

From experiment design to closed-loop control[☆]

Håkan Hjalmarsson*

Department of Signals, Sensors and Systems, Royal Institute of Technology, S-100 44 Stockholm, Sweden

Received 9 February 2004; received in revised form 23 August 2004; accepted 15 November 2004

Available online 24 January 2005

Abstract

The links between identification and control are examined. The main trends in this research area are summarized, with particular focus on the design of low complexity controllers from a statistical perspective. It is argued that a guiding principle should be to model as well as possible before any model or controller simplifications are made as this ensures the best statistical accuracy. This does not necessarily mean that a full-order model always is necessary as well designed experiments allow for restricted complexity models to be near-optimal. Experiment design can therefore be seen as the key to successful applications. For this reason, particular attention is given to the interaction between experimental constraints and performance specifications.

© 2005 Elsevier Ltd. All rights reserved.

Keywords: Identification for control; Statistical analysis; Experiment design; Model validation

1. Introduction

Ever increasing productivity demands and environmental standards necessitate more and more advanced control methods to be employed in industry. However, such methods usually require a model of the process and modeling and system identification are expensive. Quoting (Ogunnaike, 1996):

“It is also widely recognized, however, that obtaining the process model is the single most time consuming task in the application of model-based control.”

In Hussain (1999) it is reported that three quarters of the total costs associated with advanced control projects can be attributed to modeling. It is estimated that models exist for far less than one percent of all processes in regulatory control. According to Desborough and Miller (2001), one of the few instances when the cost of dynamic modeling can

be justified is for the commissioning of model predictive controllers.

It has also been recognized that models for control pose special considerations. Again quoting (Ogunnaike, 1996):

“There is abundant evidence in industrial practice that when modeling for control is not based on criteria related to the actual end use, the results can sometimes be quite disappointing.”

Hence, efficient modeling and system identification techniques suited for industrial use and tailored for control design applications have become important enablers for industrial advances. The Panel for Future Directions in Control, (Murray, Åström, Boyd, Brockett, & Stein, 2003), has identified *automatic synthesis of control algorithms, with integrated validation and verification* as one of the major future challenges in control. Quoting (Murray et al., 2003):

“Researchers need to develop much more powerful design tools that automate the entire control design process from model development to hardware-in-the-loop simulation.”

Spurred by the recognized problems, identification for control has been one of the most active areas in system identification over the last decade. Since the joint identification and control problem shares the same elements as any engineering application where system identification is involved,

[☆] This paper was not presented at the 13th IFAC symposium on system identification, SYSID-2003, Rotterdam, NL. This paper was recommended for publication in the revised form by Guest Editors Torsten Söderström, Paul van den Hof, Bo Wahlberg and Siep Weiland.

* Tel.: +46 8 7908464; fax +46 8 7907329.

E-mail address: hakan.hjalmarsson@s3.kth.se (H. Hjalmarsson).

much work under the “umbrella” of identification for control has general applicability and it seems fair to say that no other area of identification has contributed as much to the basic understanding of system identification during the last decade.

So what are the issues? Well, to get a first hint consider the following (oversimplified) problem a control engineer might be faced with:

We have this prior knowledge of the process. You are allowed to perform a closed-loop identification experiment with the existing controller in the loop. The experiment should be as short as possible and should disrupt the process minimally. We think we would like the rise-time and settling-time to be this and that, but we are not really sure. Of course, the resulting closed-loop should be stable. If you want to use anything other than a PID-controller, you need to make a very strong case for why this is necessary.

Clearly, a useful theory should be capable of handling this type of questions. Below we will try to delineate the main issues involved.

1.1. The unforgiving nature of feedback

In many applications, performance degrades gracefully as the accuracy of the model becomes worse. However, in feedback control instability may lead to disastrous consequences. By the end of the 1980s this issue prompted significant efforts to develop identification frameworks which can accommodate for various types of prior information and produce model sets to which the true system is guaranteed to belong so that it can be checked that a designed (robust) controller at least stabilizes the system. While there is still a debate on which assumptions are relevant, this line of research has definitely put a finger on the approximate nature of system identification.

1.2. The forgiving nature of feedback

There are numerous existing successful applications of PID-control to non-linear processes which are based on simple, e.g. first-order, models identified from step response tests. Hence, it is clear that simple, very crude, models often suffice to give good or, perhaps more accurately, acceptable closed-loop performance. Behind this is the rationale for feedback control: High loop gain in a frequency band makes the closed loop system insensitive to the quality of the model and the properties of the open loop system in this frequency band, provided stability can be maintained; cf. having an integrator in the controller which gives unity steady-state gain regardless of the open loop steady-state gain.

This observation translates into the fact that one would want the model set produced by system identification to be shaped such that high performance can be obtained and also such that, if desired, additional robustness may be included with little penalty on performance. This issue relates directly

to the *design of the identification experiment*. Clearly, if the experiment ensures that all features of the system relevant for control design are present in the data, this objective can be achieved. Another characteristic feature of the control application is that *it may not a priori be known in which frequency range an accurate model is required* since this depends on performance limitations of the system, some of which typically are unknown a priori.

1.3. Compatibility requirements

The uncertainty description obtained from system identification is dictated by the model structure and the prior information used; see Section 2. It may not be directly applicable to a particular control design method. For example, prediction error methods deliver an ellipsoid in the parameter space whereas robust \mathcal{H}_∞ design with unstructured uncertainty assumes frequency by frequency bounds on the uncertainty. Thus, it may be necessary to outer-bound the uncertainty description and this should be done so as to not introduce unnecessary conservatism. Another aspect of this issue is that the order of a robust controller usually depends not only on the order of the nominal model but also on the orders of the weighting filters describing bounds on model uncertainty and performance specifications. Hence, the uncertainty description may also *influence the order of the controller*.

1.4. Summary

The discussion above can be condensed as follows. The user has a number of design variables such as experimental conditions, model structure and the performance specifications available. To be able to select these in a systematic way such that stability and performance are guaranteed involves:

- ensuring that the ‘true’ system is accounted for in the set of delivered models.
- understanding which properties of the system have to be modeled accurately and which can be treated only superficially and how this relates to the performance specifications.
- designing experiments that reveal this information.
- representing this information mathematically in a way that is not overly complex.
- adjusting the performance demands such that the design becomes robust given the limitations in modeling accuracy.

The first issue is rather delicate as it involves entities that are not verifiable, and has quite naturally led to different approaches.

1.5. Outline

We will begin this paper with a part (Sections 2–6) on general modeling principles, with particular attention to issues that are of importance for control applications. Modeling frameworks are discussed in Section 2, where we highlight the similarity of algorithms ensuing from deterministic and stochastic approaches. We also elaborate on the importance of noise prior information. Many considerations in control are best handled in the frequency domain and in Section 3 we take a closer look at model uncertainty in the frequency domain. As discussed above, control applications often allow only very approximate models to be used. We therefore, in Section 4, discuss how to identify restricted complexity models without compromising statistical accuracy. Stability and performance guarantees require the true system to be accounted for in the model set delivered by the system identification. This relates to model validation which is discussed in Section 5. In Section 6 our observations are summarized.

The remaining part of the paper is directly concerned with control related issues. Robust control and its links to system identification are issues discussed in Section 7. The problem of directly identifying a restricted complexity model useful for control design is covered in Section 8. This is followed by a section on how to model non-linear systems using linear time-invariant (LTI) models. Model free tuning methods are discussed in Section 10 while Section 11 covers experiment design issues. Model validation for control is the theme in Section 12. The paper concludes with some comments.

We will restrict attention to time-domain identification, with particular focus on the prediction error method. For frequency domain methods, the reader is referred to Pintelon and Schoukens (2001), McKelvey (2000) for general treatise of the subject, and to Chen and Gu (2000) and references therein for more control oriented deterministic approaches.

2. Information content in the data

The question of what information the noisy measurement data contains regarding the system dynamics is really at the core of system identification. One can view one part of system identification as the problem of cleaning up the data with respect to noise as well as possible. In this section we will discuss the limitations of what can be achieved in this respect.

2.1. Mathematical models

We will denote the input and output of a dynamical system by $u(t)$ and $y(t)$, respectively. For a signal $x(t)$, $t = 1, 2, \dots$, we will use the shorthand notation x^N for $\{x(t)\}_{t=1}^N$. We will often omit the time-argument from signals for ease of exposition. The *observed* output/input signal sequence that is available for the system identification will be denoted $Z^N = \{y(t), u(t)\}_{t=1}^N$.

Generally, a model for such a dynamical system will denote a sequence of mappings f_t , $t = 1, 2, \dots$ such that $y(t) = f_t(u^t, e^t)$ is the model response to the input sequence $u(t)$, $t = 1, 2, \dots$ and the unmeasurable signal sequence (“noise”) $e(t)$, $t = 1, 2, \dots$

A model is *consistent with the observations* Z^N if there exists a noise sequence e^N such that $y(t) = f_t(u^t, e^t)$, $t = 1, \dots, N$ holds for Z^N .

2.2. The concept of unfalsification

That a scientific theory may be falsified by contradicting evidence but never validated by corroborating evidence was elaborated on by the philosopher Karl Popper (Popper, 1963). In system identification, it is clear that none of the models that are consistent with Z^N can be discarded unless some prior information is available. We call the set of all such consistent models *the set of unprejudiced unfalsified models*, which we denote by $\mathcal{G}(Z^N)$. This set represents the remaining uncertainty of the system dynamics given the observed data Z^N . The information contents in the observed data corresponds exactly to the set of models that are *falsified* by the observed data. The more “informative” data is, the larger is the set of candidate models that can be falsified. The set $\mathcal{G}(Z^N)$ is of course an enormous set since it includes all infinite dimensional non-linear models which are consistent with the observations. Hence, it may appear that measurement data provides very little information.

2.3. Introducing priors

The only possible way to reduce the size of the set of unprejudiced unfalsified models is to introduce prior information, both on the system dynamics and the noise. For example, we might immediately be prepared to introduce the prior that the system is causal. We shall denote by $\overline{\mathcal{G}}(Z^N)$ the set of models consistent with data and the prior information, and we shall refer to this set as *the set of unfalsified models*.

A common approach is to introduce a parametrized model structure, e.g.

$$y = G(\theta)u + H(\theta)e + \Delta(u), \quad (1)$$

where G and H are transfer functions parametrized by $\theta \in \Theta \subset \mathbf{R}^n$, where Δ is an unstructured dynamic term and where e is a noise signal. Since e may be taken such that any model in the above structure is unfalsified, assumptions on the noise are also required. This issue is crucial for the system identification problem.

2.4. Set-membership identification

The set of unfalsified models becomes manageable for the model structure (1) by imposing that $\Delta \in S_\Delta$ and $e^N \in S_e$ for some suitably chosen sets S_Δ and S_e . Common choices

for S_e are

$$S_e = \{e^N : |e(t)| \leq c, t = 1, \dots, N\}, \quad (2)$$

$$S_e = \left\{ e^N : \sum_{t=1}^N e^2(t) \leq c \right\}. \quad (3)$$

We refer to Bai, Nagpal, and Tempo (1996) for an analysis of how these and other types of noise bounds affect the size of the set of unfalsified models.

The unstructured uncertainty Δ is often taken as LTI and examples of S_Δ are

$$S_\Delta = \{\Delta : \|\Delta\|_\infty \leq \gamma\},$$

$$S_\Delta = \left\{ \Delta = \sum_{k=1}^{\infty} \delta(k)q^{-k} : |\delta(k)| \leq C\lambda^k \right\}. \quad (4)$$

Above $\|\cdot\|_\infty$ denotes the \mathcal{H}_∞ -norm, q^{-1} is the backward shift operator and $0 < \lambda < 1$. Identification methods that employ this type of constraints have become known as set-membership methods and this has been a very active area of research during the past two decades. The set of unfalsified models often becomes very complicated and one important research topic has been to find simplified characterizations using outer- and inner-bounding techniques (Milanese & Vicino, 1991). We refer to Milanese (1998) and Milanese, Norton, Piet-Lahanier, and Walter (1996) and references therein for further details on set-membership identification.

2.5. Uncertainty model unfalsification

The combination (3) and (4) together with the model structure

$$y = G(\theta)u + H(\theta)(e + \Delta u)$$

has been studied by Kosut and co-workers in a series of papers under the label *uncertainty model unfalsification*. Based on results in Poolla, Khargonekar, Tikku, Krause, and Nagpal (1994), it is shown in Kosut (1995) and Kosut (2001) that the set of unfalsified models is empty if and only if there are no $\theta \in \Theta$ and $e^N \in S_e$ that satisfies a certain matrix inequality. For ARX-models, this is a convex feasibility problem in the unknowns θ and e^N .

As pointed out in Kosut (2001), one may compute $c_{\min}(\gamma)$, the smallest c (recall that c defines the size of the noise set S_e , cf. (2) and (3)) for which there is some unfalsified model for a given bound γ on the unstructured uncertainty Δ . The graph $\gamma \rightarrow c_{\min}(\gamma)$ is referred to as the *uncertainty trade-off curve* and gives a hint on how dynamic versus noise uncertainty may be traded-off for a set of unfalsified models.

2.6. A likelihood approach to unfalsification

One way to treat the noise is to introduce a probability measure, i.e. a measure of how likely different noise sequences are. Equipped with this measure, we can in principle order *all* models we can imagine according to how likely

the corresponding noise sequences are. The model corresponding to the most likely noise sequence is the maximum likelihood (ML) estimate.

2.6.1. Gaussian prior

To be specific, let us assume that the noise is Gaussian white noise with variance λ and denote by $e(t, M)$ the noise signal that makes model M consistent with the observations. Then the negative log-likelihood for this noise signal is given by

$$V_N(M) = \frac{N}{2} \log(2\pi) + \frac{N}{2} \log(\lambda) + \frac{1}{2\lambda} \sum_{t=1}^N e^2(t, M). \quad (5)$$

It is natural to take the set of models that corresponds to noise sequences with a likelihood higher than a given level as the set of unfalsified models.

2.6.1.1. Case 1: Known noise variance. When the noise variance λ is known, (5) then leads to the set

$$\bar{\mathcal{G}}(Z^N) = \left\{ M : \frac{1}{2\lambda} \sum_{t=1}^N e^2(t, M) \leq c \right\} \quad (6)$$

for some c . Hence we see that the stochastically motivated likelihood approach leads to *set-membership identification* with the set (3) characterizing the noise under Gaussian assumptions.

The principal difference between the likelihood approach and a deterministic set-membership approach lies in how c is chosen. In the likelihood approach one could argue that the constant c should be selected such that there is only a small probability that the true system is (erroneously) falsified. To this end, consider that $e(t, M)$ really is Gaussian white noise with variance λ , i.e. M is a model that satisfies our assumptions and should rather not be falsified. Then we have that

$$\frac{1}{\lambda} \sum_{t=1}^N e^2(t, M) \sim \chi^2(N).$$

By taking $c = \frac{1}{2}\chi_\alpha^2(N)$, where $\chi_\alpha^2(N)$ is defined by $P(X \leq \chi_\alpha^2(N)) = \alpha$ for a $\chi^2(N)$ distributed random variable X , in (6) the probability of falsifying this model will be $1 - \alpha$ (e.g. 0.1%).

In practice it is, of course, computationally infeasible to order all models and some parametrization has to be introduced. In order to focus on the noise issues, we will in the remaining part of Section 2.6 restrict attention to the following simple model

$$y(t) = \varphi^T(t)\theta + e(t) \quad (7)$$

where e denotes the noise, where φ is a vector of deterministic quantities and where $\theta \in \mathbf{R}^n$ is a vector of unknown model parameters. We will comment briefly on the case where unstructured uncertainty is present in Section 2.6.1.4.

For (7) we have that

$$e(t, \theta) = y(t) - \varphi^T(t)\theta \tag{8}$$

is the noise for the model corresponding to parameter vector θ . By completing the square, cf. Ljung and Hjalmarsson (1995),

$$\sum_{t=1}^N e^2(t, \theta) = \sum_{t=1}^N e^2(t, \hat{\theta}_N) + (\theta - \hat{\theta}_N)^T R_N (\theta - \hat{\theta}_N), \tag{9}$$

where R_N and the maximum likelihood estimate $\hat{\theta}_N$ (which in this case corresponds to the least-squares estimate) are given by

$$R_N = \sum_{t=1}^N \varphi(t)\varphi^T(t), \quad \hat{\theta}_N = R_N^{-1} \sum_{t=1}^N \varphi(t)y(t). \tag{10}$$

Hence, the set of unfalsified models is given by

$$\overline{\mathcal{G}}(Z^N) = \left\{ \theta : \frac{1}{\lambda} \sum_{t=1}^N e^2(t, \hat{\theta}_N) + (\theta - \hat{\theta}_N)^T \frac{R_N}{\lambda} (\theta - \hat{\theta}_N) \leq \chi_\alpha^2(N) \right\}. \tag{11}$$

The threshold $\chi_\alpha^2(N)$ in (11) is *data-independent*. It is, in general, possible to obtain a tighter set by letting the threshold be *data-dependent*. Consider again (9), given the observations, the first term of the right-hand side is completely known whereas the second term is known to be $\chi^2(n)$ distributed when θ is the true parameter vector. Hence, the set

$$\begin{aligned} \overline{\mathcal{G}}(Z^N) &= \left\{ \theta : \frac{1}{\lambda} \sum_{t=1}^N e^2(t, \theta) \leq \frac{1}{\lambda} \sum_{t=1}^N e^2(t, \hat{\theta}_N) + \chi_\alpha^2(n) \right\} \\ &= \left\{ \theta : (\theta - \hat{\theta}_N)^T \frac{R_N}{\lambda} (\theta - \hat{\theta}_N) \leq \chi_\alpha^2(n) \right\} \end{aligned} \tag{12}$$

is the set of models having the largest likelihoods and being such that the posterior (to the observation of the data) probability that the true system is outside this set is $1 - \alpha$.

2.6.1.2. Case 2: Unknown noise variance. When λ is unknown it can be included as an unknown parameter in the model, and, hence, the set of unfalsified models consists of pairs (θ, λ) . Now, however, the choice of the threshold for the likelihood function defined by (5) becomes somewhat problematic. It is not possible to a priori set the threshold so that the true parameters (θ, λ) belong to the set of unfalsified models with a pre-specified probability, as could be done when the noise variance was known. However, given that θ is the parameter of interest, one can set the threshold such that the probability that the true θ together with *some* λ is in the set of unfalsified models with a pre-specified probability. The argument is as follows. Let θ be arbitrary. The parameter λ that minimizes $V_N(M)$ given θ

is given by

$$\lambda(\theta) = \frac{1}{N} \sum_{t=1}^N e^2(t, \theta)$$

for which

$$V_N(M) = \frac{N}{2} \log(2\pi) + \frac{N}{2} + \frac{N}{2} \log \left\{ \frac{1}{N} \sum_{t=1}^N e^2(t, \theta) \right\}.$$

This implies that the parameters θ that will belong to the set of unfalsified models, i.e. the set of unfalsified parameters θ , is given by

$$\overline{\mathcal{G}}(Z^N) = \left\{ \theta : \sum_{t=1}^N e^2(t, \theta) \leq c \right\}. \tag{13}$$

To ensure that the true θ (if there is such a parameter) is in this set with probability α , the constant c should be taken as

$$c = \left(\frac{n}{N-n} F_\alpha(n, N-n) + 1 \right) \sum_{t=1}^N e^2(t, \hat{\theta}_N),$$

where the constant $F_\alpha(n, N-n)$ is defined by $P(X \leq F_\alpha(n, N-n)) = \alpha$ where X is $F(n, N-n)$ distributed. This gives that (13) can be written as

$$\begin{aligned} \overline{\mathcal{G}}(Z^N) &= \left\{ \theta : \frac{1}{n} (\theta - \hat{\theta}_N)^T \frac{R_N}{1/(N-n) \sum_{t=1}^N e^2(t, \hat{\theta}_N)} \right. \\ &\quad \left. \times (\theta - \hat{\theta}_N) \leq F_\alpha(n, N-n) \right\}. \end{aligned} \tag{14}$$

This follows by way of (9). Now,

$$\frac{1}{n} (\theta - \hat{\theta}_N)^T \frac{R_N}{\frac{1}{N-n} \sum_{t=1}^N e^2(t, \hat{\theta}_N)} (\theta - \hat{\theta}_N)$$

is $F(n, N-n)$ distributed when θ is the true parameter vector. Thus, the set in (14) contains the true θ with probability α as desired if there is such a “true” parameter.

2.6.1.3. Relation to the prediction error method. The set (14) corresponds exactly to a confidence region for the true parameter when the least-squares estimate for the model (7) is used (Ljung, 1999b). The probability that the true parameter is outside this set is $1 - \alpha$. Hence, we can interpret least-squares identification of linear regression models under Gaussian assumptions as a likelihood-based unfalsification method. Since the least-squares estimate is identical to the maximum-likelihood (ML) estimate under these assumptions, the confidence region is also the smallest possible. That there is a close connection to the ML-method is of course natural as the likelihood approach to model unfalsification is based on the likelihood function.

For model structures beyond the linear regression type, the equivalence between likelihood-based unfalsification and

the confidence region for the prediction error method (PEM) and the ML-method can be shown to hold asymptotically. For brevity, we discuss this for the case when λ is unknown only.

Consider the LTI model structure

$$y(t) = G(q, \theta)u(t) + H(q, \theta)e(t), \tag{15}$$

where H is monic, stable and minimum phase, which is parametrized in terms of $\theta \in \Theta \subset \mathbf{R}^n$. The signal e represents noise. The noise given that the parameter vector is θ is given by

$$e(t, \theta) = H^{-1}(q, \theta)(y(t) - G(q, \theta)u(t)). \tag{16}$$

This quantity is also known as the prediction error (Ljung, 1999b) since it is the error in the one-step ahead predictor of the output when the model corresponding to the parameter vector θ is used.

The prediction error estimate with quadratic criterion is defined as

$$\hat{\theta}_N = \arg \min_{\theta \in \Theta} J_N(\theta), \quad J_N(\theta) = \frac{1}{N} \sum_{t=1}^N e^2(t, \theta).$$

Neglecting transients, this is also the maximum-likelihood estimate when the distribution is Gaussian and hence it minimizes the negative log-likelihood (5) within the model structure.

Under weak assumptions (Ljung, 1999b), it holds that

$$\lim_{N \rightarrow \infty} \hat{\theta}_N = \theta^* \triangleq \arg \min_{\theta \in \Theta \subset \mathbf{R}^n} \lim_{N \rightarrow \infty} \mathbf{E}\{J_N(\theta)\} \text{ w.p.1.} \tag{17}$$

It holds also that the prediction error estimate $\hat{\theta}_N$ converges in law to a normally distributed random variable

$$\sqrt{N}(\hat{\theta}_N - \theta^*) \xrightarrow{\mathcal{D}} \mathcal{N}(0, P) \quad \text{as } N \rightarrow \infty \tag{18}$$

for some covariance matrix P .

Suppose now that the true system is in the model set (notice that this is an assumption that we have so far not used in Section 2), i.e.

$$y(t) = G(q, \theta^\circ)u(t) + H(q, \theta^\circ)e_\circ(t), \tag{19}$$

where e_\circ is (a realization of) white noise, for some $\theta^\circ \in \Theta$. Then $\theta^* = \theta^\circ$ (under suitable identifiability conditions (Ljung, 1999b)) and the expression for the covariance matrix P is given by

$$P = \lambda R^{-1}, \quad \text{where } R = \mathbf{E}\{\psi(t, \theta^*)\psi^T(t, \theta^*)\}, \tag{20}$$

with $\psi(t, \theta) = -de(t, \theta)/d\theta$.

With

$$R_N = N \cdot R = N \cdot \mathbf{E}\{\psi(t, \theta^*)\psi^T(t, \theta^*)\}, \tag{21}$$

it follows from (18) that as $N \rightarrow \infty$,

$$(\theta^\circ - \hat{\theta}_N)^T \frac{R_N}{1/N \sum_{t=1}^N e^2(t, \hat{\theta}_N)} (\theta^\circ - \hat{\theta}_N) \xrightarrow{\mathcal{D}} \chi^2(n) \tag{22}$$

and, hence,

$$\overline{\mathcal{G}}(Z^N) = \left\{ \theta : (\theta - \hat{\theta}_N)^T \frac{R_N}{1/N \sum_{t=1}^N e^2(t, \hat{\theta}_N)} (\theta - \hat{\theta}_N) \leq \chi_\alpha^2(n) \right\} \tag{23}$$

is a confidence region which asymptotically includes the true parameter θ° with probability α .

Let us now characterize the set of unfalsified models in the likelihood-based approach given that (19) holds. This is in general a difficult task, due to that θ appears non-linearly in the prediction error. However, for small enough bound c , all models in this set will have to have parameters θ close to the prediction error estimate $\hat{\theta}_N$. In this situation, a second order Taylor expansion gives

$$\sum_{t=1}^N e^2(t, \theta) \approx \sum_{t=1}^N e^2(t, \hat{\theta}_N) + (\theta - \hat{\theta}_N)^T R_N (\theta - \hat{\theta}_N) \tag{24}$$

where R_N is given by (21). Compare this with (9). From (22) we have that, suitably normalized, the second term of the right-hand side of (24) is approximately $\chi^2(n)$ distributed for large N . Using these observations, we can now proceed as in Section 2.6.1.2 to obtain that the set of models corresponding to the most likely noise sequences and which is large enough that the probability that the true system belongs to the set is α , can be described approximately by set (23).

For a recent study on conditions for the asymptotic prediction error theory to be valid, see Bittanti, Campi, and Garatti (2002). For recent results on finite sample properties we refer to Weyer and Campi (2002), Campi and Weyer (2002). Similar conclusions can also be made when the noise is not Gaussian but we will not pursue this topic further.

2.6.1.4. Case 3: Unstructured uncertainty present. For model structures of type (1) one can also pursue a likelihood approach. The set of unfalsified models will correspond to a set of parametric uncertainty, cf. (23), together with the unstructured uncertainty. However, the parametric uncertainty set will be larger compared to (23) since it corresponds to all parameters for which there exists a $\Delta \in S_\Delta$ such that the likelihood of the corresponding noise sequence is acceptable.

2.6.2. Uniform prior

The assumption that the noise distribution has support $[-c, c]$ leads to a set of unfalsified models of the type (2) (Ninness & Goodwin, 1995). In this set, all models are equally likely.

The model set is in this case a polyhedron which can be approximated by an ellipsoid, i.e. with a set of the type (23). The problem of determining the outerbounding ellipsoid of

minimal volume is a convex optimization problem (Pronzato & Walter, 1994).

2.7. Stochastic embedding

Just as e may be modeled in a stochastic framework, the unstructured uncertainty Δ in the model structure (1) may be modeled in a stochastic framework. This leads to what is known as the stochastic embedding approach (Goodwin, Gevers, & Ninness, 1992; Goodwin, Braslavsky, & Seron, 2002).

2.8. On the value of noise priors

Let us return to the likelihood approach in Section 2.6. We started out in Section 2.6.1.1 by assuming that the noise variance λ was known. This resulted in the set (11). Notice, that for a given model structure there may be no model at all that belongs to this set—a very powerful result! Poor models will thus be rejected. Unfortunately, overly complex models will not be falsified and thus overmodelling can be a problem. One possibility of dealing with this problem is to use Occam's razor:

Pick a model with as low complexity as possible in the set of unfalsified models.

There is often no reason to favor any other model. In Section 8.11, however, we will see that there may be reasons for other choices.

In Section 2.6.1.1 we also observed that we could shrink the size of the set by looking at the data before the threshold was selected. This resulted in the set (12). Notice that for *any* given model structure, the set of unfalsified models will now be non-empty. By making the threshold data dependent we have gained in accuracy but the price paid is that the objectivity of the criterion used in (11) has been lost. When the noise variance is unknown, there is no objective criterion. As for (12), the set of unfalsified models (14) will be prejudiced on the model structure. Different model structures now have to be compared against each other in order to find the “right” structure. This leads us to model validation which we will discuss in Section 5.

To conclude, we have argued that knowledge of the noise characteristics is extremely valuable. As pointed out, e.g. in Pintelon and Schoukens (2001) and Ljung (1999a), the noise sequence itself can be estimated if periodic inputs are used, cf. the case when the input is zero, then the output equals the noise.

2.9. Summary

In this section, we have presented identification as a way of producing sets of unfalsified models and illustrated that both deterministic and stochastic modeling paradigms fit into this framework. Interestingly, it seems as if many of these different approaches result in similar unfalsified model sets, see Reinelt, Garulli, and Ljung (2002).

For stochastic models, we discussed the use of the likelihood function as a criterion for ordering models. This led to the definition of the set of unfalsified models as the set of models for which the likelihood function is above a certain threshold. The threshold was determined so that the probability that the true system is outside this set, i.e. is falsified, is smaller than some given (small) number. It was shown that this set is equivalent to a standard confidence region for the prediction error method. The likelihood-based approach to obtain a set of unfalsified models is thus just another way of viewing the prediction error method. However, an important insight is that all models in a set of unfalsified models have one thing in common: they all have the likelihood function larger than a certain threshold. As we will see in Section 4.4, this observation will be instrumental when discussing statistically accurate models of restricted complexity, i.e. models having lower order than the true system.

We noted in Section 2.8 that it is difficult to quantify the information contents in data when the model structure is uncertain. However, one recurring theme in this paper will be that it is possible to *ensure that the information required for our particular application can be made available by proper experiment design.*

There is an on-going healthy cross-fertilization of ideas between deterministic and stochastic approaches, see e.g. Ljung and Hjalmarsson (1995), Tjärnström and Garulli (2002), Partington and Mäkilä (1999), Hakvoort, Schrama, and Van den Hof (1994), de Vries and Van den Hof (1995), Milanese and Taragna (1999), Tempo, Bai, and Dabbene (1997), Bai, Ye, and Tempo (1999). For excellent overviews of different modeling frameworks we refer to Mäkilä, Partington, and Gustafsson (1995) and Ninness and Goodwin (1995).

3. Frequency domain characterization of the set of unfalsified models

In the previous section the discussion was concerned with the characterization of the set of unfalsified models in the parameter space. However, as we will see in Section 7, when the model set is to be used for control design it is of interest to characterize the model set in the frequency domain. We will in this section focus on the prediction error method.

We will assume that the true system is given by

$$y(t) = G_o(q)u(t) + H_o(q)e_o(t), \quad (25)$$

where e_o is white noise with variance λ_0 and where H_o is stable, monic and minimum phase. We will denote the power spectral densities of u and $H_o e_o$ by Φ_u and Φ_v , respectively. The model structure is given by (15) with $\theta \in \mathbf{R}^n$.

3.1. The bias error

For the PEM using the LTI model structure (15), the limit estimate (17) can be characterized indirectly in the frequency domain using Parseval's formula. When the true system,

given by (25), is stable and operating in open loop, and when G and H are *independently parametrized*, i.e. $\theta = [\eta^T, \xi^T]^T$ and $G = G(\eta)$; $H = H(\xi)$, the limit estimate is defined by

$$\eta^* = \arg \min_{\eta}$$

$$\int_{-\pi}^{\pi} |G_{\circ}(e^{j\omega}) - G(e^{j\omega}, \eta)|^2 \frac{\Phi_u(\omega)}{|H(e^{j\omega}, \xi^*)|^2} d\omega \quad (26)$$

where ξ^* is the limit estimate of ξ . The limit estimate $G(\eta^*)$ thus minimizes a weighted \mathcal{H}_2 -norm of the error between the model and the true dynamics. It is thus in general not possible to guarantee frequency wise error bounds on the bias error. This may be critical in some applications such as control design. We shall pursue this issue in Section 4.4.

When (25) holds and the system is operating in closed-loop with a stabilizing controller C and an external excitation $r(t)$ (the reference signal) such that the input signal is given by

$$u(t) = C(q)(r(t) - y(t))$$

it holds that $\theta^* = \arg \min_{\theta \in \Theta}$

$$\int_{-\pi}^{\pi} \left\{ |G_{\circ}(e^{j\omega}) - G(e^{j\omega}, \theta)|^2 |CS(G_{\circ}, C)|^2 \Phi_r + \frac{|S(G_{\circ}, C)|^2}{|S(G(\theta), C)|^2} \Phi_v \right\} \frac{1}{|H(e^{j\omega}, \theta)|^2} d\omega, \quad (27)$$

where $S(G, C) = 1/(1 + GC)$. An alternative expression which characterizes the bias introduced by an erroneous noise model can be found in Forsell and Ljung (1999).

3.2. Variance of frequency function estimates

A simple characterization of the uncertainty in the frequency domain is in terms of the variance of the frequency function estimate $\hat{G}_N(e^{j\omega}) \triangleq G(e^{j\omega}, \hat{\theta}_N)$:

$$\text{Var}(\hat{G}_N(e^{j\omega})) \triangleq E[|G(e^{j\omega}, \hat{\theta}_N) - E[G(e^{j\omega}, \hat{\theta}_N)]|^2].$$

Introducing

$$\kappa_{n,N}(\omega) = \text{Var}(\hat{G}_N(e^{j\omega})) \frac{N \cdot \Phi_u(\omega)}{\Phi_v(\omega)},$$

where we for ease of notation have chosen to indicate only the dependence of κ on n and N explicitly, one can write

$$\text{Var}(\hat{G}_N(e^{j\omega})) = \kappa_{n,N}(\omega) \frac{\Phi_v(\omega)}{N \cdot \Phi_u(\omega)}, \quad (28)$$

whenever $\text{Var}(\hat{G}_N(e^{j\omega}))$ is well-defined. The reason for this factorization of $\text{Var}(\hat{G}_N(e^{j\omega}))$ is that under certain assumptions exact expressions or accurate approximations to $\kappa_{n,N}(\omega)$ exists.

The key to such results is a first order Taylor approximation of the variance¹

$$N \cdot \text{Var}(\hat{G}_N(e^{j\omega})) \approx \left(\frac{dG(e^{j\omega}, \theta)}{d\theta} \right)^* P \frac{dG(e^{j\omega}, \theta)}{d\theta}. \quad (29)$$

where the derivatives and P are evaluated at $\theta = \theta^*$. This expression is exact as $N \rightarrow \infty$ when the model structure is linearly parameterized

$$G(q, \theta) = \Gamma^T(q)\theta, \quad (30)$$

in which case $dG/d\theta = \Gamma$.

For general parametrizations some caution is necessary when using (29) as it is an approximation, see, e.g., Vuerinckx, Pintelon, Schoukens, and Rolain (2001) where this is illustrated for confidence bounds on estimated zeros.

In the mid-eighties, the following result was derived (presented here for the case of open loop operation) (Ljung & Yuan, 1985; Ljung, 1985)

$$\lim_{m \rightarrow \infty} \lim_{N \rightarrow \infty} \frac{N}{m} \text{Var}(\hat{G}_N(e^{j\omega})) = \frac{\Phi_v(\omega)}{\Phi_u(\omega)}, \quad (31)$$

where m is the model order. Result (31) implies that

$$\kappa_{n,N}(\omega) \approx m \quad (32)$$

for large enough model order m and number of samples N .

A complicating factor in the derivation of (31) is that for certain model structures such as Box–Jenkins and output-error, pole-zero cancellations occur when the model order exceeds the underlying true system order. In order to ensure a well defined limit estimate, the cost function has to be regularized in the analysis. In Ninness and Hjalmarsson (2004a) the effect of this regularization on the variance is studied and it is shown that as the model order tends to infinity, it is the regularization only that determines the variance. In fact, the result (31) holds for these types of model structures only when the regularization term is of the form $\delta \| \theta - 0 \|^2$ for small $\delta > 0$. Choosing another regularization point than the origin will result in another variance.

An approximation of $\kappa_{n,N}(\omega)$ with, in many cases, improved accuracy was proposed in Ninness, Hjalmarsson, and Gustafsson (1999). In Xie and Ljung (2001) an expression for $\kappa_n(\omega) \triangleq \lim_{N \rightarrow \infty} \kappa_{n,N}(\omega)$ was derived for the case of a model with fixed denominator and fixed moving average noise model excited by an auto-regressive (AR) input. This represented a major step forward as no asymptotics in the number of parameters (or model order) is involved, thus avoiding the need for regularization. In Ninness and Hjalmarsson (2004c) this result was generalized and for the Box–Jenkins case of independently parametrized dynamics and noise models the result reads as follows.

¹ x^* denotes the complex conjugate transpose.

Proposition 3.1. *Suppose that the true system is operating in open loop and given by*

$$y(t) = \frac{B_o(q)}{A_o(q)} u(t) + v(t),$$

where $v(t) = H_o(q)e_0(t)$ for some white noise sequence $e_0(t)$. Assume that the system is in the model set. Let $G(q, \theta) = q^{-k}B(q)/A(q)$ with m_b parameters in $B(q)$ and m_a parameters in $A(q)$.

Under the condition that

$$A_{\dagger} \triangleq A_o^2 H_o / \Phi_u^{1/2}, \tag{33}$$

where $\Phi_u^{1/2}$ is the stable minimum-phase spectral factor of the input spectrum, is a polynomial in z^{-1} of degree at most $m_a + m_b$, it holds that

$$\lim_{N \rightarrow \infty} N \cdot \text{Var}(\hat{G}_N(e^{j\omega})) = \kappa_n(\omega) \frac{\Phi_v(\omega)}{\Phi_u(\omega)},$$

where

$$\kappa_n(\omega) \triangleq \sum_{k=1}^{m_a+m_b} \frac{1 - |\xi_k|^2}{|e^{j\omega} - \xi_k|^2}, \tag{34}$$

where $\xi_k, k=1, \dots, m_a+m_b$, are the zeros of $z^{m_a+m_b} A_{\dagger}(z)$.

Comparing (34) with (32) we see that the factor m (the model order) is replaced by a frequency dependent factor $\kappa_n(\omega)$ which is a function of the poles of A_{\dagger} .

Remark 1. Notice that $\int_{-\pi}^{\pi} \kappa_n(\omega) d\omega = m_a + m_b$ and hence that there is a “water-bed effect” in that a small κ_n in some frequency region has to be compensated for by high κ_n in some other region.

Remark 2. Notice that the result holds, e.g., if the noise model is of MA-type and the input spectrum is of AR-type and m_b is sufficiently large.

Remark 3. The result seems to hold approximately with good accuracy also for cases where the system is not in a Box–Jenkins model set but when this model structure is flexible enough that the bias error for both the system dynamics and the noise spectrum is small. A_o and H_o should then be replaced by the corresponding models in the limit $N \rightarrow \infty$. Notice that this implies a different condition on the orders m_a and m_b than if the correct model structure was used.

Remark 4. The result seems to hold approximately with good accuracy also when the input can be well approximated by an AR-process. The spectral factor corresponding to the approximating process should then be used instead of $\Phi_u^{1/2}$ in (33). Notice that this implies a different condition on the model orders m_a and m_b than if an input of AR-type was used.

Remark 5. In closed-loop identification, the variance of the frequency function estimate is independent of whether

direct, indirect or joint input–output identification is used when (32) is valid (Gevers, Ljung, & Van den Hof, 2001). As noted in Van den Hof (1998), the variances for these different methods are likely to be different for finite model orders. Using Proposition 3.1, it is shown in Ninness and Hjalmarsson (2004b) that the accuracy of different methods may indeed be very different. This has also been supported by considerations in the parameter domain (Forsell & Ljung, 1999).

Remark 6. For finite impulse response (FIR) models, an exact expression for $\kappa_{n,N}(\omega)$ can be derived when the number of spectral lines in the input equals the model order (Hjalmarsson & Ninness, 2004).

3.3. The gain error

Assume now that the true system (25) is in the model set, i.e. (19) holds. Assume further that the confidence region for the prediction error method is given by $\overline{\mathcal{G}}(Z^N)$ defined in (23). As we will see, for control design purposes it is often sufficient to be able to characterize the gain error

$$|\hat{G}_N(e^{j\omega}) - G_o(e^{j\omega})|.$$

In going from parametric uncertainty to frequency domain uncertainty the following lemma, which is a special case of Lemma 3.1 in Wahlberg and Ljung (1992) and Theorem 1 in Bombois, Anderson, and Gevers (2000a), is useful.

Lemma 3.1. *For $x, z \in \mathbf{R}^n, 0 < Q \in \mathbf{R}^{n \times n}$, it holds*

$$x^* Q^{-1} x \leq c \quad \Rightarrow \quad |z^T x|^2 \leq c z^* Q z. \tag{35}$$

When the transfer function is linearly parametrized (30), Lemma 3.1 applied to the inequality in (23) gives the following upper bound on the gain error

$$|\hat{G}_N(e^{j\omega}) - G_o(e^{j\omega})|^2 \leq \chi_{\alpha}^2(n) \frac{1}{N} \Gamma^*(e^{j\omega}) P \Gamma(e^{j\omega}), \quad \forall \omega$$

which holds with probability α . In view of (29), the inequality above is equivalent to

$$|\hat{G}_N(e^{j\omega}) - G_o(e^{j\omega})| \lesssim \sqrt{\chi_{\alpha}^2(n) \text{Var}(\hat{G}_N(e^{j\omega}))},$$

Using (28), we have the bound

$$|\hat{G}_N(e^{j\omega}) - G_o(e^{j\omega})| \lesssim \sqrt{\chi_{\alpha}^2(n) \kappa_{n,N}(\omega) \frac{\Phi_v(\omega)}{N \cdot \Phi_u(\omega)}} \tag{36}$$

of the gain error which holds with at least $\alpha \times 100\%$ probability. Because of its relative simplicity and the explicit appearance of interesting quantities such as input and noise spectra in the expression, we will use (36) as the generic description of confidence regions for the gain error in the frequency domain. See Bombois et al. (2000a) for an insightful discussion on how confidence bounds in the parameter domain and the frequency domain relate.




Fig. 1. Different possibilities of mapping data and prior information into a controller of reduced complexity.

We remark that for linearly parameterized models, the bound (36) is conservative, but at most a factor $\sqrt{2}$. A tight overbound can easily be derived. For non-linearly parametrized models, the largest bound on the gain error can be computed exactly using LMIs (Bombois, Gevers, & Scorletti, 2000b), see also Jansson and Hjalmarsson (2004a) for computation of the maximum bound over the frequency axis.

4. A statistical view on restricted complexity modeling

Models of restricted complexity are often adequate in many applications. In process control for example, first order models with an additional dead-time are often sufficient even though the true process is much more complicated. Such models can be obtained in (at least) two principally different ways: (1) Direct estimation of a restricted complexity model or (2) Estimation of a full-order model followed by model reduction. It is here of interest to know if one method is to be preferred over the other. By full-order model we here mean a model which is able to capture the true system behavior. In practice full order-models do not exist and we will spend quite some effort in this, and the next, section discussing this.

In many applications, it is not the model itself that is of interest but some quantity derived from the model. For various reasons, it is often desirable to limit the mathematical complexity of this quantity. In control design, e.g., it is the designed controller that is of interest. Fig. 1 illustrates various ways of obtaining a restricted complexity controller via identification. The same question as for the case of estimating restricted complexity models arises: Is one of the paths better than the others?

In this section we will discuss these issues from a statistical perspective.

4.1. Statistical advantages of biased models

From a statistical perspective, approximate modeling is usually motivated by examples such as the following.

Example 4.1. Consider the following high-order FIR system

$$y(t) = \sum_{k=1}^n g_k^\circ u(t - k) + e_o(t), \tag{37}$$

where $e_o(t)$ is white Gaussian noise with variance λ_e , where the order n is very large and where the input u is white Gaussian noise with variance λ_u .

Suppose that one is interested in estimating the static gain $G_o(e^{j0}) = \sum_{k=1}^n g_k^\circ$ of the system. In the ML-approach one would then use a model structure of the same type as (37) and estimate $\theta = [g_1, \dots, g_n]^T$ using least squares. The covariance matrix of $\hat{\theta}_N$ is approximately $\lambda_e/(N\lambda_u)\mathbf{I}$, where \mathbf{I} denotes the identity matrix, and hence the variance of the estimated static gain $\hat{G}(e^{j0}) = \sum_{k=1}^n \hat{g}_k$ is approximately given by

$$\frac{n\lambda_e}{N\lambda_u}.$$

Since the estimate is unbiased, the mean-square error (MSE) $E[|\hat{G}(e^{j0}) - G_o(e^{j0})|^2]$ will be the same as the variance error. We see that due to the high system order n , the uncertainty can be significant even if the input power is large.

This observation, naturally, prompts the idea that a (slightly) biased estimate of the transfer function may give an estimate of the static gain which is better. For example, using the model structure

$$y(t) = \eta u(t - 1) + e(t), \tag{38}$$

for which the static gain estimate $\hat{G}(e^{j0})$ is identical to the estimate of η , will have a mean-square error which approximately is given by

$$\frac{\lambda_e}{N\lambda_u} + \frac{\sum_{k=2}^n |g_k^\circ|^2}{N} + \left| \sum_{k=2}^n g_k^\circ \right|^2. \tag{39}$$

The first term is the variance of the parameter estimate caused by the noise e . The second term is the variance of the parameter estimate caused by the unmodeled dynamics. The last term is the bias error due to the unmodeled dynamics. The MSE of the static gain for this biased model is significantly lower than for the ML estimate if only g_1° contributes significantly to the steady state gain!

The above example indicates that the ML-approach may be unsuitable when only approximate models are required for highly complex systems. However, the issue is a bit more subtle than at first glance.

4.2. A separation principle

Let $\hat{\theta}_{ML}$ be the ML-estimate of $\theta \in \Theta \in \mathbf{R}^n$ and let $f : \Theta \rightarrow \Omega \subset \mathbf{R}^m$ with $m \leq n$. It then holds that $f(\hat{\theta}_{ML})$ is the ML-estimate of $f(\theta)$. This is the so called invariance principle for ML-estimation (Zehna, 1966) (Theorem 5.1.1 in Zacks, 1971). Hence, it follows under very general conditions on f that if $\hat{\theta}_{ML}$ is asymptotically efficient, i.e. it is consistent and its asymptotic covariance matrix reaches the Cramér–Rao lower limit (Ljung, 1999b), then $f(\hat{\theta}_{ML})$

is also asymptotically efficient. In our context this provides us with a useful principle:

The estimator of some system dependent quantity that

- (i) first estimates a full-order model using an asymptotically efficient ML estimator, and then
- (ii) uses the full-order system estimate obtained in (i), as if it is the true system, to estimate the desired quantity

is an asymptotically efficient estimator of this quantity under general conditions.

The invariance principle can thus be seen as a separation principle where the estimation problem is separated from the application dependent part of the problem. We illustrate this using Example 4.1.

Example 4.2 (Example 4.1 continued). Given the ML-estimate $\hat{\theta}_N = [\hat{g}_1, \dots, \hat{g}_n]^T$ of the full-order model in Example 4.1, one may take $y(t) = \hat{g}_1 u(t - 1)$ as model, cf. (38). This will result in a biased estimate of the static gain with the MSE approximately given by

$$\frac{\lambda_e}{N\lambda_u} + \left| \sum_{k=2}^n g_k^\circ \right|^2. \tag{40}$$

This expression is the same as the MSE (39) of the static gain for the biased estimate in Example 4.1, except that the middle term in (39) is missing. This term is the variance contribution from the unmodeled dynamics. Hence, by first using a full-order model, inflation of the variance due to unmodeled dynamics can be avoided.

4.3. Applications of the separation principle

There are many applications of the separation principle presented in Section 4.2, and below some of these will be discussed.

4.3.1. Model reduction

Suppose that it is known that the true system G_o belongs to some model structure parametrized by $\theta \in \Theta$ but that the desired quantity is a consistent estimate of the frequency function minimizing

$$\int_{-\pi}^{\pi} |G_o(e^{j\omega}) - G(e^{j\omega}, \eta)|^2 \Phi_u(\omega) d\omega, \tag{41}$$

where $G(q, \eta)$ is a low order model parametrized by η . From the open loop bias expression (26) it is clear that one way of doing this is to use an output error structure

$$y(t) = G(q, \eta)u(t) + e(t),$$

and directly estimate η . For FIR-models, it is shown in Tjörnström and Ljung (2002) that this procedure leads to a higher variance as compared to first identifying a full-order

FIR-model $G(e^{j\omega}, \hat{\theta}_N)$, say, and then performing model reduction by minimizing

$$\int_{-\pi}^{\pi} |G(e^{j\omega}, \hat{\theta}_N) - G(e^{j\omega}, \eta)|^2 \Phi_u(\omega) d\omega \tag{42}$$

with respect to η . In Tjörnström (2002) this result is extended to the case when both the true system and the model structure are of output error type. The results are proved by explicitly computing, and comparing, the asymptotic covariance matrices for the two estimates.

The same results can be obtained by appealing to the separation principle in Section 4.2. It is in fact possible to extend the result slightly: Suppose that the true system is not of output error type (e.g. of Box–Jenkins type). Then it is optimal to first estimate a full-order model and then to perform the model reduction as in (42). Directly estimating an output error model, which also in this case asymptotically minimizes (41), can never give better statistical accuracy.

4.3.2. Simulation

In Zhu (2000) identification for simulation is considered. It is shown that modeling the spectrum of the noise is better than ignoring it, even though simulation does not require a noise model. As in Tjörnström (2002) this is proved by comparing the covariance matrices of the estimated parameters.

This result also follows directly from the separation principle.

4.3.3. Identification for control

The separation principle is also useful in identification for control problems and indicates that, from an accuracy point of view, no matter what the ultimate objective is, be it modeling to tune a simple PID-controller or modeling suitable for high performance control, one should always *first try to model as well as possible*. After that, any simplifications can be performed without jeopardizing the statistical accuracy. Hence, returning to Fig. 1, taking the lower path should be avoided if accuracy is a concern. We also conclude that going from a full-order model directly to a low order controller or via a high order controller, will not significantly affect the statistical accuracy. However, there may be other reasons for taking one path or another. Some of the paths may, e.g., be computationally infeasible, cf. Codrons, Bendotti, Falinower, and Gevers (1999). We shall return to the application of the separation principle in control applications in Section 8.

4.3.4. Model validation

The separation principle also applies in model validation. We shall, however, defer this discussion to Section 5.4.

4.4. Near-optimal restricted complexity models

The issue of biased modeling versus full-order modeling has yet another twist. Let us return to Example 4.1 once more but consider now another input signal.

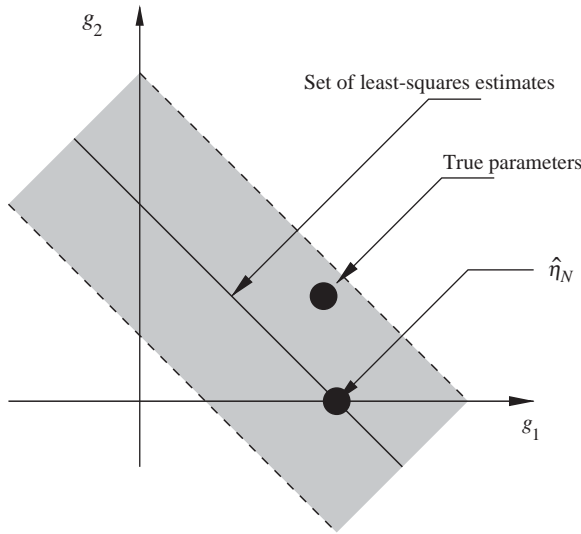


Fig. 2. Example of uncertainty region for the least-squares estimate in Example 4.3 when $n = 2$. For further explanations see text.

Example 4.3 (*Example 4.1 continued*). Suppose that the allowed input power $E[u^2(t)]$ is bounded by the constant λ_u . Then, clearly a constant input with amplitude $\sqrt{\lambda_u}$ is optimal for estimating the static gain and even though the ML criterion for the impulse response coefficients will be singular, it is easy to show that the estimate of the static gain will be well-defined and have variance approximately equal to $\lambda_e/(N\lambda_u)$, which is lower than, e.g., the minimal variance (40) for a white input. But now, the same accuracy is obtained with the simple model (38) since the unmodeled dynamics do not influence the accuracy of the estimate of η ; in fact, it is accounted for by this estimate which is now an unbiased estimate of the static gain!

The example above suggests that a judicious choice of input may allow restricted complexity models to be optimal, or near-optimal, see also Hildebrand and Gevers (2003b) for further insights. So what property of the identification problem in Example 4.3 is it that allows the biased estimate to be optimal? To answer that question, consider, for simplicity, the case when the number of impulse response coefficients $n = 2$. Since the input is not persistently exciting of sufficiently high order, the least-squares criterion will be minimized by a set of parameter vectors. This set of least-squares estimates is the solid line in Fig. 2. Also shown in this figure is a confidence region for the true parameter (the shaded region in the figure) which in this case is an ellipsoid that has degenerated into an infinite strip due to the poorly exciting input. The least squares estimate of the first order model is given by the point $\hat{\eta}_N$ in the figure. It lies in the set of least-squares estimates and is thus optimal. It may seem as a trivial observation but the conclusion is thus that if the identification experiment is performed such that the reduced order estimate belongs to the set of full-order least-squares estimates, then it will be optimal.




Fig. 3. Example of ellipsoidal uncertainty region for the least-squares estimate in Example 4.3 when $n = 2$. For further explanations see the text.

Suppose now instead that the input is such that there exists a unique least-squares estimate $\hat{\theta}_N$. The confidence region is then an ellipsoid and let us assume that it is given by the solid curve in Fig. 3. This region happens to contain a line segment of the g_1 -axis (the thick horizontal line segment in the figure). This means that any first order model with a parameter value on this segment will be as good a candidate as any second order model inside the confidence region. It also means that such a first order model will possess all the properties that models in this confidence region have. In particular, the distance to the true system in various metrics can be upper bounded. For the gain error, e.g., the triangle inequality gives

$$\begin{aligned}
 & |G(e^{j\omega}, \eta) - G_o(e^{j\omega})| \\
 & \leq |G(e^{j\omega}, \eta) - G(e^{j\omega}, \hat{\theta}_N)| + |G(e^{j\omega}, \hat{\theta}_N) - G_o(e^{j\omega})| \\
 & \leq 2\sqrt{\chi_x^2(n)\kappa_{n,N}(\omega)} \frac{\Phi_v(\omega)}{N \cdot \Phi_u(\omega)} \tag{43}
 \end{aligned}$$

for any η on this line segment. The second inequality follows from (36) in Section 3.3 (see Section 3 for the notation). This means that frequency-by-frequency error bounds can be obtained for restricted complexity models of this type. Compare this with the general \mathcal{H}_2 -norm characterizations (26) and (27) which do not allow such an error quantification.

Now, the question remains as to whether it is possible to directly identify a first order model such that it lies on the aforementioned line segment. To this end, recall the important conclusion from Section 2.6 that the ellipsoidal confidence bound actually is a level set for the least-squares criterion, cf. (9). Hence, suitable low order models are obtained by making the least-squares criterion small. In particular the least-squares estimate for the first order model is suitable. This estimate is marked as $\hat{\eta}_N$ in Fig. 3. It lies on the boundary of a scaled version of the confidence

ellipsoid, centered around the full-order least-squares estimate (marked by $\hat{\theta}_N$ in the figure). It is the ellipsoid of this type which tangents the g_1 -axis. In other words, this ellipsoid is the smallest level set of the least-squares criterion which includes a point of the g_1 -axis. We thus conclude that simple least-squares estimation of a first order model will in this case give us a model which is inside the confidence region for the full-order model and, hence, which is such that, e.g., the error bound (43) applies.

The generalization to other estimation problems is straightforward: Consider an identified full-order model with the (approximate) confidence region for the true parameter vector described by (23). If there is a point in the confidence region where some elements of θ are zero, then these parameters can be omitted in the estimation still giving an estimate which is inside the confidence region of the full-order model. We will call such models, *near-optimal reduced complexity models*. The norm of the error between this estimate and the true system parameter vector is at most a factor 2 of the norm of the error between the full-order estimate and the true system parameter vector, cf. (43).

Given that the full-order estimate $\hat{\theta}_N$ is available it is easy to test whether a reduced order estimate $\hat{\eta}_N$ is near optimal or not. Combining (23) with (24) gives that

$$\sum_{t=1}^N e^2(t, \hat{\eta}_N) - \sum_{t=1}^N e^2(t, \hat{\theta}_N) \leq \chi_{\alpha}^2(n) \frac{1}{N} \sum_{t=1}^N e^2(t, \hat{\theta}_N) \tag{44}$$

has to be satisfied for $\hat{\eta}_N$ to be a near-optimal reduced complexity estimate.

It should be clear that the bias error is of the same size as the variance error for this type of estimate. This is in line with the conclusion in Guo and Ljung (1994) that the total error is minimized by a model where the bias error does not exceed the variance error.

We conclude this section with an example which illustrates that also the noise model is important for near-optimal restricted complexity models.

Example 4.4. The true system has order 3 and is given by

$$y(t) = \frac{0.14q^{-1}u(t)}{(1 - 0.8e^{j\pi/6}q^{-1})(1 - 0.8e^{-j\pi/6}q^{-1})(1 - 0.45q^{-1}) + (1 - 0.95q^{-1})e_0(t)},$$

where the noise variance is 3. The input signal has high-pass character with variance 1.3.

Consider first the second order Box–Jenkins model structure

$$y(t) = \frac{bq^{-1}}{1 + a_1q^{-1} + a_2q^{-2}} u(t) + (1 + cq^{-1})e(t).$$

For $N = 100$ samples we have that the left-hand side of (44) evaluates to 19.2 whereas the right-hand side bound is 135. Hence, the estimated model should be near-optimal.




Fig. 4. Mean-square error of estimated frequency functions in Example 4.4. Solid line: Full-order Box–Jenkins model. Dashed line: Second order Box–Jenkins model. Dotted line: Second order output–error model.

To assess this, 100 models were estimated using different noise realizations and the sample mean-square error for the estimated frequency functions was computed. This error is shown in Fig. 4 for the reduced order Box–Jenkins model (dashed line) and for a full-order Box–Jenkins model (solid line). Clearly, the errors for the two different model structures are of the same size.

Fig. 4 also shows the sample mean-square error for an ensemble of output-error (OE) models

$$y(t) = \frac{bq^{-1}}{1 + a_1q^{-1} + a_2q^{-2}} u(t) + e(t),$$

using the same noise realizations as in the Box–Jenkins case. The sample mean-square error for the OE-models is significantly larger than for the full-order model, except at high-frequencies. The OE-model is thus far from near-optimal despite the fact that there are second-order models that can approximate the true system as well as the full-order model.² For the OE-model, however, the left-hand side of (44) evaluates to approximately 3.3×10^4 which is significantly higher than the right-hand side bound of 135. Therefore the OE-model cannot be near-optimal and it follows that it is not possible to give any frequency-wise bounds on the error for this model structure. The only characterization of the error is given by the bias expression (26) which indicates that the error should be smaller at high frequencies (as the input is of high-pass type).

4.5. The separation principle revisited

In this section we have so far argued that one should model as well as possible in order to reduce the impairing effect of measurement noise. Referring to the separation principle in Section 4.2, the full-order model can then be simplified if required without loss of statistical accuracy.

Now, as we will argue in Section 5.2, full-order models are esoteric quantities so that one always have to contend

²This we know from the Box–Jenkins case which uses the same dynamic model as the OE-model.

with restricted complexity models, not at least in control applications. To cope with this, we have introduced the concept of near-optimal restricted complexity models. Such a model is in a statistical sense almost as accurate as any identified full-order model would be, had it been possible to use a such a model. The complexity of such a model thus depends on the quality of the observed data Z^N : A first order model may be near-optimal if only poor information is available whereas it may require a 50th order model when many thousands of high signal-to-noise data samples are available.

The usefulness of near-optimal restricted complexity models lies not only in that they model the true system almost as accurately as a (thought) full-order model, but also in that *they can be used in the separation principle instead of the unattainable full-order model!* To see this, let $\xi = f(\theta)$ denote the mapping from the (thought) full-order model to the model class of interest, let $\hat{\theta}_N$ be the (full-order) ML-estimate and $\hat{\eta}_N$ a near-optimal restricted complexity model. We observed in Section 4.2 that $\hat{\xi}_N = f(\hat{\theta}_N)$ is the ML-estimate of $f(\theta^0)$. Furthermore, any point within the confidence region associated with $\hat{\xi}_N$ will have the same statistical accuracy as $\hat{\xi}_N$ within a factor 2. Now the confidence regions for $\hat{\theta}_N$ and $\hat{\xi}_N$ are related simply by the mapping f (due to the construction of $\hat{\xi}_N$) and since $\hat{\eta}_N$ belongs to the confidence region associated with $\hat{\theta}_N$, $f(\hat{\eta}_N)$ will belong to the confidence region associated with $f(\hat{\theta}_N)$. Thus, $f(\hat{\eta}_N)$ will be a near optimal estimate of $f(\theta^0)$.

5. Model validation

We saw in Section 2.8 that when a model structure has been selected, the set of unfalsified models $\overline{\mathcal{G}}(Z^N)$ can by definition not be falsified by the data Z^N when the noise variance is unknown. One could say that the model builder is trapped inside the model structure. Hence, there is a need to “look over the fence” to ensure that there are no other model structures that can represent the data in a more plausible way, or alternatively test $\overline{\mathcal{G}}(Z^N)$ on new data. This is what model validation is about!

5.1. Model error modeling

Consider the following example.

Example 5.1. Let the residuals of a nominal model be denoted by $\varepsilon(t)$ and let $\varphi(t) = [u(t - 1), \dots, u(t - n)]^T$. Then a standard test-statistic, for testing whether the cross-correlation between $\varepsilon(t)$ and $\varphi(t)$ is zero, is given by

$$\zeta_N = \frac{1}{1/N \sum_{t=1}^N \varepsilon^2(t)} \left\| \sum_{t=1}^N \varepsilon(t)\varphi(t) \right\|_{R_N^{-1}},$$

where R_N is defined in (10) and where $\|x\|_Q \triangleq \sqrt{x^T Q x}$. Normalized by $1/n$, this statistic is asymptotically

$F(n, N - n)$ - distributed when ε is white noise, so a suitable cross-correlation test is

$$\zeta_N \leq n F_{\alpha}(n, N - n). \tag{45}$$

We can express ζ_N somewhat differently. Let $\hat{\theta}_N$ be the least-squares estimate, cf. (10), for the following FIR model of the residuals:

$$\varepsilon(t) = \varphi^T(t)\theta + e(t).$$

It is easy to see that

$$\zeta_N = \hat{\theta}_N^T \frac{R_N}{1/N \sum_{t=1}^N \varepsilon^2(t)} \hat{\theta}_N,$$

and hence the test (45) corresponds to testing

$$\hat{\theta}_N^T \frac{R_N}{1/N \sum_{t=1}^N \varepsilon^2(t)} \hat{\theta}_N \leq n F_{\alpha}(n, N - n).$$

Comparing with (14) (where for this example $e(t, \hat{\theta}_N) = \varepsilon(t) - \varphi^T(t)\hat{\theta}_N$), we see that this test is closely related to testing whether the zero model $\theta = 0$ belongs to the set of unfalsified models for the above FIR-model of the residuals under the assumption that the residuals can be modeled by this FIR-model. The difference lies in the denominator which is an estimate of the variance of the residuals of the model error model. Above the estimate is conditioned on that the true parameter in the model error model is zero which gives the variance of the original residuals whereas (14) uses the least-squares estimate of the model error estimate. The difference corresponds to the difference between hypothesis testing and computing confidence regions.

The above example was used in Ljung (1999a) to point out that standard model validation tests such as cross-correlation tests between residuals and inputs can be interpreted as first modeling the residuals, with the resulting model named *model error model*, and then testing whether the zero model is included in the set of unfalsified model error models. Instead of just computing yes/no answers to tests such as (45) it was suggested that an intuitively appealing way of presenting these tests is by plotting the Bode-diagram of the model error model with uncertainty regions indicated. From this insight follows also that more complex models than finite impulse response (FIR) models can be used and it is recommended that the model structure for the model error model should be considerably richer than the nominal model. The reader is also referred to Ljung and Guo (1997) for results on how the model error is bounded by test statistics such as ζ_N .

The main message in Ljung (1999a) is that if the nominal model is unfalsified, i.e. the uncertainty region for the model error model includes the zero model, then, even though the nominal model structure (with its own uncertainty description) is unfalsified, one should use the nominal model structure together with the *uncertainty region of the model*

error model. Since the model structure for the model error model is more flexible than the model structure for the nominal model, this will give a larger, and hence “safer”, set of unfalsified models.

Notice that this conclusion is *completely in line* with the discussion in Section 4.4 if we consider the nominal model as being of restricted complexity whereas the model error model is flexible enough to capture all dynamics in the residuals. Then it is the uncertainty set associated with the model error model that is relevant.

We are here at the crux of the modeling problem—the model builder wants to be sure that his model set includes the ‘true’ system. However, we stress that

even though the confidence region associated with the model error model structure may be more conservative compared to the confidence region for the nominal model, there is still no guarantee that this set contains the true system since in general cannot guarantee that the model error model captures the remaining dynamics completely.

Nevertheless, the concept of model error modeling has helped make explicit the *necessary* leap of faith in system identification.

This brings us to the next topic on the agenda.

5.2. The true system: a mirage

Consider the following *gedanken experiment*. A continuous time true system is LTI but infinite dimensional with single poles spaced many decades apart. It is excited by a band-limited input signal which covers frequencies up to a certain frequency ω_{\max} which includes n of the true system poles. As more and more data are collected from this setup, the identification procedure appears to converge to an n th order model which seems to be a correct description of the true system as this model will pass all validation tests. This model is in fact a near-optimal restricted complexity model as the poor excitation at high frequencies make the uncertainty of a full-order model *very* high in this frequency region, cf. Section 4.4. However, as even more data samples are collected (still using the same input spectrum), eventually the small discrepancies in the system behavior below frequency ω_{\max} , as compared to an n th order model, will become detectable from data and the model order may need to be adjusted (upwards) in order for the model to be unfalsified.

From this mental exercise, we can conclude that even what we may consider as a full-order model, is indeed only a near-optimal restricted complexity model. As the signal-to-noise ratio increases, more and more details of the system can be modeled and, further, a different type of input may give a drastically different model, which may again appear to be an excellent model of the true system (for this particular input). These observations support the arguments brought forward in Skelton (1989) that any model is input dependent and that, hence, the quest for a model of the true system is futile.

The inclusion of an unstructured dynamic term such as (4) in set-membership identification can be seen as a way of incorporating this modeling limitation. In (Hjalmarsson, 1991) a linear time-varying dynamic term is used to account for unmodeled dynamics and to prevent the modeling accuracy to increase unrealistically.

The conclusion that the true system cannot be modeled to an arbitrary degree may seem disappointing. One has to accept that we are *at best working with models that are near-optimal*, cf. Section 4.4. However, the main message in this contribution will be that by carefully selecting the input, the system can be forced to reveal the properties that are relevant for the particular application and this is all that is needed for a successful application. As already pointed out in Section 2.9, this will be a recurrent theme in the paper. We will approach this topic in the next sub-section. We also refer the reader to Section 6, Section 11 and the concluding remarks in Section 13 for further elaborations on this topic.

5.3. Validating with confidence

Suppose that it is critical for the application that any unmodeled dynamics in a certain model structure \mathcal{M}_{ME} does not exceed a certain bound. In a robust control context, it could for example be that the peak gain in a certain frequency band should not exceed a certain value. In such a situation it would boost the confidence of the model builder if it could be ascertained that the model, which we denote by $\{\hat{G}_N, \hat{H}_N\}$, (and its corresponding uncertainty set) would be falsified if this is the case.

Well, let us examine the outcome if a model error model is estimated using the structure \mathcal{M}_{ME} to which the unmodeled dynamics belongs. For simplicity, let us assume that the true system is LTI, cf. (25), and that the asymptotic results for the prediction error method in Section 2.6.1.3 are valid. In this case the residuals (16) are given by

$$\varepsilon = (G_o - \hat{G}_N)u_F + \frac{H_o}{\hat{H}_N} e_o,$$

where $u_F = \hat{H}_N^{-1}u$, and the model \hat{G}_N will be falsified if the uncertainty region for the model error model does not include the zero model.

Using (36) and some simple algebra, this is guaranteed to happen if

$$|G_o(e^{j\omega}) - \hat{G}_N(e^{j\omega})| > 2\sqrt{\chi_x^2(n)\kappa_{n,N}(\omega) \cdot \frac{\Phi_v(\omega)}{N\Phi_u(\omega)}} \quad (46)$$

for some frequency ω . Above $\kappa_{n,N}$ is associated with the model structure \mathcal{M}_{ME} . Hence, if we want to ensure that model errors larger than some function $\delta(\omega)$ are detected, then the experiment should be carried out such that the right-hand side of (46) is less than $\delta(\omega)$. We emphasize, again, that this conclusion is predicated on the assumption that \mathcal{M}_{ME} is flexible enough to capture the unmodeled dynamics.




Fig. 5. Dashed line: input spectrum. Thin solid line: true system. Thick solid line: uncertainty region around estimated nominal first order OE-model.

Example 5.2. A third order system with a resonance is corrupted by white noise and excited with a low pass input, also with a resonance (see Fig. 5). The system is identified using a first order output error model. The model, together with its uncertainty region (based on that the true system is in the model set) is shown in Fig. 5. Clearly the model has missed the resonance peak and the uncertainty region is misleading. The model error is shown in Fig. 6 together with the bound from (46) based on a 10th order FIR model error model. We see that we can expect to detect the resonance in our model error model but not any model error at other frequencies. In the same figure, the uncertainty bound for a 10th order FIR model error model is shown. As predicted, the resonance peak is detected since the uncertainty region for the model error model does not include zero around the resonance, whereas the model error at other frequency bands is not detected.

Notice that condition (46) depends on the input not only through the input spectrum but also through the factor $\kappa_{n,N}(\omega)$, cf. Proposition 3.1. This has a, perhaps unexpected, implication.

Example 5.3 (Example 5.2 continued). Suppose that the order of the model error model is increased from 10 to 100. One would then expect the lower bound (46) for detecting unmodeled dynamics to increase significantly. For the case when (32) holds, it should increase by a factor of $\sqrt{100/10} \approx 3.2$. The bound is shown in Fig. 7 for the two cases. We see that there is actually an increase of approximately this factor, except at low frequencies and around $\omega = \pi/2$, which happens to be where the peak of the input spectrum is located, where there is only a minor increase.

Notice also that even though there is a peak in the input spectrum at $\omega = \pi/2$, it is a factor of 20 smaller than the input spectrum at low frequencies, cf. Fig. 5. Hence, the

small increase around $\omega = \pi/2$ cannot be explained by the magnitude of the input spectrum around this frequency.

The phenomenon is due to the factor $\kappa_{n,N}(\omega)$ in (46). The conditions in Proposition 3.1 are approximately satisfied for the present example (cf. Remark 3 after Proposition 3.1) so for large N , $\kappa_{n,N} \approx \kappa_n$ with κ_n given by (34). A plot of $\kappa_n(\omega)$ is shown in Fig. 8 for the two model orders. The poles in $\kappa_n(\omega)$ consist in this case of the poles 0.9, 0.9, $0.97e^{\pm j\pi/2}$ of the stable spectral factor of the input spectrum and the poles of the FIR-model. The double pole at 0.9 gives a large contribution to $\kappa_n(\omega)$ at low frequencies, whereas the complex poles give a large contribution around $\omega = \pi/2$. Since all the poles of an FIR-model are at the origin, $\kappa_n(\omega)$ gets a frequency independent contribution of m from an m th order FIR-model. Thus, $\kappa_n(\omega)$ increases linearly with the model order m and this is clearly seen in Fig. 8 since the solid line (corresponding to $m = 100$) is offset by $100 - 10 = 90$ above the dashed line (corresponding to $m = 10$). However, the relative increase at different frequencies is vastly different. At frequencies where the poles of the input spectrum contribute very little, the increase is a factor 10 but at frequencies where the influence from the poles of the input spectrum is significant, the relative increase is much less. Hence, the relative increase in the uncertainty bound is much less at low frequencies and, especially, around $\omega = \pi/2$.

The key observation in Example 5.3 is rather unexpected and indeed good news as it implies that the input spectrum may be designed so as to allow very flexible model error models with only minor penalty in the falsification power at certain frequency bands. This is also consistent with the fact that when periodic inputs are used, over-modeling does not result in increased variance of the estimated frequency function for frequencies corresponding to the spectral lines of the input (although problems occur at other frequencies due to inexact pole/zero cancellations). However, notice also a large $\kappa_n(\omega)$ gives larger uncertainty bounds so, near peaks of the input spectrum, the bounds can be significantly worse than the noise to signal ratio. In these frequency regions the high-order approximation (32) typically underestimates the true variance.

5.4. Validating restricted complexity quantities

Let G be a given model and suppose that we would like to find out if this model represents a certain property of the true system in a sufficiently accurate way. This can often be formulated as that some function of the error between the model and the true system should be small. To be explicit, suppose that the relative error should be smaller than some function $\gamma(\omega)$, i.e.

$$\left| \frac{G(e^{j\omega}) - G_o(e^{j\omega})}{G(e^{j\omega})} \right| \leq \gamma(\omega).$$

Suppose now that we would like to validate this property using a data set Z^N . We can now directly refer to the

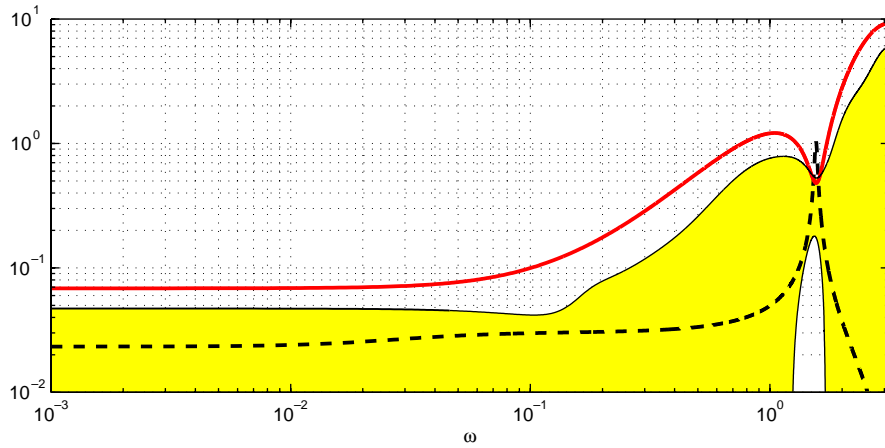


Fig. 6. Dashed line: model error. Solid line: lower bound for model errors that are guaranteed to be detected by a 10th order FIR model error model. Shaded area: uncertainty region for estimated 10th order model error model.




Fig. 7. Smallest model error magnitude guaranteed to be detected in model validation. Dashed line: $m = 10$. Solid line: $m = 100$.




Fig. 8. $\kappa_n(\omega)$. Dashed line: $m = 10$. Solid line: $m = 100$.

separation principle introduced in Section 4.2. An asymptotically efficient estimate of the relative error is

$$\hat{\Delta}_N(e^{j\omega}) = \frac{G(e^{j\omega}) - G(e^{j\omega}, \hat{\theta}_N)}{G(e^{j\omega})},$$

where $\hat{\theta}_N$ is any asymptotically efficient full-order estimate of the true system parameters. A confidence region for the relative error can be obtained by mapping the uncertainty set (23) through the linear transformation $\hat{\Delta}_N$. If the resulting confidence region for the relative error includes values larger

than $\gamma(\omega)$ for some frequency ω , the model is invalidated.

The same technique can be applied when the function is non-linear in G_\circ . However, in this case the resulting confidence region may be very complex. Then, a first order Taylor approximation may be used to obtain an approximation of the confidence region that depends linearly on the covariance matrix of $\hat{\theta}_N$. In the very interesting contribution (Ninness & Henriksen, 2003), it is discussed how to obtain exact confidence regions by way of Markov chain Monte Carlo simulations.

An alternative to the above procedure is to compute a model error model for the residuals $\varepsilon(t) = y(t) - G(q)u(t)$, and use the confidence region associated with this model error model to estimate the size of the relative error. Notice, however, that this may not give optimal accuracy as opposed to the use of the separation principle.

6. Half-time intermission

Before we proceed with the second part of the paper which is directly concerned with how identification and control interrelates, let us pause and summarize the observations

so far: We have in Section 5.2 argued that the quest for a full-order model is futile, and that *the best one can hope to obtain is a near-optimal restricted complexity model*. The objective of identifying such a model is worthwhile to pursue from a statistical accuracy point of view, leading to the pragmatic conclusion:

- (i) *Always first model as well as possible.*

As argued in Sections 4.2 and 4.5, such models are suitable, from a statistical perspective, for replacing the true system in subsequent computations of other quantities such as low order models or the parameters in a controller.

Now, the obvious question is how does one know if one has obtained a near-optimal restricted complexity model? Unfortunately, there is no precise answer to this question as this is a model validation problem, cf. the discussion in Section 5.1; verifying the condition (44) on the mean-square error requires the full-order model. It is also important to realize that *uncertainty descriptions based on near-optimal restricted complexity models are not necessarily reliable*. Compare with Example 4.3, computing uncertainty bounds based on that the system only has one impulse response coefficient will clearly not reflect the true model error. It is the confidence region associated with the full-order model that is relevant. In practice one may

- (ii) *use a very flexible model structure as benchmark for computing confidence bounds and mean-square error.*

We now come to the role played by the input, a subject we have touched on in Section 5.3 and which we will now elaborate further on. Ideally, one should select the input such that the full-order *confidence region is as large as possible, while still satisfying the uncertainty specifications required by the application*. This would mean a minimum of experimentation on the system and that the simplest possible (from the point of view of the application) model would be near-optimal and useful for the application.

In practice, using a very flexible model in lieu of the true system, as suggested in (ii), would imply a concern whether the associated confidence bounds grow so large that they become useless. To examine this issue further, let us return to Example 4.3 and notice that the accuracy of the static gain estimate is *independent* of the model order! Hence, in this example, the model builder does not have to worry at all about that the model estimate will be too uncertain should it turn out that the true system order was higher than expected.³ The reason for this lies in the choice of input signal which is concentrated to one frequency only. This indicates that (near) sinusoidal inputs may help limit the variance uncertainty for high-order models. This was also

the theme in Section 5.3 and we will now use Proposition 3.1 to make this even more explicit.

Suppose that the input spectrum is

$$\Phi_u(\omega) = \alpha^2 \frac{1 - \rho^2}{|e^{j\omega} - \rho|^2} + \tilde{\Phi}_u(\omega),$$

where $\rho < 1$ is close to 1 so that the first term approximately corresponds to a constant term in the input with amplitude α . Suppose that the conditions of Proposition 3.1 hold. Order the ξ_k in (34) such that $\xi_1 = \rho$, then it holds

$$\begin{aligned} N \text{Var}(\hat{G}_N(e^{j\omega})) \\ \approx \left(\frac{1 - \rho^2}{|e^{j\omega} - \rho|^2} + \sum_{k=2}^{m_a+m_b} \frac{1 - |\xi_k|^2}{|e^{j\omega} - \xi_k|^2} \right) \\ \times \frac{\Phi_v(\omega)}{\alpha^2(1 - \rho^2)/(|e^{j\omega} - \rho|^2) + \tilde{\Phi}_u(\omega)}. \end{aligned}$$

For $\omega \approx 0$, the above expression reduces to Φ_v/α^2 for ρ sufficiently close to 1, i.e. the noise to signal ratio at the zero frequency which is independent of the model order. It is also easy to see that at frequencies not in a neighborhood of zero, the impact from the constant part of the input will be small.

The above derivation can be generalized to sinusoidal inputs. We are thus in position to suggest the advice

- (iii) *select the input such that the model uncertainty at frequency regions of interest is insensitive to the model complexity.*

As we have indicated above (near) sinusoidal inputs are useful from this perspective, whereas wide band inputs such as white noise do not possess this property.

When the model structure is pre-specified one can notice that broad-band excitation can be difficult to handle as it then might not exist a near-optimal restricted complexity model within the pre-specified structure which may render the error quantification difficult.

Equipped with (i)–(iii), which can be viewed as generally applicable pragmatic guidelines, we are now finally in position to discuss the control application from a system identification perspective. We will continue to use (or abuse) the concept of a full-order model as it is useful as benchmark for other methods and as there exists practical implementations in terms of near-optimal restricted complexity models.

7. Links between control and identification

In this section we will discuss how control and identification interact. In Section 7.1 we give a brief review of robust control. In Section 7.2, the main issues in relation to identification are summarized. These issues are then treated in subsequent sections.

³ The attentive reader will notice that the argument assumes that the observation interval is long enough that the complete step-response has been observed.




Fig. 9. Closed-loop system.




Fig. 10. Robust control configuration.

7.1. Robust control

Fig. 9 illustrates a multivariable feedback configuration where the controller C and the true system G_o are LTI. The closed-loop system equations are

$$\begin{pmatrix} y \\ u \end{pmatrix} = \begin{bmatrix} G_o \\ I \end{bmatrix} S_I(G_o, C) \begin{bmatrix} C & I \end{bmatrix} \begin{pmatrix} r \\ w \end{pmatrix} \triangleq M(G_o, C) \begin{pmatrix} r \\ w \end{pmatrix}, \quad (47)$$

where $S_I(G_o, C) \triangleq (I + CG_o)^{-1}$ is the achieved input sensitivity function. Above, r and w are known external excitations. The (1,1)-block of $M(G_o, C)$ is the complementary sensitivity function

$$T(G_o, C) = (I + G_o C)^{-1} G_o C.$$

The closed-loop system in Fig. 9 is internally stable if and only if all four closed-loop transfer functions in $M(G_o, C)$ are in \mathcal{H}_∞ .

Extracting all uncertain elements of the true system G_o results in the configuration in Fig. 10 where Δ represents the uncertain elements. Scaling functions have been introduced in the block P so that Δ is normalized in some way. The performance is measured by some induced norm from the external input p (which may consist of r and w in Fig. 9 and/or other signals of interest) to the signal z . Also here scalings are introduced so that the desired performance is that the gain from p to z is less than, e.g., 1. The generalized plant P thus consists of a nominal model G as well as robust stability and performance weightings. The objective is to design the controller C such that the closed-loop is stable and the performance criterion is satisfied for all uncertainty blocks Δ of some pre-specified structure. We refer to

Skogestad and Postlethwaite (1996) and Zhou, Doyle, and Glover (1996) for comprehensive treatments of the robust control problem.

7.1.1. Model sets in robust control

The, perhaps, most common uncertainty description in robust control is multiplicative uncertainty

$$G_o = G(I + W_I \Delta), \quad (48)$$

where Δ is unstructured, i.e. a full complex matrix, satisfying $\|\Delta\|_\infty \leq 1$, and W_I is a frequency weighting.

Model uncertainty can also be modeled as perturbations of the coprime factors X and Y in a normalized coprime factorization (Zhou et al., 1996) $G = X^{-1}Y$ of a nominal model G :

$$\left\{ (X + \Delta_X)^{-1}(Y + \Delta_Y) : \begin{bmatrix} \Delta_Y \\ \Delta_X \end{bmatrix} \right\|_\infty < \varepsilon \right\}. \quad (49)$$

Another useful measure is the v -gap metric introduced in Vinnicombe (1993):

$$\delta_v(G_1, G_2) = \begin{cases} \max_\omega \kappa(G_1(e^{j\omega}), G_2(e^{j\omega})) & \text{if } W(G_1, G_2) = 0 \\ 1 & \text{otherwise,} \end{cases} \quad (50)$$

where κ is the chordal-distance

$$\kappa(G_1(e^{j\omega}), G_2(e^{j\omega})) = \frac{|G_1(e^{j\omega}) - G_2(e^{j\omega})|}{\sqrt{1 + |G_1(e^{j\omega})|^2} \sqrt{1 + |G_2(e^{j\omega})|^2}} \quad (51)$$

and where

$$W(G_1, G_2) = wno(1 + G_1^* G_2) + \eta(G_2) - \tilde{\eta}(G_1).$$

Above $wno(G)$ denotes the winding number about the origin of $G(z)$ as z traces the unit circle, avoiding unit circle poles and zeros by indentation into the exterior of the unit disc. Furthermore $\eta(G)$ ($\tilde{\eta}(G)$) denotes the number of poles in the complement of the closed (open) unit disc.

7.1.2. Analyzing robustness

The robust stability problem entails checking whether the feedback configuration in Fig. 10 is stable for all allowed perturbations Δ for a given controller C . Similarly, the robust performance problem entails checking that some performance criterion is satisfied for all allowable perturbations. The computational complexity of checking these conditions depends very much on the class of perturbations.

For the unstructured case, i.e. when $\|\Delta\|_\infty \leq 1$ is the only condition on Δ , robust stability can be easily checked by computing the \mathcal{H}_∞ -norm of a certain transfer function. Consider, e.g., the case of multiplicative input uncertainty (48). It is easy to show that

$$S_I(G_o, C) = [I + \bar{\Delta}(G_o, G, T_I(G, C))]^{-1} S_I(G, C) \quad (52)$$

where

$$\bar{\Delta}(G_o, G, T_I) = T_I G^{-1} (G_o - G) \quad (53)$$

is the input relative error $G^{-1}(G_o - G)$ weighted by the complementary sensitivity function seen at the input of the system $T_I = I - S_I$. Hence, if the nominal design $S_I(G, C)$ is stable, a small gain argument gives that robust stability is guaranteed if

$$\|\bar{\Delta}(G_o, G, T_I(G, C))\|_\infty < 1 \quad (54)$$

when GC and G_oC have the same number of unstable poles.

Robust performance is harder to verify and the structured singular value μ (Zhou et al., 1996) has to be computed for non-conservative results. For unstructured uncertainty, this is still a convex optimization problem.

Simple sufficient conditions for robust performance exist as well. We illustrate this for multiplicative input uncertainty (48). Suppose that it is desired that

$$\|W_p S_I(G_o, C)\|_\infty < 1, \quad \forall G_o = G(I + W_I \Delta) \quad (55)$$

with $\|\Delta\|_\infty \leq 1$ and where W_p is a scalar weighting function. From (52) it follows that

$$\begin{aligned} \bar{\sigma}(W_p S_I(G_o, C)) &\leq \bar{\sigma}(W_p S_I(G, C)) \bar{\sigma}((I + \bar{\Delta}(G_o, G, T_I(G, C)))^{-1}) \\ &\leq \frac{\bar{\sigma}(W_p S_I(G, C))}{\underline{\sigma}(I + \bar{\Delta}(G_o, G, T_I(G, C)))} \\ &\leq \frac{\bar{\sigma}(W_p S_I(G, C))}{1 - \bar{\sigma}(\bar{\Delta}(G_o, G, T_I(G, C)))}. \end{aligned}$$

Hence, robust performance in the sense (55) is guaranteed if

$$\bar{\sigma}(W_p S_I(G, C)) + \bar{\sigma}(\bar{\Delta}(G_o, G, T_I(G, C))) < 1, \quad \forall \omega \quad (56)$$

which, using (48), results in the condition

$$\bar{\sigma}(W_p S_I(G, C)) + \bar{\sigma}(T_I(G, C)W_I) < 1, \quad \forall \omega. \quad (57)$$

Remark 1. Condition (57) is satisfied if robust stability (54) and nominal performance $\bar{\sigma}(W_p S_I(G, C)) < 1$ both hold with sufficient margin (Morari & Zafriou, 1989; Skogestad & Postlethwaite, 1996). Not surprisingly, one may arrive at the same condition by considering the performance degradation from the nominal design. The triangle inequality gives

$$\begin{aligned} \bar{\sigma}(W_p S_I(G_o, C)) &\leq \bar{\sigma}(W_p S_I(G, C)) + \bar{\sigma}(W_p(S_I(G_o, C) - S_I(G, C))) \end{aligned}$$

which leads to (57).

Remark 2. For SISO (single-input/single-output) systems, (57) is equivalent to (55) since for SISO systems the left-hand side of (56) is the structured singular value for robust performance.

Robust stability for the set (49) can be checked using the generalized stability margin

$$b_{G,C} = \begin{cases} \|M(G, C)\|_\infty^{-1} & \text{if } M \text{ is stable,} \\ 0 & \text{otherwise.} \end{cases} \quad (58)$$

Stability is guaranteed only if $b_{G,C} \geq \varepsilon$ (Glover & McFarlane, 1989).

The maximum stability margin for a model with normalized coprime factorization $G = X^{-1}Y$ is given by (Glover & McFarlane, 1989)

$$b_G \triangleq \max_C b_{G,C} = \sqrt{1 - \left\| \begin{array}{c} Y \\ X \end{array} \right\|_H^2},$$

where $\|\cdot\|_H$ denotes the Hankel norm (see, e.g. Zhou et al., 1996). Thus b_G gives an upper bound on the maximum allowable normalized coprime uncertainty.

It can be shown that for $b_{G,C} > 0$,

$$b_{G,C} = \min_\omega \kappa \left(G(e^{j\omega}), -\frac{1}{C(e^{j\omega})} \right).$$

The following result was first presented in Vinnicombe (1993).

Proposition 7.1. *Let the model set be a ball in the v -gap metric (50):*

$$\{G_o : \delta_v(G, G_o) \leq \beta\}$$

where G is an arbitrary nominal model. Then it holds that a controller C stabilizes all systems in this set if and only if $b_{G,C} > \beta$.

The following result also holds (Vinnicombe, 1993).

Proposition 7.2. *For a given nominal model G and system G_o it holds that G_o is stabilized by all controllers in the set $\{C : b_{G,C} > \beta\}$ if and only if $\delta_v(G, G_o) < \beta$.*

The v -gap metric is the only metric for which Propositions 7.1 and 7.2 both hold.

Robust stability for real parametric uncertainty is treated in Rantzer and Megretski (1994). A convex parametrization of all controllers that stabilize the system for all possible parameter combinations is provided when the uncertain real parameters appear linearly in the closed-loop characteristic polynomial.

7.1.3. Synthesizing robust controllers

For general structures of the uncertainty, robust control design amounts to minimizing the structured singular value of a certain transfer matrix with respect to the controller. There are no generally applicable algorithms for such μ -synthesis problems; for complex perturbations, so called *DK*-iterations (Skogestad & Postlethwaite, 1996) may be used to, at least, decrease μ .

In Section 7.1.2 it was pointed out that the robust performance condition (55) could be expressed as (57). The latter condition is closely related to the mixed sensitivity \mathcal{H}_∞ condition

$$\left\| \begin{array}{c} W_p S(G, C) \\ W_I T(G, C) \end{array} \right\|_\infty < 1$$

where $S(G, C) = I - T(G, C)$ is the sensitivity function. Thus, \mathcal{H}_∞ control design can be used to approximately optimize robust performance in this case (Skogestad & Postlethwaite, 1996).

\mathcal{H}_∞ loop-shaping (Glover & McFarlane, 1989) combines classical loop-shaping ideas with modern robust control. First pre- and post-compensators W_1 and W_2 are added to the system to shape the loop gain W_2GW_1 of the nominal system in a desired manner. Then a controller C is designed such that the stability margin (58) of the shaped system W_2GW_1 is sufficiently large. As discussed in Section 7.1.2, this guarantees robust stability with respect to perturbations in the coprime factors of the shaped system. It also means that stability is guaranteed for all systems G_o which are such that $\delta_v(W_2GW_1, W_2G_oW_1) < b_{W_2GW_1, C}$, cf. Section 7.1.2.

A simple and very instructive design method is internal model control (IMC) (Morari & Zafriou, 1989). The controller C is parametrized in terms of the nominal model G as $C = Q(I - GQ)$ where Q is a stable transfer function. A two-step procedure is employed where in the first step Q is determined so as to meet some nominal performance specification, e.g. $\|W_2S(G, C)W_1\|_\infty < 1$. In the second step, Q is augmented with a “de-tuning” filter F giving a new “ Q ” equal to QF . The filter F is selected so as to ensure robust stability. For multiplicative input uncertainty (48), the robust stability condition (54) can be expressed as $\|QF(G_o - G)\|_\infty < 1$ since $T_I(G, C) = QFG$ in this case.

Model reference control is closely related to IMC. Here the target is to design the controller C such that a particular sensitivity function is obtained, at least nominally. For example, the nominal design may be such that

$$T(G, C(G)) = T_d$$

for some desired complementary sensitivity function T_d . Solving for C gives

$$C(G) = G^{-1}(I - T_d)^{-1}T_d. \quad (59)$$

Naturally, non-minimum phase zeros in G are not allowed to be canceled in this design. Notice, that with this design it is possible to parametrize the model G directly in terms of the optimal controller C :

$$G(C) = (I - T_d)^{-1}T_dC^{-1}. \quad (60)$$

There is thus a one-to-one relationship between models and optimal controllers. The parametrization (60) is known in adaptive control as direct parametrization.

A synthesis procedure for real parametric uncertainty is presented in Ghulchak and Rantzer (2002). It builds on Rantzer and Megretski (1994) and produces, via convex optimization, a robustly stabilizing controller which is arbitrarily close to the optimal controller.

7.2. Identification and control

In the preceding section we have tried to give a flavour of modern robust control theory. We will now turn to the problem of using models obtained from system identification in robust control design. An immediate insight is the following:

The set of unfalsified models delivered by system identification generally corresponds to very structured uncertainty descriptions from the point of view of robust control.

For example ellipsoidal sets in the parameter domain such as (23), do not readily fit into any of the model set descriptions in Section 7.1.1. As indicated in Section 7.1.3, this in general means that the robust control design problem becomes very difficult. In fact, it is often highly non-trivial to even analyse the robustness properties for a given controller. In Section 7.3 we report on some recent progress in robustness analysis tailored for ellipsoidal model sets of the type (23). Awaiting further development in the areas of robustness analysis and control synthesis for general model sets, an alternative approach has been to outerbound the set of unfalsified models in a way that is congruent with existing robust control theory. Naturally, it is desired to introduce as little conservatism as possible in this procedure, as well as keeping the complexity of the bound to the minimum. In Section 7.4 we discuss some results in this direction.

In Section 7.5 we review some control design methods that have been employed in conjunction with identified model sets.

Characteristic to existing robust control methods is that the uncertainty specifications are assumed to be pre-specified. When the identification part of the problem is taken into account, an additional freedom in the problem formulation is unleashed in that the experiment design can be used to influence the uncertainty set. Hence, for given experimental constraints there are experimental conditions for which the resulting uncertainty set is such that the robust performance specifications cannot be improved upon, i.e. the uncertainty set is optimally shaped (for robust performance). The complexity of modern robust control algorithms as well as system identification has (so far) prohibited explicit solutions to such problems. A further complicating factor is that industrial practice motivate the use of restricted complexity models as well as restricted complexity controllers. The non-transparency of modern robust control algorithms, induced by their complexity, has forced deliberations on these topics to rely on simplified arguments and in Section 7.6 we will discuss some insights that can be obtained.

7.3. Robustness analysis for identified models

A simple expression for the “real” stability margin is derived in (Rantzer, 1992). Based on this result, robust stability for identified ARX-models is considered in Kosut and Ansonson (1994), see also Raynaud (1993).

In a very interesting series of papers (Bombois et al., 2000b; Bombois, Gevers, Scorletti, & Anderson, 2001, 2001), summarized in Gevers, Bombois, Codrons, Scorletti, and Anderson (2003), Gevers, Bombois and co-workers have developed analysis tools for the set of unfalsified models obtained in prediction error identification, i.e. (23). Also building upon Rantzer (1992), given a controller C , a computable, necessary and sufficient condition for stability for all models in (23) is provided. It is also shown that the largest chordal distance (51) between any model in (23) and a nominal model G can be computed by convex optimization. Hence, by frequency gridding, the largest v -gap can be computed approximately. Referring to Proposition 7.2, this in turn means that it is possible to characterize the largest ball (in the v -gap metric) of controllers which are guaranteed to stabilize all systems in (23). It is also shown that for the performance measure $\sigma_1(W_2 M(G_\circ, C) W_1)$ (recall that M is the closed loop transfer function defined in (47)), where the frequency weightings W_1 and W_2 are block-diagonal, it is possible to compute the worst case performance over the set (23) via convex optimization.

7.4. Outerbounding the set of unfalsified models

Robust stability is the issue in Douma, Van den Hof, and Bosgra (2003). Consider a ball \mathcal{S} of systems in the gap metric (Zames & El-Sakkary, 1980) and a corresponding set \mathcal{C} of controllers characterized as being the largest ball in the gap metric (centered around some nominal controller) which is such that all controllers in this ball stabilize all systems in \mathcal{S} . It is then shown that there exists a pair of sets, one which consists of systems and one which consists of controllers, which are characterized by the double Youla parametrization and which are such that the set of systems includes \mathcal{S} and the set of controllers include those in \mathcal{C} . Hence, with respect to robust stability, the double Youla parametrization is less conservative than the gap metric for representing model sets. By similar arguments it is shown that the double Youla parametrization is less conservative than using the v -gap and the A -gap (see Bongers, 1991 for a definition) to characterize model sets.

In Van den Hof (1998) it is pointed out that choosing a model structure such that the model error is affine in the performance criterion implies considerable simplifications. When the performance is measured by some frequency weighted element of $M(G, C)$, it is pointed out that the dual Youla parametrization is one such parametrization. Further insights on the implications of outerbounding are reported in Douma and Van den Hof (2003).

Pointwise bounds in the frequency domain are derived in Wahlberg and Ljung (1992) using a set-membership setting. The effects of unknown impulse response coefficients, initial condition and an unknown-but-bounded disturbance are taken into account. Another contribution to the frequency domain bounding of ellipsoidal parameter sets is Devilbiss and Yurkovich (1998). Confidence bounds in the frequency

domain, taking undermodelling explicitly into account, for the instrumental variable method are developed in Hakvoort and Van den Hof (1997).

A recent contribution where frequency domain bounds for a set-membership method are derived in Milanese and Taragna (2002). Here prior assumptions on the decay rate of the impulse response coefficients are used to determine the grid size of frequency points so that intersample variation can be neglected.

For standard \mathcal{H}_∞ methods to be applicable, the uncertainty bound should be the magnitude of a rational transfer function. A method that produces such bounds is presented in Scheid, Bayard, and Yam (1991).

7.5. Control algorithms using identified model sets

In this section we will give some examples of control algorithms that have been developed specifically to cope with model sets produced by identification algorithms.

7.5.1. Ellipsoidal model sets

The problem of minimizing the maximum linear quadratic regulator (LQR) cost in a model set described by ellipsoidal parametric uncertainty, such as (23), is studied in Lau, Boyd, Kosut, and Franklin (1991). It is shown that the solution is the LQR corresponding to a “worst-case” plant in the model set. A heuristic algorithm for computing the worst-case plant is proposed.

In Raynaud, Pronzato, and Walter (1995) the design specification is that the closed loop poles should be inside a disc with a pre-specified radius $\rho < 1$. The objective is to find the controller that maximizes the volume of a model set of the type (23) (i.e. the constant $\gamma_x^2(n)$ in (23)) such that the closed-loop poles satisfy the design objective for all models in this set. It is remarked that, following Rantzer and Megretski (1994), the problem can be recast as an infinite dimensional convex problem. Instead a lower bound for the maximum volume, amenable to standard \mathcal{H}_∞ robust control, is used. An alternative path might be to use the method in Ghulchak and Rantzer (2002) which addresses the problem formulated in Rantzer and Megretski (1994).

There is at present no control design method that accounts for both robust stability and robust performance when the model set is defined by (23). For SISO systems it is suggested in Bombois, Scorletti, Anderson, Gevers, and Van den Hof (2002) to first determine the set of controllers for which the nominal performance is somewhat better than the desired robust performance. This is a standard \mathcal{H}_∞ problem and the set can be easily characterized. It is then tested whether all controllers in this set stabilize all systems in the model set. A similar test for the performance criterion is presented. The robust stability test boils down to the computation of the structured singular value of a certain matrix. If these two tests are passed, any controller in the set of controllers can be used.

7.5.2. Model sets from set-membership identification

IMC is employed in Malan, Milanese, Regruto, and Taragna (2001). The Q part of the controller, cf. Section 7.1.3, is obtained by solving

$$\hat{Q} = \arg \min_{Q \in \mathcal{H}_\infty} \|(1 - QG)W\|_\infty$$

subject to the robust stability condition $\|QW_A\|_\infty < 1$ (W_A is an upper bound on the additive model error. This is a standard \mathcal{H}_∞ problem provided W_A is the magnitude of a stable rational transfer function.

7.6. Linking identification and robust control

In Remark 1 in Section 7.1.2 we saw that robustness issues may be analyzed by considering the difference between the closed-loop transfer function of the nominal design and the true one. This has been a fruitful starting point for examining the links between identification and control. In particular, this type of analysis has been employed to examine what constitutes a good nominal model (Schrama, 1992; Lee, Anderson, Mareels, & Kosut, 1995; Lee, Anderson, Kosut, & Mareels, 1993; Zang, Bitmead, & Gevers, 1995; Åström, 1993). But it can also be used to analyze how the model uncertainty and the experimental conditions influence the performance. Building on Hjalmarsson and Jansson (2003), we shall pursue such an analysis in this section.

Let us introduce the weighted difference between the closed-loop system consisting of the true system G_o controlled by the controller C_o , and the designed closed-loop system using the nominal model G with some controller C

$$\begin{aligned} J(G_o, C_o, G, C) &\triangleq \\ &\begin{pmatrix} w_{11}I & 0 \\ 0 & w_{21}I \end{pmatrix} (M(G_o, C_o) - M(G, C)) \\ &\times \begin{pmatrix} w_{12}I & 0 \\ 0 & w_{22}I \end{pmatrix} \\ &= W_1(M(G_o, C_o) - M(G, C))W_2, \end{aligned} \quad (61)$$

where w_{ij} , $i, j = 1, 2$ are scalar frequency weights. As this expression may seem cumbersome, the reader is recommended to restrict attention to the case

$$\begin{aligned} J(G_o, C_o, G, C) &= G_o S_I(G_o, C_o) C_o - G S_I(G, C) C \\ &= T(G_o, C_o) - T(G, C), \end{aligned} \quad (62)$$

which corresponds to the choice $w_{11}=w_{12}=1$, $w_{21}=w_{22}=0$. Much of the analysis in the identification for control literature has focused on this (1, 1)-block. However, performance robustness is equally important for the other transfer functions.

It is natural to consider the difference between the nominal design $M(G, C(G))$, which usually has some desirable properties, and the corresponding closed-loop transfer function when the designed controller $C(G)$ is applied to the true system, i.e. $M(G_o, C(G))$. For this reason we introduce

$$V(G, C) \triangleq J(G_o, C, G, C) \quad (63)$$

suppressing for notational convenience the dependence of V on G_o . Various norms of $V(G, C(G))$ can be used to quantify the performance degradation when the designed controller is applied to the true system, as compared to the nominal design.

It is also of interest to consider performance degradation relative to the overall optimal design $M(G_o, C(G_o))$ (which would result if perfect knowledge of G_o was available in the control design). This can be measured by norms of $J(G_o, C(G_o), G_o, C(G)) = W_1(M(G_o, C(G_o)) - M(G_o, C(G)))W_2$. The average squared 2-norm of this quantity has, e.g., been used in input design which we will discuss in Section 11.

Notice that when $M(G, C(G)) = M$ is independent of G , then $J(G_o, C(G_o), G_o, C(G)) = -J(G_o, C(G), G, C(G))$ so that the performance degradation measures $V(G, C(G))$ and $J(G_o, C(G_o), G_o, C(G))$ coincide.

Returning to V , notice that it holds that (recall (53))

$$\begin{aligned} V(G, C) &= \\ &= \begin{bmatrix} w_{11}C^{-1} \\ -w_{21}I \end{bmatrix} (I + \bar{\Delta})^{-1} \bar{\Delta} S_I(G, C) [w_{12}C \quad w_{22}I], \end{aligned} \quad (64)$$

where $\bar{\Delta} = \bar{\Delta}(G_o, G, T_I(G, C))$. From this expression we can re-derive the robust stability result (54) and we conclude once more that the size of the weighted input relative error $\bar{\Delta}$ is intimately tied to robust stability.

Regarding robust performance, we can overbound the gain of V in the following way:

$$\begin{aligned} \bar{\sigma}(V(G, C)) & \\ &\leq \bar{\sigma}(\bar{\Delta}(G_o, G, T(G, C))) \\ &\quad \times \bar{\sigma}((I + \bar{\Delta}(G_o, G, T(G, C)))^{-1}) \\ &\quad \times \sqrt{\frac{w_{11}^2}{\bar{\sigma}^2(C)} + w_{21}^2} \\ &\quad \times \sqrt{w_{12}^2 \bar{\sigma}^2(S_I(G, C)C) + w_{22}^2 \bar{\sigma}^2(S_I(G, C))}. \end{aligned} \quad (65)$$

For SISO systems the bound is tight. It is also possible to derive a similar expression that is based on the weighted output relative error $\Delta_O = (G_o - G)G^{-1}T(G, C)$, see, e.g. Morari and Zafriou (1989).

Studying (65), we see that the first two factors depend only on the weighted relative model error $\bar{\Delta}$ while the last two factors depend solely on the nominal design. We shall now discuss each factor separately:

- The first factor is the gain of the weighted relative error. Hence, the relative accuracy of the model, as compared to the designed complementary sensitivity function, can be used to control robust performance
- As we have pointed out above, the second factor is related to stability. We can bound this term

according to

$$\bar{\sigma}((I + \bar{\Delta}(G_o, G, T(G, C)))^{-1}) \leq \frac{1}{1 - \bar{\sigma}(\bar{\Delta}(G_o, G, C))}.$$

The closer to 1 the gain of $\bar{\Delta}$, the poorer robustness margin and the poorer performance. Notice that by combining the above inequality with the inequality (65), a generalization of the robust performance condition (57) is obtained.

- The third factor $\sqrt{w_{11}^2/\sigma^2(C) + w_{21}^2}$ becomes large at frequencies where the controller has a low gain direction. This can occur, e.g., if the system is ill-conditioned with large differences in gain in different directions or in the rare event that the designed closed-loop bandwidth is lower than the bandwidth of the nominal model.
- The first term inside the square root in the fourth factor, i.e. $w_{12}^2\bar{\sigma}^2(S_I(G, C)C)$, is related to the gain from the reference r to the input u in Fig. 9. It is thus a measure of the control effort. Hence, this term is large at frequencies where the bandwidth of the nominal design significantly exceeds the bandwidth of the nominal model.

The second term $w_{22}^2\bar{\sigma}^2(S_I(G, C))$ is the gain of the nominal input sensitivity. This term may become large around the designed bandwidth if, e.g., the designed bandwidth exceeds the limitations imposed by non-minimum phase zeros present in the system.

Notice that $\bar{\Delta}$ is involved both in the robust stability and the robust performance conditions. The magnitude of this factor depends on both the *quality of the model* G and the *control specifications* (represented by $T_I(G, C)$). It is around $\bar{\Delta}$ where the interplay between system identification and control is staged and below we will illustrate some of the trade-offs that it induces by considering the problem of making $|\bar{\Delta}| \ll 1$. For simplicity, we will consider the SISO-case only. Let \hat{G}_N be an identified model and let $C(\hat{G}_N)$ be a controller designed based on this nominal model (and possibly also its confidence region). Since, input and output sensitivities are the same for SISO systems we shall drop the subscript I in quantities such as T_I , etc.

Trade-off 1: Performance specifications versus experimentation effort

Suppose that the set of unfalsified models is given by (36). Then

$$\|\bar{\Delta}\|_\infty \leq \sup_\omega \left| \frac{T(\hat{G}_N, C(\hat{G}_N))}{\hat{G}_N} \right| \times \sqrt{\chi_{\bar{\Delta}}^2(n)\kappa_{n,N} \frac{\Phi_v}{N \cdot \Phi_u}}, \tag{66}$$

with probability α . Here T/\hat{G}_N is the designed transfer function from the reference r to the input u in Fig. 9, which we




Fig. 11. Thick solid line: estimated model \hat{G}_N . Thin dashed line: desired complementary sensitivity function T for design 1. Lower dotted line: desired complementary sensitivity function T for design 2. Thick dashed line: designed control effort T/\hat{G}_N for design 1. Upper dotted line: designed control effort T/\hat{G}_N for design 2.

could call the designed control effort. Let SENPR denote the signal energy density to noise power density ratio $N\Phi_u/\Phi_v$. Then (66) implies roughly that SENPR has to be at least an order of magnitude larger than the squared designed control effort in order to guarantee robust performance.

In order to get some insight in what this implies, notice that in the passband of the designed complementary sensitivity function, $|T/\hat{G}_N| \approx |1/\hat{G}_N|$. Typically G_o , which \hat{G}_N tries to approximate, is of low-pass character with gain larger than one at low frequencies, and hence the designed control effort will be (significantly) less than one for frequencies up to when the system’s own gain drops below 1, i.e. the open loop cross-over frequency. Hence, in this frequency range the model uncertainty may be allowed large while still having a small $\|\bar{\Delta}\|_\infty$. This is the forgiving nature of feedback giving the model builder some slack. Above this frequency region the magnitude of $|T/\hat{G}_N|$ will increase until the bandwidth of T is reached where it will start to decrease if T rolls-off faster than \hat{G}_N . Thus it is in this frequency region where an accurate model will be required. More precisely, for a system that rolls off like $1/\omega^n$, a SENPR of the order at least $\mathcal{O}(\omega^{2n})$ is required in this frequency band. This indicates that increasing the closed-loop bandwidth becomes increasingly expensive with respect to experimentation beyond the open loop system’s own bandwidth. This is illustrated in Fig. 11 where the designed control effort is shown for two different choices of designed bandwidth (designs 1 and 2). As the bandwidth is increased above the bandwidth of the estimated model, the designed control effort increases significantly implying in turn a drastic increase in the minimum SENPR which is proportional to the square of the designed control effort as mentioned above. We shall discuss experiment design further in Section 11.

Another important conclusion is that (66) involves quantities which a priori are unknown. Hence, it may be beneficial to update the experiment design as more information about the system is obtained. This leads to *adaptive experiment design* which we shall discuss in Section 11.3.

In summary, (66) captures the trade-off between the input spectrum and the performance specifications. The

importance of adapting the performance specifications to the system uncertainty is the leading principle in the ‘wind-surfer approach’, see e.g. Lee et al. (1995). Here the performance is gradually increased along with improved confidence in the model. See also Schrama (1992), furthermore see Cadic, Weiland, and Polderman (2003) for a recent discussion on the importance of this issue.

Trade-off 2: Performance specifications versus model complexity

Suppose that in order to limit the complexity of the controller, a low order model G_{lo} is to be used. From the separation principle in Section 4.2 we have that the optimal model is obtained by first estimating a full-order model \hat{G}_{ML} using ML estimation and then solving

$$\hat{G}_{lo} = \arg \min_{G_{lo}} \sup_{\omega} \left| \frac{\hat{G}_{ML} - G_{lo}}{G_{lo}} T(G_{lo}, C(G_{lo})) \right|.$$

The total error $\hat{G}_{lo} - G_o$ can be assessed by also computing the variance for \hat{G}_{lo} . However, already the minimum of the expression above gives useful information. For a given bound on $|\bar{A}|$, a lower bound on the complexity of G_{lo} is obtained. Conversely, given the complexity of G_{lo} , an upper bound on achievable bandwidth for T is obtained.

8. Direct identification of restricted complexity models for control: LTI systems

Since the order of the controller typically depends on the order of the model, it is often desirable to restrict the complexity of the model. Thus there is a need to identify restricted complexity models that are suitable for control. This is the theme of this section. In one line of research the objective has been to develop identification methods which directly identify low order models suitable for control design. A minimum requirement in this context has been methods that work for large data sets. This means that the focus has been on the bias error $|G_o(e^{j\omega}) - G(e^{j\omega}, \theta^*)|$ and variance aspects have to a large extent been neglected, cf. Sections 8.3–8.7.

In order to facilitate the comparison of different methods only SISO systems will be considered. We define as formal objective to minimize the performance degradation as compared to the nominal control design. This means minimizing (recall (63)) $\|V(G, C(G))\|$, where $\|\cdot\|$ is any norm, with respect to G , in other words the identification is linked to the control design via a performance degradation criterion.

For ease of comparison, the simple case (62) is considered throughout Section 8. Notice that $V(G, C)$ then simplifies to

$$\begin{aligned} V(G, C) &= T(G_o, C) - T(G, C) = S(G, C) - S(G_o, C) \\ &= (1 + \bar{A}(G_o, G, T(G, C)))^{-1} \bar{A}(G_o, G, T(G, C)) \\ &\quad \times S(G, C), \end{aligned} \tag{67}$$

where \bar{A} is given by (53) with $T_l = T$ (which holds for SISO systems). We shall occasionally use γ to denote some nominal performance measure, e.g. designed bandwidth.

We will begin this exposé in Section 8.1 by elaborating on how to identify restricted complexity models in an asymptotically efficient way. We will then in subsequent sections review different approaches that have appeared in the literature. In Section 8.9 stability aspects for such performance based methods will be considered. From this discussion we conclude that it is of interest to consider methods which make the weighted relative error (53) small and in Section 8.10 we discuss how to identify near-optimal restricted complexity models with this property. Finally, in Section 8.11 we discuss preferential identification, an alternative way of connecting the identification criterion with robust control.

8.1. Asymptotically efficient identification

From Section 4.2 it follows that it is optimal, with respect to the asymptotic statistical accuracy, to first identify a full-order model using an asymptotically efficient estimator, and then reduce the complexity according to the following procedure:

Let $G(\theta)$ represent a full-order model, i.e. $\exists \theta^\circ$ s.t. $G(\theta^\circ) = G_o$, and let $\hat{\theta}_{ML}$ denote the corresponding ML-estimate. Let $G(\eta)$ represent a restricted complexity model and define

$$\hat{J}(\theta, \eta, \gamma) \triangleq J(G(\theta), C(G(\eta), \gamma), G(\eta), C(G(\eta), \gamma)),$$

where $J(G_1, C_1, G_2, C_2)$ has been defined in (61) (or the simpler (62)). Then, take $\hat{\eta}$ as the minimizer of $\|\hat{J}(\hat{\theta}_{ML}, \eta, \gamma)\|$ with respect to η . The variance of $\|\hat{J}(\theta^\circ, \hat{\eta}, \gamma)\|$ may be estimated numerically using Gauss’ approximation formula (Ljung, 1999b) and implicit differentiation of $\hat{J}(\theta, \eta, \gamma)$. Hence, for given data, the performance specification γ may be adjusted so that achieved and designed performance are guaranteed to be sufficiently close as well as to ensure robust stability.

8.2. High-order modeling and model reduction

It may be useful both from practical and computational points of view to first estimate a high-order model and then perform model reduction. The method ASYM, see Zhu (1998, 2001), is a fully integrated method for identification for control based on this approach and, hence, very close in spirit to the use of the separation principle discussed in Section 8.1. The input design is based on high-order optimal input design; see Section 11.1. A two step approach is used with a high-order ARX-model estimated in the first step, the motivation being its computational simplicity and that the high-order theory is applicable, cf. Section 3.2 and (32). Model reduction is performed using the asymptotic ML method (Wahlberg, 1989) which means minimizing a frequency weighted \mathcal{L}_2 criterion where the weighting corresponds to the inverse of an estimate of the variance of

the estimated high-order frequency function. The model reduction step is thus asymptotically efficient as the sample size *and* the model order tends to infinity. Model order selection is also based on the high order variance expression. To determine the unfalsified set of controllers, (36) is used with $\kappa_{n,N}(\omega) = m$ (=the model order) which follows from (32). The method was developed in the early 1990's and it has been successfully applied to numerous multivariable processes in process industry.

High-order ARX-modeling is also advocated in Rivera and Jun (2000).

8.3. Direct restricted complexity identification

It may not be convenient to identify a full-order model and around 1990 several ideas of how to directly identify a nominal model G which approximately minimizes $\|V(G, C(G))\|_2$ appeared. As we will outline below, the basic idea is to choose the *design variables* in the identification method such that the asymptotic bias expression, e.g. (26), approximates the \mathcal{L}_2 -norm of the error (67).

The quantity $V(G, C)$ in (67) can also be expressed as

$$\begin{aligned} V(G, C) &= T(G_o, C) - T(G, C) \\ &= \bar{A}(G_o, G, T(G, C))S(G_o, C) \\ &= S(G, C)(G_o - G)CS(G_o, C). \end{aligned} \quad (68)$$

Hence,

$$\|V(G, C)\|_2^2 = \int_{-\pi}^{\pi} |G_o - G|^2 |CS(G_o, C)|^2 |S(G, C)|^2 d\omega.$$

Comparing this expression with the closed-loop bias expression (27), we see that $\|V(G, C)\|_2^2$ corresponds to the first term in the closed-loop bias expression (27) if the noise model is taken as

$$H = \Phi_r^{1/2} S^{-1}(G, C). \quad (69)$$

Hence, the model that minimizes $\|V(G, C)\|_2^2$ with respect to G will be obtained asymptotically when the system is noise free if the identification is performed in closed-loop with the controller C using the noise model (69). We will discuss this approach further in Section 8.5.

Before we proceed, observe that the above derivation can also be done in the time domain. Let u and y , denote the input and output, respectively, of the closed-loop system in Fig. 9 with $w \equiv 0$ and the controller C in the loop. The manipulations in (68) correspond in the time domain to

$$\begin{aligned} y - T(G, C)r &= (S(G, C) + T(G, C))y - T(G, C)r \\ &= S(G, C)(y - GC(r - y)) \\ &= S(G, C)(y - Gu). \end{aligned} \quad (70)$$

The expression (70) is the prediction error (16) when data are collected in closed-loop with C as the controller and $H = S^{-1}(G, C)$ as noise model. The only difference compared to (69) is due to that the spectrum of (70) is weighted with the reference spectrum which is not the case in (68).

A simpler cost function is obtained by linearizing (62) with respect to G_o

$$\begin{aligned} J(G_o, C_o, G, C) &\approx ((G_o - G)C_o + (C_o - C)G)S(G, C_o)S(G, C) \\ &\triangleq \bar{J}(G_o, C_o, G, C). \end{aligned}$$

This gives the approximation

$$\bar{V}(G, C) \triangleq \bar{J}(G_o, C, G, C) = (G_o - G)CS^2(G, C) \quad (71)$$

to $V(G, C)$ (63) and instead of minimizing $\|V(G, C(G))\|_2$, $\|\bar{V}(G, C(G))\|_2$ could be minimized (or $\|\bar{V}(G, C)\|_2$ if the direct parametrization (60) is used).

Comparing the \mathcal{L}_2 -norm of (71) with the open loop bias expression (26), we see that open loop identification with the noise model

$$H = \Phi_u^{1/2} (S^2(G, C(G))C(G))^{-1} \quad (72)$$

corresponds to minimizing $\bar{V}(G, C(G))$ under noise free conditions. We will discuss this approach in Section 8.6.

Remark 1. Comparing (71) with (67), we see that the stability guaranteeing term $(1 + \bar{A})^{-1}$ is not present in the approximation $\bar{V}(G, C)$. Furthermore, comparing with the discussion on conditions on $|\bar{A}|$ in Section 7.6, we see that there is less emphasis on making $|\bar{A}|$ small at low frequencies. Hence, there is a possibility that $|\bar{A}| > 1$ at low frequencies when (71) is used, with a potential risk for destabilization.

In Sections 8.4–8.7 we will elaborate on how the expressions above have been used in the literature.

8.4. Model and controller reduction

Expression (71) was, perhaps, first used for model reduction purposes. Given a full-order model G_o and a desired complementary sensitivity function, the problem of finding a control relevant reduced order model such that $\|S(G_o, C(G))\|$ is as small as possible was examined in Rivera and Morari (1987). Through the use of a triangle inequality for the sensitivity function, the objective was simplified into minimizing the norm of $J(G_o, C(G), G, C(G)) = V(G, C(G))$ with respect to G , which in turn was approximated by the norm of $\bar{V}(G, C(G))$.

In Rivera and Morari (1987), the function $\|\bar{V}(G, C(G))\|$ is minimized with respect to G numerically under the assumption that

$$|S(G, C(G))T(G, C(G))|$$

is independent of G , cf. IMC and model reference control (see Section 7.1.3). From this we see that a small *relative* model error is achieved around the designed bandwidth where neither of $S(G, C(G))$ or $T(G, C(G))$ are small.

The problem of controller reduction was also studied in Rivera and Morari (1987), see also Anderson and Liu (1989). Given the high order controller C_o , the objective is to determine C such that some norm of

$$U(G_o, C) \triangleq J(G_o, C_o, G_o, C) \\ = T(G_o, C_o) - T(G_o, C) \quad (73)$$

is small. Linearizing with respect to C gives

$$U(G_o, C) \approx \frac{C_o - C}{C_o} T(G_o, C_o) S(G_o, C_o) \triangleq \bar{U}(G_o, C). \quad (74)$$

A similar procedure as for the model reduction problem is now used.

In fact these model reduction and controller reduction problems are dual to each other: With a direct parametrization of G in terms of the controller C according to

$$G = \frac{1}{C} \frac{T(G_o, C_o)}{1 - T(G_o, C_o)} \quad (75)$$

then

$$U(G_o, C) = -V(G(C), C). \quad (76)$$

Hence, the controller reduction problem $\min_C \|U(G_o, C)\|$ can be seen as the model reduction problem $\min_G \|V(G, C)\|$ where G is parametrized in terms of C according to (75).

A variation of the above controller reduction idea is presented in Landau, Karimi, and Constantinescu (2001). Here, no Taylor approximation is introduced. For the optimization, a recursive algorithm is employed which uses simulated closed-loop signals. Boundedness of the closed-loop signals of the closed loop system with the reduced order controller is guaranteed under a certain positivity condition.

8.5. Iterative methods

In Section 8.3 we have seen that $V(G, C)$ can be minimized with respect to G using closed-loop identification with the controller C in the loop when the noise model (69) is used and data are noise free. Since the objective is to minimize $V(G, C(G))$ this suggests the following iterative procedure. At iteration k :

- (i) Identify a model G_k using the noise model (69) and closed-loop data collected with controller $C_k = C(G_{k-1})$ in the loop.
- (ii) Replace C_k with $C_{k+1} = C(G_k)$ in the closed loop. Let $k = k + 1$ and go to Step (i)

The methods in Zang et al. (1995) and Schrama (1992) are examples of this approach. See also Åström (1993). Surveys of this type of methods can be found in Gevers (1993), Van den Hof and Schrama (1995) and Albertos and Sala (2002).

Despite the intuitive character, the above scheme will not converge to the minimum of $V(G, C(G))$ (Hjalmarsson,

Gunnarsson, & Gevers, 1995). The iterations above correspond to, for some given function $g(x, y)$, trying to minimize the function $f(\theta) = g(\theta, \theta)$ by the iterations

- (i) $\theta_k = \arg \min_{\theta} g(\theta, \theta_{k-1})$.
- (ii) Let $k = k + 1$ and go to Step (i)

Any convergence point θ^* satisfies that the partial derivative of g with respect to the first argument x is zero, i.e. $\partial g(\theta^*, \theta^*) / \partial x = 0$, whereas a necessary condition for optimality is $0 = f'(\theta) = \partial g(\theta, \theta) / \partial x + \partial g(\theta, \theta) / \partial y$. Thus convergence to the optimal solution cannot in general be guaranteed. This problem had earlier on been noted in adaptive control, see Ljung and Söderström (1983), Anderson et al. (1986) and Phillips, Kosut, and Franklin (1998). It has also turned out to be difficult to establish convergence of the above procedure. Despite these shortcomings, this type of method has been found useful and successful applications exist, see e.g. Partanen and Bitmead (1995), Schrama and Bosgra (1993), Albertos and Sala (2002). In Section 8.11 we will see that the uncertainty model unfalsification approach allows convergent iterative methods to be developed.

The above derivation assumed noise free data. In the case of noisy data, the bias of the direct prediction error method cannot be tuned at the user's will by the use of a prefilter/noise-model. This has spurred the development of a number of closed-loop identification methods for which this is possible. We refer to Van den Hof (1998), Forsell and Ljung (1999) and Landau, Karimi, and Constantinescu (1997) for details.

8.6. Prefiltering methods

In Section 8.3 we introduced the approximation $\bar{V}(G, C)$ (71) to $V(G, C)$ (63) and showed that under noise free conditions $\|\bar{V}(G, C(G))\|_2$ can be minimized asymptotically with prediction error identification if the noise model is chosen as (72). For noisy data, the bias expression (26) is no longer valid as the noise model (72) is not independently parametrized of G (Forsell & Ljung, 1999). However, this prefiltering method is simple to use as it requires no special experimental conditions and can be expected to work if the signal-to-noise ratio is high and some care regarding the designed bandwidth is exercised, see Remark 1 in Section 8.3.

The above idea is the basis in Rivera, Pollard, and Garcia (1992) where, for numerical reasons, an iterative procedure is proposed where the estimated model from iteration $k - 1$ is used in the noise model at iteration k .

In virtual reference feedback tuning (VRFT) (Campi, Lecchini, & Savaresi, 2002) data is used to optimize the controller such that $\|T(G_o, C) - T_d\|_2$, where T_d is some fixed reference model, is minimized. This means that $\|U(G_o, C)\|_2$, see (73), is minimized with $T(G_o, C_o) = T_d$. The approximation $\bar{U}(G_o, C)$, see (74), is used and it is observed that $\bar{U}(G_o, C)$ can be rewritten as $(1 - T_d)T_d(C(T_d^{-1}G_o - G_o) - 1)$. The \mathcal{L}_2 -norm of this expression can thus, in the limit $N \rightarrow \infty$, be minimized

by minimizing the sum of squares of

$$\varepsilon = \Phi_{uu}^{-1/2} (1 - T_d) T_d (C(r_{T_d} - y) - u) \quad (77)$$

with respect to C for an arbitrary noise free data set Z^N . The signal $r_{T_d} = T_d^{-1}y$ is called the virtual reference signal.

At first glance this procedure might seem unrelated to the prefiltering approach discussed in this section. However, the relation (76) between $U(G_o, C)$, used in VRFT, and $V(G, C(G))$, used in the prefiltering approach, suggests that there is a close connection. In fact, with the direct parametrization (60) we see that VRFT corresponds to minimizing $\|\bar{V}(G(C), C)\|_2$, i.e. VRFT can be interpreted as a prefiltering approach using the direct parametrization (60).

A variant of the prefiltering approach is proposed in Holmberg, Valentinotti, and Bonvin (2000). Data are collected in closed-loop with some controller C and from this data set a model of the complementary sensitivity function is estimated from which an estimate $\hat{S}(G_o, C)$ of the sensitivity function $S(G_o, C)$ can be obtained. Based on (68), $S(G, C)C\hat{S}(G_o, C)$ is then used as prefilter in the identification. Notice that there is no guarantee either that this procedure will find the model that minimizes $V(G, C(G))$. Based on the simplifying assumption that the model error is only due to noise, an attempt to take the model uncertainty into account in the control design is done in Holmberg et al. (2000).

8.7. Data dependent prefilers

A non-parametric estimate of G_o is given by $\hat{\Phi}_{yr}/\hat{\Phi}_{ur}$ where $\hat{\Phi}_{xr}$ is a non-parametric estimate of the cross-spectrum between r (see Fig. 9) and x . Now, replacing the true G_o by this estimate in (62) gives the cost function

$$\begin{aligned} & \left\| J \left(\frac{\hat{\Phi}_{yr}}{\hat{\Phi}_{ur}}, C(G), G, C(G) \right) \right\| \\ &= \left\| S(G, C(G))C(G) \frac{1}{1 + \hat{\Phi}_{yr}/\hat{\Phi}_{ur}C(G)} \right. \\ & \quad \left. \times \left(\frac{\hat{\Phi}_{yr}}{\hat{\Phi}_{ur}} - G \right) \right\|. \end{aligned} \quad (78)$$

It can be shown that this cost function asymptotically (in N) converges to $\|J(G_o, C(G), G, C(G))\| = \|V(G, C(G))\|$ and that consequently the minimizing G also asymptotically minimizes the desired objective, provided that the spectral estimates are consistent (Hjalmarsson & Lindqvist, 2001). The method can be seen as a simplified version of the optimal method outlined in Section 8.1 with the non-parametric estimate replacing the asymptotically efficient ML-estimate.

The cost function (78) can be expressed in the time domain as

$$\|(\mathcal{T}_{NM}^T(w)\mathcal{T}_{NM}(r))^{-1}\mathcal{T}_{NM}^T(\varepsilon_L(\theta))\mathcal{T}_{NM}(r)\|, \quad (79)$$

where

$$\varepsilon_L(\theta) = S(y - G(\theta)u), \quad w = \frac{1}{C}u + y,$$

and where $\mathcal{T}_{NM}(z)$ is defined as the $(N + M - 1) \times M$ lower Toeplitz matrix with $[z(0), \dots, z(N - 1), 0, \dots, 0]^T$ in the first column. Using the Frobenius norm in (79), results in that the \mathcal{L}_2 -norm is obtained asymptotically in (78) whereas the 2-norm will result in that the norm in (78) is the \mathcal{L}_∞ -norm (Massoumnia & Kosut, 1994).

We remark that with the cost function (79) as starting point it is not obvious that this method actually is a two step procedure involving a non-parametric estimate of the true system in the first step.

8.8. Minimizing average performance degradation

Up till now we have in this section assumed that the same controller $C(\hat{G}_N)$ is used on the true system as on the nominal design. It is possible to obtain an achieved performance closer to the nominal performance $T(\hat{G}_N, C(\hat{G}_N))$ by replacing $C(\hat{G}_N)$ with the controller C which minimizes the average performance deterioration

$$\mathbf{E}\{J(G_o, C, \hat{G}_N, C(\hat{G}_N))\},$$

where the expectation is over G_o in the set of unfalsified models. This has been explored in Goodwin, Wang, and Miller (1999), see also Goodwin, Graebe, and Salgado (2001).

8.9. Stability, performance and robustness

It may be tempting to conclude that since the nominal design $M(G, C)$ is stable and since $\|V(G, C)\|$ (67) is minimized, the achieved closed-loop system $T(G_o, C)$ is stable. The truth of this, however, depends on the norm that is used. In relation to (67) we note that \mathcal{H}_2 - or \mathcal{H}_∞ -norms are suitable. However, the methods discussed previously in this section are based on frequency domain expressions that do not take stability into account, i.e. the norms are of \mathcal{L}_2 or \mathcal{L}_∞ type. Hence, even if we, as in the case of data-dependent prefiltering, are able to generate the ideal cost function (67), we get no information whatsoever whether the minimizing controller stabilizes the true system or not, with one exception and that is when we can match $T(G_o, C)$ with $T(G, C)$ perfectly. One exception to this is the method in Landau et al. (2001) which uses signals generated by closed-loop simulation.

We will use a simple example to illustrate this using the method presented in Campi et al. (2002). It should be stressed though, that the problem is not specific to this method⁴ but *intrinsic* to all methods which solely aim at minimizing $\|V(G, C(G))\|$.

⁴This method is chosen because of the simple calculations.

Example 8.1. Let the true system be a pure time-delay $y(t) = u(t - 1)$. Suppose that a simple proportional controller $C = \rho$ is used on the time-delay system. Notice that the closed-loop is unstable for $|\rho| > 1$.

Let the desired complementary sensitivity function be $T_d(e^{j\omega}) = 1 - \delta + \delta e^{-j\omega}$ where the parameter δ will be discussed later. Clearly, T_d cannot be achieved no matter what controller is used unless $\delta = 1$.

Consider now the error signal (77). It can be shown that when the input is white noise with zero mean and unit variance, then $E[\varepsilon^2(t, \rho)]$ is minimized by

$$\rho_{\text{opt}} = \frac{4\delta - 1}{6\delta}.$$

Hence, for $0 < \delta < 0.1$, $|\rho_{\text{opt}}| > 1$ and the corresponding closed loop system will be unstable. Hence stability is not ensured when the desired closed-loop cannot be matched by the actual closed-loop.

We conclude from Example 8 that the performance specifications have to be *matched* to the approximating ability of the model and that to ensure stability some robust stability condition has to be included. This in turn means that it is not sufficient for a method to be able to tune the bias error optimally, a method also has to be able to provide bounds for the model error and these have to be accounted for in the control design.

As we saw in Section 7.6 (Trade-off 1), it becomes typically increasingly expensive to obtain sufficient process information to ensure robust performance (and stability) as the bandwidth is increased. Hence, it is natural to gradually increase the specified bandwidth. This was recognized in Lee et al. (1993) where the term “the windsurfer approach”, alluding to how a windsurfer gradually improves his/hers performance, was coined.

An early reference to iterative identification and *robust* control is Bayard, Yam, and Mettler (1992). Another contribution is de Callafon and Van den Hof (1997). The issue of ensuring that the collected information guarantees improved performance is not addressed (Eq. (4) in de Callafon & Van den Hof (1997)), instead considerable effort is spent on how to select uncertainty descriptions and the control design so as to allow computationally feasible solutions. Interesting applications of this method are reported in de Callafon and Van den Hof (2001) and Bøjstrup, Niemann, Kjølstad Poulsen, and Jørgensen (2003).

Another way to cope with the robustness issue is to introduce caution in controller updates. Assuming the present controller to be stabilizing (but providing unsatisfactory performance), the increment in the controller update is limited so as to ensure stability also with the new controller (Anderson, Bombois, Gevers, & Kulcsar, 1998; Gevers, 2000; Bitmead, 2000).

In Section 8.11 we will discuss a quite different approach based on robust control.

8.10. Control relevant near-optimal models

In Section 7.1.2 we saw that using a good nominal control design combined with a model that makes the magnitude of the weighted relative model error \bar{A} (defined in (53)) sufficiently small leads to robust performance, cf. (56). Similar conclusions were drawn in Section 7.6. We shall now pursue this idea.

For a full-order model, $|\bar{A}|$ is, frequency by frequency, bounded by (see (66))

$$\left| \frac{T(\hat{G}_N, C(\hat{G}_N))}{\hat{G}_N} \right| \sqrt{\chi_x^2(n) \kappa_{n,N} \frac{\Phi_v}{N \cdot \Phi_u}}. \quad (80)$$

For a near-optimal model⁵ the bound has to be multiplied by a factor 2.

Now consider the scenario that the data set Z^N is given and it is required that $|\bar{A}(\omega)|$ is bounded by some function. Then, firstly in order to have a handle on the model error frequency by frequency, the model complexity has to be chosen such that the model is near-optimal. Furthermore, the remaining variable at the user’s disposal is T which has to be chosen sufficiently small that the condition on $|\bar{A}|$ is satisfied.

When the input design is at the user’s disposal, there is considerably more freedom for the user. For given performance specifications and with no restrictions on the model complexity, (iii) in Section 6 gives that the input should be designed such that (80) is less than one but not smaller than necessary. This will give data that are sufficiently informative for the given performance specifications while at the same time minimizing the modeling requirements.

8.11. Preferential identification

So far in this section, the objective has been to ensure that the identification is performed such that the difference in performance between the nominal design (using the identified model) and the actual design is small, as measured by the function V . When the control design minimizes an explicit criterion, it is natural to instead study the problem of how to identify a model \hat{G}_N such that this criterion is minimized for the true system when the controller based on \hat{G}_N is used. This problem can be addressed using the separation principle in a way similar to what was done in Section 8.1. However, we will not pursue this idea further but instead discuss an interesting approach pioneered in Krause, Stein, and Khargonekar (1992) for adaptive control.

A bit simplified, the idea in Krause et al. (1992) can be summarized as follows:

- (i) *Use a robust control law*: assume that there is a robust control law $C(\eta)$ such that when the system signals

⁵ Recall, see Section 4.4, that a near-optimal model is any model within the confidence region of a full-order model.

$\{y(t), u(t)\}_{t=t_o}^N$ in a closed-loop experiment with this controller are consistent with certain modeling assumptions parametrized by some vector η (t_o is a time offset to allow for a transient phase where the assumptions are not satisfied), then a certain closed-loop performance $\mathcal{Y}(\eta)$ is guaranteed for this experiment.

One example of modeling assumptions is that

$$\left| y(t) - \sum_{k=1}^m g_k u(t-k) \right| \leq \gamma \sqrt{\sum_{k=1}^n |u(t-k)|^2} + c, \quad (81)$$

with $n \geq m$, and $\eta = [g_1, \dots, g_m, \gamma, c]^T$. Notice that the bound γ on the unmodeled dynamics and the noise bound c are not assumed known. Notice also that, *and this is important*, the assumption regarding the control law does not mean that the system signals have to satisfy the modeling assumptions for any input; they just have to be consistent with the assumptions for the experiment at hand.

- (ii) *Use uncertainty model unfalsification*: use an identification method which is able to estimate η on-line and in closed-loop such that the estimate, η_t say, converges in finite time $t = t_o < \infty$ with the estimate η_{t_o} being consistent with $\{y(t), u(t)\}_{t=t_o}^t$.

When (i) and (ii) are satisfied, it follows immediately that for the *adaptive controller* $C(\eta_t)$, performance for $\{y(t), u(t)\}_{t=t_o}^t$ is guaranteed by $\mathcal{Y}(\eta_{t_o})$.

Remark 1. Notice that the obtained η_{t_o} need not correspond to a “true” η , i.e. it does not have to be valid for all possible inputs to the system; from (i) it suffices that it holds for the particular experiment. The same η_{t_o} does thus not guarantee performance $\mathcal{Y}(\eta_{t_o})$ in a new closed-loop experiment. It may not even guarantee stability.

Remark 1 points to that this is a quite different approach from what we would normally perceive as robust control. This seems to be the price paid for not using a priori knowledge of a set to which the true system belongs.

There is considerable freedom in choosing the estimation algorithm in (ii) and this freedom can be used to achieve the same performance guarantee as if the true system was known. In *preferential* identification (a term coined in Krause et al., 1992), one chooses the model which *promises the best performance* for the controller:

$$\eta_t = \arg \min_{\eta \in M_t} \mathcal{Y}(\eta), \quad (82)$$

where M_t denotes a set which includes all η that cannot be falsified by data up to time t .

Now if (ii) holds for the estimator (82), it holds that performance is guaranteed by $\mathcal{Y}(\eta_{t_o})$. Suppose now that η_o is a valid description of the true system, meaning that (81) holds regardless of which input is applied. Then $\eta_o \in M_t$,

$t = 1, 2, \dots$, and in view of (82) it holds that $\mathcal{Y}(\eta_{t_o}) \leq \mathcal{Y}(\eta_o)$. The actual performance guarantee after the estimates have converged is thus *at least* as good as if a valid η had been known from the beginning and used for the control design. Clearly, a very powerful result.

The key issue in the approach is to ensure convergence of the algorithm. For a generalization of (81), a dead-zone is used in the parameter update in Sokolov (1996). This ensures convergence in finite time without any persistence of excitation conditions. Notice that the bounds γ and c are not known in advance! The dead-zone implies a slight degradation of the performance guarantee. One shortcoming of this algorithm is that the adaptation horizon t_o is not known in advance, so that one may have to wait a long time until convergence and during this transient phase no performance guarantees are available.

There are various ramifications of the above idea, e.g. Veres and Sokolov (1998) and Sokolov (2001). Furthermore, Veres has adapted the concept to iterative identification and control (Veres, 1999; Veres & Wall, 2000; Veres, 2001). Computational and complexity issues are the topics of Veres, Messaoud, and Norton (1999) and Xia and Veres (1999).

To conclude, in contrast to the methods in Section 8.5, the robust control/uncertainty model unfalsification/preferential identification paradigm has, subject to the limitations discussed above, proven able to produce iterative identification and control methods which converge with performance guarantees.

9. Direct identification of restricted complexity models for control: non-linear systems

There is abundant practical evidence that LTI models often are sufficient for control design for non-linear systems. In this section we will discuss some related issues. We will limit attention to SISO systems.

From a system identification perspective it is of interest to know how much uncertainty robust (and adaptive) control can handle. In a series of paper, partly summarized in Guo (2002), Guo and co-workers have explored this topic.

9.1. Performance aspects

Let us assume that the system G_o in Fig. 9 is non-linear and noise-free and also that $w \equiv 0$. Let us also assume that a linear model G is used to design an LTI model reference controller (59). We will now discuss how to find a suitable model G such that the non-linear feedback system consisting of the non-linear system and the linear controller defined by (59) responds to r with the desired response $y_d = T_d r$. In Sections 8.1–8.3 we discussed how restricted complexity models for LTI systems could be identified in closed-loop. When a non-linear full-order model is available, the method outlined in Section 8.1 can be adapted to the non-linear setting. When this is not the case, the ideas in Section 8.3 of closed-loop identification can be used. Notice that no use

of that the system is LTI was made in the derivation (70). Hence, for any system,

$$y - T_d r = y - y_d = (1 - T_d)(y - Gu), \quad (83)$$

where u and y are the closed-loop signals with the controller $C = C(G)$ defined in (59) in the loop. This expression for $y - y_d$ was derived in this non-linear setting in [Henriksson, Markusson, and Hjalmarsson \(2001\)](#) and similar observations have been made in [Horowitz \(1992\)](#). Thus, if there is a model G^* such that the corresponding model $C(G^*)$ results in $y = y_d$, then the right-hand side of (83) will be zero when $G = G^*$ and when data is collected with $C(G^*)$ as the controller in the loop. This implies that the model G^* corresponding to the desired controller will be obtained in closed-loop identification when the desired $C(G^*)$ is operating in the loop and if the prefilter $1 - T_d$ is used in the identification ([Henriksson et al., 2001](#)).

Notice that since the system is non-linear it may very well happen that G^* is non-causal in which case the above idea breaks down.

Disregarding this, the discussion above supports the intuitively appealing idea that the identification experiments should be carried out under the *desired* operating conditions. A limitation of the argument above is that it is based on studying one single trajectory and, hence, does not give any information about the behavior for other reference signals, and in particular of closed-loop stability. We will return to how to generate the desired operating conditions in Section 11.3.2 where also an example is presented.

9.2. Control relevant near-optimal models

In Section 8.10 the approach was to ensure robust stability and robust performance by making $|\bar{A}|$ small. The motivation was the arguments brought forward in Section 7.6 which in turn were based on the factorization (64) (or the simpler expression (67)). Now this factorization generalizes to the non-linear case ([Eker & Nikolaou, 2002](#)). For example, (67) still holds where now G_o is to be interpreted as a non-linear operator. Furthermore, bounds such as (65) apply if the largest singular value is replaced by some induced norm (In [Eker & Nikolaou \(2002\)](#) it is suggested to use an induced differential norm as this reduces the conservatism). For example, the norm of \bar{A} is bounded by

$$\|T(G, C)W\|_\infty \left\| W^{-1} \frac{G_o - G}{G} \right\|_{i,2} < 1,$$

where the second norm is the induced \mathcal{L}_2 -norm

$$\|G\|_{i,2} = \sup_{u \neq 0} \frac{\|G(u)\|_2}{\|u\|_2}, \quad (84)$$

where $\|u\|_2 = \sqrt{\sum_k |u(k)|^2}$. Above W is an arbitrary linear⁶ weighting filter.

⁶ Invertible non-linear operators can also be used.

This means that when a linear model and linear control design is sought for a non-linear system, the minimization of the weighted relative error is still a valid criterion.

In Section 4.4 we discussed near-optimal LTI models of systems that were themselves LTI. The conclusions from this section also extend to non-linear systems. This means that for a non-linear parametric model structure, to which the true system belongs, the simplest unfalsified model can be used with the total error being less than two times that of the full-order model structure. Thus, if there is an LTI model within the confidence region for the full-order model, this model can be used. If it exists, such an LTI model can be obtained by direct least-squares identification as described in Section 4.4. We remark that it may be non-trivial to compute the maximum error gain between the model corresponding to the center of the full-order confidence region and a model inside this region. This is, however, an issue more related to non-linear systems analysis than system identification.

From the above follows that the approach in Section 8.10 carries over to the non-linear setting.

In [Ljung \(2000\)](#) the choice of weighting filter and how to estimate the non-linear gain $\|W^{-1}(G_o - G)/G\|$ are discussed. It is pointed out that only lower bounds can be obtained from data and that periodic inputs can be useful since they allow the noise to be averaged out. Another preliminary contribution in this important area is [Schoukens, Pintelon, and Dobrowiecki \(2002\)](#) and in [Mosskull, Wahlberg, and Galic \(2003\)](#) these ideas are used to verify stability of an induction machine drive.

This issue has also spurred activities in assessing how “small” nonlinearities may influence parameter estimates based on linear models. In [Enqvist and Ljung \(2002\)](#) it is illustrated that LTI-models may be extremely sensitive to nonlinearities. In [Schoukens, Dobrowiecki, and Pintelon \(1998\)](#) a general framework is developed for analyzing how linear estimates are affected by nonlinearities of the system when the input excitation is periodic. For a concise survey of this framework see [Schoukens, Pintelon, Dobrowiecki, and Rolain \(2003\)](#). Best LTI-approximants for non-linear systems are discussed in [Mäkilä & Partington \(2003\)](#) and [Enqvist & Ljung \(2003\)](#).

9.3. High-order modeling

A three step procedure is proposed in [Ling and Rivera \(2001\)](#) for identification of control relevant non-linear models. First a non-linear ARX model is estimated. This model is then transformed into a Volterra series model since such a model is more applicable to existing control design methods. Finally, a non-linear counterpart to the model reduction procedure discussed in Section 8.4 is performed.

10. Direct controller tuning

In this section we will discuss two controller tuning methods, unfalsified control and Iterative Feedback Tuning

(IFT), where data is mapped directly into the controller. Another method with this property is virtual reference feedback tuning, cf. Section 8.6. As for VRFT, unfalsified control and IFT were derived without any explicit use of models. However, it is instructive to, as was done for VRFT, also interpret these methods in a modeling framework and this is the objective of this section.

10.1. Unfalsified control

Suppose that some apparatus is going to be constructed that, when applied to the system, performs a certain task and that, given only input/output data Z^N from the system and no prior system information whatsoever, we would like to verify if the designed system satisfies some performance specifications. To be specific, let us consider the problem of testing whether a certain controller C satisfies some given performance specifications. Well, if the controller satisfies the performance on at least one of the models in the set of unprejudiced unfalsified models $\mathcal{G}(Z^N)$ (see Section 2.2), then the controller C cannot be discarded since that particular model may correspond to the true system in which case the controller would satisfy the specifications. We would in this case say that the *controller is unfalsified*. However, since the models in $\mathcal{G}(Z^N)$ have completely arbitrary input/output behavior, except for the specific input–output trajectory defined by Z^N , which they all share, it is possible to find an unfalsified model such that the closed-loop system consisting of this model and C satisfies every specification that is not violated when the closed-loop system with C as controller exhibits the input/output behavior Z^N . As we will see, this makes it often very simple to check whether a given controller can be falsified or not by data only. The idea of unfalsified control was introduced by Safonov and co-workers (Safonov & Tsao, 1994, 1997), see also Kosut (1995).

To illustrate the machinery suppose that the input is generated by a one-degree of freedom controller

$$u = C(r - y), \quad (85)$$

where r is an external reference signal. Suppose also that a reference model T_d is given and that our performance specifications are

$$\|y - T_d r\|_2 \leq \bar{\rho} \|r\|_2, \quad \forall r : \|r\|_2 < \infty \quad (86)$$

for a given constant $\bar{\rho}$. We now ask the following question: Given arbitrary input/output data Z^N , what can be said about which controllers satisfy (86)? The key to resolving this problem is to note that it follows from (85) that for arbitrary input/output data u, y from the system, the signal

$$r_C = \frac{1}{C} u + y \quad (87)$$

is the reference signal r which with C in the loop would produce exactly the input/output data u, y . Hence, given arbitrary input/output data Z^N , we can think of the corresponding $\{r_C(t), y(t)\}_{t=1}^N$ as a (fictitious) closed-loop

data set with the controller C in the loop. The controller is falsified precisely when the specification (86) is not satisfied for the *specific* signal pair $\{r_C(t), y(t)\}_{t=1}^N$.

By, for each time instance, selecting the controller in the set of unfalsified controllers according to some selection rule, an adaptive unfalsified controller is obtained (Cabral & Safonov, 2003). It has, e.g., been suggested to pick the *most promising controller* according to the criterion function which governs the unfalsification process (Safonov & Cabral, 2001), cf. the discussion on preferential identification in Section 8.11 where the *most promising model* was selected.

Unfalsified control can be given an even firmer relation to model based control. Under LTI assumptions, the unfalsified controller that minimizes (86) restricted to $r = r_C$ (87), corresponds to the prefiltering approach described in Section 8.6 with the \mathcal{L}_∞ -norm employed (instead of the \mathcal{L}_2 -norm used in VRFT) and the direct parametrization (60) of the model (Hjalmarsson & Lindqvist, 2001).

10.2. Iterative Feedback Tuning

Consider the problem of minimizing the norm of $V(G(C), C)$ (67) when model reference control is used, i.e. we are trying to find a reduced order controller C that makes the complementary sensitivity function as close as possible to the fixed reference model T_d . For simplicity, assume that the system is noise free and $w \equiv 0$. Parseval's formula gives

$$\begin{aligned} \|V(G(C), C)\|_2^2 &= \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{t=1}^N ((T(G_\circ, C) - T_d)r(t))^2 \end{aligned} \quad (88)$$

with r being white noise. One approach is to minimize (88) numerically using some descent algorithm such as Gauss–Newton. This was a popular approach in the 1950s and early 1960s.

For this, the sensitivity, i.e. gradient, of $(T(G_\circ, C) - T_d)r$ with respect to C is required (or rather the parameters of C but we will omit this from the discussion). Straightforward differentiation gives

$$\begin{aligned} \frac{d}{dC} (T(G_\circ, C) - T_d)r &= C^{-1}T(G_\circ, C)(1 - T(G_\circ, C))r \\ &= C^{-1}T(G_\circ, C)(r - y(C)), \end{aligned} \quad (89)$$

where $y(C)$ denotes the output of the closed-loop system (47) with controller C in the loop, assuming $w \equiv 0$. Thus we see that the above sensitivity can be obtained from the closed loop system (9) under noise-free conditions with controller C in the loop by (1) first performing an experiment with r as reference and collecting the output $y(C)$ and (2) using $r - y(C)$ as reference in a new experiment whose output is filtered through C^{-1} . In Narendra and Streeter (1964) an alternative approach where (1) and (2) were done simultaneously through an outer feedback with long time-delay was proposed.

Iterative feedback tuning (IFT) (Hjalmarsson, Gevers, Gunnarsson, & Lequin, 1998; Hjalmarsson, 2002) is a generalization of the idea above. The sensitivities of the closed loop signals with respect to the controller parameters are computed from two closed-loop experiments as outlined above. It can be shown that, even in the presence of noise, the signal sensitivities are unbiased (modulo transient effects) and hence it is possible to guarantee that any convergence point of the algorithm corresponds to a *stationary point* of the desired objective function by the use of a stochastic approximation algorithm. The idea can be applied to any, differentiable, signal based objective function, i.e. not only (88). IFT has been applied by the chemical multinational Solvay SA for tuning of PID loops in distillation columns and evaporators (Hjalmarsson et al., 1998).

From (89) we may also deduce other ways to approximate the sensitivity. The separation principle in Section 4.2 gives that if a good estimate \hat{T} of $T(G_o, C)$ is available, then a good sensitivity estimate is obtained by replacing $T(G_o, C)$ by \hat{T} everywhere in the middle expression of (89).

In De Bruyne and Carrette (1997) and Kammer, Bitmead, and Bartlett (2000) it is suggested to avoid the second experiment by replacing $T(G_o, C)$ in the right-hand side expression of (89) by an estimate obtained using closed-loop data when C is operating in the loop. In De Bruyne and Carrette (1997) a parametric model is used whereas a non-parametric model is employed in Kammer et al. (2000).

The use of signal sensitivities can be seen as local modeling (around the current controller parameters) of how the closed-loop signals depend on the controller. This has the important implication that IFT is able to cope with certain nonlinearities also, cf. Hjalmarsson (1998) and Sjöberg and De Bruyne (1999). See also Sjöberg et al. (2003) for an algorithm tailored especially for non-linear systems.

10.3. De-correlation

An interesting approach is presented in Karimi, Miskovic, and Bonvin (2003) where the controller is tuned such that $y - y_d$, where y_d is the desired response, is un-correlated with an instrumental vector which is a function of lagged values of the reference r . It is shown that this can be done iteratively as in IFT but that only one experiment is required per iteration.

11. Experiment design

The reader may have noticed that input design has been a recurring theme up to now. Let us recapitulate:

- When the true system is in the model set, (36) gives a bound on the frequency responses for unfalsified models which depends on the input through the input spectrum as well as through $\kappa_{n,N}(\omega)$, cf. (34).
- In Section 5.3 we saw that (36) also determines the strength of any validation statements that can be made.

- From (iii) in Section 6 we may conclude that an ideal experiment should reveal exactly the information required for the control design.
- In Section 7.6 we saw explicitly in (66) how the, for control, important weighted model error (53) depends on the experimental conditions.
- In Section 4.4 it was illustrated that by choosing the input suitably, the statistical properties of restricted complexity estimators may be similar to those of the ML-estimate. One may think of this as that the model bias is tuned by the data. This relates to our observations in Section 9.1 where it was indicated that using data from the desired operating conditions is very useful when trying to identify a restricted complexity model for a non-linear system. Going back to Section 8.5, we see that this is exactly the reason why closed-loop identification may help to obtain “control relevant” models for LTI systems as well.

In this section we shall further discuss how experiment design can be used to improve the closed-loop performance.

In Section 11.1 we first discuss input design for optimal average performance. We argued in Section 7.2 that the very high complexity of an optimal input design problem for a typical control application makes such problems intractable (at least with existing mathematical machinery). However, by considering the weighted relative model error, introduced in Section 7.6, tractable problems can be obtained. We have already touched on this issue in Section 8.10 and we will pursue this in Section 11.2 by taking the input design into account.

Another impeding factor in the use of optimal experiment design is that the optimal solution depends on the *unknown* true system. This problem could be handled by input designs which are robust against uncertainty about the underlying system were it not that also here computational complexity puts severe restrictions on what can be achieved. It also seems as if robust input design has received very little attention. Apart from the very modest contribution in Hjalmarsson and Jansson (2003), the author is not aware of any contributions to this area and we will discuss this topic very briefly in Section 11.2.

An alternative, or complement, to robust input design is adaptive input design where the information gathered in the experiment is continuously, or batch-wise, used to update the input design. This is discussed in Section 11.3. Related to this is the use of pre-tests to gain some preliminary information for further experiment design. This is the topic in Section 11.4.

Experiment design for multivariable systems is briefly discussed in Section 11.5. Finally, Section 11.6 is concerned with certain aspects of input design for near-optimal models.

A very important part of the experiment design is to ensure that the input is “plant-friendly”. This means that, apart from generating informative experiments, the input should be in line with industrial demands. We will not venture further into

this topic but the reader is referred to Rivera, Lee, Braun, and Mittelmann (2003) and Parker, Heemstra, Doyle Pearson, and Ogunnaike (2001) and references therein.

11.1. Optimal average performance input design

One approach to input design is to minimize the average of some performance degradation measure with respect to the input, subject to constraints on the input and output spectra. The constraints have traditionally been of \mathcal{L}_2 type. One may, e.g., use (recall (63))

$$\mathbf{E}\{\|V(\hat{G}_N, C(\hat{G}_N))\|_2^2\}, \quad (90)$$

where the expectation is over \hat{G}_N , as design criterion. An approximate solution can be derived using the first order approximation (71) of $V(\hat{G}_N, C(\hat{G}_N))$, i.e. by instead minimizing

$$\mathbf{E}\{\|\bar{V}(\hat{G}_N, C(\hat{G}_N))\|_2^2\}. \quad (91)$$

This leads to a so called L -optimal input design problem where the weighted trace of the covariance matrix (20) should be minimized. Here the weighting matrix depends on the problem formulation.

It has been more common to consider

$$E\{\|J(G_o, C(G_o), G_o, C(G))\|_2^2\} \quad (92)$$

as criterion, which again via a Taylor approximation leads to an L -optimal input design problem. As mentioned in Section 7.6, the two criteria (90) and (92) coincide when $M(G, C(G))$ is independent of G . One such example is model reference control. Also minimum variance control of minimum phase systems fits into this category if, as is relevant for this problem, J in (62) is replaced by the difference in transfer functions from the white noise disturbance e in (25) to the output between the true system and the model.

Starting with the now classical reference (Gevers & Ljung, 1986), there have been a series of contributions to the above problem based on the variance expression (28) using the high-order approximation (32), and its closed-loop counterpart, with specific applications to identification for control. When there is a constraint on the output variance, the optimal experiment is to use a controller which is the solution to an LQG-problem determined solely by the constraints (Forsell & Ljung, 2000). Thus applications with different objective functions but with the same constraints share the same optimal experiment. Hence, the optimal experiment is typically in closed-loop in this case. However, it is worth noticing that when the noise is white, the solution to the associated LQG-problem is open loop operation when only the output variance is constrained.

Instead of relying on (32), the variance expression (29) can be used directly for input design (Cooley & Lee, 2001; Lindqvist & Hjalmarsson, 2001). In Cooley and Lee (2001) it is suggested to optimize directly over the input sequence whereas in Lindqvist and Hjalmarsson (2001) it is shown

that when u is the output of a FIR-filter driven by white noise, the optimal filter coefficients can be computed by convex optimization. Since any spectrum can be approximated by a FIR-process the approach is in principle generally applicable. However, when large lags are required, computational complexity becomes an issue. One can then use other basis functions (Jansson & Hjalmarsson, 2004c). Designs for FIR systems in the case of periodic or finite samples are discussed in Lee (2003) and Jansson and Hjalmarsson (2004c).

11.2. \mathcal{H}_∞ and robust input design

The designs in the preceding section are all geared towards optimizing the average performance as they are based on (91) (or (92)). By instead using the bound (66) on the weighted relative model error \bar{A} as design criterion, an \mathcal{H}_∞ type of design is obtained with guaranteed robust stability and robust performance (Hjalmarsson & Jansson, 2003).

A typical problem formulation could be (recall (21))

$$\begin{aligned} \min_{\Phi_u} \quad & \beta \\ \text{subject to} \quad & \left| T \frac{G_0 - G(\theta)}{G(\theta)} \right| \leq \gamma_P \quad \forall \omega, \\ & (\theta - \theta_o)^T \frac{R_N}{\lambda} (\theta - \theta_o) \leq \chi_\alpha^2(n), \\ & \frac{1}{2\pi} \int_{-\pi}^{\pi} \Phi_u(\omega) d\omega \leq \beta. \end{aligned} \quad (93)$$

Here the constraints state that the weighted relative error (recall (53)) $|\bar{A}|$ should be smaller than some pre-specified value γ_P for all frequencies ω and for all models in the confidence region (23) (which is given by the second constraint). Thus the minimum input energy required to meet this objective, and the corresponding input spectrum, are sought. This type of problem has been coined “least-costly identification experiments” in Bombois, Scorletti, Van den Hof, and Gevers (2004). Following (iii) in Section 6, we have that the bound γ_P should be made small, but not too small in order to limit the modeling requirements.

It turns out that this type of problem can be approximated by a convex optimization problem that gives solutions close to the optimum (Jansson & Hjalmarsson, 2004c, 2004b; Jansson, 2004). We illustrate with an example.

Example 11.1. The true system is given by

$$y(t) = \frac{0.36q^{-1}}{1 - 0.7q^{-1}} u(t) + e(t),$$

where $e(t)$ is zero-mean white noise with variance 0.1. We want to estimate a model

$$G(\theta) = \frac{bq^{-1}}{1 + aq^{-1}}, \quad \theta = [a \ b]^T,$$

based on $N = 500$ samples of input/output data. The input design problem is formulated as in (93) with $\gamma_P = 0.1$, $\chi_\alpha^2(n) = 5.99$ (which corresponds to a confidence level of 95%) and with T given in Fig. 12.




Fig. 12. Magnitude plot for Example 11.1. Thick solid line: optimal input spectrum. Dashed line: white input spectrum with same power as the optimal input. Thin solid line: desired complementary sensitivity function. Dash-dotted line: open loop system.




Fig. 13. Example 11.1. Dots: the estimated model parameters from 1000 Monte Carlo runs based on input design (93). Dashed ellipse: estimated 95% confidence bound for the parameters for design (93). Dotted ellipse: confidence bound when the input is white noise with the same energy as the design of (93). Contour lines with interval 0.025 are plotted for $\|\bar{A}\|_\infty$ and $\|\bar{A}\|_\infty = 0.1$ corresponds to the thick solid contour.

The minimum input power is $\beta = 0.28$ with the resulting input spectrum shown in Fig. 12. The “ringing” (barely visible) in the input spectrum at high frequency is due to the FIR-filter (order 20) used to shape the input.

Fig. 13 shows the parameter estimates of 1000 Monte Carlo runs. We see that the estimated models are clustered inside the contour $\|\bar{A}\|_\infty = 0.1$ as desired. In fact, 96.7% of the estimated models satisfy the performance constraint. The reason for a higher percentage compared to the specified 95% is due to that also some estimates outside the confidence ellipsoid are inside the level curve $\|\bar{A}\|_\infty = 0.1$, see Fig. 13.

It should be noted that, as usual in experiment design, in order to compute the optimal design in the example, the true system has to be known. Methods that are robust with respect to uncertainty about the system is a wide open research field.

Optimal input design for low order models may sometimes produce counterintuitive designs. This can be attributed to the strong extrapolation properties over the frequency axis for low order models. The behavior, and hence, the accuracy, at one frequency is highly coupled to the behavior at another frequency. The input spectrum obtained in Example 11.1 is quite reassuring in this aspect. Despite a first order model, the input energy is largest in the frequency region between the open loop systems cross-over frequency and the designed closed-loop system bandwidth. This is in complete consistency with the discussion of the performance specifications versus experimentation effort (Trade-off 1) in Section 7.6.

In Section 7.1.2 we saw that robust stability is tied to the v -gap. From Proposition 7.1 we see that the larger the maximum v -gap is for a model set, the larger the generalized stability margin $b_{G,C}$ has to be in order to guarantee stability for systems belonging to the model set. In particular if the maximum v -gap exceeds 1, stability cannot be guaranteed. This is the motivation for the work in Hildebrand and Gevers (2003a) which presents a method for minimizing (with respect to the input spectrum) the worst-case v -gap taken over all models in the confidence region (23) produced by prediction error identification.

11.3. Adaptive input design

As we have seen, in the early 1990s several schemes that iterated between identification and closed-loop control were proposed. In fact, it was suggested, Schrama, 1992, that high performance control based on restricted complexity models required such iterations. However, as we have indicated in Sections 7.6 and 8.1, this is not necessarily so. It has also been shown (Böling & Mäkilä, 1998), that the example in Schrama (1992) can be solved without iterations. So the question is rather *what* can be gained by iterating between identification and experimentation?

Clearly, when we are collecting data, we obtain new information about the system. Hence, iterating between identification and experimentation is beneficial if we can improve our experiment design such that more *control relevant* information becomes available than if we would have stayed with our present design. We have seen that there are two quantities to consider:

- *Minimizing variance.* The input should be chosen such as to minimize the impact from the noise induced errors in the controller $C(\hat{G}_N)$ on the achieved closed-loop performance.
- *Tuning the bias.* When the model set is restricted it is important that data contains the features relevant to control. This is embodied in that the desired operating conditions should be mimicked as closely as possible.

We remark that these two objectives may be conflicting.

11.3.1. Minimizing variance: adaptive input design for LTI systems

One of the few instances where it has been possible to prove that iterative identification and control design actually improves performance is when a full-order model is used for minimum variance control design. Under the assumption that (32) is valid, it was shown in Hjalmarsson, Gevers, and DeBruyne (1996) that iterating between identification and subsequent certainty equivalence input design, and experimentation will always improve the average performance compared to any fix input design (other than the optimal), provided the experiment time is sufficiently long. Certainty equivalence input design means that the last estimated nominal model is used instead of the true system in an optimal input design. In such an approach it would be nice if the optimal experiment design corresponded to the optimal controller since then such iterations would not conflict with the control objective. Unfortunately, this seems to hold only for minimum variance control (Gevers & Ljung, 1986).⁷

A potential problem with this approach is that instability may occur during the identification experiment if the optimal experiment is in closed-loop and the certainty equivalence design is based on a model of poor quality. An alternative which avoids this problem is to tune the input spectrum adaptively in *open loop* (Lindqvist & Hjalmarsson, 2001). Adaptive designs are also considered in, e.g., Cooley and Lee (2001), Lee (2003), Rivera et al. (2003) and Lacy, Bernstein, and Erwin (2003).

11.3.2. Tuning the bias: adaptive input design for non-linear systems

For non-linear systems, we have in Section 9.1 seen that it may be advantageous to have the system operating under the desired conditions when data are collected. It may seem as if this would require that the system operates in closed-loop using an already well-tuned controller which in turn would mean that there is no need to retune the controller. However, we would here like to draw the reader's attention to the fact that there exist open loop alternatives that are competitive. One advantage with collecting data in open loop is that stability is not an issue.

A systematic method to iteratively generate feed-forward controls that approach a given output trajectory is iterative learning control (ILC) (Moore, 1993). Given an initial "trial" input $u_0(t)$ over the time-horizon $t = 1, 2, \dots, N$, and the target output trajectory $\{y_d(t)\}_{t=1}^N$, the input is iteratively refined using the error between the achieved output and the target output as correction. A simple version of the algorithm is to update the input at iteration k according to

$$u_{k+1}(t) = u_k(t) + L_k(y_d(t) - y_k(t)), \quad t = 1, \dots, N, \quad (94)$$

⁷ As pointed out in Forssell and Ljung (2000) this does not hold for a two-degree of freedom model reference control as claimed in Hjalmarsson et al. (1996).




Fig. 14. Outputs used to identify linear models in Example 11.2. Solid line: generated by one step of the ILC algorithm. Dashed line: generated by white noise input. Dotted line: desired output (shown for comparison).

where L_k is some dynamical operator, typically LTI, which has to be chosen such that the algorithm converges. It can be shown, cf. Markusson, Hjalmarsson, and Norrlöf (2002), that L_k plays the role of a model of the inverse of the system. Observe, e.g., that if L_k is the inverse of the true system, then the desired input will be produced in one iteration of the above algorithm under noise free conditions. What makes this algorithm interesting is that the model may be changed between iterations. Hence, it may be possible to use quite simple LTI models and still be able to generate an input which drives the system sufficiently close to the desired operating conditions that data collected from these conditions subsequently can be used to identify a LTI model useful for control design (Markusson et al., 2002; Tao, Kosut, & Ekblad, 1994).

Example 11.2. The Van de Vusse system (van de Vusse, 1964) is a non-minimum phase system often used as a benchmark problem for non-linear process control algorithms (Doyle, Ogunnaike, & Pearson, 1995). Here the kinetic parameters have been chosen such that the system is described by

$$\dot{x}_1 = -50x_1 - 10x_1^2 + (10 - x_2)u,$$

$$\dot{x}_2 = 50x_1 - 100x_2 - x_2u,$$

$$y = x_2,$$

around the equilibrium point. The objective (chosen solely to illustrate the preceding discussion and not any specific chemical process control problem) is that the complementary sensitivity function has the same magnitude as

$$T_d(s) = \frac{1}{0.01s + 1}.$$

The system is sampled with sampling time 0.72 s. One iteration in the ILC-algorithm (94) was taken with the initial input being white noise. The inverse of a second order OE model, identified around the initial trajectory, was used as L_0 in the algorithm. The resulting output (shown as the solid line in Fig. 14) was used to identify a second order OE-model on which an internal model controller was designed. The closed-loop output is shown as the solid line in Fig. 15. It follows the desired output (the dotted curve in the figure)




Fig. 15. Example 11.2. Closed loop output. Solid line: controller designed using model identified from ILC-data. Dashed line: controller designed using model identified from white noise data. Dotted line: desired output.

quite well. For comparison, a model identified using a white noise input (the output is the dashed line in Fig. 14) with the same energy was used in the control design instead. The corresponding closed-loop response is shown as the dashed line in Fig. 15. The response is significantly more oscillatory.

A more spectacular example can be found in Jansson and Hjalmarsson (2002) where an inverted pendulum is controlled.

We also point out that should the process be operating in closed-loop (as is often the case) with some poorly tuned controller, ILC may still be used to generate a reference trajectory $\{r(t)\}_{t=1}^N$ such that the operating conditions become closer to the desired ones.

Another example of adaptive input design for non-linear systems is Zhao and Kanellakopoulos (2002). Here the system has known structure with linear regressions $\theta^T \psi_k(t)$ of observable, possibly, non-linear regressors that influence the states. The system structure is such that by observing the output it is possible to compute the linear combinations $\theta^T \psi_k(t)$. The input is used to drive the system states such that the regressors become linearly independent and hence the unknown parameter vector θ becomes identifiable.

11.3.3. Tuning the bias: using non-linear feedback

The relay experiment used in the relay auto-tuner method, originally proposed in Åström and Hägglund (1984), can be seen as a way of generating an experiment which provides information such that the bias in a model is “tuned” for control. In Karimi, Garcia, and Longchamp (2002) the idea of relay experiments is combined with tuning using gradients, cf. Section 10.2.

11.4. Pre-tests – Identifying performance limitations

The experiment design is eased significantly if the inherent limitations of the system are known. Often amplitude and slew-rate constraints on actuators, which limit the achievable bandwidth, are known and this knowledge should be

incorporated in the experiment design, cf. the discussion of the control effort in Section 7.6.

Non-minimum phase zeros also limit the achievable bandwidth. An explicit variance expression has been derived which can be used for designing pre-test experiments so that these zeros can be accurately identified (Mårtensson & Hjalmarsson, 2003; Jansson, 2004). An interesting aspect is that the asymptotic accuracy of the identified zeros is basically independent of the model order when the prediction-error method is used. Hence, model order selection is not a critical issue here.

11.5. Input design for multivariable systems

Input design is perhaps most important for multivariable systems, especially for ill-conditioned processes such as high-purity distillation columns. To appreciate this, notice that the low gain directions of the system will be poorly identified unless precautions are taken to ensure that the signal to noise ratios are sufficient in these directions. The problem arises as the designed controller will use high gain in these poorly identified directions which may be disastrous for the closed-loop behavior when the model is poor. The key to solving this is to use correlated inputs and both open loop and closed-loop methods have been proposed to this end. In Jacobsen (1994) an open loop grey-box experimental design is proposed as well as a closed loop method. Another closed-loop method is used in Zhu and Butoyi (2002). A general observation is that attempts to perform SISO identification on such plants will fail to deliver relevant models, cf. Zhu and Butoyi (2002) and Jacobsen (1994). The approaches in Cooley and Lee (2001) and Lee (2003), already discussed in Section 11.1, are also applicable to multivariable systems.

11.6. Input design for near-optimal restricted complexity models

In Section 4.4 we saw that near-optimal models of restricted complexity exist if the confidence region for the full-order model is large enough. For a model structure \mathcal{M} of a certain complexity and a given confidence level α , this puts an upper bound on how much information from an experiment that is valuable from a statistical point of view. As more information is added so that the least-squares estimate for \mathcal{M} is no longer inside the confidence region (corresponding to level α) for the full-order model, there is no longer any model in \mathcal{M} which is near-optimal. Thus, if a certain complexity of the model structure is pre-specified (as well as the confidence level), such an upper bound should be incorporated in the experiment design constraints.

12. Validation of control designs

In this section we will discuss some ideas for validation of control designs. We will focus on validation of stability

but the ideas carry over to performance criteria as well. The discussion will be limited to SISO systems.

12.1. Validating stability using the separation principle

Suppose that G is some model of the open loop system G_o which is LTI and that C is an LTI controller which has been designed using G such that the resulting complementary sensitivity function is $T(G, C)$. Then, as we have seen in Section 7.1.2, the achieved closed-loop is guaranteed to be stable if

$$\|\bar{A}\|_\infty < 1, \tag{95}$$

where $\bar{A} = T(G, C)G^{-1}(G_o - G)$.

Suppose that some input–output data set Z^N is available for validation of (95). In Section 5.4 we discussed general principles for validation of restricted complexity models. The advice was to use the separation principle. Here this means first forming a full-order estimate \hat{G}_N of G_o using the available data and then computing the estimate

$$\widehat{\bar{A}}_N = T(G, C)G^{-1}(\hat{G}_N - G), \tag{96}$$

and finally forming a confidence region for this estimate. Using (36) gives

$$\begin{aligned} & |\widehat{\bar{A}}_N(e^{j\omega}) - \bar{A}(e^{j\omega})| \\ & \leq \left| \frac{T(G, C)}{G} \right| \sqrt{\chi_\alpha^2(n)\kappa_{n,N}(\omega) \frac{\Phi_v(\omega)}{N \cdot \Phi_u(\omega)}} \end{aligned}$$

which gives that if

$$|\widehat{\bar{A}}_N(e^{j\omega})| + \left| \frac{T(G, C)}{G} \right| \sqrt{\chi_\alpha^2(n)\kappa_{n,N}(\omega) \frac{\Phi_v(\omega)}{N \cdot \Phi_u(\omega)}} < 1$$

stability is guaranteed (with probability α).

12.2. Validating stability using power iterations

Suppose now that, instead of using a pre-determined data set Z^N , we have the option of choosing the input. For \bar{A} LTI and stable we have that the \mathcal{H}_∞ -norm coincides with the induced \mathcal{L}_2 -norm (84) and that inputs that maximize the ratio $\|\bar{A}u\|_2/\|u\|_2$ are sinusoids with frequency where \bar{A} has maximum gain (Zhou et al., 1996; Khalil, 1996). Hence, it would be easy to check (95) if such a sinusoid is used as input. Unfortunately, this frequency is not known a priori but one may envisage procedures which adapt the frequency such that the gain is maximized. A simple method in this spirit is as follows:

- (1) Let $k = 0$ and select an arbitrary input sequence $\{u_k(t)\}_{t=1}^N$ with the constraint that $\|u_k\|_2 = \beta$ for some constant $\beta > 0$.
- (2) Perform an experiment where $u_k(t)$ is applied to G_o : $y_k(t) = G_o(q)u_k(t)$, $t = 1, \dots, N$.




Fig. 16. Magnitude curve of the frequency response in Example 12.1.

- (3) Filter the corresponding output through the inverse of G : $G^{-1}y_k(t)$.
- (4) Subtract $u_k(t)$ from this signal: $G^{-1}y_k(t) - u_k(t)$. Neglecting disturbances, Steps 2–4 give the signal $z_k(t) = \bar{A}(q)u_k(t)$, $t = 1, \dots, N$.
- (5) Let $\gamma_k = \|z_k\|_2/\beta$.
- (6) Let $u_{k+1}(t) = \frac{1}{\gamma_k} z_k(N + 1 - t)$, $t = 1, \dots, N$.
- (7) Let $k = k + 1$ and go to Step 2.

The γ_k computed in Step 5 is clearly an underbound to the induced \mathcal{L}_2 -norm (84) of \bar{A} .

When the experiments are noise free and performed under zero initial conditions it can be shown that u_k will, modulo transient effects, converge to a sinusoid with approximately the desired frequency. Furthermore,

$$\gamma_k \rightarrow \left\| \sum_{l=0}^{N-1} \delta_l q^{-l} \right\|_\infty \text{ as } k \rightarrow \infty$$

where δ_k , $k = 1, 2, \dots$ are the impulse response coefficients of \bar{A} and where the convergence is monotonic. Thus for N sufficiently large, a good estimate of $\|\bar{A}\|_\infty$ is obtained.

In each iteration, it is the frequency where \bar{A} has maximum gain that is amplified the most. Due to the normalization in Step 6, the energy at other frequencies will be damped out as the iterations proceed. The algorithm is closely related to the power method for computing the largest eigenvalue of a symmetric matrix and we therefore call the iterations for *power iterations*.

Example 12.1. Let \bar{A} be of second order with the magnitude curve of the frequency response given in Fig. 16. The maximum gain is at $\omega \approx 0.79$.

Initially a white input with variance 1 and length $N = 100$ is applied. It gives the output shown as the dashed line in Fig. 17. Also shown in this figure is the output after 10 power iterations. Clearly, the latter output is not far from a sinusoid with frequency around $\omega = 0.79$.

In Fig. 18, the estimated gain γ_k is given for the first 10 iterations together with $\|\bar{A}\|_\infty$ (dashed line). Clearly, the sequence of gain estimates converges exponentially fast to a rather accurate lower bound for the true gain of the system.




Fig. 17. Outputs in Example 12.1. Dashed curve: initial output. Solid line: output after 10 iterations.




Fig. 18. Solid line: gain versus iteration number in Example 12.1. Dashed line: \mathcal{H}_∞ -norm of \bar{A} .




Fig. 19. Output in Example 12.2 after 10 iterations.

Example 12.2. Power iterations applied to another \bar{A} gives the output in Fig. 19 after 10 iterations. It does not even remotely resemble a sinusoid.

The explanation is that \bar{A} is not LTI in this case, in fact it is given by the continuous time system

$$\dot{x}_1 = x_2,$$

$$\dot{x}_2 = -2x_1^3 - 4x_2 + u,$$

$$y = x_2,$$

sampled with a sampling time $T = 0.1$.

The maximum of the gain sequence computed in Step 5 in the power iterations still is a lower bound to the induced \mathcal{L}_2 -norm. This sequence is shown in Fig. 20. Also in this case, the gain sequence increases monotonically. In Example 6.9 in Khalil (1996) an upper bound for the induced \mathcal{L}_2 -norm of this system is given as 1/4 (marked as the dashed line in Fig. 20). From Fig. 20 we see that the lower bound produced by the power iterations is quite accurate.

There are two conclusions to be drawn from Example 12.2. Firstly, power iterations can be used to detect non-linearities in a system. If the iterations do not converge to a sinusoidal signal, the system is non-linear. Secondly, power iterations may give a useful lower bound on the induced \mathcal{L}_2 -norm of certain non-linear systems, cf. Section 9.2. There is no proof for the latter statement and the author certainly does not claim that power iterations produce monotonically increasing lower bounds, or that the produced bounds are accurate, for general non-linear systems. However, the nice behavior observed for this and some other non-linear systems is intriguing.

13. Concluding remarks

Looking back at the ground covered in this paper (not to mention all that regrettably had to be omitted), it is clear that much progress has been made in this research area since the book (Bitmead, Gevers, & Wertz, 1990) and the SYSID plenary in Budapest in 1991 (Gevers, 1991), which in many respects can be seen as triggers for activities in this field.

Looking at the problem from a statistical perspective, we first remind the reader of the guidelines in Section 6:

- (i) *Always first model as well as possible.*
- (ii) *Use a very flexible model structure as benchmark for computing confidence bounds and mean-square error.*
- (iii) *Select the input such that the model uncertainty at frequency regions of interest is insensitive to the model complexity.*

Notice that we in (i) do not refer to full-order modeling, the aim is rather to obtain a near-optimal restricted complexity model. Following advice (i) ensures good statistical properties and enables the under-modeling to be quantified. The model can subsequently be simplified without loss of accuracy. Referring to Fig. 1, this means that going from an near-optimal restricted model to a low order controller by way of controller reduction or model reduction will not significantly influence the statistical properties of the procedure. Advice (ii) ensures a reasonable assessment of the model error and




Fig. 20. Solid line: gain versus iteration number in Example 12.2. Dashed line: upper bound on \mathcal{L}_2 -norm.

as we saw in Section 6, (iii) can be ensured by inputs with large spectral peaks (even though there seems to be more to understand here).

From (i) follows that identification and control are not as intertwined as might be believed:

- For a *given* data set, producing the set of unfalsified models has little to do with the control problem. Once the set of unfalsified models is obtained, what remains is a robust control problem with, in turn, little to do with the identification problem.

There is, however, as pointed out in Section 7.4, an interface problem between the two areas that cannot be neglected in that the model set description obtained from system identification typically does not fit the present robust control design framework. Much work has been done in this area, cf. Section 7.4.

In order to use (iii), it is necessary to know the modeling requirements. For control, we have argued (Sections 7.6, 8.10, 9.2 and 11.2) that these are in principle encapsulated by the following condition on the weighted relative error:

$$\|\bar{\Delta}(G_o, \hat{G}_N, T)(\hat{G}_N, C)\| \leq \gamma_P. \quad (97)$$

Condition (97) clarifies the roles of the design variables and how these should be designed and traded off against each other. For example, for a SISO LTI system the bound (66) applies

$$\|\bar{\Delta}\|_\infty \leq \sup_\omega \left| \frac{T(\hat{G}_N, C(\hat{G}_N))}{\hat{G}_N} \right| \sqrt{\chi_\alpha^2(n) \kappa_{n,N} \frac{\Phi_v}{N \cdot \Phi_u}}. \quad (98)$$

From this expression we noted in Section 7.6 that there is a

- *control effort versus experimental information trade-off*

in force. The more aggressive controller, the more information is necessary in order to guarantee a small $\|\bar{\Delta}\|$.

Condition (98) also provides guidelines for input design, cf. Section 11.2. High input energy density is required where the designed control effort, $|T(\hat{G}_N, C(\hat{G}_N))/\hat{G}_N|$, is large, but *should* be small where the designed control effort is small in order to reduce the modeling effort.

Overall, we have argued that the experiment design is the *most important* design variable for a successful application and we can summarize our observations as follows:

- *excite where it may hurt most.*

On a principal level, the input should be chosen such that system properties that are highly detrimental to the closed-loop performance are revealed by the experiment, cf. the discussion in Section 5.3, and the discussion of the performance specifications versus experimentation effort (Trade-off 1) in Section 7.6. See also Example 11.1.

- *do not excite where it does not hurt.*

There is no benefit from knowing what one does not have to know; in this context it only means that the modeling problem becomes more complex. Of course, it is typically a priori difficult to be aware of which system information one has to know for the control design and this is one of the *key problems* in identification and control. Nevertheless, it is instructive to keep this advice in mind.

- *adaptation increases accuracy.*

Since the true system and the disturbances are unknown, efficient resolution of the previous item is facilitated by adaptive or iterative methods which are able to improve on the input design so that future data samples contain information of better quality from the control application point of view. At present, the only result in this direction for prediction error identification seems to be (Hjalmarsson et al., 1996) which in turn is based on some restrictive assumptions. Notably, preferential identification in conjunction with robust control has proven a viable route for adaptation, see Section 8.11. It should also be remarked that there are differing opinions regarding the value of adaptive input design (Veres, 2000).

In this context, we would like to draw the attention to the potential of using the system itself to (semi-)automatically generate the information of interest. This idea has been used several times in this paper. In Section 10.2, the system was used to generate sensitivity information. Iterative Learning Control was used in Section 11.3.2 to gradually move the system trajectory closer to the desired one. It was pointed out in Section 11.3.3 that relay feedback can be seen as a way of using the system itself to automatically generate a certain type of information. Finally, power iterations were used in Section 12.2 to estimate the gain of a system.

An interesting topic here is to further explore what non-linear feedback has to offer for automatic generation of suitable experimental conditions.

In a number of contributions, see e.g. Section 8.5, identification in closed-loop has been advocated. While this may certainly be motivated by many reasons, e.g. safety and the fact that most processes are already operating in closed-loop, there is from the point of statistical accuracy not obvious that closed-loop identification is preferable to open loop identification. In Section 9.1 we saw that for bias reasons, the experimental conditions should reflect the *desired* operation characteristics. A closed-loop experiment with a poorly tuned controller may in this respect give less useful data than a well designed open loop experiment, cf. Section 11.3.2. Based on the high order approximation (32) it follows (Forsell & Ljung, 2000) that when the output variance is constrained, it is optimal with closed-loop identification with a controller that depends on the *true system* and the experimental constraints to counter-effect noise induced uncertainty (variance) effectively. However, also here there is no guarantee in practice that closed-loop data are more informative.

While much of the work in the area has focused on LTI systems, cf. Section 8, there has in recent years been a shift towards more realistic problem settings where the system is considered to exhibit different kinds of non-linear behaviors, cf. Section 9. With the diversity of non-linear systems, one may expect this research area to proliferate in coming years; there are, e.g., results emerging for hybrid systems, see, e.g., Ferrari-Trecate, Muselli, Liberati, and Morari (2003). This observation also accentuates the importance of close collaboration with application fields in order to address relevant problems.

A number of interesting benchmarks have appeared throughout the years, see, e.g., Graebe (1994; 1995) and Landau, Karimi, and Hjalmarsson (2003). To facilitate the comparison of different algorithms, a set of generally acknowledged benchmark problems would be most useful.

Despite the leverage that the forgiving nature of feedback offers, the joint identification and control problem is complex with many aspects that have to be balanced against each other. Since little effort has been spent so far on packaging algorithms in a way that meets the requirements of industrial practitioners, cf. Rivera and Jun (2000) and Zhu (1998), it is not surprising that, as reported in Zhu (1998), system identification still has had relatively little impact on industrial control engineers in general. However, as hopefully evidenced by this survey, the progress in this research area is very promising and there certainly are interesting years ahead!

Acknowledgements

During the last decade I have had the privilege to collaborate with Franky De Bruyne, László Gerencsér, Michel Gevers, Svante Gunnarsson, Anders Hansson, Weng Khuen

Ho, Henrik Jansson, Kristian Lindqvist, Lennart Ljung, Ola Markusson, Jonas Mårtensson, Brett Ninness, Johan Schoukens, Jonas Sjöberg, Bo Wahlberg and Sandor Veres, on this topic; a process which has lead to many of the insights in this paper. Thanks guys for enlightening me, you're a great bunch! Special thanks goes to Brett and Bo for proof-reading early versions of this manuscript, to Johan for valuable comments and to Henrik for providing Example 11.1. Thanks also to the five referees for their many insightful comments which were instrumental in improving the quality of the paper.

This research has been supported by The Swedish Research Council.

References

- Albertos, P., & Sala, A., (Eds.). (2002). *Iterative identification and control*. Berlin: Springer.
- Anderson, B. D. O., Bitmead, R. R., Johnson, C. R., Kokotovic, P. V., Kosut, R. L., Mareels, I. M. Y., Praly, L., & Riedle, B. D. (1986). *Stability of adaptive systems: passivity and averaging analysis*. Cambridge, MA: MIT Press.
- Anderson, B. D. O., Bombois, X., Gevers, M., & Kulcsar, K. (1998). Caution in iterative modelling and control design. In *Proceedings of the IFAC workshop on adaptive control and signal processing*, Glasgow (pp. 13–19).
- Anderson, B. D. O., & Liu, Y. (1989). Controller reduction: concepts and approaches. *IEEE Transactions on Automatic Control*, 34(8), 802–812.
- Åström, K. J. (1993). Matching criteria for control and identification. In *Proceedings of the second European control conference*, Groningen, Netherlands (pp. 248–251).
- Åström, K. J., & Häggglund, T. (1984). Automatic tuning of simple regulators with specifications on phase and amplitude margin. *Automatica*, 20(5), 645–651.
- Bai, E. -W., Nagpal, K. M., & Tempo, R. (1996). Bounded-error parameter estimation: Noise models and recursive algorithms. *Automatica*, 32(7), 985–999.
- Bai, E. -W., Ye, Y. Y., & Tempo, R. (1999). Bounded error parameter estimation: a sequential analytic center approach. *IEEE Transactions on Automatic Control*, 44(6), 1107–1117.
- Bayard, D. S., Yam, Y., & Mettler, E. (1992). A criterion for joint optimization of identification and robust control. *IEEE Transactions on Automatic Control*, 37(7), 986–991.
- Bitmead, R. R. (2000). Iterative optimal control design. In P. Albertos, A. Sala (Eds.), *Iterative identification and control* (pp. 185–208). Berlin: Springer.
- Bitmead, R. R., Gevers, M., & Wertz, V. (1990). *Adaptive optimal control—the thinking man's GPC*. New York: Prentice-Hall.
- Bittanti, S., Campi, M. C., & Garatti, S. (2002). New results on the asymptotic theory of system identification for the assessment of the quality of estimated models. In *Proceedings of the 41st IEEE conference on decision and control*, Las Vegas, Nevada, USA (pp. 1814–1819).
- Bøjstrup, K., Niemann, H. H., Kjølstad Poulsen, N., & Jørgensen, S. B. (2003). Iterative identification for control and robust performance of bioreactor. In *Proceedings of the 13th IFAC symposium on system identification*, Rotterdam, Netherlands (pp. 1673–1678).
- Böling, J. M., & Mäkilä, P. M. (1998). On control relevant criteria in H_∞ identification. *IEEE Transactions on Automatic Control*, 43(5), 694–700.
- Bombois, X., Anderson, B. D. O., & Gevers, M. (2000a). Mapping parametric confidence ellipsoids to Nyquist plane for linearly parametrized transfer functions. In: G.C. Goodwin (Ed.), *Model identification and adaptive control* (pp. 53–71). Berlin: Springer.

- Bombois, X., Gevers, M., & Scorletti, G. (2000b). A measure of robust stability for an identified set of parametrized transfer functions. *IEEE Transactions on Automatic Control*, 45(11), 2141–2145.
- 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., Anderson, B. D. O., Gevers, M., & Van den Hof, P. M. J. (2002). A new robust control design procedure based on a PE identification uncertainty set. In *Proceedings of the 15th IFAC world congress*, Barcelona, Spain (pp. 2842–2847).
- Bombois, X., Scorletti, G., Van den Hof, P. M. J., & Gevers, M. (2004). Least costly identification experiment for control: a solution based on a high-order model approximation. In *Proceedings of the American control conference*.
- Bongers, P. M. M. (1991). On a new robust stability margin. In *Recent advances in mathematical theory of systems, control, networks and signal proceedings, MTNS-91* (pp. 377–382).
- Cabral, F. B., & Safonov, M. G. (2003). Unfalsified model reference adaptive control using the ellipsoid algorithm. In *Proceedings of the 42nd IEEE conference on decision and control*, Maui, Hawaii (pp. 3250–3255).
- Cadic, M., Weiland, S., & Polderman, J. W. (2003). Strong robustness measures for sets of linear SISO systems. In *Proceedings of the 13th IFAC symposium on system identification*. Rotterdam, Netherlands (pp. 39–44).
- Campi, M. C., Lecchini, A., & Savaresi, S. M. (2002). Virtual reference feedback tuning: a direct method for the design of feedback controllers. *Automatica*, 38(8), 1337–1346.
- Campi, M. C., & Weyer, E. (2002). Finite sample properties of system identification methods. *IEEE Transactions on Automatic Control*, 47(8), 1329–1334.
- Chen, J., & Gu, G. (2000). *Control-oriented system identification. An \mathcal{H}_∞ approach*. New York: Wiley.
- Codrons, B., Bendotti, P., Falinower, C.-M., & Gevers, M. (1999). A comparison between model reduction and controller reduction: Application to a PWR nuclear plant. In *Proceedings of the 38th IEEE conference on decision and control*. Phoenix, Arizona, USA (pp. 4625–4630).
- Cooley, B. L., & Lee, J. H. (2001). Control relevant experiment design for multivariable systems described by expansions in orthonormal bases. *Automatica*, 37(2), 273–281.
- De Bruyne, F., & Carrette, P. (1997). Synthetic generation of the gradient for an iterative controller optimization method. In *Proceedings of the fourth European control conference*. Bruxelles, Belgium.
- de Callafon, R. A., & Van den Hof, P. M. J. (1997). Suboptimal feedback control by a scheme of iterative identification and control. *Mathematical Modelling of Systems*, 3(1), 77–101.
- de Callafon, R. A., & Van den Hof, P. M. J. (2001). Multivariable feedback relevant system identification of a wafer stepper system. *IEEE Transactions on Control Systems Technology*, 9(2), 381–390.
- de Vries, D. K., & Van den Hof, P. M. J. (1995). Quantification of uncertainty in transfer function estimation: a mixed probabilistic-worst-case approach. *Automatica*, 31(4), 543–557.
- Desborough, L., & Miller, R. (2001). Increasing customer value of industrial control performance monitoring—Honeywell's experience. In *Chemical Process Control*, vol. VI, Tucson, AR (pp. 153–185).
- Devilbiss, S. L., & Yurkovich, S. (1998). Exploiting ellipsoidal parameter set estimates in H_∞ robust control design. *International Journal of Control*, 69(2), 271–284.
- Douma, S. G., & Van den Hof, P. M. J. (2003). Relation between uncertainty structures in identification for robust control. In *Proceedings of the 13th IFAC symposium on system identification*, Rotterdam, Netherlands (pp. 33–38).
- Douma, S. G., Van den Hof, P. M. J., & Bosgra, O. H. (2003). Controller tuning freedom under plant identification uncertainty: double Youla beats gap in robust stability. *Automatica*, 39(2), 325–333.
- Doyle, F. J., Ogunnaike, B. A., & Pearson, R. K. (1995). Nonlinear model-based control using second-order Volterra models. *Automatica*, 31(5), 697–714.
- Eker, S. A., & Nikolaou, M. (2002). Linear control of nonlinear systems: interplay between nonlinearity and feedback. *A.I.Ch.E. Journal*, 48(9), 1957–1980.
- Enqvist, M., & Ljung, L. (2002). Estimating nonlinear systems in a neighborhood of LTI-approximants. In *Proceedings of the 41st IEEE conference on decision and control*, Las Vegas, Nevada (pp. 1005–1010).
- Enqvist, M., & Ljung, L. (2003). Linear models of nonlinear FIR systems with Gaussian inputs. In *Proceedings of the 13th IFAC symposium on system identification*, Rotterdam, Netherlands (pp. 1910–1915).
- Ferrari-Trecate, G., Muselli, M., Liberati, D., & Morari, M. (2003). A clustering technique for the identification of piecewise affine systems. *Automatica*, 39(2), 205–217.
- Forsell, U., & Ljung, L. (1999). Closed-loop identification revisited. *Automatica*, 35(7), 1215–1241.
- Forsell, U., & Ljung, L. (2000). Some results on optimal experiment design. *Automatica*, 36(5), 749–756.
- Gevers, M. (1991). Connecting identification and robust control: a new challenge. In *Proceedings of the IFAC/IFORS symposium on identification and system parameter estimation*, Budapest, Hungary.
- Gevers, M. (1993). Towards a joint design of identification and control?. In: H.L. Trentelman, & J.C. Willems (Eds.), *Essays on control: perspectives in the theory and its applications* Boston: Birkhäuser.
- Gevers, M. (2000). Identification and validation for robust control. In: P. Albertos, & A. Sala (Eds.), *Iterative identification and control* (pp. 185–208). Berlin: Springer.
- Gevers, M., Bombois, X., Codrons, B., Scorletti, G., & Anderson, B. D. O. (2003). Model validation for control and controller validation in a prediction error identification framework—Part I: theory, Part II: applications. *Automatica*, 39(3), 403–427.
- Gevers, M., & Ljung, L. (1986). Optimal experiment designs with respect to the intended model application. *Automatica*, 22(5), 543–554.
- Gevers, M., Ljung, L., & Van den Hof, P. M. J. (2001). Asymptotic variance expressions for closed-loop identification. *Automatica*, 37(5), 781–786.
- Ghulchak, A., & Rantzer, A. (2002). Robust control under parametric uncertainty via primal-dual convex analysis. *IEEE Transactions on Automatic Control*, 47(4), 632–636.
- Glover, K., & McFarlane, D. (1989). Robust stabilization of normalized coprime factor plant descriptions with H_∞ -bounded uncertainty. *IEEE Transactions on Automatic Control*, 34(8), 821–830.
- Goodwin, G. C., Braslavsky, J. H., & Seron, M. M. (2002). Non-stationary stochastic embedding for transfer function estimation. *Automatica*, 38(1), 47–62.
- Goodwin, G. C., Gevers, M., & Ninness, B. (1992). Quantifying the error in estimated transfer functions with application to model order selection. *IEEE Transactions on Automatic Control*, 37(7), 913–928.
- Goodwin, G. C., Graebe, S. F., & Salgado, M. E. (2001). *Control system design*. Englewood Cliffs, NJ: Prentice-Hall.
- Goodwin, G. C., Wang, L., & Miller, D. (1999). Bias-variance trade-off issues in robust controller design using statistical confidence bounds. In *Proceedings of the 14th IFAC world congress*, Beijing, P.R. China.
- Graebe, S. F. (1994). Control of an unknown plant: a benchmark of new format. *Automatica*, 30(4), 567–575.
- Graebe, S. F. (Ed.). (1995). *European Journal of Control* 1(2), 75–173. (Special issue).
- Guo, L. (2002). Exploring the maximum capability of adaptive feedback. *International Journal of Adaptive Control and Signal Processing*, 16, 341–354.
- Guo, L., & Ljung, L. (1994). The role of model validation for assessing the size of the unmodelled dynamics. In *Proceedings of the 33rd IEEE conference on decision and control*, vol. 4, (pp. 3894–3899).

- Hakvoort, R. G., Schrama, R. J. P., & Van den Hof, P. M. J. (1994). Approximate identification with closed loop criterion and application to LQG feedback design. *Automatica*, 30(4), 679–690.
- Hakvoort, R. G., & Van den Hof, P. M. J. (1997). Identification of probabilistic system uncertainty regions by explicit evaluation of bias and variance errors. *IEEE Transactions on Automatic Control*, 42(11), 1516–1528.
- Henriksson, B., Markusson, O., & Hjalmarsson, H. (2001). Control relevant identification of nonlinear systems using linear models. In *Proceedings of the American control conference*, Arlington, USA.
- Hildebrand, R., & Gevers, M. (2003a). Identification for control: optimal input design with respect to a worst-case v -gap cost function. *SIAM Journal Control Optim.*, 41(5), 1586–1608.
- Hildebrand, R., & Gevers, M. (2003b). Quantification of the variance of estimated transfer functions in the presence of undermodelling. In *Proceedings of the 13th IFAC symposium on system identification*, Rotterdam, Netherlands (pp. 1850–1855).
- Hjalmarsson, H. (1991). Estimation of the variability of time-varying systems. In *Proceedings of the IFAC/IFORS symposium on identification and system parameter estimation* (pp. 1701–1706).
- Hjalmarsson, H. (1998). Control of nonlinear systems using Iterative Feedback Tuning. In *Proceedings of the American control conference* (pp. 2083–2087).
- Hjalmarsson, H. (2002). Iterative feedback tuning—an overview. *International Journal on Adaptive Control and Signal Processing*, 16(5), 373–395.
- Hjalmarsson, H., Gevers, M., & DeBruyne, F. (1996). For model based control design criteria, closed loop identification gives better performance. *Automatica*, 32(12), 1659–1673.
- Hjalmarsson, H., Gevers, M., Gunnarsson, S., & Lequin, O. (1998). Iterative Feedback Tuning: theory and applications. *IEEE Control Systems Magazine*, 18(4), 26–41.
- Hjalmarsson, H., Gunnarsson, S., & Gevers, M. (1995). Optimality and sub-optimality of iterative identification and control design schemes. In *Proceedings of the American control conference*, Seattle, Washington (pp. 2559–2563).
- Hjalmarsson, H., & Jansson, H. (2003). Using a sufficient condition to analyze the interplay between identification and control. In *Proceedings of the 13th IFAC symposium on system identification*. Rotterdam, Netherlands (pp. 45–50).
- Hjalmarsson, H., & Lindqvist, K. (2001). Identification for control: L_2 and L_∞ methods. In *Proceedings of the 40th IEEE conference on decision and control*, Orlando, Florida, USA.
- Hjalmarsson, H., & Ninness, B. (2004). An exact finite sample variance expression for linearly parameterized frequency function estimates. *Automatica*, submitted for publication.
- Holmberg, U., Valentinotti, S., & Bonvin, D. (2000). An identification-for-control procedure with robust performance. *Control Engineering Practice*, 8(10), 1107–1117.
- Horowitz, I. M. (1992). *Quantitative feedback design theory (QFT)*. Boulder, Colorado: QFT Publications
- Hussain, M. A. (1999). Review of the applications of neural networks in chemical process control—simulation and on-line implementation. *Artificial Intelligence in Engineering*, 13, 55–68.
- Jacobsen, E. W. (1994). Identification for control of strongly interactive plants. In *Proceedings of the A.I.Ch.E. annual meeting*, San Francisco, California, USA.
- Jansson, H. (2004). *Experiment design with applications in identification for control*. Ph.D. Thesis, Royal Institute of Technology, Stockholm.
- Jansson, H., & Hjalmarsson, H. (2002). From open loop learning to closed loop control. In *Proceedings of the 41st IEEE conference on decision and control*. IEEE, Las Vegas, USA (pp. 4209–4214).
- Jansson, H., & Hjalmarsson, H. (2004a). Convex computation of worst-case criteria with applications in identification and control. In *Proceedings of the 43rd IEEE conference on decision and control*, Bahamas.
- Jansson, H., & Hjalmarsson, H. (2004b). A framework for mixed \mathcal{H}_∞ and \mathcal{H}_2 input design. In *MTNS'04*, Leuven, Belgium.
- Jansson, H., & Hjalmarsson, H. (2004c). A general framework for mixed \mathcal{H}_∞ and \mathcal{H}_2 input design. *IEEE Transactions on Automatic Control*, submitted for publication.
- Kammer, L. C., Bitmead, R. R., & Bartlett, P. L. (2000). Direct iterative tuning via spectral analysis. *Automatica*, 36(9), 1301–1307.
- Karimi, A., Garcia, D., & Longchamp, R. (2002). Iterative controller tuning using Bode's integrals. In *Proceedings of the 41st IEEE conference on decision and control*, Las Vegas, Nevada (pp. 4227–4232).
- Karimi, A., Miskovic, L., & Bonvin, D. (2003). Iterative correlation-based controller tuning with application to a magnetic suspension system. *Control Engineering Practice*, 11(9), 1069–1078.
- Khalil, H. L. (1996). *Nonlinear systems*. 2nd ed., Englewood Cliffs, NJ: Prentice-Hall.
- Kosut, R. L. (1995). Uncertainty model unfalsification: a system identification paradigm compatible with robust control design. In *Proceedings of the 34th IEEE conference on decision and control*, vol. 4, New Orleans, LA (pp. 3492–3497).
- Kosut, R. L. (1999). Iterative unfalsified adaptive control: Analysis of the disturbance-free case. In *Proceedings of the American control conference*, San Diego, California (pp. 566–570).
- Kosut, R. L. (2001). Uncertainty model unfalsification. *Annual Reviews in Control*, 25, 65–76.
- Kosut, R. L., & Anderson, B. D. O. (1994). Least-squares parameter set estimation for robust control design. In *Proceedings of the American control conference*, vol. 3 (pp. 3002–3006).
- Krause, J. M., Stein, G., & Khargonekar, P. P. (1992). Sufficient conditions for robust performance of adaptive controllers with general uncertainty structure. *Automatica*, 28(2), 277–288.
- Lacy, S. L., Bernstein, D. S., & Erwin, R. S. (2003). Finite-horizon input selection for system identification. In *Proceedings of the 42nd IEEE conference on decision and control* (pp. 6485–6490).
- Landau, I. D., Karimi, A., & Constantinescu, A. (1997). Recursive algorithms for identification in closed loop: a unified approach and evaluation. *Automatica*, 33(8), 1499–1523.
- Landau, I. D., Karimi, A., & Constantinescu, A. (2001). Direct controller order reduction by identification in closed loop. *Automatica*, 37(11), 1689–1702.
- Landau, I. D., Karimi, A., & Hjalmarsson, H. (Eds.). (2003). Design and optimization of restricted complexity controllers. [Special issue]. *European Journal of Control*, 9(1), 1–132.
- Lau, M. K., Boyd, S., Kosut, R. L., & Franklin, G. F. (1991). Robust control design for ellipsoidal plant set. In *Proceedings of the 30th IEEE conference on decision and control*, Brighton (pp. 291–296).
- Lee, J. H. (2003). Control-relevant design of periodic test input signals for iterative open-loop identification of multivariable FIR systems. In *Proceedings of the 13th IFAC symposium on system identification*, Rotterdam, Netherlands (pp. 935–940).
- Lee, W. S., Anderson, B. D. O., Mareels, I. M. Y., & Kosut, R. L. (1995). On some key issues in the windsurfer approach to adaptive robust control. *Automatica*, 31(11), 1619–1636.
- Lee, W. S., Anderson, B. D. O., Kosut, R. L., & Mareels, I. M. Y. (1993). A new approach to adaptive robust control. *International Journal of Adaptive Control and Signal Processing*, 7(3), 183–211.
- Lindqvist, K., & Hjalmarsson, H. (2001). Identification for control: Adaptive input design using convex optimization. In *Proceedings of the 40th IEEE conference on decision and control*, Orlando, Florida, USA.
- Ling, W.-M., & Rivera, D. E. (2001). A methodology for control-relevant nonlinear system identification using restricted complexity models. *Journal of Process Control*, 11(2), 209–222.
- Ljung, L. (1985). Asymptotic variance expressions for identified black-box transfer function models. *IEEE Transactions on Automatic Control*, 30(9), 834–844.

- Ljung, L. (1999a). Model validation and model error modeling. In: B. Wittenmark, & A. Rantzer (Eds.), *The Åström symposium on control* (pp. 15–42). Sweden: Studentlitteratur. Lund.
- Ljung, L. (1999b). *System identification: theory for the user*. 2nd ed., Englewood Cliffs, NJ: Prentice-Hall.
- Ljung, L. (2000). Model error modeling and control design. In *Proceedings of the 12th IFAC symposium on system identification*, Santa Barbara, CA.
- Ljung, L., & Guo, L. (1997). The role of model validation for assessing the size of the unmodeled dynamics. *IEEE Transactions on Automatic Control*, 42(9), 1230–1239.
- Ljung, L., & Hjalmarsson, H. (1995). System identification through the eyes of model validation. In *Proceedings of the third European control conference*, Rome, Italy (pp. 949–954).
- Ljung, L., & Söderström, T. (1983). *Theory and practice of recursive identification*. Cambridge, MA: MIT Press.
- Ljung, L., & Yuan, Z. D. (1985). Asymptotic properties of black-box identification of transfer functions. *IEEE Transactions on Automatic Control*, 30(6), 514–530.
- Mäkilä, P. M., & Partington, J. R. (2003). On linear models for nonlinear systems. *Automatica*, 39(1), 1–13.
- Mäkilä, P. M., Partington, J. R., & Gustafsson, T. K. (1995). Worst-case control-relevant identification. *Automatica*, 31(12), 1799–1819.
- Malan, S., Milanese, M., Regruto, D., & Taragna, M. (2001). Robust control from data via uncertainty model sets identification. In *Proceedings of the 40th IEEE conference on decision and control*, IEEE, Orlando, Florida, USA.
- Markusson, O., Hjalmarsson, H., & Norrlöf, M. (2002). A general framework for Iterative Learning Control. In *Proceedings of the 15th IFAC world congress*, Barcelona, Spain (pp. 1697–1702).
- Mårtensson, J., & Hjalmarsson, H. (2003). Identification of performance limitations using general SISO structures. In *Proceedings of the 13th IFAC symposium on system identification*, Rotterdam, Netherlands (pp. 519–524).
- Massoumnia, M. -A., & Kosut, R. L. (1994). A family of norms for system identification problems. *IEEE Transactions on Automatic Control*, 39(5), 1027–1031.
- McKelvey, T. (2000). Frequency domain identification. In *Proceedings of the 12th IFAC symposium on system identification*, Santa Barbara, CA.
- Milanese, M. (1998). Learning models from data: the set membership approach. In *Proceedings of the American control conference*, vol. 1 (pp. 178–182).
- Milanese, M., Norton, J., Piet-Lahanier, H., & Walter, E. (1996). *Bounding approaches to system identification*. New York: Plenum Press.
- Milanese, M., & Taragna, M. (1999). H_∞ identification of “soft” uncertainty models. *Systems and Control Letters*, 37(4), 217–228.
- Milanese, M., & Taragna, M. (2002). Optimality, approximation, and complexity in set membership H_∞ identification. *IEEE Transactions on Automatic Control*, 47(10), 1682–1690.
- Milanese, M., & Vicino, A. (1991). Optimal estimation theory for dynamic systems with set membership uncertainty: an overview. *Automatica*, 27(6), 997–1009.
- Moore, K. L. (1993). *Iterative learning control for deterministic systems Advances in industrial control*. Berlin: Springer.
- Morari, M., & Zafriou, E. (1989). *Robust process control*. Englewood Cliffs, NJ: Prentice-Hall.
- Mosskull, H., Wahlberg, B., & Galic, J. (2003). Validation of stability for an induction machine drive using experiments. In *Proceedings of the 13th IFAC symposium on system identification*, Rotterdam, Netherlands (pp. 1460–1465).
- Murray, R. M., Åström, K. J., Boyd, S. P., Brockett, R. W., & Stein, G. (2003). Future directions in control in an information-rich world. *IEEE Control Systems Magazine*, 23(2), 20–33.
- Narendra, K. S., & Stretter, D. N. (1964). An adaptive procedure for controlling undefined linear processes. *IEEE Transactions on Automatic Control*, 9(4), 545–548.
- Ninness, B., & Goodwin, G. C. (1995). Estimation of model quality. *Automatica*, 31(12), 1771–1797.
- Ninness, B., & Henriksen, S. (2003). System identification via a computational Bayesian approach. In *Proceedings of the 13th IFAC symposium on system identification*, Rotterdam, Netherlands (pp. 785–790).
- Ninness, B., & Hjalmarsson, H. (2004a). The effect of regularisation on variance error. *IEEE Transactions on Automatic Control*, 49(7), 1142–1147.
- Ninness, B., & Hjalmarsson, H. (2004b). On the frequency domain accuracy of closed loop estimates. *Automatica*, to appear.
- Ninness, B., & Hjalmarsson, H. (2004c). Variance error quantifications that are exact for finite model order. *IEEE Transactions on Automatic Control*, 49(8), 1275–1291.
- Ninness, B., Hjalmarsson, H., & Gustafsson, F. (1999). The fundamental role of general orthonormal bases in system identification. *IEEE Transactions on Automatic Control*, 44(7), 1384–1406.
- Ogunnaike, B. A. (1996). A contemporary industrial perspective on process control theory and practice. *Annual Reviews in Control*, 20, 1–8.
- Parker, R. S., Heemstra, D., Doyle, F. J., Pearson, R. K., & Ogunnaike, B. A. (2001). The identification of nonlinear models for process control using tailored “plant-friendly” input sequences. *Journal Process Control*, 11(2), 237–250.
- Partanen, A. G., & Bitmead, R. R. (1995). The application of an iterative identification and controller design to a sugar cane crushing mill. *Automatica*, 31(11), 1547–1563.
- Partington, P. M., & Mäkilä, J. R. (1999). On robustness in system identification. *Automatica*, 35(5), 907–916.
- Phillips, S. M., Kosut, R. L., & Franklin, G. F. (1988). An averaging analysis of discrete-time indirect adaptive control. In *Proceedings of the American control conference*, Atlanta, USA.
- Pintelon, R., & Schoukens, J. (2001). *System identification. A frequency domain approach*. New York: Wiley-IEEE Press.
- Poolla, K., Khargonekar, P. P., Tikku, A., Krause, J., & Nagpal, K. (1994). A time-domain approach to model validation. *IEEE Transactions on Automatic Control*, 39(5), 951–959.
- Popper, K. R. (1963). *Conjectures and refutations: the growth of scientific knowledge*. London: Routledge.
- Pronzato, L., & Walter, E. (1994). Minimum-volume ellipsoids containing compact sets: application to parameter bounding. *Automatica*, 30(11), 1731–1739.
- Rantzer, A. (1992). Convex robustness specifications for real parametric uncertainty in linear systems. In *Proceedings of the American control conference* (pp. 583–587).
- Rantzer, A., & Megretski, A. (1994). A convex parametrization of robustly stable controllers. *IEEE Transactions on Automatic Control*, 39(9), 1802–1808.
- Raynaud, H.-F. (1993). A criterion for robust stabilization with respect to parameter uncertainty variation in H_∞ controller synthesis. In *Proceedings of the second European control conference*.
- Raynaud, H.-F., Pronzato, L., & Walter, E. (1995). Robust identification and control based on ellipsoidal parametric uncertainty descriptions. *European Journal of Control*, 1(6), 245–255.
- Reinelt, W., Garulli, A., & Ljung, L. (2002). Comparing different approaches to model error modeling in robust identification. *Automatica*, 38(5), 787–803.
- Rivera, D. E., & Jun, K. S. (2000). An integrated identification and control design methodology for multivariable process system applications. *IEEE Control Systems Magazine*, 20(3), 25–37.
- Rivera, D. E., Lee, H., Braun, M. W., & Mittelmann, H. D. (2003). Plant-friendly system identification: a challenge for the process industries. In *Proceedings of the 13th IFAC symposium on system identification*, Rotterdam, Netherlands (pp. 917–922).
- Rivera, D. E., & Morari, M. (1987). Control-relevant model reduction problems for SISO H_2 , H_∞ and μ -controller synthesis. *International Journal of Control*, 46(2), 505–527.

- Rivera, D. E., Pollard, J. F., & Garcia, C. E. (1992). Control-relevant prefiltering: A systematic design approach and case study. *IEEE Transactions on Automatic Control*, 37(7), 964–974.
- Safonov, M. G., & Cabral, F. B. (2001). Fitting controllers to data. *Systems and Control Letters*, 43(4), 299–308.
- Safonov, M. G., & Tsao, T.-C. (1994). The unfalsified control concept and learning. In *Proceedings of the 33rd IEEE conference on decision and control*, vol. 3 (pp. 2819–2824).
- Safonov, M. G., & Tsao, Y.-C. (1997). The unfalsified control concept and learning. *IEEE Transactions on Automatic Control*, 42(6), 843–847.
- Scheid, R. E., Bayard, D. S., & Yam, Y. (1991). A linear programming approach to characterizing norm bounded uncertainty from experimental data. In *Proceedings of the American control conference* (pp. 1956–1958).
- Schoukens, J., Dobrowiecki, T., & Pintelon, R. (1998). Parametric and nonparametric identification of linear systems in the presence of nonlinear distortions—a frequency domain approach. *IEEE Transactions on Automatic Control*, 43(2), 176–190.
- Schoukens, J., Pintelon, R., & Dobrowiecki, T. (2002). Identification of the stability of feedback systems in the presence of nonlinear distortions. In *Proceedings of the 15th IFAC world congress*, Barcelona, Spain (pp. 418–423).
- Schoukens, J., Pintelon, R., Dobrowiecki, T., & Rolain, Y. (2003). Identification of linear systems with nonlinear distortions. In *Proceedings of the 13th IFAC symposium on system identification*, Rotterdam, Netherlands (pp. 1761–1772).
- Schrama, R. J. P. (1992). Accurate models for control design: the necessity of an iterative scheme. *IEEE Transactions on Automatic Control*, 37(7), 991–994.
- Schrama, R. J. P., & Bosgra, O. H. (1993). Adaptive performance enhancement by iterative identification and control design. *International Journal on Adaptive Control and Signal Processing*, 7(5), 475–487.
- Sjöberg, J., & De Bruyne, F. (1999). On a nonlinear controller tuning strategy. In *Proceedings of the 14th IFAC world congress*, vol. 1, Beijing, P.R. China (pp. 343–348).
- Sjöberg, J., De Bruyne, F., Agarwal, M., Anderson, B. D. O., Gevers, M., Kraus, F. J., & Linard, N. (2003). Iterative controller optimization for nonlinear systems. *Control Engineering Practice*, 11(9), 1079–1086.
- Skelton, R. E. (1989). Model error concepts in control design. *International Journal of Control*, 49(5), 1725–1753.
- Skogestad, S., & Postlethwaite, I. (1996). *Multivariable feedback control—analysis and design*. New York: Wiley.
- Sokolov, V. F. (1996). Closed-loop identification for the best asymptotic performance of adaptive robust control. *Automatica*, 32(8), 1163–1176.
- Sokolov, V. F. (2001). Adaptive ℓ_1 robust control for SISO systems. *Systems and Control Letters*, 42(5), 379–393.
- Tao, K.M., Kosut, R.L., & Ekblad, M. (1994). Feedforward learning—nonlinear processes and adaptation. In *Proceedings of the 33rd IEEE conference on decision and control* (pp. 1060–1065).
- Tempo, R., Bai, E.-W., & Dabbene, F. (1997). Probabilistic robustness analysis: Explicit bounds for the minimum number of samples. *Systems and Control Letters*, 30(5), 237–242.
- Tjörnström, F. (2002). Variance aspects of L_2 model reduction when undermodeling—the output error case. In *Proceedings of the 15th IFAC world congress*, Barcelona, Spain (pp. 1255–1260).
- Tjörnström, F. & Garulli, A. (2002). A mixed probabilistic bounded-error approach to parameter estimation in the presence of amplitude bounded white noise. In *Proceedings of the 41st IEEE conference on decision and control*, Las Vegas, Nevada, USA (pp. 3422–3427).
- Tjörnström, F., & Ljung, L. (2002). L_2 model reduction and variance reduction. *Automatica*, 38(9), 1517–1530.
- van de Vusse, J. (1964). Plug-flow type reactor versus tank reactor. *Chemical Engineering Science*, 19(12), 994–997.
- Van den Hof, P. M. J. (1998). Closed-loop issues in system identification. *Annual Reviews in Control*, 22, 173–186.
- Van den Hof, P. M. J., & Schrama, R. J. P. (1995). Identification and control—closed loop issues. *Automatica*, 31(12), 1751–1770.
- Veres, S. M. (1999). Iterative identification and control redesign via unfalsified sets of models: a basic scheme. *International Journal of Control*, 72(10), 887–903.
- Veres, S. M. (2000). Self-tuning control by model unfalsification (part I). *International Journal of Control*, 73(17), 1548–1559.
- Veres, S. M. (2001). Convergence of control performance by unfalsification of models—levels of confidence. *International Journal on Adaptive Control and Signal Processing*, 15(5), 471–502.
- Veres, S. M., Messaoud, H., & Norton, J. P. (1999). Limited-complexity model-unfalsifying adaptive tracking-control. *International Journal of Control*, 72(15), 1417–1426.
- Veres, S. M., & Sokolov, V. F. (1998). Adaptive robust control under unknown plant orders. *Automatica*, 34(6), 723–730.
- Veres, S. M., & Wall, D. S. (2000). *Synergy and duality of identification and control*. London: Taylor & Francis.
- Vinnicombe, G. (1993). Frequency domain uncertainty and the graph topology. *IEEE Transactions on Automatic Control*, 38(9), 1371–1383.
- Vuerinckx, R., Pintelon, R., Schoukens, J., & Rolain, Y. (2001). Obtaining accurate confidence regions for the estimated zeros and poles in system identification problems. *IEEE Transactions on Automatic Control*, 46(4), 656–659.
- Wahlberg, B. (1989). Model reduction of high-order estimated models: The asymptotic ML approach. *International Journal of Control*, 49(1), 169–192.
- Wahlberg, B., & Ljung, L. (1992). Hard frequency-domain model error bounds from least-squares like identification techniques. *IEEE Transactions on Automatic Control*, 37(7), 900–912.
- Weyer, E., & Campi, M. C. (2002). Non-asymptotic confidence ellipsoids for the least-squares estimate. *Automatica*, 38(9), 1539–1547.
- Xia, H., & Veres, S. M. (1999). Improved efficiency of adaptive robust control by model unfalsification. *Automatica*, 35(5), 981–986.
- Xie, L.-L., & Ljung, L. (2001). Asymptotic variance expressions for estimated frequency functions. *IEEE Transactions on Automatic Control*, 46(12), 1887–1899.
- Zacks, S. (1971). *The theory of statistical inference*. New York: Wiley.
- Zames, G., & El-Sakkary, A. K. (1980). Unstable systems and feedback: The gap metric. In *Proceedings of the Allerton conference* (pp. 380–385).
- Zang, Z., Bitmead, R. R., & Gevers, M. (1995). Iterative weighted least-squares identification and weighted LQG control design. *Automatica*, 31(11), 1577–1594.
- Zehna, P. W. (1966). Invariance of maximum likelihood estimation. *Annals of Mathematical Statistics*, 37, 755.
- Zhao, J., & Kanellakopoulos, I. (2002). Active identification for discrete-time nonlinear control—part I: output-feedback systems. part II: strict-feedback systems. *IEEE Transactions on Automatic Control*, 47(2), 210–240.
- Zhou, K., Doyle, J. C., & Glover, K. (1996). *Robust and optimal control*. Englewood Cliffs, NJ: Prentice-Hall.
- Zhu, Y. (1998). Multivariable process identification for MPC: the asymptotic method and its applications. *Journal of Process Control*, 8(2), 101–115.
- Zhu, Y. (2000). Use of error criteria in identification for control. In *Proceedings of the 12th IFAC symposium on system identification*, Santa Barbara, CA.
- Zhu, Y. (2001). *Multivariable system identification for process control*. Oxford: Pergamon Press.
- Zhu, Y., & Butoyi, F. (2002). Case studies on closed-loop identification for MPC. *Control Engineering Practice*, 10(4), 403–417.



Håkan Hjalmarsson was born in 1962. He received the M.S. degree in Electrical Engineering in 1988, and the Licentiate degree and the Ph.D. degree in Automatic Control in 1990 and 1993, respectively, all from Linköping University, Sweden. He has held visiting research positions at California Institute of Technology, Louvain University and the University of Newcastle, Australia. He has served as an Associate Editor for *Automatica* and been a Guest Editor for *European*

Journal of Control and Control Engineering Practice. He is currently an Associate Editor for *IEEE Transactions on Automatic Control* and Professor in the Department of Signals, Sensors and Systems at the Royal Institute of Technology (KTH), Stockholm, Sweden. He is a member of the IFAC Technical Committee on Modeling, Identification and Signal Processing since 1996. In 2001 he received the KTH award for outstanding contribution to undergraduate education. His research interests include system identification, signal processing, control and estimation in communication networks and automated tuning of controllers.