QUANTICOL

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

quanticol
http://www.quanticol.eu

D1.3

Distributed control for Collective Adaptive Systems

Revision: 1.0; September 29, 2015

Author(s): Nicolas Gast (Ed.) (INRIA) Josu Doncel (INRIA), Farshad Shams (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 work on Task 1.2. Task 1.2 started in month 13 of the project and will run for three years. Its goal is to develop centralised and distributed control algorithms that lead to an efficient use of resources. In this deliverable, we present generic methods that are widely applicable for a large variety of CAS. We also present some specific applications to smart-grid and electricity markets.

- We first present adaptivity in terms of solving an optimisation problem by using a decentralised algorithm. We develop a general framework based on the use of mean-field game theory and Lagrangian decomposition. We show how, given a centralised optimisation problem, we can define a mean-field game whose equilibrium is an optimal allocation. We develop a distributed algorithm that is proven to converge to this equilibrium and therefore to a globally optimal allocation of resources. This is discussed in Section 2 and collects the results of [GLBT14; DGG15; CGM14; GM15].

- The previous approach is top-down: starting from an optimisation problem, we derive a distributed algorithm that solves it. This approach is well-suited to systems where a single organisational entity controls a whole system (for instance in distribution networks) and can modify the behaviour of each agent. In Section 3, we adopt a bottom-up approach and study how adaptivity can be obtained by using market mechanisms. We study two directions of how to modify these markets to incorporate more renewable energy: to incentivize the creation of coalitions of micro-producers that exchange energy locally [ST15] and to make people accountable for short-term variability [Bon+15].

- Finally, in Section 4, we present two new methods based on moment closure to study how reasonable control policies can improve or deteriorate the performance of a CAS. We obtain approximate models for two specific features: having many heterogeneous objects [GVH15] and a space-constrained load balancing policy [Gas15]. These works show that some policies that were conjectured optimal in practice deteriorate the overall performance.

Contents

1 Introduction

Collective adaptive systems (CASs) are composed of many autonomous entities that interact and cooperate to reach certain objectives. Typical objectives are the stability of the system, to optimise the efficiency while performing a certain task or to provide a service to users. In this report, we focus on the design and the evaluation of distributed algorithms that allow CAS to adapt to their external environment. We formulate the capacity of a system to adapt in terms of a stochastic optimisation problem. In this context, an *adaptive* system is a system that has the ability to change its dynamics in order to maximise a given objective function, by learning its impact on the environment. We focus on the following aspects:

- Given a centralised objective, how to automatically build and deploy local control algorithms to reach this objective. Our idea is that this can be achieved by sending coordination signals to small task-oriented entities. This idea has been used to schedule the consumption of some electric appliances, such as fridges or electrical vehicles.
- It is often hard for local entities to measure their influence on the global system. In order to take good decisions, local entities have to *learn* their environment in order to decide which decision to take. This leads to two problems: how to learn when measurements are noisy, and how to use this knowledge to take a decision.
- Real-time prices and market mechanisms are one way to design decentralised control algorithms. In such systems, the agents fix the prices at which they exchange goods and each agent tries to maximise a local utility that depends on the price at which it exchanges goods with the others. These mechanisms are inherently distributed and often simple to implement but they do not always lead to an efficient allocation of resources. Characterising their emergent behaviour allows one to design efficient market mechanisms.

In this deliverable, we model all of these problems as stochastic optimisation problems. The system is composed of a large number of possibly heterogeneous entities that can take decisions locally. Our goal is to develop distributed mechanisms such that the individual actions of each agent maximises a common objective. We achieve this by employing tools from stochastic approximation to develop mean-field methods applicable to stochastic optimal control.

Mean-field methods in performance evaluation have been originally developed to study adaptive strategies of large systems, such as wireless access protocols or load balancing techniques, *e.g.*, [BB08; Bor+13]. These methods are mostly descriptive. Our first contribution, described in Section 2, is to show how to use mean-field game theory to solve centralised optimisation problems. This technique is generic and allows one to transform a complex stochastic optimisation problem into local problems by using a Lagrangian decomposition of the problem. The Lagrangian multipliers of these problems can be viewed as virtual prices on which agents agree. Following these ideas, we study electricity markets in Section 3 and propose new mechanisms in order to incorporate the multi-scale features of renewable energy – which are the fact that (1) they change at a fast-time scales and (2) produce locally. Finally, in Section 4, we present two examples of distributed control algorithms that are designed to improve the performance of a system and that exhibit properties for which we had to develop specific approximation techniques: the presence of many heterogeneous objects and the presence of spatially constrained interactions.

2 Distributed Optimisation: Mean-Field Games and Learning Tools

In this section, we are interested in the control and optimisation of distributed systems such as smart grids. We first introduce mean-field games as a general framework to describe systems composed of a large number of rational agents. We then present how Lagrangian relaxation can be employed to transform a centralised optimisation problem into a mean-field game and we present an algorithm to solve these problems. We present some application examples of the previous techniques. Finally, we describe online learning algorithms for stochastic optimisation problems.

2.1 Mean-Field Games: General Framework

We consider a population with N objects that evolve in continuous time. Each object is a player of the game. It has a finite state space and can take actions to influence the evolution of its state. We denote by $x_i(t)$ the state of object i at time t and $(x_i(t))_t$ is the trajectory of object i. We also say that $\mathbf{x}(\mathbf{t}) = (x_1(t), \ldots, x_N(t))$ is the state of the system at time t. To each object is associated an instantaneous cost function $C_i(x_i(t), \mathbf{x}_{-i}(t))$ that depends on its own state $x_i(t)$ and also on the state of the rest of the objects $\mathbf{x}_{-i}(t)$. The total cost for an object i is the integral of this instantaneous cost over time, i.e., $\int_t C_i(x_i(t), \mathbf{x}_{-i}(t))dt$. We assume that objects are selfish and rational: given the strategy of the others, an object chooses a strategy that minimises its cost.

In this context, a crucial notion is the Nash equilibrium. A set of strategies $\mathbf{x}$ is said to be a Nash equilibrium if no object benefits from unilateral deviation [Nas51]. The concept of Nash equilibrium is natural when considering multiple competing players. As such, there is a large literature concerning the existence, uniqueness, convergence and stability of Nash equilibria, see for example the books [Nis+07; Mye13; Wei97]. Nevertheless, it is shown in [DGP09] that computing a Nash equilibrium is PPAD-complete[1], which indicates that it does not seem to be tractable when the number of objects is large. This suggests that for models with a very large number of objects, such as CAS, the notion of Nash equilibrium might not be appropriate.

As an alternative, the notion of mean-field games has been introduced by Lasry and Lions in [LL06a; LL06b; LL07]. Mean-field games provide a mathematical model to study the behaviour of a very large number of rational agents interacting. Letting the number of agents go to infinity, we obtain a system with continuous objects, which yields a simplification of the games and the equilibria computations.

The key assumption under mean-field games is to assume that (1) objects are exchangeable — *i.e.*, the identity of an object is not important, only its state is – and (2) there are many small objects – *i.e.*, alone, the action of a single object does not influence the system dynamics. Under this assumption the cost of the object i is a function $\int_0^T C(x_i(t), M(t))\ dt$, where $M(t)$ is the empirical distribution of objects at time t. Assumptions (1) and (2) imply that the choices of an object i do not affect the distribution of the mass $M(t)$ (because the number of objects is large). The object i chooses its trajectory $(x_i(t))_t$ so as to minimise its cost. A mean-field equilibrium is thus defined as a trajectory of the distribution of objects $(M(t))_t$ such that $M(t)$ minimises $\int_t C(x(t), M(t))\ dt$ over all possible trajectories $x(t)$. Interestingly, we observe that the calculation of a mean-field equilibrium is reduced to a single fixed-point problem, which is much simpler to analyse than the N-object equilibrium.

Mean-field games are considered as a good approximation to analyse models where the number of agents is large. Indeed, Lasry and Lions in [LL06a; LL06b; LL07] show that Nash equilibria converge to a mean-field equilibrium. Furthermore, they also prove that there exists a unique mean-field equilibrium, if the function $\int_t C(x(t), M(t))\ dt$ is monotone in $M(t)$. In the following section, we present how to solve centralised stochastic optimisation problem by using mean-field games.

2.2 Solving Stochastic Optimisation Problems by using Lagrangian Relaxation

We consider a stochastic optimisation problem in a system of objects that evolve in continuous time. The instantaneous cost is a function of the empirical measure of the objects, $M(t)$, and the decision rule at time t, $d(t)$. We denote by $G(M(t), d(t))$ the instantaneous cost of the system. In a centralised version of this control problem, a controller (*e.g.*, the controller of the whole electricity network)

[1] PPAD stands for "polynomial parity arguments on directed graphs". It is a complexity class that is a subclass of NP and is believed to be strictly greater than P.

establishes the actions that all the objects must take in order to minimise the cost in the system during a given time interval. When the number of possible values for $M(t)$ is small, this centralised control problem can be solved by dynamic programming [Put14]. However, this problem becomes intractable as soon as the number of objects grows.

The difficulty of this problem comes from the fact that the state space grows exponentially with the number of objects: if there are N objects that each have S possible states, the number of states is S^N (which is roughly 10^{15} for $N = 30$ objects and $S = 3$ states). A way to decompose the problem consists of decoupling all objects and assigning prices for each possible action that an object might take. Each object can then selfishly solve a small stochastic optimal control problem of dimension S. If the prices are properly chosen, the selfish responses of the objects (*i.e.*, the mean-field equilibrium) lead to decisions that minimise the selfish utility function G.

The construction of this mean-field game can be automated. The main difficulty is in deciding the prices. This can be done by using a Lagrangian decomposition of the problem, by doubling the variable $M(t)$ by a "free" variable π and adding the constraints $\pi(t) = M(t)$. Here, $\pi(t)$ represents the local state of an object and $M(t)$ is a variable whose only condition is to be equal to $\pi(t)$. Given the multipliers λ, the Lagrangian of the system $L_0(\pi(t), M(t), \lambda(t))$ is easier to solve than the original problem for the following reasons. First, the minimisation over $M(t)$ can be solved separately for each t (there are no constraints that couple $M(t)$ and $M(t + dt)$). Second, the minimisation on π corresponds to solving a Markov decision process of small dimension S. We know that if the multipliers λ are such that the selfish response π is equal to M, then they form a solution of the centralised control problem. Computing these multipliers can be solved by an iterative mechanism, for example by using the alternating direction method of multipliers (ADMM) [Boy+11]. This method works iteratively and is shown to converge whenever the function G is convex, the stochastic dynamics define a convex set and the (unaugmented) Lagrangian has a saddle point. This methodology is detailed in [GLBT14, Section 5] on an example taken from an electrical network.

2.3 Some Applications

In the previous sections, we have presented mean-field games as a general framework to analyse decentralised control of systems with a large number of objects. We have also explained how a Lagrangian relaxation can be used to solve a centralised control problem in a decentralised manner. We now expose two application examples.

2.3.1 Impact of Demand-Response on the Efficiency and Prices in Real-Time Electricity Markets

A first example is given in [GLBT14], where we study decentralised control of a population of flexible appliances in an electricity market. The system is composed of many objects (such as fridges or boilers) whose consumption can be anticipated or delayed. We model an electricity market as a mean-field game and show that, when users are price-takers, the strategy that selfish users follow coincide with a policy that maximises the social welfare, *i.e.*, that maximises the sum of the individual benefits. This means that the market is efficient. We then show how to compute numerically the equilibrium price by using the ADMM and the method described in Section 2.2. This numerical method extends our previous results presented in [Gas+13].

This paper is an illustration of the algorithm presented in Section 2.2. This method allows us to study numerically different scenarios that we parametrised by using historical wind generation data from the UK. As expected, the social welfare increases with the amount of flexible consumption. This is because our method solves the global optimisation problem. Nevertheless, we observe two counter-intuitive phenomena. First, when some players lack information, a larger quantity of flexible consumption can be detrimental to the whole system. Second, if a market based directly on the above Lagrangian decomposition leads to an efficient allocation of resources, it does not recompense the players in a fair manner. In the long term, this could discourage the installation of more flexible appliances. This suggests that the pricing mechanisms directly derived from this Lagrangian decomposition should probably be used as virtual pricing mechanisms: the virtual prices to compute a decentralised optimal control algorithm and a different mechanism should be deployed to encourage players to participate in this game (see for example Section 3.1).

2.3.2 A Mean-Field Game with Interactions for Epidemic Models

A second example is given in [DGG15], where we present a mean-field game where the objects follow an epidemic behaviour. We consider that the strategy of objects is how to vaccinate, with a cost associated with vaccination and with each infected object. The strategy of the individual object does not alter the dynamics of the population.

The objective of the individual object is to minimise its expected cost. We model the behaviour of the individual object as a continuous time Markov decision process. We show that there is a unique mean-field equilibrium and it is of threshold type, i.e., it consists of vaccinating at maximal rate from the beginning until a certain threshold and then not to vaccinate. We also consider a centralised problem, where a vaccination policy can be chosen so as to minimise the total cost in the system. We prove that the solution of the centralised problem, i.e., the global optimum, is also of threshold type.

We show that the thresholds at which the vaccination rate of the global optimum and the equilibrium change is characterised uniquely by the cost of vaccination. We build a pricing mechanism by which we force the equilibrium to coincide with the global optimum. According to our numerical experiments, the vaccination price charged to the individual objects should be always less than the actual vaccination price in order to force the individual to maximise social welfare.

2.4 Distributed Learning Algorithms for Semidefinite Programming

The approach presented in the previous sections assumes that the objects or the controller know, at least partially, the possible impact of the actions on the objective function. This is an *offline* optimisation setting and an agent can compute or estimate locally a best response to the actions of the others. In [CGM14; GM15], we consider a different approach and study an online approach. In an *online* approach, the agent does not have access to the form of the objective function and seeks to optimise its objective based on indirect (and possibly noisy) performance measurements.

Motivated by an application to wireless networks, we propose an online algorithm for stochastic multi-agent semidefinite optimisation problems. This algorithm is fully distributed and requires minimal (and possibly imperfect) gradient information and does not require any coordination between the optimising agents. This algorithm is obtained as a variable step-size stochastic approximation [Ben99] of a continuous-time matrix exponential learning scheme.

We establish the convergence of the algorithm. For our applications to wireless mobile systems, this implies the convergence of the network to a stable, optimum state. When applied to throughput maximisation in wireless multiple-input and multiple-output (MIMO) systems, the proposed algorithm retains its convergence properties under a wide array of mobility impediments such as user update asynchronicities, random delays and/or ergodically changing channels. Our theoretical analysis is complemented by extensive numerical simulations which illustrate the robustness and scalability of the proposed method in realistic network conditions. They also show that our method outperforms traditional water-filling techniques.

3 Markets Efficiency in Smart-Grids

While the previous sections were more theoretical, this section is focused only on applications to electrical grids. The current electrical grid is organised in a top-down approach. The systems are large and interconnected, *e.g.*, a single synchronous network for the whole of continental Europe.

The stability of these systems is guaranteed by a constant balance between electricity generation and consumption. This balance is maintained by a mixture of electrical markets and ancillary services[2]. Electrical markets operate at national levels and use forecasts to schedule the quantity of energy that will be generated by each power plant in advance. Deviations from forecasts are corrected in real-time by using ancillary services (mainly primary, secondary and tertiary reserve).

The increasing penetration of renewable energy sources in existing power systems leads to a number of changes in electricity market mechanisms. The wind and solar generators constitute local sources of electricity that are hard to predict, highly volatile and work at local scale. Their production can change unexpectedly at a much faster time-scale that consumption or conventional generation. There are currently no mechanisms to support the additional costs induced by power fluctuations. These costs are socialised and eventually charged to electricity customers. In this section, we present two pieces of work that aim at integrating more renewable energy sources in the markets. We first show how to use a coalition game approach to encourage consumers to exchange their energy locally. We then propose a market mechanism on how to make producers accountable for the costs induced by short-term power-variability. This section is based on [ST15; Bon+15].

3.1 Coalitions in Smart Grids

In [ST15], we investigate the problem of power trading coordination among "micro-grids" (MGs) (e.g., solar panels, wind turbines, plugged in hybrid electric vehicles, etc.) in smart power grids. With the purpose of minimising the amount of dissipated power during generation and transfer, we introduce an algorithm based on dynamic learning and coalitional game theory which allows the MGs to autonomously cooperate and self-organise into a set of coalitions. Hence, whenever some MGs have an excess of power while others have a need for power, it might be beneficial for these MGs to exchange energy with one another instead of requesting it from the main grid. The advantage of such a local exchange is twofold: 1) the energy exchange between nearby MGs can significantly reduce the amount of power that is wasted during the transmission over the distribution lines and 2) performing a local exchange of energy contributes further to the autonomy of the MG system while reducing the demand and reliance on the main electric grid. Thus, it is of interest to devise a scheme that enables such a local energy trade between MG elements that are in need of energy, i.e., "demanders", and MGs that have an excess of energy to transfer, i.e., "suppliers".

We formulate the problem using coalitional game theory in which MGs form a set of not necessarily disjoint and possibly singleton coalitions. The coordination among MGs determines the amount of power to transfer over each transmission line. We solve two problems, namely, how to motivate the MGs to form coalitions, and how to appropriately distribute the extra profits (e.g., payoff) produced through forming the coalitions. We achieve this by computing the Shapley value, which is the weight of the contribution of the MGs to their coalition. Figure 1 shows an example where two cooperative coalitions and one non-cooperative singleton coalition are formed. For instance, MG1 can decide to transfer 10% of its (surplus) quantity to MG2 and 30% to MG3 and the rest to the grid. MG4 can decide to transfer 50% of its quantity to MG3 and the rest to the grid. The rest of deficit quantities of MG3 and MG4 will be provided by the grid. MG5 will be completely served by the grid.

For achieving the best power distribution over transmission lines, we propose two dynamic learning algorithms: 1) coalition formation dynamic learning, and 2) power loss minimisation dynamic learning. Coalition formation dynamic learning achieves a coalition formation structure and then the complementary power loss minimisation dynamic learning leads the MGs to the maximum performance in terms of power saving. We show the stability (the convergence to a fixed-point) of both dynamic processes using the Kakutani fixed point theorem. As we show in [ST15], our approach enables MGs to come to a power trading coordination among themselves such that dissipated power in the proposed cooperative smart grid is only 10% of that in traditional power distribution networks.

[2]The ancillary services are services that maintain grid stability and security, such as frequency regulation. They can be provided by specific generator that can be activated quickly.

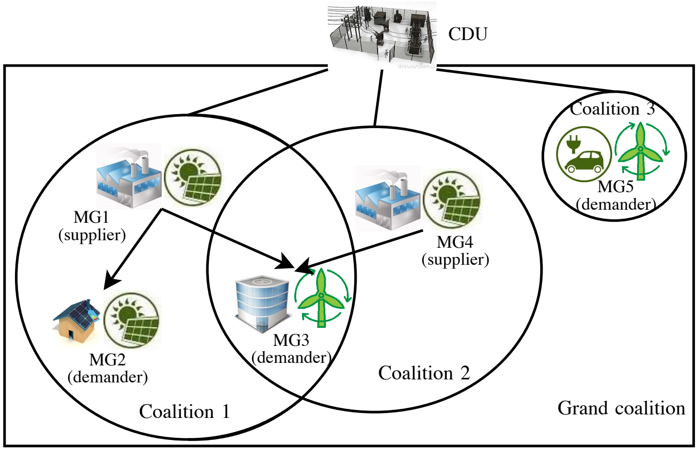

Figure 1: A sample of formed coalitions for cooperative power distribution.

3.2 Integration of Fast Time-Scale Variability in Electricity Markets

Wholesale electricity markets were proposed with the aim of making power system operation more economic and efficient. One of the key concepts is that of balance responsibility, for which both power suppliers and consumers are deemed accountable for deviations between their scheduled and actual production (resp. consumption). Renewable energy generation, with its limited predictability, is increasingly asked to take full part in such mechanisms and hence to be financially responsible for the energy regulation needs it induces. Current energy markets implicitly assume that an actor that plans to produce a certain quantity of energy will deliver this energy at a constant power. Hence, the deviations from plan are only penalised as a function of the energy mismatch of this actor. In contrast, when it comes to power-related regulation services, the costs are socialised and supported by the demand side (see for instance the ENTSO-E overview in [EE14]). However, a generator that produces power in a variable manner induces a cost, by creating the need to use ancillary services. In future power systems where the need for increased power-related services may substantially originate from the production side, power suppliers should be made accountable in a way similar to the case of energy-related regulation. Such mechanisms would also serve as an incentive for renewable energy generators to better self-regulate their own power fluctuations (for instance, using storage, coordination with demand-response, virtual power plant concepts, etc.).

In [Bon+15], we claim that the total cost of secondary reserves should also reflect the short-term power variability. Specifically, our objective is to find a way to generalise current market mechanisms to incentivize generators to deliver less fluctuating power profiles, while still respecting their energy-related commitments. Similar considerations could be made on the consumption side, there looking at fluctuations in power consumption. We emphasise here renewable energy generation since, if not coupled with any system allowing for smoothing fluctuations, e.g., using storage or demand-response in a virtual power plant setup, it is naturally expected that these renewable energy generation sources will deliver highly fluctuating power production.

In the above mentioned paper, we assess today's costs of power-related regulation services, based on the test cases of Germany and Western Denmark. We introduce a new market framework, complementing market participants revenues with a component reflecting the power regulation costs they induced, defined in various ways. We provide insights that show that our attribution mechanism is close to optimal and illustrate how revenue is shared among different market participants on real-data.

4 Moment Closure for the Analysis of Adaptive Strategies

Mean-field methods have been used mostly to study the performance of systems that are composed of a large number of identical and anonymous individuals that interact randomly with one another. In this last section, we present two papers that study two application examples that exhibit two properties that are important in CAS: heterogeneity and spatially constrained dynamics, for which we had to develop specific approximation techniques

In the first example, based on [GVH15], the difficulty arises from the non-homogeneity of the objects. In the second example, based on [FG14; Gas15], the difficulty arises from the local-interactions between objects.

4.1 Transient Analysis of Heterogeneous Systems: Application to Cache Management

Caches are used to speed up the access to a data source by adding a small but fast memory between the data source and the application or users that want to access objects from this source. Caches are omnipresent in today's computer architectures (in processors but also in content distribution networks). In [GVH15], we study the performance of a family of cache replacement algorithms. The cache is virtually decomposed into lists. When a requested item is not in the cache, it replaces an item of the first list (picked at random). When a requested item is in the cache, it moves forward one list and is exchanged with one item of the next list (picked at random). The classical policies RANDOM (one item is replaced at random) or CLIMB (at each request, an item move forward one place) can be obtained as special cases.

We study the cache content distribution and miss probability under the independent reference model (IRM model). This model assumes that there is a fixed population of objects and that the probability that an item is requested does not depend on which items were previously requested. Our focus in this paper is when objects are heterogeneous. Hence, the state-space of such a system is large (of the order of $\binom{N}{m}$ for a cache of size m and N possible objects). Despite this large state-space, we exhibit an algorithm to compute the steady-state probability for an item not to be in the cache.

The main originality of this work is an approximation method to study the transient behaviour of the system. To do so, we introduce a mean-field model that consists of a system of ordinary differential equations. We prove in [GVH15, Theorem 6] that this model becomes exact as the cache size and number of items tend to infinity. This result is illustrated in Figure 2, where we compare a simulation of the stochastic system with its ODE counterpart. This approximation allows us to provide guidelines on how to select a replacement algorithm, within the family considered, that achieves a good trade-off between the cache reactivity and its steady-state hit probability. We verify that these guidelines lead to a policy that outperforms classical policies such as least recently used (LRU) by using traces of real data of Youtube videos. This work also allowed us to disprove a natural conjecture from the 70's which was that increasing the number of lists always improves the performance of a cache [GVH15, Section 4.4].

Most of the convergence results that concern mean-field models assume that there is an homogeneous population of objects [BB08], or at least that the objects can be clustered in a finite number of clusters that all contain many similar objects. One of the major contributions of our paper [GVH15] is to extend these results to a fully heterogeneous mean-field model. Similar ideas can be used to justify the approximation that we use [Gas+15] to make predictions of available bikes in bike-sharing systems.

4.2 Incentives in Bike-Sharing Systems and Pair-Approximation

In [FG14], we introduce a model of bike-sharing systems based on a queuing network. One contribution of this paper was to propose and study an incentive mechanism for bike sharing, in which, when a user

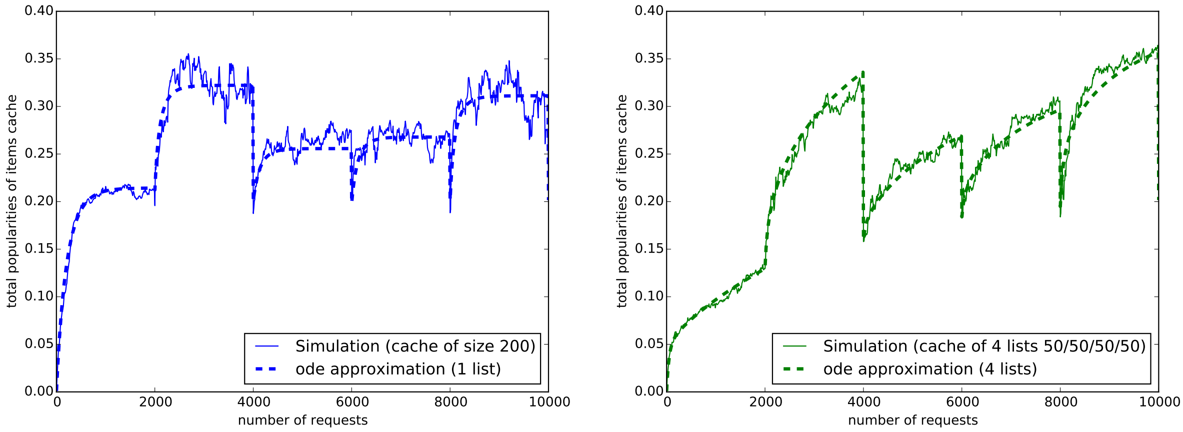

One list Four lists

Figure 2: Comparison of the transient behaviour of the ODE (dashed-line) and a stochastic simulation. In this simulation, the popularity of objects are reshuffled every 2000 request items. We observe that the ODEs predict very accurately the sum of the popularities of objects in the cache.

wants to perform a trip, he chooses two destinations at random and goes to the most empty one. This load-balancing strategy is known as the power of two-choice. When it is used to allocate jobs to server farms, it is known to dramatically improve the performance [VDK96; Mit96]. One of the contributions of [FG14] was to show that this is also the case for bike-sharing systems: when a user goes to the most empty station between two destinations chosen at random, the number of problematic stations decreases by an exponential factor. Such a system could be easily implemented in practice by offering a reward to a user that helps the system.

The theoretical analysis of [FG14] suffers from a limitation, which is that it assumes that the choices are done at random among all stations. This neglects a key aspect of the system: one might think of encouraging people to return a bike in the least-loaded of two neighbouring rental station but not to a station too far from their true destination.

In [Gas15], we study the power of two-choice in a setting where the two servers are not picked independently at random but are connected by an edge in an underlying graph. An exact analytic solution seems out of reach. Hence, we use pair-approximation, a technique in wide-spread use in biology [Oht+06]. We derive the equations and show that they describe accurately the steady-state of the system. Our results show that, even in a graph of degree 2, choosing between two neighbouring nodes improves dramatically the performance compared to a random allocation. This result has recently been verified using real-traces of data in [Des15].

5 Discussion

Main achievements In this deliverable, we have presented mathematical methods that allow one to study and enforce a CAS to be adaptive. We have shown how to deploy small distributed algorithms to each object in order to reach a global objective. These tools have been developed by considering mathematical models of CAS in terms of stochastic systems composed of many objects. What we call an *adaptive* system is a system that has the ability to change its dynamics in order to maximise a given objective function, by learning its impact on the environment.

One key contribution of this work is the development of a framework based on the use of mean-field game theory and Lagrangian decomposition to transform a centralised optimisation problem into a decentralised control algorithm. This algorithm can be proved to converge to a centralised optimum. We have shown how to apply these results in the context of smart-grids by using virtual market mechanisms. We have also studied real electrical markets and in particular, which modifications should be made to incorporate more renewable energy production. The markets try to model two aspects of the renewable energy production: the locality of the production and its high volatility. Our work shows how simple mechanisms can encourage people to build local coalitions and reduce their short-term volatility. Finally, we have shown how to use moment closure techniques to study the emergent behaviour of some adaptive systems that are heterogeneous or that have geographic constraints. One of the interesting contributions of this work was to disprove some conjectures from the 70's and show that some natural adaptive strategies might in fact decrease the performance of a system.

Future work We are currently expanding these results both by extending the theoretical results and by working on their application.

- One of our ambitions is to develop a mean-field game theory to solve optimisation problems. We aim at developing numerical methods that are widely applicable and at finding a minimal set of assumptions that guarantee the applicability of these results. This theory should be complemented by language design, in links with work packages 4 and 5.
- As we aim our theory to be applied, we are also investigating the application of these results to electrical distribution networks. This work is done in a collaboration with Schneider Electric and a co-supervised PhD student. The first application on which we focus is how to optimally curtail solar energy production to guarantee voltage constraints in a distribution network, by using robust distributed optimisation algorithms. Next, our goal is to investigate demand response and a possible implementation using the Link smart-meter development in France.
- Last, we are currently developing new analytical methods to study adaptive strategies in heterogeneous and spatially-constrained systems. We first want to extend the work reported in Section 4.1 to explore more complex strategies. The methodology developed in [TT15] and reported in Deliverable 1.2 and Deliverable 3.2 could be useful to speed-up the numerical evaluation. We also want to incorporate a spatial component in our bike-sharing prediction tool [Gas+15].

Relationship with other workpackages

WP1 Section 6 of Deliverable 1.1 contains some preliminary results about mean-field control, focusing on the application to smart-grid. In the present deliverable, we present a more generic framework application to CAS. Furthermore, the outcomes of WP1 are collected both in the current document and in Deliverable 1.2, which focuses on imprecise systems. In Deliverable 1.2, we show how to view heterogeneity as imprecision and how to use aggregation techniques to simplify the analysis.

WP2 In this deliverable, we show how the use of pair-approximation is useful to study systems in which two objects interact locally with their neighbours. WP2 also considers spatial modelling in terms of spatially heterogeneous environments. A natural extension would be to study the accuracy of pair-approximation in spatially heterogeneous environments, in particular in the case of bike-sharing systems.

WP3 The quantitative satisfaction score of spatio-temporal logic formulae is a real valued number which can be seen as a measure of robustness of the satisfaction of the property. As such, it can provide a function to be optimised. In that case, the mean-field method that we presented in this document can help producing decentralised control policies or parameters of policies that robustly maximise the satisfaction probability of a formula. The lumpability methods reported in Deliverable 3.2 would also be helpful for solving large mean field game problems, by automatically reducing the number of dimensions.

WP4 A mean-field language for optimisation – once the mean-field optimisation framework is stabilised, we aim to develop language primitives in CARMA that provide a bridge with the optimisation routines and the control mechanisms considered in the deliverable. One possibility would be to extend CARMA with constructs specifying the local goals of individual agents, with a semantics that constructs an appropriate optimisation problem from the model description.

WP5 The link with WP5 concerns two aspects. First, we are currently applying our control algorithms to the smart-grid case studies with a short-term goal of technology transfer. Second, we also want to pursue the development of numerical methods to be integrated into the tool, in collaboration with WP4.

References (from the Quanticol project within the reporting period)

[Bon+15] F. Bona, G. Nicolas, J.-Y. Le Boudec, P. Pierre, and T. Dan-Cristian. "Pricing variability in power delivery in markets with substantial renewable energy generation". In: *working paper* (2015).

[DGG15] J. Doncel, N. Gast, and B. Gaujal. "A Mean-Field Game with Interactions for Epidemic Models". In: *In preparation* (2015).

[FG14] C. Fricker and N. Gast. "Incentives and redistribution in homogeneous bike-sharing systems with stations of finite capacity". English. In: *EURO Journal on Transportation and Logistics* (2014), pp. 1–31. ISSN: 2192-4376. DOI: 10.1007/s13676-014-0053-5. URL: http://dx.doi.org/10.1007/s13676-014-0053-5.

[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.

[Gas15] N. Gast. "The Power of Two Choices on Graphs: the Pair-Approximation is Accurate". In: *ACM SIGMETRICS Performance Evaluation Review* (2015).

[GLBT14] N. Gast, J.-Y. Le Boudec, and D. Tomozei. "Impact of Demand-Response on the Efficiency and Prices in Real-Time Electricity Markets". In: *Proceedings of ACM E-ENERGY 2014.* 2014.

[GM15] B. Gaujal and P. Mertikopoulos. "A Distributed Stochastic Approximation Algorithm for Stochastic Semidefinite Programming". In: *submitted* (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.

[ST15] F. Shams and M. Tribastone. "Power Trading Coordination in Smart Grids Using Dynamic Learning and Coalitional Game Theory". In: *Quantitative Evaluation of Systems.* Springer, 2015.

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

References

[BB08] M. Benäim and J.-L. Boudec. "A Class of Mean Field Interaction Models for Computer and Communication Systems". In: *Performance Evaluation* (2008).

[Ben99] M. Benaïm. "Dynamics of stochastic approximation algorithms". In: *Seminaire de probabilites XXXIII.* Springer, 1999, pp. 1–68.

[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 (visited on 12/11/2013).

[Boy+11] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. "Distributed optimization and statistical learning via the alternating direction method of multipliers". In: *Foundations and Trends® in Machine Learning* 3.1 (2011), pp. 1–122.

[CGM14] P. Coucheney, B. Gaujal, and P. Mertikopoulos. "Distributed optimization in multi-user MIMO systems with imperfect and delayed information". In: *Information Theory (ISIT), 2014 IEEE International Symposium on.* IEEE. 2014, pp. 3097–3101.

[Des15] P. S. Dester. *The Power of Two Choices in Bike-sharing Systems.* Tech. rep. 2015.

[DGP09] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. "The complexity of computing a Nash equilibrium". In: *SIAM Journal on Computing* 39.1 (2009), pp. 195–259.

[EE14] ENTSO-E. *ENTSO-E Overview of transmission tariffs in Europe: Synthesis 2014*. Tech. rep. `https://www.entsoe.eu/publications/market-reports/Documents/SYNTHESIS_2014_Final_140703.pdf`, accessed 20/08/2015. ENTSO-E, 2014.

[Gas+13] N. Gast, J.-Y. Le Boudec, A. Proutière, and D. Tomozei. "Impact of Storage on the Efficiency and Prices in Real-Time Electricity Markets". In: *Proceedings of ACM E-ENERGY*. 2013.

[LL06a] J.-M. Lasry and P.-L. Lions. "Jeux à champ moyen. I–Le cas stationnaire". In: *Comptes Rendus Mathématique* 343.9 (2006), pp. 619–625.

[LL06b] J.-M. Lasry and P.-L. Lions. "Jeux à champ moyen. II–Horizon fini et contrôle optimal". In: *Comptes Rendus Mathématique* 343.10 (2006), pp. 679–684.

[LL07] J.-M. Lasry and P.-L. Lions. "Mean field games". In: *Japanese Journal of Mathematics* 2.1 (2007), pp. 229–260.

[Mit96] M. Mitzenmacher. "The power of two choices in randomized load balancing". PhD thesis. University of California at Berkeley, Department of Computer Science, 1996.

[Mye13] R. B. Myerson. *Game theory*. Harvard university press, 2013.

[Nas51] J. Nash. "Non-cooperative games". In: *Annals of mathematics* (1951), pp. 286–295.

[Nis+07] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani. *Algorithmic Game Theory*. New York, NY, USA: Cambridge University Press, 2007.

[Oht+06] H. Ohtsuki, C. Hauert, E. Lieberman, and M. A. Nowak. "A simple rule for the evolution of cooperation on graphs and social networks". In: *Nature* 441.7092 (2006), pp. 502–505.

[Put14] M. L. Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.

[VDK96] 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.

[Wei97] J. W. Weibull. *Evolutionary game theory*. MIT press, 1997.