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

Deterministic Hybrid Systems

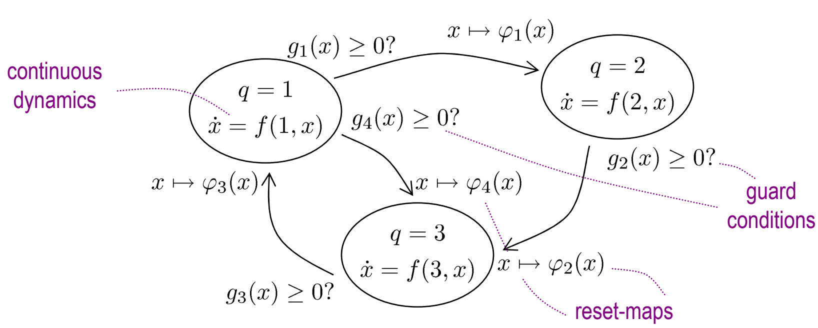

$q(t) \in \mathcal{Q}=\{1,2,\ldots\}$ $\equiv$ discrete state
$x(t) \in \mathbb{R}^n$ $\equiv$ continuous state

$\Big\}$ 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

UCSB

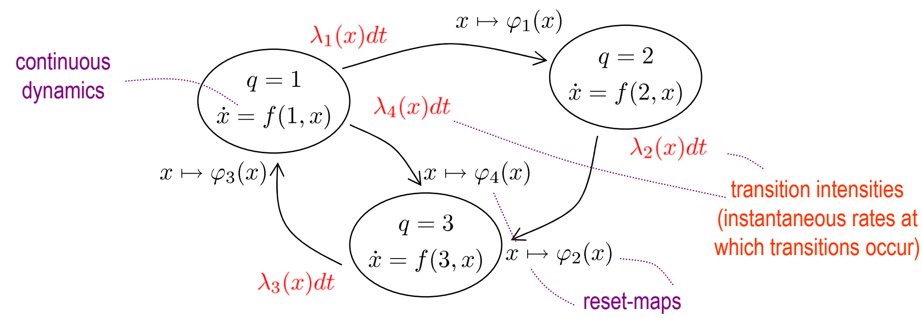

$N_\ell(t) \in \mathbb{N} \equiv$ transition counter, which is incremented by one each time the ℓth reset-map $\varphi_\ell(x)$ is "activated" } right-continuous by convention

$$\mathrm{E}\Big[\underbrace{N_\ell(t_2) - N_\ell(t_1)}_{\text{number of transitions on } (t_1, t_2]}\Big] = \mathrm{E}\Big[\int_{t_1}^{t_2} \lambda_\ell(x(t))\, dt\Big]$$

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

Stochastic Hybrid Systems

UCSB

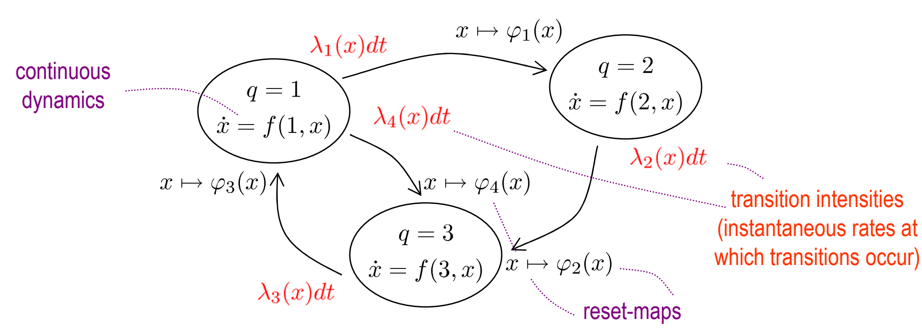

Special case: When all λ_ℓ are constant, transitions are controlled by a continuous-time Markov process

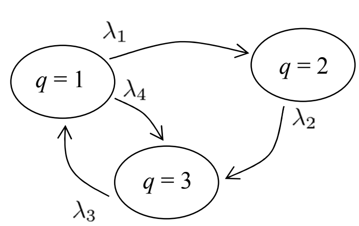

specifies q (independently of x)

Formal model—Summary

UCSB

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

Continuous dynamics:

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

Transition intensities:

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

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

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

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, 1983)

Example I: TCP congestion control

UCSB

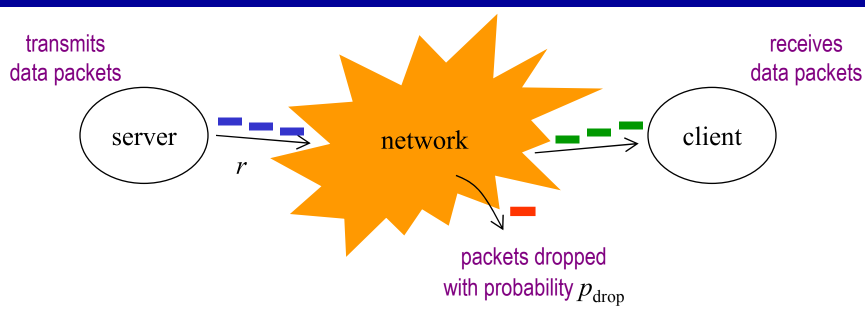

congestion control $\equiv$ selection of the rate r at which the server transmits packets
feedback mechanism $\equiv$ packets are dropped by the network to indicate congestion

Example I: TCP congestion control

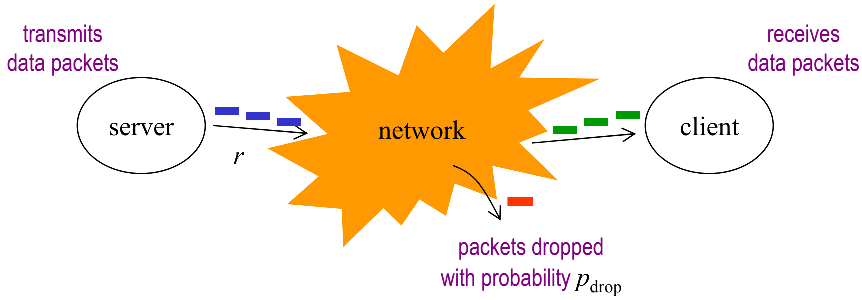

congestion control $\equiv$ selection of the rate r at which the server transmits packets
feedback mechanism $\equiv$ packets are dropped by the network to indicate congestion

TCP (Reno) congestion control: packet sending rate given by

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

congestion window (internal state of controller)
round-trip-time (from server to client and back)

- initially w is set to 1
- until first packet is dropped, w increases exponentially fast (slow-start)
- after first packet is dropped, w increases linearly (congestion-avoidance)
- each time a drop occurs, w is divided by 2 (multiplicative decrease)

Example I: TCP congestion control

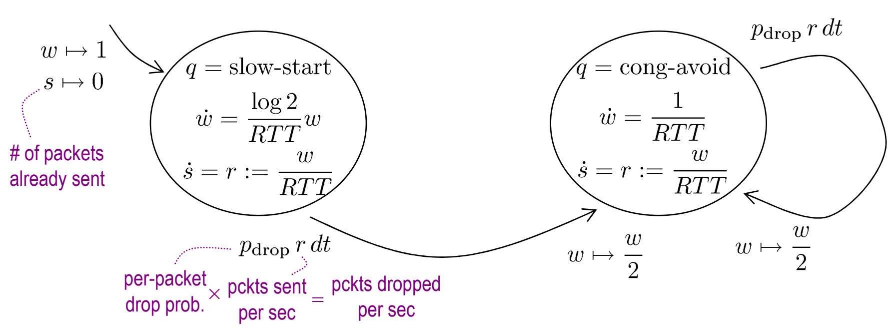

TCP (Reno) congestion control: packet sending rate given by

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

congestion window (internal state of controller)
round-trip-time (from server to client and back)

- initially w is set to 1
- until first packet is dropped, w increases exponentially fast (slow-start)
- after first packet is dropped, w increases linearly (congestion-avoidance)
- each time a drop occurs, w is divided by 2 (multiplicative decrease)

Analysis—Lie Derivative

UCSB

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

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

UCSB

$\dot{x} = f(q,x,t)$	$\lambda_\ell(q,x,t)$	$(q,x) = \phi_\ell(q^-,x^-,t)$
continuous dynamics	transition intensities	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

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

$L_f\psi$
Lie derivative of ψ

L completely defines the SHS dynamics

disclaimer: see paper for technical assumptions

Example I: TCP congestion control

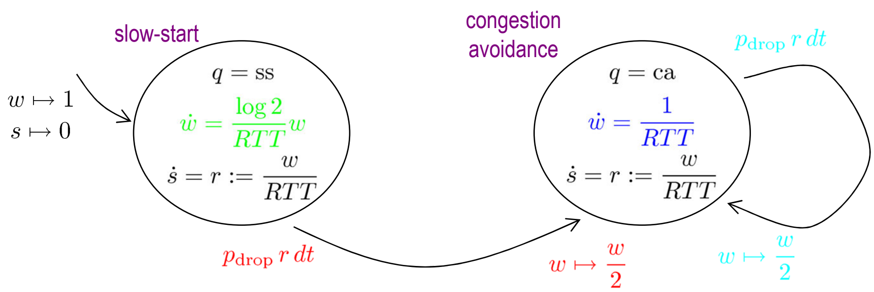

$$(L\psi)(q,w,t) = \begin{cases} \dfrac{\partial \psi}{\partial w}\dfrac{\log 2}{RTT} w + \dfrac{\partial \psi}{\partial t} + \left[\psi\left(\mathrm{ca}, \dfrac{w}{2}\right) - \psi(\mathrm{ss}, w)\right] \dfrac{p_{\mathrm{drop}} w}{RTT} & q = \mathrm{ss} \\ \dfrac{\partial \psi}{\partial w}\dfrac{1}{RTT}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

This TCP model is *always-on*

Example II: On-off TCP model

- 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

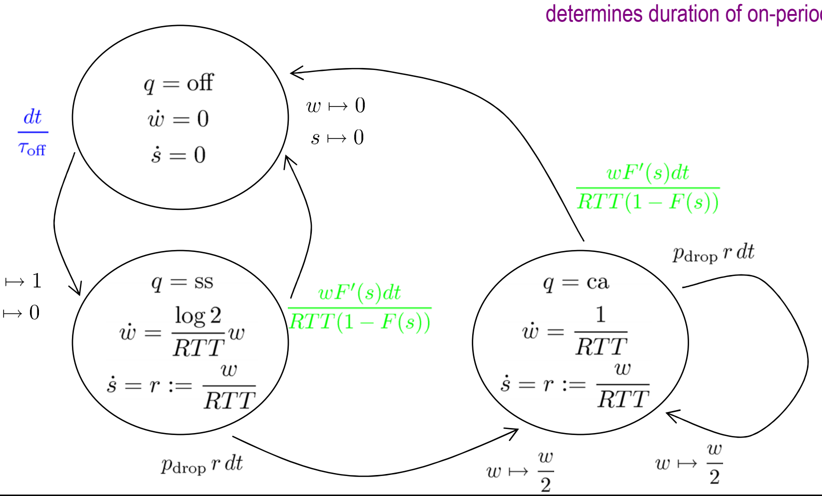

Using the SHS generator…

UCSB

Mixture of exponential distribution for transfer-sizes:

M exp. distr. with mean $\{ k_1, k_2, \ldots, k_M \}$

M probabilities $\{ p_1, p_2, \ldots, p_M \}$ $\quad \sum_{i=1}^{M} p_i = 1$

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

can approximate well distributions found in the literature …

$\mu_{q,n} \equiv$ expected value of r^n when SHS is in mode q

probability of being in off mode

$$\dot{\mu}_{\text{off},0} = -\frac{\mu_{\text{off},0}}{\tau_{\text{off}}} + \sum_{j=1}^{M} k_j^{-1}(\mu_{\text{ss}_j,1} + \mu_{\text{ca}_j,1})$$

slow-start with file-size from ith exponential distribution

$$\dot{\mu}_{\text{ss}_i,n} = \frac{p_i w_0^n \mu_{\text{off},0}}{\tau_{\text{off}} RTT^n} + n\frac{(\log 2) - n_{\text{ack}}\dot{RTT}}{n_{\text{ack}} RTT}\mu_{\text{ss}_i,n} - (p_{\text{drop}} + k_i^{-1})\mu_{\text{ss}_i,n+1}$$

$$\dot{\mu}_{\text{ca}_i,n} = \frac{n\mu_{\text{ca}_i,n-1}}{n_{\text{ack}} RTT^2} - \frac{n\dot{RTT}\mu_{\text{ca}_i,n}}{RTT} - (p_{\text{drop}} + k_i^{-1})\mu_{\text{ca}_i,n+1} + \frac{p_{\text{drop}}}{2^n}(\mu_{\text{ss}_i,n+1} + \mu_{\text{ca}_i,n+1})$$

cong.-avoidance with file-size from ith exponential distribution

Using the SHS generator…

UCSB

Mixture of exponential distribution for transfer-sizes:

M exp. distr. with mean $\{ k_1, k_2, \ldots, k_M \}$

M probabilities $\{ p_1, p_2, \ldots, p_M \}$ $\quad \sum_{i=1}^{M} p_i = 1$

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

ca … *infinite system of ODEs*

But:
distributions are well approximated by log-Normal
so one can truncate the infinite system of ODEs
(keeping only first 2 moments)

$\mu_{q,n} \equiv$ expected value of

off mode

$$\dot{\mu}_{\text{off},0} = -\frac{\mu_{\text{off},0}}{\tau_{\text{off}}} + \sum_{j=1}^{M} k_j^{-1}(\mu_{\text{ss}_j,1} + \mu_{\text{ca}_j,1})$$

slow-start w… from ith exponential … n

$$\dot{\mu}_{\text{ss}_i,n} = \frac{p_i w_0^n \mu_{\text{off},0}}{\tau_{\text{off}} RTT^n} + n\frac{(\log 2) - n_{\text{ack}}\dot{RTT}}{n_{\text{ack}} RTT}\mu_{\text{ss}_i,n} - (p_{\text{drop}} + k_i^{-1})\mu_{\text{ss}_i,n+1}$$

$$\dot{\mu}_{\text{ca}_i,n} = \frac{n\mu_{\text{ca}_i,n-1}}{n_{\text{ack}} RTT^2} - \frac{n\dot{RTT}\mu_{\text{ca}_i,n}}{RTT} - (p_{\text{drop}} + k_i^{-1})\mu_{\text{ca}_i,n+1} + \frac{p_{\text{drop}}}{2^n}(\mu_{\text{ss}_i,n+1} + \mu_{\text{ca}_i,n+1})$$

cong.-avoidance with file-size from ith exponential distribution

Case I

UCSB

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

$p_1 = 88.87\%$ $k_1 = 3.5\text{KB}$ ("mice" files)

$p_2 = 11.13\%$ $k_2 = 246\text{KB}$ ("elephant" files, but not heavy tail)

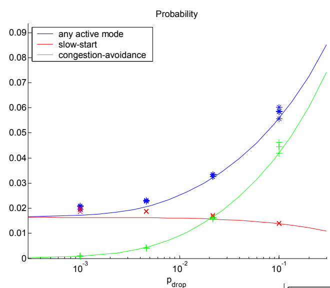

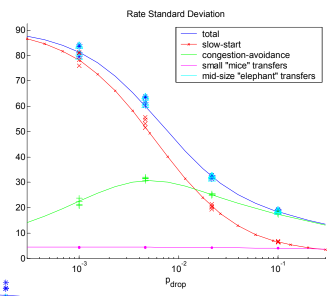

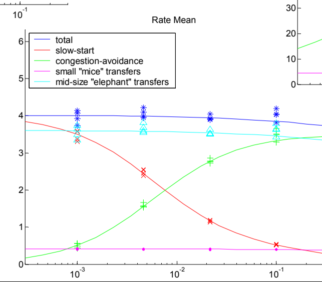

Monte-Carlo
& theoretical
steady-state values
(vs. drop probability)

Case II

UCSB

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

$p_1 = 98\%$ $k_1 = 6\text{KB}$ ("mice" files)

$p_2 = 1.7\%$ $k_2 = 400\text{KB}$ ("elephant" files)

$p_3 = .02\%$ $k_3 = 10\text{MB}$ ("mammoth" files)

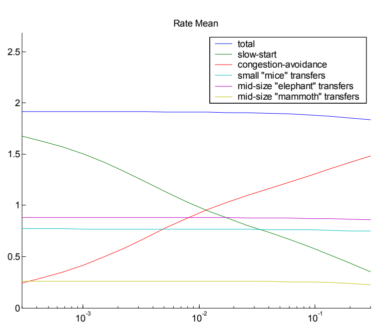

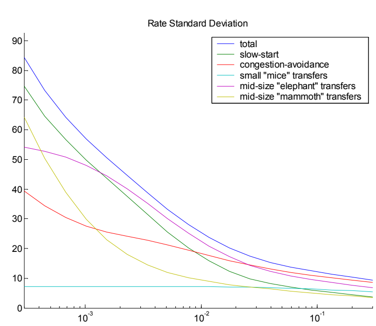

Steady-state values
vs.
drop probability

Conclusions

1. Presented a new model for SHSs with transitions triggered by stochastic events (inspired by piecewise deterministic Markov Processes)
2. Provided analysis tool—SHS's generator
3. Illustrated use by modeling and analyzing on-off TCP flows

Current/future work

1. Investigate the use of SHS to extend solutions of deterministic hybrid systems
2. Develop other tools of analyze SHS
3. Investigate the *implications of the TCP-Reno* results on active queuing mechanism
4. Construct SHS models for *other congestion control algorithms*
5. Use SHS models to study problems of distributed *control with limited communication bandwidth*

Generalizations

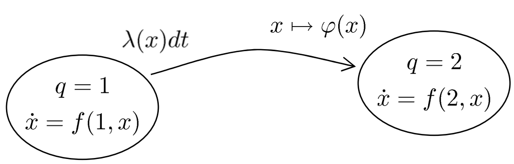

1. *Stochastic resets* can be obtained by considering multiple intensities/reset-maps

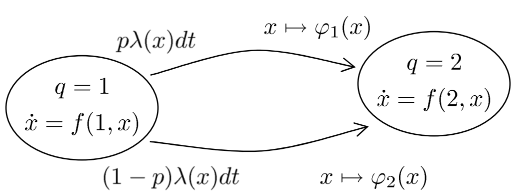

$$x \mapsto \begin{cases} \varphi_1(x) & \text{w.p. } p \\ \varphi_2(x) & \text{w.p. } 1-p \end{cases}$$

One can further generalize this to resets governed by to a continuous distribution [see paper]

$$x \sim \mu(q, x, dx)$$

Generalizations

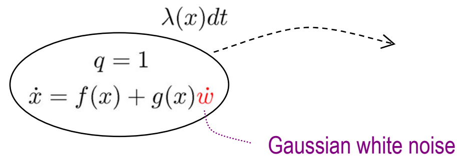

2. *Stochastic differential equations (SDE) for the continuous state* can be emulated by taking limits of SHSs

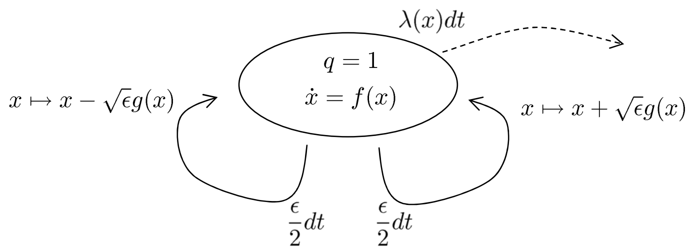

The solution to the SDE is obtained as $\epsilon \to 0^+$

more details on paper…

Generalizations

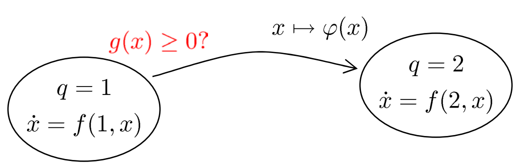

3. *Deterministic guards* can also be emulated by taking limits of SHSs

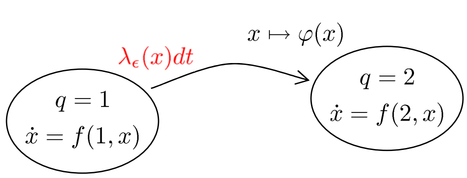

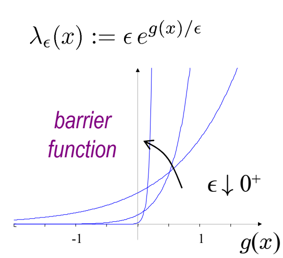

The solution to the hybrid system with a deterministic guard is obtained as $\epsilon \downarrow 0^+$

This provides a mechanism to regularize systems with chattering and/or Zeno phenomena…

Example III: Bouncing-ball

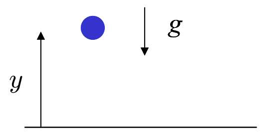

$y \leq 0, \ \dot{y} < 0?$

$\ddot{y} = -g$

$\dot{y} \mapsto -c\dot{y}$

$c \in (0,1) \equiv$ energy absorbed at impact

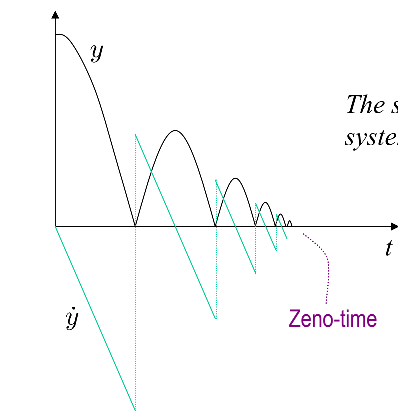

The solution of this deterministic hybrid system is only defined up to the Zeno-time

Example III: Bouncing-ball

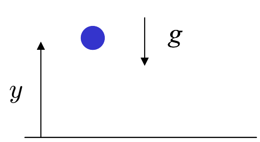

$\lambda_\epsilon(y, \dot{y})dt$

$\ddot{y} = -g$

$\dot{y} \mapsto -c\dot{y}$

$$\lambda_\epsilon(y, \dot{y}) := \begin{cases} \epsilon e^{-y/\epsilon} & \dot{y} < 0 \\ 0 & \dot{y} > 0 \end{cases}$$

$c \in (0,1) \equiv$ energy absorbed at impact

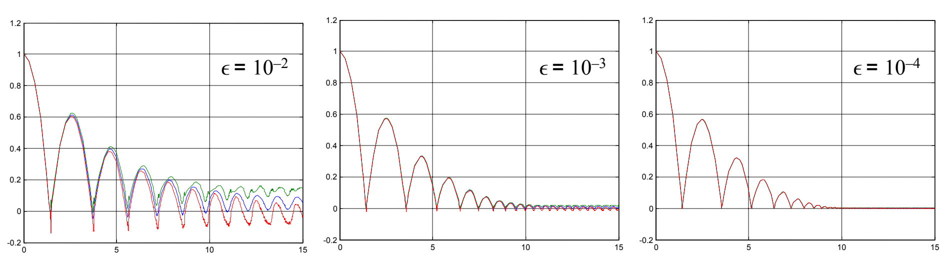

mean (blue) 95% confidence intervals (red and green)