Information and Control

Witsenhausen Revisited

Sanjoy Mitter and Anant Sahai

Motivation

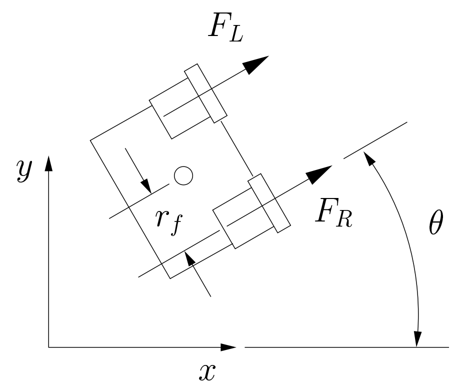

Consider the problem of trying to make the vehicle follow a certain trajectory. The vehicle is described by the following equations:

$$\begin{aligned} m\ddot{x} &= -\mu\dot{x} + (F_R + F_L)\cos\theta \\ m\ddot{y} &= -\mu\dot{y} + (F_R + F_L)\sin\theta \\ J\ddot{\theta} &= -\psi\dot{\theta} + (F_R - F_L)r_f \end{aligned}$$

Figure from T. Chung et al, "A Platform for Cooperative and Coordinated Control of Multiple Vehicles: The Caltech Multi-Vehicle Wireless Testbed," Conference on Control and Optimization, 2002.

Motivation

The equations can be linearized around the desired trajectory and the system can be put in the form

$$\dot{X} = AX + BU + w,$$

where, for instance, $X = \left[x, \dot{x}, y, \dot{y}, \theta, \dot{\theta}\right]^T$ and $U = [F_R, F_L]^T$. Given that we observe some linear combination of states

$$Y = CX + v,$$

our job is to determine u as a function of Y so as to minimize the cost function

$$E\left[\int_{t_0}^{t_1} (X^T(t)R_1X(t) + U^T(t)R_2U(t))dt\right].$$

Motivation

Under some assumptions (Gaussian white noise, positive matrices etc), the problem can 'easily' be solved. The optimal control law turns out to be linear,

$$U(t) = -F^0 \hat{X}(t),$$

where

$$F^0(t) = R_2^{-1} B^T P(t).$$

$P(t)$ is the solution to the matrix Riccati equation

$$-\dot{P}(t) = R_1 - P(t) B R_2^{-1} B^T P(t) + A^T P(t) + P(t) A.$$

Also $\hat{X}(t)$ is given by the observer

$$\dot{\hat{X}}(t) = A\hat{X}(t) + BU(t) + K(t)[y(t) - C\hat{X}(t)],$$

with

$$K^0(t) = Q(t) C^T(t) V_2^{-1}.$$

$Q(t)$ is the solution to another matrix Riccati equation

$$\dot{Q}(t) = V_1 - Q(t)C^T V_2^{-1} C Q(t) + A^T Q(t) + Q(t)A.$$

So can solve these equations and implement the controller.

Frankly, this is the only problem control theory *really* knows how to solve!

Motivation

So, given any problem, try to convert it to the standard problem.

But in the standard problem, there are standard assumptions:

- Linear system.
- Gaussian Random Variables.
- Quadratic Costs.
- Controller has access to all the observations and the controls

What if some of these fail?

Motivation

- Single vehicle, single observer. Random delays in the transfer of information due to communication channel.
- Single vehicle, multiple observers. May have different delays/ packet loss.
- Multiple vehicles, centralized controller.
- Multiple vehicles, decentralized controller.
- Multiplicative noises.

Motivation

Control problems surprisingly crop up everywhere.

- Traditional control applications.
- Biological organisms.
- Congestion control in Networks.
- Feedback loops in Electronic Amplifiers.
- Communications.
- And so on ...

This presentation would try to take a deeper look at the issue.

Problem Framework

- Consider a system described by the evolution of the state vector X.
- Time goes on in a discrete fashion from 1 to N.
- There are M sensors, each observing some function of the state variable.
- There are K control stations each supplying a control function based on some information that reaches them.
- The plant evolves according to a causal sequence of events.

Problem Framework

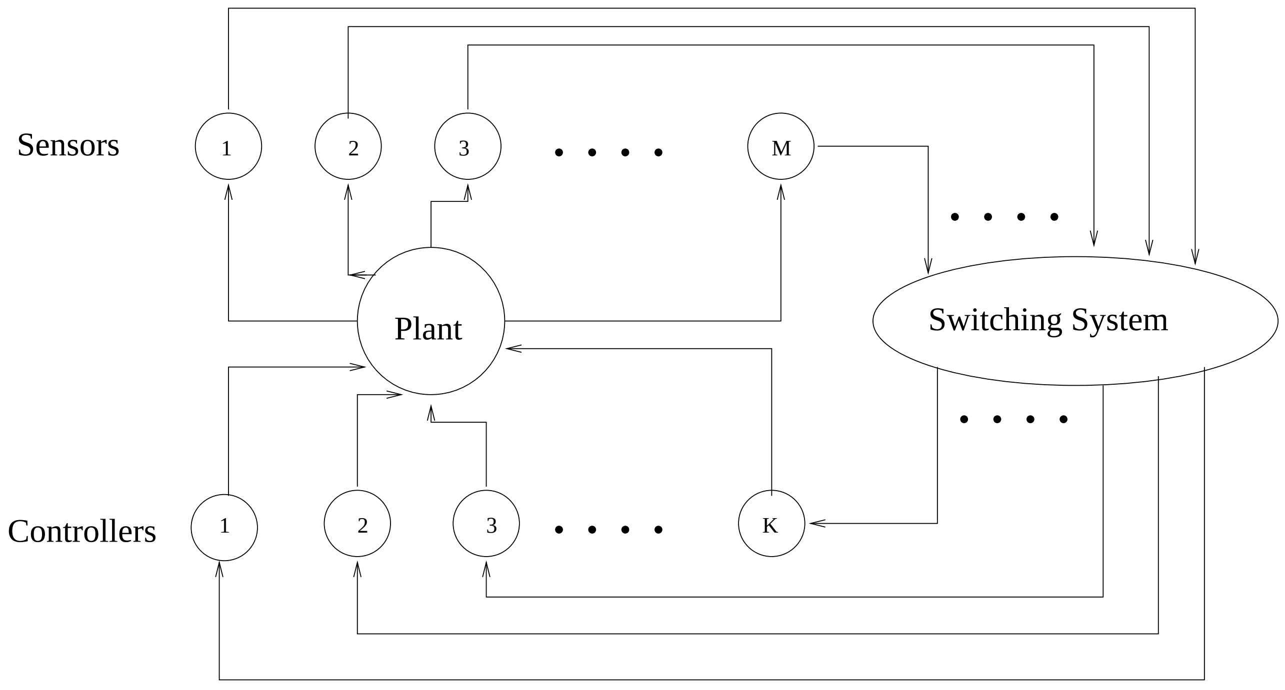

Problem Framework

- The initial state X_0 is generated (in a random fashion).
- Outputs $y_1^1, y_1^2, \ldots, y_1^M$ are observed.

$$\{y_1^i\} = \{g_1^i(X_0, w_1^i)\}, i = 1, 2, \ldots, M.$$

- Controls $\{u_1^i\}$, $i = 1, 2, \ldots, K$ are applied.
- The system state transitions to X_1.

$$X_1 = f_1(X_0, v_1, u_1^1, \ldots, u_1^K).$$

- The outputs are again observed, controls applied and the state transitions to X_2.
- This continues until final state X_N is reached.
- Uncertainity enters the system through the *primitive* random variables: initial state X_0, process noise v_k and sensor noise w_k^i.

Problem Framework

- There is a pre-specified cost function on the state and the controls

$$\sum_{k=1}^{N} h_k(X_k, u_k^1, \ldots, u_k^K).$$

- The designer picks the design γ that specifies the control laws γ_k^i for each controller i at each time step k.
- The controllers implement this control law to generate the control signals u_k^i. γ is designed to minimize the expected cost.
- The constraints on the designer are:
  - Information Pattern: Which controller has access to which observation variables?
  - Feasible Control Laws: What are the constraints on the functional form of the control laws?

Some Definitions

- **Defn**: If all the functions f_k and g_k^i are jointly linear in all their arguments, the system is called *linear*.

$$\begin{aligned} X_{n+1} &= A_{n+1} X_n + \sum_{i=1}^{K} B_{n+1}^i u_{n+1}^i + C_{n+1} v_{n+1} \\ y_{n+1}^i &= D_{n+1}^i X_n + E_{n+1}^i w_{n+1}^i . \end{aligned}$$

- Define the following sets of pairs of indices. For $k = 1, 2, \ldots, N$,

$$Y_k = \{(\tau, m) | \tau = 1, \ldots, k; m = 1, \ldots, M\}.$$

For $k = 2, \ldots, N+1$,

$$U_k = \{(\tau, m) | \tau = 1, \ldots, k-1; m = 1, \ldots, K\}.$$

Defn: A *data basis* at time k is a pair (A, B) where A is a subset of Y_k and B is a subset of U_k.

Information Patterns

- **Defn**: An *information pattern* is the assignment, to each index pair (k, m) in U_{N+1}, of a data basis at time k, denoted by $(Y_{k,m}, U_{k,m})$.
  - "What the controller m knows at the time step k."
- **Defn**: Given an information pattern, control station m is said to have a *perfect recall* when either $N = 1$ or else if $Y_{k,m} \subset Y_{k+1,m}$, $U_{k,m} \subset U_{k+1,m}$ for $k = 1, \ldots N - 1$.
  - "Once a controller gets some information, it retains it for all time to come."
- **Defn**: An information pattern is *classical* if the information pattern is independent of the subscript m and the information pattern has perfect recall.
  - "All the systems receive the same information and remember it."

Information Patterns

- **Defn**: A *delayed sharing pattern* is one in which all K control stations have perfect recall, each having an associated observation post. They receive their own data immediately and share it with a delay of n steps.
  - $K = 1$, a single station with access to all past outputs and controls: *Strictly Classical Pattern.*
  - $N = 1, K > 1$, K stations each acting only once: *Team Theory.*
  - $n = N > 1$, K stations without any 'lateral' sharing of information.

More Definitions

- Let (Y, U) be a data basis at time k. Does it make sense to talk of the conditional distribution of x_{k-1} given the data specified by the data basis (Y, U)?
- In general, information about control laws is also required.
- For L, a subset of U_{N+1}, define γ_L to be the specification of the functional form of just those control laws $\gamma_{\theta,m}$ with subscripts (θ, m) in L.
- **<u>Defn</u>**: A triple (Y, U, L) is called a *field basis* if for any two designs γ and $\hat{\gamma}$, the relation $\gamma_L = \hat{\gamma}_L$, that is, $\gamma_\theta^m(.) \equiv \hat{\gamma}_\theta^m$ for (θ, m) in L, implies that the sigma fields generated by $(Y, U; \gamma)$ and by $(Y, U, \hat{\gamma})$ are the same.
  - So the knowledge of just those control laws designated by L is sufficient to determine unambiguously the field generated by the variables designated by (Y, U).

More Definitions

- **Example**: The triple $(\emptyset, \emptyset, L)$ is a field basis for any subset L of U_{N+1}, regardless of the information pattern. Similarly for a single-input, single-output scalar system, the triple $(\{1,1\},\{1,1\},\emptyset)$ is a field basis.
- Note that a field basis does not necessarily determine a conditional distribution.
- **Defn**: The triple (Y, U, L) is called a *conditioning basis* for a variable z if there exists a function F, such that for any design γ, the conditional distribution of z given data (Y, U) is almost surely $F(y_N, u_U, \gamma_L)$.
- **Example**: For any $Y \subset Y_k$, $U \subset U_k$ and any information pattern, the triple (Y, U, U_k) is a field basis and a conditional basis for X_{k-1}.

Assertions

Witsenhausen, in his survey paper, states 11 assertions. We cover the 4 which Mitter-Sahai reproduce.

- **Assertion**: Any problem with classical information pattern is equivalent to a problem with strictly classical pattern, though additional state variables may be required.
  - Classical information pattern: "All the systems receive the same information and remember it."
  - Strictly classical information pattern: "There is only one controller with access to every observation and perfect memory."
  - Any control law possible with classical information pattern is also possible with strictly classical information pattern.

Assertions

- **Assertion**: If for every (k, m), $(Y_{k,m}, U_{k,m}, \emptyset)$ is a field basis, then the given feedback control problem is equivalent to a feedforward control problem. The data available depends only on the primitive random variables and not upon the control variables applied.
  - Feedback control problem: "Control laws use the observations from the plant output."
  - Feedforward control problem: "Control laws never use any observation from the plant, so the actual control signals can be computed ahead of time."
  - $(Y_{k,m}, U_{k,m}, \emptyset)$ being a field basis for every (k, m): "At any time, whether I know the functional form of the control law does not matter in calculating the probabilities of the various variables."
  - Under this condition, the infimum of the cost achievable by any possible control law is the same as achieved by the best feedforward control law.

Assertions

- **Assertion**: For a problem with strictly classical pattern, there exists a conditional distribution of X_{k-1} given the variables $(y_1, \ldots, y_k, u_1, \ldots, u_{k-1})$ independent of the design γ.
  - Perfectly classical information pattern: "There is only one controller with access to every observation and perfect memory."
  - Remember, the system evolved as

$$X_{k+1} = f_{k+1}(X_k, v_1, u_1^1, \ldots, u_1^K).$$

- Further, if we let F_k denote this conditional distribution, then there is no loss if we restrict the designs to be of the form $\gamma_k = \phi_k(F_k)$.
  - No loss: The infimum of costs possible with *any* control law is achievable with designs of form given above.

Assertions

So the optimum control law for the problem is of the form

$$\gamma_k = \phi_k(F_k).$$

- Evaluation of F_k is called the *filtering problem* and is independent of the criterion and of the control laws.
- ϕ_k is a function defined over the space of distributions for X_{k-1}.
- Thus the conditional law is a sufficient statistic for control purposes.
- A *one-way* separation exists between filtering and control. F_k can be computed with no concern to the design and then the design law is specified.

Assertions

- **Assertion**: For a perfectly classical linear Gaussian system, the conditional distribution of X_{k-1} has a Gaussian version with covariance independent of the data and mean affine in the data.
  - A system is Gaussian if all the primitive random variables (noises and the initial state) have Gaussian distributions.
  - Remember, for a perfectly classical system, the optimal control law was of the form $\gamma_k = \phi_k(F_k)$, where F_k was the conditional distribution of X_{k-1}.
  - In addition, if the system is linear and Gaussian, then this conditional distribution is also Gaussian. Hence it is completely specified by its mean and covariance.
  - The calculation for the covariances can be done through linear discrete time-filtering off-line, without knowledge of the design.
  - Whether or not the design γ is affine is irrelevant.

Assertions

- **Assertion**: Consider a perfectly classical linear system with a quadratic cost criteria. If we find the optimal control law $\phi_k(X_{k-1})$ for the same system assuming perfect observation and setting
  - all the primitive random variables v_k at their mean values,
  - with the initial distribution for the state x_{k-1} as a point mass at X_{k-1}.

  Then the optimal control law for the original system is given by $\phi_k(\overline{F_k})$, where $\overline{F_k}$ is the mean of the conditional distribution for X_{k-1} and can be determined by a recursive filter.

  Both ϕ_k and $\overline{F_k}$ are affine and their determination can be accomplished entirely independent of each other!

Counterexample

- There were four main assumptions for optimality of affine control laws:
  - Linear systems,
  - Gaussian primitive variables,
  - Quadratic Cost, and
  - Perfectly classical information patterns.
- Is the fourth one required?
- Yes, consider the counter-example proposed.

Problem Parameters

- N = 2 time steps.
- Scalar state variable x.
- Gaussian initial condition x_0, with zero mean and variance σ^2. Say $\sigma = 10$.
- At time step 1, the output is $y_1 = g_1(x_0) = x_0$, information to be used for calculating the control u_1 is $Y_1 = \{y_1\}$ and $U_1 = \emptyset$. The state transitions as $x_1 = f_1(x_0, u_1) = x_0 + u_1$.
- At time step 2, the output is $y_2 = g_2(x_1) = x_1 + w$ (w zero mean, unit variance Gaussian noise), information to be used for calculating the control u_2 is $Y_2 = \{y_2\}$ and $U_2 = \emptyset$. The state transitions as $x_2 = f_2(x_1, u_2) = x_1 - u_2$.
- The cost expressions are $h_1(x, u) = k^2u^2$, $h_2(x, u) = x^2$. Say $k = 0.1$.
- The information pattern is memoryless.

Problem Solutions

- Perfectly classical information pattern would have been $Y_1 = \{y_1\}$, $U_1 = \emptyset$, $Y_2 = \{y_1, y_2\}$ and $U_2 = \{u_1\}$. The solution in that case would have been the affine control laws: $\gamma_1(y_1) = 0$ and $\gamma_2(y_1, y_2, u_1) = y_1$ with minimum cost as zero.
- Under the specified information pattern, can show that the optimal affine control laws can only give an expected cost 0.99.
- Non-linear controllers $\gamma_1(y_1) = -y_1 + \sigma\ \mathrm{sgn}(y_1)$ and $\gamma_2(y_2) = \sigma\ \tanh(\sigma y_2)$ give expected cost 0.404.
- Even better non-linear controllers are possible which drive the cost down asymptotically to zero.
- Thus affine controllers are not optimal in the absence of perfectly classical information pattern.

Discussion: Why did this happen?

- Note the form of the cost function. At the second stage, all we are penalizing is x_2, which is calculated by $(x_1 - u_2)$. The controller has access only to $y_1 = x_1 + w$, and we are asking it to estimate x_1 as best as it can, so that we can set u_2 as close to x_1 as we can.
- On the other hand, at the first stage, we are not penalizing the state x_1. We are giving the controller a lot of freedom by asking it to minimize only u_1 without consideration for x_1.
- Thus there are two objectives:
  - We are asking for x_1 to be of low entropy so that it can easily be predicted.
  - We are asking for x_1 to be of high energy so w is not able to affect it much.

Discussion

- Affine controllers means gaussian states and for a gaussian random variable, high energy implies high entropy while low energy implies low entropy. Thus the two objectives are in direct opposition.
- Non-linear controllers get around this by making x_1 as a two-point distribution. Thus they can achieve both objectives.
- Note that this issue became important only when the information pattern was not perfectly classical. Only then was the problem reminiscent of a communication problem and we were involved in entropy and energy issues.

Conclusions

- Traditionally the control problem formulation has involved only estimation and control. However, if the information pattern is not perfectly classical, there is also a strong communication aspect present.
- In particular, how the different parts of the system share information amongst themselves has a strong bearing on both the estimation algorithms and the control algorithms.
- A delayed sharing pattern is not classical. Thus when communication channels are present and introduce delays, affine control laws might not be optimal.

Future Directions

- Further study of linear Gaussian quadratic control problems with general information pattern is required. This would help in recasting problems in several domains like communications as control problems and solving them.
- For specific given information patterns, existence of optimum affine control laws should be examined. Work has been done by using Markov state models on communication channels and networks present in the system.
- Algorithms for approaching an optimal solution need to be developed. Affine control laws yield several local minima, even though that class might not be complete.