Mean field for performance models with deterministically-timed transitions

Richard A. Hayden
Department of Computing, Imperial College London, United Kingdom
Email: rh@doc.ic.ac.uk

***Abstract*—We extend the mean-field (a.k.a. fluid-analysis) approach for massively-parallel continuous-time Markov chains (CTMCs) to models with both Markovian and deterministically-timed transitions. We introduce a new low-level formalism for specifying massively-parallel models with generally-timed transitions, the *population generalised semi-Markov process (PGSMP)*. We then show how systems of coupled delay differential equations (DDEs) which approximate transient component counts may be mechanistically derived from such models. This is possible not only in the case of non-competing deterministic delays, but also where the deterministic transitions may race locally with simultaneously enabled Markovian ones. For a large class of PGSMP models, we are able to prove mean-field convergence formally and to construct a second-order limit process.**

In the steady state, we show that it is not always straightforward to apply the fixed point approach to approximate the stationary behaviour of a PGSMP. To address this, we provide a new algorithm for computing steady-state mean-field approximations for PGSMP models leveraging the theory of Markov regenerative processes.

Throughout, our techniques are illustrated on a number of simple worked examples drawn from the areas of peer-to-peer networks and wireless sensor networks.

I. Introduction

In this paper, we develop a new mean-field methodology for a class of massively-parallel *generalised semi-Markov processes (GSMPs)* [1, 2, 3]. Specifically, we focus on population models where individuals can enable both Markovian and deterministically-timed transitions, which are crucial for the accurate modelling of many real-world computer and networking protocols. We encode such models in a new low-level formalism, the *population generalised semi-Markov process (PGSMP)*.

The motivation for the mean-field approach is the same as in the continuous-time Markov chain (CTMC) case — unsurprisingly, GSMP models with many components also become computationally intractable to explicit state techniques [e.g. 4, 5, 6, 7] rapidly as a result of the familiar state-space explosion problem. Our approach is based on the derivation of *delay differential equations (DDEs)*[1] [e.g. 8] from PGSMP models and generalises the traditional mean-field approach as applied to CTMC models based on ordinary differential equations (ODEs) [e.g. 9, 10, 11, 12, 13, 14, 15].

The class of models to which our approach applies is very broad — the only significant restriction we make is that at most one deterministic transition may be enabled by each individual in any given local state. However, globally, there is no restriction on the concurrent enabling of deterministic transitions by different individuals.

As in the CTMC case, the size of the system of DDEs is equal to the number of physical local states that components in the model can be in. Therefore this approach represents a significant improvement with respect to both accuracy and efficiency when compared with the traditional CTMC mean-field approach where deterministic transitions are approximated using phase-type distributions. Indeed, by the central limit theorem, a 100-phase Erlang approximation to a deterministic distribution of duration 1 still has a probability of about 32% of lying outside of $[0.9, 1.1]$. Furthermore, incorporating this approximate distribution into the behaviour of an individual in a CTMC mean-field model will increase the number of mean-field ODEs required by 100, or more, if there is allowed to be other Markovian behaviour in competition with it. On the other hand, the mean-field approach based on DDEs presented here captures deterministic distributions directly without the need for additional physical states or for approximation of the deterministic distribution itself.

In Section II-A, we introduce the PGSMP formalism. In Section III, we show how systems of coupled DDEs can be derived mechanistically from PGSMP models with deterministic transitions, under the initial assumption that they do not race locally with exponential transitions (we term this class *delay-only* PGSMPs). For this class of model, we are able to prove transient mean-field convergence and give a second-order limit result. Then in Section IV, we extend the derivation of mean-field differential equations to general PGSMP models where exponential transitions may be enabled concurrently with deterministic ones. In both cases, we illustrate the methodology on a peer-to-peer software update model. Finally, in Section V, we show that the traditional ODE fixed point approach to approximating steady-state component counts in the mean-field regime is not always as straightforward in the case of DDEs. To resolve this, we provide an algorithm for computing the long-time mean-field limits by observing that individuals in a PGSMP model behave as Markov regenerative processes under mean-field and stationarity assumptions.

A. Related work

The closest related work can be found in the biology and chemistry literature. Systems of DDEs have been derived to

[1]Functional differential equations where the right-hand side can depend on the history of the unknown function, not just its value at the current time.

approximate stochastic models of reaction networks where deterministic delays are possible after reactions occur [e.g. 16, 17, 18, 19, 20]. However, these models differ from those considered here in a number of critical ways.

Firstly, they are restricted to models with a very rigid structure in that deterministic transitions may not be enabled concurrently with Markovian transitions. This suits the kinds of models encountered in biology and chemistry but is restrictive from a performance modelling point of view — for example, deterministic timeouts cannot be captured. The techniques presented in this paper allow the derivation of DDEs from a much more general class of models. Indeed, in this paper, the Markovian behaviour that is allowed to occur concurrently with an enabled deterministic transition is unconstrained as long as each individual component enables at most one deterministic transition in any given local state.

Secondly, the semantics of the models from which DDEs are derived in the literature often appears to be specific to chemical models. Here, it appears that the state of a component is not necessarily updated after every Markovian delay, but only after the proceeding deterministic delay [e.g. 16, 19]. Also, this appears to be the case in the so-called *purely delayed* approach of [18]. That is, the system only records that a Markovian transition has fired after the successive deterministic delay has expired — the reaction together with its deterministic delay is considered as one atomic action. This means that transition rates cannot be defined with the full knowledge of the state of each individual and, for our purposes, would significantly restrict the kind of performance models we could consider.

Furthermore, to the best of our knowledge, the mean-field convergence result of Theorem 1, the second-order limit process in Section III-C and the steady-state approximation algorithm of Section V are also new and have not appeared in the literature before, even in the less general context of chemical reaction networks.

Finally, we note that this paper builds upon a very preliminary extended abstract that has already appeared [21].

II. Population generalised semi-Markov processes

In this section, we introduce a new low-level formalism which captures the class of models with which this paper is concerned. Models specified in this formalism are termed *population generalised semi-Markov processes (PGSMPs)*. The descriptor *population* refers to the fact we are interested in models consisting of large collections of interacting, but stochastically-identical components, the setting in which mean-field approaches are naturally applied. Where our formalism differs from existing low-level approaches to specifying Markovian population models [e.g. 14, 22] is that we allow some of the transitions between the states of the constituent individuals to be timed according to distributions other than exponential. Indeed, the underlying stochastic process will be a *generalised semi-Markov process (GSMP)* [1, 2, 3].

In the chemistry and biology literature, there do exist stochiometric formalisms which allow the specification of reaction systems where deterministic delays between reaction initiation and completion are possible [e.g. 16, 18, 19, 20, 23]. However, as discussed in Section I-A, these representations are designed solely with molecular reaction systems in mind and thus impose significant structural limitations on how the deterministic delays can be used and how they may interact with exponential delays enabled concurrently by a given individual. For example, it does not seem possible using these formalisms to specify systems consisting of individuals where a deterministic timeout can cancel a simultaneously enabled exponential transition and vice-versa.

A. Definition of PGSMPs

A PGSMP model consists of many interacting components each belonging to one of a finite set of *component classes* $\mathcal{K}$. Each component in a given class $k \in \mathcal{K}$ can inhabit one of a finite set of *local states* $\mathcal{S}_k$. The global state space, say $\mathcal{X}$, of a PGSMP model then consists of elements $\mathbf{x} = (\mathbf{x}_k)_{k\in\mathcal{K}}$ where each $\mathbf{x}_k = (x_{k,s})_{s\in\mathcal{S}_k}$ and $x_{k,s} \in \mathbb{Z}_+$ tracks the number of class-k components currently in the local state s.

Exponential transitions are specified by a finite set of *Markovian transitions* $\mathcal{C}$. Each $c \in \mathcal{C}$ specifies a finite *change multiset* L_c, which consists of three-tuples (k, s, s') such that $k \in \mathcal{K}$ and $s, s' \in \mathcal{S}_k$ each of which specifies that as part of a c-transition, a distinct, uniformly randomly selected class-k component in local state s moves to local state s'. Write also $\mathbf{l}^c = (\mathbf{l}^c_k)_{k\in\mathcal{K}}$ where each $\mathbf{l}^c_k = (l^c_{k,s})_{s\in\mathcal{S}_k}$ and $l^c_{k,s} := |\{|(k, s', s) \in L_c|\}| - |\{|(k, s, s') \in L_c|\}|$, which represents the total change in class-k components that are in local state s when a c-transition occurs. The aggregate rate of c-transitions is given by a *rate function* $r_c : \mathcal{X} \to \mathbb{R}_+$. We assume that the rate function is defined such that it is zero whenever a transition is not possible due to there not being enough distinct components in the required local states.

Deterministically-timed transitions are specified by *event clocks* in a similar fashion to standard GSMPs [2]. Specifically, each class of components $k \in \mathcal{K}$ has associated a finite set of event clocks $\mathcal{E}_k$. Each event clock $e \in \mathcal{E}_k$ is specified by its *set of active states* $\mathcal{A}_e \subseteq \mathcal{S}_k$, its *event transition probability function* $p_e : \mathcal{S}_k \times \mathcal{S}_k \to [0, 1]$ and the *clock duration* $d_e \in \mathbb{R}_{>0}$. When a component enters a state in $\mathcal{A}_e$ for the first time, the clock is initialised to the deterministic value d_e. If the clock time elapses before the component leaves the set of active states, it moves immediately to a new local state by sampling from the discrete probability distribution $p_e(s, \cdot)$, where s is the local state of the component when the clock reaches zero.[2] Otherwise, the clock is disabled when the component leaves the set of active states and is reset if it later returns to the set of active states.

As mentioned above, the key restriction we make for all PGSMP models considered in this paper is that at most one event clock may be active in any local state. That is, for each $k \in \mathcal{K}$ and $s \in \mathcal{S}_k$, $|\{e \in \mathcal{E}_k : s \in \mathcal{A}_e\}| \leq 1$. We will see that this restriction is necessary for the mean-field analyses presented in the sequel. This restriction also means

[2] If this transition is to another state in $\mathcal{A}_e$, the clock is then reset to d_e.

that, with probability one, it is not possible for two transitions (Markovian or generally-timed) to occur simultaneously within a single component. Finally, we write $\mathbf{x}^0$ for the initial state of the model.

III. Mean-field for delay-only PGSMPs with deterministic clocks

We will focus initially on a class of PGSMPs with the further structural restriction that, within a given component, deterministic transitions may not be enabled concurrently with Markovian ones. For this class, we are also able to prove transient mean-field convergence formally and construct a second-order limit process. We refer to such models as *delay only* since the deterministic transitions in the constituent components then serve only to introduce deterministic delays between periods of Markovian behaviour.

Formally, the class of delay-only PGSMPs is specified by two restrictions: for all $k \in \mathcal{K}$ and $e \in \mathcal{E}_k$; if $s \in \mathcal{A}_e$ then there can be no $c \in \mathcal{C}$ with $(k, s, s') \in L_c$; and for each $s \in \mathcal{A}_e$, there must exist some $s' \in \mathcal{S}_k$ such that $s' \notin \mathcal{A}_{e'}$ for any $e' \in \mathcal{E}_k$ and $p_e(s, s') = 1$. The first restriction guarantees that no Markovian transitions are enabled concurrently with deterministic transitions, as above. The second restriction guarantees, firstly, that after any deterministic transition completes, the component jumps into a unique state.[3] Secondly, it also ensures that the completion of a deterministic transition cannot immediately enable another.

For each local state $s \in \mathcal{S}_k$ of some component class $k \in \mathcal{K}$, we write $v_{k,s}(t)$ for the mean-field approximation to the number of class-k components in state s at time $t \in \mathbb{R}_+$. We also let $\mathbf{v}(t) = (\mathbf{v}_k(t))_{k\in\mathcal{K}}$ where each $\mathbf{v}_k(t) = (v_{k,s}(t))_{s\in\mathcal{S}_k}$. The corresponding system of DDEs then takes the following form:

$$\dot{v}_{k,s}(t) = \sum_{c\in\mathcal{C}} l^c_{k,s} r_c(\mathbf{v}(t)) + \sum_{e\in\mathcal{E}_k}\sum_{c\in\mathcal{C}} \mathbf{1}_{\{t\geq d_e\}} \left(\sum_{s'\in\mathcal{A}_e} \underbrace{p_e(s',s) l^c_{k,s'} r_c(\mathbf{v}(t-d_e))}_{\substack{\text{Rate of } e \text{ clocks starting at } t-d_e \\ \text{due to expo. transits. into } s' \\ \text{that end in } s}} - \underbrace{\mathbf{1}_{\{s\in\mathcal{A}_e\}} l^c_{k,s} r_c(\mathbf{v}(t-d_e))}_{\substack{\text{Rate of } e \text{ clocks starting at } t-d_e \\ \text{due to expo. transits. into } s}} \right) \tag{1}$$

The first term represents the standard contribution of the Markovian transitions to the mean-field equation. The second term captures transitions due to deterministic delays which were initiated by a preceding exponential transition.

Initial conditions are given by $\mathbf{v}(0) := \mathbf{x}^0$.[4] Note that Eq. (1) only accounts for deterministic transitions which are initiated due to a transition into a state by a Markovian transition so it is incorrect if the initial conditions are such that any component begins in a local state that enables a deterministic transition. This is easily resolved by adding the contribution of such transitions separately and instantaneously as jumps in the mean-field solution at the point at which they occur. The additional jump terms $v^0_{k,s}(t)$ have a form very similar to the second term of Eq. (1):

$$v^0_{k,s}(t) := \sum_{e\in\mathcal{E}_k} \mathbf{1}_{\{t\geq d_e\}} \left(\underbrace{\sum_{s'\in\mathcal{A}_e} p_e(s',s) x^0_{k,s'}}_{\substack{e \text{ clocks starting} \\ \text{at 0 in } s' \text{ that end in } s}} - \underbrace{\mathbf{1}_{\{s\in\mathcal{A}_e\}} x^0_{k,s}}_{\substack{e \text{ clocks} \\ \text{starting at 0 in } s}} \right)$$

In order to introduce the jump in the mean-field equation we express the DDE of Eq. (1) as an integral equation and the jump term may then simply be added to it:[5]

$$\begin{aligned} v_{k,s}(t) = x^0_{k,s} &+ \sum_{c\in\mathcal{C}} l^c_{k,s} \int_0^t r_c(\mathbf{v}(u))\,\mathrm{d}u \\ &+ \sum_{e\in\mathcal{E}_k}\sum_{c\in\mathcal{C}} \left(\sum_{s'\in\mathcal{A}_e} p_e(s',s) l^c_{k,s'} - \mathbf{1}_{\{s\in\mathcal{A}_e\}} l^c_{k,s} \right) \\ &\times \int_0^{(t-d_e)\vee 0} r_c(\mathbf{v}(u))\,\mathrm{d}u + v^0_{k,s}(t) \end{aligned} \tag{2}$$

Due to the presence of the indicator functions in the right-hand sides of the DDEs, we do not require that solutions are differentiable at time points $t = d_e$ for $e \in \mathcal{E}_k$, $k \in \mathcal{K}$. Further, if there are components starting in local states enabling deterministic transitions, then the solutions to Eq. (2) may not even be continuous at these time points due to the jump terms.

A. Worked example: delay-only peer-to-peer software update model

In this section, we derive the system of DDEs as defined in the previous section for a simple example model of a peer-to-peer software update process. We consider two general types of nodes in this model which we term *old* and *updated*. Old nodes are those running an old software version and new nodes are those which have been updated to a new version. Both types of nodes alternate between being *on* and *off*. When an updated node is on, an old node may locate it and subsequently update itself in a peer-to-peer fashion. Whenever an old node comes on, it polls the network for new nodes (so it can be updated) before giving up if it does not find one after a timeout.

In this model, we require just one component class which captures both old and updated nodes, so we omit the k parameter on the notation of the previous sections. Nodes have five possible local states: updated nodes can be on and off, which we denote by `a` and `b`, respectively. Old nodes can be on (`c`), off (`e`) or in a state representing an old node which is on but has given up seeking updates (`d`). In the notation of Section II-A, the set of local states is thus $\mathcal{S} := \{\mathtt{a}, \mathtt{b}, \mathtt{c}, \mathtt{d}, \mathtt{e}\}$.

Our eventual intention is to allow the timeout transition to be deterministic reflecting the situation in many real-world protocols. However, this would not result in a *delay-only*

[3] This is a technical but not, in fact, a modelling restriction. Probabilistic behaviour immediately after a deterministic transition completes as encoded by p_e, can, for this class of models, be moved to before the deterministic transition by a suitable reconfiguration of the state space and Markovian transition rates.

[4] Note that the indicator function appearing in Eq. (1) means that initial data for $t < 0$ does not need to be provided.

[5] We use the notation $a \vee b := \max(a, b)$.

Transition, $c \in \mathcal{C}$	Change multiset, L_c	Rate function, r_c
1	$\{\!\{(\mathtt{c},\mathtt{e})\}\!\}$	$\rho x_{\mathtt{c}}$
2	$\{\!\{(\mathtt{d},\mathtt{e})\}\!\}$	$\rho x_{\mathtt{d}}$
3	$\{\!\{(\mathtt{c},\mathtt{a})\}\!\}$	$\frac{\beta x_{\mathtt{c}} x_{\mathtt{a}}}{N}$
4	$\{\!\{(\mathtt{c},\mathtt{d})\}\!\}$	$\kappa x_{\mathtt{c}}$
5	$\{\!\{(\mathtt{a},\mathtt{b})\}\!\}$	$\rho x_{\mathtt{a}}$

TABLE I: *Markovian transitions for the delay-only software update model.*

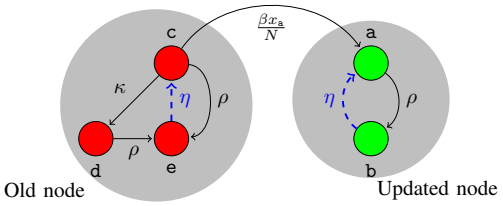

Fig. 1: *Representation of the behaviour of a single node in the delay-only software update model.*

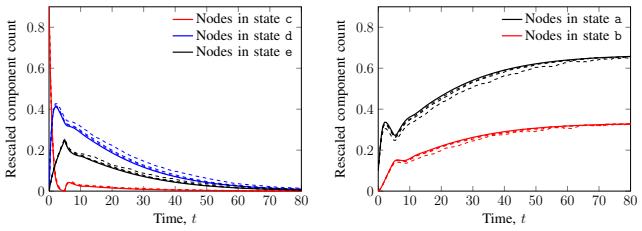

(a) *Old nodes* **(b)** *Updated nodes*

Fig. 2: *Delay-only software update model rescaled DDE approximation (solid lines) compared with rescaled actual means for $N = 20$, 50 and 100 (dashed lines). Initial component proportions are $(0.1, 0, 0.9, 0, 0)$ and parameters are $\eta = 5.0$, $\beta = 2.0$, $\rho = 0.1$ and $\kappa = 0.67$.*

PGSMP since this transition is enabled concurrently with other transitions. We therefore defer the presentation of the version of the model with deterministic timeouts to Section IV-A after we have introduced the general mean-field analysis. However, to give an example of a delay-only PGSMP, we consider in this section a version of the model where all transitions are Markovian except for the transitions bringing nodes from their off state into their on state, which is deterministic with duration η. Formally, there are two event clocks $\mathtt{t}_{\circ}$ and $\mathtt{t}_{\mathtt{u}}$ with $\mathcal{A}_{\mathtt{t}_{\circ}} := \{\mathtt{e}\}$, $d_{\mathtt{t}_{\circ}} := \eta$, $p_{\mathtt{t}_{\circ}}(\mathtt{e},\mathtt{c}) = 1$, $\mathcal{A}_{\mathtt{t}_{\mathtt{u}}} := \{\mathtt{b}\}$. $d_{\mathtt{t}_{\mathtt{u}}} := \eta$ and $p_{\mathtt{t}_{\mathtt{u}}}(\mathtt{b},\mathtt{a}) := 1$.

The Markovian transitions are given in Table I where N is the total component population. We follow the usual methodology for the stochastic modelling of epidemics in that the chance of an old node finding an updated one in a small period of time is proportional to the number of available updated nodes. The local behaviour of a node is also depicted graphically in Figure 1.

The DDEs corresponding to this model are then:

$$
\begin{aligned}
\dot{v}_{\mathtt{a}}(t) = & -\rho v_{\mathtt{a}}(t) + \frac{\beta v_{\mathtt{c}}(t) v_{\mathtt{a}}(t)}{N} + \mathbf{1}_{\{t\geq\eta\}}\rho v_{\mathtt{a}}(t-\eta) \\
\dot{v}_{\mathtt{b}}(t) = & -\mathbf{1}_{\{t\geq\eta\}}\rho v_{\mathtt{a}}(t-\eta) + \rho v_{\mathtt{a}}(t) \\
\dot{v}_{\mathtt{c}}(t) = & -\rho v_{\mathtt{c}}(t) - \frac{\beta v_{\mathtt{c}}(t) v_{\mathtt{a}}(t)}{N} - \kappa v_{\mathtt{c}}(t) \\
& + \mathbf{1}_{\{t\geq\eta\}}(\rho v_{\mathtt{d}}(t-\eta) + \rho v_{\mathtt{c}}(t-\eta)) \\
\dot{v}_{\mathtt{d}}(t) = & -\rho v_{\mathtt{d}}(t) + \kappa v_{\mathtt{c}}(t) \\
\dot{v}_{\mathtt{e}}(t) = & -\mathbf{1}_{\{t\geq\eta\}}(\rho v_{\mathtt{d}}(t-\eta) + \rho v_{\mathtt{c}}(t-\eta)) + \rho v_{\mathtt{d}}(t) + \rho v_{\mathtt{c}}(t)
\end{aligned}
$$

This system of DDEs can be integrated numerically by adapting existing ODE solvers or specialised DDE routines such as the `dde23` routine in MATLAB®. The solution of these DDEs (rescaled by dividing by N) for one set of parameters is shown in Figure 2 compared with the corresponding rescaled component-count expectations as computed by many stochastic simulation replications. We observe that the means do appear to converge to the mean-field solutions in line with Theorem 1 to be presented in the next section.

B. Mean-field convergence for delay-only PGSMPs

In this section of the paper, we prove transient mean-field convergence for delay-only PGSMPs. We begin by constructing a sequence of delay-only PGSMP models indexed by $N \in \mathbb{Z}_+$ with increasing total component population size. Their underlying stochastic processes are denoted $\{\mathbf{x}^N(t) \in \mathbb{R}_+^n\}_{N\in\mathbb{Z}_+}$, where $n := |\mathcal{S}|$ for $\mathcal{S} := \{(k,s) : k \in \mathcal{K}, s \in \mathcal{S}_k\}$; and $\mathbf{x}^N(t) = (\mathbf{x}_k^N(t))_{k\in\mathcal{K}}$ where each $\mathbf{x}_k^N(t) = (x_{k,s}^N(t))_{s\in\mathcal{S}_k}$ and $x_{k,s}^N(t) \in \mathbb{Z}_+$ tracks the number of class-k components currently in the local state s for the Nth model.

We assume that the set of component classes $\mathcal{K}$; the sets of local states $\mathcal{S}_k$; the set of transitions $\mathcal{C}$ and the change multisets L_c; the sets of event clocks $\mathcal{E}_k$, the sets of active states $\mathcal{A}_e$, the transition probability functions p_e and the deterministic delays d_e are all fixed for all elements of the sequence. The rate functions r_c^N are allowed to vary with N and the initial conditions for the Nth model in the sequence are given by $N\mathbf{x}^0$ for some $\mathbf{x}^0 \in \mathbb{Z}_+^n$. Write $\mathcal{X}^N \subseteq \mathbb{Z}_+^n$ for the reachable state space of the Nth model.

Similarly to the case of *density-dependent Markov chains* [24, 25], we assume that we may define $r_c(\mathbf{x}) := (1/N)r_c^N(N\mathbf{x})$ for all $\mathbf{x} \in \mathbb{R}^n$ independently of N. Furthermore, we assume that r_c satisfies a local Lipschitz condition on $\mathbb{R}^n$ and that for all $c \in \mathcal{C}$, $r_c^N(\mathbf{x}) \leq R(\|\mathbf{x}\| + 1)$ for all $\mathbf{x} \in \mathcal{X}^N$ where $R \in \mathbb{R}_+$ is independent of N. Define the rescaled processes $\bar{\mathbf{x}}^N(t) := (1/N)\mathbf{x}^N(t)$, then the mean-field approximation $\mathbf{v}(t)$ to $\bar{x}_{k,s}^N(t)$, derived following Section III, satisfies the integral equation of Eq. (2). We assume that a unique solution to this system of coupled equations exists globally, in which case it is clear that there is some compact set $S \subset \mathbb{R}^n$ that contains both the mean-field solution and all of the rescaled processes $\bar{\mathbf{x}}^N(t)$ for all $t \in \mathbb{R}_+$.

The useful property of delay-only PGSMPs over general PGSMPs is that their underlying stochastic process admits a representation in terms of random time changes of Poisson processes [23, 24]. In particular, let $\{P_c(t) : c \in \mathcal{C}\}$ be a set of mutually-independent rate-1 Poisson processes, then we

may write:

$$\bar{x}^N_{k,s}(t) = x^0_{k,s} + \frac{1}{N}\sum_{c\in\mathcal{C}} l^c_{k,s} P_c\left(\int_0^t r^N_c(\mathbf{x}^N(u))\,\mathrm{d}u\right) + \frac{1}{N}\sum_{e\in\mathcal{E}_k}\sum_{c\in\mathcal{C}}\left(\sum_{s'\in\mathcal{A}_e} p_e(s',s) l^c_{k,s'} - \mathbf{1}_{\{s\in\mathcal{A}_e\}} l^c_{k,s}\right) \times P_c\left(\int_0^{(t-d_e)\vee 0} r^N_c(\mathbf{x}^N(u))\,\mathrm{d}u\right) + v^0_{k,s}(t) \quad (3)$$

The mean-field convergence theorem now follows.

Theorem 1. *Under the assumptions and setup given above, we have, for any $T > 0$ and $\epsilon > 0$:*

$$\lim_{N\to\infty} \mathbb{P}\left\{\sup_{t\in[0,T]} \|\bar{\mathbf{x}}^N(t) - \mathbf{v}(t)\| > \epsilon\right\} = 0$$

Proof: On the compact set S, the functions r_c are all Lipschitz continuous; let K be a Lipschitz constant for all of these functions with respect to the maximum norm $\|\mathbf{x}\|_\infty := \max\{|x_i|\}$ on $\mathbb{R}^n$. Now define:

$$D^N_{k,s}(T) := \sup_{t\in[0,T]}\Bigg|\bar{x}^N_{k,s}(t) - x^0_{k,s} - \sum_{c\in\mathcal{C}} l^c_{k,s}\int_0^t r_c(\bar{\mathbf{x}}^N(u))\,\mathrm{d}u - \sum_{e\in\mathcal{E}_k}\sum_{c\in\mathcal{C}}\left(\sum_{s'\in\mathcal{A}_e} p_e(s',s) l^c_{k,s'} - \mathbf{1}_{\{s\in\mathcal{A}_e\}} l^c_{k,s}\right) \times \int_0^{(t-d_e)\vee 0} r_c(\bar{\mathbf{x}}^N(u))\,\mathrm{d}u - v^0_{k,s}(t)\Bigg|$$

Then we have, for $t \in [0,T]$:

$$\begin{aligned}|\bar{x}^N_{k,s}(t) - v_{k,s}(t)| &\le D^N_{k,s}(T) + \sum_{c\in\mathcal{C}} |l^c_{k,s}| \int_0^t |r_c(\bar{\mathbf{x}}^N(u)) - r_c(\mathbf{v}(u))|\,\mathrm{d}u \\ &\quad + \sum_{e\in\mathcal{E}_k}\sum_{c\in\mathcal{C}}\left|\sum_{s'\in\mathcal{A}_e} p_e(s',s) l^c_{k,s'} - \mathbf{1}_{\{s\in\mathcal{A}_e\}} l^c_{k,s}\right| \times \int_0^{(t-d_e)\vee 0} |r_c(\bar{\mathbf{x}}^N(u)) - r_c(\mathbf{v}(u))|\,\mathrm{d}u \\ &\le D^N_{k,s}(T) + ZK\int_0^t \epsilon^N(u)\,\mathrm{d}u\end{aligned}$$

where $\epsilon^N(t) := \|\bar{\mathbf{x}}^N(t) - \mathbf{v}(t)\|_\infty$ and:

$$Z := \sum_{c\in\mathcal{C}} |l^c_{k,s}| + \sum_{e\in\mathcal{E}_k}\sum_{c\in\mathcal{C}}\left|\sum_{s'\in\mathcal{A}_e} p_e(s',s) l^c_{k,s'} - \mathbf{1}_{\{s\in\mathcal{A}_e\}} l^c_{k,s}\right|$$

independent of t and N. Then we have:

$$\epsilon^N(t) \le \sum_{(k,s)\in\mathcal{S}} |\bar{x}^N_{k,s}(t) - v_{k,s}(t)| \le \sum_{(k,s)\in\mathcal{S}} D^N_{k,s}(T) + nZK\int_0^t \epsilon^N(u)\,\mathrm{d}u$$

An application of Grönwall's lemma [e.g. 24, Page 498] then yields $\epsilon^N(t) \le \sum_{(k,s)\in\mathcal{S}} D^N_{k,s}(T)\exp(nZKT)$. Now note that:

$$D^N_{k,s}(T) \le \sup_{t\in[0,T]}\sum_{c\in\mathcal{C}}\frac{Y_c}{N}\left|P_c\left(\int_0^t r^N_c(\mathbf{x}^N(u))\,\mathrm{d}u\right) - \int_0^t r^N_c(\mathbf{x}^N(u))\,\mathrm{d}u\right|$$

where:

$$Y_c := |l^c_{k,s}| + \sum_{e\in\mathcal{E}_k}\left|\sum_{s'\in\mathcal{A}_e} p_e(s',s) l^c_{k,s'} - \mathbf{1}_{\{s\in\mathcal{A}_e\}} l^c_{k,s}\right|$$

Therefore $D^N_{k,s}(T)$ can be bounded above by:

$$\sum_{c\in\mathcal{C}} Y_c \sup_{t\in[0,T]}\left|\frac{1}{N}P_c(NR(S+1)t) - R(S+1)t\right|$$

where $S := \sum_{(k,s)\in\mathcal{S}} x^0_{k,s}$. The result then follows by the strong law of large numbers for the Poisson process, which is equivalent to the functional strong law of large numbers [e.g. 26, Section 3.2], that is, for all $S \in \mathbb{R}_+$, $\sup_{s\in[0,S]} \|P_c(Ns)/N - s\| \to 0$ as $N \to \infty$ with probability 1. ∎

C. Second-order approximations for delay-only PGSMPs

From [24, Corollary 5.5 and Remark 5.4], we see that it is possible to construct, on the same probability space as the Poisson processes $\{P_c(t) : c \in \mathcal{C}\}$, mutually-independent standard Brownian motions $\{B_c(t) : c \in \mathcal{C}\}$, such that:

$$Z_c := \sup_{t\in\mathbb{R}_+}\frac{|P_c(t) - t - B_c(t)|}{\log(2\vee t)} < \infty \quad \text{almost surely}$$

We note that an intuitive generalisation of strong approximation results for density-dependent CTMCs [25], would then suggest the second-order approximation $x^N_{k,s}(t) \approx Nv_{k,s}(t) + \sqrt{N}e_{k,s}(t)$ where the $e_{k,s}(t)$ satisfy the following set of coupled stochastic delay differential equations:

$$\begin{aligned}e_{k,s}(t) = &\sum_{c\in\mathcal{C}} l^c_{k,s}\left[\int_0^t Dr_c(\mathbf{v}(u))\cdot\mathbf{e}(u)^T\,\mathrm{d}u + B_c\left(\int_0^t r_c(\mathbf{v}(u))\,\mathrm{d}u\right)\right] \\ &+ \sum_{e\in\mathcal{E}_k}\sum_{c\in\mathcal{C}}\left(\sum_{s'\in\mathcal{A}_e} p_e(s',s) l^c_{k,s'} - \mathbf{1}_{\{s\in\mathcal{A}_e\}} l^c_{k,s}\right) \\ &\times\left[\int_0^{(t-d_e)\vee 0} Dr_c(\mathbf{v}(u))\cdot\mathbf{e}(u)^T\,\mathrm{d}u + B_c\left(\int_0^{(t-d_e)\vee 0} r_c(\mathbf{v}(u))\,\mathrm{d}u\right)\right]\end{aligned}$$

where Dr_c is the Jacobian matrix of first-order partial derivatives of r_c. The process $\mathbf{e}(t)$ is an Ornstein–Uhlenbeck process with time delay. We save the formal proof of second-order convergence to this limit process and the investigation of the approximation to a future paper.

IV. Mean-field for general PGSMPs with deterministic clocks

In this section, we extend the derivation of mean-field equations to a more general class of PGSMP with deterministic clocks. Specifically, we now allow deterministic transitions to

race locally with concurrently-enabled Markovian transitions. The only restriction we impose is that we disallow components with cycles of uninterrupted deterministic behaviour in their state space. Formally, we do not allow components whose structure would admit cycles in their local state space (by either Markovian or deterministic transitions) $s_1 \to s_2 \to \ldots \to s_n$ for $s_1, s_2, \ldots, s_n \in \mathcal{S}_k$ such that: a deterministic transition is always enabled in each state ($s_i \in \mathcal{A}_e$ for some $e \in \mathcal{E}_k$); at least one of the transitions in the path is deterministic; the component only leaves an active set along the path by a deterministic transition ($s_i \in \mathcal{A}_e$ and $s_{i+1} \notin \mathcal{A}_e$ implies that the transition $s_i \to s_{i+1}$ in the path is deterministic); and $s_1, s_n \in \mathcal{A}_e$ for some $e \in \mathcal{E}_k$. We call models with no cycles of this kind *cycle-free* PGSMPs and it will become clear when we define the general mean-field equations why this restriction is necessary. Note that a delay-only PGSMP is automatically cycle free by virtue of the assumptions made at the start of Section III.

In order to define the system of DDEs in this more general context, we require, for each $k \in \mathcal{K}$ and $e \in \mathcal{E}_k$, the $|\mathcal{S}_k| \times |\mathcal{S}_k|$-matrix-valued function $\mathbf{Q}_{k,e}(\mathbf{v}(t))$ defined for $t \in \mathbb{R}_+$ and $s, s' \in \mathcal{S}_k$ by:

$$[Q_{k,e}(\mathbf{v}(t))]_{s,s'} := \begin{cases} \sum_{c \in \mathcal{C}} \frac{|\{(k,s,s') \in L_c\}|}{v_{k,s}(t)} r_c(\mathbf{v}(t)) & \\ \qquad \text{if } s \in \mathcal{A}_e \text{ and } s \neq s' & \\ -\sum_{x \in \mathcal{S}_k} [Q_{k,e}(\mathbf{v}(t))]_{s,x} & \\ \qquad \text{if } s = s' & \\ 0 \qquad \text{otherwise} \end{cases}$$

Informally, $\mathbf{Q}_{k,e}(\mathbf{v}(t))$ is a time-dependent generator matrix which tracks the local evolution of a class-k component while the deterministic transition corresponding to event clock e is enabled. More specifically, let $\mathbf{v} := (\mathbf{v}(s))_{s \in [0,T]}$ represent the entire solution history on the (implicit) domain $[0, T]$ on which it is being solved and write $\mathbf{Y}^{t_0}_{k,e}(t, \mathbf{v})$ for the *generalised matrix exponential* of $\mathbf{Q}_{k,e}(\mathbf{v}(t))$ — that is, the solution at time t of the non-autonomous linear ordinary initial value problem defined by:

$$\begin{aligned} \mathbf{Y}^{t_0}_{k,e}(0, \mathbf{v}) &= \mathbf{I} \\ \dot{\mathbf{Y}}^{t_0}_{k,e}(u, \mathbf{v}) &= \mathbf{Y}^{t_0}_{k,e}(u, \mathbf{v}) \mathbf{Q}_{k,e}(\mathbf{v}(t_0 + u)) \end{aligned} \tag{4}$$

for $u \in [0, t]$. For $s, s' \in \mathcal{A}_e$, the quantity $[Y^{t_0}_{k,e}(d_e, \mathbf{v})]_{s,s'}$ then represents the mean-field approximation of the conditional probability that given that a class-k component entered state s at time t_0 and initiated the deterministic transition corresponding to e, this transition *was not* disabled before time $t_0 + d_e$ and the component is in state s' just before time $t_0 + d_e$.

In terms of $\mathbf{Y}^{t_0}_{k,e}(t, \mathbf{v})$, we define the function $g^e_{k,s,s'}(t, \mathbf{v})$, which will capture the aggregate rate at which class-k components enter state $s' \in \mathcal{S}_k$ from state $s \in \mathcal{A}_e$ due to deterministic transitions corresponding to e completing at time t:[6]

$$\begin{aligned} g^e_{k,s,s'}(t, \mathbf{v}) := \mathbf{1}_{\{t \geq d_e\}} & \\ \times \sum_{z \in \mathcal{A}_e} \Bigg[& \underbrace{\sum_{c \in \mathcal{C}} |\{(k, y, z) \in L_c : y \notin \mathcal{A}_e\}| r_c(\mathbf{v}(t - d_e))}_{\text{Rate of } e \text{ clocks starting at } t - d_e \text{ due to expo. transits. into } z} \\ + & \underbrace{\sum_{e' \in \mathcal{E}_k} \sum_{y \in \mathcal{A}_{e'}} g^{e'}_{k,y,z}(t - d_e, \mathbf{v})}_{\text{Rate of } e \text{ clocks starting at } t - d_e \text{ due to other clocks completing}} \Bigg] \\ \times & \underbrace{[Y^{t-d_e}_{k,e}(d_e, \mathbf{v})]_{z,s} \times p_e(s, s')}_{\text{Probability that } e \text{ completes at time } t \text{ ending in state } s' \text{ via } s} \end{aligned} \tag{5}$$

Finally, the DDE corresponding to $v_{k,s}(t)$ is then constructed by summing the flux due to exponential transitions and that due to deterministic transitions as captured by the function $g^e_{k,s,s'}(t, \mathbf{v})$:

$$\begin{aligned} \dot{v}_{k,s}(t) = & \sum_{c \in \mathcal{C}} l^c_{k,s} r_c(\mathbf{v}(t)) \\ & + \sum_{e \in \mathcal{E}_k} \left[\sum_{s' \in \mathcal{A}_e} g^e_{k,s',s}(t, \mathbf{v}) - \sum_{s' \in \mathcal{S}_k} g^e_{k,s,s'}(t, \mathbf{v}) \right] \end{aligned} \tag{6}$$

with initial conditions again given by $\mathbf{v}(0) := \mathbf{x}^0$. At this stage it becomes clear why this mean-field equation makes sense necessarily only for cycle-free PGSMPs — Eq. (5) is a recursive definition and is guaranteed to terminate finitely in the cycle-free case.

As in the delay-only case, we also require the addition of an instantaneous jump term for the equation to be correct in the case of initial conditions with components starting in local states enabling deterministic transitions. In this case, the jump term $v^0_{k,s}(t)$ again has a similar form to the second term of Eq. (6), except for that we compute possible deterministic paths forward from time zero rather than backward from time t:

$$v^0_{k,s}(t) := \sum_{s' \in \mathcal{S}_k} \sum_{e \in \mathcal{E}_k} \sum_{z \in \mathcal{A}_e} x^0_{k,z} \left(g^{e,z}_{k,s',s}(t, 0, \mathbf{v}) - g^{e,z}_{k,s,s'}(t, 0, \mathbf{v}) \right)$$

and:

$$\begin{aligned} g^{e,z}_{k,s,s'}(t, w, \mathbf{v}) := \mathbf{1}_{\{t \geq w + d_e\}} \Bigg(& [Y^w_{k,e}(d_e, \mathbf{v})]_{z,s} \times p_e(s, s') \\ + \sum_{y \in \mathcal{A}_e} [Y^w_{k,e}(d_e, \mathbf{v})]_{z,y} & \sum_{e' \in \mathcal{E}_k} \sum_{f \in \mathcal{A}_{e'}} p_e(y, f) g^{e',f}_{k,s,s'}(t, w + d_e, \mathbf{v}) \Bigg) \end{aligned}$$

Then writing Eq. (6) as an integral equation as for the delay-

[6] We define $g^e_{k,s,s'}(t, \mathbf{v})$ to be identically zero if $s \notin \mathcal{A}_e$.

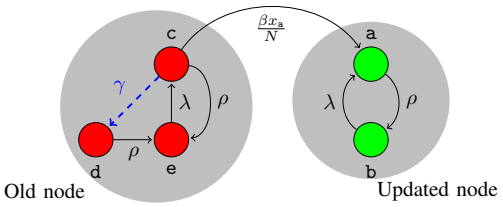

Fig. 3: *Representation of the behaviour of a single node in the software update model.*

only case allows us to incorporate the jump term:

$$v_{k,s}(t) = \sum_{c\in\mathcal{C}} l^c_{k,s} \int_0^t r_c(\mathbf{v}(u))\,\mathrm{d}u + v^0_{k,s}(t) + \sum_{e\in\mathcal{E}_k} \int_0^t \left(\sum_{s'\in\mathcal{A}_e} g^e_{k,s',s}(u,\mathbf{v}) - \sum_{s'\in\mathcal{S}_k} g^e_{k,s,s'}(u,\mathbf{v}) \right) \mathrm{d}u \tag{7}$$

As in the delay-only case, solutions may not be smooth, or even continuous, at time points $t = d_e$ for $e \in \mathcal{E}_k$, $k \in \mathcal{K}$. Furthermore, if there are deterministic transitions linking states in a chain, this non-smoothness will also extend to the times that are the sums of the deterministic delays in the chain.

Although we only prove mean-field convergence for delay-only deterministically-timed PGSMP models in this paper, we fully expect that such a result holds also in the more general case considered in this section. It appears that stronger proof devices will be required since it is not clear that PGSMPs that are not delay only admit a straightforward representation in terms of Poisson processes as in Eq. (3). These issues will be addressed in a paper that is currently in preparation.

A. Worked example: peer-to-peer software update model

In this section, we consider a version of the software update model introduced in Section III-A with deterministic timeout transitions. This is possible now that we have introduced the general mean-field analysis for PGSMPs. Specifically, we now consider the same model where all transitions are Markovian except for the timeout transition, which has deterministic duration γ. The local behaviour of an individual node in this version of the model is given in Figure 3. The Markovian transitions are the same as those in Table I except for the omission of $c = 4$ and the addition of one with change multiset $\{\!|(\mathtt{e},\mathtt{c})|\!\}$ and rate function $\lambda x_{\mathtt{e}}$, and one with change multiset $\{\!|(\mathtt{b},\mathtt{a})|\!\}$ and rate function $\lambda x_{\mathtt{b}}$. There is now only one event clock $\mathtt{t}$ with $\mathcal{A}_{\mathtt{t}} := \{\mathtt{c}\}$, $d_{\mathtt{t}} := \gamma$ and $p_{\mathtt{t}}(\mathtt{c},\mathtt{d}) := 1$.

In order to derive the corresponding system of DDEs, we first construct the matrix $\mathbf{Q}_{\mathtt{t}}(\mathbf{v}(t))$. Its only non-zero row is that corresponding to local state $\mathtt{c}$ and, if we assume that the local states are ordered $\mathtt{c},\mathtt{a},\mathtt{b},\mathtt{d},\mathtt{e}$ in the matrix, it is given by:

$$\left(-\rho - \frac{\beta v_{\mathtt{a}}(t)}{N}, \frac{\beta v_{\mathtt{a}}(t)}{N}, 0, 0, \rho\right)$$

The only element of $\mathbf{Y}^{t_0}_e(t,\mathbf{v})$ that we require in this case can be computed directly in terms of the mean-field solutions:

$$[Y^{t_0}_e(t,\mathbf{v})]_{\mathtt{c},\mathtt{c}} = \exp\left(-\int_{t_0}^{t_0+t} \frac{\beta v_{\mathtt{a}}(s)}{N}\,\mathrm{d}s - \rho t\right)$$

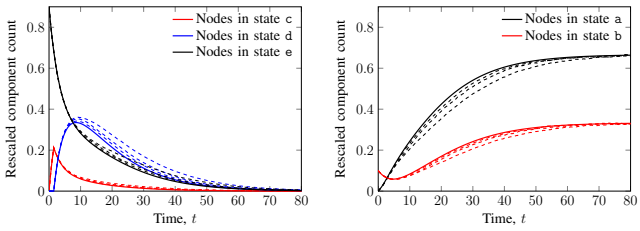

(a) *Old nodes* **(b)** *Updated nodes*

Fig. 4: *Software update model rescaled DDE approximation (solid lines) compared with rescaled actual means for $N = 20$, 50 and 100 (dashed lines). Initial component proportions are $(0, 0.1, 0, 0, 0.9)$ and parameters are $\lambda = 0.2$, $\beta = 2.0$, $\rho = 0.1$ and $\gamma = 1.5$.*

From this we may compute:

$$g^{\mathtt{t}}_{\mathtt{c},\mathtt{d}}(t,\mathbf{v}) = \mathbf{1}_{\{t\geq\gamma\}} \lambda v_{\mathtt{e}}(t-\gamma) \exp\left(-\int_{t-\gamma}^t \frac{\beta v_{\mathtt{a}}(s)}{N}\,\mathrm{d}s - \rho\gamma\right)$$

Finally, the system of DDEs is then:

$$\begin{aligned}
\dot{v}_{\mathtt{a}}(t) =& -\rho v_{\mathtt{a}}(t) + \frac{\beta v_{\mathtt{c}}(t) v_{\mathtt{a}}(t)}{N} + \lambda v_{\mathtt{b}}(t)\\
\dot{v}_{\mathtt{b}}(t) =& -\lambda v_{\mathtt{b}}(t) + \rho v_{\mathtt{a}}(t)\\
\dot{v}_{\mathtt{c}}(t) =& -\rho v_{\mathtt{c}}(t) - \frac{\beta v_{\mathtt{c}}(t) v_{\mathtt{a}}(t)}{N} + \lambda v_{\mathtt{e}}(t)\\
&- \mathbf{1}_{\{t\geq\gamma\}} \lambda v_{\mathtt{e}}(t-\gamma) \exp\left(-\int_{t-\gamma}^t \frac{\beta v_{\mathtt{a}}(s)}{N}\,\mathrm{d}s - \rho\gamma\right)\\
\dot{v}_{\mathtt{d}}(t) =& -\rho v_{\mathtt{d}}(t)\\
&+ \mathbf{1}_{\{t\geq\gamma\}} \lambda v_{\mathtt{e}}(t-\gamma) \exp\left(-\int_{t-\gamma}^t \frac{\beta v_{\mathtt{a}}(s)}{N}\,\mathrm{d}s - \rho\gamma\right)\\
\dot{v}_{\mathtt{e}}(t) =& -\lambda v_{\mathtt{e}}(t) + \rho v_{\mathtt{d}}(t) + \rho v_{\mathtt{c}}(t)
\end{aligned}$$

The solution for one set of parameters is shown in Figure 4 compared with the corresponding component-count expectations as computed by many stochastic simulation replications. We observe that the means do appear to converge to the mean-field solutions.

In this case, we note that it is not necessary to entirely recompute the integral $\int_{t-\gamma}^t \frac{\beta v_{\mathtt{a}}(s)}{N}\,\mathrm{d}s$ at each time step of the numerical procedure. That is, we may write, for small δt:

$$\int_{t+\delta t-\gamma}^{t+\delta t} \frac{\beta v_{\mathtt{a}}(s)}{N}\,\mathrm{d}s \approx \int_{t-\gamma}^t \frac{\beta v_{\mathtt{a}}(s)}{N}\,\mathrm{d}s + \delta t \left(\frac{\beta v_{\mathtt{a}}(t)}{N} - \frac{\beta v_{\mathtt{a}}(t-\gamma)}{N}\right)$$

This means that the additional time complexity for solving this system of five DDEs over a system of five ordinary differential equations is of the order of the cost of one additional ODE.

For more general models, where the Markovian behaviour that can occur concurrently with enabled deterministic transitions is more complicated, it may not be possible to construct the required elements of $\mathbf{Y}^{t-d_e}_{k,e}(d_e,\mathbf{v})$ explicitly in terms of the mean-field solutions. In such cases, the auxiliary linear ODE of Eq. (4) can be solved numerically at each time step, up to time d_e. For deterministic delays which are small compared to the total time period of interest (such as in the model of this section), this also does not represent a significant additional

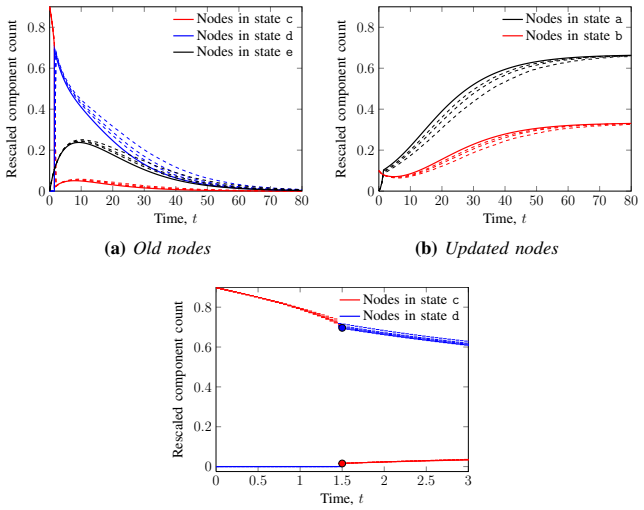

(a) *Old nodes* **(b)** *Updated nodes*

(c) *Old nodes — around* $t = \gamma$

Fig. 5: *Software update model rescaled DDE approximation (solid lines) with a jump term compared with rescaled actual means for* $N = 20$*,* 50 *and* 100 *(dashed lines). Initial component proportions are* $(0, 0.1, 0.9, 0, 0)$ *and parameters are* $\lambda = 0.2$*,* $\beta = 2.0$*,* $\rho = 0.1$ *and* $\gamma = 1.5$*.*

overhead. Another alternative is to introduce additional DDEs into the system so that $\mathbf{Y}_{k,e}^{t-d_e}(d_e, \mathbf{v})$ is integrated simultaneously with $\mathbf{v}(t)$.

We now consider a set of initial conditions where there are components initially in state c, so that the deterministic timeout transition is enabled at time zero. This means that we must use the integral form of the mean-field equations with the jump term given in Eq. (7). In the case of the example of this section, the jump term is zero for all local states except for c and d:

$$v_{\mathtt{c}}^0(t) = -\mathbf{1}_{\{t \geq \gamma\}} x_{\mathtt{c}}^0 \exp\left(-\int_0^\gamma \frac{\beta v_{\mathtt{a}}(s)}{N}\,\mathrm{d}s - \rho\gamma\right)$$
$$v_{\mathtt{d}}^0(t) = \mathbf{1}_{\{t \geq \gamma\}} x_{\mathtt{c}}^0 \exp\left(-\int_0^\gamma \frac{\beta v_{\mathtt{a}}(s)}{N}\,\mathrm{d}s - \rho\gamma\right)$$

Figure 5 shows a comparison of the mean-field solutions and simulated expectations with such an initial condition and the resulting jump at time $t = \gamma$.

V. Mean-field approximations in the stationary regime

The stationary component counts of a CTMC amenable to the mean-field approach can be approximated inexpensively by solving the algebraic fixed point equation obtained by setting the right-hand side of the mean-field ODEs to zero, which is usually much less expensive than simply integrating the ODEs over a suitably large time horizon. For theoretical mean-field convergence as $N \to \infty$ to hold necessarily in the limit $t \to \infty$, it is required that the unique fixed point of the system of ODEs is globally asymptotically stable (attractive) [14, 27], that is, that all trajectories starting from CTMC initial states converge to this fixed point in the limit $t \to \infty$.

However, in the case of PGSMPs, it does not appear that the situation is so straightforward. To see why, consider the very simple component class depicted in Figure 6. An instance of

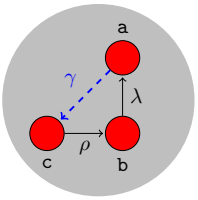

Fig. 6: *A very simple PGSMP component class.*

this class repeatedly performs two Markovian transitions with a deterministic delay in between. The system of DDEs derived according to Section III are:

$$\begin{aligned}
\dot{v}_{\mathtt{a}}(t) &= -\mathbf{1}_{\{t \geq \gamma\}} \lambda v_{\mathtt{b}}(t-\gamma) + \lambda v_{\mathtt{b}}(t) \\
\dot{v}_{\mathtt{b}}(t) &= -\lambda v_{\mathtt{b}}(t) + \rho v_{\mathtt{c}}(t) \\
\dot{v}_{\mathtt{c}}(t) &= -\rho v_{\mathtt{c}}(t) + \mathbf{1}_{\{t \geq \gamma\}} \lambda v_{\mathtt{b}}(t-\gamma)
\end{aligned} \tag{8}$$

If we assume that, as $t \to \infty$, a solution to this system of DDEs converges to a constant value, say $v_{\mathtt{a}}$, $v_{\mathtt{b}}$ and $v_{\mathtt{c}}$, we have the following system of algebraic fixed point equations:

$$0 = -\lambda v_{\mathtt{b}} + \lambda v_{\mathtt{b}}, \quad 0 = -\lambda v_{\mathtt{b}} + \rho v_{\mathtt{c}}, \quad 0 = -\rho v_{\mathtt{c}} + \lambda v_{\mathtt{b}}$$

This reduces simply to $\lambda v_{\mathtt{b}} = \rho v_{\mathtt{c}}$ and we see that this system of DDEs has infinitely many fixed points, even when the population constraint $v_{\mathtt{a}} + v_{\mathtt{b}} + v_{\mathtt{c}} = N$ is included. The immediate conclusion from this would therefore be that not one of these fixed points can be globally attractive for each PGSMP initial condition, and thus, that the mean-field approach should not be applied here in the stationary regime. However, this ignores the fact that for given initial conditions, the DDEs have an initial period of time-inhomogeneous dynamics caused by indicator functions on time, such as those in Eq. (8), and, in the case of initial conditions enabling determinitistic transitions, additional jump terms activated at particular time instants. This means that fixed points of the DDEs are not necessarily stationary solutions when supplied as initial conditions, and the existence of infinitely many fixed points does not then necessarily prohibit a single one of the fixed points from being a globally stable attractor of the system of mean-field equations.

For example, in the PGSMP model of Figure 6, all agents operate independently, and, for initial conditions that are a probability vector, the equations of Eq. (8) thus capture the evolution of probability mass for a single component exactly. Since a single component in this model has a unique stationary distribution, it thus follows that in this case, the equations of Eq. (8) must have a global attractor which is this stationary distribution.

In the more general mean-field case, we will consider PGSMP models which we assume, firstly, have a unique stationary distribution, and, secondly, their mean-field equations have a unique globally attractive fixed point in the above sense. However, this still leaves open the issue of isolating and computing this specific fixed point efficiently. In the next section, we present a method aimed at achieving this based on Markov regenerative processes. It is beyond the scope of this paper to prove that the method always succeeds in

computing the attractive fixed point, and, further, that the stationary distribution converges to it, although, our conjecture is that this is true under the above assumptions.

A. Individual components as Markov regenerative processes

Write $\mathbf{v}$ for the mean-field approximation to the stationary component counts of the PGSMP model under consideration, where $\mathbf{v} = (\mathbf{v}_k)_{k\in\mathcal{K}}$ and each $\mathbf{v}_k = (v_{k,s})_{s\in\mathcal{S}_k}$. In this section, we construct a system of algebraic equations which can then be solved for the $v_{k,s}$. In order to do so, we will show how, under mean-field and stationarity assumptions, the dynamics of an individual class-k component can be approximated by that of a *Markov regenerative process (MRGP)* [4, 28, 29].[7] The definition of an MRGP is as follows [28, 29].

Definition 1 (Markov regenerative process (MRGP))**.** *A continuous-time stochastic process $Z(t)$ is called a Markov regenerative process if there exists a Markov renewal sequence $\{(Y_m, T_m)\}_{m\in\mathbb{Z}_+}$ such that the conditional distribution of $Z(T_m + t)$ given $Y_m = i$ and the information $\{Z(u)\}_{0\le u\le T_m}$ is equal to that of $Z(t)$ given $Y_0 = i$.*

For a given $\mathbf{v} \in \mathbb{R}^n_+$, we can consider redefining all Markovian transition rate functions independently of actual component counts by setting them equal to $r_c(\mathbf{v})$. Under such a modification, each individual component in the PGSMP can be analysed independently of other components but with Markovian transition rates that depend on the choice of $\mathbf{v}$. It turns out that, for a given $\mathbf{v}$, a class-k component evolves as an MRGP, which is a direct consequence of the fact that each individual component may only ever enable at most one deterministically-timed transition in a given local state.

Specifically, let $X^{\mathbf{v}}_k(t) \in \mathcal{S}_k$ be the stochastic process tracking the state of a class-k component in the case that its transition rates are defined in terms of the particular quantity $\mathbf{v} \in \mathbb{R}^n_+$. Then we may construct a Markov renewal sequence $\{(Y^{k,\mathbf{v}}_m, T^{k,\mathbf{v}}_m)\}_{m\in\mathbb{Z}_+}$ in accordance with Definition 1 by letting $T^{k,\mathbf{v}}_0 := 0$ and defining recursively:

$$T^{k,\mathbf{v}}_{m+1} := \begin{cases} \inf\{t > T^{k,\mathbf{v}}_m : X^{\mathbf{v}}_k(t) \notin \mathcal{A}_e\} \\ \qquad \text{if } X^{\mathbf{v}}_k(T^{k,\mathbf{v}}_m) \in \mathcal{A}_e \text{ for some } e \in \mathcal{E}_k \\ \inf\{t > T^{k,\mathbf{v}}_m : X^{\mathbf{v}}_k(t) \neq X^{\mathbf{v}}_k(T^{k,\mathbf{v}}_m)\} \\ \qquad \text{otherwise} \end{cases}$$

and $Y^{k,\mathbf{v}}_m := X^{\mathbf{v}}_k(T^{k,\mathbf{v}}_m)$.

Since then $X^{\mathbf{v}}_k(t) \in \mathcal{S}_k$ is an MRGP for each $\mathbf{v} \in \mathbb{R}^n_+$, we may compute its stationary distribution using known formulae for MRGPs [28, 29], to be described below, resulting in stationary probabilities, say $\pi_{k,s}(\mathbf{v})$. Under the mean-field assumptions, it then makes sense to formulate the system of equations $v_{k,s} = N_k\pi_{k,s}(\mathbf{v})$ where N_k is the total population of class-k components in the model. This system of algebraic equations can then be solved directly for the mean-field approximation to the stationary component counts $\mathbf{v}$.

It remains to show how to compute $\pi_{k,s}(\mathbf{v})$. An MRGP $X(t)$ is characterised by its *kernel* $\mathbf{K}(t) := [K_{ij}(t)]_{ij}$ where $K_{ij}(t) := \mathbb{P}\{Y_1 = j, T_1 \le t | Y_0 = i\}$, and its *local kernel* $\mathbf{E}(t) := [E_{ij}(t)]_{ij}$, where $E_{ij}(t) := \mathbb{P}\{X(t) = j, T_1 > t | Y_0 = i\}$. In the case of steady-state analysis, we are in fact only interested in the quantities $\mathbf{P} := \mathbf{K}(\infty)$, $\mathbf{A} := \int_0^\infty \mathbf{E}(t)\,\mathrm{d}t$ and $\mu := [\mu_i]_i$ where $\mu_i := \mathbb{E}[T_1 | Y_0 = i]$. In the case of the MRGP $X^{\mathbf{v}}_k(t)$, $\mathbf{P}_k(\mathbf{v})$ is given for $s, s' \in \mathcal{S}_k$ by:

$$[P_k(\mathbf{v})]_{s,s'} = \begin{cases} \sum_{c\in\mathcal{C}} \frac{|\{(k,s,s')\in L_c\}|}{v_{k,s}} r_c(\mathbf{v})/R_s(\mathbf{v}) \\ \qquad \text{if } s \notin \mathcal{A}_e \text{ for any } e \in \mathcal{E}_k \text{ and } R_s(\mathbf{v}) \neq 0 \\ \sum_{x\in\mathcal{A}_e} [\exp(\mathbf{Q}_{k,e}(\mathbf{v})d_e)]_{s,x}\, p_e(x, s') \\ \qquad \text{if } s, s' \in \mathcal{A}_e \text{ for some } e \in \mathcal{E}_k \\ \sum_{x\in\mathcal{A}_e} [\exp(\mathbf{Q}_{k,e}(\mathbf{v})d_e)]_{s,x}\, p_e(x, s') \\ \quad + [\exp(\mathbf{Q}_{k,e}(\mathbf{v})d_e)]_{s,s'} \\ \qquad \text{if } s \in \mathcal{A}_e \text{ and } s' \notin \mathcal{A}_e \text{ for some } e \in \mathcal{E}_k \\ 0 \qquad \text{otherwise} \end{cases}$$

and $R_s(\mathbf{v}) := \sum_{c\in\mathcal{C}} \frac{|\{(k,s,\cdot)\in L_c\}|}{v_{k,s}} r_c(\mathbf{v})$. Furthermore, $\mathbf{A}_k(\mathbf{v})$ is given by:

$$[A_k(\mathbf{v})]_{s,s'} = \begin{cases} 1/R_s(\mathbf{v}) \qquad \text{if } s = s' \notin \mathcal{A}_e \text{ for any } e \in \mathcal{E}_k \\ \int_0^{d_e} [\exp(\mathbf{Q}_{k,e}(\mathbf{v})t)]_{s,s'}\,\mathrm{d}t \\ \qquad \text{if } s, s' \in \mathcal{A}_e \text{ for some } e \in \mathcal{E}_k \\ 0 \qquad \text{otherwise} \end{cases}$$

and finally:

$$[\mu_k(\mathbf{v})]_s = \begin{cases} 1/R_s(\mathbf{v}) \qquad \text{if } s \notin \mathcal{A}_e \text{ for any } e \in \mathcal{E}_k \\ \sum_{x\in\mathcal{A}_e} \int_0^{d_e} [\exp(\mathbf{Q}_{k,e}(\mathbf{v})t)]_{s,x}\,\mathrm{d}t \\ \qquad \text{if } s \in \mathcal{A}_e \text{ for some } e \in \mathcal{E}_k \end{cases}$$

Then if $\mathbf{p}(\mathbf{v})$ is a probability vector solution to the linear system $\mathbf{p}(\mathbf{v}) = \mathbf{p}(\mathbf{v})\mathbf{P}_k(\mathbf{v})$, we have:

$$\pi_{k,s}(\mathbf{v}) = \sum_{x\in\mathcal{S}_k} \beta_x \frac{[A_k(\mathbf{v})]_{xs}}{[\mu_k(\mathbf{v})]_x}$$

where $\beta_x := \frac{p_x(\mathbf{v})[\mu_k(\mathbf{v})]_x}{\sum_{y\in\mathcal{S}_k} p_y(\mathbf{v})[\mu_k(\mathbf{v})]_y}$.

It is important to note that computation of $\pi_{k,s}(\mathbf{v})$ requires only standard matrix exponentials as opposed to the generalised matrix exponentials which were required in Section IV. In some cases, where the Markovian behaviour that can occur concurrently with deterministic transitions has a suitable structure, it is even possible to compute the matrix exponentials $\exp(\mathbf{Q}_{k,e}(\mathbf{v})t)$ symbolically in terms of $\mathbf{v}$, t and the model parameters using, for example, the function `mexp` in MATLAB®. This is particularly beneficial if the model is to be analysed for many different parameter combinations. The worked example in the next section is one such model.

B. Worked example

In this section we illustrate the technique of the previous section with a simple wireless sensor network worked example. This model has two component classes: class-`n` components are wireless sensor nodes, and class-`g` components are gateway nodes. There are N components in total. The local state space of a class-`n` component is $\mathcal{S}_{\mathtt{n}} := \{\mathtt{c}, \mathtt{d}, \mathtt{e}\}$ and of a class-`g` component is $\mathcal{S}_{\mathtt{g}} := \{\mathtt{a}, \mathtt{b}\}$. Wireless sensor

[7] Note that MRGPs are sometimes also referred to in the literature as *semi-regenerative processes* [28].

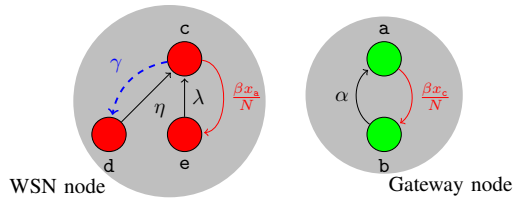

Fig. 7: *Representation of the behaviour of wireless sensor network (WSN) nodes and gateway nodes in the wireless sensor network model.*

Transition, $c \in \mathcal{C}$	Change multiset, L_c	Rate function, r_c
1	$\{\!\{(\mathtt{a},\mathtt{b}),(\mathtt{c},\mathtt{e})\}\!\}$	$\frac{\beta x_{\mathtt{a}} x_{\mathtt{c}}}{N}$
2	$\{\!\{(\mathtt{b},\mathtt{a})\}\!\}$	$\alpha x_{\mathtt{b}}$
3	$\{\!\{(\mathtt{d},\mathtt{c})\}\!\}$	$\eta x_{\mathtt{d}}$
4	$\{\!\{(\mathtt{e},\mathtt{c})\}\!\}$	$\lambda x_{\mathtt{e}}$

TABLE II: *Markovian transitions for wireless sensor network model.*

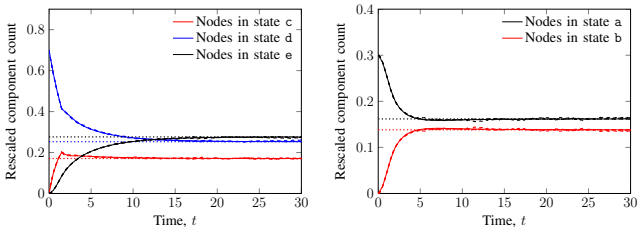

(a) *WSN nodes* **(b)** *Gateway nodes*

Fig. 8: *WSN model rescaled DDE approximation (solid lines) and stationary rescaled component-count approximations computed using the method of Section V (dotted horizontal lines) compared with rescaled actual means for* $(N_{\mathtt{n}}, N_{\mathtt{g}}) = (14, 6)$, $(35, 15)$ *and* $(70, 30)$ *(dashed lines). Initial component counts for the DDEs and simulation are* $(N_g, 0, 0, N_n, 0)$ *and parameters are* $\alpha = 0.4$, $\lambda = 0.2$, $\beta = 2.0$, $\gamma = 1.5$ *and* $\eta = 0.35$.

nodes begin in state e where they may detect an event of interest, after which they move into state c. Here they attempt to communicate observed data to a gateway node, but give up after a deterministic timeout, moving to state d. If they timeout, they try again to transmit the data after a further (Markovian) delay. After a gateway node has received data from a wireless sensor node, it enters state b where it is temporarily unable to receive further communication from wireless nodes, reflecting a processing delay. The deterministic timeout is specified by the single event clock $\mathtt{t} \in \mathcal{E}_{\mathtt{n}}$ with $\mathcal{A}_{\mathtt{t}} := \{\mathtt{c}\}$, $d_{\mathtt{t}} := \gamma$ and $p_{\mathtt{t}}(\mathtt{c},\mathtt{d}) := 1$. The local behaviour of both types of node is depicted graphically in Figure 7 and the Markovian transitions are given in Table II. We drop explicit reference to the component class in notation where there is no ambiguity.

Following the last section, we construct the quantities $\mathbf{P}_{\mathtt{n}}(\mathbf{v})$, $\mathbf{P}_{\mathtt{g}}(\mathbf{v})$, $\mathbf{A}_{\mathtt{n}}(\mathbf{v})$, $\mathbf{A}_{\mathtt{g}}(\mathbf{v})$, $\mu_{\mathtt{n}}(\mathbf{v})$ and $\mu_{\mathtt{g}}(\mathbf{v})$ as follows, where the enumeration of local states in matrices and vectors is in the order $\mathtt{c},\mathtt{d},\mathtt{e}$ and $\mathtt{a},\mathtt{b}$, respectively:

$$\mathbf{P}_{\mathtt{n}}(\mathbf{v}) = \begin{pmatrix} 0 & \exp(\frac{-\gamma\beta v_{\mathtt{a}}}{N}) & 1-\exp(\frac{-\gamma\beta v_{\mathtt{a}}}{N}) \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}$$

$$\mathbf{A}_{\mathtt{n}}(\mathbf{v}) = \begin{pmatrix} (\beta v_{\mathtt{a}}/N)^{-1}(1-\exp(\frac{-\gamma\beta v_{\mathtt{a}}}{N})) & 0 & 0 \\ 0 & \eta^{-1} & 0 \\ 0 & 0 & \lambda^{-1} \end{pmatrix}$$

$$\mathbf{P}_{\mathtt{g}}(\mathbf{v}) = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \qquad \mathbf{A}_{\mathtt{g}}(\mathbf{v}) = \begin{pmatrix} (\beta v_{\mathtt{c}}/N)^{-1} & 0 \\ 0 & \alpha^{-1} \end{pmatrix}$$

$$\mu_{\mathtt{n}}(\mathbf{v}) = ((\beta v_{\mathtt{a}}/N)^{-1}(1-\exp(-\gamma\beta v_{\mathtt{a}}/N)), \eta^{-1}, \lambda^{-1})^T$$
$$\mu_{\mathtt{g}}(\mathbf{v}) = ((\beta v_{\mathtt{c}}/N)^{-1}, \alpha^{-1})^T$$

For given $N_{\mathtt{n}}$ and $N_{\mathtt{g}}$ with $N = N_{\mathtt{n}} + N_{\mathtt{g}}$, we may then construct the coupled system of algebraic equations described in the previous section: $\mathbf{v}_{\mathtt{n}} = N_{\mathtt{n}}\pi_{\mathtt{n}}(\mathbf{v})$ and $\mathbf{v}_{\mathtt{g}} = N_{\mathtt{g}}\pi_{\mathtt{g}}(\mathbf{v})$, which can be solved efficiently for $\mathbf{v}$, subject to the population conservation constraints $v_{\mathtt{c}} + v_{\mathtt{d}} + v_{\mathtt{e}} = N_{\mathtt{n}}$ and $v_{\mathtt{a}} + v_{\mathtt{b}} = N_{\mathtt{g}}$, by any standard iterative numerical procedure such as that implemented by `fmincon` in MATLAB®. The DDE solution for one set of parameters is shown in Figure 8 compared with the stationary rescaled component-count approximations computed using the method of the last section. We see that in the long-time limit, the DDEs, the simulated means (in the limit of large populations), and the stationary rescaled component counts computed using the method of the last section all agree.

Another interesting aspect of the approach of this section is that the *cycle free* restriction required for the derivation of DDEs in Section IV can be dropped if only stationary component-count approximations are required. This is because the algorithm of the last section does not actually require the construction of the associated system of DDEs and can still be applied to a non-cycle-free PGSMP model since the MRGP construction is still correct. To illustrate this, we consider in Figure 9 a version of the wireless sensor network model where the Markovian transition $c = 3$ is replaced with another deterministic transition of duration τ. Additionally, we add a further Markovian transition with change multiset $\{\!\{(a,b),(d,e)\}\!\}$ and rate function $\frac{\beta_2 x_{\mathtt{a}} x_{\mathtt{d}}}{N}$. This captures the situation where after the initial timeout of duration γ, the wireless node enters a state where it communicates its observed data at a potentially reduced transmission level to save power. If, after a further deterministic timeout of duration τ, this has still not been successful, the node reverts to attempting transmission at the original level. This version of the model is clearly not cycle free since there is an unbroken cycle of deterministic transitions between local states c and d.

The quantities $\mathbf{P}_{\mathtt{n}}(\mathbf{v})$, $\mathbf{P}_{\mathtt{g}}(\mathbf{v})$, $\mathbf{A}_{\mathtt{n}}(\mathbf{v})$, $\mathbf{A}_{\mathtt{g}}(\mathbf{v})$, $\mu_{\mathtt{n}}(\mathbf{v})$ and $\mu_{\mathtt{g}}(\mathbf{v})$ in this case are then:

$$\mathbf{P}_{\mathtt{n}}(\mathbf{v}) = \begin{pmatrix} 0 & \exp(\frac{-\gamma\beta v_{\mathtt{a}}}{N}) & 1-\exp(\frac{-\gamma\beta v_{\mathtt{a}}}{N}) \\ \exp(\frac{-\tau\beta_2 v_{\mathtt{a}}}{N}) & 0 & 1-\exp(\frac{-\tau\beta_2 v_{\mathtt{a}}}{N}) \\ 1 & 0 & 0 \end{pmatrix}$$

$$\mathbf{A}_{\mathtt{n}}(\mathbf{v}) = \begin{pmatrix} \frac{1-\exp(\frac{-\gamma\beta v_{\mathtt{a}}}{N})}{\beta v_{\mathtt{a}}/N} & 0 & 0 \\ 0 & \frac{1-\exp(\frac{-\tau\beta_2 v_{\mathtt{a}}}{N})}{\beta_2 v_{\mathtt{a}}/N} & 0 \\ 0 & 0 & \lambda^{-1} \end{pmatrix}$$

$$\mathbf{P}_{\mathtt{g}}(\mathbf{v}) = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \qquad \mathbf{A}_{\mathtt{g}}(\mathbf{v}) = \begin{pmatrix} (\frac{\beta v_{\mathtt{c}} + \beta_2 v_{\mathtt{d}}}{N})^{-1} & 0 \\ 0 & \alpha^{-1} \end{pmatrix}$$

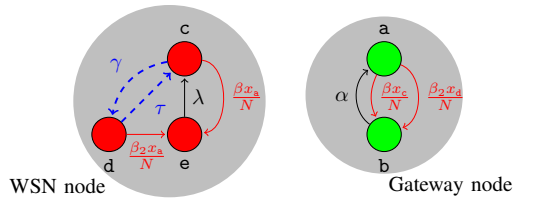

Fig. 9: *Representation of the behaviour of wireless sensor network (WSN) nodes and gateway nodes in the non-cycle-free wireless sensor network model.*

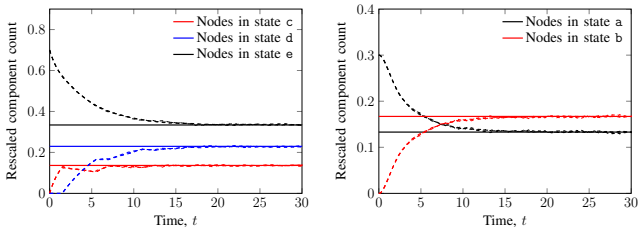

(a) *WSN nodes* **(b)** *Gateway nodes*

Fig. 10: *Non-cycle-free WSN model stationary rescaled component-count approximations computed using the method of Section V (solid horizontal lines) compared with rescaled actual means for* $(N_{\mathtt{n}}, N_{\mathtt{g}}) = (14, 6)$, $(35, 15)$ *and* $(70, 30)$ *(dashed lines). Initial component counts for the simulation are* $(N_g, 0, 0, 0, N_n)$ *and parameters are* $\alpha = 0.4$, $\lambda = 0.2$, $\beta = 2.0$, $\beta = 1.0$, $\gamma = 1.5$ *and* $\tau = 4.0$.

$$\mu_{\mathtt{n}}(\mathbf{v}) = \left(\frac{1 - \exp(\frac{-\gamma\beta v_{\mathtt{a}}}{N})}{\beta v_{\mathtt{a}}/N}, \frac{1 - \exp(\frac{-\tau\beta_2 v_{\mathtt{a}}}{N})}{\beta_2 v_{\mathtt{a}}/N}, \lambda^{-1} \right)^T$$

$$\mu_{\mathtt{g}}(\mathbf{v}) = \left(\left(\frac{\beta v_{\mathtt{c}} + \beta_2 v_{\mathtt{d}}}{N} \right)^{-1}, \alpha^{-1} \right)^T$$

Numerical results for this version of the model are given in Figure 10, and, even in the non-cycle-free case, we see that the approach yields very accurate approximations.

VI. Conclusion

In this paper we have presented a new mean-field framework for the scalable performance analysis of massively-parallel stochastic systems with both Markovian and deterministically-timed transitions. Our approach is based on the derivation of systems of delay differential equations from models specified in the new low-level formalism: population generalised semi-Markov processes (PGSMPs). The class of models to which our approach applies is very broad and, in particular, there are no restrictions on the Markovian behaviour which may be enabled simultaneously with a deterministic transition.

In terms of theoretical results, we present a mean-field convergence result and second-order limit process for a large class of PGSMP models. Additionally, we construct a tractable system of algebraic equations which can be solved efficiently for mean-field approximations to component counts in the steady state. Finally, we have verified the effectiveness and accuracy of our techniques on a variety of worked examples.

With respect to future work, we are currently investigating DDE-based mean-field approximations for the case of PGSMPs with transitions that are more generally timed than deterministic [21]. Furthermore, we also intend to prove mean-field convergence and construct second-order limit processes for more general classes of PGSMPs than considered in this paper.

Acknowledgments

The author is funded by EPSRC on the Analysis of Massively Parallel Stochastic Systems (AMPS) project (reference EP/G011737/1). The author also wishes to thank the anonymous referees for their detailed comments which helped to improve the paper.

References

[1] K. Matthes, "Zur Theorie der Bedienungsprozesse," in *Transactions of the 3rd Prague Conference on Information Theory, Statistical Decision Functions and Random Processes*, 1962, pp. 512–528.
[2] P. W. Glynn, "A GSMP formalism for discrete event systems," *Proceedings of the IEEE*, vol. 77, no. 1, pp. 14–23, 1989.
[3] P. G. Harrison and B. Strulo, "SPADES - a process algebra for discrete event simulation," *Journal of Logic and Computation*, vol. 10, no. 1, pp. 3–42, Feb. 2000.
[4] H. Choi, V. Kulkarni, K. Trivedi, and M. Ajmone Marsan, "Transient analysis of deterministic and stochastic Petri nets," in *Application and Theory of Petri Nets*, vol. 691. Springer Berlin / Heidelberg, 1993, pp. 166–185.
[5] C. Lindemann and A. Thümmler, "Transient analysis of deterministic and stochastic Petri nets with concurrent deterministic transitions," *Performance Evaluation*, vol. 36-37, no. 1, pp. 35–54, Aug. 1999.
[6] A. Puliafito, M. Scarpa, and K. S. Trivedi, "Petri nets with k simultaneously enabled generally distributed timed transitions," *Performance Evaluation*, vol. 32, no. 1, pp. 1–34, Feb. 1998.
[7] D. R. Cox, "The analysis of non-Markovian stochastic processes by the inclusion of supplementary variables," *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 51, no. 03, pp. 433–441, Jul. 1955.
[8] H. Smith, *An Introduction to Delay Differential Equations with Applications to the Life Sciences*. Springer, 2010.
[9] J. Anselmi and M. Verloop, "Energy-aware capacity scaling in virtualized environments with performance guarantees," *Performance Evaluation*, vol. 68, no. 11, pp. 1207–1221, Nov. 2011.
[10] M. Gribaudo, "Analysis of Large Populations of Interacting Objects with Mean Field and Markovian Agents," in *6th European Performance Engineering Workshop (EPEW)*, ser. Lecture Notes in Computer Science, J. T. Bradley, Ed., vol. 5652. Springer Berlin Heidelberg, 2009, pp. 218–219.
[11] N. Gast and G. Bruno, "A mean field model of work stealing in large-scale systems," in *ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, vol. 38, no. 1. New York, USA: ACM Press, Jun. 2010, p. 13.
[12] J. Hillston, "Fluid flow approximation of PEPA models," in *Second International Conference on the Quantitative Evaluation of Systems (QEST)*. IEEE, Sep. 2005, pp. 33–42.
[13] R. Bakhshi, L. Cloth, W. Fokkink, and B. R. Haverkort, "Mean-field framework for performance evaluation of push-pull gossip protocols," *Performance Evaluation*, vol. 68, no. 2, pp. 157–179, Feb. 2011.
[14] M. Benaïm and J.-Y. Le Boudec, "A class of mean field interaction models for computer and communication systems," *Performance Evaluation*, vol. 65, no. 11-12, pp. 823–838, Nov. 2008.
[15] R. A. Hayden and J. T. Bradley, "A fluid analysis framework for a Markovian process algebra," *Theoretical Computer Science*, vol. 411, no. 22-24, pp. 2260–2297, May 2010.
[16] R. Barbuti, G. Caravagna, A. Maggiolo-Schettini, and P. Milazzo, "Delay stochastic simulation of biological systems: a purely delayed approach," *Transactions on Computational Systems Biology XIII*, vol. 6575, pp. 61–84, Jan. 2011.
[17] G. Caravagna and J. Hillston, "Bio-PEPAd: A non-Markovian extension of Bio-PEPA," *Theoretical Computer Science*, vol. 419, pp. 26–49, Feb. 2012.
[18] G. Caravagna, "Formal Modeling and Simulation of Biological Systems with Delays," Ph. D. thesis, Università di Pisa, 2011.
[19] R. Schlicht and G. Winkler, "A delay stochastic process with applications in molecular biology," *Journal of Mathematical Biology*, vol. 57, no. 5, pp. 613–48, Nov. 2008.
[20] M. Barrio, K. Burrage, A. Leier, and T. Tian, "Oscillatory regulation of Hes1: discrete stochastic delay modelling and simulation," *PLoS Computational Biology*, vol. 2, no. 9, pp. 1017–30, 2006.
[21] R. A. Hayden, "Mean-field approximations for performance models with generally-timed transitions," *ACM SIGMETRICS Performance Evaluation Review*, vol. 39, no. 3, pp. 119–121, 2011.
[22] T. A. Henzinger and V. Wolf, "Formalisms for Specifying Markovian Population Models," *International Journal of Foundations of Computer Science*, vol. 22, no. 4, pp. 823–841, 2011.
[23] D. F. Anderson, "A modified next reaction method for simulating chemical systems with time dependent propensities and delays," *The Journal of Chemical Physics*, vol. 127, no. 21, p. 214107, Dec. 2007.
[24] S. N. Ethier and T. G. Kurtz, *Markov Processes: Characterization and Convergence*. Wiley, 2005.
[25] T. G. Kurtz, "Strong approximation theorems for density dependent Markov chains," *Stochastic Processes and their Applications*, vol. 6, no. 3, pp. 223–240, Feb. 1978.
[26] W. Whitt, "Internet supplement to Stochastic-Process Limits," 2002.
[27] R. A. Hayden, A. Stefanek, and J. T. Bradley, "Fluid computation of passage time distributions in large Markov models," *Theoretical Computer Science*, vol. 413, no. 1, pp. 106–141, 2012.
[28] E. Çinlar, *Introduction to Stochastic Processes*. Prentice-Hall, 1975.
[29] V. G. Kulkarni, *Modeling and Analysis of Stochastic Systems*. Chapman & Hall/CRC, 1996.