

Delay margin analysis of uncertain linear control systems using probabilistic μ

F. Somers¹, Clement Roos¹, Jean-Marc Biannic¹, F. Sanfedino², Valentin Preda³, S. Bennani³, and H. Evain⁴

¹ONERA Toulouse

²ISAE-SUPAERO

³European Space Research and Technology Centre

⁴Centre spatial de Toulouse

October 30, 2023

Abstract

Monte Carlo simulations have long been a widely used method in the industry for control system validation. They provide an accurate probability measure for sufficiently frequent phenomena, but are often time-consuming and may fail to detect very rare events. Conversely, deterministic techniques such as μ or IQC-based analysis allow fast calculation of worst-case stability margins and performance levels, but in the absence of a probabilistic framework, a control system may be invalidated on the basis of extremely rare events. Probabilistic μ -analysis has therefore been studied since the 1990s to bridge this analysis gap by focusing on rare but nonetheless possible situations that may threaten system integrity. The solution adopted in this paper implements a branch-and-bound algorithm to explore the whole uncertainty domain by dividing it into smaller and smaller subsets. At each step, sufficient conditions involving μ upper bound computations are used to check whether a given requirement – related to the delay margin in the present case – is satisfied or violated on the whole considered subset. Guaranteed bounds on the exact probability of delay margin satisfaction or violation are then obtained, based on the probability distributions of the uncertain parameters. The difficulty here arises from the exponential term classically used to represent a delay, which must be replaced by a rational expression to fit into the Linear Fractional Representation (LFR) framework imposed by μ -analysis. Two different approaches are proposed and compared in this paper. First, an equivalent representation using a rational function of degree 2 with the same gain and phase as the real delay, which results into an LFR with frequency-dependent uncertainty bounds. Then, a Padé approximation, whose order should be chosen carefully to handle the trade-off between conservatism and complexity. A constructive way to derive minimal LFR from Padé approximations of any order is also provided as an additional contribution. The whole method is first assessed on a simple satellite benchmark, and its applicability to realistic problems involving a larger number of states and uncertainties is then demonstrated.

Delay margin analysis of uncertain linear control systems using probabilistic μ

F. Somers¹, C. Roos¹, J-M. Biannic¹, F. Sanfedino², V. Preda³, S. Bennani^{3,*} and H. Evain⁴

Abstract

Monte Carlo simulations have long been a widely used method in the industry for control system validation. They provide an accurate probability measure for sufficiently frequent phenomena, but are often time-consuming and may fail to detect very rare events. Conversely, deterministic techniques such as μ or IQC-based analysis allow fast calculation of worst-case stability margins and performance levels, but in the absence of a probabilistic framework, a control system may be invalidated on the basis of extremely rare events. Probabilistic μ -analysis has therefore been studied since the 1990s to bridge this analysis gap by focusing on rare but nonetheless possible situations that may threaten system integrity. The solution adopted in this paper implements a branch-and-bound algorithm to explore the whole uncertainty domain by dividing it into smaller and smaller subsets. At each step, sufficient conditions involving μ upper bound computations are used to check whether a given requirement – related to the delay margin in the present case – is satisfied or violated on the whole considered subset. Guaranteed bounds on the exact probability of delay margin satisfaction or violation are then obtained, based on the probability distributions of the uncertain parameters. The difficulty here arises from the exponential term classically used to represent a delay, which must be replaced by a rational expression to fit into the Linear Fractional Representation (LFR) framework imposed by μ -analysis. Two different approaches are proposed and compared in this paper. First, an equivalent representation using a rational function of degree 2 with the same gain and phase as the real delay, which results into an LFR with frequency-dependent uncertainty bounds. Then, a Padé approximation, whose order should be chosen carefully to handle the trade-off between conservatism and complexity. A constructive way to derive minimal LFR from Padé approximations of any order is also provided as an additional contribution. The whole method is first assessed on a simple satellite benchmark, and its applicability to realistic problems involving a larger number of states and uncertainties is then demonstrated.

I. INTRODUCTION

Time delays are an integral part of almost every control engineering problem. They introduce linear phase shifts that limit the control bandwidth and affect closed-loop stability. They can be constant or time-varying, although constant ones are considered in most studies, including this one. The objective of this paper is to quantify the probability that a given delay margin requirement is satisfied or violated for a (potentially high-dimensional) uncertain linear system. Moreover, the focus is on detecting rare but nonetheless possible uncertainty combinations that may threaten system integrity, *i.e.* on studying cases where the probability of violation is very low but not zero. As highlighted in the abstract, probabilistic μ -analysis has been identified for many years as a serious alternative to Monte Carlo simulations to analyze such rare events [1], [2], [3]. It combines efficient μ -based algorithms with a branch-and-bound scheme to explore the whole uncertainty domain, also considering probability distributions on the uncertain parameters [4]. Some efficient algorithms already exist to study robust stability, performance, as well as gain, phase and disk margins, see *e.g.* [5], [6]. It therefore seems natural to extend this approach to delay margin analysis.

The most common techniques to analyze the impact of delays on stability or performance in the time domain are based on Lyapunov-Krasovskii functionals and Lyapunov-Razumikhin functions [7]. Driven by the classical Lyapunov theory, they aim at constructing energy-storage function(al)s for time-delay closed-loop systems. The Lyapunov-Krasovskii theorem requires the time derivative of the functionals to be negative definite along the solutions of the system, while the Lyapunov-Razumikhin condition requires negative definiteness only for the solutions that tend to escape the neighborhood of the equilibrium [8]. Research on Lyapunov methods and their applications to time-delay systems are numerous. A lot of literature deals with deterministic Linear Time Invariant (LTI) problems either in the presence or absence of uncertainty, see *e.g.* [9], [10]. The use of Lyapunov function(al)s usually leads to criteria that can be expressed in terms of Linear Matrix Inequalities (LMI) and solved with dedicated solvers. For small-scale problems, LMI-based computations are very suitable and fast. Nevertheless, they become quite time-consuming when realistic high-dimensional systems with multiple uncertainties are considered. There are also contributions dedicated to the analysis of nonlinear, time-varying or stochastic systems, as in [11], [12]. But their application to high-order systems also remains tricky.

Most frequency domain approaches are based on the small-gain theorem and generally rely on μ -analysis or Integral Quadratic Constraints (IQC). μ -analysis makes use of the structured singular value to study robust stability and H_∞ performance of linear systems, while accounting for the structure of the uncertainties [13]. Guaranteed and very accurate

This work was carried out under ESA contract 4000134724/21/NL/GLC/my and CNES contract R&T R-S22/BS-0005-073

¹ONERA, The French Aerospace Lab, Toulouse, France

²ISAE-SUPAERO, Toulouse, France

³ESA-ESTEC, Noordwijk, The Netherlands

⁴CNES, Toulouse, France

*Currently not affiliated with ESA

bounds on stability margins and performance levels can usually be obtained quickly, even for real-world problems [14]. Examples of how μ -analysis can be used to study time-delay systems are proposed in [15], [16], [17]. IQC analysis can be interpreted as the combination of (scaled) small gain techniques and positivity/passivity techniques which study the interconnection of a linear operator with a non-linearity (Lur'e problem) [18]. It can simultaneously deal with uncertainties, time-varying parameters and non-linearities, and therefore allows to analyze a broader range of systems. However, it can be conservative and requires the use of an LMI solver (unlike μ -analysis for which efficient LMI-free techniques exist), leading to significant computational times for high-order systems. Applications of IQC to the robustness analysis of time-delay systems can be found in [18], [19], [20].

Given the objective of this study and this brief literature review, frequency domain analysis techniques are naturally preferred to time domain ones in the sequel. The commonly accepted representation of a constant delay in the frequency domain is the exponential term $e^{-s\tau}$, where s is the Laplace variable and τ is the considered time delay in seconds. Nevertheless, most approaches require it to be replaced by a polynomial or rational expression to fit into the Linear Fractional Representation (LFR) framework. This step has a major influence on the conservatism of the analysis. In some studies, time delays are eliminated by covering their value sets with unit disks, which leads to overly conservative results [21]. Various other approaches replace the exponential term by Padé approximations [22], Taylor series expansions [23] or the Laguerre formula [24]. Introducing more complexity by increasing the order of the approximation improves accuracy, at the price of a higher computational cost. Alternatively, an equivalent replacement function is presented in [16], which has the same properties (unit gain and phase varying linearly with frequency) as the actual time delay. However, a parameter with frequency-dependent bounds is introduced, which makes the use of standard robustness analysis techniques more difficult, as explained in Section IV-B. In this context, two different approaches – introduced in Section IV-A – are investigated and compared in this paper. First, a new equivalent representation using a rational function of degree 2, simpler than that of [16], is introduced and transformed into a minimal LFR. Second, a classical Padé approximation is considered, and a new constructive way to derive a minimal LFR whatever the order is provided.

The main contribution of this paper is then to propose a branch-and-bound algorithm, which explores the whole uncertainty domain and uses the aforementioned characterizations to compute tight bounds on the probability of delay margin violation for linear systems in the presence of parametric uncertainties with given probability distributions. In particular, sufficient conditions involving μ upper bound computations are introduced in Sections IV-B and IV-C to check whether the delay margin requirement is satisfied or violated on an entire subset of the uncertainty domain. The paper is organized as follows. The considered problem is first formally stated in Section II. Section III gives a brief overview to the proposed solutions and highlights the contributions. The main theoretical results are then detailed in Section IV, as well as all practical algorithms. The performance of the method in terms of accuracy and computational time is finally evaluated in Section V, by applying it to a set of benchmarks of increasing complexity.

II. PROBLEM STATEMENT

Let us consider the following continuous-time uncertain LTI system:

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

The real uncertain parameters $\delta = (\delta_1, \dots, \delta_N)$ are bounded and without loss of generality normalized, so that the whole set of admissible uncertainties is covered when $\delta \in \mathcal{B}_\delta = [-1 \ 1]^N$. They are independent random variables, whose probability density functions f are supported on the bounded interval $[-1 \ 1]$. It is assumed that $A(\delta)$, $B(\delta)$, $C(\delta)$, $D(\delta)$ are polynomial or rational functions of the δ_i with real coefficients, and that system (1) can be transformed into a Linear Fractional Representation $y = \mathcal{F}_u(M(s), \Delta)u$, as defined in [13] and shown in Fig. 1. In simple words, the uncertainties are separated from the nominal LTI system $M(s)$ and isolated in a block-diagonal operator $\Delta = \text{diag}(\delta_1 I_{n_1}, \dots, \delta_N I_{n_N}) \in \mathbb{R}^{p \times p}$, 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 $\mathbf{\Delta}$. Let $\mathcal{B}_\Delta = \{\Delta \in \mathbf{\Delta} : \delta \in \mathcal{B}_\delta\}$ and $\mathcal{D}_\Delta = \{\Delta \in \mathbf{\Delta} : \delta \in \mathcal{D}\}$ be the subsets of $\mathbf{\Delta}$ corresponding to \mathcal{B}_δ and to a given box $\mathcal{D} \subset \mathcal{B}_\delta$ respectively.

In this work, system (1) and Fig. 1 describe a control loop opened at the place where the delay margin should be computed. The closed-loop interconnection is therefore recovered by applying a unit negative feedback between y and u , *i.e.* by setting $G(\tau s) = 1$ in Fig. 2. Using these notations, probabilistic delay margin analysis can be formalized as follows:

Problem 2.1: Compute the probability $\overline{P}_{\Delta, f}^\phi(M(s))$ that the delay margin is smaller than a given threshold ϕ when $\Delta \in \mathcal{B}_\Delta$ for the negative feedback loop obtained by connecting y to u in Fig. 1.

In other words, $\overline{P}_{\Delta, f}^\phi(M(s))$ is the probability that the interconnection of Fig. 2 becomes unstable for some $G(\tau s) = e^{-\tau s}$, $\tau \in [0 \ \phi]$. A control system can then be rejected or validated depending on whether $\overline{P}_{\Delta, f}^\phi(M(s))$ does or does not exceed a given threshold ϵ . Two approaches are presented in Section IV to compute tight bounds on $\overline{P}_{\Delta, f}^\phi(M(s))$ for Single-Input Single-Output (SISO) systems, *i.e.* when $u \in \mathbb{R}$ and $y \in \mathbb{R}$. The global underlying idea is first summarized in Section III.

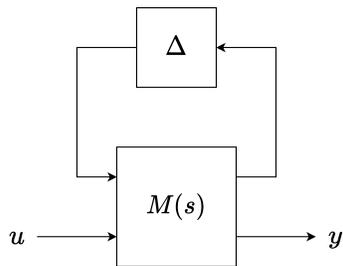


Fig. 1: Linear Fractional Representation $y = \mathcal{F}_u(M(s), \Delta)u$

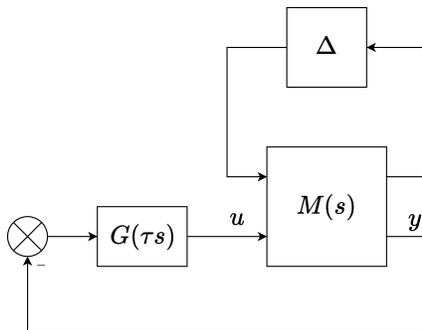


Fig. 2: Feedback loop for delay margin analysis

Remark 2.1: Most of the theoretical considerations in this paper rely on the structured version of the small gain theorem, see Theorem 11.8 of [13]. So strictly speaking, Δ should belong to an open set, *i.e.* \mathcal{B}_δ should be equal to $] -1 \ 1[^N$. The deliberate choice was made to define it as a closed set for two reasons. Firstly, this is relevant from an engineering point of view. And secondly, the main objective here is to compute probabilities, and the probability that δ belongs to the boundary of \mathcal{B}_δ (which is a space of dimension $N - 1$) is zero, so considering open or closed sets does not affect the results.

III. OVERVIEW OF THE PROPOSED SOLUTION

In recent years, a significant effort has been put in the development of probabilistic μ theory and its implementation in the STOchastic Worst Case Analysis Toolbox (STOWAT) [5]. Probabilistic μ -analysis is not intended to replace either its deterministic counterpart or Monte Carlo simulations. It actually finds its place between the two as shown in Fig. 14 of [4]. If an excessive degradation in stability or performance is detected by deterministic worst-case μ -analysis, probabilistic μ can provide more insight into the probability of failure. It allows for very fast iterations between probability levels and uncertainty characterization, and makes uncertainty resizing possible at medium cost. Moreover, it provides useful information to the Monte Carlo analysis carried out during the final stage of the Verification & Validation (V&V) process. Probabilistic stability and H_∞ performance were studied first [25], [4], followed by gain/phase/disk margins, in the SISO case initially [6], and then for Multiple-Input Multiple-Output (MIMO) systems [26]. Lately, a first delay margin analysis algorithm was introduced in [27]. The present work follows on from [27] and aims to provide consolidated results on probabilistic delay margin analysis.

As with the other probabilistic stability margins, the proposed approach combines μ -based tools with a branch-and-bound algorithm to explore the whole uncertainty domain \mathcal{B}_δ . The delay margin is first computed at the center of \mathcal{B}_δ , *i.e.* for $\Delta = 0$. If it is larger (resp. smaller) than the desired threshold ϕ , a satisfaction (resp. violation) test is then performed to check whether the delay margin requirement is satisfied (resp. violated) on the entire domain \mathcal{B}_δ using sufficient conditions involving μ upper bound computations (see Sections IV-B and IV-C). If this cannot be guaranteed, \mathcal{B}_δ is finally partitioned into smaller boxes and this process is repeated until each box has guaranteed sufficient/insufficient delay margin, or is small enough to be neglected (see Section IV-D). Guaranteed upper and lower bounds on the exact probability of delay margin violation $\overline{P}_{\Delta, f}^\phi(M(s))$ are finally obtained, based on the probability distributions f of the uncertain parameters δ , thus solving Problem 2.1.

As highlighted in Section I, a preliminary step to using μ -analysis is to transform the interconnection of Fig. 2 into an LFR, which requires to replace the exponential delay $G(\tau s) = e^{-\tau s}$ with a rational function. In [27], a simple function of degree 2 was proposed with exactly the same characteristics as the delay, *i.e.* unitary gain and linear phase. A major difficulty with respect to the gain/phase/disk margins is that this equivalent representation introduces a parameter with frequency-dependent bounds that adds to the uncertainties already present in Δ . An algorithm was proposed in [27] to take this dependence

on frequency into account, but only for the satisfaction test. In contrast, a second-order Padé approximation was used for the violation test, which allows to implement the same approach as for the gain/phase/disk margins, but introduces some conservatism as highlighted in [27]. In this context, the main contributions of this paper are twofold:

- 1) Higher-order Padé approximations are considered in Section IV-C for the violation test. It is shown in Section V-A that very accurate results can be obtained with almost no increase in computational time. In addition, a method to transform Padé approximations of any order into minimal LFR is introduced in Section IV-A-2, which is a new result.
- 2) An alternative method based on the aforementioned equivalent characterization of the delay is proposed for the violation test in Section IV-C. A thorough comparison of both approaches is performed in terms of conservatism and computational efficiency in Sections IV-C and V.

Moreover, all theoretical results are carefully demonstrated, the algorithms are thoroughly documented, and the applicability of the proposed approach to realistic systems with a larger number of states and uncertainties is studied, which was not carried out in [27].

IV. PROBABILISTIC DELAY MARGIN

A. Interconnection for delay margin analysis

The negative feedback loop of Fig. 2 is first built from the LFR of Fig. 1. $G(\tau s)$ represents a time delay $\tau \in [0 \phi]$:

$$G(\tau s) = e^{-\tau s} \quad (2)$$

and the nominal closed-loop uncertain system is obtained for $\tau = 0$, *i.e.* $G(\tau s) = 1$. $G(\tau s)$ does not have a rational dependence on τ . Therefore, it should be replaced with a rational function to be able to build an LFR and apply robustness tools such as μ -analysis. Two solutions are proposed in this paper. They are introduced below and their use is discussed further in Sections IV-B and IV-C.

1) *Equivalent representation:* The first approach was introduced in [27], and is inspired by [16] while being significantly simpler. It relies on the fact that $G(\tau s)$ has a unitary gain and introduces a frequency-dependent phase shift equal to $-\tau\omega$ at the frequency ω . The delay margin is thus guaranteed to be larger than ϕ if the system remains stable for any phase shift in the interval $\mathcal{I}(\omega) = [-\min(\phi\omega, 2\pi) \ 0]$ at the frequency ω . $G(\tau s)$ can therefore be equivalently replaced with any rational expression $\Phi(\alpha)$ such that $|\Phi(\alpha)| = 1$, and whose phase $\angle\Phi(\alpha)$ covers the interval $\mathcal{I}(\omega)$ when α covers $[0 \ \alpha_{max}(\omega)]$. The following choice is made:

$$\Phi(\alpha) = \frac{2j\alpha^2 - 2(1+j)\alpha + 1}{-2j\alpha^2 - 2(1-j)\alpha + 1} \quad (3)$$

for which:

$$\alpha_{max}(\omega) = \begin{cases} \frac{1 - t(\omega) - \sqrt{1 + t(\omega)^2}}{2} & \text{if } \omega < \pi/\phi \\ 0.5 & \text{if } \omega = \pi/\phi \\ \frac{1 - t(\omega) + \sqrt{1 + t(\omega)^2}}{2} & \text{if } \omega > \pi/\phi \end{cases} \quad (4)$$

and $t(\omega) = \tan \frac{-\min(\phi\omega, 2\pi)}{2}$. To get equation (4), just note that:

$$\angle\Phi(\alpha) = 2 \arctan \frac{2\alpha(\alpha - 1)}{1 - 2\alpha} \quad (5)$$

which gives for $\alpha = \alpha_{max}(\omega)$:

$$\tan \frac{\angle\Phi(\alpha_{max}(\omega))}{2} = t(\omega) = \frac{2\alpha_{max}(\omega)(\alpha_{max}(\omega) - 1)}{1 - 2\alpha_{max}(\omega)} \quad (6)$$

and finally $2\alpha_{max}(\omega)^2 + 2(t(\omega) - 1)\alpha_{max}(\omega) - t(\omega) = 0$. The motivation for this choice of $\Phi(\alpha)$ is that $\angle\Phi(\alpha)$ and $\alpha_{max}(\omega)$ are simple functions with a smooth behavior, as shown in Fig. 3.

The function $\Phi(\alpha)$ being rational in α , it can be written as an LFR $\mathcal{F}_u(T, \Delta_m)$, where T is a static matrix defined as:

$$T = \begin{bmatrix} 2 & -j & | & 1 \\ 2 & -2j & | & 2 \\ \hline 0 & -2j & | & 1 \end{bmatrix} \quad (7)$$

and $\Delta_m = \alpha I_2$, $\alpha \in [0 \ \alpha_{max}(\omega)]$, contains a single parameter with frequency dependent bounds. Standard matrix manipulations based on the Redheffer star product [13] finally allow to equivalently transform the interconnection of Fig. 2, where $G(\tau s)$ is replaced with $\Phi(\alpha)$, into that of Fig. 4. Note that a rational function $\Phi(\alpha)$ of degree 1 in α would lead to a simpler LFR, where α is not repeated in Δ_m . This is however not possible here: α would indeed tend to $\pm\infty$ to allow $\angle\Phi(\alpha)$ to cover $[-2\pi \ 0]$, but the μ -based tools used in the sequel require the Δ block of the LFR to be bounded.

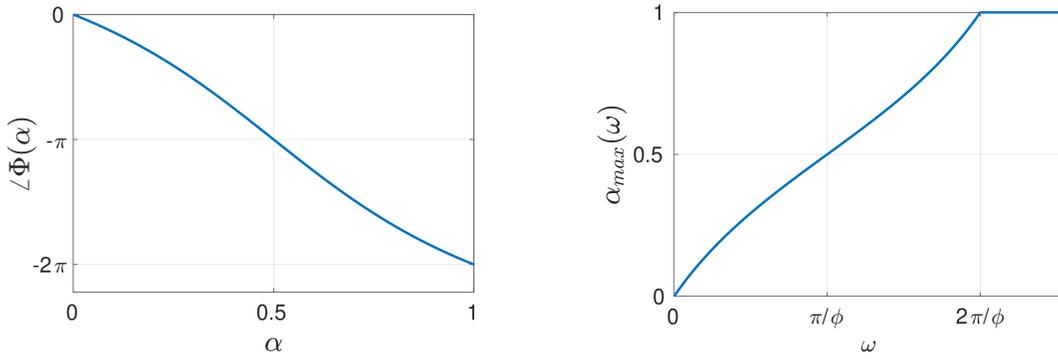


Fig. 3: Functions $\angle\Phi(\alpha)$ and $\alpha_{max}(\omega)$ involved in the equivalent representation of the delay function

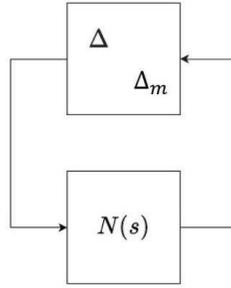


Fig. 4: Transformation of the interconnection of Fig. 2

2) *Padé approximation*: A more common approach to deal with the non-rational dependence on τ of the exact time-delay (2) is to replace it with a Padé approximation $\Psi_n(\tau s)$ [23]. The latter is characterized by a unitary gain as for the delay. However, its phase is not a linear function of the frequency, but becomes closer to $-\tau\omega$ when the order n of the approximation increases, as shown in Fig. 5 (left). The choice of n therefore allows to handle the trade-off between accuracy and complexity. On the one hand, it has been observed in [27] that a simple second-order approximation sometimes leads to conservative results when used in a probabilistic μ framework, which was to be expected given the discrepancy between the phase plot of the delay and that of the second-order approximation. On the other hand, it will be shown in Section V-A that higher-order approximations usually solve this problem, with an almost negligible increase in computational time. A fifth-order approximation typically appears to be a good candidate, since its phase plot almost coincides with that of the delay over the frequency interval $[0\ 2\pi]$.

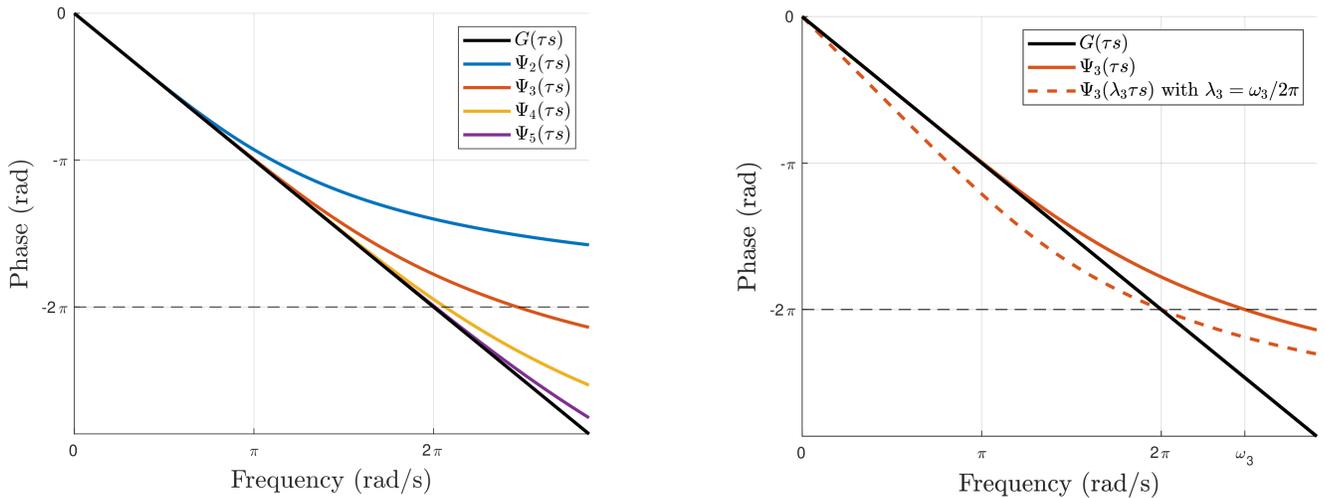


Fig. 5: Phase plots of the delay function and its (scaled) Padé approximations for $\tau = 1s$

An important property of Padé approximations in the context of this paper is that their phase is always greater than that of the delay, regardless of the order. This is illustrated in Fig. 5 (left) for $n = 2 \dots 5$. It is also possible to define scaled

versions $\Psi_n(\lambda_n \tau s)$ for $n \geq 3$, whose phase is, this time, always lower than that of the delay over the interval $[0 \ 2\pi]$. It is shown in [28] that $\lambda_n = \omega_n/2\pi$, where ω_n is the phase crossover frequency of $\Psi_n(\tau s)$ at the -2π line. An example is shown in Fig. 5 (right) for $n = 3$. This is formalized in Lemma 4.1, which primarily builds upon Lemma 5 of [28].

Lemma 4.1: The phase functions $\angle\Psi_n(\tau\omega)$ and $\angle\Psi_n(\lambda_n\tau\omega)$ of the Padé approximation and its scaled version satisfy:

$$-\tau\omega \leq \angle\Psi_n(\tau\omega) \leq 0 \quad \forall\omega \in [0 \ \omega_n] \quad \text{and} \quad \forall n \geq 0 \quad (8)$$

$$-2\pi \leq \angle\Psi_n(\lambda_n\tau\omega) \leq -\tau\omega \quad \forall\omega \in [0 \ 2\pi] \quad \text{and} \quad \forall n \geq 3 \quad (9)$$

Proof: Lemma 5 of [28] proves the left-hand side of (8) for $n \geq 1$ and the right-hand side of (9) for $n \geq 3$. $\Psi_0(\tau\omega) = 1$, i.e. $\angle\Psi_0(\tau\omega) = 0$, which proves (8) for $n = 0$. For $n \geq 1$, it can be shown using Lemma 4 of [28] that:

$$\frac{d\angle\Psi_n(\lambda_n\tau\omega)}{d\omega} = \frac{-\lambda_n\tau V_{n-1}((\lambda_n\tau\omega)^2)}{V_{n-1}((\lambda_n\tau\omega)^2) + (\lambda_n\tau\omega)^{2n}} \quad (10)$$

where V_{n-1} is a polynomial of degree $n-1$ with strictly positive coefficients. The fact that $\angle\Psi_n(0) = 0$ and $\frac{d\angle\Psi_n(\tau\omega)}{d\omega} \leq 0$ (obtained by setting λ_n to 1 in (10)) shows that $\angle\Psi_n(\tau\omega) \leq 0$, i.e. proves the right-hand side of (8). Finally, the fact that $\angle\Psi_n(\lambda_n\tau 2\pi) = \angle\Psi_n(\tau\omega_n) = -2\pi$ and $\frac{d\angle\Psi_n(\lambda_n\tau\omega)}{d\omega} \leq 0$ shows that $\angle\Psi_n(\lambda_n\tau\omega) \geq -2\pi$ for all $\omega \in [0 \ 2\pi]$, i.e. proves the left-hand side of (9). ■

Like $\Phi(\alpha)$ above, $\Psi_n(\tau s)$ is a rational function. It can therefore be written as an LFR $\mathcal{F}_u(P_n(s), \Delta_m)$, where $P_n(s)$ is a n th-order LTI system and $\Delta_m = \beta I_n$ contains a single normalized parameter $\beta = -1 + 2\tau/\phi \in]-1 \ 1]$, as stated in Proposition 4.1.

Proposition 4.1: The Padé approximation $\Psi_n(\tau s)$ of any order $n \geq 1$ of the time delay $\tau > 0$:

$$\Psi_n(\tau s) = \frac{\sum_{i=0}^n (-1)^i a_i (\tau s)^i}{\sum_{i=0}^n a_i (\tau s)^i}, \quad a_i = \frac{(2n-i)!}{i!(n-i)!} \quad (11)$$

can be written as a Linear Fractional Representation $\mathcal{F}_u(P_n(s), \beta I_n)$, where $\beta = -1 + 2\tau/\phi$ and $\phi > 0$. A state-space representation of $P_n(s)$ is given by:

$$A_P = \frac{2}{\phi} \begin{bmatrix} -\frac{a_{n-1}}{a_n} & -\frac{a_{n-2}}{a_{n-1}} & 0 & \dots & 0 \\ \frac{a_{n-1}}{a_n} & 0 & -\frac{a_{n-3}}{a_{n-2}} & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & 0 \\ \vdots & \vdots & & \ddots & -\frac{a_0}{a_1} \\ (-1)^n \frac{a_{n-1}}{a_n} & 0 & \dots & \dots & 0 \end{bmatrix} \in \mathbb{R}^{n \times n} \quad (12)$$

$$B_P = \begin{bmatrix} 1 & 1 & 0 & \dots & 0 & 2 \\ -1 & 0 & 1 & \ddots & \vdots & 0 \\ \vdots & \vdots & \ddots & \ddots & 0 & \vdots \\ \vdots & \vdots & & \ddots & 1 & \vdots \\ (-1)^{n+1} & 0 & \dots & \dots & 0 & 1 + (-1)^{n+1} \end{bmatrix} \in \mathbb{R}^{n \times (n+1)} \quad (13)$$

$$C_P = \frac{2}{\phi} \begin{bmatrix} \frac{a_{n-1}}{a_n} & 0 & \dots & 0 \\ 0 & \frac{a_{n-2}}{a_{n-1}} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & \frac{a_0}{a_1} \\ \hline (-1)^{n+1} \frac{a_{n-1}}{a_n} & 0 & \dots & 0 \end{bmatrix} \in \mathbb{R}^{(n+1) \times n} \quad (14)$$

$$D_P = \begin{bmatrix} -1 & 0 & \dots & 0 & \vdots & 0 \\ 0 & -1 & \ddots & \vdots & \vdots & \vdots \\ \vdots & \ddots & \ddots & 0 & \vdots & \vdots \\ 0 & \dots & 0 & -1 & \vdots & 0 \\ \hline (-1)^n & 0 & \dots & 0 & \vdots & (-1)^n \end{bmatrix} \in \mathbb{R}^{(n+1) \times (n+1)} \quad (15)$$

where $\frac{a_{i-1}}{a_i} = \frac{(2n-i+1)i}{n-i+1}$. This representation is normalized in the sense that $\beta \in]-1 \ 1] \Leftrightarrow \tau \in]0 \ \phi]$, and minimal since the order of $P_n(s)$ and the number of repetitions of β are as small as possible (both equal to n).

Proof: See appendix. ■

Standard matrix manipulations based on the Redheffer star product [13] finally allow to equivalently transform the interconnection of Fig. 2, where $G(\tau s)$ is replaced with $\Psi_n(\tau s)$, into that of Fig. 4, where $\Delta_m = \beta I_n$.

Let us now consider a box $\mathcal{D} \subset \mathcal{B}_\delta$. In both cases – equivalent representation and Padé approximation – the interconnection $N(s) - \text{diag}(\Delta, \Delta_m)$, $\Delta \in \mathcal{D}_\Delta$, of Fig. 4 is normalized and replaced with $\tilde{N}(s) - \text{diag}(\tilde{\Delta}, \Delta_m)$, $\tilde{\Delta} \in \mathcal{B}_\Delta$. Several methods are then introduced in Sections IV-B and IV-C to check whether the delay margin requirement ϕ introduced in Section II is satisfied or violated on the entire box \mathcal{D} .

Remark 4.1: Both the equivalent representation (3) and the Padé approximation (11) of the delay (2) lead to the interconnection of Fig. 4, with the same Δ , but with different $N(s)$ and Δ_m . In the sequel, Δ_m is made explicit when necessary to avoid confusion: $\tilde{N}(s) - \text{diag}(\tilde{\Delta}, \alpha I_2)$ (resp. $\tilde{N}(s) - \text{diag}(\tilde{\Delta}, \beta I_n)$) corresponds to the equivalent (resp. Padé-based) LFR.

B. Satisfaction test

Based on the **equivalent representation** (3) of the delay function, Proposition 4.2 provides a necessary and sufficient condition to verify whether the delay margin requirement is satisfied on an entire box $\mathcal{D} \subset \mathcal{B}_\delta$.

Proposition 4.2: The delay margin is larger than ϕ on a given box $\mathcal{D} \subset \mathcal{B}_\delta$ if and only if the interconnection $\tilde{N}(s) - \text{diag}(\tilde{\Delta}, \alpha I_2)$ is stable $\forall \tilde{\Delta} \in \mathcal{B}_\Delta$ and $\forall \alpha \in [0 \ \alpha_{max}(\omega)]$.

Proof: The delay margin is larger than ϕ on \mathcal{D} iff the interconnection of Fig. 2, where $G(\tau s) = e^{-\tau s}$, is stable $\forall \Delta \in \mathcal{D}_\Delta$ and $\forall \tau \in [0 \ \phi]$. This condition remains necessary and sufficient first when $G(\tau s)$, $\tau \in [0 \ \phi]$, is replaced with $\Phi(\alpha)$, $\alpha \in [0 \ \alpha_{max}(\omega)]$, then when the resulting interconnection is replaced with that of Fig. 4, where $\Delta_m = \alpha I_2$, $\alpha \in [0 \ \alpha_{max}(\omega)]$, and finally when Δ is normalized. Indeed, these successive interconnections are rigorously equivalent. ■

As stated in Lemma 4.2, Proposition 4.2 boils down to solving a non-standard μ -analysis problem, where the bound on α is frequency-dependent.

Lemma 4.2: The following two conditions are equivalent:

- 1) $\tilde{N}(s) - \text{diag}(\tilde{\Delta}, \alpha I_2)$ is stable $\forall \tilde{\Delta} \in \mathcal{B}_\Delta$ and $\forall \alpha \in [0 \ \alpha_{max}(\omega)]$
- 2) $\mu(\tilde{N}_\omega(j\omega)) \leq 1 \quad \forall \omega \geq 0$

where the interconnection $\tilde{N}_\omega(s) - \text{diag}(\tilde{\Delta}, \xi I_2)$ is obtained from $\tilde{N}(s) - \text{diag}(\tilde{\Delta}, \alpha I_2)$ by normalizing α at the frequency ω , i.e. by setting $\alpha = \frac{\alpha_{max}(\omega)}{2}(\xi + 1)$, $\xi \in [-1 \ 1]$.

Proof: According to the definition of the structured singular value μ and the structured version of the small gain theorem, 1) \Rightarrow 2). The converse is not straightforward, since $\tilde{N}_\omega(s)$ has some complex coefficients coming from the transformation matrix T introduced in (7). It must be shown that there is no need to consider negative frequencies in condition 2). Assume that $\mu(\tilde{N}_\omega(j\omega)) \leq 1$ for some $\omega > 0$. This means that:

$$\det \left(I - \tilde{N}_\omega(j\omega) \text{diag} \left(\tilde{\Delta}, \xi I_2 \right) \right) \neq 0 \quad (16)$$

$\forall \tilde{\Delta} \in \mathcal{B}_\Delta$ and $\forall \xi \in [-1 \ 1]$. This is equivalent to:

$$\det (I - M(j\omega) \text{diag} (\Delta, -G(j\omega))) \neq 0 \quad (17)$$

Noting that $\det(\overline{X}) = \overline{\det(X)}$ for any complex square matrix X , that Δ is a real matrix, and that $M(s)$ has real coefficients, condition (17) is in turn equivalent to:

$$\det \left(I - M(-j\omega) \text{diag} \left(\Delta, -\overline{G(j\omega)} \right) \right) \neq 0 \quad (18)$$

According to equation (2), $\overline{G(j\omega)} = e^{-j\tau\omega} = e^{j\tau\omega} = G(-j\omega)$. Therefore, condition (18) is equivalent to:

$$\det (I - M(-j\omega) \text{diag} (\Delta, -G(-j\omega))) \neq 0 \quad (19)$$

This shows that the non-singularity condition (17) – and therefore (16) – is also satisfied when ω is replaced with $-\omega$, and that condition 2) is sufficient to cover both positive and negative frequencies, which means that 2) \Rightarrow 1). ■

A classical approach to solve the μ -analysis problem of Lemma 4.2 is to replace the whole frequency range with a finite grid $(\omega_k)_k$, and check if $\mu(\tilde{N}_{\omega_k}(j\omega_k)) \leq 1$ for all k (due to NP-hardness, upper bounds $\bar{\mu}(\tilde{N}_{\omega_k}(j\omega_k))$ are computed in practice instead of the exact values, typically using the (D, G) -scaling formulation of [29]). But this strategy is doomed to fail. When the box \mathcal{D} gets closer to the limit between the domains of delay margin satisfaction and violation, the peak value of μ tends to 1 and the risk of missing critical frequencies where $\mu > 1$ increases. Adding more frequencies to the grid can be tempting, but the computational time quickly becomes prohibitive. Moreover, several tests have shown that some frequencies where $\mu > 1$ always end up being missed, thus leading to the erroneous claim that the delay margin requirement is satisfied.

A better approach relies on the Hamiltonian-based algorithm of [30], which computes an entire frequency interval on which a μ upper bound remains valid, assuming that all uncertainties have constant bounds. This algorithm can be used as is for $\omega \geq 2\pi/\phi$, since in this case $\alpha_{max}(\omega) = 1$, see Fig. 3 (right). However, a three-step adaptation formalized in Algorithm 1 has to be introduced for $\omega \in [0, 2\pi/\phi]$, where $\alpha_{max}(\omega)$ is frequency-dependent:

- 1) The method of [30] is first applied at a given frequency ω_k , assuming that all uncertainties (including α) have constant bounds. This leads to an interval $[\underline{\omega}_k^0, \bar{\omega}_k^0]$ including ω_k on which $\bar{\mu}(\tilde{N}_{\omega_k}(j\omega)) \leq 1$ with the D, G scalings computed at ω_k . But the upper bound $\alpha_{max}(\omega)$ on α actually increases with frequency for the considered delay margin problem, as shown in Fig. 3 (right). The real objective is therefore to compute the largest possible interval $[\underline{\omega}_k, \bar{\omega}_k]$ including ω_k on which $\bar{\mu}(\tilde{N}_{\omega_k}(j\omega)) \leq 1$. The increase of $\alpha_{max}(\omega)$ with frequency implies that $\underline{\omega}_k < \omega_k$. On the other hand, $\bar{\omega}_k^0$ is not an admissible upper bound, but $\bar{\omega}_k$ can be sought in the interval $[\omega_k, \bar{\omega}_k^0]$. At this stage, all that is known is that $[\underline{\omega}_k^0, \bar{\omega}_k^0]$, where $\bar{\omega}_k^0 = \omega_k$, is a valid interval. The next two steps respectively consist of decreasing its lower bound and increasing its upper bound as much as possible.
- 2) As illustrated in Fig. 6 (left), a way to decrease the lower bound is to normalize α at $\omega_k^1 = \omega_k^0$, and compute using [30] an interval $[\underline{\omega}_k^1, \bar{\omega}_k^1]$ on which $\bar{\mu}(\tilde{N}_{\omega_k^1}(j\omega)) \leq 1$, still considering the D, G scalings previously computed at ω_k . $\underline{\omega}_k^1$ is necessary lower than $\underline{\omega}_k^0$, since the range of variation of α is lower than before, $[0, \alpha_{max}(\omega_k^1)]$ instead of $[0, \alpha_{max}(\omega_k)]$. The same is repeated at frequencies $\omega_k^2 = \omega_k^1, \dots, \omega_k^i = \omega_k^{i-1}$ until the decrease becomes marginal, which finally leads to $\underline{\omega}_k = \omega_k^i$. Stability is then proved for the gray region in Fig. 6 (left), which encompasses the area under the $\alpha_{max}(\omega)$ curve over the interval $[\underline{\omega}_k, \omega_k]$.
- 3) The upper bound can be increased using a dichotomic search on the interval $[\bar{\omega}_k^0, \omega_k^0]$. To do this, α is normalized at the central frequency ω_k^1 , and it is first checked if $\bar{\mu}(\tilde{N}_{\omega_k^1}(j\bar{\omega}_k^0)) \leq 1$ with the D, G scalings previously computed at ω_k . If no, the dichotomy continues over the interval $[\bar{\omega}_k^1, \omega_k^1] = [\bar{\omega}_k^0, \omega_k^1]$. Otherwise, an interval $[\bar{\omega}_k^0, \omega_m^1]$ on which $\bar{\mu}(\tilde{N}_{\omega_k^1}(j\omega)) \leq 1$ is computed, as depicted in Fig. 6 (right). If $\omega_m^1 \geq \omega_k^1$ (case illustrated in the figure), it can be concluded that $\bar{\omega}_k^1 = \omega_k^1$ is a lower bound on $\bar{\omega}_k$, and the dichotomy continues over the interval $[\bar{\omega}_k^1, \omega_k^1] = [\omega_k^1, \omega_k^0]$. The same process is then repeated at the central frequency ω_k^2 . In the example of Fig. 6 (right), $\omega_m^2 \leq \omega_k^2$. ω_m^2 is thus a lower bound on $\bar{\omega}_k$, and the dichotomy continues over the interval $[\bar{\omega}_k^2, \omega_k^2] = [\omega_m^2, \omega_k^2]$. The frequencies $\bar{\omega}_k^3, \dots, \bar{\omega}_k^i$ are determined in the same way until the increase becomes marginal, which finally leads to $\bar{\omega}_k = \bar{\omega}_k^i$ ($i = 3$ in the figure). Stability is then proved for the gray region in Fig. 6 (right), which encompasses the area under the $\alpha_{max}(\omega)$ curve over the interval $[\omega_k, \bar{\omega}_k]$.

In practice, Algorithm 1 is applied repeatedly at frequencies $\omega_1, \omega_2, \dots$ for which no μ upper bound is known yet, until:

- the union of all validity intervals $[\underline{\omega}_1, \bar{\omega}_1], [\underline{\omega}_2, \bar{\omega}_2], \dots$ covers the whole frequency range, *i.e.* $\cup_k [\underline{\omega}_k, \bar{\omega}_k] = \mathbb{R}_+$, which guarantees that the delay margin is larger than ϕ on \mathcal{D} ,
- or a frequency ω_k is found such that $\bar{\mu}(\tilde{N}_{\omega_k}(j\omega_k)) > 1$, which means that either the delay margin is not larger than ϕ on the entire box \mathcal{D} , or the μ upper bound is too conservative to prove the converse (see Remark 4.2).

Remark 4.2: Although the equivalent representation (3) of the delay function is used, the condition of Proposition 4.2 becomes only sufficient if μ upper bounds are computed instead of the exact values. Fortunately, conservatism remains low in most cases, as highlighted in [14]. Moreover, [31] shows that it tends to zero when the size of \mathcal{D} tends to zero.

Remark 4.3: Algorithm 1 requires a single μ upper bound computation, at step 1.2. The associated D, G scalings are then simply reused when applying the Hamiltonian-based technique of [30], which is quite fast since it boils down to an eigenvalue computation. Therefore, computing the interval $[\underline{\omega}_k, \bar{\omega}_k]$ around ω_k on which $\bar{\mu}(\tilde{N}_{\omega_k}(j\omega)) \leq 1$ is not much longer than checking whether $\bar{\mu}(\tilde{N}_{\omega_k}(j\omega_k)) \leq 1$. This allows to cover the whole frequency range \mathbb{R}_+ quickly.

Remark 4.4: It would be possible to use the scaled Padé approximation $\Psi_n(\lambda_n \tau s)$ instead of the equivalent representation (3) of the delay function for the satisfaction test. Indeed, the phase variations of $\Psi_n(\lambda_n \tau s)$ for all $\tau \in [0, \phi]$ encompass those of $G(\tau s)$, as seen in Fig. 5 (right) and proved in Lemma 4.1. But unlike the equivalent representation, $\Psi_n(\lambda_n \tau s)$ is only an approximation. Its order n can be increased to improve accuracy, and it should in any case be greater than 3 for the right-hand side of inequality (9) to be satisfied. As a consequence, the parameter β is repeated $n \geq 3$ times in the resulting LFR, whereas α (and therefore ξ) is only repeated twice with the equivalent representation. The latter is therefore both more accurate and computationally less involving, that is why it is privileged here.

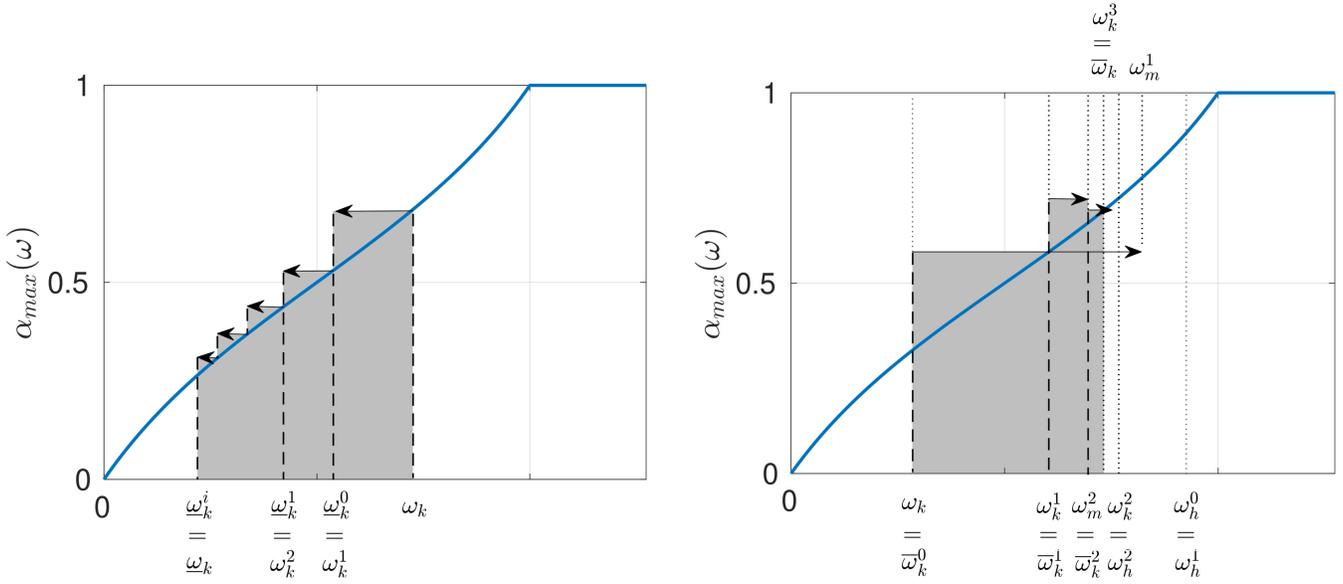


Fig. 6: Determination of the frequency interval $[\underline{\omega}_k \ \overline{\omega}_k]$ on which $\overline{\mu}(\tilde{N}_\omega(j\omega)) \leq 1$ (left = lower bound, right = upper bound)

Algorithm 1 Validity interval $[\underline{\omega}_k \ \overline{\omega}_k]$ around a given frequency ω_k

Step 1 - Initialization:

- 1) normalize α at ω_k to get the fully normalized interconnection $\tilde{N}_{\omega_k}(s) - \text{diag}(\tilde{\Delta}, \xi I_2)$, $\tilde{\Delta} \in \mathcal{B}_\Delta$, $\xi \in [-1 \ 1]$
- 2) check whether $\overline{\mu}(\tilde{N}_{\omega_k}(j\omega_k)) \leq 1$ using the (D, G) -scaling formulation of [29]
- 3) **if not then STOP else** use [30] with the D, G scalings of step 1.2 to compute an interval $[\underline{\omega}_k^0 \ \omega_h^0]$ around ω_k on which $\overline{\mu}(\tilde{N}_{\omega_k}(j\omega)) \leq 1$

Step 2 - Lower bound improvement:

initialization: set $i = 0$ and $\omega_k^0 = \omega_k$

while $\omega_k^i - \underline{\omega}_k^i > \epsilon$ **do**

- 1) set $i \leftarrow i + 1$ and $\omega_k^i = \underline{\omega}_k^{i-1}$
- 2) normalize α at ω_k^i and compute $\tilde{N}_{\omega_k^i}(s)$
- 3) use [30] with the D, G scalings of step 1.2 to compute an interval $[\underline{\omega}_k^i \ \omega_k^i]$ on which $\overline{\mu}(\tilde{N}_{\omega_k^i}(j\omega)) \leq 1$

end while

set $\underline{\omega}_k = \omega_k^i$

Step 3 - Upper bound improvement:

initialization: set $i = 0$ and $\overline{\omega}_k^0 = \omega_k$

while $\omega_h^i - \overline{\omega}_k^i > \epsilon$ **do**

- 1) set $i \leftarrow i + 1$ and $\omega_k^i = \frac{(\overline{\omega}_k^{i-1} + \omega_h^{i-1})}{2}$
- 2) normalize α at ω_k^i and compute $\tilde{N}_{\omega_k^i}(s)$
- 3) **if** $\overline{\mu}(\tilde{N}_{\omega_k^i}(j\overline{\omega}_k^{i-1})) \leq 1$ with the D, G scalings of step 1.2 **then**
 - use [30] with the D, G scalings of step 1.2 to compute an interval $[\overline{\omega}_k^{i-1} \ \omega_m^i]$ on which $\overline{\mu}(\tilde{N}_{\omega_k^i}(j\omega)) \leq 1$
 - **if** $\omega_m^i > \omega_k^i$ **then** set $[\overline{\omega}_k^i \ \omega_h^i] = [\omega_k^i \ \omega_h^{i-1}]$ **else** set $[\overline{\omega}_k^i \ \omega_h^i] = [\omega_m^i \ \omega_k^i]$
 - else** set $[\overline{\omega}_k^i \ \omega_h^i] = [\overline{\omega}_k^{i-1} \ \omega_k^i]$

end while

set $\overline{\omega}_k = \overline{\omega}_k^i$

C. Violation test

Based on the **Padé approximation** (11) of the delay function, Proposition 4.3 provides a sufficient condition to verify whether the delay margin requirement is violated on an entire box $\mathcal{D} \subset \mathcal{B}_\delta$.

Proposition 4.3: The delay margin is lower than ϕ on a given box $\mathcal{D} \subset \mathcal{B}_\delta$ if $\forall \tilde{\Delta} \in \mathcal{B}_\Delta$, $\exists \hat{\beta} \in [-1 \ 1]$ such that the interconnection $\tilde{N}(s) - \text{diag}(\tilde{\Delta}, \hat{\beta} I_n)$ is unstable.

Proof: The delay margin is lower than ϕ on \mathcal{D} iff $\forall \Delta \in \mathcal{D}_\Delta, \exists \tau \in [0, \phi]$ such that the interconnection of Fig. 2, where $G(\tau s) = e^{-\tau s}$, is unstable. The fact that the left-hand side of equation (8) is not an equality means that there is a region between the phase plots of $G(\phi s)$ and $\Psi_n(\phi s)$ that cannot be reached by any n th-order Padé approximation $\Psi_n(\tau s), \tau \in [0, \phi]$. So replacing $G(\tau s)$ with $\Psi_n(\tau s)$ makes the condition of Proposition 4.3 only sufficient. Finally, equivalently replacing the interconnection of Fig. 2 with that of Fig. 4, and normalizing Δ and τ , does not change the result. ■

The condition of Proposition 4.3 cannot be directly evaluated by means of standard μ -based tools and is therefore replaced by that of Proposition 4.4, which amounts to choosing the same value of $\hat{\beta}$ for all $\tilde{\Delta} \in \mathcal{B}_\Delta$.

Proposition 4.4: The delay margin is lower than ϕ on a given box $\mathcal{D} \subset \mathcal{B}_\delta$ if $\exists \hat{\beta} \in [-1, 1]$ such that the interconnection $\tilde{N}(s) - \text{diag}(\tilde{\Delta}, \hat{\beta}I_n)$ is unstable $\forall \tilde{\Delta} \in \mathcal{B}_\Delta$.

Proof: The condition of Proposition 4.4 being obtained by imposing the same value of $\hat{\beta}$ for all $\tilde{\Delta} \in \mathcal{B}_\Delta$, it implies that of Proposition 4.3. ■

As $\hat{\beta}$ remains constant in Proposition 4.4, it can be integrated into $\tilde{N}(s)$ to form a reduced normalized interconnection $\tilde{N}_r(s) - \tilde{\Delta}$, where $\tilde{N}_r(s)$ is unstable. It is then sufficient to check whether this reduced interconnection remains unstable $\forall \tilde{\Delta} \in \mathcal{B}_\Delta$, which can be done easily by determining with standard μ -based tools if $\bar{\mu}(\tilde{N}_r(j\omega)) \leq 1 \forall \omega \geq 0$. Note that it is sufficient to consider positive frequencies here. Indeed, $\tilde{N}_r(s)$ being the normalized interconnection of system (1) and the Padé approximation (11), it only has real coefficients.

As suggested in [6], $\hat{\beta}$ can be determined by studying the stability of the nominal interconnection $\tilde{N}(s) - \text{diag}(0_{p \times p}, \beta I_n)$. A finite number of values $(\beta^k)_k$ which grid the interval $[-1, 1]$ are considered, and the one which moves a pole of the interconnection the farthest in the right half-plane is selected to be $\hat{\beta}$. One drawback is that it does not prevent a pole of $\tilde{N}_r(s)$ from being very close to the imaginary axis, which might make the aforementioned μ -test fail. An alternative is to select the β^k which makes $\tilde{N}_r(s)$ unstable and maximizes the smallest distance between a pole of $\tilde{N}_r(s)$ and the imaginary axis. It has been observed in practice that this strategy increases the chances of the aforementioned μ -test being successful.

Unlike the delay margin satisfaction test of Section IV-B, some conservatism is introduced here:

- 1) The Padé approximation (11) is used instead of the exact delay (2). This is admissible, since its phase is smaller in absolute value than that of $G(\tau s)$ for all frequencies, as stated in Lemma 4.1. But unlike Proposition 4.2, this makes the condition of Proposition 4.3 non-necessary. Fortunately, high-order Padé approximations can be used to mitigate this conservatism with almost no increase in computational time, see Remark 4.5 and Section V.
- 2) $\hat{\beta}$ is fixed to a constant value in Proposition 4.4, which makes the condition even more non-necessary. But as the size of \mathcal{D} decreases along the iterations of the branch-and-bound algorithm (see Algorithm 2 in Section IV-D), there is more and more chance that a single value of $\hat{\beta}$ will suit.
- 3) The way to choose $\hat{\beta}$ is intuitive, but it may not be optimal. It may even happen that no $\beta \in (\beta^k)_k$ makes $\tilde{N}_r(s)$ unstable, whereas there exists one in $[-1, 1]$ which does. In the latter case, the satisfaction test of Proposition 4.2 is performed on the box \mathcal{D} instead of the violation test of Proposition 4.4, and it inevitably fails due to the aforementioned unstable configuration. Failing to detect a destabilizing $\hat{\beta}$ therefore has no other consequence than splitting \mathcal{D} unnecessarily when applying the branch-and-bound algorithm of Section IV-D, which might slightly increase the computational time but is not a critical issue. Moreover, it has been observed in [6] when working on the gain and phase margins that the grid-based approach to select $\hat{\beta}$ usually gives satisfactory results.

It is also possible to propose a sufficient condition based on the **equivalent representation** (3) instead of the Padé approximation $\Psi_n(\tau s)$ to check whether the delay margin requirement is violated on an entire box $\mathcal{D} \subset \mathcal{B}_\delta$. $\hat{\beta}$ is first computed as above, giving a delay $\hat{\tau} \in [0, \phi]$ such that the interconnection of Fig. 2 is unstable for $G(\hat{\tau} s)$ and $\Delta = 0$. This conversion from $\hat{\beta}$ into $\hat{\tau}$ is always possible, since the phase variations of $G(\tau s)$ for all $\tau \in [0, \phi]$ encompass those of $\Psi_n(\tau s)$, as seen in Fig. 5 and proved in Lemma 4.1. The interconnection $\tilde{N}(s) - \text{diag}(0_{p \times p}, \hat{\alpha}(\omega)I_2)$ is therefore unstable, where $\hat{\alpha}(\omega)$ is obtained from equation (4) by replacing ϕ with $\hat{\tau}$.

Proposition 4.5: The delay margin is lower than ϕ on a given box $\mathcal{D} \subset \mathcal{B}_\delta$ if the interconnection $\tilde{N}(s) - \text{diag}(\tilde{\Delta}, \hat{\alpha}(\omega)I_2)$ is unstable $\forall \tilde{\Delta} \in \mathcal{B}_\Delta$.

Unlike α which belongs to a (frequency-dependent) interval in Proposition 4.2, $\hat{\alpha}(\omega)$ is fixed here at a given frequency. It should therefore be integrated into $\tilde{N}(s)$ before a μ -based test can be performed. However, this results in frequency-dependent state-space matrices, which prevents the use of the Hamiltonian-based algorithm of [30]. The problem could be solved on a frequency grid, but as already highlighted in Section IV-B, this is not reliable. Consequently, the sufficient condition of Proposition 4.5 cannot be evaluated easily and is replaced with that of Proposition 4.6. The idea is to check instability for all values of α in a small interval $[\hat{\alpha}(\omega) - \epsilon_\alpha, \hat{\alpha}(\omega) + \epsilon_\alpha]$ around $\hat{\alpha}(\omega)$, and not only for $\alpha = \hat{\alpha}(\omega)$. α is thus considered as an uncertainty as in Section IV-B.

Proposition 4.6: The delay margin is lower than ϕ on a given box $\mathcal{D} \subset \mathcal{B}_\delta$ if the interconnection $\tilde{N}(s) - \text{diag}(\tilde{\Delta}, \alpha I_2)$ is unstable $\forall \tilde{\Delta} \in \mathcal{B}_\Delta$ and $\forall \alpha \in [\hat{\alpha}(\omega) - \epsilon_\alpha, \hat{\alpha}(\omega) + \epsilon_\alpha]$, where ϵ_α introduces a small deviation around $\hat{\alpha}(\omega)$.

Proof: The condition of Proposition 4.6 clearly implies that of Proposition 4.5. ■

Similarly to Proposition 4.2, Proposition 4.6 boils down to solving a non-standard μ -analysis problem, where the bounds on α are frequency-dependent. A variant of Algorithm 1 can be used for this purpose, with the following key differences:

- The normalization at a given frequency ω is done by setting $\alpha = \hat{\alpha}(\omega) + \epsilon_\alpha \xi$, $\xi \in [-1, 1]$.
- At the end of step 1, $\underline{\omega}_k^0$ is the largest between the lower bound of the interval determined by the Hamiltonian-based method of [30] (case shown in Fig. 7), and the frequency ω for which $\hat{\alpha}(\omega) = \hat{\alpha}(\omega_k) - \epsilon_\alpha$. Similarly, $\bar{\omega}_k^0$ is the smallest between the upper bound of the interval determined by [30], and the frequency ω for which $\hat{\alpha}(\omega) = \hat{\alpha}(\omega_k) + \epsilon_\alpha$ (case shown in Fig. 7).
- Step 2 is similar to Algorithm 1, except that the frequency ω for which $\hat{\alpha}(\omega) = \hat{\alpha}(\omega_k^i) - \epsilon_\alpha$ must also be considered as in the previous item when computing $\underline{\omega}_k^i$.
- Step 3 is not a dichotomic search anymore, but becomes similar to Step 2.

In the example of Fig. 7, instability is proven for the (light and dark) gray region, which encompasses the $\hat{\alpha}(\omega)$ curve over the interval $[\underline{\omega}_k^1, \bar{\omega}_k^1]$. By letting the algorithm run until progress becomes marginal, larger and larger intervals $[\underline{\omega}_k^2, \bar{\omega}_k^2]$, $[\underline{\omega}_k^3, \bar{\omega}_k^3]$, \dots are obtained, the last one being the desired validity interval $[\underline{\omega}_k, \bar{\omega}_k]$.

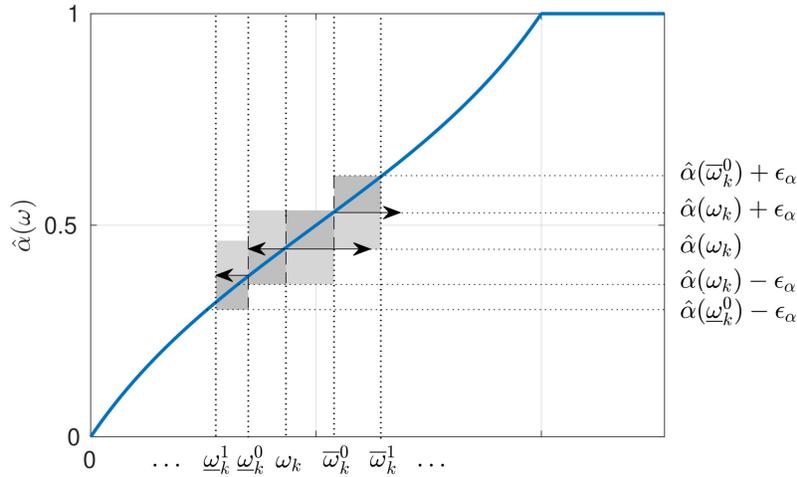


Fig. 7: Adaptation of Algorithm 1 for the delay margin violation test

On the one hand, ϵ_α should be sufficiently small to maximize the chances of the μ -test being successful at step 1.2, and thus to limit conservatism. On the other hand, a very small value makes the length of the intervals $[\underline{\omega}_k^i, \bar{\omega}_k^i]$ increase quite slowly, which results in a large computational time. In practice, setting $\epsilon_\alpha \approx 0.01$ seems to be a good compromise. Several variations of this algorithm can be implemented. A skew- μ computation can for example be performed at step 1.2 instead of a μ -test to determine the largest ϵ_α which makes the interconnection $\tilde{N}(s) - \text{diag}(\tilde{\Delta}, \alpha I_2)$ unstable $\forall \tilde{\Delta} \in \mathcal{B}_\Delta$. It is also possible to consider the interval $[\hat{\alpha}(\omega) - \epsilon_\alpha^1, \hat{\alpha}(\omega) + \epsilon_\alpha^2]$, where typically $\epsilon_\alpha^1 < \epsilon_\alpha^2$, to take into account the fact that a system usually tends to become more unstable when increasing the delay than when reducing it. Finally, another option is to consider two different intervals $[\hat{\alpha}(\omega) - \epsilon_\alpha, \hat{\alpha}(\omega)]$ and $[\hat{\alpha}(\omega), \hat{\alpha}(\omega) + \epsilon_\alpha]$ depending on whether ω is smaller or larger than ω_k , *i.e.* for step 2 and step 3 respectively. Values below (resp. above) $\hat{\alpha}(\omega)$ are indeed useless when computing $\bar{\omega}_k$ (resp. $\underline{\omega}_k$), and even penalizing. This last variant only requires instability to be proven for the dark gray region in Figure 7. It has proved to be the most effective in practice, and this is the one evaluated in Section V.

The equivalent representation (3) is more accurate than Padé approximations. Nevertheless, some conservatism is introduced in Proposition 4.6. First, the function $\hat{\alpha}(\omega)$ is obtained from $\hat{\beta}$ and therefore suffers from the same drawbacks (see items 2 and 3 above). Then, there is a region between the phase plots of $G(\phi s)$ and $\Psi_n(\phi s)$ that cannot be reached by any n th-order Padé approximation $\Psi_n(\tau s)$, $\tau \in [0, \phi]$. This means that there may be delays in $[0, \phi]$ that destabilize the system, whereas no n th-order Padé approximation does. But this risk decreases as the order of the approximation increases. Moreover, as already highlighted above, failing to detect a destabilizing delay has no other consequence than splitting \mathcal{D} unnecessarily when applying the branch-and-bound algorithm of Section IV-D. Finally, instability is investigated over the whole interval $[\hat{\alpha}(\omega) - \epsilon_\alpha, \hat{\alpha}(\omega) + \epsilon_\alpha]$ instead of just for $\hat{\alpha}(\omega)$. This justifies the choice of a sufficiently small ϵ_α , but nevertheless not too small, as explained above.

Remark 4.5: As in Section IV-B, using a Padé approximation (11) is more conservative than using the equivalent representation (3) of the delay function. But unlike the satisfaction test, β is set to $\hat{\beta}$ for the violation test, and no longer appears in the LFR. The Padé-based condition of Proposition 4.4 is therefore faster than its satisfaction-oriented counterpart mentioned in Remark 4.4, but also faster than the condition of Proposition 4.6, where α cannot be replaced with $\hat{\alpha}$ and still appears in the LFR. Moreover, it is possible to choose a high-order approximation without any significant increase in the computational time, as illustrated in Section V.

D. Algorithmic issues

The conditions for determining whether a given delay margin requirement is satisfied (Proposition 4.2) or violated (Proposition 4.4 or 4.6) on an entire box $\mathcal{D} \subset \mathcal{B}_\delta$ are embedded into the branch-and-bound scheme of Algorithm 2, which is similar to the one of [6]. The uncertainty domain is divided into smaller and smaller boxes until each box has guaranteed sufficient/insufficient delay margin, or has a probability lower than a user-defined threshold p_{min} . Before that, a preliminary stability analysis should be performed with Algorithm 1 of [25], leading to:

$$\mathcal{B}_\delta = D_s \cup D_{\bar{s}} \cup D_{s_u} \quad (20)$$

where D_s , $D_{\bar{s}}$ and D_{s_u} are the domains of guaranteed stability, guaranteed instability and undetermined stability respectively. The delay margin can indeed only be evaluated for stable systems, so the application of Algorithm 2 is restricted to D_s , leading to:

$$D_s = D_m \cup D_{\bar{m}} \cup D_{m_u} \quad (21)$$

where D_m , $D_{\bar{m}}$ and D_{m_u} are the domains of guaranteed delay margin satisfaction, guaranteed delay margin violation and undetermined delay margin, with corresponding probabilities $p(D_m)$, $p(D_{\bar{m}})$ and $p(D_{m_u})$. Combing (20) and (21) leads to the following partitioning of \mathcal{B}_δ :

$$\mathcal{B}_\delta = D_m \cup D_{\bar{m}} \cup D_{m_u} \cup D_{\bar{s}} \cup D_{s_u} \quad (22)$$

Guaranteed bounds on the exact probability $\bar{P}_{\Delta, f}^\phi(M(s))$ of delay margin violation are finally obtained, which solves Problem 2.1:

$$p(D_{\bar{m}}) \leq \bar{P}_{\Delta, f}^\phi(M(s)) \leq p(D_{\bar{m}}) + p(D_{m_u}) + p(D_{s_u}) = p(D_s) - p(D_m) + p(D_{s_u}) \quad (23)$$

Algorithm 2 Probabilistic delay margin analysis

```

 $\mathcal{L} \leftarrow \{\mathcal{B}_\delta\}$   $\triangleright$  list of all boxes left to investigate
 $D_m, D_{\bar{m}}, D_{m_u} \leftarrow \emptyset$ 
 $p(D_m), p(D_{\bar{m}}), p(D_{m_u}) \leftarrow 0$ 
while  $\mathcal{L} \neq \emptyset$  do
  extract the box  $\mathcal{D} \in \mathcal{L}$  with the highest probability
  compute the interconnection  $\tilde{N}(s) - \text{diag}(\tilde{\Delta}, \beta I_n)$ 
  check if  $\hat{\beta}$  exists which makes  $\tilde{N}_r(s)$  unstable
  if not then  $\triangleright$  nominal delay margin on  $\mathcal{D}$  is  $\geq \phi$ 
    check margin satisfaction on  $\mathcal{D}$  with Proposition 4.2
    if guaranteed then
      add  $\mathcal{D}$  to  $D_m$  and set  $p(D_m) \leftarrow p(D_m) + p(\mathcal{D})$ 
    else
      declare current iteration as inconclusive
    end if
  else  $\triangleright$  nominal delay margin on  $\mathcal{D}$  is  $< \phi$ 
    check margin violation on  $\mathcal{D}$  with Proposition 4.4 or 4.6
    if guaranteed then
      add  $\mathcal{D}$  to  $D_{\bar{m}}$  and set  $p(D_{\bar{m}}) \leftarrow p(D_{\bar{m}}) + p(\mathcal{D})$ 
    else
      declare current iteration as inconclusive
    end if
  end if
  if current iteration is inconclusive then
    if  $p(\mathcal{D}) > p_{min}$  then
      select a direction for cutting  $\mathcal{D}$ , e.g. using the  $\mu$ -sensitivities [4]
      partition  $\mathcal{D}$  and add the resulting boxes into  $\mathcal{L}$ 
    else
      add  $\mathcal{D}$  to  $D_{m_u}$  and set  $p(D_{m_u}) \leftarrow p(D_{m_u}) + p(\mathcal{D})$ 
    end if
  end if
end while

```

V. NUMERICAL RESULTS AND COMPARATIVE STUDY

The two variants of Algorithm 2 (using either Proposition 4.4 or 4.6 for the delay margin violation test) have been implemented and incorporated into the STOchastic Worst Case Analysis Toolbox (STOWAT). To assess their capabilities, a simple satellite model with two uncertainties, adapted from [13], is first analyzed in Section V-A. The low number of uncertainties allows for graphical representation of the results, which enhances clarity. Similar to the other probabilistic stability margins [6], the analysis of more advanced models with more states and uncertainties is also possible, as shown in Section V-B. All computational times reported in this paper were obtained using Matlab R2022b running serially on a single core on a Windows 10 laptop from 2021 with an Intel Core i7-1165G7 CPU running at 3 GHz and 16 GB of RAM.

A. Spinning satellite benchmark

The satellite is represented as a symmetric cylinder spinning around the symmetry axis z with a constant angular rate Ω . The angular rates ω_x and ω_y around the x and y axes are controlled using torques T_x and T_y . Let I_x , $I_y = I_x$ and I_z be the inertia of the satellite with respect to the x , y and z axes respectively. The rotational motion can be described by:

$$\begin{cases} T_x = I_x \dot{\omega}_x - \omega_y \Omega (I_x - I_z) \\ T_y = I_x \dot{\omega}_y - \omega_x \Omega (I_z - I_x) \end{cases} \quad (24)$$

Uniformly distributed uncertain parameters $\delta_1 \in [-0.5 \ 2.5]$ and $\delta_2 \in [0 \ 2]$ are introduced, leading to:

$$\begin{bmatrix} \dot{\omega}_x \\ \dot{\omega}_y \end{bmatrix} = \begin{bmatrix} 0 & a \\ -a & 0 \end{bmatrix} \begin{bmatrix} \omega_x \\ \omega_y \end{bmatrix} + \begin{bmatrix} \delta_1 & 0 \\ 0 & \delta_2 \end{bmatrix} \begin{bmatrix} u_x \\ u_y \end{bmatrix} \quad (25)$$

where $a = \Omega \left(1 - \frac{I_z}{I_x}\right)$, $u_x = \frac{T_x}{I_x}$ and $u_y = \frac{T_y}{I_x}$. Two measures ν_x and ν_y are available:

$$\begin{bmatrix} \nu_x \\ \nu_y \end{bmatrix} = \begin{bmatrix} 1 & a \\ -a & 1 \end{bmatrix} \begin{bmatrix} \omega_x \\ \omega_y \end{bmatrix} \quad (26)$$

and a static controller K is applied:

$$\begin{bmatrix} u_x \\ u_y \end{bmatrix} = -K \begin{bmatrix} \nu_x \\ \nu_y \end{bmatrix} = - \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \nu_x \\ \nu_y \end{bmatrix} \quad (27)$$

It is assumed in the sequel that $a = 10$. A preliminary probabilistic stability analysis is performed on the uncertain closed-loop system (25)-(27) using Algorithm 1 of [25], so as to obtain partition (20). The first channel of the control-loop model is then opened, the second one remaining closed. An LFR is obtained as in Fig. 1, where $u = u_x$ and $y = \nu_x$. Probabilistic delay margin analysis is finally carried out for $\phi = 0.2s$ with Algorithm 2, where the violation test is performed using either a second-order Padé approximation as in [27] (variant #1) or the equivalent representation of the delay function (variant #2). All results are summarized in Table I for different values of the stopping criterion p_{min} . The number of iterations and the CPU time correspond to the delay margin analysis only, *i.e.* the stability analysis is not taken into account.

p_{min} (%)	$p(D_m)$ (%)		$p(D_{\bar{m}})$ (%)		$p(D_{m_u})$ (%)		Iterations		CPU time (s)		$p(D_{\bar{s}})$ (%)	$p(D_{s_u})$ (%)
	#1	#2	#1	#2	#1	#2	#1	#2	#1	#2		
10^{-2}	20.13	20.10	26.80	26.62	1.46	1.67	586	610	4.4	9.3	48.74	2.87
10^{-3}	21.44	21.42	27.51	27.39	0.46	0.60	1913	2133	12.6	26.9	49.62	0.97
10^{-4}	21.73	21.72	27.80	27.71	0.18	0.29	5517	6853	34.1	87.1	49.93	0.36
10^{-5}	21.89	21.89	27.91	27.87	0.05	0.09	22500	33248	145.8	428.9	50.06	0.09

TABLE I: Results for $\phi = 0.2s$ using both variants of Algorithm 2 and different values of the stopping criterion p_{min}

As expected, the domains of undetermined stability D_{s_u} and undetermined delay margin D_{m_u} become smaller as p_{min} decreases. Moreover, it can be checked that most boxes of D_{m_u} include combinations of uncertainties for which the delay margin is higher than ϕ , and others for which it is lower. This means that they could not be assigned to D_m or $D_{\bar{m}}$ anyway, which shows that the various sources of conservatism highlighted in Section IV-C are negligible in this example. Another positive feature is that the average CPU time per iteration is quite low – about 6.5ms (resp. 12.8ms) for variant #1 (resp. #2) – given the number of tasks performed, including the determination of $\hat{\beta}$, a μ upper bound computation on the whole frequency range, but also a number of other calculations, as described in Algorithm 2. Finally, it appears that variant #1 is slightly more efficient than variant #2, both in terms of accuracy and computational time. A likely explanation is that the second-order Padé approximation used instead of the exact delay in variant #1 is almost non-conservative in this example, while the threshold ϵ_α in variant #2 is a little less so. This is confirmed by decreasing ϵ_α , which improves accuracy but at the cost of a significant increase in computational time. The results obtained with variant #1 and $p_{min} = 10^{-2}\%$ are plotted in Fig. 8 as an example, with the following color code:

- green: the delay margin is guaranteed to be $\geq \phi$,
- red: the delay margin is guaranteed to be $< \phi$,
- blue: the delay margin is undetermined,
- orange: instability is guaranteed,
- gray: stability is undetermined.

A comparison is made with a classical grid-based approach, where the delay margin is computed for a finite number of points in the uncertainty domain:

- magenta: the delay margin is guaranteed to be $\geq \phi$,
- yellow: the delay margin is guaranteed to be $< \phi$,
- black: the system is unstable.

It can be observed that both analyses are in good agreement, but Algorithm 2 provides guaranteed results on the whole uncertainty domain, and not only for a finite number of points.

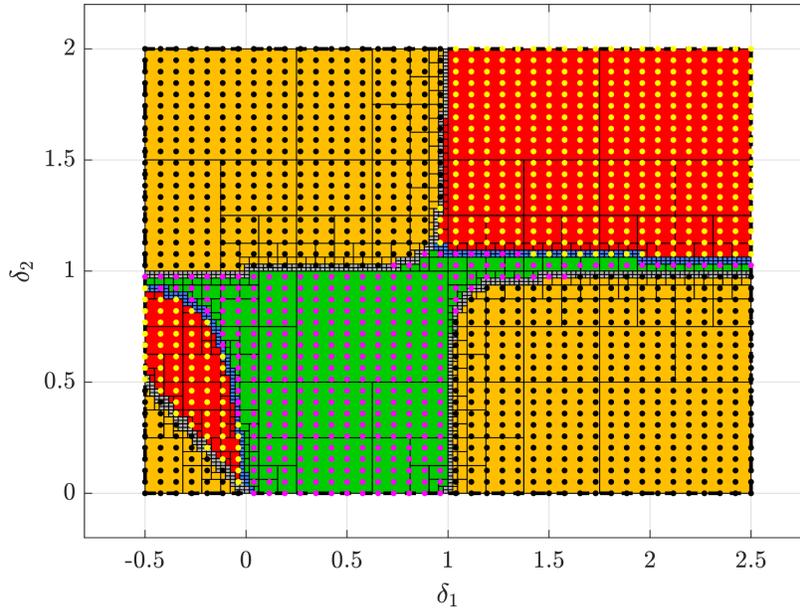


Fig. 8: Delay margin analysis for $\phi = 0.2\text{s}$ using variant #1 with $p_{min} = 10^{-2}\%$ and a second-order Padé approximation

As p_{min} decreases, the gray and blue areas shrink, becoming almost indistinguishable when $p_{min} \leq 10^{-4}\%$. However, if variant #1 is run again for $\phi = 0.4\text{s}$, a substantial part of the uncertainty domain remains undetermined, whatever the value of p_{min} . This can be observed in Fig. 9 (obtained with $p_{min} = 10^{-3}\%$), which focuses on the reduced uncertainty domain $\delta_1 \in [0, 1]$ and $\delta_2 \in [0, 0.75]$. To find out why, these results are compared with two grid-based approaches, where the system is perturbed by either an exact time delay (Fig. 9a) or a second-order Padé approximation (Fig. 9b). The two grids differ significantly in the undetermined area (blue), where the Padé approximation does not destabilize the system, while the true time delay does. Conservatism of variant #1 therefore results from the use of a second-order Padé approximation. It cannot be reduced by changing the stopping criterion of Algorithm 2, but rather by considering higher-order approximations. This is confirmed by Table II and Fig. 10 (obtained with $p_{min} = 10^{-3}\%$), where a significant decrease in $p(D_{mu})$ can be observed when the order of the Padé approximation increases. The number of iterations and the computational time follow the same trend. A costly iterative process, which consists of dividing the undetermined domain into smaller and smaller boxes until the stopping criterion is reached, is indeed alleviated. It can be seen that a fourth-order Padé approximation is the best compromise here, since there is no improvement with a fifth-order one. Moreover, the increase in computational time resulting from the choice of a higher-order approximation is marginal, as explained in Remark 4.5. As shown in Table II and Fig. 11, variant #2 gives intermediate results, better than variant #1 with a Padé approximation of order 2, but worse than with an approximation of order greater than or equal to 3. This illustrates the fact that variant #1 becomes better than variant #2 as soon as using a Padé approximation instead of a real delay for the violation test is not too conservative, which is the case if the order of the approximation is sufficiently high. In addition, since increasing the order has no real impact on computational time, it can be assumed that in the majority of cases, variant #1 will give better results.

Variant	Order	$p(D_m)$ (%)	$p(D_{\bar{m}})$ (%)	$p(D_{m_u})$ (%)	Iterations	CPU time (s)
#1	2	10.98	36.39	2.04	4431	30.7
#1	3	10.98	37.61	0.82	2839	18.1
#1	4	10.98	37.69	0.74	2731	17.5
#1	5	10.98	37.69	0.74	2731	17.6
#2	NA	10.98	37.24	1.19	3303	41.8

TABLE II: Results for $\phi = 0.4s$ using both variants of Algorithm 2 with increasing-order Padé approximations

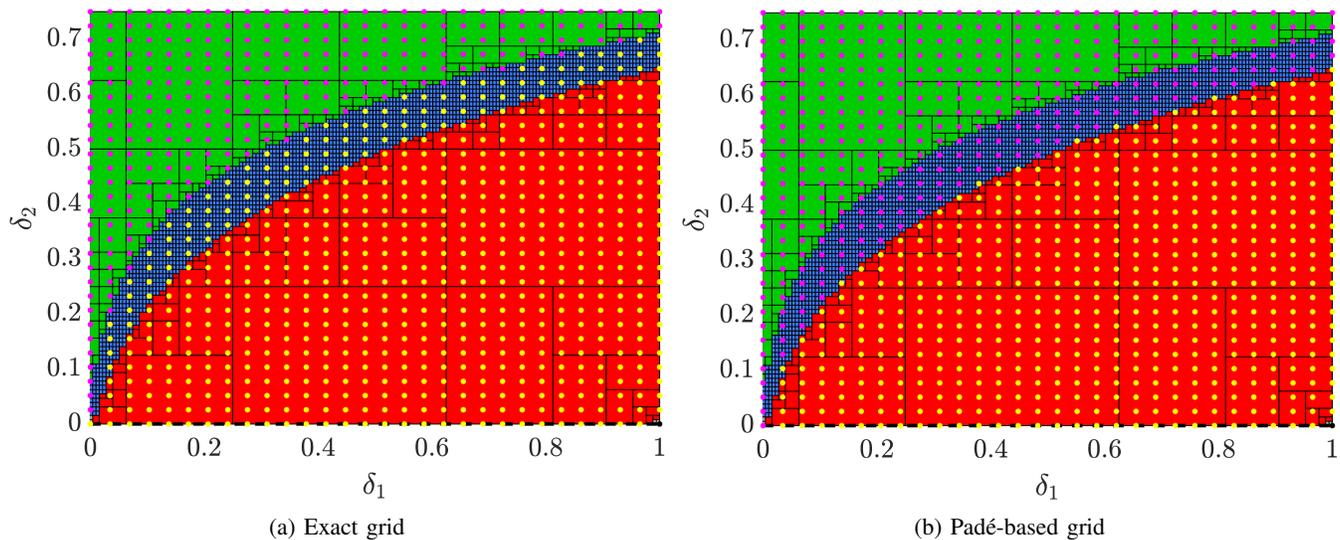


Fig. 9: Comparison for $\phi = 0.4s$ between variant #1 (second-order Padé) and two grid-based methods

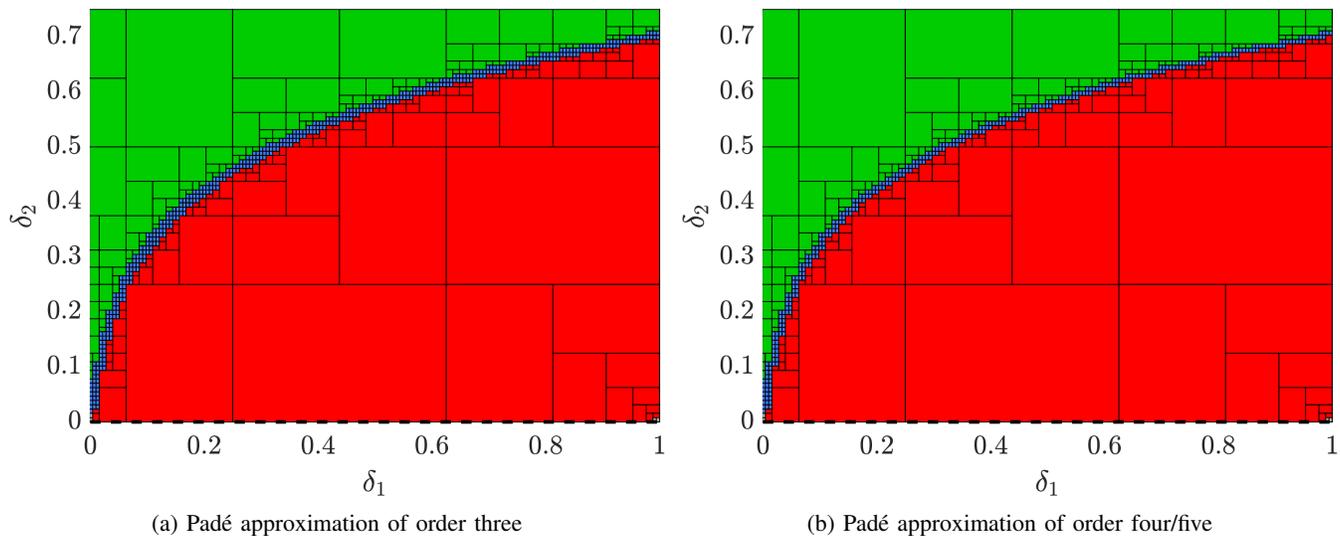


Fig. 10: Delay margin analysis for $\phi = 0.4s$ using variant #1 with increasing-order Padé approximations

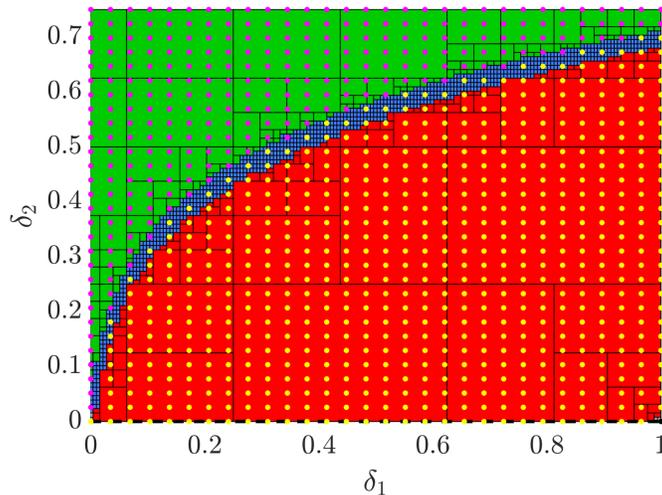


Fig. 11: Comparison for $\phi = 0.4s$ between variant #2 and an exact grid-based method

B. Extensive benchmark analysis

To show the applicability of Algorithm 2 to real-world problems, five models of increasing complexity are analyzed: an academic example, a mass-spring-damper system, a bus steering system, a flexible aircraft and a hard disk drive. All these benchmarks are documented in published papers and implemented in the SMART library of the SMAC toolbox [32]. Table III presents for each of them the input/output structure, the number of states, and the size and structure of the uncertainty block Δ ($n \times p$ means that there are n uncertainties repeated p times each).

Benchmark	Description	Reference	I/O structure	States	Size of Δ	Structure of Δ
1	Academic example	[33]	SISO	4	3	3×1
2	Bus steering system	[34]	SISO	9	5	$1 \times 2 + 1 \times 3$
3	Mass-spring-damper system	[35]	SISO	7	6	6×1
4	Flexible aircraft	[15]	MIMO	46	20	20×1
5	Hard disk drive	[36]	SISO	29	27	$19 \times 1 + 4 \times 2$

TABLE III: Benchmarks used for algorithm assessment

The framework for this study is as follows:

- Benchmarks 4 and 5 are not stable on \mathcal{B}_δ . It can indeed be seen in [14] that the peak value of μ is equal to 4.47 and 1.21 respectively. The focus being here on probabilistic delay margin assessment, the uncertainty domain is shrunk by a factor of 5 (resp. 1.4) for benchmark 4 (resp. 5) to ensure robust stability, *i.e.* $p(D_s) = 100\%$.
- All uncertainties follow a normal distribution of mean value 0 and standard deviation 1/3 truncated on $[-1, 1]$.
- As already highlighted in Section I, probabilistic μ -analysis is particularly suitable to detect rare events. To properly assess the performance of Algorithm 2, the threshold ϕ is therefore set to $1.1\phi_{wc}$, where ϕ_{wc} denotes the worst-case delay margin on \mathcal{B}_δ computed with the SMART Library of the SMAC Toolbox [32]. This ensures that $\overline{P}_{\Delta, f}^\phi(M(s))$ – the probability that the delay margin is smaller than ϕ – is sufficiently low.
- The delay margin is analyzed at the plant input. All benchmarks have a single control input, except the flexible aircraft, which has two. Algorithm 2 being dedicated to the analysis of SISO systems, a loop-at-a-time analysis is performed for benchmark 4. In the sequel, loop 1 (resp. 2) simply means that a delay is introduced in loop 1 (resp. 2).
- The stopping criterion is not given in terms of p_{min} as in Section V-A. Here, the branch-and-bound algorithm is interrupted as soon as the probability of undetermined delay margin $p(D_{m_u})$ becomes lower than $10^{-3}\%$, *i.e.* when $p(D_m) + p(D_{\overline{m}}) \geq 99.999\%$.

All numerical results obtained with variant #1 (resp. variant #2) are included in Table IV (resp. Table V), using in all cases a fifth-order Padé approximation to compute $\hat{\beta}$. CPU_m (resp. $CPU_{\overline{m}}$) denotes the average CPU time per iteration when the delay margin satisfaction (resp. violation) test is performed.

Benchmark (loop)	ϕ (s)	$p(D_m)$ (%)	$p(D_{\bar{m}})$ (%)	Iterations	CPU time (s)	CPU _m (ms)	CPU _{\bar{m}} (ms)
1	0.0407	99.99897	3.8×10^{-5}	396	4.3	11.8	5.4
2	0.0340	99.99827	7.4×10^{-4}	364	8.7	24.9	18.2
3	0.0337	99.99911	0	17	0.7	32.0	NA
4 (1)	0.0315	99.99970	0	21	7.5	278.8	NA
4 (2)	0.0163	99.99931	0	27	11.8	378.0	NA
5	4.051×10^{-5}	99.99905	0	196	95.7	486.2	NA

TABLE IV: Probabilistic delay margin analysis results with variant #1

Benchmark	ϕ (s)	$p(D_m)$ (%)	$p(D_{\bar{m}})$ (%)	Iterations	CPU time (s)	CPU _m (ms)	CPU _{\bar{m}} (ms)
1	0.0407	99.99893	7.2×10^{-5}	459	6.5	11.9	20.6
2	0.0340	99.99846	5.4×10^{-4}	971	33.9	24.8	65.3

TABLE V: Probabilistic delay margin analysis results with variant #2

In all cases, the guaranteed probability $p(D_m)$ of delay margin satisfaction is close to 100%, which confirms that delay margin violation is a rare event. Algorithm 2 is therefore being used properly. The number of iterations differs greatly from one benchmark to another, the general trend being that the more rare an event is, the fewer iterations are needed to obtain a high value of $p(D_m)$. The computational time always remains very low, even for benchmarks with many states and uncertainties, which is a nice result. But it should be kept in mind that it would increase significantly for larger values of ϕ , *i.e.* for larger values of $p(D_{\bar{m}})$, where Monte Carlo simulations would probably be more efficient.

A comparison of the last two columns of Table IV shows that the violation test of Proposition 4.4 is faster than the satisfaction test of Proposition 4.2, which is in line with Remark 4.5. On the opposite, the last two columns of Table V show that the violation test of Proposition 4.6 is slower than the satisfaction test of Proposition 4.2. This is again consistent, since more iterations are required to compute the validity interval $[\underline{\omega}_k, \bar{\omega}_k]$ due to the threshold ϵ_α , as explained in Section IV-C. As a result, it can be observed that both variants of Algorithm 2 have roughly the same accuracy, but as already noticed in Section V-A, variant #2 is slower.

Finally, when the number of uncertainties is large, Algorithm 2 fails to compute a non-zero value of $p(D_{\bar{m}})$, even though there are uncertainties for which the delay margin requirement is not satisfied. The violation test is never called, so both variants of Algorithm 2 give exactly the same results, that is why benchmarks 3-6 do not appear in Table V. By letting the algorithm run longer, it can be observed that $p(D_m)$ tends towards 100%, which shows that the probability of delay margin violation is in fact extremely small. This phenomenon is a combination of two factors. Firstly, the most critical uncertainty combinations are generally located close to the vertices of the uncertainty domain. Secondly, all uncertainties follow a normal distribution here, so the more uncertainties there are, the lower the probability of being close to a vertex.

VI. CONCLUSION

The objective of this paper is to quantify the probability that a given delay margin requirement is violated for a (potentially high-dimensional) uncertain linear system. Moreover, the focus is on detecting rare but nonetheless possible uncertainty combinations, *i.e.* on studying cases where the probability of violation is very low but not zero. A probabilistic approach mixing branch-and-bound and μ -analysis is proposed, and two algorithms are developed. The first one replaces the delay with a Padé approximation. Introduced in [27], it was originally limited to second-order approximations, which sometimes made the analysis very conservative. But thanks to a new constructive method to transform Padé approximations of any order into minimal LFR, a generalization is performed in this paper, allowing to significantly increase accuracy with almost no impact on the computational time. The second contribution is an alternative method based on an equivalent representation of the delay function using a rational function of degree 2 involving a parameter with frequency-dependent bounds. A thorough comparison of both algorithms in terms of conservatism and computational efficiency is performed. The final contribution consists of demonstrating their applicability to realistic problems with a potentially large number of states and uncertainties.

REFERENCES

- [1] 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*, 1996, pp. 3412–3417.
- [2] S. Khatri and P. Parrilo, "Guaranteed bounds for probabilistic μ ," in *Proceedings of the IEEE Conference on Decision and Control*, 1998, pp. 3349–3354.
- [3] G. Balas, P. Seiler, and A. Packard, "Analysis of an UAV flight control system using probabilistic μ ," in *Proceedings of the AIAA AIAA Guidance, Navigation, and Control Conference*, 2012.
- [4] 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*, vol. 62, pp. 120–129, 2021.
- [5] C. Roos, J.-M. Biannic, and H. Evain, "A new step towards the integration of probabilistic μ in the aerospace V&V process," *CEAS Space Journal*, 2023, doi:10.1007/s12567-023-00487-y.
- [6] F. Somers, S. Thai, C. Roos, J.-M. Biannic, S. Bennani, V. Preda, and F. Sanfedino, "Probabilistic gain, phase and disk margins with application to AOCS validation," in *Proceedings of the IFAC Symposium on Robust Control Design*, 2022.

- [7] K. Gu, V. Kharitonov, and J. Chen, *Stability of time-delay systems*. Springer Science + Business Media, 2003.
- [8] S. Niculescu, E. Verriest, L. Dugard, and J. Dion, “Stability and robust stability of time-delay systems: a guided tour,” in *Stability and Control of Time-delay Systems*, L. Dugard and E. Verriest, Eds. Springer, 1998.
- [9] V. Kharitonov, “Robust stability analysis of time delay systems: a survey,” *Annual Reviews in Control*, vol. 23, pp. 185–196, 1999.
- [10] F. Gouaisbaut and D. Peaucelle, “Delay-dependent robust stability of time delay systems,” in *Proceedings of the IFAC Symposium on Robust Control Design*, 2006.
- [11] B. Zhou and A. Egorov, “Razumikhin and Krasovskii stability theorems for time-varying time-delay systems,” *Automatica*, vol. 71, p. 281–291, 2016.
- [12] G. Portilla, I. Alexandrova, and S. Mondié, “Lyapunov–Krasovskii functionals for a class of homogeneous perturbed nonlinear time delay systems,” in *Proceedings of the IEEE Conference on Decision and Control*, 2021, pp. 4743–4748.
- [13] K. Zhou, J. Doyle, and K. Glover, *Robust and optimal control*. Prentice-Hall, 1996.
- [14] 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*, vol. 44, pp. 219–230, 2015.
- [15] G. Ferreres, *A practical approach to robustness analysis with aeronautical applications*. Springer, 1999.
- [16] F. Lescher and C. Roos, “Robust stability of time-delay systems with structured uncertainties: a μ -analysis based algorithm,” in *Proceedings of the IEEE Conference on Decision and Control*, 2011, pp. 4955–4960.
- [17] Y. Huang and K. Zhou, “Robust stability of uncertain time-delay systems,” *IEEE Transactions on Automatic Control*, vol. 45, no. 11, pp. 2169–2173, 2000.
- [18] A. Megretski and A. Rantzer, “System analysis via integral quadratic constraints,” *IEEE Transactions on Automatic Control*, vol. 42, no. 6, pp. 819–830, 1997.
- [19] M. Jun and M. Safonov, “IQC robustness analysis for time-delay systems,” *International Journal of Robust and Nonlinear Control*, vol. 11, no. 15, pp. 1455–1468, 2001.
- [20] M. Fu, H. Li, and S. Niculescu, “Robust stability and stabilization of time-delay systems via integral quadratic constraint approach,” in *Stability and Control of Time-delay Systems*, L. Dugard and E. Verriest, Eds. Springer, 1998.
- [21] J. Zhang, C. Knopse, and P. Tsiotras, “Stability of time-delay systems: equivalence between Lyapunov and scaled small-gain conditions,” *IEEE Transactions on Automatic Control*, vol. 46, no. 3, pp. 482–486, 2001.
- [22] S. Zhang and V. Vittal, “Design of wide-area damping control robust to transmission delay using μ -synthesis approach,” in *Proceedings of the IEEE PES General Meeting Conference & Exposition*, 2014.
- [23] V. Hanta and A. Procházka, “Rational approximation of time delay,” in *Proceedings of the International Conference on Technical Computing*, 2009.
- [24] J. Lam, “Analysis on the Laguerre formula for approximating delay systems,” *IEEE Transactions on Automatic Control*, vol. 39, no. 7, pp. 1517–1521, 1994.
- [25] S. Thai, C. Roos, and J.-M. Biannic, “Probabilistic μ -analysis for stability and \mathcal{H}_∞ performance verification,” in *Proceedings of the American Control Conference*, 2019, pp. 3099–3104.
- [26] F. Somers, C. Roos, F. Sanfedino, S. Bennani, and V. Preda, “Extension of probabilistic gain, phase, disk and delay margins for multi-input, multi-output space control systems,” in *Proceedings of the 12th International Conference on Guidance, Navigation & Control Systems (ESA GNC)*, 2023.
- [27] —, “A μ -analysis based approach to probabilistic delay margin analysis of uncertain linear systems,” in *Proceedings of the IEEE Conference on Control Technology and Applications*, 2023.
- [28] J. Zhang, C. Knopse, and P. Tsiotras, “New results for the analysis of linear systems with time-invariant delays,” *International Journal of Robust and Nonlinear Control*, vol. 13, no. 12, pp. 1149–1175, 2003.
- [29] P. Young, M. Newlin, and J. Doyle, “Computing bounds for the mixed μ problem,” *International Journal of Robust and Nonlinear Control*, vol. 5, no. 6, pp. 573–590, 1995.
- [30] C. Roos and J.-M. Biannic, “Efficient computation of a guaranteed stability domain for a high-order parameter dependent plant,” in *Proceedings of the American Control Conference*, 2010, pp. 3895–3900.
- [31] M. Newlin and P. Young, “Mixed μ problems and branch and bound techniques,” *International Journal of Robust and Nonlinear Control*, vol. 7, pp. 145–164, 1997.
- [32] 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*, 2013, pp. 176–181, available with the SMAC Toolbox at w3.onera.fr/smac/smart.
- [33] R. De Gaston and M. Safonov, “Exact calculation of the multiloop stability margin,” *IEEE Transactions on Automatic Control*, vol. 33, no. 2, p. 156–171, 1988.
- [34] J. Ackermann and W. Sielén, “Robust control for automatic steering,” in *Proceedings of the American Control Conference*, 1990, p. 795–800.
- [35] —, “What is a large number of parameters in robust systems?” in *Proceedings of the IEEE Conference on Decision and Control*, 1990, pp. 3496–3497.
- [36] D. Gu, P. Petkov, and M. Konstantinov, *Robust control design with Matlab*. Springer-Verlag, 2005.

APPENDIX PROOF OF PROPOSITION 4.1

Dividing $\Psi_n(\tau s)$ by $(\tau s)^n$ in the numerator and denominator leads to:

$$\Psi_n(\tau s) = \frac{Y(s)}{U(s)} = \frac{\sum_{i=0}^n (-1)^i a_i (\tau s)^{i-n}}{\sum_{i=0}^n a_i (\tau s)^{i-n}} \quad (28)$$

Let us now take the term $i = n$ out of the sum:

$$Y(s) = (-1)^n U(s) + \left(\sum_{i=0}^{n-1} (-1)^i \frac{a_i}{a_n} (\tau s)^{i-n} \right) U(s) - \left(\sum_{i=0}^{n-1} \frac{a_i}{a_n} (\tau s)^{i-n} \right) Y(s) \quad (29)$$

Then the term $i = n - 1$:

$$Y(s) = (-1)^n \left(U(s) - \frac{a_{n-1}}{a_n} (\tau s)^{-1} \left(U(s) + (-1)^n Y(s) + \left(\sum_{i=0}^{n-2} (-1)^{i+n-1} \frac{a_i}{a_{n-1}} (\tau s)^{i-n+1} \right) U(s) - \left(\sum_{i=0}^{n-2} (-1)^{n-1} \frac{a_i}{a_{n-1}} (\tau s)^{i-n+1} \right) Y(s) \right) \right) \quad (30)$$

And so on with all the terms:

$$Y(s) = (-1)^n \left(U(s) - \frac{a_{n-1}}{a_n} (\tau s)^{-1} \left(U(s) + (-1)^n Y(s) - \frac{a_{n-2}}{a_{n-1}} (\tau s)^{-1} \left(U(s) + (-1)^{n-1} Y(s) - \frac{a_{n-3}}{a_{n-2}} (\tau s)^{-1} \left(U(s) + (-1)^n Y(s) - \dots - \frac{a_0}{a_1} (\tau s)^{-1} \left(U(s) - Y(s) \right) \dots \right) \right) \right) \right) \quad (31)$$

The next step is to transform equation (31) into an LFR $\mathcal{F}_u(\tilde{P}_n(s), \tau^{-1} I_n)$:

$$\begin{cases} \dot{x} &= \tilde{A}x + [\tilde{B}_1 \ \tilde{B}_2] \begin{bmatrix} w \\ u \end{bmatrix} \\ \begin{bmatrix} z \\ y \end{bmatrix} &= \begin{bmatrix} \tilde{C}_1 \\ \tilde{C}_2 \end{bmatrix} x + \begin{bmatrix} \tilde{D}_{11} & \tilde{D}_{12} \\ \tilde{D}_{21} & \tilde{D}_{22} \end{bmatrix} \begin{bmatrix} w \\ u \end{bmatrix} \\ w &= \tau^{-1} I_n z \end{cases} \quad (32)$$

To do so, the inputs $w \in \mathbb{R}^n$, outputs $z \in \mathbb{R}^n$ and states $x \in \mathbb{R}^n$ are progressively introduced in (31). Firstly:

$$y = (-1)^n (u - w_1) \quad (33)$$

$$w_1 = \tau^{-1} z_1 \quad (34)$$

$$z_1 = \frac{a_{n-1}}{a_n} x_1 \quad (35)$$

$$\dot{x}_1 = u + (-1)^n y - w_2 = 2u - w_1 - w_2 \quad (36)$$

By continuing in this way, the following relationships are then established:

$$w_i = \tau^{-1} z_i \quad (37)$$

$$z_i = \frac{a_{n-i}}{a_{n-i+1}} x_i \quad (38)$$

$$\dot{x}_i = 2u - w_1 - w_{i+1} \quad \text{for } i = 1, 3, \dots \quad (39)$$

$$= w_1 - w_{i+1} \quad \text{for } i = 2, 4, \dots \quad (40)$$

$$= (1 + (-1)^{n+1}) u + (-1)^n w_1 \quad \text{for } i = n \quad (41)$$

and the state-space matrices of equation (32) are finally obtained as follows:

$$\tilde{A} = 0_{n \times n}, \quad [\tilde{B}_1 \ \tilde{B}_2] = \begin{bmatrix} -1 & -1 & 0 & \dots & 0 & \vdots & 2 \\ 1 & 0 & -1 & \ddots & \vdots & \vdots & 0 \\ \vdots & \vdots & \ddots & \ddots & 0 & \vdots & \vdots \\ \vdots & \vdots & & \ddots & -1 & \vdots & \vdots \\ (-1)^n & 0 & \dots & \dots & 0 & \vdots & 1 + (-1)^{n+1} \end{bmatrix} \quad (42)$$

$$\begin{bmatrix} \tilde{C}_1 \\ \tilde{C}_2 \end{bmatrix} = \begin{bmatrix} \frac{a_{n-1}}{a_n} & 0 & \dots & 0 \\ 0 & \frac{a_{n-2}}{a_{n-1}} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & \frac{a_0}{a_1} \\ \hline 0 & \dots & \dots & 0 \end{bmatrix}, \quad \begin{bmatrix} \tilde{D}_{11} & \tilde{D}_{12} \\ \tilde{D}_{21} & \tilde{D}_{22} \end{bmatrix} = \begin{bmatrix} 0 & \dots & \dots & 0 & \vdots & 0 \\ \vdots & & & \vdots & \vdots & \vdots \\ \vdots & & & \vdots & \vdots & \vdots \\ 0 & \dots & \dots & 0 & \vdots & 0 \\ \hline (-1)^{n+1} & 0 & \dots & 0 & \vdots & (-1)^n \end{bmatrix} \quad (43)$$

The next step is to set $\tau^{-1}I_n = \mathcal{F}_u(\tilde{T}_n, \beta I_n)$, where $\tilde{T}_n = \begin{bmatrix} -I_n & \frac{2}{\phi}I_n \\ -I_n & \frac{2}{\phi}I_n \end{bmatrix}$. In other words, $\beta = -1 + 2\tau/\phi$ is normalized in the sense that $\beta \in]-1, 1[\Leftrightarrow \tau \in]0, \phi[$. Then $\Psi_n(\tau s) = \mathcal{F}_u(\tilde{P}_n(s), \tau^{-1}I_n) = \mathcal{F}_u(\tilde{P}_n(s), \mathcal{F}_u(\tilde{T}_n, \beta I_n)) = \mathcal{F}_u(\mathcal{S}(\tilde{T}_n, \tilde{P}_n(s)), \beta I_n)$, where $\mathcal{S}(\cdot, \cdot)$ denotes the Redheffer star product defined in Section 10.4 of [13]. Noting that $\tilde{A} = 0$, $\tilde{C}_2 = 0$, $\tilde{D}_{11} = 0$ and $\tilde{D}_{12} = 0$, a state-space representation of $P_n(s) = \mathcal{S}(\tilde{T}_n, \tilde{P}_n(s))$ is given by:

$$A_P = \frac{2}{\phi} \tilde{B}_1 \tilde{C}_1 \quad , \quad B_P = \begin{bmatrix} -\tilde{B}_1 & \tilde{B}_2 \end{bmatrix} \quad , \quad C_P = \frac{2}{\phi} \begin{bmatrix} \tilde{C}_1 \\ \tilde{D}_{21} \tilde{C}_1 \end{bmatrix} \quad , \quad D_P = \begin{bmatrix} -I_n & 0 \\ -\tilde{D}_{21} & \tilde{D}_{22} \end{bmatrix} \quad (44)$$

Combining (42), (43) and (44) finally leads to equations (12)-(15). Moreover, the LFR $\mathcal{F}_u(P_n(s), \beta I_n)$ is minimal. Indeed, $\Psi_n(\tau s)$ is a transfer function of degree n , so $P_n(s)$ cannot have less than n states. And τ^n appears in both the numerator and denominator of $\Psi_n(\tau s)$ in (11), so β (which is an affine function of τ) cannot be repeated less than n times. ■