Multi-Agent Decision Making with Limited Information Exchange

TAMER BAŞAR

Dept ECE, CAS and CSL, UIUC

tbasar@control.csl.uiuc.edu

Caltech, Pasadena

March 22, 2006

Outline

- Interactions between Communication & Control
- Joint Sensor-Controller Design

  -- What/how to transmit & how to control
- Estimation and Control with "Power-Limited Communication"

  -- When to transmit & when to control
- Extensions to Teams and Games

Communication & Control

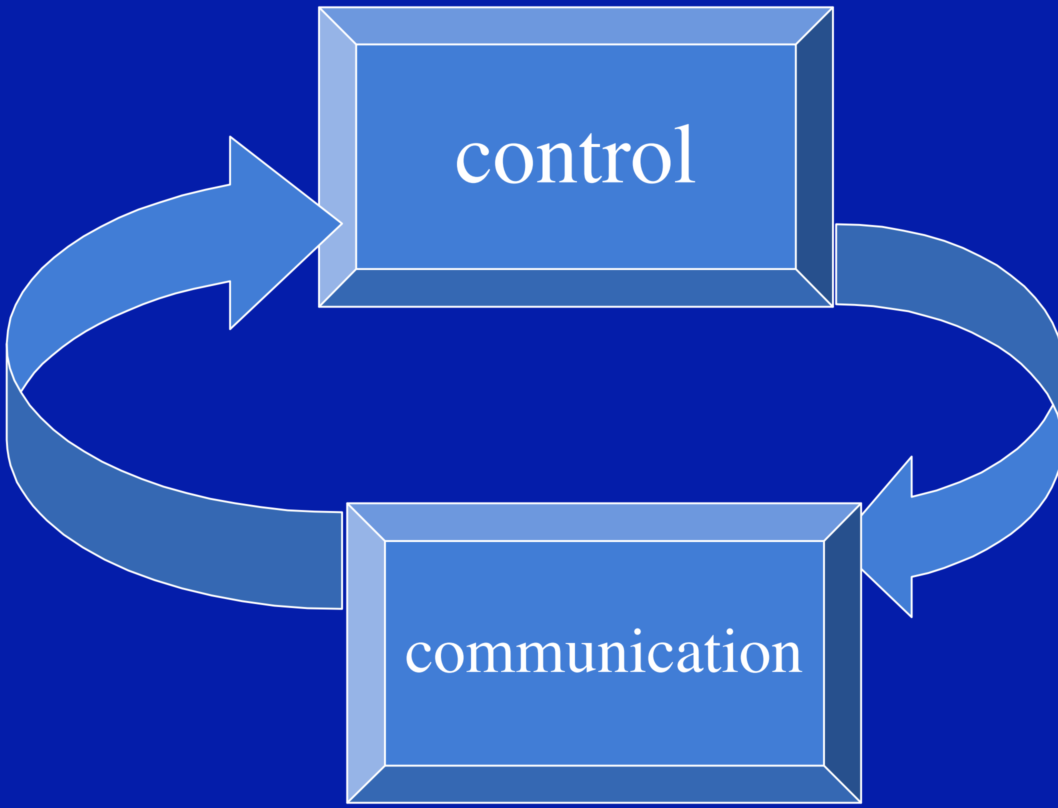

Standard Setup for Control

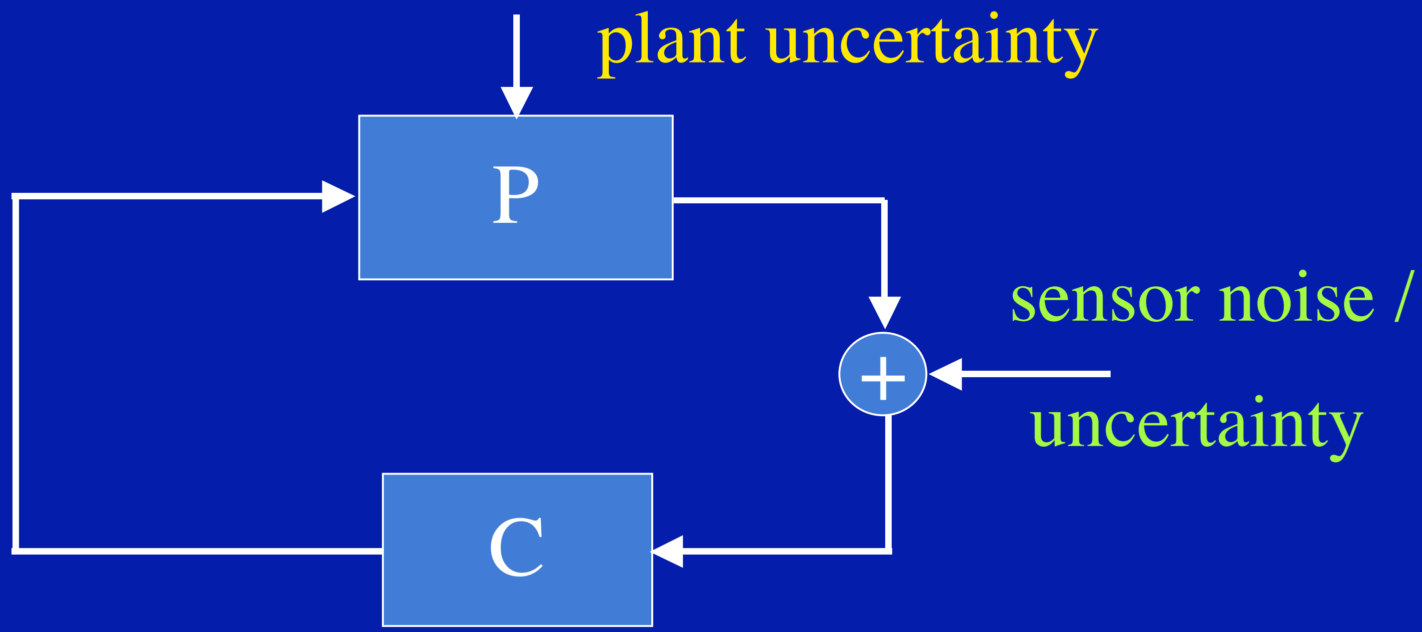

dual role of control: action & probing

generally not aligned

Separation / Neutrality

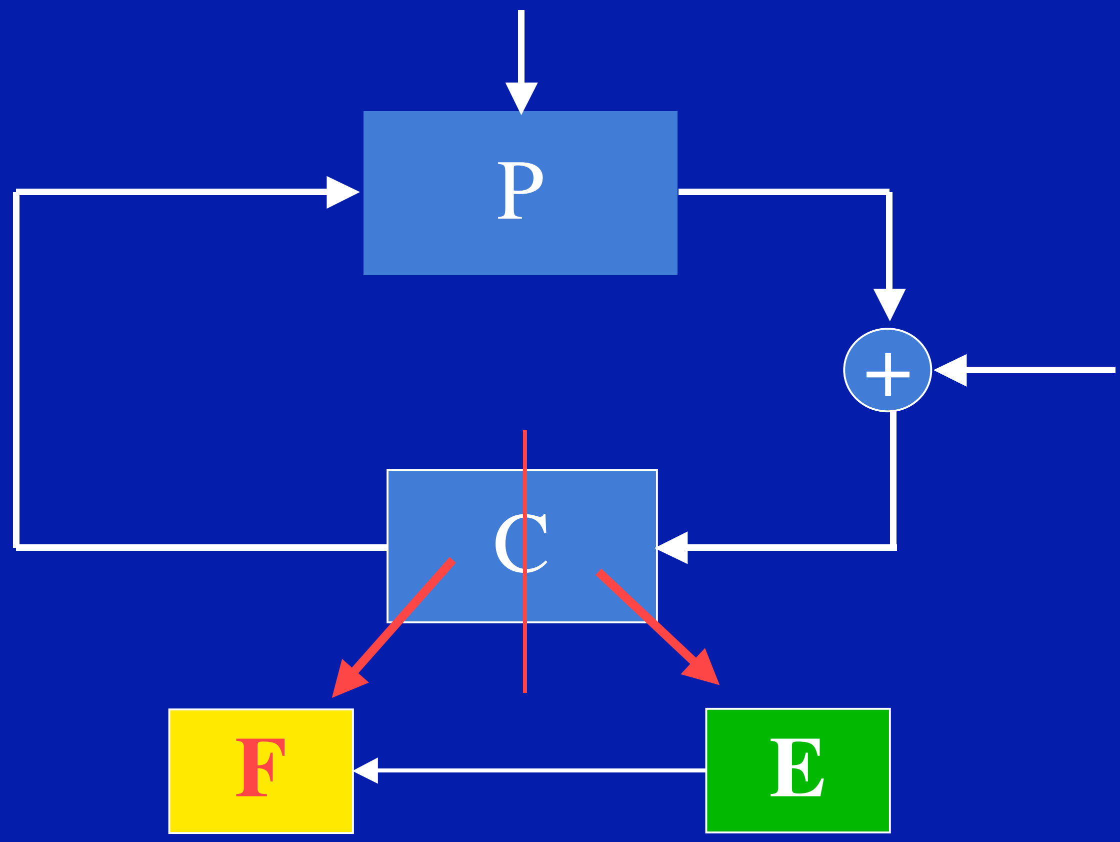

Non-classical

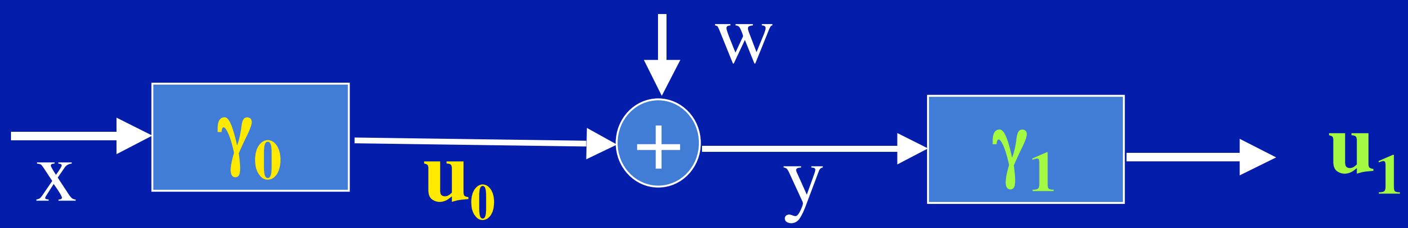

$x \sim N(0, \sigma_x^2)$ $w \sim N(0, \sigma_w^2)$

$J(\gamma_0 , \gamma_1) = E\,[\,Q(x, u_0, u_1) \mid \gamma_0 , \gamma_1\,]$

$J^* = \min \min J(\gamma_0 , \gamma_1)$

Witsenhausen (1968)

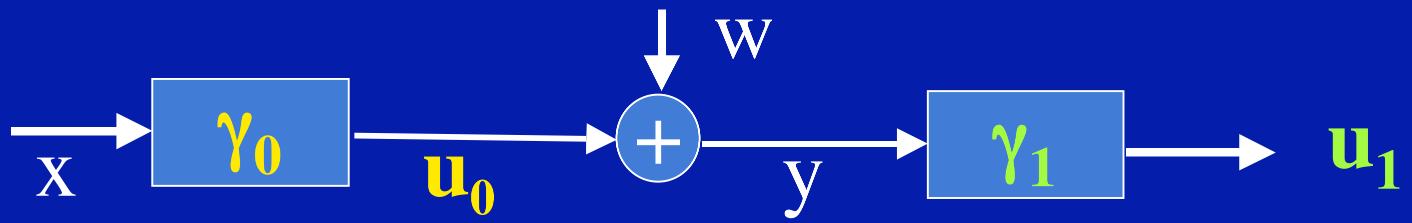

$$Q(x, u_0, u_1) = k_0 (u_0 - x)^2 + (u_0 - u_1)^2$$

⇒ optimal control law exists, but its structure is not known

-- roles of γ_0 and γ_1 are not aligned

Gaussian Test Channel

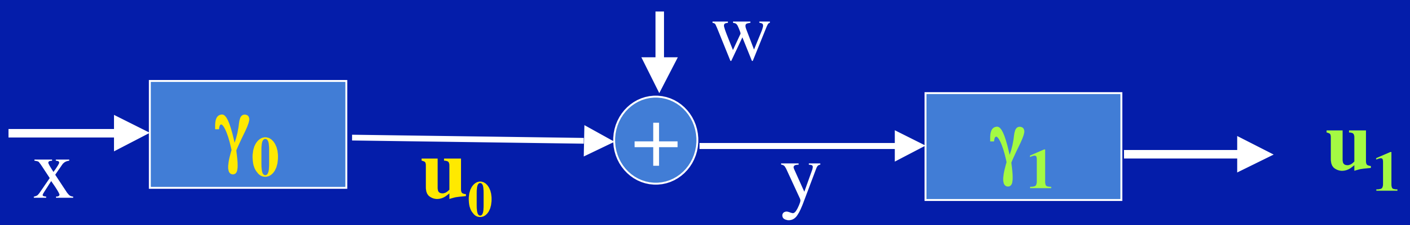

$$Q(x, u_0, u_1) = k_0 (u_0)^2 + (u_1 - x)^2$$

⇒ optimal control law (encoder/decoder) exists, and is linear
-- roles of γ_0 and γ_1 are aligned

Gaussian Test Channel

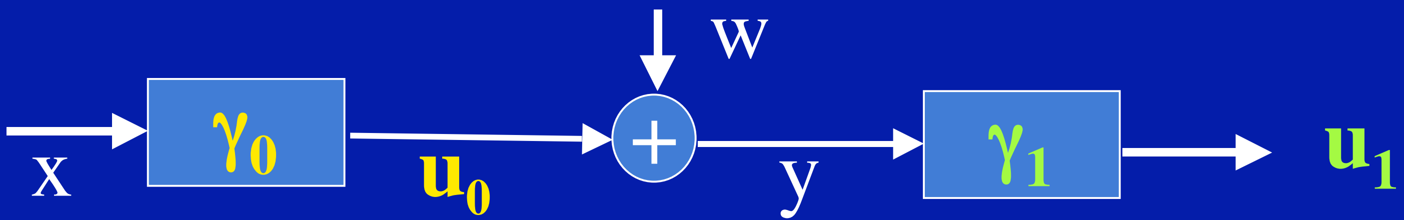

$$Q(x, u_0, u_1) = k_0 (u_0)^2 + (u_1 - x)^2 + b_0 u_0 x$$

⇒ optimal control law (encoder/decoder) exists, and is linear

-- roles of γ_0 and γ_1 are aligned

Generalized Gaussian Test Channel

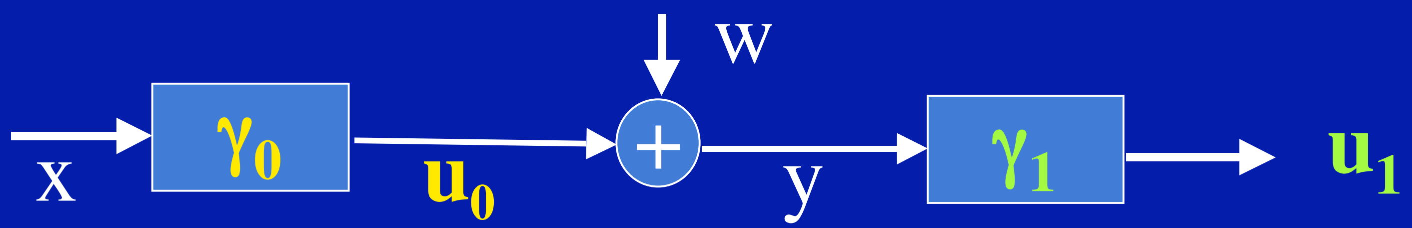

$$Q(x, u_0, u_1) = k_0 (u_0)^2 + (u_1 - x)^2 + b_0 u_0 x$$

$$E[Q]=F(\gamma_0, \gamma_1) \geq k_0 \alpha + \beta + \inf_\gamma b_0 E[\gamma_0(x)x]$$

$$\geq k_0 \alpha + \beta - \text{sgn}(b_0)\, \sigma_x \sqrt{\alpha}$$

Generalized Gaussian Test Channel

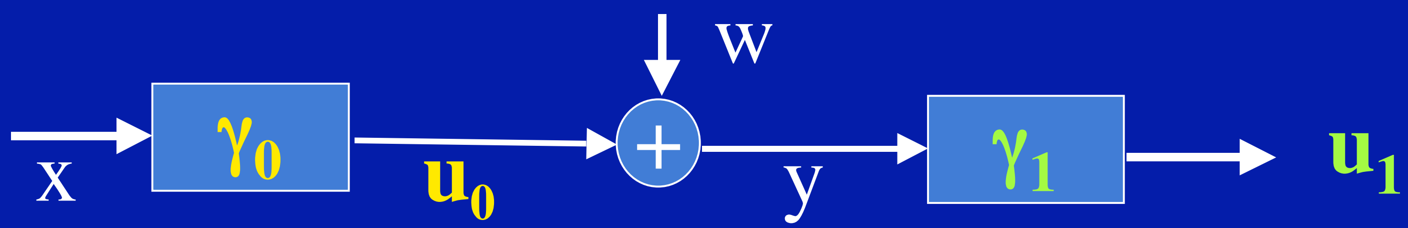

$$Q(x, u_0, u_1) = k_0 (u_0)^2 + (u_1 - x)^2 + b_0 u_0 x$$

$$E[Q]=F(\gamma_0, \gamma_1) \geq k_0 \alpha + \beta + \inf_\gamma b_0 E[\gamma_0(x)x]$$

$$\geq k_0 \alpha + \beta - \text{sgn}(b_0) \sigma_x \sqrt{\alpha}$$

DPT: $I(X;Y) \geq I(X;U_1)$

Generalized Gaussian Test Channel

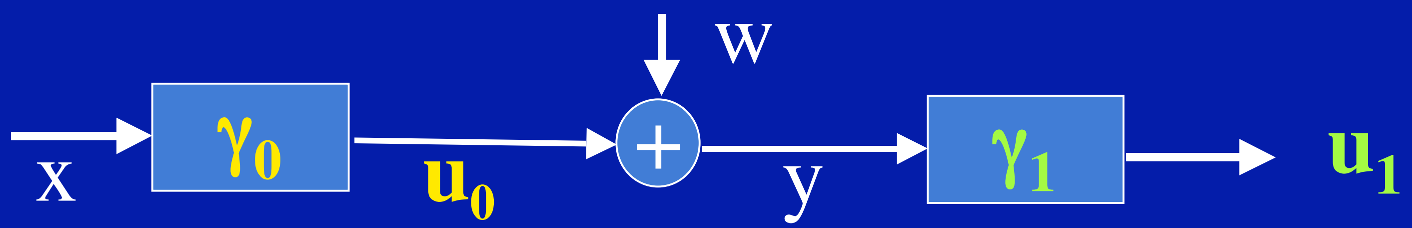

$$Q(x, u_0, u_1) = k_0 (u_0)^2 + (u_1 - x)^2 + b_0 u_0 x$$

$$E[Q]=F(\gamma_0, \gamma_1) \geq k_0 \alpha + \beta + \inf_\gamma b_0 E[\gamma_0(x)x]$$

$$\geq k_0 \alpha + \beta - \mathrm{sgn}(b_0) \sigma_x \sqrt{\alpha}$$

$$(1/2)\log (1+(\alpha/ \sigma_w^2)) \geq I(X;Y) \geq I(X;U_1) \geq (1/2)\log (\sigma_x^2/ \beta)$$

$C(\alpha)$ $R(\beta)$

Generalized Gaussian Test Channel

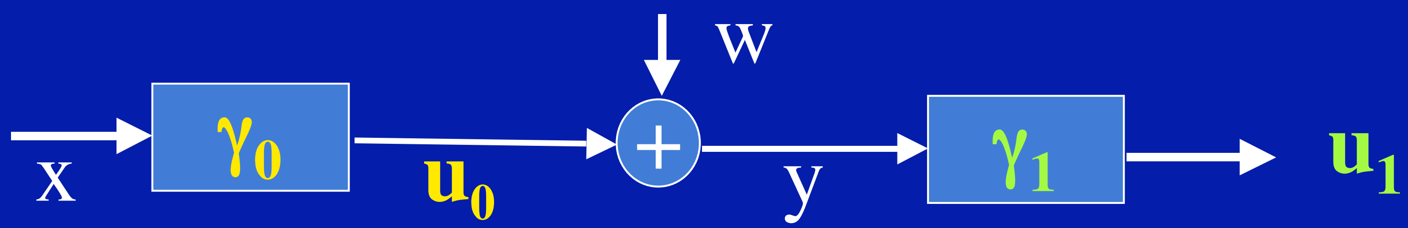

$$Q(x, u_0, u_1) = k_0 (u_0)^2 + (u_1 - x)^2 + b_0 u_0 x$$

$$E[Q]=F(\gamma_0, \gamma_1) \geq k_0 \alpha + \beta + \inf_\gamma b_0 E[\gamma_0(x)x]$$

$$\geq k_0 \alpha + \beta - \mathrm{sgn}(b_0) \sigma_x \sqrt{\alpha}$$

$$(1/2)\log (1+(\alpha/ \sigma_w^2)) \geq I(X;Y) \geq I(X;U_1) \geq (1/2)\log (\sigma_x^2/ \beta)$$

==> $\beta \geq \sigma_w^2 \sigma_x^2 / (\sigma_w^2 + \alpha)$

Generalized Gaussian Test Channel

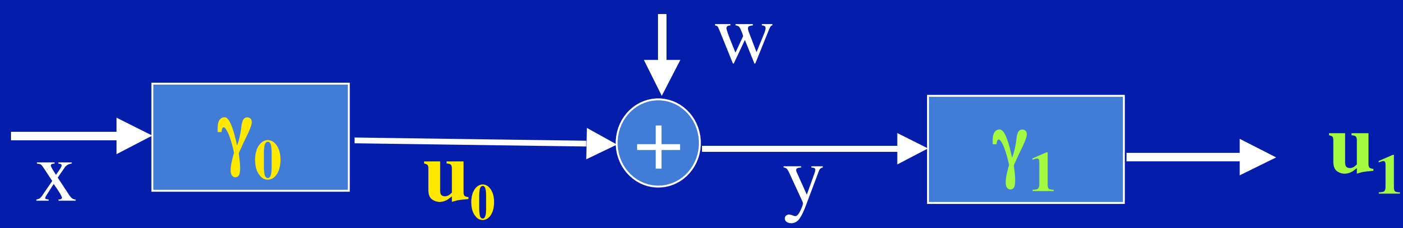

$$Q(x, u_0, u_1) = k_0 (u_0)^2 + (u_1 - x)^2 + b_0 u_0 x$$

$$E[Q]=F(\gamma_0, \gamma_1) \geq k_0 \alpha + \beta + \inf_{\gamma} b_0 E[\gamma_0(x) x]$$

$$\geq k_0 \alpha + \beta - \mathrm{sgn}(b_0)\, \sigma_x \sqrt{\alpha}$$

==> $\beta \geq \sigma_w^2 \sigma_x^2 / (\sigma_w^2 + \alpha)$

Inequality is tight with $\gamma_0(x) = -\mathrm{sgn}(b_0)(\sqrt{\alpha} / \sigma_x)\, x$

Generalized Gaussian Test Channel

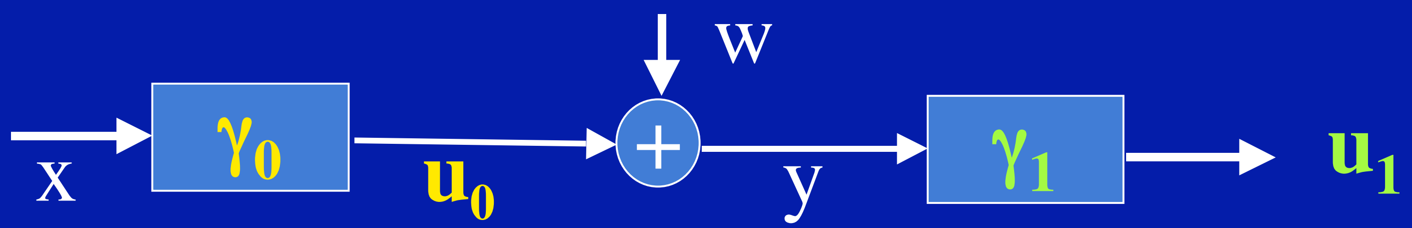

$$Q(x, u_0, u_1) = k_0 (u_0)^2 + (u_1 - x)^2 + b_0 u_0 x$$

$$E[Q]=F(\gamma_0, \gamma_1) \geq k_0 \alpha + \beta - |b_0| \sigma_x \sqrt{\alpha}$$

$$\geq \mathbf{k_0 \alpha + \sigma_w^2 \sigma_x^2 / (\sigma_w^2 + \alpha) - |b_0| \sigma_x \sqrt{\alpha}}$$

Obtain the α that minimizes the bound --> α^*

Then, $\gamma_0^*(x) = -\text{sgn}(b_0)(\sqrt{\alpha^*} / \sigma_x)\, x$, $\gamma_1^*(y) = E[x|y]$

Generalized Gaussian Test Channel

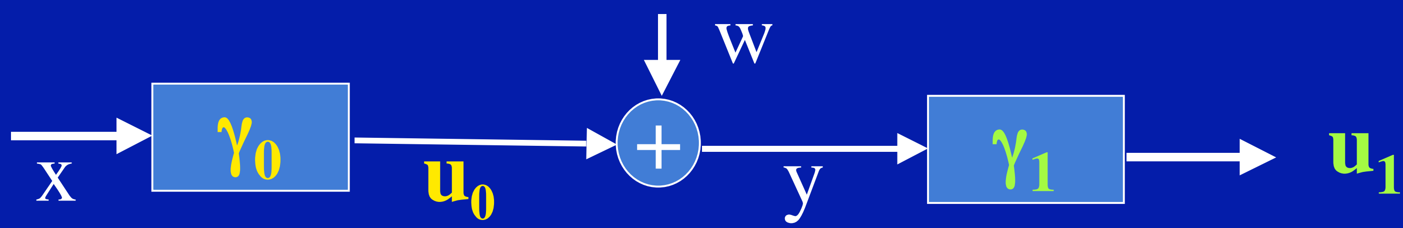

$$Q(x, u_0, u_1) = k_0 (u_0)^2 + (u_1 - x)^2 + b_0 u_0 x$$

One of the *few* instances when static/causal coding (and linear in this case) leads to attainment of equality in $C(\alpha) \geq R(\beta)$

Revisit: Witsenhausen (1968)

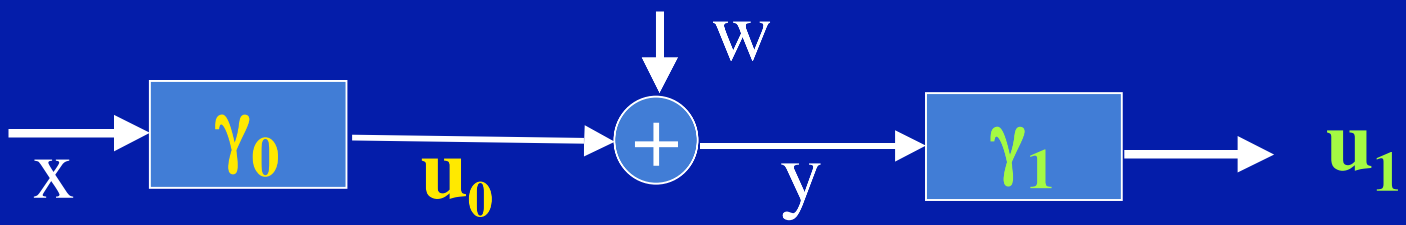

$$Q(x, u_0, u_1) = k_0 (u_0 - x)^2 + (u_0 - u_1)^2$$

Because of the product term $u_0 u_1$ the preceding analysis does not apply here

Revisit: Witsenhausen (1968)

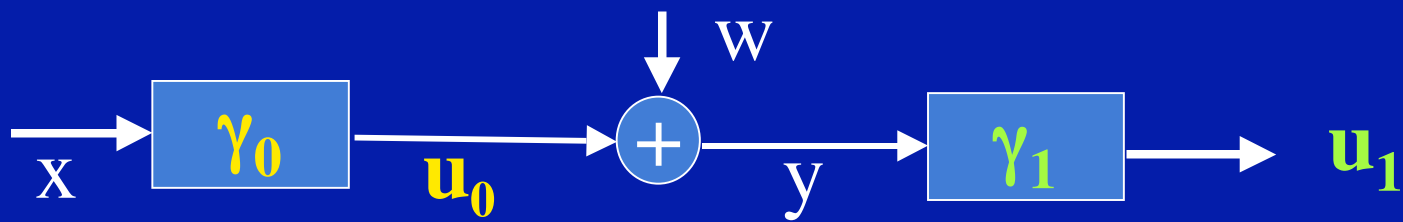

$$Q(x, u_0, u_1) = k_0 (u_0 - x)^2 + (u_0 - u_1)^2$$

Note that this is equivalent to a 2-stage LQG problem with restricted memory:

State equations: $x_1 = x_0 + v_0$ $\quad$ $x_2 = x_1 - v_1$

Measurement: $y_0 = x_0$ $\quad$ $y_1 = x_1 + w$

Controls: $v_0 = \mu_0(y_0)$, $v_1 = \mu_1(y_1)$; Cost: $(x_2)^2 + (v_0)^2$

Pick $u_0 = x_0 + v_0$ $\quad$ $u_1 = v_1$ $\quad$ $x = x_0$

However, with Conflicting Objectives

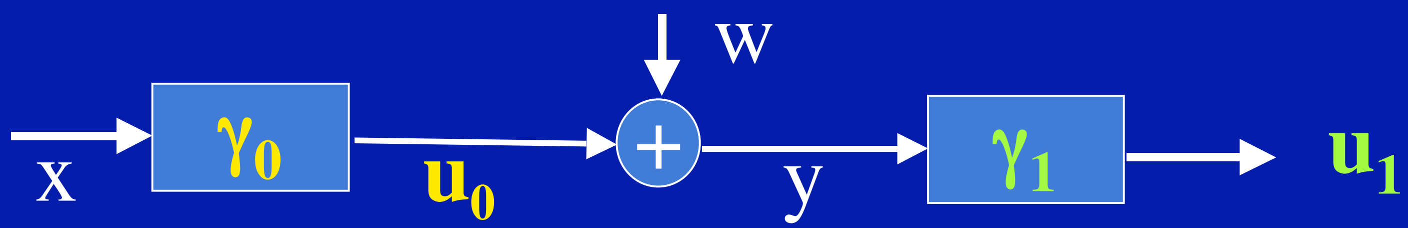

$$Q(x, u_0, u_1) = -k_0 (u_0 - x)^2 + (u_0 - u_1)^2$$

$$J_* = \min_{\gamma_1} \max_{\gamma_0} J(\gamma_0, \gamma_1)$$

⇒ Unique saddle-point solution, control laws are linear **(TB'71)**

Recap

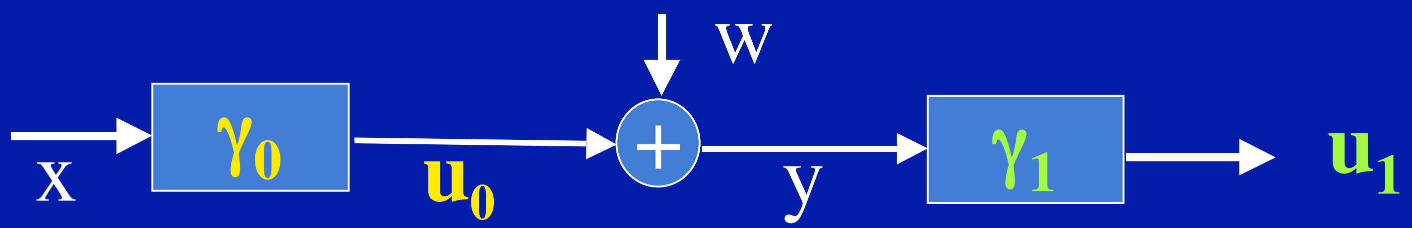

$Q_W = \; k_0 (u_0 - x)^2 + (u_0 - u_1)^2$ conflicting roles

$Q_G = \; - k_0 (u_0 - x)^2 + (u_0 - u_1)^2$ aligned roles

$Q_{TC} = \; k_0 (u_0)^2 + (u_1 - x)^2$ aligned roles

⇨ Not only the information structure but also the cost function is a determining factor

Classical/Neutral

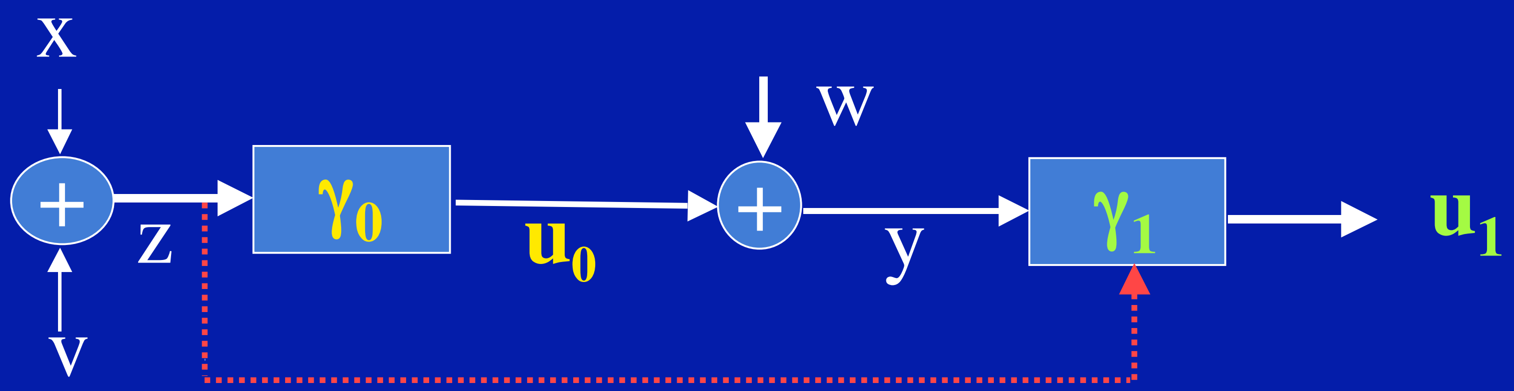

Solution is linear for all quadratic Q (x, u_0, u_1)

Separation / certainty equivalence

Quasi-Classical Team

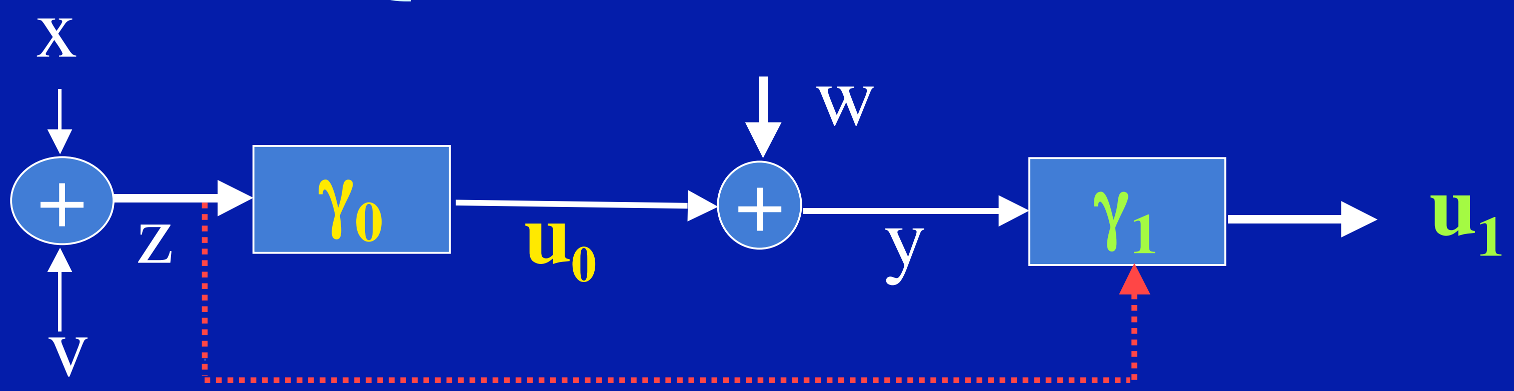

u_0 multi-dimensional, decentralized -- $\gamma_{0i}(z_i)$

u_1 multi-dimensional, decentralized -- $\gamma_{1i}(y_i, z)$

Solution is linear for all quadratic Q (x, u_0, u_1)

(based on Radner, 1962)

Quasi-Classical NZS game

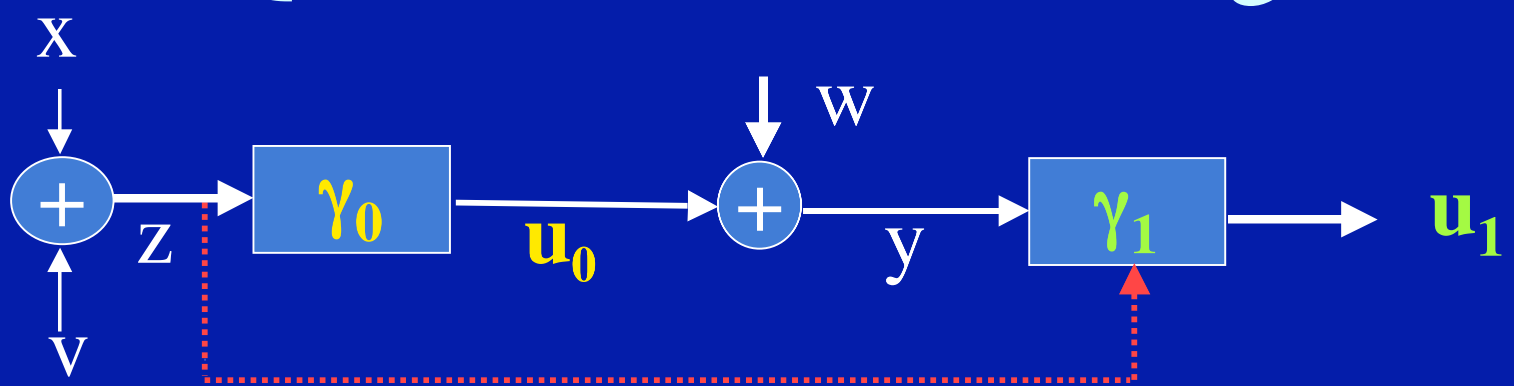

u_0 multi-dimensional, decentralized -- $\gamma_{0i}(z_i)$

u_1 multi-dimensional, decentralized -- $\gamma_{1i}(y_i, z)$

NZS game with Pi's quadratic cost $Q_i(x, u_0, u_1)$

Nash solution is unique and linear, if exists (TB'74)

Networks providing the communication medium

-- wired and wireless links

Networks & Control

- Traditional Control Systems have dedicated data paths for communication

Controller-Plant Communication

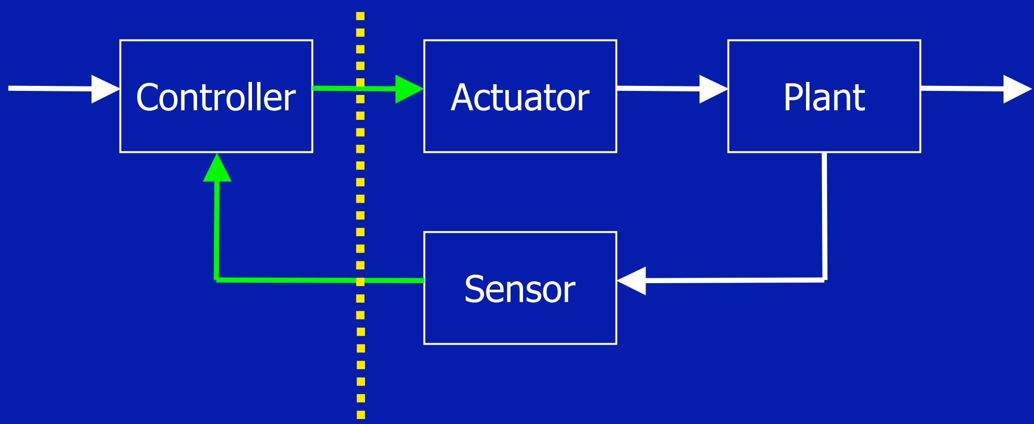

Networks & Control

- In Networked Control Systems (NCS) controller-plant communication takes place over a network

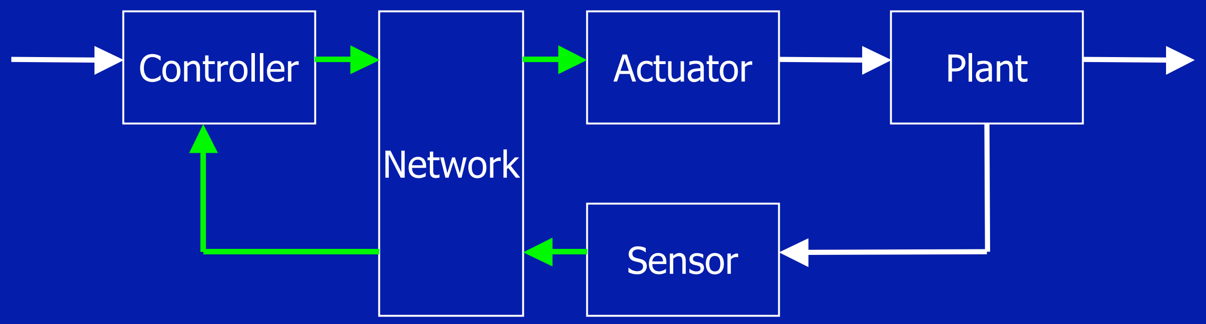

Remote Control Paradigm

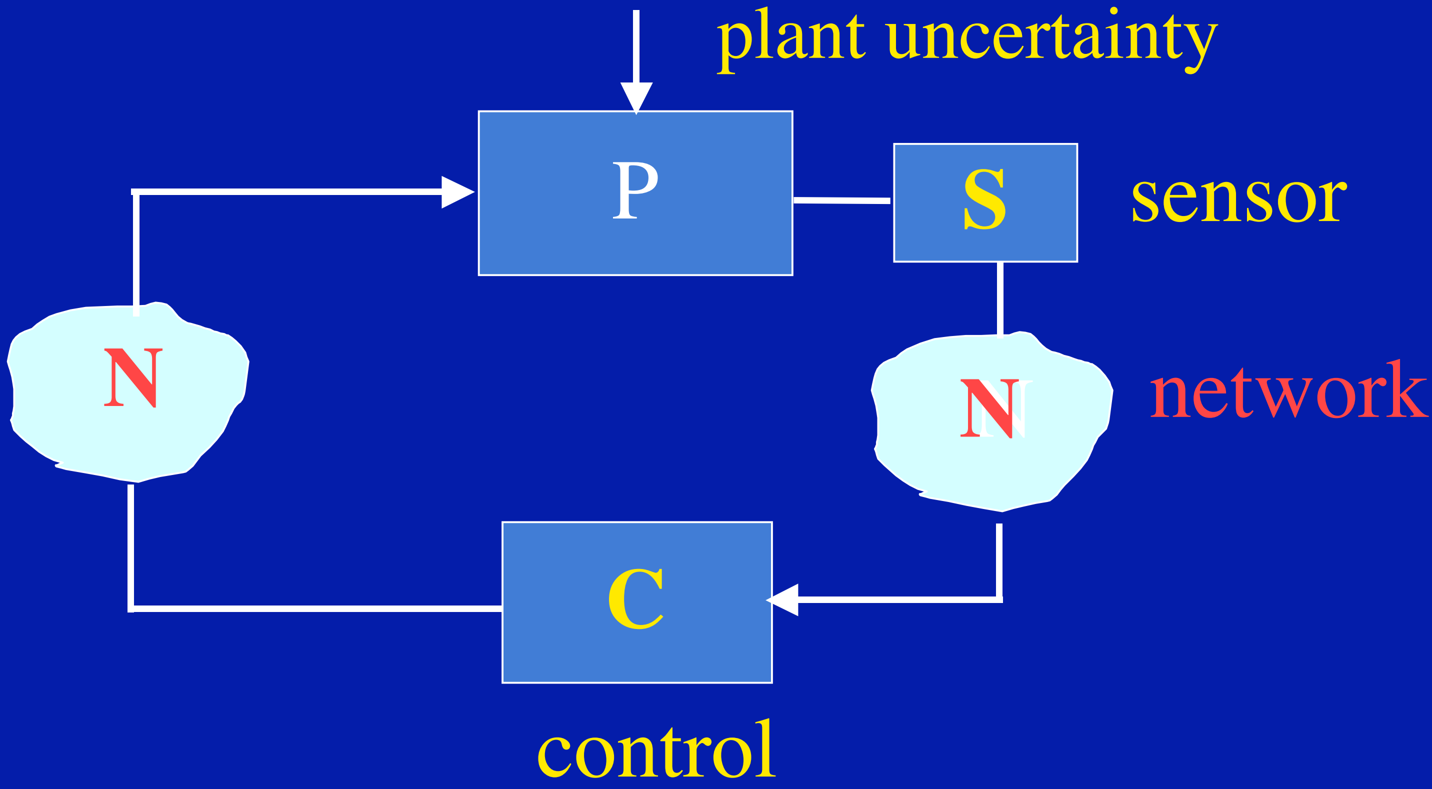

Remote Control Paradigm

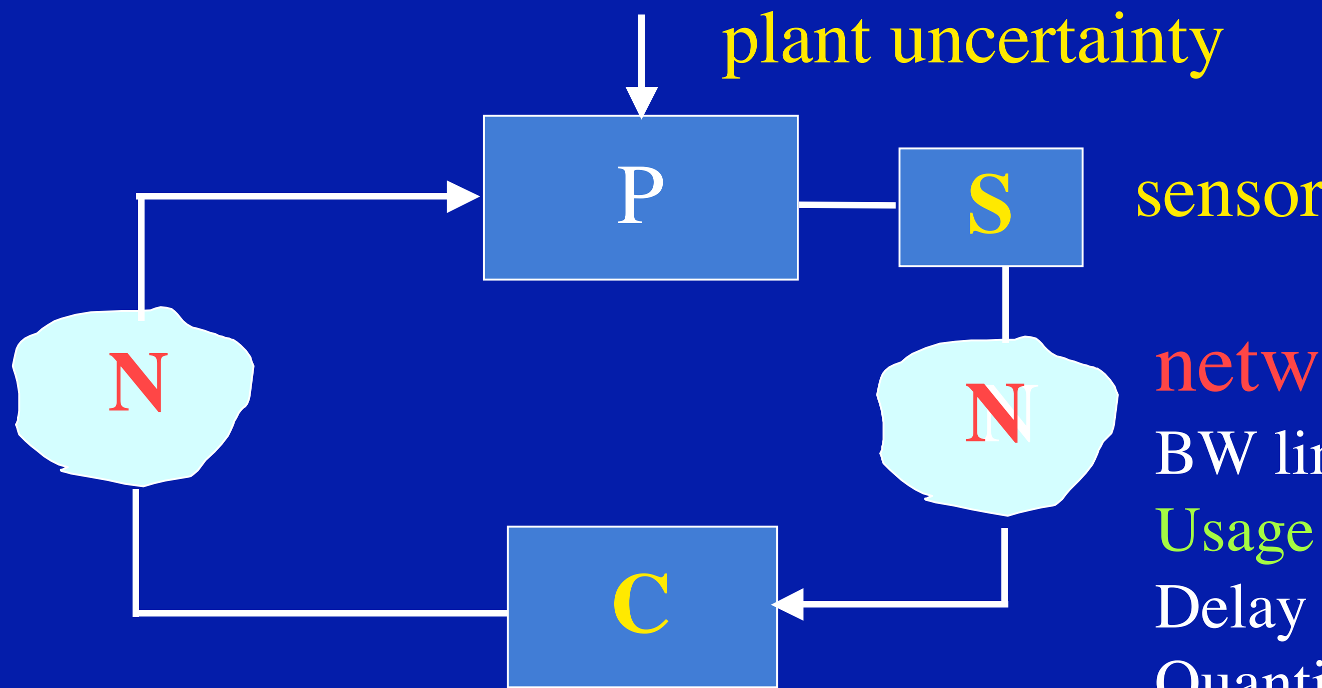

network
BW limitations
Usage limitations
Delay
Quantization (Q)
Failure

PI (S, Q, C) → optimize

Remote Control Paradigm

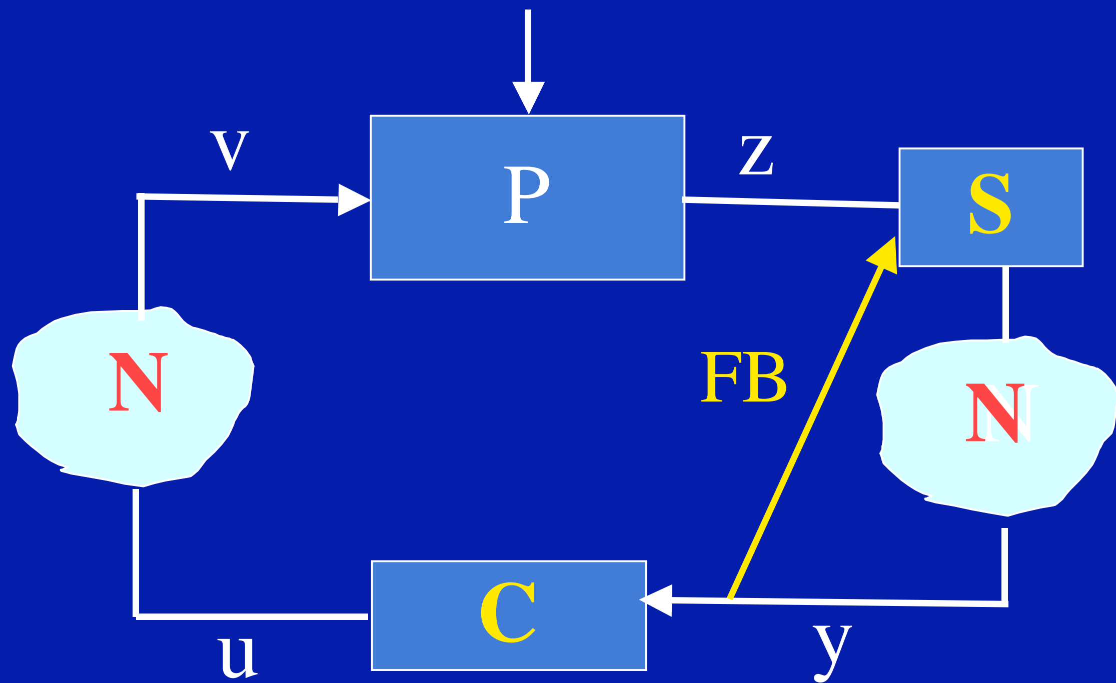

sensor

network

BW limitations
Usage limitations
Delay
Quantization (Q)
Failure

PI (S, Q, C) → optimize

Non-classical information!

Remote Control Paradigm

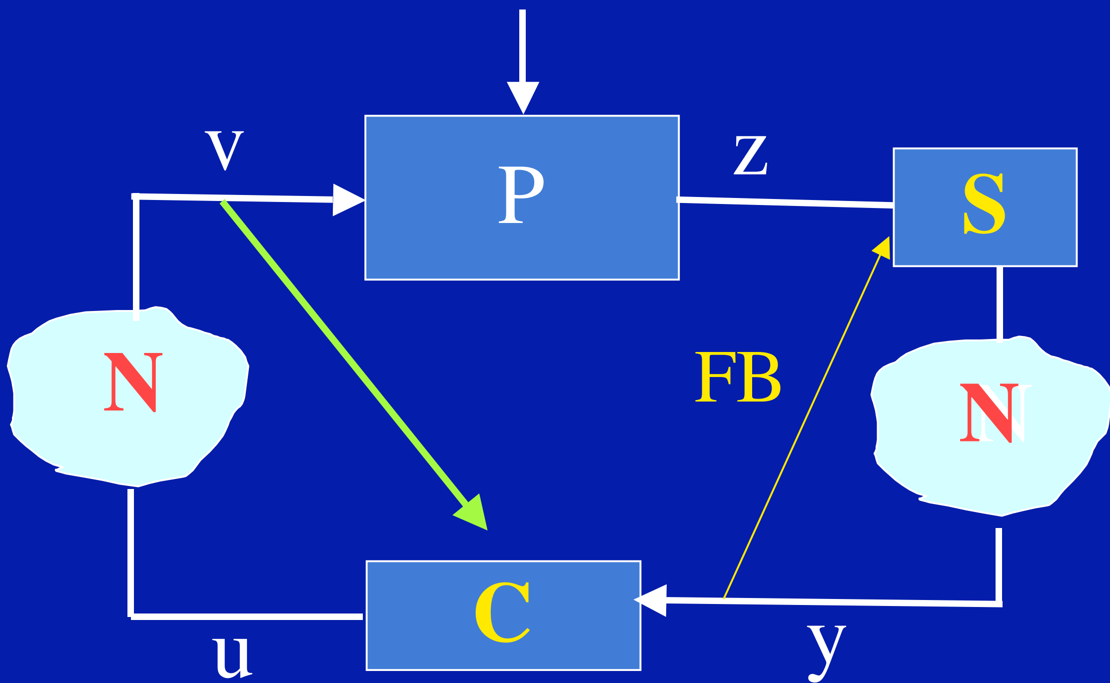

sensor

network

BW limitations
Usage limitations
Delay
Quantization (Q)
Failure

PI (S, Q, C) → optimize

Non-classical information!

Why non-classical?

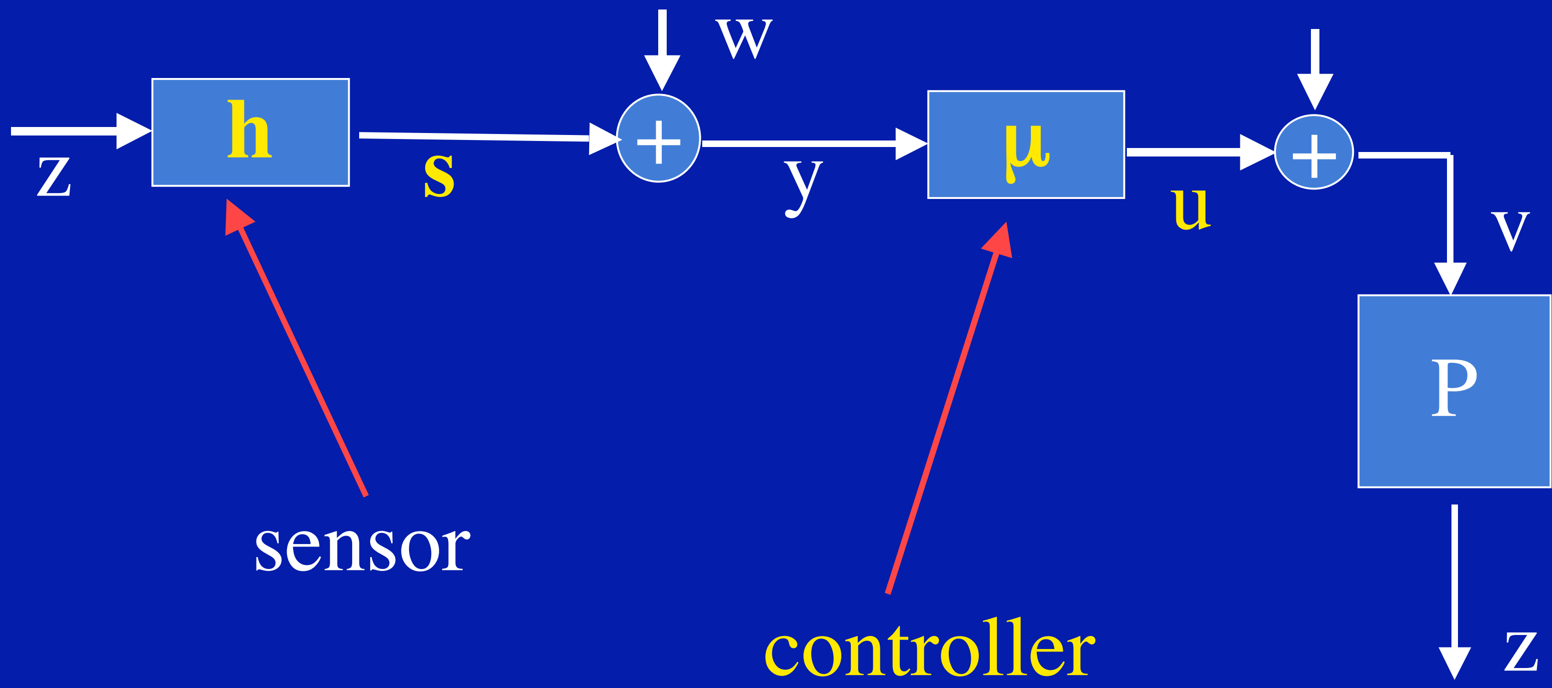

Why non-classical?

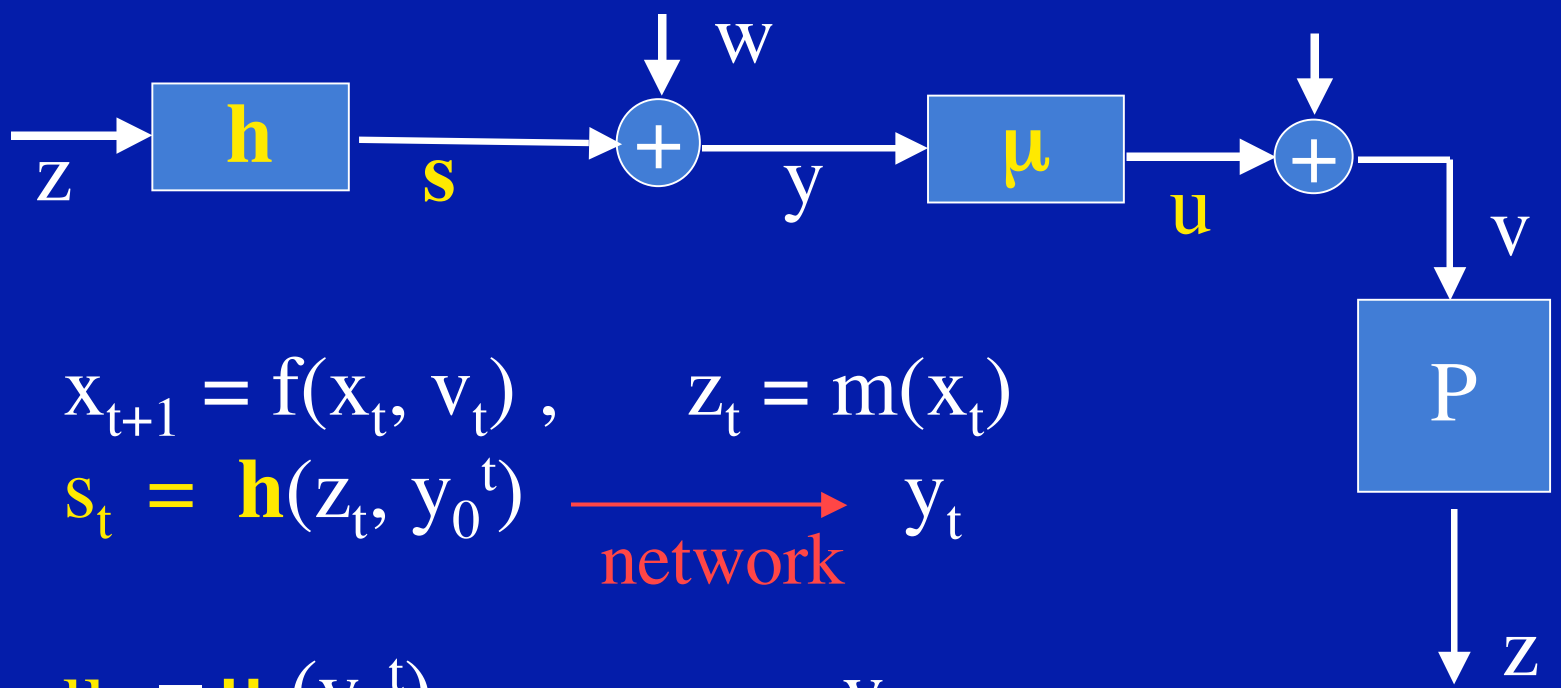

$x_{t+1} = f(x_t, v_t)$, $z_t = m(x_t)$

$s_t = \mathbf{h}(z_t, y_0^t)$ —network→ y_t

$u_t = \boldsymbol{\mu}(y_0^t)$ —network→ v_t

$PI = E\{\|x\|^2 + \|u\|^2 + \|s\|^2\}$

Wireless Networks & Control

- Estimation & Control when controller-plant communication takes place over a Wireless Network

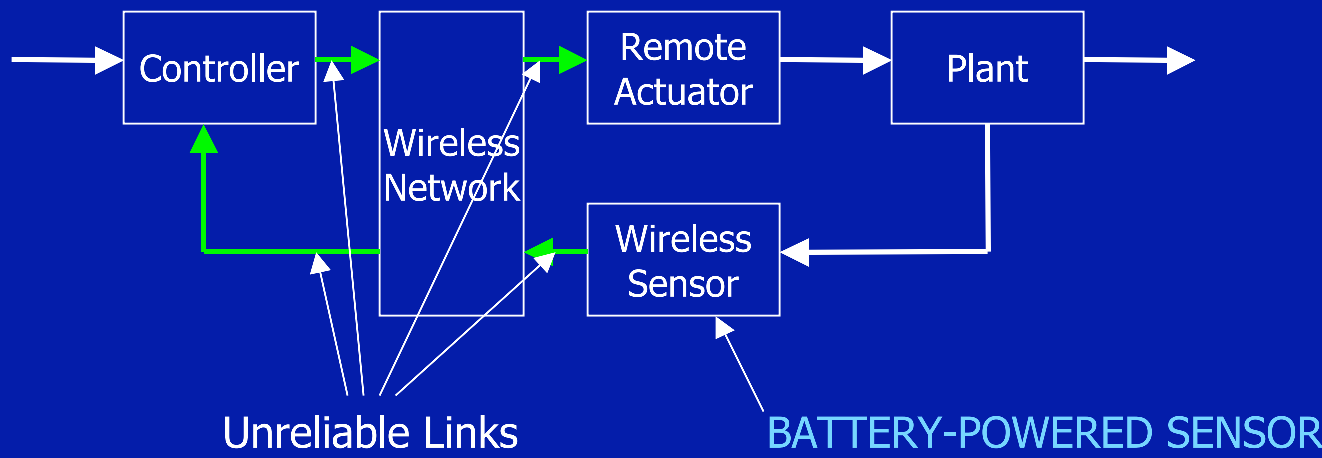

Estimation and Control with Power-Limited Communication

March 22, 2006 -- CalTech

Information Transmission over a Limited-Use Channel

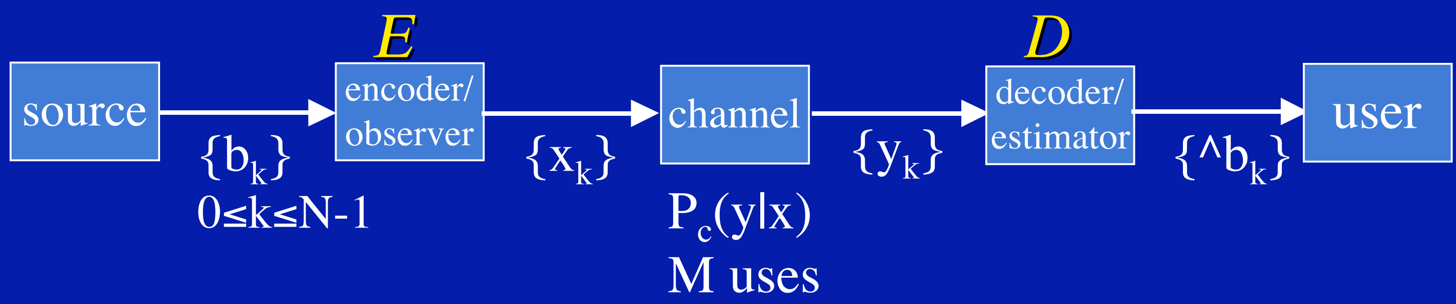

$x_k = E(z_k)$ $\quad$ $x \in X$ $\quad$ $y \in Y$

$z_k = b_k + v_k$ $\quad$ $M < N$

Given a "source" and a "memoryless channel", for a given message length N, and number of channel uses M, what is the minimum attainable value of the average distortion $D_{(M,N)}$ and a corresponding ***E & D*** *pair?*

Information Transmission over a Limited-Use Channel

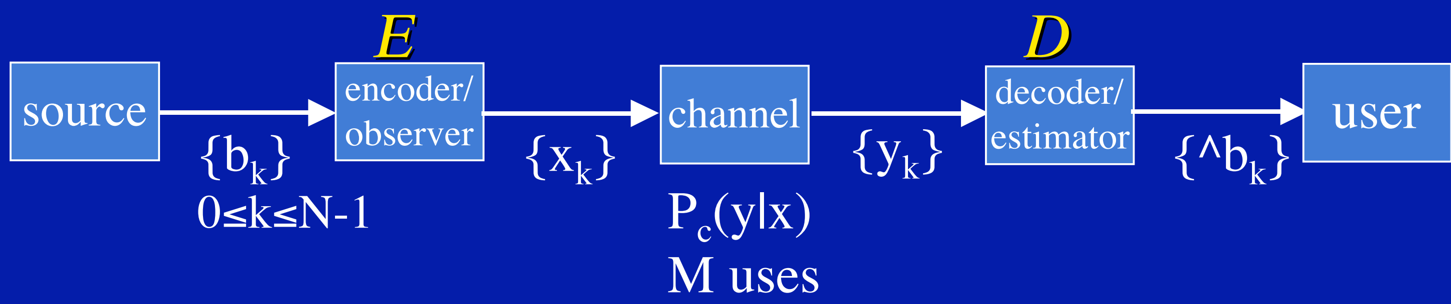

$x_k = \mathcal{E}(z_k)$ $\quad\quad$ $x \in \mathcal{X}$ $\quad\quad$ $y \in \mathcal{Y}$

$z_k = b_k + v_k$ $\quad\quad$ $M < N$

*Given a "source" and a "memoryless channel", for a given message length N, and number of channel uses M, what is the minimum attainable value of the average distortion $D_{(M,N)}$ and a corresponding **E & D** pair?*

e.g. $D_{(M,N)} = E\{(1/N) \sum_0^{N-1} (b_k - \hat{b}_k)^2\}$ $\quad\quad$ MSE

or the probability of error measure

Information Transmission over a Limited-Use Channel

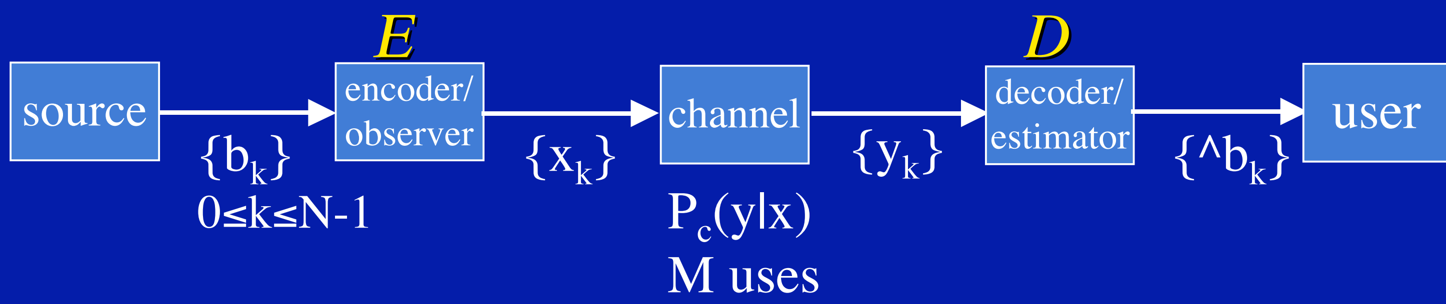

$x_k = E(z_k)$ $\quad x \in X$ $\quad y \in Y$

$z_k = b_k + v_k$ $\quad M < N$

Order of actions at time k:

1. b_k (or z_k) becomes available to the sensor
2. Sensor makes a decision: transmit/shape or not
3. Estimator acts by generating $\hat{b}_k$
4. Estimation error is incurred and we move to k+1

Information Transmission over a Limited-Use Channel

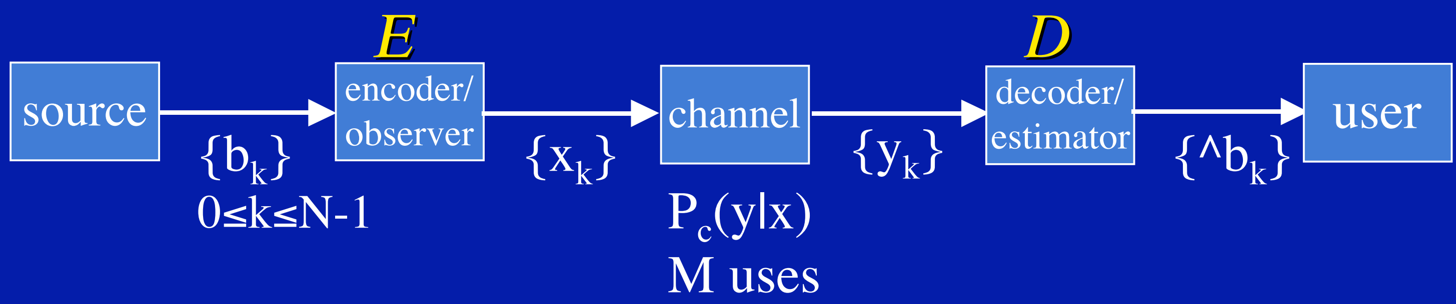

$x_k = \mathcal{E}(z_k)$ $\quad x \in \mathcal{X} \quad y \in \mathcal{Y}$

$z_k = b_k + v_k$ $\quad M < N$

Dynamic & Non-classical

Order of actions at time k:

1. b_k (or z_k) becomes available to the sensor
2. Sensor makes a decision: transmit/shape or not
3. Estimator acts by generating $\hat{b}_k$
4. Estimation error is incurred and we move to k+1

Control over a Limited-Use Channel

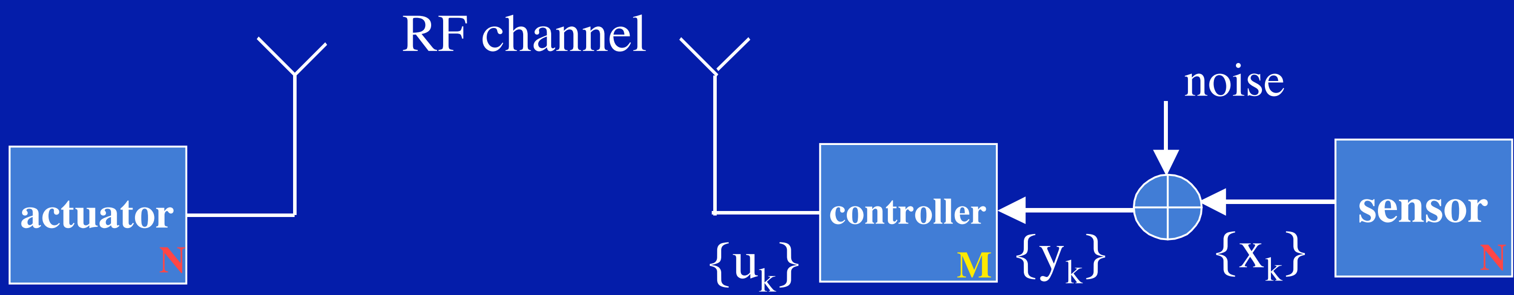

$x_{k+1} = f(x_k, u_k, w_k)$ $\qquad x_k \in X_k$, $u_k \in U_k$, $y \in Y_k$

$y_k = h_k(x_k) + v_k$ $\qquad M < N$

$u_k = \mu_k (I_k)$, $I_k := \{y_{[0,k]}, u_{[0,k-1]}\}$, $\mu_k \rightarrow U_k$ only M times

Given a horizon of N units, and with controller allowed to transmit for only M < N times, what is the minimum attainable value of a performance index J, and a corresponding controller?

Control over a Limited-Use Channel

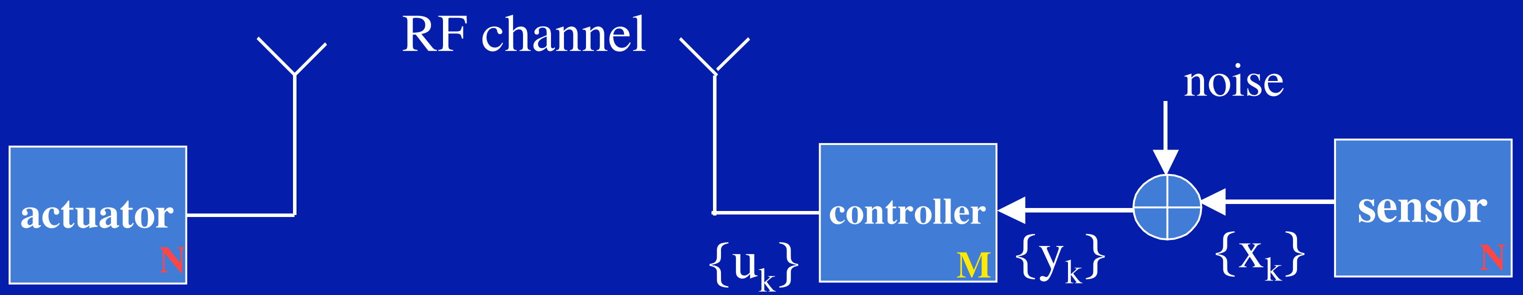

$x_{k+1} = f(x_k, u_k, w_k)$ $x_k \in X_k$, $u_k \in U_k$, $y \in Y_k$

$y_k = h_k(x_k) + v_k$ $M < N$

$u_k = \mu_k(I_k)$, $I_k := \{y_{[0,k]}, u_{[0,k-1]}\}$, $\mu_k \rightarrow U_k$ only M times

Given a horizon of N units, and with control allowed to transmit for only M < N times, what is the minimum attainable value of a performance index J, and a corresponding controller?

e.g. $J = E\{q(x_N) + \sum_0^{N-1} g(x_k,u_k)\}$

Estimation: A Special Case

N=2, M=1, b_0, b_1 i.i.d. Gaussian, 0-mean, variance σ^2

Perfect channel, no noise

Estimation error: $e = E\left\{(b_0 - \hat{}b_0)^2 + (b_1 - \hat{}b_1)^2\right\}$

Open-loop sensor policy:

Arbitrarily picks transmission time ==> $e_{OL} = \sigma^2$

Closed-loop sensor policy:

Transmit b_0 if it lies outside $[\alpha, \beta]$, $\alpha < 0 < \beta$; otherwise b_1

Minimization problem faced by sensor:

$$e_{(\alpha,\beta)} = \int_\alpha^\beta (b - E[b \mid b \in [\alpha, \beta]])^2 f(b)\, db + \sigma^2 P\{b_0 \notin [\alpha, \beta]\}$$

Special Case: Solution

$(\alpha^*, \beta^*) = (-\sigma, \sigma)$

$\Rightarrow$

$$e_{CL}^* = e_{(\alpha^*, \beta^*)} = [1 - \sqrt{(2 / \pi e)}]\, \sigma^2$$
$$\approx 0.52\, \sigma^2$$

$\Rightarrow$

48% improvement over the OL policy

Special Case: Solution

$(\alpha^*, \beta^*) = (- \sigma, \sigma)$

$\Rightarrow$

The knowledge of no action is useful information !!

$\Rightarrow$

48% improvement over the OL policy

Information Transmission over a Limited-Use Channel

-- i.i.d. case, perfect channel --

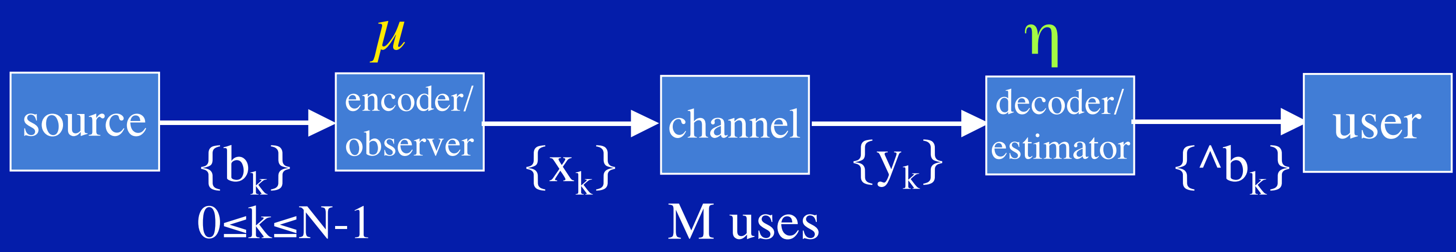

$x_k = \mu_k(I^e{}_k)$ $\hat{b}_k = \eta_k(I^d{}_k)$ $z_k = b_k + v_k$

s_k : # channel uses left at time k

t_k : # decision slots left at time k

$I^e{}_k = \{(s_k, t_k); z_{[0,k]}, x_{[0,k-1]}\}$, $1 \leq k \leq N-1$; $I^e{}_0 = \{(s_0, t_0); z_0\}$

$I^d{}_k = \{(s_k, t_k); y_{[0,k]}\}$, $0 \leq k \leq N-1$

ρ_k : 0-1 quantity; 1 if sensor transmits

$s_{k+1} = s_k - \rho_k$, $s_0 = M$ *s-dynamics*

Information Transmission over a Limited-Use Channel

-- i.i.d. case, perfect channel --

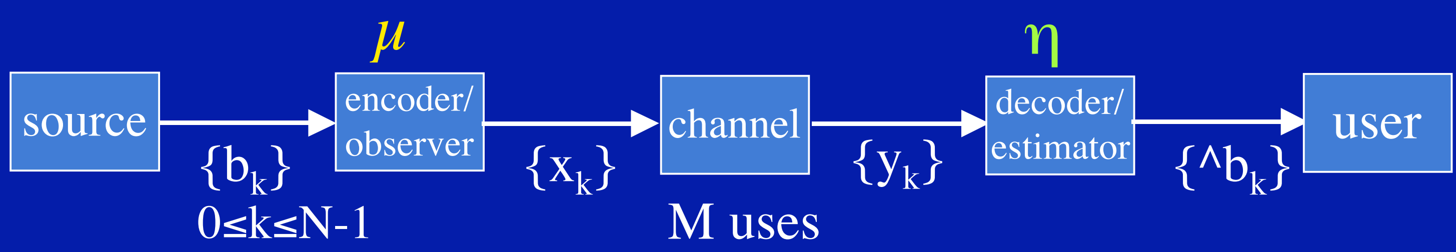

$x_k = \mu_k(I^e_k)$ $\hat{b}_k = \eta_k(I^d_k)$ $z_k = b_k + v_k$

Find $\mu_{[0, N-1]}$ and $\eta_{[0, N-1]}$ that minimize a given PI: MSE or Probability of error

$I^d_k = \{(s_k, t_k); y_{[0,k]}\}, \quad 0 \leq k \leq N-1$

ρ_k : 0-1 quantity; 1 if sensor transmits

$s_{k+1} = s_k - \rho_k$, $s_0 = M$ *s-dynamics*

Information Transmission over a Limited-Use Channel

-- i.i.d. case, perfect channel --

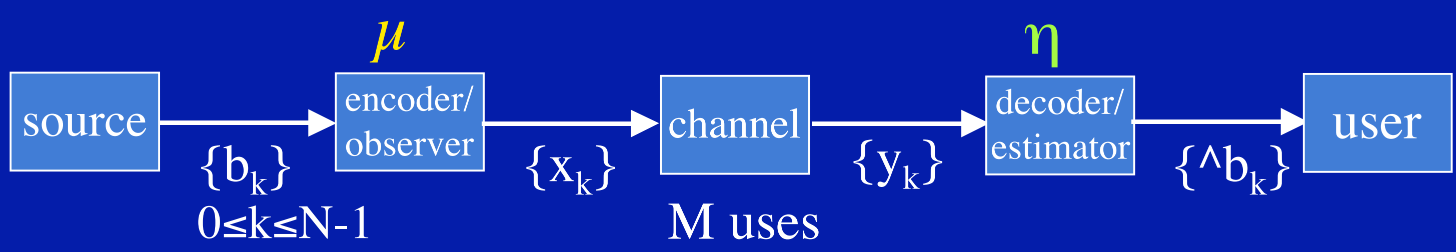

$x_k = \mu_k(I^e_k)$ $\hat{b}_k = \eta_k(I^d_k)$ $z_k = b_k + v_k$

Find $\mu_{[0, N-1]}$ and $\eta_{[0, N-1]}$ that minimize a given PI: MSE or Probability of error

- Best choice for η_k is $E[b_k \mid (s_k,t_k); x_k]$ or MAP
- Sufficient statistics for optimum μ_k is
  $S^e_k := \{(s_k,t_k); z_k\}$

Solution

Best sensor policy is of the form:

At time k transmit z_k if it is in a measurable set ¥(s_k,t_k), otherwise do not

¥(s,t) obtained offline as the minimizer in a recursive equation satisfied by accumulated optimum error, $e^*(s,t)$, at each point (s,t):

$$e^*(s,t) = \min_{¥(s,t)} \{e^*(s-1, t-1)\ \mathrm{Prob}(z \in ¥) + e^*(s, t-1)\ \mathrm{Prob}(z \notin ¥)$$
$$\text{average error at } (s,t) \text{ due to decision at } (s,t)\}$$

$e^*(t,t) = 0$, $e^*(0,t) = t\ \mathrm{var}(b)$

Specific structure of ¥(s,t) depends on the pdf/pmf and PI.

Offline Computations

The recursion can be solved offline, and the optimal sets ¥(s,t) tabulated starting with small values of (s,t)

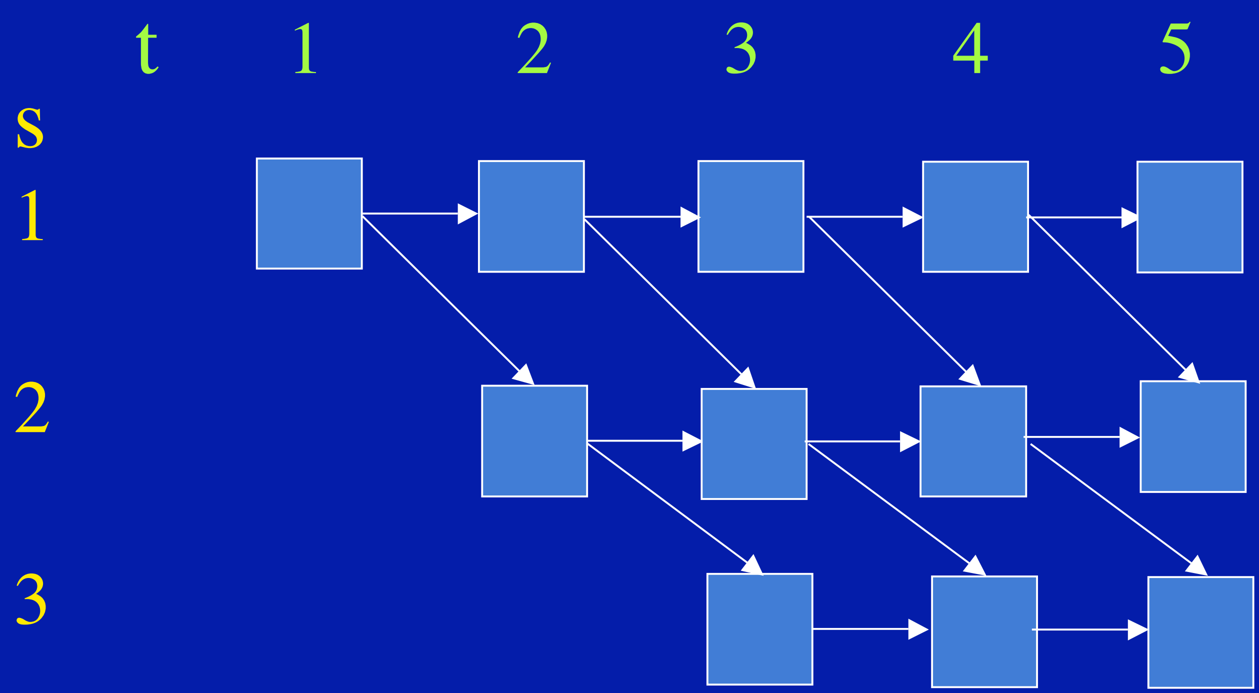

Explicit Solution in a Special Case

Continuous distribution f for b, f(-b) = f(b)

No noise v (from source to sensor) --n.l.o.g

==> $¥(s,t) = [-\beta_{(s,t)}, \beta_{(s,t)}]$

$$\beta_{(s,t)} = \sqrt{\{e^*(s-1, t-1) - e^*(s, t-1)\}}$$

Gaussian: $\varepsilon_{(s,t)} := e^*(s,t) / \mathrm{var}(b)$

$$\varepsilon_{(s,\,t)} = \varepsilon_{(s-1,\,t-1)} - [(\beta_{(s,t)})^2 - 1][2\Phi(\beta_{(s,t)}) - 1]$$

$$-(2/\sqrt{2\pi})\, \beta_{(s,t)} \exp\left((\beta_{(s,t)})^2 / 2\right)$$

$$\varepsilon_{(t,\,t)} = 0, \quad \varepsilon_{(0,\,t)} = t$$

The Best Interval in Gaussian Case

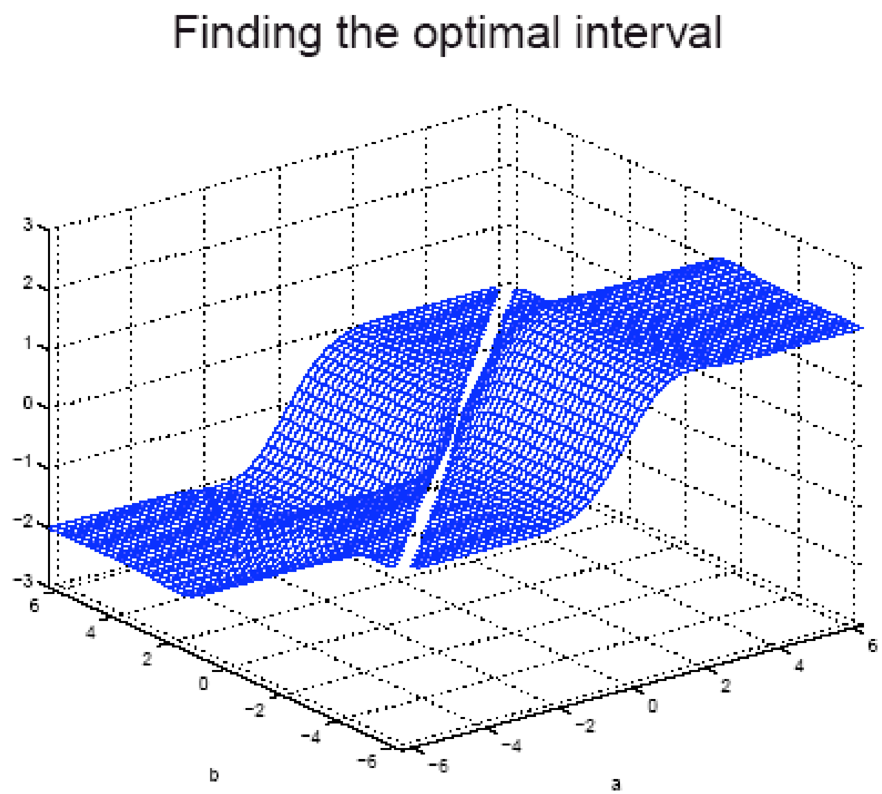

Plot of the objective function in the Gaussian case with $\mathcal{T}_{(s,t)} = [a, b]$ when $\sigma_x^2 = 1$, $e^*_{(s-1,t-1)} = 3$, $e^*_{(s,t-1)} = 1$

An Illustrative Example

- Problem: Given a time-horizon of fixed length $N = 100$, estimate the state of a zero-mean *i.i.d.* Gaussian process with unit variance $\sigma_x^2 = 1$.
- Design Criterion: The cumulative estimation error should not exceed $20\sigma_x^2 = 20$.
- Solution I: Make 80 sensor transmissions picked at arbitrary times.
- Solution II: Use the optimal sensor transmission and estimation policies:

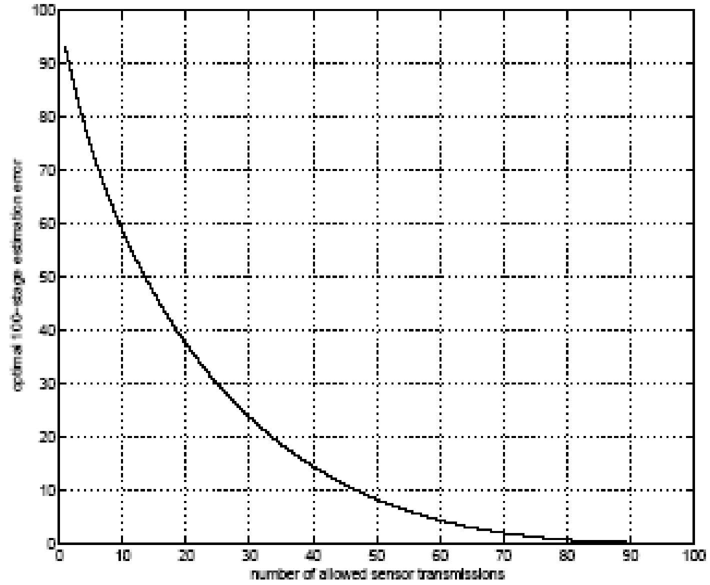

An Illustrative Example (cont.)

- Estimation error of 20 can be achieved with only 34 sensor transmissions! This is approximately a $\frac{80-34}{80} \times 100 \approx$ 58% improvement, which could save a lot of battery-power if the sensor is power-limited, or it could save a lot of transmission slots if the sensor is time-slot limited.
- Design Verification:
  1. Plot of optimal sample paths of number of sensor transmissions left starting with $(N, M) = (100, 34)$
  2. Comparison of the sample path cumulative estimation errrors between the optimal and an arbitrary transmission policy

An Illustrative Example (cont.)

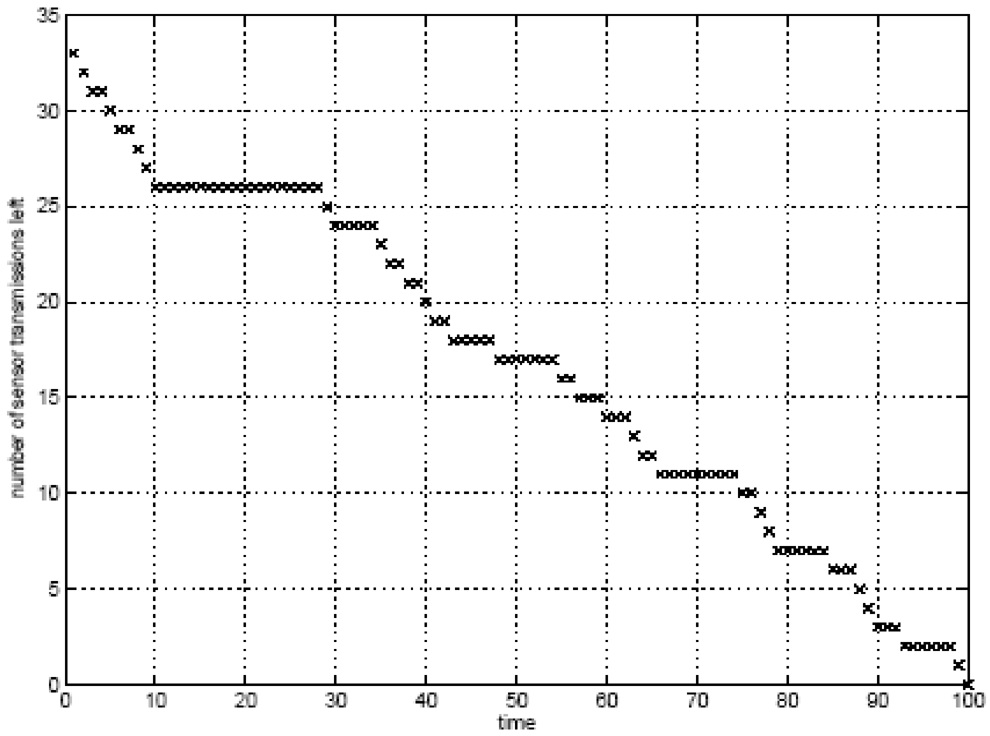

Typical sample path of the number of sensor transmissions left under the optimal transmission policy of the sensor

$$(N, M) = (100, 34)$$

An Illustrative Example (cont.)

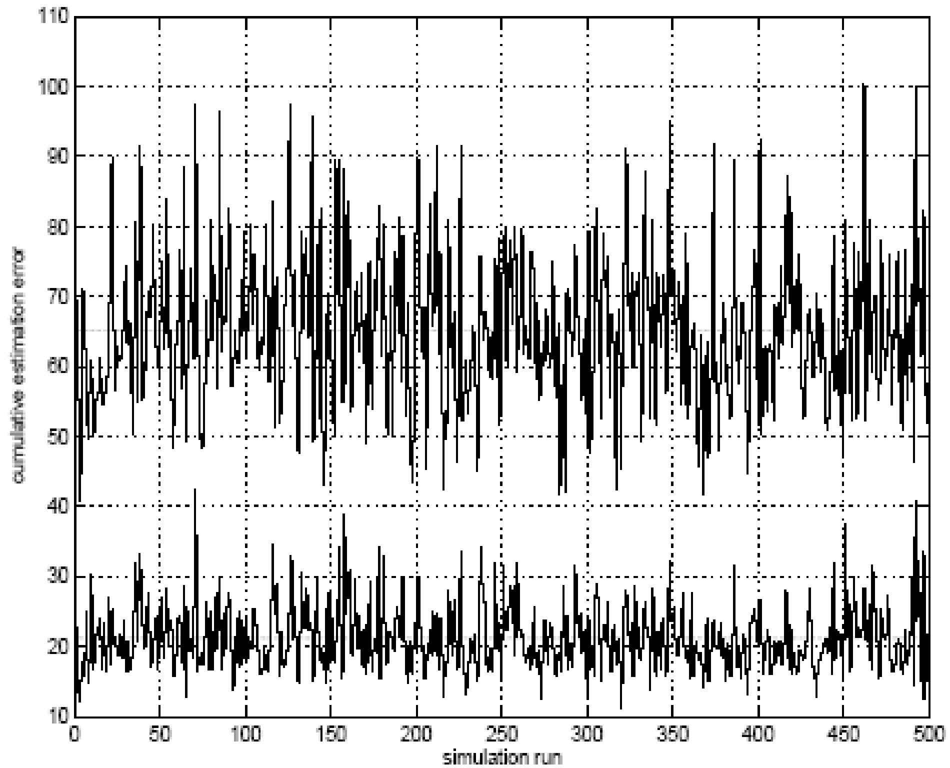

Comparison of the sample path cumulative estimation errors between the optimal and an arbitrary transmission policy

$$(N, M) = (100, 34)$$

Extensions: Markov Case

- The underlying process is Markov:

$$x_{k+1} = Ax_k + w_k$$

- The solution is similar in the Gaussian case. Now, a function of (L, M, N), where L is the number of time units passed since the last transmission.
- Let

$$\Sigma_L = (1 + A^2 + \ldots + A^{2L})$$

- The DP recursion, where $x \sim \mathcal{N}(0, \Sigma_L \sigma_w^2)$

$$e^*_{(r,s,t)} = \min_{\mathcal{T}_{(r,s,t)}} \left\{ e^*_{(0,s-1,t-1)} \int_{x \in \mathcal{T}_{(r,s,t)}} f(x)dx + e^*_{(r+1,s,t-1)} \int_{x \in \mathcal{T}^c_{(r,s,t)}} f(x)dx + \int_{x \in \mathcal{T}^c_{(r,s,t)}} \left[x - E\{x | x \in \mathcal{T}^c_{(r,s,t)}\} \right]^2 f(x)dx \right\}$$

Solution: Markov Case

Initialize $r_0 = 0,\ s_0 = M,\ t_0 = N$. For each k in $0 \leq k \leq N-1$ do

1. Lookup the optimal interval $\mathcal{T}^*_{(r_k,s_k,t_k)}$ from the table that was determined offline.
2. Observe x_k, and apply the sensor policy

$$\mu_k^*(x_k, (r_k, s_k, t_k)) = \begin{cases} (1, x_k) & \text{if } x_k \in \mathcal{T}^*_{(r_k,s_k,t_k)} \\ (0, \emptyset) & \text{if } x_k \in \mathcal{T}^{*c}_{(r_k,s_k,t_k)} \end{cases}$$

3. Apply the estimator, where $x_k \sim \mathcal{N}(0, \Sigma_{r_k}\sigma_w^2)$

$$\hat{\mu}_k^*(\mathcal{T}^*_{(r_k,s_k,t_k)}) = E\{x_k | x_k \in \mathcal{T}^{*c}_{(r_k,s_k,t_k)}\} = \frac{\int_{x_k \in \mathcal{T}^{*c}_{(r_k,s_k,t_k)}} x_k f(x_k) dx_k}{\int_{x_k \in \mathcal{T}^{*c}_{(r_k,s_k,t_k)}} f(x_k) dx_k}$$

4. Update

$$r_{k+1} = (1-\sigma_k)(r_k+1),\ s_{k+1} = s_k - \sigma_k,\ t_{k+1} = t_k - 1$$

Recap: Smart Sensing

Input parameters:

1. Planning horizon: Desired length of time the sensor will be in operation (N)
2. Battery size: Size of the power source installed into the sensor (M)
3. Process model (data): This is application specific, and can be gathered from a variety of sources including, measurement/historic data, modeling, etc. (Source Statistics)
4. Performance criterion: (Distortion Measure)

Smart Sensing

How?

- Given the input parameters and the channel compute the optimum transmission and decoding policies for the wireless sensor & the receiver
- The computation is "power aware" (M), and "planning-horizon aware" (N)
- Minimal online computational complexity (lookup table, binary comparisons,…)
- Offline computational complexity will depend on the source and channel models

Control over a Limited-Use Channel

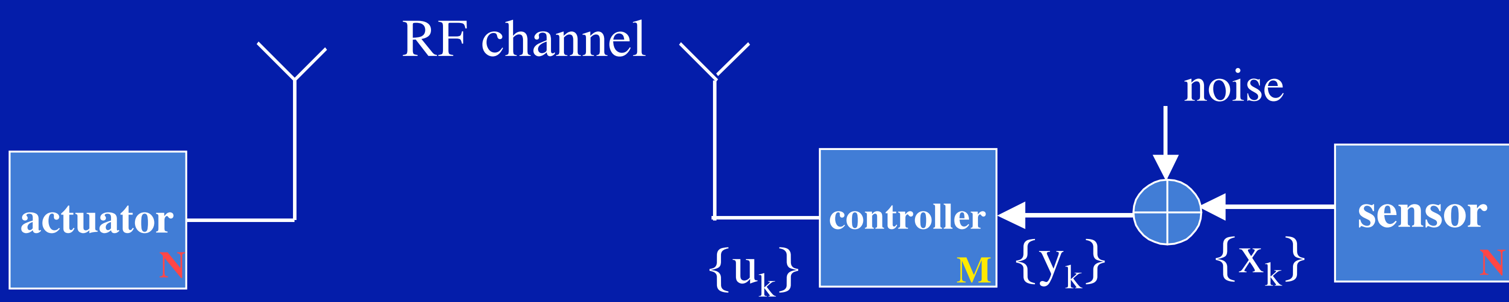

$$x_{k+1} = Ax_k + u_k + w_k \qquad x_k \in \mathbb{R}\ ,\ u_k \in \mathbb{R}\ ,\ y \in \mathbb{R}$$

$$y_k = x_k + v_k \qquad M < N$$

$$u_k = \mu_k(I_k),\ \ I_k := \{y_{[0,k]}, u_{[0,k-1]}\},\ \mu_k \rightarrow \mathbb{R} \text{ only M times}$$

Given a horizon of N units, and with control allowed to transmit for only $M < N$ times, what is the minimum attainable value of a performance index J, and a corresponding controller?

$$J = E\{ (x_N)^2 + \sum_0^{N-1} (x_k)^2 \}$$

Control over a Limited-Use Channel

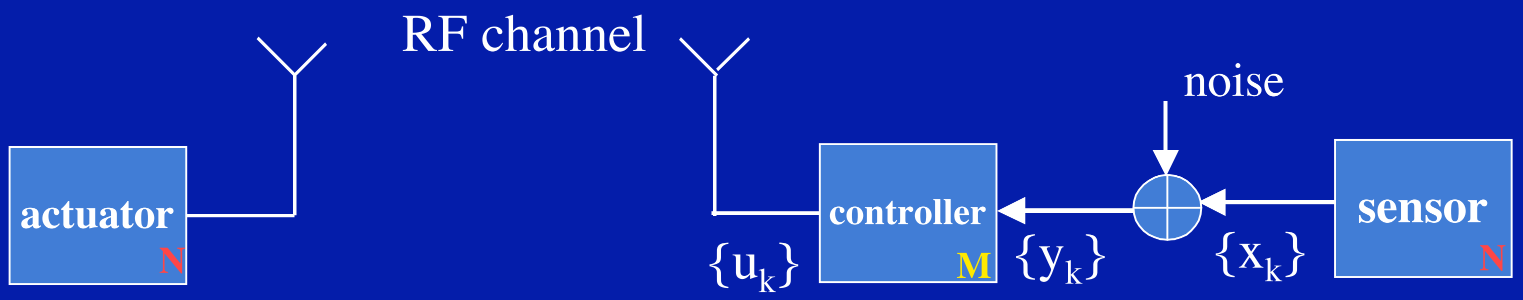

"Open-loop controller schedule": Determine *a priori* M time slots (out of N) when to apply control, and find the best such M slots as a result of combinatorial optimization.

Result: $u_k = -A\, E[x_k \mid I_k]$ for $0 \le k \le M-1$

$= 0$ for $M \le k \le N$

Closed-Loop Controller Schedule

Scheduling of control is on-line/*a posteriori*

- s : # control actions left (a total of M)
- t : # decision instances left (a total of N)
- Given (M,N), in retrograde time, t : 1 --> N

  $$\max\{0, M-(N-t)\} \leq s \leq \min\{t, M\}$$
- With t fixed, the maximum interval for t is [0, t]
- Approach is DP, moving forward in t

  (s,t) can lead to either (s-1, t-1) or (s, t-1)

Closed-Loop Controller Schedule (continued)

Online Computation

- For every k, x_k - $E[x_k \mid I_k]$ is independent of control
- If, at k, decision is to control, the optimum choice is

$$u_k = -A\, E[x_k \mid I_k]$$

where $E[x_k \mid I_k]$ is generated by Kalman filter.

- Decision is to control, if $\mid E[x_k \mid I_k] \mid$ exceeds a certain threshold, $\tau(s_k, t_k)$

$$\Longrightarrow \quad s_{k+1} = s_k - 1; \qquad t_{k+1} = t_k - 1$$

Closed-Loop Controller Schedule (continued)

Offline Computation

- $\tau(s_k, t_k)$ is a function of N, s_k, t_k, $\sigma^2_{k|k-1}$
- Compute optimum cost-to-go's at (s,t), based on whether a control was applied, $J_{(s,t)}^{(1)}$, or not, $J_{(s,t)}^{(0)}$
- $\Delta_{(s,t)} := J_{(s,t)}^{(0)} - J_{(s,t)}^{(1)} = 0$ determines the threshold (it has two roots -- one positive and one negative)

$$\Delta_{(s,t)}(\tau(s,t)) = \Delta_{(s,t)}(-\tau(s,t)) = 0$$

Numerical Solutions

Numerical integration was used to compute the recursions $\Delta_{(s,t)}$, which led to thresholds $\tau(s, t)$

Implemented the optimal control with M actions for an N-stage problem (N=20). Computed $J_{(M,N)}^*$ based on sample paths, and for different M values

Times of control action ==>

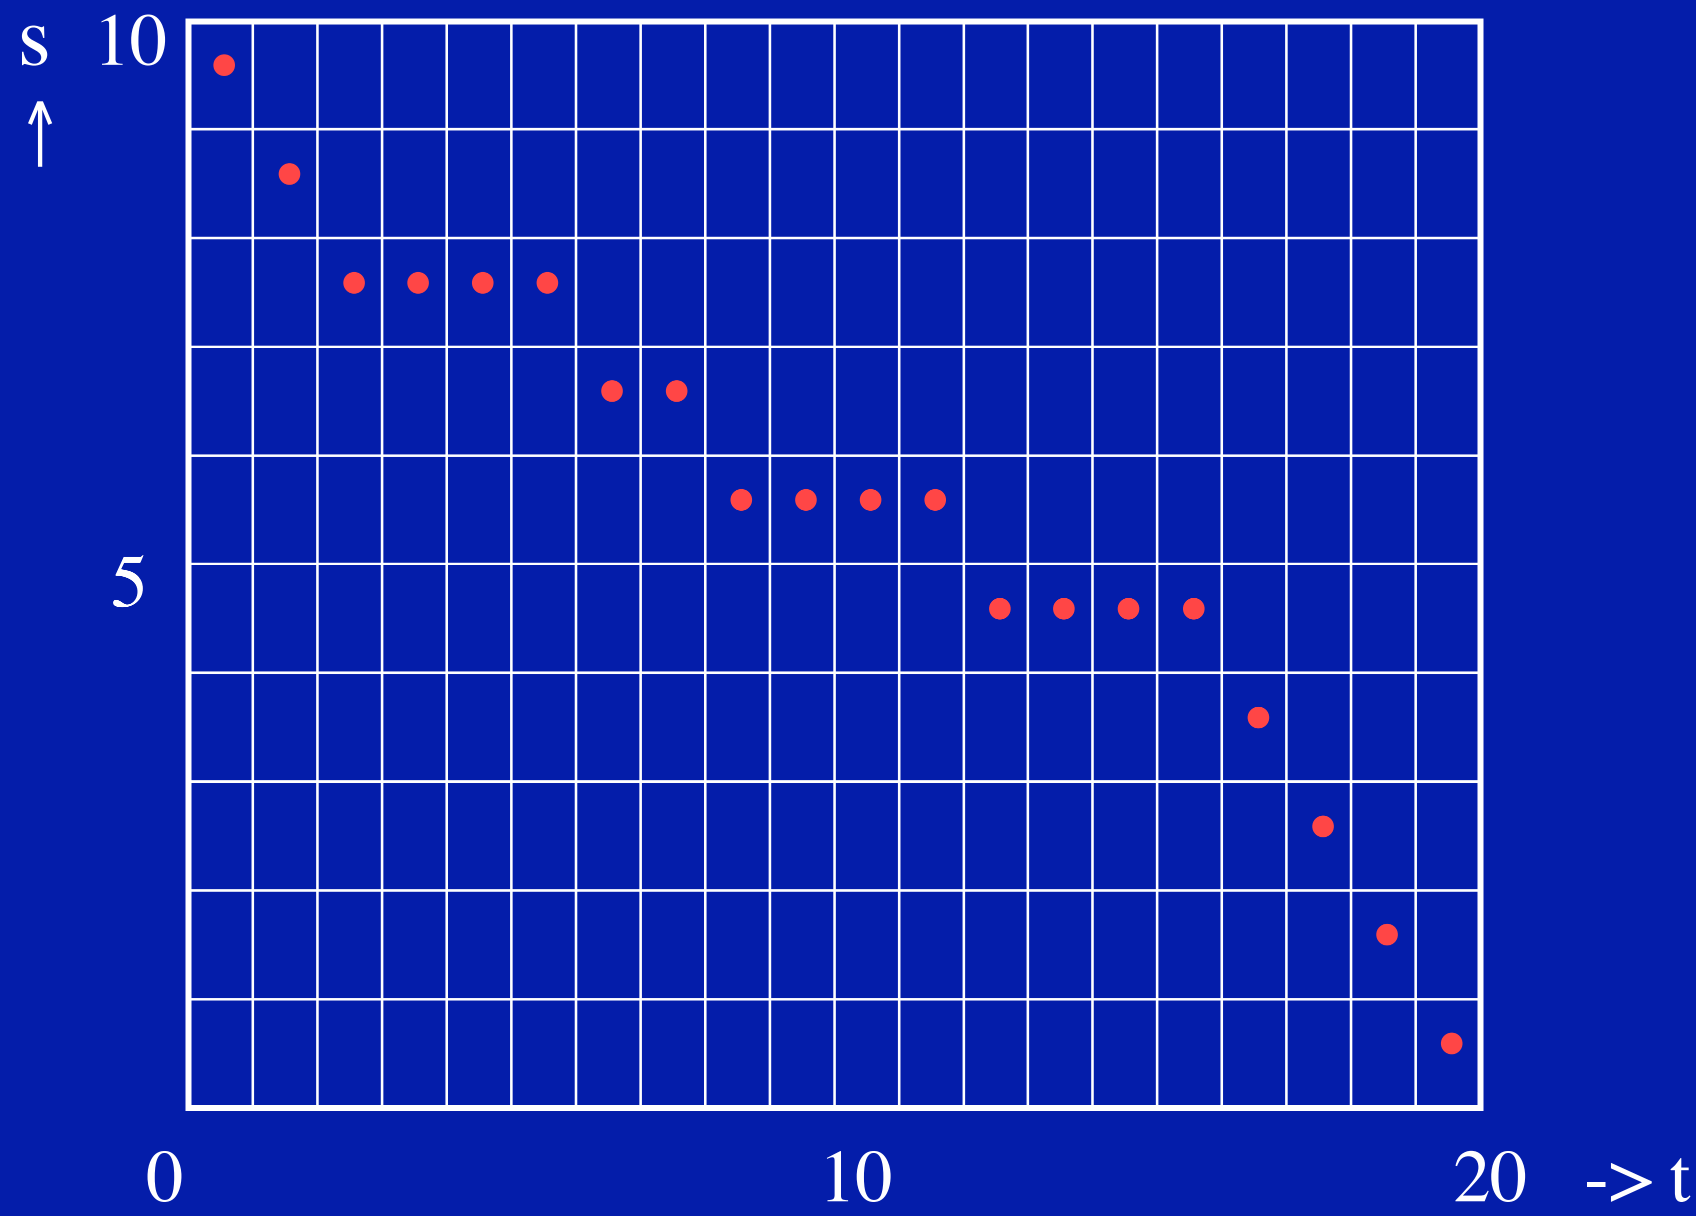
s
10
5
0
10
20
-> t

Numerical Solutions (continued)

$N = 20, A = 1, \sigma_w^2 = \sigma_v^2 = 1$

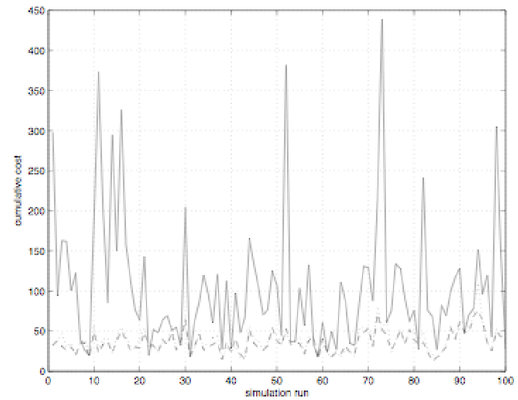

Sample path N-stage costs, for M = 1, 10, 20
100 simulation runs

M	$J^*_{(M,N)}$	%
1	96.4266	203.9327
2	68.1907	114.9343
3	47.2060	48.7914
4	44.0160	38.7366
5	39.8642	25.6503
6	37.1557	17.1132
7	35.6168	12.2627
8	34.1551	7.6555
9	33.6935	6.2005
10	33.6913	6.1936

M	$J^*_{(M,N)}$	%
11	33.2445	4.7853
12	32.9262	3.7820
13	32.8267	3.4684
14	32.4936	2.3249
15	32.1082	1.2037
16	31.9824	0.8072
17	31.8822	0.4914
18	31.8417	0.3637
19	31.7337	0.0233
20	31.7263	0

Other Settings

- Noisy measurements (KF with uncertain observations)
- Multidimensional case (conceptually similar)
- Cost on control (no longer linear)
- Unreliable/failure prone links for sensing and control transmission

Other Settings

- Noisy measurements (KF with uncertain observations)
- Multidimensional case (conceptually similar)
- Cost on control (no longer linear)
- Unreliable/failure prone links for sensing and control transmission
- **Multiple agents with infrequent exchange of information and decision**

THANKS !

QUESTIONS

DISCUSSION