Data-driven model improvement for model-based control

Citation for published version (APA):

Forgione, M., Bombois, X., & Hof, Van den, P. M. J. (2015). Data-driven model improvement for model-based control. *Automatica*, 52, 118-124. <https://doi.org/10.1016/j.automatica.2014.11.006>

DOI:

[10.1016/j.automatica.2014.11.006](https://doi.org/10.1016/j.automatica.2014.11.006)

Document status and date:

Published: 01/01/2015

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Brief paper

Data-driven model improvement for model-based control[☆] Marco Forgone^{a,1}, Xavier Bombois^b, Paul M.J. Van den Hof^c^a Laboratoire Ampère, Ecole Centrale de Lyon, 36 Avenue Guy de Collongue, 69134 Ecully Cedex, France^b Delft Center for Systems and Control, Mekelweg 2 (3mE building) 2628 CD, Delft, The Netherlands^c Department of Electrical Engineering, P.O. Box 513, 5600 MB, Eindhoven, The Netherlands

ARTICLE INFO

Article history:

Received 17 May 2013

Received in revised form

5 June 2014

Accepted 28 October 2014

Available online 4 December 2014

Keywords:

Identification for Control

Experiment Design

Iterative schemes

Dual control

ABSTRACT

We present a framework for the gradual improvement of model-based controllers. The total time of the learning procedure is divided into a number of learning intervals. After a learning interval, the model is refined based on the measured data. This model is used to synthesize the controller that will be applied during the next learning interval. Excitation signals can be injected into the control loop during each of the learning intervals. On the one hand, the introduction of an excitation signal worsens the control performance during the current learning interval since it acts as a disturbance. On the other hand, the informative data generated owing to the excitation signal are used to refine the model using a closed-loop system identification technique. Therefore, the control performance for the next learning interval is expected to improve. In principle, our objective is to maximize the overall control performance taking the effect of the excitation signals explicitly into account. However, this is in general an intractable optimization problem. For this reason, a convex approximation of the original problem is derived using standard relaxations techniques for Experiment Design. The approximated problem can be solved efficiently using common optimization routines. The applicability of the method is demonstrated in a simulation study.

© 2014 Elsevier Ltd. All rights reserved.

1. Introduction

It is well known that the performance of a model-based controlled system largely depends on the quality of the model that is used to synthesize the controller. System identification provides tools that can be used to construct models using measured input/output data. Together with the model, most of the identification methods also provide a measure of the model uncertainty.

The relation between the uncertainty of an identified model and the expected control performance has been studied in the field of Identification for Control (Gevers, 2005). An important finding was that often few model features determine the performance of

the controller to a large extent. Therefore, the identification experiments have to be designed in such a way that these control-relevant features can be accurately identified (Gevers, 2002). In the earliest contributions, open-loop identification experiments were used for this purpose (Gevers & Ljung, 1986).

During the nineties, closed-loop identification techniques (Van Den Hof & Schrama, 1995) gained an increasing attention in the Identification for Control community. A clear advantage of closed-loop identification is that the normal (closed-loop) operation can continue while the identification data are collected. In this sense, the closed-loop identification procedures are intrinsically less intrusive than the ones based on open-loop data.

The use of closed-loop identification was also supported by the intuition that the control-relevant features of the system are naturally emphasized when the system is operating in a condition which resembles the desired controlled behaviour. In specific cases, it was formally proven that the optimal experimental condition for control-oriented identification are met when the optimal controller (with respect to the same control objective) is present in the loop (Hjalmarsson, Gevers, & De Bruyne, 1996).

Note however that the optimal controller is always unknown. In fact, its determination is the ultimate goal of the user. Therefore, the optimal experimental conditions can only be approached by adopting iterative schemes consisting of repeated closed-loop

[☆] The financial support of the Institute for Sustainable Process Technology (ISPT) under the grant ISPT PH-00-04 is gratefully acknowledged. The material in this paper was partially presented at the 2013 American Control Conference (ACC13), June 17–19 2013, Washington, DC, USA. This paper was recommended for publication in revised form by Associate Editor Wei Xing Zheng under the direction of Editor Torsten Söderström.

E-mail addresses: marco.forgione@ec-lyon.fr (M. Forgone), x.j.a.bombois@tudelft.nl (X. Bombois), p.m.j.vandenhof@tue.nl (P.M.J. Van den Hof).

¹ Tel.: +33 472 186042; fax: +33 478433717.

identification and model-based control design steps (De Callafon & Van den Hof, 1997; Schrama, 1992).

An aspect that was not thoroughly investigated in these contributions was the choice of the *excitation signals* fed to the system during the closed-loop experiments, often simply taken as white noise signals for ease of analysis. However, a careful choice of the excitation signals can be beneficial both to improve the model accuracy and to limit the identification cost (Bombois, Scorletti, Gevers, Van den Hof, & Hildebrand, 2006). On the one hand, a high level of excitation leads to informative data sets which can be used to identify accurate models. On the other hand, the excitation signals also act as disturbances on the controlled system and consequently leads to a (temporary) performance degradation when they are applied. Therefore, there is a trade-off between the performance degradation due to the application of the excitation signals and the improvement that is expected due to the increased model accuracy. In the literature, the excitation is said to have a *dual effect* on the control performance (Tse & Bar-Shalom, 1973) for this reason.

In this paper, we consider the problem of designing the excitation signals in an iterative identification/control scheme aiming to maximize the overall control performance, while guaranteeing a minimum performance level at all time. In our framework, the total time of the learning procedure is divided into a number of *learning intervals*. Excitation signals can be injected into the control loop during each of these intervals. After an interval, the measured data are used to refine the model using closed-loop identification. Based on that model, a new controller is designed. The controller is applied during the next learning interval, and so on and so forth for the following ones.

We define the cost $\mathcal{T}_k$ of one interval k as the sum of the performance degradation due to the difference between the true system and the identified model (modelling error cost), and the one due to the presence of the excitation signal (excitation cost) during that interval. Following, we determine the excitation signals to be applied in each interval by minimizing the sum $\sum_k \mathcal{T}_k$ of the cost over all the intervals, subject to a constraint $\mathcal{T}_k \leq \bar{\mathcal{T}}_k$ on the cost for each interval. Note that the modelling error cost in one interval depends on the excitation signals applied during all the previous intervals since the model is identified based on the previous data, while the excitation cost depends on the excitation signal applied during the current interval. Thus, we are here taking the dual effect of the excitation signals explicitly into account.

It has to be mentioned that the optimization problem that we would like to solve in order to find these excitation signal is intractable as such. However, using established relaxation techniques and tools developed in the field of Experiment Design (ED), we can derive an approximation of the original optimization problem that is convex and can be solved efficiently.

In fact, the problem of designing excitation signals which guarantee desired properties for the identified model has been extensively studied in the ED field. The classic ED approaches consider only two distinct phases: an identification phase in which the excitation signal is fed to the system and a model is identified, and a control phase in which a controller based on the identified model is applied. The objective is to find compromise between the excitation cost in the identification phase and the modelling error cost in the control phase (Bombois et al., 2006; Gevers & Ljung, 1986).

A limit of the classic approaches is that if the identification phase is too short, it might not be possible to satisfy the performance requirement for the control phase without violating constraints on the excitation in the identification phase. One could circumvent this issue by extending the duration of the identification phase, but this implies that he/she has to wait a longer time before having any improvement in the control performance.

Our approach can be seen as an extension of the classic ED to a situation with several phases (i.e. the learning intervals), and in

which the dual effect of the excitation signal is considered altogether for all the learning intervals. As we will show in the numerical example, by considering several learning intervals we can gradually improve the controller and achieve a better overall performance than in a classic two-phase ED framework.

Our approach has also a certain analogy with the *actively adaptive* learning algorithms discussed in Pronzato, Kulcsár, and Walter (1996), since it takes explicitly into account the dual effect of the excitation. However, the approach in Pronzato et al. (1996) leads to stochastic dynamic optimization problems that can be solved only for very specific model structures and control objectives, while our approach can be applied to almost any LTI model structure and control objective. There are indeed important differences between our approach and the one in Pronzato et al. (1996). First, in Pronzato et al. (1996) the input is optimized for the identification and the control objective altogether, while in our framework the controller takes care of the control objective and the superposed excitation signal takes care of the identification objective. Second, the model is updated at each time instant in Pronzato et al. (1996), while we perform the identification only at the end of a learning interval. These simplifications allow us to use the classical ED tools to tackle this complicate problem in a wider range of cases.

The rest of this paper is organized as follows. In Section 2 the framework is discussed in detail. In Section 3 the ED problem is introduced and the approximated convex optimization problem is derived. The framework is applied to a simulation study in Section 5 and conclusions are drawn in Section 6.

2. The framework

The *true system* S_o is the linear time-invariant system

$$y = G_o(q^{-1})u + H_o(q^{-1})e \quad (1)$$

where u is the input, y is the output, e is white Gaussian noise with variance σ_e^2 , and q^{-1} is the unit-delay operator. G_o and H_o are stable discrete-time transfer functions; H_o is monic and minimum phase.

S_o is known to belong to a *model structure* $\mathcal{M} = \{M(\theta), \theta \in \mathbb{R}^p\}$ where θ is the model parameter. The true system S_o is described in $\mathcal{M}$ by a (unique) *true parameter* θ_o , i.e. $\exists! \theta_o \mid S_o = M(\theta_o)$.

In order to reject the disturbance $H_o(q^{-1})e$, we would like to operate (1) in closed loop. We assume that an initial model $M_1 = M(\hat{\theta}_1)$ and controller $C_1 = C(\hat{\theta}_1)$ are available. However, M_1 is a relatively poor representation of the true system S_o , and the performance of the initial loop $[C_1 \ S_o]$ is rather poor.

In the sequel, we will present an iterative model and controller update procedure whose objective is to gradually improve the accuracy of the model and the control performance. For this purpose, the total time of the learning procedure is divided into n learning intervals of duration N . During each interval, a specially tailored excitation signal r_k is applied to the loop. At the end of the interval, input and output data are collected and are used to improve the accuracy of the available model using prediction error identification. Based on the improved model, the controller is updated and applied to the true system for the next interval. The procedure is illustrated in Fig. 1 and will be presented in more detail in the next subsection.

Remark 1. If an initial model M_1 is not available, or the uncertainty of M_1 is so large that the controller C_1 designed using such model could lead to an unstable loop $[C_1 \ S_o]$, the first interval can be performed in open loop (i.e. $C_1 = 0$). Indeed, the true system S_o is per se stable.

Remark 2. It is possible that the closed-loop system will be operated for a total time that is longer than the learning procedure itself. By setting a constraint $\mathcal{T}_N \leq \bar{\mathcal{T}}_N$ on the total cost in the last interval of the learning procedure, one can guarantee that the performance $\bar{\mathcal{T}}_N$ will be maintained for the rest of the closed-loop operation.

Fig. 1. The n learning intervals with successive model and controller updates.

2.1. Iterative identification and control design

Suppose that we are about to start the learning interval k . At this moment, the true system is operated with a controller $C(\hat{\theta}_k)$ which has been designed with an identified model $M_k = M(\hat{\theta}_k)$. As will become clear in the sequel, the identified parameter vector $\hat{\theta}_k$ is normally distributed around θ_o with a covariance matrix P_k .

During the N time samples of interval k , an excitation signal r_k is applied (see Fig. 1) to the closed-loop system: $u_k = r_k - C(\hat{\theta}_k)y_k$. Even though other choices are possible, we will assume in the sequel that r_k is chosen as a white noise signal filtered by an arbitrary FIR filter. At the end of interval k , the data set $Z_k = \{u_k(t), y_k(t) \mid t = 1, \dots, N\}$ is collected and is used to obtain a more accurate model $M_{k+1} = M(\hat{\theta}_{k+1})$. For this purpose, the parameter vector $\hat{\theta}_{k+1}$ is identified using not only the new data set Z_k , but also all previous data sets $Z_{k-1}, Z_{k-2}, \dots, Z_1$. Since $\hat{\theta}_k$ has been determined using the previous data sets, this can be done by using $\hat{\theta}_k$ and its covariance matrix P_k as a regularization term. Hence, $\hat{\theta}_{k+1}$ is determined as

$$\hat{\theta}_{k+1} = \arg \min_{\theta \in \mathbb{R}^p} \frac{1}{\sigma_e^2} \sum_{t=1}^N \varepsilon_k^2(t, \theta) + (\theta - \hat{\theta}_k)^\top P_k^{-1} (\theta - \hat{\theta}_k) \quad (2)$$

where $\varepsilon_k(t, \theta) = H(q^{-1}, \theta)^{-1} (y_k(t) - G(q^{-1}, \theta)u_k(t))$. Since $\hat{\theta}_{k+1}$ is determined based on kN data and N has been chosen relatively large, it is acceptable to use the asymptotic properties of the prediction error estimate (Ljung, 1999). Consequently, the parameter vector $\hat{\theta}_{k+1}$ identified in this way is (assumed) normally distributed around θ_o with a covariance matrix P_{k+1} given by

$$P_{k+1}^{-1} = I_k + P_k^{-1} \quad (3)$$

where $I_k > 0$ is the so-called *information matrix* corresponding to the data Z_k : $I_k(\theta_o) = \frac{N}{\sigma_e^2} \bar{E} (\psi_k(t, \theta_o) \psi_k(t, \theta_o)^\top)$ with $\psi_k(t, \theta_o) = -\frac{\partial \varepsilon_k(t, \theta)}{\partial \theta} \Big|_{\theta=\theta_o}$ (Ljung, 1999). Using the recursive nature of (3), we have also that

$$P_{k+1}^{-1} = I_k + I_{k-1} + \dots + I_1 + P_1^{-1} \quad (4)$$

where P_1 is the covariance matrix of the initial parameter vector $\hat{\theta}_1$. This covariance matrix is available for instance if the initial model had been previously identified using the prediction error framework. If that is not the case, P_1^{-1} can be set to 0 in (4).

Since the updated parameter is estimated with an increasing number of data after each interval, its accuracy will increase for increasing k . In fact, we can define for each estimate $\hat{\theta}_k$ an *uncertainty ellipsoid* $\mathcal{D}_k$ where the modelling error $\theta_o - \hat{\theta}_k$ lies with a certain probability α as

$$\mathcal{D}_k \triangleq \{\delta \in \mathbb{R}^p \mid \delta^\top P_k^{-1} \delta \leq \chi_p^2(\alpha)\} \quad (5)$$

where $\chi_p^2(\alpha)$ is the α -percentile of the chi-squared distribution having p degrees of freedom. Due to (3) and since $I_k > 0$, the volume of the uncertainty region decreases after each interval. Note

that this property always holds if r_k is designed as filtered white noises and does not require r_k to be designed in an optimal sense.

The spectrum Φ_{r_k} of the signal r_k can be parametrized as follows (Jansson & Hjalmarsson, 2005): $\Phi_{r_k}(\omega) = R_k(0) + 2 \sum_{j=1}^m R_k(j) \cos(j\omega)$ where $R_k(j)$, $j = 0, 1, \dots, m$ is the autocorrelation sequence of the signal r_k . The coefficients $R_k(j)$ will be used as decision variables of the optimal experiment design problem. Using this linear parametrization of Φ_{r_k} , the information matrix I_k defined below (3) can be rewritten as an affine function of the coefficients R_k (Bombois et al., 2006). We will hereafter use the notation $I_k(\theta_o, R_k)$ for the information matrix. Note that since $I_k(\theta_o, R_k)$ is affine in R_k , P_{k+1}^{-1} is affine in $R_1, R_2, \dots, R_k$ due to (4).

Using the data Z_k , we have obtained a more accurate model $M_{k+1} = M(\hat{\theta}_{k+1})$. This new model can be used to design an updated controller $C(\hat{\theta}_{k+1})$. We will suppose that a control design method has been fixed a-priori and thus the controller is a function $C(\cdot)$ of the parameter vector. The choice for the controller design criterion is not important in the development of our framework.

Since the model $M(\hat{\theta}_{k+1})$ is more accurate than $M(\hat{\theta}_k)$, the controller $C(\hat{\theta}_{k+1})$ is very likely to perform better than the controller $C(\hat{\theta}_k)$ that was in the loop during interval k . However, before applying this new controller to the true system, it is safer to verify whether $C(\hat{\theta}_{k+1})$ does not destabilize any of the loops $[C(\hat{\theta}_{k+1}) \ M(\hat{\theta}_{k+1} + \delta)]$ with $\delta \in \mathcal{D}_{k+1}$. A necessary and sufficient condition to perform this robust stability test can be found in Bombois, Gevers, Scorletti, and Anderson (2001). If the controller $C(\hat{\theta}_{k+1})$ passes this test, we have the guarantee that the loop $[C(\hat{\theta}_{k+1}) \ S_o]$ is stable (with at least the probability α related to $\mathcal{D}_{k+1}$) and the controller $C(\hat{\theta}_{k+1})$ can be applied during the interval $k+1$.

Remark 3. In the eventuality that robust stability is not validated for the controller $C(\hat{\theta}_{k+1})$, the controller $C(\hat{\theta}_k)$ can be kept for the interval $k+1$. This situation is further discussed in Remark 7.

2.2. Total cost, modelling error cost, excitation cost

As mentioned in the introduction, one of the objectives of the optimal experiment design problem in this paper is to minimize the overall cost over the n intervals. The cost $\mathcal{T}_k$ of interval k can be evaluated by comparing the output y_k obtained during this interval and the output that would have been obtained in an ideal situation (i.e. a situation where we would perfectly know the true system S_o). The output y_k during interval k is given by

$$y_k(t) = \frac{\overbrace{H_o}^{=y_{e,k}}}{1 + C(\hat{\theta}_k)G_o} e(t) + \frac{\overbrace{G_o}^{=y_{r,k}}}{1 + C(\hat{\theta}_k)G_o} r_k \quad (6)$$

while the output y_o of the ideal loop is

$$y_o(t) = \frac{H_o}{1 + C(\theta_o)G_o} e(t). \quad (7)$$

The differences between (6) and (7) are the presence of the term $y_{r,k}$ (which is due to the excitation signal r_k) in (6) and the different controllers present in the two loops. Indeed, in the ideal loop the controller $C(\theta_o)$ based on the true parameter θ_o is present, while the controller $C(\hat{\theta}_k)$ actually present in the loop during the interval k is designed based on the parameter $\hat{\theta}_k$.

The cost $\mathcal{T}_k$ can now be defined as the power of the difference $y_k - y_o$ between these two outputs. Since r_k is independent of e , the total cost $\mathcal{T}_k$ can be split up into the sum of a *modelling error cost* $\mathcal{V}_k$

(i.e. the power of $y_{e,k} - y_0$) and an *excitation cost* $\mathcal{E}_k$ (i.e. the power of $y_{r,k}$). By applying the Parseval theorem, we get

$$\mathcal{T}_k = \underbrace{\left\| \frac{H(\theta_0)}{1 + C(\theta_0)G(\theta_0)} - \frac{H(\theta_0)}{1 + C(\hat{\theta}_k)G(\theta_0)} \right\|_{\mathcal{H}_2}^2}_{=\mathcal{V}_k} \sigma_e^2 + \underbrace{\frac{1}{2\pi} \int_{-\pi}^{\pi} \left| \frac{G_0}{1 + C(\hat{\theta}_k)G_0} \right|^2 \Phi_{r_k} d\omega}_{=\mathcal{E}_k}. \quad (8)$$

The modelling error cost $\mathcal{V}_k$ represents the performance degradation caused by the use of the controller $C(\hat{\theta}_k)$ instead of the optimal controller $C(\theta_0)$. The excitation cost $\mathcal{E}_k$ represents the performance degradation caused by the introduction of the excitation signal r_k .

The modelling error cost $\mathcal{V}_k = \mathcal{V}_k(\theta_0, \hat{\theta}_k)$ is a nonlinear function of θ_0 and $\hat{\theta}_k$ and has a global minimum equal to 0 in $\theta_0 = \hat{\theta}_k$. If we parametrize the spectrum Φ_{r_k} of r_k as done in the previous subsection, the excitation cost $\mathcal{E}_k$ is linear in the coefficients $R_k(j)$ (Bombois et al., 2006). We denote the excitation cost as $\mathcal{E}_k(\theta_0, \hat{\theta}_k, R_k) = R_k^\top c(\theta_0, \hat{\theta}_k)$ where $R_k \in \mathbb{R}^{(m+1) \times 1}$ is a vector containing the coefficients $R_k(j)$, $j = 0, 1, \dots, m$ and $c(\theta_0, \hat{\theta}_k) \in \mathbb{R}^{(m+1) \times 1}$ is a nonlinear vector function of θ_0 and $\hat{\theta}_k$ (Bombois et al., 2006).

Note that the costs $\mathcal{V}_k$ and $\mathcal{E}_k$ cannot be evaluated since they both depend on the unknown true parameter vector θ_0 . However, we can consider these quantities in a worst-case sense by computing their maximum value over the ellipsoid $\mathcal{D}_k$:

$$\mathcal{V}_k^{\text{wc}} \triangleq \max_{\delta \in \mathcal{D}_k} \mathcal{V}_k(\hat{\theta}_k + \delta, \hat{\theta}_k), \quad \mathcal{E}_k^{\text{wc}} \triangleq \max_{\delta \in \mathcal{D}_k} \mathcal{E}_k(\hat{\theta}_k + \delta, \hat{\theta}_k, R_k). \quad (9)$$

For the computation of the worst-case terms (9), we will here use the approach introduced in Hjalmarsson (2009) and based on a second-order Taylor approximation of the functions $\mathcal{V}_k(\cdot, \hat{\theta}_k)$ and $\mathcal{E}_k(\cdot, \hat{\theta}_k, R_k)$ around $\hat{\theta}_k$:

$$\mathcal{V}_k(\hat{\theta}_k + \delta, \hat{\theta}_k) \approx \frac{1}{2} \delta^\top V''(\hat{\theta}_k) \delta \quad (10)$$

$$\mathcal{E}_k(\theta_k + \delta, \hat{\theta}_k, R_k) \approx R_k^\top c(\hat{\theta}_k, \hat{\theta}_k) + R_k^\top J_c(\hat{\theta}_k) \delta + \frac{1}{2} \delta^\top \left(\sum_{j=0}^m E_j(\hat{\theta}_k) R_k(j) \right) \delta \quad (11)$$

where $V''(\hat{\theta}_k)$ and $E_j(\hat{\theta}_k)$ are the Hessian matrices of $\mathcal{V}_k(\cdot, \hat{\theta}_k)$ and of the j th entry of $c(\cdot, \hat{\theta}_k)$, respectively, and where $J_c(\hat{\theta}_k)$ is the Jacobian of $c(\cdot, \hat{\theta}_k)$. Using (10), $\mathcal{V}_k^{\text{wc}}$ can be computed as Hjalmarsson (2009)

$$\mathcal{V}_k^{\text{wc}} = \min_{\lambda_k} \lambda_k \quad \text{such that} \quad P_k^{-1} \geq \frac{V''(\hat{\theta}_k) \chi_\alpha^2(p)}{2\lambda_k}. \quad (12)$$

Using (11) and applying S-procedure (Boyd & Vandenberghe, 2004), $\mathcal{E}_k^{\text{wc}}$ can be computed as

$$\mathcal{E}_k^{\text{wc}} = \min_{\gamma_k, \tau_k} \gamma_k \quad \text{such that} \quad (13)$$

$$\begin{aligned} \tau_k &\geq 0 \\ \left[\begin{array}{cc} \frac{1}{2} \sum_j E_j R_k(j) - \tau_k \frac{P_k^{-1}}{\chi^2} & \frac{1}{2} (R_k^\top J_c)^\top \\ \frac{1}{2} R_k^\top J_c & R_k^\top c(\hat{\theta}_k, \hat{\theta}_k) + \tau_k - \gamma_k \end{array} \right] &\leq 0. \end{aligned}$$

Note that the second-order approximations (10)–(11) will become more accurate when $\hat{\theta}_k$ is closer to θ_0 . Consequently, the effects of this approximation will automatically decrease when k becomes larger.

Remark 4. Another approach for obtaining the worst-case terms $\mathcal{V}_k^{\text{wc}}$ and $\mathcal{E}_k^{\text{wc}}$ is to compute the maximum values of $\mathcal{V}_k(\hat{\theta}_k + \delta_l, \hat{\theta}_k)$ and $\mathcal{E}_k(\hat{\theta}_k + \delta_l, \hat{\theta}_k)$, respectively over a number of randomly extracted sample points δ_l , $l = 1, 2, \dots, n_s$. It is possible to determine the minimum value of n_s which guarantees a given worst-case probability using the theory of Randomized Algorithms. Since $\mathcal{V}_k^{\text{wc}}$ and $\mathcal{E}_k^{\text{wc}}$ will be computed inside a convex optimization problem (see next section), n_s can be chosen according to the classic *scenario approach* (Campi, Garatti, & Prandini, 2009). If the number of points n_s determined according to Campi et al. (2009) is prohibitively large, the more recent results of the so-called *sequential approach* (Chamanbaz, Dabbene, Tempo, Venkataramanan, & Qing-Guo, 2013) can also be applied.

3. Experiment Design

Before interval $k = 1$, the objective of the ED is to determine the spectra of the excitation signals for all the n learning intervals which will minimize (in a worst case sense) the overall cost over the n intervals while guaranteeing that the cost of each interval remains below a given threshold. Let us define the worst-case total cost $\mathcal{T}_k^{\text{wc}} \triangleq \mathcal{V}_k^{\text{wc}} + \mathcal{E}_k^{\text{wc}}$. The ED problem can be expressed as (Forgione, Bombois, & Van den Hof, 2013)

$$R^{\text{opt}} = \arg \min_R \sum_{k=1}^n \mathcal{T}_k^{\text{wc}} \quad \text{such that} \quad (14)$$

$$\mathcal{T}_k^{\text{wc}} \leq \bar{\mathcal{T}}_k, \quad \text{for } k = 1, 2, \dots, n \quad (15)$$

where the variable $R \triangleq \{R_1, R_2, \dots, R_n\}$ contains the coefficients parametrizing the excitation spectra for all the learning intervals. The dual effect of the excitation signal is incorporated in the problem formulation (14)–(15). On the one hand, if the excitation spectrum Φ_{r_k} during the interval k is “large”, the worst-case total cost $\mathcal{T}_k^{\text{wc}}$ is also large due to the contribution of $\mathcal{E}_k^{\text{wc}}$. On the other hand, this large excitation will lead to a “small” covariance matrix P_{k+1} , which in turn gives a small modelling error cost $\mathcal{V}_{k+1}^{\text{wc}}$ for the next interval (see (12)).

In order to formulate the above optimization problem as a convex problem, we will use the affine relation existing between the decision variable R and P_k^{-1} and the fact that P_k^{-1} appears linearly in the constraint of (12). In (13), this it is unfortunately not the case since P_k^{-1} appears in a product with the decision variable τ_k . To convexify the constraint in (13), we redefine $\mathcal{E}_k^{\text{wc}}$ as the worst case cost over the initial uncertainty ellipsoid $\mathcal{D}_1$ (instead of $\mathcal{D}_k$), which is equivalent to replacing P_k by the initial covariance matrix² P_1 in (13). This introduces a conservatism, but $\mathcal{E}_k^{\text{wc}}$ remains an upper bound on $\mathcal{E}_k$. It is furthermore an acceptable approximation since the actual cost $\mathcal{E}_k$ (unlike $\mathcal{V}_k$) is not related to the modelling error and thus it is not expected to reduce after each interval. Finally, we will also have to tackle the so-called “chicken-and-egg” problem, that is a characteristic of most optimal experiment design frameworks (Ljung, 1999). Indeed, the constraints in (12) and (13) are functions of the identified parameter vectors $\hat{\theta}_k$, $k \geq 2$ that are not available before the first interval i.e. at the moment when (14)–(15) has to be solved. As is generally done, $\hat{\theta}_k$ for all $k > 1$ will be replaced by an initial estimate in the optimization problem. In this case, the initial estimate is $\hat{\theta}_1$. We also use $\hat{\theta}_1$ as an estimate of θ_0 in the expression of the information matrix $I_k : I_k(\theta_0, R_k) \approx I_k(\hat{\theta}_1, R_k)$. It is also to be noted that, as opposed to other ED frameworks,

² If P_1 is not available, the nominal excitation cost $R_k^\top c(\hat{\theta}_k, \hat{\theta}_k)$ can be used instead of $\mathcal{E}_k^{\text{wc}}$.

the effects of the approximations introduced to tackle the chicken-and-egg problem (and of the replacement of P_k by P_1 in (13)) will be mitigated by the receding horizon mechanism proposed later in this paper (see Section 4).

Proposition 5. Let us introduce n scalar variables $\lambda \triangleq \{\lambda_1, \lambda_2, \dots, \lambda_n\}$, n scalar variables $t \triangleq \{t_1, t_2, \dots, t_n\}$, n scalar variables $\gamma \triangleq \{\gamma_1, \gamma_2, \dots, \gamma_n\}$, n scalar variables $\tau \triangleq \{\tau_1, \tau_2, \dots, \tau_n\}$ and n matrix variables $Q \triangleq \{Q_1, Q_2, \dots, Q_n\}$, $Q_k \in \mathbb{R}^{m \times m}$. Using (12) (with $\hat{\theta}_k = \hat{\theta}_1$) and (13) (with $\hat{\theta}_k = \hat{\theta}_1$ and $P_k = P_1$) to compute γ_k^{wc} and $\mathcal{E}_k^{\text{wc}}$ respectively, the ED Problem (14)–(15) corresponds to the following convex semidefinite problem

$$R^{\text{opt}} = \arg \min_{R, \lambda, \gamma, \tau, Q} \sum_{k=1}^n \lambda_k + \sum_{k=1}^n \gamma_k \quad \text{such that} \quad (16)$$

$$\lambda_k + \gamma_k \leq \bar{\mathcal{J}}_k \quad (17)$$

$$\overbrace{P_{k-1}^{-1} + I_k(\hat{\theta}_1, R_k)}^{P_k^{-1}} \geq t_k \frac{\chi_{\alpha}^2(p)V''(\hat{\theta}_1)}{2} \quad (18)$$

$$\tau_k \geq 0$$

$$\begin{bmatrix} \frac{1}{2} \sum_j E_j R_k(j) - \tau_k \frac{P_1^{-1}}{\chi^2} & \frac{1}{2} (R_k^T J_c)^T \\ \frac{1}{2} R_k^T J_c & R_k^T c(\hat{\theta}_1, \hat{\theta}_1) + \tau_k - \gamma_k \end{bmatrix} \leq 0$$

$$\begin{bmatrix} \lambda_k & 1 \\ 1 & t_k \end{bmatrix} \geq 0 \quad (19)$$

$$\begin{bmatrix} Q_k - A^T Q_k A & C_k^T - A^T Q_k B \\ C_k - B^T Q_k A & D_k + D_k^T - B^T Q_k B \end{bmatrix} \geq 0 \quad (20)$$

for $k = 1, \dots, n$

with $A = \begin{bmatrix} 0 & 0 \\ I_{m-1} & 0 \end{bmatrix}$, $B = [1 \ 0 \ \dots \ 0]$, $C_k = [R_k(1) \ R_k(2) \ \dots \ R_k(m)]$ and $D_k = \frac{R_k(0)}{2}$.

Proof. The objective function (16) is the sum of the worst case modelling error costs λ_k and excitation costs γ_k . The LMIs (19) guarantee the conditions $\lambda_k \geq 0$, $t_k \geq 0$ and $t_k \geq \frac{1}{\lambda_k}$ simultaneously. This in turn implies that (18) is equivalent to the constraint in (12). Finally, (20) guarantees that $\Phi_{r_k}(\omega) \geq 0$, $\forall \omega$ as an application of the positive real lemma (see Jansson & Hjalmarsson, 2005).

4. Receding horizon

By solving the optimization problem (16)–(20) before the first interval we can design not only the spectrum $\Phi_{r_1}^{\text{opt}}$ of the signal r_1 that will be applied during this first interval, but also the spectra of the excitation signals $r_2, r_3, \dots, r_n$ for all the following intervals. During the first interval, we generate a signal r_1 having the desired spectrum $\Phi_{r_1}^{\text{opt}}$ and apply this signal as excitation to the loop $[C(\hat{\theta}_1) \ S_0]$. After the execution of this first interval, the data Z_1 are collected and a new parameter vector $\hat{\theta}_2$ is identified using (2). Based on $\hat{\theta}_2$, a new controller $C(\hat{\theta}_2)$ is designed and applied to the true system S_0 (after the robust stability check). We could then proceed with interval 2 by applying to the closed loop $[C(\hat{\theta}_2) \ S_0]$ a signal r_2 having the spectrum $\Phi_{r_2}^{\text{opt}}$ obtained from the solution of the previous optimization problem. However, we here propose to redesign the spectra Φ_{r_k} for $k = 2, 3, \dots, n$ using the newly identified parameter vector $\hat{\theta}_2$. This parameter vector is indeed a more accurate estimate of θ_0 than the initial estimate $\hat{\theta}_1$ since it has been estimated with twice more data (and consequently $P_2 < P_1$).

Consequently, evaluating $\mathcal{E}_k^{\text{wc}}$ at interval $k \geq 2$ as the worst case of $\mathcal{E}_k$ over the new uncertainty ellipsoid $\mathcal{D}_2$ is less conservative than doing it over $\mathcal{D}_1$. Moreover, replacing θ_0 and $\hat{\theta}_k$ (for $k > 2$) by $\hat{\theta}_2$ instead of $\hat{\theta}_1$ is also more appropriate in order to tackle the chicken and egg problem since $\hat{\theta}_2$ will be generally closer to θ_0 and $\hat{\theta}_k$ (for $k > 2$) than $\hat{\theta}_1$.

For the reasons above, the spectra Φ_{r_k} , $k = 2, 3, \dots, n$ will be redesigned solving a similar ED problem as the one presented in Section 3, but using the new estimate $\hat{\theta}_2$ and its covariance matrix P_2 to evaluate $\mathcal{E}_k^{\text{wc}}$ and to deal with the chicken-and-egg problem. This spectrum redesign procedure, inspired by the receding horizon mechanism in MPC control (Maciejowski & Huzmezan, 1997), will be performed after each interval. Note that a similar mechanism for the adaptive solution of an ED problem was adopted in Stigter, Vries, and Keesman (2006).

Owing to the spectrum redesign procedure, the effects of the approximations on the obtained spectra will become smaller after each interval and the obtained spectra will become increasingly more effective to achieve the ED objectives. Note that even though the effect of the approximation may be significant for the very first intervals, our approach will still lead to models having and increasing accuracy. This is indeed guaranteed whatever the excitation signal, as long as it is chosen (like in our case) as filtered white noise.

Remark 6. In principle, a better performance could always be obtained by reducing the interval length N . The controller and the excitation signals could be updated more frequently using a more recent parameter estimate. In this sense, the best choice would be to set N to just one time step. However, in the ED procedure we made use of properties of prediction error identification which are asymptotic in N . Therefore, we need to choose N sufficiently large for these properties to hold. Dropping this condition would lead to the same (generally untractable) formulations obtained in the actively adaptive learning algorithms (Pronzato et al., 1996). It is hard to quantify the effect of the interval length N on the overall performance in analytical form. Simulation results with different values of N will be presented in the next section.

Remark 7. The situation described in Remark 3 where the robust stability check fails for the controller C_{k+1} and C_k is kept for the interval $k + 1$ may be a symptom that the identified model M_{k+1} is for some unexpected reason rather poor. In this circumstance, the result of the ED procedure (which is based on such model) may also be very inaccurate. The excitation signal r_{k+1} can be set in this case to a white noise signal having the maximum allowed power. By doing so, the accuracy of the model M_{k+2} will increase in all the directions and the problem most likely will not occur at the next iteration.

5. Simulation study

The framework is applied in this section to a simulation example. We consider a Box–Jenkins (BJ) model structure $\mathcal{M} = \{M(\theta), \theta \in \mathbb{R}^6\}$. A model $M(\theta)$ in this structure has $G(q^{-1}, \theta) = \frac{\theta_1 q^{-1} + \theta_2 q^{-2}}{1 + \theta_3 q^{-1} + \theta_4 q^{-2}}$, $H(q^{-1}, \theta) = \frac{1 + \theta_5 q^{-1}}{1 + \theta_4 q^{-1}}$. The true system $S_0 = M(\theta_0)$ is described by $\theta_0 = [0.8 \ 0 \ 0 \ -0.6 \ 0.985 \ 0.819]^T$ and the variance of e is $\sigma_e^2 = 1$. The initial model is $M(\theta_1)$ with $\hat{\theta}_1 = [0.676 \ 0.464 \ 0.099 \ 0.6 \ 1.24 \ 0.858]^T$. The initial covariance matrix P_1 is

$$P_1 = \begin{bmatrix} 0.044 & -0.022 & 0 & 0 & 0.007 & -0.009 \\ -0.022 & 0.056 & 0 & 0 & 0.008 & 0.003 \\ 0 & 0 & 0.0006 & 0.0004 & 0 & 0 \\ 0 & 0 & 0.0004 & 0.0004 & 0 & 0 \\ 0.007 & 0.008 & 0 & 0 & 0.007 & -0.003 \\ -0.009 & 0.003 & 0 & 0 & -0.003 & 0.005 \end{bmatrix}.$$

Fig. 2. Spectra Φ_k^f solution of the first ED problem in Case 1.

Fig. 3. Spectra Φ_k^r of the excitation signals actually applied to the system in Case 1.

The controllers design function is the $\mathcal{H}_2$ criterion $C(\hat{\theta}_k) = \arg \min_K \left\| \begin{bmatrix} \frac{H(\hat{\theta}_k)}{1 + KG(\hat{\theta}_k)} & \frac{\sqrt{\beta}KH(\hat{\theta}_k)}{1 + KG(\hat{\theta}_k)} \end{bmatrix}^T \right\|_{\mathcal{H}_2}^2$ with $\beta = 0.1$. The initial controller is thus $C_1 = C(\hat{\theta}_1)$. The worst-case terms are computed with probability $\alpha = 0.99$.

For the first case considered (Case 1), the total time N_{tot} is divided into $n = 12$ intervals having equal length $N = 200$. The constraints on the worst case total cost are set to $\tilde{\mathcal{T}}_k = \tilde{\mathcal{T}}^h = 0.7$ for $k = 1, \dots, 6$ and $\tilde{\mathcal{T}}_k = \tilde{\mathcal{T}}^l = 0.005$ for $k = 7, \dots, 12$.

A first ED problem based on the initial model is performed before the execution of interval 1 and the optimal sequence of excitation spectra $\{\Phi_{r_1}, \Phi_{r_2}, \dots, \Phi_{r_n}\}$ is found. The first five spectra of this sequence are reported in Fig. 2. The following spectra are zero up to numerical precision. An excitation signal r_1 with spectrum Φ_{r_1} is generated and the interval 1 is performed using this excitation signal.

After the execution of the first interval, the data Z_1 are collected and used to estimate the parameter $\hat{\theta}_2$ and its covariance P_2 . Following, the controller $C_2 = C(\hat{\theta}_2)$ is also designed and the robust stability of the uncertain closed loop system $[C_2 M(\hat{\theta}_2 + \delta)]$ with $\delta \in \mathcal{D}_2$ is verified using the robust stability tools in Bombois et al. (2001). Subsequently, a new ED problem involving the remaining intervals is formulated and solved. The result is a new sequence of excitation spectra $\{\Phi_{r_2}, \Phi_{r_3}, \dots, \Phi_{r_n}\}$. The first element of this sequence is used to realize the excitation signal r_2 implemented in the interval 2 and the procedure is iterated for all the following intervals.

The spectra of the excitation signals actually fed to the system in the RH implementation are reported in Fig. 3. Note that these spectra are significantly different from the ones computed before the first interval (Fig. 2). Furthermore, only the spectra relative to the first 3 intervals are non-zero. Owing to the RH mechanism, the algorithm detects that the excitation in the intervals 4 and 5 is no longer required. By removing the excitation in these intervals, the overall performance improves significantly.

Fig. 4. Total cost $\mathcal{T}_k$ vs. time t in Case 1 ($n = 12$), Case 2 ($n = 6$), Case 3 ($n = 4$) and Case 4 ($n = 2$) intervals. The green circles denote the time instants corresponding to the start of a learning interval.

Fig. 5. Bode plot of G_o , $G(\hat{\theta}_1)$, $G(\hat{\theta}_2)$, $G(\hat{\theta}_3)$, $G(\hat{\theta}_4)$.

In the top left plot of Fig. 4 the experimental total cost $\mathcal{T}_k^e$ is reported,³ together with the constraint $\tilde{\mathcal{T}}_k$ and the worst case total cost $\mathcal{T}_k^{wc,RH}$ of interval k dynamically constructed during the operation.⁴ $\mathcal{T}_k^{wc,RH}$ reaches the constraint $\tilde{\mathcal{T}}_k$ for the first three intervals and decreases in the following ones, so that it can satisfy the new level of the constraint for $k = 7, \dots, 12$. The experimental total cost $\mathcal{T}_k^e$ is always below $\tilde{\mathcal{T}}_k$ as expected. $\mathcal{T}_k^e$ increases from interval 1 to 2, since more excitation is applied in the second interval. In the following intervals, $\mathcal{T}_k^e$ decreases and is close to zero for $k \geq 4$.

In Fig. 5 the Bode diagrams of G_o and of the identified models $G(\hat{\theta}_1)$, $G(\hat{\theta}_2)$, $G(\hat{\theta}_3)$, $G(\hat{\theta}_4)$ are reported. For the identified models, the 99% uncertainty region is indicated by the coloured area. The improvement of the models and the reduction of their uncertainty regions over the first four learning intervals is evident in this plot.

In order to investigate the influence of the interval length, we applied the ED framework to the same system in three more cases dividing the total time $N_{\text{tot}} = 2400$ into $n = 6$ intervals having equal length $N = 400$ (Case 2), into $n = 4$ intervals having equal length $N = 600$ (Case 3), and into $n = 2$ intervals having lengths $N_1 = 1800$, $N_2 = 600$, respectively (Case 4). The constraints $\tilde{\mathcal{T}}_k$ are set to $\tilde{\mathcal{T}}^h = 0.7$ in the first $n/2$ intervals and to $\tilde{\mathcal{T}}^l = 0.05$ for the last $n/2$ intervals.

³ $\mathcal{T}_k^e$ is the sample-based approximation of total cost $\mathcal{T}_k$.

⁴ $\mathcal{T}_k^{wc,RH}$ is the worst-case total cost corresponding to the excitation signal actually applied in the interval k .

The experimental total cost $\mathcal{T}_k^e$, the worst case $\mathcal{T}_k^{\text{wc,RH}}$ and the constraint $\bar{\mathcal{T}}_k$ for all the cases are reported in Fig. 4. It appears that the average over time of both $\mathcal{T}_k^e$ and $\mathcal{T}_k^{\text{wc,RH}}$ is lower when a larger number of shorter intervals are selected. The best performance is obtained in Case 1, followed by the Cases 2, 3, 4.

Note that Case 4, where only $n = 2$ intervals are considered, corresponds to classical experiment design problems made up of an identification phase (interval 1) and a control phase (interval 2). In this case, the constraint $\bar{\mathcal{T}}^h = 0.7$ is kept for a longer time than in the other cases since the length of the first interval is chosen as $N_1 = 1800$. This was done because it was not possible to satisfy the constraint $\mathcal{T}_2^{\text{wc}} \leq 0.005$ without violating $\mathcal{T}_1^{\text{wc}} \leq 0.7$ having two intervals of equal length $N = 1200$ samples.

6. Conclusions

We have presented a new procedure for the gradual update of a model-based controller. Our framework is based on an iterative control procedure consisting of successive closed-loop identification and model-based controller design steps. Tools from Experiment Design are used to design the excitation signals for all the intervals aiming to maximize a measure of the overall performance, while guaranteeing a minimum performance at all time. In the future, we would like to extend the framework to nonlinear dynamical systems.

References

- Bombois, X., Gevers, M., Scorletti, G., & Anderson, B. D. O. (2001). Robustness analysis tools for an uncertainty set obtained by prediction error identification. *Automatica*, 37(10), 1629–1636.
- Bombois, X., Scorletti, G., Gevers, M., Van den Hof, P. M. J., & Hildebrand, R. (2006). Least costly identification experiment for control. *Automatica*, 42(10), 1651–1662.
- Boyd, S., & Vandenberghe, L. (2004). *Convex optimization*. Cambridge University Press.
- Campi, M. C., Garatti, S., & Prandini, M. (2009). The scenario approach for systems and control design. *Annual Reviews in Control*, 33(2), 149–157.
- Chamanbaz, M., Dabbene, F., Tempo, R., Venkataramanan, V., & Qing-Guo, Wang (2013). Sequential randomized algorithms for sampled convex optimization. In *IEEE conference on computer aided control system design, CACSD, August* (pp. 182–187).
- De Callafon, R. A., & Van den Hof, P. M. J. (1997). Suboptimal feedback control by a scheme of iterative identification and control design. *Mathematical Modelling of Systems*, 3(1), 77–101.
- Forgione, M., Bombois, X., & Van den Hof, P. M. J. (2013). Experiment design for batch-to-batch model-based learning control. In *Proc. American control conference* (pp. 3918–3923). IEEE.
- Gevers, M. (2002). A decade of progress in iterative process control design: from theory to practice. *Journal of Process Control*, 12(4), 519–531.
- Gevers, M. (2005). Identification for control: from the early achievements to the revival of experiment design. *European Journal of Control*, 11, 1–18.
- Gevers, M., & Ljung, L. (1986). Optimal experiment designs with respect to the intended model application. *Automatica*, 22(5), 543–554.
- Hjalmarsson, H. (2009). System identification of complex and structured systems. *European Journal of Control*, 15(3), 4.
- Hjalmarsson, H., Gevers, M., & De Bruyne, F. (1996). For model-based control design, closed-loop identification gives better performance. *Automatica*, 32(12), 1659–1673.
- Jansson, J., & Hjalmarsson, H. (2005). Input design via LMIs admitting frequency-wise model specifications in confidence regions. *IEEE Transactions on Automatic Control*, 50(10), 1534–1549.
- Ljung, L. (1999). *System identification*. Wiley Online Library.
- Maciejowski, J. M., & Huzmezan, M. (1997). *Predictive control*. Springer.
- Pronzato, L., Kulcsár, C., & Walter, E. (1996). An actively adaptive control policy for linear models. *IEEE Transactions on Automatic Control*, 41(6), 855–858.
- Schrama, R. J. P. (1992). Accurate identification for control: the necessity of an iterative scheme. *IEEE Transactions on Automatic Control*, 37(7), 991–994.
- Stigter, J. D., Vries, D., & Keesman, K. J. (2006). On adaptive optimal input design: a bioreactor case study. *AIChE Journal*, 52(9), 3290–3296.
- Tse, E., & Bar-Shalom, Y. (1973). An actively adaptive control for linear systems with random parameters via the dual control approach. *IEEE Transactions on Automatic Control*, 18(2), 109–117.
- Van Den Hof, P. M. J., & Schrama, R. J. P. (1995). Identification and control—closed-loop issues. *Automatica*, 31(12), 1751–1770.

Marco Forgione was born in 1986 in Varese, Italy and he obtained his M.Sc. degree in Computer Engineering from the University of Pavia, Italy in 2009. Soon afterwards, he joined the Delft Center for Systems and Control department of the Delft University of Technology, The Netherlands as a Ph.D. candidate. He obtained his Ph.D. degree in 2014. He is currently a postdoctoral researcher in the Laboratoire Ampère of the Ecole Centrale de Lyon, France. His main research interests are iterative control, experiment design, and robust control with applications in process engineering and microelectronic system design.

Xavier Bombois was born in Brussels in 1974. He received the Electrical Engineering and Ph.D. degrees from the Université Catholique de Louvain (Belgium) in 1997 and 2000, respectively. Currently, he is an assistant professor at the Delft Center for Systems and Control (The Netherlands). His main interests are identification for control, prediction error identification, uncertainty bounding and robust control.

Paul M.J. Van den Hof received the M.Sc. and Ph.D. degrees from the Department of Electrical Engineering, Eindhoven University of Technology, Eindhoven, The Netherlands, in 1982 and 1989, respectively. In 1986 he moved to Delft University of Technology, where he was appointed as Full Professor in 1999. Since 2003, he has been founding co-director of the Delft Center for Systems and Control (DCSC), with appointments in the faculty of Mechanical, Maritime, and Materials Engineering, and the faculty of Applied Sciences. As of 2011, he holds a Full Professor position with the Electrical Engineering Department, Eindhoven University of Technology. Since 2005, he has also been a Scientific Director of the National Research and Graduate School “Dutch Institute of Systems and Control” (DISC), and National Representative of the Dutch NMO in IFAC. His research interests include issues of system identification, identification for control, and model-based control and optimization, with applications in industrial process control systems, including petroleum reservoir engineering systems, and high-tech systems.

He was General Chair of the 13th IFAC Symposium on System Identification, held in Rotterdam, The Netherlands, in 2003. He has been a member of the IFAC Council (1999–2005), and Associate Editor and Editor of *Automatica* (1992–2005). He is a Fellow of IFAC and of IEEE.