LQ optimal control for partially specified input noise

Alexander Erreygers Jasper De Bock Gert de Cooman
Arthur Van Camp

Ghent University

28th European Conference on Operational Research

Scalar linear systems

The *controller* is interested in the system

$$X_{k+1} = aX_k + bu_k + W_k, \tag{1}$$

for $k \in N = \{0, 1, \ldots, n\}$, where $n \in \mathbb{N}$, $a \in \mathbb{R}$ and $b \in \mathbb{R} \setminus \{0\}$, where

X_{k+1} is the real-valued *state*,

u_k is the real-valued *control input*,

W_k is the real-valued stochastic *noise*.

In general, system parameters a and b can be time dependent.

Scalar linear systems

The *controller* is interested in the system

$$X_{k+1} = aX_k + bu_k + W_k. \tag{1}$$

Observation assumptions

1. Before applying u_k, the controller observes the actual value x_k of X_k (hence $X_0 \equiv x_0$).
2. The controller has perfect recall.

Scalar linear systems

The *controller* is interested in the system

$$X_{k+1} = aX_k + b\phi_k(X^k) + W_k. \qquad (1)$$

Observation assumptions

1. Before applying u_k, the controller observes the actual value x_k of X_k (hence $X_0 \equiv x_0$).
2. The controller has perfect recall.

Controller determines u_k from state history $x^k \coloneqq (x_0, \ldots, x_k)$:

$$u_k = \phi_k(x^k).$$

ϕ_k : $\mathbb{R}^{k+1} \to \mathbb{R}$ is a feedback function,

ϕ $\coloneqq (\phi_0, \ldots, \phi_n)$ is a *control policy*,

Φ denotes the set of all control policies.

Scalar linear systems

The *controller* is interested in the system

$$X_{k+1} = aX_k + b\phi_k(X^k) + W_k. \tag{1}$$

Observation assumptions

1. Before applying u_k, the controller observes the actual value x_k of X_k (hence $X_0 \equiv x_0$).
2. The controller has perfect recall.

Controller knows x^k and $\phi \rightarrow$ can calculate w^{k-1}.

Optimality of a control policy

For any control policy $\phi \in \Phi$, any $k \in N$ and any state history $x^k \in \mathbb{R}^{k+1}$ we define the *quadratic cost functional* as

$$J[\phi|x^k] \coloneqq \sum_{\ell=k}^{n} r\phi_\ell(x^k, X_{k+1:\ell})^2 + qX_{\ell+1}^2,$$

where $q \geq 0$ and $r > 0$ are real-valued coefficients.

Precise noise model

Definition (Precise noise model or PNM)

The controller's beliefs about the noise $W_0, \ldots, W_n$ are modelled using a linear expectation operator E.

Optimality of a control policy

For any control policy $\phi \in \Phi$, any $k \in N$ and any state history $x^k \in \mathbb{R}^{k+1}$ we define the *quadratic cost functional* as

$$J[\phi|x^k] := \sum_{\ell=k}^{n} r\phi_\ell(x^k, X_{k+1:\ell})^2 + qX_{\ell+1}^2,$$

where $q \geq 0$ and $r > 0$ are real-valued coefficients.

Definition (Optimality)

A control policy $\hat{\phi}$ is *optimal* if for all x_0

$$\hat{\phi} \in \arg\min_{\phi \in \Phi} \mathrm{E}(J[\phi|x_0]).$$

Optimality of a control policy

Assume that at time k the controller knows the state history x^k and noise history w^{k-1}.
We should only compare control policies $\phi \in \Phi$ that could have resulted in x^k and w^{k-1}, i.e. such that x^k, w^{k-1} and ϕ are a solution of the system dynamics.

$$\Phi(x^k, w^{k-1}) \coloneqq \big\{\phi \in \Phi\colon \phi, x^k \text{ and } w^{k-1} \text{ are a solution of the system dynamics.}\big\}$$

Optimality of a control policy

Assume that at time k the controller knows the state history x^k and noise history w^{k-1}.
We should only compare control policies $\phi \in \Phi$ that could have resulted in x^k and w^{k-1}, i.e. such that x^k, w^{k-1} and ϕ are a solution of the system dynamics.

$$\Phi(x^k, w^{k-1}) \coloneqq \big\{\phi \in \Phi \colon \phi, x^k \text{ and } w^{k-1} \text{ are} \\ \text{a solution of the system dynamics.}\big\}$$

Definition (Optimality)

A control policy $\hat{\phi}$ is *optimal* for the state history x^k and the noise history w^{k-1} if

$$\hat{\phi} \in \operatorname*{arg\,min}_{\phi \in \Phi(x^k, w^{k-1})} \mathrm{E}(J[\phi | x^k] | w^{k-1}).$$

The principle of optimality

Principle of optimality

A control policy that is “optimal” for the “current state” should also be optimal for the “remaining states” it can end up in.

The principle of optimality

Principle of optimality

A control policy that is "optimal" for the "current state" should also be optimal for the "remaining states" it can end up in.

Assume that $\hat{\phi}$ is optimal for all $x_0 \in \mathbb{R}$.
The controller

1. observes x_0,
2. applies $u_0 = \phi_0(x_0)$,
3. observes x_1 and computes w_0.

Is $\hat{\phi}$ *optimal for* (x_0, x_1) *and* w_0*?*

The principle of optimality

Principle of optimality

A control policy that is “optimal” for the “current state” should also be optimal for the “remaining states” it can end up in.

Assume that $\hat{\phi}$ is optimal for all $x_0 \in \mathbb{R}$.
The controller

1. observes x_0,
2. applies $u_0 = \phi_0(x_0)$,
3. observes x_1 and computes w_0.

Is $\hat{\phi}$ *optimal for* (x_0, x_1) *and* w_0*?* Not necessarily!

Definition (Complete optimality)

If for all $k \in N$ the control policy $\phi \in \Phi$ is optimal for all x^k and w^{k-1} such that x^k, w^{k-1} and ϕ are compatible, then it is *completely optimal*.

Unique optimal control policy

Theorem

The unique completely optimal control policy $\hat{\phi}$ *is given by*

$$\hat{\phi}_k(x^k) := -\tilde{r}_k b \left(m_{k+1} a x_k + h_{k|w^{k-1}} \right).$$

$\tilde{r}_k$ and m_{k+1} are derived from backwards recursive relations.

Feedforward $h_{k|w^{k-1}}$ is derived from $h_{n+1|w^n} := 0$ and

$$h_{k|w^{k-1}} := a\tilde{r}_{k+1} r \mathrm{E}(h_{k+1|w^{k-1},W_k} | w^{k-1}) + m_{k+1} \mathrm{E}(W_k | w^{k-1}).$$

Unique optimal control policy

Theorem

The unique completely optimal control policy $\hat{\phi}$ is given by

$$\hat{\phi}_k(x^k) \coloneqq -\tilde{r}_k b \left(m_{k+1} a x_k + h_{k|w^{k-1}} \right).$$

$\tilde{r}_k$ and m_{k+1} are derived from backwards recursive relations.

Feedforward $h_{k|w^{k-1}}$ is derived from $h_{n+1|w^n} \coloneqq 0$ and

$$h_{k|w^{k-1}} \coloneqq a\tilde{r}_{k+1} r \mathrm{E}(h_{k+1|w^{k-1},W_k}|w^{k-1}) + m_{k+1}\mathrm{E}(W_k|w^{k-1}).$$

- − Precise specification of noise model is necessary.
- − Calculating the feedforward is intractable.
- − Backwards recursive calculations

- + Almost immediately generalisable to time-dependent a_k, b_k, r_k and q_{k+1} and/or multi-dimensional systems.

Unique optimal control policy

Disadvantages

- Calculating the feedforward is intractable.

  Feedforward $h_{k|w^{k-1}}$ is derived from $h_{n+1|w^n} := 0$ and

  $$h_{k|w^{k-1}} := a\tilde{r}_{k+1} r \mathrm{E}(h_{k+1|w^{k-1},W_k}|w^{k-1}) + m_{k+1}\mathrm{E}(W_k|w^{k-1}).$$

Unique optimal control policy

Disadvantages

- Calculating the feedforward is intractable.
- S White noise model: $W_0, \ldots, W_n$ are mutually independent. Feedforward h_k is derived from $h_{n+1} \coloneqq 0$ and

$$h_k \coloneqq a\tilde{r}_{k+1} r h_{k+1} + m_{k+1} \mathrm{E}(W_k).$$

Unique optimal control policy

Disadvantages

- − Calculating the feedforward is intractable.
- S White noise model: $W_0, \dots, W_n$ are mutually independent.
  Feedforward h_k is derived from $h_{n+1} \coloneqq 0$ and

$$h_k \coloneqq a\tilde{r}_{k+1} r h_{k+1} + m_{k+1} \mathrm{E}(W_k).$$

- − Backwards recursive calculations
- S White noise model & stationarity simplify these calculations.
  If $\mathrm{E}(W_k) \equiv \mathrm{E}(W)$ for all $k \in N$, then

$$m_{k+1} \xrightarrow[n\to\infty]{} m, \qquad \tilde{r}_k \xrightarrow[n\to\infty]{} \tilde{r}, \qquad h_k \xrightarrow[n\to\infty]{} h.$$

Partially specified noise model

- Precise specification of noise model is necessary.

Partially specified noise model

- Precise specification of noise model is necessary.

Definition (Partially specified noise model or PSNM)

The *partially specified noise model* $\mathcal{E}$ is the largest subset of the set of all precise noise models such that for all $\mathrm{E} \in \mathcal{E}$, all $k \in N$ and all w^{k-1}

$$\underline{\mathrm{E}}(W_k) \leq \mathrm{E}(W_k | w^{k-1}) \leq \overline{\mathrm{E}}(W_k).$$

Note: $\mathcal{E}$ does not assume independence!

Partially specified noise model

- Precise specification of noise model is necessary.

Definition (Partially specified noise model or PSNM)

The *partially specified noise model* $\mathcal{E}$ is the largest subset of the set of all precise noise models such that for all $\mathrm{E} \in \mathcal{E}$, all $k \in N$ and all w^{k-1}

$$\underline{\mathrm{E}}(W_k) \leq \mathrm{E}(W_k | w^{k-1}) \leq \overline{\mathrm{E}}(W_k).$$

Note: $\mathcal{E}$ does not assume independence!

Definition (E-admissibility)

A control policy is *E-admissible* if it is completely optimal for at least one precise noise model in the partially specified noise model.

E-admissible control policies

From the definition of E-admissibility, it follows immediately that any E-admissible control policy has the form

$$\phi_k(x^k) = -\tilde{r}_k b \left(m_{k+1} a x_k + h_{k|w^{k-1}} \right).$$

E-admissible control policies

Theorem

For any E-admissible control policy, the feedfworward term $h_{k|_{w^{k-1}}}$ is bounded: for all $k \in N$ and for all noise histories w^{k-1},

$$\underline{h}_k \leq h_{k|_{w^{k-1}}} \leq \overline{h}_k.$$

Moreover, any $h_{k|_{w^{k-1}}} \in [\underline{h}_k, \overline{h}_k]$ is reached by some $\mathrm{E} \in \mathcal{E}$.

Strict bounds $\underline{h}_k$ and $\overline{h}_k$ are derived from $[\underline{h}_{n+1}, \overline{h}_{n+1}] \coloneqq 0$ and

$$[\underline{h}_k, \overline{h}_k] \coloneqq a\tilde{r}_{k+1} r[\underline{h}_{k+1}, \overline{h}_{k+1}] + m_{k+1}[\underline{\mathrm{E}}(W_k), \overline{\mathrm{E}}(W_k)].$$

E-admissible control policies

Theorem

For any E-admissible control policy, the feedfworward term $h_{k|w^{k-1}}$ is bounded: for all $k \in N$ and for all noise histories w^{k-1},

$$\underline{h}_k \leq h_{k|w^{k-1}} \leq \overline{h}_k.$$

Moreover, any $h_{k|w^{k-1}} \in [\underline{h}_k, \overline{h}_k]$ is reached by some $\mathrm{E} \in \mathcal{E}$.

- + Imprecise specification
- + Computation of $\underline{h}_k$ and $\overline{h}_k$ is tractable.
- + Easily generalised to a_k, b_k, r_k and q_{k+1}.

- ? Which control policy to apply?
- − Backwards recursive calculations
- ? Generalisation to multi-dimensional systems is not immediate.

- Backwards recursive calculations
- *Stationarity* of bounds on expectation simplifies these calculations.

  If $\underline{\mathrm{E}}(W_k) \equiv \underline{\mathrm{E}}(W)$ and $\overline{\mathrm{E}}(W_k) \equiv \overline{\mathrm{E}}(W)$ for all $k \in N$, then

$$m_{k+1} \xrightarrow[n\to\infty]{} m, \quad \tilde{r}_k \xrightarrow[n\to\infty]{} \tilde{r}, \quad \underline{h}_k \xrightarrow[n\to\infty]{} \underline{h}, \quad \overline{h}_k \xrightarrow[n\to\infty]{} \overline{h}.$$

- Backwards recursive calculations
- *Stationarity* of bounds on expectation simplifies these calculations.
  If $\underline{\mathrm{E}}(W_k) \equiv \underline{\mathrm{E}}(W)$ and $\overline{\mathrm{E}}(W_k) \equiv \overline{\mathrm{E}}(W)$ for all $k \in N$, then

$$m_{k+1} \xrightarrow[n\to\infty]{} m, \quad \tilde{r}_k \xrightarrow[n\to\infty]{} \tilde{r}, \quad \underline{h}_k \xrightarrow[n\to\infty]{} \underline{h}, \quad \overline{h}_k \xrightarrow[n\to\infty]{} \overline{h}.$$

- Which control policy to apply?
- Possibility of using a secondary decision criterion.

Summary

- The partially specified noise model only assumes bounds on the conditional expectation of the noise.

Summary

- The *partially specified* noise model only assumes bounds on the conditional expectation of the noise.
- Every E-admissible control policy is a combination of the same *state feedback* and possibly different noise feedforward.

Summary

- The *partially specified* noise model only assumes bounds on the conditional expectation of the noise.
- Every E-admissible control policy is a combination of the same *state feedback* and possibly different *noise feedforward*.
- Tight bounds on E-admissible noise feedforward can be easily calculated.
  How to choose which element in the feedforward interval to apply remains an open question.

Summary

- The *partially specified* noise model only assumes bounds on the conditional expectation of the noise.
- Every E-admissible control policy is a combination of the same *state feedback* and possibly different *noise feedforward*.
- *Tight bounds* on E-admissible noise feedforward can be easily calculated.

  How to choose which element in the feedforward interval to apply remains an open question.
- Unfortunately, these results are not immediately generalised to multi-dimensional systems.