RMM 19 Mar 07

◎

Distributed estimation & control

A. Distributed estimation

1. Problem formulation, examples
2. Information filter (centralized)
3. Information filter (on graph)

B. Distributed control

1. Witzenhausen example
2. Constrained LQR
3. Interconnected systems (quickly)

C. Sensor networks

1. Vijay max-cut
2. Ling trees (quickly)

No time

Distributed Sensor Fusion

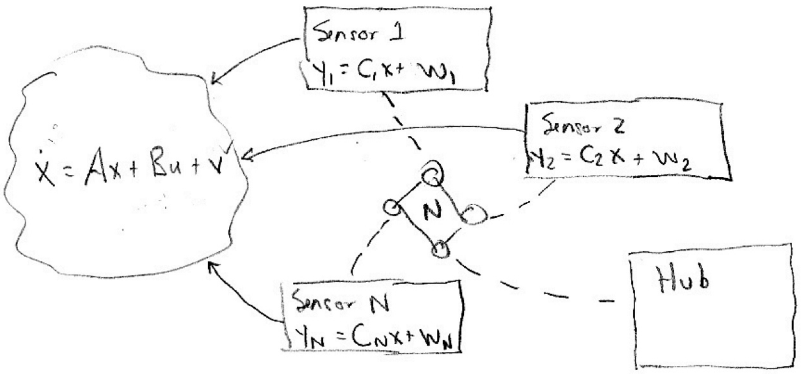

Consider a single process with multiple sensors, networked together
Would like to form estimate $\hat{x}$, either at each sensor or a central hub

Examples:

1. Alice - multiple sensors looking at environment plus possible need for different information at different points (eg. urban planning)

2. RoboFlag - each robot needs estimate of (local) environment plus players need to know entire environment

Case 1: centralized hub

Easy approach: sensors send data to hub, the hub runs (large) KF

Alternative approach: information filter

- Start with static case (no dynamics)

Minimum mean square error (MMSE) measurements

Let Y be a random variable that depends on X. (eg. $y = Cx + w$)

The minimum mean square error is given by minimizing

$$\min E\{(X-\hat{X})(X-\hat{X})^T\} \quad \text{over } X \ \& \ Y \text{ where } \hat{X} = g(Y)$$

Prop The MMSE estimate is given by $E\{X \mid Y = y\}$

pf Let $p_{X,Y}(x,y)$ denote the joint density function of X & Y.
For simplicity, we consider the scalar case

$$C = E\{(X-\hat{X})^2\} = \int_x \int_y (x - g(y))^2 \, p_{XY}(x,y) \, dx \, dy$$

$$= \int_y \left(\int_x (x-g(y))^2 \, p_{X|Y}(x|y) \, dx \right) p_Y(y) \, dy$$

check

$$\frac{\partial C}{\partial g(y)} = \int_y \left(\int_x 2(x-g(y)) \, p_{X|Y}(x|y) \, dx \right) p_Y(y) \, dy$$

why is this OK?

$$= 2\int_y \left(g(y) - \int_x x \, p_{X|Y}(x|y) \, dx \right) p_Y(y) \, dy$$

$$= 2\int_y \left(g(y) - E\{X \mid Y = y\} \right) p_Y(y) \, dy$$

This is minimized when $g(y) = E\{X \mid Y = y\}$

Remarks

1. Can extend to vector random variables
2. Can extend to dynamic process (eg Kalman filter)

Static Sensor Fusion

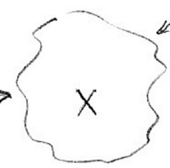

Random variable. Zero mean, cov R_X

$$\left.\begin{aligned} Y_1 &= H_1 X + W_1 \\ &\vdots \\ Y_q &= H_q X + W_q \end{aligned}\right\} \quad Y = HX + W$$

<u>Prop</u> Let W be zero mean, Gaussian noise with covariance R_W, independent of X. Then the MMSE estimate of X given $Y = y$ is

$$P^{-1}\hat{X} = H^T R_W^{-1} Y \qquad \text{where} \quad P^{-1} = \left(R_X^{-1} + H^T R_W^{-1} H\right)$$

<u>Pf</u> Matrix inversion formula (see notes)

<u>Prop</u> Suppose we have a set of independent sensors with

$$Y_i = H_i X + W_i \qquad i = 1, \ldots, n$$

and $\{W_i\}$ and X are uncorrelated. Let $\hat{X}_i$ & P_i be the estimates for individual sensors. Then

$$P^{-1}\hat{X} = \sum_{i=1}^{q} P_i^{-1}\hat{X}_i \qquad P^{-1} = \sum_{i=1}^{q} P_i^{-1} - (q-1) R_X^{-1}$$

<u>Pf</u> Exploit fact that R_W is block diagonal.

<u>Remarks</u>

1. This is a static version of the information filter
2. Each term in sum computed locally $\Rightarrow$ complexity of computation is shifted to sensors. Central hub just has to invert $n \times n$ matrix
3. Especially useful when $q \gg n$

Static Sensor Fusion on a graph

Suppose we have no central hub & want each sensor to converge to a global estimate.

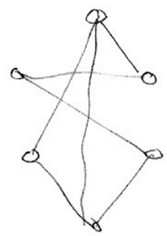

Weighted consensus algorithm

$$x_i[k+1] = x_i[k] + hW_i^{-1}\sum_{j\in N_i}\left(x_j[k] - x_i[k]\right)$$

$$= (I - hL)\,x[k]$$

$$\Rightarrow \quad x_i[\infty] = \sum W_i^{-1} x[0] \qquad \text{check}$$

To use for static sensor fusion, set $x[0] = \hat{x}_i$ and $W_i = P_i^{-1}$

Remarks

1. Can extend this simple approach to handle changing topologies; see Xiao, Boyd, Lall (2005)
2. XBL also look at properties of local estimates before final convergence & what happens when you have finite-bit communications (eg quantization)

Distributed Kalman Filtering

Suppose now that X is a random process with

$$x[k+1] = Ax[k] + v[k] \qquad v, w \text{ Gaussian, uncorrelated}$$

$$y_i[k] = C_i x[k] + w_i[k] \qquad x[0] \quad ''$$

Prop (Info filter) The Kalman filter equations for the centralized Kalman filter are

$$P^{-1}[k|k]\hat{x}(k|k) = P^{-1}(k|k-1)\hat{x}(k|k-1) - C^T R_w^{-1} y[k]$$

$$P^{-1}(k|k) = P^{-1}(k|k-1) - C^T R_w^{-1} C$$

Pf Matrix algebra (see notes)

Prop (Distributed info filter) For independent measurements

$$P^{-1}(k|k) = P^{-1}(k|k-1) + \sum_{i=1}^{N}\left(P_i^{-1}(k|k) - P_i^{-1}(k|k-1)\right)$$

$$P^{-1}(k|k)\hat{x}(k|k) = P^{-1}(k|k-1)\hat{x}(k|k-1) + \sum\left(P_i^{-1}(k|k)\hat{x}_i(k|k) - P_i^{-1}(k|k-1)\hat{x}_i(k|k-1)\right)$$

Proof Block diagonal R_w plus matrix manipulation

Remarks

1. Sensors compute local updates & send to fusion node; fusion node updates global estimate; similar to static case
2. Can extend to fully distributed, but need to be careful about rates of convergence

RMM 19 Mar 08

6a

Distributed Control

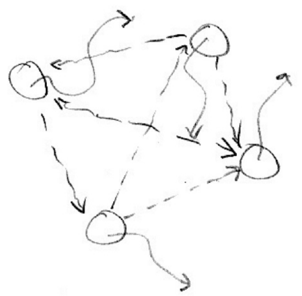

Given N dynamical systems plus a shared task, find decentralized controller that use local info + neighbors and stabilizes system to desired point.

Interconnected systems: suppose dynamics are coupled in a regular way

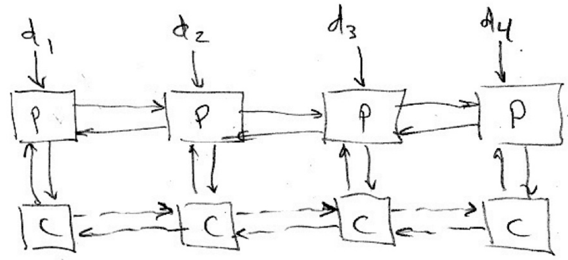

D'Andrea, Dullerud, Lall et al:

H_∞ techniques where control structure mirrors process interconnection

Useful for tightly coupled systems

Decoupled systems: agent dynamics are not coupled, but cost function (task) is coupled:

$$\dot{x}^i = f(x^i, u^i) \qquad J = \int_0^T L(x, u)\,dt + V(x(T))$$

Would like controller such that $u = \alpha(x^i, x^{-i})$ where $x^{-i} = \{x^j : j \in N_i\}$.

Naive approach: LQR with constraints on dependence of feedback

Get nonconvex problem; solution not necessarily a linear feedback

RMM 19 Mar 08

66

Witsenhausen example

Consider a very simple two stage process:

Stage 1: $x[1] = x[0] + u[0]$ $\quad u[0] = \gamma_0(y[0])$

$y[0] = x[0]$ $\quad R_{x_0} = \sigma^2$

Stage 2: $x[2] = x[1] - u[1]$ $\quad u[1] = \gamma_1(y[1])$

$y[1] = x[1] + w$ $\quad R_w = 1$

Goal: minimize $J = k^2 u^2[0] + x^2[2]$

Optimal answer with full info: $\gamma_0 = 0,\ \gamma_1 = x[0] \Rightarrow J = 0$

Linear answer: $\gamma_0 = a\,y[0] = a\,x[0],\ \gamma_2 = b\,y[1]$

$x[1]$ = Gaussian with variance $(1+a)^2\sigma^2$

Optimal $\gamma_1 = \hat{x}_1 = E\{x_1 \mid y_1\}$. Can show $b = \dfrac{(1+a)^2\sigma^2}{1+(1+a)^2\sigma^2}$ why?

$$E\{x[2]\} = \frac{(1+a^2)\sigma^2}{1+(1+a^2)\sigma^2} \Rightarrow J = k^2a^2\sigma^2 + \frac{(1+a)^2\sigma^2}{1+(1+a)^2\sigma^2}$$

Can numerically optimize. Eg $k = 0.1,\ \sigma = 10 \Rightarrow a = -0.0101$ and optimal cost is 0.99

Nonlinear answer: $u[0] = -y_1 + B\left\lfloor \frac{y_1}{B} + \frac{1}{2} \right\rfloor$ ← quantizer (sends info to stage 2)

$u[1] = B\left\lfloor \frac{y_2}{B} + \frac{1}{2} \right\rfloor$ ← noise small relative to bin size

Let n = # bins and set $k = \frac{1}{n^2},\ \sigma = n^2,\ B = n$

$\lim_{n\to\infty} E\{J_{nl}\} = 0 \qquad \lim_{n\to\infty} E\{J_{lin}\} = 1$

⇒ Linear solution is not optimal!

∴ Linear optimal control for distributed systems died.

RMM 19 Mar 08

9

Suboptimal Distributed Control

Ref: Gupta, Hassibi & M, IJC 2005

Consider a collection of linear systems w/ structured controller

$$\dot{x}^i = A^i x^i + B^i u^i \qquad u^i = K^i x^i + \sum_{j \in N_i} K^{ij}(x^i - x^j) \quad (*)$$

"constrained control law"

Questions:

1. Is it possible to stabilize a system using information from other vehicles if an individual system is unstable?
2. Are there information topologies that make system unstabilizable even if individual systems are stable?

Prop Consider the system & feedback structure $(*)$

(i) A formation is controllable $\Leftrightarrow$ Each individual agent is controllable

(ii) A formation is stabilizable $\Leftrightarrow$ " is stabilizable

Pf Check subspaces

Discrete-time optimal control

$$x^i[k+1] = A^i x^i[k] + B^i u^i[k] + v^i[k] \qquad u^i = \text{same}$$

Try to optimize finite horizon cost with $E\{x^T[0]\,x[0]\} = R[0]$

$$J_T = E\left\{ \sum_{k=0}^{T-1} x^T[k] Q x[k] + u^T[k] Q_u u[k] + x^T[T] P_T x[T] \right\}$$

Brute force: minimize over NT free parameters (might work)

RMM 19 Mar 08

8

Simpler approach (Gupta, 2004): Consider the infinite horizon quadratic cost ($T=\infty$, P_T dropped). Can show that the unconstrained optimal cost is given by

$$J_\infty = E\{x^T[0]\,P\,x[0]\} \qquad P = (Q_x + K^T Q_u K) + (A+BK)^T P (A+BK)$$

$$= \text{trace}\{P R[0]\} \qquad \text{(discrete Algebraic Ricatti equation)}$$

Let $R[k]$ be the covariance of $x[k]$, which evolves as

$$R[k+1] = (A+BK)\,R[k]\,(A+BK)^T$$

Assume controller is of the form $u = Kx$, $K = \sum_{i=1}^{n} \alpha_i \Phi_i$ where Φ_i span allowable fbks.

Skip. Just say "Solve numerically"

Compute gradiant of cost wrt α to find $\text{trace}\left(\frac{\partial P}{\partial \alpha} R[0]\right) = 0$:

$$\frac{\partial P}{\partial \alpha_i} = (A+BK)^T \frac{\partial P}{\partial \alpha_i}(A+BK) + \Sigma_i + \Sigma_i^T \qquad (*)$$

$$\Sigma_i = \Phi_i^T \left(Q_u K + B^T P (A+BK)\right)$$

By definition of $R[k+1]$

replace w/ (*)

$$\text{trace}\left((A+BK)^T \frac{\partial P}{\partial \alpha_i}(A+BK)\,R[0]\right) = \text{trace}\left(\frac{\partial P}{\partial \alpha_i} R[1]\right)$$

$$\vdots$$

$$\text{trace}\left(\frac{\partial P}{\partial \alpha_i} R[0]\right) = \text{trace}\left(\left(\frac{\partial P}{\partial \alpha} R(k) + \Sigma_i X[k] + \Sigma_i^T X[k]\right)\right)$$

$$X[k] = R[0] + R[1] + \cdots + R[k]$$

Eventually, $R[k] \to 0$ and $X \to$ steady state: $X = R[0] + (A+BK)X(A+BK)^T$

Finally, we need to solve

$$\text{trace}(\Sigma_i X + \Sigma_i^T X) = 0 \qquad K = \sum_{i=1}^{n} \alpha_i \Phi_i$$

$\therefore$ Get minimum value. Can extend to allow random disturbances as well...

Final comments on distributed estimation & control

- Distributed estimation is relatively straightforward and essentially solved for linear Gaussian
- Distributed control is much harder & optimality becomes difficult quickly (but RHC will still work...)