research supported by NSF

Stochastic Hybrid Systems:
Applications to Communication Networks

João P. Hespanha

Center for Control Engineering
and Computation

University of California
at Santa Barbara

UCSB

Talk outline

UCSB

1. A (simple) model for stochastic hybrid systems (SHSs)
2. SHSs models for network traffic under TCP
3. Analysis tools for SHSs
4. Dynamics of TCP

Collaborators:
Stephan Bohacek, Junsoo Lee, Katia Obraczka, Abhyudai Singh, Yonggang Xu

Acknowledgements:
Mustafa Khammash, John Lygeros

Deterministic Hybrid Systems

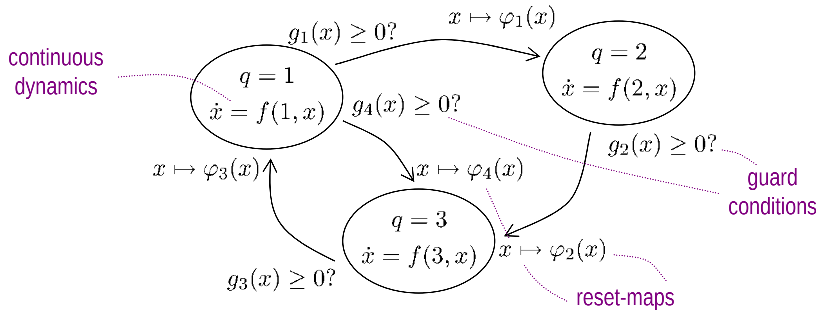

$q(t) \in \mathcal{Q}=\{1,2,\ldots\}$ $\equiv$ discrete state

$x(t) \in \mathbb{R}^n$ $\equiv$ continuous state

} right-continuous by convention

we assume here a deterministic system so the invariant sets would be the exact complements of the guards

Stochastic Hybrid Systems

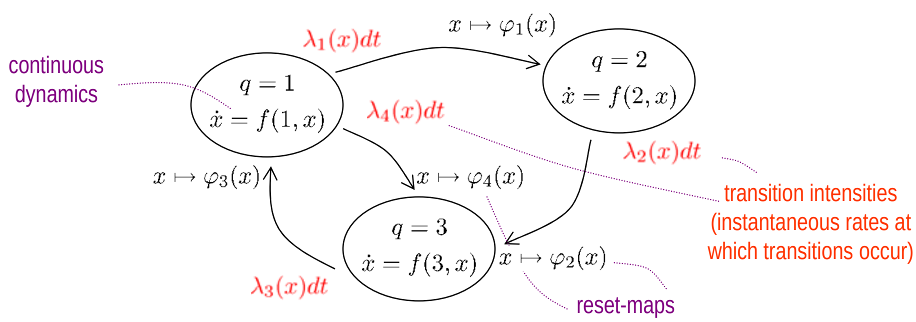

$N_\ell(t) \in \mathbb{N} \equiv$ transition counter, which is incremented by one each time the ℓth reset-map $\varphi_\ell(x)$ is "activated"

} right-continuous by convention

$$\underbrace{\mathrm{P}\Big(N_\ell(t+dt) > N_\ell(t)\Big)}_{\text{at least one transition on } (t,\, t+dt]} \xrightarrow[dt \to 0]{} \mathrm{E}\Big[\lambda_\ell\big(x(t)\big)dt\Big]$$

at least one transition on $(t, t+dt]$

proportional to "elementary" interval length dt and transition intensity $\lambda_\ell(x)$

Stochastic Hybrid Systems

UCSB

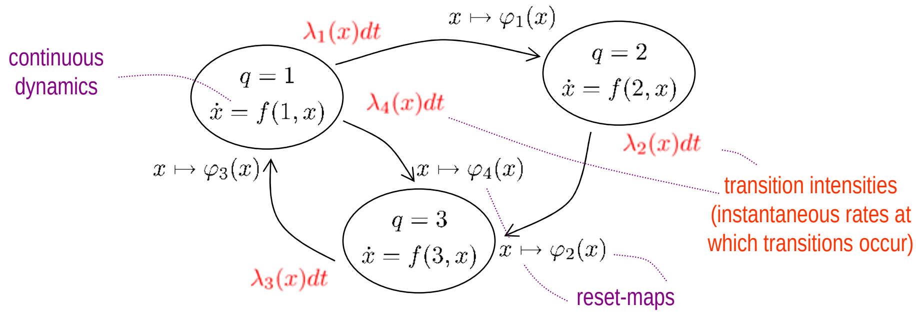

$N_\ell(t) \in \mathbb{N} \equiv$ transition counter, which is incremented by one each time the ℓth reset-map $\varphi_\ell(x)$ is "activated"

} right-continuous by convention

$$\mathrm{E}\left[N_\ell(t_2) - N_\ell(t_1)\right] = \mathrm{E}\left[\int_{t_1}^{t_2} \lambda_\ell(x(t))\,dt\right]$$

number of transitions on $(t_1, t_2]$

equal to integral of transition intensity $\lambda_\ell(x)$ on $(t_1, t_2]$

Stochastic Hybrid Systems

UCSB

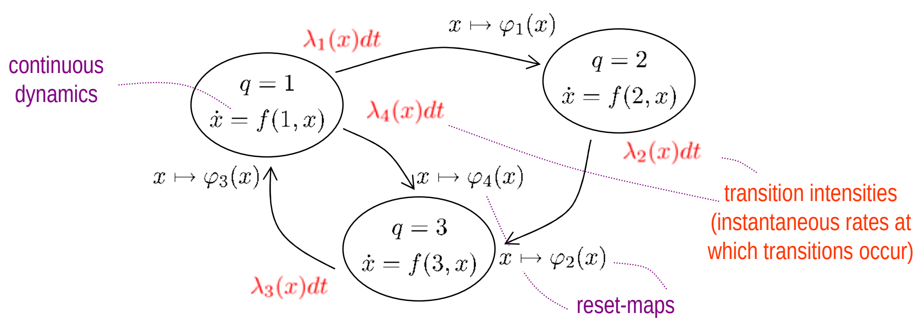

Special case: When all λ_ℓ are constant, transitions are controlled by a continuous-time Markov process

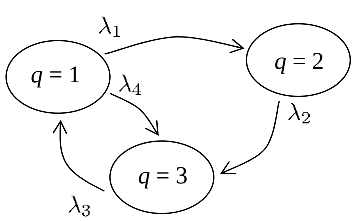

specifies q (independently of x)

Formal model—Summary

UCSB

State space: $q(t) \in \mathcal{Q}=\{1,2,\ldots\}$ ≡ discrete state
$x(t) \in \mathbb{R}^n$ ≡ continuous state

Continuous dynamics:

$\dot{x} = f(q, x, t)$ $\qquad f : \mathcal{Q} \times \mathbb{R}^n \times [0, \infty) \to \mathbb{R}^n$

Transition intensities:

$\lambda_\ell(q, x, t)$ $\qquad \lambda_\ell : \mathcal{Q} \times \mathbb{R}^n \times [0, \infty) \to [0, \infty) \quad \ell \in \{1, \ldots, m\}$

Reset-maps (one per transition intensity): # of transitions

$(q, x) \mapsto \phi_\ell(q, x, t)$ $\qquad \phi_\ell : \mathcal{Q} \times \mathbb{R}^n \times [0, \infty) \to \mathcal{Q} \times \mathbb{R}^n \quad \ell \in \{1, \ldots, m\}$

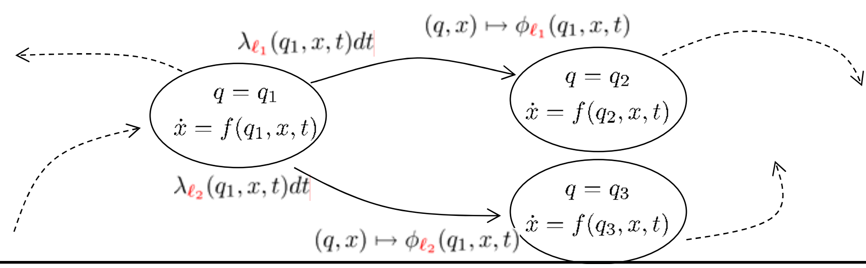

Formal model—Summary

UCSB

State space: $q(t) \in \mathcal{Q}=\{1,2,\ldots\}$ ≡ discrete state
$x(t) \in \mathbb{R}^n$ ≡ continuous state

Continuous dynamics:

$\dot{x} = f(q, x, t)$ $\qquad f : \mathcal{Q} \times \mathbb{R}^n \times [0, \infty) \to \mathbb{R}^n$

Transition intensities:

$\lambda_\ell(q, x, t)$ $\qquad \lambda_\ell : \mathcal{Q} \times \mathbb{R}^n \times [0, \infty) \to [0, \infty) \quad \ell \in \{1, \ldots, m\}$

Reset-maps (one per transition intensity): # of transitions

$(q, x) \mapsto \phi_\ell(q, x, t)$ $\qquad \phi_\ell : \mathcal{Q} \times \mathbb{R}^n \times [0, \infty) \to \mathcal{Q} \times \mathbb{R}^n \quad \ell \in \{1, \ldots, m\}$

Results:

1. [*existence*] Under appropriate regularity (Lipschitz) assumptions, there exists a measure "consistent" with the desired SHS behavior
2. [*simulation*] The procedure used to construct the measure is constructive and allows for efficient generation of *Monte Carlo sample paths*
3. [*Markov*] The pair ($q(t)$, $x(t)$) $\in \mathcal{Q} \times \mathbb{R}^n$ is a (Piecewise-deterministic) Markov Process (in the sense of M. Davis, 1993)

Hespanha. Stochastic Hybrid Systems: Applications to Communication Networks. HSCC'04

Transmission Control Protocol

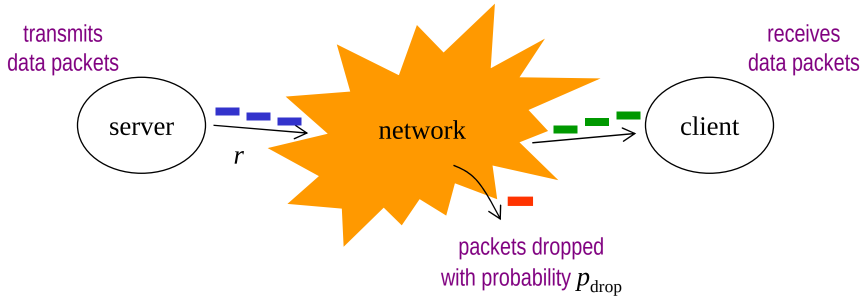

congestion control $\equiv$ selection of the rate r at which the server transmits packets

feedback mechanism $\equiv$ packets are dropped by the network to indicate congestion

TCP window-based control

w (window size) $\equiv$ number of packets that can remain unacknowledged for by the destination

e.g., $w = 3$

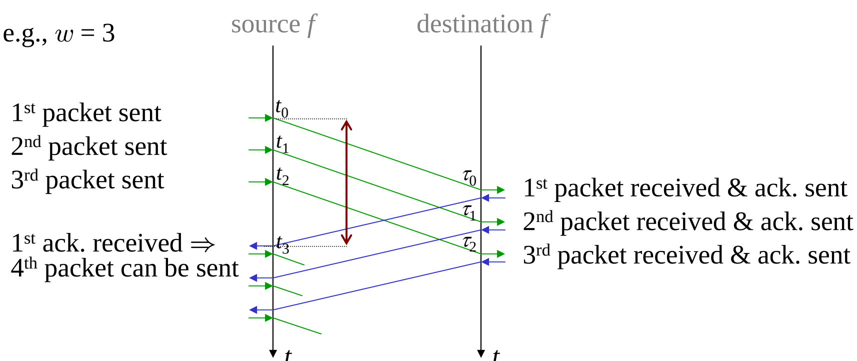

w effectively determines the sending rate r:

$$r(t) = \frac{w(t)}{RTT(t)}$$

round-trip time

Built-in negative feedback:

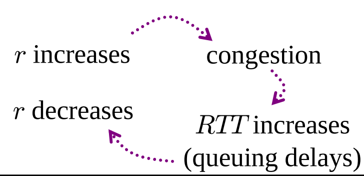

TCP Reno *congestion avoidance*

UCSB

1. While there are no drops, increase w by 1 on each RTT (additive increase)
2. When a drop occurs, divide w by 2 (multiplicative decrease)

   (congestion controller constantly probes the network for more bandwidth)

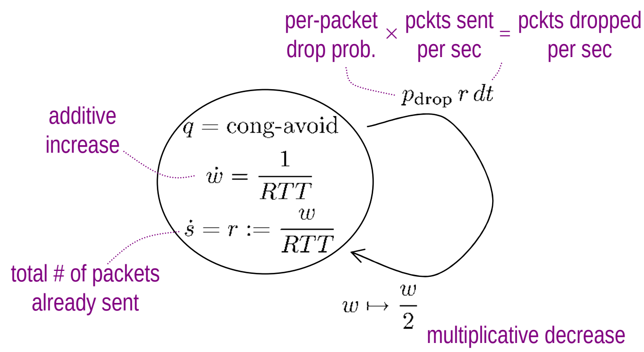

Three feedback mechanisms: As rate r increases ⇒ …

a) congestion ⇒ RTT increases (queuing delays) ⇒ r decreases

b) congestion ⇒ p_{drop} increases (active queuing) ⇒ r decreases

c) multiplicative decrease more likely ⇒ r decreases

Delay in drop detection

UCSB

3. In general there is some delay between drop occurrence and detection (assumed exponentially distributed with mean τ_{delay})

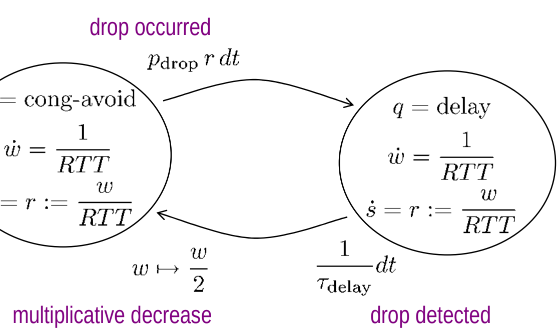

TCP Reno *slow start*

UCSB

3. Until a drop occurs double w on each RTT
4. When a drop occurs divide w by 2 and transition to congestion avoidance

   (get to full bandwidth utilization fast)

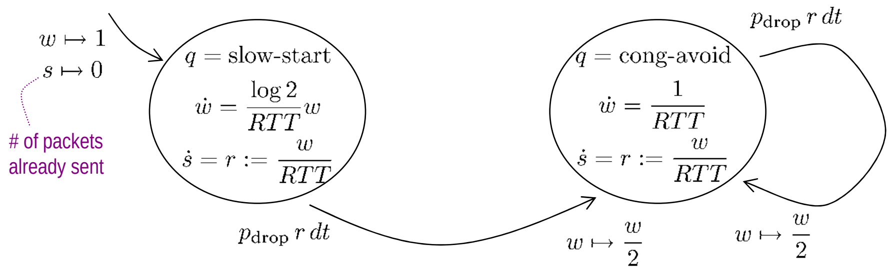

for simplicity this diagram ignores delay

TCP has several other modes that can be modeled by hybrid systems

[Bohacek, Hespanha, Lee, Obraczka. A Hybrid Systems Modeling Framework for Fast and Accurate Simulation of Data Communication Networks. In SIGMETRICS'03]

On-off TCP model

UCSB

- Between transfers, server remains *off* for an exponentially distributed time with mean τ_{off}
- Transfer-size is a random variable with cumulative distribution $F(s)$

determines duration of on-periods

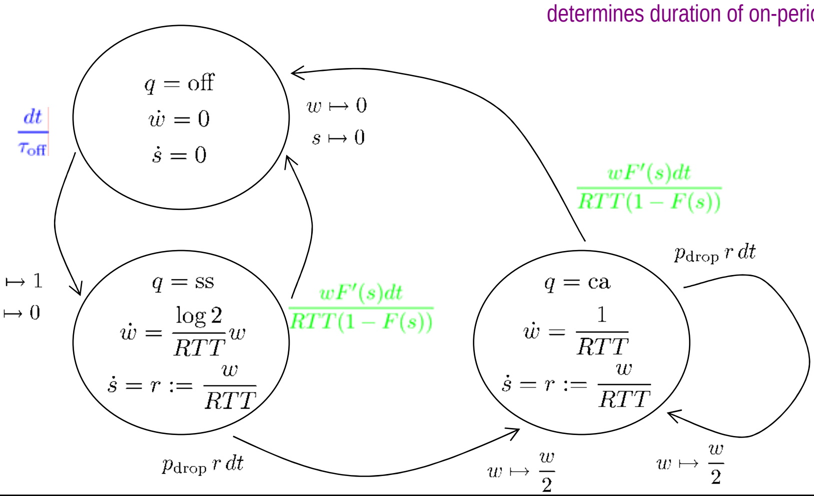

On-off TCP model

- Between transfers, server remains *off* for an exponentially distributed time with mean τ_{off}
- Transfer-size is exponentially distributed with average k (packets)

simple but not very realistic…

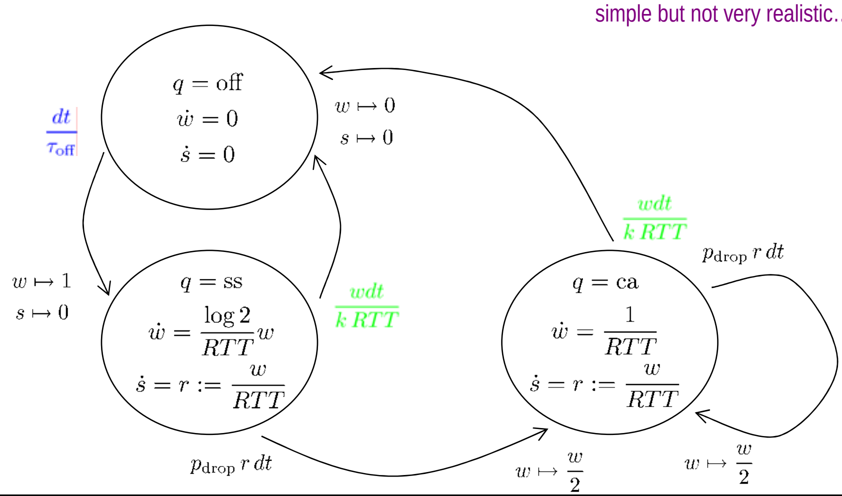

SHS models for TCP

long-lived TCP flows
(no delays)

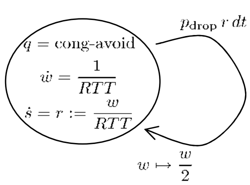

on-off TCP flows
(no delays)

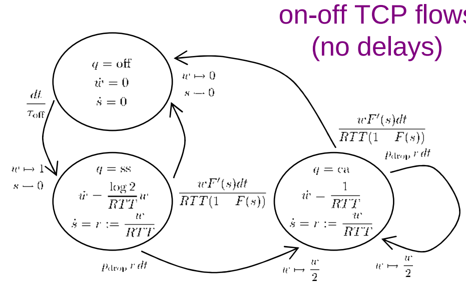

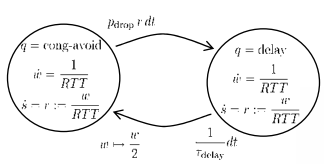

long-lived TCP
flows with delay

on-off TCP
flows with delay

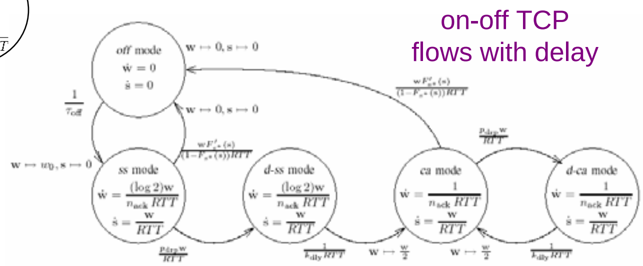

Analysis—Lie Derivative

$$\dot{x} = f(x,t) \qquad x \in \mathbb{R}^n$$

Given function $\psi : \mathbb{R}^n \times [0,\infty) \to \mathbb{R}$

$$\dot{\psi}(x,t) = \underbrace{\frac{\partial \psi}{\partial x} f(x,t) + \frac{\partial \psi}{\partial t}}_{L_f \psi}$$

derivative along solution to ODE

$L_f \psi$ Lie derivative of ψ

One can view L_f as an operator

space of scalar functions on $\mathbb{R}^n \times [0,\infty)$ $\to$ space of scalar functions on $\mathbb{R}^n \times [0,\infty)$

$\psi(x, t) \mapsto L_f \psi(x, t)$

L_f completely defines the system dynamics

Analysis—Generator of the SHS

$\dot{x} = f(q,x,t)$ continuous dynamics

$\lambda_\ell(q,x,t)$ transition intensities

$(q,x) = \phi_\ell(q^-, x^-, t)$ reset-maps

Given function $\psi : \mathcal{Q} \times \mathbb{R}^n \times [0,\infty) \to \mathbb{R}$

$$\frac{\mathrm{d}}{\mathrm{d}t} \mathrm{E}[\psi(q,x,t)] = \mathrm{E}\Big[(L\psi)(q,x,t)\Big]$$

generator for the SHS

Dynkin's formula (in differential form)

where

$$(L\psi)(q,x,t) := \underbrace{\frac{\partial \psi}{\partial x} f(q,x,t) + \frac{\partial \psi}{\partial t}}_{L_f \psi} + \sum_{\ell=1}^{m} \underbrace{\Big(\psi\big(\phi_\ell(q,x,t),t\big) - \psi(q,x,t)\Big)}_{\text{reset instantaneous variation}} \lambda_\ell(q,x,t)$$

$L_f \psi$ Lie derivative of ψ

reset instantaneous variation

intensity

L completely defines the SHS dynamics

Disclaimer: see following paper for technical assumptions

Hespanha. Stochastic Hybrid Systems: Applications to Communication Networks. HSCC'04

Long-lived TCP flows

long-lived
TCP flows

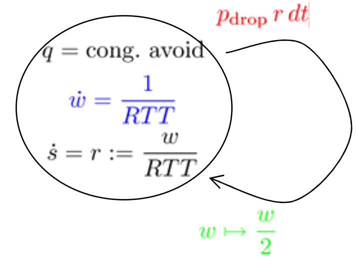

$$(L\psi)(w,t) = \frac{\partial \psi}{\partial w}\frac{1}{RTT} + \frac{\partial \psi}{\partial t} + \left[\psi\left(\frac{w}{2},t\right) - \psi(w,t)\right] p_{\text{drop}} \frac{w}{RTT}$$

Long-lived TCP flows

long-lived TCP flows
(with slow start)

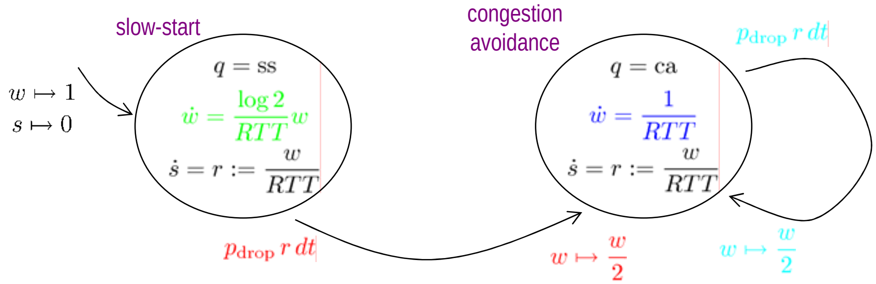

$$(L\psi)(q,w,t) = \begin{cases} \frac{\partial \psi}{\partial w}\frac{\log 2}{RTT} w + \frac{\partial \psi}{\partial t} + \left[\psi\left(\text{ca}, \frac{w}{2}\right) - \psi(\text{ss}, w)\right] \frac{p_{\text{drop}} w}{RTT} & q = \text{ss} \\ \frac{\partial \psi}{\partial w}\frac{1}{RTT} + \frac{\partial \psi}{\partial t} + \left[\psi\left(\text{ca}, \frac{w}{2}\right) - \psi(\text{ss}, w)\right] \frac{p_{\text{drop}} w}{RTT} & q = \text{ca} \end{cases}$$

Analysis—PDF for SHS state

$$\dot{x} = f(q, x, t) \qquad \lambda_\ell(q, x, t) \qquad (q, x) = \phi_\ell(q^-, x^-, t)$$

continuous dynamics — transition intensities — reset-maps

Let ρ(· ; t) be the *probability density function* (PDF) for the state (x,q) at time t :

$$\mathrm{E}[\psi(q(t), x(t), t)] = \sum_{q \in \mathcal{Q}} \int_{\mathbb{R}^n} \psi(q, x, t)\rho(q, x; t)dx$$

When ρ(· ; t) is smooth, one can deduce from the generator that

$$\frac{\partial \rho(q, x; t)}{\partial t} = -\frac{\partial f(q, x, t)\rho(q, x; t)}{\partial x} + \sum_{\ell=1}^{m} \Bigg(- \lambda_\ell(q, x, t)\rho(q, x; t) +$$

$$+ \sum_{\substack{\bar{q} \in \mathcal{Q}: \\ \phi_\ell^{\bar{q}}(\bar{q}, \phi_\ell^{-x}(\bar{q}, x, t)) = q}} \lambda_\ell(\bar{q}, \phi_\ell^{-x}(\bar{q}, x, t), t)\rho(\bar{q}, \phi_\ell^{-x}(\bar{q}, x); t) \left| \frac{\partial \phi_\ell^{-x}(\bar{q}, x, t)}{\partial z} \right| \Bigg)$$

inverse of $x \mapsto \phi_\ell^x(q, x, t)$

functional PDE
very difficult to solve

Analysis—Moments for SHS state

$$\dot{x} = f(q, x, t) \qquad \lambda_\ell(q, x, t) \qquad (q, x) = \phi_\ell(q^-, x^-, t)$$

continuous dynamics — transition intensities — reset-maps

often a few low-order moments suffice to study a SHS…

z (scalar) random variable with mean μ and variance σ^2

$$\mathrm{P}\left(\mathbf{z} \geq \epsilon \mid \mathbf{z} \geq 0\right) \leq \frac{\mu}{\epsilon}$$

Markov inequality
(∀ $\epsilon > 0$)

$$\mathrm{P}\left(|\mathbf{z} - a| \geq \epsilon\right) \leq \frac{\mathrm{E}[|\mathbf{z} - a|^n]}{\epsilon^n}$$

Bienaymé inequality
(∀ $\epsilon > 0$, $a \in \mathbb{R}$, $n \in \mathbb{N}$)

$$\mathrm{P}\left(|\mathbf{z} - \mu| \geq \epsilon\right) \leq \frac{\sigma^2}{\epsilon^2}$$

Tchebychev inequality

Polynomial SHSs

$\dot{x} = f(q, x, t)$ — continuous dynamics

$\lambda_\ell(q, x, t)$ — transition intensities

$(q, x) = \phi_\ell(q^-, x^-, t)$ — reset-maps

Given function $\psi : \mathcal{Q} \times \mathbb{R}^n \times [0,\infty) \to \mathbb{R}$

generator for the SHS

$$\frac{\mathrm{d}}{\mathrm{d}t}\,\mathrm{E}[\psi(q, x, t)] = \mathrm{E}\Big[(L\psi)(q, x, t)\Big]$$

Dynkin's formula (in differential form)

where

$$(L\psi)(q, x, t) := \frac{\partial \psi}{\partial x} f(q, x, t) + \frac{\partial \psi}{\partial t} + \sum_{\ell=1}^{m} \Big(\psi\big(\phi_\ell(q, x, t), t\big) - \psi(q, x, t)\Big)\lambda_\ell(q, x, t)$$

A SHS is called a *polynomial SHS* (pSHS) if its generator maps finite-order polynomial on x into finite-order polynomials on x

Typically, when

$$x \mapsto f(q, x, t) \quad x \mapsto \lambda_\ell(q, x, t) \quad x \mapsto \phi_\ell(q, x, t)$$

are all polynomials $\forall\, q, t$

Moment dynamics for pSHS

$x(t) \in \mathbb{R}^n$ — continuous state

$q(t) \in \mathcal{Q}=\{1,2,\ldots\}$ — discrete state

Test function: Given $\bar{q} \in \mathcal{Q}$ $\quad m = (m_1, m_2, \ldots, m_n) \in \mathbb{N}^n_{\geq 0}$

$$\psi_{\bar{q}}^{(m)}(q, x) := \begin{cases} x_1^{m_1} x_2^{m_2} \cdots x_n^{m_n} & q = \bar{q} \\ 0 & q \neq \bar{q}, \end{cases}$$

for short $x^{(m)}$

Uncentered moment:

$$\mu_{\bar{q}}^{(m)}(t) := \mathrm{E}\left[\psi_{\bar{q}}^{(m)}(q(t), x(t))\right]$$

E.g,

$$\mu_1^{(0,0,\ldots,0)}(t) = \mathrm{P}\big(q(t) = 1\big) \qquad \mu_1^{(1,1,\ldots,0)}(t) = \mathrm{E}\left[x_1(t) x_2(t) I_{q(t)=1}\right]$$

$$\mu_1^{(0,1,\ldots,0)}(t) = \mathrm{E}\left[x_2(t) I_{q(t)=1}\right] \qquad \mu_1^{(2,0,\ldots,0)}(t) = \mathrm{E}\left[x_1(t)^2 I_{q(t)=1}\right]$$

Moment dynamics for pSHS

$x(t) \in \mathbb{R}^n$ continuous state

$q(t) \in \mathcal{Q}=\{1,2,\ldots\}$ discrete state

Test function: Given $\bar{q} \in \mathcal{Q}$ $\quad m = (m_1, m_2, \ldots, m_n) \in \mathbb{N}^n_{\geq 0}$

$$\psi_{\bar{q}}^{(m)}(q,x) := \begin{cases} x_1^{m_1} x_2^{m_2} \cdots x_n^{m_n} & q = \bar{q} \\ 0 & q \neq \bar{q}, \end{cases}$$

for short $x^{(m)}$

Uncentered moment:

$$\mu_{\bar{q}}^{(m)}(t) := \mathrm{E}\left[\psi_{\bar{q}}^{(m)}\big(q(t), x(t)\big)\right]$$

For polynomial SHS…

$\psi_{\bar{q}}^{(m)}(q,x)$ monomial on x $\Rightarrow$ $L\psi_{\bar{q}}^{(m)}(q,x)$ polynomial on x $\Rightarrow$ $L\psi_{\bar{q}}^{(m)}(q,x)$ linear comb. of test functions

$$\dot{\mu}_{\bar{q}}^{(m)} = \frac{\mathrm{d}}{\mathrm{d}t} \mathrm{E}[\psi_{\bar{q}}^{(m)}(q,x)] = \mathrm{E}\left[(L\psi_{\bar{q}}^{(m)})(q,x)\right] = \sum_{i=1}^{k} \alpha_i \mu_{q_i}^{(m_i)}$$

linear moment dynamics

Moment dynamics for TCP

long-lived TCP flows

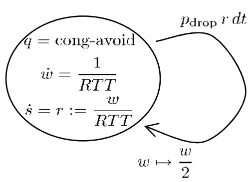

$$\mu^{(k)} := \mathrm{E}[r^k] = \mathrm{E}\left[\frac{w^k}{RTT^k}\right]$$

$$\dot{\mu}^{(1)} = \frac{1}{RTT^2} - \frac{p}{2}\mu^{(1)}$$
$$\dot{\mu}^{(2)} = \frac{2}{RTT^2}\mu^{(1)} - \frac{3p}{4}\mu^{(3)}$$
$$\dot{\mu}^{(3)} = \frac{3}{RTT^2}\mu^{(2)} - \frac{7p}{8}\mu^{(4)}$$
$$\vdots$$

Moment dynamics for TCP

UCSB

long-lived
TCP flows

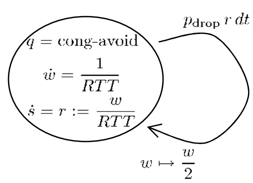

on-off TCP flows
(exp. distr. files)

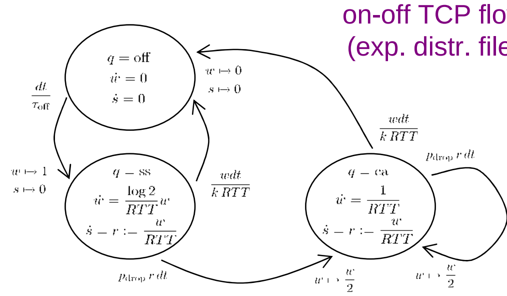

$$\mu^{(k)} := \mathrm{E}[r^k] = \mathrm{E}\left[\frac{w^k}{RTT^k}\right]$$

$$\mu^{(k)}_{\text{ca}} := \mathrm{E}[r^k I_{q=\text{ca}}] = \mathrm{E}\left[\frac{w^k}{RTT^k} I_{q=\text{ca}}\right] \cdots$$

$$\dot{\mu}^{(1)} = \frac{1}{RTT^2} - \frac{p}{2}\mu^{(1)}$$

$$\dot{\mu}^{(2)} = \frac{2}{RTT^2}\mu^{(1)} - \frac{3p}{4}\mu^{(3)}$$

$$\dot{\mu}^{(3)} = \frac{3}{RTT^2}\mu^{(2)} - \frac{7p}{8}\mu^{(4)}$$

$$\vdots$$

$$\begin{bmatrix} \dot{\mu}^{(0)}_{\text{off}} \\ \dot{\mu}^{(0)}_{\text{ss}} \\ \dot{\mu}^{(0)}_{\text{ca}} \\ \dot{\mu}^{(1)}_{\text{ss}} \\ \dot{\mu}^{(1)}_{\text{ca}} \\ \dot{\mu}^{(2)}_{\text{ss}} \\ \dot{\mu}^{(2)}_{\text{ca}} \\ \vdots \end{bmatrix} = \begin{bmatrix} -\tau_{\text{off}}^{-1}\mu^{(0)}_{\text{off}} + \frac{1}{k}\mu^{(1)}_{\text{ss}} + \frac{1}{k}\mu^{(1)}_{\text{ca}} \\ \tau_{\text{off}}^{-1}\mu^{(0)}_{\text{off}} - (\frac{1}{k}+p)\mu^{(1)}_{\text{ss}} \\ p\mu^{(1)}_{\text{ss}} - \frac{1}{k}\mu^{(1)}_{\text{ca}} \\ \frac{\tau_{\text{off}}^{-1} w_0}{RTT}\mu^{(0)}_{\text{off}} + \frac{\log 2}{RTT}\mu^{(1)}_{\text{ss}} - (\frac{1}{k}+p)\mu^{(2)}_{\text{ss}} \\ \frac{1}{RTT^2}\mu^{(0)}_{\text{ca}} + \frac{p}{2}\mu^{(2)}_{\text{ss}} - (\frac{1}{k}+\frac{p}{2})\mu^{(2)}_{\text{ca}} \\ \frac{\tau_{\text{off}}^{-1} w_0^2}{RTT^2}\mu^{(0)}_{\text{off}} + \frac{\log 4}{RTT}\mu^{(2)}_{\text{ss}} - (\frac{1}{k}+p)\mu^{(3)}_{\text{ss}} \\ \frac{2}{RTT^2}\mu^{(1)}_{\text{ca}} + \frac{p}{4}\mu^{(3)}_{\text{ss}} - (\frac{1}{k}+\frac{3p}{4})\mu^{(3)}_{\text{ca}} \\ \vdots \end{bmatrix}$$

Truncated moment dynamics

UCSB

Experimental evidence indicates that the sending rate is well approximated by a Log-Normal distribution

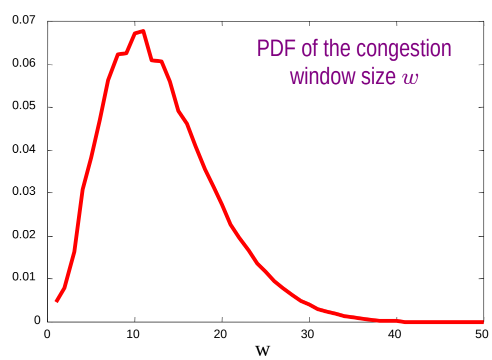

z Log-Normal $\Rightarrow$ $\mathrm{E}[\mathbf{z}^3] = \frac{\mathrm{E}[\mathbf{z}^2]^3}{\mathrm{E}[\mathbf{z}]^3}$

r Log-Normal (on each mode) $\Rightarrow$ $\mu_{q,3} = \mathrm{E}[\mathbf{r}^3 \mid \mathbf{q} = q]\mu_{q,0} \approx \frac{\mathrm{E}[\mathbf{r}^2 \mid \mathbf{q} = q]^3}{\mathrm{E}[\mathbf{r} \mid \mathbf{q} = q]^3}\mu_{q,0} = \frac{\mu^3_{q,2}\,\mu_{q,0}}{\mu^3_{q,1}}$

Data from: Bohacek, A stochastic model for TCP and fair video transmission. INFOCOM'03

Moment dynamics for TCP

UCSB

long-lived
TCP flows

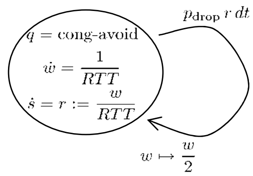

on-off TCP flows
(exp. distr. files)

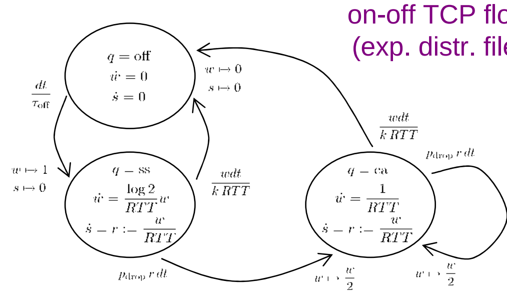

$$\mu^{(k)} := \mathrm{E}[r^k] = \mathrm{E}\left[\frac{w^k}{RTT^k}\right]$$

$$\mu^{(k)}_{ca} := \mathrm{E}[r^k I_{q=ca}] = \mathrm{E}\left[\frac{w^k}{RTT^k} I_{q=ca}\right] \cdots$$

$$\dot\mu^{(1)} = \frac{1}{RTT^2} - \frac{p}{2}\mu^{(1)}$$

$$\dot\mu^{(2)} = \frac{2}{RTT^2}\mu^{(1)} - \frac{3p}{4}\left(\frac{\mu^{(2)}}{\mu^{(1)}}\right)^3$$

$$\begin{bmatrix} \dot\mu^{(0)}_{off} \\ \dot\mu^{(0)}_{ss} \\ \dot\mu^{(0)}_{ca} \\ \dot\mu^{(1)}_{ss} \\ \dot\mu^{(1)}_{ca} \\ \dot\mu^{(2)}_{ss} \\ \dot\mu^{(2)}_{ca} \end{bmatrix} = \begin{bmatrix} -\tau_{off}^{-1}\mu^{(0)}_{off} + \frac{1}{k}\mu^{(1)}_{ss} + \frac{1}{k}\mu^{(1)}_{ca} \\ \tau_{off}^{-1}\mu^{(0)}_{off} - (\frac{1}{k}+p)\mu^{(1)}_{ss} \\ p\mu^{(1)}_{ss} - \frac{1}{k}\mu^{(1)}_{ca} \\ \frac{\tau_{off}^{-1} w_0}{RTT}\mu^{(0)}_{off} + \frac{\log 2}{RTT}\mu^{(1)}_{ss} - (\frac{1}{k}+p)\mu^{(2)}_{ss} \\ \frac{1}{RTT^2}\mu^{(0)}_{ca} + \frac{p}{2}\mu^{(2)}_{ss} - (\frac{1}{k}+\frac{p}{2})\mu^{(2)}_{ca} \\ \frac{\tau_{off}^{-1} w_0^2}{RTT^2}\mu^{(0)}_{off} + \frac{\log 4}{RTT}\mu^{(2)}_{ss} - (\frac{1}{k}+p)\mu^{(0)}_{ss}(\frac{\mu^{(2)}_{ss}}{\mu^{(1)}_{ss}})^3 \\ \frac{2}{RTT^2}\mu^{(1)}_{ca} + \frac{p}{4}\mu^{(0)}_{ss}(\frac{\mu^{(2)}_{ss}}{\mu^{(1)}_{ss}})^3 - (\frac{1}{k}+\frac{3p}{4})\mu^{(0)}_{ca}(\frac{\mu^{(2)}_{ca}}{\mu^{(1)}_{ca}})^3 \end{bmatrix}$$

finite-dimensional
nonlinear ODEs

Long-lived TCP flows

UCSB

long-lived TCP
flows with delay
(one RTT average)

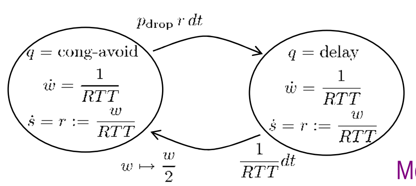

Monte-Carlo & theoretical
steady-state distribution
(vs. drop probability)

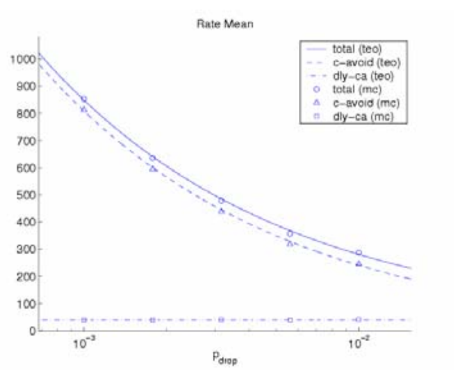

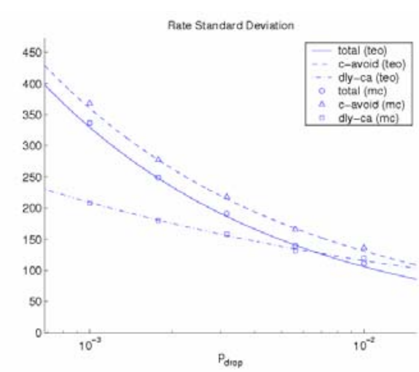

Long-lived TCP flows

moment dynamics in response to abrupt change in drop probability

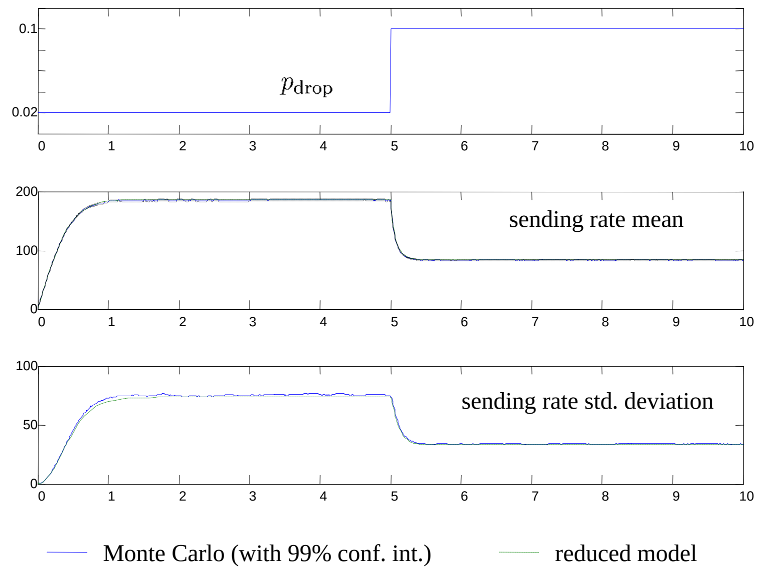

On-off TCP flows

SHS for on-off TCP flows

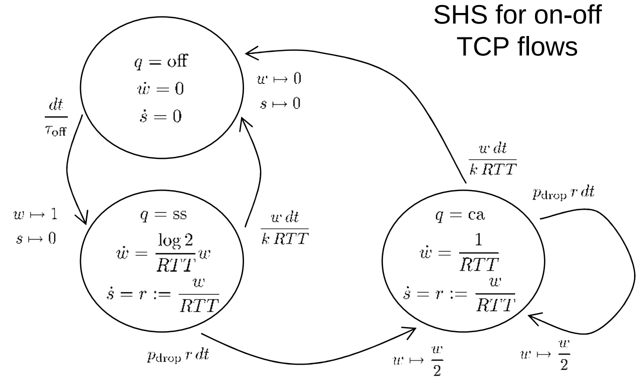

Transfer-sizes exponentially distributed with mean k = 30.6KB (Unix'93 files)
Off-time exponentially distributed with mean τ_{off} = 1 sec
Round-trip-time RTT = 50msec

On-off TCP flows

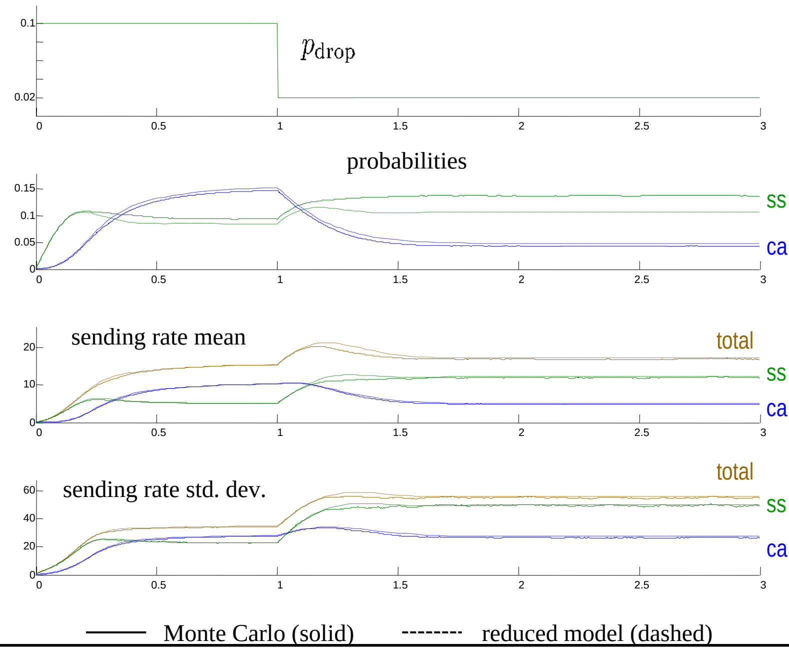

On-off TCP flows

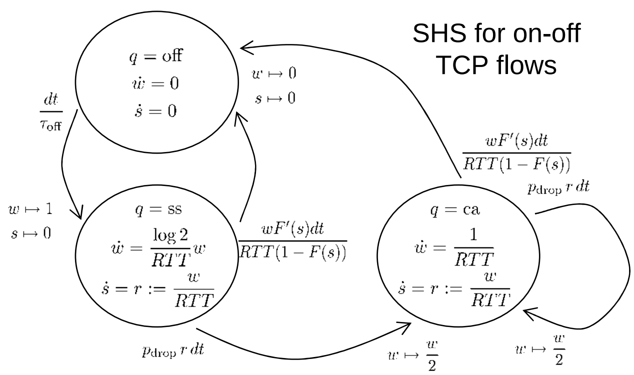

Mixture of exponential distribution for transfer-sizes:

m exp. distr. with mean $\{ k_1, k_2, \ldots, k_m \}$ m probabilities $\{ p_1, p_2, \ldots, p_m \}$ $\sum_{i=1}^{m} p_i = 1$

transfer-sizes are extracted from an exponential distribution with mean k_i w.p. p_i

can approximate well distributions found in the literature …

On-off TCP flows: Case I

Transfer-size approximating the distribution of file sizes in the UNIX file system

p_1 = 88.87%	k_1 = 3.5KB	("mice" files)
p_2 = 11.13%	k_2 = 246KB	("elephant" files, but not heavy tail)

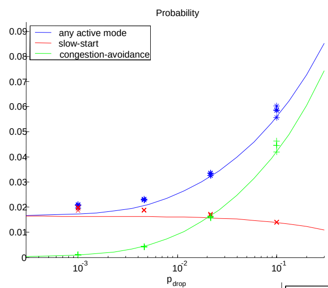

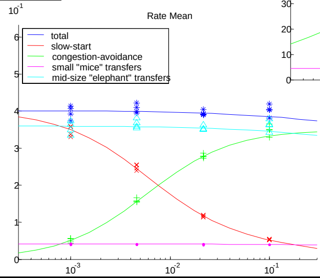

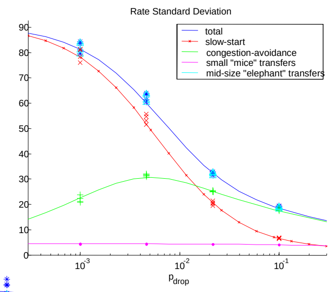

Monte-Carlo & theoretical
steady-state distribution
(vs. drop probability)

On-off TCP flows: Case II

Transfer-size distribution at an ISP's www proxy [Arlitt *et al.*]

p_1 = 98%	k_1 = 6KB	("mice" files)
p_2 = 1.7%	k_2 = 400KB	("elephant" files)
p_3 = .02%	k_3 = 10MB	("mammoth" files)

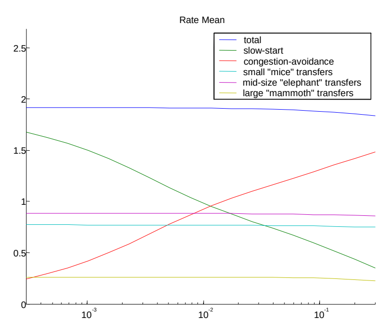

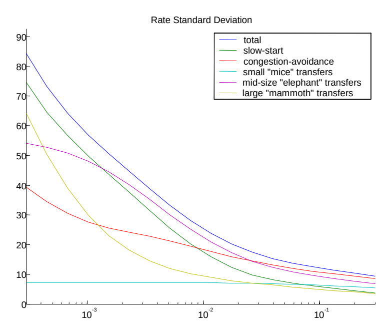

Theoretical
steady-state distribution
(vs. drop probability)

Truncated moment dynamics (revisited)

For polynomial SHS…

$$\dot\mu^{(m)}_{\bar q} = \frac{\mathrm{d}}{\mathrm{d}t}\,\mathrm{E}[\psi^{(m)}_{\bar q}(q,x)] = \mathrm{E}\Big[(L\psi^{(m)}_{\bar q})(q,x)\Big] = \sum_{i=1}^{k}\alpha_i\mu^{(m_i)}_{q_i}$$

linear moment dynamics

Stacking all moments into an (infinite) vector μ_∞

$$\dot\mu_\infty = A_\infty\mu_\infty$$

infinite-dimensional linear ODE

In TCP analysis…

$$\mu_\infty = \begin{bmatrix}\mu^{(0)}_{\mathrm{ss}}\\ \mu^{(0)}_{\mathrm{ca}}\\ \mu^{(1)}_{\mathrm{ss}}\\ \vdots\\ \mu^{(3)}_{\mathrm{ss}}\\ \mu^{(3)}_{\mathrm{ca}}\\ \vdots\end{bmatrix}$$

μ: lower order moments of interest

$\bar\mu$: moments of interest that affect μ dynamics

$$\dot\mu = A\mu + B\bar\mu$$

approximated by nonlinear function of μ

Truncation by derivative matching

$$\dot\mu_\infty = A_\infty\mu_\infty$$

infinite-dimensional linear ODE

$$\dot\mu = A\mu + B\bar\mu$$

truncated linear ODE (nonautonomous, not nec. stable)

$$\dot\nu = A\nu + B\varphi(\nu)$$

nonlinear approximate moment dynamics

Assumption: 1) μ and ν remain bounded along solutions to

$$\dot\mu_\infty = A_\infty\mu_\infty \quad \text{and} \quad \dot\nu = A\nu + B\varphi(\nu)$$

2) $\dot\mu_\infty = A_\infty\mu_\infty$ is asymptotically stable

Theorem: $\forall\,\delta > 0\ \exists\,N$ s.t. if $\dfrac{\mathrm{d}^k\mu}{\mathrm{d}t^k} = \dfrac{\mathrm{d}^k\nu}{\mathrm{d}t^k}, \quad \forall k\in\{1,\dots,N\}$

then

$$\|\mu(t)-\nu(t)\| \le \beta(\|\mu(t_0)-\nu(t_0)\|, t-t_0) + \delta, \qquad \forall t\ge t_0\ge 0$$

class $\mathcal{KL}$ function

Hespanha. Polynomial Stochastic Hybrid Systems. HSCC'05

Truncation by derivative matching

UCSB

infinite-dimensional linear ODE

$$\dot{\mu}_\infty = A_\infty \mu_\infty$$

$$\dot{\mu} = A\mu + B\bar{\mu}$$

truncated linear ODE
(nonautonomous, not nec. stable)

$$\dot{\nu} = A\nu + B\varphi(\nu)$$

nonlinear approximate
moment dynamics

Assumption: 1) μ and ν remain bounded along solutions to

$$\dot{\mu}_\infty = A_\infty \mu_\infty \quad \text{and} \quad \dot{\nu} = A\nu + B\varphi(\nu)$$

2) $\dot{\mu}_\infty = A_\infty \mu_\infty$ is asymptotically stable

Theorem: $\forall\, \delta > 0 \;\exists\, N$ s.t. if $\frac{\mathrm{d}^k \mu}{\mathrm{d}t^k} = \frac{\mathrm{d}^k \nu}{\mathrm{d}t^k}, \quad \forall k \in \{1, \dots, N\}$

then

$$\|\mu(t) - \nu(t)\| \le \beta(\|\mu(t_0) - \nu(t_0)\|, t - t_0) + \delta, \qquad \forall t \ge t_0 \ge 0$$

Proof idea:
1) N derivative matches ⇒ μ & ν match on compact interval of length T
2) stability of A_∞ ⇒ matching can be extended to [0,∞)

Truncation by derivative matching

UCSB

infinite-dimensional linear ODE

☹ Given δ finding N is very difficult

☺ In practice, small values of N (e.g., N=2) already yield good results

☺ Can use $\frac{\mathrm{d}^k \mu}{\mathrm{d}t^k} = \frac{\mathrm{d}^k \nu}{\mathrm{d}t^k}, \quad \forall k \in \{1, \dots, N\}$

to determine φ(·): $k = 1$ → boundary condition on φ

$k = 2$ → PDE on φ

Theorem: $\forall\, \delta > 0 \;\exists\, N$ s.t. if $\frac{\mathrm{d}^k \mu}{\mathrm{d}t^k} = \frac{\mathrm{d}^k \nu}{\mathrm{d}t^k}, \quad \forall k \in \{1, \dots, N\}$

then

$$\|\mu(t) - \nu(t)\| \le \beta(\|\mu(t_0) - \nu(t_0)\|, t - t_0) + \delta, \qquad \forall t \ge t_0 \ge 0$$

Proof idea:
1) N derivative matches ⇒ μ & ν match on compact interval of length T
2) stability of A_∞ ⇒ matching can be extended to [0,∞)

and now something completely different…

Decaying-dimerizing molecular reactions (DDR):

$$S_1 \xrightarrow{c_1} 0 \qquad S_2 \xrightarrow{c_2} 0 \qquad 2\,S_1 \underset{c_4}{\overset{c_3}{\rightleftarrows}} S_2$$

pSHS model

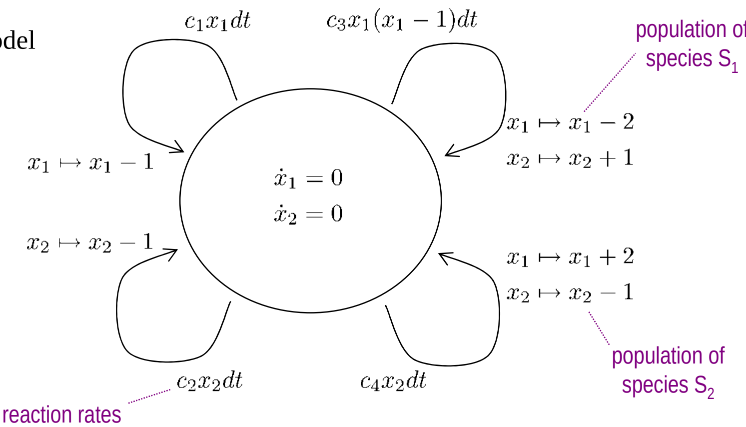

Moment dynamics for DDR

Decaying-dimerizing molecular reactions (DDR):

$$S_1 \xrightarrow{c_1} 0 \qquad S_2 \xrightarrow{c_2} 0 \qquad 2\,S_1 \underset{c_4}{\overset{c_3}{\rightleftarrows}} S_2$$

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

$$+ \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ -2c_3 & 0 \\ 0 & c_3 \\ \frac{c_3}{2} & -c_3 \end{bmatrix} \begin{bmatrix} \mu^{(3,0)} \\ \mu^{(2,1)} \end{bmatrix}$$

$\mu^{(1,0)} := \mathrm{E}[x_1]$ $\quad \mu^{(2,0)} := \mathrm{E}[x_1^2]$ $\quad \mu^{(3,0)} := \mathrm{E}[x_1^3]$

$\mu^{(0,1)} := \mathrm{E}[x_2]$ $\quad \mu^{(0,2)} := \mathrm{E}[x_2^2]$ $\quad \mu^{(2,1)} := \mathrm{E}[x_1^2 x_2$

$\mu^{(1,1)} := \mathrm{E}[x_1 x_2]$

Truncated DDR model

Decaying-dimerizing molecular reactions (DDR):

$$S_1 \xrightarrow{c_1} 0 \qquad S_2 \xrightarrow{c_2} 0 \qquad 2\,S_1 \underset{c_4}{\overset{c_3}{\rightleftarrows}} S_2$$

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

$$+ \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ -2c_3 & 0 \\ 0 & c_3 \\ \frac{c_3}{2} & -c_3 \end{bmatrix} \begin{bmatrix} \left(\frac{\mu^{(2,0)}}{\mu^{(1,0)}}\right)^3 \\ \frac{\mu^{(2,0)}}{\mu^{(0,1)}}\left(\frac{\mu^{(1,1)}}{\mu^{(1,0)}}\right)^2 \end{bmatrix}$$

by matching

$$\frac{d^k\mu}{dt^k} = \frac{d^k\nu}{dt^k}, \qquad \forall k \in \{1,2\}$$

for deterministic distributions

Hespanha. Polynomial Stochastic Hybrid Systems. HSCC'05

Monte Carlo vs. truncated model

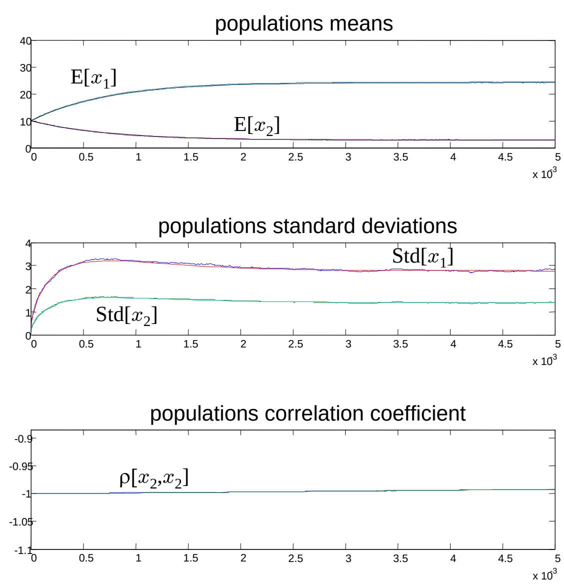

Fast time-scale transient

(lines essentially undistinguishable at this scale)

Monte Carlo vs. truncated model

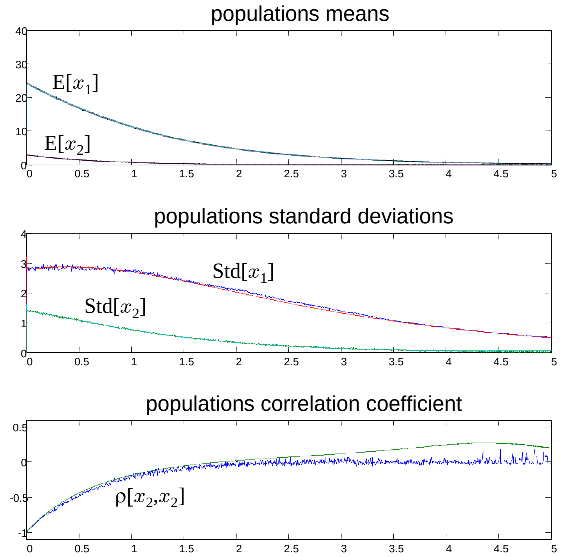

Slow time-scale evolution

only noticeable error when populations become very small (a couple of molecules)

Conclusions

1. A simple SHS model (inspired by piecewise deterministic Markov Processes) can go a long way in modeling network traffic
2. The analysis of SHSs is generally difficult but there are tools available (generator, Dynkin's equation, moment dynamics, truncations)
3. This type of SHSs (and tools) finds use in several areas (traffic modeling, networked control systems, molecular biology)

Bibliography

UCSB

- M. Davis. Markov Models & Optimization. Chapman & Hall/CRC, 1993
- S. Bohacek, J. Hespanha, J. Lee, K. Obraczka. Analysis of a TCP hybrid model. In *Proc. of the 39th Annual Allerton Conf. on Comm., Contr., and Computing*, Oct. 2001.
- S. Bohacek, J. Hespanha, J. Lee, K. Obraczka. 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)*, June 2003.
- J. Hespanha. Stochastic Hybrid Systems: Applications to Communication Networks. In Rajeev Alur, George J. Pappas, *Hybrid Systems: Computation and Control*, number 2993 in Lect. Notes in Comput. Science, pages 387-401, Mar. 2004.
- Hespanha. Polynomial Stochastic Hybrid Systems. To be presented at the HSCC'05.
- J. Hespanha, A. Singh. Stochastic Models for Chemically Reacting Systems Using Polynomial Stochastic Hybrid Systems. Submitted to the *Int. J. on Robust Control Special Issue on Control at Small Scales*, Nov. 2004.

All papers (and some ppt presentations) available at
http://www.ece.ucsb.edu/~hespanha