Experiment Design for Batch-To-Batch Model-Based Learning Control

Marco Forgione[1], Xavier Bombois[1], Paul Van den Hof[2]

[1]Delft University of Technology
Delft Center for Systems and Control

[2]Eindhoven University of Technology
Department of Electrical Engineering

American Control Conference 2013

Overview

We combine ideas from Identification for Control and the tools for Experiment Design in order to develop and actively adaptive control algorithm.

"A model-based controller is progressively improved using closed-loop system identification. Excitation is provided to the system when this is convenient."

Overview

We combine ideas from Identification for Control and the tools for Experiment Design in order to develop and actively adaptive control algorithm.

"A model-based controller is progressively improved using closed-loop system identification. Excitation is provided to the system when this is convenient."

Identification for Control

System running in closed loop, but the control performance is not optimal.

"Improve the control performance while limiting the excitation cost."

An identification experiment followed by the "normal operation"

- Control performance $\mathcal{V}$ depends on the parameter covariance P.
- The parameter covariance P depends on the excitation signal r.
- Excitation cost $\mathcal{E}$ depends on excitation signal r.

A trade-off between the excitation cost $\mathcal{E}$ and the control performance $\mathcal{V}$.

Identification for Control

System running in closed loop, but the control performance is not optimal.

"Improve the control performance while limiting the excitation cost."

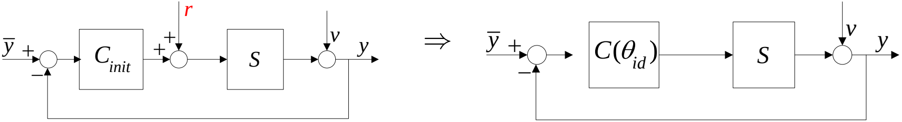

An identification experiment followed by the "normal operation"

- Control performance $\mathcal{V}$ depends on the parameter covariance P.
- The parameter covariance P depends on the excitation signal r.
- Excitation cost $\mathcal{E}$ depends on excitation signal r.

A trade-off between the excitation cost $\mathcal{E}$ and the control performance $\mathcal{V}$.

Identification for Control

System running in closed loop, but the control performance is not optimal.

"Improve the control performance while limiting the excitation cost."

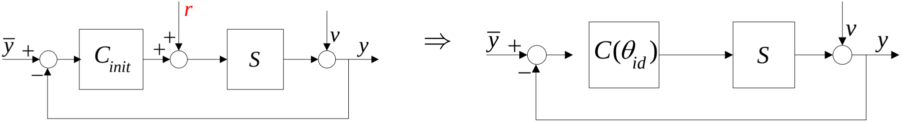

An identification experiment followed by the "normal operation"

- Control performance $\mathcal{V}$ depends on the parameter covariance P.
- The parameter covariance P depends on the excitation signal r.
- Excitation cost $\mathcal{E}$ depends on excitation signal r.

A trade-off between the excitation cost $\mathcal{E}$ and the control performance $\mathcal{V}$.

Identification for Control

System running in closed loop, but the control performance is not optimal.

"Improve the control performance while limiting the excitation cost."

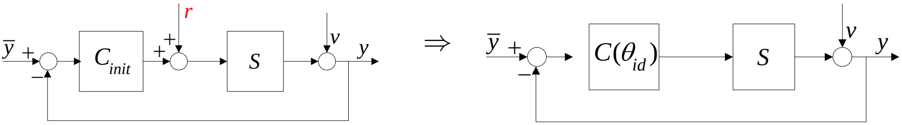

An identification experiment followed by the "normal operation"

- Control performance $\mathcal{V}$ depends on the parameter covariance P.
- The parameter covariance P depends on the excitation signal r.
- Excitation cost $\mathcal{E}$ depends on excitation signal r.

A trade-off between the excitation cost $\mathcal{E}$ and the control performance $\mathcal{V}$.

Identification for Control

System running in closed loop, but the control performance is not optimal.

"Improve the control performance while limiting the excitation cost."

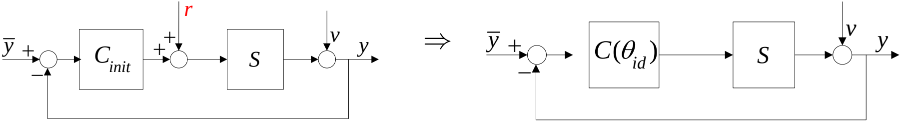

An identification experiment followed by the "normal operation"

- Control performance $\mathcal{V}$ depends on the parameter covariance P.
- The parameter covariance P depends on the excitation signal r.
- Excitation cost $\mathcal{E}$ depends on excitation signal r.

A trade-off between the excitation cost $\mathcal{E}$ and the control performance $\mathcal{V}$.

Experiment Design

For LTI systems

- The covariance P is a nonlinear, nonconvex function of the excitation signal (time domain).
- The information matrix $F = P^{-1}$ is a linear function of the excitation power spectrum (frequency domain).

Input design in the frequency domain using a two-step procedure:

1. Determine an optimal spectrum $\Phi_r(\omega)$ (convex optimization).
2. Find a signal $r(t)$ with spectrum $\Phi_r(\omega)$ (stochastic realization).

Experiment Design

For LTI systems

- The covariance P is a nonlinear, nonconvex function of the excitation signal (time domain).
- The information matrix $F = P^{-1}$ is a linear function of the excitation power spectrum (frequency domain).

Input design in the frequency domain using a two-step procedure:

1. Determine an optimal spectrum $\Phi_r(\omega)$ (convex optimization).
2. Find a signal $r(t)$ with spectrum $\Phi_r(\omega)$ (stochastic realization).

Experiment Design

For LTI systems

- The covariance P is a nonlinear, nonconvex function of the excitation signal (time domain).
- The information matrix $F = P^{-1}$ is a linear function of the excitation power spectrum (frequency domain).

Input design in the frequency domain using a two-step procedure:

1. Determine an optimal spectrum $\Phi_r(\omega)$ (convex optimization).
2. Find a signal $r(t)$ with spectrum $\Phi_r(\omega)$ (stochastic realization).

Experiment Design

For LTI systems

- The covariance P is a nonlinear, nonconvex function of the excitation signal (time domain).
- The information matrix $F = P^{-1}$ is a linear function of the excitation power spectrum (frequency domain).

Input design in the frequency domain using a two-step procedure:

1. Determine an optimal spectrum $\Phi_r(\omega)$ (convex optimization).
2. Find a signal $r(t)$ with spectrum $\Phi_r(\omega)$ (stochastic realization).

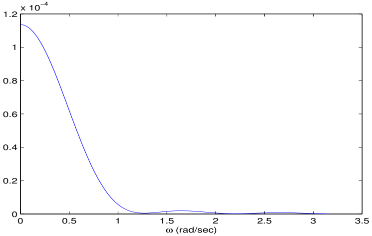

$\Longrightarrow$

Experiment Design

For LTI systems

- The covariance P is a nonlinear, nonconvex function of the excitation signal (time domain).
- The information matrix $F = P^{-1}$ is a linear function of the excitation power spectrum (frequency domain).

Input design in the frequency domain using a two-step procedure:

1. Determine an optimal spectrum $\Phi_r(\omega)$ (convex optimization).
2. Find a signal $r(t)$ with spectrum $\Phi_r(\omega)$ (stochastic realization).

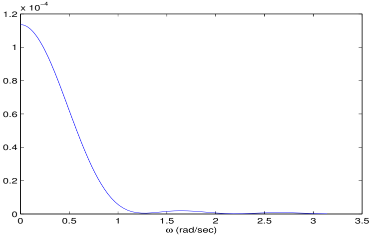

$\Rightarrow$

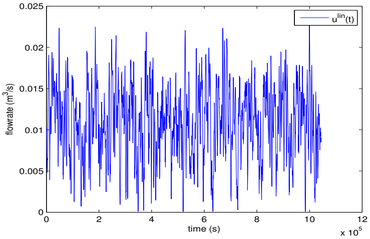

Itendification for Control

Classical: "Given a maximum allowed perturbation, find the excitation signal that gives the best control performance."

$$\max \mathcal{V} \qquad \text{such that} \qquad \mathcal{E} \leq \bar{\mathcal{E}}.$$

M. Gevers and L.Ljung.
Optimal experiment designs with respect to the intended model application.
Automatica, 22(5):543-554, 1986

Least costly: "Given a minimum allowed performance level, find the excitation signal that minimizes the perturbation."

$$\min \mathcal{E} \qquad \text{such that} \qquad \mathcal{V} \geq \bar{\mathcal{V}}.$$

X.Bombois, G.Scorletti, M.Gevers, P.M.J. Van den Hof and R.Hildebrand.
Least costly identification experiment for control.
Automatica, 42(10):1651-1662, 2006

Limitations:

- Two distinct phases: identification and normal operation.
- $\mathcal{V}$ and $\mathcal{E}$ considered separately.

"Can we design experiments in such a way that the overall performance is optimized during the whole time of operation?"

Itendification for Control

Classical: "Given a maximum allowed perturbation, find the excitation signal that gives the best control performance."

$$\max \mathcal{V} \qquad \text{such that} \qquad \mathcal{E} \leq \bar{\mathcal{E}}.$$

M. Gevers and L.Ljung.
Optimal experiment designs with respect to the intended model application.
Automatica, 22(5):543-554, 1986

Least costly: "Given a minimum allowed performance level, find the excitation signal that minimizes the perturbation."

$$\min \mathcal{E} \qquad \text{such that} \qquad \mathcal{V} \geq \bar{\mathcal{V}}.$$

X.Bombois, G.Scorletti, M.Gevers, P.M.J. Van den Hof and R.Hildebrand.
Least costly identification experiment for control.
Automatica, 42(10):1651-1662, 2006

Limitations:

- Two distinct phases: identification and normal operation.
- $\mathcal{V}$ and $\mathcal{E}$ considered separately.

"Can we design experiments in such a way that the overall performance is optimized during the whole time of operation?"

Itendification for Control

Classical: "Given a maximum allowed perturbation, find the excitation signal that gives the best control performance."

$$\max \mathcal{V} \qquad \text{such that} \qquad \mathcal{E} \leq \bar{\mathcal{E}}.$$

M. Gevers and L.Ljung.
Optimal experiment designs with respect to the intended model application.
Automatica, 22(5):543-554, 1986

Least costly: "Given a minimum allowed performance level, find the excitation signal that minimizes the perturbation."

$$\min \mathcal{E} \qquad \text{such that} \qquad \mathcal{V} \geq \bar{\mathcal{V}}.$$

X.Bombois, G.Scorletti, M.Gevers, P.M.J. Van den Hof and R.Hildebrand.
Least costly identification experiment for control.
Automatica, 42(10):1651-1662, 2006

Limitations:

- Two distinct phases: identification and normal operation.
- $\mathcal{V}$ and $\mathcal{E}$ considered separately.

"Can we design experiments in such a way that the overall performance is optimized during the whole time of operation?"

Itendification for Control

Classical: "Given a maximum allowed perturbation, find the excitation signal that gives the best control performance."

$$\max \mathcal{V} \qquad \text{such that} \qquad \mathcal{E} \leq \bar{\mathcal{E}}.$$

M. Gevers and L.Ljung.
Optimal experiment designs with respect to the intended model application.
Automatica, 22(5):543-554, 1986

Least costly: "Given a minimum allowed performance level, find the excitation signal that minimizes the perturbation."

$$\min \mathcal{E} \qquad \text{such that} \qquad \mathcal{V} \geq \bar{\mathcal{V}}.$$

X.Bombois, G.Scorletti, M.Gevers, P.M.J. Van den Hof and R.Hildebrand.
Least costly identification experiment for control.
Automatica, 42(10):1651-1662, 2006

Limitations:

- Two distinct phases: identification and normal operation.
- $\mathcal{V}$ and $\mathcal{E}$ considered separately.

"Can we design experiments in such a way that the overall performance is optimized during the whole time of operation?"

The Framework

Linear system operated in closed-loop over n consecutive batches.

- Before a batch, identification and controller re-design.
- Excitation signal r_k in each batch.

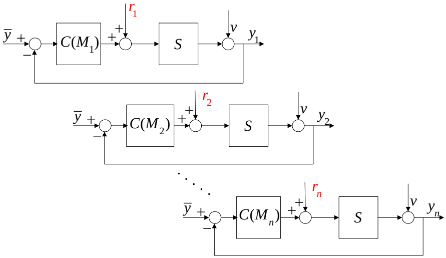

Excitation r_k has a dual effect. Worsens performance during the batch k, but can improve performance for batch $k+1$.

"Design the signals r_k to optimize the performance over n batches."

The Framework

Linear system operated in closed-loop over n consecutive batches.

- Before a batch, identification and controller re-design.
- Excitation signal r_k in each batch.

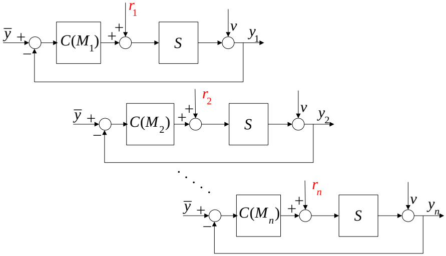

Excitation r_k has a dual effect. Worsens performance during the batch k, but can improve performance for batch $k+1$.

"Design the signals r_k to optimize the performance over n batches."

The Framework

Linear system operated in closed-loop over n consecutive batches.

- Before a batch, identification and controller re-design.
- Excitation signal r_k in each batch.

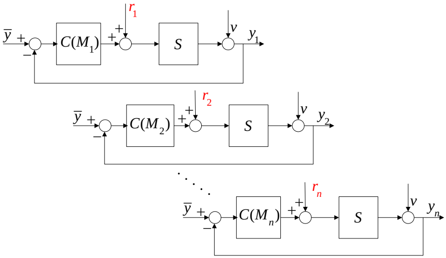

Excitation r_k has a dual effect. Worsens performance during the batch k, but can improve performance for batch $k + 1$.

"Design the signals r_k to optimize the performance over n batches."

The Framework

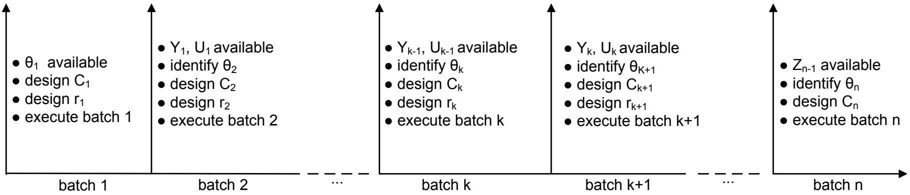

Steps:

- Identification
- Controller design
- Experiment design
- Execute batch k

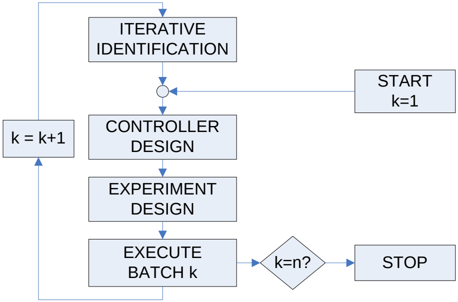

It is an actively adaptive learning control algorithm.

The Framework

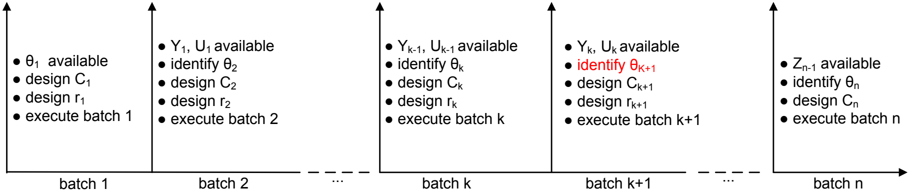

Steps:

- Identification
- Controller design
- Experiment design
- Execute batch k

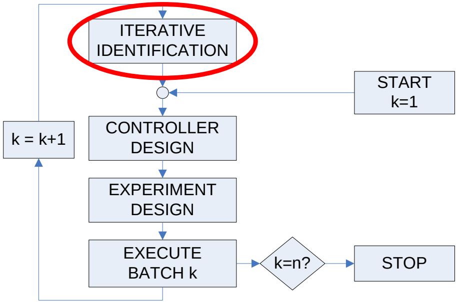

It is an actively adaptive learning control algorithm.

The Framework

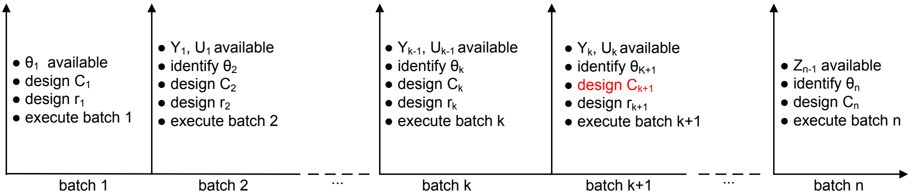

Steps:

- Identification
- Controller design
- Experiment design
- Execute batch k

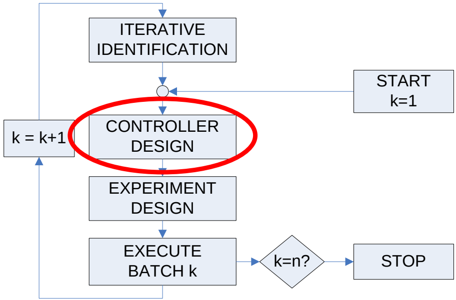

It is an actively adaptive learning control algorithm.

The Framework

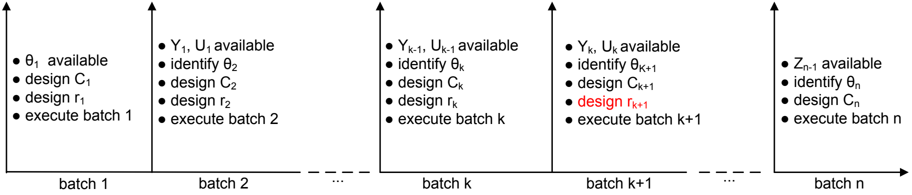

Steps:

- Identification
- Controller design
- Experiment design
- Execute batch k

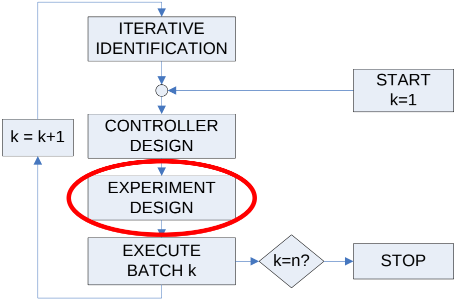

It is an actively adaptive learning control algorithm.

The Framework

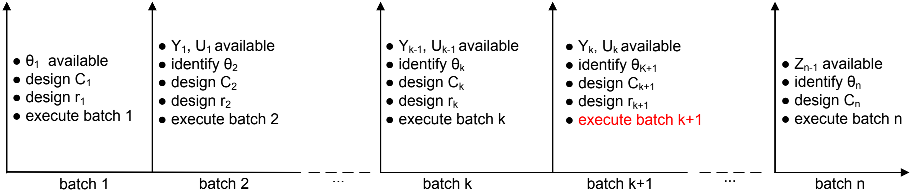

Steps:

- Identification
- Controller design
- Experiment design
- Execute batch k

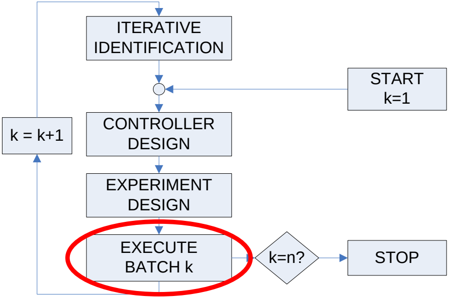

It is an actively adaptive learning control algorithm.

The Framework

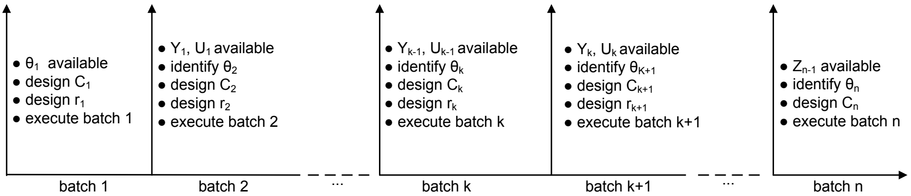

Steps:

- Identification
- Controller design
- Experiment design
- Execute batch k

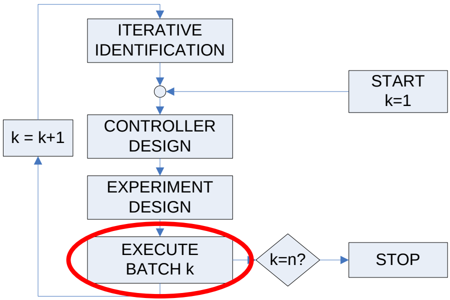

It is an actively adaptive learning control algorithm.

Iterative Identification

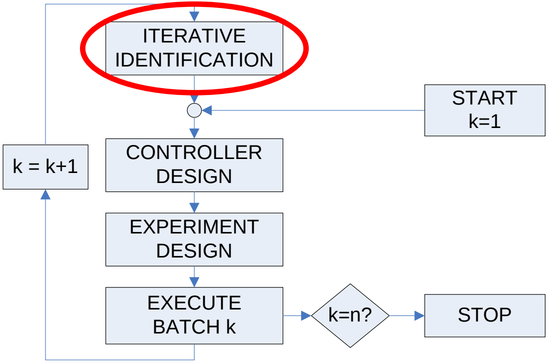

Iterative Identification

When the batch k is executed

- Data (Y_k, U_k) are collected.
- Previous estimate $\hat{\theta}_k \sim \mathcal{N}(\theta_o, R_k^{-1})$ is available.

The updated parameter estimate $\hat{\theta}_{k+1}$ is computed as

$$\hat{\theta}_{k+1} = \arg\min_{\theta} \frac{1}{\sigma_e^2} \left\| Y_k - \hat{Y}(U_k, \theta) \right\|_2^2 + \left\| \theta - \hat{\theta}_k \right\|_{R_k}^2 .$$

The parameter $\hat{\theta}_{k+1} \sim \mathcal{N}(\theta_o, R_{k+1}^{-1})$ with $R_{k+1} = R_k + F_k$.

Information Matrix and excitation spectrum

The information matrix F_k is a linear function of the spectrum $\Phi_r(\omega)$.

Iterative Identification

When the batch k is executed

- Data (Y_k, U_k) are collected.
- Previous estimate $\hat{\theta}_k \sim \mathcal{N}(\theta_o, R_k^{-1})$ is available.

The updated parameter estimate $\hat{\theta}_{k+1}$ is computed as

$$\hat{\theta}_{k+1} = \arg\min_{\theta} \frac{1}{\sigma_e^2} \left\| Y_k - \hat{Y}(U_k, \theta) \right\|_2^2 + \left\| \theta - \hat{\theta}_k \right\|_{R_k}^2 .$$

The parameter $\hat{\theta}_{k+1} \sim \mathcal{N}(\theta_o, R_{k+1}^{-1})$ with $R_{k+1} = R_k + F_k$.

Information Matrix and excitation spectrum

The information matrix F_k is a linear function of the spectrum $\Phi_r(\omega)$.

Iterative Identification

When the batch k is executed

- Data (Y_k, U_k) are collected.
- Previous estimate $\hat{\theta}_k \sim \mathcal{N}(\theta_o, R_k^{-1})$ is available.

The updated parameter estimate $\hat{\theta}_{k+1}$ is computed as

$$\hat{\theta}_{k+1} = \arg\min_{\theta} \frac{1}{\sigma_e^2} \left\| Y_k - \hat{Y}(U_k, \theta) \right\|_2^2 + \left\| \theta - \hat{\theta}_k \right\|_{R_k}^2 .$$

The parameter $\hat{\theta}_{k+1} \sim \mathcal{N}(\theta_o, R_{k+1}^{-1})$ with $R_{k+1} = R_k + F_k$.

Information Matrix and excitation spectrum

The information matrix F_k is a linear function of the spectrum $\Phi_r(\omega)$.

Iterative Identification

When the batch k is executed

- Data (Y_k, U_k) are collected.
- Previous estimate $\hat{\theta}_k \sim \mathcal{N}(\theta_o, R_k^{-1})$ is available.

The updated parameter estimate $\hat{\theta}_{k+1}$ is computed as

$$\hat{\theta}_{k+1} = \arg\min_{\theta} \frac{1}{\sigma_e^2} \left\| Y_k - \hat{Y}(U_k, \theta) \right\|_2^2 + \left\| \theta - \hat{\theta}_k \right\|_{R_k}^2 .$$

The parameter $\hat{\theta}_{k+1} \sim \mathcal{N}(\theta_o, R_{k+1}^{-1})$ with $R_{k+1} = R_k + F_k$.

Information Matrix and excitation spectrum

The information matrix F_k is a linear function of the spectrum $\Phi_r(\omega)$.

Controller Design

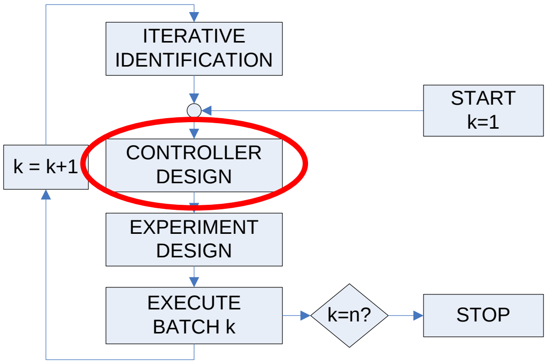

Here we use an $\mathcal{H}_2$ criterion. Different choices of $C(\hat{\theta}_k)$ possible. . .

Experiment Design

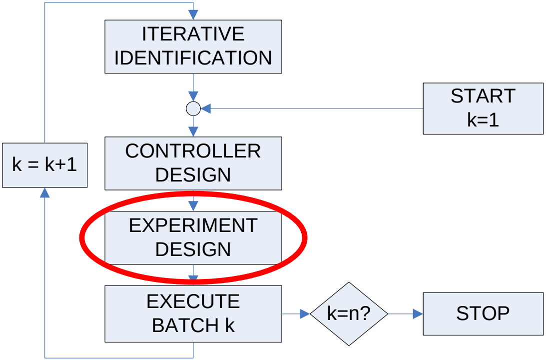

Experiment Design

Overview

Let us define:

Optimal Loop

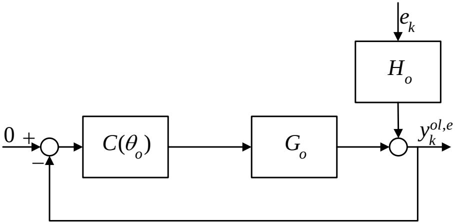

Experimental Loop

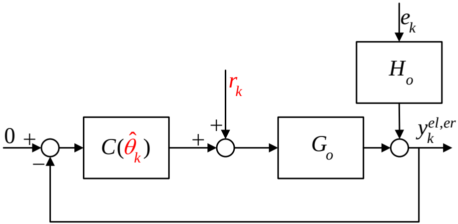

Experiment Design
Objective

Define the total cost for a batch as

$\mathcal{T}_k \triangleq$

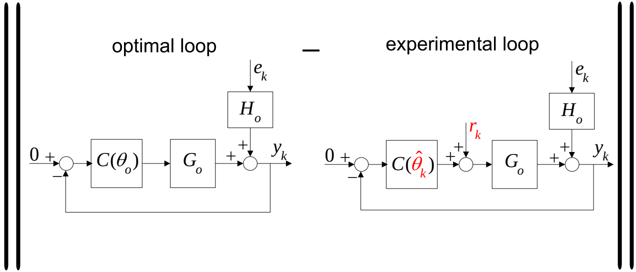

Excitation signals r_k are designed in order to

- minimize $\sum_{k=1}^{n} \mathcal{T}_k$.
- satisfy constraints $\mathcal{T}_k \leq \bar{\mathcal{T}}_k$ for each batch.

Experiment Design

Objective

Define the total cost for a batch as

$\mathcal{T}_k \triangleq$

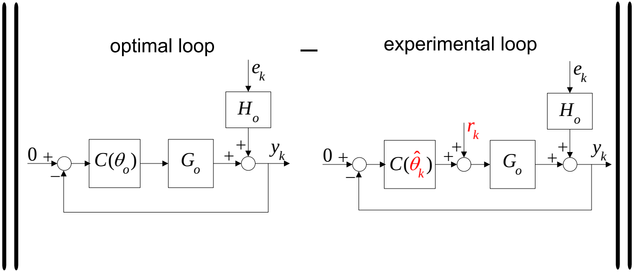

Excitation signals r_k are designed in order to

- minimize $\sum_{k=1}^{n} \mathcal{T}_k$.
- satisfy constraints $\mathcal{T}_k \leq \bar{\mathcal{T}}_k$ for each batch.

Total Cost, Application Cost & Excitation Cost

Total Cost: power of output difference between the two loops:

$$\mathcal{T}_k \triangleq E[(y_k^{ol,e} - y_k^{el,er})^2].$$

Optimal Loop

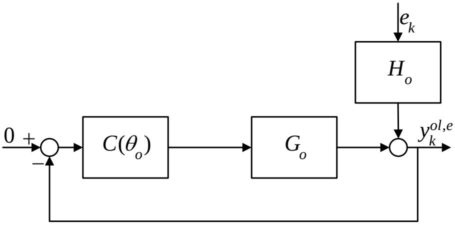

Experimental Loop

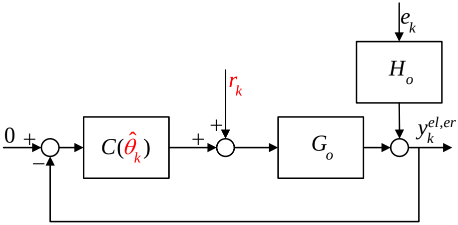

Since $r_k \perp e_k$:

Experiment Design

Total Cost, Application Cost & Excitation Cost

Total Cost: power of output difference between the two loops:

$$\mathcal{T}_k \triangleq E[(y_k^{ol,e} - y_k^{el,er})^2].$$

Optimal Loop

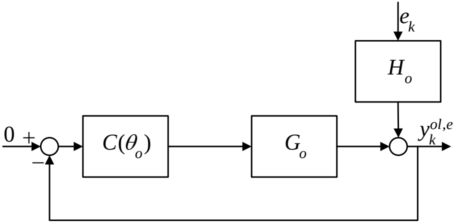

Experimental Loop

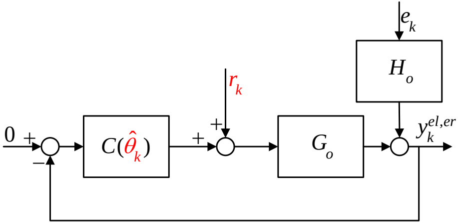

Since $r_k \perp e_k$:

$$\overbrace{E[(y_k^{ol,e} - y_k^{el,er})^2]}^{\text{Total Cost } \mathcal{T}_k} = \overbrace{E[(y_k^{ol,e} - y_k^{el,e})^2]}^{\text{Control Cost } \mathcal{V}_k} + \overbrace{E[(y_k^{el,r})^2]}^{\text{Excitation Cost } \mathcal{E}_k}.$$

Experiment Design

Total Cost, Application Cost & Excitation Cost

Total Cost: power of output difference between the two loops:

$$\mathcal{T}_k \triangleq E[(y_k^{ol,e} - y_k^{el,er})^2].$$

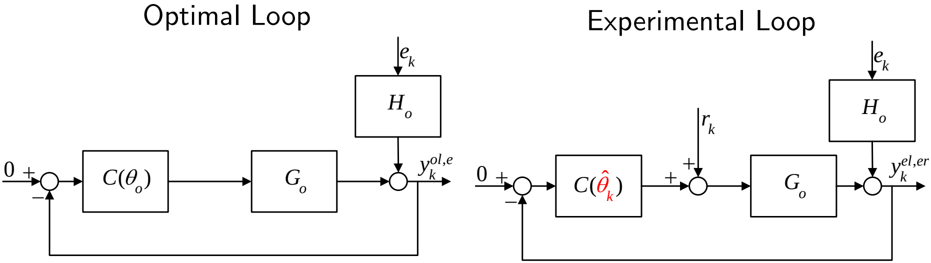

Since $r_k \perp e_k$:

$$\overbrace{E[(y_k^{ol,e} - y_k^{el,er})^2]}^{\text{Total Cost } \mathcal{T}_k} = \overbrace{E[(y_k^{ol,e} - y_k^{el,e})^2]}^{\text{Control Cost } \mathcal{V}_k} + \overbrace{E[(y_k^{el,r})^2]}^{\text{Excitation Cost } \mathcal{E}_k} .$$

Experiment Design

Total Cost, Application Cost & Excitation Cost

Total Cost: power of output difference between the two loops:

$$\mathcal{T}_k \triangleq E[(y_k^{ol,e} - y_k^{el,er})^2].$$

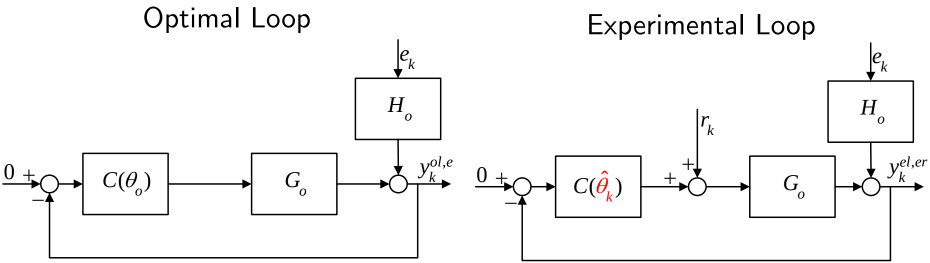

Since $r_k \perp e_k$:

$$\overbrace{E[(y_k^{ol,e} - y_k^{el,er})^2]}^{\text{Total Cost } \mathcal{T}_k} = \overbrace{E[(y_k^{ol,e} - y_k^{el,e})^2]}^{\text{Control Cost } \mathcal{V}_k} + \overbrace{E[(y_k^{el,r})^2]}^{\text{Excitation Cost } \mathcal{E}_k} \quad .$$

Experiment Design

Objective

- Experiment Design Problem (for $k = 1$):
  minimize the summation of the total cost over the future n batches

$$\min \sum_{k=1}^{n} \mathcal{T}_k \qquad \text{subject to}$$

$$\mathcal{T}_k \leq \bar{\mathcal{T}}_k,\ k = 1, 2, \ldots, n.$$

- Optimization variables: (spectra of) excitation signals $r_1, r_2, \ldots, r_n$.
- $\mathcal{T}_k = \mathcal{V}_k + \mathcal{E}_k$ random variables $\Rightarrow$ minimization in a worst-case sense.

  Approximations required to compute $\mathcal{T}_k^{\mathrm{wc}} = \mathcal{V}_k^{\mathrm{wc}} + \mathcal{E}_k^{\mathrm{wc}} \ldots$

  Leads to a convex SDP optimization (LMIs with linear objective)

Experiment Design

Objective

- Experiment Design Problem (for $k = 1$):
  minimize the summation of the total cost over the future n batches

$$\min_{\Phi_1, \Phi_2, \ldots, \Phi_n} \sum_{k=1}^{n} \mathcal{T}_k \qquad \text{subject to}$$

$$\mathcal{T}_k \leq \bar{\mathcal{T}}_k, \; k = 1, 2, \ldots, n.$$

- Optimization variables: (spectra of) excitation signals $r_1, r_2, \ldots, r_n$.
- $\mathcal{T}_k = \mathcal{V}_k + \mathcal{E}_k$ random variables $\Rightarrow$ minimization in a worst-case sense.

  Approximations required to compute $\mathcal{T}_k^{\mathrm{wc}} = \mathcal{V}_k^{\mathrm{wc}} + \mathcal{E}_k^{\mathrm{wc}} \ldots$

  Leads to a convex SDP optimization (LMIs with linear objective)

Experiment Design

Objective

- Experiment Design Problem (for $k = 1$):
  minimize the summation of the total cost over the future n batches

$$\min_{\Phi_1, \Phi_2, \ldots, \Phi_n} \sum_{k=1}^{n} \mathcal{T}_k^{\mathrm{wc}} \qquad \text{subject to}$$
$$\mathcal{T}_k^{\mathrm{wc}} \leq \bar{\mathcal{T}}_k, \; k = 1, 2, \ldots, n.$$

- Optimization variables: (spectra of) excitation signals $r_1, r_2, \ldots, r_n$.
- $\mathcal{T}_k = \mathcal{V}_k + \mathcal{E}_k$ random variables $\Rightarrow$ minimization in a worst-case sense.

Approximations required to compute $\mathcal{T}_k^{\mathrm{wc}} = \mathcal{V}_k^{\mathrm{wc}} + \mathcal{E}_k^{\mathrm{wc}} \ldots$

Leads to a convex SDP optimization (LMIs with linear objective)

Experiment Design

Objective

- Experiment Design Problem (for $k = 1$):
  minimize the summation of the total cost over the future n batches

$$\min_{\Phi_1,\Phi_2,\ldots,\Phi_n} \sum_{k=1}^{n} \mathcal{T}_k^{\mathrm{wc}} \qquad \text{subject to}$$

$$\mathcal{T}_k^{\mathrm{wc}} \leq \bar{\mathcal{T}}_k,\ k = 1, 2, \ldots, n.$$

- Optimization variables: (spectra of) excitation signals $r_1, r_2, \ldots, r_n$.
- $\mathcal{T}_k = \mathcal{V}_k + \mathcal{E}_k$ random variables $\Rightarrow$ minimization in a worst-case sense.

  Approximations required to compute $\mathcal{T}_k^{\mathrm{wc}} = \mathcal{V}_k^{\mathrm{wc}} + \mathcal{E}_k^{\mathrm{wc}} \ldots$

  Leads to a convex SDP optimization (LMIs with linear objective)

Experiment Design

Objective

- Experiment Design Problem (for $k = 1$):
  minimize the summation of the total cost over the future n batches

$$\min_{\Phi_1, \Phi_2, \ldots, \Phi_n} \sum_{k=1}^{n} \mathcal{T}_k^{\mathrm{wc}} \qquad \text{subject to}$$

$$\mathcal{T}_k^{\mathrm{wc}} \leq \bar{\mathcal{T}}_k, \; k = 1, 2, \ldots, n.$$

- Optimization variables: (spectra of) excitation signals $r_1, r_2, \ldots, r_n$.
- $\mathcal{T}_k = \mathcal{V}_k + \mathcal{E}_k$ random variables $\Rightarrow$ minimization in a worst-case sense.

  Approximations required to compute $\mathcal{T}_k^{\mathrm{wc}} = \mathcal{V}_k^{\mathrm{wc}} + \mathcal{E}_k^{\mathrm{wc}} \ldots$

  Leads to a convex SDP optimization (LMIs with linear objective)

Experiment Design

Chicken and the egg approximation

We will need to evaluate $\mathcal{V}_k^{\text{wc}}, \mathcal{E}_k^{\text{wc}}$ for $k = 1, \dots, n$ before the execution of the first batch.

- For the Control Cost

$$\mathcal{V}_k^{\text{wc}} = \min_{\lambda_k} \frac{1}{\lambda_k} \qquad \text{s.t } R_k(\theta_o, \hat{\theta}_k, \dots, \hat{\theta}_2, \hat{\theta}_1) \geq \lambda_k \frac{V''(\hat{\theta}_k)\chi^2_\alpha(n)}{2}.$$

- For the Excitation Cost

$$\mathcal{E}_k^{\text{wc}} = \max_s \; \mathcal{E}_k(\tilde{\theta}_s, \hat{\theta}_k).$$

Quantities in red are not known!

Typical chicken & the egg issue of Experiment Design.

They are all replaced with $\hat{\theta}_1$.

In order to alleviate the chicken & the egg issue, the Experiment Design is implemented in Receding Horizon over the batches.

Experiment Design

Chicken and the egg approximation

We will need to evaluate $\mathcal{V}_k^{\text{wc}}, \mathcal{E}_k^{\text{wc}}$ for $k = 1, \ldots, n$ before the execution of the first batch.

- For the Control Cost

$$\mathcal{V}_k^{\text{wc}} = \min_{\lambda_k} \frac{1}{\lambda_k} \qquad \text{s.t } R_k(\hat{\theta}_1, \hat{\theta}_1, \ldots, \hat{\theta}_1, \hat{\theta}_1) \geq \lambda_k \frac{V''(\hat{\theta}_1)\chi^2_\alpha(n)}{2}.$$

- For the Excitation Cost

$$\mathcal{E}_k^{\text{wc}} = \max_s \ \mathcal{E}_k(\tilde{\theta}_s, \hat{\theta}_1).$$

Quantities in red are not known!
Typical chicken & the egg issue of Experiment Design.
They are all replaced with $\hat{\theta}_1$.

In order to alleviate the chicken & the egg issue, the Experiment Design is implemented in Receding Horizon over the batches.

Experiment Design

Chicken and the egg approximation

We will need to evaluate $\mathcal{V}_k^{\mathrm{wc}}, \mathcal{E}_k^{\mathrm{wc}}$ for $k = 1, \dots, n$ before the execution of the first batch.

- For the Control Cost

$$\mathcal{V}_k^{\mathrm{wc}} = \min_{\lambda_k} \frac{1}{\lambda_k} \qquad \text{s.t } R_k(\hat{\theta}_1, \hat{\theta}_1, \dots, \hat{\theta}_1, \hat{\theta}_1) \geq \lambda_k \frac{V''(\hat{\theta}_1)\chi^2_\alpha(n)}{2}.$$

- For the Excitation Cost

$$\mathcal{E}_k^{\mathrm{wc}} = \max_s \ \mathcal{E}_k(\tilde{\theta}_s, \hat{\theta}_1).$$

Quantities in red are not known!
Typical chicken & the egg issue of Experiment Design.
They are all replaced with $\hat{\theta}_1$.

In order to alleviate the chicken & the egg issue, the Experiment Design is implemented in Receding Horizon over the batches.

Experiment Design

Receding Horizon Implementation

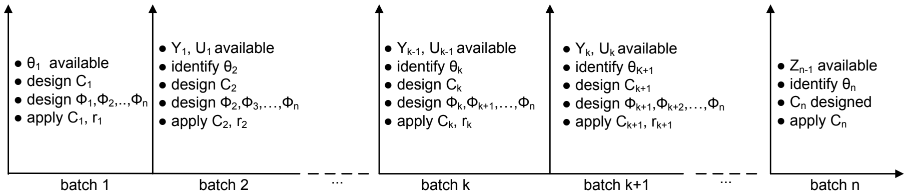

1. ED(1) for batch 1 based on $\hat{\theta}_1$. Spectra $(\Phi_1, \ldots, \Phi_n)$ found. r_1 applied in batch 1. Batch 1 executed, data (Y_1, U_1) collected.
2. Parameter $\hat{\theta}_2$ identified from the data. ED(2) for batch 2 based on $\hat{\theta}_2$. New spectra $(\Phi_2, \ldots, \Phi_n)$ found. Signal r_2 applied in batch 2. Batch 2 executed, data (Y_2, U_2) collected.
3. . . .

Experiment Design

Receding Horizon Implementation

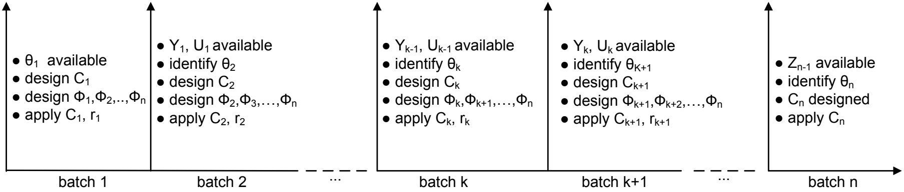

1. ED(1) for batch 1 based on $\hat{\theta}_1$. Spectra $(\Phi_1, \ldots, \Phi_n)$ found. r_1 applied in batch 1. Batch 1 executed, data (Y_1, U_1) collected.
2. Parameter $\hat{\theta}_2$ identified from the data. ED(2) for batch 2 based on $\hat{\theta}_2$. New spectra $(\Phi_2, \ldots, \Phi_n)$ found. Signal r_2 applied in batch 2. Batch 2 executed, data (Y_2, U_2) collected.
3. ...

Experiment Design

Receding Horizon Implementation

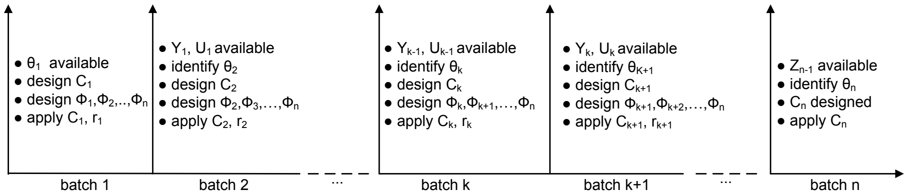

1. ED(1) for batch 1 based on $\hat{\theta}_1$. Spectra $(\Phi_1, \ldots, \Phi_n)$ found. r_1 applied in batch 1. Batch 1 executed, data (Y_1, U_1) collected.
2. Parameter $\hat{\theta}_2$ identified from the data. ED(2) for batch 2 based on $\hat{\theta}_2$. New spectra $(\Phi_2, \ldots, \Phi_n)$ found. Signal r_2 applied in batch 2. Batch 2 executed, data (Y_2, U_2) collected.
3. ...

Numerical Simulation

Second-order system $\mathcal{S}_o$ in a BJ model structure.

- $N = 2400$ total samples.
- $n = 12$ batches of length 200.
- Constraints:
  - $\mathcal{T}_k \leq 0.7$ for $k = 1, \ldots, 6$.
  - $\mathcal{T}_k \leq 0.05$ for $k = 7, \ldots, 12$.

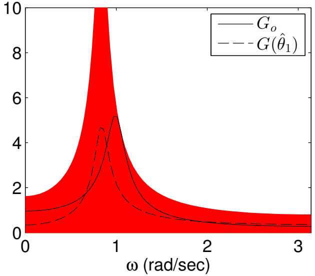

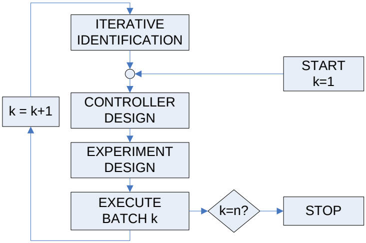

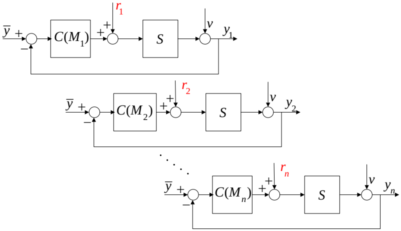

Numerical Simulation

Excitation Spectra

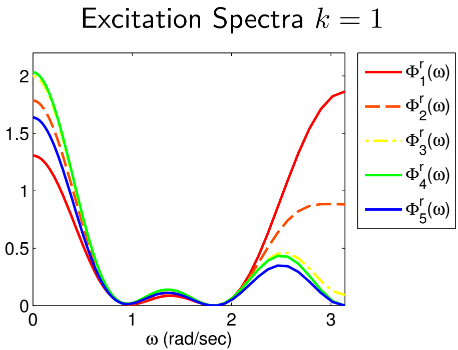

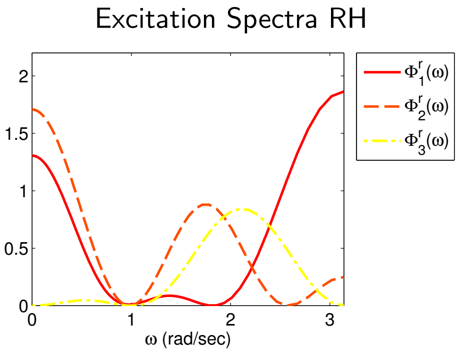

Numerical Simulation

Total cost

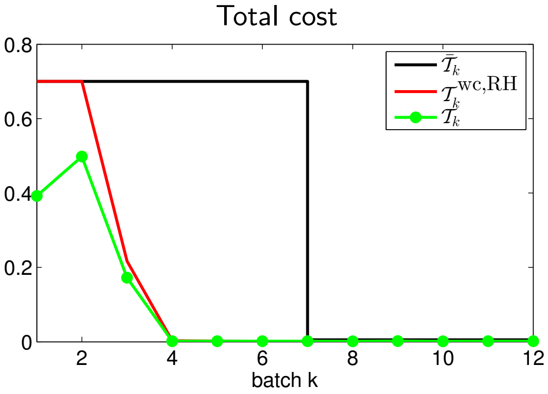

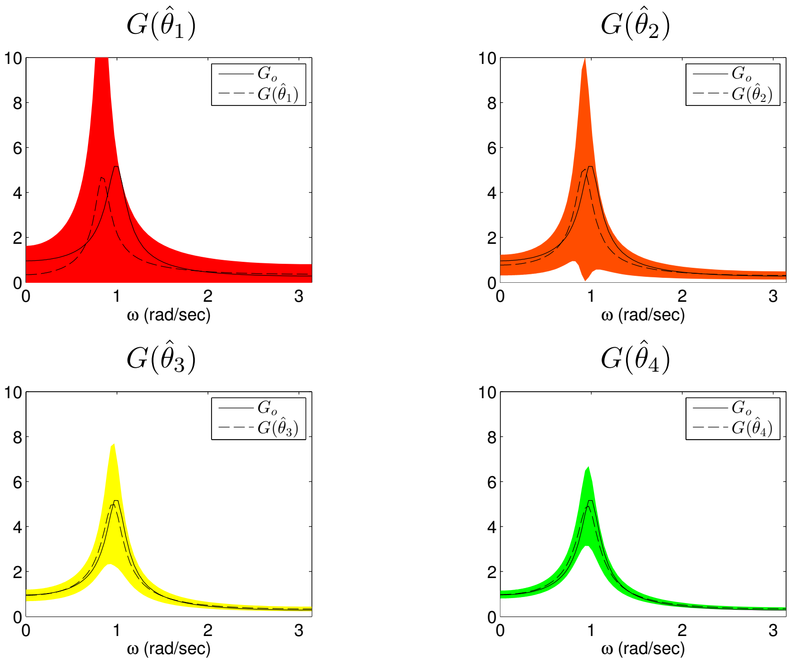
$G(\hat{\theta}_1)$
G_o
$G(\hat{\theta}_1)$
ω (rad/sec)
$G(\hat{\theta}_2)$
G_o
$G(\hat{\theta}_2)$
ω (rad/sec)
$G(\hat{\theta}_3)$
G_o
$G(\hat{\theta}_3)$
ω (rad/sec)
$G(\hat{\theta}_4)$
G_o
$G(\hat{\theta}_4)$
ω (rad/sec)

Numerical Simulation

Total cost

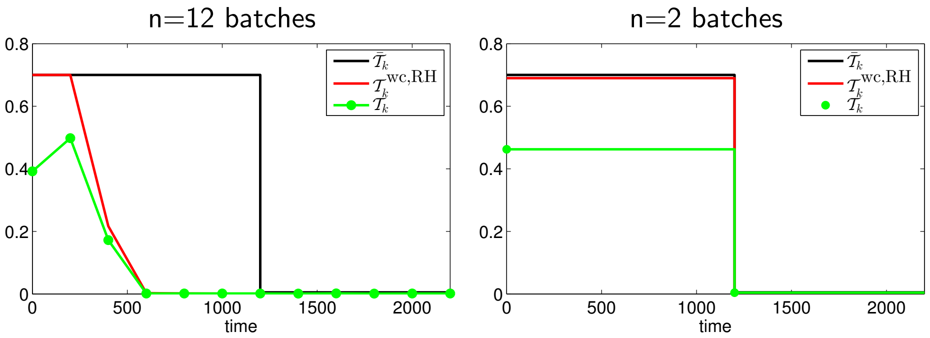

$n = 2$ corresponds to a least costly identification.

Numerical Simulation

Total cost

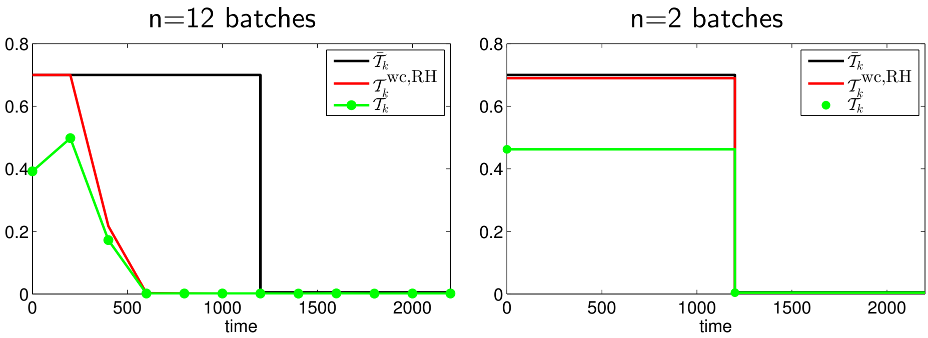

$n = 2$ corresponds to a least costly identification.

Conclusions

An actively adaptive control algorithm based on Experiment Design tools.

- Optimization of the overall performance.
- No distinction between identification and control batches.
- Excitation only when it pays back.

Some open issues:

- Approximations to compute the worst-case. Analysis?
- Batch systems are often nonlinear.
- Initial conditions plays a significative role.

On-going work for nonlinear experiment design.

Conclusions

An actively adaptive control algorithm based on Experiment Design tools.

- Optimization of the overall performance.
- No distinction between identification and control batches.
- Excitation only when it pays back.

Some open issues:

- Approximations to compute the worst-case. Analysis?
- Batch systems are often nonlinear.
- Initial conditions plays a significative role.

On-going work for nonlinear experiment design.

Conclusions

An actively adaptive control algorithm based on Experiment Design tools.

- Optimization of the overall performance.
- No distinction between identification and control batches.
- Excitation only when it pays back.

Some open issues:

- Approximations to compute the worst-case. Analysis?
- Batch systems are often nonlinear.
- Initial conditions plays a significative role.

On-going work for nonlinear experiment design.

Thank you.
Questions?

Experiment Design

Worst-case control cost

- From Parseval relation $\mathcal{V}_k = E[(y_k^{ol,e} - y_k^{el,e})^2] \ =$

$$\mathcal{V}_k(\theta_o, \hat{\theta}_k) = \frac{1}{2\pi}\int_{-\pi}^{\pi} \left| \frac{1}{1 + C(\hat{\theta}_k)G(\theta_o)} - \frac{1}{1 + C(\theta_o)G(\theta_o)} \right|^2 |H|^2 (\theta_o)\sigma_e^2 \, d\omega$$

- We approximate $\mathcal{V}_k(\theta_o, \hat{\theta}_k)$ as a quadratic function of θ_o locally around $\hat{\theta}_k$

$$\mathcal{V}_k(\theta_o, \hat{\theta}_k) \approx \frac{1}{2}(\theta_o - \hat{\theta}_k)^\top V''(\hat{\theta}_k)(\theta_o - \hat{\theta}_k).$$

- Since $\theta_o - \hat{\theta}_k \sim \mathcal{N}(0, R_k^{-1})$, using standard ellipsoids we can find the worst-case $\mathcal{V}_k$ with probability α as

$$\mathcal{V}_k^{\text{wc}} = \min_{\lambda_k} \frac{1}{\lambda_k} \qquad \text{s.t } R_k(\Phi_1, \Phi_2, \ldots, \Phi_{k-1}) \geq \lambda_k \frac{V''\chi_\alpha^2(n)}{2}.$$

- $R_k(\Phi_1, \Phi_2, \ldots, \Phi_{k-1})$ linear $\Rightarrow$ convex optimization!

Experiment Design

Worst-case control cost

- From Parseval relation $\mathcal{V}_k = E[(y_k^{ol,e} - y_k^{el,e})^2] \; =$

$$\mathcal{V}_k(\theta_o, \hat{\theta}_k) = \frac{1}{2\pi}\int_{-\pi}^{\pi}\left|\frac{1}{1+C(\hat{\theta}_k)G(\theta_o)} - \frac{1}{1+C(\theta_o)G(\theta_o)}\right|^2 |H|^2(\theta_o)\sigma_e^2 \, d\omega$$

- We approximate $\mathcal{V}_k(\theta_o, \hat{\theta}_k)$ as a quadratic function of θ_o locally around $\hat{\theta}_k$

$$\mathcal{V}_k(\theta_o, \hat{\theta}_k) \approx \frac{1}{2}(\theta_o - \hat{\theta}_k)^\top V''(\hat{\theta}_k)(\theta_o - \hat{\theta}_k).$$

- Since $\theta_o - \hat{\theta}_k \sim \mathcal{N}(0, R_k^{-1})$, using standard ellipsoids we can find the worst-case $\mathcal{V}_k$ with probability α as

$$\mathcal{V}_k^{\text{wc}} = \min_{\lambda_k} \frac{1}{\lambda_k} \qquad \text{s.t } R_k(\Phi_1, \Phi_2, \ldots, \Phi_{k-1}) \geq \lambda_k \frac{V''\chi_\alpha^2(n)}{2}.$$

- $R_k(\Phi_1, \Phi_2, \ldots, \Phi_{k-1})$ linear $\Rightarrow$ convex optimization!

Experiment Design

Worst-case control cost

- From Parseval relation $\mathcal{V}_k = E[(y_k^{ol,e} - y_k^{el,e})^2] \ =$

$$\mathcal{V}_k(\theta_o, \hat{\theta}_k) = \frac{1}{2\pi}\int_{-\pi}^{\pi}\left|\frac{1}{1 + C(\hat{\theta}_k)G(\theta_o)} - \frac{1}{1 + C(\theta_o)G(\theta_o)}\right|^2 |H|^2(\theta_o)\sigma_e^2 \, d\omega$$

- We approximate $\mathcal{V}_k(\theta_o, \hat{\theta}_k)$ as a quadratic function of θ_o locally around $\hat{\theta}_k$

$$\mathcal{V}_k(\theta_o, \hat{\theta}_k) \approx \frac{1}{2}(\theta_o - \hat{\theta}_k)^\top V''(\hat{\theta}_k)(\theta_o - \hat{\theta}_k).$$

- Since $\theta_o - \hat{\theta}_k \sim \mathcal{N}(0, R_k^{-1})$, using standard ellipsoids we can find the worst-case $\mathcal{V}_k$ with probability α as

$$\mathcal{V}_k^{\text{wc}} = \min_{\lambda_k} \frac{1}{\lambda_k} \qquad \text{s.t } R_k(\Phi_1, \Phi_2, \ldots, \Phi_{k-1}) \geq \lambda_k \frac{V''\chi_\alpha^2(n)}{2}.$$

- $R_k(\Phi_1, \Phi_2, \ldots, \Phi_{k-1})$ linear $\Rightarrow$ convex optimization!

Experiment Design

Worst-case control cost

- From Parseval relation $\mathcal{V}_k = E[(y_k^{ol,e} - y_k^{el,e})^2] \;=$

$$\mathcal{V}_k(\theta_o, \hat{\theta}_k) = \frac{1}{2\pi}\int_{-\pi}^{\pi}\left|\frac{1}{1 + C(\hat{\theta}_k)G(\theta_o)} - \frac{1}{1 + C(\theta_o)G(\theta_o)}\right|^2 |H|^2(\theta_o)\sigma_e^2 \, d\omega$$

- We approximate $\mathcal{V}_k(\theta_o, \hat{\theta}_k)$ as a quadratic function of θ_o locally around $\hat{\theta}_k$

$$\mathcal{V}_k(\theta_o, \hat{\theta}_k) \approx \frac{1}{2}(\theta_o - \hat{\theta}_k)^\top V''(\hat{\theta}_k)(\theta_o - \hat{\theta}_k).$$

- Since $\theta_o - \hat{\theta}_k \sim \mathcal{N}(0, R_k^{-1})$, using standard ellipsoids we can find the worst-case $\mathcal{V}_k$ with probability α as

$$\mathcal{V}_k^{\text{wc}} = \min_{\lambda_k} \frac{1}{\lambda_k} \qquad \text{s.t } R_k(\Phi_1, \Phi_2, \ldots, \Phi_{k-1}) \geq \lambda_k \frac{V''\chi_\alpha^2(n)}{2}.$$

- $R_k(\Phi_1, \Phi_2, \ldots, \Phi_{k-1})$ linear $\Rightarrow$ convex optimization!

Experiment Design

Worst-case excitation cost

The excitation cost

$$\mathcal{E}_k(\theta_o, \hat{\theta}_k) = \frac{1}{2\pi} \int_{-\pi}^{\pi} \left| \frac{G(\theta_o)}{1 + C(\hat{\theta}_k) G(\theta_o)} \right|^2 \Phi_k^r(\omega)\, d\omega.$$

depends on the decision variables!

Solution based on Randomized Algorithms. . .
Using the initial estimate $\hat{\theta}_1 \sim \mathcal{N}(\theta_o, R_1^{-1})$:

1. Draw q samples $\tilde{\theta}_s$.
2. Compute $\mathcal{E}_{k,s} = \mathcal{E}_k(\tilde{\theta}_s, \theta_k)$ for $s = 1, \ldots, q$.
3. Extract the empirical maximum $\mathcal{E}_k^{\mathrm{wc}} = \max_s \ \mathcal{E}_{k,s}$.

The number of samples q can be tuned such that $\mathcal{E}_k^{\mathrm{wc}}$ is the Worst Case Excitation Cost with probability α (randomized algorithms).

Experiment Design

Worst-case excitation cost

The excitation cost

$$\mathcal{E}_k(\theta_o, \hat{\theta}_k) = \frac{1}{2\pi} \int_{-\pi}^{\pi} \left| \frac{G(\theta_o)}{1 + C(\hat{\theta}_k) G(\theta_o)} \right|^2 \Phi_k^r(\omega) \, d\omega.$$

depends on the decision variables!

Solution based on Randomized Algorithms. . .
Using the initial estimate $\hat{\theta}_1 \sim \mathcal{N}(\theta_o, R_1^{-1})$:

1. Draw q samples $\tilde{\theta}_s$.
2. Compute $\mathcal{E}_{k,s} = \mathcal{E}_k(\tilde{\theta}_s, \theta_k)$ for $s = 1, \ldots, q$.
3. Extract the empirical maximum $\mathcal{E}_k^{\mathrm{wc}} = \max_s \; \mathcal{E}_{k,s}$.

The number of samples q can be tuned such that $\mathcal{E}_k^{\mathrm{wc}}$ is the Worst Case Excitation Cost with probability α (randomized algorithms).

Experiment Design

Worst-case excitation cost

The excitation cost

$$\mathcal{E}_k(\theta_o, \hat{\theta}_k) = \frac{1}{2\pi} \int_{-\pi}^{\pi} \left| \frac{G(\theta_o)}{1 + C(\hat{\theta}_k) G(\theta_o)} \right|^2 \Phi_k^r(\omega)\, d\omega.$$

depends on the decision variables!

Solution based on Randomized Algorithms. . .
Using the initial estimate $\hat{\theta}_1 \sim \mathcal{N}(\theta_o, R_1^{-1})$:

1. Draw q samples $\tilde{\theta}_s$.
2. Compute $\mathcal{E}_{k,s} = \mathcal{E}_k(\tilde{\theta}_s, \theta_k)$ for $s = 1, \ldots, q$.
3. Extract the empirical maximum $\mathcal{E}_k^{\text{wc}} = \max_s \ \mathcal{E}_{k,s}$.

The number of samples q can be tuned such that $\mathcal{E}_k^{\text{wc}}$ is the Worst Case Excitation Cost with probability α (randomized algorithms).