Performance Evaluation 109 (2017) 52–76

Contents lists available at ScienceDirect

Performance Evaluation

journal homepage: www.elsevier.com/locate/peva

Spatial fluid limits for stochastic mobile networks

Max Tschaikowski, Mirco Tribastone *

IMT School for Advanced Studies Lucca, Italy

ARTICLE INFO

Article history:
Received 25 April 2014
Received in revised form 6 December 2016
Accepted 29 December 2016
Available online 16 January 2017

Keywords:
Stochastic networks
Mobility models
Fluid limits
Reaction–diffusion

ABSTRACT

We consider Markov models of large-scale networks where nodes are characterized by their local behavior and by a mobility model over a two-dimensional lattice. By assuming random walk, we prove convergence to a system of partial differential equations (PDEs) whose size depends neither on the lattice size nor on the population of nodes. This provides a macroscopic view of the model which approximates discrete stochastic movements with continuous deterministic diffusions. This result may be used in practice to justify a space coarsening induced by the numerical PDE solver. We illustrate this by modeling a network of mobile nodes with on/off behavior performing file transfers with connectivity to 802.11 access points, empirically validated against discrete-event simulation. We show high quality of the PDE approximation in the case of moderate lattice granularity and population sizes.

1. Introduction

This paper is concerned with the modeling and analysis of large-scale computer and communication networks with mobile nodes. We consider an expressive model of a node which is assumed to evolve over a set of *local states*. These may represent, for instance, distinct phases of a node's behavior that distinguish between periods of activity and inactivity in a wireless network or, with a richer granularity, they may be related to different states in a communication protocol. We are interested in understanding how the interplay between the node's local behavior and its mobility pattern has an impact on the overall system's performance.

In this paper, we consider *Markov population processes* as the starting point of our analysis. These represent a general class of continuous-time Markov chains (CTMCs), characterized by a state descriptor which is given by a vector where each element gives the population of nodes in a specific local state. In general, such stochastic models do not scale well with increasing system sizes because they are based on an explicit discrete-state representation. This makes the analysis very difficult computationally and poses a detriment to model-assisted parameter-space exploration and capacity planning. However, under mild assumptions, a Markov population process may yield a *fluid approximation* [1–7] as a system of ordinary differential equations (ODEs). This can be shown to be the deterministic limit behavior that holds when the number of nodes goes to infinity; see, e.g., [8]. Fluid limits have been developed for a wide range of models of distributed and networked systems, including, for instance, load balancing [9,10], optical switches [11], virtualized environments [12], and peer-to-peer networks [13]. In [14] Benaïm and Le Boudec offer a general framework for a Markov model of *interacting agents* evolving through a set of *local states*. In all cases, while the Markov process has a state space with cardinality that is exponential in the number of agents (in the worst case), the fluid limit only depends linearly on the number of local states. This procedure

* Corresponding author.
E-mail addresses: max.tschaikowski@imtlucca.it (M. Tschaikowski), mirco.tribastone@imtlucca.it (M. Tribastone).

http://dx.doi.org/10.1016/j.peva.2016.12.005
0166-5316/

is very attractive in practice because it enables computationally efficient solutions, when the local state space is small, that are typically very accurate for large populations.

In the aforementioned models, however, spatial effects are not explicitly present because the system is inherently static, or because it is a convenient simplification to abstract away from them for the purposes of model tractability [15]. Unfortunately, explicitly taking into account mobility and space in a fluid model generally leads to a rapid growth of the local state space. In this paper, for instance, we deal with population processes defined by a CTMC with a random walk (RW) model over a two-dimensional topology. We consider a typical partitioning of the spatial domain into *regions* (or *patches*, as usually called in theoretical biology [16]). Within each region an assumption of homogeneity is made, and spatial effects are incorporated as transitions across regions [17]. If the dynamics of each region admits a fluid limit, then the overall system behavior is given by $L(K+1)^2$ ODEs, where, according to the notation of this paper, L is the local state space size and $(K+1)^2$ is the total number of regions. Such dependence is simply due to the fact that the model must keep track of the total population of nodes in each of the L local states, *at each region*.

We consider this problem in the context of a generic framework for *mobile reaction networks*, a concise model to describe Markov population processes. This is inspired by stochastic reaction networks, which are frequently used for the analysis of systems of (bio-)chemical reactions (e.g., [18]). We assume that nodes may interact with each other within the same region, and undergo unbiased RW with a parameter, the *migration rate*, which can be dependent on their local state. We also consider *absorbing* and *reflective* boundary conditions for the spatial domain. In the presence of the former, each node that hits the boundary exits the domain without the possibility of ever re-entering it. Instead, if the latter is considered, the boundaries build a barrier of the domain that cannot be crossed by nodes. Those assumptions do not exclude the possibility of having exogenous arrivals into the system, e.g., by means of a Poisson process in each region.

The purpose of this paper is to develop a technique to effectively analyze models when K and the total population of nodes are large. It is natural to consider an approach with two dimensions of scaling: the first one is the celebrated *density-dependent* form that allows us to obtain a continuous deterministic limit for the concentrations (or densities) of the node populations [8], while keeping the regions discrete. This is achieved by constructing a family of population processes indexed by a parameter, hereafter denoted by N, such that the initial state of each element of the family is a population of nodes that grows linearly with N. Then, the result of Kurtz [8] ensures that the sequence converges, as $N \to \infty$, to the solution of an ODE system of size $L(K+1)^2$. The second scaling occurs on the spatial dimension. We define a suitable sequence of migration rates such that increasing K means considering regions which are closer and closer to each other on a regular mesh in the unit square in $\mathbb{R}^2$ (i.e., the regions are at a distance $1/K$ from each other). After setting the scene, we show that the sequence of ODE systems converges, as $K \to \infty$, to a system of L partial differential equations (PDEs) of reaction–diffusion type. In these PDEs, the diffusive terms model the continuous migration across regions, whereas the reactive terms describe the local interactions between nodes. To the best of our knowledge, such a convergence result has not been proven before. We also argue that our limit may allow for a more efficient numerical solution. This may seem surprising because analytical solutions of PDE systems are scarce and numerical solutions of PDE systems rely on a discretization of space. However while in the case of the limit ODE system of size $L(K+1)^2$ the discretization is dictated by the stochastic model, in our limit PDE system the coarseness of the discrete mesh depends on the PDE solver. In other words, the PDE solver may in effect give rise to a coarsening of the original spatial domain.

A typical situation of practical interest to which our framework can be applied is that of large-scale mobile networks such as personal communication services (e.g., [19]): there are many base stations (e.g., in a wide-area cellular network) and the area served by a base station can be modeled as a region, which can contain potentially many mobile nodes that may migrate across regions. Here, the modeler may wish to predict how nodes distribute across the network over a given time horizon [20]. In this paper, we demonstrate the applicability of our method by studying a network of mobile nodes with connectivity provided by an 802.11 access point located in each region. First, we successfully validate our CTMC model against discrete-event simulation using the JiST/SWANS discrete-event simulation framework [21]. Then, we show that the PDE solution provides an excellent estimate of the network's performance for relatively small population sizes and moderate lattice granularity.

Paper outline. The remainder of the paper is organized as follows. Section 2 overviews related work. In order to fix notation and provide intuition on the scaling limits considered in the paper, Section 3 introduces *stationary reaction networks*, which describe *static* networks with no mobility, whilst allowing nodes to be described by a local state space. Here we discuss how such models admit a classic fluid limit as a system of ODEs. Section 4 presents mobile reaction networks, which are defined as a conservative extension of stationary networks with an explicit mobility model. We introduce a straightforward spatial ODE limit that depends upon the lattice granularity. In Section 5 we discuss the main contribution of this paper, namely the convergence of the spatial ODE limit to a system of reaction–diffusion PDEs by assuming that nodes undergo unbiased random walk during their evolution. Section 6 discusses the numerical tests on our validation model. Finally, Section 7 concludes the paper.

2. Related work

Reaction–diffusion PDEs. PDEs of reaction–diffusion type are very well understood in many disciplines, such as biology [22], ecology [23], and chemistry [24]. It is beyond the scope of this paper to provide a general overview of the literature. Instead,