Research Supported by NSF & ARO

Stochastic
Modeling, Analysis, and Design
of
Networked Control Systems

João Hespanha

UC SANTA BARBARA
engineering

Networked control systems

UC SANTA BARBARA
engineering

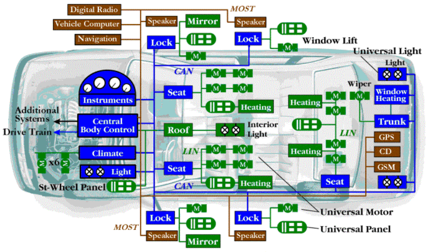

Efficiency and safety in cars depend on a network of hundreds of ECUs (power train, ABS, stability control, speed control, transmission, …)

Robotic agents free humans from unpleasant, dangerous, and/or repetitive tasks in which human performance would degrade over time due to fatigue

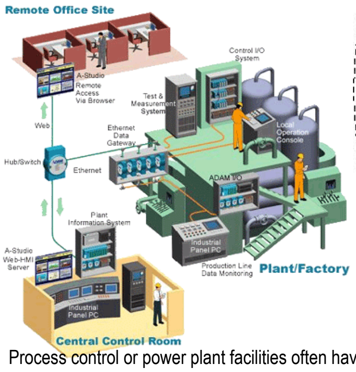

Process control or power plant facilities often have between several thousand of coupled control loops

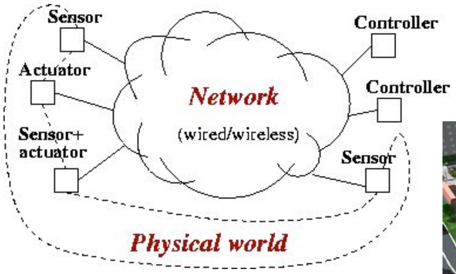

Buildings consume 72% of electricity, 40% of all energy, and produce close to 50% of U.S. carbon emissions

Three Key Ideas

- Importance of Stochastic Modeling versus Worst Case in NCSs
- Analysis & Design Results Available for Stochastic NCSs
  - Time-driven SHSs
  - Lyapunov-based methods
- (Stochastic) Control Tools for NCS Protocol-Design

(ex) students: D. Antunes (IST), A. Mesquita (UCSB), Y. Xu (Advertising.com)

collaborators: C. Silvestre (IST)

acknowledgements: NSF, Institute for Collaborative bio-technologies (ARO)

disclaimer: This is an overview, technical details in papers referenced in bottom right corner… http://www.ece.ucsb.edu/~hespanha

---

Example #1: Networked Control System

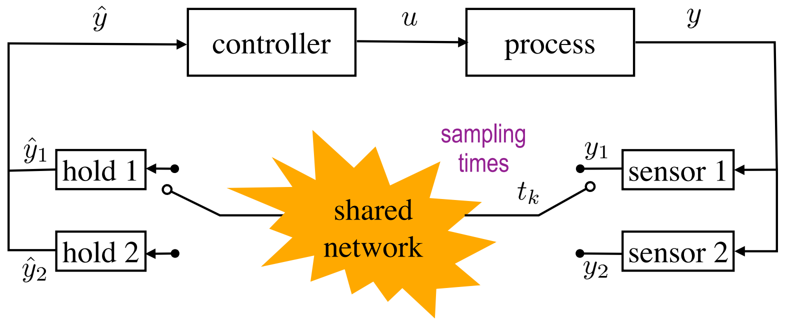

process: $$\dot{x}_P = A_P x_P + C_P u$$ $$y = C_P x_P + D_P u$$

controller: $$\dot{x}_C = A_C x_C + C_C \hat{y}$$ $$\hat{y} = C_C x_C + D_C \hat{y}$$

round-robin network access:

$$\dot{\hat{y}} = 0$$

hold

$$\begin{bmatrix} \hat{y}_1(t_k) \\ \hat{y}_2(t_k) \end{bmatrix} = \begin{cases} \begin{bmatrix} y_1(t_k^-) \\ \hat{y}_2(t_k^-) \end{bmatrix} & k \text{ odd} \\ \begin{bmatrix} \hat{y}_1(t_k^-) \\ y_2(t_k^-) \end{bmatrix} & k \text{ even} \end{cases}$$

Example #1: Networked Control System

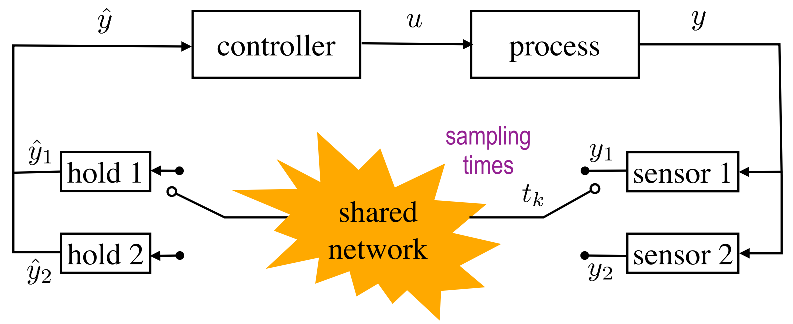

process: $\dot{x}_P = A_P x_P + C_P u$

$y = C_P x_P + D_P u$

controller: $\dot{x}_C = A_C x_C + C_C \hat{y}$

$\hat{y} = C_C x_C + D_C \hat{y}$

What if the network is not available at a sample time t_k ?

1st wait until network becomes available

2nd send (old) data from original sampling of continuous-time output

or

2nd send (latest) data from current sampling of continuous-time output

⇒ intersampling times $t_{k+1} - t_k$ typically become random variables

---

Example #1: Networked Control System

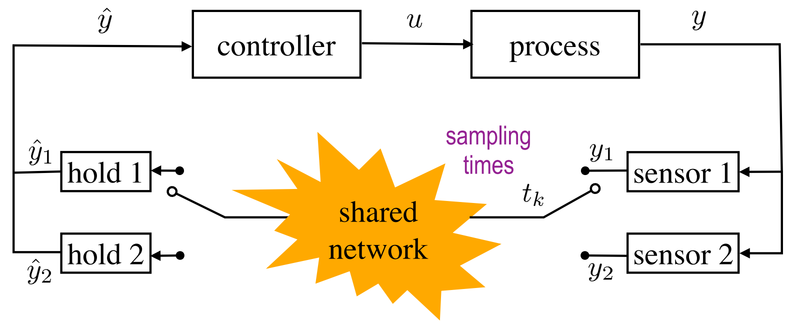

Typical results:

Deterministic Modeling	Stochastic Modeling
• $t_{k+1} - t_k \in [0, T], \forall k$	• $t_{k+1} - t_k$ iid random variables
• worst-case sequence	• $t_{k+1} - t_k$ unif. distri. in $[0, T]$
• stability guaranteed for $T < .0279$	• stability guaranteed for $T < .112$
(suff. condition for stability)	(nec. & suff. condition for stability)

Example #1: Networked Control System

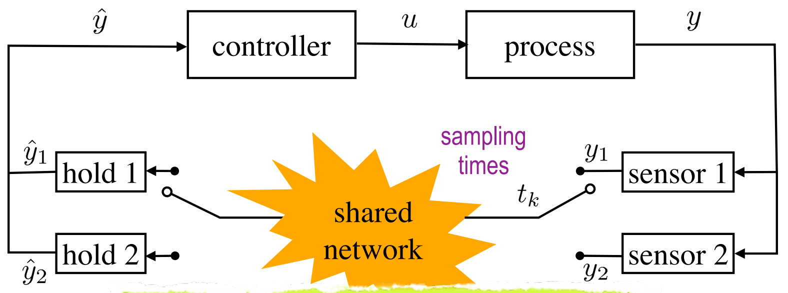

How to model/analyze such systems?

Stochastic Hybrid Systems

Deterministic

- $t_{k+1} - t_k \in [0, T]$, $\forall k$
- worst-case sequence
- stability guaranteed for $T < .0279$

(suff. condition for stability)

deling

- $t_{k+1} - t_k$ iid random variables
- $t_{k+1} - t_k$ unif. distri. in $[0, T]$
- stability guaranteed for $T < .112$

(nec. & suff. condition for stability)

Deterministic Hybrid Systems

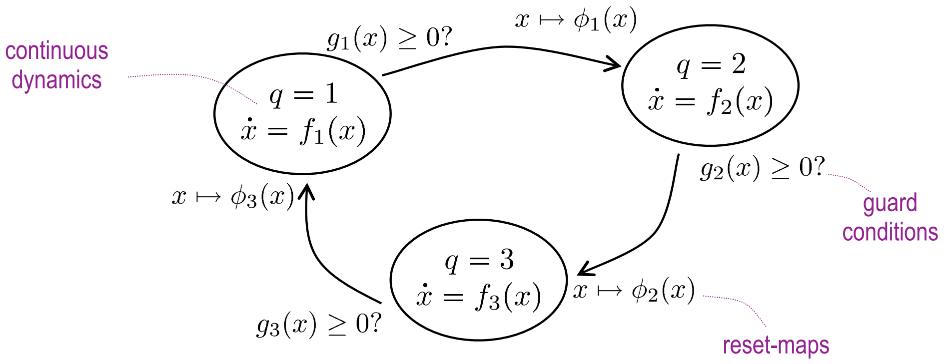

$q(t) \in \mathcal{Q}=\{1,2,\ldots\}$ $\equiv$ discrete state

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

Deterministic Hybrid Systems

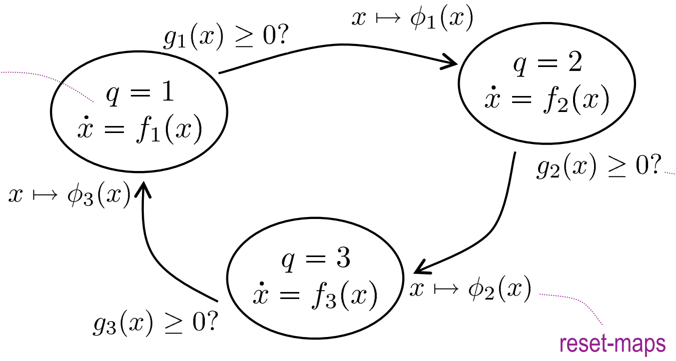

continuous dynamics

guard conditions

reset-maps

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

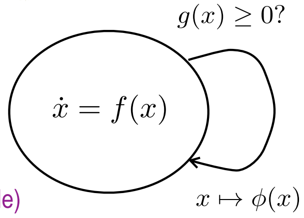

impulsive system
(single discrete mode)

Stochastic Hybrid Systems (time-triggered)

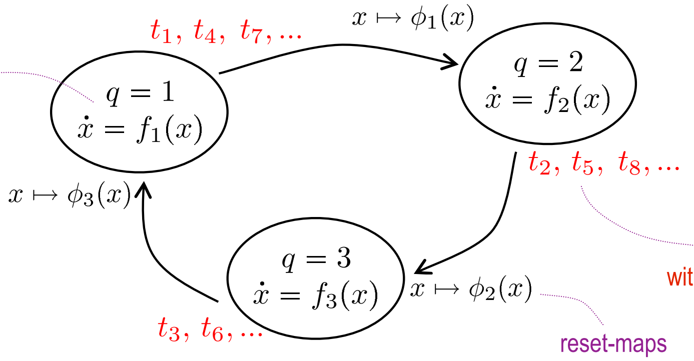

continuous dynamics

transition times
$t_{k+1} - t_k$ i.i.d.
with given distribution

reset-maps

Also known as SHSs *driven by renewal processes*:

$N(t) \equiv$ # of transitions before time t

renewal process
(iid inter-increment times)

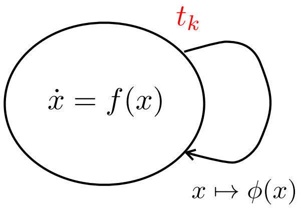

stochastic impulsive system (SIS)
(single discrete mode)

Stochastic Hybrid Systems (time-triggered)

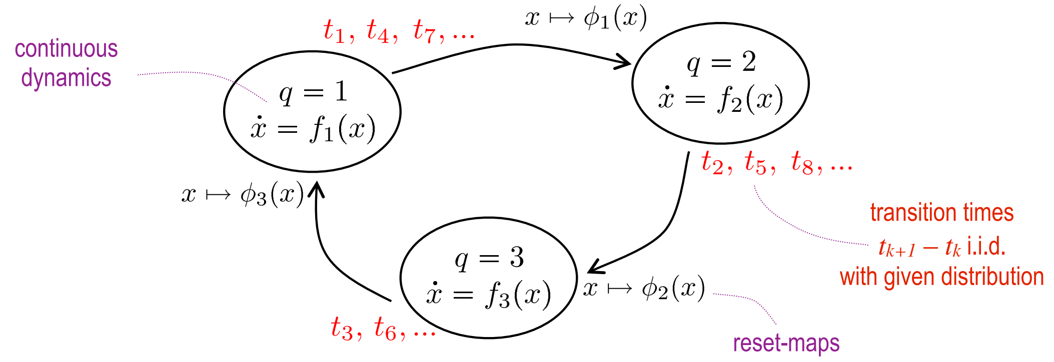

Also kno

$N(t) \equiv$

re
(iid int

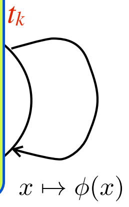

Special case: when $t_{k+1} - t_k$ i.i.d. exponentially distributed

- called Markovian Jump Systems
- in this case $x(t)$ is a Markov Process
- well developed theory (analysis & design)

[Costa, Fragoso, Boukas, Loparo, Lee, Dullerud]

Example #1: Networked Control System

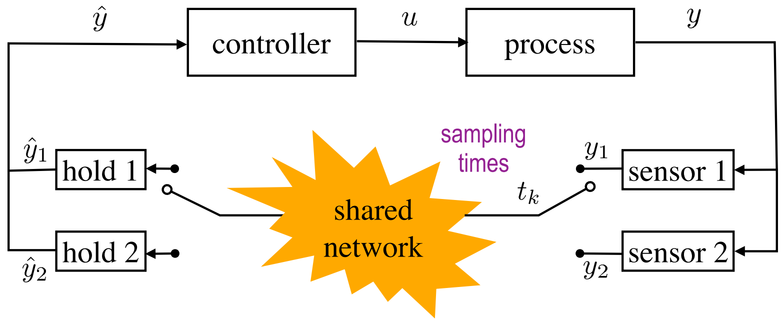

round-robin network access:

$$\dot{\hat y} = 0$$

hold

$$\begin{bmatrix}\hat y_1(t_k)\\ \hat y_2(t_k)\end{bmatrix} = \begin{cases} \begin{bmatrix} y_1(t_k^-)\\ \hat y_2(t_k^-)\end{bmatrix} & k \text{ odd}\\[1em] \begin{bmatrix} \hat y_1(t_k^-)\\ y_2(t_k^-)\end{bmatrix} & k \text{ even}\end{cases}$$

- $t_{k+1} - t_k$ iid random variables
- $t_{k+1} - t_k$ unif. distri. in $[0, T]$

Example #1: Networked Control System

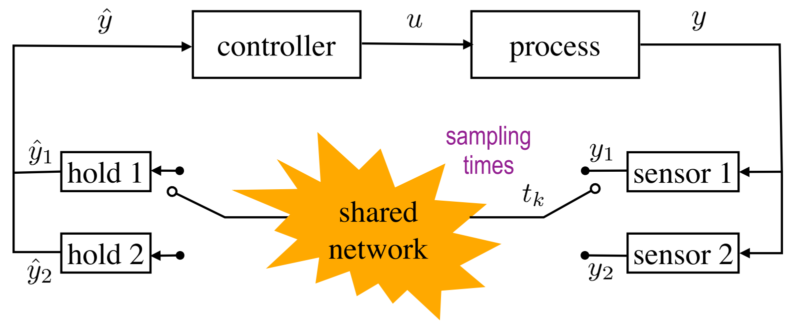

$t_{k+1} - t_k$ ~ time-interval between successive transmissions

$$\begin{bmatrix}\hat{y}_1(t_k)\\ \hat{y}_2(t_k)\end{bmatrix} = \begin{bmatrix}y_1(t_k^-)\\ \hat{y}_2(t_k^-)\end{bmatrix}$$

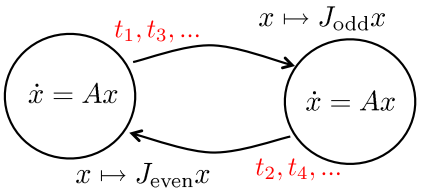

$$x := \begin{bmatrix}x_P\\ x_C\\ \hat{y}\end{bmatrix}$$

$$\begin{bmatrix}\hat{y}_1(t_k)\\ \hat{y}_2(t_k)\end{bmatrix} = \begin{bmatrix}\hat{y}_1(t_k^-)\\ y_2(t_k^-)\end{bmatrix}$$

Stability of Linear Time-triggered SIS

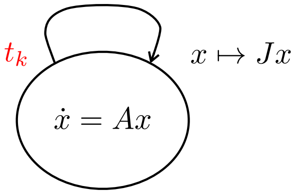

stochastic impulsive system
(single discrete mode)

$t_{k+1} - t_k$ ~ i.i.d., with cumulative distribution function $F(\cdot)$

Defining $x_k := x(t_k)$

state at jump times

$$x_{k+1} = \underbrace{J}_{\text{reset}}\underbrace{e^{A(t_{k+1}-t_k)}}_{\text{continuous evolution}}x_k$$

Stability of Linear Time-triggered SIS

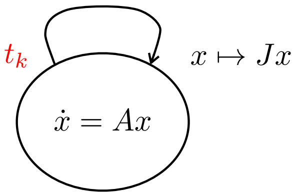

stochastic impulsive system
(single discrete mode)

$t_{k+1} - t_k \sim$ i.i.d., with cumulative distribution function $F(\cdot)$

Defining $x_k := x(t_k)$

state at jump times

$$x_{k+1} = \underbrace{J}_{\text{reset}} \underbrace{e^{A(t_{k+1}-t_k)}}_{\text{continuous evolution}} x_k$$

For a given $P = P' > 0$

$$\mathrm{E}[x_{k+1}' P x_{k+1} \mid x_k] = x_k' \, \mathrm{E}_{F(\boldsymbol{\Delta})} \left[\underbrace{e^{A'\boldsymbol{\Delta}} J' P J e^{A\boldsymbol{\Delta}}}_{} \right] x_k$$

expectation w.r.t. $\boldsymbol{\Delta} = t_{k+1} - t_k$
(cumulative distribution F)

Stability of Linear Time-triggered SIS

stochastic impulsive system
(single discrete mode)

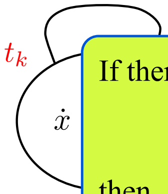

If there exists

LMI on $P_{n \times n}$

$$P > 0, \quad \mathrm{E}_{F(\boldsymbol{\Delta})} \left[e^{A'\boldsymbol{\Delta}} J' P J e^{A\boldsymbol{\Delta}} \right] < P$$

then

$$\lim_{k \to \infty} \mathrm{E}[\|x_k\|^2] = 0 \quad \text{(exp. fast in index } k\text{)}$$

What about $x(t)$ between jumps?

Defining x_k

For a given $P = P' > 0$

$$\mathrm{E}[x_{k+1}' P x_{k+1} \mid x_k] = x_k' \, \mathrm{E}_{F(\boldsymbol{\Delta})} \left[\underbrace{e^{A'\boldsymbol{\Delta}} J' P J e^{A\boldsymbol{\Delta}}}_{} \right] x_k$$

expectation w.r.t. $\boldsymbol{\Delta} = t_{k+1} - t_k$
(cumulative distribution F)

Stability of Linear Time-triggered SIS

All stability notions require $\lim_{k\to\infty} \|x_k\| = 0$ exponentially fast

the nec. & suff. conditions only differ on the requirements on the tail of distribution

$$1 - F(s) = \mathrm{P}(t_{k+1} - t_k > s)$$

- Mean-square exponential stability, i.e., $\lim_{t\to\infty} \mathrm{E}[\|x(t)\|^2] \overset{\text{exp. fast}}{=} 0$

- Mean-square asymptotic stability, i.e., $\lim_{t\to\infty} \mathrm{E}[\|x(t)\|^2] = 0$

- Mean-square stochastic stability, i.e., $\int_0^\infty \mathrm{E}[\|x(t)\|^2]dt < \infty$

Stability of Linear Time-triggered SIS

All stability notions require $\lim_{k\to\infty} \|x_k\| = 0$ exponentially fast

the nec. & suff. conditions only differ on the requirements on the tail of distribution

$$1 - F(s) = \mathrm{P}(t_{k+1} - t_k > s)$$

(versions of these results for multiple discrete modes are available)

Theorem:

- Mean-square exponential stability, i.e., $\lim_{t\to\infty} \mathrm{E}[\|x(t)\|^2] \overset{\text{exp. fast}}{=} 0$

$$\Leftrightarrow \quad \exists P > 0,\ \mathrm{E}_{F(\boldsymbol{\Delta})}\left[e^{A'\boldsymbol{\Delta}} J' P J e^{A\boldsymbol{\Delta}}\right] < P \text{ and } \lim_{s\to\infty} e^{A's} e^{As} \big(1 - F(s)\big) \overset{\text{exp. fast}}{=} 0$$

- Mean-square asymptotic stability, i.e., $\lim_{t\to\infty} \mathrm{E}[\|x(t)\|^2] = 0$

$$\Leftrightarrow \quad \exists P > 0,\ E_{F(\boldsymbol{\Delta})}[e^{A'\boldsymbol{\Delta}} J' P J e^{A\boldsymbol{\Delta}}] < P \text{ and } \lim_{s\to\infty} e^{A's} e^{As} \big(1 - F(s)\big) = 0$$

- Mean-square stochastic stability, i.e., $\int_0^\infty \mathrm{E}[\|x(t)\|^2]dt < \infty$

$$\Leftrightarrow \quad \exists P > 0,\ \mathrm{E}_{F(\boldsymbol{\Delta})}\left[e^{A'\boldsymbol{\Delta}} J' P J e^{A\boldsymbol{\Delta}}\right] < P \text{ and } \int_0^\infty e^{A's} e^{As} F(ds) < \infty$$

[Antunes *et al*, 2009]

Three Key Ideas

- Importance of Stochastic Modeling versus Worst Case in NCSs
- Analysis & Design Results Available for Stochastic NCSs
  - Time-driven SHSs
  - Event-driven SHSs
- (Stochastic) Control Tools for NCS Protocol-Design

(ex) students: D. Antunes (IST), A. Mesquita (UCSB), Y. Xu (Advertising.com)

collaborators: C. Silvestre (IST)

acknowledgements: NSF, Institute for Collaborative bio-technologies (ARO)

disclaimer: This is an overview, technical details in papers referenced in bottom right corner… http://www.ece.ucsb.edu/~hespanha

---

So far time-driven SHSs…

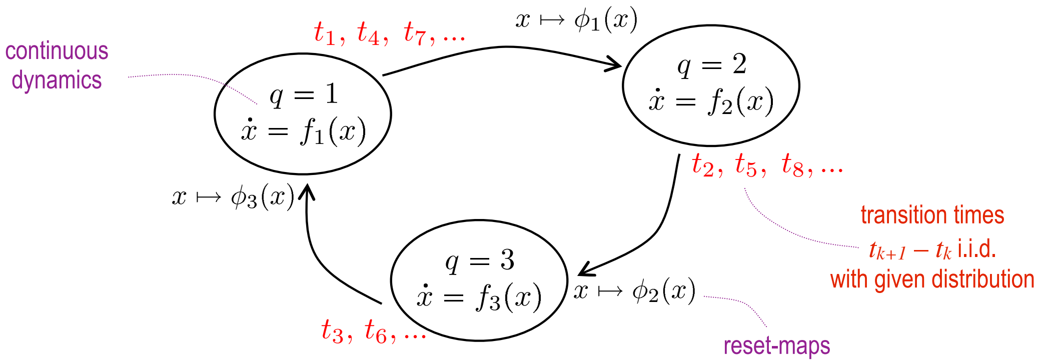

continuous dynamics

transition times $t_{k+1} - t_k$ i.i.d. with given distribution

reset-maps

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

Stochastic Hybrid Systems

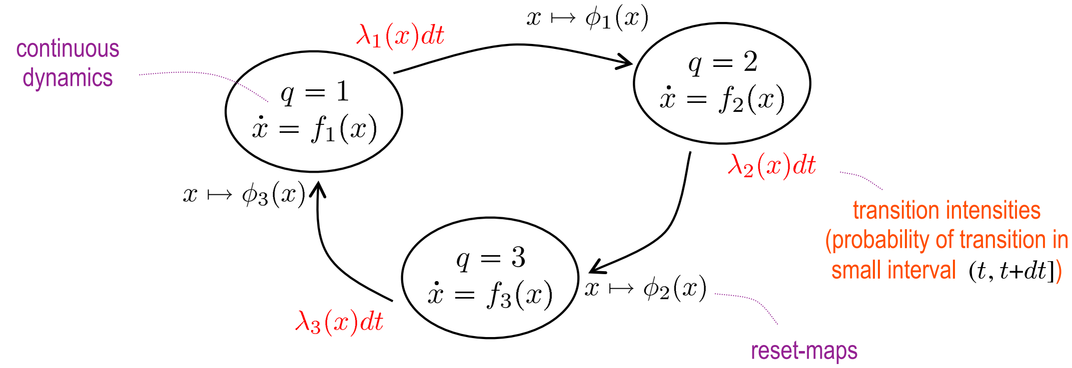

continuous dynamics

$q = 1$, $\dot{x} = f_1(x)$

$\lambda_1(x)dt$ $x \mapsto \phi_1(x)$

$q = 2$, $\dot{x} = f_2(x)$

$\lambda_2(x)dt$ $x \mapsto \phi_2(x)$

$q = 3$, $\dot{x} = f_3(x)$

$\lambda_3(x)dt$ $x \mapsto \phi_3(x)$

transition intensities (probability of transition in small interval $(t, t+dt]$)

reset-maps

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

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

Special case: Markovian jump system

When all λ_ℓ are constant $\Rightarrow$ time triggered SHS with exponential $t_{k+1} - t_k$

Stochastic Hybrid Systems with Diffusion

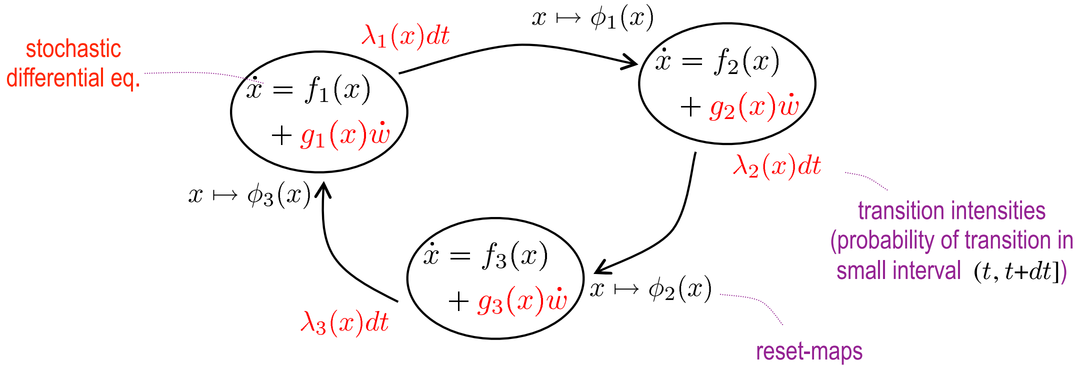

stochastic differential eq.

$\dot{x} = f_1(x) + g_1(x)\dot{w}$

$\lambda_1(x)dt$ $x \mapsto \phi_1(x)$

$\dot{x} = f_2(x) + g_2(x)\dot{w}$

$\lambda_2(x)dt$ $x \mapsto \phi_2(x)$

$\dot{x} = f_3(x) + g_3(x)\dot{w}$

$\lambda_3(x)dt$ $x \mapsto \phi_3(x)$

transition intensities (probability of transition in small interval $(t, t+dt]$)

reset-maps

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

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

Example #2: Estimation through network

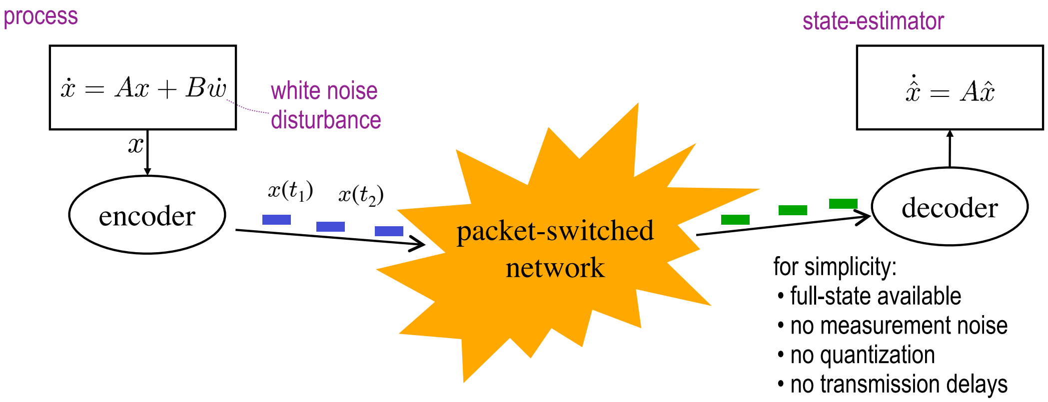

process

$\dot{x} = Ax + B\dot{w}$

white noise
disturbance

x

encoder

$x(t_1)$ $x(t_2)$

packet-switched
network

decoder

state-estimator

$\dot{\hat{x}} = A\hat{x}$

for simplicity:
- full-state available
- no measurement noise
- no quantization
- no transmission delays

encoder logic $\equiv$ determines *when* to send measurements to the network
decoder logic $\equiv$ determines *how* to incorporate received measurements

---

Stochastic communication logic

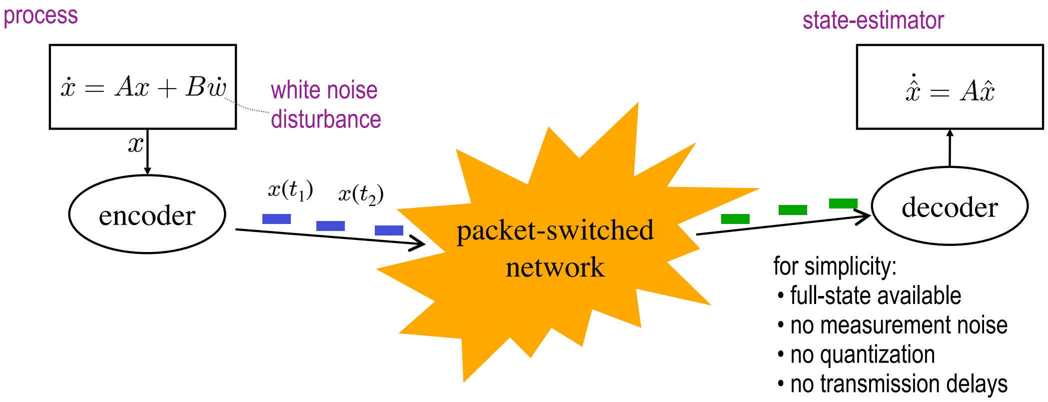

process

$\dot{x} = Ax + B\dot{w}$

white noise
disturbance

x

encoder

$x(t_1)$ $x(t_2)$

packet-switched
network

decoder

state-estimator

$\dot{\hat{x}} = A\hat{x}$

for simplicity:
- full-state available
- no measurement noise
- no quantization
- no transmission delays

decoder logic $\equiv$ determines *how* to incorporate received measurements

1. upon reception of $x(t_k)$, reset $\hat{x}(t_k)$ to $x(t_k)$

encoder logic $\equiv$ determines *when* to send measurements to the network

1. keep track of remote estimate $\hat{x}$
2. send measurements stochastically
3. probability of sending data increases as $\hat{x}$ deviates from x

[related ideas pursued by Astrom, Basar, Hristu, Kumar, Tilbury]

Error Dynamics

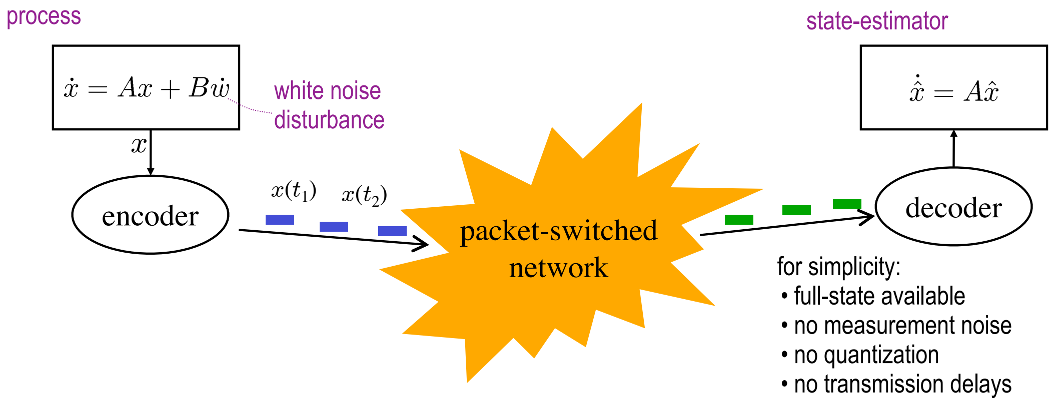

process

$\dot{x} = Ax + B\dot{w}$

white noise disturbance

x

encoder

$x(t_1)$ $x(t_2)$

packet-switched network

decoder

state-estimator

$\dot{\hat{x}} = A\hat{x}$

for simplicity:

- full-state available
- no measurement noise
- no quantization
- no transmission delays

Error dynamics: $e := x - \hat{x}$

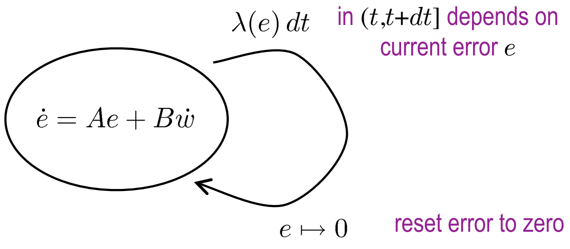

$\lambda(e)\, dt$

prob. of sending data in $(t,t+dt]$ depends on current error e

$\dot{e} = Ae + B\dot{w}$

$e \mapsto 0$

reset error to zero

ODE – Lie Derivative

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

Given scalar-valued function $V : \mathbb{R}^n \to \mathbb{R}$

$$\boxed{\underbrace{\frac{dV\big(x(t)\big)}{dt}}_{\text{derivative along solution to ODE}} = \underbrace{\frac{\partial V\big(x(t)\big)}{\partial x} f\big(x(t)\big)}_{\substack{L_f V \\ \text{Lie derivative of } V}}}$$

Basis of “Lyapunov” formal arguments to establish boundedness and stability…

E.g., picking $V(x) := \|x\|^2$

$$\frac{dV\big(x(t)\big)}{dt} = \frac{\partial V}{\partial x} f(x) \leqslant 0 \quad \Rightarrow \quad V\big(x(t)\big) = \|x(t)\|^2 \leqslant \|x(0)\|^2$$

$\|x\|^2$ remains bounded along trajectories !

Generator of a Stochastic Hybrid System

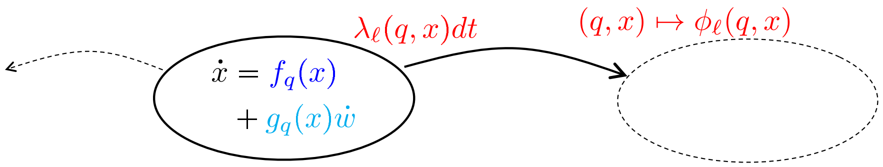

Given scalar-valued function $V : \mathcal{Q} \times \mathbb{R}^n \to \mathbb{R}$

$$\frac{d}{dt}\,\mathrm{E}\Big[V\big(q(t),x(t)\big)\Big] = E\Big[(LV)\big(q(t),x(t)\big)\Big]$$

x & q are discontinuous, but the expected value is differentiable

Dynkin's formula (in differential form)

where

$$(LV)(q,x) \coloneqq \frac{\partial V}{\partial x}(q,x) f_q(x)$$

Lie derivative

(extended) generator of the SHS

$$+ \sum_{\ell=1}^{m} \underset{\text{intensity}}{\lambda_\ell(q,x)} \Big(\overbrace{V\big(\phi_\ell(q,x)\big) - V(q,x)}^{\text{instantaneous variation}}\Big)$$

Reset term

$$+ \frac{1}{2}\,\mathrm{trace}\left(g_q(x)' \frac{\partial^2 V}{\partial x^2} g_q(x)\right)$$

Diffusion term

Lyapunov Analysis – SHSs

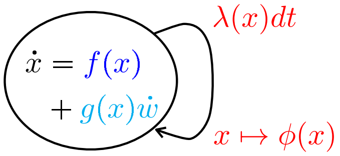

$$\frac{d}{dt}\,\mathrm{E}\Big[V\big(x(t)\big)\Big] = E\Big[(LV)\big(x(t)\big)\Big]$$

sample-path notions

class-K functions: (zero at zero & mon. increasing)

probability of $\|x(t)\|$ exceeding any given bound M, can be made arbitrarily small by making $\|x_0\|$ small

$$\begin{cases} \alpha_1(\|x\|) \le V(x) \le \alpha_2(\|x\|) \\ LV(x) \le -\alpha_3(\|x\|) \end{cases} \Rightarrow \begin{cases} \mathrm{P}\big(\exists t : \|x(t)\| \ge M\big) \le \frac{\alpha_2(\|x_0\|)}{\alpha_1(M)} \\ \mathrm{P}\big(x(t) \to 0\big) = 1 \end{cases}$$

almost sure (a.s.) asymptotic stability

expected-value notions

$$\begin{cases} V(x) \ge 0 \\ LV(x) \le -W(x) \end{cases} \Rightarrow \int_0^\infty \mathrm{E}\Big[W\big(x(t)\big)\Big]dt < \infty$$

stochastic stability (mean square when $W(x) = \|x\|^2$)

$$\begin{cases} V(x) \ge W(x) \ge 0 \\ LV(x) \le -\mu V + c \end{cases} \Rightarrow \mathrm{E}\Big[W\big(x(t)\big)\Big] \le e^{-\mu t} V(x_0) + \frac{c}{\mu}$$

exponential stability (mean square when $W(x) = \|x\|^2$)

Example #2: Remote estimation

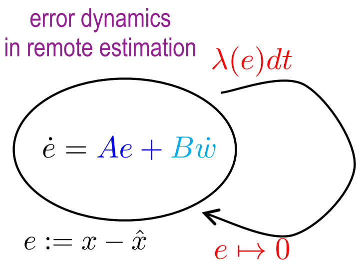

Dynkin's formula

$$\frac{d}{dt}\mathrm{E}\Big[V\big(e(t)\big)\Big] = E\Big[(LV)\big(e(t)\big)\Big]$$

$$(LV)(e) := \frac{\partial V}{\partial e}Ae + \lambda(e)\Big(V(0) - V(e)\Big) + \frac{1}{2}\,\mathrm{trace}\Big(B'\frac{\partial^2 V}{\partial e^2}B\Big)$$

For constant rate: $\lambda(e) = \gamma$ (exp. distributed inter-jump times) using $V(e) = e'Pe$

1. $\mathrm{E}[\,e\,] \to 0$ if and only if $\gamma > \Re\big[\lambda_i(A)\big],\ \forall i$
2. $\mathrm{E}[\,\|e\|^m\,]$ bounded if and only if $\gamma > m\,\Re\big[\lambda_i(A)\big],\ \forall i$

getting more moments bounded requires higher comm. rates

For radially unbounded rate: $\lambda(e)$ (reactive transmissions) using $V(e) = \|e\|^2$

3. $\mathrm{E}[\,e\,] \to 0$ (always)
4. $\mathrm{E}[\,\|e\|^m\,]$ bounded $\forall\, m$

Moreover, one can achieve the same $\mathrm{E}[\,\|e\|^2\,]$ with less communication than with a constant rate or periodic transmissions…

[Xu *et al*, 2006]

Back to Time-triggered SIS…

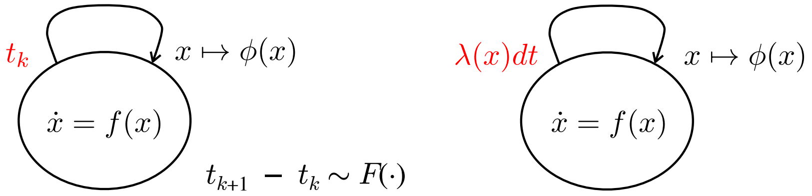

Can we pick an intensity $\lambda(\cdot)$ to obtain the desired distribution for the t_k ?

YES

Back to Time-triggered SIS…

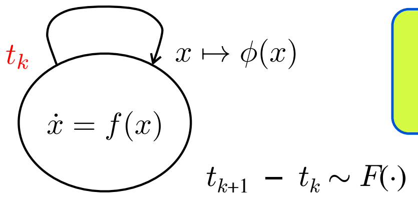

This representation allows one to *combine in the same SHS* time- and event-triggered transitions!

Can we pick an intensity $\lambda(\cdot)$ to obtain the desired distribution for the t_k ?

YES

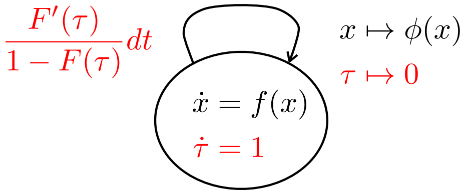

time since last reset

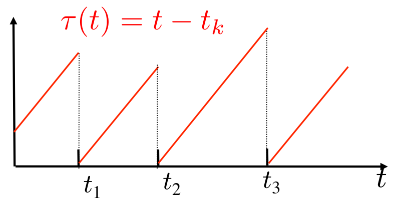

the aggregate state (x,τ) is a Markov process

Converse Lyapunov Stability

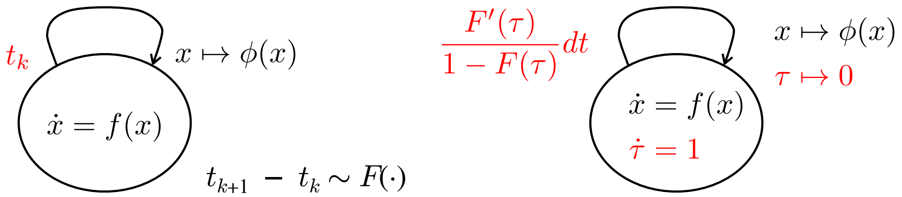

Theorem:

System is mean-square exponentially stable, i.e., $\lim_{t\to\infty} \mathrm{E}[\|x(t)\|^2] \overset{\text{exp. fast}}{=} 0$

$$\Updownarrow$$

$\exists P(\tau)$ such that defining $V(x,\tau) = x'P(\tau)x$

Lyapunov-like function quadratic on x for fixed τ

$$\begin{cases} c_1 I \le P(\tau) \le c_2 I & \Rightarrow \quad V \text{ is positive definite} \\ (LV)(x,\tau) \le -\epsilon V(x,\tau) & \Rightarrow \quad \frac{\mathrm{d}}{\mathrm{d}t}\mathrm{E}[V(x,\tau)] \le -\epsilon\,\mathrm{E}[V(x,\tau)] \end{cases}$$

(motivates choices for Lyapunov function for nonlinear systems)

Three Key Ideas

- Importance of Stochastic Modeling versus Worst Case in NCSs
- Analysis & Design Results Available for Stochastic NCSs
  - Time-driven SHSs
  - Event-driven SHSs
- (Stochastic) Control Tools for NCS Protocol-Design

(ex) students: D. Antunes (IST), A. Mesquita (UCSB), Y. Xu (Advertising.com)

collaborators: C. Silvestre (IST)

acknowledgements: NSF, Institute for Collaborative bio-technologies (ARO)

disclaimer: This is an overview, technical details in papers referenced in bottom right corner… http://www.ece.ucsb.edu/~hespanha

---

Transport layer protocols

Most common (general purpose) protocols:

high drop rates can lead to poor performance and eventually instability

UDP

- no attempt at error correction
- no attempt to control data rate

TCP

- error correction
  - all packets sent should be acknowledged by receiver
  - lack of acknowledgement of packet *n* leads to retransmission of same packet after packet *n* + 3 (triple duplicate ack mechanism)
- congestion control
  - packet drops are taken as a sign of con rease

delayed retransmissions are essentially useless; too much overhead in ack every packet

Illustrative 1-D problem

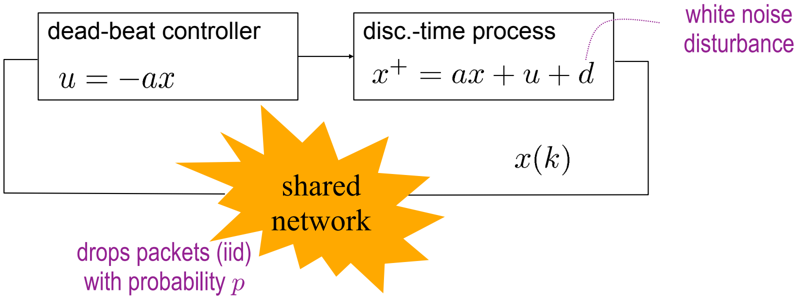

dead-beat controller

$u = -ax$

disc.-time process

$x^+ = ax + u + d$

white noise
disturbance

$x(k)$

shared
network

drops packets (iid)
with probability p

The closed-loop is *mean-square stable* (i.e., E[$x(k)^2$] <∞) if and only if

$$p < \frac{1}{|a|^2}$$

(it is also straightforward to compute a tight asymptotic bound on E[$x(k)^2$])

Intuition: ignoring the disturbance d

$$x(k+1)^2 = \begin{cases} 0 & \text{with probability } 1-p \\ |a|^2 x(k)^2 & \text{with probability } p \end{cases} \quad \Rightarrow \quad \mathrm{E}[x(k+1)^2] = p\,|a|^2\,\mathrm{E}[x(k)^2]$$

Illustrative 1-D problem

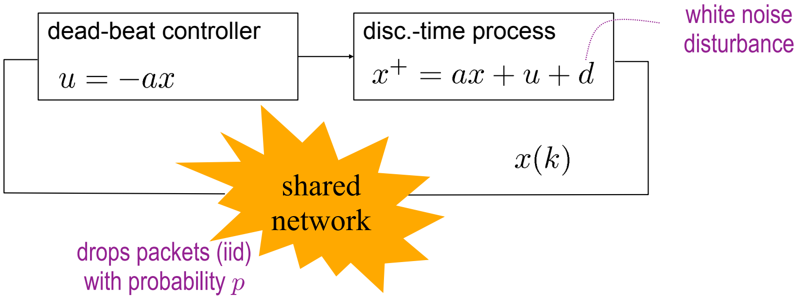

dead-beat controller

$u = -ax$

disc.-time process

$x^+ = ax + u + d$

white noise
disturbance

$x(k)$

shared
network

drops packets (iid)
with probability p

The closed-loop is *mean-square stable* (i.e., E[$x(k)^2$] <∞) if and only if

$$p < \frac{1}{|a|^2}$$

(it is also straightforward to compute a tight asymptotic bound on E[$x(k)^2$])

But what if |a|>1 and the probability of drop is larger than this bound?

Redundant transmissions

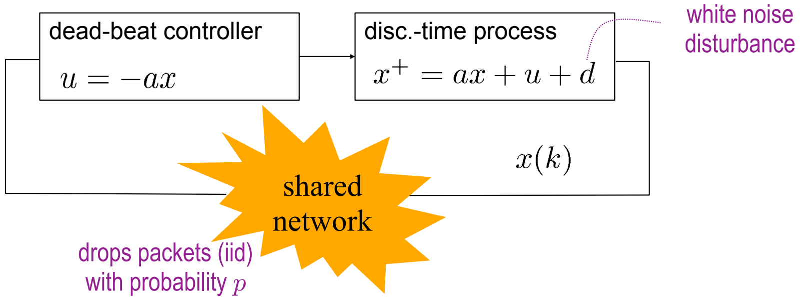

redundant transmissions ≡ at each time step one sends N copies of $x(k)$ through independent channels (time, frequency, or spatial diversity), each with drop probability p

The closed-loop is *mean-square stable* (i.e., E[$x(k)^2$] <∞) if and only if

$$p^N < \frac{1}{|a|^2} \quad \Leftrightarrow \quad p < \frac{1}{|a|^{\frac{2}{N}}}$$

any drop probability can be accommodated by choosing N sufficiently large

but transmission rate is *N times larger*

A simple "error-correction" protocol

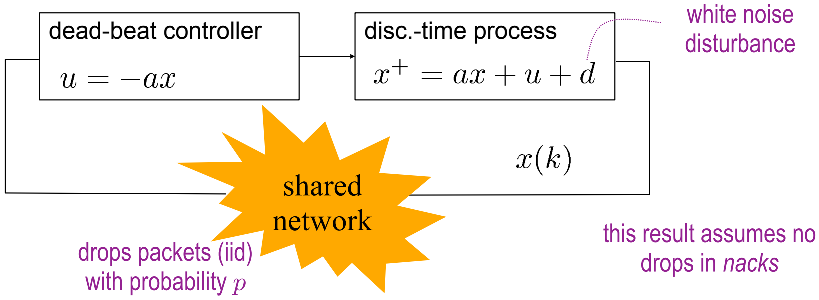

this result assumes no drops in *nacks*

1. when a packet is lost, receiver sends a "negative acknowledgement" (*nack*)
2. transmitter generally sends *one* packet at each sampling time, however…
3. upon reception of *nack*, transmitter sends *two* copies of the following packet

The closed-loop is *mean-square stable* (i.e., E[$x(k)^2$] <∞) if and only if

$$p < \frac{1}{|a|}$$

similar bound as if always sending two packets

but average transmission rate is only *1+O(p) times larger*

[ACC'09]

Even better…

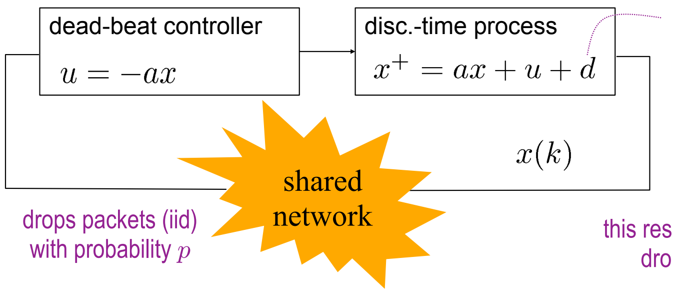

white noise disturbance

drops packets (iid) with probability p

this result assumes no drops in *nacks*

Pick a function $v : \mathbb{N} \to \mathbb{N}$, with $v(0) = 1$

1. when a packet is lost, receiver sends a "negative acknowledgement" (*nack*)
2. transmitter keeps track of number $\ell(k)$ of consecutive drops prior to time k

- transmitter sends $v\big(\ell(k)\big)$ copies of each packet

For every p, a, and N, one can find a function $v : \mathbb{N} \to \mathbb{N}$ such that

- closed-loop is *mean-square stable* (i.e., E[$x(k)^2$] < ∞)
- average transmission rate is only *1+O(p^N) times larger*
- requires at least N independent channels

stabilizes any system

arbitrarily small increase in the transmission rate

all but one channel are rarely utilized

Even better…

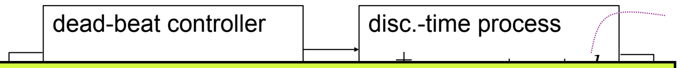

white noise disturbance

- can stabilize any system for any drop probability
- with arbitrarily small increase in the transmission rate

no (completely) free lunch… E[$x(k)^2$] will be large

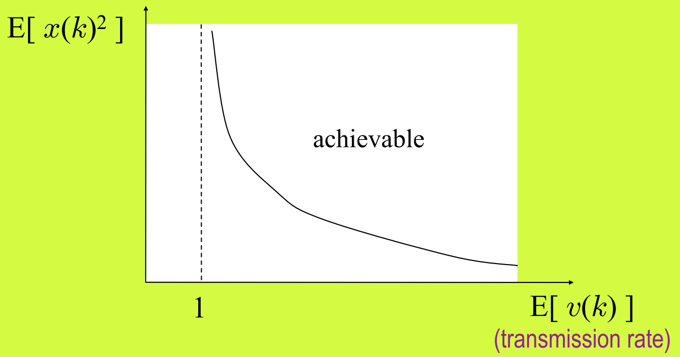

(transmission rate)

ult assumes no
ps in *nacks*

Pick a fu
1. when a … *nack*)
2. receiv … k

- transm

For ever

- closed
- average transmission rate is only *1+O(p^N) times larger*
- requires at least N independent channels

ny system

arbitrarily small increase in the transmission rate

all but one channel are rarely utilized

Optimal “error-correction” protocols

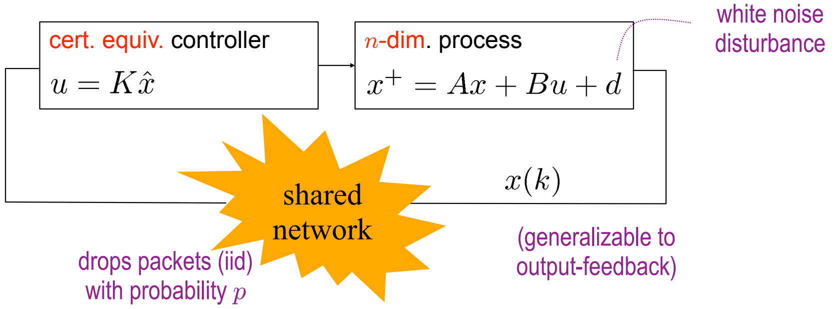

choose $v(k) \equiv$ number of copies of $x(k)$ to send at time instant k

to minimize

$$\lim_{N\to\infty} \underbrace{\left(\frac{1}{N}\,\mathrm{E}\left[\sum_{k=0}^{N-1}\|x(k)-\hat{x}(k)\|^2\right]\right)}_{\substack{\text{state estimation error}\\ \text{(performance)}}} + \lambda \underbrace{\left(\frac{1}{N}\,\mathrm{E}\left[\sum_{k=0}^{N-1} v(k)\right]\right)}_{\substack{\text{transmission rate}\\ \text{(communication)}}}$$

average-cost optimal control of a Markov process on $\mathbb{R}^n$

Optimal “error-correction” protocols

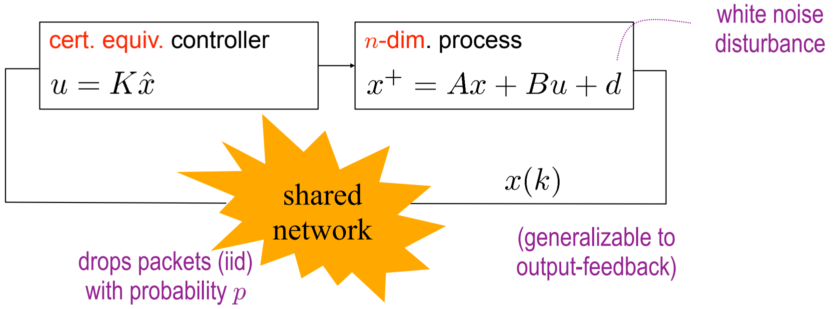

$$\lim_{N\to\infty} \left(\frac{1}{N}\,\mathrm{E}\left[\sum_{k=0}^{N-1}\|x(k)-\hat{x}(k)\|^2\right]\right) + \lambda \left(\frac{1}{N}\,\mathrm{E}\left[\sum_{k=0}^{N-1} v(k)\right]\right)$$

Theorem:

- optimal $v(k)$ is generated by a memoryless policy of the form

$$v(k) = \pi^*\big(x(k) - \hat{x}(k)\big)$$

transmitter must construct a state estimate to determine optimal $v(k)$

- optimal policy π^* can be computed using dynamic programming and value-iteration

computationally difficult for large n

Example

UC SANTA BARBARA
engineering

$$x^+ = 2x + u + w$$
$$w \sim N(0,3)$$
$$p = .15$$

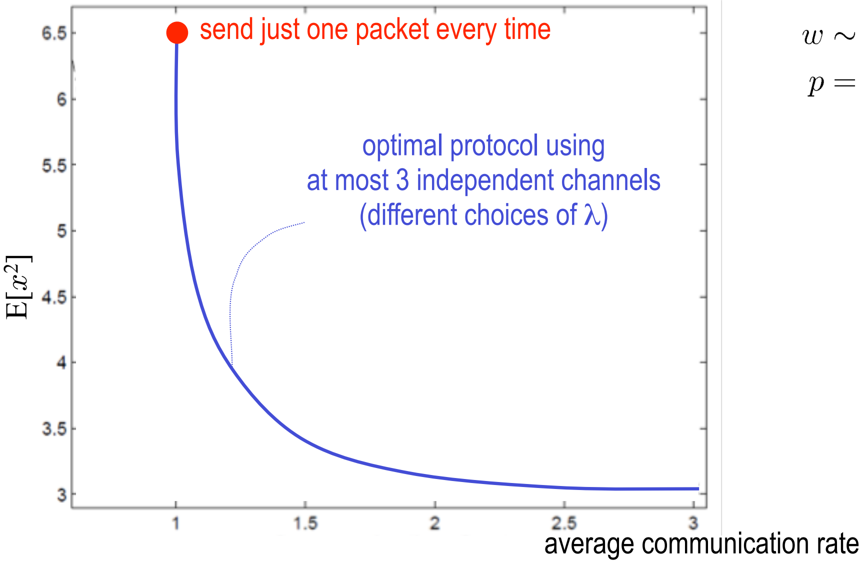

send just one packet every time

optimal protocol using
at most 3 independent channels
(different choices of λ)

Optimal “simplified” protocols

UC SANTA BARBARA
engineering

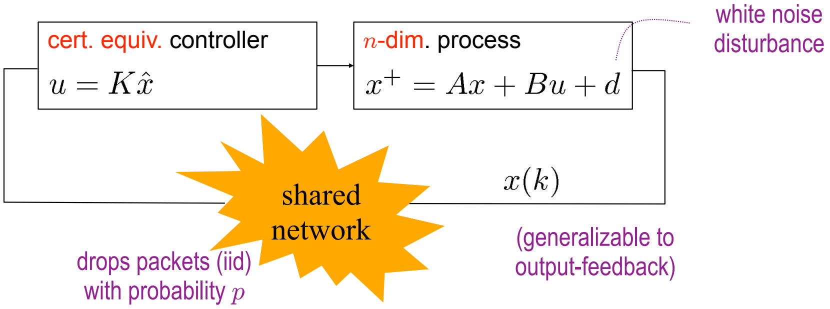

cert. equiv. controller
$u = K\hat{x}$

n-dim. process
$x^+ = Ax + Bu + d$

white noise
disturbance

shared
network

$x(k)$

drops packets (iid)
with probability p

(generalizable to
output-feedback)

choose $v(k) \equiv$ number of copies of $x(k)$ to send at time instant k

to minimize

$$\lim_{N\to\infty} \underbrace{\left(\frac{1}{N} \mathrm{E}\left[\sum_{k=0}^{N-1} \|x(k) - \hat{x}(k)\|^2 \right] \right)}_{\text{state estimation error (performance)}} + \lambda \underbrace{\left(\frac{1}{N} \mathrm{E}\left[\sum_{k=0}^{N-1} v(k) \right] \right)}_{\text{transmission rate (communication)}}$$

but transmitter must choose $v(k)$ based only on # of consecutive drops (from nacks)

Optimal “simplified” protocols

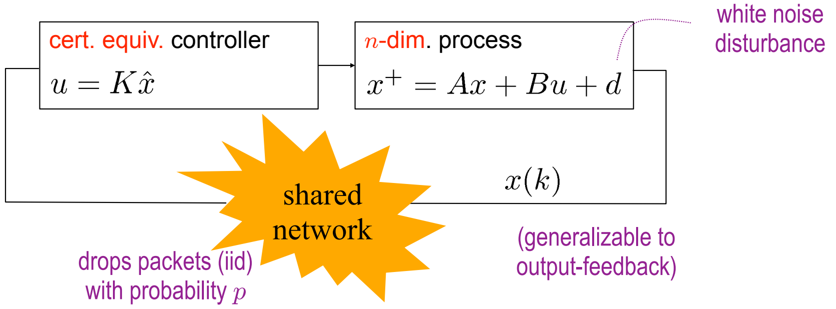

white noise disturbance

drops packets (iid) with probability p

(generalizable to output-feedback)

$$\lim_{N\to\infty}\left(\frac{1}{N}\,\mathrm{E}\left[\sum_{k=0}^{N-1}\|x(k)-\hat{x}(k)\|^2\right]\right)+\lambda\left(\frac{1}{N}\,\mathrm{E}\left[\sum_{k=0}^{N-1}v(k)\right]\right)$$

Theorem:

- optimal $v(k)$ is generated by a memoryless policy of the form

$$v(k) = \pi^*\big(\ell(k)\big)$$

transmitter only needs to keep track of $\ell(k) \equiv$ # of consecutive drops (from nacks)

- optimal policy π^* can be computed using dynamic programming and value-iteration

computationally much easier (optimization on countable-state MDP with size independent of n)

Example

$$x^+ = 2x + u + w$$
$$w \sim N(0,3)$$
$$p = .15$$

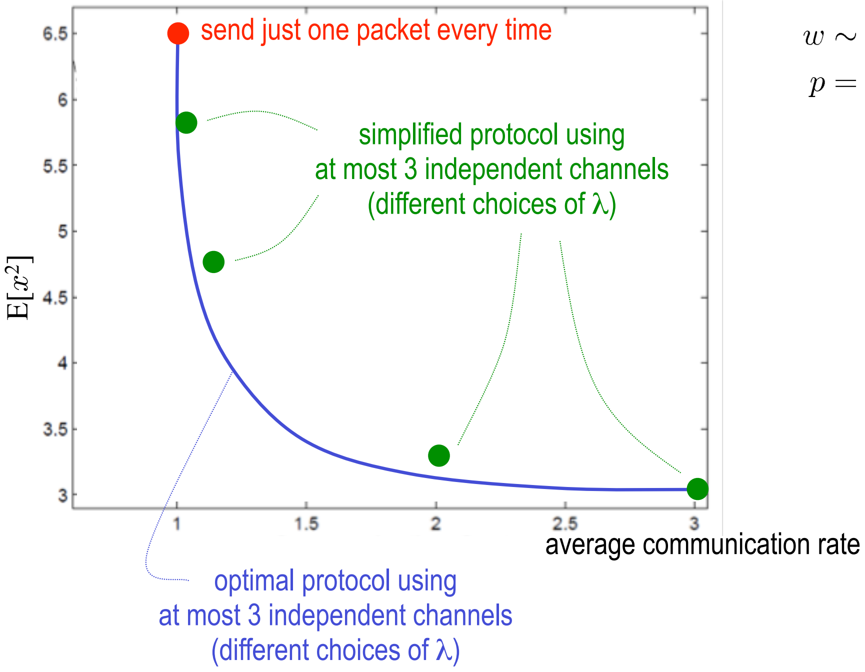

send just one packet every time

simplified protocol using at most 3 independent channels (different choices of λ)

optimal protocol using at most 3 independent channels (different choices of λ)

Example

$$x^+ = 2x + u + w$$
$$w \sim N(0,3)$$
$$p = .15$$

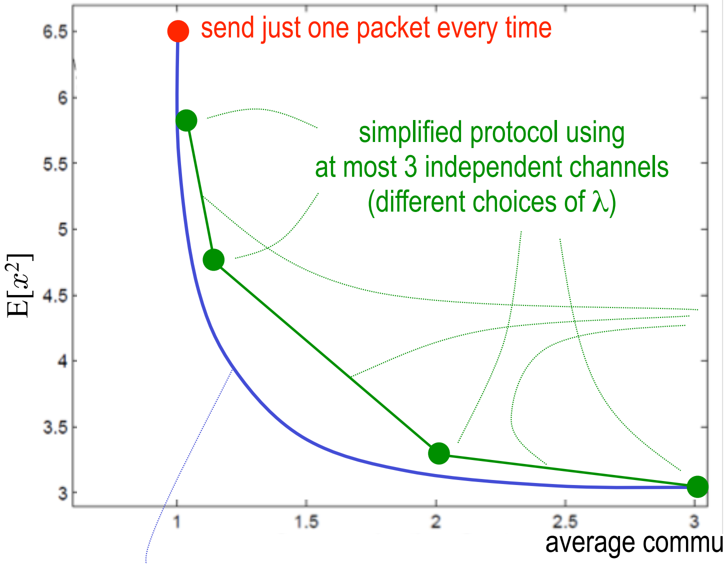

send just one packet every time

simplified protocol using
at most 3 independent channels
(different choices of λ)

stochastic protocol using
at most 3 independent channels
(different choices of λ)

optimal protocol using
at most 3 independent channels
(different choices of λ)

Three Key Ideas

- Importance of Stochastic Modeling versus Worst Case in NCSs
- Analysis & Design Results Available for Stochastic NCSs
  - Time-driven SHSs
  - Lyapunov-based methods
- (Stochastic) Control Tools for NCS Protocol-Design

(ex) students: D. Antunes (IST), A. Mesquita (UCSB), Y. Xu (Advertising.com)

collaborators: C. Silvestre (IST)

acknowledgements: NSF, Institute for Collaborative bio-technologies (ARO)

disclaimer: This is an overview, technical details in papers referenced in bottom right corner… http://www.ece.ucsb.edu/~hespanha