Polynomial Stochastic Hybrid Systems (Extended Version)

Technical Report

João Pedro Hespanha*

Center for Control Engineering and Computation
University of California, Santa Barbara, CA 93101
hespanha@ece.ucsb.edu
http://www.ece.ucsb.edu/~hespanha

Abstract. This paper deals with polynomial stochastic hybrid systems (pSHSs), which generally correspond to stochastic hybrid systems with polynomial continuous vector fields, reset maps, and transition intensities. For pSHSs, the dynamics of the statistical moments of the continuous states evolve according to infinite-dimensional linear ordinary differential equations (ODEs). We show that these ODEs can be approximated by finite-dimensional nonlinear ODEs with arbitrary precision. Based on this result, we provide a procedure to build this type of approximations for certain classes of pSHSs. We apply this procedure for several examples of pSHSs and evaluate the accuracy of the results obtained through comparisons with Monte Carlo simulations. These examples include: the modeling of TCP congestion control both for long-lived and on-off flows; state-estimation for networked control systems; and the stochastic modeling of chemical reactions.

1 Introduction

Hybrid systems are characterized by a state-space that can be partitioned into a continuous domain (typically $\mathbb{R}^n$) and a discrete set (typically finite). For the stochastic hybrid systems considered here, both the continuous and the discrete components of the state are stochastic processes. The evolution of the continuous-state is determined by a stochastic differential equation and the evolution of the discrete-state by a transition or reset map. The discrete transitions are triggered by stochastic events much like transitions between states of a continuous-time Markov chains. However, the rate at which these transitions occur may depend on the continuous-state. The model used here for SHSs, whose formal definition can be found in Sec. 2, was introduced in [1] and is heavily inspired by the Piecewise-Deterministic Markov Processes (PDPs) in [2]. Alternative models can be found in [3,4,5].

The extended generator of a stochastic hybrid system allows one to compute the time-derivative of a "test function" of the state of the SHS along solutions to the system, and can be viewed as a generalization of the Lie derivative for deterministic systems [1,2]. Polynomial stochastic hybrid systems (pSHSs) are characterized by extended

* Supported by the National Science Foundation under grants CCR-0311084, ANI-0322476.

generators that map polynomial test functions into polynomials. This happens, e.g., when the continuous vector fields, the reset maps, and the transition intensities are all polynomial functions of the continuous state. An important property of pSHSs is that if one creates an infinite vector containing the probabilities of all discrete modes, as well as all the multi-variable statistical moments of the continuous state, the dynamics of this vector are governed by an infinite-dimensional linear ordinary differential equation (ODE), which we call the *infinite-dimensional moment dynamics* and define formally in Sec. 3.

SHSs can model large classes of systems but their formal analysis presents significant challenges. Although it is straightforward to write partial differential equations (PDEs) that express the evolution of the probability distribution function for their states, generally these PDEs do not admit analytical solutions. The infinite-dimensional moment dynamics provides an alternative characterization for the distribution of the state of a pSHS. Although generally statistical moments do not provide a description of a stochastic process as accurate as the probability distribution, results such as Tchebycheff, Markoff, or Bienaymé inequalities [6, pp. 114–115] can be used to infer properties of the distribution from its moments.

In general, the infinite-dimensional *linear* ODEs that describe the moment dynamics for pSHSs are still not easy to solve analytically. However, sometimes they can be accurately approximated by a finite-dimensional *nonlinear* ODE, which we call the *truncated moment dynamics.* We show in Sec. 3 that, under suitable stability assumptions, it is in principle possible for a finite-dimensional nonlinear ODE to approximate the infinite-dimensional moment dynamics, up to an error that can be made arbitrarily small. This result is based on a Taylor series expansion that is used to prove that the difference between the solutions to the finite- and infinite-dimensional ODEs can be made arbitrarily small on a compact time interval $[0,T]$. Subsequently, this is extended to the unbounded time interval $[0,\infty)$ using an argument similar to that found in [7] to analyze two time scale systems.

Aside from its theoretical interest, the above mentioned result motivates a procedure to actually construct these finite-dimensional approximations for certain classes of pSHSs. This procedure, which is described in Sec. 4, is applicable to pSHS for which the (infinite) matrix that characterizes the moment dynamics exhibits a certain diagonal-band structure and appropriate decoupling between certain moments of distinct discrete modes. The details of this structure can be found in Lemma 1.

To illustrate the applicability of the results we consider several systems that appeared in the literature and that can be modeled by pSHSs. For each example, we construct in Sec. 5 truncated moment dynamics and evaluate how they compare with estimates for the moments obtained from a large number of Monte Carlo simulations. The examples considered include:

1. The modeling of the sending rate for network traffic under TCP congestion control. We consider two distinct cases: long-lived traffic corresponding to the transfer of files with infinite length; and on-off traffic consisting of file transfers with exponentially distributed lengths, alternated by times of inactivity (also exponentially distributed). These examples are motivated by [1,8].

2. The modeling of the state-estimation error in a networked control system that occasionally receives state measurements over a communication network. The rate at which the measurements are transmitted depends on the current estimation error. This type of scheme was shown to out-perform periodic transmission and can actually be used to approximate an optimal transmission scheme [9,10].
3. Gillespie's stochastic modeling for chemical reactions [11], which describes the evolution of the number of particles involved in a set of reactions. The reactions analyzed were taken from [12,13] and are particularly difficult to simulate due to the existence of two very distinct time scales.

2 Polynomial Stochastic Hybrid Systems

A SHS is defined by a stochastic differential equation (SDE)

$$\dot{\mathbf{x}} = f(\mathbf{q},\mathbf{x},t) + g(\mathbf{q},\mathbf{x},t)\dot{\mathbf{n}}, \qquad \begin{array}{l} f : \mathscr{Q} \times \mathbb{R}^n \times [0,\infty) \to \mathbb{R}^n, \\ g : \mathscr{Q} \times \mathbb{R}^n \times [0,\infty) \to \mathbb{R}^{n\times k}, \end{array} \tag{1}$$

a family of *m discrete transition/reset maps*

$$(\mathbf{q},\mathbf{x}) = \phi_\ell(\mathbf{q}^-,\mathbf{x}^-,t), \qquad \phi_\ell : \mathscr{Q} \times \mathbb{R}^n \times [0,\infty) \to \mathscr{Q} \times \mathbb{R}^n, \tag{2}$$

$\forall \ell \in \{1,\ldots,m\}$, and a family of *m transition intensities*

$$\lambda_\ell(\mathbf{q},\mathbf{x},t), \qquad \lambda_\ell : \mathscr{Q} \times \mathbb{R}^n \times [0,\infty) \to [0,\infty), \tag{3}$$

$\forall \ell \in \{1,\ldots,m\}$, where $\mathbf{n}$ denotes a k-vector of independent Brownian motion processes and $\mathscr{Q}$ a (typically finite) discrete set. A SHS characterizes a jump process $\mathbf{q} : [0,\infty) \to \mathscr{Q}$ called the *discrete state*; a stochastic process $\mathbf{x} : [0,\infty) \to \mathbb{R}^n$ with piecewise continuous sample paths called the *continuous state*; and m stochastic counters $\mathbf{N}_\ell : [0,\infty) \to \mathbb{N}_{\geq 0}$ called the *transition counters*.

In essence, between transition counter increments the discrete state remains constant whereas the continuous state flows according to (1). At transition times, the continuous and discrete states are reset according to (2). Each transition counter $\mathbf{N}_\ell$ counts the number of times that the corresponding discrete transition/reset map ϕ_ℓ is "activated." The frequency at which this occurs is determined by the transition intensities (3). In particular, the probability that the counter $\mathbf{N}_\ell$ will increment in an "elementary interval" $(t,t+dt]$, and therefore that the corresponding transition takes place, is given by $\lambda_\ell(\mathbf{q}(t),\mathbf{x}(t),t)dt$. In practice, one can think of the intensity of a transition as the instantaneous rate at which that transition occurs. The reader is referred to [1] for a mathematically precise characterization of this SHS.

The following result can be used to compute expectations on the state of a SHS. For briefness, we omit a few technical assumptions that are straightforward to verify for the SHSs considered here:

Theorem 1 ([1]). *Given a function $\psi : \mathscr{Q} \times \mathbb{R}^n \times [0,\infty) \to \mathbb{R}$ that is twice continuously differentiable with respect to its second argument and once continuously differentiable with respect to the third one, we have that*

$$\frac{\partial \, \mathrm{E}[\psi(\mathbf{q}(t),\mathbf{x}(t),t)]}{\partial t} = \mathrm{E}[(L\psi)(\mathbf{q}(t),\mathbf{x}(t),t)], \tag{4}$$

where $\forall (q,x,t) \in \mathscr{Q} \times \mathbb{R}^n \times [0,\infty)$

$$\begin{aligned}(L\psi)(q,x,t) := \frac{\partial \psi(q,x,t)}{\partial x} f(q,x,t) + \frac{\partial \psi(q,x,t)}{\partial t} + \\ + \frac{1}{2}\,\mathrm{trace}\Big(\frac{\partial^2 \psi(q,x)}{\partial x^2} g(q,x,t) g(q,x,t)'\Big) + \\ + \sum_{\ell=1}^{m} \Big(\psi\big(\phi_\ell(q,x,t),t\big) - \psi(q,x,t)\Big)\lambda_\ell(q,x,t),\end{aligned} \tag{5}$$

and $\frac{\partial \psi(q,x,t)}{\partial t}$, $\frac{\partial \psi(q,x,t)}{\partial x}$, *and* $\frac{\partial^2 \psi(q,x)}{\partial x^2}$ *denote the partial derivative of* $\psi(q,x,t)$ *with respect to* t, *the gradient of* $\psi(q,x,t)$ *with respect to* x, *and the Hessian matrix of* ψ *with respect to* x, *respectively. The operator* $\psi \mapsto L\psi$ *defined by* (5) *is called the* extended generator *of the SHS.* □

We say that a SHS is *polynomial (pSHS)* if its extended generator L is closed on the set of finite-polynomials in x, i.e., $(L\psi)(q,x,t)$ is a finite-polynomial in x for every finite-polynomial $\psi(q,x,t)$ in x. By a *finite-polynomials in x* we mean a function $\psi(q,x,t)$ such that $x \mapsto \psi(q,x,t)$ is a (multi-variable) polynomial of finite degree for each fixed $q \in \mathscr{Q}$, $t \in [0,\infty)$. A pSHS is obtained, e.g., when the vector fields f and g, the reset maps ϕ_ℓ, and the transition intensities λ_ℓ are all finite-polynomials in x.

Examples of Polynomial Stochastic Hybrid Systems

Example 1 (Random walk). Let $\mathbf{x} \in \mathbb{R}$ denote the velocity of a moving particle that is instantaneously increased or decrease by a fixed amount $\sqrt{\varepsilon}$ every so often. The interval of time between increases (as well as between decreases) is exponentially distributed with mean 2ε. The process $\mathbf{x}$ can be generated by a SHS with continuous dynamics $\dot{\mathbf{x}} = -a\mathbf{x}$ and two reset maps $\mathbf{x} \mapsto \phi_{1,2}(\mathbf{x}) := \mathbf{x} \pm b\sqrt{\varepsilon}$ both with constant transitions intensities $\lambda_{1,2}(\mathbf{x}) := \frac{1}{2\varepsilon}$. The constant $a \geq 0$ accounts for viscous friction. This SHS has a single discrete mode that we omitted for simplicity. The generator for this SHS is given by

$$(L\psi)(x) = -ax\frac{\partial \psi(x)}{\partial x} + \frac{\psi(x + b\sqrt{\varepsilon}) + \psi(x - b\sqrt{\varepsilon}) - 2\psi(x)}{2\varepsilon},$$

which is closed on the set of finite-polynomials in x. It is well known that as $\varepsilon \to 0$, $\mathbf{x}$ approaches the solution to the SDE $\dot{\mathbf{x}} = a\mathbf{x} + b\dot{\mathbf{n}}$, where $\mathbf{n}$ is a Brownian motion process. This SDE is easy to solve analytically, at least when the initial distribution of $\mathbf{x}$ is Gaussian. Our (somewhat academic) goal here would be to study what happens when ε is not necessarily close to zero and for an arbitrary initial distribution. □

Example 2 (TCP long-lived [14]). The congestion window size $\mathbf{w} \in [0,\infty)$ of a long-lived TCP flow can be generated by a SHS with continuous dynamics $\dot{\mathbf{w}} = \frac{1}{RTT}$ and a reset map $\mathbf{w} \mapsto \frac{\mathbf{w}}{2}$, with intensity $\lambda(\mathbf{w}) := \frac{p\mathbf{w}}{RTT}$. The round-trip-time RTT and the drop-rate p are parameters that we assume constant. This SHS has a single discrete mode that we omitted for simplicity. The generator for this SHS is given by

$$(L\psi)(w) = \frac{1}{RTT}\frac{\partial\psi(w)}{\partial w} + \frac{pw\big(\psi(w/2)-\psi(w)\big)}{RTT},$$

which is closed on the set of finite-polynomials in w. □

Example 3 (TCP on-off [14]). The congestion window size $\mathbf{w} \in [0,\infty)$ for a stream of TCP flows separated by inactivity periods can be generated by a SHS with three discrete modes $\mathscr{Q} := \{\mathrm{ss},\mathrm{ca},\mathrm{off}\}$, one corresponding to slow-start, another to congestion avoidance, and the final one to flow inactivity. Its continuous dynamics are defined by

$$\dot{\mathbf{w}} = \begin{cases} \frac{(\log 2)\mathbf{w}}{RTT} & \mathbf{q} = \mathrm{ss} \\ \frac{1}{RTT} & \mathbf{q} = \mathrm{ca} \\ 0 & \mathbf{q} = \mathrm{off}; \end{cases}$$

the reset maps associated with packet drops, end of flows, and start of flows are given by $(\mathbf{q},\mathbf{w}) \mapsto \phi_1(\mathbf{q},\mathbf{w}) := \big(\mathrm{ca},\frac{\mathbf{w}}{2}\big)$, $(\mathbf{q},\mathbf{w}) \mapsto \phi_2(\mathbf{q},\mathbf{w}) := (\mathrm{off},0)$, and $(\mathbf{q},\mathbf{w}) \mapsto \phi_3(\mathbf{q},\mathbf{w}) := (\mathrm{ss},w_0)$, respectively; and the corresponding intensities are

$$\lambda_1(\mathbf{q},\mathbf{w}) := \begin{cases} \frac{p\mathbf{w}}{RTT} & \mathbf{q} \in \{\mathrm{ss},\mathrm{ca}\} \\ 0 & \mathbf{q} = \mathrm{off} \end{cases} \qquad \lambda_2(\mathbf{q},\mathbf{w}) := \begin{cases} \frac{\mathbf{w}}{kRTT} & \mathbf{q} \in \{\mathrm{ss},\mathrm{ca}\} \\ 0 & \mathbf{q} = \mathrm{off} \end{cases}$$

$$\lambda_3(\mathbf{q},\mathbf{w}) := \begin{cases} \frac{1}{\tau_{\mathrm{off}}} & \mathbf{q} = \mathrm{off} \\ 0 & \mathbf{q} \in \{\mathrm{ss},\mathrm{ca}\}. \end{cases}$$

The round-trip-time RTT, the drop-rate p, the average file size k (exponentially distributed), the average off-time τ_{off} (also exponentially distributed), and the initial window size w_0 are parameters that we assume constant. The generator for this SHS is given by

$$(L\psi)(q,w) = \begin{cases} \frac{(\log 2)w}{RTT}\frac{\partial\psi(\mathrm{ss},w)}{\partial w} + \frac{pw\left(\psi(\mathrm{ca},w/2)-\psi(\mathrm{ss},w)\right)}{RTT} + \frac{w\left(\psi(\mathrm{off},0)-\psi(\mathrm{ss},w)\right)}{kRTT} & q = \mathrm{ss} \\ \frac{1}{RTT}\frac{\partial\psi(\mathrm{ca},w)}{\partial w} + \frac{pw\left(\psi(\mathrm{ca},w/2)-\psi(\mathrm{ca},w)\right)}{RTT} + \frac{w\left(\psi(\mathrm{off},0)-\psi(\mathrm{ca},w)\right)}{kRTT} & q = \mathrm{ca} \\ \frac{\psi(\mathrm{ss},w_0)-\psi(\mathrm{off},w)}{\tau_{\mathrm{off}}} & q = \mathrm{off}, \end{cases}$$

which is closed on the set of finite-polynomials in w. □

Example 4 (Networked control system [9]). Suppose that the state of a stochastic scalar linear system $\dot{\mathbf{x}} = a\mathbf{x} + b\dot{\mathbf{n}}$ is estimated based on state-measurements received through a network. For simplicity we assume that state measurements are noiseless and delay free. The corresponding state estimation error $\mathbf{e} \in \mathbb{R}$ can be generated by a SHS with continuous dynamics $\dot{\mathbf{e}} = a\mathbf{e} + b\dot{\mathbf{n}}$ and one reset map $\mathbf{e} \mapsto 0$ that is activated whenever a state measurement is received. It was conjectured in [9] and later shown in [10] that

transmitting measurements at a rate that depends on the state-estimation error is optimal if one wants to minimize the variance of the estimation error, while penalizing the average rate at which messages are transmitted. This motivates considering the following reset intensity $\lambda(\mathbf{e}) := \mathbf{e}^{2\rho}$, $\rho \in \mathbb{N}_{\geq 0}$. This SHS has a single discrete mode that we omitted for simplicity and its generator is given by

$$(L\psi)(e) := a e \frac{\partial \psi(e)}{\partial e} + \frac{b^2}{2} \frac{\partial^2 \psi(e)}{\partial e^2} + \big(\psi(0) - \psi(e)\big) e^{2\rho},$$

which is closed on the set of finite-polynomials in e. □

Example 5 (Decaying-dimerizing reaction set [12,13]). The number of particles $\mathbf{x} := (\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3)$ of three species involved in the following set of decaying-dimerizing reactions

$$S_1 \xrightarrow{c_1} 0, \qquad 2S_1 \xrightarrow{c_2} S_2, \qquad S_2 \xrightarrow{c_3} 2S_1, \qquad S_2 \xrightarrow{c_4} S_3 \tag{6}$$

can be generated by a SHS with continuous dynamics $\dot{\mathbf{x}} = 0$ and four reset maps

$$\mathbf{x} \mapsto \phi_1(\mathbf{x}) := \begin{bmatrix} \mathbf{x}_1 - 1 \\ \mathbf{x}_2 \\ \mathbf{x}_3 \end{bmatrix} \qquad \mathbf{x} \mapsto \phi_2(\mathbf{x}) := \begin{bmatrix} \mathbf{x}_1 - 2 \\ \mathbf{x}_2 + 1 \\ \mathbf{x}_3 \end{bmatrix}$$
$$\mathbf{x} \mapsto \phi_3(\mathbf{x}) := \begin{bmatrix} \mathbf{x}_1 + 2 \\ \mathbf{x}_2 - 1 \\ \mathbf{x}_3 \end{bmatrix} \qquad \mathbf{x} \mapsto \phi_4(\mathbf{x}) := \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 - 1 \\ \mathbf{x}_3 + 1 \end{bmatrix}$$

with intensities $\lambda_1(\mathbf{x}) := c_1 \mathbf{x}_1$, $\lambda_2(\mathbf{x}) := \frac{c_2}{2} \mathbf{x}_1(\mathbf{x}_1 - 1)$, $\lambda_3(\mathbf{x}) := c_3 \mathbf{x}_2$, and $\lambda_4(\mathbf{x}) := c_4 \mathbf{x}_2$, respectively. Since the numbers of particles take values in the discrete set of integers, we can regard the $\mathbf{x}_i$ as either part of the discrete or continuous state. We choose to regard them as continuous variables because we are interested in studying their statistical moments. In this case, the SHS has a single discrete mode that we omit for simplicity. The generator for this SHS is given by

$$\begin{aligned}(L\psi)(x) = {} & c_1 x_1 \big(\psi(x_1 - 1, x_2, x_3) - \psi(x)\big) + \frac{c_2}{2} x_1(x_1 - 1)\big(\psi(x_1 - 2, x_2 + 1, x_3) - \psi(x)\big) \\ & + c_3 x_2 \big(\psi(x_1 + 2, x_2 - 1, x_3) - \psi(x)\big) + c_4 x_2 \big(\psi(x_1, x_2 - 1, x_3 + 1) - \psi(x)\big),\end{aligned}$$

which is closed on the set of finite-polynomials in x. □

3 Moment Dynamics

To fully characterize the dynamics of a SHS one would like to determine the evolution of the probability distribution for its state $(\mathbf{q}, \mathbf{x})$. In general, this is difficult so a more reasonable goal is to determine the evolution of (i) the probability of $\mathbf{q}(t)$ being on each mode and (ii) the moments of $\mathbf{x}(t)$ conditioned to $\mathbf{q}(t)$. In fact, often one can even get away with only determining a few low-order moments and then using results such as Tchebycheff, Markoff, or Bienaymé inequalities [6, pp. 114–115] to infer properties of the overall distribution.

Given a discrete state $\bar{q} \in \mathscr{Q}$ and a vector of n integers $m = (m_1, m_2, \ldots, m_n) \in \mathbb{N}^n_{\geq 0}$, we define the *test-function associated with $\bar{q}$ and m* to be

$$\psi^{(m)}_{\bar{q}}(q,x) := \begin{cases} x^{(m)} & q = \bar{q} \\ 0 & q \neq \bar{q}, \end{cases} \qquad \forall q \in \mathscr{Q}, x \in \mathbb{R}^n$$

and the *(uncentered) moment associated with $\bar{q}$ and m* to be

$$\mu^{(m)}_{\bar{q}}(t) := \mathrm{E}\left[\psi^{(m)}_{\bar{q}}\big(\mathbf{q}(t), \mathbf{x}(t)\big)\right] \qquad \forall t \geq 0. \tag{7}$$

Here and in the sequel, given a vector $x = (x_1, x_2, \ldots, x_n)$, we use $x^{(m)}$ to denote the monomial $x_1^{m_1} x_2^{m_2} \cdots x_n^{m_n}$.

PSHSs have the property that if one stacks all moments in an infinite vector μ_∞, its dynamics can be written as

$$\dot{\mu}_\infty = A_\infty(t)\mu_\infty \qquad \forall t \geq 0, \tag{8}$$

for some appropriately defined infinite matrix $A_\infty(t)$. This is because $\forall \bar{q} \in \mathscr{Q}, m = (m_1, \ldots, m_n) \in \mathbb{N}^n_{\geq 0}$, the expression $(L\psi^{(m)}_{\bar{q}})(q,x,t)$ is a finite-polynomial in x and therefore can be written as a finite linear combination of test-functions (possibly with time-varying coefficients). Taking expectations on this linear combination and using (4), (7), we conclude that $\dot{\mu}^{(m)}_{\bar{q}}$ can be written as linear combination of uncentered moments in μ_∞, leading to (8). In the sequel, we refer to (8) as the *infinite-dimensional moment dynamics*. Analyzing (and even simulating) (8) is generally difficult. However, as mentioned above one can often get away with just computing a few low-order moments. One would therefore like to determine a finite-dimensional system of ODEs that describes the evolution of a few low-order models, perhaps only approximately.

When the matrix A_∞ is lower triangular (e.g., as in Example 1 and Example 4 with $\rho = 0$), one can simply truncate the vector μ_∞ by dropping all but its first k elements and obtain a finite-dimensional system that exactly describes the evolution of the moments[1]. However, in general A_∞ has nonzero elements above the main diagonal and therefore if one defines $\mu \in \mathbb{R}^k$ to contain the top k elements of μ_∞, one obtains from (8) that

$$\dot{\mu} = I_{k\times\infty} A_\infty(t)\mu_\infty = A(t)\mu + B(t)\bar{\mu}, \qquad \bar{\mu} = C\mu_\infty, \tag{9}$$

where $I_{k\times\infty}$ denotes a matrix composed of the first k rows of the infinite identity matrix, $\bar{\mu} \in \mathbb{R}^r$ contains all the moments that appear in the first k elements of $A_\infty(t)\mu_\infty$ but that do not appear in μ, and C is the projection matrix that extracts $\bar{\mu}$ from μ_∞. Our goal is to approximate the infinite dimensional system (8) by a finite-dimensional nonlinear ODE of the form

$$\dot{\nu} = A(t)\nu + B(t)\bar{\nu}(t), \quad \bar{\nu} = \varphi(\nu, t), \tag{10}$$

where the map $\varphi : \mathbb{R}^k \times [0,\infty) \to \mathbb{R}^r$ should be chosen so as to keep $\nu(t)$ close to $\mu(t)$. We call (10) the *truncated moment dynamics* and φ the *truncation function*. We need the following two stability assumptions to establish sufficient conditions on φ for the approximation to be valid.

[1] Other truncations are possible, but it is often convenient to keep the low-order moments of $\mathbf{x}$ as the states of the truncated system because these are physically meaningful.

Assumption 1 (Boundedness). *There exist sets Ω_μ and Ω_ν such that all solutions to* (8) *and* (10) *starting at some time $t_0 \geq 0$ in Ω_μ and Ω_ν, respectively, exist and are smooth on $[t_0,\infty)$ with all derivatives of their first k elements uniformly bounded by the same constant. The set Ω_ν is assumed to be forward invariant.* □

Assumption 2 (Incremental Stability). *There exists a function*[2] *$\beta \in \mathscr{KL}$ such that, for every solution μ_∞ to* (8) *starting in Ω_μ at some time $t_0 \geq 0$, and every $t_1 \geq t_0$, $\nu_1 \in \Omega_\nu$ there exists some $\hat{\mu}_\infty(t_1) \in \Omega_\mu$ whose first k elements match ν_1 and*

$$\|\mu(t) - \hat{\mu}(t)\| \leq \beta(\|\mu(t_1) - \hat{\mu}(t_1)\|, t - t_1), \qquad \forall t \geq t_1, \tag{11}$$

where $\mu(t)$ and $\hat{\mu}(t)$ denote the first k elements of the solutions to (8) *starting at $\mu_\infty(t_1)$ and $\hat{\mu}_\infty(t_1)$, respectively.* □

Remark 1. Assumption 2 is essentially a requirement of uniform stability for (8), but it was purposely formulated without requiring Ω_μ to be a subset of a normed space. This avoids having to choose a norm under which the (infinite) vectors of moments are bounded. Note, e.g., that the usual ℓ_p-norms, $p \in [1,\infty]$ would not work for most "interesting" distributions. It is also important to emphasize that a certain "richness" of Ω_μ is implicit in this assumption. Indeed, for the existence of $\hat{\mu}_\infty(t_1) \in \Omega_\mu$ whose first components match an arbitrary $\nu_1 \in \Omega_\nu$ one needs Ω_ν to be contained in the projection of Ω_μ into its first k components. □

The result that follows establishes that the difference between solutions to (8) and (10) converges to an arbitrarily small ball, provided that a sufficiently large but *finite* number of derivatives of these signals match point-wise. To state this result, the following notation is needed: We define the matrices $C^i(t)$, $i \in \mathbb{N}_{\geq 0}$ recursively by

$$C^0(t) = C, \qquad C^{i+1}(t) = C^i(t)A_\infty(t) + \dot{C}^i(t), \qquad \forall t \geq 0,\ i \in \mathbb{N}_{\geq 0},$$

and the family of functions $L^i\varphi : \mathbb{R}^k \times [0,\infty) \to \mathbb{R}^r$, $i \in \mathbb{N}_{\geq 0}$ recursively by

$$(L^0\varphi)(\nu,t) = \varphi(\nu,t), \quad (L^{i+1}\varphi)(\nu,t) = \frac{\partial(L^i\varphi)(\nu,t)}{\partial\nu}\big(A(t)\nu + B(t)\varphi(\nu,t)\big) + \frac{\partial(L^i\varphi)(\nu,t)}{\partial t},$$

$\forall t \geq 0$, $\nu \in \mathbb{R}^k$, $i \in \mathbb{N}_{\geq 0}$. These definitions allow us to compute time derivatives of $\bar{\mu}(\tau)$ and $\bar{\nu}(\tau)$ along solutions to (8) and (10), respectively, because

$$\frac{\mathrm{d}^i\bar{\mu}(t)}{\mathrm{d}t^i} = C^i(t)\mu_\infty(t), \qquad \frac{\mathrm{d}^i\bar{\nu}(t)}{\mathrm{d}t^i} = (L^i\varphi)(\nu(t),t), \qquad \forall t \geq 0,\ i \in \mathbb{N}_{\geq 0}. \tag{12}$$

Theorem 2 (cf. Appendix). *For every $\delta > 0$, there exists an integer N sufficiently large for which the following result holds: Assuming that for every $\tau \geq 0$, $\mu_\infty \in \Omega_\mu$*

$$C^i(\tau)\mu_\infty = (L^i\varphi)(\mu,\tau), \qquad \forall i \in \{0,1,\ldots,N\}, \tag{13}$$

[2] A function $\beta : [0,\infty)\times[0,\infty) \to [0,\infty)$ is said to be of *class* $\mathscr{KL}$ if $\beta(0,t) = 0$, $\forall t \geq 0$; $\beta(s,t)$ is continuous and strictly increasing on s for each fixed $t \geq 0$; and $\lim_{t\to\infty}\beta(s,t) = 0$, $\forall s \geq 0$.

where μ denotes the first k elements of μ_∞, then

$$\|\mu(t) - \nu(t)\| \leq \beta(\|\mu(t_0) - \nu(t_0)\|, t - t_0) + \delta, \qquad \forall t \geq t_0 \geq 0, \tag{14}$$

along all solutions to (8) *and* (10) *with initial conditions $\mu_\infty(t_0) \in \Omega_\mu$ and $\nu(t_0) \in \Omega_\nu$, respectively, where $\mu(t)$ denotes the first k elements of $\mu_\infty(t)$.* □

The proof of this theorem is technical and therefore we leave it for the appendix. The idea behind it is that the conditions (13) guarantees that when the solution μ_∞ to the infinite-dimensional system (8) starts in Ω_μ, the $N+1$ derivatives of its first k elements μ match those of the solution ν to the truncated system (10). Using a standard argument based on a Taylor series expansion, we thus conclude that these two solutions will remain close in a bounded interval of length T. To extend this to the unbounded interval $[0,\infty)$, we use the stability of the infinite-dimensional (8): After the first interval $[0,T]$, we can start another solution $\hat{\mu}_\infty$ to (8) on Ω_μ, whose first k elements $\hat{\mu}$ match ν at time T, and conclude from (13) that $\hat{\mu}$ and ν will remain close on $[T,2T]$. But because of stability μ and $\hat{\mu}$ will converge to each other and therefore μ and ν also remain close. This argument can be repeated to obtain the bound (14) on the unbounded time interval $[0,\infty)$. A similar argument was used in the proof of [7, Theorem 1] to analyze two time scale systems. It is important to emphasize that for this argument to hold, Ω_μ need not be forward invariant because we only use (13) for the solutions $\hat{\mu}$ that are always "re-started" in Ω_μ. This explains why one may get good matches between the original and the truncated systems, even if (13) only holds for fairly small subsets Ω_μ of the overall infinite-dimensional state-space.

4 Construction of Approximate Truncations

Given a constant $\delta > 0$ and sets Ω_μ, Ω_ν, it may be very difficult to determine the integer N for which the approximation bound (14) holds. This is because, although the proof of Theorem 2 is constructive, the computation of N requires explicit knowledge of the function $\beta \in \mathscr{KL}$ in Assumption 2 and, at least for most of the examples considered here, this assumption is difficult to verify. Nevertheless, Theorem 2 is still useful because it provides the explicit conditions (13) that the truncation function φ should satisfy for the solution to the truncated system to approximate the one of the original system. For the problems considered here we require (13) to hold for $N = 1$, which corresponds to a second Taylor expansions of the solution. Note that for $N = 1$, (13) simply becomes

$$C\mu_\infty = \varphi(\mu,\tau), \qquad CA_\infty(\tau)\mu_\infty = \frac{\partial\varphi(\mu,\tau)}{\partial\mu} I_{k\times\infty} A_\infty(\tau)\mu_\infty + \frac{\partial\varphi(\mu,\tau)}{\partial t}, \tag{15}$$

$\forall \mu_\infty \in \Omega_\mu,\ \tau \geq 0$. Lacking knowledge of β, we will not be able to explicitly compute for which values of δ (14) will hold, but we will show by simulation that the truncation obtained provides a very accurate approximation to the infinite-dimensional system (8), even for such a small choice of N. We restrict our attention to functions φ and sets Ω_μ for which it is simple to use (15) to explicitly compute truncated systems.

Separable truncation functions: For all the examples considered, we consider functions φ of the form

$$\varphi(\nu,t) = \Lambda\, \nu^{(\Gamma)} := \Lambda \begin{bmatrix} \nu_1^{\gamma_{11}} \nu_2^{\gamma_{12}} \cdots \nu_k^{\gamma_{1k}} \\ \nu_1^{\gamma_{21}} \nu_2^{\gamma_{22}} \cdots \nu_k^{\gamma_{2k}} \\ \vdots \\ \nu_1^{\gamma_{r1}} \nu_2^{\gamma_{r2}} \cdots \nu_k^{\gamma_{rk}} \end{bmatrix}, \tag{16}$$

for appropriately chosen constant matrices $\Gamma := [\gamma_{ij}] \in \mathbb{R}^{r\times k}$ and $\Lambda \in \mathbb{R}^{r\times r}$, with Λ diagonal. In this case,

$$\begin{aligned} \frac{\partial \varphi(\nu)}{\partial \nu} &= \Lambda \begin{bmatrix} \nu_1^{\gamma_{11}} \nu_2^{\gamma_{12}} \cdots \nu_k^{\gamma_{1k}} \gamma_{11} \nu_1^{-1} & \nu_1^{\gamma_{11}} \nu_2^{\gamma_{12}} \cdots \nu_k^{\gamma_{1k}} \gamma_{12} \nu_2^{-1} & \cdots & \nu_1^{\gamma_{11}} \nu_2^{\gamma_{12}} \cdots \nu_k^{\gamma_{1k}} \gamma_{1k} \nu_k^{-1} \\ \nu_1^{\gamma_{21}} \nu_2^{\gamma_{22}} \cdots \nu_k^{\gamma_{2k}} \gamma_{21} \nu_1^{-1} & \nu_1^{\gamma_{21}} \nu_2^{\gamma_{22}} \cdots \nu_k^{\gamma_{2k}} \gamma_{22} \nu_2^{-1} & \cdots & \nu_1^{\gamma_{21}} \nu_2^{\gamma_{22}} \cdots \nu_k^{\gamma_{2k}} \gamma_{2k} \nu_k^{-1} \\ \vdots & \vdots & \ddots & \vdots \\ \nu_1^{\gamma_{r1}} \nu_2^{\gamma_{r2}} \cdots \nu_k^{\gamma_{rk}} \gamma_{r1} \nu_1^{-1} & \nu_1^{\gamma_{r1}} \nu_2^{\gamma_{r2}} \cdots \nu_k^{\gamma_{rk}} \gamma_{r2} \nu_2^{-1} & \cdots & \nu_1^{\gamma_{r1}} \nu_2^{\gamma_{r2}} \cdots \nu_k^{\gamma_{rk}} \gamma_{rk} \nu_k^{-1} \end{bmatrix} \\ &= \Lambda \operatorname{diag}[\nu^{(\Gamma)}] \Gamma \operatorname{diag}[\nu_1^{-1}, \nu_2^{-1}, \ldots, \nu_k^{-1}] \end{aligned}$$

and therefore (15) becomes

$$C\mu_\infty = \Lambda\, \mu^{(\Gamma)}, \tag{17a}$$

$$CA_\infty(\tau)\mu_\infty = \Lambda \operatorname{diag}[C\mu_\infty] \Gamma \operatorname{diag}[\mu_1^{-1}, \mu_2^{-1}, \ldots, \mu_k^{-1}] I_{k\times\infty} A_\infty(\tau)\mu_\infty. \tag{17b}$$

This particular form for φ is convenient because (17a) often uniquely defines Λ and then (17b) provides a linear system of equations on Γ, for which it is straightforward to determine if a solution exists.

Deterministic distributions: A set Ω_μ that is particularly tractable corresponds to deterministic distributions $\mathscr{F}_{\text{det}} := \{P(\cdot\,; q, x) : x \in \Omega_x, q \in \mathscr{Q}\}$, where $P(\cdot\,; q, x)$ denotes the distribution of $(\mathbf{q}, \mathbf{x})$ for which $\mathbf{q} = q$ and $\mathbf{x} = x$ with probability one; and Ω_x a subset of the continuous state space $\mathbb{R}^n$. For a particular distribution $P(\cdot\,; q, x)$, the (uncentered) moment associated with $\bar{q}$ and $m \in \mathbb{N}^n_{\geq 0}$ is given by

$$\mu_{\bar{q}}^{(m)} := \int \psi_{\bar{q}}^{(m)}(\tilde{q}, \tilde{x}) P(d\tilde{q}\, d\tilde{x}; q, x) = \psi_{\bar{q}}^{(m)}(q, x) := \begin{cases} x^{(m)} & q = \bar{q} \\ 0 & q \neq \bar{q}, \end{cases}$$

and therefore the vectors μ_∞ in Ω_μ have this form. Although this family of distributions may seem very restrictive, it will provide us with truncations that are accurate even when the pSHSs evolve towards very "nondeterministic" distributions, i.e., with significant variance. For this set Ω_μ, (17) takes a particularly simple form and the following result provides a simple set of conditions to test if a truncation is possible.

Lemma 1 (cf. Appendix). *Let Ω_μ be the set of deterministic distributions $\mathscr{F}_{\text{det}}$ with Ω_x containing some open ball in $\mathbb{R}^n$ and consider truncation functions φ of the form* (16)*. The matrices Λ and Γ must satisfy:*

1. *Λ must be the identity matrix.*

2. *For every moment* $\mu_{q_\ell}^{(m_\ell)}$ *in* $\bar{\mu}$ *and every moment* $\mu_{q_i}^{(m_i)}$ *in* μ *for which* $q_i \neq q_\ell$*, one must have* $\gamma_{\ell,i} = 0$*.*
3. *For every moment* $\mu_{q_\ell}^{(m_\ell)}$ *in* $\bar{\mu}$*, one must have* $\sum_{\substack{i=1 \\ q_i=q_\ell}}^{k} \gamma_{\ell,i} m_i = m_\ell$*.*

Moreover, the following conditions are necessary for the existence of a function φ *of the form* (16) *that satisfies* (15)*:*

4. *For every moment* $\mu_{q_\ell}^{(m_\ell)}$ *in* $\bar{\mu}$ *the polynomial*[3] $\sum_{\substack{i=1 \\ q_i=q_\ell}}^{\infty} a_{\ell,i} x^{(m_i)}$ *must belong to the linear subspace generated by the polynomials*

$$\Big\{ \sum_{\substack{i=1 \\ q_i=q_\ell}}^{\infty} a_{j,i} x^{(m_\ell - m_j + m_i)} : 1 \leq j \leq k,\ q_j = q_\ell \Big\}.$$

5. *For every moment* $\mu_{q_\ell}^{(m_\ell)}$ *in* $\bar{\mu}$ *and every moment* $\mu_{q_i}^{(m_i)}$ *in* μ_∞ *with* $q_i \neq q_\ell$*, we must have* $a_{\ell,i} = 0$*. Here we are denoting by* $a_{j,i}$ *the* j*th row,* i*th column entry of* A_∞*.* □

Condition 4 imposes a diagonal-band-like structure on the submatrices of A_∞ consisting of the rows/columns that correspond to each moment that appears in $\bar{\mu}$. This condition holds for Examples 2, 3, and 4, but not for Example 5. However, we will see that the moment dynamics of this example can be simplified so as to satisfy this condition without introducing a significant error.

Condition 5 imposes a form of decoupling between different modes in the equations for $\dot{\bar{\mu}}$. This condition holds trivially for all examples that have a single discrete mode. It does not hold for Example 3, but also here it is possible to simplify the moment dynamics to satisfy this condition without introducing a significant error.

5 Examples of Truncations

We now present truncated systems for the several examples considered before and discuss how the truncated models compare to estimates of the moments obtained from Monte Carlo simulations. All Monte Carlo simulations were carried out using the algorithm described in [1]. Estimates of the moments were obtained by averaging a large number of Monte Carlo simulations. In most plots, we used a sufficiently large number of simulations so that the 99% confidence intervals for the mean cannot be distinguished from the point estimates at the resolution used for the plots. Ir is worth to emphasize that the results obtained through Monte Carlo simulations required computational efforts orders of magnitude higher than those obtained using the truncated systems.

Example 1 (Random walk). For this system $\psi^{(m)}(x) = x^m$, $\forall m \in \mathbb{N}_{\geq 0}$ and we conclude that

$$(L\psi^{(m)})(x) = -ma x^m + \frac{(x + b\sqrt{\varepsilon})^m + (x - b\sqrt{\varepsilon})^m - 2x^m}{2\varepsilon}$$

[3] We are considering polynomials with integer (both positive and negative) powers.

$$= -max^m + \sum_{\substack{i=2 \\ i \text{ even}}}^{m} \binom{m}{i} b^i \varepsilon^{\frac{i-2}{2}} \psi^{(m-i)}(x).$$

Therefore

$$\begin{bmatrix} \dot{\mu}^{(1)} \\ \dot{\mu}^{(2)} \\ \dot{\mu}^{(3)} \\ \dot{\mu}^{(4)} \\ \dot{\mu}^{(5)} \\ \vdots \end{bmatrix} = \begin{bmatrix} 0 \\ b \\ 0 \\ b^2\varepsilon \\ 0 \\ \vdots \end{bmatrix} + \begin{bmatrix} -a & 0 & \cdots & & & & \\ 0 & -2a & 0 & \cdots & & & \\ 3b & 0 & -3a & 0 & \cdots & & \\ 0 & 6b & 0 & -4a & 0 & \cdots & \\ 5b^2\varepsilon & 0 & 10b & 0 & -5a & 0 & \cdots \\ \vdots & \vdots & \ddots & \ddots & \ddots & \ddots & \ddots \end{bmatrix} \begin{bmatrix} \mu^{(1)} \\ \mu^{(2)} \\ \mu^{(3)} \\ \mu^{(4)} \\ \mu^{(5)} \\ \vdots \end{bmatrix}$$

Since $\mu^{(0)}(t) = 1$, $\forall t \geq 0$ we excluded $\mu^{(0)}$ from the vector μ, resulting in an affine system (as opposed to linear).

Because the infinite-dimensional moment dynamics are lower-triangular, this system can be truncated exactly. For example, a third order truncation yields

$$\begin{bmatrix} \dot{\mu}^{(1)} \\ \dot{\mu}^{(2)} \\ \dot{\mu}^{(3)} \end{bmatrix} = \begin{bmatrix} -a & 0 & 0 \\ 0 & -2a & 0 \\ 3b & 0 & -3a \end{bmatrix} \begin{bmatrix} \mu^{(1)} \\ \mu^{(2)} \\ \mu^{(3)} \end{bmatrix} + \begin{bmatrix} 0 \\ b \\ 0 \end{bmatrix}$$

which can be solved explicitly:

$$\mu^{(1)}(t) = e^{-at}\mu^{(1)}(0), \quad \mu^{(2)}(t) = \frac{b}{2a}(1 - e^{-2at}) + e^{-2at}\mu^{(2)}(0),$$
$$\mu^{(3)}(t) = \frac{3b}{2a}(e^{-at} - e^{-3at})\mu^{(1)}(0) + e^{-3at}\mu^{(3)}(0),$$

for $a \neq 0$, and

$$\mu^{(1)}(t) = \mu^{(1)}(0), \quad \mu^{(2)}(t) = bt + \mu^{(2)}(0), \quad \mu^{(3)}(t) = 3bt\mu^{(1)}(0) + \mu^{(3)}(0),$$

for $a = 0$. It is interesting to note that these equations do not depend on ε because this parameter only appears in the equation for $\dot{\mu}^{(4)}$. This means that up to the third order moments this pSHS is indistinguishable from the SDE $\dot{\mathbf{x}} = a\mathbf{x} + b\dot{\mathbf{n}}$. □

Example 2 (TCP long-lived). Since for this system it is particularly meaningful to consider moments of the packet sending rate $\mathbf{r} := \frac{\mathbf{w}}{RTT}$, we choose $\psi^{(m)}(w) = \frac{w^m}{RTT^m}$, $\forall m \in \mathbb{N}_{\geq 0}$ and conclude that

$$(L\psi^{(m)})(w) = \frac{mw^{m-1}}{RTT^{m+1}} - \frac{2^m - 1}{2^m}\frac{pw^{m+1}}{RTT^{m+1}} = \frac{m\psi^{(m-1)}(w)}{RTT^2} - p\frac{2^m - 1}{2^m}\psi^{(m+1)}(w).$$

Therefore

$$\begin{bmatrix} \dot{\mu}^{(1)} \\ \dot{\mu}^{(2)} \\ \dot{\mu}^{(3)} \\ \dot{\mu}^{(4)} \\ \vdots \end{bmatrix} = \begin{bmatrix} \frac{1}{RTT^2} \\ 0 \\ 0 \\ 0 \\ \vdots \end{bmatrix} + \begin{bmatrix} 0 & -\frac{p}{2} & 0 & \cdots & & & \\ \frac{2}{RTT^2} & 0 & -\frac{3p}{4} & 0 & \cdots & & \\ 0 & \frac{3}{RTT^2} & 0 & -\frac{7p}{8} & 0 & \cdots & \\ 0 & 0 & \frac{4}{RTT^2} & 0 & -\frac{15p}{16} & 0 & \cdots \\ \vdots & \vdots & \ddots & \ddots & \ddots & \ddots & \ddots \end{bmatrix} \begin{bmatrix} \mu^{(1)} \\ \mu^{(2)} \\ \mu^{(3)} \\ \mu^{(4)} \\ \mu^{(5)} \\ \vdots \end{bmatrix}$$

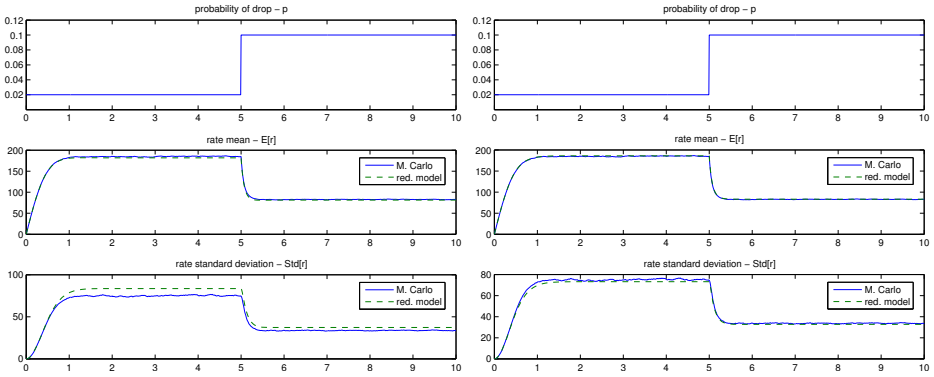

(a) Truncated model (18), (19) (b) Truncated model (18), (20)

Fig. 1. Comparison between Monte Carlo simulations and two truncated models for Example 2, with $RTT = 50$ms and a step in the drop-rate p from 2% to 10% at time $t = 5$sec.

Since $\mu^{(0)}(t) = 1$, $\forall t \geq 0$ we excluded $\mu^{(0)}$ from the vector μ, resulting in an affine system (as opposed to linear).

We consider a truncation whose state contains the first and second moments of the sending rate. In this case, (9) can be written as follows:

$$\begin{bmatrix} \dot{\mu}^{(1)} \\ \dot{\mu}^{(2)} \end{bmatrix} = \begin{bmatrix} 0 & -\frac{p}{2} \\ \frac{2}{RTT^2} & 0 \end{bmatrix} \begin{bmatrix} \mu^{(1)} \\ \mu^{(2)} \end{bmatrix} + \begin{bmatrix} \frac{1}{RTT^2} \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ -\frac{3p}{4} \end{bmatrix} \bar{\mu}, \tag{18}$$

where $\bar{\mu} := \mu^{(3)}$ evolves according to $\dot{\mu}^{(3)} = \frac{3\mu^{(2)}}{RTT^2} - \frac{7p\mu^{(4)}}{8}$. In this case, (17) has a unique solution φ, resulting in a truncated system given by (18) and

$$\bar{\mu} = \varphi(\mu^{(1)}, \mu^{(2)}) := \frac{(\mu^{(2)})^{\frac{5}{2}}}{(\mu^{(1)})^2}. \tag{19}$$

Figure 1(a) shows a comparison between Monte Carlo simulations and this truncated model. The dynamics of the first moment of the sending rate is very accurately predicted by the truncated model and there is roughly a 10% error in the steady-state value of the standard deviation. In the simulations, a step in the drop probability was introduced at time $t = 5$sec to show that the truncated model also captures well transient behavior. Figure 1(b) refers to a more accurate truncated model proposed in [1], which corresponds to

$$\bar{\mu} = \varphi(\mu^{(1)}, \mu^{(2)}) := \left(\frac{\mu^{(2)}}{\mu^{(1)}}\right)^3. \tag{20}$$

Its construction relied on the experimental observation that the steady-state distribution of the sending rate appears to be approximately Log-Normal. It turns out that this model

could also have been obtained using the procedure described in Sec. 3 by requiring (13) to hold only for $N = 0$ but for a set Ω_μ consisting of all Log-Normal distributions, which is richer than the set $\mathscr{F}_{\text{det}}$ of deterministic distributions used to construct (19). □

Example 3 (TCP on-off). For this system we also consider moments of the sending rate $\mathbf{r} := \frac{\mathbf{w}}{RTT}$ on the ss and ca modes, and therefore we use

$$\psi^{(0)}_{\text{off}}(q,w) = \begin{cases} 1 & q = \text{off} \\ 0 & q \in \{\text{ca},\text{ss}\} \end{cases} \qquad \psi^{(m)}_{\text{ss}}(q,w) = \begin{cases} \frac{w^m}{RTT^m} & q = \text{ss} \\ 0 & q \in \{\text{ca},\text{off}\} \end{cases}$$

$$\psi^{(m)}_{\text{ca}}(q,w) = \begin{cases} \frac{w^m}{RTT^m} & q = \text{ca} \\ 0 & q \in \{\text{ss},\text{off}\} \end{cases}$$

which yields

$$(L\psi^{(0)}_{\text{off}})(q,w) = -\tau^{-1}_{\text{off}}\psi^{(0)}_{\text{off}}(q,w) + \frac{1}{k}\psi^{(1)}_{\text{ss}}(q,w) + \frac{1}{k}\psi^{(1)}_{\text{ca}}(q,w)$$

$$(L\psi^{(m)}_{\text{ss}})(q,w) = \frac{\tau^{-1}_{\text{off}} w_0^m}{RTT^m}\psi^{(0)}_{\text{off}}(q,w) + \frac{m\log 2}{RTT}\psi^{(m)}_{\text{ss}}(q,w) - (p+\frac{1}{k})\psi^{(m+1)}_{\text{ss}}(q,w)$$

$$(L\psi^{(m)}_{\text{ca}})(q,w) = \frac{p}{2^m}\psi^{(m+1)}_{\text{ss}}(q,w) + \frac{m}{RTT^2}\psi^{(m-1)}_{\text{ca}}(q,w) - (\frac{1}{k}+p-\frac{p}{2^m})\psi^{(m+1)}_{\text{ca}}(q,w).$$

Therefore

$$\begin{bmatrix} \dot{\mu}^{(0)}_{\text{off}} \\ \dot{\mu}^{(0)}_{\text{ss}} \\ \dot{\mu}^{(0)}_{\text{ca}} \\ \dot{\mu}^{(1)}_{\text{ss}} \\ \dot{\mu}^{(1)}_{\text{ca}} \\ \dot{\mu}^{(2)}_{\text{ss}} \\ \dot{\mu}^{(2)}_{\text{ca}} \\ \dot{\mu}^{(3)}_{\text{ss}} \\ \dot{\mu}^{(3)}_{\text{ca}} \\ \vdots \end{bmatrix} = \begin{bmatrix} -\tau^{-1}_{\text{off}} & 0 & 0 & \frac{1}{k} & \frac{1}{k} & 0 & \cdots & & & & \\ \tau^{-1}_{\text{off}} & 0 & 0 & -\frac{1}{k}-p & 0 & 0 & \cdots & & & & \\ 0 & 0 & 0 & p & -\frac{1}{k} & 0 & \cdots & & & & \\ \frac{\tau^{-1}_{\text{off}} w_0}{RTT} & 0 & 0 & \frac{\log 2}{RTT} & 0 & -\frac{1}{k}-p & 0 & \cdots & & & \\ 0 & 0 & \frac{1}{RTT^2} & 0 & 0 & \frac{p}{2} & -\frac{1}{k}-\frac{p}{2} & 0 & \cdots & & \\ \frac{\tau^{-1}_{\text{off}} w_0^2}{RTT^2} & 0 & 0 & 0 & 0 & \frac{\log 4}{RTT} & 0 & -\frac{1}{k}-p & 0 & \cdots & \\ 0 & 0 & 0 & 0 & \frac{2}{RTT^2} & 0 & 0 & \frac{p}{4} & -\frac{1}{k}-\frac{3p}{4} & 0 & \cdots \\ \frac{\tau^{-1}_{\text{off}} w_0^3}{RTT^3} & 0 & 0 & 0 & 0 & 0 & 0 & \frac{\log 8}{RTT} & 0 & -\frac{1}{k}-p & 0 \cdots \\ 0 & 0 & 0 & 0 & 0 & 0 & \frac{3}{RTT^2} & 0 & 0 & \frac{p}{8} & -\frac{1}{k}-\frac{7p}{8} \ 0 \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \ddots & \ddots \end{bmatrix} \begin{bmatrix} \mu^{(0)}_{\text{off}} \\ \mu^{(0)}_{\text{ss}} \\ \mu^{(0)}_{\text{ca}} \\ \mu^{(1)}_{\text{ss}} \\ \mu^{(1)}_{\text{ca}} \\ \mu^{(2)}_{\text{ss}} \\ \mu^{(2)}_{\text{ca}} \\ \mu^{(3)}_{\text{ss}} \\ \mu^{(3)}_{\text{ca}} \\ \mu^{(4)}_{\text{ss}} \\ \mu^{(4)}_{\text{ca}} \\ \vdots \end{bmatrix}$$

Since $\mu^{(0)}_{\text{off}}$, $\mu^{(0)}_{\text{ss}}$, and $\mu^{(0)}_{\text{ca}}$ are the probabilities that $\mathbf{q}$ is equal to off, ss, and ca, respectively, these quantities must add up to one. We could therefore exclude one of them from μ.

We consider a truncation whose state contains the zeroth, first, and second moments of the sending rate. In this case, (9) can be written as follows:

$$\begin{bmatrix} \dot{\mu}^{(0)}_{\text{off}} \\ \dot{\mu}^{(0)}_{\text{ss}} \\ \dot{\mu}^{(0)}_{\text{ca}} \\ \dot{\mu}^{(1)}_{\text{ss}} \\ \dot{\mu}^{(1)}_{\text{ca}} \\ \dot{\mu}^{(2)}_{\text{ss}} \\ \dot{\mu}^{(2)}_{\text{ca}} \end{bmatrix} = \begin{bmatrix} -\tau^{-1}_{\text{off}} & 0 & 0 & \frac{1}{k} & \frac{1}{k} & 0 & 0 \\ \tau^{-1}_{\text{off}} & 0 & 0 & -\frac{1}{k}-p & 0 & 0 & 0 \\ 0 & 0 & 0 & p & -\frac{1}{k} & 0 & 0 \\ \frac{\tau^{-1}_{\text{off}} w_0}{RTT} & 0 & 0 & \frac{\log 2}{RTT} & 0 & -\frac{1}{k}-p & 0 \\ 0 & 0 & \frac{1}{RTT^2} & 0 & 0 & \frac{p}{2} & -\frac{1}{k}-\frac{p}{2} \\ \frac{\tau^{-1}_{\text{off}} w_0^2}{RTT^2} & 0 & 0 & 0 & 0 & \frac{\log 4}{RTT} & 0 \\ 0 & 0 & 0 & 0 & \frac{2}{RTT^2} & 0 & 0 \end{bmatrix} \begin{bmatrix} \mu^{(0)}_{\text{off}} \\ \mu^{(0)}_{\text{ss}} \\ \mu^{(0)}_{\text{ca}} \\ \mu^{(1)}_{\text{ss}} \\ \mu^{(1)}_{\text{ca}} \\ \mu^{(2)}_{\text{ss}} \\ \mu^{(2)}_{\text{ca}} \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ -\frac{1}{k}-p & 0 \\ \frac{p}{4} & -\frac{1}{k}-\frac{3p}{4} \end{bmatrix} \bar{\mu}, \quad (21)$$

where $\bar{\mu} := [\,\mu_{\rm ss}^{(3)}\ \mu_{\rm ca}^{(3)}\,]'$ evolves according to

$$\dot{\mu}_{\rm ss}^{(3)} = \frac{\tau_{\rm off}^{-1} w_0^3}{RTT^3}\mu_{\rm off}^{(0)} + \frac{\log 8}{RTT}\mu_{\rm ss}^{(3)} - (\frac{1}{k}+p)\mu_{\rm ss}^{(4)}, \tag{22a}$$

$$\dot{\mu}_{\rm ca}^{(3)} = \frac{3}{RTT^2}\mu_{\rm ca}^{(2)} + \frac{p}{8}\mu_{\rm ss}^{(4)} - (\frac{1}{k}+\frac{7p}{8})\mu_{\rm ca}^{(4)}. \tag{22b}$$

However, (22) does not satisfy condition 5 in Lemma 1 because the different discrete modes do not appear decoupled: $\dot{\mu}_{\rm ss}^{(3)}$ depends on $\mu_{\rm off}^{(0)}$, and $\dot{\mu}_{\rm ca}^{(3)}$ depends on $\mu_{\rm ss}^{(4)}$. For the purpose of determining φ, we ignore the cross coupling terms and approximate (22) by

$$\dot{\mu}_{\rm ss}^{(3)} \approx \frac{\log 8}{RTT}\mu_{\rm ss}^{(3)} - (\frac{1}{k}+p)\mu_{\rm ss}^{(4)}, \qquad \dot{\mu}_{\rm ca}^{(3)} \approx \frac{3}{RTT^2}\mu_{\rm ca}^{(2)} - (\frac{1}{k}+\frac{7p}{8})\mu_{\rm ca}^{(4)}. \tag{23}$$

The validity of these approximations generally depends on the network parameters. When (23) is used, it is straightforward to verify that (17) has a unique solution φ, resulting in a truncated system given by (21) and

$$\bar{\mu} = \varphi(\mu) = \left[\ \frac{\mu_{\rm ss}^{(0)}(\mu_{\rm ss}^{(2)})^3}{(\mu_{\rm ss}^{(1)})^3}\ \ \frac{(\mu_{\rm ca}^{(0)})^{\frac{1}{2}}(\mu_{\rm ca}^{(2)})^{\frac{5}{2}}}{(\mu_{\rm ca}^{(1)})^2}\ \right]'. \tag{24}$$

Since $\mu_{\rm off}^{(0)}$, $\mu_{\rm ss}^{(0)}$, and $\mu_{\rm ca}^{(0)}$ correspond to the probabilities that $\mathbf{q}$ is equal to off, ss, and ca, respectively, these quantities must add up to one and we can exclude one of them from the state.

Figure 2 shows a comparison between Monte Carlo simulations and the truncated model (21), (24). The dynamics of the first and second order moments are accurately predicted by the truncated model. There is some steady-state error in the probabilities of each mode, especially after the drop-rate p dropped to 2%. This corresponds to a situation in which the sending rate in the ss mode exceeds that of the ca mode and therefore the $\mu_{\rm ss}^{(4)}$ term dropped from (23) in the $\dot{\mu}_{\rm ca}^{(3)}$ equation may actually have a significant effect. Somewhat surprisingly, even in this case the first and second order moments are very well approximated. In preparing this paper, several simulation were executed for different network parameters and initial conditions. Figure 2 shows typical best-case (before $t = 1$) and worst-case (after $t = 1$) results.

The average file size used to produce Fig. 2 is consistent with the one observed in the UNIX file system reported in [15]. It should be emphasized that the overall distribution of the UNIX file system and more importantly of files transmitted over the Internet is not well approximated by an exponential distribution. More reasonable distributions are Pareto or mixtures of exponentials. Both could be modeled by pSHSs, but for simplicity here we focused our attention on the simplest case of a single exponential distribution. □

Example 4 (Networked control system). For this system $\psi^{(m)}(e) = e^m$, $\forall m \in \mathbb{N}_{\geq 0}$ and

$$(L\psi^{(m)})(e) = a\,m\,e^m + \frac{m(m-1)b^2}{2}e^{m-2} - e^{m+2\rho}$$

Fig. 2. (left) Comparison between Monte Carlo simulations (solid lines) and the truncated model (21), (24) (dashed lines) for Example 3, with $RTT = 50$ms, $\tau_{\text{off}} = 1$sec, $k = 20.39$ packets (corresponding to 30.58KB files broken into 1500bytes packets, which is consistent with the file-size distribution of the UNIX file system reported in [15]), and a step in the drop-rate p from 10% to 2% at $t = 1$sec.

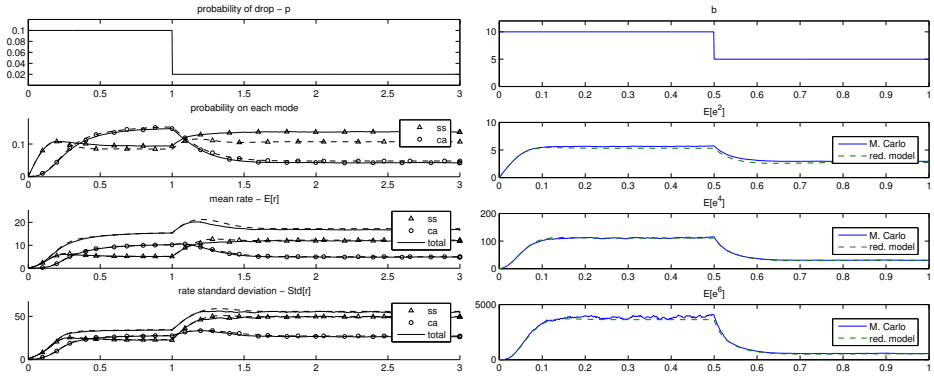

Fig. 3. (right) Comparison between Monte Carlo simulations (solid lines) and the truncated model (25), (26) (dashed lines) for Example 4, with $a = 1$, $q = 1$, and a step in the parameter b from 10 to 2 at time $t = 0.5$sec.

$$= am\,\psi^{(m)}(e) + \frac{m(m-1)b^2}{2}\psi^{(m-2)}(e) - \psi^{(m+2\rho)}(e).$$

Therefore

$$\begin{bmatrix} \dot{\mu}^{(1)} \\ \dot{\mu}^{(2)} \\ \dot{\mu}^{(3)} \\ \dot{\mu}^{(4)} \\ \dot{\mu}^{(5)} \\ \vdots \end{bmatrix} = \begin{bmatrix} a\mu^{(1)} - \mu^{(1+2\rho)} \\ b^2 + 2a\mu^{(2)} - \mu^{(2+2\rho)} \\ 3b^2\mu^{(1)} + 3a\mu^{(3)} - \mu^{(3+2\rho)} \\ 6b^2\mu^{(2)} + 4a\mu^{(4)} - \mu^{(4+2\rho)} \\ 10b^2\mu^{(3)} + 5a\mu^{(5)} - \mu^{(5+2\rho)} \\ \vdots \end{bmatrix}$$

In particular, for $\rho = 0$ we get

$$\begin{bmatrix} \dot{\mu}^{(1)} \\ \dot{\mu}^{(2)} \\ \dot{\mu}^{(3)} \\ \dot{\mu}^{(4)} \\ \dot{\mu}^{(5)} \\ \vdots \end{bmatrix} = \begin{bmatrix} 0 \\ -b^2 \\ 0 \\ 0 \\ 0 \\ \vdots \end{bmatrix} + \begin{bmatrix} a-1 & 0 & \cdots & & & \\ 0 & 2a-1 & 0 & \cdots & & \\ -3b^2 & 0 & 3a-1 & 0 & \cdots & \\ 0 & -6b^2 & 0 & 4a-1 & 0 & \cdots \\ 0 & 0 & -10b^2 & 0 & 5a-1 & 0 \ \cdots \\ \vdots & \vdots & \ddots & \ddots & \ddots & \ddots \ \ddots \end{bmatrix} \begin{bmatrix} \mu^{(1)} \\ \mu^{(2)} \\ \mu^{(3)} \\ \mu^{(4)} \\ \mu^{(5)} \\ \vdots \end{bmatrix}$$

and for $\rho = 1$ we get

$$\begin{bmatrix} \dot{\mu}^{(1)} \\ \dot{\mu}^{(2)} \\ \dot{\mu}^{(3)} \\ \dot{\mu}^{(4)} \\ \dot{\mu}^{(5)} \\ \vdots \end{bmatrix} = \begin{bmatrix} 0 \\ -b^2 \\ 0 \\ 0 \\ 0 \\ \vdots \end{bmatrix} + \begin{bmatrix} a & 0 & -1 & 0 & \cdots & & & & \\ 0 & 2a & 0 & -1 & 0 & \cdots & & & \\ -3b^2 & 0 & 3a & 0 & -1 & 0 & \cdots & & \\ 0 & -6b^2 & 0 & 4a & 0 & -1 & 0 & \cdots & \\ 0 & 0 & -10b^2 & 0 & 5a & 0 & -1 & 0 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots & \ddots & \ddots & \ddots & \ddots \end{bmatrix} \begin{bmatrix} \mu^{(1)} \\ \mu^{(2)} \\ \mu^{(3)} \\ \mu^{(4)} \\ \mu^{(5)} \\ \vdots \end{bmatrix}$$

Since $\mu^{(0)}(t) = 1$, $\forall t \geq 0$ we excluded $\mu^{(0)}$ from the vector μ, resulting in an affine system (as opposed to linear).

For $\rho = 0$, the infinite-dimensional dynamics have a lower-triangular structure and therefore an exact truncation is possible. However, this case is less interesting because it corresponds to a reset-rate that does not depend on the continuous state and is therefore farther from the optimal [9,10]. We consider here $\rho = 1$. In this case, the odd and even moments are decoupled and can be studied independently. It is straightforward to check that if the initial distribution of $\mathbf{e}$ is symmetric around zero, it will remain so for all times and therefore all odd moments are constant and equal to zero. Regarding the even moments, the smallest truncation for which condition 4 in Lemma 1 holds is a third order one, for which (9) can be written as follows:

$$\begin{bmatrix} \dot{\mu}^{(2)} \\ \dot{\mu}^{(4)} \\ \dot{\mu}^{(6)} \end{bmatrix} = \begin{bmatrix} 2a & -1 & 0 \\ 6b^2 & 4a & -1 \\ 0 & 15b^2 & 6a \end{bmatrix} \begin{bmatrix} \mu^{(2)} \\ \mu^{(4)} \\ \mu^{(6)} \end{bmatrix} + \begin{bmatrix} b^2 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ -1 \end{bmatrix} \bar{\mu}, \tag{25}$$

where $\bar{\mu} := \mu^{(8)}$ evolves according to $\dot{\mu}^{(8)} = -28b^2\mu^{(6)} + 8a\mu^{(8)} - \mu^{(10)}$. It is straightforward to verify that (17) has a unique solution φ, resulting in a truncated system given by (25) and

$$\bar{\mu} = \varphi(\mu^{(2)}, \mu^{(4)}, \mu^{(6)}) = \mu^{(2)} \Big(\frac{\mu^{(6)}}{\mu^{(4)}} \Big)^3. \tag{26}$$

Figure 3 shows a comparison between Monte Carlo simulations and the truncated model (25), (26). The dynamics of the all the moments are accurately predicted by the truncated model. The nonlinearity of the underlying model is apparent by the fact that halving b at time $t = 0.5$sec, which corresponds to dividing the variance of the noise by 4, only results in approximately dividing the variance of the estimation error by 2. □

Example 5 (Decaying-dimerizing reaction set). For this system the test functions are of the form $\psi^{(m_1,m_2,m_2)}(x) = x_1^{m_1} x_2^{m_2} x_3^{m_3}$, $\forall m_1, m_2, m_3 \in \mathbb{N}_{\geq 0}$ and we conclude that

$$\begin{aligned} (L\psi^{(m_1,m_2,m_2)})(x) &= c_1 x_1 \big((x_1-1)^{m_1} - x_1^{m_1} \big) x_2^{m_2} x_3^{m_3} \\ &\quad + \frac{c_2}{2} x_1 (x_1-1) \big((x_1-2)^{m_1} (x_2+1)^{m_2} - x_1^{m_1} x_2^{m_2} \big) x_3^{m_3} \\ &\quad + c_3 x_2 \big((x_1+2)^{m_1} (x_2-1)^{m_2} - x_1^{m_1} x_2^{m_2} \big) x_3^{m_3} + c_4 x_2 \big((x_2-1)^{m_2} (x_3+1)^{m_3} - x_2^{m_2} x_3^{m_3} \big) x_1^{m_1} \\ &= c_1 \sum_{i=0}^{m_1-1} \tbinom{m_1}{i} (-1)^{m_1-i} \psi^{(i+1,m_2,m_3)}(x) \end{aligned}$$

$$+\frac{c_2}{2}\sum_{\substack{i,j=0\\(i,j)\neq(m_1,m_2)}}^{m_1,m_2}\binom{m_1}{i}\binom{m_2}{j}(-2)^{m_1-i}\left(\psi^{(i+2,j,m_3)}(x)-\psi^{(i+1,j,m_3)}(x)\right)$$
$$+c_3\sum_{\substack{i,j=0\\(i,j)\neq(m_1,m_2)}}^{m_1,m_2}\binom{m_1}{i}\binom{m_2}{j}2^{m_1-i}(-1)^{m_2-j}\psi^{(i,j+1,m_3)}(x)$$
$$+c_4\sum_{\substack{i,j=0\\(i,j)\neq(m_2,m_3)}}^{m_2,m_3}\binom{m_2}{i}\binom{m_3}{j}(-1)^{m_2-i}\psi^{(m_1,i+1,j)}(x), \quad (27)$$

where the summations result from the power expansions of the terms $(x_i-c)^{m_i}$. Therefore

$$\begin{bmatrix}\dot\mu^{(1,0,0)}\\ \dot\mu^{(0,1,0)}\\ \dot\mu^{(0,0,1)}\\ \dot\mu^{(2,0,0)}\\ \dot\mu^{(0,2,0)}\\ \dot\mu^{(1,1,0)}\\ \dot\mu^{(3,0,0)}\\ \dot\mu^{(2,1,0)}\\ \vdots\end{bmatrix}=\begin{bmatrix} -c_1+c_2 & 2c_3 & 0 & -c_2 & 0 & \cdots\\ -c_2/2 & -c_3-c_4 & 0 & c_2/2 & 0 & \cdots\\ 0 & c_4 & 0 & 0 & 0 & \cdots\\ c_1-2c_2 & 4c_3 & 0 & -2c_1+4c_2 & 0 & 4c_3\\ -c_2/2 & c_3+c_4 & 0 & c_2/2 & -2c_3-2c_4 & -c_2\\ c_2 & -2c_3 & 0 & -3c_2/2 & 2c_3 & -c_1+c_2-c_3-c_4\\ -c_1+4c_2 & 8c_3 & 0 & 3c_1-10c_2 & 0 & 12c_3\\ -2c_2 & -4c_3 & 0 & 4c_2 & 4c_3 & c_1-2c_2-4c_3\\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots\end{bmatrix}$$
$$\begin{bmatrix} \cdots & & & & & & & & \\ \cdots & & & & & & & & \\ \cdots & & & & & & & & \\ & -2c_2 & 0 & \cdots & & & & & \\ 0 & c_2 & 0 & \cdots & & & & & \\ c_2/2 & -c_2 & 0 & \cdots & & & & & \\ -3c_1+9c_2 & 6c_3 & 0 & -3c_2 & 0 & \cdots & & & \\ -5c_2/2 & 0 & -2c_1+4c_2-c_3-c_4 & 4c_3 & c_2/2 & -2c_3 & 0 & \cdots & \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots\end{bmatrix}\begin{bmatrix}\mu^{(1,0,0)}\\ \mu^{(0,1,0)}\\ \mu^{(0,0,1)}\\ \mu^{(2,0,0)}\\ \mu^{(0,2,0)}\\ \mu^{(1,1,0)}\\ \mu^{(3,0,0)}\\ \mu^{(2,1,0)}\\ \mu^{(1,2,0)}\\ \mu^{(4,0,0)}\\ \mu^{(3,1,0)}\\ \vdots\end{bmatrix}.$$

For this example we consider a truncation whose state contains all the first and second order moments for the number of particles of the first and second species. To keep the formulas small, we omit from the truncation the second moments of the third species, which does not appear as a reactant in any reaction and therefore its higher order statistics do not affect the first two. In this case, (9) can be written as follows:

$$\begin{bmatrix}\dot\mu^{(1,0,0)}\\ \dot\mu^{(0,1,0)}\\ \dot\mu^{(0,0,1)}\\ \dot\mu^{(2,0,0)}\\ \dot\mu^{(0,2,0)}\\ \dot\mu^{(1,1,0)}\end{bmatrix}=\begin{bmatrix} -c_1+c_2 & 2c_3 & 0 & -c_2 & 0 & 0\\ -\frac{c_2}{2} & -c_3-c_4 & 0 & \frac{c_2}{2} & 0 & 0\\ 0 & c_4 & 0 & 0 & 0 & 0\\ c_1-2c_2 & 4c_3 & 0 & 4c_2-2c_1 & 0 & 4c_3\\ -\frac{c_2}{2} & c_3+c_4 & 0 & \frac{c_2}{2} & -2c_3-2c_4 & -c_2\\ c_2 & -2c_3 & 0 & -\frac{3c_2}{2} & 2c_3 & c_2-c_1-c_3-c_4\end{bmatrix}\begin{bmatrix}\mu^{(1,0,0)}\\ \mu^{(0,1,0)}\\ \mu^{(0,0,1)}\\ \mu^{(2,0,0)}\\ \mu^{(0,2,0)}\\ \mu^{(1,1,0)}\end{bmatrix}+\begin{bmatrix} 0 & 0\\ 0 & 0\\ 0 & 0\\ -2c_2 & 0\\ 0 & c_2\\ \frac{c_2}{2} & -c_2\end{bmatrix}\bar\mu,$$

where $\bar{\mu} := \left[\mu^{(3,0,0)}\ \mu^{(2,1,0)}\right]'$ evolves according to

$$\begin{aligned}\dot{\mu}^{(3,0,0)} &= (-c_1+4c_2)\mu^{(1,0,0)} + 8c_3\mu^{(0,1,0)} + (3c_1-10c_2)\mu^{(2,0,0)} + 12c_3\mu^{(1,1,0)} \\ &\quad + (-3c_1+9c_2)\mu^{(3,0,0)} + 6c_3\mu^{(2,1,0)} - 3c_2\mu^{(4,0,0)} \qquad (29a)\\ \dot{\mu}^{(2,1,0)} &= -2c_2\mu^{(1,0,0)} - 4c_3\mu^{(0,1,0)} + 4c_2\mu^{(2,0,0)} + 4c_3\mu^{(0,2,0)} + (c_1-2c_2-4c_3)\mu^{(1,1,0)} \\ &\quad - \frac{5c_2\mu^{(3,0,0)}}{2} + (4c_2-2c_1-c_3-c_4)\mu^{(2,1,0)} + 4c_3\mu^{(1,2,0)} + \frac{c_2\mu^{(4,0,0)}}{2} - 2c_2\mu^{(3,1,0)}. \qquad (29b)\end{aligned}$$

This system does not satisfy condition 4 in Lemma 1 because the $\mu^{(1,0,0)}$, $\mu^{(0,1,0)}$ terms in the right-hand sides of (29) lead to monomials in x_1 and x_2 in $\sum_{i=1}^{\infty} a_{\ell,i}x^{(m_i)}$ that do not exist in any of the polynomials $\left\{\sum_{i=1}^{\infty} a_{j,i}x^{(m_\ell - m_j + m_i)} : 1 \le j \le 6\right\}$. These terms can be easily traced back to the lowest-order terms in power expansions in (27) and disappear if we only keep the three highest powers of x_1 in the expansion of $(x_1-1)^{m_1}$ that corresponds to the first reaction and the two highest powers of x_1 and x_2 in the expansions of $(x_1 \pm 2)^{m_1}$ and $(x_2 \pm 1)^{m_2}$ that correspond to the second and third reactions. In practice, this leads to the following simplified version of (28)

$$\begin{bmatrix}\dot{\mu}^{(1,0,0)}\\ \dot{\mu}^{(0,1,0)}\\ \dot{\mu}^{(0,0,1)}\\ \dot{\mu}^{(2,0,0)}\\ \dot{\mu}^{(0,2,0)}\\ \dot{\mu}^{(1,1,0)}\end{bmatrix} = \begin{bmatrix} -c_1+c_2 & 2c_3 & 0 & -c_2 & 0 & 0\\ -\frac{c_2}{2} & -c_3-c_4 & 0 & \frac{c_2}{2} & 0 & 0\\ 0 & c_4 & 0 & 0 & 0 & 0\\ c_1 & 0 & 0 & 4c_2-2c_1 & 0 & 4c_3\\ 0 & c_4 & 0 & 0 & -2c_3-2c_4 & -c_2\\ c_2 & -2c_3 & 0 & -\frac{3c_2}{2} & 2c_3 & c_2-c_1-c_3-c_4\end{bmatrix}\begin{bmatrix}\mu^{(1,0,0)}\\ \mu^{(0,1,0)}\\ \mu^{(0,0,1)}\\ \mu^{(2,0,0)}\\ \mu^{(0,2,0)}\\ \mu^{(1,1,0)}\end{bmatrix} + \begin{bmatrix} 0 & 0\\ 0 & 0\\ 0 & 0\\ -2c_2 & 0\\ 0 & c_2\\ \frac{c_2}{2} & -c_2\end{bmatrix}\bar{\mu}, \qquad (30)$$

where now $\bar{\mu} := \left[\mu^{(3,0,0)}\ \mu^{(2,1,0)}\right]'$ evolves according to

$$\begin{aligned}\dot{\mu}^{(3,0,0)} &= 3c_1\mu^{(2,0,0)} + (-3c_1+3c_2)\mu^{(3,0,0)} + 6c_3\mu^{(2,1,0)} - 3c_2\mu^{(4,0,0)}\\ \dot{\mu}^{(2,1,0)} &= 2c_2\mu^{(2,0,0)} + (c_1-4c_3)\mu^{(1,1,0)} - \frac{5c_2\mu^{(3,0,0)}}{2}\\ &\quad + (-2c_1+2c_2-c_3-c_4)\mu^{(2,1,0)} + 4c_3\mu^{(1,2,0)} + \frac{c_2\mu^{(4,0,0)}}{2} - 2c_2\mu^{(3,1,0)},\end{aligned}$$

for which condition 4 in Lemma 1 does hold, allowing us to find a unique solution φ to (17), resulting in a truncated system given (28) and

$$\bar{\mu} = \varphi(\mu) = \left[\left(\frac{\mu^{(2,0,0)}}{\mu^{(1,0,0)}}\right)^3\ \frac{\mu^{(2,0,0)}}{\mu^{(0,1,0)}}\left(\frac{\mu^{(1,1,0)}}{\mu^{(1,0,0)}}\right)^2\right]'. \qquad (31)$$

Ignoring the lowest-order powers of x_1 and x_2 in the power expansions is valid when the populations of these species are high. In practice, the approximation still seems to yield good results even when the populations are fairly small as shown by the Monte Carlo simulations.

Figure 4 shows a comparison between Monte Carlo simulations and the truncated model (28), (31). The coefficients used were taken from [13, Example 1]: $c_1 = 1$, $c_2 = 10$, $c_3 = 1000$, $c_4 = 10^{-1}$. In Fig. 4(a), we used the same initial conditions as in [13, Example 1]:

$$\mathbf{x}_1(0) = 400, \qquad \mathbf{x}_2(0) = 798, \qquad \mathbf{x}_3(0) = 0. \qquad (32)$$

Table 1. Comparison between estimates obtained from Monte Carlo simulations and the truncated model for Example 5. The Monte Carlo simulation data was taken from [13].

Source for the estimates	$\mathrm{E}[\mathbf{x}_1(0.2)]$	$\mathrm{E}[\mathbf{x}_2(0.2)]$	$\mathrm{StdDev}[\mathbf{x}_1(0.2)]$	$\mathrm{StdDev}[\mathbf{x}_2(0.2)]$
10,000 MC. simul.	387.3	749.5	18.42	10.49
model (28), (31)	387.2	749.6	18.54	10.60

The match is very accurate, as can be confirmed from Table 1. The values of the parameters chosen result in a pSHS with two distinct time scales, which makes this pSHS computationally difficult to simulate ("stiff" in the terminology of [13]). The initial conditions (32) start in the "slow manifold" $\mathbf{x}_2 = \frac{5}{1000}\mathbf{x}_1(\mathbf{x}_1 - 1)$ and Fig. 4(a) essentially shows the evolution of the system on this manifold. Figure 4(b) zooms in on the interval $[0, 5\times 10^{-4}]$ and shows the evolution of the system towards the manifold when it starts away from it at

$$\mathbf{x}_1(0) = 800, \qquad \mathbf{x}_2(0) = 100, \qquad \mathbf{x}_3(0) = 200. \tag{33}$$

Figures 4(c)–4(d) shows another simulation of the same reactions but for much smaller initial populations:

$$\mathbf{x}_1(0) = 10, \qquad \mathbf{x}_2(0) = 10, \qquad \mathbf{x}_3(0) = 5. \tag{34}$$

The truncated model still provides an extremely good approximation, with significant error only in the covariance between $\mathbf{x}_1$ and $\mathbf{x}_2$ when the averages and standard deviation of these variables get below one. This happens in spite of having used (30) instead of (29) to compute φ. □

6 Conclusions and Future Work

In this paper, we showed that the infinite-dimensional linear moment dynamics of a pSHS can be approximated by a finite-dimensional nonlinear ODE with arbitrary precision. Moreover, we provided a procedure to build this type of approximation. The methodology was illustrated using a varied pool of examples, demonstrating its wide applicability. Several observations arise from these examples, which point to directions for future research:

1. Surprisingly, in all the examples considered the truncation function φ does not depend on the parameters of the vector fields or the transition intensities. Clearly, this cannot be a general property of pSHSs and apparently we were simply "lucky" in the examples for which we applied our procedure. However, there may be deeper reasons for this independence that require further investigation.
2. The computation of the truncation function does not directly take into account the steady-state distribution of the pSHS. This is desirable because for all the examples presented it is difficult to compute the steady-state values of the moments. However, when steady-state distributions are available, they can improve the steady-state accuracy of the truncated system. This was done in Example 2, for which

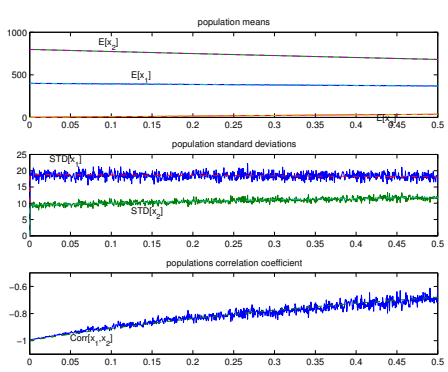

(a) Large population, over a long time scale with initial condition (32)

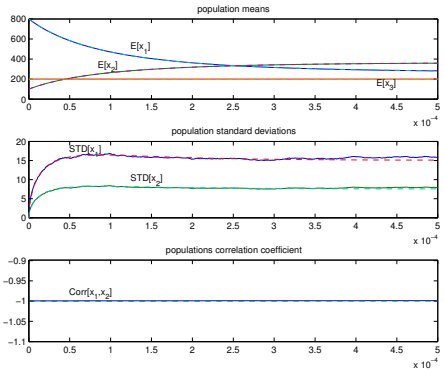

(b) Large population, over a short time scale with initial condition (33)

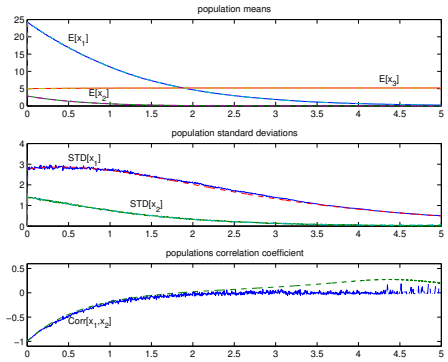

(c) Small population, over a long time scale with initial condition (34)

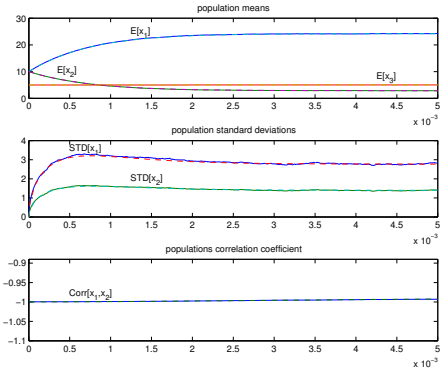

(d) Small population, over short time scale with initial condition (34)

Fig. 4. Comparison between Monte Carlo simulations (solid lines) and the truncated model (28), (31) (dashed lines) for Example 5.

experimental observations indicate that the steady-state distribution of the sending rate appears to be approximately Log-Normal. However, it would be desirable to devise constructions that improve the steady-state accuracy even when this type of insight is not readily available.

3. In all the examples presented, we restricted our attention to truncation functions φ of the form (16) and we only used deterministic distributions to compute φ. Mostly likely, better results could be obtained by considering more general distributions, which may require more general forms for φ.

4. The truncation of pSHSs that model chemical reactions proved especially accurate. This motivates the search for systematic procedures to automatically construct a truncated system from chemical equations such as (6). Another direction for future research consists of comparing the truncated models obtains here with those in [16].

An additional direction for future research consists of establishing computable bounds on the error between solutions to the infinite-dimensional moments dynamics and to its finite-dimensional truncations.

Appendix

Proof (Theorem 2). For a given $\delta > 0$, let $T > 0$ and $N > 0$ be sufficiently large so that for every solution μ_∞ and ν to (8) and (10) starting at some $t_0 \geq 0$ in Ω_μ and Ω_ν, respectively, we have[4]

$$\beta(\|\mu(t_1) - \nu(t_1)\|, \tau) \leq \frac{\delta}{2}, \qquad \forall \tau \geq T,\ t_1 \geq t_0 \tag{35}$$

$$\frac{(2T)^{N+2}}{(N+2)!}\left(\sup_{\tau \geq t_0}\left\|\frac{\mathrm{d}^{N+2}\mu(\tau)}{\mathrm{d}t^{N+2}}\right\| + \sup_{\tau \geq t_0}\left\|\frac{\mathrm{d}^{N+2}\nu(\tau)}{\mathrm{d}t^{N+2}}\right\|\right) \leq \frac{\delta}{2}. \tag{36}$$

The existence of T is guaranteed by the fact that $\beta \in \mathscr{K}\mathscr{L}$ and by the boundedness of ν and μ, as per Assumption 1. For the chosen T, the existence of N is guaranteed by Assumption 1 and the fact that $\lim_{N\to\infty} \frac{(2T)^{N+2}}{(N+2)!} = 0$. The justification for these definitions will become clear shortly.

Let μ_∞ and ν be solutions to (8) and (10) starting in Ω_μ and Ω_ν, respectively, at some time $t_0 \geq 0$. For a given $t_1 \geq t_0$, since Ω_ν is forward invariant, $\nu_1 := \nu(t_1) \in \Omega_\nu$ and therefore, because of Assumption 2, there exists some $\hat{\mu}_\infty(t_1) \in \Omega_\mu$ whose first k elements match ν_1 and for which (11) holds. Therefore

$$\begin{aligned}\|\mu(t) - \nu(t)\| &\leq \|\mu(t) - \hat{\mu}(t)\| + \|\hat{\mu}(t) - \nu(t)\| \\ &\leq \beta(\|\mu(t_1) - \nu(t_1)\|, t - t_1) + \|\hat{\mu}(t) - \nu(t)\|, \qquad \forall t \geq t_1. \end{aligned} \tag{37}$$

We now show by induction on i that

$$\frac{\mathrm{d}^i \hat{\mu}(t_1)}{\mathrm{d}t^i} = \frac{\mathrm{d}^i \nu(t_1)}{\mathrm{d}t^i}, \qquad \forall i \in \{0, 1, \ldots, N+1\}. \tag{38}$$

The base of induction ($i = 0$) holds trivially by the way $\hat{\mu}(t_1)$ was selected. Suppose now that the statement hold for some $i \leq N$. Taking $i+1$ derivatives of $\hat{\mu}(t) - \nu(t)$ with respect to t, one obtains

$$\begin{aligned}\frac{\mathrm{d}^{i+1}\big(\hat{\mu}(t) - \nu(t)\big)}{\mathrm{d}t^{i+1}} &= \frac{\mathrm{d}^i}{\mathrm{d}t^i}\Big(\dot{\hat{\mu}}(t) - \dot{\nu}(t)\Big) \\ &= \frac{\mathrm{d}^i}{\mathrm{d}t^i}\Big(A(t)\big(\hat{\mu}(t) - \nu(t)\big)\Big) + \frac{\mathrm{d}^i}{\mathrm{d}t^i}\Big(B(t)\big(C\hat{\mu}_\infty(t) - \bar{\nu}(t)\big)\Big)\end{aligned}$$

[4] The left-hand-sides of (35) and (36) need to add up to δ but other options are possible.

$$= \sum_{j=0}^{i} \binom{i}{j} \frac{\mathrm{d}^{i-j} A(t)}{\mathrm{d}t^{i-j}} \frac{\mathrm{d}^{j}\left(\hat{\mu}(t)-\nu(t)\right)}{\mathrm{d}t^{j}} + \sum_{j=0}^{i} \binom{i}{j} \frac{\mathrm{d}^{i-j} B(t)}{\mathrm{d}t^{i-j}} \frac{\mathrm{d}^{j}\left(C\hat{\mu}_\infty(t)-\bar{\nu}(t)\right)}{\mathrm{d}t^{j}}. \quad (39)$$

All the terms $\frac{\mathrm{d}^{j}\left(\hat{\mu}(t)-\nu(t)\right)}{\mathrm{d}t^{j}}$ in the first summation are equal to zero at $t=t_1$ because of the induction hypothesis. Moreover, from (12) and (13) we conclude that

$$\frac{\mathrm{d}^{j} C\hat{\mu}_\infty(t_1)}{\mathrm{d}t^{j}} = C^{j}(t_1)\hat{\mu}_\infty(t_1) = (L^{j}\varphi)(\hat{\mu}(t_1),t_1) = (L^{j}\varphi)(\nu(t_1),t_1) = \frac{\mathrm{d}^{j}\bar{\nu}(t_1)}{\mathrm{d}t^{j}}.$$

This means that all the terms $\frac{\mathrm{d}^{j}\left(C\hat{\mu}_\infty(t)-\bar{\nu}(t)\right)}{\mathrm{d}t^{j}}$ in (39) are also equal to zero at $t=t_1$, which finishes the proof of (38).

Expanding $\hat{\mu}(t)$ and $\nu(t)$ as an Nth-order Taylor series around the point $t_1 \geq 0$ and using (38) and (36), conclude that $\|\hat{\mu}(t)-\nu(t)\| \leq \frac{\delta}{2}$, $\forall t \in [t_1, t_1+2T]$. From this and (37), we obtain

$$\|\mu(t)-\nu(t)\| \leq \beta(\|\mu(t_1)-\nu(t_1)\|, t-t_1) + \frac{\delta}{2}, \qquad \forall t \in [t_1, t_1+2T]. \quad (40)$$

We show next by induction that for every $i \in \mathbb{N}_{\geq 1}$,

$$\|\mu(t)-\nu(t)\| \leq \delta, \qquad \forall t \in [t_0+iT, t_0+(i+1)T]. \quad (41)$$

The basis of induction ($i=1$) is a direct consequence of (40) with $t_1 := t_0$ and (35). Assuming now that (41) holds for some $i \geq 1$, we conclude from (40) with $t_1 := t_0+iT$ that

$$\|\mu(t)-\nu(t)\| \leq \beta(\|\mu(t_0+iT)-\nu(t_0+iT)\|, t-t_0-iT) + \frac{\delta}{2} \leq \delta,$$

$\forall t \in [t_0+(i+1)T, t_0+(i+2)T]$, which finishes the proof of (41). From (40) with $t_1 := t_0$ and (41) if follows that

$$\|\mu(t)-\nu(t)\| \leq \beta(\|\mu(t_0)-\nu(t_0)\|, t-t_0) + \delta, \qquad \forall t \geq t_0,$$

which completes the proof. ■

Proof (Lemma 1). Take a vector $\mu_\infty \in \Omega_\mu$ corresponding to the distribution $P(\cdot; q, x)$ and let $\mu_{q_i}^{(m_i)}$ be the moment corresponding to the ith row of μ_∞. Considering a particular row of C that has a one at the ℓth columns and zero at every other column, the corresponding row in the vector equality (17a) can be written as

$$\mu_{q_\ell}^{(m_\ell)} = \lambda_\ell \prod_{i=1}^{k} (\mu_{q_i}^{(m_i)})^{\gamma_{\ell,i}} \quad \Leftrightarrow \quad \begin{cases} x^{(m_\ell)} & q_\ell = q \\ 0 & q_\ell \neq q \end{cases} = \lambda_\ell \prod_{i=1}^{k} \begin{cases} 1 & \gamma_{\ell,i} = 0 \\ x^{(\gamma_{\ell,i} m_i)} & \gamma_{\ell,i} \neq 0,\ q_i = q \\ 0 & \gamma_{\ell,i} \neq 0,\ q_i \neq q \end{cases}$$

where $\lambda_\ell \neq 0$ is the diagonal entry of Λ that corresponds this particular row of C. For $q_\ell = q$ and $x \neq 0$ the left-hand-side is nonzero and therefore we must have $\gamma_{\ell,i} = 0$

whenever $q_i \neq q$, because otherwise the right-hand-side would be zero, which proves 2. Moreover, we must also have $x^{(m_\ell)} = \lambda_\ell x^{\left(\sum_{i=1,q_i=q_\ell}^{k} m_i \gamma_{\ell,i}\right)}$, $\forall x \in \Omega_x$, which proves 1 and 3. Consider now the corresponding row in the vector equality (17b), which can be written as

$$\sum_{i=1}^{\infty} a_{\ell,i} \mu_{q_i}^{(m_i)} = \mu_{q_\ell}^{(m_\ell)} \sum_{j=1}^{k} \frac{\gamma_{\ell,j}}{\mu_{q_j}^{(m_j)}} \sum_{i=1}^{\infty} a_{j,i} \mu_{q_i}^{(m_i)} \tag{42}$$

For $q_\ell = q$, since $\gamma_{\ell,j} = 0$ whenever $q_j \neq q = q_\ell$ (because of 2), we conclude that

$$\sum_{\substack{i=1 \\ q_i=q}}^{\infty} a_{\ell,i} x^{(m_i)} = \sum_{\substack{j=1 \\ q_j=q}}^{k} \gamma_{\ell,j} \sum_{\substack{i=1 \\ q_i=q}}^{\infty} a_{j,i} x^{(m_\ell - m_j + m_i)}, \qquad \forall x \in \Omega_x,$$

which proves 4. For $q_\ell \neq q$, the right-hand-side of (42) is zero and therefore $\sum_{\substack{i=1 \\ q_i=q}}^{\infty} a_{\ell,i} x^{(m_i)} = 0$, $\forall x \in \Omega_x$ which proves 5. ■

References

1. Hespanha, J.P.: Stochastic hybrid systems: Applications to communication networks. In Alur, R., Pappas, G.J., eds.: Hybrid Systems: Computation and Control. Number 2993 in Lect. Notes in Comput. Science. Springer-Verlag, Berlin (2004) 387–401
2. Davis, M.H.A.: Markov models and optimization. Monographs on statistics and applied probability. Chapman & Hall, London, UK (1993)
3. Hu, J., Lygeros, J., Sastry, S.: Towards a theory of stochastic hybrid systems. In Lynch, N.A., Krogh, B.H., eds.: Hybrid Systems: Computation and Control. Volume 1790 of Lect. Notes in Comput. Science., Springer (2000) 160–173
4. Pola, G., Bujorianu, M.L., Lygeros, J., Benedetto, M.D.D.: Stochastic hybrid models: An overview. In: Proc. of the IFAC Conf .on Anal. and Design of Hybrid Syst. (2003)
5. Bujorianu, M.L.: Extended stochastic hybrid systems and their reachability problem. In: Hybrid Systems: Computation and Control. Lect. Notes in Comput. Science. Springer-Verlag, Berlin (2004) 234–249
6. Papoulis, A.: Probability, Random Variables, and Stochastic Processes. 3rd edn. McGraw-Hill Series in Electrical Engineering. McGraw-Hill, Inc., New York (1991)
7. Teel, A.R., Moreau, L., Nešić, D.: A unified framework for input-to-state stability in systems with two time scales. IEEE Trans. on Automat. Contr. **48** (2003) 1526–1544
8. Bohacek, S., Hespanha, J.P., Lee, J., Obraczka, K.: A hybrid systems modeling framework for fast and accurate simulation of data communication networks. In: Proc. of the ACM Int. Conf. on Measurements and Modeling of Computer Systems (SIGMETRICS). (2003)
9. Xu, Y., Hespanha, J.P.: Communication logics for networked control systems. In: Proc. of the 2004 Amer. Contr. Conf. (2004)
10. Xu, Y., Hespanha, J.P.: Optimal communication logics for networked control systems. In: Proc. of the 43rd Conf. on Decision and Contr. (2004) To appear.
11. Gillespie, D.T.: A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. J. of Computational Physics **22** (1976) 403–434
12. Gillespie, D.T., Petzold, L.R.: Improved leap-size selection for accelerated stochastic simulation. J. of Chemical Physics **119** (2003) 8229–8234

13. Rathinam, M., Petzold, L.R., Cao, Y., Gillespie, D.T.: Stiffness in stochastic chemically reacting systems: The implicit tau-leaping method. J. of Chemical Physics **119** (2003) 12784–12794
14. Hespanha, J.P.: A model for stochastic hybrid systems with application to communication networks. Submitted to the *Int. Journal of Hybrid Systems* (2004)
15. Irlam, G.: Unix file size survey – 1993. Available at `http://www.base.com/gordoni/ufs93.html` (1994)
16. Kampen, N.G.V.: Stochastic Processes in Physics and Chemistry. Elsevier Science (2001)