions that will continue after this project. Another important achievement is that we developed tools that use theoretical results obtained previously in the project, for instance to compute bounds on uncertain systems. This makes the results available to users who are not specialists in the theory behind them. Last, we applied the ideas learned during this project to build concrete control mechanisms. The work about the control of distribution networks, that started thanks to the QUANTICOL project, has now become an active collaboration between the company Schneider Electric and Inria.

Deviations from plan – Most of the objectives that were identified at the beginning of the project have been achieved. In some cases, the results even exceeded the expectation. One example is the “refined approximation” presented in Section 2: In the description of work, we planned to “extend the theory to second-order mean-field approximations” but we did not expect the results to be so powerful whilst still easy to use.

During the course of this project, we also produced strong contributions to the case studies – *e.g.*, [16, 27] – although we must acknowledge that compared with what we envisioned four years ago, our work has more emphasis on the theoretical work than on the case studies. We explain this by two reasons. The first is that EPFL, who was the leader on the “smart-grid” case study withdrew from the project after the movement of Nicolas Gast from EPFL to Inria. The second is that some theoretical results exceeded our expectations which led us to dedicate more of our time to these.

Future work and collaborations. As written in the Description of Work, "this topic will require further substantial investigations beyond the end of the project". In what follows, we list a few questions that are left open on which we are currently working or on which we plan to work.

Our results about a refined mean-field approximation suggest that this approximation has good potential. Our first step in this direction, in collaboration with WP4 and WP5, is to develop a tool that automates its numerical computation. We expect two impacts of this: to increase the number potential users and to create a list of comprehensive benchmarks that assess the accuracy of the approximation in practical cases. We are also working on more theoretical questions. One of them is to apply this refinement in the context of heterogeneous systems. One application, linked with WP2, is to study spatial dynamics. Our approach could be more precise than moment closure approaches in this context. We are also trying to generalise our results to fluid model checking and PCTL formulas for mean-field models, such as developed within WP3.

Another direction of research that we are currently pursuing is to develop control mechanisms for electrical distribution networks. The methods that we developed in this project allowed us to define learning approximation techniques that we use to develop robust and automatic control mechanisms. We will use the work reported [16] as a first step towards a fully automated control architecture in a distribution network. This raises theoretical questions such as the optimal control of loads and the fairness between users but also very practical ones of software engineering and robust implementation. We are currently developing a prototype implementation of these ideas in collaboration with the company Schneider Electric.

References (Publications of the Quanticol project during the reporting period)

Note: the technical reports mentioned in this part have been submitted and are currently under review.

[1] E. Bartocci, L. Bortolussi, T. Brázdil, D. Milios, and G. Sanguinetti. "Policy Learning for Time-Bounded Reachability in Continuous-Time Markov Decision Processes via Doubly-Stochastic Gradient Ascent". In: *Quantitative Evaluation of Systems - QEST 2016*. 2016, pp. 244–259. DOI: 10.1007/978-3-319-43425-4_17.

[2] F. Bona, N. Gast, J.-Y. Le Boudec, P. Pinson, and D.-C. Tomozei. "Attribution mechanisms for ancillary service costs induced by variability in power delivery". In: *IEEE Transactions on Power Systems* (2016), p. 10. DOI: 10.1109/TPWRS.2016.2598760. URL: https://hal.inria.fr/hal-01403913.

[3] L. Bortolussi and N. Gast. "Mean Field Approximation of Uncertain Stochastic Models". In: *46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN 2016)*. Toulouse, France, June 2016. URL: https://hal.inria.fr/hal-01302416.

[4] L. Bortolussi and N. Gast. "Mean-Field Limits Beyond Ordinary Differential Equations". In: *Formal Methods for the Quantitative Evaluation of Collective Adaptive Systems*. Ed. by M. Bernardo, R. De Nicola, and J. Hillston. Vol. Programming and Software Engineering. 16th International School on Formal Methods for the Design of Computer, Communication, and Software Systems, SFM 2016, Bertinoro, Italy, June 20-24, 2016, Advanced Lectures. June 2016. DOI: 10.1007/978-3-319-34096-8_3. URL: https://hal.inria.fr/hal-01334358.

[5] J. Doncel, N. Gast, and B. Gaujal. "Are mean-field games the limits of finite stochastic games?" In: *The 18th Workshop on MAthematical performance Modeling and Analysis*. Nice, France, June 2016. URL: https://hal.inria.fr/hal-01321020.

[6] J. Doncel, N. Gast, and B. Gaujal. *Discrete Mean-Field Games: Existence of Equilibria and Convergence*. Tech. rep. 2017.

[7] J. Doncel, N. Gast, and B. Gaujal. *Optimal v.s. selfish vaccination: a mean-field game model*. Tech. rep. 2017.

[8] J. Doncel, N. Gast, M. Tribastone, M. Tschaikowski, and A. Vandin. *UTOPIC: Under-Approximation Through Optimal Control*. Tech. rep. 2017.

[9] N. Gast. *Can we beat mean-field approximation?* Tech. rep. 2017.

[10] N. Gast. "Construction of Lyapunov functions via relative entropy with application to caching". In: *The 18th Workshop on MAthematical performance Modeling and Analysis*. Nice, France, June 2016. URL: `https://hal.inria.fr/hal-01321017`.

[11] N. Gast. "Expected values estimated via mean-field approximation are 1/N-accurate". In: *Proceedings of ACM Sigmetrics 2017 Conference*. 2017.

[12] N. Gast and B. Van Houdt. "Asymptotically Exact TTL-Approximations of the Cache Replacement Algorithms LRU(m) and h-LRU". In: *28th International Teletraffic Congress (ITC 28)*. Würzburg, Germany, Sept. 2016. URL: `https://hal.inria.fr/hal-01292269`.

[13] N. Gast and B. Van Houdt. "Transient and Steady-state Regime of a Family of List-based Cache Replacement Algorithms". In: *Queueing Systems*. Special Issue: SIGMETRICS 2015 (June 2016). This paper is an extended version of the ACM SIGMETRICS 2015 paper that is accessible at https://hal.inria.fr/hal-01143838. DOI: `10.1007/s11134-016-9487-9`. URL: `https://hal.inria.fr/hal-01334354`.

[14] A. Georgoulas, J. Hillston, and G. Sanguinetti. "Unbiased Bayesian inference for population Markov jump processes via random truncations". In: *Statistics and Computing* (2016), pp. 1–12. ISSN: 1573-1375. DOI: `10.1007/s11222-016-9667-9`. URL: `http://dx.doi.org/10.1007/s11222-016-9667-9`.

[15] E. Incerto, M. Tribastone, and C. Trubiani. "Symbolic Performance Adaptation". In: *Proceedings of the 11th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS'16)*. ACM, 2016, pp. 140–150. DOI: `10.1145/2897053.2897060`.

[16] B. Vinot, F. Cadoux, N. Gast, and R. Heliot. *Controlling a Low-Voltage Distribution Network by Learning its Sensitivity Matrix*. Tech. rep. 2017.

References

[17] E. Asarin, T. Dang, and A. Girard. "Reachability Analysis of Nonlinear Systems Using Conservative Approximation". In: *HSCC*. 2003.

[18] C. Baier, H. Hermanns, J.-P. Katoen, and B. R. Haverkort. "Efficient computation of time-bounded reachability probabilities in uniform continuous-time Markov decision processes". In: *Theor. Comput. Sci.* 345.1 (2005), pp. 2–26.

[19] S. Balsamo and A. Marin. "Queueing Networks". In: *International School on Formal Methods for Performance Evaluation SFM*. 2007, pp. 34–82.

[20] S. Balsamo and M. Marzolla. "Performance evaluation of UML software architectures with multiclass Queueing Network models". In: *International Workshop on Software and Performance*. 2005, pp. 37–42.

[21] C. Barrett, R. Sebastiani, S. A. Seshia, S. A. S. Cesare Tinelli Clark Barrett Roberto Sebastiani, and C. Tinelli. "Handbook of Satisfiability: Volume 185 Frontiers in Artificial Intelligence and Applications". In: ed. by A. Biere, A. Biere, M. Heule, H. van Maaren, and T. Walsh. 2009. Chap. Satisfiability Modulo Theories.

[22] 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.

[23] S. Bogomolov, C. Mitrohin, and A. Podelski. "Composing Reachability Analyses of Hybrid Systems for Safety and Stability". In: *ATVA*. 2010. DOI: `10.1007/978-3-642-15643-4_7`.

[24] X. Chen, S. Sankaranarayanan, and E. Ábrahám. "Under-approximate flowpipes for non-linear continuous systems". In: *FMCAD*. 2014, pp. 59–66. DOI: `10.1109/FMCAD.2014.6987596`.

[25] J.-w. Cho, J.-Y. Le Boudec, and Y. Jiang. "On the validity of the fixed point equation and decoupling assumption for analyzing the 802.11 mac protocol". In: *ACM SIGMETRICS Performance Evaluation Review* 38.2 (2010), pp. 36–38.

[26] V. Cortellessa and R. Mirandola. "Deriving a queueing network based performance model from UML diagrams". In: *International Workshop on Software and Performance*. 2000, pp. 58–70.

[27] 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.

[28] A. Girard. "Reachability of Uncertain Linear Systems Using Zonotopes". In: *HSCC*. Springer, 2005.

[29] M. Kowal, M. Tschaikowski, M. Tribastone, and I. Schaefer. "Scaling size and parameter spaces in variability-aware software performance models". In: *International Conference on Automated Software Engineering*. 2015. URL: https://dl.dropboxusercontent.com/u/13100903/papers/ase2015.pdf.

[30] T. G. Kurtz. "Solutions of Ordinary Differential Equations as Limits of Pure Jump Markov Processes". In: *Journal of Applied Probability* 7 (1970), pp. 49–58.

[31] R. Lal and P. Prabhakar. "Bounded error flowpipe computation of parameterized linear systems". In: *EMSOFT*. 2015, pp. 237–246. DOI: 10.1109/EMSOFT.2015.7318279.

[32] J.-M. Lasry and P.-L. Lions. "Mean field games". In: *Japanese Journal of Mathematics* 2.1 (2007), pp. 229–260.

[33] R. Majumdar and V. S. Prabhu. "Computing Distances between Reach Flowpipes". In: *HSCC*. 2016, pp. 267–276. DOI: 10.1145/2883817.2883850.

[34] M. D. Mitzenmacher. "The Power of Two Random Choices in Randomized Load Balancing". PhD thesis. PhD thesis, Graduate Division of the University of California at Berkley, 1996.

[35] J. K. Scott and P. I. Barton. "Bounds on the reachable sets of nonlinear control systems". In: *Automatica* 49.1 (2013), pp. 93 –100.

[36] C. U. Smith, C. M. Lladó, and R. Puigjaner. "Performance Model Interchange Format (PMIF 2): A comprehensive approach to Queueing Network Model interoperability". In: *Perform. Eval.* 67.7 (2010), pp. 548–568.

[37] M. Tschaikowski and M. Tribastone. "Approximate Reduction of Heterogenous Nonlinear Models With Differential Hulls". In: *IEEE Transactions Automatic Control* 61.4 (2016), pp. 1099–1104. DOI: 10.1109/TAC.2015.2457172.

[38] N. D. Vvedenskaya, R. L. Dobrushin, and F. I. Karpelevich. "Queueing system with selection of the shortest of two queues: An asymptotic approach". In: *Problemy Peredachi Informatsii* 32.1 (1996), pp. 20–34.

[39] B. Xue, Z. She, and A. Easwaran. "Under-Approximating Backward Reachable Sets by Polytopes". In: *CAV*. Ed. by S. Chaudhuri and A. Farzan. Springer International Publishing, 2016, pp. 457–476. DOI: 10.1007/978-3-319-41528-4_25.

[40] L. Ying. "On the Approximation Error of Mean-Field Models". In: *Proceedings of the 2016 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Science*. ACM. 2016, pp. 285–297.