

Probabilistic robust design of control systems for high-fidelity cyber-physical testing

Thomas Sauder^{a,c}, Stefano Marelli^b, Asgeir J. Sørensen^a

^aCentre for Autonomous Marine Operations and Systems, Department of Marine Technology, NTNU, 7491 Trondheim, Norway

^bChair of Risk, Safety and Uncertainty Quantification, Institute of Structural Engineering, ETH Zürich, Stefano-Franscini-Platz 5, CH-8093 Zürich, Switzerland

^cSINTEF Ocean, P.O. Box 4762 Torgard, 7465 Trondheim, Norway

Abstract

Cyber-physical empirical methods consist in partitioning a dynamical system under study into a set of physical and numerical substructures that interact in real-time through a control system. In this paper, we define and investigate the *fidelity* of such methods, that is their capacity to generate systems whose outputs remain close to those of the original system under study. In practice, fidelity is jeopardized by uncertain and heterogeneous artefacts originating from the control system, such as actuator dynamics, time delays and measurement noise. We present a computationally efficient method, based on surrogate modeling and active learning techniques, to (1) verify that a cyber-physical empirical setup achieves probabilistic robust fidelity, and (2) to derive fidelity bounds, which translate to absolute requirements to the control system. For verification purposes, the method is first applied to the study of a simple mechanical system. Its efficiency is then demonstrated on a more complex problem, namely the active truncation of slender marine structures, in which the substructures' dynamics cannot be described by an analytic solution.

Key words: cyber-physical empirical method, fidelity, artefacts, probabilistic robustness, adaptive kriging

1 Introduction

Empirical methods constitute the cornerstone of most scientific fields. They consist in generating knowledge about a system *through observation*, to verify hypotheses and build models of the reality that surrounds us. Cyber-physical empirical methods (CPEMs) have been developed and applied in diverse engineering fields such as civil engineering [McCrum and Williams, 2016], aerospace [Wallace et al., 2007], automotive [Misselhorn et al., 2006], engine development [Filipi et al., 2006], electrical power engineering [Edrington et al., 2015], thermomechanics [Whyte et al., 2015], and floating wind energy [Sauder et al., 2016, Bayati et al., 2017]. CPEM aim at studying a dynamical system by partitioning it into physical and numerical *substructures* that interact in real-time through a control system. While the behaviour of the *physical substructures* is partly unknown, the *nu-*

merical substructures are described by state-of-the-art, validated and, in general, non-analytic computational models. CPEM therefore augment classical empirical methods with validated numerical models, to address problems that classical empirical methods alone, or models alone, can not conveniently or reliably address. This is for example the case: (1) *when the system under study is “ill-conditioned”*, i.e. when it contains a large span of characteristic spatial dimensions and/or time constants. In that case, the part of the system that does not fit in the laboratory, or whose dynamic is slow, can advantageously be replaced by a numerical model [Sauder et al., 2018]. (2) *when scaling effects should be tackled*, such as in the case of model testing of floating wind turbines [Allen and Goupee, 2017]. (3) *when component testing should be performed*, that is when the focus is on the performance of a specific uncertain substructure, that is interacting with the other substructures as part of a complex system.

Email addresses: thomas.sauder@sintef.no (Thomas Sauder), marelli@ibk.baug.ethz.ch (Stefano Marelli), asgeir.sorensen@ntnu.no (Asgeir J. Sørensen).

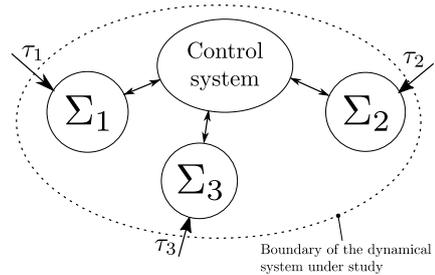
The control system at the heart of a CPEM plays a crucial role for its validity as an empirical method. Indeed, the dynamical properties of a cyber-physical empirical

setup must reflect those of the real system under study, and this, in spite of heterogeneous *artefacts* such as time delay or sensor noise, which are inevitably introduced by the components of the control system. This performance requirement, which we denote *high fidelity*, is the subject of this paper. Even if CPEM are developed in a controlled laboratory environment, some uncertainty is entailed to the artefacts: sensor noise variance, or the interconnection delays between the substructures, remain for example uncertain at the design stage, and can be quantified accurately only when the setup is realized. However, the *amount of uncertainty* on these quantities can be estimated from expert judgment [Jaynes, 1957] or dedicated surveys [Schwarz, 1978], and modeled within a probabilistic framework.

In this setting, our objective is to analyze whether a CPEM achieves *probabilistic robust fidelity*, i.e. high fidelity with a probability $1 - P_f$ larger than a defined admissible threshold $1 - \varepsilon_{\text{adm}}$, with $\varepsilon_{\text{adm}} \in (0, 1)$. An account of probabilistic robust analysis and design methods for control systems has been given in [Calafiore et al., 2011]. Their advantages in terms of conservatism and complexity when compared to the classical *worst-case* robustness approaches have been discussed in details in [Tempo et al., 2013, Chapter 5], and are illustrated by a simple case study in [Chen et al., 2005]. The main drawback of probabilistic approaches is that estimating P_f implies the evaluation of a multidimensional integral, which is challenging in high-dimensions and when non-analytic models are involved. As we will see, this is a typical situation encountered when dealing with CPEMs. To address such problems, randomized algorithms have been developed [Tempo et al., 2013], which rely on *sampling* the performance function. The required number of samples is however shown to be large for low P_f , i.e for highly reliable systems [Alamo et al., 2015].

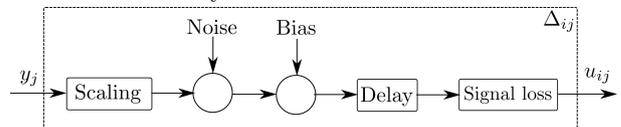
The main scientific contribution of this paper is the development of a new method to verify the *probabilistic robust fidelity* of CPEMs, and to derive *fidelity bounds*, which can be used as specifications to the control system. The devised method (1) is non-intrusive, and thus not limited to analytic models, which allows its application to the wide class of dynamical systems studied with CPEMs, (2) allows handling an arbitrary *number* and *type* of artefacts, which exhibit parametric uncertainty, and (3) is based on surrogate modelling and active learning techniques to achieve unprecedented computational efficiency, even for high-dimensional and high-reliability problems.

The paper is organized as follows. In Section 2, we define quantitatively the *fidelity* φ of a CPEM, and formulate the robust fidelity and fidelity bounds problems. The solution to these problems is described in Section 3, and illustrated by two example cases in Section 4. Finally, conclusions are given in Section 5.



(a) Generic representation of a CPEM with three substructures Σ_i , and exogenous excitation τ .

(b) Interconnected set of three substructures Σ_i subjected to artefacts Δ_{ij} . Note that the y_i output signals are doubled here for clarity.



(c) Synthesis of each Δ_{ij} from *elementary* artefacts, examples of which are listed in Table 1.

Figure 1. Model of the cyber-physical experimental setup.

Notations. We let $\mathbb{N}_Q = \{0, 1, \dots, Q\}$, $\mathbb{N}^* = \mathbb{N} \setminus \{0\}$ and $\mathbb{N}_Q^* = \mathbb{N}_Q \setminus \{0\}$. The uniform distribution with support $[a, b]$ is denoted $\mathcal{U}(a, b)$, and $\mathcal{N}(\mu, \Sigma)$ denotes the multivariate Gaussian distribution with mean μ and covariance matrix Σ . The symbol \sim means “distributed as”. For $i \in \mathbb{N}^*$, $\theta^{(i)}$ refers to a sample of the random vector Θ , while θ_i refers to the i^{th} component of θ .

2 Problem statement and definition of fidelity

Let s denote the total number of substructures, and let $\tau(t) = (\tau_1(t), \dots, \tau_s(t))$ represent an exogenous excitation signal, with support $[0, T]$, acting on the substructural partition. A generic way of modelling CPEM, suitable for all the applications listed in Section 1, is

the following interconnected system: $\forall i \in \mathbb{N}_s^*$, $\dot{x}_i = f_i(x_i, u_{ij}, \tau_i)$ and $y_i = h_i(x_i)$, where x_i is the internal state of substructure Σ_i , and y_i its output. u_{ij} is the input to Σ_i originating from y_j . Following e.g. bond graph terminology, if y_i is a *flow*, then $\forall j \neq i$, u_{ij} must be *efforts*, and vice versa. The interconnections between the substructures are subjected to heterogeneous artefacts, such as measurement noise or communication-induced time delays: $\forall i \in \mathbb{N}_s^*$, $\forall j \in \mathbb{N}_s^* \setminus \{i\}$, $\forall t > 0$, $u_{ij}(t) = \Delta_{ij}((y_j(t'))_{t' \in [0, t]}, \theta)$, where θ is a M -tuple parameter defined on a domain $\mathcal{D} \subset \mathbb{R}^M$, that describes the artefact's properties for the whole system. The components of θ may include sensor noise variance, or the amount of communication-induced time delay between two substructures. Fig. 1 presents this structure for $s = 3$. The effects of given artefacts on a reference signal are illustrated in [Sauder et al., 2018].

Owing to the presence of *numerical* substructures among the Σ_i , studying the robustness of such an interconnection would in principle imply the use of digital or hybrid control theory, see e.g. [Tabuada et al., 2014, Goebel et al., 2012]. In the particular case of CPEM, however, the “cyber” parts of the system (the numerical substructures) represent continuous dynamical systems. Assuming that adequate time-stepping algorithms are employed, see e.g. [Bursi et al., 2013], we can make abstraction of their digital nature, and treat the problem as a classical probabilistic robust control problem. Note, however, that some aspects related to the digital nature of the hardware, such as quantization, sampling, and zero-order hold, can be approached by the present method, by modeling them as artefacts.

We will in the following now introduce the definition of the fidelity of a CPEM. For $i \in \mathbb{N}_s^*$, let \bar{x}_i denote the state of Σ_i when $\forall j \neq i$, $u_{ij} \equiv y_j$ (i.e. without artefacts). For a given $Q \in \mathbb{N}^*$, let $(\gamma_q)_{q \in \mathbb{N}_Q^*}$, be a family of cost functions satisfying

$$\forall q \in \mathbb{N}_Q^*, (\forall i \in \mathbb{N}_s, x_i |_{\theta} \rightarrow \bar{x}_i) \Rightarrow (\gamma_q \rightarrow 0) \quad (1)$$

Definition 1 *The fidelity φ is defined on \mathcal{D} as*

$$\varphi(\theta) := -\frac{1}{2} \log \sum_{q \in \mathbb{N}_Q^*} \gamma_q^2(\{x_i |_{\theta}(t), \bar{x}_i(t)\}_{i \in \mathbb{N}_s^*, t \in [0, T]}) \quad (2)$$

The rationale behind Definition 1 is the following. (1) Each γ_q function compares selected quantities of interest (QoI) derived from the states $(x_1, x_2, \dots, x_s) |_{\theta}$ with the corresponding QoI derived from $(\bar{x}_1, \bar{x}_2, \dots, \bar{x}_s)$. If all states x_i converge towards \bar{x}_i , then all γ_q tend to zero, and $\varphi \rightarrow \infty$. Fidelity quantifies therefore the capability of the CPEM to generate QoI that are similar to the real system, when subjected to same excitation. (2) The reciprocal is however not true: high fidelity can be achieved

even if some states x_i which are *not* of interest, i.e. not included in the calculation of any γ_q , differ from \bar{x}_i . This is a major difference with the concept of *resilient* cyber-physical systems, see e.g. [Fawzi et al., 2014]. A high fidelity value does not imply a correct estimation of the complete state x_i in presence of artefacts, but rather a correct estimation of selected state-derived quantities of interest. This will be illustrated in Example case II. (3) If the cyber-physical empirical setup becomes unstable because of the introduced artefacts, some γ_q may blow up in some domains of \mathcal{D} . On the other hand, when studying *high-fidelity* setups, we may be interested in emphasizing the difference between small values of the γ_q . The logarithm is introduced for this reason. (4) A sum of the squares, rather a *maximum* function, is used to combine the cost functions γ_q , which preserves the smoothness properties of the functions $\theta \mapsto \gamma_q(\theta)$. Using a *maximum* function instead would have compromised the differentiability of φ even if the γ_q were smooth functions. This choice will prove convenient when analyzing the problem.

As a probabilistic robust approach to fidelity is considered, we assume the artefact parameter θ to be the realization of a random vector Θ with a known (but *arbitrary*) distribution f_{Θ} . Given a minimum admissible fidelity $\varphi_{\text{adm}} \in \mathbb{R}$, the two problems addressed in the in the present paper can then be formulated as follows. (1) *Robust fidelity*: does $P[\varphi(\Theta) < \varphi_{\text{adm}}] < \varepsilon_{\text{adm}}$ hold? In other words, given a dynamical system and a substructural partition, can artefacts that are likely to be introduced by the control system lead to an unacceptable loss of fidelity? (2) *Fidelity bounds*. What is the set $\{\theta \in \mathcal{D} | \varphi(\theta) \geq \varphi_{\text{adm}}\}$, i.e. the absolute constraints on the control system to guarantee sufficiently high fidelity?

3 Solution

Let us first pinpoint that solving problems (1) and (2) above is equivalent to identifying the *domain of failure* $\mathcal{D}_f \subset \mathbb{R}^M$ defined by

$$\mathcal{D}_f = \{\theta \in \mathcal{D} | \varphi(\theta) < \varphi_{\text{adm}}\}. \quad (3)$$

Indeed, letting $\mathbb{I}_{\mathcal{D}_f}(\theta)$ be the indicator function for the set \mathcal{D}_f , the probability of failure $P_f := P[\varphi(\Theta) < \varphi_{\text{adm}}]$ is simply

$$P_f = \int_{\mathcal{D}_f} f_{\Theta}(\theta) d\theta = \int_{\mathcal{D}} \mathbb{I}_{\mathcal{D}_f}(\theta) f_{\Theta}(\theta) d\theta \quad (4)$$

Because of the non-analytic character of the f_i functions, a sampling-based approach must be employed to determine \mathcal{D}_f . A naive Monte-Carlo simulations (MCS) approach, would consist in drawing N samples of Θ , calculate $\varphi(\theta^{(i)})$ for each of them, and then compare the result to φ_{adm} . P_f could then be estimated by

Table 1

Some examples of elementary artefacts, their possible sources, and the corresponding describing parameters (component of θ). Fourth and fifth columns: probabilistic description of the artefacts involved in example cases I and II (Section 4). The signal loss artefact is parametrized by a probability of occurrence $\zeta_1 \in [0, 1)$ and an inverse duration parameter $\zeta_2 > 0$. When it occurs, the duration d of the signal loss is distributed as $f_D(d) = \zeta_2 e^{-\zeta_2 d}$. For example case II, x and z refer to the horizontal and vertical components of the force measurement, and *sens.* and *act.* refer to the artefacts modelling the sensing and actuation part of the loop, respectively. For more details, see Section 4.2.

Type of artefact	Example of sources	Describing parameter(s)	Example case I	Example case II
Gain	sensor calibration error	scaling factor	$\Theta_1 \sim \mathcal{N}(1, 0.005^2)$ [-]	$\Theta_1 \sim \mathcal{N}(1, 0.015^2)$ [-] (x)
	actuator limitations		$\Theta_3 \sim \mathcal{N}(1, 0.005^2)$ [-]	$\Theta_2 \sim \mathcal{N}(1, 0.015^2)$ [-] (z)
Bias	zero measurement	bias value	-	$\Theta_3 \sim \mathcal{N}(1, 0.05^2)$ [N] (x)
	sensor drift		-	$\Theta_4 \sim \mathcal{N}(1, 0.05^2)$ [N] (z)
Noise	measurement noise	variance	-	$\Theta_5 \sim \mathcal{U}(0.025^2, 0.05^2)$ [N ²] (x)
			-	$\Theta_6 \sim \mathcal{U}(0.025^2, 0.05^2)$ [N ²] (z)
Constant time delay	computational time	duration	$\Theta_2 \sim \mathcal{U}(0, 20)$ [ms]	$\Theta_7 \sim \mathcal{U}(0, 5)$ [ms] (sens.)
	communication time		$\Theta_4 \sim \mathcal{U}(0, 20)$ [ms]	$\Theta_{10} \sim \mathcal{U}(0, 5)$ [ms] (act.)
Signal loss	unfinished iteration process	probability of occurrence	-	$\Theta_8 \sim \mathcal{U}(1\%, 10\%)$ [-] (sens.)
			-	$\Theta_{11} \sim \mathcal{U}(1\%, 10\%)$ [-] (act.)
	faulty measurement	inverse duration parameter	-	$\Theta_9 \sim \mathcal{U}(0.77, 3.87)$ [s ⁻¹] (sens.)
			-	$\Theta_{12} \sim \mathcal{U}(0.77, 3.87)$ [s ⁻¹] (act.)
Quantization	analog to digital conversion	resolution	-	-
Zero-order hold	sampling	sampling period	-	-
Saturation	actuator limitations	limit(s)	-	-

$N^{-1} \sum_{i \in \mathbb{N}_N^*} \mathbb{I}_{\mathcal{D}_f}(\theta^{(i)})$. However, since $\mathbb{I}_{\mathcal{D}_f}(\theta^{(i)})$ follows a binomial distribution, estimating P_f with a confidence interval of δ (typically 5%), requires performing at least $N \geq (1 - P_f)(\delta^2 P_f)^{-1}$ simulations (see also [Chernoff, 1952] for bounds not depending directly on P_f). This is practically unfeasible when the f_i functions are computationally expensive, and when P_f is small, which is *in fine* what is expected from the control system design.

Our strategy to characterize \mathcal{D}_f and evaluate P_f is to replace $\varphi(\theta)$ by a suitable *surrogate* model, which can be interrogated, and with which MCS can be performed at a negligible cost. The formulation and identification of this model will be outlined in Sections 3.1 and 3.2, respectively. Using active learning, the surrogate model is refined by parsimonious and targeted sampling of $\varphi(\Theta)$. This will be detailed in Section 3.3. Before proceeding, let us note that sampling $\varphi(\theta)$ implies that the system described in Section 2 must be *co-simulated* for $t \in [0, T]$, accounting for the presence of artefacts. To do so, an iterative procedure must ensure flow- and effort-consistency between the substructures and artefacts at each time step. Providing details regarding the co-simulation is not within the scope of the present paper, but the interested reader may consult [Sauder et al., 2018] for details and examples.

3.1 Polynomial-chaos kriging (PCK) model of $\varphi(\theta)$

In the following, we assume that an initial set \mathcal{E} of N samples of Θ has been generated using a *space-filling* sampling method, such as latin hypercube sampling. The corresponding values of φ are evaluated, and gathered in $\mathcal{F} = [\varphi(\theta^{(1)}), \varphi(\theta^{(2)}), \dots, \varphi(\theta^{(N)})]^\top$. We will now introduce the chosen surrogate model for φ , which combines the benefits of polynomial chaos-based and kriging-based models to surrogate complex functions [Schöbi et al., 2015].

Definition 2 [Schöbi et al., 2016] *A Polynomial-Chaos Kriging (PCK) surrogate model is a Gaussian random process $\mathcal{K}(\theta, \omega)$ of the form:*

$$\mathcal{K}(\theta, \omega) = \sum_{\alpha \in \mathcal{A}} a_\alpha \psi_\alpha(\theta) + \sigma^2 Z(\theta, \omega, l) \quad (5)$$

where Z is a zero-mean, unit variance stationary Gaussian process with variance σ^2 , ω is the underlying probability space represented by a parametric correlation function $\mathcal{R}(\theta, \theta', l)$, while the $(\psi_\alpha)_{\alpha \in \mathcal{A}}$ form a sparse polynomial chaos expansion, i.e a wisely selected finite set of orthonormal polynomials with respect to Θ [Blatman and Sudret, 2011]. The (ψ_α) are used as regressors, and are weighted by a set of coefficients (a_α) , gathered in a $|\mathcal{A}|$ -tuple which we denote β . Z is assumed in this work to be described by the general Matérn correlation function¹.

A PCK constitutes a surrogate model, since its coefficients β , σ and l can be tuned based on \mathcal{E} and \mathcal{F} , so that the mean and most probable value of \mathcal{K} at a point θ , i.e. $\mathbb{E}[\mathcal{K}(\theta, \omega)]$, surrogates $\varphi(\theta)$. Furthermore, the key aspect of this model is that the *uncertainty of the PCK estimate can be quantified* through $\text{Var}[\mathcal{K}(\theta, \omega)]$. Indeed, assuming that the coefficients β , σ , and l , involved in (5) are known, the following statement can be made.

¹ The general Matérn correlation function is defined by:

$$\mathcal{R}(\theta, \theta', l) = 2^{1-\nu} \Gamma^{-1}(\nu) (2\nu h^2)^{\nu/2} K_\nu(\sqrt{2\nu h^2}) \quad (6)$$

where $h^2 = \sum_{i=1}^M \left(\frac{|\theta_i - \theta'_i|}{l_i} \right)^2$, and where $l \in (\mathbb{R}^+)^M$ contains scale parameters describing the amount of correlation between neighbours for each component θ_i . Γ and K_ν are the Gamma and the modified Bessel function of the second kind, respectively. ν is the shape parameter related to the smoothness properties of the process. See e.g. [Dubourg, 2011].

Claim 3 $\varphi(\theta)$ can be estimated at any point in \mathcal{D} by

$$\varphi(\theta) \approx \mu_{\mathcal{K}}(\theta) := f(\theta)^\top \beta + r(\theta)^\top R^{-1}(\mathcal{F} - F\beta) \quad (7)$$

with an uncertainty

$$\sigma_{\mathcal{K}}^2(\theta) := \sigma^2(1 - r(\theta)^\top R^{-1}r(\theta)^\top) \quad (8)$$

where $f(\theta) \in \mathbb{R}^{|\mathcal{A}| \times 1}$ contains the $\psi_\alpha(\theta)$ polynomials evaluated at θ , $F \in \mathbb{R}^{|\mathcal{A}| \times N}$ contains, in each column the $\psi_\alpha(\theta)$ polynomials evaluated at $\theta^{(i)}$, $r \in \mathbb{R}^{N \times 1}$ is the cross-correlation vector with $\forall i \in \mathbb{N}_N^*$, $r_i(\theta) = \mathcal{R}(\theta, \theta^{(i)}, l)$, and $R \in \mathbb{R}^{N \times N}$ is the correlation matrix with $\forall i, j \in \mathbb{N}_N^*$, $R_{ij} = \mathcal{R}(\theta^{(i)}, \theta^{(j)}, l)$.

PROOF. By definition, a Gaussian process is a collection of random variables, any finite number of which have a joint Gaussian distribution. From Definition 2, $[\varphi(\theta), \mathcal{F}^\top]^\top$ is then the realization of a random vector $V \in \mathbb{R}^{N+1}$ distributed as:

$$V \sim \mathcal{N} \left(\begin{bmatrix} f(\theta)^\top \beta \\ F\beta \end{bmatrix}, \sigma^2 \begin{bmatrix} 1 & r(\theta)^\top \\ r(\theta) & R \end{bmatrix} \right) \quad (9)$$

According to [Santner et al., 2003, Theorem 3.2.1], the estimator of $\varphi(\theta)$ with minimum mean-squared prediction error is $\mathbb{E}[V_1 | [V_2, \dots, V_N]^\top]$. As derived in [Eaton, 2007, p.117], this conditional expectation is a Gaussian random variable, whose mean $\mu_{\mathcal{K}}(\theta)$, and thus most probable value, is equal to the right hand side of (7). Its variance is given by (8). \square

It is important to make the distinction between the randomness of Θ and the randomness introduced through ω in Definition 2. While the former is a reality of our robust fidelity problem, the latter gives us a *tool to model the epistemic uncertainty* on $\varphi(\theta)$ at locations where it has not been evaluated, and leads to the result stated in Claim 3.

3.2 Determination of the PCK model parameters

The PCK model coefficients in (5) are found based on the available data $(\mathcal{E}, \mathcal{F})$ as follows. In a first step, an optimal sparse set of polynomials $(\psi_\alpha)_{\alpha \in \mathcal{A}}$ is determined through solution of an ℓ_1 -regularized sparse least squares problem with the so-called *least angle regression* method. The exact procedure is outlined in [Blatman and Sudret, 2011], and not detailed here due to space constraints. Then,

Claim 4 the coefficients β and σ^2 can be estimated by:

$$\hat{\beta} = (F^\top R^{-1}F)^{-1} F R^{-1} \mathcal{F} \quad (10)$$

$$\hat{\sigma}^2 = \frac{1}{N} (\mathcal{F} - F\hat{\beta})^\top R^{-1} (\mathcal{F} - F\hat{\beta}) \quad (11)$$

PROOF. Let us assume in a first stage that l in (6) is fixed. The optimal coefficients $\hat{\beta}$ and $\hat{\sigma}^2$ in (5) are those which lead to a process \mathcal{K} whose *most likely realization* is $\varphi(\theta)$. Since by construction, \mathcal{K} is a Gaussian process, the optimal coefficients $\hat{\beta}$ and $\hat{\sigma}^2$ in (5) are those which maximize the likelihood, or log-likelihood of \mathcal{F} , expressed as

$$\begin{aligned} \log L(\mathcal{F} | \beta, \sigma^2) = \\ -\frac{N}{2} \log 2\pi\sigma^2 - \frac{1}{2} \log |R| - \frac{1}{2\sigma^2} (\mathcal{F} - F\beta)^\top R^{-1} (\mathcal{F} - F\beta) \end{aligned} \quad (12)$$

The values of β and σ^2 maximizing this expression necessarily satisfy the first-order optimality conditions $\nabla_\beta(\log L) = 0$ and $\nabla_{\sigma^2}(\log L) = 0$. Since R is symmetric, they read

$$F^\top R^{-1}(\mathcal{F} - F\beta) = 0 \quad (13)$$

$$-\frac{N}{2} \frac{1}{\sigma^2} + \frac{1}{2\sigma^4} (\mathcal{F} - F\beta)^\top R^{-1} (\mathcal{F} - F\beta) = 0 \quad (14)$$

which leads to (10) and (11). \square

The hyperparameter l in \mathcal{R} is determined as the one minimizing the *leave-one-out cross validation error* E_{LOO} of the PCK model. In principle, E_{LOO} is calculated as follows: a PCK model is established as described above, using *all but one* samples in \mathcal{E} , and the error between the PCK model and the actual value of φ is evaluated *at the sample that was not included*. This operation is repeated for each sample in \mathcal{E} , and the E_{LOO} is obtained from the arithmetic mean of the N results, see [Bachoc, 2013]. In the present work, we used a genetic algorithm [Goldberg, 1989] to find the global optimum l .

3.3 Adaptive Kriging (AK)

We recall that our objective is to characterize \mathcal{D}_f . At the initial stage, N is in general chosen large enough to obtain good estimates of the expected fidelity $\mathbb{E}[\varphi(\Theta)]$, the variance $\text{Var}[\varphi(\Theta)]$, and some useful sensitivity information [Sauder et al., 2018], but as we will see in the Example cases, it is generally too small to properly characterize \mathcal{D}_f . However, by using the PCK model, estimates of $\varphi(\theta)$ and the associated uncertainty are available anywhere in \mathcal{D} from (7) and (8). Taking advantage of this information, the surrogate model \mathcal{K} can be refined by active learning as suggested in [Schöbi et al., 2016], and summarized in the following.

Let the *limit state* hypersurface be the set $\mathcal{L} = \{\theta \in \mathcal{D} | \varphi(\theta) = \varphi_{\text{adm}}\}$, and Φ be the standard normal cumulative distribution function. First, a 5% confidence

bound about the estimated \mathcal{L} is defined through the sets $\mathcal{D}_f^+ = \{\theta \in \mathcal{D} | \mu_{\mathcal{K}}(\theta) - k\sigma_{\mathcal{K}}(\theta) < \varphi_{\text{adm}}\}$, and $\mathcal{D}_f^- = \{\theta \in \mathcal{D} | \mu_{\mathcal{K}}(\theta) + k\sigma_{\mathcal{K}}(\theta) < \varphi_{\text{adm}}\}$, with $k = \Phi^{-1}(97.5\%) \approx 1.96$. The relative complement $\mathbb{M}_f = \mathcal{D}_f^+ \setminus \mathcal{D}_f^-$ is denoted the *limit state margin* \mathbb{L} [Dubourg, 2011]: it defines a location where (1) φ is probably close to φ_{adm} , and/or (2) there is a large uncertainty on the value of φ . This is where $K \in \mathbb{N}^*$ new relevant samples of Θ should be generated, and added to \mathcal{E} to refine the PCK model.

These K samples are selected among a large set of *auxiliary* samples of Θ (with typical cardinality 10^5 - 10^6) generated using MCS. The *probability of misclassification* P_m of each sample located in \mathbb{L} is evaluated as

$$P_m(\theta) := \Phi\left(-\frac{|\mu_{\mathcal{K}}(\theta) - \varphi_{\text{adm}}|}{\sigma_{\mathcal{K}}(\theta)}\right) \quad (15)$$

which corresponds to the probability of the PCK model to predict that θ is in \mathcal{D}_f while it is actually not, or vice-versa. If the enrichment is performed with $K = 1$, \mathcal{E} is simply enriched with the sample θ^{new} featuring the largest P_m . If $K > 1$, K clusters of points in \mathbb{L} with high P_m values are identified, and a sample representative of each cluster is selected using a K -means clustering algorithm [Zaki and Meira, 2014, Chapter 13]. Both approaches will be illustrated in the example cases. The enriched sets \mathcal{E} and \mathcal{F} are then used to define a new, refined, PCK model as described in Section 3.2. As it will be demonstrated in the following example cases, the locations of \mathcal{L} and \mathcal{D}_f become, step by step, more exact, and the probability of failure P_f can be estimated from (4) by performing MCS *using the PCK model* and the auxiliary sample set.

The uncertainty on the estimated P_f is quantified by comparing the probabilities of failure estimated from \mathcal{D}_f^+ and \mathcal{D}_f^- , respectively. If the difference between them is less than 5% of P_f (evaluated from the estimated \mathcal{D}_f), convergence is considered to be achieved. Convergence towards the true value of P_f when $N \rightarrow \infty$ is guaranteed by the fact that (1) the learning function (15) weights uncertain areas of \mathcal{D} , and thus eventually explores the whole space \mathcal{D} , and (2) the resulting PCK (5) will interpolate exactly $\varphi(\theta)$ since it is an universal approximation function. Note that no *formal* proof of optimal convergence, as compared to other established methods, is available for this algorithm. It has, however, been shown to be more computationally efficient (see e.g. [Echard et al., 2011]) and more precise (see e.g. [Schöbi et al., 2016]) than existing methods, and is able to tackle non-connected failure regions \mathcal{D}_f .

3.4 Summary of the method

We have outlined a procedure that generates a PCK model \mathcal{K} surrogating $\varphi(\theta)$, with particularly high accu-

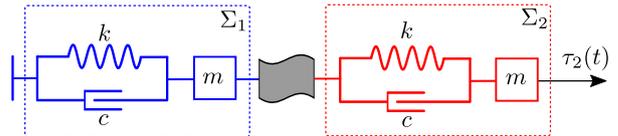


Figure 2. Case I. A two mass-spring-damper system. The grey flag-shaped block corresponds to the location of the control system.

racy near the limit state \mathcal{L} . As a result, the probabilistic robust fidelity of a cyber-physical empirical setup (problem 1) can simply be analyzed by interrogating the PCK model on a large set of auxiliary samples and using (4). The PCK model can also be interrogated to verify that the *as-built* setup, characterized by an estimated θ , satisfies the fidelity bounds (problem 2). Finally, in a design phase, the optimal control system can be found as the one minimizing some cost function $c(\theta)$, while ensuring high-enough fidelity:

$$\underset{\theta \in \mathcal{D}}{\text{minimize}} c(\theta) \text{ subject to } \mu_{\mathcal{K}}(\theta) - \varphi_{\text{adm}} \geq 0 \quad (16)$$

4 Example cases

The devised method will be demonstrated on two examples, which differ by the type of dynamical system under study, but also by the number and nature of the considered artefacts. Case I serves as a verification of the method. The mechanical system is simple, analytically described, and subjected to a small number of artefacts. It can therefore conveniently be analyzed by classical methods. Case II is an industrial problem originating from the field of ocean engineering, arguably difficult to approach with classical methods. In both cases, the cyber-physical empirical setup contains $s = 2$ substructures, interconnected as shown in Fig. 1b. Σ_1 is numerical, and Σ_2 is physical (but also simulated here, for the purpose of the analysis). Σ_2 is subjected to an exogenous load τ_2 . It is decided to *measure* the force (effort) y_2 at the interface between the substructures, and to *control* the linear velocity (flow) y_1 of the interface. The γ_q functions involved in the definition of the fidelity (2) will be detailed for each case. The UQLab framework is used to perform the probabilistic analysis, see [Marelli and Sudret, 2014].

4.1 Case I: coupled oscillators

We first consider the linear oscillators coupled and substructured as represented in Fig. 2. We set $m = 1\text{kg}$, $k = 1\text{N/m}$, and $c = 1\% \times 2\sqrt{km}$, leading to a lightly damped system. We investigate its response to a step force τ_2 acting on mass m_2 , from 0 to 1N with $T = 100\text{s}$. The *velocity* V_2 of the mass m_2 is the QoI, and the following cost function ($Q = 1$) is used in definition of the

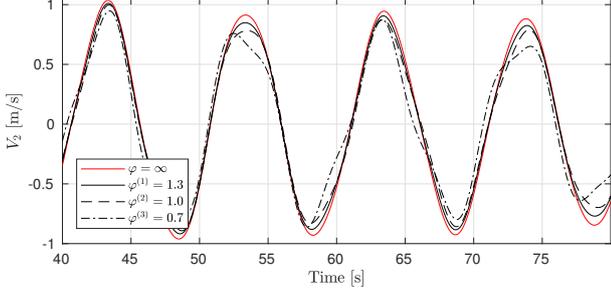


Figure 3. Case I. Time series of the QoI, in the absence and presence of artefacts (red and black curves, respectively). For the latter case, three different samples of Θ , leading to three values of φ are shown: $\theta^{(1)} = (1.007, 4.23, 0.993, 6.75)^\top$ (solid line), $\theta^{(2)} = (1.004, 10.33, 0.998, 12.31)^\top$ (dashed line), and $\theta^{(3)} = (0.990, 18.53, 0.993, 17.17)^\top$ (dash-dot line).

fidelity (2):

$$\gamma_1 = \left(\frac{\int_0^T (V_2(t)|_\theta - \bar{V}_2(t))^2 dt}{\int_0^T \bar{V}_2(t)^2 dt} \right)^{\frac{1}{2}} \quad (17)$$

This cyber-physical empirical setup is subject to the following set of artefacts. A calibration error (constant gain α_s) and a data processing-induced time delay τ_s are affecting the force measurement. Similar artefacts affect the velocity actuation (parametrized by α_a and τ_a , respectively), representing a first approximation of the actuator dynamics. The artefact parameters are gathered in $\theta = (\alpha_s, \tau_s, \alpha_a, \tau_a)^\top$. The probabilistic properties of each component of Θ are detailed in the fourth column of Table 1. Fig. 3 shows how artefacts influence the QoI, and how this translates into a loss of fidelity. As a design choice, we require a minimum fidelity $\varphi_{\text{adm}} = 0.7$ (worst case in Fig. 3), with reliability $1 - \varepsilon_{\text{adm}} = 0.999$.

In the present case, some properties of the cyber-physical empirical setup can be established from frequency domain analysis. The transfer function, in the Laplace domain, between τ_2 and the QoI V_2 reads

$$H(s) = \frac{s [P(s) + \alpha e^{-\tau s} (k + cs)]}{P(s)^2 + \alpha e^{-\tau s} (k + cs) m s^2} \quad (18)$$

where $P(s) = m s^2 + c s + k$, $\alpha = \alpha_s \alpha_a$ and $\tau = \tau_s + \tau_a$. Equation (18) indicates that α_s and α_a (resp. τ_s and τ_a) play symmetric roles, and that the *total* scaling factor α and the *total* delay τ are of importance for this system.

The PCK model \mathcal{K} is calibrated, as outlined in Section 3.2, from an initial set \mathcal{E} containing $N = 100$ samples of Θ , see Fig. 5. Note that none of the artefacts' realization in the initial set \mathcal{E} leads to unacceptable fidelity. \mathcal{E} is step-wise enriched with new samples as described in Section 3.3 and illustrated in Fig. 4. The full set of samples generated during the enrichment process is represented by square markers in Fig. 5. They resolve the

fidelity bounds for this setup. P_f is evaluated from (4), and is found to converge after ca 100 iterations (see Fig. 9) towards $1.6 \cdot 10^{-3} > \varepsilon_{\text{adm}}$, which means that robust fidelity is *not* achieved for this setup.

The following remarks are in order when considering Fig. 5. (1) The fact that failure clearly occurs in well-defined regions of the $\alpha - \tau$ plane is consistent with (18): the individual values of τ_s and τ_a (resp. α_s and α_a) are irrelevant for this system, it is their sum (resp. product) that matters. Note, however, that in general such knowledge is not available, and that such structures are not easily identifiable, especially when dealing with a high-dimensional Θ . (2) As expected, loss of fidelity occurs when τ becomes large and when α deviates (i.e. increases or decreases) from unity. In contrast, by inspection of (18), it can be established that the dynamical system becomes *unstable* for $\tau \geq 40$ ms when $\alpha = 1$, and that the stability margin may be increased only by *decreasing* α , when $\tau = 40$ ms. While stability is in general necessary to ensure high fidelity of a cyber-physical empirical setup, it is obviously not a *sufficient* condition. (3) While most of the samples selected by the enrichment algorithm are located near \mathcal{L} , some of them seem inadequately generated in regions with high fidelity. The reason is that they typically feature an unlikely large value of α_s , combined with an unlikely small value of α_a (or vice-versa). Even if $\alpha = \alpha_s \alpha_a$ is close to unity, such a combination falls in an unexplored region of \mathcal{D} , leading to a large value of $\sigma_{\mathcal{K}}$ and thus P_m , in (15). The corresponding sample is therefore selected by the AK algorithm. This is commonly referred to as sample space exploration.

4.2 Case II: active truncation of a slender structure

We will now show how the proposed method allows tackling more complex robust fidelity problems, in higher dimension spaces, and in which non-analytic description of the substructures must be used. We revisit the case study detailed in [Sauder et al., 2018], in which a floating oil production platform is considered, moored in 1200m water depth with a polyester mooring line. The objective is to perform hydrodynamic model testing of this system at a scale $\lambda = 1/60$. Due to space limitations in the hydrodynamic laboratories, the mooring line is truncated 240m below the free surface, as shown in Fig. 6. The upper portion of the line (Σ_2), is represented by a reduced-scale physical model in the hydrodynamic laboratory, while its lower portion (Σ_1) is simulated with the finite element method. The control strategy is as indicated in the introduction of Section 4. The mooring line is subjected to current and wave loads. The latter are represented by a force $(F_x(t), F_z(t))^\top$ acting on the top of the line, see the upper plot in Fig. 7.

The following artefacts are assumed to affect the setup. The force sensors (one for the horizontal component F_x , and one for the vertical component F_z), located at

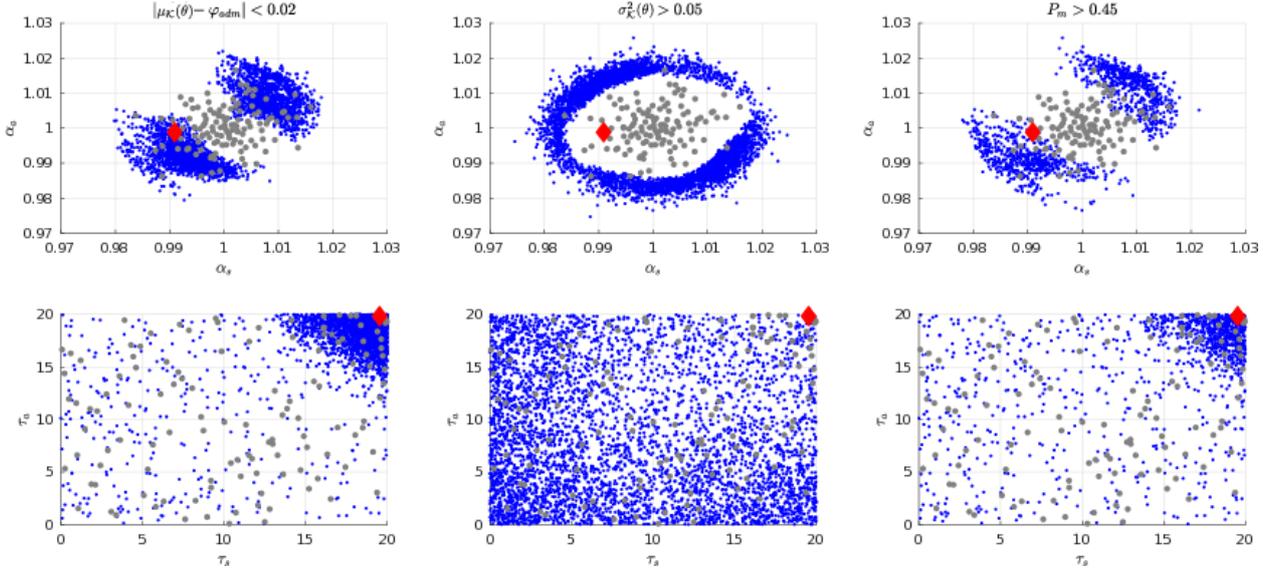


Figure 4. Case I. Illustration of the enrichment procedure. In all plots, the grey points correspond to the initial space-filling set \mathcal{E} . In the first column, the blue points correspond to all auxiliary samples for which $E[\mathcal{K}]$ is close to φ_{adm} (projection in the (α_s, α_a) plane on top, and in the (τ_s, τ_a) plane on the bottom). The second column shows the regions of large uncertainty $\text{Var}[\mathcal{K}]$, and the third column to resulting high values of the learning function P_m . The red diamond marker represents θ^{new} .

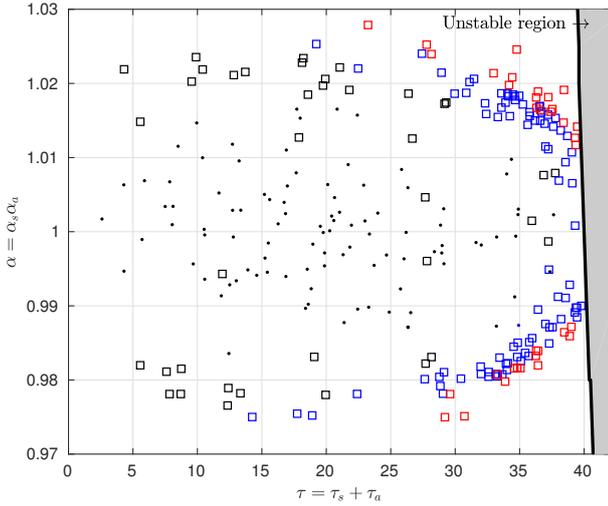


Figure 5. Case I. Samples of $\theta = (\alpha_s, \tau_s, \alpha_a, \tau_a)^\top$, represented in the α - τ plane. Dots represent samples from the initial set \mathcal{E} (100 samples), while squares correspond to the enrichment process (160 samples). Black markers corresponds to $\varphi(\theta) > 0.75$, blue markers to $\varphi(\theta) \in (0.70, 0.75]$ and red markers correspond to $\varphi(\theta) \leq 0.70$.

the truncation point, suffers from calibration error (parameters θ_1 and θ_2 , respectively), bias (θ_3 and θ_4 , respectively), and measurement noise (θ_5 and θ_6 , respectively). Time delays are induced by communication and data processing (parametrized by θ_7 on the sensor side, and θ_{10} on the actuator side). Signal loss and subsequent jump, occur both due to communication errors with the underwater force sensors (θ_8 and θ_9 parametrize the probability and inverse duration parameters, respec-

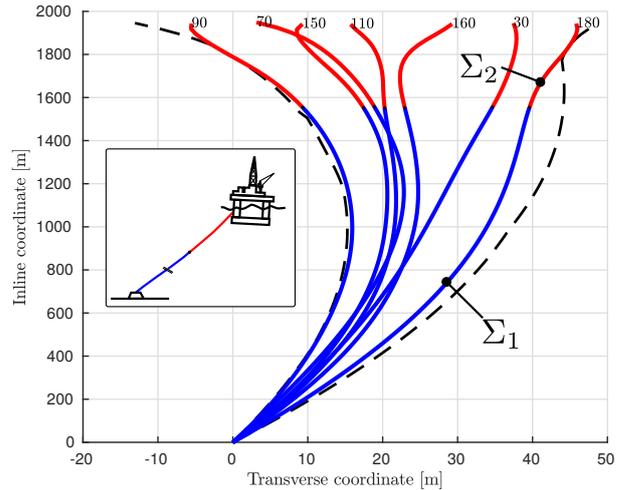


Figure 6. Case II. Frame on the bottom left: general illustration of the problem (not at scale). Main plot: mooring line configuration at given time instants when subjected to the external load $\tau_2(t)$. An inline/transverse coordinate system is used to highlight the deformations of the mooring line. The dashed lines represent the motion envelope of the line.

tively), and due to nonlinear iterations of the numerical substructure that do not complete on time (θ_{11} and θ_{12}). Fig. 1c shows how the elementary artefacts are composed to create Δ_{12} and Δ_{21} . The two middle plots in Fig. 7 show how the artefacts affect the force $y_2 = (f_x, f_z)^\top$ and velocity $y_1 = (v_x, v_z)^\top$ at the truncation point. The probabilistic description of Θ is given in the last column of Table 1. For this Case, $Q = 2$ and the QoI are the velocity components of the top of the line, which we denote V_x and V_z . The fidelity φ is calculated based on the

following γ_q functions.

$$\gamma_1 = \left(\frac{\int_0^T (V_x(t)|_\theta - \bar{V}_x(t))^2 dt}{\int_0^T \bar{V}_x(t)^2 dt} \right)^{\frac{1}{2}} \quad (19)$$

$$\gamma_2 = \left(\frac{\int_0^T (V_z(t)|_\theta - \bar{V}_z(t))^2 dt}{\int_0^T \bar{V}_z(t)^2 dt} \right)^{\frac{1}{2}} \quad (20)$$

We set the minimum admissible fidelity to $\varphi_{\text{adm}} = 0.8$ with a reliability of $1 - \varepsilon_{\text{adm}} = 0.999$.

Fig. 8 illustrates the first step of the enrichment procedure. The interesting regions (associated to large P_m) are detected where the ratio θ_2/θ_1 deviates significantly from unity (distortion of the force angle), θ_8 is small (long periods of signal loss of the force sensor), and, to a some extent, when significant delays θ_8 and θ_9 occur. These findings are consistent with the physics of the problem, discussed in details in [Sauder et al., 2018]. A key difference with the analysis performed in Case I, is that $K = 16$ new samples of θ are added at each step of the enrichment process. Since the new samples are chosen by a clustering algorithm, they tend to be nicely spread over the limit state margin. This leads to an efficient identification of all portions of \mathcal{D}_f , which can be non-connected, and thus to a smoother convergence of P_f . Furthermore, the co-simulations and evaluations of φ for the batch of 16 new samples can be performed at the same time, in parallel. The evolution of the estimated probability of failure P_f during the enrichment process is shown in Fig. 9. From the initial sample set (256 samples), P_f is estimated to $2.5 \times 10^{-2} > \varepsilon_{\text{adm}}$, but after about 60 steps, P_f stabilizes around $3.6 \cdot 10^{-4} < \varepsilon_{\text{adm}}$. Small oscillations are still visible in the 20 last steps, which can be suppressed (this may be necessary if P_f is very close to ε_{adm}) by increasing the cardinality of the *auxiliary* sample sets, used to compute (4). Note that even if $\dim(\mathcal{D}) = 12$, the total number of required steps remained of the order of 10^2 . [Schöbi et al., 2016] also showed that this number was relatively insensitive to the order of magnitude of P_f .

5 Conclusion

We have presented a method for the design and analysis of control systems orchestrating CPEMs. The objective was to require probabilistic robust fidelity of the setup, despite the presence of parametric and uncertain artefacts, inevitably introduced by the components of the control system. In particular we showed how fidelity bounds could efficiently be established by using a PCK surrogate model, gradually enhanced by Adaptive Kriging. Example case II demonstrated that the proposed method could be used in the verification of ultra-deepwater floating systems in hydrodynamic laboratories. This complex problem, of great interest to

the offshore industry, involves a variety of heterogeneous artefacts, and its analysis using existing analytical or sampling-based methods would be practically unfeasible. The presented method is part of a larger framework, in which the sensitivity of fidelity to each component of the control system can be efficiently established [Sauder et al., 2018]. It is expected that the combination of surrogate modelling and active learning techniques, applied here to the analysis of CPEM, can contribute to solve, in an efficient and pragmatic manner, a much wider class of probabilistic robust control problems.

Acknowledgements

This work was supported by the Research Council of Norway through the Centres of Excellence funding scheme, Project number 223254 - AMOS, and through the project 254845/O80 “Real-Time Hybrid Model Testing for Extreme Marine Environments”.

References

- [Alamo et al., 2015] Alamo, T., Tempo, R., Luque, A., and Ramirez, D. R. (2015). Randomized methods for design of uncertain systems: Sample complexity and sequential algorithms. *Automatica*, 52:160–172.
- [Allen and Goupee, 2017] Allen, C. K. and Goupee, A. J. (2017). Assessment of Wind/wave basin capability for emulating active blade pitch and generator control influence on floating wind turbine response. In *Proceedings of the 27th International Offshore and Polar Engineering Conference (ISOPE 2017)*, San Francisco.
- [Bachoc, 2013] Bachoc, F. (2013). Cross Validation and Maximum Likelihood estimations of hyper-parameters of Gaussian processes with model misspecification. *Computational Statistics & Data Analysis*, 66:55–69.
- [Bayati et al., 2017] Bayati, I., Belloli, M., and Facchinetti, A. (2017). Wind Tunnel 2-DoF Hybrid/HIL Tests on the OC5 Floating Offshore Wind Turbine. page V010T09A076. ASME.
- [Blatman and Sudret, 2011] Blatman, G. and Sudret, B. (2011). Adaptive sparse polynomial chaos expansion based on least angle regression. *Journal of Computational Physics*, 230(6):2345–2367.
- [Bursi et al., 2013] Bursi, O. S., Wang, Z., Jia, C., and Wu, B. (2013). Monolithic and partitioned time integration methods for real-time heterogeneous simulations. *Computational Mechanics*, 52(1):99–119.
- [Calafiore et al., 2011] Calafiore, G. C., Dabbene, F., and Tempo, R. (2011). Research on probabilistic methods for control system design. *Automatica*, 47(7):1279–1293.
- [Chen et al., 2005] Chen, X., Aravena, J. L., and Zhou, K. (2005). Risk analysis in robust control-making the case for probabilistic robust control. In *Proceedings of the American Control Conference*, pages 1533–1538.
- [Chernoff, 1952] Chernoff, H. (1952). A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Annals of Mathematical Statistics*, 23(4):493–507.
- [Dubourg, 2011] Dubourg, V. (2011). *Adaptive Surrogate Models for Reliability Analysis and Reliability-Based Design Optimization*. PhD thesis, Université Blaise Pascal-Clermont-Ferrand II.

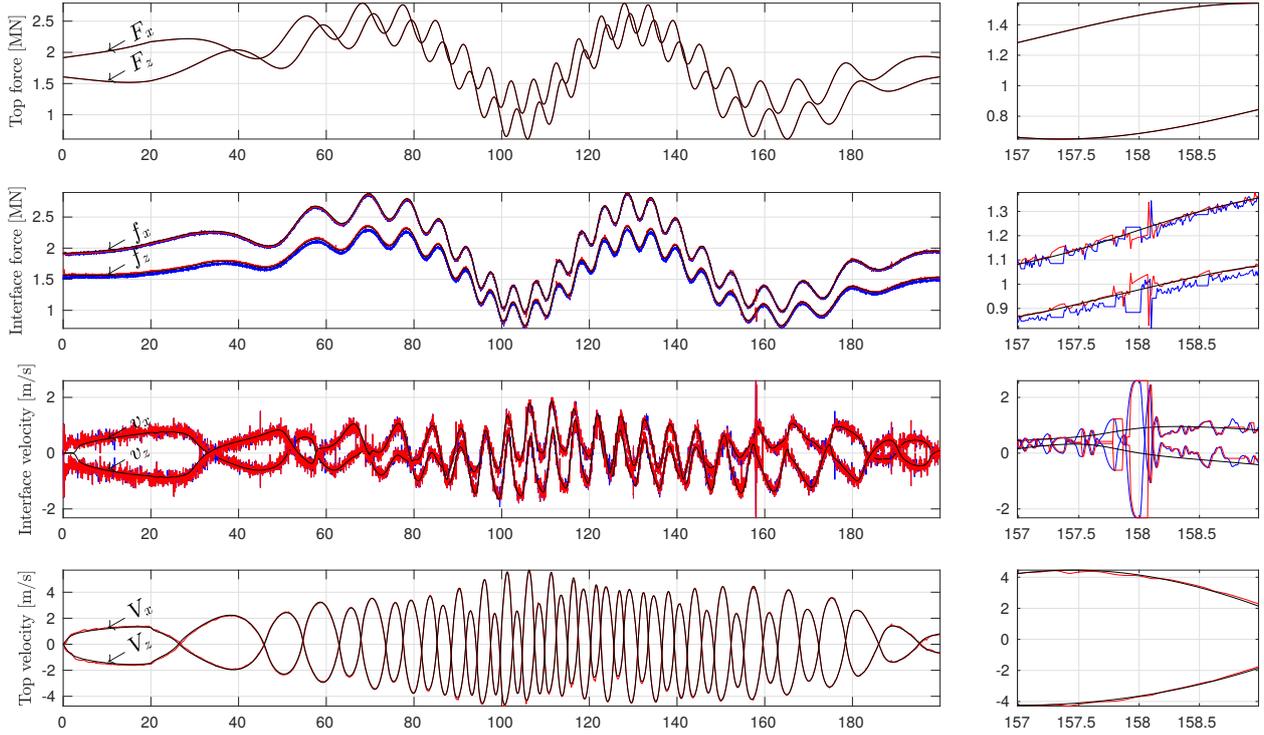


Figure 7. Case II. From top to bottom: times series of the (1) force $\tau_2(t)$ applied to the top of the line, (2) force at the truncation point (3) velocity of the truncation point and (4) velocity of the top of the line, that is the QoI for this problem. The black lines correspond to the system without artefacts ($\varphi = \infty$). The other curves correspond to a random realization of Θ . The blue curves correspond to quantities evaluated on Σ_1 side (y_1 and u_{12} in Fig. 1b), and red curves to quantities evaluated on Σ_2 side (y_2 and u_{21}). Note that despite the significant discrepancies at the truncation point, induced by the artefacts, the QoI are correctly reproduced by this setup, which illustrates the difference between resilience and fidelity (see discussion in Section 2), and motivates the use of selection/comparison functions γ_q in (2).

[Eaton, 2007] Eaton, M. L. (2007). Multivariate Statistics: A Vector Space Approach. *Lecture Notes, Monograph Series, Institute of Mathematical Statistics*, 53.

[Echard et al., 2011] Echard, B., Gayton, N., and Lemaire, M. (2011). AK-MCS: An active learning reliability method combining Kriging and Monte Carlo Simulation. *Structural Safety*, 33(2):145–154.

[Edrington et al., 2015] Edrington, C. S., Steurer, M., Langston, J., El-Mezyani, T., and Schoder, K. (2015). Role of Power Hardware in the Loop in Modeling and Simulation for Experimentation in Power and Energy Systems. *Proceedings of the IEEE*, 103(12):2401–2409.

[Fawzi et al., 2014] Fawzi, H., Tabuada, P., and Diggavi, S. (2014). Secure Estimation and Control for Cyber-Physical Systems Under Adversarial Attacks. *IEEE Transactions on Automatic Control*, 59(6):1454–1467.

[Filipi et al., 2006] Filipi, Z., Fathy, H., Hagen, J., Knafl, A., Ahlawat, R., Liu, J., Jung, D., Assanis, D. N., Peng, H., and Stein, J. (2006). Engine-in-the-loop testing for evaluating hybrid propulsion concepts and transient emissions-HMMWV case study. Technical report, SAE Technical Paper.

[Goebel et al., 2012] Goebel, R., Sanfelice, R. G., and Teel, A. R. (2012). *Hybrid Dynamical Systems: Modeling, Stability, and Robustness*. Princeton University Press, Princeton, N.J.

[Goldberg, 1989] Goldberg, D. E. (1989). *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley.

[Jaynes, 1957] Jaynes, E. T. (1957). Information Theory and

Statistical Mechanics. *The Physical Review*, 106(4):620–630.

[Marelli and Sudret, 2014] Marelli, S. and Sudret, B. (2014). UQLAB : A framework for Uncertainty Quantification in MATLAB. SIAM Conference on Uncertainty Quantification (ICVRAM), USA, 2014.

[McCrum and Williams, 2016] McCrum, D. and Williams, M. (2016). An overview of seismic hybrid testing of engineering structures. *Engineering Structures*, 118:240–261.

[Misselhorn et al., 2006] Misselhorn, W. E., Theron, N. J., and Els, P. S. (2006). Investigation of hardware-in-the-loop for use in suspension development. *Vehicle System Dynamics*, 44(1):65–81.

[Santner et al., 2003] Santner, T. J., Williams, B. J., and Notz, W. I. (2003). *The Design and Analysis of Computer Experiments*. Springer New York, New York, NY. OCLC: 851741845.

[Sauder et al., 2016] Sauder, T., Chabaud, V., Thys, M., Bachynski, E. E., and Sæther, L. O. (2016). Real-time Hybrid Model Testing of a Braceless Semi-submersible Wind turbine. Part I: The Hybrid Approach. In *ASME 2016 35th International Conference on Ocean, Offshore and Arctic Engineering*, No OMAE2016-54435.

[Sauder et al., 2018] Sauder, T., Marelli, S., Larsen, K., and Sørensen, A. J. (2018). Active truncation of slender marine structures: Influence of the control system on fidelity. *Applied Ocean Research*, 74:154–169.

[Schöbi et al., 2016] Schöbi, R., Sudret, B., and Marelli, S. (2016).

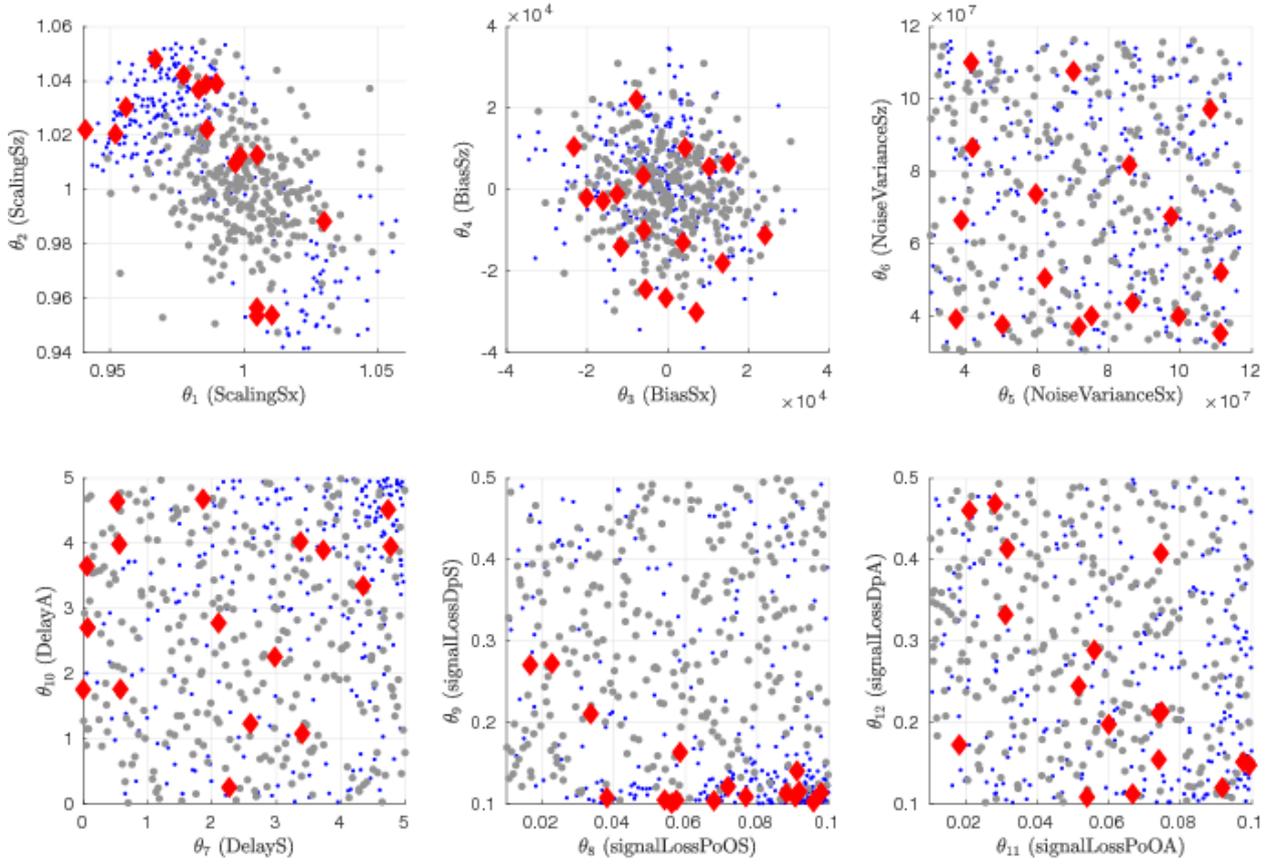


Figure 8. Case II. The grey dots represent the samples in the initial set \mathcal{E} . The blue dots corresponds to areas of the twelve dimensional space with large probability of misclassification ($P_m > 0.45$) at the first step of the enrichment process. The red diamonds represent the $K = 16$ samples selected by the clustering algorithm, at this step.

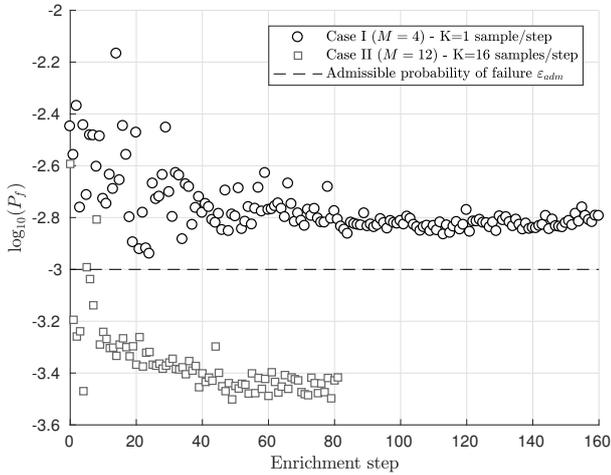


Figure 9. Evolution of the estimated probability of failure during the enrichment process for example cases I and II.

Rare Event Estimation Using Polynomial-Chaos Kriging. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering*, page D4016002.

[Schöbi et al., 2015] Schöbi, R., Sudret, B., and Wiart, J. (2015). Polynomial-chaos-based Kriging. *International Journal for*

Uncertainty Quantification, 5(2).

[Schwarz, 1978] Schwarz, G. (1978). Estimating the Dimension of a Model. *Annals of Statistics*, 6(2):461–464.

[Tabuada et al., 2014] Tabuada, P., Caliskan, S. Y., Rungger, M., and Majumdar, R. (2014). Towards Robustness for Cyber-Physical Systems. *IEEE Transactions on Automatic Control*, 59(12):3151–3163.

[Tempo et al., 2013] Tempo, R., Calafiore, G., and Dabbene, F. (2013). *Randomized Algorithms for Analysis and Control of Uncertain Systems: With Applications*. Communications and control engineering. Springer, London, second edition edition.

[Wallace et al., 2007] Wallace, M. I., Wagg, D. J., Neild, S. A., Bunniss, P., Lieven, N. A. J., and Crewe, A. J. (2007). Testing coupled rotor blade-lag damper vibration using real-time dynamic substructuring. *Journal of Sound and Vibration*, 307(3):737–754.

[Whyte et al., 2015] Whyte, C. A., Mackie, K. R., and Stojadinovic, B. (2015). Hybrid Simulation of Thermomechanical Structural Response. *Journal of Structural Engineering*, 142(2):04015107.

[Zaki and Meira, 2014] Zaki, M. J. and Meira, W. (2014). *Data Mining and Analysis: Fundamental Concepts and Algorithms*. Cambridge University Press, New York, NY.