Geometric Insights into Scenario-Based Stochastic MPC

Maria M. Seron

(Joint work with Graham C. Goodwin)

Journée GT – CPNL (Commande Prédictive Non Linéaire)
CentraleSupélec, 13 June 2019

Maria M. Seron
Centre for Complex Dynamic Systems and Control
School of Electrical Engineering and Computing
University of Newcastle, Australia

Outline

INTRODUCTION

Introduction: Why MPC

Nice Features

- Widely acknowledged as one of the most important developments in systems and control in the second half of the 20^{th} century.
- Many books have been written on the topic and MPC has been used in thousands of industrial applications.
- Key advantages include the ability to deal with multivariable systems and hard constraints in a systematic manner.

Main Principle

The idea of MPC is to generate a stabilising feedback controller, $\mathcal{K}_N(x)$, by repeatedly solving a finite horizon open loop optimal control problem and implementing it in a receding horizon manner.

Introduction: MPC as a Feedback Controller in RHC Form

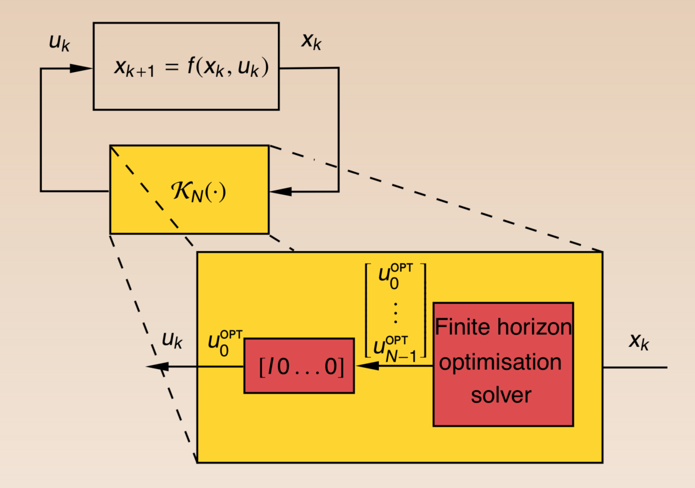

Figure: Receding horizon control

Introduction: MPC and Future Disturbances

Prediction

- Central to the concept of MPC is the idea of predicting the system response over a finite future horizon.
- The way in which MPC deals with the uncertainty associated with future disturbances is of major importance and has led to significant research effort.
- Two general approaches that have been proposed in contemporary literature are Robust MPC and Stochastic MPC.
- Although they share many common attributes, the core difference between Robust and Stochastic MPC lies in how the disturbances are described.

Introduction: Robust and Stochastic MPC

Robust MPC (e.g., min-max MPC, tube-based MPC)

- Disturbances take values in a compact set.
- Every possible disturbance has to be accounted for and has equal importance.
- Guaranteed to always be safe $\longrightarrow$ conservative.

Stochastic MPC

- Disturbances are considered as random processes, not necessarily bounded.
- A 'value function' is assigned to each possible disturbance realisation.
- The value function is described mathematically by embedding the disturbance in a probability space.
- Guaranteed to be safe 'in probability' $\longrightarrow$ less conservative.

Introduction: Stochastic MPC

'Curse' of Dimensionality

- Stochastic MPC implementations are significantly more complex than others since it optimises over *control policies* (POMPC) rather than *control sequences* (SOMPC).
- If one assumes the disturbance can take W different values over a horizon of length N, then a full solution to the associated policy optimisation problem leads to $\sum_{k=0}^{N-1} W^k$ decision variables.
- Thus, it is important to examine under what circumstances the added complexity associated with Stochastic MPC based on policy optimisation brings noteworthy performance gains.

Introduction: Stochastic MPC, Explicit View

Equivalence between different forms of MPC

- We will see that the control laws associated with the different algorithms (POMC, SOMPC, traditional MPC) can indeed be equivalent under certain circumstances.
- In cases where the solutions are different, we can obtain ‘geometric’ insights into the performance gap.
- The tool to use is the geometric features of stochastic MPC, revealed by analysing explicit forms of the solution.
- Explicit MPC, instead of repeatedly finding the online solution of an optimisation problem, solves the optimisation problem offline, obtaining an explicit function that maps the state (and reference/disturbances) to the optimal control input.

MOTIVATION

Motivational Application: Diabetes Management

The problem

- Regulation of blood glucose levels (BGL) by administering insulin in individuals having Type I diabetes.
- Constraints: one-sided control (insulin can be added but not removed), hard state constraints (BGL cannot be lower than a certain level that results in hypoglycaemia, putting the person's life at risk), soft state constraints (it is desirable to keep BGL below a certain level most of the time).
- Uncertainty: in both model parameters and the nature of the disturbance inputs (food consumption, exercise and stress).
- Although future disturbances cannot be known in advance, it is possible to *consider a number of disturbance scenarios with different probabilities* according to the individual's lifestyle.

BGL 4 day-response under Stochastic MPC (blue solid line) ‘traditional’ MPC (red dashed line), with ‘disturbances’: dinner at 7pm, exercise at 8pm, breakfast at 7am and lunch at 12 pm.

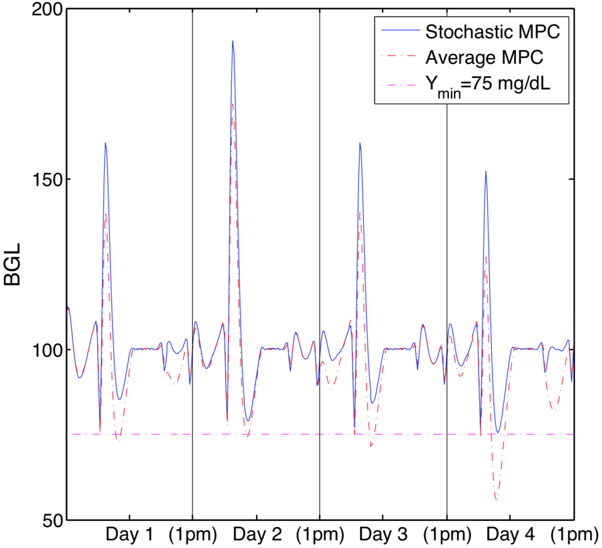

Probability distribution of BGL for each algorithm. Striking result: POMPC has only a slightly increased probability that BGL lies near the target value of 100 mg/dL !!

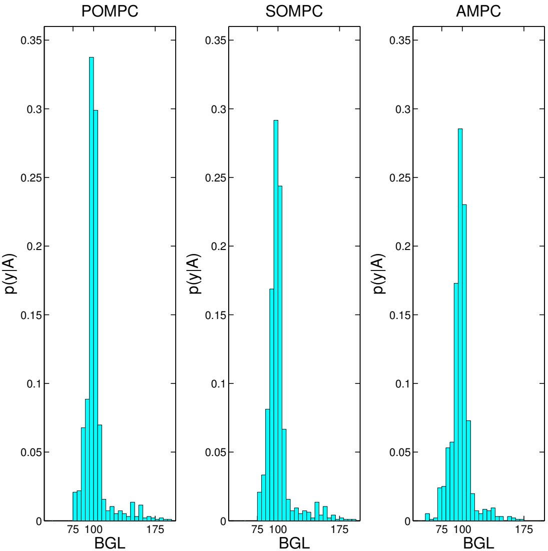

PROBLEM FORMULATION

Overview of (State-Feedback) Stochastic MPC

We consider a linear system of the form:

$$x_{k+1} = Ax_k + Bu_k + Ew_k \qquad (1)$$

where x_k is the system state, u_k is the control input and $\{w_k\}$ is a random disturbance sequence whose past values are known but future values have a probabilistic description.

Core idea

Find a functional mapping from state $\{x_k\}$ to input u_k which (a) respects constraints and (b) optimises a performance goal, J_N, assumed to take the following form: Given x_0,

$$J_N = \mathbb{E}[V_N | x_0], \qquad V_N \doteq \sum_{k=0}^{N-1} \ell(x_k, u_k) + Q(x_N), \qquad (2)$$

where $\ell(x_k, u_k)$ is the stage cost, $Q(x_N)$ a terminal cost; $\mathbb{E}$ denotes expectation (conditional on x_0 and over future disturbances).

Disturbance Quantisation and Scenarios

Complexity

The above is a function optimisation problem. Such problems are computationally intractable. Hence approximations are required.

- Restrict the disturbances to a finite set $\mathcal{W} \doteq \{\bar{w}_1, \ldots, \bar{w}_W\}$, with associated probabilities $\bar{p}_1, \ldots, \bar{p}_W$.
- Draw S scenarios s_i for the future values of process $\{w_k\}$:

$$\{w_k(s_i), \quad k = 0, \ldots, N-1, \quad i = 1, \ldots, S\}$$

  (There exists recent theory (R. Tempo et al) which allows one to bound the number of scenarios needed to achieve a certain level of probabilistic performance.)

- For each scenario sequence, pose a different input sequence:

$$\{u_k(s_i), \quad k = 0, \ldots, N-1, \quad i = 1, \ldots, S\}$$

Policy Constraints: What the Predicted Future Control Sequence is Allowed to Depend Upon

We refine the notation to represent a predicted control sequence associated with a disturbance scenario $w_0(s_i), \ldots, w_{N-1}(s_i)$:

$$\bar{u}^{\mathrm{P}}(k, s_i) = u^{\mathrm{P}}[k, w_0(s_i), \ldots, w_{N-1}(s_i)] \tag{3}$$

The associated state trajectory is described by

$$x_{k+1}(s_i) = Ax_k(s_i) + B\bar{u}^{\mathrm{P}}(k, s_i) + Ew_k(s_i); \quad x_0(s_i) = x \tag{4}$$

To complete the formulation we resolve the expectation in the cost function (2) as follows:

$$J_N = \sum_{k=0}^{N-1} \sum_{i=1}^{S} \left[\ell(x_k(s_i), \bar{u}^{\mathrm{P}}(k, s_i), w_k(s_i)) + Q(x_N(s_i))\right] p_i, \tag{5}$$

where $\{w_k(s_i)\}$ takes the values $w_0(s_i), \ldots, w_{N-1}(s_i)$ with associated probability p_i.

POMPC: Policy-Optimisation MPC

Recall:

$$\bar{u}^{\mathrm{P}}(k, s_i) = u^{\mathrm{P}}[k, w_0(s_i), \ldots, w_{N-1}(s_i)]$$

- In POMPC, the predicted input at time k is constructed to be a function of the (predicted) state at time k.
- Equivalently, the predicted input at time k is a function of x_0 and the disturbances that will have already appeared before time k, namely w_0, ..., w_{k-1}.
- Mathematically: for $k = 0, \ldots, N-1$,

$$\bar{u}^{\mathrm{P}}(k, s_i) = u^{\mathrm{POMPC}}[k, w_0(s_i), \ldots, w_{k-1}(s_i), \bullet, \ldots, \bullet]$$

The '$\bullet$' denotes that the function takes the same value for all possible arguments in that location.

SOMPC: Sequence-Optimisation MPC

Recall:

$$\bar{u}^{\mathrm{P}}(k, s_i) = u^{\mathrm{P}}[k, w_0(s_i), \ldots, w_{N-1}(s_i)]$$

- In SOMPC one optimises a single (predicted) open loop input sequence for all scenarios.
- Mathematically, for $k = 0, \ldots, N-1$:

$$\bar{u}^{\mathrm{P}}(k, s_i) = u^{\mathrm{SOMPC}}[k, \bullet, \ldots, \bullet]$$

  As before, '$\bullet$' denotes that the function takes the same value for all possible arguments in that location.
- Note that the problem formulation remains stochastic since the cost function is the expected value over all possible disturbance scenarios.

Special Case: Traditional (Average) MPC

Recall:

$$\bar{u}^{\mathrm{P}}(k, s_i) = u^{\mathrm{P}}[k, w_0(s_i), \ldots, w_{N-1}(s_i)]$$

- This is a variant of SOMPC where only one disturbance scenario is considered, namely the 'average' disturbance.
- In this case the future disturbances are replaced by their average value

$$w_k(s_i) = \bar{w} \doteq \sum_{j=1}^{W} p_j \bar{w}_j, \quad k \geq 0, \ \forall i = 1, \ldots, S.$$

- The resulting problem is purely deterministic.

Other Constraints

- Input (hard) constraints:

$$u_k \in \mathcal{U} \qquad \text{for all } k \geq 0$$

- State (hard) constraints:

$$x_k \in \mathcal{X} \qquad \text{for all } k \geq 0$$

- Probabilistic (or chance) constraints:

$$P(x_k \in \mathcal{X}) \geq p$$

  with $p \in (0, 1)$.

MAIN RESULTS

SOMPC and Traditional (Average) MPC are Equivalent for Linear Systems with Input Constraints and Quadratic Cost

Consider the previously defined linear system. The disturbance $w_k \in \mathbb{R}^q$ is a random variable with arbitrary distribution, and mean

$$\mathbb{E}[w_k] = \bar{w}_k.$$

The cost function is specialised to the quadratic form

$$J_N = \mathbb{E}[V_N | x_0], \qquad V_N = x_N^T Q_N x_N + \sum_{k=0}^{N-1} (x_k^T Q x_k + u_k^T R u_k),$$

for some matrices Q, R, Q_N of compatible dimensions.

Theorem

For a linear system and quadratic cost function, without state or output constraints, SOMPC is equivalent to MPC using the average disturbance value, that is, both optimisation problems yield identical optimal control and state trajectories.

Scalar Case: Equivalence Between POMPC, SOMPC and Average MPC

Consider the special (scalar) case:

$$x_{k+1} = ax_k + bu_k + ew_k, \qquad J_N = \frac{1}{2}\mathbb{E}\left[\sum_{k=1}^{N} x_k^2 \middle| x_0\right].$$

The constraints are

$$x_0 \text{ given}, \quad |u_k| \le \Delta, \quad \Delta > 0, \quad w_k \in \mathcal{W}, \quad \forall k = 0, 1, \ldots$$

$\mathcal{W}$ contains a finite number of elements with probabilities in $\mathcal{P}$.
Both $\mathcal{W}$ and $\mathcal{P}$ have symmetric properties with respect to zero:

$$\mathcal{W} = \{\bar{w}_1, \ldots, \bar{w}_W\} = \{-\sigma_1, \ldots, -\sigma_{\lfloor W/2 \rfloor}, \sigma_0, \sigma_{\lfloor W/2 \rfloor}, \ldots, \sigma_1\}$$
$$\mathcal{P} = \{\bar{p}_1, \ldots, \bar{p}_W\} = \{\rho_1, \ldots, \rho_{\lfloor W/2 \rfloor}, \rho_0, \rho_{\lfloor W/2 \rfloor}, \ldots, \rho_1\}$$

Theorem

Consider the scalar system with the above formulation. Then POMPC yields the same control input as Traditional (average) MPC and SOMPC.

More General Cases: Geometric Insights

Revision: In the deterministic case, for linear systems, quadratic cost and linear constraints, the problem to solve in MPC is a quadratic programme (QP):

$$\mathbf{u}^{\mathrm{OPT}}(x) = \operatorname*{argmin}_{\mathbf{u}\in R_{\mathrm{UC}}} \frac{1}{2}\mathbf{u}^T H\mathbf{u} + \mathbf{u}^T Fx$$

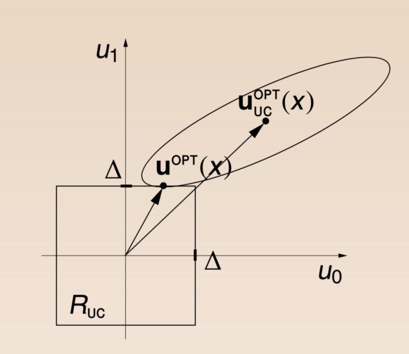

$\frac{1}{2}\mathbf{u}^T H\mathbf{u} + \mathbf{u}^T Fx =$ constant defines *ellipsoids* centred at the unconstrained optimum $\mathbf{u}_{\mathrm{UC}}^{\mathrm{OPT}}(x) = -H^{-1}Fx$.

Solving the QP amounts to finding the *smallest* ellipsoid that intersects the boundary of R_{UC}, and $\mathbf{u}^{\mathrm{OPT}}(x)$ is the *point of intersection*.

Explicit Solution as Minimum Distance Problem

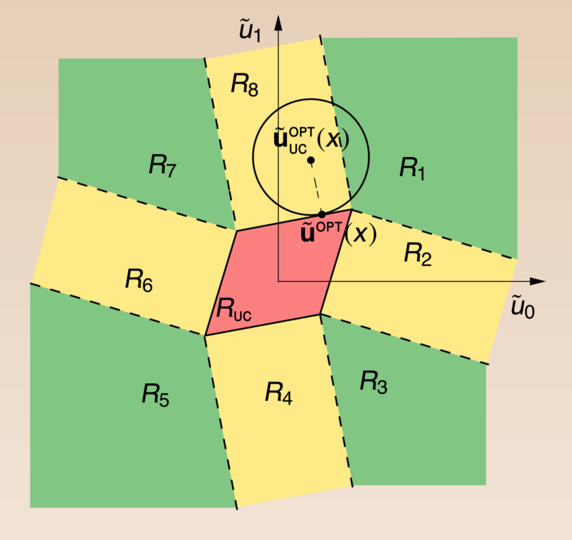

The solution of the QP is obtained by partitioning $\mathbb{R}^2$ into nine regions.

The first region is the polytope R_{UC}.

Regions R_1 to R_8 are delimited by lines that are normal to the faces of R_{UC} and pass through its vertices.

Explicit Solution as Minimum Distance Problem

The *optimal constrained* solution $\tilde{\mathbf{u}}^{\mathrm{OPT}}(x)$ is determined by the region in which the *optimal unconstrained* solution $\tilde{\mathbf{u}}^{\mathrm{OPT}}_{\mathrm{UC}}(x)$ lies.

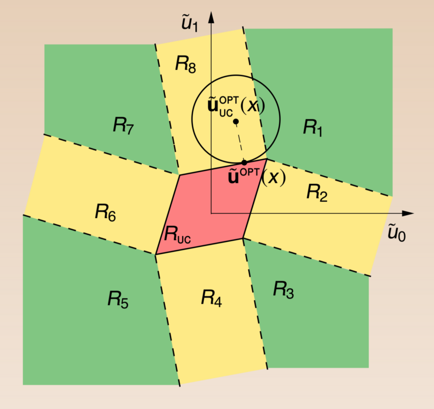

- The solution in R_{UC} is $\tilde{\mathbf{u}}^{\mathrm{OPT}}(x) = \tilde{\mathbf{u}}^{\mathrm{OPT}}_{\mathrm{UC}}(x)$;
- The solution in regions R_1, R_3, R_5 and R_7 is simply equal to the vertex that is contained in the region.
- The solution in regions R_2, R_4, R_6 and R_8 is defined by the orthogonal projection of $\tilde{\mathbf{u}}^{\mathrm{OPT}}_{\mathrm{UC}}(x)$ onto the faces of R_{UC}.

As a result, we obtain a characterisation of the QP solution as

$$\mathbf{u}^{\mathrm{OPT}}(x) = H^{-1/2}\tilde{\mathbf{u}}^{\mathrm{OPT}}(x) \quad \text{if} \quad \tilde{\mathbf{u}}^{\mathrm{OPT}}_{\mathrm{UC}}(x) = -H^{-1/2}Fx \in R_i,$$

where $\tilde{\mathbf{u}}^{\mathrm{OPT}}(x)$ is an affine function.

Geometric Insights in the Stochastic Case

The idea can be extended to the stochastic case if the future disturbances are described in a tree structure.

Letting $w_k^i = \bar{w}_i \in \mathcal{W}$, for $k = 0, \ldots, N-1$, $i \in \{1, \ldots, W\}$, the disturbance sequences and associated probabilities are:

$$\{\underbrace{w_0^1}_{\bar{w}_1}, \underbrace{w_1^1}_{\bar{w}_1} \ldots, \underbrace{w_{N-3}^1}_{\bar{w}_1}, \underbrace{w_{N-2}^1}_{\bar{w}_1}\}, \quad \text{with probability } p_1 \doteq \bar{p}_1^{N-2}\bar{p}_1$$

$$\vdots$$

$$\{\underbrace{w_0^1}_{\bar{w}_1}, \underbrace{w_1^1}_{\bar{w}_1} \ldots, \underbrace{w_{N-3}^1}_{\bar{w}_1}, \underbrace{w_{N-2}^W}_{\bar{w}_W}\}, \quad \text{with probability } p_W \doteq \bar{p}_1^{N-2}\bar{p}_W$$

$$\vdots$$

$$\{w_0^W, w_1^W \ldots, w_{N-3}^W, w_{N-2}^1\}, \quad \text{with probability } p_{W^{N-1}-W+1} \doteq \bar{p}_W^{N-2}\bar{p}_1$$

$$\vdots$$

$$\{w_0^W, w_1^W \ldots, w_{N-3}^W, w_{N-2}^W\}, \quad \text{with probability } p_{W^{N-1}} \doteq \bar{p}_W^{N-2}\bar{p}_W$$

Propagation of control inputs

- The control input in POMPC depends on previous values of disturbances. Then it has a tree-like structure according to the above disturbance patterns.
- There is:
  - A single value u_0 at time 0;
  - W values for the input u_1 at time 1, denoted by $u_1^1, \ldots, u_1^W$, which correspond to $w_0^1, \ldots, w_0^W$.
  - W^2 values for the input u_2 at time 2: u_2^1, u_2^2, ..., $u_2^{W^2}$, etc.
  - The final control move, u_{N-1} has W^{N-1} possible values, denoted by $u_{N-1}^1, \ldots, u_{N-1}^{W^{N-1}}$.
- This enumeration gives a total of $\sum_{k=0}^{N-1} W^k$ possibilities over the prediction horizon.
- Only u_0 is applied to the plant in an RHC fashion.

Propagation of States and Outputs

The state x_k and output $y_k = Cx_k$ can be propagated in sequences corresponding to the disturbance and control sequences:

$$
\begin{aligned}
y_1 &= Cx_1 = C(Ax_0 + Bu_0 + Ew_0),\\
y_2^1 &\doteq Cx_2^1 \doteq C\big[A(Ax_0 + Bu_0 + Ew_0^1) + Bu_1^1 + Ew_1\big]\\
&\vdots\\
y_2^W &\doteq Cx_2^W \doteq C\big[A(Ax_0 + Bu_0 + Ew_0^W) + Bu_1^W + Ew_1\big]\\
y_3^1 &\doteq Cx_3^1 \doteq C\big[A^2(Ax_0 + Bu_0 + Ew_0^1) + ABu_1^1 + AEw_1^1 + Bu_2^1 + Ew_2\big]\\
&\vdots\\
y_3^{W^2} &\doteq Cx_3^{W^2} \doteq C\big[A^2(Ax_0 + Bu_0 + Ew_0^W) + ABu_1^W + AEw_1^W +\\
&\qquad Bu_2^{W^2} + Ew_2\big]\\
&\vdots
\end{aligned}
$$

Vector Notation for Inputs, Outputs and Disturbances

Stage vectors:

$$\mathbf{u}_0 \doteq u_0, \quad \mathbf{u}_{k-1} \doteq \text{col}\,(u_{k-1}^1, \ldots, u_{k-1}^{W^{k-1}}) \in \mathbb{R}^{W^{k-1}}, \quad k = 2, \ldots, N$$
$$\mathbf{y}_1 \doteq y_1, \quad \mathbf{y}_k \doteq \text{col}\,(y_k^1, \ldots, y_k^{W^{k-1}}) \in \mathbb{R}^{W^{k-1}}, \quad k = 2, \ldots, N,$$

Total vectors:

$$\mathbf{U}_N \doteq \text{col}\,(\mathbf{u}_0, \mathbf{u}_1 \ldots, \mathbf{u}_{N-1}) \in \mathbb{R}^{d_N}$$
$$\mathbf{Y}_N \doteq \text{col}\,(\mathbf{y}_1, \mathbf{y}_2 \ldots, \mathbf{y}_N) \in \mathbb{R}^{d_N}$$

where $d_N \doteq \sum_{j=0}^{N-1} W^j$
Disturbance value vector:

$$\mathbf{s} \doteq \text{col}\,(\sigma_1, \ldots, \sigma_{\lfloor W/2 \rfloor})$$

Cost Function Evaluation

Consider for simplicity a scalar output y_k and the 'cheap' control cost function

$$J_N = \frac{1}{2}\mathbb{E}\left[\sum_{k=1}^{N} y_k^2 \middle| x_0\right]$$

The above cost can be expressed as

$$J_N = \frac{1}{2}\mathbf{Y}_N^T \mathbf{P}_N \mathbf{Y}_N,$$

with

$$\mathbf{P}_N \doteq \text{blkdiag}\ (\pi_1^N, \ldots, \pi_N^N),$$

where each diagonal matrix $\pi_k^N = \mathrm{diag}\big(\otimes_{k-1} \vec{P}\big) \in \mathbb{R}^{W^{k-1} \times W^{k-1}}$ corresponds to each stage output vector $\mathbf{y}_k$.
The vector $\vec{P}$ contains the probabilities

$$\vec{P} \doteq \begin{bmatrix} \bar{p}_1 & \ldots & \bar{p}_W \end{bmatrix}^T = \begin{bmatrix} \rho_1 & \ldots & \rho_{\lfloor W/2 \rfloor} & \rho_0 & \rho_{\lfloor W/2 \rfloor} & \ldots & \rho_1 \end{bmatrix}^T.$$

Formulation of a QP

We write the QP problem

$$\min_{\mathbf{U}_N} \frac{1}{2} \underbrace{\mathbf{U}_N^T \Gamma_N^T \mathbf{P}_N \Gamma_N \mathbf{U}_N}_{\tilde{\mathbf{U}}_N^T \tilde{\mathbf{U}}_N} + \underbrace{\mathbf{U}_N^T \Gamma_N^T \mathbf{P}_N [\Lambda_N \;\; \Omega_N]}_{\tilde{\mathbf{U}}_N^T F_N} \begin{bmatrix} x_0 \\ \mathbf{s} \end{bmatrix}.$$

Including input constraints $|u_k| \leq \Delta$, the associated QP in the new coordinates takes the form

$$\min_{\tilde{\mathbf{U}}_N} \frac{1}{2} \tilde{\mathbf{U}}_N^T \tilde{\mathbf{U}}_N + \tilde{\mathbf{U}}_N^T F_N \begin{bmatrix} x_0 \\ \mathbf{s} \end{bmatrix}$$

$$\text{subject to} : |\tilde{\Phi}_N \tilde{\mathbf{U}}_N| = |\bar{\Gamma}_N^{-1} \tilde{\mathbf{U}}_N| \leq \Delta \begin{bmatrix} 1 \\ \vdots \\ 1 \end{bmatrix}$$

The unconstrained (UC) solution is

$$\tilde{\mathbf{U}}_{N,\mathrm{UC}}^{\mathrm{OPT}} = -F_N \begin{bmatrix} x_0 \\ \mathbf{s} \end{bmatrix}$$

Explicit Solution

As for the deterministic case, the explicit solution is obtained by partitioning the space of the decision variable, $\tilde{\mathbf{U}}_N$, into regions. Each region, $\mathcal{R}_{N,\ell}$, of the partition has associated

- 'active set' ℓ, containing the indices of the elements of $\mathbf{U}_N$ that hit the constrains; and
- 'active vector' $\bar{\Delta}_{N,\ell}$, containing the value attained by each constrained element (Δ or $-\Delta$).

Then the optimal constrained solution whenever $\tilde{\mathbf{U}}_{N,\mathrm{UC}}^{\mathrm{OPT}} \in \mathcal{R}_{N,\ell}$ is given by

$$\tilde{\mathbf{U}}_N^{\mathrm{OPT}} = \underbrace{\tilde{\Phi}_{N,\ell}^T[\tilde{\Phi}_{N,\ell}\tilde{\Phi}_{N,\ell}^T]^{-1}\left[\bar{\Delta}_{N,\ell} - \tilde{\Phi}_{N,\ell}\,\tilde{\mathbf{U}}_{N,\mathrm{UC}}^{\mathrm{OPT}}\right]}_{\text{constraint correction term}} + \tilde{\mathbf{U}}_{N,\mathrm{UC}}^{\mathrm{OPT}}$$

where $\tilde{\Phi}_{N,\ell}$ is the matrix formed by selecting the rows of the constraint matrix with indices in ℓ.

Geometric Interpretation of Deterministic-Stochastic Gap

Consider $N = 2$ and the disturbance taking 2 values: σ and $-\sigma$. The predicted controls, u_0, u_1^1 and u_1^2, are associated with the two scenarios $s_1 = (u_0, u_1^1)$ and $s_2 = (u_0, u_1^2)$.

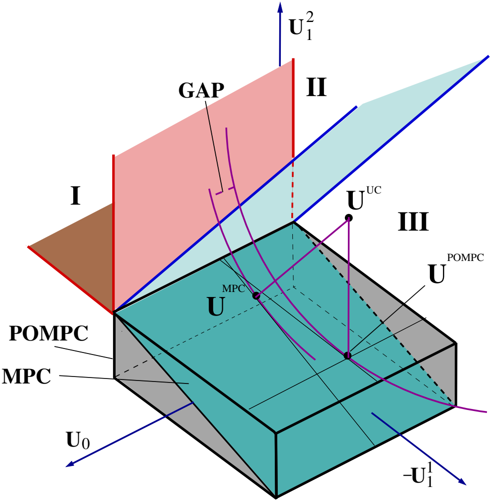

The grey ‘box’ is the constraint set for POMPC.

The green surface inside the box is the constraint set for traditional MPC. (It corresponds to adding the constraint $u_1^2 = u_1^1$.)

The cost level sets are spheres centred at the unconstrained optimum, $\mathbf{U}^{\mathrm{UC}}$.

Geometric Interpretation of Deterministic-Stochastic Gap

The *constrained* optimum correspond to the orthogonal projection of $\mathbf{U}^{\mathrm{UC}}$ on the facets of the constraint set.

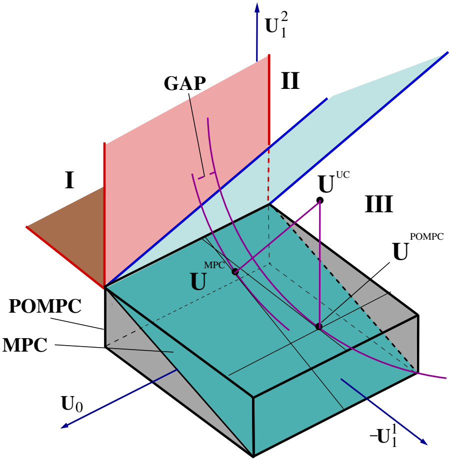

Region I projects onto a common edge between the two constraint sets, so for all initial conditions and disturbance values that yield an unconstrained optimum in region $\implies$ POMPC =MPC

Region II projects on the edge for MPC and on the top face of the box (away from the edge) for POMPC $\implies$ POMPC $\neq$MPC

Region III projects on the green plane (inside the box) for MPC and on the box's top face for POMPC $\implies$ POMPC $\neq$MPC

Geometric Interpretation of Deterministic-Stochastic Gap

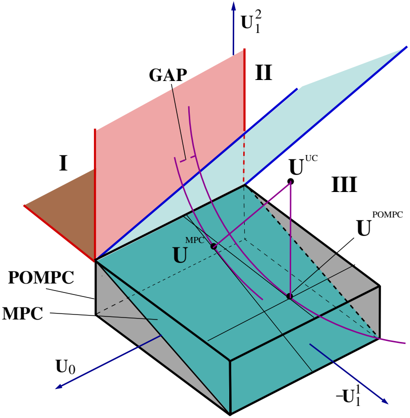

The *performance gap* is the difference in distance from the unconstrained optimum, as indicated by the lines between the cost level surfaces (only curves are shown for clarity).

Note that the gap can be interpreted as POMPC 'bringing the constraint set closer' to the unconstrained optimum.

Thus, the geometric interpretation of the explicit solution helps to directly quantify the performance gap between the two strategies.

Even for this simple case, POMPC, SOMPC and MPC can give either equal or different solutions depending on the initial conditions and disturbance values.

Numerical Example

The system matrices are taken as

$$A = \begin{bmatrix} -a_2 & a_2 \\ 0 & -a_1 \end{bmatrix}, \quad B = \begin{bmatrix} 0 \\ b \end{bmatrix}, \quad E = \begin{bmatrix} 0 \\ e \end{bmatrix}, \quad C = \begin{bmatrix} 1 & 0 \end{bmatrix}$$

The values are $a_1 = 1.4$, $a_2 = 2.2$, $b = 1.5$, $e = 1$, and control constraint $\Delta = 1.6$. The disturbance takes the values $-\sigma$ or σ with equal probability. We take $N = 3$.

To illustrate, consider the first control move in one of the regions

$$\tilde{\mathcal{U}}_0^{\mathrm{OPT}} = \frac{a_2\, b(a_1 + a_2)}{2 + (a_1 + a_2)^2}\, \Delta + \begin{bmatrix} L_1(a_1, a_2) & L_2(a_1, a_2) \end{bmatrix} x_0 + \frac{-e(a_1 + a_2)^2 a_2}{2 + (a_1 + a_2)^2}\, \sigma$$

The above expression holds for all (x_0, σ) for which the unconstrained optimal vector belongs to the considered region. Note the presence of a term depending on σ, which distinguishes this solution from Traditional (average) MPC.

Numerical Example

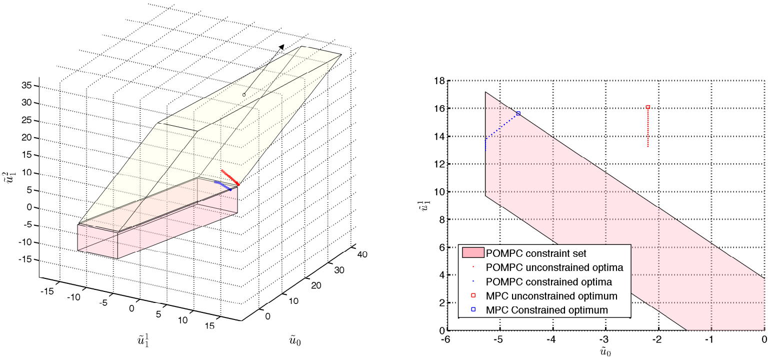

Figure: Constraint set (pink polytope centred at zero) and region corresponding to the third control saturated (yellow polytope). The unconstrained optimal solution $\tilde{\mathcal{U}}_{\mathrm{UC}}^{\mathrm{OPT}}$ (red points) and the constrained optimal solution $\tilde{\mathcal{U}}^{\mathrm{OPT}}$ (blue points) are plotted for $x_0 = -\begin{bmatrix} 6.3271 & 4.1454 \end{bmatrix}^T$ and $\sigma \in [0, 0.5]$. Right: cut on $(\tilde{u}_0, \tilde{u}_1^1)$ plane.

The constraint set for POMPC (pink) contains the constraint set for Traditional MPC (blue) in such way that the edges of the latter set are contained in the boundary of the POMPC set.

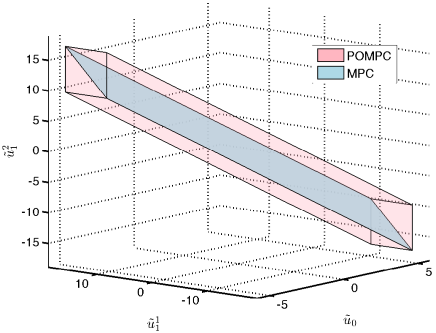

Thus, for those values of (x_0, σ) for which the POMPC solution lies on the common edges, there is no cost gap between POMPC and traditional MPC. In all other cases, there is a gap.

Concluding Remarks

- In its general form, stochastic MPC amounts to solving a functional optimisation problem. This problem is generally intractable, except for some particular cases (e.g., LQG).
- Approximations and simplifications are needed to make the problem tractable. For example: scenarios, disturbance quantisation, controller parameterisation.
- Using scenarios the problem formulation becomes deterministic. Easier to solve but usually the problem has large dimension.
- Thus, it is relevant to identify conditions under which deterministic and stochastic MPC give solutions that are close, or even identical, to avoid the burden of trying to solve a more difficult problem.
- The explicit solution to MPC and related geometric interpretation is useful to compare the different strategies and quantify the performance gap.