



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Spatial representations and analysis techniques

Citation for published version:

Galpin, V 2016, Spatial representations and analysis techniques. in *Formal Methods for the Quantitative Evaluation of Collective Adaptive Systems: 16th International School on Formal Methods for the Design of Computer, Communication, and Software Systems, SFM 2016, Bertinoro, Italy, June 20-24, 2016, Advanced Lectures*. Lecture Notes in Computer Science, vol. 9700, Springer International Publishing, QUANTICOL Summer School: 16th International School on Formal Methods for the Design of Computer, Communication and Software Systems: Quantitative Evaluation of Collective Adaptive Systems, Bertinoro, Italy, Italy, 20/06/16. https://doi.org/10.1007/978-3-319-34096-8_5

Digital Object Identifier (DOI):

[10.1007/978-3-319-34096-8_5](https://doi.org/10.1007/978-3-319-34096-8_5)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Formal Methods for the Quantitative Evaluation of Collective Adaptive Systems

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Spatial Representations and Analysis Techniques

Vashti Galpin

Laboratory for Foundations of Computer Science
School of Informatics, University of Edinburgh
Vashti.Galpin@ed.ac.uk

Abstract. Space plays an important role in the dynamics of collective adaptive systems (CAS). There are choices between representations to be made when we model these systems with space included explicitly, rather than being abstracted away. Since CAS often involve a large number of agents or components, we focus on *scalable* modelling and analysis of these models, which may involve approximation techniques. Discrete and continuous space are considered, for both models of individuals and models of populations. The aim of this tutorial is to provide an overview that supports decisions in modelling systems that involve space.

1 Introduction

Collective adaptive systems (CAS) are systems which consist of a number of components which interact (directly or indirectly) to achieve goals, by collaboration, and in some instances, by competition. These components may be static or mobile, as in the case of a robot swarm. Various smart transport systems provide examples of CAS; for example, bike-sharing schemes and ride sharing. Because movement is fundamental in these systems, space and spatial aspects are important characteristics and influence the behaviour that these systems demonstrate. Therefore, we wish to understand the dynamics of these systems and how these may vary with changes in the implementation of the system, and changes in use of the system. In the bike-sharing example, incentives can be offered to users to influence their behaviour in terms of the station a bike is returned to, or alternatively a system may suddenly show very poor performance when the user base grows beyond a certain size. Alterations to timetables of other public transport such as trains, could also impact the effectiveness of a bike-sharing scheme. Furthermore, roadworks or new lane markings can modify the space that the bike users travel through, affecting performance.

We model these systems to understand their behaviour because it is frequently not possible to experiment with the actual systems, either because of the disruption this will cause, or because the systems have not yet been constructed. In this chapter, we focus here on modelling dynamic systems (which we also refer to as time-based) that involve some notion of space. These are systems where the behaviour of the system is observed as time passes¹. When trying to understand the behaviour of a collective

¹ Another approach to space is to consider it topologically, that is to consider the relationships between points in space. This can be applied to both discrete and continuous space. Details can be found elsewhere in this volume [19] in the context of spatial and spatio-temporal logics.

adaptive system by developing a model of the system, it can be moderately straightforward to programmatically construct an agent-based model where the agents move in a representation of real space. But often, for a realistic number of agents, it is not computationally feasible to simulate this model a sufficient number of times to understand its overall behaviour through the use of descriptive statistics. Additionally, an agent-based model is likely to have a very large states space because it considers individuals separately. The computational costs of many other analysis techniques are often dependent on the number of distinct states that the system can take on, and hence cannot be applied to these individual-based models.

Thus, detailed agent-based models may lead to more precision but at the cost of choices for analysis. Typically in modelling, one wishes to retain the details that the model is designed to answer, and to abstract from everything else. Therefore, carefully chosen abstractions are crucial, and this tutorial provides details about a particular type of abstraction and associated approximation of results, that of population-based modelling, rather than solely modelling individuals. These abstractions contribute to *scalable* analysis. By this, we mean that when modelling large systems with many components, our analysis can be computed 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 becomes infeasible. In that case, possible solutions are then to consider whether size can be reduced by working with a more abstract model, or to consider a different approximation technique which is more scalable.

Furthermore, we focus on stochastic models. Stochasticity allows model behaviour to vary, and hence captures the variation we observe in the systems we wish to model. Specifically, we use random length durations drawn from exponential distributions. The exponential distribution is suitable and convenient for modelling because it has a single parameter (which is the inverse of the average duration), it is memoryless (which means that what happens next is only dependent on the current state, as opposed to any previous states, and this negates the need when simulating to keep track of prior states or amount of time elapsed), and other distributions can be approximated by combinations of exponential distributions. In their most basic form, our models are continuous-time Markov chains (CTMCs) and their discrete version, where probabilities are used to determine the next state, discrete-time Markov chains (DTMCs). We also consider extensions and variations of these models, but in general, any stochasticity in our models occurs because of exponentially-distributed durations or probabilistic choices. One extension that we may use in some cases is allowing the exponential rate (and probabilities) to be functional and depend on time or other aspects of the model. This introduces time inhomogeneity into our models, and this is often important to capture variations in behaviour at different times of day, for example. The disadvantage of allowing time inhomogeneity is that it can reduce the number of analysis techniques that are applicable.

This presentation does not consider any languages for specifying models but instead focusses on mathematical representations of systems (which we will refer to as models) to which analysis techniques can be applied. The choice of representation for a model is often influenced by the type of analysis and approximation techniques that are avail-

able, and the aim of this tutorial to support such decisions when modelling space. This chapter starts with a discussion of the type of mathematical representations and analysis techniques that can be used if space is not considered explicitly, and then moves onto consider these with the addition of space. Techniques for discrete space are considered in detail in Section 3, followed by those for continuous space in Section 4. In these two sections, general concepts are introduced for the type of space, followed by a high-level discussion of the basic model and analysis techniques. Details are given of techniques that have relevance to CAS, followed by a brief review of how they have been used in different disciplines. Finally in Section 5, techniques that can be applied to both types of space, or to models containing both types of space are considered.

2 Representations for dynamic modelling

Before considering the role of space, we introduce a number of dimensions that we consider germane to our modelling, so that we can develop a classification of dynamic modelling techniques relating to the modelling context described in the introduction. Even without considering space, there are already a number of choices that lead to different ways in which to model dynamic systems in a quantified manner. We consider the dimensions and the choices on each dimension. For example, the time dimension considers how time is treated in different types of Markov chains. There are other aspects of time such as non-determinism and causality, but these are not a strong focus of our general modelling approach, and so are not included in the classification.

Time: Time is non-negative, strictly increasing and infinite, and can either be a non-negative real or integer. In some models, a finite end-point may be used to delimit the period of interest.

discrete: In the context of this tutorial, discrete time is used in those modelling approaches where choices are probabilistic. At each clock tick (which can be associated with an integer if useful for the specific model), the next state is chosen probabilistically from all possible next states. For example, discrete time Markov chains (DTMCs) use this approach [53, 70].

continuous: In this case, time is represented by the non-negative real numbers. Actions such as changing state have a duration associated with them. In the case of continuous time Markov chains (CTMCs), stochasticity is introduced by having random durations that are drawn from exponential distributions [70].

State: States can be viewed as capturing a quality or attribute of an individual. An individual is assumed to be in a single state at each point in time².

discrete: Usually when the states associated with an individual are discrete, there are a finite number of them. However, in the case of an attribute like *year-of-birth*, there may be a countably infinite number of values.

continuous: A continuous-valued state can be interpreted as measurement of some quantity associated with the individual. An example of this would be *temperature* or *height*.

² An individual could have more than one attribute, and then the individual's state is multidimensional with a value for each attribute. In this case, the individual's state is a tuple of values.

TIME	discrete			
AGGR	none (individuals)		state (populations)	
STATE	discrete	continuous	discrete	continuous
	DTMC [70]	LMP [73]	population DTMC [10]	difference equations, ODEs [10, 67]

TIME	continuous			
AGGR	none (individuals)		state (populations)	
STATE	discrete	continuous	discrete	continuous
	CTMC [70]	CTMP [24]	population CTMC [10, 58]	population ODEs [10, 58]

Fig. 1. Classification of mathematical models in terms of time, aggregation and state (DTMC: discrete time Markov chain, LMP: labelled Markov process, ODE: ordinary differential equation, CTMC: continuous time Markov chain, CTMP: continuous-time Markov process)

Aggregation: Individuals can be considered separately, or the focus can be on the number of individuals in each state. This is more relevant to discrete state approaches than continuous state. In the continuous case, aggregation can be described by a function, or discretisation can be applied to obtain frequency data.

none: Behaviour of each individual is considered separately. This is often referred to as *agent-based* or *individual-based*.

state-based: The behaviour of groups of individuals is considered by counting the number of individuals in each state over time (giving a non-negative integer value), or by having a non-negative real-valued approximation to this number. This approach appears under a number of different names in the literature including *population-based*, *state frequency data*, *numerical vector form*, and *counting abstraction*. The term *occupancy measure* is used when counts are normalised by the population size.

These dimensions can be expressed in a table, which can then be populated with mathematical modelling techniques from the literature. Figure 1 illustrates this and describes the modelling techniques that fit each possible combination of elements for each dimension. There is the possibility of hybrid approaches for the state and aggregation dimensions and we discuss these briefly in Section 5.2.

2.1 Scalable modelling and analysis techniques

As mentioned in the introduction, we focus on Markov chain models. Basic definitions can be found in the appendix. An important aspect of our modelling approach is the application of the mean-field technique where the analysis of a population CTMC or DTMC can be approximated by an analysis using ordinary differential equations

(ODEs) [58, 10]. As the number of states of a Markov chain increases (the “state-space explosion” problem), the analysis of the Markov chain becomes intractable. Modelling a large number of individuals can lead to a very large Markov chain. This can be mitigated by using a population Markov chain where behaviour is considered at a population level rather than at an individual level. The choice of a population Markov chain means we are interested in how many individuals from a population P_A are in each local state A_i , given by N_{A_i} , and the states in the Markov chain have the form $(N_{A_1}, \dots, N_{A_n})$. However, for large systems this may still not be sufficient to obtain reasonable analysis times, and an approximation using ODEs obtained from the population Markov chain can be used. This gives a system of ODEs for the variables $(X_{A_1}, \dots, X_{A_n})$. The population Markov chain considers non-negative integer-valued population counts whereas the ODEs take a fluid approach and population quantities are non-negative real values X_{A_i} . Considering the modelling techniques in Figure 1 for both discrete time and continuous time, the Markov chain obtained by considering many individuals (in the first column) can be transformed into a smaller Markov chain (in the third column) which can then be approximated by ODEs (in the fourth column).

This transformation uses the mean-field approximation technique which comes from physics, where it refers to the approach where movement of an individual particle is considered in the field generated by other particles rather than trying to solve the more complex problem of many particles interacting [68]. In modelling of systems, it has come to mean an approach where it is assumed that when the number of individuals in a stochastic system becomes very large, the population-level behaviour of the system can be expressed as ODEs which provide an “average” behaviour. Results such as those proved by Kurtz [58] demonstrate that under certain conditions, convergence occurs, namely as the number of individuals tends to infinity, the difference between the stochastic trajectories of the subpopulation sizes and the deterministic trajectories of the subpopulation sizes tends to zero. Practically, in many cases, good approximations using the ODE approach over the stochastic approach can be achieved at relatively low numbers of individuals [85] and there are error bounds on the approximations [21]. The mean-field approach is discussed in more detail elsewhere in this volume [9].

Additionally, we will consider moment closure approaches to approximation. For a PCTMC, it is possible to obtain ODEs that describe how the moments (expected values) of variables and products of variables vary over time. Typically, this results in an infinite system of ODEs, because the ODE for each moment is dependent on higher moments. For example, the ODE for $E[X]$ may involve not only $E[X]$ and $E[Y]$ but also $E[X^2]$, $E[Y^2]$ and $E[XY]$. Likewise, the ODE for $E[XY]$ may involve expectations of the product of three variables. Moment closure techniques provide approximations for these higher-order moments through a number of techniques that will be described later in this tutorial, thus providing ODEs that give an approximation for the moments. The mean-field approach described above can be seen as a specific instance of moment closure where second order moments are replaced by the products of expectations ($E[XY]$ is approximated by $E[X]E[Y]$, for example) under certain conditions relating to mass actions and pairwise interactions. This is equivalent to assuming that variances and covariances are zero, and is a reasonable assumption to make if they are likely to be small enough to be safely abstracted from. Typically, in the spatial case, we wish to consider

covariances and other higher moments to ensure that spatial variation is included and not abstracted from.

Returning to Figure 1, Markov processes (in the second column of the figure) do not fit into this work flow (of transforming an individual-based model to a population-based model and then using an ODE approximation) and seem different from the other modelling techniques, as they are characterised by a continuous state space which can also be interpreted as any continuous aspect of a model, including space. We do not consider labelled Markov processes (LMPs) further in this chapter, but we will comment further on continuous-time Markov processes (CTMPs) in Section 4.1.

The research surveyed in this chapter involves transformation and analysis techniques. Transformations of models may be necessary for a different analysis to be applied. The counting abstraction as described above is an aggregation technique, and treating population sizes as being real-valued rather than integral, is fluidisation. Another form of aggregation is when multiple locations are considered as a single location. Finally, discretisation happens when some continuous value is transformed to a discrete value, such as transforming real space to discrete space. Hybridisation which can involve fluidisation to make some parts of a discrete model continuous, or discretisation to make parts of a continuous model discrete, is discussed in Section 5.2.

2.2 Introducing space

In this tutorial, **Space** will be considered in two different ways.

continuous: Here, space is represented by real values in the case of one-dimensional space, pairs of real values in the two-dimensional case and triples of real values in the three-dimensional case. It is always (uncountably) infinite but may be bounded in extent. Continuous space used in this way can be seen as an exact representation of actual physical space.

discrete: Approaches that use discrete space assume a number (usually finite) of distinct locations where connectivity between locations is described by an adjacency relation³. At each location, there may be multiple individuals, although in some cases, such as cellular automata [49], this may be restricted to a single individual. A location may be an abstraction or aggregation of actual space.

The table in Figure 2 shows the mathematical models for the different combinations of time, aggregation, state and space. Here, we have chosen to focus on continuous time models; however there are discrete time models of various approaches, for example, some variants of interacting particle systems (IPSs) use probabilities [29]. We now consider each entry of the table in Figure 2 briefly together with illustrative diagrams.

2.3 Discrete space illustrated

The approaches in the discrete-space category consider space to consist of a (usually) finite number of locations that have connections between them. The most straightfor-

³ For CAS, we are usually interested in the adjacency of different regions of space, and as we will see later, we use graphs to describe this relationship. Another approach is where space has a nested arrangement, as seen in biological modelling. This containment relationship can be represented graphically by trees, but we do not focus on this arrangement of space further.

TIME	continuous			
AGGR	none (individuals)		state (populations)	
STATE	discrete	continuous	discrete	continuous
SPACE				
discrete	CTMC, IPS [29]	TDSHA [12] PDMP [22]	patch population CTMC [17]	patch population ODEs [17]
continuous	molecular dynamics [20] agents	CTMP [24]	spatio-temporal point processes [78]	PDEs [46]

Fig. 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)

ward way is to consider these models as graphs with the locations as nodes and the links as edges. Discrete space is illustrated in Figures 3 and 4, showing the general case of an arbitrary graph, and the case of a more regular graph structure, respectively. Regular space models are those that have a regular pattern of locations [28, 29]. For example, the locations could be laid out in the rectangular grid, or a hexagonal tiling. The locations that represent space can be situated at the nodes of the regular graphs or in the spaces (faces) created by the regular graph as shown in Figure 4. Regular space will be more formally defined in Section 3.

In the diagrams, we assume individuals are from two populations. The first, P_A consists of red and white tokens, and has states A_1 and A_2 . The second, P_B consists of blue and white tokens with states B_1 , B_2 and B_3 . The current state of an individual is indicated on the top of the token. The four diagrams in each figure represent four single points in time and do not show change over time⁴.

Figures 3(1) and 4(1) show discrete-space models of individuals with discrete states, hence there is no aggregation into populations. Some models only allow one individual in each location, such as interacting particle systems (IPs) [29] and cellular automata (CA) [49], but others may allow multiple individuals. In the case of single individuals at a node, this can be indicated by a flat token as illustrated in Figure 6.

Models of discrete space without aggregation and with continuous state are shown in Figures 3(2) and 4(2). The continuous state is indicated by a solid token where the height indicates the value of a single continuous state. This is an inherently continuous value rather than the notion of population size approximation by continuous values described earlier in this section, and could be a measurement such as strength of radio

⁴ For two-dimensional and three-dimensional space, the best visualisation method for change over time is video. For one-dimensional space, a graph with two axes can be used.




Fig. 3. Discrete space: (1) no aggregation, discrete state; (2) no aggregation, continuous state; (3) aggregation of state, possible aggregation of space, discrete state (4) aggregation of state, possible aggregation of space, continuous state

signal or length of battery life. Different colours have been used in the diagram to make it clear that the values are continuous but not a population approximation. In Figures 3(2) and 4(2), there is an assumption of at most one individual per node and face, respectively, and two values associated with that individual.

Next we consider discrete-state aggregation in the context of discrete space, as illustrated in Figures 3(3) and 4(3) by the fact that individual tokens are grouped into stacks at nodes in the network, and it is the size of the stack that is relevant rather than the location of each individual. Finally, in the case of continuous state aggregation in discrete space, each region or point is associated with approximations to the discrete population shown in Figures 3(3) and 4(3). These are illustrated in Figures 3(4) and 4(4). At each node, for each state in each population, there is a real number that approximates the number of individuals in that state. This is illustrated by a token with a real-valued height for each state in each population. Note that in Figure 3(4), the lowest node has a non-zero value for blue tokens in state B_3 although there were none in the CTMC model in Figure 3(3), illustrating that approximation can occur.

2.4 Continuous space illustrated

We first consider continuous space with no aggregation and discrete state. This covers approaches where each individual's location and state are modelled separately from those of other individuals. An example of this type of model is where the movement and interaction of each molecule is modelled individually, as in molecular dynamics [7]. Agent-based models take a similar approach. Figure 5(1) illustrates this. The continuous space is indicated by a bounded area and each individual is shown at its own location. These models are typically computationally expensive to simulate.

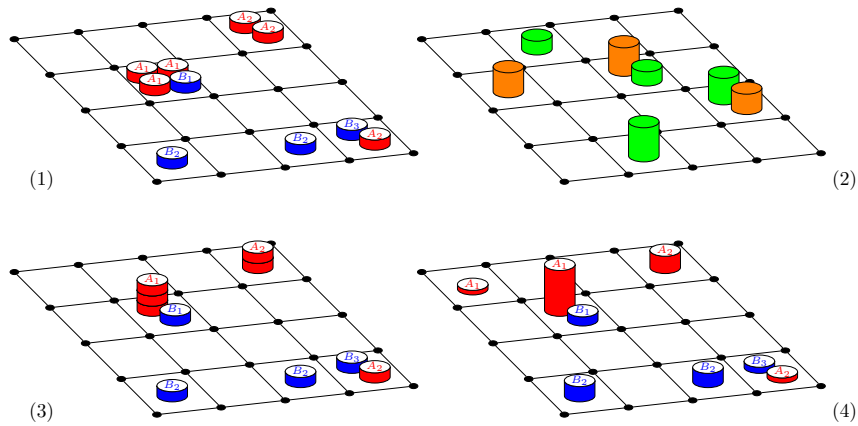


Fig. 4. Regular discrete space: (1) no aggregation, discrete state; (2) no aggregation, continuous state; (3) aggregation of state, possible aggregation of space, discrete state (4) aggregation of state, possible aggregation of space, continuous state

Moving on to state that is continuous rather than discrete, leads to continuous-time Markov processes (CTMPs) [24], if we assume some of the continuous dimensions relate to space and the others to state. Since there is no state-based aggregation, this approach models individuals rather than populations. The continuous space is indicated by a bounded area and each individual is shown at its own location. The continuous state is indicated by the varying heights of the tokens, and in Figure 5(2), it is assumed that there is only one (non-spatial) measurement per individual, although multiple different measurements are possible.

For the case of aggregation with discrete state, each point in space can be filled by zero, one or more individuals [78]. Hence for each point in space, it is possible to aggregate the number of individuals in each state. Figure 5(3) shows a fairly sparse number of individuals but much denser arrangements are also possible. Finally, when aggregation is continuous in nature, then at each point in space, there is a real value describing an approximation to the number of individuals at that point [71, 20]. In the case of two-dimensional space, the population of each state can be represented in three-dimensions by surfaces as defined by partial differential equations (PDEs). Figure 5(4) illustrates a surface describing the number of individuals at each point for state A_1 . In contrast to Figure 5(3), this figure illustrates a very dense situation.

2.5 Summary

As is the case with techniques that do not include space, presented in Figure 1, the techniques using continuous state without aggregation (the second column of models in Figure 2) seem distinctly different to the other approaches. The techniques that can be applied to models without space described earlier in this chapter (approximation by ODEs of a population DTMC or CTMC) can be applied to discrete space since the




Fig. 5. Continuous space: (1) no aggregation, discrete state; (2) no aggregation, continuous state; (3) aggregation, discrete state; (4) aggregation, continuous state.

Markov chain involved is a population Markov chain that takes location into account. Furthermore, taking the hydrodynamic limit of IPS (which are discrete space models without aggregation) models provides PDEs [23].

In all of the models described in the previous section, there may be interaction between individuals (even if this interaction is expressed at the population level). Opportunity for interaction is often related to colocation or proximity (which requires some notion of neighbourhood or distance). Many models capture movement of individuals explicitly and then use colocation or proximity to determine the possibility of interaction, although there are some models that only use proximity without movement such as IPSs and CA. We discuss movement in more detail when we consider the analysis techniques for the two different kinds of space.

3 Discrete-space modelling techniques

We now consider discrete space in more detail and formality, so we introduce both notation and concepts relevant to discrete space. We will focus here on the continuous-time models, with pointers to the discrete-time models where appropriate.

In the most general case, we assume a finite (or at most countably infinite) set of points or locations L with some naming convention [41]. Most generally, the set of locations L can be taken as the vertices of an undirected graph, and the connections between locations (the adjacency relation) can be defined as edges in that graph. The

edges of the graph E_L are drawn from the subsets of size two of the location set $\mathcal{P}_2(L)$, so $E_L \subseteq \mathcal{P}_2(L)$. Each edge has the form $\{l_1, l_2\}$, and edges of the form $\{l, l\}$ are permitted. We have chosen to use an undirected graph which is to be understood as allowing movement or interaction in at least one direction between the two locations. The absence of an edge means that movement and interaction can never take place, in either direction. Parameters associated with an edge express (possibly in a time-varying manner) the propensity for movement or interaction in either direction. If a parameter is zero at a particular time for a particular direction, it means that no active interaction or movement can take place at that time point. Hence, the graph of locations provides a skeleton for describing what movement or interaction is possible.

Locations in discrete space models can have two main sources, either they are essentially locations on a map, such as bike-stations or bus stops, or alternatively each location represents a region on a two-dimensional map, and space is aggregated. These are called patch-based models. The edges of the graph can be determined by various factors. Adjacency of regions is an obvious choice, but there may be other context-specific elements, for example, presence of connections between regions such as railway lines or similar. A topic whose exploration is beyond the scope of this chapter is that of how to divide a map in regions. A simple approach is to base it on a tiling of the plane using triangles, quadrilaterals or hexagons. More complex approaches involve taking local information into account and creating irregular patches. Computer networks can be seen as being located in discrete space, either physically or logically.

An issue for discrete space (and continuous space) is determining what happens at the boundaries of the space. One approach is to ensure there are none by working with infinite structures such as infinite graphs, or alternatively boundaryless structures such as tori. A rectangular region can be transformed into a torus by joining the top and bottom edges (to form a cylinder) and then joining the left and right ends (by curving the tube). Other approaches work with boundaries and either choose to keep individuals inside the region (by reflection or other techniques) or to treat boundary locations as sources and/or sinks.

The discrete space approach as described above is very general as it allows arbitrary graphs over locations, as well as heterogeneity for parameters. In the literature there are modelling techniques that are defined for specific graph subclasses and we will discuss some of these below.

3.1 Spatial parameters and regularity

A modelling technique with discrete space will have parameters that depend on locations, or links between locations. 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 define two functions to describe these parameter sets as follows

- $\lambda(l)$ for $l \in L$, and
- $\eta(l_1, l_2)$ and $\eta(l_2, l_1)$ for $\{l_1, l_2\} \in E_L$.

The range of these functions will remain abstract for the purposes of this discussion. Note that although the edges of the graph are not directed, the function η is sensitive

to direction. Movement is obviously directional but interaction can be undirected when considering an abstract view of effect 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 in parameters. In the rest of this chapter, the term *transfer* will be used to refer to both movement and interaction.

We present the following definitions, leading to a definition of spatial homogeneity (a term which is used in the literature but not formally defined), by considering the location-related parameters. A spatial model is

- *location homogeneous* if $\lambda(l_i) = \lambda(l_j)$ for all locations $l_i, l_j \in L$.
- *transfer homogeneous* if $\eta(l_i, l_j) = \eta(l_j, l_i) = \eta(l_{i'}, l_{j'}) = \eta(l_{j'}, l_{i'})$ for all edges $\{l_i, l_j\}, \{l_{i'}, l_{j'}\} \in E_L$.
- *(spatially) parameter homogeneous* if it is both location and transfer homogeneous.
- *spatially homogeneous* if it is parameter homogeneous, and its location graph is complete⁵. Regular connections between locations which do not give total connectivity are discussed below.

Models with spatial homogeneity have a symmetry that can allow for analyses that are not possible for more complex models. Examples are the bike-sharing system considered in [39] where the metrics of interest are the number of empty and full bike stations.

Spatial inhomogeneity/heterogeneity 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. Note that if a parameter $\rho_{i,j} \in \eta(l_i, l_j)$ is constant for all i and j but other parameters vary by locations, then the model is spatially inhomogeneous.

Regular discrete space covers those discrete space models where the organisation of space is regular (rather than an arbitrary graph where each vertex may have an arbitrary number of edges). By contrast to spatial homogeneity, regularity of space is more difficult to define formally when starting from a graph (and we do not give details here), although it is very straightforward to identify visually [72]. Terms such as lattice, grid or mesh are frequently used to describe a graph based on a square or rectangular tiling of the plane. The other two regular tiling possibilities are equilateral triangles and regular hexagons. Alternatively, a graph with regular structure can be constructed by identifying points in $\mathbb{Z} \times \mathbb{Z}$ or $\mathbb{R} \times \mathbb{R}$, and adding links. We will not attempt that level of generality for discrete space beyond saying that regular space should have the property that at each location (except possibly at boundary locations) there is a uniform way to determine the immediate neighbours⁶. One-dimensional regular space can be represented simply as an undirected path. We do not tackle a formal definition of three-dimensional regular space.

⁵ A complete undirected graph has an edge $\{l, l'\}$ between each pair of vertices l and l' .

⁶ We exclude from this definition n -hop neighbours in an arbitrary graph (see definition of n -hop in the next subsection).

3.2 Neighbours and neighbourhoods

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 usually 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, with obvious adaptations.

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 capabilities. Other approaches use a (perception) function that determines the *de facto* neighbours of an individual by specifying the other individuals with which it can interact.

3.3 Techniques for individual discrete-space models

We now consider the different modelling techniques that have been applied to discrete space starting with those that do not involve aggregation of state. When there is no aggregation and state is discrete, the focus is on individuals and an example is an agent-based system over discrete space. Each individual has some state and is located at exactly one location. There may be a restriction to one individual per location. To describe these models in their most general form, we assume that each individual I (where I is a unique name for the individual) has associated time-based information:

- $\text{loc}(I, t) \in L$ which is its location at time t
- $\text{state}(I, t) \in \{A_1, \dots, A_n\}$ which is its state⁷ at time t

Additionally there are rules that describe how an individual can change location or change state. Since this is a continuous time model, these rules specify rates to describe how long it takes on average for the changes to occur. Each rate defines an exponential distribution, and may be constant or the rates may be functions that depend on the presence of others at that location, the characteristics of the location or the current time (thus introducing time-inhomogeneity). The behaviour of the agents in this modelling technique is thus described as they individually change state and/or location. Assuming a fixed population size, we can model this system as a CTMC, where each state in the CTMC is a tuple consisting of information about each individual in the system. If we assume N individuals then a state has the following form

$$((\text{loc}(I_1, t), \text{state}(I_1, t)), \dots, (\text{loc}(I_N, t), \text{state}(I_N, t))))$$

⁷ If the population P_A has multiple attributes $A^{[1]}, \dots, A^{[p]}$, then $\text{state}(I, t) = (A_i^{[1]}, \dots, A_i^{[p]})$ represents a tuple of attributes.




Fig. 6. A graph-transformation rule applied to an individual discrete-space model with discrete state

There are $(L \times n)^N$ states in this Markov chain if there are L locations, n states and every combination of state and location is possible for all individuals.

Simulation suits this type of model, and techniques for simulating systems where behaviour is based on functional exponential rates are well understood [43]. They can also be analysed using standard numerical CTMC techniques for steady state and transient behaviour. However, a large number of individuals can make this computationally infeasible.

Next, we consider discrete space modelling techniques without aggregation but where the state is continuous. Therefore, instead of having a rule describing how (discrete) state change can happen, there needs to be a rule describing how the continuous state changes over time. A good candidate for this type of rule is an ODE. These techniques are hybrid in that they exhibit both continuous behaviour with respect to state and stochastic behaviour with respect to space. Transition-driven stochastic hybrid automata (TDSHAs) [12] and piecewise deterministic Markov chains (PDMPs) [22] are suitable modelling techniques. Both of these also introduce the possibility of instantaneous behaviour.

3.4 Pair approximation: spatial moment closure based on structure

The technique called pair approximation, which we will refer to as structure-based moment closure, provides ODEs which describe the changes over time in the probabilities of certain pairs (adjacent locations) in the model [91, 66]. From these ODEs, the proportion of locations in a particular state can be determined.

It is applied to a specific class of discrete-space models of individuals with discrete state, namely graph-transformation models. In these models, each node either represents a single individual or a single position in space or location which can take on exactly one of a small number of states. 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 using continuous time). A graph transformation rule describes how a small subgraph or pattern can be transformed in another pattern. There

are two possible 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. Here, we investigate a static model of space and so we only consider the first type of transformation in this chapter. An example of such a rule is given in Figure 6. The lack of distinction between location and agent is indicated by the fact that the disks are flat rather than raised tokens, as mentioned earlier.

As an example, consider a graph-based SIR model⁸ where each node is an individual who can be in one of a number of states (susceptible, infected, recovered, hence the abbreviation SIR) and the edges of the graph link individuals that can affect each other. The graph-transformation rules include a linked pair consisting of one susceptible and one infected being modified to a linked pair consisting of two infected nodes (as illustrated in Figure 6), and a infected node being modified to a recovered node. In ecological modelling, 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 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 [91, 66], but one can also 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 generation of ODEs is similar to that of the moment ODEs described in Section 2.1 and leads to an infinite system of ODEs. This system of ODEs can be closed using certain closure techniques (which will be discussed in more detail in Section 3.6) and thereby give an approximation to the true value. Structure-based moment closure has also been considered as a multi-scale technique [31]. In this case, different sizes of neighbourhood are used for different types of interaction.

3.5 Techniques for population discrete-space models

We now move on to consider discrete space when aggregation of state occurs resulting in populations, whose sizes are either integral or real-valued. It is assumed that we have many individuals to whom the same set of rules apply with the same parameters, and we choose to view them as a population and to reason about them as a population.

⁸ This is different to the population SIR model that appears in another chapter in this volume [9] because there is at most one individual at each node in the graph, and that individual has an associated state, rather than subpopulations in each state.




Fig. 7. Behaviour in a population discrete-space model with discrete state

These models are population CTMCs where subpopulations in different locations are viewed as separate subpopulations. These are also called patch-based models and there are various examples in the literature [17, 93].

We consider a population P_A . At each point in time, each individual in P_A is in exactly one of its local states A_1, \dots, A_n . Let $N_{A_i}(t)$ refer to the number of individuals in population P_A that are in state A_i at time t . These are called *subpopulations*. The total number of individuals in the population at time t can be expressed as $N_A(t) = \sum_{i=1}^n N_{A_i}(t)$. Furthermore, if no births or deaths are assumed, and an individual must be in one of the available states⁹, then $N_A(t_1) = N_A(t_2)$ for all times t_1 and t_2 and the size of P_A is a constant N_A . We use $X_{A_i}(t) \in \mathbb{R}_{\geq 0}$ to represent a non-negative real-valued description of the population P_A which is an approximation to $N_{A_i}(t)$.

If we assume that we have a fixed number of locations, l_1, \dots, l_L , we can now obtain the counts of subpopulations at each location. So for P_A , we have a value $N_{A_i}^{(k)}$ which is the number of individuals at location k in state i . Additionally

$$N_{A_i} = \sum_{k=1}^L N_{A_i}^{(k)} \quad \text{and} \quad N_A^{(k)} = \sum_{i=1}^n N_{A_i}^{(k)} \quad \text{and} \quad N_A = \sum_{i=1}^n N_{A_i} = \sum_{k=1}^L N_A^{(k)}$$

We can create a continuous time Markov chain smaller than that of the previous section consisting of at most $(N_A + 1)^{L \times n}$ states where each state has the form

$$(N_{A_1}^{(1)}, \dots, N_{A_n}^{(1)}, \dots, N_{A_1}^{(k)}, \dots, N_{A_n}^{(k)}, \dots, N_{A_1}^{(L)}, \dots, N_{A_n}^{(L)})$$

This provides a discrete aggregated representation of individuals in space where for each location, we know how many individuals are in each state without knowing exactly which individual at that location is in which state. An example of behaviour in such a model is illustrated in Figure 7 where an individual in state B_3 moves from one location to another and the population sizes at those locations change as a result of this movement.

In the case of continuous state aggregation, the notation $X_{A_i}^{(k)}$ is used for the real value that describes the quantity of individuals in state i at location k . Since this can

⁹ In some models, births and deaths can be included for a fixed size population by introducing a “dead” state. However, this requires that there is a finite maximum population size.

be a non-integer value, it is an approximation to the actual count $N_{A_i}^{(k)}$. Since the subpopulation sizes are treated as continuous values, a standard modelling technique is to express the change in this quantity in terms of an ODE.

$$\frac{dX_{A_i}^{(k)}}{dt} = F_{i,k}((X_{A_1}^{(1)}, \dots, X_{A_n}^{(1)}, \dots, X_{A_1}^{(k)}, \dots, X_{A_n}^{(k)}, \dots, X_{A_1}^{(L)}, \dots, X_{A_n}^{(L)}), t)$$

This is a population ODE because it tracks the changes in subpopulation sizes over time. There are $L \times n$ variables in total; one for each combination of state and location. The inclusion of t as an argument to $F_{i,j}$ indicates that it can be a time-inhomogeneous ODE. This ODE often has the following form

$$\begin{aligned} \frac{dX_{A_i}^{(k)}}{dt} = & f_{i,k}(X_{A_1}^{(k)}, \dots, X_{A_n}^{(k)}) + \\ & \sum_{j=1, j \neq k}^L (g_{i,k,j}(X_{A_1}^{(k)}, \dots, X_{A_n}^{(k)}, X_{A_1}^{(j)}, \dots, X_{A_n}^{(j)}) - h_{i,k,j}(X_{A_1}^{(k)}, \dots, X_{A_n}^{(k)}, X_{A_1}^{(j)}, \dots, X_{A_n}^{(j)})) \end{aligned}$$

where $f_{i,k}$ captures the local behaviour which only depends on the subpopulation sizes locally, $g_{i,k,j}$ describes the inflow of population from location j to location k , $h_{i,k,j}$ describes the outflow of population from location k to location j , and these flows depend only on the subpopulation sizes in location k and location j . This is a time-homogeneous ODE since change over time is only dependent on subpopulation sizes (that are dependent on time) rather than on time directly. For both the general and regular space cases and assuming only movement/interaction between 1-hop neighbours, then a term $X_{A_i}^{(j)}$ should only appear in the right hand side of the ODE if $\{l_k, l_j\}$ is an edge in the location graph.

In both models, discrete population and continuous population, 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.

Since PCTMCs with locations are PCTMCs then the usual 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. The computational feasibility is limited by the size of the state space.

Simulation is also applicable to simulate individual trajectories of behaviour using an algorithm such as that proposed by Gillespie [44]. 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. In the case of PCTMCs with locations, the assumption of well-mixedness must be made for each location.

Finally, the techniques based on Kurtz's result [58] that express the average behaviour of a PCTMC as ODEs also apply to the fluidisation of a PCTMC with locations. The assumption of well-mixedness also applies, as with Gillespie simulation. Although the ODEs provide an approximation to the true values, this is achieved much faster as it is easier to calculate the trajectory of a set of coupled ODEs than it is to do multiple simulations for statistical analysis. Techniques such as exact fluid lumpability and related

approximation techniques [87, 88] identify when it is possible to apply an aggregation when dealing with ODEs and these techniques are discussed further elsewhere in this volume [90].

We can also consider homogeneity of parameters. In the case of spatial homogeneity, the fact that parameters are identical may make the model amenable to an analytic approach, rather than requiring simulation [39]. However, variations in parameters and rates do not affect the speed of 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).

Another issue to consider that relates to spatial heterogeneity is that of dynamic space where nodes can leave and join a network and links can be added or removed. Although we do consider that time-homogeneity may be a feature of our PCTMCs and associated ODEs because rates are dependent on time, we do not consider dynamic location graphs here, because of the complexity introduced by this additional change in behaviour over time.

3.6 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 will use the term *aggregate moment closure* for the techniques that are applicable to population discrete-space models because it is more descriptive. Aggregate moment closure requires fluidisation of the population model, derivation of moment ODEs, and application of an approximation technique to close the moment ODEs.

In this approach, moment ODEs (see the appendix for a definition) 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 basic approach is to obtain an ODE for each subpopulation for the ensemble¹⁰ 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 variables (a higher order moment). The ODE for this can then be derived and this again is likely to contain even higher order moments. In most cases, the system of ODEs 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. Earlier it was mentioned that the mean-field approximation (in the sense of Kurtz) is given by the first moment ODEs with approximations for variances and covariances based on an assumption that these were zero or negligible (see also [9] in this volume). Because the covariance captures spatial variation, we must have ODEs for at least second moments but third and higher moments can be approximated. There are four ways to approach this approximation.

¹⁰ The mean (at time t) over all stochastic realisations (at time t).

- 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 [65].
- 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 mentioned above, hence this technique can only be applied to third and higher order moments [61]. The modified mean-field approach from ecology takes a similar approach by approximating higher moments with powers of first order moments [74].
- 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 [61, 62].
- Apply a Taylor expansion of moments, as used in scale transition theory [18] which formalises how local dynamics relate to global dynamics, particularly in the case of nonlinearity.

Most applications of this technique assume a complete graph, or alternatively when neighbourhood is used in an incomplete graph, approximate the results with those obtained from a complete graph [62].

Another approach to moment closure is language-based where information from the model specification language is used to determine which moments are likely to be negligible [36]. A neighbourhood relation is derived from the (language-based) model to determine when it is appropriate to approximate the expectation of a product with the product of expectations. This relation could also use spatial information to determine approximation.

3.7 Multi-scale techniques based on differences in rates

As mentioned previously, 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 [45, 79]. This can be done both within a stochastic approach and a deterministic approach using ODEs. Another technique is timescale decomposition applied to CTMCs 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 [80].

In ecological modelling, spatial aggregation methods consider the combination of different time scales that are location-based [1]. 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 [37, 92] can be applied to obtain a reduced system. This technique can perform much better than the spatial moment technique when there is substantial demographic variation across patches but it does require differences in rates.

In other models outside of ecology, particularly those involving computer systems, it is likely to be the case 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 and then techniques based on QSSA is more appropriate.

3.8 Applications of discrete space models

In this section, some applications of the discrete space models that have been presented are now discussed briefly. For a detailed survey of the applications of discrete space models, the reader is referred to [41].

Ecology: Space plays a crucial role in many ecological models and ecologists are interested in global qualities of the whole space such as whether species persist or can co-exist, as well as dynamic patterns such as stationarity, oscillatory behaviour, chaos or multistability [68]. Berc [5] provides a classification of spatial models where he considers the time, space and population as different dimensions. Reaction-dispersal networks (also called metapopulation models) are continuous-time, discrete-space, continuous-population models that describe change over time by a system of ODEs over species in locations. They are the same as ODE patch models in our terminology. Coupled-map lattices are a discrete-time model defined systems of difference equations [51] and regular discrete space, and allow continuous population sizes. Morozov and Poggiale [68] highlight that the term "mean-field" can be used in the ecology literature to both describe the non-spatial Kurtz-based approximation technique as well spatial approaches.

Biology: Bittig and Uhrmacher [7] identify five distinct methods for spatial modelling in cell biology that offer different granularities in their approximation of physical reality. Two of these are continuous space approaches and are discussed in Section 4.5. The discrete-space models are those that use compartments as a nested arrangement of space, discrete-space lattice approaches with a single molecule at each face of the lattice, and discrete-space lattice approaches where multiple molecules are permitted at each face. For an overview of techniques to model diffusion, both stochastically and continuously, see [33]. Patch models are also used to model biochemical reaction systems [61]. Pattern formation is also important in biology and Turing's paper gave an initial insight into this process [89]. Pattern formation is considered in [19] in this volume.

Epidemiology: Riley [75] identifies four distinct approaches to disease spread modelling that considers different levels of interaction: patch-based, distance, multi-group and individual. Patch-based or *metapopulation* models¹¹ are used extensively in modelling of epidemics [27]. These models often focus on the calculation of the basic reproduction number, R_0 , which determines whether a disease will die out or spread to the whole population. Individual discrete space models have also been used for disease modelling [59], as illustrated in Section 3.4.

Networking: Computer networks, in particular ad hoc networks and mobile networks, often require spatial modelling for evaluation. For example, computer and mobile phone virus spread modelling involves spatial aspects and much of this research draws on epidemiological approaches [48, 55]. Routing protocols may have spatial aspects that can be discrete or continuous [95]. Patch models have been used to model information transmission between mobile nodes [17, 94, 35].

Forest fires: Propagation of forest fires is investigated using Multi-class Multi-type Markovian Agent Model (M^2MAM) [16]. The approach models individual agents in discrete space and from this, a patch ODE model is derived. Forest fires have also been modelled using stochastic cellular automata in a climate model [60].

Robotics: A robotics case study consists of a swarm of robots that have to collectively identify a shortest path [63]. The division of a path into separate sections which are considered as discrete locations provides a way to approximate the traversal time by real robots and the convergence on the shortest path.

Emergency egress: The modelling of evacuation from a multi-story building [64] involves a multi-story building with building elements such as rooms, corridors and stairwells, doors and exits. To model the movement of people and the time to evacuate the building, a discrete-space model using patches was developed.

Crowd behaviour: Spontaneous drinking parties are a common phenomenon in cities in the south of Spain [76]. A model shows that the introduction of small variations that break symmetry, both in space and in the degree of connectivity between locations and in the behaviour of the individuals can lead to new behaviour [11]. This example is considered elsewhere in this volume [90].

Bike sharing: Bike sharing systems have been modelled with homogeneous discrete space using a population CTMC approach with an associated mean-field model [39]. When space is not homogeneous, a clustering approach has been used to group similar locations together [40]. This example is also considered in this volume [90].

A number of the above examples are CAS. Other CAS examples where discrete-space techniques are applicable include smart transport and smart grids. The next section considers modelling with continuous space.

4 Continuous-space modelling techniques

Continuous space is more straightforward to define than discrete space. In this section, we will focus on two-dimensional space; however, both one- and three-dimensional

¹¹ The basic epidemiological SIR model is called the compartment model [13] and this consists of a single population with no spatial aspects. It should not be confused with the compartment models in biology which are patch-based models.

space may be useful in various contexts. Continuous space can either be the Euclidean plane extending infinitely in all directions, $\mathbb{R} \times \mathbb{R}$, or it can be a bounded connected (contiguous) subset of this plane. Points in the plane can be referred to by their coordinates $(x, y) \in \mathbb{R} \times \mathbb{R}$. As with discrete space, we can consider two cases, depending on whether we focus on individuals or populations.

This section starts with considering individual-based continuous-space models. Next, population continuous-space models are presented, followed by two techniques that are relevant for population discrete-space modelling, but involve continuous-space models or techniques as well. The section ends with examples of the application of continuous-space techniques in various disciplines.

4.1 Techniques for individual continuous-space models

In these models, we consider identifiable individuals. There are many different models of the movement of individuals through two-dimensional space, such as models of animal movement and models for ad hoc and opportunistic networks [14]. These are often stochastic and capture the probability of movement in a particular direction at a certain speed. Additionally, it may be necessary to determine what happens at the boundary of the space. Often, it is assumed that the space is the surface of a torus and hence has no boundaries – this is more common than assuming the surface of a sphere, as it is hard to map subsets of $\mathbb{R} \times \mathbb{R}$ to the sphere. There are also models to describe the movement of a related group of individuals through the space [14]. Connectivity models on the other hand, describe interaction (for example, contact duration and time between contacts) rather than location [52] so they are implicit movement models. Interaction can be interpreted as dynamic graphs with the individuals as the nodes.

Next, we consider the form that these models can take. If I is an individual, then it has associated information, similar to the discrete state case.

- $\text{loc}(I, t) \in \mathbb{R} \times \mathbb{R}$ which is its location at time t , and
- $\text{state}(I, t) \in \{A_1, \dots, A_n\}$ which is its state¹² at time t .

There are rules which describe how the individual changes state and these may take into account the individual's current location, and rules that describe an individual's movement through space which may take into account the individual's state. As with discrete space, the rates for state change are exponential and can be functional. Unlike with discrete space, it is not useful to construct a Markov chain whose states are obtained from the locations and states of each individual. Discrete event simulation can be used to explore the behaviour of these systems [38].

In the case that the state is continuous, then

- $\text{state}(I, t) \in \mathbb{R}^n$ for $n \geq 1$, which is continuous and represents its state at time t .

As with the discrete space case, some way is required that describes the change of state over time, and an ODE can be used for this. Some models require both discrete and continuous non-aggregated states and this requires a hybrid solution. Agent-based models

¹² As with discrete space, if the population P_A has multiple attributes $A^{[1]}, \dots, A^{[p]}$, then $\text{state}(I, t) = (A_i^{[1]}, \dots, A_i^{[p]})$ representing a tuple of attributes.

in continuous space are examples of an individual continuous-space model where individuals can take on discrete states or continuous values.

A different approach to modelling continuous state with continuous time is that of continuous time Markov processes (CTMP) [24]. A CTMP is a tuple (S, Σ, R, L) where (S, Σ) forms a specific type of topological manifold and $R : S \times \Sigma \rightarrow \mathbb{R}_{\geq 0}$ is a rate function which is measurable in its first coordinate and a measure on its second coordinate. L is a state labelling function. Applying this in the context of space, the manifold is $(\mathbb{R} \times \mathbb{R}, \Sigma)$ where Σ consists of the open sets of $\mathbb{R} \times \mathbb{R}$, hence defining a σ -algebra. A notion of path through this space can be defined describing the behaviour of an individual. Furthermore, if there are additional continuous quantities associated with the individual then additional dimensions of \mathbb{R} can be used.

4.2 Techniques for population continuous-space models

When individuals are aggregated into populations, there is no need to keep track of them individually and *densities* become more important. In spatio-temporal point processes¹³, each point in space (x, y) has an associated integral count for a state in a population at a specific point in time t . We can denote this as $N_{A_i}((x, y), t)$ and its behaviour is described by a function $\lambda((x, y), t)$. In general, λ can depend on all preceding events, but in the case of a Poisson process, it only depends on (x, y) and t [78]. If λ is a constant, then there is no spatial heterogeneity. If the equation defining λ includes comparison with other points, then either clustering or inhibitory behaviour can be defined. If time and space are independent then λ can be defined by $\lambda((x, y), t) = \lambda_1(x, y)\lambda_2(t)$. The form of λ may also describe a reduction in the population at a specific point (x, y) and dispersal of that population to other points, thus capturing movement.

For continuous aggregation of populations, we now consider the classical model of movement in continuous space, that of partial differential equations. For populations described by $X_{A_i}((x, y), t)$, the general form is

$$F_i(x, y, t, X_{A_1}, \dots, X_{A_n}, \frac{\partial X_{A_i}}{\partial x}, \frac{\partial X_{A_i}}{\partial y}, \frac{\partial X_{A_i}}{\partial t}, \frac{\partial^2 X_{A_i}}{\partial x^2}, \frac{\partial^2 X_{A_i}}{\partial xy}, \frac{\partial^2 X_{A_i}}{\partial y^2}) = 0$$

if we assume that we are interested in second order partial derivatives over space only for the population $X_{A_i}((x, y), t)$. Note that writing the PDE in this form simply allows it to be described as a function over all the derivatives of interest rather than as a single partial derivative being equal to a function of other derivatives. When interactions between populations are to be modelled, diffusion-reaction PDEs are used since they can express movement as diffusion and interaction as reactions [20, 89]. The diffusion terms can also capture drift which accounts for obstacles or external stimuli such as wind, the likelihood of continuing in the same direction, the effect of the density of other individuals, and the impact of environmental characteristics. The reaction term describes interactions between individuals. Examples are given in the following sections. There are various techniques for solving PDEs which we will not consider here, many of which involve discretising the plane into a mesh [81].

¹³ In contrast to spatio-temporal point processes, spatial point processes describe distributions in space, and do not include a notion of change over time [3] and hence are not relevant in this context.

We now consider two approaches to modelling discrete space where continuous space plays an important role, in the sense that transformation from one type of space to another is involved.

4.3 PDE-based analysis of discrete-space models

Tschaikowski and Tribastone [88] have considered an approach which involves taking a discrete space model with random walks to continuous space through spatial 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, \dots, A_L while being at any point in \mathcal{R} , meaning that the CTMC state

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

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 a hostile environment, while the latter account for closed environments. It can be 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 $N \rightarrow \infty$. 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 [42] which solves the ODE system of size $\mathcal{O}(L \cdot K^2)$ and that can be also interpreted as a finite difference scheme [84] of a PDE system of size L . By combining this with the former result, one then proves that the solutions of the ODE system of size $\mathcal{O}(L \cdot K^2)$ converge, as $K \rightarrow \infty$, 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 [86, 88], thus showing that a characterization of mobile systems in terms of PDEs gives rise to shorter calculation times.

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

The use of fluid approximation of population and spatial discretisation has been applied in an ad hoc manner to a 2-dimensional space model of delay-tolerant networks [35]. A general approach based on Markovian agents has been proposed for 1-dimensional space which aggregates and fluidises individuals and discretises space.

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. 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 [35]. This enabled the population-based modelling of systems with many more nodes and still provided good approximations.

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 [15]. A detailed definition of Markovian agents is beyond the scope of this tutorial.

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 movement speed of MAs solely depends on the current state of the agents. A new term to describe the agents that move can be derived from the Taylor expansion of the movement term. The change in value can 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 [47] 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.

4.5 Applications of continuous space

As with the case for discrete space, the aim of this section is to briefly consider various applications and a survey can be found in [41].

Ecology: Spatio-temporal point processes have been used to model plant growth and dispersal [8] and other applications [25]. Markov random graphs on continuous space over continuous time can also be considered as spatio-temporal point processes [50]. Holmes *et al* [46] review the use of PDEs in ecological applications, and consider the different forms of PDEs that are used for different models including Brownian (random) motion, drift and the *telegraph* equation.

Biology: Bittig and Uhrmacher [7] describe two continuous space approaches for cellular modelling: particle space and PDEs. In the former, each molecule is modelled separately and these models can be simulated more efficiently by assuming that each particle is only affected by nearby events. When using PDEs, often only simple diffusion based on Brownian motion is required. Fange *et al* [34] describe three different techniques for spatially heterogeneous stochastic kinetics as *microscopic* when each individual particle is considered in terms of its position (continuous-space), as *mesoscopic* when the Reaction Diffusion Master Equation (RDME) is used (discrete space) and as *macroscopic* when PDEs are used. PDEs can also be obtained by taking the hydrodynamic limit of IPSs, namely as the number of particles tends to infinity [23, 30]. Pattern formation is important in biology and an important PDE in this context is the Swift-Hohenberg equation [82].

Epidemiology: Spatial point processes have been used to model the spread of foot and mouth disease [26]. Kendall [54] proposed the first spatial epidemic PDE model

based on the Kermack-McKendrick nonspatial compartment model, and this has been extended to the Diekmann-Thieme model where traits of individuals affect both their susceptibility to infection and their infectiveness to other individuals [77].

Networking: There is a substantial amount of work on mobility models, both at the analytical level and experimentally through traces in the domain of networking [14, 69]. Connectivity models provide an abstraction of mobility models in that they provide information about intercontact time [52]. Stochastic geometry has been applied to wireless networks [2] and epidemiological approaches using PDEs have been used for routing in networks [57].

Continuous-space techniques can be applied to CAS modelling when individual movement is to be tracked, or when it is possible to aggregate movement using PDEs because of the large subpopulation sizes. However, any techniques that tracks individuals is unlikely to be scalable. In the next section, hybrid approaches are considered that can be used to mitigate this problem.

5 Other approaches to modelling space

The 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. Using logic-based approaches, spatial and spatio-temporal model checking can be applied to either sort of model and are addressed in another chapter in this volume [19].

5.1 Crowding

In biological modelling of cells, crowding (occupation of space) is an important issue, because 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 [56, 83]. The lattice-based approaches can be similar to individual discrete-space models but use regular graph rather than arbitrary graphs.

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 maximum population count for a location has been reached. This can lead to discontinuous rate functions. Crowding can be important in CAS, as we may want to impose occupation limits, such as the number of people in a shared taxi, or the capacity of a bike station in a bike-sharing scheme.

5.2 Hybrid approaches

Hybridness is a ubiquitous feature in many models of real systems. As far as space is concerned, there are many ways in which one can construct hybrid models. Here we list some possibilities for future research, with CAS examples from smart transport.

- 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 [6]. This produces a model combining individual objects moving in space with discretised stochastic diffusion process. If we consider models of interaction of pedestrians with public transportation, we can investigate 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 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. 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. An example of collective adaptive system of this kind may be a crowd movement scenario, in which different squares of a city are described in detail, and the flow of people in and out of each square is represented in a location-based style.
- 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 imagine hybrid models in space where small and large populations are both present [9]. 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.

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) [12] or PDMPs (piecewise deterministic Markov processes) [22] can be used.

6 Conclusion

To conclude, this tutorial has provided information about the choices that can be made when modelling space in a quantified manner, focussing on the modelling of CAS. Scalability of techniques have been considered, with specific references to moving away from individual-based modelling to population modelling, using both exact and approximate techniques. There has been an exploration of techniques for both discrete and continuous space, as well a review of how techniques have been applied in the literature, and specific details of techniques that have been considered for CAS.

Acknowledgements This work is supported by the EU project QUANTICOL, 600708. The author thanks Jane Hillston and Mieke Massink for their useful comments.

Appendix: Discrete and continuous time Markov chains

This section briefly introduces these concepts, as they would be used in stochastic modelling both without aggregation of state and with aggregation of state (population-based Markov chains) [4, 10].

Definition 1. A discrete time Markov chain (DTMC) is a tuple $\mathcal{M}_D = (\mathcal{S}, \mathbf{P})$ where

- \mathcal{S} is a finite set of states, and
- $\mathbf{P} : \mathcal{S} \times \mathcal{S} \rightarrow [0, 1]$ is a probability matrix satisfying $\sum_{S' \in \mathcal{S}} \mathbf{P}(S, S') = 1$ for all $S \in \mathcal{S}$.

A DTMC is *time-abstract* [4] in the sense that time is viewed as a sequence of discrete steps or clock ticks. It describes behaviour as follows: if an entity or individual is currently in state $S \in \mathcal{S}$ then the probability of the entity being in state S' at the next time step is defined by $\mathbf{P}(S, S')$. Under certain conditions, the steady state of the DTMC can be determined and this describes when the DTMC is at equilibrium and gives the (unchanging) probability of being in any of the states of \mathcal{S} . By contrast, transient state probabilities can be determined at each point in time before steady state is achieved.

Definition 2. 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

- $\mathbf{R} : \mathcal{S} \times \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$ is a rate matrix.

CTMCs are *time-aware* [4] since they use continuous time. 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. As with DTMCs and 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.

Let $E(S) = \sum_{S' \in \mathcal{S}} \mathbf{R}(S, S')$ be the exit rate of state S . Then the embedded DTMC of a CTMC has entries in its probability matrix of the form $\mathbf{P}(S, S') = \mathbf{R}(S, S')/E(S)$ if $E(S) > 0$ and $\mathbf{P}(S, S') = 0$ otherwise. DTMCs and CTMCs can be state-labelled (usually with propositions) or transition-labelled (usually with actions). The research in QUANTICOL focusses on transition-labelled Markov chains. We next consider population Markov chains, both discrete time and continuous time. Instead of considering an entity with states, we now 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. Our definition in the continuous-time case is slightly simpler than that appearing in another chapter in this volume [9] since transitions do not have guards and we do not parameterise the Markov chain with the population size.

Definition 3. A population discrete time Markov chain (PDTMC) is a tuple $\mathcal{X}_D = (\mathbf{X}, \mathcal{D}, \mathcal{T})$ where

- $\mathbf{X} = (X_1, \dots, X_n)$ is a vector of variables
- \mathcal{D} is a countable set of states defined as $\mathcal{D} = \mathcal{D}_1 \times \dots \times \mathcal{D}_n$ where each $\mathcal{D}_i \subseteq \mathbb{N}$ represents the domain of X_i
- $\mathcal{T} = \{\tau_1, \dots, \tau_m\}$ is the set of transitions of the form $\tau_j = (\mathbf{v}, p)$ where
 - $\mathbf{v} = (v_1, \dots, v_n) \in \mathbb{N}^n$ 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
 - $p : \mathcal{D} \rightarrow \mathbb{R}_{\geq 0}$ is the probability function of transition τ_j that defines a sub-probability distribution, namely $\sum_{\tau \in \mathcal{T}} p_{\tau}(\mathbf{d}) \leq 1$ for all $\mathbf{d} \in \mathcal{D}$, such that $p(\mathbf{d}) = 0$ whenever $\mathbf{d} + \mathbf{v} \notin \mathcal{D}$

Definition 4. A population continuous time Markov chain (PCTMC) is a tuple $\mathcal{X}_C = (\mathbf{X}, \mathcal{D}, \mathcal{T})$ where

- \mathbf{X} and \mathcal{D} are defined as in the previous definition,
- $\mathcal{T} = \{\tau_1, \dots, \tau_m\}$ is the set of transitions of the form $\tau_j = (\mathbf{v}, r)$ where
 - \mathbf{v} is defined as in the previous definition,
 - $r : \mathcal{D} \rightarrow \mathbb{R}_{\geq 0}$ is the rate function of transition τ_j with $r(\mathbf{d}) = 0$ whenever $\mathbf{d} + \mathbf{v} \notin \mathcal{D}$.

In both types of population Markov chain, the associated Markov chain can be obtained. In both cases, \mathcal{D} is the state space \mathcal{S} . For the population DTMC, the probability matrix of its associated DTMC is defined as

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

and since probability functions define sub-probabilities then the rest of the probability mass must be accounted for by defining

$$\mathbf{P}(\mathbf{d}, \mathbf{d}) = 1 - \sum_{\tau \in \mathcal{T}, \mathbf{v}_\tau \neq \mathbf{0}} p_\tau(\mathbf{d}).$$

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 [58] 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 $\mathbf{X}(t) = (X_1(t), \dots, X_n(t))$ defining a trajectory over time. The ODEs can be expressed in terms of a single vector ODE as

$$\dot{\mathbf{X}} = \frac{d\mathbf{X}}{dt} = \mathbf{f}(\mathbf{X})$$

where $\mathbf{f}(\mathbf{X})$ is a function derived from the specifics of the PCTMC (see [9] in this volume for details). It is also possible to approximate the moments of a PCTMC using the ODEs [32]

$$\frac{d}{dt} E[M(\mathbf{X}(t))] = \sum_{\tau \in \mathcal{T}} E[(M(\mathbf{X}(t) + \mathbf{v}_\tau) - M(\mathbf{X}(t)))r_\tau(\mathbf{X}(t))]$$

where $M(\mathbf{X})$ denotes the moment to be calculated, \mathbf{v}_τ and $r_\tau(\mathbf{X}(t))$ represents the update vector and the rate of a transition τ , respectively.

References

1. Auger, P., Poggiale, J., Sánchez, E.: A review on spatial aggregation methods involving several time scales. *Ecological Complexity* 10, 12–25 (Jun 2012)
2. Baccelli, F., Błaszczyszyn, B.: *Stochastic Geometry and Wireless Networks: Volume I and II*. NOW Publishers (2009)
3. Baddeley, A., Bárány, I., Schneider, R.: Spatial point processes and their applications. *Stochastic Geometry: Lectures given at the CIME Summer School 2004* pp. 1–75 (2007)
4. Baier, C., Katoen, J.P., Hermanns, H., Wolf, V.: Comparative branching-time semantics for Markov chains. *Information and Computation* 200, 149–214 (2005)
5. Berec, L.: Techniques of spatially explicit individual-based models: construction, simulation, and mean-field analysis. *Ecological Modelling* 150, 55–81 (2002)
6. Bittig, A., Haack, F., Maus, C., Uhrmacher, A.: Adapting rule-based model descriptions for simulating in continuous and hybrid space. In: *Proceedings of CMSB 2011*. pp. 161–170. ACM (2011)
7. Bittig, A., Uhrmacher, A.: Spatial modeling in cell biology at multiple levels. In: *Winter Simulation Conference (WSC 2010)*. pp. 608–619. IEEE (2010)

8. Bolker, B., Pacala, S.: Using moment equations to understand stochastically driven spatial pattern formation in ecological systems. *Theoretical Population Biology* 52, 179–197 (1997)
9. Bortolussi, L., Gast, N.: Mean-field limits beyond ordinary differential equations. In: *Formal Methods for the Quantitative Evaluation of Collective Adaptive Systems (SFM-16)*. LNCS (2016)
10. Bortolussi, L., Hillston, J., Latella, D., Massink, M.: Continuous approximation of collective systems behaviour: a tutorial. *Performance Evaluation* 70, 317–349 (2013)
11. Bortolussi, L., Latella, D., Massink, M.: Stochastic process algebra and stability analysis of collective systems. In: *Proceedings of COORDINATION 2013*. pp. 1–15. LNCS 7890 (2013)
12. Bortolussi, L., Policriti, A.: Hybrid dynamics of stochastic programs. *Theoretical Computer Science* 411, 2052–2077 (2010)
13. Brauer, F.: Compartmental models in epidemiology. In: Allen, L., Brauer, F., van den Driessche, P., Wu, J. (eds.) *Mathematical Epidemiology*, pp. 19–80. Springer (2008)
14. Camp, T., Boleng, J., Davies, V.: A survey of mobility models for ad hoc network research. *Wireless Communications and Mobile Computing* 2, 483–502 (2002)
15. Cerotti, D., Gribaudo, M., Bobbio, A.: Markovian agents models for wireless sensor networks deployed in environmental protection. *Reliability Engineering & System Safety* 130, 149–158 (2014)
16. Cerotti, D., Gribaudo, M., Bobbio, A., Calafate, C., Manzoni, P.: A Markovian agent model for fire propagation in outdoor environments. In: *Proceedings of EPEW 2010*. pp. 131–146. LNCS 6342, Springer (2010)
17. Chaintreau, A., Le Boudec, J.Y., Ristanovic, N.: The age of gossip: spatial mean field regime. In: *Proceedings of SIGMETRICS/Performance 2009*. pp. 109–120. ACM (2009)
18. Chesson, P.: Scale transition theory: its aims, motivations and predictions. *Ecological Complexity* 10, 52–68 (2012)
19. Ciancia, V., Latella, D., Loreti, M., Massink, M.: Spatial logic and spatial model checking for closure spaces. In: *Formal Methods for the Quantitative Evaluation of Collective Adaptive Systems (SFM-16)*. LNCS (2016)
20. Codling, E., Plank, M., Benhamou, S.: Random walk models in biology. *Journal of the Royal Society Interface* 5, 813–834 (2008)
21. Darling, R., Norris, J.: Differential equation approximations for Markov chains. *Probability Surveys* 5, 37–79 (2008)
22. Davis, M.: *Markov Models and Optimization*. Chapman & Hall (1993)
23. De Masi, A., Presutti, E.: *Mathematical Methods for Hydrodynamic Limits*. Lecture Notes in Mathematics, Springer-Verlag (1991)
24. Desharnais, J., Panangaden, P.: Continuous stochastic logic characterizes bisimulation of continuous-time markov processes. *Journal of Logic and Algebraic Programming* 56, 99–115 (2003)
25. Diggle, P.: *Spatio-temporal point processes: methods and applications*. Working paper, Department of Biostatistics, Johns Hopkins University (2005)
26. Diggle, P.: Spatio-temporal point processes, partial likelihood, foot and mouth disease. *Statistical Methods in Medical Research* 15, 325–336 (Aug 2006)
27. van den Driessche, P.: Spatial structure: Patch models. In: Allen, L., Brauer, F., van den Driessche, P., Wu, J. (eds.) *Mathematical Epidemiology*, pp. 179–190. Springer (2008)
28. Durrett, R., Levin, S.: The importance of being discrete (and spatial). *Theoretical Population Biology* 46, 363–394 (1994)
29. Durrett, R., Levin, S.: Stochastic spatial models: a user’s guide to ecological applications. *Philosophical Transactions of the Royal Society B: Biological Sciences* 343, 329–350 (1994)
30. Durrett, R., Neuhauser, C.: Particle systems and reaction-diffusion equations. *The Annals of Probability* 22, 289–333 (1994)

31. Ellner, S.: Pair approximation for lattice models with multiple interaction scales. *Journal of Theoretical Biology* 210, 435–447 (Jun 2001)
32. Engblom, S.: Computing the moments of high dimensional solutions of the master equation. *Applied Mathematics and Computation* 180, 498–515 (2006)
33. Erban, R., Chapman, J., Maini, P.: A practical guide to stochastic simulations of reaction-diffusion processes. arXiv preprint arXiv:0704.1908 (2007)
34. Fange, D., Berg, O., Sjöberg, P., Elf, J.: Stochastic reaction-diffusion kinetics in the microscopic limit. *Proceedings of the National Academy of Sciences* 107, 19820–19825 (2010)
35. Feng, C.: Patch-based hybrid modelling of spatially distributed systems by using stochastic HYPE – ZebraNet as an example. In: *Proceedings of QAPL 2014* (2014)
36. Feng, C., Hillston, J., Galpin, V.: Automatic moment-closure approximation of spatially distributed collective adaptive systems. *ACM TOMACS* (2015), accepted for publication
37. Fenichel, N.: Persistence and smoothness of invariant manifolds for flows. *Indiana University Mathematics Journal* 21, 1972 (1971)
38. Fishman, G.: *Discrete-event Simulation*. Springer-Verlag (2001)
39. Fricker, C., Gast, N.: Incentives and regulations in bike-sharing systems with stations of finite capacity. arXiv preprint arXiv:1201.1178 (2012)
40. Fricker, C., Gast, N., Mohamed, H.: Mean field analysis for inhomogeneous bike sharing systems. *DMTCS Proceedings* 01, 365–376 (2012)
41. Galpin, V., Feng, C., Hillston, J., Massink, M., Tribastone, M., Tschaikowski, M.: Review of time-based techniques for modelling space. Tech. Rep. TR-QC-05-2014, QUANTICOL (2014)
42. Gear, C.W.: *Numerical Initial Value Problems in Ordinary Differential Equations*. Prentice Hall (1971)
43. Gillespie, D.: Exact stochastic simulation of coupled chemical reactions. *Journal of Physical Chemistry* 81, 2340–2361 (1977)
44. Gillespie, D.: Stochastic simulation of chemical kinetics. *Annual Review of Physical Chemistry* 58, 35–55 (2007)
45. Gorban, A., Radulescu, O., Zinovyev, A.: Asymptotology of chemical reaction networks. *Chemical Engineering Science* 65, 2310–2324 (Apr 2010)
46. Holmes, E., Lewis, M., Banks, J., Veit, R.: Partial differential equations in ecology: spatial interactions and population dynamics. *Ecology* 75, 17–29 (1994)
47. Horton, G., Kulkarni, V., Nicol, D., Trivedi, K.: Fluid stochastic Petri nets: Theory, applications, and solution techniques. *European Journal of Operational Research* 105, 184–201 (1998)
48. Hu, H., Myers, S., Colizza, V., Vespignani, A.: WiFi networks and malware epidemiology. *Proceedings of the National Academy of Sciences* 106, 1318–1323 (Jan 2009)
49. Ilachinski, A.: *Cellular Automata: A Discrete Universe*. World Scientific (2001)
50. Isham, V.: An introduction to spatial point processes and Markov random fields. *International Statistical Review/Revue Internationale de Statistique* 41, 21–43 (1981)
51. Kaneko, K.: Diversity, stability, and metadynamics: Remarks from coupled map studies. In: Bascompte, J., Solé, R. (eds.) *Modeling Spatiotemporal Dynamics in Ecology*, pp. 27–45. Springer-Verlag (1998)
52. Kathiravelu, T., Pears, A.: Reproducing opportunistic connectivity traces using connectivity models. In: *2007 ACM CoNEXT conference*. p. 34. ACM (2007)
53. Kemeny, J., Snell, J.: *Finite Markov Chains*. Springer (1976)
54. Kendall, D.: Mathematical models of the spread of infection. In: *Mathematics and Computer Science in Biology and Medicine*, pp. 213–225. Medical Research Council London (1965)
55. Kephart, J., White, S.: Directed-graph epidemiological models of computer viruses. In: *Proceedings of the IEEE Computer Society Symposium on Research in Security and Privacy*. pp. 343–359. IEEE (1991)

56. Klann, M., Koepl, H.: Spatial simulations in systems biology: from molecules to cells. *International Journal of Molecular Sciences* 13, 7798–7827 (2012)
57. Klein, D., Hespánha, J., Madhoo, U.: A reaction-diffusion model for epidemic routing in sparsely connected MANETs. In: *Proceedings of INFOCOM 2010*. pp. 1–9. IEEE (2010)
58. Kurtz, T.: *Approximation of population processes*. SIAM (1981)
59. Levin, S., Durrett, R.: From individuals to epidemics. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences* 351, 1615–1621 (1996)
60. Lichtenegger, K., Schappacher, W.: A carbon-cycle-based stochastic cellular automata climate model. *International Journal of Modern Physics C* 22, 607–621 (2011)
61. Marion, G., Mao, X., Renshaw, E., Liu, J.: Spatial heterogeneity and the stability of reaction states in autocatalysis. *Physical Review E* 66, 051915 (Nov 2002)
62. Marion, G., Swain, D., Hutchings, M.: Understanding foraging behaviour in spatially heterogeneous environments. *Journal of Theoretical Biology* 232, 127–142 (2005)
63. Massink, M., Brambilla, M., Latella, D., Dorigo, M., Birattari, M.: On the use of Bio-PEPA for modelling and analysing collective behaviors in swarm intelligence. *Swarm Intelligence* 7, 201–228 (2013)
64. Massink, M., Latella, D., Bracciali, A., Harrison, M., Hillston, J.: Scalable context-dependent analysis of emergency egress models. *Formal Aspects of Computing* 24, 267–302 (2012)
65. Matis, J., Kiffe, T.: Effects of immigration on some stochastic logistic models: a cumulant truncation analysis. *Theoretical Population Biology* 56, 139–161 (1999)
66. Matsuda, H., Ogita, N., Sasaki, A., Satō, K.: Statistical mechanics of population: The lattice Lotka-Volterra model. *Progress of Theoretical Physics* 88, 1035–1049 (1992)
67. McCaig, C., Norman, R., Shankland, C.: From individuals to populations: a mean field semantics for process algebra. *Theoretical Computer Science* 412, 1557–1580 (2011)
68. Morozov, A., Poggiale, J.C.: From spatially explicit ecological models to mean-field dynamics: the state of the art and perspectives. *Ecological Complexity* 10, 1–11 (2012)
69. Musolesi, M., Mascolo, C.: Mobility models for systems evaluation. In: Garbinato, B., Miranda, H., Rodrigues, L. (eds.) *Middleware for Network Eccentric and Mobile Applications*, pp. 43–62. Springer (2009)
70. Norris, J.: *Markov Chains*. Cambridge University Press (1998)
71. Okubo, A., Levin, S.A.: *Diffusion and Ecological Problems: Modern Perspectives*. Springer (2001)
72. Othmer, H., Scriven, L.: Instability and dynamic pattern in cellular networks. *Journal of Theoretical Biology* 32, 507–537 (1971)
73. Panangaden, P.: *Labelled Markov Processes*. Imperial College Press (2009)
74. Pascual, M., Roy, M., Laneri, K.: Simple models for complex systems: exploiting the relationship between local and global densities. *Theoretical Ecology* 4, 211–222 (2011)
75. Riley, S.: Large-scale spatial-transmission models of infectious disease. *Science* 316, 1298–1301 (2007)
76. Rowe, J.E., Gomez, R.: El Botellón: Modeling the movement of crowds in a city. *Complex Systems* 14, 363–370 (2003)
77. Ruan, S.: Spatial-temporal dynamics in nonlocal epidemiological models. In: *Mathematics for Life Science and Medicine*, pp. 97–122. Springer (2007)
78. Schoenberg, F., Brillinger, D., Guttorp, P.: Point processes, spatial-temporal. In: *Encyclopedia of Environmetrics*, vol. 3, pp. 1573–1577. Wiley Online Library (2002)
79. Segel, L., Slemrod, M.: The quasi-steady-state assumption: a case study in perturbation. *SIAM review* 31, 446–477 (1989)
80. Simon, H.A., Ando, A.: Aggregation of variables in dynamic systems. *Econometrica* 29, 111–138 (1961)
81. Slepchenko, B., Schaff, J., Macara, I., Loew, L.: Quantitative cell biology with the virtual cell. *Trends in Cell Biology* 13, 570–576 (2003)

82. Swift, J., Hohenberg, P.: Hydrodynamic fluctuations at the convective instability. *Physical Review A* 15, 319 (1977)
83. Takahashi, K., Arjunan, S., Tomita, M.: Space in systems biology of signaling pathways—towards intracellular molecular crowding in silico. *FEBS Letters* 579, 1783–1788 (2005)
84. Thomas, J.W.: *Numerical Partial Differential Equations: Finite Difference Methods*. Springer-Verlag (1995)
85. Tribastone, M., Gilmore, S., Hillston, J.: Scalable differential analysis of process algebra models. *IEEE Transactions on Software Engineering* 38, 205–219 (2012)
86. Tschaikowski, M., Tribastone, M.: A Partial-differential Approximation for Spatial Stochastic Process Algebra. In: *Proceedings of VALUETOOLS 2014* (2014)
87. Tschaikowski, M., Tribastone, M.: Exact fluid lumpability in markovian process algebra. *Theoretical Computer Science* 538, 140–166 (2014)
88. Tschaikowski, M., Tribastone, M.: Spatial fluid limits for stochastic mobile networks. *Performance Evaluation* (2015), under minor revision
89. Turing, A.: The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society (Part B)* 237, 37–72 (1953)
90. Vandin, A., Tribastone, M.: Quantitative abstractions for collective adaptive systems. In: *Formal Methods for the Quantitative Evaluation of Collective Adaptive Systems (SFM-16)*. LNCS (2016)
91. Webb, S., Keeling, M., Boots, M.: 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)
92. Wiggins, S.: *Normally hyperbolic invariant manifolds in dynamical systems*. Springer Science & Business Media (1994)
93. Wu, J., Loucks, O.: From balance of nature to hierarchical patch dynamics: a paradigm shift in ecology. *The Quarterly Review of Biology* 70, 439–466 (1995)
94. Zhou, X., Ioannidis, S., Massoulié, L.: On the stability and optimality of universal swarms. *ACM SIGMETRICS Performance Evaluation Review* 39, 301–312 (2011)
95. Zungeru, A., Ang, L.M., Seng, K.P.: Classical and swarm intelligence based routing protocols for wireless sensor networks: A survey and comparison. *Journal of Network and Computer Applications* 35, 1508–1536 (Sep 2012)