FRONT PAGE

PROJECT FINAL REPORT

Grant Agreement number: 600708

Project acronym: QUANTICOL

Project title: A Quantitative Approach to Management and Design of Collective and Adaptive Behaviours

Funding Scheme: Small or Medium scale focussed research project (STREP)

Topic: ICT-2011 9.10: FET-Proactive `Fundamentals of Collective Adaptive Systems'

Period covered: from 01/04/2014 to 31/03/2017

Name of the scientific representative of the project's co-ordinator[1], Title and Organisation:

Prof Jane Hillston, University of Edinburgh

Tel: +44 131 650 5199

Fax: +44 131 651 1426

E-mail: quanticol-coordination@ed.ac.uk

Project website[7] address:

www.quanticol.eu

[1] Usually the contact person of the coordinator as specified in Art. 8.1. of the Grant Agreement.

QUANTICOL

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

quantiCOL

http://www.quanticol.eu

D7.4

Final project report

Revision: 1.0; May 5, 2017

Author(s): Jane Hillston (UEDIN)

Due date of deliverable: Month 48 (March 2017)
Actual submission date: 5th May 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

1 Executive Summary

The QUANTICOL project has sought to support designers and operators of collective adaptive systems (CAS) through the construction and analysis of models. Systems are considered to fall into this category when their functionality is derived from the interaction of a large number of agents. The agents involved may be computational (e.g. pieces of software, sensors, etc.), physical devices (e.g. RFID tags, buses, bikes, electrical generating elements or energy consuming units) or human (passengers, householders or users). Through these complex interactions, agents may both compete and cooperate to achieve both individual and global goals. Understanding and predicting the behaviour of such systems is extremely difficult as the outcomes may be *emergent* and non-intuitive with respect to the behaviour and goals of individual agents.

Often modelling focusses solely on the functional behaviour — aiming to understand what desirable behaviour is possible and what undesirable behaviour may be prevented. However in the context of CAS notions of correctness go beyond ensuring that agent behaviour satisfies some functional constraints. We are often concerned to know the likelihood of possible behaviours, the timescales over which behaviours develop, whether access to resources will be equitable, and how behaviour and resources may be distributed over space. The *quantitative* questions place additional demands on how models are constructed and analysed, but facilitate questions about a richer set of properties related to the behaviour of systems. Thus the modelling framework developed in QUANTICOL supports the construction and analysis of *quantitative models*. Given the size and complexity of CAS, a particular challenge is with respect to scalability and we have placed emphasis on this, achieving results well beyond the state-of-the-art at the start of the project.

The project has successfully initiated several areas of research resulting in a rich suite of software tools underpinned by highly innovative techniques for representing and analysing CAS. We have substantially extended the framework of mean-field analysis in order to analyse systems in which agents operator on different scales, operationally, temporally or spatially. For example, techniques of continuous approximation have been developed even for cases when some aspects of behaviour must be kept discrete, resulting in hybrid models. We have studied how best to capture spatial aspects of behaviour and investigated when it is essential to retain information about the locations of agents, and when such information can be safely abstracted. Even when spatial approximation is applied, scalability can remain a challenge; consequently we have established a number of equivalence relations which reduce models in an exact manner in order to form more compact and efficiently solved models. We have also developed a number of approximation techniques which can also substantially reduce the computational effort needed to analyse models. Straightforward execution of models is often not enough, as more sophisticated forms of interrogation are required in order to derive all the necessary information about systems behaviour. For this purpose we have investigated the use of a number of temporal, spatial and spatio-temporal logics and their associated model checking algorithms, greatly expanding the types of queries that can be automatically applied to CAS. Whilst analysing models is a significant challenge, it is also important within a modelling framework to establish a suitable language in which to express models, and for this purpose we have designed the Carma language and CaSL, the Carma Specification Language. Carma has a number of novel features including attributed-based communication, a quantitative semantics, and a distinct environment which exercises influence and imposes constraints on individual agents. CaSL offers a more programmatic-interface to building models in Carma and greatly enhances the accessibility of the language. Throughout the project we have been motivated by, and validated our results against, models of smart urban transport and smart energy networks, clear examples of CAS.

Results of the project have strong potential impact both for those concerned with building and operating CAS, but also for the users of CAS who will benefit from systems better tailored to their needs. All results have been disseminated through standard academic channels, as witnessed by an impressive list of publications, but also through engagement with stakeholders. Furthermore we are actively exploring exploitation routes for at least two results from the project.

2 Summary of Project Context and Objectives

Project Context The design of collective adaptive systems (CAS for short) must be supported by a powerful, well-founded framework that offers formal modelling and quantitative analysis. CAS consist of a large number of spatially distributed heterogeneous entities with decentralised control and varying degrees of complex autonomous behaviour. These entities or agents may be competing for shared resources even when collaborating to reach common goals. Often humans are both agents within such systems and end-users standing outside them. As end-users, they may be completely unaware of the sophisticated underlying technology needed to fulfil critical socio-technical goals such as effective transportation, communication, and work. The pervasive but transparent nature of CAS, together with the importance of the goals that they address, mean that it is imperative that thorough *a priori* analysis of their design is carried out to investigate all aspects of their behaviour, including quantitative and emergent aspects, before they are put into operation. We want to have high confidence that, once operational, they can adapt to changing requirements autonomously without operational disruption. Unfortunately, the defining characteristics of these systems mean that their (possibly non-linear) behaviour is often highly unpredictable or counter-intuitive. Formal, scalable, quantitative analysis, which provides multiple perspectives on system behaviour while being based on well-established reasoning techniques, is imperative to master such complex systems.

However quantitative analysis requires models which are enriched with information about the timing and probability of events within the system. Whilst existing formalisms had been designed to capture and reason about the quantitative aspects of behaviour of agent-based systems (e.g., stochastic process algebras such as PEPA, EMPA and IMC) at the time the project started these offer only limited support for collective behaviour and little or no support for adaptation. Conversely process algebras which represented adaptation, for example by incorporating a knowledge store that influenced agent behaviour as in SCEL, did not capture quantitative reasoning. Thus care would be needed to design a high-level modelling language which is able to capture the possible actions that a system might undertake as well as the likelihood and duration of those actions, and allow those actions to be influenced by the state of the system.

When the systems are spatially distributed, and the behaviour of agents within the system may depend on their location within the system, then spatial information must also be captured. Some formalisms existed which would represent the location of agents, but many of these were based on hierarchical or organisation notions of space rather than physical space which is often an element in CAS. Synchronous communication is often the key form of interaction in process algebra models but the open nature of CAS, in which agents may enter or leave the system throughout execution, makes synchronous communication unsuitable. Specifically, agents within the model may not know who are the possible recipients of a message at any given time, or even if such recipients exist.

Furthermore enriching the language to capture all these aspects of behaviour enlarges the potential state space of CAS models, a problem further exacerbated by the sheer scale of the systems considered. Thus once constructed such models pose significant challenges for analysis because of the number of possible states that the system might find itself in, the so-called *curse of dimensionality* or *state space explosion* problem. In order to carry out quantitative analysis the language must be given a formal semantics in terms of a suitable mathematical model amenable to solution in order to reason about the quantitative behaviour. The most common choice for this underlying mathematical model is a continuous time Markov chain (CTMC) but numerical solution of such CTMCs will be infeasible for all but the smallest CAS and stochastic simulation can require considerable computational resource.

Thus the objective of the QUANTICOL project — to develop an innovative formal design framework that provides a unique specification language for CAS and a large variety of tool-supported, scalable quantitative analysis and verification techniques — was well-motivated but extremely ambitious, posing both technical and engineering challenges.

Project Objective The developed design framework was intended to not only support quantitative analysis and verification, but also to enable and facilitate experimentation and discovery of new design patterns for emergent behaviour and control over spatially distributed CAS. It would support both agent-based modelling techniques and equation-based techniques starting from system specifications at the individual (micro) level.

The work in the project has been driven by smart city applications which can be very large scale CAS comprised of heterogeneous entities with spatially inhomogeneous distribution. These demand computationally scalable approaches and provided an ideal domain in which to develop and test the QUANTICOL framework. Since scalability was a major issue a primary focus of the project was to avoid the construction of full discrete state spaces (as required for CTMCs) but instead extend and exploit mean-field and fluid flow techniques, expressed as ordinary differential equations (ODEs) rather than simulation-based approaches. Nevertheless, we recognised that the latter would still play an important role in the exploration of smaller systems, in the validation of new techniques and as the only alternative for models which do not satisfy the conditions to make them suitable for continuous approximation.

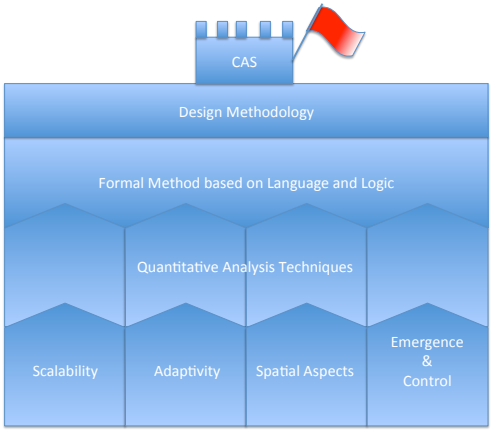

Figure 1: Quanticol research vision

Figure 1 shows the QUANTICOL vision of the development of the research in this project. Starting from the basement of the castle, the principal case studies have driven the development of theoretical extensions of mean field and fluid approaches tackling four distinct challenges: 1) scalability 2) (self-)adaptivity 3) spatial inhomogeneity and 4) design and control of emergent behaviour. Through the work of the project we have developed a deeper understanding of these areas and established several foundational results, greatly enhancing the scope of analysis techniques based on mean-field or fluid approximation, characterising cases in which spatial abstractions can be made to improve scalability of solution and establishing new algorithms for control in the context of smart grid systems and more broadly. Meanwhile, higher up in the castle, process algebraic verification techniques have developed and integrated into a uniform design framework supported by analysis and verification tools exploiting existing and new results in mean field and quantitative formal methods based on process algebra and temporal logics. Subsequently, the new theoretical results have been embodied in the software tools that populate the framework, making the results available in a software engineering context. The techniques and tools have been tested and validated on various case studies ranging from small exemplifying systems (for the purpose of illustration of the techniques and for tutorials), to variants of the smart city applications that were driving the theoretical parts of the project. Furthermore, the experience in the design and analysis of a range of case studies has been used to develop a pathways for design and analysis that will guide software engineers in the development of CAS for smart cities.

Research Objectives The structure of the project reflects the major QUANTICOL objectives. The main objectives and the work packages, as originally defined, are listed below.

1. Development of a rigorous mathematical framework to engineer emergent behaviour in multi-scale systems.

2. Definition of a formal framework enabling the description of systems with spatial aspects with a process algebraic language, and supporting different kinds of representation of space and abstractions based on mean field techniques in an automatic way.

3. Development of stochastic temporal logics and novel forms of scalable model-checking approaches exploiting mean-field/fluid flow approximations that can address the verification of properties at the macro and micro level and their combination.

4. Development of a generic language offering linguistic support for high-level CAS modelling, expressing adaptivity and spatial awareness and enabling emergent behaviour.

5. Design and implementation of a software tool suite will be pursued which will provide a unified formal framework for the specification, analysis, and verification of models of CAS.

The target outcomes of the project included a scalable, quantitative design and verification framework and tools for large heterogeneous CAS to make model-based design of CAS of unprecedented size and complexity feasible where traditional agent-based approaches, reliant on extensive simulations, would be prohibitively costly.

Major Achievements of the Project The project has made substantial progress towards the research objectives identified in the proposal, and exceeded our expectations with respect to several of them. In particular:

1. We have thoroughly explored rigorous mathematical analysis of the emergent behaviour in multi-scale systems, through both modelling and data analysis, and proposed a number of novel techniques which can be used in practice by designers, operators and users to better understand and predict the behaviour of systems. For example we have shown how data from existing bike sharing systems can be used to parameterise predictive models, how these models can be suitably abstracted so that predictions can be made in real time, and offered algorithms for making predictions.

   Key results in this area have substantially extended the class of models for which mean-field approximations can be meaningfully derived. We have established new mean-field approximations for CAS with multiple scale behaviour (either with respect to time or population sizes) and imprecise and uncertain behaviour. We have also established a new refined approximation that depends on the system size N. This offers much more accurate results for moderately-sized systems than classic mean-field approximations. Furthermore we have developed numerical methods to analyse the class of limit models for uncertain and imprecise population models allowing the properties of such systems to be studied quantitatively.

   Understanding and predicting the behaviour of CAS is not enough — we need also to be able to control the behaviour so that it stays within acceptable bounds. We have developed generic methods for constructing control algorithms that lead to efficient use of resources. These methods are applicable to a wide variety of CAS, but we have demonstrated them on specific smart-grid and electricity market applications.

2. Much of this work is centred on the formal process algebraic language, CARMA (Collective Adaptive Resource-sharing Markovian Agents), developed during the project. This richly expressive language has been shown to be capable of capturing a wide range of CAS, dealing with problems such as open-ness, scale and spatial restrictions through attribute-based communication, explicit representation of locations and an explicit representation of the environment in which

agents operate. The language is equipped with a structured operational semantics giving rise to an underlying CTMC, which forms the basis of quantitative analysis both through agent-based and population-based simulation and mean-field approximations. We have sought to make this language accessible to a wide range of users, including those not familiar with process-algebra based formalisms:

- models are specified in CaSL, the CARMA specification language, which provides a more programmatic-style of modelling;
- we offer a rich set of spatial constructs which ease the description of systems in which space plays a major role;
- the language is equipped with a rich set of syntactic and static analysis checks which help to ensure that models are free from trivial errors;

3. Throughout the project we have investigated a number of scalable verification techniques and developed innovative and efficient fluid and on-the-fly mean field model checking algorithms that address local, global and local-to-global properties. These ODE-based approaches are now additionally supported by formally-defined and automated model reduction techniques which can substantially reduce the size of the system of ODEs that must be considered. Since the conditions for fluid approximation are not always satisfied, we have also enhanced and integrated the statistical model checker MultiVeStA and developed a number of other model reduction techniques suitable for CAS. We initiated a study of the logical approaches to the treatment of spatial information based on closure spaces, and this has resulted in the topochecker tool which has been demonstrated to be applicable to a wide range of case studies.

   Designers are often concerned with a set of alternative configurations of systems which could be deployed. Thus another form of scalable analysis is based on approaches to family-based verification of many alternative (but closely related) models to be carried out without resorting to complete enumeration of the alternatives to be verified individually. Within the project we have developed a number of new approaches to this problem including the quantitative feature-oriented language QFLan which may be used to study properties such as quality of service, reliability or performance of dynamically reconfigurable product lines.

4. The CARMA specification language, CaSL, provides a generic language offering linguistic support for high-level CAS modelling. The inclusion of an environment, with functions to determine rates and probabilities within the model and allowing those functions to depend on the current *state* of the components undertaking the actions allows a rich form of adaptivity to be incorporated into models (here *state* is comprised of both the logical state and attribute values or knowledge within a component). The attributes associated with components, readily lend themselves to capturing spatial awareness, and coupled with attribute-based communication and functional rates and probabilities, means that *what* a component can do, depends on *where* the component is located, a common feature of CAS. The adaptive nature of rates, probabilities, and even which actions are possible means that the behaviour of the system is emergent upon the behaviour of the individual components and their interactions. Moreover models are *open* in the sense that components can join or leave during the system evolution.

5. The QUANTICOL framework is supported by an extensive software tool suite, much of it developed as Eclipse plug-ins, offering portability and familiarity to many potential users. Many of the newly developed techniques have been implemented in software as a proof of concept, allowing ourselves and others to thoroughly experiment with the techniques on realistic case studies. We have implemented an Eclipse plug-in to support the construction and analysis of CARMA models, using CaSL. This is publicly available and demonstrated on a number of case studies.

3 Main Science and Technology Results

3.1 Work Package 1: Emergent Behaviour and Adaptivity

Throughout the project, work package 1 sought to extend the applicability of mean-field approximation techniques for studying the dynamic behaviour of CAS. Mean-field approximation refers to a collection of techniques that make the analysis of large collective stochastic systems easier, through taking advantage of asymptotic results about the behaviour of the collective when the population tends to infinity. 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 this work package we aimed to extend the classical mean field approach in order to deal with specific features of CAS, like their multi-scale nature and their intrinsic uncertainty, further tailoring these methods to provide algorithms to define control policies. Moreover, we hoped to be able embed this new capability into the language being designed for modelling CAS in work package 4. The intention was to test our new approach on CAS case studies, particularly focusing on smart electric grids in this work package, as their scale and need for control made them especially well-matched to the techniques to be developed in this work package.

The focus of Work Package 1 during the first period was on investigating issues related to multi-scale modelling, particularly in the context of modelling smart grids. Multiple scales are inherent in many CAS. The different scales may be *temporal*. For example, adjusting the production of a large power plant takes hours while appliances can be switched on or off the instant a signal is received. They may also be *organisational* and represent a hierarchical structure; for example a smart building that contains many rooms. Our initial investigation into multi-scale modelling resulted in a number of publications and a detailed technical report surveying the techniques that were currently available for scalable analysis of multi-scale systems, assessing their feasibility for the future work of the project.

Within temporal multi-scale modelling generally two timescales are assumed, decomposing the model into *fast* and *slow* behaviour. This situation emerges when the states of some entities evolve at a much faster time scale than others. A common tactic is to abstract from the behaviour on one time-scale whilst representing the other in detail. As we consider CAS systems, our emphasis is on scalable analysis techniques related to mean-field approximation. Thus our review of existing techniques focussed on time-scale reduction techniques which can be readily applied to mean-field models. We found that the reduction techniques for deterministic dynamical systems, are much more mature than on-going work on stochastic systems. Thus our review found a number of time-scale reduction techniques that could be readily applied to mean-field models and suggested a number of approaches towards developing analogous techniques for stochastic systems. Initial results were published in a paper in which we established a fluid analysis technique for models with immediate transitions, representing systems in which the fast time scale is so fast that state transitions on that time scale have negligible delay. This technique allows, for example, the analysis of smart grid scenarios with many generators and providers (see for example, the model in Figure 2, and is most effective in optimisation problems such as trying to maximise electricity production, subject to environmental constraints such as on pollution or the carbon footprint.

Within organisational multi-scale modelling we identified two cases. Firstly we considered systems in which there are multiple population scales. For example, within a smart grid there might be one centralised controller which interacts with many appliances (see for example Figure 3). We established that within such systems, when the population grows, the limit is naturally described by a stochastic hybrid system. Project partners EPFL/INRIA and CNR had both previously worked on this area and we developed a stronger understanding of how these previous results are related.

In the second case we considered *systems of systems*, systems whose entities may themselves be regarded as systems, to arbitrary depths of hierarchical organisation. Here we developed a formalism based on nested automata, encompassing rich forms of interaction between automata at the same (horizontal) or different (vertical) levels in the hierarchy. This formalism allowed us to automatically

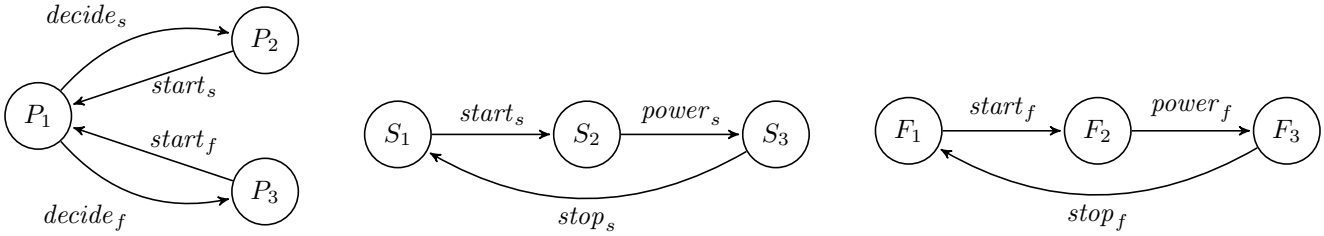

Figure 2: Fast and slow generator model in the interacting automata formalism of *Differential Analysis of Interacting Automata with Immediate Actions*. (Luca Bortolussi and Mirco Tribastone, 7th International Conference on Performance Evaluation Methodologies and Tools, Torino, Italy, December 2013.)

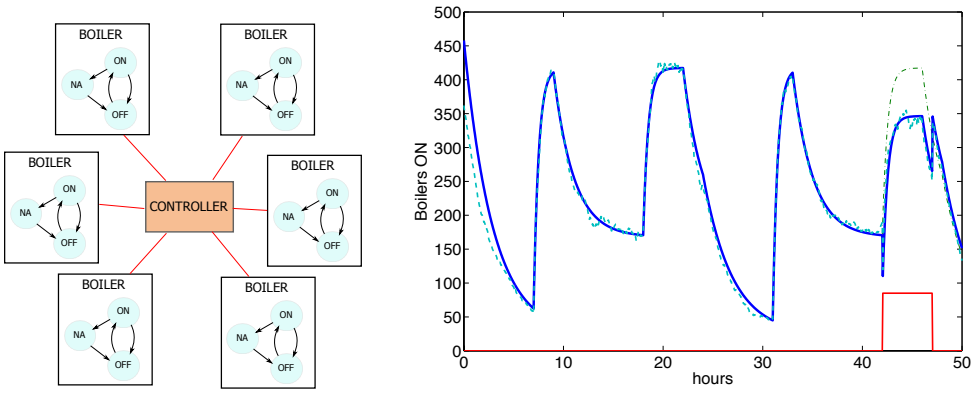

Figure 3: Simple model of a smart water boiler, which can be switched off by a central controller to act as a virtual energy reservoir. Each boiler (left picture) has three states: on, off, and not available. This last state corresponds to the boiler being switched off by the central controller. Boilers switch randomly between states on and off. The on rate, in particular, depends on the period of the day, and is higher during the morning and evening, when people usually take showers. The controller, in response to high energy demands, can turn off only boilers that are in the on state. We assume it turns off a fraction of active boilers, here 50%. Reactivation of boilers happens after a random time, connected with the duration of the energy demand peak. If the number of boilers is sufficiently high, then we can approximate it as continuous. The controller, instead, remains a discrete agent, giving rise to a hybrid model. In the right figure, we show a trajectory of the hybrid model (boilers on, solid blue line), compared with a trajectory of the fully stochastic model (dashed cyan line) for a small system with 500 boilers. Already in this case the accuracy is good. The red square line at the bottom shows the period in which the controller turned the boilers off. The thin blue dash-dotted line shows the number of active boilers as if the controller would have not acted, visually depicting the amount of energy saved. In this example, we have noticed a speed-up in the simulation of one order of magnitude with respect to a wholly discrete simulation.

reduce the complexity of the mean-field equations representing the full system, by exploiting symmetries in the model. This method can be applied iteratively, to construct hierarchical abstractions of systems. We illustrated our method to describe the behaviour of a collection of smart buildings.

The final piece of work reported for this work package during the first year concerned control algorithms for electrical systems in the context of renewable energy supplies, illustrated in Figure 4. In this work, we demonstrated the use of optimization tools for building such algorithms. Specifically we modelled and dealt with two specific challenges: large forecast uncertainties and the presence of

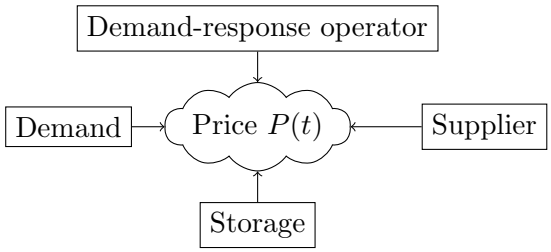

(a) Schematic model of the real-time electricity market.

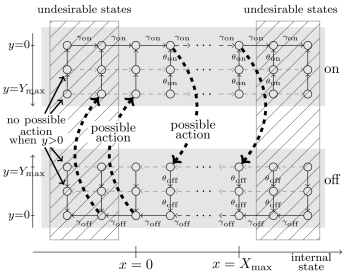

(b) Each flexible load is modelled by a Markov chain.

Figure 4: The real-time electricity market model and the demand-response model.

delays due to multiple time scales. The study considered two directions, based on centralized and distributed control respectively. We developed storage and demand/response management policies, where a central controller sends signals to smart users to adapt the consumption to the production. These policies have been shown to be more robust to forecast errors than existing strategies.

The work of work package 1 during the second period developed across the three distinct tasks within the work package. Within Task 1.1 the investigation of multi-scale modelling, which was initiated during the first period was consolidated with several novel results related to the analysis of systems which exhibit multi-scale behaviour in a number of ways. In particular new results were obtained for mean field limits of systems with heterogeneous population sizes. This enlarges the class of models amenable to this scalable approximation technique because previous results relied on all populations scaling in the same way.

Other significant results were obtained for mean field limits in the presence of uncertainty and when there are multiple time scales in operation in the system. Here the type of uncertainty considered is related to the rates at which the events in the system occur and encompasses both the case that the rate is constant but unknown (termed *uncertain*) and the case that the rate may fluctuate during the operation of the system (termed *imprecise*). We were able to demonstrate how the mean field limits can greatly simplify the study of uncertain and imprecise population models, and Markov Decision Processes (see for example, Figure 5). Moreover we expressed the limit in terms of differential inclusions and developed numerical methods to analyse such models. For models with multiple time scales we established conditions which guarantee the correctness of exchanging the order in which time scale reduction and mean field limits are taken. All these results represent a substantial enlargement of the class of CAS models which are amenable to scalable analysis based on mean field approximation.

In Task 1.3 first steps were taken to identify how these results can be incorporated in the analysis techniques associated with the modelling language developed during the project in work package 4, CARMA. Specifically in internal report IR1.1 we have shown how to map a CARMA model to an appropriate Markov population process and taken initial steps towards determining syntactic conditions in CARMA to identify which mean field approximation approach is suitable for a given model.

In the second period work package 1 also encompassed work on Task 1.2, which was initiated in month 13. The focus of this task was on control and adaptivity within CAS. Both centralised and distributed control algorithms have been investigated with the objective of finding policies which will lead to efficient use of resources. Whilst generic methods have been developed, with wide applicability for CAS, specific examples related to smart grids and electricity markets were taken as motivating examples and demonstrations. A number of significant results were achieved in this task during the second period. Specifically:

- Considering adaptivity in terms of solving an optimisation problem, using a decentralised algorithm, we developed a general framework based on the use of mean-field game theory and Lagrangian decomposition and shown how, given a centralised optimisation problem, we can

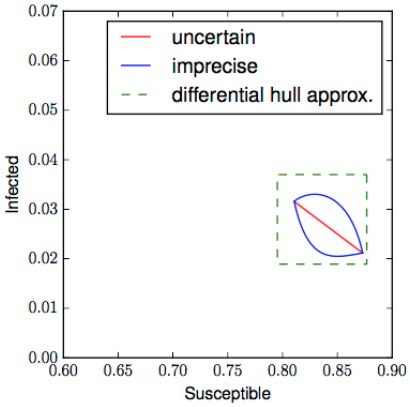

Figure 5: Steady-state regime of imprecise and uncertain SIR (susceptible/infected/recovered) models investigated via differential inclusions.

define a mean-field game whose equilibrium is an optimal allocation. We also developed a distributed algorithm that is proven to converge to this equilibrium and therefore to a globally optimal allocation of resources. This work has been reported in a number of publications.

- We also devised a method to start from an optimisation problem and 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 electricity distribution networks) and can modify the behaviour of each agent. We adopt a bottom-up approach and study how adaptivity can be obtained by using market mechanisms and study two approaches to modifying these markets to incorporate more renewable energy. These are to incentivise the creation of coalitions of micro-producers that exchange energy locally and to make people accountable for short-term variability.

- We investigated two new methods based on moment closure to study how reasonable control policies can improve or deteriorate the performance of a CAS. We obtained approximate models for two specific features: having many heterogeneous objects and a space-constrained load balancing policy. These works showed that some policies that were previously conjectured to be optimal in practice reduce the overall performance.

During the final reporting period, we continued the work of work package 1, by developing both the theory related to mean-field approximation and control, and its application, including tool development in collaboration with work package 5. Several significant contributions have been made in this final reporting period.

We have made a major contribution through analysing 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, whereas previously it had been assumed to decrease as $1/\sqrt{N}$. This allowed us to define 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 than the original mean-field approximation. Indeed it allows mean-field approximation to be applied to population sizes that would previously not have been considered. Our experimentation has shown that even for relatively small systems ($N = 10$) the refined approximation is very accurate, and a significant improvement on the classical mean-field approximation. This result opens many interesting questions that will continue after this project.

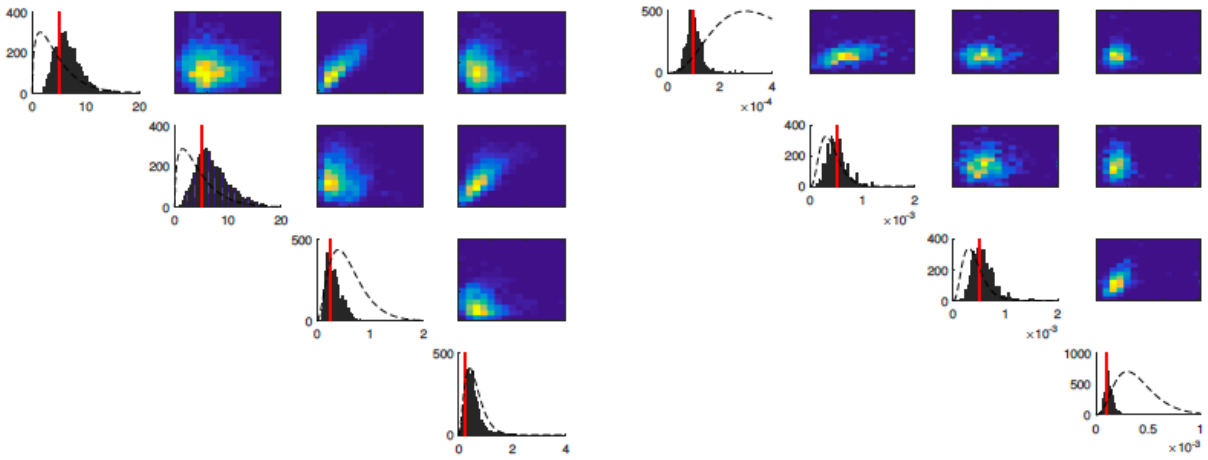

Figure 6: An example of results generated by novel estimation methods developed during the QUANTICOL project: histograms represent the posterior distribution for each parameter, and heat maps showing pairwise correlations. True parameter values are shown in red.

The second major contribution in this work package in the final period concerns the development of numerical algorithms to quantify the effect of uncertainties on the parameters of the systems. Within the QUANTICOL project, we construct and study stochastic models of CAS. 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. We also developed novel machine learning algorithms which allowed the degree of uncertainty associated with parameters within a model to be refined based on data observations (see Figure 6).

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. As in earlier periods, another strand of the work within this work package has been related to exerting control over systems to achieve desirable outcomes. Our third contribution in this period concerns the use of mean-field approximations to create efficient control policies. Using this approach we have compared centralised and decentralised policies. In particular, we obtained a counter-intuitive result, 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 have developed 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 smooth their consumption. Figure 7 shows the effectiveness of the real-time control algorithm in the distribution network.

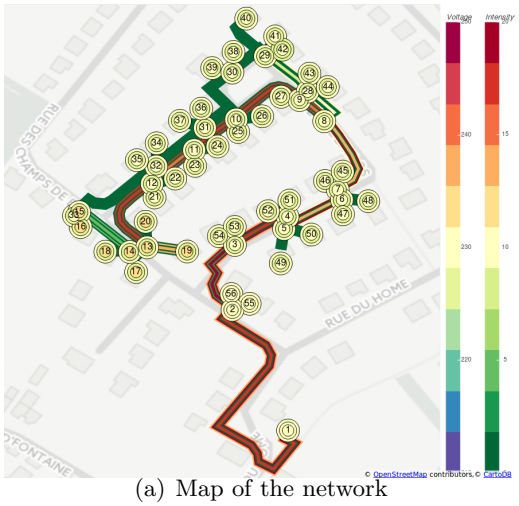

(a) Map of the network

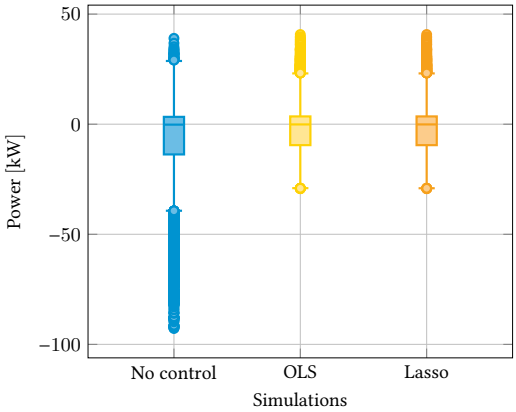

(b) Power flow at the transformer : the transformer should operate between -30 and +40 kW. Without control, this limit is often violated. This is no more the case with our control mechanism (OSL and Lasso correspond to two variants of our algorithm).

Figure 7: Map of the electrical distribution network and power flow at the transformer in the model.

3.2 Work Package 2: Collective Adaptive Behaviour in Space

Work package 2 aimed to develop a methodological framework for spatial modelling based on sound mathematical foundations. Thus initial work looked to identify those approaches to spatial representation that are most amenable to use in the context of modelling CAS when efficient evaluation is a prime consideration. Once such approaches were identified, this work package supported the embedding of these formal approaches into the languages and logics developed within QUANTICOL, as well as linking them with approximation results available in the literature and from the work of the project. A third strand of work was concerned with offering guidelines in practical ways to make use of spatial data when constructing models of CAS, particularly those related to transportation in smart cities.

Thus during the first year of the project the primary focus of work package 2 was on surveying the existing approaches to incorporating spatial aspects within dynamic mathematical models and developing a deeper understanding of which of these approaches may be useful for the further development of the modelling languages and analysis techniques later in the project. So during the first period the major output of this work package was a technical report (QUANTICOL TR-QC-05-2014) which reviews a diverse literature including mathematics, ecology, biology and computer science.

Space is important in the QUANTICOL project because the project case studies, especially those related to smart urban transport have strong spatial elements. Therefore the review of the existing work, its scope, applicability and limitations, provided an essential foundation for the future development of QUANTICOL. In particular we established a classification of the different approaches reviewed. The objective of the classification is to make clear what approaches are available and how they differ from each other in order to guide future work on spatial approaches within the project. Furthermore, the classification was applied to the initial work that has been done within the project on smart transport case studies, providing valuable insight into the most promising approaches for further study as shown in Figure 8.

The review also considered the availability of mean-field and other approximation techniques for the surveyed spatial modelling techniques, bearing in mind the scalability issues that inevitably arise in CAS.

Essentially three distinct approaches to capturing space within dynamic models were identified:

1. Space can be discrete and described by a graph of locations. Depending on the structure of the

Time	continuous or discrete			
Aggr	none		state and/or space	
State	discrete	continuous	discrete	continuous
Space				
discrete				
general	PEPA bus HYPE bus HYPE1 bike StoKlaim bike		PEPA bus PALOMA bike EPFL2 bike StocS bike PEPA-S bike Bio-PEPA bike CTMC bike	
regular	CTMC route		EPFL1 bike	
homo-geneous			EPFL1 bike DTMC bike	
continuous			HYPE2 bike StocS bike extension	

Figure 8: Classification of QUANTICOL case studies in terms of time, aggregation, state and space

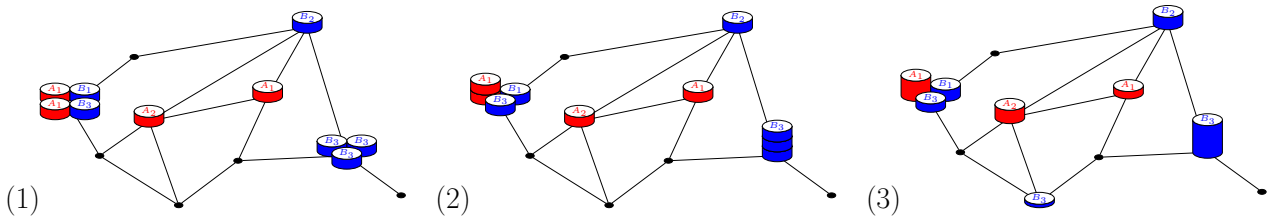

Figure 9: Discrete space: (1) discrete state for each individual; (2) individuals aggregated by state, population count is discrete; (3) individuals aggregated by state, population count is continuous and approximates the discrete population count

graph and the parameters associated with locations and movement between locations, discrete space can be classified as regular or homogeneous. Figure 9 illustrates different treatments of individuals and populations in discrete space.

2. Space can be seen as continuous: as Euclidean space in one, two or three dimensions.

3. Space can also be considered abstractly as topological space, whether discrete or continuous, and this approach allows for reasoning about concepts such as adjacency and neighbourhoods.

For the QUANTICOL project the first and last approaches were the most promising and we developed preliminary guidelines arising from the review and classification which were carried out. These can be summarised as follows.

- Existing mean-field techniques for discrete space often focus on global measures rather than an understanding of spatial heterogeneity. Furthermore, they are often specific to the scenario being modelled, so different specific techniques or more general techniques were likely to be required for QUANTICOL.

- The existing literature uses standard ODE techniques. The ODEs from QUANTICOL models were anticipated to be more complex and require more advanced techniques, particularly when space is inhomogeneous with respect to time. Moreover, no existing approaches for describing the time and/or space dependencies of a model were identified.
- There appeared to be few uses of hybrid modelling techniques for modelling space and this seemed to be an area worthy of further investigation.

These topics were pursued in the subsequent periods.

During the second period of the project the focus of work package 2 shifted from the capture of spatial information in models, to techniques to analyse such models, particularly those with heterogeneous spatial characteristics. The result of this survey was presented in an internal report IR2.1. The objective was to identify the most promising techniques for the systems considered in the QUANTICOL case studies, especially those which are complementary to modelling in Carma, the language under development in work package 4. To this end, some exploratory models were developed to investigate the extent to which techniques like spatial moment closures could usefully be deployed in QUANTICOL. The results were mixed with good approximations obtained in some cases but less good results in other cases, suggesting that further work was needed to characterise those models in which techniques such as spatial moment closures and pair approximation can be safely applied. This work was carried out under the auspices of Task 2.1.

In addition in Task 2.2, and supporting the design of Carma, we surveyed the existing language structures used to capture spatial aspects of behaviour in process algebras and related formalisms. In addition to serving the purposes of the QUANTICOL project, this comprehensive survey provided a valuable contribution to the field. It was reported in an internal report and is now in the final stages of preparation for submission for publication. It was a useful resource in work package 4 when work was undertaken to develop the space syntax of Carma and its specification language.

The final task of work package 2, Task 2.3, is concerned with parameter and model fitting from spatial data. Work continued within this task throughout the second and third period, investigating a number of different techniques for fitting the parameters governing the time to move from one location to another in models of both bus operations and bike sharing systems, working with real data available from system operators. These models have been used in case study models developed to solve specific planning problems, developed under the auspices of work package 5. The results of this task were reported in D2.2 in month 36 of the project.

The research in Task 2.1 into spatial analysis techniques proved to be more challenging than originally envisaged because the spatial issues in CAS we wish to analyse have a high level of mathematical complexity and additionally because existing techniques typically focus on assessing spatial heterogeneity at the global level rather than locally. This is not a good match to the usual objectives when modelling CAS, when local assessment is often important. For example it is important to quantify the availability of bikes at a particular bike station rather than focusing on aggregated measures of availability across all bike stations — average availability in the entire system will not help a user decide on whether it is feasible to make a journey between two given stations. Despite the difficulties a number of techniques for spatial analysis of CAS were developed and demonstrated in the final reporting period, although these were perhaps not as widely applicable as we initially hoped.

The work on the use of spatial data for spatial modelling was completed in month 36 of the project. Space plays an important role in QUANTICOL because of the nature of the case studies and we were able to report results on the use of spatial data both in model construction and in parameter estimation.

When considering public transport systems there is often a wealth of data because the service operators collect GPS data of vehicle locations. We have developed a methodology for taking such vehicle data and developing parameterised patch-based (discrete-space) models of vehicle routes. The methodology includes a phase of patch identification once automated map generation has been used to identify routes. GPS data can then be used to derive parameters for the time taken to cross a

patch. In particular we investigated the use of hyper-Erlang and shifted Erlang distributions for these dwell times. The models developed and parameterised with this approach have been used to assess the impact of speed limit reductions on timetable adherence and for statistical model checking of temporal properties, such as those describing punctuality.

Space can also be considered in a static manner, and we have characterised the spatial distribution of various bike sharing schemes. Furthermore, we have investigated which distributions can be used to generate topologies of bike stations that are closest to the actual arrangement for a number of cities. The approaches under consideration were the regular grid, the Poisson spatial point process, the Ginibre spatial point process, and a placement based on rating of amenities. Using street map data to exclude areas where stations cannot be placed, pixels are weighted as to their likelihood of being the position of a bike station. Two different approaches to the assessment of the coverage of the generated topologies when compared with the actual topologies have been presented. The rating-weighted scheme was able to create realistic patterns for several cities. Having such a technique for generating "artificial" bike sharing systems could be a valuable planning tool for town planners exploring the provision of a bike sharing scheme.

Journey data can also be used in model development and parameterisation. This data can be synthetic or real. An example of the use of real journey data from a bike-sharing scheme is in determining which bike stations have an influence on other stations, thus capturing the spatial layout and interactions of these stations. By considering only the stations with a high level of influence, the model can be reduced in size and thus its analysis becomes possible. As reported in Deliverable 2.2 synthetic journey data has been used in the transformation of an individual continuous-space model to a population discrete-space model using patches. Here, a description of movement was used to create a small simulation of individual movement, and from this parameters were estimated for the population model.

As mentioned previously, research into mathematical and computational methods for scalable spatial modelling for CAS was consolidated during the final period. In the initial investigation phase of QUANTICOL, we decided to focus our attention on discrete models of space, in which space is described by a set of locations, connected by edges representing the connections between locations. However, such a discrete representation of space poses formidable computational challenges, as complexity of model analysis scales at least linearly with the number of locations. We developed several techniques to abstract and transform space in order to ease the computational burden. These can be grouped into three classes of approach:

- **Pragmatic approximations**: a spatial model can be simplified using some heuristic criterion which is shown to work well in practice. Within this umbrella, we reported on a method based on behavioural distances to aggregate locations, thus simplifying the space structure. We also discussed approaches to simplify second order moment equations in the presence of localities, by setting to zero moments of population variables in locations which are too distant. In particular, we described an approach, illustrated on bike sharing, to identify the most important locations influencing a given one, and a method based on a syntactic distance between agents interacting in space.

- **Transformation of space**: here the spatial attributes of a model are kept, but space is transformed into a qualitatively different mathematical object, which makes the analysis simpler. In particular, we reported on techniques that take a grid describing a discretised 2-dimensional space, and approximate it with a continuous state space. The model itself changes from a patch-based population Markov chain to a set of partial differential equations, that can be solved more efficiently. Limit theorems were proved to establish the correctness of the method. The approach was also extended to a fast simulation result, allowing us to track the behaviour of a single agent, approximating it as a switched Brownian motion, which is faster to analyse than the full stochastic model. We also reported on a method to construct patches from an agent based model in which agents move in continuous space.

- **Full abstraction from space**: sometimes a spatial model can be replaced by a simpler one without an explicit representation of space. In this class, we developed three methods: the use of mean-field techniques that guarantee the asymptotic decoupling of the behaviour in different locations, allowing us to analyse each location independently; spatial moment closures, that can be used to capture mean and variance of the total population over all locations though a system of differential equations which is independent of localities; and pair approximation, a technique to capture local structure in spatially non-homogeneous models, by counting pairs of neighbour individuals. Figure 10 demonstrates how this approach is able to accurately approximate the behaviour of a full model, in a situation where mean-field approximation fails.

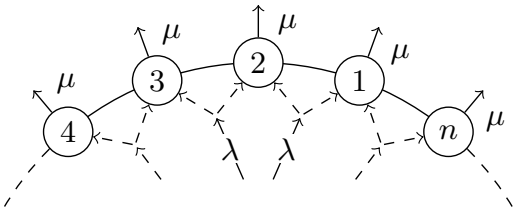

(a) Model: Jobs arrive at rate $N\lambda$ in the system. For each arrival, $i \in \{1 \dots n\}$ is picked at random and the job is allocated to the server i or $i+1$ (mod n) that has the least number of jobs. The size of each job is exponentially distributed of mean $1/\mu$.

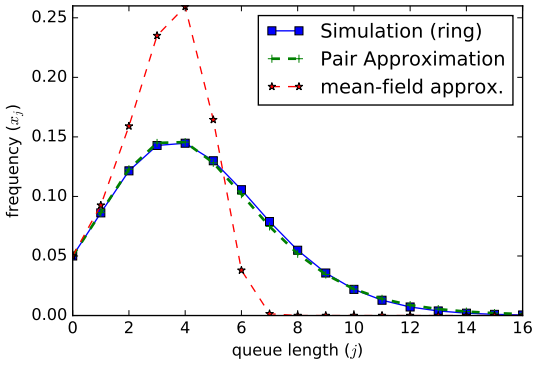

(b) Simulation results

Figure 10: Pair-approximation: Model and simulation results.

3.3 Work Package 3: Logic and Scalable Verification combining Micro and Macro Perspectives

Work package 3 was at the heart of the QUANTICOL project, building on the foundational work carried out within work packages 1 and 2, and underpinning the developments in work packages 4 and 5. Our goal was to develop novel, scalable formal analysis techniques and the underlying theories to support the design of CAS consisting of very large numbers of autonomous heterogeneous components. The kind of abstractions applied in the mean-field approach provided the key to overcoming long standing scalability problems in otherwise very successful formal verification approaches such as stochastic model-checking. A further goal was to develop logic based analysis techniques for spatial and spatio-temporal aspects of systems.

The first main objective was to develop both a stochastic temporal logic and novel forms of scalable model-checking approaches, exploiting mean-field/fluid-flow approximations, and to develop spatial logics and related model-checking techniques. In practice a number of different logics and model-checking algorithms were developed, reflecting the reality that different instances of CAS have different characteristics which must be captured and scrutinised, and no one approach is suitable for all cases.

The second objective concerned the use of abstractions. Mean-field approximations are ultimately based on numerically solving sets of ODE. The computational complexity to do so increases with the number of equations that are needed to model the system. It is therefore of utmost importance to keep this number as low as possible, in particular when the equations are derived in an automatic way from higher level process algebraic specifications. Thus we have focused on developing appropriate forms of abstraction with the aim to reduce models while preserving certain behavioural relations. Furthermore, we have also investigated other forms of abstractions at the level of CTMC.

The third objective of this work package has been the study of the relation (in terms of properties) between local and global views of a system by using established techniques from variability analysis, as applied in fields like product family engineering. This has involved the study of relations between

(representations of) small populations and a compact (family) representation of a large population 'built' from these smaller populations, by indicating the commonalities and variability of single entities in their overall environment. We have investigated both the conditions under which properties can be preserved top-down as well as bottom-up, i.e. from local to global and vice versa. This involved extending state-of-the-art variability analysis techniques to encompass quantitative behaviour.

Work in the first year of work package 3 mostly focused on the development of novel scalable approaches to the verification of CAS based on mean-field and fluid techniques, corresponding to the main objective of the first phase of Task 3.1. We had notable success in this area extending previous results, and developing novel model checking techniques which are highly scalable. In addition we also investigated spatial logics, which allow spatial aspects of dynamic systems to be the focus of verification. Model checking is a powerful technique for the automatic verification of concurrent and distributed systems, so developing efficient and scalable forms of model checking is central to the QUANTICOL framework for CAS.

The use of *fluid* model checking was first proposed by project partners in the context of properties of single agents within a large population expressed in a subset of the Continuous Stochastic Logic (CSL). During the first year of the project the technique was extended to the whole time-bounded fragment of CSL, and the steady state operator when appropriate. In addition a novel approach was developed for verifying collective properties expressed in terms of one clock deterministic timed automata. In this work a linear noise approximation (or central limit approximation) is used to verify properties at the global system level. For each individual, an agent model is combined with an automaton expressing the property and the population model is formed by counting pairs of state-property configurations. The evolution of the counting process is given a Gaussian approximation from which the probabilities of interest can be readily computed. This work is in the continuous time population CTMC setting.

Related to this work on verification of specifications of continuous time models, was some further work on the stability of such models. Stability analysis provides important information about the predictability of the dynamics of systems and their sensitivity to parameter values. A new numerical method, in the form of a tool-chain, was proposed for stability analysis starting from a stochastic process algebra specification of agent coordination in a collective dynamic system.

Additionally, we developed a new verification method for CAS in the DTMC setting, based on an on-the-fly model checking of DTMC models against probabilistic logic (PCTL) formulas. Again consideration is limited to *bounded* PCTL formulae expressing properties of a single agent (or small set of agents) in a large population combined with a preliminary formulation of *global* properties. The developed on-the-fly model checking algorithm is parametric with respect to the semantic interpretation that is placed on the model, meaning that both an exact probabilistic and approximate mean-field semantics can be used, taking advantage of fast simulation results in the latter case. The approach was supported by a prototype tool, FlyFast, which has been further elaborated in the later periods of the project. The use of the approximate model checking technique was demonstrated on the QUANTICOL bike sharing example. In the discrete setting, we also worked on estimating the approximation error incurred by the mean-field approach for some classes of models.

We also considered statistical model checking within the suite of approaches for our framework. In this context we developed a novel efficient technique and tool to deal with uncertainty in the values of model parameters in a statistically sound way. The approach was based also on recent advances in machine learning and pattern recognition.

As experience of spatial aspects of models, particularly with respect to model checking had been limited in previous work, we also conducted an extensive literature study on spatial logics during the first period of the work package. Concluding that there was a good fit between the requirements for modelling our CAS case studies and the abstract topological and closure spaces considered in Task 2.1 on spatial representations in work package 2.

Initial work on tasks 3.2 and 3.3 was also undertaken during the first period. In particular work started on a theory of aggregation of ODEs induced by an equivalence relation over the local states of

a process algebra model that captures symmetries in the fluid semantics according to the well-known notion of exact lumpability for ODEs. During the second period significant results were obtained on the reduction and comparison of systems of ODEs based on novel behavioural equivalence relations. A number of such relations were designed and demonstrated to achieve significant reductions on some models. This work has continued to be refined and elaborated throughout the project to its current state where there are a number of outstanding publications and a tool implementation, ERODE, which supports a very general form of the ODE reduction technique. However, mean-field and fluid approximation techniques are not suitable for the analysis of *all* large scale CAS. This is for example the case in systems that involve uncertain parameter values, which occurs frequently in performance analysis. During the third period exploiting a technique involving simulation-based model checking has been shown to be a promising solution in such cases. This technique, also known as *smoothed model checking*, exploits properties of Gaussian processes to verify properties of models with several parameters. Furthermore, during the second period we also established a method to identify and remove populations from discrete representations of CAS models when those populations have no significant impact on a measure of interest during simulation-based analysis and a method which automatically derives ODEs to capture the moments of population CTMCs and carries out moment-closure analysis to reduce the number of ODEs which must be considered.

As exploratory work towards the relation between local and global system views and variability analysis in the context of software product lines, an initial feature model of a family of bike-sharing systems was developed in the first period. It was annotated with attributes and global quantitative constraints aiming to minimise the total cost of a chosen configuration and at the same time maximise customer satisfaction and capacity of docking stations. This proved to be particularly useful when the employment of a bike sharing system is considered in a particular city with its own geography and usage patterns.

During the second period in work package 3 we worked on a spatial logic with basic spatial operators that is general enough to be applied in the diverse spatial settings that were identified in the context of work package 2 in the first period. Continuing to work in the theoretical framework of Closure Spaces (a generalisation of Topological Spaces) which provides a set of useful basic abstract spatial operators (closure, interior, boundary and many derived ones) that provide a structured way to define higher level spatial logical operators. Moreover, we found that they are also suitable for the development of efficient spatial and (branching time) spatio-temporal model-checking algorithms in which these same closure space-based operators play a role as well. Furthermore, metrics and distance functions were added in an orthogonal way providing further spatial richness. The initially theoretical approach proved useful in applications such as the combination of the closure space-based model checking algorithms with metric interval temporal logic and monitoring algorithms for a linear time spatio-temporal logic to verify interesting spatial-temporal properties of reaction-diffusion models, such as their robustness to perturbations (see Figs. 11 and 12). The logic has qualitative and quantitative semantics, and monitoring algorithms have been designed and implemented. Other developed applications of spatio-temporal model checking were the analysis of emergent spatio-temporal phenomena, such as the formation of spatial clusters of full stations in bike sharing systems and the phenomenon of clumping in public urban bus transportation systems. Prototypes of both the linear time and branching time versions of spatio-temporal model checkers have been developed over the course of the final two periods.

During the second period, we also improved and extended the techniques for scalable verification of temporal properties exploiting stochastic approximation, particularly for a class of global properties and for timed properties of individual agents expressed by deterministic timed automata, with a single clock. We also continued the investigation of efficient statistical model checking approaches for systems with uncertainty of parameters, leveraging machine learning ideas. In particular, we extended these approaches to spatio-temporal properties. These techniques have been implemented and tested on a number of case studies.

Work in Task 3.3 made significant achievements during period 2 of the project. We developed a

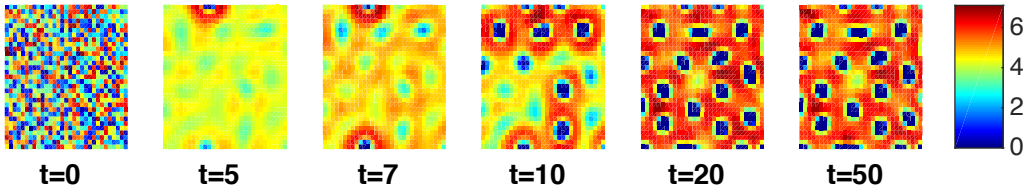

Figure 11: Simulation of a Turing reaction-diffusion system modelling a process of morphogenesis underlying pattern formation in animal fur as a result of two interacting chemical substances A and B. The initial condition has been set randomly. The colour map for the concentration of the substances is specified in the legend on the right.

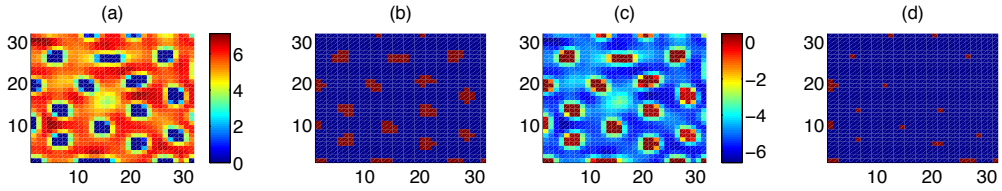

Figure 12: Automatic identification of spatial patterns using spatio-temporal logics. Areas with a high level of substance A surrounded by those with a low level of A and of a certain diameter, that persist over a certain time interval. Snapshot of a simulation (a), related spatial analysis result of the property (red points are those that satisfy the property) (b), robustness of the validity of the property (c), similar property identifying smaller and more round spots.

number of extensions to software product line engineering (SPLE) approaches to make them suitable for CAS. In particular, family-based verification of behavioural aspects of CAS were pursued. In family-based analysis, the system model covers both behaviour that is common to all products of the family as well as variation points used to differentiate among the individual products that can be derived from the family. This way, logical properties can be analysed at the family level using variability knowledge to deduce results for products, rather than having to generate and analyse each single product, which is often very costly and does not scale well. Some of the techniques were applied to the project's bike sharing case study. Fig. 13 illustrates the Pareto front for variants of a bike sharing system taking quantitative aspects into consideration. Within Task 3.3 an SPLE modelling and analysis framework was developed and implemented in the Variability Model Checker VMC that provides family-based verification of qualitative state and action based properties in an on-the-fly fashion. Furthermore, an alternative approach was presented to reason about software product lines focussing on behavioural relations. For this purpose, the Variant Process Algebra VPA that has a family-based semantics in which variants can be explicitly labelled was developed. Several quantitative extensions of variability analysis were developed to handle variability in software performance models. Finally, a first proof-of-concept for feature-oriented modular verification was developed and, for a restricted notion of coherent branching feature bisimulation, a minimisation algorithm was developed and its correctness shown.

During the final period further work has been done on the family-based approach to model checking, achieving good results. In particular, an Eclipse-based tool for the quantitative feature-oriented language QFlan, introduced during the second reporting period, has been developed, using statistical model checking, to address properties such as quality of service, reliability, or performance of dynamically reconfigurable product lines.

As explained above, throughout the project a common objective of all three tasks within work package 3 has been to make substantial contributions to the theoretical foundations of scalable and spatial formal analysis methods in order to underpin the development of a formal verification framework for Collective Adaptive Systems (CAS). This theoretical work has been expanded in the final

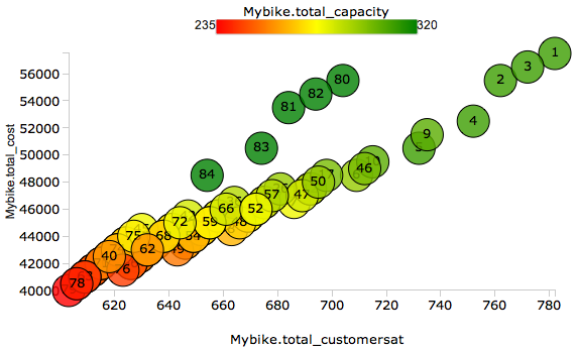

Figure 13: Bubble graphs of Pareto fronts for a Bike Sharing System showing the costs of variants involving a range of bikes between 235 and 320 taking other (not visualised) aspects into consideration such as expected customer satisfaction based on model based performance measures.

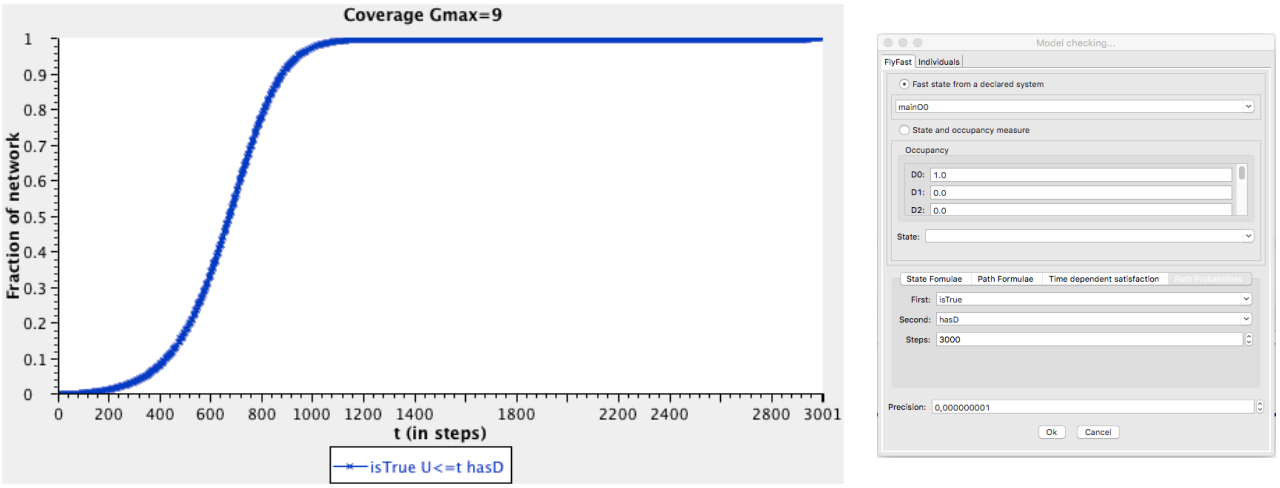

Figure 14: `FlyFast` model checker showing mean field model checking results for network coverage (i.e. the fraction of nodes that have seen a newly inserted data-element D within time bound t) for a gossip protocol with 2500 nodes for time bounds t ranging from 0 to 3000 time units. The 3000 model-checking runs are performed in a total time of only several seconds due to an on-the-fly mean field approach and an efficient use of memoization

reporting period and consolidated with tool development and consideration of applications, including some outside the domain of CAS.

Specifically, we extended fluid model checking and on-the-fly mean-field model checking previously developed for the analysis of *local* reachability properties of an individual object in the context of one or more large populations, to consider fluid model checking to address time-bounded *global* reachability properties of the system, for example, properties concerning the probability that the size of a particular population exceeds certain levels within a certain time interval. Moreover, we have also shown how fluid model checking can be enhanced to deal with various types of *reward* properties. This allows more sophisticated queries to be used to interrogate the behaviour of models.

The discrete time, probabilistic mean field model checking approach, that led to the open source model checking tool FlyFast, has been further improved and used for a number of case studies, such as the analysis of a benchmark gossip protocol (see Figure 14) as well as a bike sharing model. In this latter case a CaSL model is reduced to a FlyFast model, linking the work in work package 4 with work in work package 3 and demonstrating the overall vision for how the theoretical foundations established in this work package can be incorporated into the QUANTICOL framework. The bike sharing example also illustrates that, under some suitable conditions, the discrete time approach can be used to approximate fluid model checking results. As a further step towards the integration of FlyFast with specification languages based on a predicate-based communication paradigm, a front-end for the FlyFast modelling language has been developed that incorporates components and predicate-based communication inspired by the CARMA language.

In the previous reporting periods several spatial and spatio-temporal logics have been proposed and efficient related model checking and monitoring algorithms have been developed. In the third reporting period this work has proceeded by developing novel spatial logic operators on one hand, and by the development of more efficient prototypes. Some of these, such as the `jSSTL` tool, have been made available as an Eclipse plug-in and their front-end made compatible with the CARMA language. Others, in particular `topochecker`, have been developed as a stand-alone model checker and have been shown to be applicable not only to CAS but also to unforeseen areas such as medical imaging (see Figure 15). Spatio-temporal model checking has also been combined with statistical model checking exploiting a feature of the MultiVeStA tool and a simulation tool for bike sharing systems, to analyse a property on all points in space simultaneously for each simulation run, making it feasible to apply stochastic spatio-temporal model checking on large CAS such as a bike sharing system of the size of that of London. Furthermore MultiVeStA itself has been extended to incorporate checking of steady state properties in addition to transient properties.

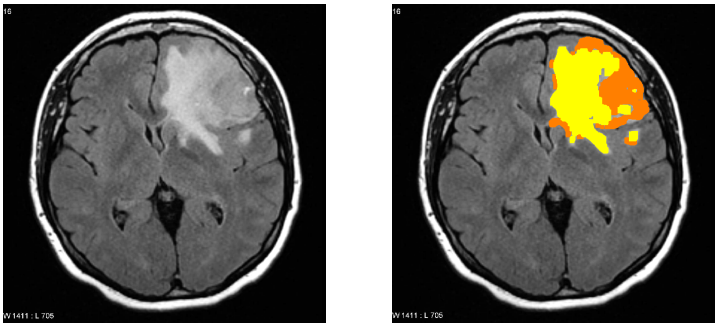

Figure 15: (left) A slice of a FLAIR MR acquisition of a brain affected by a glioblastoma. (right) Final result of application of threshold, distance, and the texture operator. The yellow area is the identified oedema; the orange area is the tumor (case courtesy of A.Prof Frank Gaillard, Radiopaedia.org, rID: 5292).

3.4 Work Package 4: Language and Design Methodology

The main objective of this work package was to support the development of CAS in a specifically designed language called CARMA[1]. This language is the extension and evolution of the SCEL language under the auspices of another FET project called ASCENS[2]. The language was concerned with issues of collective adaptation in the presence of a huge number of components that can be spatially distributed. It was to incorporate linguistic abstractions for modelling and manipulating knowledge, behaviour, aggregation, locations and interactions. The intention was to support the description of individual components, their location and their interactions but also to model the global behaviour. The language has its basis in process algebra and supports the development of a comprehensive framework for the development, verification and analysis of CAS.

During the first period of the project the focus of work package 4 was on the exploratory work to support the development of the modelling language. In particular we considered design principles for the language and worked on the identification of language *primitives* able to address the complex requirements of developing a language rich enough to capture essential features of CAS and yet remain amenable to scalable analysis techniques. To identify these abstractions and linguistic primitives, we relied on various formalisms that QUANTICOL partners had previously developed and experimented with them to model simple CAS. At the end of this work a general consensus was reached in the project that, to be effective, any language for CAS should provide:

- Separation of knowledge and behaviour;
- Control over abstraction levels;
- Bottom-up design;
- Mechanisms to take into account the environment;
- Support for both global and local views; and
- Automatic derivation of the underlying mathematical model.

This exploratory work has considerably deepened our appreciation of the trade-offs between language expressiveness and scalable analysis and the choices involved in capturing adaptation and an environment. This provided a firm foundation for the development of the QUANTICOL modelling language. Moreover, the exploratory formalisms are each interesting in their own right and have spawned several publications.

Starting from the identified requirements, during the second reporting period we designed a new language to support the specification and analysis of CAS, with the particular objective of supporting quantitive evaluation and verification. This language was named CARMA, Collective Adaptive Resource-sharing Markovian Agents.

CARMA combines the lessons we learnt from other stochastic process algebras such as PEPA, EMPA, MTIPP and MoDEST, with those learnt from languages specifically designed to model CAS, such as SCEL, the AbC calculus, PALOMA, and the Attributed Pi calculus, which feature attribute-based communication and explicit representation of locations.

A distinctive contribution of the language CARMA is the rich set of communication primitives that are offered. This new language supports both unicast and broadcast communication. This richness is important to enable the spatially distributed nature of CAS, where agents may have only local awareness of the system, yet the design objectives and adaptation goals are often expressed in terms of global behaviour. Representing these rich patterns of communication in classical process algebras or traditional stochastic process algebras would be difficult, and would require the introduction of

[1] Originally called CAS-SCEL in the proposal.

[2] `http://www.ascens-ist.eu`

additional model components to represent buffers, queues, and other communication structures. Another feature of CARMA is the explicit representation of the environment in which processes interact, allowing rapid testing of a system under different open world scenarios. The environment in CARMA models can evolve at runtime, due to the feedback from the system, and it further modulates the interaction between components, by shaping rates and interaction probabilities.

At the end of the second period Deliverable D4.2 presented the description of CARMA and its operational semantics together with a set of tools supporting analysis of CARMA models. In this toolset was included a Java library for simulating CARMA models and an Eclipse plug-in for supporting specification and analysis of CAS in CARMA. In this plug-in, CARMA systems are specified using an appropriate high-level language for designers of CAS, named the CARMA *Specification Language*. This is mapped to the process algebra, and hence will enable qualitative and quantitive analysis of CAS during system development by enabling a design workflow and analysis pathway. CaSL was not introduced to add to the expressiveness of CARMA, which we believe to be well-suited to capturing the behaviour of CAS, but rather to ease the task of modelling for users who are unfamiliar with process algebra and similar formal notations.

CARMA is currently equipped with a structured operational semantics in the FuTS style, which allows compact representation of the complex rules needed to fully specify the behaviour of a CARMA model. This semantics gives rise to an underlying continuous time Markov chain (CTMC) and evaluation of models can be carried out by stochastic simulation. An initial version of these semantics were developed during the second reporting period and they were subsequently simplified during the third period as part of the work on developing a population level semantics and mean-field approximation based on ordinary differential equations.

The first version of the software tools necessary to support programming of CAS were released during the second period of the project, but they have undergone continual improvement and refinement during the final period.

During the final period CARMA was regarded as stable, and the focus was on improving the accessibility of the language through a variety of mechanisms, including a programmatic specification language and a graphic language for specifying spatial aspects of models. Thus modelling in CARMA is now supported by CaSL, the CARMA specification language, providing a less mathematically-oriented, programming-style language. Developing this specification language is intended to make modelling with CARMA accessible to a wide audience of potential users interested in CAS, not just those already familiar with formal modelling with process algebras. We have also developed exemplar models that are available to potential users, and extended the suite of tools to offer a modeller different approaches to model analysis.

Space plays a key role in many CAS and we have revisited the support that is offered to faithfully capture the spatial aspects within a model, resulting in improved syntax to assist the modeller and a graphical front end which can be used to automatically generate the spatial aspects of models. Throughout this development process we have informed our design and improved our implementation by developing a number of different exemplar models, even going beyond the smart cities application domain in order to ensure that our languages are sufficiently expressive to capture a wide range of systems. Within CaSL the spatial aspect of a model is seen as underpinning (and constraining) the behaviour of the components, in a similar way to the role of the environment, as depicted in Figure 16.

Of course, to be practically useful a modelling language must be implemented in a robust set of software tools to allow the modeller to construct and analyse the model with confidence. In addition to the editing and model simulation facilities previously presented in the Eclipse plug-in tool, the CARMA tool suite now also incorporates a graphical plug-in to assist in the specification of the spatial aspects of models (see Figure 17), links to the MultiVeStA tool for statistical model checking, and a command line interface with support for experimentation. The software tool suite for CARMA is available at `https://quanticol.github.io`. To ease the access to CaSL and its tool, a specific web site has been instantiated where users can access tool documentation and examples. Moreover, a bug reporting system is also available.

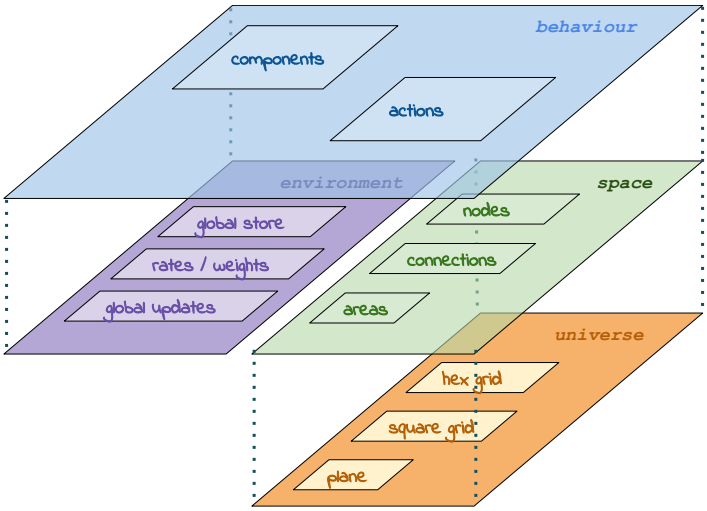

Figure 16: The layers of CaSL

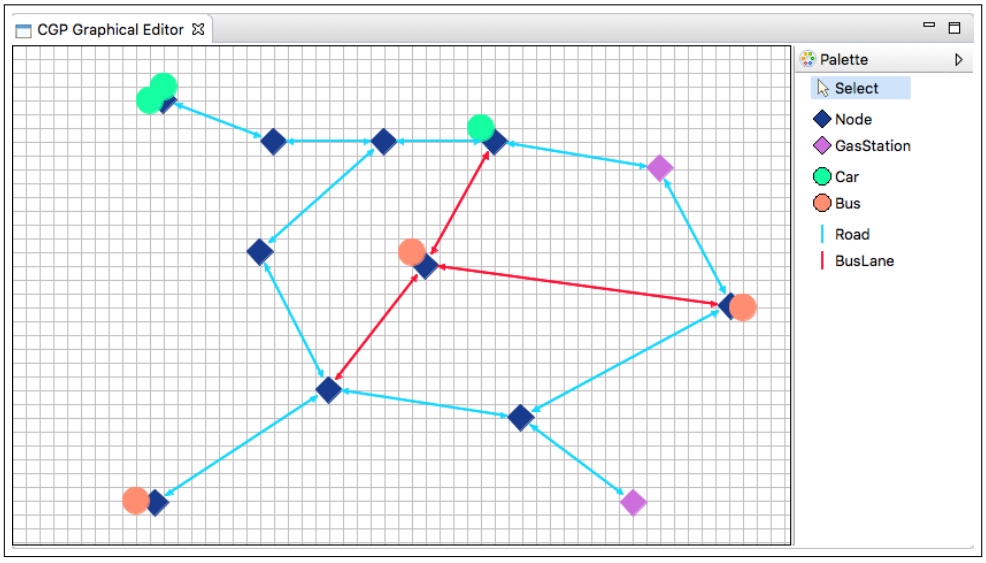

Figure 17: A screenshot of the graphical interface for path and components layout.

We have also elucidated a design workflow and analysis pathway, which is supported in the tools. This takes into account the different phases that a model goes through, from initial design, elaboration, parameterisation and then use as a tool to investigate the behaviour of the system under study. The modeller needs different support at each of these stages and we have sought to provide what is appropriate for each stage, as far as is feasible within the limited time and resource of the project. Building within Eclipse has allowed us to provide many consistency-checking features which greatly enhance the support for the modeller. These internal checks seek to ensure that CaSL models are free from the type of minor error that can be frustrating and time-consuming during model development. However, once a modeller is fully confident of their model, a graphical user interface can become cumbersome and inconvenient. Thus we also provide a command line interface to support efficient exploitation of models under different experimental frames. Moreover, the results of model analysis are automatically enhanced with metadata to assist with their interpretation and reproducibility.

Whilst we have explored other applications ranging from epidemiological models to food security models, our primary focus for applications and case studies has remained smart cities, and in this context we have developed a demonstration of the analysis of CaSL models of two of the scenarios from our smart city case studies (see Deliverable 4.3). Specifically we consider a mesoscale model of buses within a city, particularly paying attention to the congestion that occurs when multiple routes share the same bus stops, and issues related to regulatory compliance and appropriate spacing on frequent bus services. In the second example, we consider the key issue related to user satisfaction within urban bike sharing systems — whether a user will find a bike or a slot at a convenient location when they want one.

To study the expressive power of attributed based communication, a simple calculus named *AbC* (a calculus for Attribute-Based Communication) has been introduced. The expressiveness and effectiveness of attribute based communication are demonstrated both in terms of modeling scenarios featuring collaboration, reconfiguration, and adaptation and of the possibility of encoding channel-based interactions and other interaction patterns. Behavioral equivalences for *AbC* are introduced for establishing formal relationships between different descriptions of the same system. Finally, a Java run-time environment, named $Ab^{a}C$uS, has been developed to support deployment of systems based on attribute-based communication. This runtime environment enables programming of collective adaptive systems by relying on the communication primitives of the *AbC*/Carma. By means of a number of examples, we also show how opportunistic behaviors, achieved by run-time attribute updates, can be exploited to express different communication and interaction patterns and to program challenging case studies.

3.5 Work Package 5: Models Validation and Tool Support

The objective for work package 5 was to provide a development environment hosting the implementation of the modelling language at the core of QUANTICOL (Carma) as well as a wide range of applications and libraries for the analysis and verification of Carma and associated models. Smart city case studies were also envisaged to be core part of the work package, providing a testbed for the techniques established in other work packages, especially work package 3, and ultimately demonstrating the success of the developed framework.

The activities in work package 5 during the first period of the project were primarily concerned with initial data-gathering activities regarding the case studies on smart public transportation (bus networks and bike sharing systems) and smart grids. Through this work we have been able to establish a critical validation of the quality of the information that is publicly available, and an analysis of the real data offered to us by project collaborators.

During the first period we developed a particularly fruitful interaction with Lothian Buses, an Edinburgh-based company which operates a large bus network. We have obtained large sample data sets related to bus network structure and measurements related to the City of Edinburgh both through data available to the general public and confidential data from Lothian Buses. This data from the

Figure 18: Screenshot from the live map at `https://web.barclayscyclehire.tfl.gov.uk/maps` providing live bike availability at London's Waterloo station.

operator offered high-quality information, consisting of detailed records of GPS locations, bus speeds, timestamps, and bus identifiers. With this data we were able to carry out exploratory modelling studies (resulting in two papers accepted for publication) further cementing our relationship with Lothian Buses and establishing that the quantity and quality of the data were sufficient to calibrate and validate certain spatial models of bus networks.

For bike-sharing systems we focussed on publicly available data. Most systems offer web access to live information on the availability of bikes at parking stations (see Figure 18 for an example). These websites can be systematically queried to reconstruct real availability traces at any desired granularity. Some operators also provided historical traces consisting of journey details, including start/end dates and source/destination stations. Using the publicly available datasets for the bike sharing system of the City of London as a prototypical example, our data analysis confirmed that these measurements could be used for spatial models with an explicit representation of the network topology. Models of smart grids can benefit from the availability of measurements and forecasts at the transmission level, offering aggregated information about energy production, consumption, and market prices at the national electricity-network scale. Load measurements to calibrate models at a smaller scale (i.e., neighbourhood or building level) are also available in the research community. These are also accompanied by network benchmarks consisting of prototypical layouts that have been extensively studied in the literature. Overall, the available data was sufficient for the QUANTICOL purposes, as demonstrated by a number of related papers published already in the first period and suggested that there would be no difficulties in conducting cast studies throughout the project.

Another aspect of the work of this work package during the first period was concerned with identifying use-case requirements for building the QUANTICOL tool integration platform. An initial requirements-elicitation activity highlighted a predominance of Java as the platform of choice for tool development by the project partners. This choice, however could still accommodate non-Java contributions, due to specific tool requirements or to the partners' expertise. Such an integration is possible, for instance, by exploiting a number of already available language bridging mechanisms such as Java Native Interfaces or Matlab's javabuilder. The emergence of a consensus around Java as the platform of reference for tool development within the project gave a solid basis on which to plan our future work on tools and their integration.

During the second period work package 5 focused on the development of case studies of smart cities in public transportation and smart grid, assisted by a development environment hosting the implementation of CARMA, our reference modelling language, as well as "satellite" software tools that support further research results developed in the project. The work package is divided into three main tasks: Task 5.1 is concerned with smart public transportation systems; Task 5.2 deals with smart grids; Task 5.3 is concerned with tools integration. Tasks 5.1 and 5.2 are divided into two

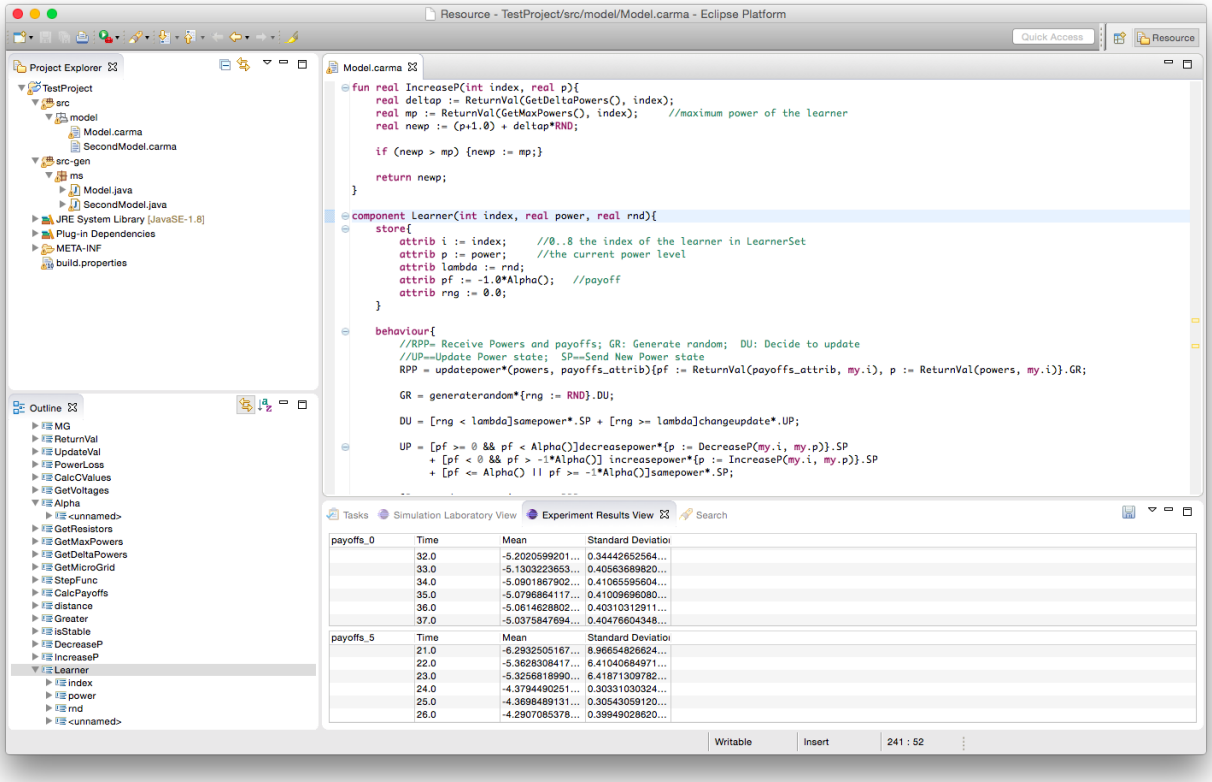

Figure 19: A screen-shot of the CARMA Eclipse plug-in.

subtasks, *a* and *b*, which focus on problem scoping/data gathering and model building and method validation, respectively. All these tasks were active during the second reporting period of the project. and progressed well.

In Task 5.1.a work continued on problem scoping and data gathering activities as we further evaluated other sources of public-transport data, most notably those of the bike sharing system of London. The use of Paris' bike sharing system data, to which we committed in Task 5.1.b, was instrumental for starting a promising line of research on predictive models of availability in bike sharing systems. We continued to tackle questions related to smart bus transportation systems, leading to a number of publications.

The work on smart grids (in Tasks 5.2.a and 5.2.b) also progressed well, with a slight shift in emphasis due to a partner's change of affiliation from EPFL to INRIA, leading to less ready access to significant smart-grid expertise. This was counterbalanced by effort on the part of other partners to tackle smart-grid challenges. These efforts led to unplanned contributions from the field of dynamic learning and coalitional game theory for power trading and coordination in smart grids.

During the second period a significant amount of work was devoted to Task 5.3 on tools integration. This work was centred on establishing software support for the CARMA process algebra and specification language, developed in work package 4. This resulted in a major outcome, the CARMA Eclipse plug-in, shown in Figure 19. Developing within the Eclipse framework was motivated by the possibility offered by its plug-in system to maintain loose coupling between components, easing the development and maintenance by different partners, and to reuse powerful available plug-ins.

In addition, all project partners have been very active throughout the project in developing software tools of different degrees of maturity in order to support further research related to the project, but not directly (at least in the original intentions) to CARMA. As such they are currently standalone, either due to their exploratory nature or because they are linked to other software.

The final period saw substantial activity on software tool development as we sought to consolidate the results of other work packages in a form which is useful both within the project, and to other researchers. These software tools, the QUANTICOL software tool suite, form a valuable dissemination channel to a variety of communities. Several of the tools developed, and their application to smart city scenarios, were reported in Deliverable 5.3 of the project.

Specifically, the presented tools are:

- FlyFast, a first-of-its-kind, *on-the-fly mean-field* probabilistic model checker for bounded PCTL (Probabilistic Computation Tree Logic) properties of a *selected individual* in the context of systems that consist of a *large number* of independent, *interacting objects*. The underlying on-the-fly mean field model checking algorithm was developed and proven correct in the context of work package 3, task 3.1. FlyFast is provided within the jSAM (java StochAstic Model Checker) framework which is an open source Eclipse plug-in[3] integrating a set of tools for stochastic analysis of concurrent and distributed systems specified using process algebras.

- *jSSTL*, a Java tool for the specification and the verification of Signal Spatio-Temporal Logic (SSTL) properties also developed in work package 3, task 3.1. It consists of a library (the *jSSTL* API) and a front-end developed as an Eclipse plug-in. The plug-in provides a user friendly interface to the tool, whereas the library can be used to integrate *jSSTL* within other applications and tools. The objective is to explore a question pertaining to bike sharing systems: *if I do not find a bike/free slot, is there another station at the distance less than a certain value where I can find a bike/free slot?* A station ℓ satisfies ϕ_1 if and only if it is always true that, between 0 and T_{end} minutes, there exists a station at a distance less than or equal to d, where there is at least one bike and a station at a distance less than or equal to d where there is at least one free slot.

- *ERODE*, a tool for the evaluation and reduction of ordinary differential equations implemented as an Eclipse plug-in, and consolidating work carried out in work package 3, task 3.2, its architecture is shown in Figure 21.

- *UTOPIC*, which supports an under-approximation technique for the reachability analysis of nonlinear systems of ordinary differential equations (ODEs). It implements an algorithm based on control-theoretic principles of optimal control, developed in the context of WP1. *UTOPIC* itself is implemented as an extension of *ERODE*.

- *topochecker*, a spatio-temporal model checker based on closure spaces and Kripke frames. Currently it checks a spatial extension of Computation Tree Logic named STLCS (Spatio-Temporal Logic for Closure Spaces). The underlying theory has been developed in the context of work package 3.

[3] http://quanticol.github.io/jSAM/

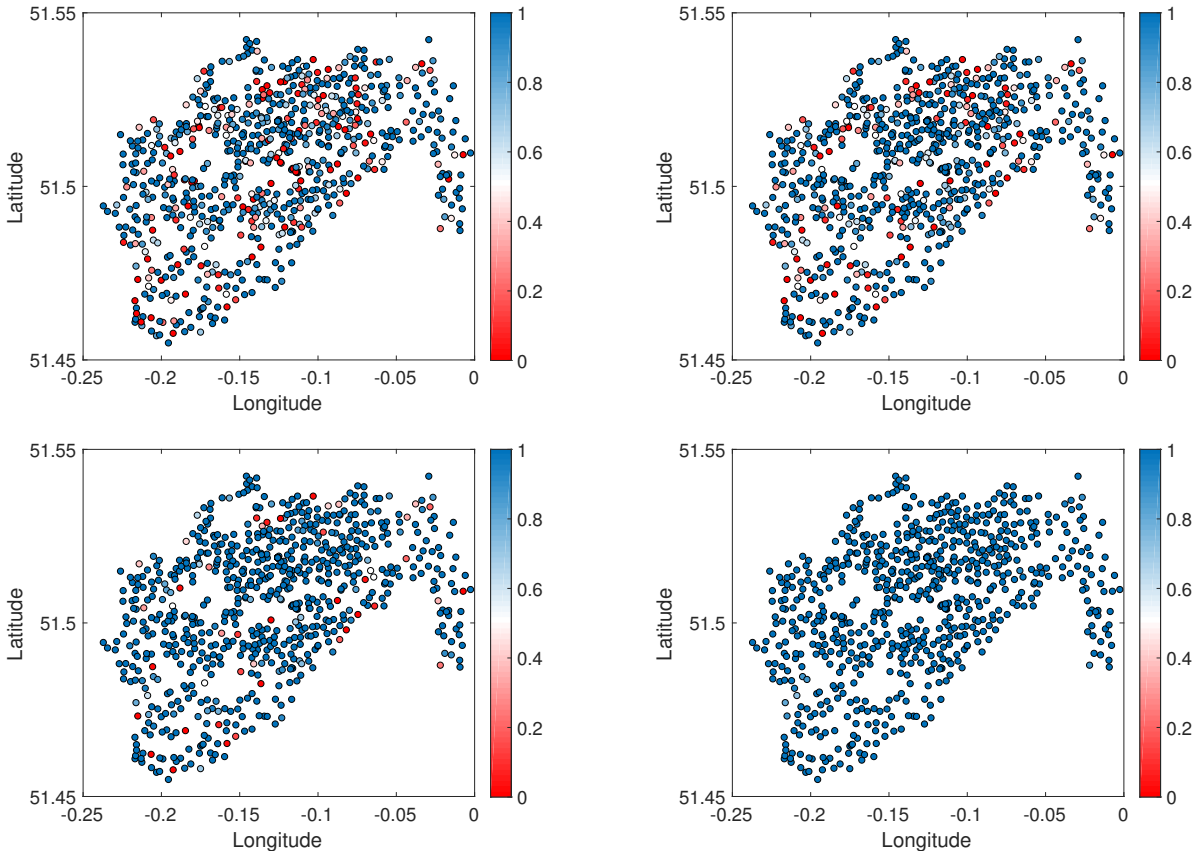

Figure 20: Approximate probability satisfaction degree of formula ϕ_1 for 1000 runs for each BSS station for (a) $d = 0$, (b) $d = 0.2$, (c) $d = 0.3$ and (d) $d = 0.5$. The value of the degree is given by the colour legend.

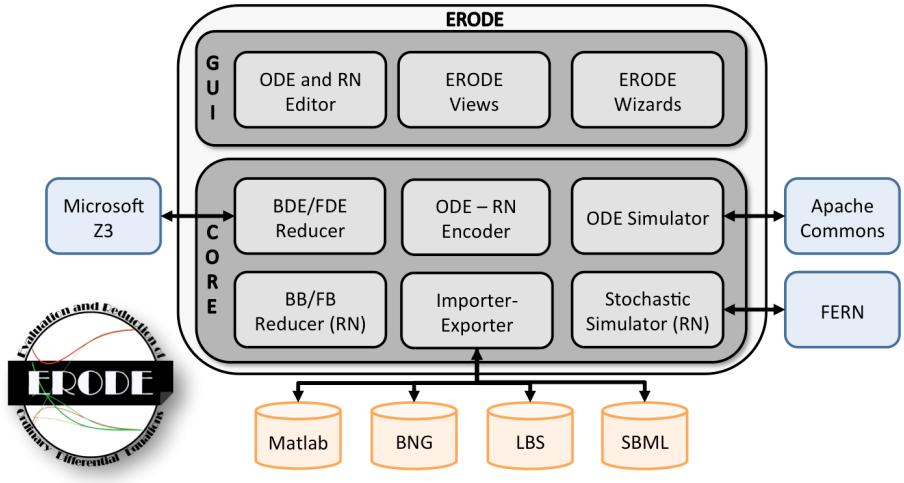

Figure 21: *ERODE*'s Architecture.

4 Impact of the Project

Our long term research vision is the development of a comprehensive software engineering environment, supporting a model-based design methodology for the development of CAS for smart city applications taking non-functional properties into account. The QUANTICOL project has contributed substantially to this vision, establishing a formal quantitative modelling framework. Our original vision of the potential impact of the project is shown in Figure 22.

4.1 Potential Impact

In our vision we aimed to reach three distinct communities relating to science, technology and society. The contributions that we have made with respect to the scientific and technology communities have been described in the previous section; here we consider only the efforts that we have invested to ensure that our contributions are visible and to encourage their take-up as broadly as possible. Furthermore, our original plan included efforts to have an impact more widely on society, both in socio-economic and societal terms.

In the following we outline the transformational impacts we originally we aimed to produce and give a brief assessment of our success in each area:

- **Transformational impact on researchers working on quantified analysis techniques.** Our objective was to bring foundational developments in the mathematical analysis of CAS together and enshrine them into a modelling methodology for widespread use. We have been moderately successful in achieving this, although we have not yet realised a fully integrated tool which encompasses all the methodological advances that we have made during the project. Nevertheless we have made a substantial first step in this direction. The QUANTICOL tool suite allows the scope of the mathematical results to be explored and exploited.

- **Transformational impact on researchers working on formal methods.** We believed that there was an opportunity to have significant impact on the formal methods community through the alternative approach to state space explosion offered by our scalable modelling techniques. We have certainly generated a lot of interest in this community as witnessed by the numerous invited presentations that have been given by member of the QUANTICOL team. Moreover we are starting to see other groups adopting mean-field based approaches to tackle a wide class of large scale distributed systems.

- **Transformational impact on researchers and developers of CAS.** The design methodology for CAS developed in QUANTICOL was intended to be our major route to creating transformational impact on researchers and developers of CAS. Whilst there has been interest from and engagement with stakeholder groups within CAS, we have not yet achieved independent use of our methodology by researchers within that community. However this is unsurprising given that many of the tools are still under development and case study results are only just being published. Nevertheless there has been sufficient interest from some CAS developers to generate follow-on projects exploring the use further (see INRIA/Schneider Electric project described below).

- **Transformational impact on users of ICT systems based on CAS.** Our major opportunity of socio-economic and societal impact was through the uses of systems influenced and improved by our methodology. CAS underlie many of the socio-technical systems being envisaged today for smart cities and beyond. By their very nature, these systems are pervasive and often largely transparent to the user. The crucial role that they will play in people's daily lives combined with this near invisibility makes it particularly vital that their functional and non-functional correctness should be strongly established before they are deployed. The impact

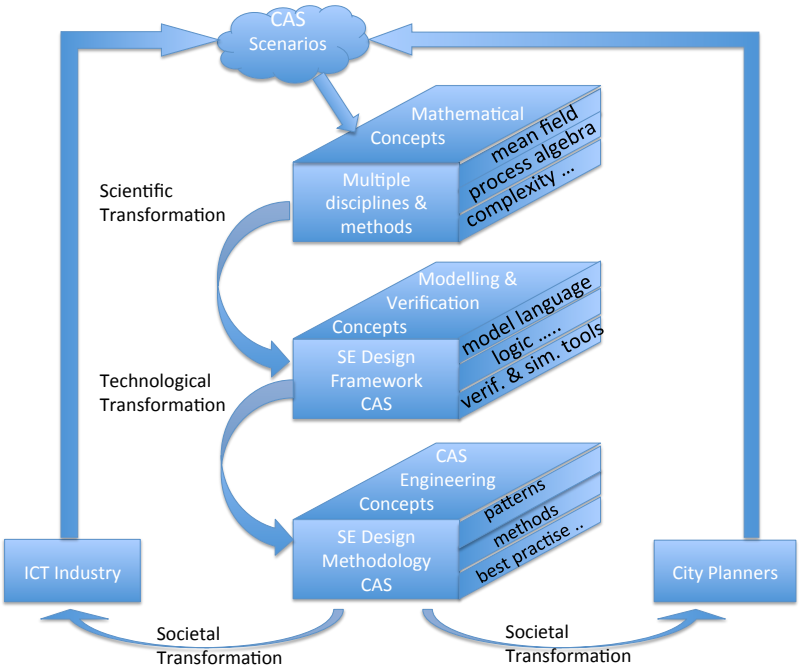

Figure 22: Our vision of the potential impact of QUANTICOL in the project proposal

of moving towards large scale decentralised interactive and adaptive systems cannot be overestimated even on a local city based scale, not to mention when seen from a European or global perspective. Though they hold the promise of advantages in cost-effectiveness and organisational flexibility, decentralised architectures are also exposed to risks such as instability and security. When shifting to large scale distributed adaptive systems, we need to ensure that a) the society continues to function in the ways we have grown accustomed to, and b) that such systems remain manageable both in terms of cost and resource usage, but also in terms of control.

We remain optimistic that this long-term objective will be reached through the use of our open source software suite and published case studies, to encourage the adoption of the QUANTICOL methodology, allowing designers and other stakeholders to model, understand and control CAS appropriately and therefore benefit their users.

The two CAS scenarios adopted as case studies within the project — smart transformation and smart grid — have been the most immediate vehicles for achieving societal impact. Through these examples we have been able to showcase the scientific and technological achievements of the projects in an accessible way to potential users and other stakeholders. We have used a number of vehicles to make this impact:

- Key stakeholders in smart urban transport in Edinburgh, Lothian Buses and Edinburgh City Council, have been represented on the Advisory Board, and have regularly interacted with UEDIN members of the project team on issues related to bus modelling.
- Meanwhile the CNR-ISTI team have established strong links with PisaMo S.p.A, Pisa and Bicincittà, Turin to support the work on bike sharing systems.
- INRIA have engaged with member of the energy industry to explore issues related to control in smart grid scenarios. This has led to an on-going collaboration and further funding for INRIA from Schneider Electric.

- Team members have participated in public-facing events such as FIAB-Pisa (Federation Italiana Amici della Bicicletta – Italian Federation of Bicycle Friends) meetings, meetings of Sustrans, a UK charity whose objective it is to encourage sustainable transport, and a "transport hack" organised by the City of Edinburgh.
- Vincenzo Ciancia gave a presentation entitled "La bici condivisa (The shared bike)" for a general public audience in the context of the European "Night of Researchers" event called BRIGHT at the Area della Ricerca of CNR in Pisa.
- Print articles appeared aimed at more general audiences in ERCIM news and EU Researcher.

Finally, Figure 22 shows how the technological results of the project contribute to the **transformational impact on society** by providing a tool supported design methodology for CAS. Such a methodology has both great relevance for the ICT industry that will be involved in the development of reliable and safe CAS based systems for smart cities, as well as for citizens and city councils that need to be able to require the right level of guarantees for the functioning of such systems.

4.2 Main Dissemination Activities

Academic publications and associated presentations The QUANTICOL project has been highly successful in disseminating the results of its research in leading academic journals, conferences, workshops, and other venues. A record of the publications from the project is listed on the QUANTICOL web site and is broken down by year of publication. As of the time of submission of this document, the publication counts per year are as shown below.

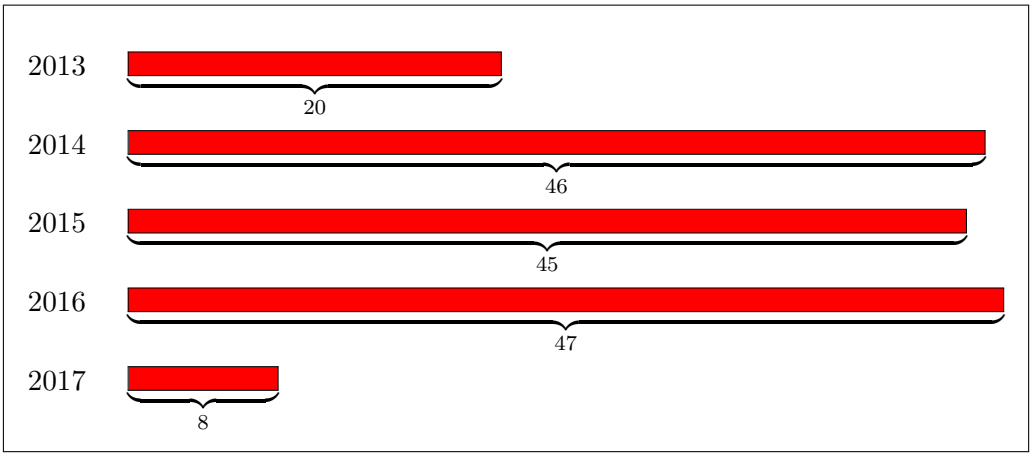

This gives a total of 166 publications over the years of the QUANTICOL project, approaching an impressive average of one publication per week of the project's duration. This total paper count is set to increase for 2017 because of papers which are currently under review (or, in some cases, already accepted) for publication. We have chosen not to include these in this summary (or feature them on our website) because they have not yet appeared in print. Except where publisher's requirements prohibit this, the project publications are available under open-access terms from archival sources in the form of the author's final submitted copy.

We can break down the publications of the project by category as journal papers, conference papers, workshop papers, summer school papers, and others (such as theses). The project has made the strategic decision to prioritise timely dissemination of work in high-quality conferences in order to increase awareness in the scientific community of our ongoing programme of work. This decision brings the benefit that research can be disseminated within a year, rather than after the typical lead time of a scientific journal, which could be several years.

Presenting work at a conference (or workshop) has the added dimension that one of the authors prepares a presentation on the work, aiming for a coherent and comprehensible summary of the work which can be appreciated by researchers who are not necessarily currently working in the same area. This requirement to actively motivate interest in the research for the audience at a conference has the beneficial effect of reminding the authors themselves of the practical motivations of the project and bringing their attention back to the purpose behind the work. In our experience submitting a paper to a journal does not necessarily have the same requirement to motivate the work.

The breakdown of the project publications by type is given below.

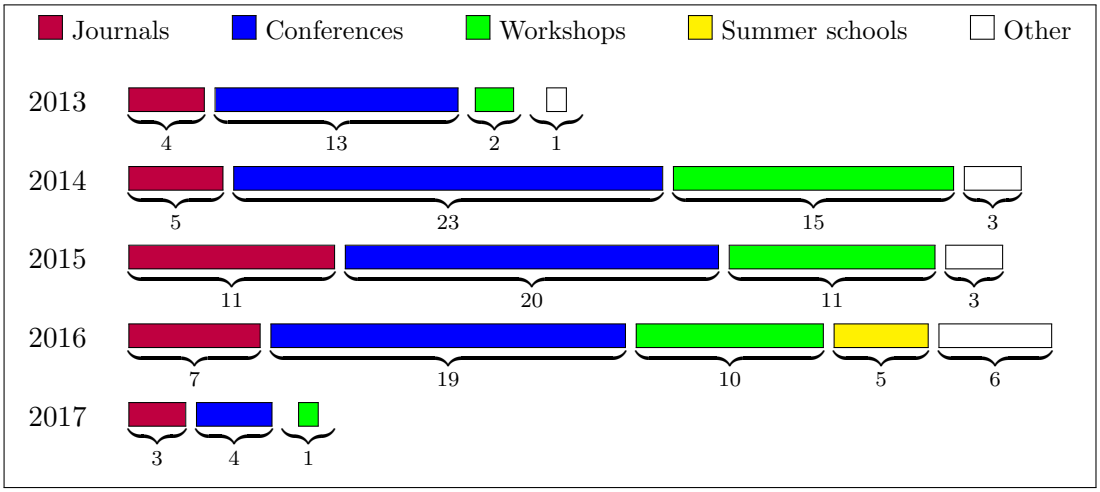

Summarising these figures gives a total of 30 journal papers, 79 conference papers, 39 workshop papers, 5 summer school papers, and 13 other publications such as theses and papers in festschrifts and other symposia over the years of the QUANTICOL project.

In order to summarise the project publications thematically, we have made the following "word cloud" from the titles of the papers published by the project. As usual in the formation of word clouds, more frequently-occurring words are displayed more prominently. From the image below we can confirm that the focus of the project has actually been (as planned) on *quantitative modelling* and *analysis* of *collective adaptive systems* using *probabilistic* and *stochastic models*.

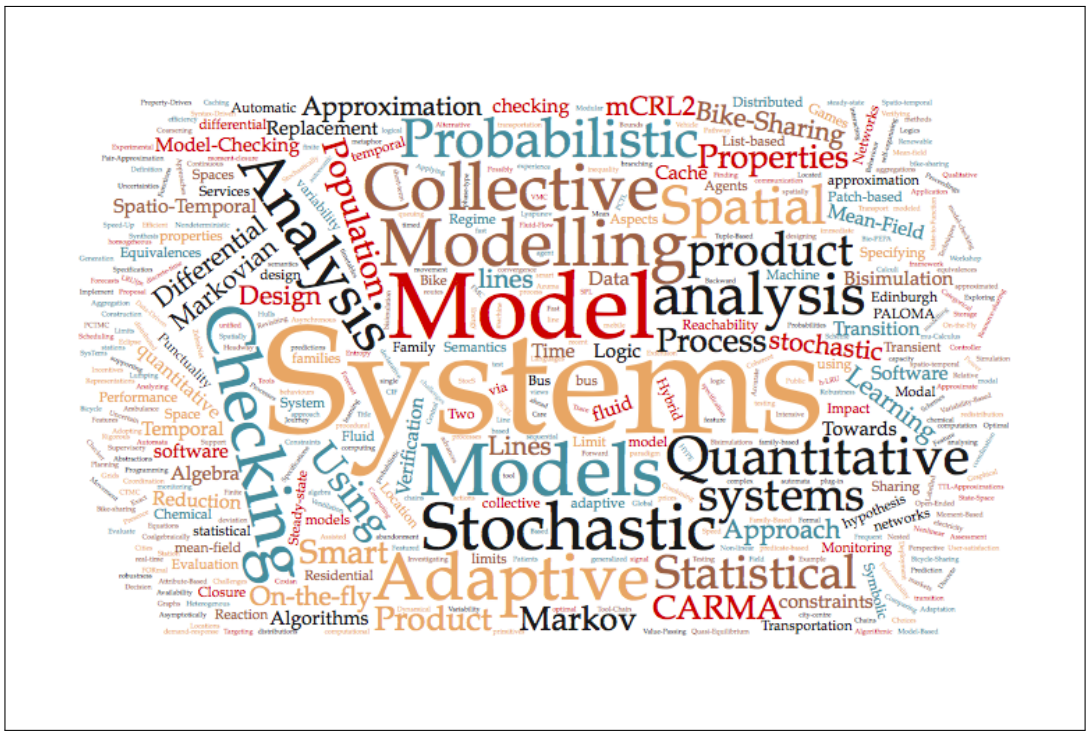

Non-academic publications In addition to technical publications which target the scientific community the QUANTICOL project has also prepared non-technical summaries on the methods and vision of the project for publications such as newsletters and periodicals which are not necessarily targeting a computer science audience.

The FoCAS project (*Fundamentals of Collective Adaptive Systems*) coordinating action on collective adaptive systems provided a useful service to the QUANTICOL project through its FoCAS newsletter which encouraged us to concentrate on non-technical writing which provides an accessible overview of our technical achievements.

Among the venues for our publications of non-academic summaries of our work was ERCIM news, a newsletter where we published:

- *Adopting a machine learning approach in the design of smart transportation systems.* Davide Bacciu, Antonio Carta, Stefania Gnesi and Laura Semini. ERCIM News 105: Special theme Planning and Logistics. April 2016.
- *A quantitative approach to the design and analysis of collective adaptive systems for smart cities.* Maurice ter Beek, Luca Bortolussi, Vincenzo Ciancia, Stefania Gnesi, Jane Hillston, Diego Latella and Mieke Massink. ERCIM news 98 — special theme: Smart Cities. July 2014.

Open source software There has been substantial effort on software development during the third period of the project, resulting in a multifaceted suite of tools. The software development has been informed by both theory and practice, and a number of case studies have been undertaken in order to ensure the applicability of the techniques and their implementation. We have also placed some emphasis on ensuring the framework is accessible and suitable to users. For example, the process algebra-based language CARMA is now supported by the CARMA Specification Language, CASL, providing a more programmatic style of model-building, a graphical interface for model construction and a command-line tool for model execution and experimentation. These are all reported in D4.3. Moreover, a design workflow and analysis pathway has been developed, embedded in the tool, to guide the user to apply appropriate methods to their model. In a further step to enhance usability we have considered how available spatial data can be usefully deployed in spatial models (D2.2) particularly in the context of smart urban transportation systems. In addition to the software tools centred on CARMA, a number of auxiliary tools, supporting various techniques developed in other work packages have also been developed and many of these are reported in D5.3.

Summer Schools, Workshops and Tutorials Another major objective of the final period has been the wide dissemination of the result of the project. Our work on this is detailed in Section 4.6, but we particularly highlight here the Summer School (reported in D6.3) on *Formal Methods for the Design of Computer, Communication and Software Systems: Quantitative Evaluation of Collective Adaptive Systems* which took place in June 2016, and the accompanying book which was published by Springer in the well-known LNCS series, and the FORECAST workshop which was organised at the STAF multi conference in Vienna in July 2016, specifically to target the software engineering community. A journal special issue is currently in preparation with extended versions of the best papers from the workshop.

4.3 Exploitation of Results

At the end of the second period we took the opportunity to carry out an audit of the exploitable knowledge items which had been created in the project so far. The project partners are all academic institutions and project members are keen to see their results used and exploited by others, there are few cases where we believe this could lead directly to commercialisation activities. Nevertheless we

identified many opportunities for potential indirect exploitation by third parties, including commercialisation in some cases, and we would envisage consulting with developers in these cases. At month 30 we identified the following knowledge items:

1. **Distributed algorithm for optimising the use of electric resources**

   Innovation: Distributed algorithm for optimising the use of electric resources.

   Target group: Distribution network operators or the companies who build control systems for such networks

   Exploitation: Indirect

   Contact person: Nicolas Gast, INRIA

2. **A generic framework for the continuous approximation of stochastic systems**

   Innovation: Framework for the continuous approximation of stochastic systems

   Target group: Researchers and modelling tool developers

   Exploitation: Indirect

   Contact person: Luca Bortolussi, CNR

3. **Model transformation from continuous space to discrete space for scalability**

   Innovation: The basis of a new technique for model transformation from continuous space representations with individual agents to discrete space representations with populations.

   Target group: Researchers

   Exploitation: Direct or indirect

   Contact person: Cheng Feng, UEDIN

4. **Model transformation from discrete space to continuous space for efficiency**

   Innovation: An efficient approach to analysing population grid-based discrete-space models through transformation into continuous space

   Target group: Researchers

   Exploitation: Direct or indirect

   Contact person: Mirco Tribastone, IMT

5. **A process algebra for population discrete-space modelling**

   Innovation: The PALOMA modelling language and associated moment-closure techniques techniques with freely available software

   Target group: Researchers and modellers

   Exploitation: Direct or indirect

   Contact person: Jane Hillston, UEDIN

6. **Spatio-Temporal Model Checking**

   Innovation: Software tool for spatio-temporal model checking

   Target group: Developers of industry-strength model checking tools and their users.

   Exploitation: Direct and indirect

Contact person: Vincenzo Ciancia, CNR

7. **Bisimulation-based model order reduction**

   Innovation: Software tools implementing algorithms for the reduction of dynamical systems, based on bisimulation minimisation and program verification techniques.

   Target group: Companies involved with large-scale simulations of dynamical systems.

   Exploitation: Indirect

   Contact person: Mirco Tribastone, IMT

8. **A language for modelling and programming collective adaptive systems**

   Innovation: The CARMA language for modelling, analysing and programming collective adaptive systems

   Target group: System designers, decision makers, end users

   Exploitation: Direct or indirect

   Contact person: Michele Loreti, IMT

9. **Techniques for bus route modelling**

   Innovation: A procedure for using existing AVL bus data to build and parameterise models of dynamic bus system performance

   Target group: Researchers, city planners, transport operators

   Exploitation: Direct or indirect

   Contact person: Daniël Reijsbergen/Stephen Gilmore, UEDIN

10. **A prediction algorithm for bike-sharing systems**

    Innovation: A stochastic prediction algorithm for bike-sharing systems.

    Target group: Third party app developers or bike-sharing system operators

    Exploitation: Indirect

    Contact person: Daniël Reijsbergen, UEDIN, Guillaume Massonnet, INRIA

We believe that each of these represents a genuine innovation which offers substantial enhancements for their potential beneficiaries.

The final period has seen the enrichment of some of these knowledge items and the emergence of a new item:

- Work in work package 1 on the control of distribution networks which led to item 1 above, has been further developed and has now become the subject of an active collaboration between the company Schneider Electric and QUANTICOL partner INRIA. An exploratory prototype is being developed and this collaboration will continue beyond the end of the QUANTICOL project.

- Further work on scalable model solution based on moment closures, has led to an improvement in the prediction algorithm for bike-sharing systems reported in item 10 above. This algorithm has been demonstrated on data from the London bike sharing system to establish both its accuracy and its efficiency, showing that it would be plausible to use it in real-time, for example within a mobile phone application. Although none of the partners are actively working on such an application (software development of that nature is outside our domain of expertise) the algorithm has been openly published and is available for exploitation.

- The knowledge item 7 is now fully supported in the ERODE software tool. However the primary beneficiaries of this work are expected to be other academics, and so there are no immediate plans for commercialisation of either the techniques nor the associated tool.
- A refinement of knowledge item 6 has been specifically targeted at medical image recognition. Use of formal modelling and logic-based property querying is highly novel in this area but initial studies show that it is very promising. More detailed evaluation work is now underway in collaboration with Azienda Ospedaliera Universitaria Senese, Fisica Sanitaria. The team responsible, in CNR-ISTI are actively seeking routes to exploitation. They applied for a LaunchPad project in September 2016 (which unfortunately was not selected for funding despite obtaining above threshold scores for all evaluation criteria) and are continuing to investigate the options to exploit this development.

4.4 Project Public Website

Further details of the project can be found on our website `www.quanticol.eu`. This website has been continually developed and enhanced throughout the project as we regard it as the main interface between the project and our scientific peers. In addition software developed during the project and suitable for use by third parties is available from download from `quanticol.github.io` and number of other sites as detailed on `www.quanticol.eu`.

QUANTICOL list of scientific (peer-reviewed) publications, starting with the most important ones

Title	Main author	Title of the periodical or the series	Number, date or frequency	Publisher	Place of publication	Year of publication	Relevant pages	Permanent identifiers	Is/will open access be provided for this publication
Automatic moment-closure approximation of spatially distributed collective adaptive systems.	Cheng Feng (with Jane Hillston and Vashti Galpin)	ACM Transactions on Modelling and Computer Simulation	Volume 26, Issue 4, Article 26 (March 2016)	ACM	New York, USA	2016	pp. 26:1—26:22	http://doi.org/10.1145/2883608 and http://bit.ly/2pEvrCu	Yes
Model checking single agent behaviours by fluid approximation	Luca Bortolussi (with Jane Hillston)	Information and Computation	Volume 242, June 2015	Elsevier	Amsterdam	2015	pp. 183–226	http://doi.org/10.1016/j.ic.2015.03.002 and http://bit.ly/2pA1dlV	Yes
Spatial fluid limits for stochastic mobile networks	M. Tschaikowski (with M. Tribastone)	Performance Evaluation	Volume 109, March 2017	Elsevier	Amsterdam	2017	pp. 52—76	http://dx.doi.org/10.1016/j.peva.2016.12.005 and https://arxiv.org/pdf/1307.4566.pdf	Yes
An experience in using machine learning for short-term predictions in smart transportation systems	Davide Bacciu (with Antonio Carta, Stefania Gnesi and Laura Semini)	Journal of Logical and Algebraic Methods in Programming	Volume 87, February 2017	Elsevier	Amsterdam	2017	pp. 52—66	http://doi.org/10.1016/j.jlamp.2016.11.002	No
Proceedings of the Workshop on FORmal methods for the quantitative Evaluation of Collective Adaptive SysTems	Maurice H. ter Beek (with Michele Loreti)	Electronic Proceedings in Theoretical Computer Science	EPTCS 217, 2016	Open Publishing Association	Waterloo, Australia	2016	pp. 1—92	https://arxiv.org/abs/1607.02001	Yes
Model Checking Spatial Logics for Closure Spaces	Vincenzo Ciancia (with Diego Latella, Michele Loreti and Mieke Massink)	Logical Methods in Computer Science,	Volume 12, Issue 4. October 11, 2016	episciences.org		2016	pp. 1—51	https://arxiv.org/abs/1609.06513	Yes
Approximate Reduction of Heterogenous Nonlinear Models With Differential Hulls	Max Tschaikowski (with Mirco Tribastone)	IEEE Transactions on Automatic Control	Volume 61, Number 4, April 2016	IEEE	New York, USA	2016	pp. 1099—1104	http://doi.org/10.1109/TAC.2015.2457172 and http://bit.ly/2paf4zT	Yes

Title	Main author	Title of the periodical or the series	Number, date or frequency	Publisher	Place of publication	Year of publication	Relevant pages	Permanent identifiers	Is/will open access be provided for this publication
Incentives and redistribution in homogeneous bike-sharing systems with stations of finite capacity	Christine Fricker (with Nicolas Gast)	EURO Journal on Transportation and Logistics	August 2016, Volume 5, Issue 3.	Springer	Berlin, Heidelberg	2016	pp. 261–291	http://doi.org/10.1007/s13676-014-0053-5 and https://arxiv.org/abs/1201.1178	Yes
ERODE: A Tool for the Evaluation and Reduction of Ordinary Differential Equations	Luca Cardelli (with Mirco Tribastone, Max Tschaikowski and Andrea Vandin)	TACAS'17: Tools and Algorithms for the Construction and Analysis of Systems	LNCS 10206	Springer	Berlin, Heidelberg	2017	pp. 310–328	http://doi.org/10.1007/978-3-662-54580-5_19 and http://bit.ly/2paZ9xf	Yes
Modelling and Analysis of Collective Adaptive Systems with CARMA and its Tools	Michele Loreti (with Jane Hillston)	Formal Methods for the Quantitative Evaluation of Collective Adaptive Systems	LNCS 9700, 11 June 2016	Springer	Berlin, Heidelberg	2016	pp. 83–119	http://doi.org/10.1007/978-3-319-34096-8_4 and http://www.research.ed.ac.uk/portal/en/publications/modelling-and-analysis-of-collective-adaptive-systems-with-carma-and-its-tools(298d104e-d255-4c2b-9c74-6708755eaa39).html	Yes
Quantitative Abstractions for Collective Adaptive Systems	Andrea Vandin (with Mirco Tribastone)	Formal Methods for the Quantitative Evaluation of Collective Adaptive Systems	LNCS 9700, 11 June 2016	Springer	Berlin, Heidelberg	2016	pp. 202—232	http://doi.org/10.1007/978-3-319-34096-8_7	No
Efficient Syntax-Driven Lumping of Differential Equations	Luca Cardelli (with Mirco Tribastone, Max Tschaikowski and Andrea Vandin)	TACAS 2016: Tools and Algorithms for the Construction and Analysis of Systems	LNCS 9636, 09 April 2016	Springer	Berlin, Heidelberg	2016	pp. 93–111	http://doi.org/10.1007/978-3-662-49674-9_6 and http://bit.ly/2o3rupn	Yes

Title	Main author	Title of the periodical or the series	Number, date or frequency	Publisher	Place of publication	Year of publication	Relevant pages	Permanent identifiers	Is/will open access be provided for this publication
Symbolic computation of differential equivalences	Luca Cardelli (with Mirco Tribastone, Max Tschaikowski and Andrea Vandin)	POPL'16: Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages	January 20 - 22, 2016	ACM	New York, USA	2016	pp. 137–150	http://doi.org/10.1145/2837614.2837649 and http://bit.ly/2omyZZw	Yes
A unified framework for differential aggregations in Markovian process algebra	Max Tschaikowski (with Mirco Tribastone)	Journal of Logical and Algebraic Methods in Programming	Volume 84, Issue 2, March 2015	Elsevier	Amsterdam	2015	pp. 238–258	http://doi.org/10.1016/j.jlamp.2014.10.004 and http://bit.ly/2oqMBTr	Yes
On-the-fly PCTL fast mean-field approximated model-checking for self-organising coordination	Diego Latella (with Michele Loreti and Mieke Massink)	Science of Computer Programming	Volume 110, 15 October 2015	Elsevier	Amsterdam	2015	pp. 23—50	http://doi.org/10.1016/j.scico.2015.06.009 and http://bit.ly/2px4Vea	Yes
System design of stochastic models using robustness of temporal properties	Ezio Bartoccia (with Luca Bortolussi, Laura Nenzi, Guido Sanguinetti)	Theoretical Computer Science	Volume 587, 4 July 2015	Elsevier	Amsterdam	2015	pp. 3–25	http://doi.org/10.1016/j.tcs.2015.02.046 and http://bit.ly/2pxthEu	Yes
Probabilistic Forecasts of Bike-Sharing Systems for Journey Planning	Nicolas Gast (with Guillaume Massonnet, Daniël Reijsbergen and Mirco Tribastone)	CIKM '15: Conference on Information and Knowledge Management	October 18–23, 2015	ACM	New York, USA	2015	pp. 703–712	https://doi.org/10.1145/2806416.2806569 and https://hal.inria.fr/hal-01185840/	Yes
Speed-Up of Stochastic Simulation of PCTMC Models by Statistical Model Reduction	Cheng Feng (and Jane Hillston)	EPEW '15: European Performance Engineering Workshop	LNCS 9272, 22 August 2015	Springer	Berlin, Heidelberg	2015	pp. 291–305	http://doi.org/10.1007/978-3-319-23267-6_19 and http://bit.ly/2plpehA	Yes
On-the-fly Fluid Model Checking via Discrete Time Population Models	Diego Latella (with Michele Loreti and Mieke Massink)	EPEW '15: European Performance Engineering Workshop	LNCS 9272, 22 August 2015	Springer	Berlin, Heidelberg	2015	pp. 193-207	http://doi.org/10.1007/978-3-319-23267-6_13 and http://bit.ly/2otz7Wr	Yes

Title	Main author	Title of the periodical or the series	Number, date or frequency	Publisher	Place of publication	Year of publication	Relevant pages	Permanent identifiers	Is/will open access be provided for this publication
Qualitative and Quantitative Monitoring of Spatio-Temporal Properties	Laura Nenzi (with Luca Bortolussi, Vincenzo Ciancia, Michele Loreti and Mieke Massink)	RV '15: Runtime Verification	LNCS 9333, 15 November 2015	Springer	Berlin, Heidelberg	2015	pp. 21–37	http://doi.org/10.1007/978-3-319-23820-3_2 and http://bit.ly/2otIV2k	Yes
Optimal Generation and Storage Scheduling in the Presence of Renewable Forecast Uncertainties	Nicolas Gast (with Dan-Cristian Tomozei and Jean-Yves Le Boudec)	IEEE Transactions on Smart Grid	Volume 5, Issue 3, May 2014	IEEE	New York, USA	2014	pp. 1328–1339	http://doi.org/10.1109/TSG.2013.2285395 and http://bit.ly/2oX76Wu	Yes
A Quantitative Approach to the Design and Analysis of Collective Adaptive Systems for Smart Cities	Maurice ter Beek (with Luca Bortolussi, Vincenzo Ciancia, Stefania Gnesi, Jane Hillston, Diego Latella and Mieke Massink)	ERCIM News 98. Special theme: Smart cities	July 2014	ERCIM		2014		http://bit.ly/2oXIRsy	Yes
Spatial fluid limits for stochastic mobile networks	Max Tschaikowski (with Mirco Tribastone)	Performance Evaluation	Volume 109, March 2017	Elsevier	Amsterdam	2017	pp. 52–76	https://doi.org/10.1016/j.peva.2016.12.005 and https://arxiv.org/abs/1307.4566	Yes
An experience in using machine learning for short-term predictions in smart transportation systems.	Davide Bacciua (with Antonio Carta, Stefania Gnesi and Laura Semini)	Journal of Logical and Algebraic Methods in Programming	Volume 87, February 2017	Elsevier	Amsterdam	2017	pp. 52–66	https://doi.org/10.1016/j.jlamp.2016.11.002 and http://bit.ly/2pL3s41	Yes
A computational approach to steady-state convergence of fluid limits for Coxian queuing networks with abandonment	M. Tschaikowski (with M. Tribastone)	Annals of Operations Research	23 Apr 2016	Springer	Berlin, Heidelberg	2016	pp. 1—20	http://dx.doi.org/10.1007/s10479-016-2193-5	No

Title	Main author	Title of the periodical or the series	Number, date or frequency	Publisher	Place of publication	Year of publication	Relevant pages	Permanent identifiers	Is/will open access be provided for this publication
A Tool-Chain for Statistical Spatio-Temporal Model Checking of Bike Sharing Systems	Vincenzo Ciancia (with Diego Latella, Mieke Massink, Rytis Paškauskas and Andrea Vandin)	ISoLA 2016: Leveraging Applications of Formal Methods, Verification and Validation: Foundational Techniques	LNCS 9952, 5 October 2016	Springer	Berlin, Heidelberg	2016	pp. 657–673	http://doi.org/10.1007/978-3-319-47166-2_46	No
Rigorous Graphical Modelling of Movement in Collective Adaptive Systems	Natalia Zoń (with Stephen Gilmore and Jane Hillston)	ISoLA 2016: Leveraging Applications of Formal Methods, Verification and Validation: Foundational Techniques	LNCS 9952, 5 October 2016	Springer	Berlin, Heidelberg	2016	pp. 674–688	http://doi.org/10.1007/978-3-319-47166-2_47 and http://www.research.ed.ac.uk/portal/en/publications/rigorous-graphical-modelling-of-movement-in-collective-adaptive-systems(6b90bd13-be7a-4979-9e96-1253a2b1bacf).html	Yes
Property-Driven State-Space Coarsening for Continuous Time Markov Chains	Michalis Michaelides (with Dimitrios Milios, Jane Hillston, and Guido Sanguinetti)	QEST 2016: International Conference on Quantitative Evaluation of Systems	LNCS 9826, 3 August 2016	Springer	Berlin, Heidelberg	2016	pp. 3—18	http://dx.doi.org/10.1007/978-3-319-43425-4_1 and http://www.research.ed.ac.uk/portal/en/publications/propertydriven-statespace-coarsening-for-continuous-time-markov-chains(5ec44eb6-8535-489b-9f89-fb933173b8f1).html	Yes

Title	Main author	Title of the periodical or the series	Number, date or frequency	Publisher	Place of publication	Year of publication	Relevant pages	Permanent identifiers	Is/will open access be provided for this publication
CARMA Eclipse plug-in: A tool supporting design and analysis of Collective Adaptive Systems	Jane Hillston (with Michele Loreti)	QEST 2016: International Conference on Quantitative Evaluation of Systems	LNCS 9826, 3 August 2016	Springer	Berlin, Heidelberg	2016	pp. 167—171	http://doi.org/10.1007/978-3-319-43425-4_12 and http://www.research.ed.ac.uk/portal/en/publications/carma-eclipse-plugin-a-tool-supporting-design-and-analysis-of-collective-adaptive-systems(ae0e9fc1-d15b-4e0d-b370-413c86867258).html	Yes
Stochastic and Spatial Equivalences for PALOMA	Paul Piho (with Jane Hillston)	Proceedings of the Workshop on FORmal methods for the quantitative Evaluation of Collective Adaptive SysTems	EPTCS 217, 2016	episciences.org		2016	pp. 69–80	https://arxiv.org/abs/1607.02234	Yes
Statistical Model Checking for Product Lines	Maurice H. ter Beek (with Axel Legay, Alberto Lluch Lafuente and Andrea Vandin)	ISoLA 2016: Leveraging Applications of Formal Methods, Verification and Validation: Foundational Techniques	LNCS 9952, 5 October 2016	Springer	Berlin, Heidelberg	2016	pp. 114–133	http://doi.org/10.1007/978-3-319-47166-2_8 and https://hal.inria.fr/hal-01406531	Yes
Adopting a Machine Learning Approach in the Design of Smart Transportation Systems	Davide Bacciu (with Antonio Carta, Stefania Gnesi and Laura Semini)	ERCIM News 105: Special theme Planning and Logistics.	April 2016	ERCIM	Sophia Antipolis, France	2016		http://ercim-news.ercim.eu/en105/special/adopting-a-machine-learning-approach-in-the-design-of-smart-transportation-systems	Yes

Title	Main author	Title of the periodical or the series	Number, date or frequency	Publisher	Place of publication	Year of publication	Relevant pages	Permanent identifiers	Is/will open access be provided for this publication
Probabilistic Modelling of Station Locations in Bicycle-Sharing Systems	Daniël Reijsbergen	STAF 2016: Software Technologies: Applications and Foundations	LNCS 9946, 1 December 2016	Springer	Berlin, Heidelberg	2016	pp. 83–97	http://doi.org/10.1007/978-3-319-50230-4_7 and http://bit.ly/2oy4fGY	Yes
Moment-Based Probabilistic Prediction of Bike Availability for Bike-Sharing Systems	Cheng Feng (with Jane Hillston and Daniël Reijsbergen)	QEST 2016: Quantitative Evaluation of Systems	LNCS 9826, 3 August 2016	Springer	Berlin, Heidelberg	2016	pp. 139–155	http://doi.org/10.1007/978-3-319-43425-4_9 and http://bit.ly/2o4lpqD	Yes
Location Aggregation of Spatial Population CTMC Models	Luca Bortolussi (with Cheng Feng)	QAPL'16: Quantitative Aspects of Programming Languages and Systems	EPTCS 227, 25 Oct 2016	episciences.org		2016	pp. 30–43	http://doi.org/10.4204/EPTCS.227.3	Yes
Spatial Representations and Analysis Techniques	Vashti Galpin	Formal Methods for the Quantitative Evaluation of Collective Adaptive Systems	LNCS 9700, 11 June 2016	Springer	Berlin, Heidelberg	2016	pp. 120–155	http://dx.doi.org/10.1007/978-3-319-34096-8_5 and http://bit.ly/2p5vKbv	Yes
Mean Field Approximation of Uncertain Stochastic Models	Luca Bortolussi (with Nicolas Gast)	DSN'16: Dependable Systems and Networks	3 October 2016	IEEE	New York, USA	2016	pp. 287—298	http://doi.org/10.1109/DSN.2016.34 and https://hal.inria.fr/hal-01302416	Yes
On-the-Fly Mean-Field Model-Checking for Attribute-Based Coordination	Vincenzo Ciancia (with Diego Latella and Mieke Massink)	COORDINATION 2016: Coordination Models and Languages	LNCS 9686, 24 May 2016	Springer	Berlin, Heidelberg	2016	pp. 67–83	http://doi.org/10.1007/978-3-319-39519-7_5 and http://bit.ly/2oa0PpZ	Yes
Spatial Logic and Spatial Model Checking for Closure Spaces	Vincenzo Ciancia (with Diego Latella, Michele Loreti and Mieke Massink)	Formal Methods for the Quantitative Evaluation of Collective Adaptive Systems	LNCS 9700, 11 June 2016	Springer	Berlin, Heidelberg	2016	pp. 156–201	http://doi.org/10.1007/978-3-319-34096-8_6	No

Title	Main author	Title of the periodical or the series	Number, date or frequency	Publisher	Place of publication	Year of publication	Relevant pages	Permanent identifiers	Is/will open access be provided for this publication
Variability-Based Design of Services for Smart Transportation Systems	Maurice H. ter Beek (with Alessandro Fantechi, Stefania Gnesi and Laura Semini)	ISoLA 2016: Leveraging Applications of Formal Methods, Verification and Validation: Foundational Techniques	LNCS 9953, 5 October 2016	Springer	Berlin, Heidelberg	2016	pp. 465–481	http://doi.org/10.1007/978-3-319-47169-3_38	No
Modelling Ambulance Deployment with CARMA	Vashti Galpin	COORDINATION 2016: Coordination Models and Languages	LNCS 9686, 24 May 2016	Springer	Berlin, Heidelberg	2016	pp. 121–137	http://doi.org/10.1007/978-3-319-39519-7_8 and http://bit.ly/2o96yNG	Yes
Modelling and analysing variability in product families: Model checking of modal transition systems with variability constraints	Maurice H. ter Beek (with Alessandro Fantechia, Stefania Gnesia and Franco Mazzantia)	Journal of Logical and Algebraic Methods in Programming	Volume 85, Issue 2, February 2016	Elsevier	Amsterdam	2016	pp. 287–315	http://dx.doi.org/10.1016/j.jlamp.2015.11.006	Yes
Towards Automatic Decision Support for Bike-Sharing System Design	Maurice ter Beek (with Stefania Gnesi, Diego Latella and Mieke Massink)	SEFM ’15: Software Engineering and Formal Methods	LNCS 9509, 12 January 2016	Springer	Berlin, Heidelberg	2016	pp. 266–280	http://doi.org/10.1007/978-3-662-49224-6_22 and http://sysma.imtlucca.it/publications/cina/161/	Yes
An Experimental Spatio-Temporal Model Checker	Vincenzo Ciancia (with Gianluca Grilletti, Diego Latella, Michele Loreti and Mieke Massink)	SEFM ’15: Software Engineering and Formal Methods	LNCS 9509, 12 January 2016	Springer	Berlin, Heidelberg	2016	pp. 297–311	http://doi.org/10.1007/978-3-662-49224-6_24 and http://bit.ly/2oxm6u9	Yes
Quantitative Modelling of Residential Smart Grids	Vashti Galpin	SEFM ’15: Software Engineering and Formal Methods	LNCS 9509, 12 January 2016	Springer	Berlin, Heidelberg	2016	pp. 161–175	http://doi.org/10.1007/978-3-662-49224-6_14 and http://bit.ly/2obUsr5	Yes

Title	Main author	Title of the periodical or the series	Number, date or frequency	Publisher	Place of publication	Year of publication	Relevant pages	Permanent identifiers	Is/will open access be provided for this publication
Mean-Field Approximation and Quasi-Equilibrium Reduction of Markov Population Models	Luca Bortolussi (with Rytis Paškauskas)	QEST 2014: Quantitative Evaluation of Systems	LNCS 8657	Springer	Berlin, Heidelberg	2014	pp. 106–121	http://doi.org/10.1007/978-3-319-10696-0_9 and https://arxiv.org/abs/1405.4200	Yes
Family-Based Model Checking with mCRL2	Maurice H. ter Beek (with Erik P. de Vink and Tim A. C. Willemse)	FASE 2017: Fundamental Approaches to Software Engineering	22 March 2017	Springer	Berlin, Heidelberg	2017	pp. 387–405	http://dx.doi.org/10.1007/978-3-662-54494-5_23 and http://www.win.tue.nl/~evink/research/Postscript/fase17.pdf	Yes
Specification and Analysis of Open-Ended Systems with CARMA	Jane Hillston (with Michele Loreti)	Agent Environments for Multi-Agent Systems IV	LNCS 9068, 27 November 2015	Springer	Berlin, Heidelberg	2015	pp. 95–116	http://doi.org/10.1007/978-3-319-23850-0_7 and http://bit.ly/2o8SROY	Yes
Probabilistic Modelling of the Impact on Bus Punctuality of a Speed Limit Proposal in Edinburgh	Daniël Reijsbergen (with Rajeev Ratan)	VALUETOOLS'15: Proceedings of the 9th EAI International Conference on Performance Evaluation Methodologies and Tools	December 14–16, 2015	ACM	New York, USA	2015	pp. 234–241	http://doi.org/10.4108/eai.14-12-2015.2262703 and https://arxiv.org/abs/1511.05363	Yes
Model-Based Assessment of Aspects of User-satisfaction in Bicycle Sharing Systems	Mieke Massink (with Rytis Paškauskas)	ITSC'15: International Conference on Intelligent Transportation Systems	2 November 2015	IEEE	New York, USA	2015	pp. 1363—1370	http://doi.org/10.1109/ITSC.2015.224 and http://bit.ly/2pmAgSn	Yes
Applying the product lines paradigm to the quantitative analysis of collective adaptive systems	Maurice H. ter Beek (with Alessandro Fantechi and Stefania Gnesi)	SPLC '15: International Conference on Software Product Lines	July 20–24, 2015	ACM	New York, USA	2015	pp. 321—326	http://doi.org/10.1145/2791060.2791100 and http://bit.ly/2qkLwjs	Yes

Title	Main author	Title of the periodical or the series	Number, date or frequency	Publisher	Place of publication	Year of publication	Relevant pages	Permanent identifiers	Is/will open access be provided for this publication
Statistical analysis of probabilistic models of software product lines with quantitative constraints	Maurice H. ter Beek (with Axel Legay, Alberto Lluch Lafuente and Andrea Vandin)	SPLC '15: International Conference on Software Product Lines	July 20–24, 2015	ACM	New York, USA	2015	pp. 11—15	http://doi.org/10.1145/2791060.2791087 and http://bit.ly/2qwBVIX	Yes
Forward and Backward Bisimulations for Chemical Reaction Networks	Luca Cardelli (with Mirco Tribastone, Max Tschaikowski and Andrea Vandin)	CONCUR'15: International Conference on Concurrency Theory	Leibniz International Proceedings in Informatics (LIPIcs), Volume 42.	Schloss Dagstuhl--Leibniz-Zentrum fuer Informatik	Dagstuhl, Germany	2015	pp. 226—239	http://dx.doi.org/10.4230/LIPIcs.CONCUR.2015.226	Yes
Differential Bisimulation for a Markovian Process Algebra	Giulio Iacobelli (with Mirco Tribastone and Andrea Vandin)	MFCS 2015: Mathematical Foundations of Computer Science 2015	LNCS 9234, 11 August 2015	Springer	Berlin, Heidelberg	2015	pp. 293—306	http://doi.org/10.1007/978-3-662-48057-1_23 and http://bit.ly/2p1pnmu	Yes
Investigating Fluid-Flow Semantics of Asynchronous Tuple-Based Process Languages for Collective Adaptive Systems	Diego Latella (with Michele Loreti and Mieke Massink)	COORDINATION 2015: Coordination Models and Languages	LNCS 9037	Springer	Berlin, Heidelberg	2015	pp. 19—34	http://doi.org/10.1007/978-3-319-19282-6_2	No
On StocS: A Stochastic Extension of SCEL	Diego Latella (with Michele Loreti, Mieke Massink and Valerio Senni)	Software, Services, and Systems	LNCS 8950	Springer	Berlin, Heidelberg	2015	pp. 619–640	http://doi.org/10.1007/978-3-319-15545-6_35	No
Fluid Performability Analysis of Nested Automata Models	Luca Bortolussi (with Mirco Tribastone and Jane Hillston)	PASM'14: Practical Application of Stochastic Modelling	Electronic Notes in Theoretical Computer Science. Volume 310, 5 January 2015	episciences.org		2015	pp. 27–47	https://doi.org/10.1016/j.entcs.2014.12.011	Yes
Patch-based modelling of city-centre bus movement with phase-type distributions.	Daniël Reijsbergen (with Stephen Gilmore and Jane Hillston)	PASM'14: Practical Application of Stochastic Modelling	Electronic Notes in Theoretical Computer Science. Volume 310, 5 January 2015	episciences.org		2015	pp. 157–177	https://doi.org/10.1016/j.entcs.2014.12.017	Yes

Title	Main author	Title of the periodical or the series	Number, date or frequency	Publisher	Place of publication	Year of publication	Relevant pages	Permanent identifiers	Is/will open access be provided for this publication
Finding optimal timetables for Edinburgh bus routes	Ludovica Luisa Vissat (with Allan Clark and Stephen Gilmore)	PASM'14: Practical Application of Stochastic Modelling	Electronic Notes in Theoretical Computer Science. Volume 310, 5 January 2015	episciences.org		2015	pp. 179–199	https://doi.org/10.1016/j.entcs.2014.12.018	Yes
On hypothesis testing for statistical model checking	Daniël Reijsbergen (with Pieter-Tjerk de Boer, Werner Scheinhardt and Boudewijn Haverkort)	International Journal on Software Tools for Technology Transfer	Volume 17, Issue 4, August 2015	Springer	Berlin, Heidelberg	2015	pp. 377–395	http://doi.org/10.1007/s10009-014-0350-1	No
Symbolic Performance Adaptation	Emilio Incerto (with Mirco Tribastone and Catia Trubiani)	IEEE/ACM 11th International Symposium on Software Engineering for Adaptive and Self-Managing Systems	SEAMS'16, May 16-17, 2016	Association for Computing Machinery	New York, USA	2016	pp. 140—150	http://dx.doi.org/10.1145/2897053.2897060 and https://www.researchgate.net/publication/303413497_Symbolic_performance_adaptation	Yes
Comparing Chemical Reaction Networks: A Categorical and Algorithmic Perspective	Luca Cardelli (with Mirco Tribastone, Max Tschaikowski and Andrea Vandin)	LICS'16: Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science	July 05 - 08, 2016	ACM	New York, USA	2016	pp. 485–494	http://doi.org/10.1145/2933575.2935318 and http://bit.ly/2omtp9t	Yes
Supervisory Controller Synthesis for Product Lines Using CIF 3	Maurice H. ter Beek (with Michel A. Reniers and Erik P. de Vink)	ISoLA 2016: Leveraging Applications of Formal Methods, Verification and Validation: Foundational Techniques	LNCS 9952, 5 October 2016	Springer	Berlin, Heidelberg	2016	pp. 856–873	http://doi.org/10.1007/978-3-319-47166-2_59 and http://bit.ly/2p5742S	Yes
Impact of demand-response on the efficiency and prices in real-time electricity markets	Nicolas Gast (with Jean-Yves Le Boudec and Dan-Cristian Tomozei)	e-Energy'14: International conference on Future energy systems	Cambridge, United Kingdom. June 11–13, 2014	ACM	New York, USA	2014	pp. 171–182	http://doi.org/10.1145/2602044.2602052 and https://hal.inria.fr/hal-01086036	Yes

Title	Main author	Title of the periodical or the series	Number, date or frequency	Publisher	Place of publication	Year of publication	Relevant pages	Permanent identifiers	Is/will open access be provided for this publication
Specifying and Verifying Properties of Space	Vincenzo Ciancia (with Diego Latella, Michele Loreti and Mieke Massink)	TCS'14: IFIP International Conference on Theoretical Computer Science	LNCS 8705	Springer	Berlin, Heidelberg	2014	pp. 222–235	http://doi.org/10.1007/978-3-662-44602-7_18 and https://arxiv.org/abs/1406.6393	Yes
CARMA: Collective Adaptive Resource-sharing Markovian Agents	Luca Bortolussi (with Rocco De Nicola, Vashti Galpin, Stephen Gilmore, Jane Hillston, Diego Latella, Michele Loreti, and Mieke Massink)	QAPL'15: Workshop on Quantitative Aspects of Programming Languages and Systems	EPTCS 194, 2015	episciences.org		2015	pp. 16–31	http://doi.org/10.4204/EPTCS.194.2	Yes
A Statistical Approach for Computing Reachability of Non-linear and Stochastic Dynamical Systems	Luca Bortolussi (with Guido Sanguinetti)	QEST 2014: Quantitative Evaluation of Systems	LNCS 8657	Springer	Berlin, Heidelberg	2014	pp. 41—56	http://doi.org/10.1007/978-3-319-10696-0_5	No
Model Checking Value-Passing Modal Specifications	Maurice H. ter Beek (with Stefania Gnesi and Franco Mazzanti)	PSI 2014: Perspectives of System Informatics	LNCS 8974	Springer	Berlin, Heidelberg	2014	pp. 304–319	http://doi.org/10.1007/978-3-662-46823-4_25 and http://bit.ly/2pOEeo8	Yes
VMC: recent advances and challenges ahead	Maurice H. ter Beek (with Franco Mazzanti)	SPLC '14: International Conference on Software Product Lines	September 15–19, 2014	ACM	New York, USA	2014	pp. 70–77	http://doi.org/10.1145/2647908.2655969 and http://bit.ly/2p0U7nt	Yes
Two Possibly Alternative Approaches to the Semantics of Stochastic Process Calculi	Rocco De Nicola (with Diego Latella, Michele Loreti and Mieke Massink)	Microsoft Research Technical Report Series	MSR-TR-2014-104	Microsoft	Cambridge, United Kingdom	2014	pp. 95—108	http://bit.ly/2pEvOwW	Yes
On-the-fly Probabilistic Model Checking	Diego Latella (with Michele Loreti and Mieke Massink)	ICE'14: Interaction and Concurrency Experience	EPTCS 166, 2014	episciences.org		2014	pp. 45–59	http://doi.org/10.4204/EPTCS.166.6	Yes

Title	Main author	Title of the periodical or the series	Number, date or frequency	Publisher	Place of publication	Year of publication	Relevant pages	Permanent identifiers	Is/will open access be provided for this publication
Challenges in Modelling and Analyzing Quantitative Aspects of Bike-Sharing Systems	Maurice H. ter Beek (with Alessandro Fantechi and Stefania Gnesi)	ISoLA 2014: Leveraging Applications of Formal Methods, Verification and Validation. Technologies for Mastering Change	LNCS 8802	Springer	Berlin, Heidelberg	2014	pp. 351–367	http://doi.org/10.1007/978-3-662-45234-9_25 and http://bit.ly/2qrziFJ	Yes
Specifying and monitoring properties of stochastic spatio-temporal systems in signal temporal logic	Luca Bortolussi (with Laura Nenzi)	VALUETOOLS '14: Conference on Performance Evaluation Methodologies and Tools	December 09–11, 2014	ACM	New York, USA	2014	pp. 66–73	http://doi.org/10.4108/icst.Valuetools.2014.258183	No
Probabilistic Programming Process Algebra	Anastasis Georgoulas (with Jane Hillston, Dimitrios Milios and Guido Sanguinetti)	QEST 2014: Quantitative Evaluation of Systems	LNCS 8657	Springer	Berlin, Heidelberg	2014	pp. 249–264	http://doi.org/10.1007/978-3-319-10696-0_21 and http://bit.ly/2p6cVkt	Yes
PALOMA: A Process Algebra for Located Markovian Agents	Cheng Feng (with Jane Hillston)	QEST 2014: Quantitative Evaluation of Systems	LNCS 8657	Springer	Berlin, Heidelberg	2014	pp. 265–280	http://doi.org/10.1007/978-3-319-10696-0_22 and http://bit.ly/2pZd0LE	Yes
An Analysis Pathway for the Quantitative Evaluation of Public Transport Systems	Stephen Gilmore (with Mirco Tribastone and Andrea Vandin)	IFM 2014: Integrated Formal Methods	LNCS 8739	Springer	Berlin, Heidelberg	2014	pp. 71–86	http://doi.org/10.1007/978-3-319-10181-1_5 and http://bit.ly/2p69J8s	Yes
Revisiting Trace and Testing Equivalences for Nondeterministic and Probabilistic Processes	Marco Bernardo (with Rocco De Nicola and Michele Loreti)	Logical Methods in Computer Science	Volume 10, Issue 1 (March 3, 2014) lmcs:1137	episciences.org		2014	pp. 1—42	http://doi.org/10.2168/LMCS-10(1:16)2014	Yes
Hybrid semantics for Bio-PEPA	Vashti Galpin	Information and Computation	Volume 236, August 2014	Elsevier	Amsterdam	2014	pp. 122–145	https://doi.org/10.1016/j.ic.2014.01.016 and http://bit.ly/2pISwps	Yes

Title	Main author	Title of the periodical or the series	Number, date or frequency	Publisher	Place of publication	Year of publication	Relevant pages	Permanent identifiers	Is/will open access be provided for this publication
A computational approach to steady-state convergence of fluid limits for Coxian queuing networks with abandonment	Max Tschaikowski (with Mirco Tribastone)	Annals of Operations Research	May 2017, Volume 252, Issue 1	Springer	Berlin, Heidelberg	2017	pp. 101–120	http://doi/org/10.1007/s10479-016-2193-5 and http://bit.ly/2p2dsJb	Yes
Are Mean-field Games the Limits of Finite Stochastic Games?	Josu Doncel (with Nicolas Gast and Bruno Gaujal)	ACM SIGMETRICS Performance Evaluation Review	Volume 44 Issue 2, September 2016	ACM	New York, USA	2016	pp. 18—20	http://doi.org/10.1145/3003977.3003984 and https://hal.inria.fr/hal-01321020	Yes
Asymptotically Exact TTL-Approximations of the Cache Replacement Algorithms LRU(m) and h-LRU	Nicolas Gast (with Benny Van Houdt)	ITC 28: 28th International Teletraffic Congress	9 January 2017	IEEE	New York, USA	2017	pp. 157—165	http://doi.org/10.1109/ITC-28.2016.128 and https://hal.inria.fr/hal-01292269	Yes
From Featured Transition Systems to Modal Transition Systems with Variability Constraints	Maurice H. ter Beek (with Ferruccio Damiani, Stefania Gnesi, Franco Mazzanti and Luca Paolini)	SEFM'15: Software Engineering and Formal Methods	LNCS 9276, 21 August 2015	Springer	Berlin, Heidelberg	2015	pp. 344–359	http://doi.org/10.1007/978-3-319-22969-0_24 and http://bit.ly/2pOsPoo	Yes
Using FMC for family-based analysis of software product lines	Maurice H. ter Beek (with Alessandro Fantechi, Stefania Gnesi and Franco Mazzanti)	SPLC '15: International Conference on Software Product Lines	July 20–24, 2015	ACM	New York, USA	2015	pp. 432—439	http://doi.org/10.1145/2791060.2791118 and http://bit.ly/2oPk39P	Yes
Model Checking Markov Population Models by Central Limit Approximation	Luca Bortolussi (with Roberta Lanciani)	QEST 2013: Quantitative Evaluation of Systems	LNCS 8054	Springer	Berlin, Heidelberg	2013	pp. 123–138	http://doi.org/10.1007/978-3-642-40196-1_9 and http://bit.ly/2qsXfwr	Yes
Bounds on the deviation of discrete-time Markov chains from their mean-field model	Luca Bortolussi (with Richard A. Hayden)	Performance Evaluation	Volume 70, Issue 10, October 2013	Elsevier	Amsterdam	2013	pp. 736–749	https://doi.org/10.1016/j.peva.2013.08.012 and http://bit.ly/2pyaLO3	Yes

Title	Main author	Title of the periodical or the series	Number, date or frequency	Publisher	Place of publication	Year of publication	Relevant pages	Permanent identifiers	Is/will open access be provided for this publication
On-the-fly Fast Mean-Field Model-Checking	Diego Latella (with Michele Loreti and Mieke Massink)	TGC 2013: Trustworthy Global Computing	LNCS 8358	Springer	Berlin, Heidelberg	2013	pp. 297–314	http://doi.org/10.1007/978-3-319-05119-2_17 and http://bit.ly/2p6BKwH	Yes
Stochastic Process Algebra and Stability Analysis of Collective Systems	Luca Bortolussi (with Diego Latella and Mieke Massink)	COORDINATION 2013: Coordination Models and Languages	LNCS 7890	Springer	Berlin, Heidelberg	2013	pp. 1–15	http://doi/org/10.1007/978-3-642-38493-6_1	No
Reachability Computation for Switching Diffusions: Finite Abstractions with Certifiable and Tuneable Precision	Luca Laurenti (with Alessandro Abate, Luca Bortolussi, Luca Cardelli, Milan Ceska and Marta Kwiatkowska)	HSCC '17: Conference on Hybrid Systems: Computation and Control	April 18 – 20, 2017	ACM	New York, USA	2017	pp. 55–64	http://doi.org/10.1145/3049797.3049812 and http://bit.ly/2pEp2J5	Yes
Construction of Lyapunov Functions via Relative Entropy with Application to Caching	Nicolas Gast	ACM SIGMETRICS Performance Evaluation Review	Volume 44 Issue 2, September 2016	ACM	New York, USA	2016	pp. 6—8	http://doi.org/10.1145/3003977.3003980 and https://hal.inria.fr/hal-01321017v2	Yes
Bisimulation of Labelled State-to-Function Transition Systems Coalgebraically	Diego Latella (with Mieke Massink and Eric P de Vink)	Logical Methods in Computer Science	Volume 11, Issue 4. December 22, 2015.	episciences.org		2015	pp. 23–43	https://arxiv.org/abs/1511.05866	Yes
Using a Machine Learning Approach to Implement and Evaluate Product Line Features	Davide Bacciu (with Stefania Gnesi and Laura Semini)	WWV'15: Workshop on Automated Specification and Verification of Web Systems	EPTCS 188, 2015	episciences.org		2015	pp. 75-83	http://dx.doi.org/10.4204/EPTCS.188.8	Yes
Exploring Spatio-temporal Properties of Bike-sharing Systems	Vincenzo Ciancia (with Diego Latella, Mieke Massink and Rytis Paškauskas)	SCOPES'15: Workshop on Spatial and COllective PErvasive Computing Systems	29 October 2015	IEEE	New York, USA	2015	pp. 74—79	http://doi.org/10.1109/SASOW.2015.17	No

Title	Main author	Title of the periodical or the series	Number, date or frequency	Publisher	Place of publication	Year of publication	Relevant pages	Permanent identifiers	Is/will open access be provided for this publication
Stochastic Approximation of Global Reachability Probabilities of Markov Population Models	Luca Bortolussi (with Roberta Lanciani)	EPEW 2014: Computer Performance Engineering	LNCS 8721	Springer	Berlin, Heidelberg	2014	pp. 224–239	http://doi.org/10.1007/978-3-319-10885-8_16 and http://bit.ly/2qkNWi5	Yes
Software product line analysis with mCRL2	Maurice H. ter Beek (with Erik P. de Vink)	SPLC '14: International Conference on Software Product Lines	September 15–19, 2014	ACM	New York, USA	2014	pp. 78–85	http://doi.org/10.1145/2647908.2655970 and http://bit.ly/2qkEjjC	Yes
Towards Modular Verification of Software Product Lines with mCRL2	Maurice H. ter Beek (with Erik P. de Vink)	ISoLA 2014: Leveraging Applications of Formal Methods, Verification and Validation. Technologies for Mastering Change	LNCS 8802	Springer	Berlin, Heidelberg	2014	pp. 368–385	http://doi.org/10.1007/978-3-662-45234-9_26 and http://bit.ly/2qrjT8z	Yes
Performance Analysis of Collective Adaptive Behaviour in Time and Space	Cheng Feng (with Marco Gribaudo and Jane Hillston)	Electronic Notes in Theoretical Computer Science	ENTCS 318	Elsevier	Amsterdam	2015	pp. 53—68	http://dx.doi.org/10.1016/j.entcs.2015.10.019 and http://bit.ly/2qxVUSe	Yes
Distributed statistical analysis of complex systems modeled through a chemical metaphor	Danilo Pianini (with Stefano Sebastio snd Andrea Vandin)	HPCS'14: High Performance Computing and Simulation	22 September 2014	ACM	New York, USA	2014	pp. 416—423	http://doi.org/10.1109/HPCSim.2014.6903715	No
Transient and Steady-state Regime of a Family of List-based Cache Replacement Algorithms	Nicolas Gast (with Benny van Houdt)	ACM SIGMETRICS Performance Evaluation Review	Volume 43 Issue 1, June 2015	ACM	New York, USA	2015	pp. 123–136	http://doi.org/10.1145/2796314.2745850 and https://hal.inria.fr/hal-01143838/	Yes
Towards a Feature mu-Calculus Targeting SPL Verification	Maurice H. ter Beek (with Erik P. de Vink and Tim A. C. Willemse)	FMSPLE'16: Formal Methods and Analysis in Software Product Line Engineering	EPTCS 206, 2016	episciences.org		2016	pp. 61–75	https://arxiv.org/abs/1604.00350	Yes

Title	Main author	Title of the periodical or the series	Number, date or frequency	Publisher	Place of publication	Year of publication	Relevant pages	Permanent identifiers	Is/will open access be provided for this publication
The Power of Two Choices on Graphs: the Pair-Approximation is Accurate	Nicolas Gast	Performance Evaluation Review	Vol. 43, No. 2, September 2015	ACM	New York, USA	2015	pp. 69—71	http://doi.org/10.1145/2825236.2825263 and https://hal.inria.fr/hal-01199271	Yes
Transient and Steady-state Regime of a Family of List-based Cache Replacement Algorithms	Nicolas Gast (with Benny van Houdt)	ACM SIGMETRICS Performance Evaluation Review	Volume 43 Issue 1, June 2015	ACM	New York, USA	2015	pp. 123–136	http://doi.org/10.1145/2745844.2745850 and https://hal.inria.fr/hal-01143838	Yes
A Definition Scheme for Quantitative Bisimulation	Diego Latella (with Mieke Massink and Eric P de Vink)	QAPL'15: Workshop on Quantitative Aspects of Programming Languages and Systems	EPTCS 194, 2015	episciences.org		2015	pp. 63–78	http://doi.org/10.4204/EPTCS.194.5	Yes
Quantitative Analysis of Probabilistic Models of Software Product Lines with Statistical Model Checking	Maurice H. ter Beek (with Axel Legay, Alberto Lluch Lafuente and Andrea Vandin)	FMSPLE'15: Workshop on Formal Methods and Analysis in SPL Engineering	EPTCS 182, 2015	episciences.org		2015	pp. 56–70	http://doi.org/10.4204/EPTCS.182.5	Yes
A sequential hypothesis test based on a generalized Azuma inequality	Daniël Reijsbergen (with Werner Scheinhardt and Pieter-Tjerk de Boer)	Statistics and Probability Letters	Volume 97, February 2015	Elsevier	Amsterdam	2014	pp. 192–196	https://doi.org/10.1016/j.spl.2014.11.018 and http://bit.ly/2oQikAY	Yes
Data Verification for Collective Adaptive Systems: Spatial Model-Checking of Vehicle Location Data	Vincenzo Ciancia (with Stephen Gilmore, Diego Latella, Michele Loreti and Mieke Massink)	SASOW'14: Self-Adaptive and Self-Organizing Systems Workshops	8–12 Sept. 2014	ACM	New York, USA	2014	pp. 32—37	http://doi.org/10.1109/SASOW.2014.16 and http://bit.ly/2pHceSC	Yes
Modelling Residential Smart Energy Schemes	Vashti Galpin	SASOW'14: Self-Adaptive and Self-Organizing Systems Workshops	8–12 Sept. 2014	ACM	New York, USA	2014	pp. 49—54	http://doi.org/10.1109/SASOW.2014.19 and http://bit.ly/2qxDbpT	Yes

Title	Main author	Title of the periodical or the series	Number, date or frequency	Publisher	Place of publication	Year of publication	Relevant pages	Permanent identifiers	Is/will open access be provided for this publication
Formal Punctuality Analysis of Frequent Bus Services Using Headway Data	Daniël Reijsbergen (with Stephen Gilmore)	EPEW 2014: Computer Performance Engineering	LNCS 8721	Springer	Berlin, Heidelberg	2014	pp. 164–178	http://doi.org/10.1007/978-3-319-10885-8_12 and http://bit.ly/2pvVBJa	Yes
Stochastically timed predicate-based communication primitives for autonomic computing	Diego Latella (with Michele Loreti, Mieke Massink and Valerio Senni)	QAPL'14: Workshop on Quantitative Aspects of Programming Languages and Systems	EPTCS 154, June 2014	episciences.org		2014	pp. 1-16	http://doi.org/10.4204/EPTCS.154.1	Yes
Learning and designing stochastic processes from logical constraints	Luca Bortolussi (with Guido Sanguinetti)	QEST 2013: Quantitative Evaluation of Systems	LNCS 8054	Springer	Berlin, Heidelberg	2013	pp. 89–105	http://doi.org/10.1007/978-3-642-40196-1_7 and http://bit.ly/2p9qi4p	Yes
Differential analysis of interacting automata with immediate actions	Luca Bortolussi (with Mirco Tribastone)	ValueTools '13: Performance Evaluation Methodologies and Tools	Torino, Italy. December 10–12, 2013	ACM	New York, USA	2013	pp. 49–58	http://doi.org/10.4108/icst.valuetools.2013.254364 and http://bit.ly/2p6CH8n	Yes
Patch-based Hybrid Modelling of Spatially Distributed Systems by Using Stochastic HYPE - ZebraNet as an Example	Cheng Feng	QAPL'14: Workshop on Quantitative Aspects of Programming Languages and Systems	EPTCS 154, June 2014	episciences.org		2014	pp. 64–77	http://doi.org/10.4204/EPTCS.154.5	Yes
ABC–Fun: A Probabilistic Programming Language for Biology	Anastasis Georgoulas (with Jane Hillston and Guido Sanguinetti)	CMSB'13: Computational Methods in Systems Biology	LNCS 8130	Springer	Berlin, Heidelberg	2013	pp. 150–163	http://doi.org/10.1007/978-3-642-40708-6_12 and http://bit.ly/2pL0c8W	Yes

Title	Main author	Title of the periodical or the series	Number, date or frequency	Publisher	Place of publication	Year of publication	Relevant pages	Permanent identifiers	Is/will open access be provided for this publication
A Quantitative Approach to the Design and Analysis of Collective Adaptive Systems	Luca Bortolussi (with Rocco De Nicola, Nicolas Gast, Stephen Gilmore, Jane Hillston, Mieke Massink and Mirco Tribastone)	1st Workshop on Fundamentals of Collective Adaptive Systems	ECAI	ACM	New York, USA	2013	pp. 1—2	http://bit.ly/2pKD7TN	Yes
Coherent branching feature bisimulation	Tessa Belder, Maurice H. ter Beek and Erik P. de Vink	FMSPLE'15: Workshop on Formal Methods and Analysis in SPL Engineering	EPTCS 182, 2015	episciences.org		2015	pp. 14–30	http://doi.org/10.4204/EPTCS.182.2	Yes
Temporal Logic Based Monitoring of Assisted Ventilation in Intensive Care Patients	Sara Bufo (with Ezio Bartocci, Guido Sanguinetti, Massimo Borelli, Umberto Lucangelo and Luca Bortolussi)	ISoLA 2014: Leveraging Applications of Formal Methods, Verification and Validation. Specialized Techniques and Applications	LNCS 8803	Springer	Berlin, Heidelberg	2014	pp. 391–403	http://doi.org/10.1007/978-3-662-45231-8_30 and http://bit.ly/2oPcfEZ	Yes
Data-Driven Statistical Learning of Temporal Logic Properties	Ezio Bartocci (with Luca Bortolussi and Guido Sanguinetti)	FORMATS 2014: Formal Modeling and Analysis of Timed Systems	LNCS 8711	Springer	Berlin, Heidelberg	2014	pp. 23—37	http://doi.org/10.1007/978-3-319-10512-3_3 and http://bit.ly/2pn8vZT	Yes
Using mCRL2 for the analysis of software product lines	Maurice H. ter Beek (with Erik P. de Vink)	FormaliSE 2014: Workshop on Formal Methods in Software Engineering	June 03, 2014	ACM	New York, USA	2014	pp. 31—37	http://doi.org/10.1145/2593489.2593493 and http://bit.ly/2pB2Vnn	Yes
On the Robustness of Temporal Properties for Stochastic Models	Ezio Bartocci (with Luca Bortolussi, Laura Nenzi and Guido Sanguinetti)	HSB'13: Workshop on Hybrid Systems and Biology	EPTCS 125, September 2013	episciences.org		2013	pp. 3–19	http://dx.doi.org/10.4204/EPTCS.125.1	Yes
Interaction and Observation: Categorical Semantics of Reactive Systems Through Dialgebras	Vincenzo Ciancia	CALCO 2013: Algebra and Coalgebra in Computer Science	LNCS 8089	Springer	Berlin, Heidelberg	2013	pp. 110-125	http://doi.org/10.1007/978-3-642-40206-7_10 and https://arxiv.org/abs/1209.5903	Yes

Title	Main author	Title of the periodical or the series	Number, date or frequency	Publisher	Place of publication	Year of publication	Relevant pages	Permanent identifiers	Is/will open access be provided for this publication
Combining declarative and procedural views in the specification and analysis of product families	Maurice H. ter Beek (with Alberto Lluch Lafuente and Marinella Petrocchi)	FMSPLE'13: Formal Methods and Analysis in Software Product Line Engineering	August 26–30, 2013	ACM	New York, USA	2013	pp. 10—17	http://doi/org/10.1145/2499777.2500722 and http://bit.ly/2qwBbNT	Yes

List of Dissemination Activities

YEAR 1

On April 4, 2013, Jane Hillston gave a presentation which included information about the QUANTICOL project at the Edinburgh International Science Festival in the event E2: The Edinburgh Enlightenment Exchange: Futureshock. The audience was made up of members of the general public and there were approximately 100 people present.

On May 15, 2013, Mieke Massink posted a brief description of QUANTICOL to the blog of the related EU-IP ASCENS project: http://www.ascens-ist.eu/

In June 2013, Stephen Gilmore and Luca Bortolussi attended the FoCAS coordination action in Barcelona. All FoCAS sites were represented. There were 35 delegates, all academic. Stephen presented a short introduction to the project which was featured on the FoCAS website at http://focas.eu/focas-inter-project-meeting-2013/. Stephen and Luca participated in workshop activities with the other project members. Stephen recorded a short video interview which is available on the FoCAS website at http://focas.eu/stephen-gilmore/.

On June 3, 2013, Mieke Massink gave a brief presentation of the QUANTICOL project at the 15th International Conference on Coordination Models (COORDINATION 2013) that took place in Florence as part of the 8th International Federated Conference on Distributed Computing Techniques (DisCoTec 2013). There were appr. 40 international delegates, mostly academic and mostly from the Formal Methods community.

On June 25, 2013, Jane Hillston gave a brief presentation of the QUANTICOL project to the LFCS Lab Lunch meeting. There were approximately 30 people in the audience, predominantly from the University of Edinburgh but several from other universities in Scotland.

On June 22, 2013, Luca Bortolussi gave an invited lecture about Fluid Model Checking techniques, partially developed within QUANTICOL, at the Bertinoro summer school on Formal Methods (http://www.sti.uniurb.it/events/sfm13ds/). There were approximately 30 people in the audience, predominantly PhD students coming from all over Europe.

On July 17, 2013, Luca Bortolussi gave an invited talk at the annual meeting of the APS INFORMS society (applied probability society) in San José, Costa Rica, on fluid model checking (http://aps2013.eie.ucr.ac.cr/). There were approximatively 25 people in the audience, coming from the APS academic community, which is not much awareness of model checking techniques.

On 31st August, 2013, Jane Hillston gave an invited talk at TGC in Buenos Aires which was essentially about the QUANTICOL project. There were approximately 30 people in the academic audience from around the world.

On 2nd September 2013, Mirco Tribastone gave a presentation about the QUANTICOL project at the 1st FoCAS workshop on the Fundamentals of Collective Adaptive Systems, co-located with ECAL 2013. There were ca 30 participants, some of whom were not related to FoCAS.

On 1st October 2013, Stephen Gilmore gave a short presentation about the QUANTICOL project to the CASSTING project, another FoCAS project, at their second project meeting in Aalborg, Denmark. There were 38 delegates at the meeting, including representatives from the industrial partners.

On October 16, 2013, Vincenzo Ciancia gave a presentation which included information about the QUANTICOL project in the Dagstuhl Seminar 13422, "Nominal Computation Theory". The

audience was made up of computer scientists interested in the recently born field of nominal automata, and a preliminary link was established during the presentation between this very abstract field and topological reasoning about (possibly infinite) behaviour of collective systems.

On 5th November 2013, Stephen Gilmore and Daniël Reijsbergen presented the QUANTICOL project at Demofest 2013, an industry/academia meet-up in Glasgow.

On 7th November 2013, Vashti Galpin took part in a panel on FoCAS on multi-disciplinary approaches for smart, green and integrated transport at ICT 2013 in Vilnius, Lithuania.

On 18th November 2013, Maurice ter Beek gave an invited presentation at the Copenhagen Meeting on Variability Analysis at the IT University of Copenhagen, Denmark. This was the kick-off meeting of the Danish project VARIETE (Variability in Portfolios of Dependable Systems). The audience consisted of 15 computer scientists (one from industry, a few visitors from abroad and the participants of the project, predominantly from the IT University of Copenhagen) with affinity to formal variability analysis. The talk included general information about the QUANTICOL project, as well as a report of two recently submitted initial investigations performed by Maurice ter Beek, Stefania Gnesi and others from ISTI into (quantitative) variability analyses of a product family of bike-sharing systems.

On 22nd November 2013, Stefania Gnesi gave an invited seminar at the University of Waterloo, Canada. The audience consisted of 20 computer scientists with affinity to software engineering in general and model checking and software product line engineering in particular, predominantly from the University of Waterloo. This talk included general information about the QUANTICOL project, as well as a report of two recently submitted initial investigations performed by Maurice ter Beek, Stefania Gnesi and others from ISTI into (quantitative) variability analyses of a product family of bike-sharing systems.

On 22nd November 2013 Mieke Massink attended a meeting of FIAB-Pisa (Federazione Italiana Amici della Bicicletta - Italian Federation of Bicycle Friends) where Marco Bertini, responsible for the "Bikes Office" of PISAMO running the shared bikes system in Pisa, gave a presentation on the bike-sharing system in Pisa. The meeting was attended by about 30 people, many of them interested in the new bike-sharing system CicloPi in Pisa that started its activities in May 2013. The idea for attending this meeting from the perspective of the Quanticol project was to get in touch with the user-community of bike-sharing systems and hear about their experiences. In a later stage the idea is to organise a meeting in which some results obtained within the Quanticol project concerning bike-sharing can be presented at a FIAB meeting. The Bike sharing system in Pisa currently consists of 14 bike parking places involving 140 bikes. There are currently (Nov. 3013) 500 bike trips per day of on average 13 minutes each. Two small trucks are used to redistribute the bikes when needed and to repair broken bikes. The system runs 24 hours per day and 7 days a week. The largest user community in Pisa (60%) consists of university students. During the presentation it was also briefly mentioned that there is an ongoing research collaboration with CNR.

On 10th February 2014 Mirco Tribastone (SOTON) gave a seminar at the Dipartimento di Elettronica, Informazione, and Ingegneria of Politecnico of Milano, where he discussed some recent QUANTICOL results about the analysis of large-scale software systems.
On 14th February 2014, Jane Hillston, Allan Clark, Daniël Reijsbergen, and Stephen Gilmore (UEDIN) met members of the Smart Society FoCAS project including Michael Rovatsos to learn about each others' projects and learn about potential areas of collaboration.

YEAR 2

On 12th of April 2014, Nicolas Gast gave an invited talk at the CASSTING workshop on "Efficiency and Prices in Real-Time Electricity Market". The audience was composed of 20 persons, most of them were members of the CASSTING project.

On 3rd June 2014, Maurice ter Beek gave a presentation at FormaliSE 2014, the 2nd FME Workshop on Formal Methods in Software Engineering, affiliated with ICSE 2014, the 36th International Conference on Software Engineering, held in Hyderabad, India. The audience consisted of 30 computer scientists (including some from industry) with affinity to formal methods and software engineering. The talk started with general information about the QUANTICOL project and T3.3 in particular, before turning to the contents of the workshop contribution.

On the 12th of June 2014, Nicolas Gast gave a presentation Impact of Demand-Response on the Efficiency and Prices in Real-Time Electricity Markets, at the conference ACM e-Energy2014, Cambridge, UK.

On 25th June 2014, Maurice ter Beek gave a presentation at PSI 2014, the 9th International Conference on Perspectives of System Informatics, held in St. Petersburg, Russia, on computer science research and applications. The audience consisted of almost 100 computer scientists, including a reasonable number of them from industry and two Turing award winners (C.A.R. Hoare and E.M. Clarke) who are renown experts in process algebra and model checking, respectively. The talk started with general information about the QUANTICOL project and T3.3 in particular, before turning to the contents of the conference contribution.

On 10th of July, 2014, Maurice ter Beek, Luca Bortolussi Vincenzo Ciancia, Stefania Gnesi, Jane Hillston, Diego Latella and Mieke Massink published a brief article on page 32 in the ERCIM News No. 85 in the context of the Special Theme of "Smart Cities". The title of the article is "A Quantiative Approach to the Design and Analysis of Collective Adaptive Systems for Smart Cities" and presents the Quanticol project and some of the work performed at the CNR-ISTI partner during the first year of the project. ERCIM News reaches a broad audience in the European Community in Information Technology and Applied Mathematics with more than 7,000 readers subscribed to the online edition and about 6,000 printed copies. The online edition has excellent visibility with a Google ranking of PR8.

On 23rd July 2014, Jane Hillston gave a PUMA (Program and Model Analysis) seminar at TUM in Munich. The title of the seminar was "Fluid Approximation for the Analysis of Collective Systems". There were approximately 20 people in audience, a mix of PhD students and faculty from TUM and LMU.

On 2nd September 2014, Jane Hillston gave an invited presentation at CONCUR 2014 and IFIP TCS 2014 in Rome. The title of the talk was "The Benefits of Sometimes Not Being Discrete". There were over 100 people in the international audience of theoretical computer scientists.

On Sept. 5, 2014, Stephen Gilmore, Diego Latella, Michele Loreti and Mieke Massink published a brief article (in italian) in ScienzainRete (ScienceontheNet), the ``portal dedicated to Italian research in Europe and in the world. The website is congenial to researchers, politics, journalists, teachers and all those curious about science and technology. On Scienceonthenet it is possible to find news and information about the many Italian research center and their international collaborations, together with articles, video, infographics, features, comments and debates about research policies in Italy and abroad. Scienceonthenet also provides information about European announcements and funds available, an updated database of research-related jobs, press reviews, documents and much more. Scienceonthenet is a project supported by a grant from the Italian Minister for University and Research (MIUR), in a partnership with the National Council of Research (CNR) and the Gruppo 2003". [Facciamo in modo ``smart'' le smart city ("Making Smart Cities Smart")]

On November the 4th, 2014: Nicolas Gast was invited to gave a presentation Asymptotic properties of object-sharing systems, at the international workshop YEQT, Eindhoven

On Nov. 6, 2014 the QUANTICOL Project has been nominated, on proposal by M. Massink and D. Latella, "Focus" of CNR-ISTI and a short description of the project activities so far has been inserted in the official CNR web site, in the CNR FOCUS section: "The Focus page shows the most important researches made by CNR Institutes, told by the researchers themselves. This way, CNR gives its public the chance of knowing deeply the research activity carried out by the Body and of verifying directly the work carried out." FOCUSes are shown in the home page of CNR, on a one per day basis.

On December the 9th, Nicolas Gast gave a presentation Volatility in Real-Time Electricity Markets: efficiency or manipulation, at the Grenoble Laboratory of Electrical Engineering, G2ELab, France. On Dec 15th 2014, Nicolas Gast gave a presentation Impact of Demand-Response on the Efficiency and Prices in Real-Time Electricity Markets, at the university Paris 6, during a CNRS workshop "COS" (contrôle optimal stochastique).

On Dec 15th 2014, Jane Hillston gave a presentation Challenges for Quantitative Analysis of Collective Adaptive Systems, at Dagstuhl seminar 14512, Collective Adaptive Systems: Qualitative and Quantitative Modelling and Analysis. Approximately 30 international researchers were present at the seminar.

On Dec 17th 2014, Mieke Massink gave a presentation On-the-fly Fast Mean Field Model Checking for Collective Adaptive Systems, at Dagstuhl seminar 14512, Collective Adaptive Systems: Qualitative and Quantitative Modelling and Analysis. Approximately 30 international researchers were present at the seminar.

On Dec 18th 2014, Jane Hillston did a video interview about CAS and QUANTICOL for Giacomo Cabri of the FoCAS Coordinating Action for inclusion on the FoCAS website.

On Jan. 16 2015, Daniël Reijsbergen gave a presentation about the QUANTICOL research involving Edinburgh buses at a SICSA workshop in Stirling about modelling and optimisation of real-world transportation problems. The presentation was attended by about 30 people from universities across Scotland, including Emma Hart from the FoCAS coordination team. The presentation focussed on the PASM and EPEW papers from 2014.

on Jan 28th 2015, Jane Hillston gave an invited seminar to the Stochastic Modelling group of the School of Mathematics at the University of Edinburgh, on A Fluid Approach to Model Checking. There were approximately 15 researchers in the audience, staff and PhD students from the School of Mathematics and the Business School of the University of Edinburgh.

On Feb. 24 2015, Daniël Reijsbergen gave a presentation about the QUANTICOL research involving Edinburgh buses (similar to the one given in Stirling) at the FATA Seminar at the University of Glasgow. About 20 researchers were present.

On March 11th 2015, Jane Hillston gave an invited seminar to the Computer Science Department at the University of York, on "Fluid Approximation for the Analysis of Collective Systems". There were about 25 people in the audience, a mix of PhD students and faculty, including some members of YCCSA (York Centre for Complex Systems Analysis).

YEAR 3

On April 9 2015, Nicolas Gast gave an invited talk at the "French Brazilian Spring School On Big Data and Smart Cities" on "Stochastic modeling, mean-field and smart cities". There was about 30 participants in the audience (faculty and PhD students).

On April 22nd 2015, Jane Hillston gave an invited seminar to the Department of Electronics and Computer Science at the University of Southampton, on "Fluid Approximation for the Analysis of Collective Systems". There were about 20 people in the audience, a mix of PhD students and faculty.

On April 27th 2015, Jane Hillston gave an invited distinguished lecture to the Department of Informatics at Kings College, London, entitled "Choosing not to be discrete --- The benefits of fluid approximations in dynamic modelling". This was a public lecture with an audience of approximately 35-40 people made up of students and staff from KCL and other institutions, and members of the general public.

On May 1st 2015, an article about QUANTICOL was published in the FoCAS reading room.
On May 7 2015, Nicolas Gast was invited to give a seminar in the math department in Nancy (France). He gave a talk on "mean-field models for large systems with application to bike-sharing systems". There was about 40 person in the audience.

On June 2 2015, Nicolas Gast gave an invited talk at the "Workshop on the Mathematics of Demand Side Management and Energy Storage" at the Open University (Milton Keynes, UK). The title was "are energy market efficient?". The audience was around 50 person (half academic, half from the industry).

On June 11th 2015, Daniël Reijsbergen gave a presentation about the QUANTICOL research involving Edinburgh buses (similar to those given in Glasgow and Stirling) at the DACS colloquium at the University of Twente. About 15 researchers from the DACS and SOR groups were present.

On 25th August 2015, Jane Hillston gave a keynote presentation at the Ershov Informatics Conference (PSI'15) in Kazan, Russia, entitled "Quantitative Analysis of Collective Adaptive Systems". After a general introduction to the area this talk introduced the CARMA modelling language. There were about 50 people present, PhD students and staff from Universities around the world.

On 21st September 2015, Rytis Paskauskas gave a presentation entitled Exploring Spatio-temporal Properties of Bike-sharing Systems at the SCOPES 2015 workshop of the SASO 2015 conference at MIT, Cambridge, MA, USA. The presentation was attended by about 30 people. It illustrated recent results on combining spatial model-checking techniques with analysis of simulation traces containing spatial information in the context of the bike-sharing case study.
On 25th September 2015, Vincenzo Ciancia gave a presentation entitled "La bici condivisa (The shared bike)" for a general public audience in the context of the European "Night of Researchers" event called BRIGHT at the Area della Ricerca of CNR in Pisa. The event attracted hundreds of visitors of all ages from Pisa and the surroundings, including many children. The presentation was one of the many interdisciplinary activities illustrating research at CNR as shown by the rich programme. The topic of the presentation was the bike sharing case study that is part of the Quanticol project and was used to illustrate some of the recent results obtained.

From the 4th of October to the 10th of October 2015 Rytis Paskauskas attended the Summer School on "Large Fluctuations and Extreme Events - Theory and Applications" in Dresden, for which he had applied. He has presented work on bike sharing disseminating the results that have been previously presented at the ITSC15 and the SCOPES15 conferences. Title of the presentation: "Model-based assessment of Aspects of User satisfaction in bicycle sharing systems"

(Oct. 5, 2015). The presentation was attended by 26 international participants and additional Lecturers of the school.

On 15th February 2016, Jane Hillston gave an invited seminar to the Programming Languages and Systems research group in the School of Computing at the University of Kent, entitled "Quantitative Analysis of Collective Adaptive Systems". There were approximately 16 people in the audience, a mix of PhD students, postdoctoral researchers and faculty.

YEAR 4

On 6th April 2016, Jane Hillston took part in a sustainable transport event organised by the charity Sustrans and the Edinburgh City Council which took place at the Edinburgh Centre for Carbon Innovation. She took the opportunity to talk to several people about the QUANTICOL project, including Sally Kerr (leader of the Digital Team at Edinburgh City Council), Katie Swann (Partnership Development Officer at Edinburgh City Council) and Tracy McKen (Transport Scotland, a branch of the Scottish Government). Tracy McKen in particular was interested in the modelling work that has been done within QUANTICOL and Jane sent her a link to the webpage and several papers the following day.

On 18th April 2016, Jane Hillston gave an invited seminar in the Mathematics Department of the University of Strathclyde, entitled "High-level languages for fluid approximation of agent-based models". There were approximately 25 people in the audience, a mix of PhD students, postdoctoral researchers and (mostly) faculty from the Departments of Mathematics and Computer Science.

On 11th May 2016, a SICSA-sponsored workshop organised by Daniël Reijsbergen was held in Edinburgh. The theme of the workshop was probabilistic modelling of urban transportation systems; the workshop featured several presentations by QUANTICOL researchers, mostly concerning work done on WP5. A particular focus was on the modelling of bike-sharing systems. The workshop was attended by roughly 25 people, including QUANTICOL researchers from three sites (Edinburgh, Grenoble, Lucca), non-QUANTICOL Informatics researchers from Edinburgh, researchers from other Scottish universities (including Heriot Watt and the University of the West of Scotland) and several people from industry (including Scottish Enterprise and Transport for Edinburgh). The workshop featured two invited speakers: Oliver O'Brien from UCL and Güneş Erdoğan from the University of Bath.

On 17th of June 2016, Mieke Massink took part in the DeCPS 2016 workshop in Pisa, co-located with the ADA-Europe conference and presented a position paper on the QUANTICOL project focussing on Spatial and Spatio-temporal model checking. In that international workshop on ``Challenges and New Approaches for Dependable and Cyber-Physical Systems Engineering'' the following European research projects were presented: INTO-CPS, U-TEST H2020, AXIOM, Industry 4.0, PROXIMA, CONCERTO, AMASS, ASSUME and QUANTICOL. The workshop provided a good occasion to exchange ideas and for networking purposes. The workshop was attended by approximately 15 to 20 people.

On July 6, 2016, Nicolas Gast gave a tutorial on mean-field methods at the French Summer School of Operations Research. There was about 40 PhD students plus 10-20 faculty.

On 8th of July 2016, Maurice ter Beek and Michele Loreti organized the Workshop on FORmal methods for the quantitative Evaluation of Collective Adaptive SysTems (FORECAST 2016) in Vienna, Austria, as a satellite event of the 4th federated event on Software Technologies: Applications and Foundations (STAF 2016) which includes the 14th International Conference on Software Engineering and Formal Methods (SEFM 2016). Part of QUANTICOL's dissemination plan, FORECAST's primary goal was to raise awareness in the Software Engineering and Formal

Methods communities of the particularities of Collective Adaptive Systems and the design and control problems which they bring. The workshop definitely succeeded in this dissemination task of the project, with workshop participation reaching peaks of over 20 attendants. The proceedings of FORECAST 2016 have been published as volume 217 of the Electronic Proceedings in Theoretical Computer Science (EPTCS). The whole proceedings can be downloaded as PDF or browsed via its DOI.

On 16th July 2016, Vashti Galpin met with Jake Beal of Raytheon BBN Technologies to discuss research carried out in the QUANTICOL project with a particular focus on the CARMA language and tools with applications to aggregate programming and synthetic biology.

On August the 30th 2016, Nicolas Gast gave an invited talk at the "Journées MAS" about electricity markets. The audience was about 50 persons (mostly academic).

On 21st September 2016, Jane Hillston gave an invited talk at the Computational Methods in Systems Biology conference in Cambridge UK on Embedding Machine Learning in Formal Stochastic Models of Biological Processes to an audience of about 100 international researchers at all levels (PhD students to full Professors).

On 15th of September 2016, Diego Latella and Mieke Massink attended a meeting with representatives of ALES-UTRC (United Technologies Research Center), including the General Manager of ALES Alberto Ferrari. The meeting was organised by Prof. Rocco De Nicola of IMT Lucca. Mieke Massink gave a 15 min presentation introducing the CNR-ISTI Team and sketching briefly recent research on spatio-temporal model checking developed within the Quanticol project. Mirco Tribastone presented work on model reduction. Laura Nenzi presented work on her PhD thesis on spatial signal temporal logic that was developed in the context of the Quanticol project at IMT Lucca.

On October 4th, 2016, Nicolas Gast gave an invited talk in the math department of the University of Grenoble about "the use of mean-field method for performance evaluation". There ware approximately 15 people in the audience (all academic)

On 11th October 2016, Vashti Galpin gave a talk about the CARMA modelling language titled from "PEPA to CARMA" at the Lab Lunch of the Laboratory for Foundations of Computer Science in the School of Informatics at the University of Edinburgh. The audience consisted of academics, researchers and doctoral students.

On 13th October 2016, Jane Hillston gave a Collective Adaptive Systems seminar to the group of Scott Smolka at State University of New York - Stony Brook, on Quantitative Analysis of Collective Adaptive Systems. There were approximately 12 researchers in the audience, mostly PhD students and postdocs.

On 26th October 2016, Jane Hillston gave the Logical Structures in Computation Seminar at the Simons Institute, Berkeley University, on Model checking single agent behaviours by fluid approximation. The talk was attended by approximately 20 international researchers participating in the Logical Structures in Computation program, (mostly senior researchers and leaders in their field).

On 19th November 2016, Jane Hillston gave an invited talk at the Christopher Strachey Centenary Celebration in Oxford. The talk was entitled A modelling language approach to defining mathematical structures via semantics, and was attended by approximately 100 researchers (mostly senior academics with some PhD students and postdocs).

On 22th November 2016, Mieke Massink gave a 15 minutes presentation at central CNR in Rome at the yearly Conference of the CNR DIITET Department of Engineering, ICT and Technologies for Energy and Transportation. The 15 minutes presentation was on a selection of the results of the

CNR-ISTI contribution to the QUANTICOL project. The presentation was attended by approximately 150 researchers (mostly senior researchers from CNR institutes nationwide and invited researchers from Industry). The presentation was also broadcasted in streaming. The title of the talk was A Quantitative Approach to Management and Design of Collective and Adaptive Behaviours.

On December 2016, Nicolas Gast gave a talk at Inria (Paris) about "what is the accuracy of mean-field approximation?". The audience was composed of 15 researchers.

On 20th January 2017, a QUANTICOL Meeting took place between ISTI-CNR, PisaMo and BicinCittà at the CNR premises in Pisa organised by Mieke Massink. PisaMo (Ufficio Bici) is the in-house public mobility company of the city of Pisa's administration responsible for the CicloPi bike-sharing system in Pisa (currently ± 200 bikes and 24 stations), supplied and maintained by BicinCittà. The audience consisted of Marco Bertini from PisaMo, Manuela Quario and Marco Giuppone from BicinCittà, and about 10 researchers associated with ISTI-CNR, among which the participants in the QUANTICOL project Mieke Massink, Maurice ter Beek, Stefania Gnesi, Rytis Paskausas and Vincenzo Ciancia. The meeting was centered around the following presentations: Model-based Assessment of Aspects of User-satisfaction in Bicycle Sharing Systems (by Rytis Paskausas), A Tool-chain for Statistical Spatio-temporal Model-checking of Bike Sharing Systems (by Vincenzo Ciancia), Analyzing the Performance of Bike-Sharing Systems (by Maurice ter Beek), and Services for Smart Transportation Systems (by Laura Semini).

On 15th March 2017 Jane Hillston gave a keynote talk at the 10th EAI International Conference on Bio-inspired Information and Communications Technologies. The talk was entitled Studying Smart Cities as Collective Adaptive Systems.

On 22nd March 2017, Vashti Galpin gave a talk titled "CARMA: A quantitative formal language for modelling collective adaptive systems" at the PPar Lunch of the Pervasive Parallelism Centre for Doctoral Training in the School of Informatics at the University of Edinburgh. The audience consisted of academics, researchers and doctoral students.

On 28th March 2017, Maurice ter Beek gave an invited presentation at the VARIETE Closing Workshop at the IT University of Copenhagen, Denmark. This workshop served as the final meeting of the Danish project VARIETE (Variability in Portfolios of Dependable Systems) for which Maurice was also invited to speak at the kick-off meeting in 2013 (see above). The audience consisted of 13 computer scientists (about half of them visitors from abroad and the others participants of the project from ITU Copenhagen) with affinity to formal (variability) modelling and analysis. The talk included material that was presented by Maurice at the QUANTICOL Plenary Meeting in Pisa and which will also be presented at the final review in Lucca in May, mainly based on a FASE'17 paper by Maurice ter Beek et al. on family-based model checking.