CEAS-GNC-2022-054

CEAS EuroGNC 2022
"Conference on Guidance, Navigation and Control"
3-5 May 2022 @ Technische Universität Berlin, Germany

A new step towards the integration of probabilistic μ in the aerospace V&V process

Clément Roos Senior research scientist, Information Processing and Systems Department, ONERA The French Aerospace Lab, Toulouse, France. clement.roos@onera.fr

Jean-Marc Biannic Research director, Information Processing and Systems Department, ONERA The French Aerospace Lab, Toulouse, France. jean-marc.biannic@onera.fr

Hélène Evain AOCS engineer, Flight Dynamics Subdirectorate, CNES, Toulouse, France. helene.evain@cnes.fr

ABSTRACT

Probabilistic μ-analysis was introduced 20 years ago as a control system validation means able to quantify the probability of rare and potentially critical events. But for a long time, no practical tool offering both good reliability and reasonable computational time was available, making this technique hardly usable in an industrial context. The STOchastic Worst-case Analysis Toolbox (STOWAT) was introduced a few years ago to bridge this gap between theory and practice. It has been significantly improved since then, thanks to the addition of new features, but above all to increasingly efficient implementations, resulting in a dramatic reduction in CPU time. However, until recently, it could only be applied to small-scale models, with up to 4 or 5 uncertainties. In the perspective of analyzing systems with a larger number of uncertain parameters, a time-consuming and tedious process was carried out in the past months. This led to a complete rewrite of the STOWAT, which is now optimized down to the sub-function level, and whose performance is assessed in this paper on a series on benchmarks of increasing complexity with up to about 20 states and 20 uncertainties. This work represents a new step towards the development of a consolidated tool that could reasonably be integrated in the aerospace V&V process in a near future, finding its place between Monte Carlo simulations – useful for quantifying the probability of sufficiently frequent phenomena – and worst-case μ-analysis – relevant for detecting extremely rare events.

Keywords: Probabilistic μ-analysis; Computational tool; Aerospace V&V process

Nomenclature

AOCS	=	Attitude and Orbit Control System
B&B	=	Branch and Bound
LTI	=	Linear Time Invariant
MC	=	Monte Carlo
SMAC	=	Systems Modeling Analysis and Control
SMART	=	Skew Mu Analysis Robustness Tools
STOWAT	=	STOchastic Worst-case Analysis Toolbox
V&V	=	Verification and Validation

1 Introduction

Due to their simplicity, MC simulations [1, 2] have long been the preferred validation means in the aerospace industry. No analytical representation of the system is needed, and the probability that a whole set of requirements are satisfied is easily computed from time-domain simulations, where the uncertain parameters of the system are sampled based on their probability distributions. But such an approach is generally time-consuming, provide only soft bounds [3] and may fail in detecting rare events. In contrast with MC simulations, optimization-based techniques intelligently search the parameter space to find an optimum, which makes them suitable for worst-case analysis problems where the aim is to identify rare but possibly critical events. In this context, stochastic optimization methods such as Genetic Algorithms or Cross Entropy based optimization [4] have been shown to perform well in a wide variety of complex aerospace problems [5–7], although there are no formal convergence proofs. But their main drawback is the lack of probabilistic bounds, such as the Chernoff bound, to characterize a confidence level.

On the other hand, less expensive deterministic and simulation-free alternatives exist and have reached a good level of maturity, as is the case for μ-analysis [8–10]. But unlike MC simulations, if worst-case scenarios are no longer missed, their probability of occurrence is also not measured. Thus, for many problems the worst-case paradigm based on μ can be overly conservative [11] because of the implicit assumption that each of the real uncertain parameters in the model has uniform probability distribution. This can invalidate a control system on the basis of very rare and therefore extremely unlikely events [12–14]. So there is a real need to fill the gap between MC simulations (able to quantify the probability of sufficiently frequent phenomena) and worst-case μ-analysis (relevant to detect extremely rare events), while providing guaranteed bounds on the calculated probabilities which is not possible with Genetic Algorithms or Cross Entropy based optimization. The ambition is to improve the current industrial standard and to fasten the V&V process, which currently accounts for up to 80% of the AOCS total development time in the space industry.

Researchers started in the 1990s to investigate probabilistic μ-analysis [3, 11, 15], which seeks to combine worst-case bounds determined by μ-analysis with probabilistic information. But although attractive from a theoretical perspective, this approach was far from being applicable at that time, as acknowledged by [11]: *It is still not clear how feasible probabilistic μ is.* The first dedicated software was developed more than twenty years later [12]. It was a major improvement, although still not sufficient to address challenging industrial applications, as highlighted in [13]: *In terms of algorithmic implementation, the situation is similar to when first appeared the preliminary implementations of deterministic μ. That is, in some cases it takes prohibitively long and on average takes too long for standard use in the control design cycle – although possibly still acceptable for limited model complexity during analysis in conjunction with Monte Carlo campaigns.* The last round began a few years ago, when [16] developed a new Matlab software based on the SMART Library of the SMAC Toolbox, that implements state-of-the-art μ-analysis based algorithms to compute robustness margins and performance levels [10]. Shortly afterwards, the first version of the STOchastic Worst-case Analysis Toolbox (STOWAT) was released by [14]. Better integrated with the SMART library, it has been significantly improved since then by [17], thanks to the addition of new features, but above all to increasingly efficient implementations, resulting in a dramatic reduction in CPU time. However, until recently, it could only be reasonably applied to small-scale models, with up to 4 or 5 uncertainties. In the perspective of analyzing systems with a larger number of uncertain parameters, a time-consuming and tedious process was carried out in the past months. This led to a complete rewrite of the STOWAT, which is now highly optimized down to the subfunction level, and whose performance is assessed in this paper on a series on benchmarks of increasing complexity. More generally, this work represents a new step towards the development of a consolidated tool that could reasonably be integrated into the aerospace V&V process in a near future. To the best of our knowledge, this is indeed the first time a realistic system with about 20 states and 20 uncertainties – a flexible satellite with sloshing effects in the present case – can be analyzed by a probabilistic μ approach in only a few minutes.

The paper is organized as follows. Section 2 briefly outlines the μ-analysis framework, both from a deterministic and probabilistic point of view. Section 3 presents the latest software tools and shows on a series of benchmarks of increasing complexity how a prototype version limited to academic applications has become in a few years mature enough to integrate an industrial V&V process. These tools are finally applied in Section 4 to the aforementioned satellite benchmark and compared to MC simulations.

2 μ-analysis framework

Let us consider a continuous-time uncertain LTI system (usually including control laws):

$$\begin{cases} \dot{x} = A(\delta)x + B(\delta)u \\ y = C(\delta)x + D(\delta)u \end{cases} \tag{1}$$

where the parametric uncertainties $\delta = (\delta_1, \ldots, \delta_N) \in \mathbb{R}^N$ are independent random variables with probability density functions $f = (f_1, \ldots, f_N)$. It is assumed that $A(\delta)$, $B(\delta)$, $C(\delta)$, $D(\delta)$ are polynomial or rational functions of the δ_i. As a result, system (1) can be transformed into a Linear Fractional Representation (LFR) as in Fig. 1 (right): the uncertainties are separated from the nominal (closed-loop) system $M(s)$ and isolated in a block-diagonal operator $\Delta = \mathrm{diag}(\delta_1 I_{n_1}, \ldots, \delta_N I_{n_N})$, where I_{n_i} is the $n_i \times n_i$ identity matrix. The set of matrices with the same block-diagonal structure as Δ is denoted $\mathbf{\Delta}$. The subset $k\mathscr{B}_{\mathbf{\Delta}} \subset \mathbf{\Delta}$ is then defined as $k\mathscr{B}_{\mathbf{\Delta}} = \{\Delta \in \mathbf{\Delta} : \overline{\sigma}(\Delta) < k\} = \{\Delta \in \mathbf{\Delta} : |\delta_i| < k, i \in [1,N]\}$, and the notation $\mathscr{B}_{\mathbf{\Delta}}$ is used in the particular case where $k = 1$.

Remark 1 *This paper focuses on real parametric uncertainties, but complex parametric uncertainties and neglected dynamics can be considered as well.*

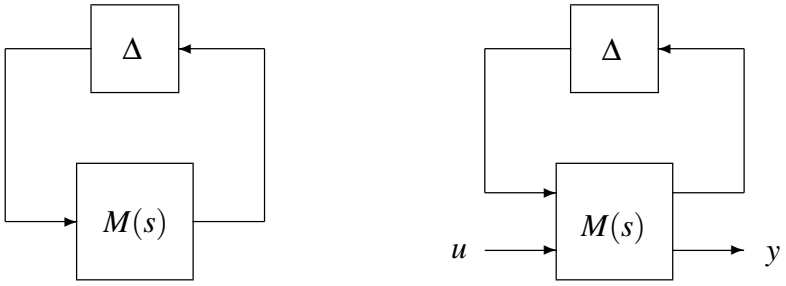

Fig. 1 Standard interconnections for robust stability (left) and performance (right) analysis

With these notations in mind, two main problems can be solved either in a deterministic framework using classical μ-analysis or in a probabilistic framework using probabilistic μ-analysis. Stability is introduced first.

Problem 1.1 (Deterministic worst-case stability) *Compute the largest value k_{wc} such that the interconnection of Fig. 1 (left) is stable for all $\Delta \in k_{wc}\mathscr{B}_{\mathbf{\Delta}}$.*

Problem 1.2 (Probabilistic robust stability) *Given a desired stability margin $k \geq k_{wc}$, compute the probability $\overline{P}^k_{\mathbf{\Delta},f}(M(s))$ that the interconnection of Fig. 1 (left) is unstable when $\Delta \in k\mathscr{B}_{\mathbf{\Delta}}$.*

Remark 2 *The δ_i are often **normalized** in practice, so that $\mathscr{B}_{\mathbf{\Delta}}$ corresponds to the set of physically meaningful uncertainties. In this case, a natural choice is $k = 1$ in Problem 1.2 and the associated probability is simply denoted $\overline{P}_{\mathbf{\Delta},f}(M(s))$. This assumption is supposed to be verified in the sequel.*

Assuming now that $k_{wc} > 1$, H_∞ performance is considered next, where $\mathscr{T}_{u\to y}(s,\Delta)$ denotes the transfer from u to y.

Problem 2.1 (Deterministic worst-case H_∞ performance) *Compute the smallest value $\gamma_{wc} > 0$ such that $\|\mathscr{T}_{u\to y}(s,\Delta)\|_\infty < \gamma_{wc}$ on Fig. 1 (right) for all $\Delta \in \mathscr{B}_{\mathbf{\Delta}}$.*

Problem 2.2 (Probabilistic robust H_∞ performance) *Given a desired performance level $\gamma \in [0\ \gamma_{wc}]$, compute the probability $\overline{P}^{\gamma}_{\mathbf{\Delta},f}(M(s))$ that $\|\mathscr{T}_{u\to y}(s,\Delta)\|_\infty > \gamma$ on Fig. 1 (right) when $\Delta \in \mathscr{B}_{\mathbf{\Delta}}$.*

Once computed, the probabilities $\overline{P}^{k}_{\mathbf{\Delta},f}(M(s))$ and $\overline{P}^{\gamma}_{\mathbf{\Delta},f}(M(s))$ can be confronted to a given tolerance level ε, so as to validate or reject the considered control system, depending on whether they are lower or higher than ε.

Remark 3 *The uncertainties being bounded, their probability distributions must be supported on a bounded interval. Uniform and truncated normal distributions are often used in practice.*

The theory behind μ-analysis is not presented in this paper due to space limitations, but the interested reader can for example refer to [8, 9, 18] and [14, 17] for the classical and the probabilistic versions respectively. Only a few facts are briefly recalled below to facilitate the understanding of Sections 3 and 4. First, computing the structured singular value μ is NP-hard in general, so bounds on k_{wc} and γ_{wc} are usually determined in Problems 1.1 and 2.1 instead of the exact values. Much work has been done in the past decades to reduce the gap between these bounds, and (almost) exact values of k_{wc} and γ_{wc} are now obtained in most cases with a reasonable computational time [19]. Then, probabilistic μ-analysis relies on B&B algorithms to explore the whole uncertainty domain. For stability analysis, this leads to the following partition of the normalized (see Remark 2) uncertainty domain $\mathscr{B}_{\boldsymbol{\delta}} = [-1\ 1]^N$:

$$\mathscr{B}_{\boldsymbol{\delta}} = D_s \cup D_{\bar{s}} \cup D_{s_u} \tag{2}$$

where D_s, $D_{\bar{s}}$ and D_{s_u} are three sets of disjoint N-cubes corresponding to the domains where stability is guaranteed, instability is guaranteed and stability is undetermined respectively, with probabilities $p(D_s)$, $p(D_{\bar{s}})$ and $p(D_{s_u})$. The domain D_{s_u} stems from the aforementioned NP-hardness issue, but also from the fact that B&B can only approximate D_s and $D_{\bar{s}}$, and not compute them exactly. The probability $p(D_{s_u})$ can be reduced by increasing the number of iterations of the algorithm, at the price of an increase in the CPU time. Guaranteed lower and upper bounds on the exact probability $\overline{P}_{\mathbf{\Delta},f}(M(s))$ of instability are finally obtained as follows, thus solving Problem 1.2:

$$p(D_{\bar{s}}) \leq \overline{P}_{\mathbf{\Delta},f}(M(s)) \leq 1 - p(D_s) = p(D_{\bar{s}}) + p(D_{s_u}) \tag{3}$$

Their accuracy depends on the chosen stopping criterion of the B&B algorithm, which allows to handle the trade-off between accuracy and computational time. Performance analysis can be done in the same way, leading to the following partition:

$$D_s = D_\gamma \cup D_{\bar{\gamma}} \cup D_{\gamma_u} \tag{4}$$

where D_γ, $D_{\bar{\gamma}}$ and D_{γ_u} correspond to the domains where performance is guaranteed ($\|\mathscr{T}_{u\to y}(s,\Delta)\|_\infty < \gamma$), non-performance is guaranteed ($\|\mathscr{T}_{u\to y}(s,\Delta)\|_\infty > \gamma$) and performance is undetermined respectively, with probabilities $p(D_\gamma)$, $p(D_{\bar{\gamma}})$ and $p(D_{\gamma_u})$. The following guaranteed bounds on the exact probability $\overline{P}^{\gamma}_{\mathbf{\Delta},f}(M(s))$ of non-performance follow:

$$p(D_{\bar{\gamma}}) \leq \overline{P}^{\gamma}_{\mathbf{\Delta},f}(M(s)) \leq p(D_s) - p(D_\gamma) = p(D_{\bar{\gamma}}) + p(D_{\gamma_u}) \tag{5}$$

The only difference with stability is that the investigated domain is limited to the domain of guaranteed stability D_s, since performance analysis only makes sense for stable systems. The following partition of the normalized uncertainty domain $\mathscr{B}_{\boldsymbol{\delta}}$ is finally obtained by combining (2) and (4):

$$\mathscr{B}_{\boldsymbol{\delta}} = D_\gamma \cup D_{\bar{\gamma}} \cup D_{\gamma_u} \cup D_{\bar{s}} \cup D_{s_u} \tag{6}$$

3 Computational tools assessment

Problems 1.1 and 2.1 can be solved using the SMART Library of the SMAC Toolbox [10, 18], which was introduced in 2013 and implements state-of-the-art μ-analysis based algorithms to compute robustness margins and performance levels. Based on this library, the STOchastic Worst-case Analysis Toolbox (STOWAT) allows to solve Problems 1.2 and 2.2 by computing *guaranteed* lower and upper bounds on the probabilities $\overline{P}_{\mathbf{\Delta},f}(M(s))$ and $\overline{P}^{\gamma}_{\mathbf{\Delta},f}(M(s))$ with the desired accuracy. Initially developed under ESA contract RFP/3-16071/19/NL/CRS/hh since 2019, it will be integrated in a forthcoming version of the SMAC Toolbox and made available on the SMAC website `w3.onera.fr/smac`.

Many improvements have been brought to the STOWAT since the first version presented in [14]. But until recently, it could only be reasonably applied to small-scale models. It generally gave good results in the presence of 4 or 5 uncertainties, but beyond that the computation time increased rapidly if a satisfactory accuracy was required. In the perspective of analyzing systems with a larger number of uncertain parameters, a time-consuming and tedious process was carried out in the past months. This led to a complete rewrite of the Matlab code, which is now optimized down to the sub-function level. This tedious work is the price to pay to be competitive with, or even outperform, simulation-based approaches. The result seems to be well worth the effort, as evidenced by the following comparison between the successive versions of the tool.

Let us consider the following simple example extracted from [16]:

$$\begin{cases} \dot{x} = \begin{bmatrix} 0 & 1 \\ -a_1(\delta_1) & -a_2(\delta_2) \end{bmatrix} x + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u \\ y = \begin{bmatrix} 1 & 0 \end{bmatrix} x \end{cases} \tag{7}$$

where $a_1(\delta_1)$, $a_2(\delta_2)$ are two uncertain parameters defined as:

$$\begin{cases} a_1(\delta_1) = 1 + 2\delta_1 \\ a_2(\delta_2) = 0.8 + \delta_2 \end{cases} \tag{8}$$

and δ_1, δ_2 are two normalized real parametric uncertainties with a uniform distribution on $[-1\ 1]$. Four probabilistic μ codes are compared, all implemented with Matlab:

1) the tool described in [16], which is the first to rely on the SMART Library of the SMAC Toolbox,
2) the prototype version (V0) of the STOWAT introduced in [14], which proposes a totally new implementation as well as additional features such as the determination of $D_{\bar{s}}$ and $D_{\bar{\gamma}}$ in addition to D_s and D_γ,
3) the first consolidated version (V1) of the STOWAT used in [17], directly inspired from the prototype one, but further optimized and robustified,
4) the most recent version (V2) of the STOWAT assessed in the present paper, which consists of a brand new code optimized down to the sub-function level as highlighted above.

Codes #1 and #2 have been kindly provided by D. Alazard and S. Thai respectively, while codes #3 and #4 have been implemented by C. Roos and J-M. Biannic.

The four codes are first applied to solve Problem 1.2 (probabilistic robust stability), and results are presented in Fig. 2 and Table 1, where it can be checked that $p(D_s) + p(D_{\bar{s}}) + p(D_{s_u}) = 100\%$ according to equation (2). A drastic reduction of CPU time is observed. The curve in Fig. 2 is indeed almost linear on a logarithmic scale and reveals that the computational effort is divided by about 10 from one version to the next. In particular, the CPU time has been divided by:

- 20 between codes #3 and #4, which shows the relevance of the complete rewrite,

- 1200 since the introduction of the STOWAT in 2019, which makes the most recent version now applicable in an industrial context (see also Table 3 and Section 4).

In the meantime, accuracy has also been improved. Code #1 is not able to quantify $p(D_{\bar{s}})$ and $p(D_{s_u})$. Code #2 can do it, but $p(D_{s_u})$ remains a bit large for such a simple example. Finally, codes #3 to #4 allow to drastically reduce $p(D_{s_u})$ from 0.47% to 0.01%, which means that 99.99% of the uncertainty domain can now be categorized in less than 1s.

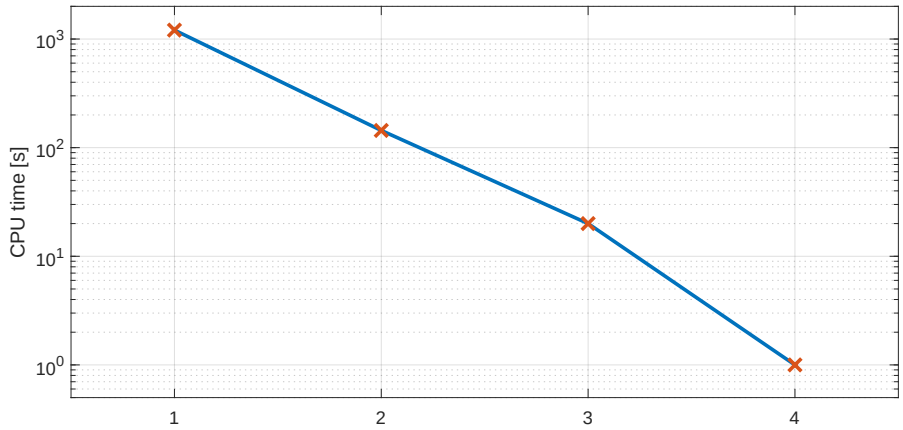

Fig. 2 CPU time for the successive versions of the tool (stability)

Table 1 Numerical results for the successive versions of the tool (stability)

Code #	**1**	**2**	**3**	**4**
CPU time [s]	1208	144	20	<1
$p(D_s)$ **[%]**	92.52	92.59	92.79	92.79
$p(D_{\bar{s}})$ **[%]**	-	6.94	7.20	7.20
$p(D_{s_u})$ **[%]**	-	0.47	0.01	0.01

The four codes are applied next to solve Problem 2.2 (probabilistic robust performance), and results are summarized in Table 2, where it can be checked that $p(D_\gamma) + p(D_{\bar{\gamma}}) + p(D_{\gamma_u}) = p(D_s)$ according to equation (4). The trend is the same as before, and a drastic reduction of CPU time is observed. In the meantime, accuracy is also improved, with $p(D_{\gamma_u})$ decreasing from one version to the next. The decrease in computational time allows to go further in the analysis by choosing a refined stopping condition for the B&B algorithm. This is illustrated in the last column of Table 2, which shows that $p(D_{\gamma_u})$ can be divided by almost 6 with a CPU time of only 18s, and in Fig. 3, where the domains of undetermined stability D_{s_u} or performance D_{γ_u} are hardly visible.

Table 2 Numerical results for the successive versions of the tool (performance)

Code #	**1**	**2**	**3**	**4**	
CPU time [s]	3370	132	27	<2	18
$p(D_\gamma)$ **[%]**	41.30	41.30	41.33	41.57	42.30
$p(D_{\bar{\gamma}})$ **[%]**	-	48.68	49.40	49.39	50.18
$p(D_{\gamma_u})$ **[%]**	-	2.61	2.06	1.82	0.32

Codes #3 (V1) and #4 (V2) of the STOWAT are now applied to a series of study cases of increasing complexity extracted from the SMAC Toolbox [19]. Only probabilistic stability analysis (Problem 1.2)

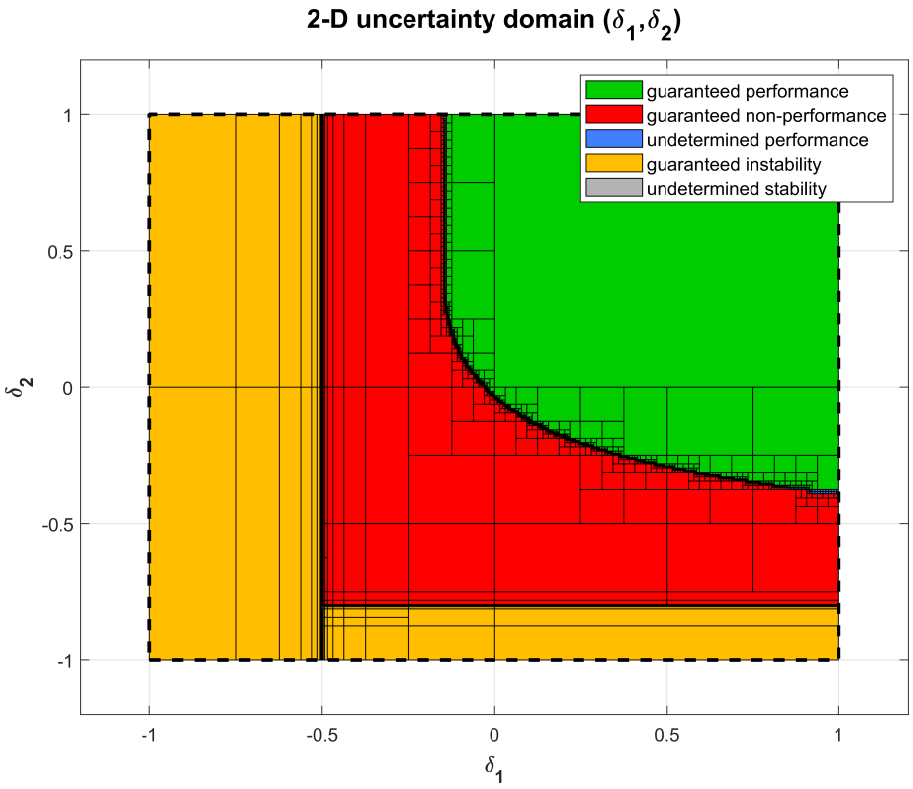

Fig. 3 Partition of the uncertainty domain with a refined stopping criterion

is investigated here, but the same trend can be observed for performance analysis. These test cases are characterized by various fields of application, system dimensions and structures of the uncertainties. Some of them contain poorly damped modes, which usually produce extremely sharp peaks on the μ plot, while others have large state vectors as well as numerous and/or highly repeated uncertainties. All results are gathered in Table 3, whose second column deserves some explanation:

- $3/3$ means that Δ is composed of 3 non-repeated uncertain parameters,
- $4/5$ means that Δ is composed of 3 non-repeated uncertain parameters and 1 repeated twice,
- the drive-by-wire vehicle has 2 non-repeated uncertain parameters and 7 repeated twice,
- the reentry vehicle has 3 non-repeated uncertain parameters and 2 repeated 4 and 6 times,
- the hard-disk drive has 19 non-repeated uncertain parameters and 4 repeated twice.

In each case, the stopping criterion is a requested value of $p(D_{s_u})$ given in column 4, below which the uncertainty domain is no longer divided and the B&B algorithm is interrupted. Note that this threshold is not the smallest that can be reached. It is just selected to have a CPU time of about 10s with code #4. Lower values could probably be reached at the price of an increase of the computational effort, as illustrated for example in Table 2 (see also below).

The results are consistent with those obtained previously on the academic model (7)-(8). Columns 5 and 6 indeed show that the CPU time is divided by about 10 on average thanks to the complete rewrite. It remains low even for high-order systems, and a good accuracy – measured here by the value of $p(D_{s_u})$ in column 4 – is obtained in all cases but one. The case of the bank-to-turn missile should be further investigated and reminds us that improvements are probably still possible. Nevertheless, additional tests reveal that values of 6%, 4% and 2% for $p(D_{s_u})$ can be reached in 11.2s, 24.1s and 112.3s respectively, which remains reasonable. In view of all these results, it appears that the new version of the STOWAT has reached a good level of maturity, which makes it possible to use it on realistic industrial applications, as shown in Section 4.

Table 3 Application of the STOWAT to study cases of increasing complexity

Benchmark	Nb of uncertainties and size of Δ	Nb of states	Requested $p(D_{s_u})$[%]	CPU time [s] Code #3	CPU time [s] Code #4
Inverted pendulum	3/3	4	0.8	32.3	1.9
Anti-aliasing filter	4/5	2	0.05	116.7	9.9
Bank-to-turn missile	4/4	6	8	102.8	7.4
Cassini spacecraft	4/4	17	0.06	35.1	3.2
Drive-by-wire vehicle	9/16	4	0	17.5	3.9
Reentry vehicle	5/13	7	0	19.7	2.5
Space shuttle	9/9	34	0.5	50.7	5.3
Hard-disk drive	23/27	29	1	102.3	19.1

Let us now analyze Table 3 a little further. As expected, good results are obtained even for systems with large state vectors. μ-analysis is indeed a frequency-domain approach, and the size of the frequency response does not depend on the order of the system. It is more surprising that there is no significant increase of the computational time or decrease of the accuracy when the number of uncertainties increases. The most likely explanation is that the B&B algorithm implemented in the STOWAT divides the uncertainty domain along the directions corresponding to the highest μ-sensitivities [20], *i.e.* to the uncertainties with the greatest influence on stability. It also shows that the limiting factor of the proposed approach is not the total number of uncertainties as one might think, but the number of uncertainties that most impact stability or performance. It is therefore preferable to have a system with 20 uncertainties, only 4 of which have a high sensitivity, than a system with 8 uncertainties of equivalent sensitivity. Fortunately, it seems that the first case is more frequent than the second one in practice, even if it must be kept in mind that an unfavorable configuration will inevitably occur at some point, which will put the analysis tools in trouble.

4 Application to a challenging satellite benchmark

4.1 Open-loop satellite modeling

Without a significant loss of generality, since coupling effects are reduced in our context, a single-axis satellite model is considered next. Let us then denote θ the attitude angle of the satellite about the considered axis and $J = 1000\,kg.m^2$ its main moment of inertia. The rigid dynamic equation simply reduces to:

$$J\ddot{\theta} = T_W + T_S + T_F + T_D \tag{9}$$

where T_W is the control input torque (produced by the reaction wheel system), T_S is the torque produced by the propellant slosh effects, T_F is induced by the flexible modes of the solar arrays attached to the main body and finally, T_D captures all remaining disturbance torques (possibly produced by the solar pressure but also by an embedded robotic arm dedicated to on-orbit services). In this work, as is usual in the literature, slosh and flexible effects are represented by poorly damped second-order linear models. The following generic expression is then obtained:

$$T_S + T_F = \sum_i \frac{L_i s^2}{s^2 + 2\xi_i \omega_i s + \omega_i^2} \ddot{\theta} \tag{10}$$

where the parameters L_i, ω_i and ξ_i respectively denote the magnitude, the frequency and the damping of each mode. The nominal values are presented in Table 4 and the corresponding open-loop Bode plot of the transfer from T_W to θ is visualized in Fig. 4.

Table 4 Slosh & flexible modes coefficients (nominal values)

Mode #i	1	2	3	4	5
L_i **[kg.m^2]**	30	40	50	300	100
ω_i **[rad/s]**	0.1	0.2	0.3	0.6	1
ξ_i	5×10^{-3}	4×10^{-3}	3×10^{-3}	1×10^{-3}	1×10^{-3}

Remark 4 *As is further detailed in [21], a more realistic control-oriented representation of the slosh effects is obtained with implicitly time-varying frequency and damping characteristics. The latter generally depend indeed on the angular velocity and acceleration of the satellite. Such variations have been omitted here to generate an uncertain but invariant model as imposed by the μ-analysis framework.*

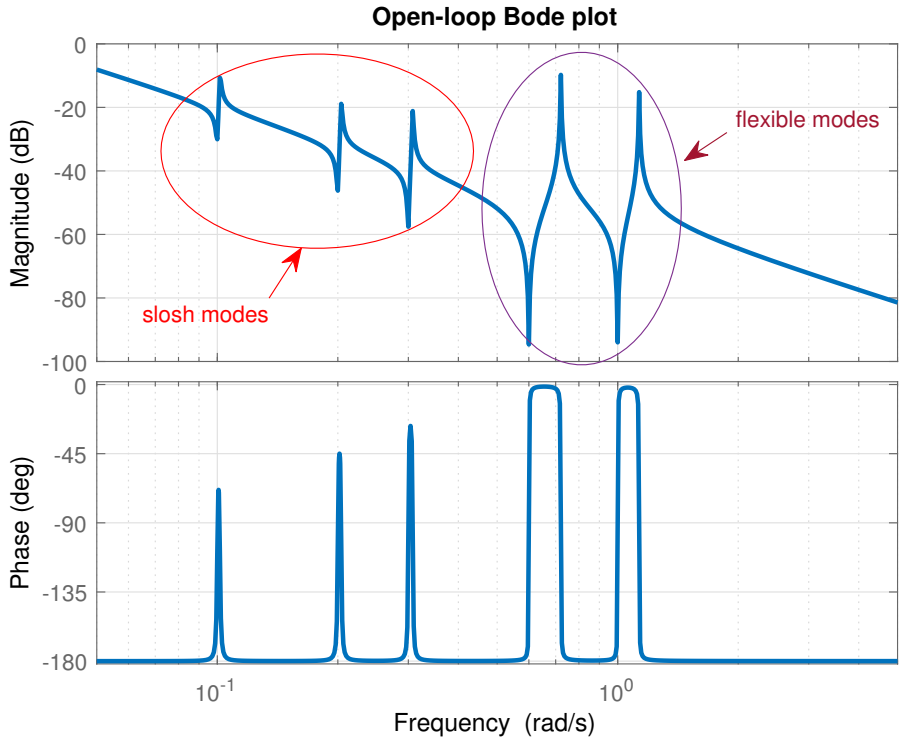

Fig. 4 Bode plot of the nominal open-loop transfer function of the satellite model from the control input torque T_W to the attitude angle θ (with 3 slosh modes and 2 flexible modes)

Finally, as mentioned above, the control input torque T_W is generated by a reaction wheel, which can be approximated by a first-order linear model whose uncertain time constant $\tau \in [0.445\,,\,0.555]$ will be fixed to its nominal value 0.5 in the design process:

$$T_W = \frac{1}{1+\tau s} T_C \tag{11}$$

Note that this actuator is also rate and magnitude limited. These properties are taken into account in the design phase but not introduced here since the considered maneuvers are assumed to be saturation free.

4.2 Structured H_∞-based robust controller design

A robust attitude tracking controller is designed for the above satellite model. As illustrated in Fig. 5, a second-order reference model $R(s)$ with desired frequency $\omega_r = 0.05\,rad/s$ and damping $\xi_r = 0.75$ is introduced. This model delivers reference signals θ_r and $\dot{\theta}_r$ which are compared to the outputs of the plant. As is clearly visible on the diagram, an extended PID structure is imposed on the controller $K(s)$ which minimizes the H_∞ norm of the weighted transfer from the exogenous inputs $w_1 = \theta_c$ and $w_2 = T_D$ to the exogenous outputs $z_1 = \varepsilon = \theta_r - \theta$, its integral z_2 and $z_3 = T_C$. The weighting functions are tuned so that:

- the optimized controller exhibits good robustness properties against actuator delay (delay margin should be above $0.45\,s$),
- the H_∞ norm of the transfer $\mathscr{T}_D(s) = \mathscr{T}_{T_D \to \varepsilon}(s)$ from disturbance torque inputs T_D to the attitude error ε should ideally not exceed 0.01.

The last objective which ensures good tracking performance despite perturbations is hard to achieve simultaneously with the delay margin constraint. After a reasonably short tuning process, a fifth-order controller $K(s)$ was yet obtained with the `systune` routine from the Matlab control toolbox. This controller exhibits a very good delay margin ($0.45\,s$) and acceptably meets the performance requirement with $\|\mathscr{T}_D(s)\|_\infty = 1.01 \times 10^{-2}$.

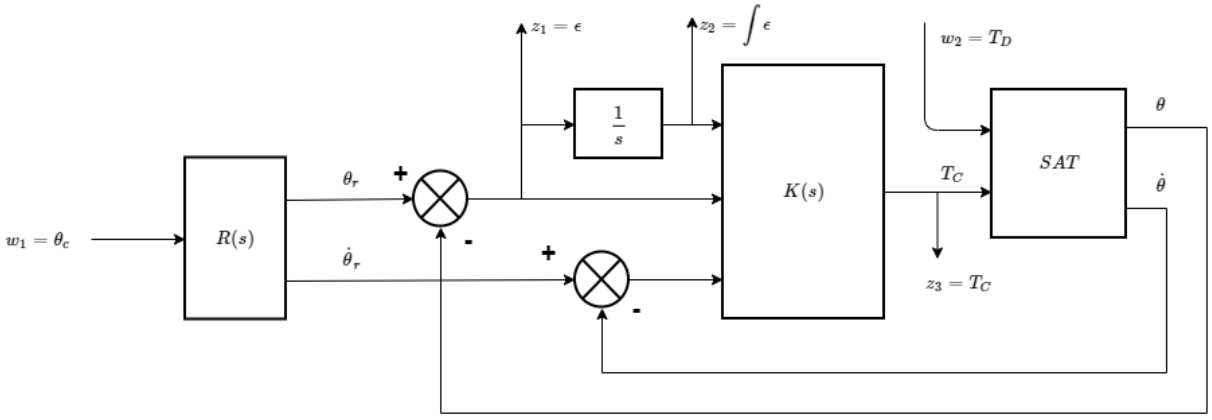

Fig. 5 H_∞ **design model**

Moreover, parametric robustness properties against possible variations of the inertia, the actuator time constant and the 15 parameters of Table 4 are partly ensured via a multi-model design approach. There is however no guarantee that the whole uncertainty domain is cleared. It is indeed numerically impossible here to include all parametric variations in the design process since there are two many parameters. A validation phase to be performed *a posteriori* is then required.

4.3 Control system validation

Let us now evaluate the pointing performance properties of the closed-loop system by confronting the latest version of the STOWAT (code #4 in Section 3) to a classical MC-based approach. Real parametric uncertainties are introduced on the inertia J of the satellite and the time constant τ of the reaction wheel, as well as the magnitudes L_i, damping factors ξ_i and frequencies ω_i of the 5 slosh and flexible modes. An uncertain system with 21 states and 17 uncertainties is therefore obtained.

Problem 1.1 is solved first using the SMART Library, which shows that stability is guaranteed when all uncertainties are allowed to vary independently of each other by $\pm 11.34\%$ around their nominal values. The main objective is now to determine if the pointing performance degradation with respect

to the nominal case is acceptable when the uncertain parameters vary by $\pm 11\%$. For this, the transfer $\mathscr{T}_D(s,\Delta)$ between the external disturbance torque T_D and the attitude error ε is considered. The nominal H_∞ norm without uncertainties is 1.01×10^{-2}, which indicates a good rejection of the perturbation. Robust performance is first analyzed using random sampling. 100000 samples are generated for each uncertain parameter, following a uniform distribution on an interval of $\pm 11\%$ around their nominal value. The highest H_∞ norm among these 100000 samples is 1.11×10^{-2} and is achieved near the frequency $\omega_s = 0.5\,rad/s$. This suggests that performance is only minimally affected by the uncertainties. The Bode plots of the nominal system and the 100 samples with the highest H_∞ norm are shown in red and blue respectively in Fig. 6. There is a fairly large dispersion on the frequency, but much less on the magnitude, which is consistent with the previous results. All computations are performed in $1300\,s$.

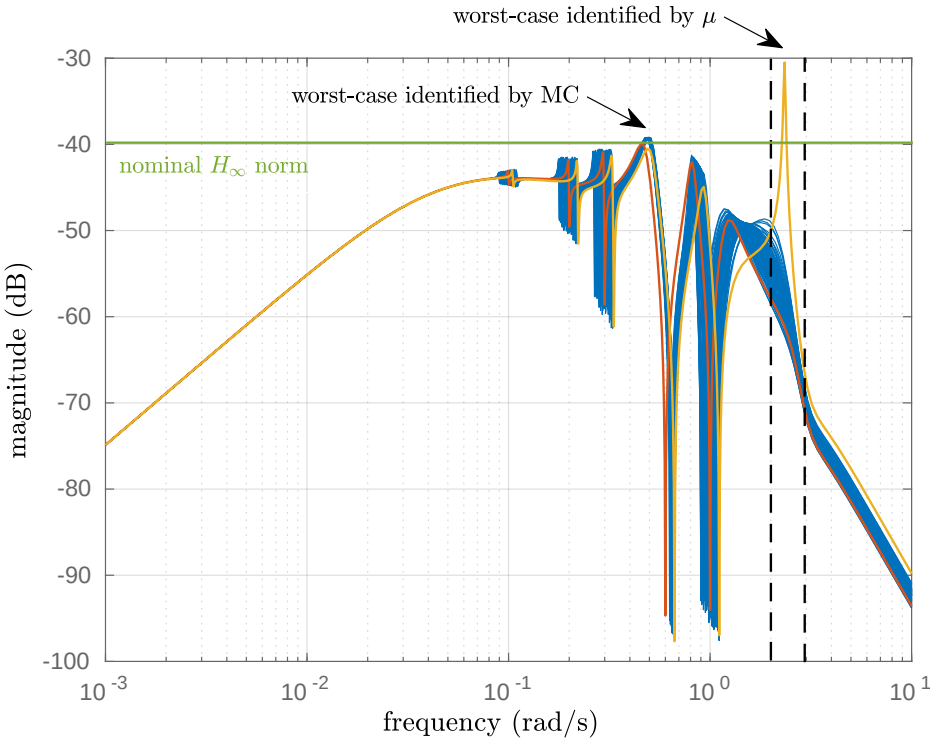

Fig. 6 Bode plot of the transfer $\mathscr{T}_D(s,\Delta)$ between the external disturbance torque and the attitude error (red = nominal, blue = 100 worst samples, yellow = worst-case identified by μ-analysis)

The SMART Library is now used to solve Problem 2.1, *i.e.* to compute a lower bound on the worst-case H_∞ norm of the transfer $\mathscr{T}_D(s,\Delta)$. 100 initial frequencies are considered to cover the whole frequency range, and are then optimized during the analysis process to find the highest possible lower bounds corresponding to worst-case uncertainty combinations. The results are obtained after only $9\,s$ and are partially reproduced in Table 5. It can be seen that in most cases, a lower bound of about 1.1×10^{-2} is obtained close to ω_s, which is in line with the results obtained previously by sampling. But iteration 93 reveals a much higher bound of 2.89×10^{-2} at the frequency $\omega_c = 2.346\,rad/s$, which is missed by the sampling approach despite the fact that 100000 cases are considered. The Bode plot of this worst-case configuration is plotted in yellow in Fig. 6. It is worth noting that ω_c is quite different from ω_s. Moreover, it does not coincide with any frequency of the slosh and flexible modes (0.1, 0.2, 0.3, 0.6 and $1\,rad/s$). It was therefore not possible before applying μ-analysis to anticipate this phenomenon, which can have a significant impact on the pointing performance. Let us indeed consider a sinusoidal disturbance $T_D(t) = \sin(w_c t)$. The resulting pointing error is plotted in Fig. 7 with the same color code as above. For the worst-case configuration identified by the SMART Library, it is about 15 times and 30 times larger than for the worst random sample and the nominal configuration respectively.

Table 5 Worst-case performance analysis results

Grid point	Initial frequency [rad/s]	Final frequency [rad/s]	Lower bound on $\|\mathscr{T}_D(s,\Delta)\|_\infty$
1	0.000	0.475	1.10×10^{-2}
2	0.004	0.483	1.10×10^{-2}
91	1.772	0.489	1.10×10^{-2}
92	2.070	0.484	1.10×10^{-2}
93	2.418	2.346	2.89×10^{-2}
94	2.825	0.476	1.10×10^{-2}
99	6.141	0.491	1.11×10^{-2}
100	7.173	0.499	1.11×10^{-2}

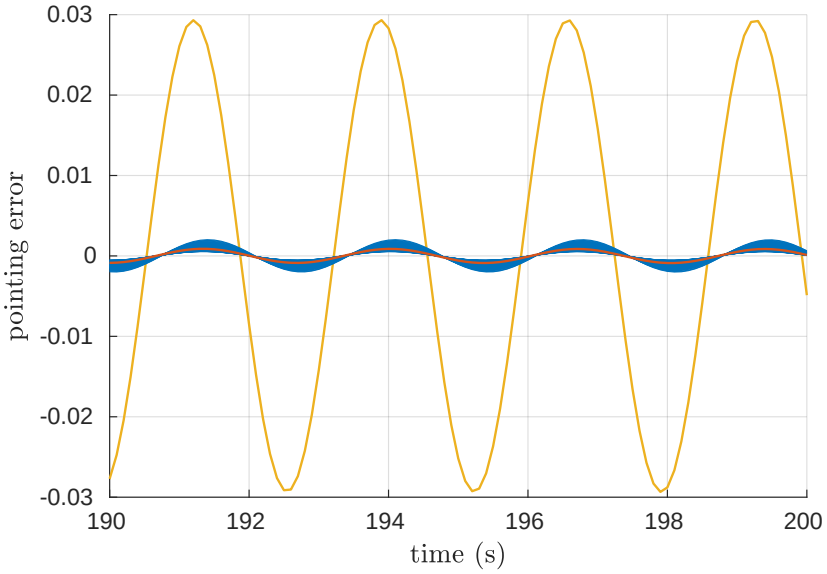

Fig. 7 Pointing error in the presence of a sinusoidal disturbance $T_D(t) = \sin(w_c t)$ (red = nominal, blue = 100 worst samples, yellow = worst-case identified by μ-analysis)

In this context, a critical issue is to compute the probability of observing such a performance degradation, so as to decide whether the proposed control system is reliable enough or not. The latest version of the STOWAT introduced in Section 3 is therefore applied. As with the previous sampling approach, all uncertainties follow a uniform distribution. But the frequency range on which performance is investigated is now limited to the interval $[2\ 3]\ rad/s$ (represented by vertical dashed lines in Fig. 6) to focus on the worst-case peak identified above. Problem 2.2 is solved with a desired performance level $\gamma = 1.01 \times 10^{-2}$, which corresponds to the H_∞ norm of the nominal system. In other words, the objective is to compute guaranteed probabilities that the H_∞ norm is lower or higher between 2 and $3\,rad/s$ than its nominal value on the whole frequency range (represented by a horizontal green line in Fig. 6). The probability of undetermined performance below which the B&B algorithm should be interrupted is fixed to $p(D_{\gamma_u}) = 0.001\%$. After about $100\,s$, the probabilities $p(D_\gamma) = 99.999\%$, $p(D_{\bar{\gamma}}) = 0\%$ and $p(D_{\gamma_u}) = 0.001\%$ of guaranteed performance, guaranteed non-performance and undetermined performance are obtained, which means that the probability that the H_∞ norm is higher than 1.01×10^{-2} between 2 and $3\,rad/s$ is upper bounded by 0.001%. A very large part of the uncertainty domain has therefore been cleared in a very reasonable computational time. The results are even better if we consider a more realistic case where the uniform distribution is replaced with a normal distribution with zero mean and variance σ^2. To illustrate this, all probabilities are quickly recomputed using the STOWAT

for different values of σ^2 between 0.1 and 5, without performing the whole analysis again. The results are plotted in Fig. 8 and show that $p(D_{\gamma_u})$ decreases when σ^2 decreases, the uniform case being recovered when $\sigma \to \infty$. For example, if $\sigma^2 = 0.2$, the probability of non-performance is no larger than $1.8 \times 10^{-6}\,\%$. And results are even better if $\sigma^2 = 0.1$, since the probability of non-performance is bounded by $1.3 \times 10^{-9}\,\%$, which is negligible. The performance degradation shown in Fig. 6 and 7 is therefore very unlikely in these cases. To sum up, if the system is well enough identified to ensure that σ remains quite small, there is certainly no need to invalidate the control system on the basis of this extremely rare worst-case event. But if σ is larger, performance degradation cannot be ignored, and either a better controller should be designed or the system should be better identified. This realistic example clearly shows the added value of the probabilistic μ approach implemented in the STOWAT, which is here faster and more reliable than MC simulations, but also less conservative than deterministic μ-analysis. Moreover, the computational time remains particularly low – a few minutes to build Fig. 8 – considering the complexity of the model, which contains almost 20 uncertainties.

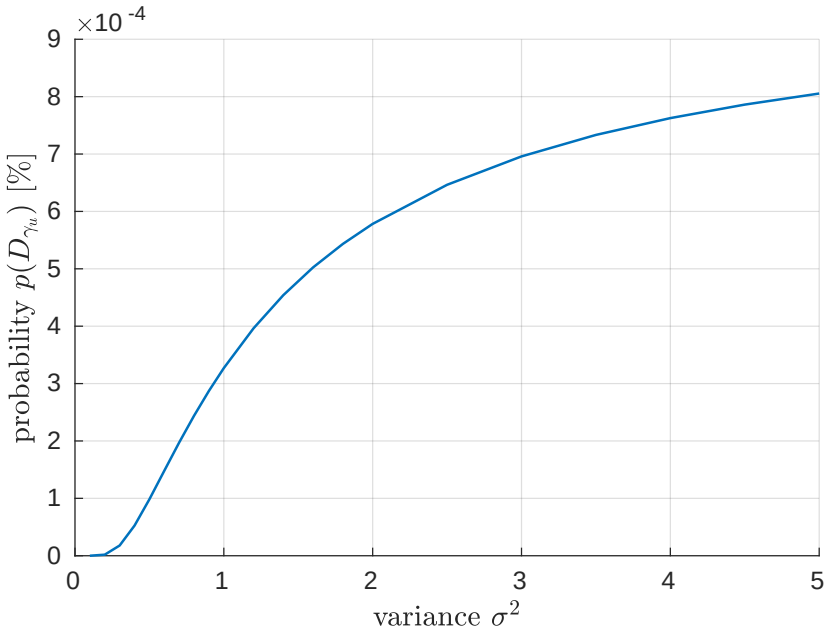

Fig. 8 Probability of undetermined performance as a function of the variance σ^2

5 Conclusion

The combination of state-of-the-art algorithms and a highly optimized implementation reveals the interesting potential of probabilistic μ to evaluate the robustness properties of complex systems. The latest version of the Matlab STOchastic Worst-case Analysis Toolbox indeed allowed to assess the pointing performance of a flexible satellite subject to slosh effects in the presence of about twenty parametric uncertainties, whereas previously existing tools could hardly handle more than 4 or 5. The implementation effort reported in this paper paves the way to the integration of probabilistic μ in the aerospace V&V process in complement to traditional MC-based techniques.

While it is always possible to further optimize the code, the effort tends to become inversely proportional to the benefit, and it is questionable whether it is worth it. More promising directions to make the analysis faster would be to use parallel computing, and to improve the B&B strategies to explore the uncertainty domain more efficiently. Research is currently underway and suggests that by combining the two, it should be possible to further reduce the computational time by a significant amount, leading to a third generation of the STOWAT capable of handling even more complex systems.

References

[1] J. Helton, J. Johnson, C. Sallaberry, and C. Storlie. Survey of sampling-based methods for uncertainty and sensitivity analysis. *Reliability Engineering and System Safety*, 91:1175–1209, 2006.

[2] D. Landau and K. Binder. *A guide to Monte Carlo simulations in statistical physics*. Cambridge University Press, 2005.

[3] X. Zhu, Y. Huang, and J. Doyle. Soft vs. hard bounds in probabilistic robustness analysis. In *Proceedings of the IEEE Conference on Decision and Control*, pages 3412–3417, 1996.

[4] P. de Boer, D. Kroese, S. Mannor, and R. Rubinstein. A tutorial on the Cross-Entropy method. *Annals of Operations Research*, 134:19–67, 2005.

[5] J. Morio and M. Balesdent. *Estimation of rare event probabilities in complex aerospace and other systems: a practical approach*. Woodhead Publishing, 2015.

[6] W. Wang, P. Menon, D. Bates, S. Ciabuschi, N. Gomes Paulino, E. Di Sotto, A. Bidaux, A. Kron, S. Salehi, and S. Bennani. Verification and validation framework for autonomous rendezvous systems in terminal phase. *Journal of Spacecraft and Rockets*, 52(2):625–629, 2015.

[7] A. Mujumdar, P. Menon, C. Roux, and S. Bennani. Cross-entropy based probabilistic analysis of VEGA launcher performance. In J. Bordeneuve-Guibé, A. Drouin, and C. Roos, editors, *Advances in Aerospace Guidance, Navigation and Control*, pages 719–737. Springer International Publishing, 2015.

[8] K. Zhou, J. Doyle, and K. Glover. *Robust and optimal control*. Prentice-Hall, New Jersey, 1996.

[9] G. Ferreres. *A practical approach to robustness analysis with aeronautical applications*. Kluwer Academic, 1999.

[10] C. Roos. Systems Modeling, Analysis and Control (SMAC) toolbox: an insight into the robustness analysis library. In *Proceedings of the IEEE International Symposium on Computer-Aided Control System Design*, pages 176–181, 2013. Available with the SMAC Toolbox at w3.onera.fr/smac/smart.

[11] S. Khatri and P. Parrilo. Guaranteed bounds for probabilistic μ. In *Proceedings of the IEEE Conference on Decision and Control*, pages 3349 – 3354, 1998.

[12] G. Balas, P. Seiler, and A. Packard. Analysis of an UAV flight control system using probabilistic μ. In *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, 2012.

[13] A. Marcos, S. Bennani, and C. Roux. Stochastic μ-analysis for launcher thrust vector control systems. In *Proceedings of the CEAS EuroGNC Conference*, 2015.

[14] S. Thai, C. Roos, and J-M. Biannic. Probabilistic μ-analysis for stability and H_∞ performance verification. In *Proceedings of the American Control Conference*, pages 3099–3104, 2019.

[15] X. Zhu. Improved bounds computation for probabilistic μ. In *Proceedings of the American Control Conference*, pages 4336–4340, 2000.

[16] A. Falcoz, D. Alazard, and C. Pittet. Probabilistic μ-analysis for system performances assessment. *IFAC-PapersOnLine*, 50(1):399 – 404, 2017. Proceedings of the 20th IFAC World Congress.

[17] J-M. Biannic, C. Roos, S. Bennani, F. Boquet, V. Preda, and B. Girouart. Advanced probabilistic μ-analysis techniques for AOCS validation. *European Journal of Control*, 2021. doi.org/10.1016/j.ejcon.2021.06.019.

[18] C. Roos, F. Lescher, J-M. Biannic, C. Döll, and G. Ferreres. A set of μ-analysis based tools to evaluate the robustness properties of high-dimensional uncertain systems. In *Proceedings of the IEEE Multiconference on Systems and Control*, pages 644–649, 2011.

[19] C. Roos and J-M. Biannic. A detailed comparative analysis of all practical algorithms to compute lower bounds on the structured singular value. *Control Engineering Practice*, 44:219–230, 2015.

[20] R. Braatz and M. Morari. μ-sensitivities as an aid for robust identification. In *Proceedings of the American Control Conference*, pages 231–236, 1991.

[21] J-M. Biannic, A. Bourdelle, H. Evain, S. Moreno, and L. Burlion. On robust LPV-based observation of fuel slosh dynamics for attitude control design. *IFAC-PapersOnLine*, 52(28):170–175, 2019. Proceedings of the 3rd IFAC Workshop on Linear Parameter-Varying Systems.