

# Location Aggregation of Spatial Population CTMC Models

Luca Bortolussi

University of Trieste, Trieste, Italy  
CNR-ISTI, Pisa, Italy  
luca@dmi.units.it

Cheng Feng

University of Edinburgh  
Edinburgh, UK  
s1109873@sms.ed.ac.uk

In this paper we focus on spatial Markov population models, describing the stochastic evolution of populations of agents, explicitly modelling their spatial distribution, representing space as a discrete, finite graph. More specifically, we present a heuristic approach to aggregating spatial locations, which is designed to preserve the dynamical behaviour of the model whilst reducing the computational cost of analysis. Our approach combines stochastic approximation ideas (moment closure, linear noise), with computational statistics (spectral clustering) to obtain an efficient aggregation, which is experimentally shown to be reasonably accurate on two case studies: an instance of epidemic spreading and a London bike sharing scenario.

## 1 Introduction

Population processes are widely used to describe a large variety of systems, including systems in biological [17], ecology [1], performance engineering [26], smart cities facilities like bike sharing [16, 18], and the spreading of epidemics [2]. Most of these systems are spatially distributed, and space has to be explicitly modelled to properly capture the relevant features of their dynamics. For example, interactions may only be allowed for entities which are co-located or within a certain physical distance of each other, or space may be segmented in such a way that even physically close entities are unable to communicate or interact. Furthermore, movement in space can be a crucial aspect of the behaviour of entities within the system. Epidemic spreading, in which infection can pass only by physical contact, is a clear example in this sense [2]. Bike sharing is another one, as the geographical location of bike stations influences the travelling time, and different stations have different demands for acquisition or deposit of bikes at different times of the day [16, 18].

Population processes are often modelled as Markovian stochastic processes, mostly as a subclass of continuous-time Markov chains (CTMCs) known as Population CTMCs (PCTMCs). These models have a very large or even infinite state space, a fact resulting in a lot of work in the computer science community to craft efficient algorithms for their analysis, ranging from specialised stochastic simulation to the use of mean-field [6] and moment closure techniques [20, 19]. These methods, in particular, approximate the large or infinite set of linear Kolmogorov equations by a much smaller set of non-linear differential equations, capturing the mean, variance, and possibly other higher order moments.

Spatial extensions of PCTMCs typically introduce discrete representations of space in terms of locality, connecting them in a general topology represented by a (weighted) graph [7]. One of the effects of explicit modelling of space is the increase of the computational cost of analysis of the system. For instance, a model with  $l$  localities will increase the number of equations for the variance in any moment closure approach by a factor of  $O(l^2)$ .

However, although space shapes the behaviour of the system, our interest is often in space-free properties, like the total number of infected individuals in an epidemic scenario, or in localised properties,

like the number of available bikes in a given station or geographic area. In these cases, a full representation of space may not be needed to compute such quantities with a reasonable accuracy, and model simplification and abstraction can be a viable strategy.

In this paper, we pursue this direction by developing an efficient approach to aggregate locations showing a similar dynamical profile. This operation will preserve the dynamical behaviour of the model, but lower the computational cost: for instance, if the number of locations decreases from  $l$  to  $k$ , then the number of differential equations for the variance is reduced by a factor  $O(k^2/l^2)$ . In particular, here we will focus on the reduced cost of the analysis of the full stochastic model by using standard simulation algorithms [17], as this is the most expensive but also most informative computational technique (excluding numerical integration of the Kolmogorov equations, which for PCTMC is unfeasible due to state space explosion).

Our approach is based on aggregating locations using state-of-the-art spectral clustering approaches, based on metrics between locations that take the steady state mean or distribution into account. Rather than working with exact solution or estimations of these quantities by simulation, which would be computationally expensive, we obtain them applying stochastic approximation ideas, either by solving mean-field equations for the mean or using a linear-noise-like approximation of the distribution, obtaining mean and variance from moment closure equations. The feasibility, effectiveness, and accuracy of our method is discussed on two case studies: an epidemic model and a London bike sharing scenario.

The paper is organised as follows: Section 2 introduces the formalism of spatial PCTMCs we will use afterwards and briefly introduces exact and approximate analysis techniques. Section 3 discusses in detail our aggregation approach, and Section 4 presents the two case studies. Conclusions are drawn in Section 5.

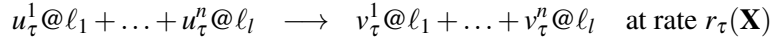
## 2 Spatial PCTMC Models

A Population Continuous Time Markov Chain (PCTMC) is a Markovian stochastic process evolving in continuous time. It consists of a number of individually indistinguishable interacting agents of different types which can be in different internal states, so that the state of the system can be described by counting how many agents of each kind are in the system. Agents' interactions are described by a set of transitions, which will change the internal state of one or more agents. In this paper, we specifically consider *spatial* PCTMC models in which agents are distributed in a finite set of discrete locations, with interactions typically happening either in the same location or in neighbouring ones. Formally, a spatial PCTMC can be expressed as a tuple  $\mathcal{P} = (\mathcal{L}, \mathbf{X}, \mathcal{T})$ :

- $\mathcal{L} = (\ell_1, \dots, \ell_l)$  is the set of discrete locations in the model, where  $|\mathcal{L}| = l$  denotes the total number of locations.
- $\mathbf{X} = (\mathbf{X}@_{\ell_1}, \dots, \mathbf{X}@_{\ell_l}) \in \mathbb{Z}_{\geq 0}^N$  is an integer vector representing the agent populations distributed over all the locations in the model, where  $|\mathbf{X}| = N = l \times n$  is the total number of distinct agent populations in the model with  $n$  representing the number of agent types.  $\mathbf{X}@_{\ell_i}$  is a sub-vector with the  $j$ th component,  $X_j@_{\ell_i}$  representing the current population of the agent type  $j$  at location  $i$ . We use  $\mathbf{X}_0$  to denote the initial state of the model.
- $\mathcal{T} = \{\tau_1, \dots, \tau_m\}$  is the set of transitions with size  $|\mathcal{T}| = m$ , of the form  $\tau = (r_\tau(\mathbf{X}), \mathbf{D}_\tau)$ , where:
  1.  $r_\tau(\mathbf{X}) \in \mathbb{R}_{\geq 0}$  is the rate function, associating with each transition the rate of an exponential distribution, depending on the global state of the model.

2.  $\mathbf{D}_\tau \in \mathbb{Z}^N$  is the update vector which gives the net change for each element of  $\mathbf{X}$  caused by transition  $\tau$ . Similarly, we let  $\mathbf{D}_\tau @ \ell_i$  denote the update vector of agent populations at location  $i$ .

Transition rules can be easily visualised in the chemical reaction style, as



where the net change of agents of type  $i$  in location  $\ell_j$  due to transition  $\tau$  is given by  $D_\tau^i @ \ell_j = v_\tau^i @ \ell_j - u_\tau^i @ \ell_j$  ( $1 \leq i \leq n$ ,  $1 \leq j \leq l$ ). This is a general format, encompassing local interactions, interactions between agents in neighbouring locations, and movement in space. The tuple  $\mathcal{P}$  contains all the information that is needed to build a CTMC on the state space  $\mathbb{Z}_{\geq 0}^N$ : its infinitesimal generator  $Q$  is given by  $Q(\mathbf{x}, \mathbf{y}) = \sum \{r_\tau(\mathbf{x}) | \mathbf{D}_\tau = \mathbf{y} - \mathbf{x}\}$ , for  $\mathbf{x} \neq \mathbf{y}$ . Spatial PCTMC can be simulated using standard stochastic simulation algorithms [17].

### 2.1 Moment Closure, Mean Field, and Linear Noise

The analysis of the stochastic model underlying a spatial PCTMC is not an easy task. Numerical methods are hindered by the state space explosion, and even stochastic simulation suffers from the presence of localities and large populations. Furthermore, it is known that for large populations the behaviour of the stochastic model becomes deterministic and converges to the solution of a the mean field differential equation [6, 22], which in the formalism defined above takes the form

$$\dot{\mathbf{X}} = F(\mathbf{X}), \quad \text{with } F(\mathbf{X}) = \sum_{\tau} \mathbf{D}_\tau r_\tau(\mathbf{X})$$

If populations in each location are in the order of tens or hundreds of thousands, mean-field equations are generally very accurate [6]. However, when populations are smaller, in the order of hundreds, stochastic fluctuations still have an important role, and the linear-noise (central limit) approximation performs better [22, 21]. The idea behind linear noise is to approximate the original stochastic model by a linearized Markov process in continuous states, whose solution is a Gaussian process. The distribution of this process at a given time is thus characterized by solving equations for the mean (incidentally, the mean-field equations above) and equations for the covariance.

An alternative strategy for the analysis of the spatial PCTMC is to derive equations for the moments of the population variables [20, 19], up to order  $k$ . Due to non-linearities in the rates, there is no exact closed form of these equations, and the differential equations for moments of order  $k$  depend on higher-order ones. Hence, equations are closed by relying to some heuristic [20, 19]. Typically, moment closure equations give a better estimation of mean and variance than linear noise, as knowledge of higher-order moments introduce correction terms in the equations of lower order ones. The first two moments can still be used to build a Gaussian approximation to the true distribution at a given time  $t$ , formally invoking a maximum entropy argument [3, 9].

## 3 Aggregation of Locations of Spatial PCTMC

In this section we present the computational methodology to reduce a spatial model by aggregating locations. The main motivation underlying this approach, as discussed in the introduction, is to reduce the computational effort in the analysis of the model. This effort, in fact, is proportional to the number of locations. This is the case for stochastic models, which need to simulate and keep track of the state of

populations in each location, and for approximate analysis techniques like moment closure. In fact, even when considering only equations for the second order moments, the number of such equations grows quadratically with the number of locations. However, if our final goal is to capture the behaviour of the system at the global level (i.e. to know how many agents of each kind are in the system), then dealing with the full set of locations may incur excessive work which can be reduced by grouping together locations having a similar overall behaviour. To achieve this goal while keeping the error low, we propose the following heuristic scheme, sketched here and detailed in the following subsections.

1. Define distance metrics between locations taking into account the dynamical behaviour of the system (Section 3.1). More specifically, we want locations showing a similar steady state behaviour to be clustered together. We will define two distances, described below in increasing order of precision:
  - (a) Mean field distance. We will consider the distance between the mean of populations at steady state, approximated by first order moment closure, which corresponds to the mean field abstraction of the PCTMC model (see Section 2.1).
  - (b) Linear noise distance. We will consider the distance between a Gaussian approximation of the steady state distribution, computed from moment closure equations for mean and variance.
2. Cluster the locations using the previously computed distance (Section 3.2). We will use a spectral clustering algorithm on graphs, exploiting the eigengap heuristic to identify the number of clusters.
3. Given a clustering of locations, construct the reduced model by suitably aggregating together the PCTMC transitions and variables (Section 3.3).

### 3.1 Distance between Locations

We will consider distances between populations of different locations at a given time. The choice of this time is important, and should not be taken too small, in order to minimise the effect of the initial state of the model. In fact, if we took information about the whole trajectory into account, different initial conditions between the populations of two locations would contribute to the distance, often resulting in a separation between the two locations. However, especially if we consider aggregated quantities, like the total number of agents of a certain kind across all locations, this difference is not very relevant, and better results can be obtained by comparing the behaviour after a finite but large time. We will see the experimental validation of this choice in the next Section. Note that we do not consider steady state behaviour (though for a very large time steady state would be reached), as we will use mean-field or linear noise approximations, which do not necessarily converge at steady state. We will consider two distances between each pair of locations  $\ell_i$  and  $\ell_j$ , with increasing levels of accuracy:

1. *Mean field distance*  $d_E(\ell_i, \ell_j)$ . This is just the Euclidean distance between the average value of populations of locations  $\ell_i$  and  $\ell_j$  at steady state, i.e.  $d_E(\ell_i, \ell_j) = \|\mu_{\mathbf{X}@\ell_i} - \mu_{\mathbf{X}@\ell_j}\|$ , where  $\mu_{\mathbf{X}@\ell}$  is the mean of the population (vector)  $\mathbf{X}@\ell$  at steady state. To compute this distance efficiently, we resort to the mean-field approximation or to any first order moment closure, see Section 2.1. This reduces the problem to the numerical integration of  $N$  differential equations ( $N = n \times l$ ), until they reach equilibrium (or until their temporal average stabilise in case they oscillate).<sup>1</sup>

---

<sup>1</sup>We are implicitly assuming that the mean field or moment closure equations will not show chaotic behaviour, which is typically the case for models of interacting agents.

2. *Linear noise distance*  $d_L(\ell_i, \ell_j)$ . To capture more accurately the steady state behaviour, we can consider a distance between the full distributions. To this end, we will resort to the Bhattacharyya distance  $d_B(Y_1, Y_2)$  [4] which measures the distance of the distribution of two one-dimensional variables  $Y_1, Y_2$ . more specifically, we will compute the Bhattacharyya distance between each population  $k$  in locations  $\ell_i$  and  $\ell_j$ , denoted by  $d_B(X_k @ \ell_i, X_k @ \ell_j)$ , and aggregate over all populations by taking the average:

$$d_L(\ell_i, \ell_j) = \frac{1}{n} (d_B(X_1 @ \ell_i, X_1 @ \ell_j) + \dots + d_B(X_n @ \ell_i, X_n @ \ell_j))$$

In order to compute  $d_B(X_k @ \ell_i, X_k @ \ell_j)$  without the need to estimate the full distribution, we make a linear noise assumption, i.e. that the steady state distribution is approximately Gaussian, see Section 2.1. Under this hypothesis,  $d_B(X_k @ \ell_i, X_k @ \ell_j)$  can be calculated by the following equation:

$$D_B(X_k @ \ell_i, X_k @ \ell_j) = \frac{1}{4} \ln \left( \frac{1}{4} \left( \frac{\sigma_{X_k @ \ell_i}^2}{\sigma_{X_k @ \ell_j}^2} + \frac{\sigma_{X_k @ \ell_j}^2}{\sigma_{X_k @ \ell_i}^2} + 2 \right) \right) + \frac{1}{4} \left( \frac{(\mu_{X_k @ \ell_i} - \mu_{X_k @ \ell_j})^2}{\sigma_{X_k @ \ell_i}^2 + \sigma_{X_k @ \ell_j}^2} \right)$$

where  $\mu_{X @ \ell}$  and  $\sigma_{X @ \ell}$  denote the mean and variance of the population variable  $X @ \ell$  at the steady state. To numerically compute the values of  $\mu_{X @ \ell}$  and  $\sigma_{X @ \ell}$ , we resort to the normal moment-closure approximation of [19], which can be obtained at a much lower computational cost than by simulating the PCTMC, by integrating  $O(N^2)$  differential equations.

We observe that the cost of computing  $d_L(\ell_i, \ell_j)$  is significantly higher than the cost of computing  $d_E(\ell_i, \ell_j)$ , as we need to integrate  $O(N^2)$  differential equations rather than  $O(N)$ . However, this cost is balanced by a higher accuracy in the reduced system, though for certain models (essentially those having similar variance for the same agent kind at different locations in a cluster) accuracy is comparable. In practice,  $d_E$  should be used when the cost of solving moment closure equations for the variance is very high due to the very large number of locations. Hence, aggregation with respect to  $d_E$  can be seen as a reduction of the number of moment closure equations. When the goal is to reduce the cost of stochastic simulation, and solving moment closure equations for the variance is cheap, it is better to rely on the metric  $d_L$ . In general, our method should be applied when stochastic simulation of the spatial PCTMC model requires excessive computational time, and computing  $d_E$  or  $d_L$  is much cheaper than the cost of stochastic simulation.

### 3.2 Spectral Clustering

Spectral clustering methods are common graph-based approaches to (unsupervised) clustering of data [28]. The dataset is composed of  $n$  objects  $S = s_1, \dots, s_n$ , among which some local symmetric and non-negative similarity measure  $A_{i,j}$  is defined.  $A_{i,j}$  is often obtained from a distance or difference measure  $d(s_i, s_j)$  between the objects. This information is then arranged in a weighted graph  $G = (S, A)$ . Within this framework, clustering is translated into a graph partitioning problem. The most common class of spectral approaches for graph partitioning (in  $k$  subsets) is to map the original data into the first  $k$  eigenvectors of some normalized version of the similarity matrix  $A$  and then apply a standard clustering algorithm such as k-means on these new coordinates.

Among the most commonly used spectral clustering algorithms are the unnormalized spectral clustering [28], normalized spectral clustering according to Shi and Malik [25], and normalized spectral clustering according to Ng, Jordan, and Weiss [24]. The three algorithms are very similar; their main distinguishing feature is the fact that they use three different graph Laplacians. As an illustration, the

normalized spectral clustering algorithm according to Ng, Jordan, and Weiss is given as follows. Given a set of objects  $S = s_1, \dots, s_n$  that we want to cluster into  $k$  subsets:

1. Form the similarity matrix  $A \in \mathbb{R}^{n \times n}$  defined by  $A_{ij} = \exp(-\frac{d(s_i, s_j)^2}{2\sigma^2})$  if  $i \neq j$ , and  $A_{ii} = 0$ .
2. Define  $D$  to be the diagonal matrix whose  $(i, i)$ -element is the sum of  $A$ 's  $i$ -th row, and construct the normalized Laplacian matrix  $L = I - D^{-1/2}AD^{-1/2}$
3. Find  $U_1, U_2, \dots, U_k$ , the first  $k$  eigenvectors of  $L$  with  $k$  smallest eigenvalues, and form the matrix  $U = [U_1, U_2, \dots, U_k] \in \mathbb{R}^{n \times k}$  by stacking the eigenvectors in columns.
4. Form the matrix  $Z$  from  $U$  by renormalizing each of  $U$ 's rows to have unit length (i.e.  $Z_{ij} = U_{ij}/(\sum_j U_{ij}^2)^{1/2}$ ).
5. Treating each row of  $Z$  as a point in  $\mathbb{R}^k$ , cluster them into  $k$  clusters via k-means.
6. Finally, assign the original object  $s_i$  to cluster  $j$  if and only if row  $i$  of the matrix  $Z$  was assigned to cluster  $j$ .

We will use the above algorithm hereafter for the clustering of locations in spatial PCTMCs .

### 3.2.1 Application to Aggregation of Locations of spatial PCTMC

Given a spatial PCTMC model we wish to aggregate its locations  $\mathcal{L} = (\ell_1, \dots, \ell_l)$ . The first step is to compute one of the two distances  $d_E$  or  $d_L$  of the previous section. From this metric, we can derive the similarity matrix  $A \in \mathbb{R}^{l \times l}$  for the locations  $\mathcal{L}$  by using Gaussian kernel with width  $\sigma$ :

$$A_{ij} = \exp\left(-\frac{d_*(\ell_i, \ell_j)^2}{2\sigma^2}\right)$$

Then, the standard spectral clustering algorithms can be applied to cluster the locations with a specific choice of the number of clusters  $k$ . We will refer to the locations' clusters by  $\hat{\mathcal{L}} = (\hat{\ell}_1, \dots, \hat{\ell}_k)$ .

In order to select  $k$ , we rely on the eigengap heuristic [23]. Specifically, we choose the number of clusters  $k$  such that all eigenvalues  $\lambda_1, \dots, \lambda_k$  are very small, but  $\lambda_{k+1}$  is relatively large. A realisation of this heuristic will be shown while discussing case studies.

## 3.3 Model Reduction

In this section, we show how to generate a reduced version of a spatial PCTMC model once the locations  $\mathcal{L}$  in the original model have been clustered into the aggregated locations  $\hat{\mathcal{L}}$ . Suppose we have clustered locations in a spatial PCTMC into  $k$  subsets. Formally, we want to map the original spatial PCTMC model  $\mathcal{P} = (\mathcal{L}, \mathbf{X}, \mathcal{T})$  to a reduced one  $\hat{\mathcal{P}} = (\hat{\mathcal{L}}, \hat{\mathbf{X}}, \hat{\mathcal{T}})$ , where  $|\hat{\mathcal{L}}| = k < l = |\mathcal{L}|$ ,  $|\hat{\mathbf{X}}| < |\mathbf{X}|$ , and  $|\hat{\mathcal{T}}| \leq |\mathcal{T}|$ . Hence, we need to construct both the aggregated vector of the populations of agents and the reduced set of transitions.

### 3.3.1 Generating the aggregated vector of agent populations

The algorithm for generating the aggregated vector of agent populations is fairly straightforward. The basic idea is to treat agents of the same type in the same cluster of locations as identical agents. Thus, we only need to sum up the populations of those identical agents. Algorithm 1 gives the corresponding pseudo code for generating the aggregated vector of agent populations. Note that in the pseudo code, we use the same notations for the original spatial PCTMC model as in Section 2.

---

**Algorithm 1** The Algorithm for generating the aggregated vector of agent populations

---

**Require:**  $\mathbf{X}, \hat{\mathbf{X}}, \mathcal{L}, \hat{\mathcal{L}}$

**Ensure:**  $|\hat{\mathbf{X}}| = n \times k \wedge \hat{X}_i @ \hat{\ell}_j = 0 \quad \forall \hat{X}_i @ \hat{\ell}_j \in \hat{\mathbf{X}} \quad \{\text{all elements in } \hat{\mathbf{X}} \text{ are set to zero initially}\}$

```

1: for  $i = 1$  to  $l$  do
2:   for  $j = 1$  to  $c$  do
3:     if  $\ell_i$  belongs to cluster  $j$  then
4:       let  $\hat{\mathbf{X}} @ \hat{\ell}_j = \hat{\mathbf{X}} @ \hat{\ell}_j + \mathbf{X} @ \ell_i$ 
5:     end if
6:   end for
7: end for
8: return  $\hat{\mathbf{X}}$ 

```

---

### 3.3.2 Generating the reduced set of transitions

Three steps are taken to generate the reduced set of transitions. The first step is to copy the transitions in the original spatial PCTMC model to  $\hat{\mathcal{T}}$ . Meanwhile, the update of agent populations in the transition should be replaced by update of corresponding aggregated populations, the agent populations appearing in the rate function should also be replaced accordingly. The first step may generate many redundant transitions in which there is no agent population being updated. Thus, the second step is to remove redundant transitions from  $\hat{\mathcal{T}}$ . The last step is de-duplication, in which we combine all transitions with the same update vector into one transition. Algorithm 2 gives the pseudo code for the three steps, in which we use  $|\hat{\ell}_i|$  to denote the number of locations belonging to cluster  $\hat{\ell}_i$ . Note that this construction will produce an approximate model with respect to the original one, the reason being the treatment of rates. Approximation stems from the fact that in the rates of the aggregated model we replace all occurrences of each  $X @ \ell$  with  $X @ \hat{\ell} / |\hat{\ell}|$ , where  $\ell \in \hat{\ell}$ . Hence we assume that for each  $\ell_i, \ell_j \in \hat{\ell}$ ,  $X @ \ell_i = X @ \ell_j$ . However, the heuristics we use to construct the aggregated model guarantee that this condition should be roughly satisfied (at steady state).

## 4 Case Studies

In this section we test our method on two case studies: a benchmark model of epidemic spreading, Section 4.1, and a realistic model of a portion of the London bike-sharing system, Section 4.2.

### 4.1 Spatial Epidemic Spreading Model

We first consider a classical epidemiological SIS model of individuals partitioned into  $m$  communities, where individuals move between communities but infections only take place within communities. Each individual is considered to be susceptible (S) or infected (I) with respect to the disease. A continuous-time SIS epidemiological model is then applied to this population as follows: each individual, regardless of susceptible or infected, can move to his/her connected communities with a specific rate ( $r_{ij}$ ). Each community is connected with three other randomly chosen communities. Each infected individual can randomly make contact with a susceptible individual in the same community, and infect her with rate  $\beta_i$  in community  $i$ . Finally, infected individuals independently recover to the susceptible state at rate  $\mu$ .

**Algorithm 2** The Algorithm for generating the reduced set of transitions**Require:**  $\mathcal{L}, \hat{\mathcal{L}}, \mathcal{T}, \hat{\mathcal{T}}$ **Ensure:**  $|\hat{\mathcal{T}}| = 0$ 


---

```

1: for all  $\tau$  in  $\mathcal{T}$  do
2:   create a new  $\hat{\tau}$ , set  $r_{\hat{\tau}} = r_{\tau}$ ,  $|\mathbf{D}_{\hat{\tau}}| = n \times k$   {all elements in  $\mathbf{D}_{\hat{\tau}}$  are set to zero initially}
3:   for  $i = 1$  to  $l$  do
4:     for  $j = 1$  to  $k$  do
5:       if  $\ell_i$  belongs to cluster  $j$  then
6:         let  $\mathbf{D}_{\hat{\tau}}@{\hat{\ell}}_j = \mathbf{D}_{\hat{\tau}}@{\ell}_j + \mathbf{D}_{\tau}@{\ell}_i$ 
7:       end if
8:     end for
9:   end for
10:  for all  $X@{\ell}$  appears in  $r_{\hat{\tau}}$  do
11:    for  $i = 1$  to  $k$  do
12:      if  $\ell$  belongs to cluster  $i$  then
13:        replace  $X@{\ell}$  with  $X@{\hat{\ell}}_i/|\hat{\ell}_i|$ 
14:      end if
15:    end for
16:  end for
17: end for
18: for all  $\hat{\tau}$  in  $\hat{\mathcal{T}}$  do
19:   if All elements in  $\mathbf{D}_{\hat{\tau}}$  equal zero then
20:     remove  $\hat{\tau}$  from  $\hat{\mathcal{T}}$ 
21:   end if
22: end for
23: for all  $\hat{\tau}_i, \hat{\tau}_j$  ( $i \neq j$ ) in  $\hat{\mathcal{T}}$  do
24:   if  $\mathbf{D}_{\hat{\tau}_i} = \mathbf{D}_{\hat{\tau}_j}$  then
25:     create a new  $\hat{\tau}$ , set  $\mathbf{D}_{\hat{\tau}} = \mathbf{D}_{\hat{\tau}_i}$  and  $r_{\hat{\tau}} = r_{\hat{\tau}_i} + r_{\hat{\tau}_j}$ 
26:     remove  $\hat{\tau}_i, \hat{\tau}_j$  from  $\hat{\mathcal{T}}$ , add  $\hat{\tau}$  to  $\hat{\mathcal{T}}$ 
27:   end if
28: end for
29: return  $\hat{\mathcal{T}}$ 

```

---

The model can be studied by a spatial PCTMC containing the following transitions:

$$\begin{array}{lll}
S(\ell_i) \rightarrow I(\ell_i) & \text{at } \beta_i \#(S(\ell_i)) \#(I(\ell_i)) & \forall i \in m \\
I(\ell_i) \rightarrow S(\ell_i) & \text{at } \mu \#(I(\ell_i)) & \forall i \in m \\
S(\ell_i) \rightarrow S(\ell_j) & \text{at } r_{ij} \#(S(\ell_i)) & \forall i, j \in m, \ell_i \text{ and } \ell_j \text{ are connected} \\
I(\ell_i) \rightarrow I(\ell_j) & \text{at } r_{ij} \#(I(\ell_i)) & \forall i, j \in m, \ell_i \text{ and } \ell_j \text{ are connected}
\end{array}$$

where  $S(\ell_i), I(\ell_i)$  represent a susceptible, infected individual in Community  $i$ ;  $\beta_i$  represents the contact rate in community  $i$ ;  $r_{ij}$  denotes the rate for an individual to travel from Community  $i$  to Community  $j$ .

In our experiment, we consider a model with  $m = 30$ ,  $\beta_i$  and  $r_{ij}$  to be random values between zero and one,  $\mu = 0.1$ . We first report the analysis using the Linear-Noise metric  $d_L$ . Computing the distance and running the spectral clustering algorithm, according to the previous section, we obtained the spectre of the normalized Laplacian matrix shown in Figure 1 (left), for the smallest 10 eigenvalues. As can be seen from the figure, the first four eigenvalues are very close to zero, and there is a large gap between the 4th and 5th eigenvalues. Thus, we set the number of clusters to four. Figure 2 (left) shows the



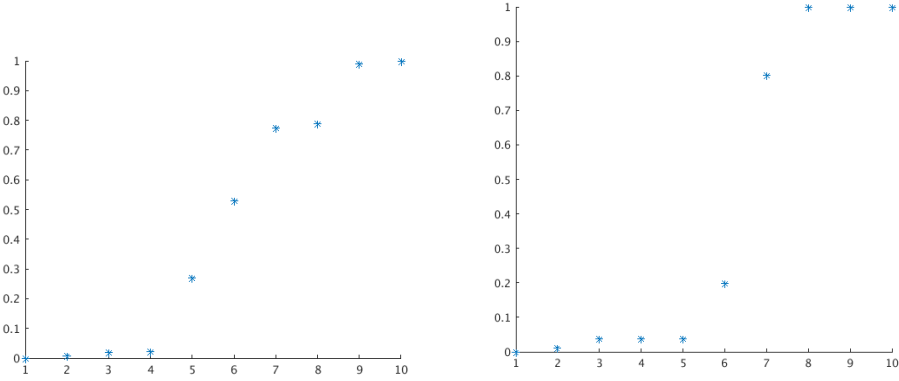


Figure 1: (left) The smallest 10 eigenvalues of the normalized Laplacian matrix of the SIS model. (right) The smallest 10 eigenvalues of the normalized Laplacian matrix of the bike-sharing model.

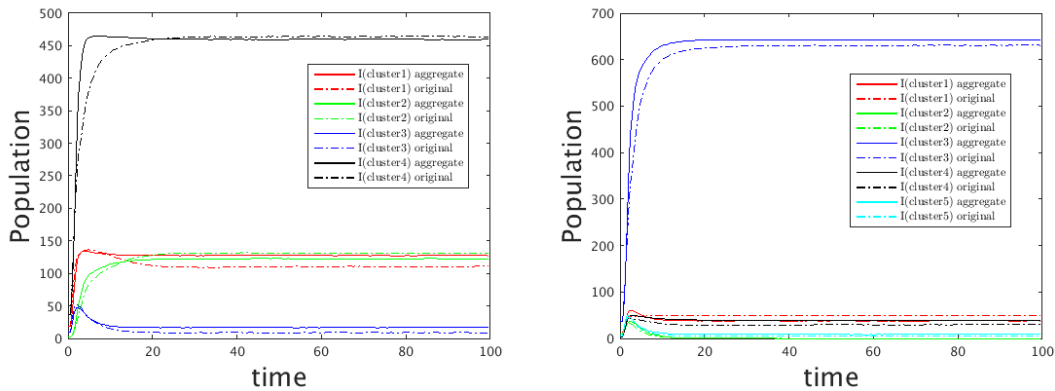


Figure 2: (left) Comparison of the infected populations before and after aggregation for the SIS model, using the  $d_L$  distance. (right) Comparison of the infected populations before and after aggregation for the SIS model, using the  $d_E$  distance.

trajectories of the infected population generated by stochastic simulation before and after aggregation. Table 1 compares the number of transitions, simulation time of 1000 runs of the SIS model before and after aggregation, as well as the average error ratio of the trajectories in Figure 2 (left) after aggregation compared with the counterpart before aggregation. As can be readily seen, our method considerably reduced the simulation cost of the model, at the price of a reasonably small loss of accuracy, the relative error being less than 10% (the average error ratio is computed as the average along the trajectory). In Figure 2 (right), instead, we show the result of the aggregation when using the mean-field distance  $d_E$ . In this case, the spectral clustering identifies 5 clusters, with a comparable overall accuracy with respect to the linear-noise distance as can be seen in Table 1.

## 4.2 Public Bike-sharing Model

The second example is a spatial PCTMC which models a public bike-sharing system. Bike-sharing systems are becoming more and more important for urban transportation. In such systems, users arrive

SIS model	No. of transitions	simulation time (1000 runs)	Avg error ratio
Before aggregation	240	8.81 mins	N/A
After aggregation ( $d_L$ )	28	59 secs	8.67%
After aggregation ( $d_E$ )	34	3.67 mins	10.44%

Table 1: Size, simulation cost (including the aggregation cost) of the SIS model before and after aggregation, and error introduced by the aggregation.

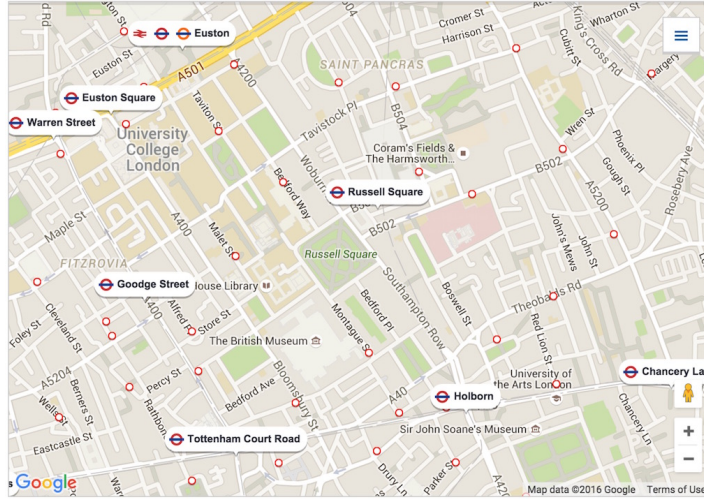


Figure 3: The map of bike-sharing stations near Russell Square in London in which red circles representing bike stations

at a station, pick up a bike, use it for a while, and then return it to another station of their choice. Recently, PCTMCs have been used to model bike-sharing systems [16, 18]. Here, we consider a map which consists of  $N$  zones. There is one public bike station in each zone. Each station has several bike slots. The pickup rate of bikes in a station is governed by an exponential distribution. When a user picks up a bike, the available number of bikes in the station will decrease by one whereas the available number of slots in that station will increase by one. The user will choose another zone in the city as their destination. When the user arrives at the destination zone, they will return their bike to the bike station in that zone. We use a spatial PCTMC containing the following transitions to represent the model:

$$\begin{aligned}
 \text{Bike}(\ell_i) &\rightarrow \text{Slot}(\ell_i) + \text{BikeTo}_j(\ell_i) && \text{at } \lambda_i p_j^i \\
 \text{BikeTo}_j(\ell_i) + \text{Slot}(\ell_j) &\rightarrow \text{Bike}(\ell_j) && \text{at } \#(\text{BikeTo}_j(\ell_i)) \mu_j^i
 \end{aligned}$$

where  $\text{Bike}(\ell_i)$  and  $\text{Slot}(\ell_i)$  denote an available bike or slot in the bike station in Zone  $i$ , respectively;  $\text{BikeTo}_j(\ell_i)$  denotes a bike in transit from Zone  $i$  to Zone  $j$ .  $\lambda_i$  is the pickup rate of bikes in the bike station in Zone  $i$ ,  $p_j^i$  is the probability to choose Zone  $j$  as the destination of a trip when picking up a bike from Zone  $i$ .  $1/\mu_j^i$  is the mean trip time from Zone  $i$  to Zone  $j$ .

Specifically, we use the above spatial PCTMC to model 30 bike stations near Russell Square in London which is illustrated in Figure 3. All the rates in the model are calculated by journey data which is available online<sup>2</sup>.

<sup>2</sup><https://tfl.gov.uk/info-for/open-data-users/our-feeds?intcmp=3671#on-this-page-4>

Bike model	No. of transitions	simulation time (1000 runs)	Avg error ratio
Before aggregation	1800	15.6 mins	N/A
After aggregation ( $d_L$ )	300	2.5 mins	11.94%
After aggregation ( $d_E$ )	180	1.6 mins	23.06%

Table 2: Size, simulation cost (including the aggregation cost) of the bike-sharing model before and after aggregation, and error introduced by the aggregation.

We apply our method with the linear noise distance to cluster the 30 bike stations. The smallest 10 eigenvalues of the normalized Laplacian matrix, computed according to the spectral clustering recipe, are shown in Figure 1 (right). According to the eigengap heuristic, there are 5 well separated clusters. Figure 4 (left) shows the trajectories of the number of available bikes in the 5 clusters generated by stochastic simulation before and after aggregation. Table 2 shows the number of transitions, simulation time of 1000 runs of the bike-sharing model before and after aggregation, as well as the average error ratio of the trajectories in Figure 4 (left) after aggregation compared with the counterpart before aggregation. In Figure 4 (right), instead, we compare the trajectories of the original model and the reduced one according to the mean-field metric  $d_E$ . In this case, we have only three clusters and the accuracy decreases considerably, as can be numerically seen in Table 2.

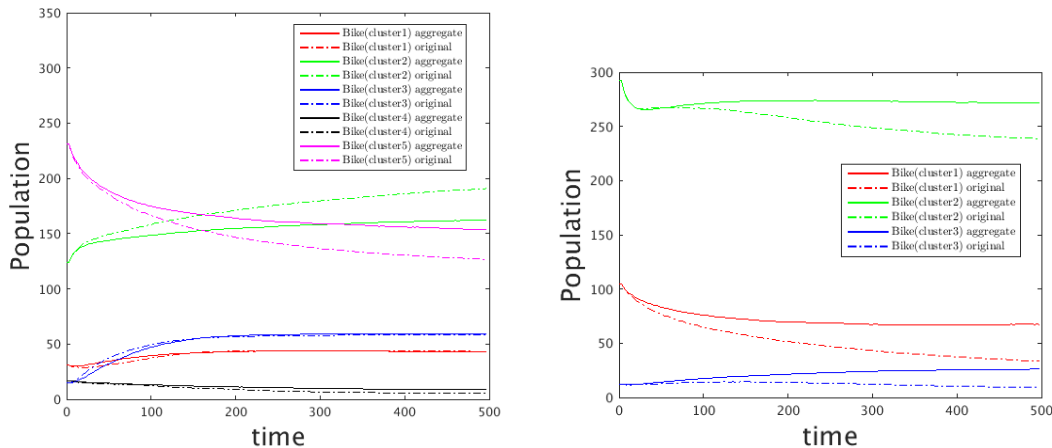


Figure 4: (left) Comparison of the available number of bikes before and after aggregation using the  $d_L$  distance. (right) Comparison of the available number of bikes before and after aggregation using the  $d_E$  distance.

In a bike-sharing scenario, we are often interested in tracking the number of bikes in some specific locations. This can be achieved in our framework by forcing some locations to be a singleton cluster. In order to understand the influence of the aggregation of remaining stations on some isolated ones, we choose, in each experiment, one station from one of the five clusters, and treat that station as a single cluster. Figure 5 shows the trajectories of the number of available bikes of the five chosen stations in five different experiments, comparing results of stochastic simulation of the original and the aggregated model. We can see the population dynamics of available bikes in the five stations still achieve good accuracy (10.84% average error ratio) even if other bike stations are aggregated.

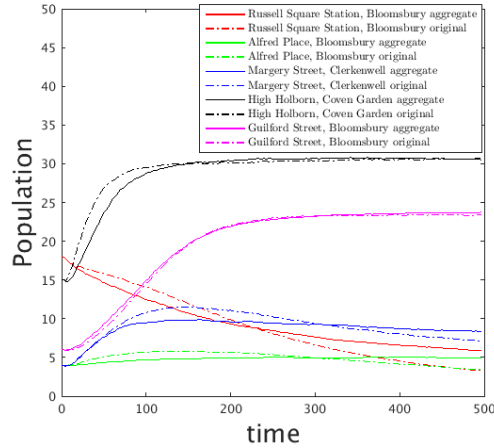


Figure 5: Comparison of the trajectories of the number of available bikes in the five chosen stations from the original and the aggregated model.

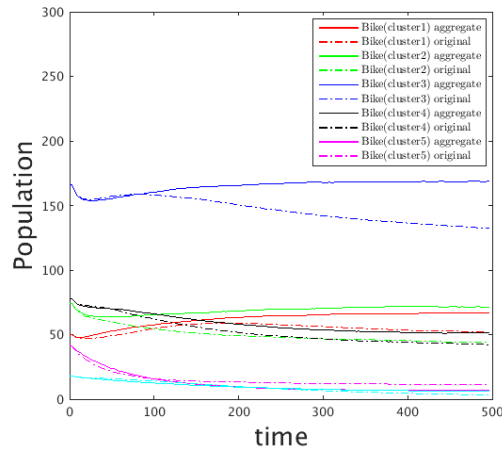


Figure 6: Comparison of the available number of bikes before and after aggregation according to physical position of stations

Additionally, we also apply a naive approach in which we use spectral clustering to cluster the 30 bike stations according to their physical positions on the map. Figure 6 shows the trajectories of the number of available bikes generated by stochastic simulation before and after aggregation in this case. The average error ratio for the trajectories is 24.51%. Thus, location clustering according to agents' population distribution outperforms clustering according to physical positions.

## 5 Related Work and Conclusion

Spatiality has been shown to be an important factor for modelling different classes of complex systems [10, 11]. As a result, there has been some effort to study the effects of spatial aggregation on stochastic models in order to achieve a balance point between model complexity and accuracy. Most of this work

focuses on how to discretise continuous space or on the effect of compartment size to represent the space, see for example [13, 12, 5]. Our work can be thought of as above these approaches, as our goal is to reduce the number of patches (locations). We show that by clustering locations in a spatial PCTMC according to linear noise or mean field distance, the size of the model can be significantly reduced but still retain reasonable accuracy. The experiment on the bike-sharing model shows that our method outperforms aggregation of locations by their physical distance. Moreover, our recent work shows that the cost of deriving higher moments based on fluid approximation can be significantly reduced using correlation heuristics [15]. This means that computing linear noise distance can also be largely reduced, which makes our method more scalable. In the future, we plan to apply this approach to models in spatial process algebras, such as PALOMA [14], CARMA [8] and MELA [27].

## 6 Acknowledgments

The authors would like to thank Jane Hillston for her helpful comments on an earlier draft of this work. This work is supported by the EU project QUANTICOL, 600708.

## References

- [1] Linda JS Allen (2010): *An introduction to stochastic processes with applications to biology*. CRC Press.
- [2] H. Andersson & T. Britton (2000): *Stochastic Epidemic Models and Their Statistical Analysis*. Springer-Verlag, doi:10.1007/978-1-4612-1158-7.
- [3] A. Andreychenko, L. Mikeev & C. Wolf (2015): *Model Reconstruction for Moment-Based Stochastic Chemical Kinetics*. *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 25(2), p. 12, doi:10.1145/2699712.
- [4] A Bhattacharyya (1943): *On a measure of divergence between two statistical population defined by their population distributions*. *Bulletin Calcutta Mathematical Society* 35, pp. 99–109.
- [5] Juan A Bonachela, Miguel A Muñoz & Simon A Levin (2012): *Patchiness and demographic noise in three ecological examples*. *Journal of Statistical Physics* 148(4), pp. 724–740, doi:10.1007/s10955-012-0506-x.
- [6] L. Bortolussi, J. Hillston, D. Latella & M. Massink (2013): *Continuous approximation of collective system behaviour: A tutorial*. *Performance Evaluation* 70(5), pp. 317–349, doi:10.1016/j.peva.2013.01.001.
- [7] L. Bortolussi & L. Nenzi (2014): *Specifying and monitoring properties of stochastic spatio-temporal systems in signal temporal logic*. In: *Proceedings of the 8th International Conference on Performance Evaluation Methodologies and Tools*, pp. 66–73, doi:10.4108/icst.Valuetools.2014.258183.
- [8] Luca Bortolussi, Rocco De Nicola, Vashti Galpin, Stephen Gilmore, Jane Hillston, Diego Latella, Michele Loreti & Mieke Massink (2015): *CARMA: Collective Adaptive Resource-sharing Markovian Agents*. In: *Proceedings Thirteenth Workshop on Quantitative Aspects of Programming Languages and Systems, QAPL 2015, London, UK.*, pp. 16–31, doi:10.4204/EPTCS.194.2.
- [9] Luca Bortolussi & Roberta Lanciani (2014): *Stochastic approximation of global reachability probabilities of Markov population models*. In: *Computer Performance Engineering*, Springer, pp. 224–239, doi:10.1007/978-3-319-10885-8\_16.
- [10] Richard Durrett & Simon Levin (1994): *The importance of being discrete (and spatial)*. *Theoretical population biology* 46(3), pp. 363–394, doi:10.1006/tpbi.1994.1032.
- [11] Rick Durrett (1999): *Stochastic spatial models*. *SIAM review* 41(4), pp. 677–718, doi:10.1137/S0036144599354707.
- [12] Radek Erban & S Jonathan Chapman (2009): *Stochastic modelling of reaction–diffusion processes: algorithms for bimolecular reactions*. *Physical biology* 6(4), p. 046001, doi:10.1088/1478-3975/6/4/046001.

- [13] Cheng Feng (2014): *Patch-based Hybrid Modelling of Spatially Distributed Systems by Using Stochastic HYPE - ZebraNet as an Example*. In: *Proceedings Twelfth International Workshop on Quantitative Aspects of Programming Languages and Systems, QAPL 2014, Grenoble, France.*, pp. 64–77, doi:10.4204/EPTCS.154.5.
- [14] Cheng Feng & Jane Hillston (2014): *PALOMA: A process algebra for located markovian agents*. In: *Quantitative Evaluation of Systems*, Springer, pp. 265–280, doi:10.1007/978-3-319-10696-0\_22.
- [15] Cheng Feng, Jane Hillston & Vashti Galpin (2016): *Automatic Moment-Closure Approximation of Spatially Distributed Collective Adaptive Systems*. *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 26(4), doi:10.1145/2883608.
- [16] Christine Fricker & Nicolas Gast (2014): *Incentives and redistribution in homogeneous bike-sharing systems with stations of finite capacity*. *EURO Journal on Transportation and Logistics*, pp. 1–31, doi:10.1007/s13676-014-0053-5.
- [17] Daniel T Gillespie (1977): *Exact stochastic simulation of coupled chemical reactions*. *The Journal of Physical Chemistry* 81(25), pp. 2340–2361, doi:10.1021/j100540a008.
- [18] Marcel C Guenther & Jeremy T Bradley (2013): *Journey data based arrival forecasting for bicycle hire schemes*. In: *Analytical and Stochastic Modeling Techniques and Applications*, Springer, pp. 214–231, doi:10.1007/978-3-642-39408-9\_16.
- [19] Marcel C Guenther, Anton Stefanek & Jeremy T Bradley (2013): *Moment closures for performance models with highly non-linear rates*. In: *Computer Performance Engineering*, Springer, pp. 32–47, doi:10.1007/978-3-642-36781-6\_3.
- [20] J. Hasenauer, V. Wolf, A. Kazeroonian & F. J. Theis (2013): *Method of conditional moments (MCM) for the Chemical Master Equation: A unified framework for the method of moments and hybrid stochastic-deterministic models*. *Journal of Mathematical Biology*, doi:10.1007/s00285-013-0711-5.
- [21] N. G. van Kampen (2007): *Stochastic processes in physics and chemistry*. Elsevier.
- [22] T. G. Kurtz (1981): *Approximation of population processes*. SIAM, doi:10.1137/1.9781611970333.
- [23] Bojan Mohar (1997): *Some applications of Laplace eigenvalues of graphs*. In: *Graph Symmetry, NATO ASI Series 497*, Springer Netherlands, pp. 225–275, doi:10.1007/978-94-015-8937-6\_6.
- [24] Andrew Y Ng, Michael I Jordan, Yair Weiss et al. (2002): *On spectral clustering: Analysis and an algorithm*. *Advances in neural information processing systems* 2, pp. 849–856, doi:10.1.1.19.8100.
- [25] Jianbo Shi & Jitendra Malik (2000): *Normalized cuts and image segmentation*. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22(8), pp. 888–905, doi:10.1109/34.868688.
- [26] Anton Stefanek, Richard A Hayden & Jeremy T Bradley (2011): *Fluid computation of the performance: energy tradeoff in large scale Markov models*. *ACM SIGMETRICS Performance Evaluation Review* 39(3), pp. 104–106, doi:10.1145/2160803.2160817.
- [27] Ludovica Luisa Vissat, Jane Hillston, Glenn Marion & Matthew J Smith (to appear): *MELA: Modelling in Ecology with Location Attributes*. *Proceedings of the fourth international symposium on Modelling and Knowledge Management applications: Systems and Domains, MoKMaSD 2015, York, UK*.
- [28] Ulrike Von Luxburg (2007): *A tutorial on spectral clustering*. *Statistics and computing* 17(4), pp. 395–416, doi:10.1007/s11222-007-9033-z.