QUANTICOL

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

http://www.quanticol.eu

Internal Report 2.1

A unified view of spatial representation and analysis techniques

Revision: 1.0

Author(s): Vashti Galpin (UEDIN), Luca Bortolussi (CNR), Vincenzo Ciancia (CNR), Cheng Feng (UEDIN), Nicolas Gast (INRIA), Jane Hillston (UEDIN), Mieke Massink (CNR), Diego Latella (CNR), Mirco Tribastone (IMT), Max Tschaikowski (IMT)

Publication date: September 30, 2015

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

Abstract

This report presents an overview of spatial representations and analysis techniques within the context of the QUANTICOL project. It first recaps the spatial classification of mathematical models from Deliverable 2.1 and compares this with an abstract data type approach to space proposed for the CARMA language in the technical report TR-QC-01-2015.

Based on the guidelines presented in Deliverable 2.1, together with recent research within the project, the report focuses on two types of discrete-space models: population discrete-space models and individual discrete-space models, both of which involve a graph over locations. In both types, we restrict our focus to time-homogeneous systems where graph structure is static.

When considering analysis techniques, transformations between mathematical models are important as these permit the application of different analysis techniques. The relevant transformations within the QUANTICOL project are aggregation, disaggregation, fluidisation, discretisation and hybridisation. These mostly support abstraction but some can also lead to more detailed models.

The report considers analysis techniques that apply specifically to population discrete-space models, and those that apply specifically to individual discrete-space models. A technique for transforming individual (1-dimensional) continuous-space models to population discrete-space models is presented. Some analyses are applicable to both types of model including spatial model checking, modelling of crowding and analyses that combine techniques and work with hybrid models.

The report concludes by considering whether space is a special attribute, the relevance of the various analysis techniques for QUANTICOL and CARMA, and proposing research topics for the remainder of Task 2.1. These include application of existing results in the context of the QUANTICOL case studies, a location aggregation technique, further investigation of techniques already considered in the project and the possibility of developing hybrid techniques that abstract from parts of the model but consider other parts in a detailed fashion.

Contents

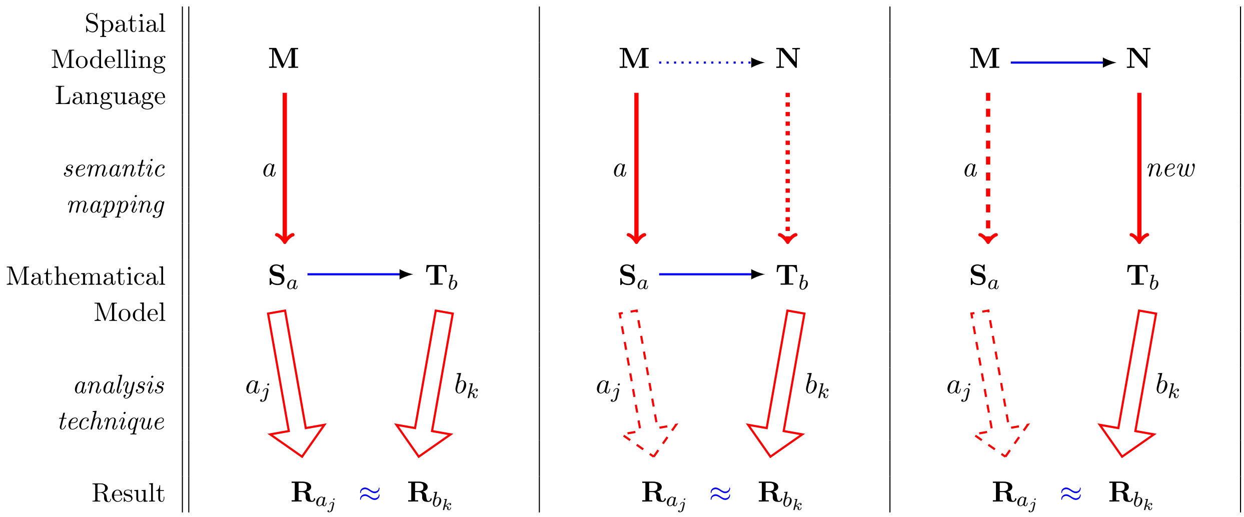

Figure 1: Model transformation: at the level of mathematical model (left), lifted to the level of modelling language (centre), at the level of modelling language

1 Introduction

Space plays an important role in the QUANTICOL case studies and hence analysing and reasoning about space is included in a number of work packages. This document considers the spatial representations that are available and the relevant analysis techniques. By contrast with Deliverable 2.1, this report focusses on analysis techniques for the relevant spatial mathematical models in some detail and makes recommendations for how research directions in Task 2.1 could develop.

We restrict our focus to time-homogeneous systems, thereby excluding the possibility that any behaviour defined is directly dependent on time. Related to this, is the restriction that structure over locations are static. Characteristics of locations may change (as long as this change is not directly time-dependent) but the actual structure that captures the relationships between locations remains static.

An important aspect of the QUANTICOL project is the use of techniques that lead to *scalable* analysis. By this, we mean that when modelling large systems with many components, our analysis can be completed programmatically in a reasonable time (with reasonable memory requirements), and as the system becomes larger, this analysis remains feasible. Concomitant with the scalability is a requirement that any analysis technique that involves approximation remains within reasonable distance from the true value. Obviously, there will be a system size at which the analysis become infeasible. In that case, possible solutions are then to consider whether the model size can be reduced by working with a more abstract model, or to consider an different approximation technique which is more scalable.

An analysis technique is defined with respect to a specific mathematical representation, so an important step in being able to apply a certain technique is transformation from one mathematical model to another. Considering the overall aims of the QUANTICOL project, the workflow involving the CARMA modelling language [CNH+15] can be described at a very high level as follows: develop a CARMA model using the CARMA language, use the CARMA semantics to obtain a mathematical model, perform analysis on this mathematical model and obtain results. The second and third step can possibly be broken down into more detail as follows: choose the appropriate CARMA semantics for the mathematical model required; if necessary, transform this mathematical model to another mathematical model for the analysis required, and then perform the analysis on the mathematical model. These details are shown in the left most diagram in Figure 1. The path from M to S_a to R_{a_j} describes the process without transformation, and the path from M to S_a to T_b to R_{b_k} represents the process with transformation. Frequently, the results obtained after transformation are an approximation to

TIME	continuous			
AGGR	none		state and/or space	
STATE	discrete	continuous	discrete	continuous
SPACE				
discrete	CTMC, IPS [DL94]	TDSHA [BP10] PDMP [Dav93]	patch population CTMC [CLBR09]	patch population ODEs [CLBR09]
continuous	molecular dynamics [CPB08] agents	CTMP [DP03]	spatio-temporal point processes [SBG02]	PDEs [HLBV94]

Figure 2: Classification of mathematical models in terms of time, aggregation, state and space (CTMC: continuous time Markov chain, IPS: interacting particle systems, TDSHA: transition-driven stochastic hybrid automata, PDMP: piecewise deterministic Markov process, ODE: ordinary differential equation, CTMP: continuous-time Markov process, PDE: partial differential equation

the results obtained more directly, as indicated in Figure 1. In this document, we consider analysis techniques and relevant transformations that apply at the mathematical model level. Beyond the scope of this current document are the lifting of mathematical model transformations to language level as illustrated by the central and rightmost diagrams of Figure 1, and they will be addressed elsewhere within the QUANTICOL project.

This report starts with a discussion of spatial representations, comparing the classification of spatial models with the spatial abstract data type proposed for CARMA. Next, the two types of discrete-space model are defined followed by an overview of analysis techniques and transformations of mathematical models. Population discrete-space models are considered in more detail together with the applicable analysis techniques and the same is done for discrete-space models. Transformations between population discrete-space and individual discrete-space models are investigated. Analysis techniques that are applicable to either model or to combinations of models are detailed. The report ends with a section of the relevance of the various analysis techniques to the project and a conclusion. There are number of appendices providing more detail about topics discussed in the report.

2 Spatial representations

In Deliverable 2.1 [GBV+14], a classification of mathematical models was presented which considered spatial aspects together with state and aggregation. A modified version of this appears in Figure 2. The original table distinguished three different types of discrete space: general, regular and homogeneous but here only one category is used. All the models given for discrete space in Figure 2 are general space models apart from IPSs (interacting particle systems) which are defined over grids/lattices which are forms of regular space. As will be seen later in this document, we will consider stochastic models similar to IPSs (interacting particle systems) defined over general space in the form of individual discrete-space models.

Space has also been considered as an abstract data type in the QUANTICOL project [CLM15]. Starting with a definition of closure space which is defined over a set of points and provides a closure operator satisfying three basic axioms, refinements and enrichments can be defined to obtain richer space models. Refinements are obtained by additional axioms and enrichment by adding functions or operators.

Topological space: This can be obtained from closure space by requiring that the closure operator is idempotent.

Graphs: Given a set and a binary relation over this set (which defines an undirected graph), a closure operator can be defined in terms of the binary relation. These closure spaces can be shown to be identical to quasi-discrete closure spaces. Infinite grids over integer values can be defined in this way and they can be enriched by introducing direction functions that allow the description of movement. Furthermore, directed graphs can be obtained by enrichment with pre and post functions.

Metric and distance spaces: Addition of metrics or pseudo-metrics is an enrichment which allows one to assess closeness of points.

Relating this back to the different types of space in Figure 2, continuous space can be obtained from closure space defined over real values in one or more dimensions, that are refined to be topological space and enriched with a metric. When considering discrete space in its most general sense, a graph is the most appropriate representation and hence quasi-discrete closure spaces provide these structures. We can then conclude that these two approaches to spatial representation are mutually consistent.

The preliminary guidelines in Deliverable 2.1 [GBV+14] suggest that population discrete-space models are most appropriate for the QUANTICOL case studies. There may also be a role for individual-based continuous space models in modelling individual bus progress as well as techniques for transforming individual continuous-space models to discrete-space models. Additionally, recent research into individual discrete-space models [Gas15] has relevance to the modelling of bike-sharing, as discussed in Section 6.1.

Thus this document focusses on analysis techniques for discrete-space models for individuals and populations.

3 Discrete-space models

The most general definition of discrete space consists of a finite set of locations and a graph over these locations and this has been identified most appropriate for the QUANTICOL project. We limit our focus to graphs that are static, in the sense that the edges of the graph do not change over time and the number of nodes remain fixed. We consider two different types of discrete-space models as now described.

Population discrete-space models: These are population CTMCs (continuous time Markov chains) where subpopulations in different locations are viewed as separate subpopulations. Appendix B defines both CTMCs and PCTMCs (population CTMCs). The latter is obtained from the former using the counting abstraction which is an aggregation technique and Appendix C describes how space can be added to a population CTMC. These are also called patch-based models and there are various examples in the literature [JL00, AP02, ADH+05, CLBR09, WL95].

The addition of discrete space requires no substantive changes to the definition of a PCTMC and the same analysis techniques can be applied. The main differences are a larger number of subpopulations and a higher likelihood of low population numbers in a particular location; both of which may affect applicability and feasibility of analysis techniques. We assume the location set is finite and the graph does not change over time. We consider only functional rates that are dependent on subpopulation, subpopulation sizes and locations, but exclude direct time dependence, ensuring time-homogeneous CTMCs and other models.

Individual discrete-space models: These are node-labelled graph-based models where each location in the graph represents the state of a single individual at that location or a state of the location. Model dynamics are defined in terms of graph-transformation rules with associated exponential rates that modify node labels (but not edges), thereby providing continuous-time models. Analysis techniques involve counting specific subgraphs of interest over time to understand the proportion of nodes in the graph. The number of locations is finite and the edges do not change over time. The rates can be functional but are not directly time dependent.

In both of these models, rates are functional and there is no specific requirement for them to be continuous, although discontinuities in rate functions may affect the applicability of certain analysis techniques. Later in this document, we consider how one can move between a population discrete-space model and an individual discrete-space model. The challenge is to determine the rates associated with each graph transformation rule from the rates of the population model.

4 Transformations and analysis techniques

As mentioned in the introduction, transformations play an important role in analysing mathematical models. When considering analysis techniques, there are those that can give exact results such as steady state analysis of a CTMC, and those that give approximate results, obtained either by calculating statistics over many stochastic simulations, or by a transformation that allows another analysis technique to be applied. Transformations can be described in terms of what they achieve. Many result in a more abstract model but this is not always the case. In some cases, when referring to aggregation and disaggregation, what is meant is a technique that applies aggregation of a model which is then solved and the solution is disaggregated – we use different meanings here.

Aggregation: This term applies to transformations where states are grouped together to form macro-states. This may or may not involve similar entities being grouped together. An important example of aggregation is the counting abstraction which transforms a Markov chain that considers the state of each entity in the system, into a population Markov chain which considers the number of processes in each state. The resulting Markov chain can have substantially fewer states which makes steady state analysis and simulation faster. This step is required for the fluidisation of population counts. In the case of space, aggregation could also be used to group together locations that have similar characteristics.

Disaggregation: This does not strictly fit into the QUANTICOL workflow described above and is included for completeness and because it is a transformation that relates to the introduction of space to models. This term describes transformations which lead to identical states being distinguished from each other. An example is the addition of a new attribute that can have a number of values. For example, if we assume a collection of entities that have a small set of states, then we can have a mathematical model such as a PCTMC which counts the number of entities in each state. Adding a colour attribute to each entity, and allowing entities to be blue or green leads to a mathematical model where we count the number of entities in each state for each colour. Assuming that only a small number of attributes with few values are added, the new model is likely to still be amenable to the same analysis techniques as the original. Addition of a location attribute can be done in the same way, and a question under consideration in this document is whether there is something intrinsically different about space, that makes addition of a location attribute distinct from that of most other attributes.

Fluidisation: In this case, the transformational aspect is changing from considering the counts of subpopulations (such as the number of processes in a particular state) as integer values, to real values. This can also be described as moving from counts of individuals to a view of a mass of individuals as a flow, hence the term fluidisation. The obvious example of this is using ODEs to describe the average behaviour of a population Markov chain. This reduces the computational cost from simulation of multiple runs of the Markov chain behaviour to the simulation of a single trajectory of each ODE. These ODEs provide the expectation (or first moment) of the population sizes. Related to fluidisation is the derivation of ODEs to describe how the first and higher moments change over time.

Discretisation: This term describes scenarios where some continuous aspect of a model is made discrete. Within the QUANTICOL project, we can consider changing from a continuous space model (for example, $\mathbb{R}^2$) to a discrete-space model where each entity is given a location attribute based on their spatial coordinates. This enables the transformation from a model of individuals moving in two-dimensional space to a model of individuals moving over a graph of locations. If the number of locations is small then this immediately leads to more scalable analysis for those techniques that use the spatial information. Application of aggregation can lead to a population CTMC model over a graph of locations. If this is followed by fluidisation applied to population counts, the final result is a population ODE model over a graph of locations. Each step leads to more scalable analysis techniques.

Hybridisation: This term refers to a model transformation that takes some aspects of a continuous model and discretises them resulting in both continuous and discrete aspects; or alternatively, where fluidisation is applied to aspects of a discrete model, resulting in a model with both continuous and discrete aspects. An example of such a model is a discrete-space model that considers individuals from some subpopulations discretely and other subpopulations in a fluid manner, for example.

We now describe analysis techniques that consider space for the two types of models given above, and mention the transformations involved. For population discrete-space models, we have the following possibilities.

Steady state and transient analysis of PCTMCs: The linear algebra numerical techniques that can be applied to PCTMCs to understand the probability of being in a specific state at steady state, or at a particular time during transient behaviour, can be applied to PCTMCs with locations.

Stochastic simulation of PCTMCs: A different approach to stochastic analysis of PCTMCs is to simulate individual trajectories using an algorithm such as that proposed by Gillespie [Gil07]. A basic assumption is that the model has the property of being well-mixed, that is the entities in the model are evenly distributed throughout space and hence there is no spatial heterogeneity. If sufficient trajectories are simulated, statistical measures can be calculated across all trajectories.

Fluid approximation of PCTMCs: The techniques based on Kurtz's result [Kur81] that express the average behaviour of a PCTMC as ODEs (ordinary differential equations) apply to the fluidisation of a PCTMC with locations. The assumption of well-mixedness also applies.

Spatial moment closure based on spatial averages: This requires fluidisation of the population model, derivation of moment ODEs, and application of an approximation technique to close the moment ODEs.

PDE-analysis of population discrete-space models: A discrete-space population ODE model with random walks is transformed by spatial fluidisation to a PDE (partial differential equation) model which is then analysed using a discretisation most suited to the specific model [TT15b].

Fluidisation and discretisation applied to agent-based continuous space models: This has been applied in an ad hoc manner to a 2-dimensional space model of delay-tolerant networks [Fen14]. A general approach based on Markovian agents has been proposed for 1-dimensional space which aggregates and fluidises individuals and discretises space.

Scale transition theory: This technique provides ODEs which approximate the global population change over time using a growth function that takes the spatial average of environmental parameters as arguments together with the relative population density in each location.

Multi-scale techniques based on difference in rates: Various multi-scale techniques can be applied to a PCTMC defined with locations, or to the equivalent ODE model. Some of these are of general applicability and some are specific to the model characteristics.

The major technique for individual discrete-space models is that of pair approximation, which we will refer to as spatial moment closure based on structure. This technique requires a fluid approximation of the count of specific subgraphs of interest, together with closure techniques. Simulation is also a technique that can be applied to individual discrete-space models, and statistical analysis of multiple trajectories can be done. Furthermore, we can consider spatial model checking, crowding and hybrid techniques in the context of both models. We will consider the applicability of these techniques in the context of Carma in the conclusion.

In the rest of the document, we consider population discrete-space models in more detail together with analysis techniques for these models that take space into account, followed by a similar section for individual discrete-space models, and finishing with analysis techniques that could be useful for both types of model or a combination of model types.

5 Population discrete-space models

As mentioned in the previous section, given a graph over a finite set of locations, it is straightforward to construct a PCTMC model that takes these locations into account. Appendix C describes this process in detail and it can be summarised as follows.

1. For each location and each subpopulation, create a subpopulation at that location. This increases the number of subpopulations by a factor of n where n is the number of locations.

2. Define rates for interaction of subpopulations in a specific location. These are typically the rates for interactions from a non-spatial model, a scaled version of the rates for interaction from a non-spatial model, new rates that are location specific, or some combination of these choices.

3. Define transition rates between locations, with the proviso that movement can only occur between two locations if there is an edge between these locations in the location graph. These rates can be the same for all pairs of locations or they can vary.

4. Define initial subpopulations for each location.

The same exact and approximate analysis techniques can be applied to a PCTMC that uses discrete space as can be applied to a PCTMC. The scalability of these techniques may be affected by the increase in the number of subpopulations and the accuracy of approximation may be reduced if some locations have small subpopulations.

Constructing a PCTMC that takes discrete space into account is essentially the same as adding an extra dimension to the state space which represents location and has one value for each possible location. Each member of the population now has an additional attribute which describes the location of that member. The location graph which determines when movement is possible between two locations plays the same role as deciding what state changes are possible. When considering entities with attributes as we do in QUANTICOL, we can see that introducing discrete space is very similar to introducing a new attribute, with the location graph replaced with a graph describing what changes in attribute value are possible. In this graph, a link between two attribute values implies that an individual with the first attribute value is to be transformed to be an individual with the second attribute value. This is discussed further in Appendix D.

Hence, we might conclude that space can be treated in the same way as an arbitrary attribute in modelling, and in many cases this is indeed the case. We consider this question further in the conclusion of this document.

5.1 Spatial heterogeneity

The issue of spatial homogeneity and inhomogeneity/heterogeneity was addressed formally in Deliverable 2.1 and these concepts are reviewed in Appendix A of the current document. A model is spatially homogeneous if rates (parameters) are independent of the location or locations where they are applied and if the graph of locations is complete.

The population discrete-space model described above can have identical rates for subpopulations inside each location, as well as identical rates for movement between locations, and have a complete location graph, so that every location is adjacent to every other location. In the context of the QUANTICOL project, we wish to consider greater spatial variation than this in our models to capture that characteristics can vary within space, so we need to go beyond these constraints. This involves both variation in rates by location, as well as considering neighbourhoods. A neighbourhood of a location (or set of locations) is a set which is defined in terms of the location graph, and considers the adjacency of other locations to the location of interest. An example is all immediately adjacent locations of a location (including the location itself) which matches the closure operator as defined on quasi-discrete closure spaces [CLM15, GBV^{+}14]. In the case of a complete location graph, any such definition includes all nodes. Hence to work with neighbourhoods in a non-trivial manner, location graphs that are not complete are required.

To model spatial heterogeneity requires capturing something about a location that makes it different to another location. A way in which this can arise is if movement from location to location differs. Considering the simple model above, possibilities are as follows, assuming we now allow parameters associated with edges to be directed (although the edges remain undirected themselves)[1], we have a few choices in how to vary parameters depending on location.

- Allow rates that determine the behaviour of a subpopulation in a location to be dependent on the location.
- Allow rates that determine the movement of a subpopulation between locations to be dependent on the source and/or target location.
- Allow rates that determine movement from a location to be dependent on the number of locations that are adjacent to that location, or on some other characteristic of that location such as centrality.

In the context of analysis of a PCTMC or its related system of ODEs, as discussed above, introducing these variations in rates does not affect the analysis, although it may make the description of the PCTMC more complex. This is because these analyses consider each possible transition (or term in the ODEs) individually and have no way to speed up analysis by considering transitions with the same rate (or identical terms in the ODEs) together (either as a group or to reduce calculation). Techniques such as exact fluid lumpability and related approximation techniques [TT14c, TT15b] identify when it is possible apply an aggregation when dealing with ODEs.

Another issue to consider that relates to spatial heterogeneity is that of dynamic space. We do not consider this in this document for two reasons. The first is that our techniques work well where we have conservation of total population sizes and other quantities, and the second is that spatial structure that changes over time moves us into the area of time inhomogeneity which we have excluded from consideration within this document. Hence, we assume a fixed number of locations, and rates that are not directly dependent on time.

[1]See the discussion in Appendix A to see how this has been defined in previous documents.

5.2 Aggregate moment closure: spatial moment closure based on averages

We now consider existing techniques from the literature referred to as *spatial moment closure* that can abstract from the details of space but still provide a spatially based approach. We suggest new names for the two most relevant techniques. Some techniques are applicable to population discrete-space models and some to individual discrete-space models, as discussed in the next section. To differentiate, we will use the term *aggregate moment closure* for the techniques that are applicable to population discrete-space models because it is more descriptive and because this technique could be applied for attributes other then spatial ones.

In this approach, moment ODEs (derived from the Chapman-Kolmogorov equations) are obtained for averages over all locations (or values for a specific attribute) for various subpopulations. When applied to spatial models, it is a spatial abstraction technique because information about what happens in individual locations is lost. The technique requires the construction of moment ODEs from a discrete-space PCTMC. These ODEs characterise how moments of the subpopulation sizes change over time.

The basic approach is to obtain an ODE for each subpopulation for the ensemble mean[2] of the average over all locations for that subpopulation. This will then (in most cases) be expressed in terms of the expectation of the product of two averages (a higher order moment). The ODE for this can then be derived and this again is likely to contain even higher order moments. In some cases, this system of ODEs eventually becomes closed, in the sense that there are maximal higher order moments that appear in the ODEs. In other cases, the system is not closed (or it is not reasonable to determine whether it is closed), and it can be closed by approximating higher order moments after a certain level. There are four ways to approach this approximation.

- Assume that the higher order moments above this level provide negligible contributions and ignore them by approximating them with zero. A related approach is to assume that higher order cumulants are zero [MKP98, MK99].
- Use the technique of stochastic linearisation which approximates the expectations of products with the product of expectations for higher order moments above this level. It is not sufficient to express second order terms as the product of first order terms as this will exclude effects caused by spatial heterogeneity, hence this technique can only be applied to third and higher order moments [MMRL02]. The modified mean-field approach from ecology takes a similar approach by approximating higher moments with powers of first order moments [PRL11].
- Assume that the data has a particular distribution and use that distribution to determine the values of the higher order moments above this level. The log normal distribution is frequently used because of its positive support which makes it suitable for population modelling [MMRL02, MSH05, KCMG05].
- Apply a Taylor expansion of moments, as used in scale transition theory [Che12].

Most applications of this technique assume a complete graph, or alternatively when neighbourhood is used, approximate the results with those obtained from a complete graph [MSH05]. A slightly different approach is taken by Levin and Durrett where the ODEs for the subpopulation in locations are dependent on the average population values over all locations [LD96]. As above, this then requires the calculation of moment ODEs and moment closure.

Moment closure has been considered in the QUANTICOL project in the context of the PALOMA process algebra [FH14], a formalism that allows the expression of agent location, and whose underlying mathematical model is a population CTMC with locations or population ODEs with locations. Furthermore, ODEs can be derived to express how the moments of the model change over time, such as the expected value of a subpopulation size. As is the case with the moments in Section 5.2 above,

[2] The mean (at time t) over all stochastic realisations (at time t).

typically these ODEs involve higher-order moments, and a technique to close these equations is required. In this research [FHG15], a number of techniques are used to reduce the ODEs so that at most they involve expectations of the product of two subpopulation sizes. Since the models involved can have many subpopulations and locations, this can still require a large number of ODEs to be solved, and a neighbourhood relation is introduced to determine when it is appropriate to approximate the expectation of a product with the product of expectations. The neighbourhood relation is derived from the process algebra model and describes the degree of correlation between two subpopulations taking into account the level of interactions between these subpopulations. Hence, it is not an explicitly spatial technique, although depending on the model, distance between subpopulations may affect the degree of correlation. As we will describe in the section on future research, different neighbourhood relations that are more explicitly spatial can be considered and the quality of their approximations can be compared to the original technique.

5.3 PDE-based analysis of discrete-space models

Tschaikowski and Tribastone [TT15b] have considered an approach which involves taking a discrete model with random walks to continuous space through fluidisation and then using PDE analysis techniques to get good approximation results.

They studied population-based CTMCs where agents are subject to a random walk on the uniform lattice $\mathcal{R} := \{(i\Delta s, j\Delta s) \mid 0 \leq i, j \leq K\}$ in the unit square $[0;1]^2$ with $\Delta s := 1/K$ and $K \geq 1$. Each agent may attain one of the local states $A_1, \ldots, A_L$ while being at any point in $\mathcal{R}$, meaning that the CTMC state

$$\vec{A} := (A_1^{(x,y)}, \ldots, A_L^{(x,y)})_{(x,y)\in\mathcal{R}}$$

provides the agent populations in each local state at each region. Agents in the same region may cooperate with each other by performing local interactions from a rich class of functions. The spatial domain is assumed to have absorbing or reflective boundary conditions. The former can be used to model hostile environment, while the latter account for closed environments.

By setting the initial conditions of the CTMC $(\vec{A}_N(t)/N)_{t\geq 0}$ to $A_l(0) = \lfloor N\alpha_l^0 \rfloor$, where $\alpha_0^l : [0;1]^2 \to [0;\infty)$ are differentiable functions, $1 \leq l \leq L$ and $N \geq 1$, it is then shown that the CTMC of size $\mathcal{O}(N^{L\cdot K^2})$ converges to the solution of an ODE system of size $\mathcal{O}(L \cdot K^2)$ as $N \to \infty$. While this is a major improvement because the complexity drops from exponential to polynomial, the ODE system may be hard to solve if K is large.

Fortunately, it is possible to identify a finite difference scheme [Gea71] which solves the ODE system of size $\mathcal{O}(L \cdot K^2)$ and that can be also interpreted as a finite difference scheme [Tho95] of a PDE system of size L. By combining this with the former result, one then proves that the solutions of the ODE systems of size $\mathcal{O}(L \cdot K^2)$ converge, as $K \to \infty$, to the solution of a PDE system of size L. This is not a purely theoretical result because one solves PDE systems by discretising them to large ODE systems and the discretization induced by a PDE solver is purely dependent on the PDE system itself and thus may be substantially coarser than the one induced by the spatial domain $\mathcal{R}$ which can be arbitrarily fine. Indeed, substantial speed-ups have been reported in [TT14a, TT15b], thus showing that a characterization of mobile systems in terms of PDEs gives rise to shorter calculation times.

5.4 Fluid approximation and spatial discretisation applied to agent-based continuous space models

Feng developed a continuous-space model with individual agents (using the process algebra stochastic HYPE) for a delay-tolerant network which used wild animals as nodes [Fen12]. Due to computational limitations, the analysis was restricted in terms of how many nodes could be modelled. The model was then transformed to a discrete-space model by dividing up space according to waterhole locations, and using the continuous space model to derive parameters for movement [Fen14]. This enabled the population-based modelling of systems with many more nodes and still provided good approximations. However, this is an ad hoc approach.

More recently, a proposal has been made to apply this process in a general way to 1-dimensional space. Specifically, it considers models which consist of Markovian agents (MAs) moving on a bounded one-dimensional continuous space. Markovian agents are a formalism that involves message-passing between agents, and whose overall behaviour can be expressed as a CTMC or a set of ODEs [CGB14, BSB+15]. A detailed definition of Markovian agents is presented in Appendix E. We assume that the movement speed of MAs solely depends on the current state of the agents. This technique is relevant within the QUANTICOL project since bus routes can be viewed as 1-dimensional journeys.

The analysis of interest is the transient evolution of the state density distribution of agents of class c in state i at position ℓ and at time t. The change in this value over a small amount of time can be expressed in terms of those agents at location ℓ who change state and those agents who move to ℓ. The term describing the agents that move can be derived using the Taylor expansion. The change in value can be then be expressed as a PDE in terms of both time and distance (in 1-dimension). Assuming upper and lower bounds, the upwind semi-discretisation technique [HKNT98] can be applied to discretise the distance aspect of the PDE leading to a set of ODEs expressing the change of state density at each discretised location.

5.5 Scale transition theory

Scale transition theory is an approach used within ecology to understand how local dynamics relate to global dynamics, particularly in the case of nonlinear population dynamics [CDMS05, Che12]. It can be applied to either space or time, and is applicable to both continuous and discrete models. Considering this theory in the context of continuous populations and discrete space, the population density for a species X in location i is defined by the ODE

$$dX_i/dt = f(\mathbf{W}_i)X_i$$

where $\mathbf{W}_i$ is a vector of spatially varying fitness factors which can include population densities as well as environmental factors. The average population density (over all locations) can be expressed as

$$d\overline{X}/dt = \overline{f(\mathbf{W})X} = \widetilde{f(\mathbf{W})} \cdot \overline{X} = [\overline{f(\mathbf{W})} + Cov(f(\mathbf{W}), V)] \cdot \overline{X}$$

where $\overline{YZ} = 1/n \sum_{i=1}^{n} Y_i Z_i$ is the average of the product over all locations and $\widetilde{f(\mathbf{W})}$ is the global level fitness and can be expressed as the spatial average of the product of fitness and the relative density of population

$$\widetilde{f(\mathbf{W})} = \overline{f(\mathbf{W})V} = \overline{f(\mathbf{W})} + Cov(f(\mathbf{W}), V)$$

where $V_i = X_i/\overline{X}$ and hence $\overline{V} = 1$. Then using various approximation techniques, the following equation is obtained

$$d\overline{X}/dt = \overline{f(\mathbf{W})X} \approx [(f(\overline{\mathbf{W}}) + 1/2 \cdot f''(\overline{\mathbf{W}})Var(\mathbf{W}) + f'(\overline{\mathbf{W}})Cov(\mathbf{W}, V)] \cdot \overline{X}$$

which provides a straighforward way to approximate the ODE for the change of population at the global level, assuming that f' and f'' are easy to calculate. The authors note that the approximation works best when spatial variation in the arguments to f is low [CDMS05].

5.6 Multi-scale techniques based on differences in rates

As mentioned in the section on spatial heterogeneity, rates can vary, and it may be possible to exploit this variation in the analysis techniques. There are well-known techniques that use differences in interaction rates between entities, such as the Quasi-Steady-State Assumption (QSSA) which assumes an equilibrium for the parts of the system that have fast interaction rates and then derives expressions for the slower parts of the system [GRZ10, SS89]. This can be done both within a stochastic approach and a deterministic approach using ODEs. Another technique is timescale decomposition applied to CMTCs which have the characteristic that its states can be partitioned into groups such that transitions between group members are fast, and transitions between groups are slow. This permits an approximation technique that allows for the CTMC represented by each group of states to be solved separately and then combined into a solution for the whole CTMC [SA61].

In QUANTICOL, such techniques have been considered for model reduction when there are multiple time scales and reported on in Deliverables 1.1 and 1.2 [GBH^{+}14, BGH15] but they have not yet been considered in the context of spatial models. In ecological modelling, spatial aggregation methods consider the combination of different time scales that are location-based [APS12]. Starting with an assumption that interactions that occur at a location are slow and movement between locations is fast, the usual ODEs for a population model can be derived, consisting of terms for migration and terms for local interaction. It is assumed that the terms for migration are multiplied by the inverse of the scale parameter, a value much smaller than 1. This expresses the difference between the fast migration and slow local interaction. Through a change of variables from subpopulation size at a location to a pair consisting of density at a location and total subpopulation over all locations, with a related change in the time variable that divides time by the scale parameter, a slow-fast system can be obtained to which either the quasi-steady-state assumption or Fenichel's theorem [Fen71, Wig94] can be applied to obtain a reduced system.

In the case of QUANTICOL models, it is likely that the pattern will be the opposite as movement between locations is typically physical, whereas interaction within locations may be computer-based and much faster than physical movement.

In some spatial models, there could be a spatial pattern in the variation that could be investigated for the development of alternative analysis techniques. In the examples listed below, the rates are not independent of each other as they vary with respect to any aspect of space, and it may be possible to leverage this fact. Examples are as follows.

Spatial variation depending on location: Consider a grid-structured graph of locations over which there is a gradient. The leftmost locations experience full intensity and the rightmost locations experience half intensity, with a stepped decrease in intensity for locations from left to right. This intensity determines the duration of some action. Assuming a single distribution, we can draw a single duration and then apply the intensity to modify the duration for a specific location. Alternatively, the distribution could be modified for each intensity and then a duration could be drawn from the distribution for the location.

Spatio-temporal variation depending on location: In this case, the pattern occurs over time and space. Assume that an event (that occurs at the end of a duration) happens in all locations but it is time-shifted from left to right, possibly representing that the event (or its effect, as in the case of a sound wave) is moving across the locations, starting from the left. Then the event will occur at time t in the leftmost location, and at time $t + x$ in the rightmost locations where x is the time taken for the event or its effect to move from the leftmost to the rightmost locations. This pattern requires a more expressive model than CTMCs.

Spatial variation depending on location structure: It appears that the only other spatial aspect that could be considered relates to characteristics of the graph. As suggested above, movement rates for movement from a location could be dependent on the number of locations to which subpopulations could move, and in this case the rate would be dependent on the degree of the location in the location graph. If we consider a finite 2-dimensional grid, we could define an overall rate for leaving any location, and the actual rate for moving to a different location would be obtained by dividing the overall rate by the number of locations to which moving was possible. For a non-corner edge location, the divisor would be three, for a corner location, it would be two and for any other location, it would be four.

Rates could also depend on the distance of the location from a special location, or the centrality of the location. Movement rates with are associated which edges could depend on the centrality of the edge, the degree of the vertices of the source location or the target location and other measures that can be defined. Although these can all be seen as different types of spatial variation, it is not clear that there is any suitable application of them.

6 Individual discrete-space models

A different approach to using discrete space is a graph-based model where each node either represents a single individual or a single position in space or location which can have one of a small number of characteristics. Whether the node itself is modelled or an individual at the node is modelled, the node is the agent in the model. Hence there is no distinction between location and agent, unlike in population discrete-space models.

The dynamics of the model are defined in terms of graph transformation rules with associated exponential rates (when working with continuous time). A graph transformation rule describes how a small subgraph or pattern can be transformed in another pattern. There are two types of transformation: those that change the state of the nodes in the graph and those that modify the graph by removing or adding nodes or edges. As described previously, we consider a static model of space and hence we only consider the first type of transformation in this document.

In a graph-based SIR model, each node is an individual who can be in one of a number of states (susceptible, infected, recovered) and the edges of the graph link individuals that can affect each other. This structure can represent space-based interaction of individuals. The graph-transformation rules include a linked pair consisting of one susceptible and one infected being modified to a linked pair consisting two infected; and a infected node being modified to a recovered node.

In ecology, nodes may represent a patch of ground which can be in a number of states including filled by a plant of a specific species, empty but suitable for growth or infertile. Often the nodes are laid out in a grid pattern, and the transformation rules describe how plants spread, and how nodes become fertile or infertile.

The technique called pair approximation which we will refer to as spatial moment closure based on structure or structure-based moment closure (to give it a more explanatory name) provides ODEs which describe how the counts of various patterns change over time. From these values, the number of nodes in each state can be calculated at each point in time. This technique is discussed in more detail in the next section.

6.1 Pair approximation: spatial moment closure based on structure

The stochastic graph transformation model is used to obtain ODEs which describe the change in how often each pattern appears over time. By patterns, we mean small graphs consisting of nodes with states of interest. The reason this technique is called pair approximation is because one can consider the patterns of interest to be a graph consisting of two linked vertices, with the two vertices having specific states, and one wants to know how many times this pattern appears in the graph of the model. Much of the existing research assumes a finite grid/lattice [WKB07, MOSS92], but we consider the more general case of arbitrary graphs rather than regular ones.

Deriving the ODE for a particular pattern may involve understanding how often a different pattern occurs (because the one pattern is transformed into the other by the stochastic process). Typically, to understand the various pair patterns that can occur, the number of certain triplet patterns must be known, and at the next step of obtaining ODEs for triplet patterns, the number of specific quadruplets must be known. This leads to an infinite system of ODEs. Using similar techniques to those described above, this system of ODEs can be closed by approximation. Structure-based moment closure has also been considered as a multi-scale technique [Ell01]. In this case, different sizes of neighbourhood are used for different types of interaction. Structure-based moment closure seems inherently spatial and it is not obvious how it can be applied to other attributes.

We can also examine the quality of the approximation. Consider the two following models where each model is composed of a large number of objects that are connected by a graph of interactions. Objects of the model can interact locally (with their neighbours) or globally.

SIR model on a graph: In this model, we start with a population of nodes that can be susceptible or infected. The disease propagates locally where each infected node infects its susceptible neighbours at a given rate. Infected nodes become immune to the disease at a given rate.

Choices in a bike-sharing model: In a bike-sharing system, bikes are parked at stations and each station can store a finite number of bikes. We assume that when a user performs a trip, she chooses two stations close to her destination and returns her bike to the station that has the least number of parked bikes.

These models are two examples of systems where interaction between local objects are important: for the epidemic model, the disease only propagates locally. For the bike-sharing model, the choices between neighbours induces local interactions between stations. Hence, they provide good candidates to test when structure-based moment closure provides accurate results.

To this end, we simulated the stochastic models corresponding to this model and compared the results with two quantities: (a) a numerical solution of the structure-based moment closure ODEs and (b) a numerical solution of the mean-field ODEs that would correspond to a space-free model of the same systems. For both models, the structure-based moment closure exhibits a behaviour that is closer to the stochastic simulation than the mean-field model. However, for the epidemic model, the structure-based moment closure results are still very far from the stochastic simulation while the structure-based moment closure results are accurate[3] for the bike-sharing example, as shown in [Gas15].

A tentative explanation is the role of initial conditions in the dynamics of such systems. In the considered SIR model, the dynamics of the model is dominated by the transient behaviour of the model and the only propagation of disease is via neighbours. In particular, the stationary distribution is not unique. This might explain why an approximation of the spatial interactions can lead to incorrect prediction. On the contrary, in the bike-sharing example of [Gas15], there are some local interactions that modify the balance of bikes between neighbours, but a large part of the dynamics comes from the movement of bikes throughout the city. The fixed point of the structure-based moment closure is unique and stable and the unique stationary distribution of the stochastic system seems to be close to this fixed point.

7 Transformations between discrete-space models

Individual discrete-space models can be seen both as less abstract and more abstract than population discrete-space models, and this can be illustrated by transformations in both directions. These transformations can be used to obtain smaller models that are amenable to different analysis techniques.

Individual discrete-space models can be less abstract because they consider individuals rather than populations. To support this view, one can consider transforming an SIR individual discrete-space model where the nodes in the graph represent individuals with one at each location, to a population discrete-space model. This transformation involves aggregating locations and thereby aggregating subpopulations. The edges in the aggregated location graph can be determined in a number of ways, but the most obvious is to have an edge between two aggregate locations whenever it is possible to find two individuals that were linked by an edge in the original location graph, such that each individual

[3]Note that the results reported in [Gas15] are numerical evidence that the structure-based moment closure is close to the unique fixed point of the structure-based moment closure equation. However, it can be shown that the fixed point of the structure-based moment closure equation is not the steady-state distribution of the original stochastic system.

is in a different one of each of the two aggregate locations. The resulting population discrete-space model differs from those described earlier in this document because it has no movement between aggregate locations since individual discrete-space models capture how individuals at each location change state rather than move. Since the graph transformation rules describe how interaction with other individuals causes state change for a specific individual, a mechanism must be added to describe how subpopulations in different aggregate locations interact, otherwise the behaviour will only be that of isolated subpopulations that evolve separately. A major challenge of the transformation is to determine the rates for the population discrete-space model so that the two models represent the same behaviour. A secondary challenge is to identify what it means for the two models to have the same behaviour.

An individual discrete-space model can also be more abstract than a population discrete-space model because each location can represent a characteristic of a location. To transform a population discrete-space model to a individual discrete-space model, it is necessary to map each location to a state that describes it. For example, in a model with two subpopulations A and B, a location could be in one of a number of (mutually-exclusive) states, such as *similar sized subpopulations for A and B*, *A is much larger than B*, *B is much larger than A*, *only A*, *only B* and *empty*. As with the previous transformation the challenges are to obtain rates for the transformed model and to determine what it means to have the same behaviour.

8 Analysis technique for both types of discrete-space model

The analysis techniques discussed in this section are not specific to whether a model is an individual or a population model and may also apply to models that have characteristics of both.

8.1 Spatial and spatio-temporal model checking

An important aspect when reasoning about agents or populations *in* space, is the possibility of expressing properties *of* space. For instance, agent/population properties may depend also on space: it might be of interest to be able to express the fact that a certain agent is surrounded (in space) by other agents, or that there is a route from (the position in space of) an agent to (the position in space of) another agent. In addition, it might be useful to describe specific aspects of the space where agents are acting, for example, obstacles. In [CLLM14], a logic for specifying properties of space has been presented which has been developed within the context of QUANTICOL. The logic is also equipped with a spatial model-checking algorithm for the verification of such properties. The approach has been applied to bus transport scenarios [CGL+14] and has been extended to a simple spatio-temporal logic with related model-checking algorithm [CGG+14] which has been applied to bike sharing [CLMP15].

Spatial and spatio-temporal properties are also important for the verification of the spatio-temporal behaviour of complex systems. Examples of such systems are the emerging patterns in animal furs that are a consequence of the local reaction between two different chemical species. Such behaviours can be modelled by reaction-diffusion equations, that, when discretised, lead to a rectangular grid in which each node is associated with the quantities of the chemical substances that are present. These amounts change over time as modelled by the equations and can be seen as signals. These all together form patterns in the grid that develop over time. In this approach, Signal Temporal Logic is extended with spatial operators developed in [CLLM14, NB14] which are capable of specifying such patterns and efficient monitoring algorithms have been developed to analyse such pattern formation.

These two approaches will be described in more detail in the WP3 Internal Report "Scalable verification for spatial-stochastic logics" due at the end of March 2016.

8.2 Spatial simulation with crowding

In biological modelling of cells, crowding is an important part of simulation involving space, because of cells have limited volume and it can be important to consider how much space various molecules take up, and how this may affect reactions, as well as the health of the cell. Models range from those that model continuous space in which each entity has a volume and collision between molecules are explicitly modelled, to grid-based approaches where there is space for only one entity in each location [KK12, TAT05]. The grid-based approaches are similar to individual discrete-space models but use a graph with regularity rather than an arbitrary graph.

For population discrete-space models, crowding can be modelled by imposing maximum quantities on locations. Functional rates for movement into a location can be defined to be zero when the location is full; either because of a maximum population count for a location or a maximum area or volume, where each subpopulation has an associated area for an individual so that occupied area can be calculated. This may lead to discontinuous rate functions.

In the context of the QUANTICOL case studies, crowding is most relevant when considering bus movement at bus stops, since if a bus cannot physically get to a bus stop because of the presence of other buses, it will be delayed in collecting passengers. This could be modelled in discrete space using either population or individual models.

8.3 Hybrid approaches

Hybridness is an ubiquitous feature in many models of real systems. Here, we take a high level view of the concept of a model being hybrid, in the sense that we label as hybrid any model that combines in itself different formalisms or mathematical descriptions or that may involve combinations of techniques. This is in contrast with a classical view of hybridness as the coexistence within a model of discrete and continuous variables.

As far as space is concerned, there are many possible ways in which one can construct hybrid models. Here we list a few, including their relevance to the QUANTICOL project. A full classification of multi-scale and hybrid representations of space is, however, outside the scope of this document.

- Space may be seen or modelled differently depending on which kind of agent we are considering in the model. An example taken from biology is in the description of large and small molecules. The former are often modelled as individual objects having a precise position in continuous space. The latter are described as populations, and hence represented by counting variables, in subregions of space [BHMU11]. This produces a model combining individual objects moving in space with a kind of discretised stochastic diffusion process. A similar combination may be useful in the QUANTICOL project, by considering models of interaction of pedestrians with public transportation. In particular, we can consider a scenario in which buses are modelled as individual entities moving in continuous space, while pedestrians or bus users are modelled as populations moving from one discrete location in the city to another, or on and off a bus. Alternatively, buses outside the city centre could be modelled as moving in continuous space, whereas those within the city centre are modelled as a population with movement rates that are determined by the number of buses.

- Another source of hybridness in spatial modelling can be related to different representations of space at different scales or in different locations. The simplest scenario to consider is a high level representation of space in terms of locations, and a low level description of space inside each location in terms of a grid or as a continuous space. In this case, one has to define appropriate interfaces between the dynamics at the two scales, in terms of abstraction and concretisation functions mapping the low level into the high level and vice versa. By contrast to the previous example, one may wish to model details of the bus movement within the city centre but represent the flow of buses in and out of the centre to different suburbs in a discrete-space style.

- A similar situation to the previous one is a scenario in which one special location of interest is treated in detail, while the rest of the system is approximated in a coarser manner as a single component. This can be helpful in developing some form of spatial fluid model checking. More substantially, the detailed model of a region may be either continuous or grid-based, while the rest of the system can be abstracted as a location-based model, possibly homogeneous, hence resorting to some kind of aggregate moment closure technique.
- Similarly, there may be situations in which different locations require a different level of detail in their treatment. For instance, in a crowd movement scenario, we may be interested in tracking the density of people on bikes in the streets or in a square, which calls for a continuous space representation and a PDE dynamics, but coupling this model with a model describing the number of people at bike stations, in order to keep track of the inflow and outflow of people from the streets or the square.
- From a more classical perspective, we can still imagine hybrid models in space where small and large populations are both present. This may be location specific, and change as the system evolves. Then, we can construct hybrid models in which some populations are kept discrete in some locations, but are approximated continuously in other ones. The research into hybrid limits from Work Package 1 is applicable in this case.

Analysing hybrid spatial models can be challenging, but also opens new ways of using locally different forms of spatial abstraction techniques. As an example, consider a multi-scale scenario where the local space is described as a fine grid, while globally space is represented by a collection of locations. In such a situation, we may use structure-based moment closure approximation locally (if that is accurate enough), de facto reducing the model to a standard location population ODE. In the case of the hybrid treatment of populations, simulation of TDSHA (transition-driven stochastic hybrid automata) [BP10] or PDMPs (piecewise deterministic Markov processes) [Dav93] can be used.

9 Space as an attribute

As mentioned earlier, the issue of whether space is different to any other attribute that could be added to a population-based model is now discussed. In the case of individual discrete-space, it appears that space is very different from adding an arbitrary attribute because it provides the structure over which analysis is done. Considering population discrete-space models, there remains the possibility that for certain models and/or certain modelling techniques, space does provide something richer, because the underlying data type representing space is likely to be equipped with more structure than an arbitrary attribute. For example, in addition to the attribute graph that represents the connectivity of distinct attribute values associated with different locations in space, our notion of space may also be equipped with an implicit metric space which allows us to define not just connectivity, but also "closeness". Thus actions in the model may have a predicate that is based on an influence range rather than direct connection — we have seen this, for example, in SCEL [DLPT14], PALOMA [FH14] and Carma [CNH+15]. Similarly, for these richer notions of space, model reduction techniques or approximate analysis may also be able to exploit the richer structure. The most abstract representation of space, in terms of closure spaces using the basic definition without enrichment or refinement, seems to offer little beyond an arbitrary attribute with an attribute graph, but richer representations of space offer more possibilities.

10 Relevance to QUANTICOL

This section considers the relevance of the various analysis techniques discussed both in the context of the QUANTICOL case studies and the Carma language [CNH+15] for describing collective adaptive systems.

As discussed in Deliverable 2.1, discrete-space models have the most relevance for the QUANTICOL smart transport case studies. Discrete space has also been used in the investigation of residential smart grids [Gal15] where surplus renewable energy is shared with neighbours and a model of power trading in smart grids [ST15].

When considering CARMA, the language that is being developed as part of the QUANTICOL project [CNH+15], we need to understand the mathematical models that provide the semantics of the language. Currently, CARMA only has individual-based semantics. Population-based semantics are still under development. Note that this does not mean that CARMA is limited to individual space models, since many individuals can be located at a single location.

Starting with individual discrete-space models, we consider how they can be expressed in CARMA and how the relevant analysis technique can be applied. In CARMA, locations can be represented by enumerated types ($L1$, $L2$, $\ldots L10$) or by records of basic or enumerated types (such as integer pairs to represent locations on a grid or lattice: $(0,0)$, $(0,1)$, $\ldots(9,9)$ for a 10×10 grid). To represent a location graph, it would seem most suitable to define it at the environment level and let it be part of the global store. Another approach is to define functions that capture the graph structure. Alternatively, each node in the graph could be an agent in the collective and have a record of its neighbours. A further idea is to supply the CARMA language with an abstract data type for closure spaces as discussed in CARMA in Deliverable 4.2 [CNH+15].

Expressing the graph transformation rules required for an individual discrete-model could be done by requiring interaction with neighbours to establish the state of the neighbours and to determine what attribute change is possible for the agent. This would then define the individual discrete-space model, and from this it may be possible to extract the pattern-based ODEs and apply moment closure at the appropriate level thus allowing a structure-based moment closure analysis. This would provide a trajectory for each pattern type, from which trajectories for each subpopulation can be obtained. Note that an individual discrete-space model with graph transformation rules does not fit the CARMA modelling paradigm and hence this is an inelegant solution for combining the two.

Population discrete-space models also require a location graph and the same ideas as discussed above apply again. Population discrete-space models require a population-based semantics for CARMA and it is reasonable to assume that this would result in either a population CTMC or population ODEs with locations. In that case, the model would be such that various techniques could be applied, including aggregate moment closure, PDE-based analysis of discrete-space, scale transition theory or various multi-scale techniques (if the rates have the appropriate characteristics). Note that these techniques mostly provide a global-level analysis of averages over all locations, so they abstract from spatial detail but the results do take account of the affect of spatial heterogeneity on the global outcomes.

CARMA can represent locations as points in continuous space [CNH+15], and it may be possible to represent movement in continuous space in an approximate manner. If individual continuous-space models can be represented in this fashion, then population fluidisation with spatial discretisation could be applied at the level of the mathematical model to achieve a discrete-space-based analysis of the model.

In terms of techniques that can be used for both types of discrete-space model, both spatial model-checking and approaches to expressing crowding should be applicable. For combination of techniques, further development of these techniques are required first. To be able to apply them to CARMA models via the software, a suitable interface is required to identify which parts of the models are to be treated in which manner.

11 Future research

This section proposes potential research topics for the remainder of Task 2.1.

- As mentioned in Section 5.2, aggregate moment closure techniques based on a semantic neighbourhood relation have already been investigated in the context of the process algebra PALOMA

[FH14]. We will also consider neighbourhood relations based on different types of spatial distance and compare the results we obtain.

- If a population discrete-space model becomes too large as a result of many subpopulations and many locations, it may be possible to aggregate locations that are similar in nature. This requires techniques to construct the model with aggregated locations from the original model and to decide how to group locations together. We will investigate whether the research from Work Package 3 on differential aggregation [TT14b] can be used for this clustering task. Differential inclusion [BGH15] and differential hulls [TT15a] could be used to assess the quality of the clustering by comparing upper and lower bounds.

- In Section 6.1 the quality of structure-based moment closure (pair approximation) was discussed and a hypothesis for the different outcome for different models was presented. We will consider whether more definite answers can be obtained when looking at the quality of this approximation.

- We will investigate the application of existing aggregate moment closure techniques as discussed in Section 5.2 to spatial models expressed in Carma.

- If it is possible to identify case studies where there are spatially varying rates so rates within locations are orders of magnitude different than those between locations, then the existing multi-scale techniques mentioned in Section 5.6 as well as those proposed in Work Package 1 [GBH+14, BP14] will be applied to models of the case studies.

- We will consider how to combine approaches to provide a hybrid technique with an initial focus on computing the average (movement) trajectory of an individual within a PDE model or population ODE model with locations. This has similarities to fluid model-checking [BH15] but without nesting of operators should be technically less complex.

- We will consider the development of a hybrid analysis technique that treats one location as distinct and uses aggregate moment closure on the other locations to derive an expression to capture behaviour outside of the distinct region. The research from Work Package 1 on uncertainty [BGH15] may provide techniques for expressing the spatial heterogeneity of the model.

12 Conclusion

This internal report has considered space and location in the context of Task 2.1 "Scalable Representations of Space". The aim of the report is to consider mathematical representations for modelling space and analysis techniques that can be applied to these representations. The report has focussed on discrete-space models as the most relevant to QUANTICOL and considered techniques for population discrete-space models, of which there are a number, and techniques for individual discrete-space models for which there is one main technique. Some techniques such as logic-based approaches apply to both types of model. Hybrid approaches can combine different types of models and be analysed with a combination of techniques.

References

[ADH+05] J. Arino, J.R. Davis, D. Hartley, R. Jordan, J.M. Miller, and P. van den Driessche. A multi-species epidemic model with spatial dynamics. *Mathematical Medicine and Biology*, 22:129–142, 2005.

[AP02] F. Arrigoni and A. Pugliese. Limits of a multi-patch SIS epidemic model. *Journal of Mathematical Biology*, 45:419–440, November 2002.

[APS12] P. Auger, J.C. Poggiale, and E. Sánchez. A review on spatial aggregation methods involving several time scales. *Ecological Complexity*, 10:12–25, June 2012.

[BBC+14] A. Bobbio, D. Bruneo, D. Cerotti, M. Gribaudo, and M. Scarpa. An intelligent swarm of markovian agents. Technical Report TR-INF-2014-06-01-UNIPMN, University of Piedmont, 2014.

[BGH15] L. Bortolussi, N. Gast, and J. Hillston. A framework for hybrid limits under uncertainty, 2015. QUANTICOL Deliverable D1.2.

[BH15] L. Bortolussi and J. Hillston. Model checking single agent behaviours by fluid approximation. *Information and Computation*, 242:183–226, 2015.

[BHLM13] L. Bortolussi, J. Hillston, D. Latella, and M. Massink. Continuous approximation of collective systems behaviour: a tutorial. *Performance Evaluation*, 70(5):317–349, 2013.

[BHMU11] A.T. Bittig, F. Haack, C. Maus, and A.M. Uhrmacher. Adapting rule-based model descriptions for simulating in continuous and hybrid space. In *Proceedings of CMSB 2011*, pages 161–170. ACM, 2011.

[BKHW05] C. Baier, J.-P. Katoen, H. Hermanns, and V. Wolf. Comparative branching-time semantics for markov chains. *Information and Computation*, 200(2):149–214, 2005.

[BP10] L. Bortolussi and A. Policriti. Hybrid dynamics of stochastic programs. *Theoretical Computer Science*, 411:2052–2077, 2010.

[BP14] L. Bortolussi and R. Paškauskas. Mean-field approximation and Quasi-Equilibrium reduction of Markov Population Models. In *Proceedings of QEST 2014*, LNCS 8657, 2014.

[BSB+15] D. Bruneo, M. Scarpa, A. Bobbio, D. Cerotti, and M. Gribaudo. An intelligent swarm of markovian agents. In *Handbook of Computational Intelligence*, pages 1345–1359. Springer, 2015.

[CDMS05] P. Chesson, M.J. Donahue, B.A. Melbourne, and A.L.W. Sears. Scale transition theory for understanding mechanisms in metacommunities. In M. Holyoak, M.A. Leibold, and R.D. Holt, editors, *Metacommunities: spatial dynamics and ecological communities*, pages 279–306. The University of Chicago Press, 2005.

[CGB09] D. Cerotti, M. Gribaudo, and A. Bobbio. Presenting dynamic Markovian agents with a road tunnel application. In *Proceedings of MASCOTS'09*, pages 1–4. IEEE, 2009.

[CGB14] D. Cerotti, M. Gribaudo, and A. Bobbio. Markovian agents models for wireless sensor networks deployed in environmental protection. *Reliability Engineering & System Safety*, 130:149–158, 2014.

[CGG+14] V. Ciancia, S. Gilmore, G. Grilletti, D. Latella, M. Loreti, and M. Massink. Spatio-Temporal Model-Checking of Vehicular Movement in Public Transport Systems, 2014. (Submitted for journal publication).

[CGL+14] V. Ciancia, S. Gilmore, D. Latella, M. Loreti, and M. Massink. Data verification for collective adaptive systems: spatial model-checking of vehicle location data. In *Proceedings of FoCAS Workshp at IEEE@ SASO 2014*, pages 32–37. Computer Society Press, 2014.

[Che12] P. Chesson. Scale transition theory: its aims, motivations and predictions. *Ecological Complexity*, 10:52–68, 2012.

[CLBR09] A. Chaintreau, J.-Y. Le Boudec, and N. Ristanovic. The age of gossip: spatial mean field regime. In *Eleventh International Joint Conference on Measurement and Modeling of Computer Systems (SIGMETRICS/Performance 2009)*, pages 109–120. ACM, 2009.

[CLLM14] V. Ciancia, D. Latella, M. Loreti, and M. Massink. Specifying and Verifying Properties of Space. In J. Diaz, I. Lanese, and D. Sangiorgi, editors, *Theoretical Computer Science (TCS 2014)*, LNCS 8705, pages 222–235. Springer, 2014.

[CLM15] V. Ciancia, D. Latella, and M. Massink. On Space in CARMA. Technical Report TR-QC-01-2015, QUANTICOL, 2015.

[CLMP15] V. Ciancia, D. Latella, M. Massink, and R. Paskauskas. Exploring spatio-temporal properties of bike-sharing systems. In *Proceedings of SCOPES Workshop at IEEE SASO 2014*, 2015.

[CNH+15] V. Ciancia, R. De Nicola, J. Hillston, D. Latella, M. Loreti, and M. Massink. CAS-SCEL semantics and implementation, 2015. QUANTICOL Deliverable D4.2.

[CPB08] E.A Codling, M.J Plank, and S. Benhamou. Random walk models in biology. *Journal of the Royal Society Interface*, 5(25):813–834, 2008.

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

[DL94] R. Durrett and S.A. Levin. Stochastic spatial models: a user's guide to ecological applications. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 343(1305):329–350, 1994.

[DLPT14] R. De Nicola, M. Loreti, R. Pugliese, and F. Tiezzi. A formal approach to autonomic systems programming: The SCEL language. *ACM Transactions on on Autonomous and Adaptive Systems*, 9(2):7:1–7:29, 2014.

[DP03] J. Desharnais and P. Panangaden. Continuous stochastic logic characterizes bisimulation of continuous-time markov processes. *The Journal of Logic and Algebraic Programming*, 56(1-2):99–115, 2003.

[Ell01] S.P. Ellner. Pair approximation for lattice models with multiple interaction scales. *Journal of Theoretical Biology*, 210:435–447, June 2001.

[Fen71] N. Fenichel. Persistence and smoothness of invariant manifolds for flows. *Indiana University Mathematics Journal*, 21:1972, 1971.

[Fen12] C. Feng. *Modelling opportunistic networks with HYPE.* MSc dissertation, School of Informatics, University of Edinburgh, 2012.

[Fen14] C. Feng. Patch-based hybrid modelling of spatially distributed systems by using stochastic HYPE – ZebraNet as an example. In *Proceedings of QAPL 2014*, 2014.

[FH14] C. Feng and J. Hillston. PALOMA: A process algebra for located markovian agents. In *Proceedings of QEST 2014*, LNCS 8657, pages 265–280. Springer, 2014.

[FHG15] C. Feng, J. Hillston, and V. Galpin. Automatic moment-closure approximation of spatially distributed collective adaptive systems. 2015. Submitted for publication.

[Gal15] V. Galpin. Quantitative modelling of residential smart grids. In *Proceedings of MoKMaSD 2015*, 2015.

[Gas15] Nicolas Gast. The power of two choices on graphs: the pair-approximation is accurate. In *Proceedings of MAMA Workshop*. ACM, 2015.

[GBH+14] N. Gast, L. Bortolussi, J. Hillston, R. Paškauskas, and M. Tribastone. Multiscale modelling informed by smart grids, 2014. QUANTICOL Deliverable D1.1.

[GBV+14] V. Galpin, L. Bortolussi, Ciancia V., A. Clark, R. De Nicola, C. Feng, S. Gilmore, N. Gast, J. Hillston, A Lluch-Lafuente, M. Loreti, M. Massink, L. Nenzi, D. Reijsbergen, V. Senni, F. Tiezzi, M. Tribastone, and M. Tschaikowski. A preliminary investigation of capturing spatial information for CAS, 2014. QUANTICOL Deliverable D2.1.

[Gea71] C. William Gear. *Numerical Initial Value Problems in Ordinary Differential Equations*. Prentice Hall PTR, 1971.

[Gil07] D.T. Gillespie. Stochastic simulation of chemical kinetics. *Annual Review of Physical Chemistry*, 58:35–55, 2007.

[GRZ10] A.N. Gorban, O. Radulescu, and A.Y. Zinovyev. Asymptotology of chemical reaction networks. *Chemical Engineering Science*, 65(7):2310–2324, April 2010.

[HKNT98] G. Horton, V.G. Kulkarni, D.M. Nicol, and K.S. Trivedi. Fluid stochastic Petri nets: Theory, applications, and solution techniques. *European Journal of Operational Research*, 105:184–201, 1998.

[HLBV94] E.E. Holmes, M.A. Lewis, J.E. Banks, and R.R. Veit. Partial differential equations in ecology: spatial interactions and population dynamics. *Ecology*, 75(1):17–29, 1994.

[JL00] V.A.A. Jansen and A.L. Lloyd. Local stability analysis of spatially homogeneous solutions of multi-patch systems. *Journal of Mathematical Biology*, 41:232–252, September 2000.

[KCMG05] I. Krishnarajah, A. Cook, G. Marion, and G. Gibson. Novel moment closure approximations in stochastic epidemics. *Bulletin of Mathematical Biology*, 67:855–873, 2005.

[KK12] M. Klann and H. Koeppl. Spatial simulations in systems biology: from molecules to cells. *International Journal of Molecular Sciences*, 13:7798–7827, 2012.

[Kur81] T.G. Kurtz. *Approximation of population processes*. SIAM, 1981.

[LD96] S.A. Levin and R. Durrett. From individuals to epidemics. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 351:1615–1621, 1996.

[MK99] J.H. Matis and T.R. Kiffe. Effects of immigration on some stochastic logistic models: a cumulant truncation analysis. *Theoretical Population Biology*, 56:139–161, 1999.

[MKP98] J.H. Matis, T.R. Kiffe, and P.R. Parthasarathy. On the cumulants of population size for the stochastic power law logistic model. *Theoretical Population Biology*, 53:16–29, 1998.

[MMRL02] G. Marion, X. Mao, E. Renshaw, and J. Liu. Spatial heterogeneity and the stability of reaction states in autocatalysis. *Physical Review E*, 66:051915, November 2002.

[MOSS92] H. Matsuda, N. Ogita, A. Sasaki, and K. Satō. Statistical mechanics of population: The lattice Lotka-Volterra model. *Progress of Theoretical Physics*, 88:1035–1049, 1992.

[MSH05] G. Marion, D.L. Swain, and M.R. Hutchings. Understanding foraging behaviour in spatially heterogeneous environments. *Journal of Theoretical Biology*, 232:127–142, 2005.

[NB14] L. Nenzi and L. Bortolussi. Specifying and monitoring properties of stochastic spatio-temporal systems in signal temporal logic. In *Proceedings of VALUETOOLS 2014*, 2014.

[PRL11] M. Pascual, M. Roy, and K. Laneri. Simple models for complex systems: exploiting the relationship between local and global densities. *Theoretical Ecology*, 4:211–222, 2011.

[SA61] Herbert A. Simon and Albert Ando. Aggregation of variables in dynamic systems. *Econometrica*, 29(2):111–138, 1961.

[SBG02] F.P. Schoenberg, D.R. Brillinger, and P. Guttorp. Point processes, spatial–temporal. In *Encyclopedia of environmetrics*, volume 3, pages 1573–1577. Wiley Online Library, 2002.

[SS89] L.A. Segel and M. Slemrod. The quasi-steady-state assumption: a case study in perturbation. *SIAM review*, 31(3):446–477, 1989.

[ST15] F. Shams and M. Tribastone. Power trading coordination in smart grids using dynamic learning and coalitional game theory. In *Proceedings of QEST 2015*, 2015.

[TAT05] K. Takahashi, S.N.V. Arjunan, and M. Tomita. Space in systems biology of signaling pathways–towards intracellular molecular crowding in silico. *FEBS Letters*, 579:1783–1788, 2005.

[Tho95] J. W. Thomas. *Numerical Partial Differential Equations: Finite Difference Methods.* Springer-Verlag, 1995.

[TT14a] M. Tschaikowski and M. Tribastone. A Partial-differential Approximation for Spatial Stochastic Process Algebra. In *Proceedings of VALUETOOLS 2014*, 2014.

[TT14b] M. Tschaikowski and M. Tribastone. A unified framework for differential aggregations in Markovian process algebra. *Journal of Logic and Algebraic Methods in Programming*, 84:238–258, 2014.

[TT14c] M. Tschaikowski and M. Tribastone. Exact fluid lumpability in markovian process algebra. *Theoretical Computer Science*, 538:140–166, 2014.

[TT15a] M. Tschaikowski and M. Tribastone. Approximate reduction of heterogenous nonlinear models with differential hulls. *IEEE Transactions on Automatic Control*, 2015.

[TT15b] M. Tschaikowski and M. Tribastone. Spatial fluid limits for stochastic mobile networks. *Performance Evaluation*, 2015. Under minor revision.

[Wig94] S. Wiggins. *Normally hyperbolic invariant manifolds in dynamical systems.* Springer Science & Business Media, 1994.

[WKB07] S.D. Webb, M.J. Keeling, and M. Boots. Host-parasite interactions between the local and the mean-field: how and when does spatial population structure matter? *Journal of Theoretical Biology*, 249:140–152, 2007.

[WL95] J. Wu and O.L. Loucks. From balance of nature to hierarchical patch dynamics: a paradigm shift in ecology. *The Quarterly Review of Biology*, 70:439–466, 1995.

A Appendix: Formal definitions from Deliverable 2.1

We assume a set of locations $\mathcal{L}$ and an undirected graph over locations $(\mathcal{L}, E_{\mathcal{L}})$ with $E_{\mathcal{L}} \subseteq \mathcal{P}_2(\mathcal{L})$. Each edge has the form $\{l_1, l_2\}$, and loops such as $\{l, l\}$ are allowed. We can consider two groups of parameters; those that are associated with locations, namely with vertices of the graph and those that are associated with interaction or movement, namely the edges of the graph, and we use the following notation.

- Λ_l for $l \in \mathcal{L}$, and
- Γ_{l_1,l_2} and Γ_{l_2,l_1} for $\{l_1, l_2\} \in E_{\mathcal{L}}$.

Although the edges of the graph are not directed, movement is inherently directional (although in many cases, the movement parameter may have same value for both directions), therefore two sets of parameters for edges are required. Interaction can be undirected when considering an abstract view of interaction or communication. Alternatively, it can be directed if one party is the sender and the other the recipient. Our choice of an undirected graph allows these details to be expressed at the parameter level.

A spatial model is

- *location homogeneous* if $\Lambda_{l_i} = \Lambda_{l_j}$ for all locations $l_i, l_j \in \mathcal{L}$.
- *transfer homogeneous* if $\Gamma_{l_i,l_j} = \Gamma_{l_j,l_i} = \Gamma_{l_{i'},l_{j'}} = \Gamma_{l_{j'},l_{i'}}$ for all edges $\{l_i, l_j\}, \{l_{i'}, l_{j'}\} \in E_{\mathcal{L}}$.
- *(spatially) parameter homogeneous* if it is both location and transfer homogeneous.
- *spatially homogeneous* if it is parameter homogeneous, and its location graph is complete[4].

Spatial inhomogeneity can be introduced in two ways: the first involves connectivity where equal accessibility is no longer assumed, and the second where all locations are still accessible from all locations, but parameters vary between locations. These are not necessarily distinct concepts. Consider the case where there is a parameter $\rho_{i,j} \in \Gamma_{l_i,l_j}$ which describes the rate of movement from location i to location j. If $\rho_{i,j}$ is the same for all i and j and no other parameters vary by location then the model is spatially homogeneous. However, if $\rho_{i,j}$ can vary and possibly be zero then not only does a specific parameter vary by location but additionally, equal accessibility no longer holds (either because on average it takes longer depending on the rate, or if the rate is zero there is no accessibility). However, if $\rho_{i,j}$ is constant for all i and j but other parameters vary by locations, then the model is spatially inhomogeneous.

In an undirected graph of locations representing discrete space, the links between locations are used to define neighbours. Given a location l, its *immediate neighbours* are those vertices l' such that $\{l, l'\}$ is an edge in the graph. Its *n-hop neighbours* are those that can be reached through a path in the location graph of at most n steps (but excluding the location l itself). In the case of a regular grid graph, the immediate neighbours (west, north, east and south) are referred to as the Von Neumann neighbourhood. The larger neighbourhood that includes the northwest, northeast, southeast and southwest points as well as the immediate neighbourhood, is known as the Moore neighbourhood. Both types of neighbourhoods can be extended to n-hop neighbours and also applied to hexagonal and triangular regular location graphs.

This is a purely spatial approach to defining neighbourhoods. However, in some cases, it can be the entity or process itself that defines its neighbourhood depending on its capability. Other approaches use a (perception) function or Boolean formula that determines the neighbours of an individual, or the others with whom an individual can interact, by specifying the attributes or properties of the other individuals with which it can interact, such as the attribute-based communication in CARMA [CNH+15].

[4] A complete undirected graph has an edge $\{l, l'\}$ between each pair of vertices l and l'. Graphs with regularity such as grids are not complete (except for very small examples).

B Appendix: CTMCs and population CTMCs

This section briefly introduces these two concepts, as they would be used in stochastic modelling [BKHW05, BHLM13].

Definition 1. *A* continuous time Markov chain (CTMC) *is a tuple* $\mathcal{M}_C = (\mathcal{S}, \mathbf{R})$ *where*

- $\mathcal{S}$ *is a finite set of* states*, and*
- $\mathbf{R} : \mathcal{S} \times \mathcal{S} \to \mathbb{R}_{\geq 0}$ *is a* rate matrix.

If an entity is currently in state S, then $\mathbf{R}(S, S')$ is a non-negative number that defines an exponential distribution from which the duration of the time taken to transition from state S to state S' can be drawn. Under certain conditions, transient and steady state probabilities can be calculated which describe the probability of being in each state at a particular time t or in the long run, respectively. CTMCs can be state-labelled (usually with propositions) or transition-labelled (usually with actions).

For population Markov chains, instead of considering an entity with states, we consider a vector of counts $\mathbf{X}$ that describes how many entities are in each state; thus it is a population view rather than an individual view.

Definition 2. *A* population continuous time Markov chain (PCTMC) *is a tuple* $\mathcal{X}_C = (\mathbf{X}, \mathcal{D}, \mathcal{T})$ *where*

- $\mathbf{X} = (X_1, \ldots, X_n)$ *is a vector of* variables
- $\mathcal{D}$ *is a countable set of* states *defined as* $\mathcal{D} = \mathcal{D}_1 \times \ldots \times \mathcal{D}_n$ *where each* $\mathcal{D}_i \subseteq \mathbb{N}$ *represents the domain of* X_i
- $\mathcal{T} = \{\tau_1, \ldots \tau_m\}$ *is the set of transitions of the form* $\tau_j = (\mathbf{v}, r)$ *where*
  - $\mathbf{v} = (v_1, \ldots, v_n) \in \mathbb{N}^n$ *is the* state change or update vector *where* v_i *describes the change in number of units of* X_i *caused by transition* τ_j*, and*
  - $r : \mathcal{D} \to \mathbb{R}_{\geq 0}$ *is the* rate function *of transition* τ_j *with* $r(\mathbf{d}) = 0$ *whenever* $\mathbf{d} + \mathbf{v} \notin \mathcal{D}$.

From a population Markov chain, the associated Markov chain can be obtained. $\mathcal{D}$ is the state space $\mathcal{S}$. For the population CTMC, the rate matrix of its associated CTMC is

$$\mathbf{R}(\mathbf{d}, \mathbf{d}') = \sum_{\tau \in \mathcal{T}, \mathbf{v}_\tau = \mathbf{d}' - \mathbf{d}} r_\tau(\mathbf{d}) \text{ whenever } \mathbf{d} \neq \mathbf{d}'$$

and if the summation is empty, then $\mathbf{R}(\mathbf{d}, \mathbf{d}') = 0$.

As the size of the population increases, it has been shown [Kur81] under specific conditions that cover a large range of models that the behaviour of an (appropriately normalised) population CTMC at time t is very close to the solution of a set of ODEs, expressed in the form $\mathbf{X}(t) = (X_1(t), \ldots, X_n(t))$ defining a trajectory over time. The ODEs can be expressed in terms of a single vector ODE as $\dot{\mathbf{X}} = d\mathbf{X}/dt = \mathbf{f}(\mathbf{X})$.

C Appendix: Adding discrete space to a population model

A spatial mode can be constructed by adding discrete space to a non-spatial model. Consider a model consisting of a number of subpopulations, $A, B, C, \ldots$, with a description of transitions between subpopulations where each transition has an associated functional rate defining an exponential distribution. The rate may be dependent on current subpopulation sizes but nothing else, in particular, not time. Additionally, the model has initial quantities for each population. This defines a (non-spatial) population continuous-time Markov chain (CTMC).

This model can be analysed by exact techniques for CTMCs, and its behaviour can be investigated using the Gillespie stochastic simulation algorithm or fluid approximation which provides a system of ordinary differential equations (ODEs) to describe the change in size of each subpopulation over time). Neither the CTMC techniques nor stochastic simulation scale well with increased subpopulation sizes, whereas the fluid approximation complexity is determined by number of subpopulations rather than subpopulation sizes.

To convert this model to a spatial one, n distinct locations can be introduced, and each member of a subpopulation is allocated to a location. We are only interested in which location an individual is in, but nothing further – we do not consider their position within a specific location and this allows the application of techniques that assume well-mixed subpopulations at the location level such as stochastic simulation and fluid approximation.

A location graph is introduced to describe how members of subpopulations can move between locations. Vertices are the locations and edges are the connections between locations. Edges between locations indicate those locations which are physically adjacent or locations for which there is some mechanism for moving between them. Some models are very abstract and assume adjacency of all locations, thus giving a location graph which is complete.

Transition rates can be taken from the original non-spatial model. In some spatial models, the rates are still dependent on the those non-spatial subpopulations sizes. More usually, they are modified to be dependent on the subpopulations in the current location and may involve scaling of the rates because of the introduction of space. An initial quantity is required for each subpopulation in each location, and this can be achieved by dividing the initial quantities of the non-spatial model across the locations.

Movement transitions are also required otherwise the subpopulations in each location will evolve independently of those in other locations. We only consider movement between two locations when there is an edge between the two locations in the location graph. These transitions require rates. We assume for this basic model that these rates are dependent on the subpopulation that is moving, as well as on the size of subpopulations. Again, we do not consider the case that time affects the rate calculation.

The analysis of this basic spatial model can be done in the same way as the original model. However, the introduction of space has increased the number of subpopulations by at most a factor of n where n is the number of locations. The analyses provide information about what happens in each location, but they also allow for calculations of averages and other statistics over all the locations.

To summarise, we have taken a basic non-spatial model, and in the most minimal manner, extended it to a model in discrete space. We next consider adding an attribute to a model, and investigate how adding discrete space differs from adding an attribute.

D Appendix: Adding a discrete attribute to a population model

We can extend a population model in a more general way by considering the addition of an arbitrary attribute. If we start with the same basic (non-spatial) model with a number of subpopulations, and we wish to add a new attribute to all individuals, which will increase the number of subpopulations. The new attribute that can take one of m values. Hence we increase the number of subpopulations by a factor of m. We also assume that we have a graph (called the attribute graph) that describes how an agent can change attribute values (which is equivalent to an agent leaving one subpopulation and joining the subpopulation with a different value for this new attribute). For example, if the attribute represents battery level, it may be possible to transition from *empty* to *charging* but not directly from *empty* to *full*. It is clear that this basic construction is the same as for adding discrete space.

In terms of analysis, as before, variants of Gillespie simulation, fluid approximation and exact techniques for CTMCs can all be applied since dividing the total population up into more subpopulations does not affect the applicability of these techniques, although it could affect how practical their application is.

E Appendix: Markovian agents

Modelling individual agents in continuous space is likely to be very computationally expensive and being able to transform such models to ones that model populations in discrete space is a useful transformation that supports scalability. Appendix F describes a 1-dimensional version of this technique that can be applied to Markovian agents, and the current appendix introduces this formalism.

Formally, a Markovian agent (MA) of class c at position ℓ is defined as a set of matrices and vectors [CGB09]: $\mathrm{MA}^c(\ell) = \{Q^c(\ell), \Lambda^c(\ell), G^c(\ell, m), A^c(\ell, m), V^c, \pi_0^c(\ell)\}$, in which:

- $Q^c(\ell) = [q_{ij}^c(\ell)]$ is a $n_c \times n_c$ matrix, in which each element $q_{ij}^c(\ell)$ represents the rate of the local transition from state i to state j, with $q_{ii}^c(\ell) = -\sum_{j\neq i}^{n_c} q_{ij}^c(\ell)$ where n_c is the number of states of a MA of class c.
- $\Lambda^c(\ell) = [\lambda_i^c(\ell)]$ is a vector, in which each element $\lambda_i^c(\ell)$ denotes the rate of a self-jump transition which reenters the same state i, for a MA of class c.
- $G^c(\ell, m) = [g_{ij}^c(\ell, m)]$ is a $n_c \times n_c$ matrix in which each element $g_{ij}^c(\ell, m)$ describes the probability of $\mathrm{MA}^c(\ell)$ generating a message of type m during a local transition from state i to state j.
- $A^c(\ell, m) = [a_{ij}^c(\ell, m)]$ is a $n_c \times n_c$ matrix, in which each element $a_{ij}^c(\ell, m)$ $(i \neq j)$ describes the acceptance probability of message type m for the $\mathrm{MA}^c(\ell)$, with induced transition from state i to state j whereas $a_{ii}^c(\ell, m)$ denotes the probability of dropping this message, and $a_{ii}^c(\ell, m) = 1 - \sum_{j\neq i} a_{ij}^c(\ell, m)$.
- $V^c = [diag(v_{ii}^c)]$ is a diagonal matrix where each element v_{ii}^c gives the speed of the agent in state i,
- $\pi_0^c(\ell)$ is the initial state probability distribution of an agent of class c at position ℓ.

Figure E shows two MAs in two different positions but which can communicate with each other, where $u_m(\ell, c, i, \ell', c', i')$ is the perception function of message m, whose value represents the probability that an agent of class c, in state i, and at position ℓ perceives a message m sent by an agent of class c' in state i' and at position ℓ'.

F Appendix: Fluid approximation and spatial discretisation in 1-dimensional space

This section considers how an 1-dimensional continuous space model of individuals expressed as Markovian agents (MAs) can be transformed to a 1-dimensional discrete-space population model. MAs are presented in Appendix E. The analysis of interest is the transient evolution of $\rho^c(\ell, t) = [\rho_i^c(\ell, t)]$, the state density distribution of agents of class c at position ℓ and at time t. Let $K^c(\ell, t) = [k_{ij}^c(\ell, t)]$ be the infinitesimal generator matrix describing the rate of stochastic transitions of a MA at position ℓ and at time t which can be easily computed from the definition of MAs. We can compute the density of agents in state i at time $t + \Delta t$ by the following equation:

$$\rho_i^c(\ell, t + \Delta t) = \rho_i^c(\ell - v_{ii}^c \Delta t, t)[1 + k_{ii}^c(\ell, t)\Delta t] + \sum_{j\neq i} \rho_j^c(\ell, t) k_{ji}^c(\ell, t)\Delta t \qquad (1)$$

where the first term in the right hand side of the above equation takes into account the case that the agents do not change their state due to stochastic transitions (assuming Δt is short enough that at most one transition can happen within the period). The second term in the right hand side takes

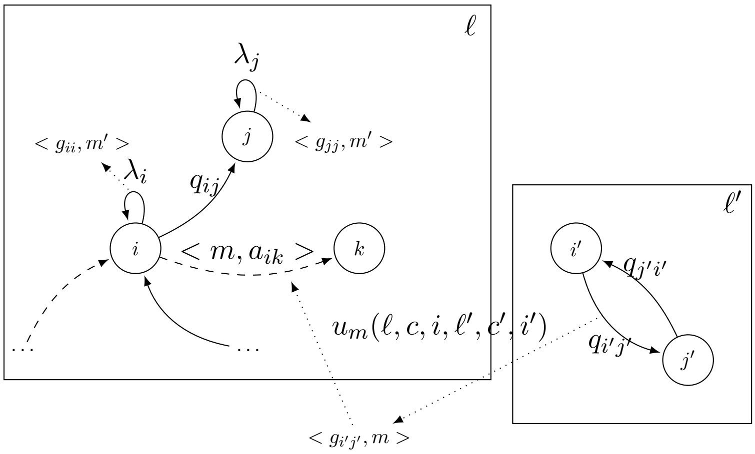

Figure 3: A schematic diagram of two MAs (after [BBC+14, BSB+15])

into account the jump from other states to state i due to a stochastic transition (assuming the jump happens at the end of the Δt slot).

Furthermore, by using Taylor's expansion, the term $\rho_i^c(\ell - v_{ii}^c \Delta t, t)$ can be expanded as:

$$\rho_i^c(\ell - v_{ii}^c \Delta t, t) = \rho_i^c(\ell, t) - \frac{\partial \rho_i^c(\ell, t)}{\partial \ell} v_{ii}^c + O(\Delta t^2) \tag{2}$$

Substituting Equation (2) into Equation (1), then we obtain:

$$\rho_i^c(\ell, t + \Delta t) = \rho_i^c(\ell, t) - \frac{\partial \rho_i^c(\ell, t) v_{ii}^c}{\partial \ell} \Delta t + \sum \rho_j^c(\ell, t) k_{ji}^c(\ell, t) \Delta t + O(\Delta t^2) \tag{3}$$

Rearranging the above equation, letting $\Delta t \to 0$ and ignoring $O(\Delta t^2)$, the transient evolution of $\rho^c(\ell, t)$ can be computed by a set of coupled PDEs using vector notation:

$$\frac{\partial \rho^c(\ell, t)}{\partial t} + \frac{\partial (\rho^c(\ell, t) V^c(\ell))}{\partial \ell} = \rho^c(\ell, t) K^c(\ell, t) \tag{4}$$

where $\frac{\partial (\rho^c(\ell,t) V^c(\ell))}{\partial \ell}$ captures agents' movement on the space, $\rho^c(\ell, t) K^c(\ell, t)$ can be thought as the reaction term due to stochastic transitions.

Then, if we consider the space being bounded between $(\ell_{\min}, \ell_{\max})$. We let the space be discretized into regular intervals of length $\Delta \ell$. By applying upwind semi-discretization technique [HKNT98], for a discretized location ℓ_k, with $\ell_k = k \times \Delta \ell$,

$$\frac{\partial (\rho_i^c(\ell_k, t) v_{ii}^c)}{\partial \ell} \approx \begin{cases} \dfrac{\rho_i^c(\ell_k, t) v_{ii}^c - \rho_i^c(\ell_{k-1}, t) v_{ii}^c}{\Delta \ell} & v_{ii}^c > 0 \\ 0 & v_{ii}^c = 0 \\ \dfrac{\rho_i^c(\ell_k, t) |v_{ii}^c| - \rho_i^c(\ell_{k+1}, t) |v_{ii}^c|}{\Delta \ell} & v_{ii}^c < 0 \end{cases}$$

Combining the above equation with Equation (4) we get an ODE system that can be solved by standard methods like Euler or Runge-Kutta.

$$\frac{\mathrm{d}\rho^c(\hat{\ell}, t)}{\mathrm{d}t} = \rho^c(\hat{\ell}, t) K^c(\hat{\ell}, t) + \rho^c(\hat{\ell}, t) D^c(\hat{\ell}, t) \tag{5}$$

where $\hat{\ell}$ is a discretized location, $D^c(\hat{\ell}, t)$ captures agents' movement between nearby locations.