Optimal control of linear systems with quadratic cost and imprecise forward irrelevant input noise

Alexander Erreygers, Jasper De Bock, Gert de Cooman and Arthur Van Camp

{alexander.erreygers, jasper.debock, gert.decooman, arthur.vancamp}@UGent.be

SYSTeMS research group, Ghent University, Belgium

UNIVERSITEIT GENT

Linear systems

We consider a finite-state, discrete-time scalar linear system with a deterministic (known) current state $X_k = x_k$. For all $\ell \in \{k, \dots, k_1\}$, the dynamics of the system is described by

$$X_{\ell+1} = a_\ell X_\ell + b_\ell u_\ell + W_\ell. \qquad \text{(DYN)}$$

In this expression, a_ℓ and b_ℓ are real-valued parameters and the state X_ℓ and noise W_ℓ at time ℓ are real-valued random variables. The control input u_ℓ at time ℓ is also real-valued.

State feedback Usually the control input u_ℓ is taken to be some real-valued function ψ_ℓ of the previous states $x_{k+1:\ell} := (x_{k+1}, x_{k+2}, \dots, x_\ell)$, called a *feedback function*. As the current state x_k is known, ψ_k is a constant. We call a tuple of feedback functions $\psi_{k:k_1} := (\psi_k, \psi_{k+1}, \dots, \psi_{k_1})$ a *control policy*. We use $\Psi_{k:k_1}$ to denote the set of all control policies $\psi_{k:k_1}$.

LQ cost functional We measure the performance of a control policy $\psi_{k:k_1}$ by means of the associated cost. For all $k \in \{k_0, \dots, k_1\}$, all $\psi_{k:k_1} \in \Psi_{k:k_1}$ and all $x_k \in \mathbb{R}$ we define the linear-quadratic (LQ) cost functional η as

$$\eta\left[\psi_{k:k_1} \middle| x_k\right] := \sum_{\ell=k}^{k_1} r_\ell \psi_\ell(X_{k+1:\ell})^2 + q_{\ell+1} X_{\ell+1}^2,$$

where $q_\ell \geq 0$ and $r_\ell > 0$ are real coefficients.

Precise noise model P

In order to model the noise $W_{k:k_1} := (W_k, W_{k+1}, \dots, W_{k_1})$, we consider an initial time k_0, let $k_0 \leq k \leq k_1$, and focus on modelling $W_{k_0:k_1}$.

Precise noise model We model our beliefs about $W_{k_0:k_1}$ using conditional probability density functions: for all $k \in \{k_0, \dots, k_1\}$ and all $w_{k_0:k-1} \in \mathbb{R}^{k-k_0}$, we are given a conditional probability density function $f_k(\cdot | w_{k_0:k-1})$, and we use $\mathrm{P}_k(\cdot | w_{k_0:k-1})$ to denote the corresponding *conditional linear prevision* operator (expectation operator). It then follows from the *law of iterated expectation* that for any gamble g on $\mathbb{R}^{k_1-k+1}$:

$$\mathrm{P}_{k:k_1}(g | w_{k_0:k-1}) = \mathrm{P}_k(\mathrm{P}_{k+1}(\cdots \mathrm{P}_{k_1}(g | w_{k_0:k-1}, W_{k:k_1-1}) \cdots | w_{k_0:k-1}, W_k) | w_{k_0:k-1}).$$

We assume that our conditional probability density functions are sufficiently well-behaved in order for the previsions in this expression to exist. We denote the set of all such precise noise models P by $\mathbb{P}$.

White noise model In the literature, it is often assumed that the noise is *independent*. This means that all the conditional probability density functions (and associated linear previsions) are equal to marginal ones.

Imprecise noise model $\mathscr{P}$

Imprecise noise model Our beliefs about $W_{k_0:k_1}$ are modelled by a set $\mathscr{P} \subseteq \mathbb{P}$ of precise noise models. This definition allows us to use the results obtained in the precise LQ problem.

Forward irrelevant noise model $\mathscr{P}$ is said to be a forward irrelevant product if there are sets of marginal probability density functions $\mathscr{Q}_k$, $k \in \{k_0, \dots, k_1\}$, such that $\mathscr{P}$ is the largest subset of $\mathbb{P}$ for which it holds that

$$f_k(\cdot | w_{k_0:k-1}) \in \mathscr{Q}_k$$

for all precise models P in $\mathscr{P}$, all k in $\{k_0, \dots, k_1\}$ and all $w_{k_0:k-1}$ in $\mathbb{R}^{k-k_0}$.

The precise LQ problem

Local optimality A control policy $\hat{\psi}_{k:k_1}$ is *locally optimal* for $x_k \in \mathbb{R}$ and $w_{k_0:k-1} \in \mathbb{R}^{k-k_0}$ if

$$\hat{\psi}_{k:k_1} \in \text{loc-opt}^{\mathrm{P}}_{k:k_1}\left(\Psi_{k:k_1} \middle| x_k, w_{k_0:k-1}\right) := \underset{\psi_{k:k_1} \in \Psi_{k:k_1}}{\arg\min}\ \mathrm{P}_{k:k_1}\left(\eta\left[\psi_{k:k_1} \middle| x_k\right] \middle| w_{k_0:k-1}\right).$$

Optimality A control policy $\hat{\psi}_{k:k_1}$ is *optimal* for $x_k \in \mathbb{R}$ and $w_{k_0:k-1} \in \mathbb{R}^{k-k_0}$ if, for all $\ell \in \{k, \dots, k_1\}$ and all $x_{k+1:\ell} \in \mathbb{R}^{\ell-k}$:

$$\hat{\psi}_{\ell:k_1}(x_{k+1:\ell}, \cdot) \in \text{loc-opt}^{\mathrm{P}}_{\ell:k_1}\left(\Psi_{\ell:k_1} \middle| x_\ell, w_{k_0:\ell-1}\right),$$

where $w_{k:\ell-1}$ is derived from (DYN) and $x_{k:\ell}$. The set of all such optimal control policies is denoted by $\text{opt}^{\mathrm{P}}_{k:k_1}\left(\Psi_{k:k_1} \middle| x_k, w_{k_0:k-1}\right)$.

Precise noise solution For any current state $x_k \in \mathbb{R}$ and noise history $w_{k_0:k-1} \in \mathbb{R}^{k-k_0}$, the set $\text{opt}^{\mathrm{P}}_{k:k_1}\left(\Psi_{k:k_1} \middle| x_k, w_{k_0:k-1}\right)$ consists of a *single* optimal control policy. For any $\ell \in \{k, \dots, k_1\}$ and $x_{k+1:\ell} \in \mathbb{R}^{\ell-k}$, it is given by

$$\hat{\psi}_\ell(x_{k+1:\ell}) = -\tilde{r}_\ell b_\ell \left(m_{\ell+1} a_\ell x_\ell + h_{\ell | w_{k_0:\ell-1}}\right). \qquad \text{(OCP)}$$

The parameters $m_{\ell+1}$ and $\tilde{r}_\ell$ are obtained from the initial condition $m_{k_1+1} := q_{k_1+1}$ and the recursive *Riccati* equation $m_\ell := q_\ell + a_\ell^2 m_{\ell+1} - \tilde{r}_\ell a_\ell^2 b_\ell^2 m_{\ell+1}^2$, with $\tilde{r}_\ell := \left(r_\ell + b_\ell^2 m_{\ell+1}\right)^{-1}$. The *noise feedforward* $h_{\ell | w_{k_0:\ell-1}}$ is obtained from the initial condition $h_{k_1+1 | w_{k_0:k_1}} := 0$ and the recursive expression

$$h_{\ell | w_{k_0:\ell-1}} := \mathrm{P}_\ell(m_{\ell+1} W_\ell + \tilde{r}_{\ell+1} a_{\ell+1} r_{\ell+1} h_{\ell+1 | w_{k_0:\ell-1}, W_\ell} | w_{k_0:\ell-1}).$$

Calculating this feedforward is intractable!

White noise solution For white noise, the recursive feedforward relation simplifies to

$$h_\ell := m_{\ell+1} \mathrm{P}_\ell(W_\ell) + \tilde{r}_{\ell+1} a_{\ell+1} r_{\ell+1} h_{\ell+1},$$

with initial condition $h_{k_1+1} := 0$.

The imprecise LQ problem

E-admissibility A control policy $\hat{\psi}_{k:k_1}$ is *E-admissible* for $x_k \in \mathbb{R}$ and $w_{k_0:k-1} \in \mathbb{R}^{k-k_0}$ if

$$\hat{\psi}_{k:k_1} \in \text{opt}^{\mathscr{P}}_{k:k_1}\left(\Psi_{k:k_1} \middle| x_k, w_{k_0:k-1}\right) := \bigcup_{\mathrm{P} \in \mathscr{P}} \text{opt}^{\mathrm{P}}_{k:k_1}\left(\Psi_{k:k_1} \middle| x_k, w_{k_0:k-1}\right).$$

Imprecise noise solution Every $\mathrm{P} \in \mathscr{P}$ corresponds to a single E-admissible control policy $\hat{\psi}_{k:k_1}$—see Equation (OCP)—that is a combination of the same state feedback and a possibly different noise feedforward. *Calculating all possible feedforwards is intractable!*

Forward irrelevant noise solution If $\mathscr{P}$ is a forward irrelevant product, then for all $\ell \in \{k_0, \dots, k_1\}$ and all $w_{k_0:\ell-1} \in \mathbb{R}^{\ell-k_0}$

$$h_{\ell | w_{k_0:\ell-1}} \in [\underline{h}_\ell, \overline{h}_\ell],$$

where $\underline{h}_{k_1+1} := 0$, $\overline{h}_{k_1+1} := 0$ and, for $a_{\ell+1} \geq 0$:

$$\underline{h}_\ell := m_{\ell+1} \underline{\mathrm{P}}_\ell(W_\ell) + \tilde{r}_{\ell+1} a_{\ell+1} r_{\ell+1} \underline{h}_{\ell+1} \text{ and } \overline{h}_\ell := m_{\ell+1} \overline{\mathrm{P}}_\ell(W_\ell) + \tilde{r}_{\ell+1} a_{\ell+1} r_{\ell+1} \overline{h}_{\ell+1},$$

with $\underline{\mathrm{P}}_\ell(W_\ell)$ and $\overline{\mathrm{P}}_\ell(W_\ell)$ the lower and upper prevision (expectation) of W_ℓ, respectively. For $a_{\ell+1} \leq 0$, $\underline{h}_{\ell+1}$ and $\overline{h}_{\ell+1}$ switch places.

Convergence For stationary linear systems (constant a_ℓ, b_ℓ, r_ℓ, q_ℓ and $\mathscr{Q}_\ell$) and large $k_1 - k$, the parameters m_k, $\tilde{r}_k$, $\underline{h}_k$ and $\overline{h}_k$ converge to easily calculable limit values.

Simulations

How do we choose which element of $[\underline{h}_\ell, \overline{h}_\ell]$ to apply? We propose two possible options:

1. use the control policy that corresponds to a white noise model
2. lazily choose the $h_\ell \in [\underline{h}_\ell, \overline{h}_\ell]$ that minimises $|u_\ell|$.

We ran two simulations to compare their performance

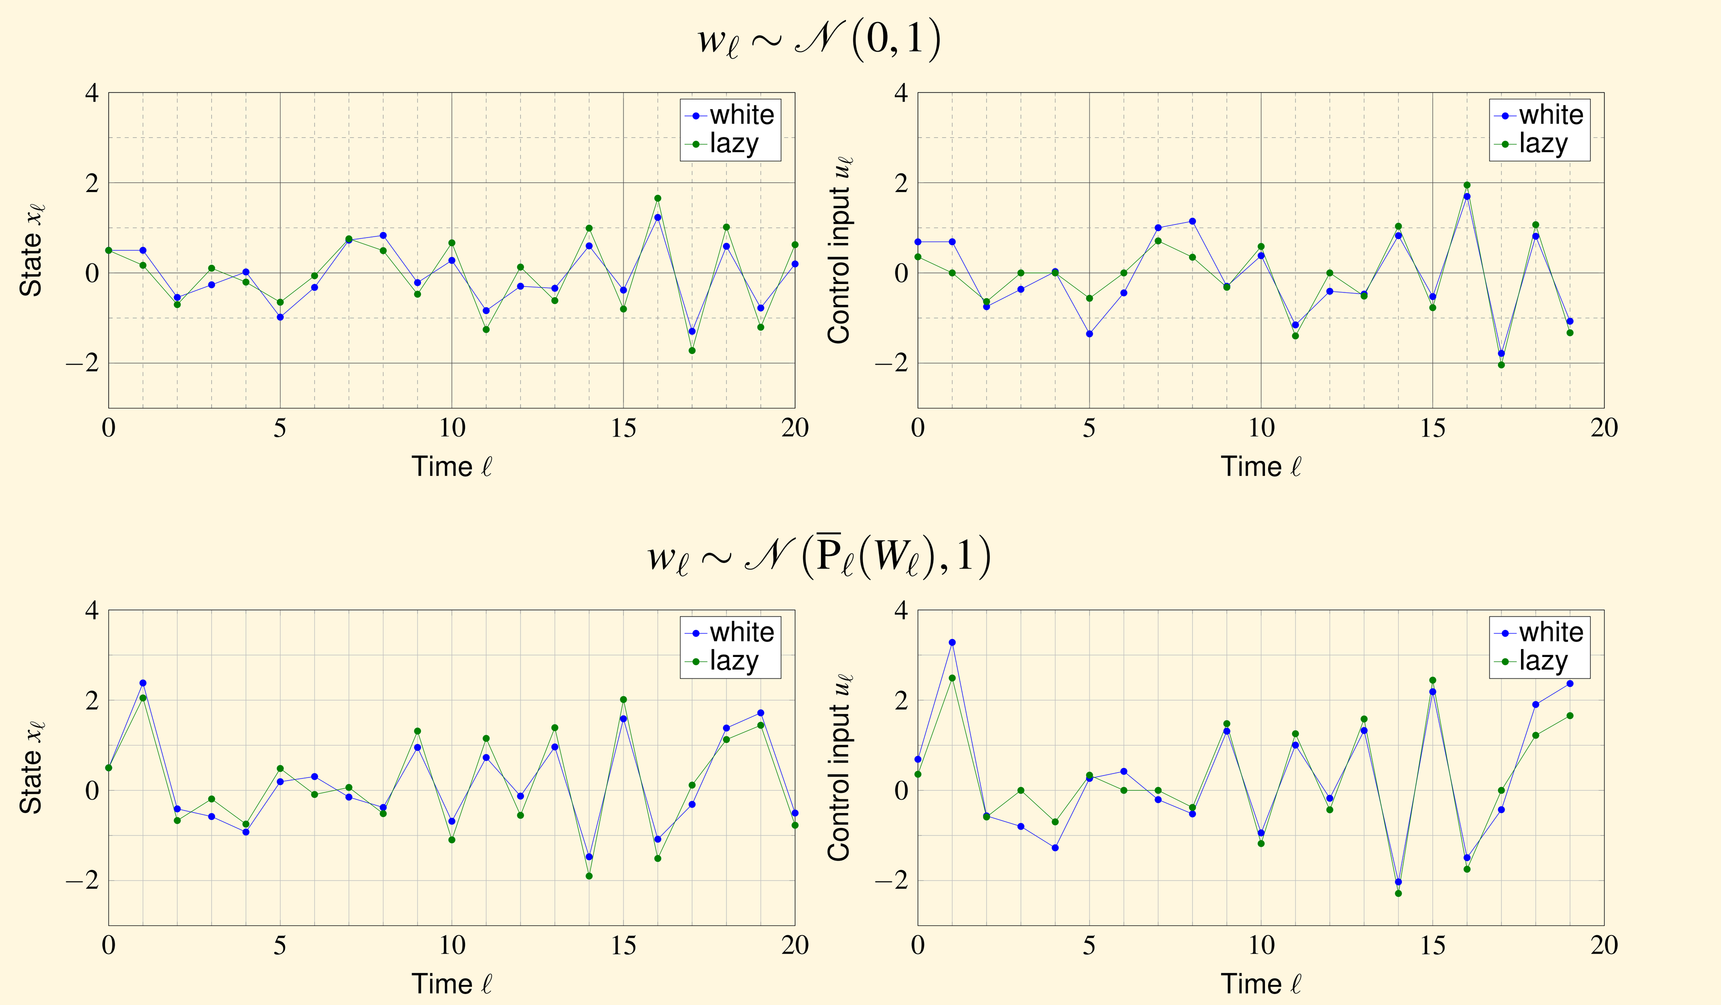

Small difference in cost, but the lazy control has more zero inputs
→ more research is definitely necessary