Set propagation in dynamical systems with generalised polynomial algebra and its computational complexity

Massimiliano Vasile[1]

Department of Mechanical and Aerospace Engineering, University of Strathclyde, 75 Montrose Street, G1 1XJ Glasgow, United Kingdom

Carlos Ortega Absil

Department of Mechanical and Aerospace Engineering, University of Strathclyde, 75 Montrose Street, G1 1XJ Glasgow, United Kingdom

Annalisa Riccardi

Department of Mechanical and Aerospace Engineering, University of Strathclyde, 75 Montrose Street, G1 1XJ Glasgow, United Kingdom

Abstract

This paper presents an approach to propagate sets of initial conditions and model parameters through dynamical systems. It is assumed that the dynamics is dependent on a number of model parameters and that the state of the system evolves from some initial conditions. Both model parameters and initial conditions vary within a set Ω. The paper presents an approach to approximate the set Ω with a polynomial expansion and to propagate, under some regularity assumptions, the polynomial representation through the dynamical system. The approach is based on a generalised polynomial algebra that replaces algebraic operators between real numbers with operators between polynomials. The paper first introduces the concept of generalised polynomial algebra and its use to propagate sets through dynamical systems. Then it analyses, both theoretically and experimentally, its time complexity and compares it against the time complexity of a non-intrusive counterpart.

Email addresses: `massimiliano.vasile@strath.ac.uk` (Massimiliano Vasile),
`carlos.ortega@strath.ac.uk` (Carlos Ortega Absil),
`annalisa.riccardi@strath.ac.uk` (Annalisa Riccardi)

[1]Corresponding author

Finally, the paper provides an empirical convergence analysis on two illustrative examples of linear and non-linear dynamical systems.

---

1. Introduction

When studying dynamical systems it is often interesting to analyse the evolution of extended regions of the state space or to evaluate the sensitivity to the variation of some model parameters. In the specific case of orbital mechanics the interest could be to study the evolution of a set of particles (dust or debris), or a set of initial conditions for a single object. This problem is of particular relevance in the assessment of conjunctions and collisions.

A common approach is to use massive Monte Carlo simulations but in recent times alternative approaches based on Polynomial Chaos Expansion, Jones et al. (2013), Chebyshev expansions, Tardioli et al. (2015), Gauss Mixture Models, De Mars and Jah (2013), and High Dimensional Model Representations Tardioli et al. (2015) have been proven to be a valid alternative to Monte Carlo simulations. These methods can be collectively classified as non intrusive because they do not require any access to the dynamic equations and treat the problem as a black box.

Another class of approaches that has gained popularity as a valid alternative to direct Monte Carlo simulations is based on a Taylor series expansion of the dynamics, see Park and Scheeres (2006). In this case the propagation can be achieved by introducing an algebra, on the space of Taylor polynomials, that replaces the standard computer algebra on real numbers. A wide range of applications of Taylor polynomial algebra in orbital mechanics can be found in the work of Di Lizia et al. (2008), Di Lizia et al. (2014) and Armellin et al. (2010) and in the work of Jorba and Zou (2005). These techniques are based on the so called Truncated Power Series Algebra (TPSA) introduced by Berz (1986, 1987) for the computation of transfer maps in particle optics and extended to rigorous numerics in 1997 with the introduction of Taylor Models, see Berz (1997).

In the case of Taylor series, the basis functions are monomials defining the i^{th}-order variation with respect to a reference point in the state space $\mathbf{x}_0$. The residual error is proportional to the $(i+1)^{th}$ derivative and the $(i+1)^{th}$-order monomial. Numerical integration and evaluation of the dynamics are

performed in the TPSA, hence at each integration step the full polynomial representation of the current state is available.

The same idea can be generalised to a different set of basis functions provided that the corresponding algebraic rules between their basis can be defined. The idea of a more generic TPSA dates back to Epstein et al. (1982) with the name of Ultra Arithmetic. However, in 2003 the possibility of using a TPSA alternative to the one built on Taylor basis was discarded by Makino and Berz (2003) because of several drawbacks related to polynomial multiplication and growth of the magnitude of the coefficients. It was with the work of Brisebarre and Joldes (2010) that a comparison of TPSAs on Taylor, Chebyshev and Newton bases was formalised in more general terms. The results proved that, for the univariate case, the approximations derived from algebrae defined on Chebyshev bases had smaller reminders than Taylor models, albeit requiring more computational time for the same order of expansion. One of the main advantages of using Chebyshev series expansions is the uniform convergence over the interval of expansion for Lipschitz continuous functions. In fact, the series converges also when a finite number of discontinuities in f and its derivatives are present J.C. Mason (2002); Amparo Gil (2008); Trefethen (2013). On the contrary, Taylor series require the function f to be $n+1$ times differentiable in a neighborhood of the expansion point and, even when they have a positive radius of convergence, convergence pointwise to the function f and uniformly on every compact subset of the convergence neighborhood. Furthermore, when Taylor series converge, they might not converge to the correct function f except at the expansion point George B. Arfken (2013).

This paper extends the work of Brisebarre and Joldes (2010) to the multivariate case and to generic polynomial expansions. The paper provides a rigorous investigation of the time complexity of a TPSA based on polynomial bases, other than Taylor, in comparison to its non-intrusive counterpart.

The extension to generic polynomial expansions is of particular interest in the case in which the set Ω, over which initial conditions and model parameters take values, is not a hypercube and is not defined with respect to a reference solution $\mathbf{x}_0$. The latter point is particularly important as a reference solution might not be available or might not be unique (for example in the case of sets of particles).

The methodology proposed hereafter takes any generic polynomial representation of a set, converts the polynomials in monomials and applies an algebra on the resulting monomial basis expansion. In this paper, the ini-

tial polynomial representation is based on Chebyshev expansions on finite intervals. The complexity and accuracy of the proposed Generalised Polynomial Algebra (GPA) are compared to a non-intrusive counterpart, based on Chebyshev multivariate interpolation of a set of propagated states.

The next two sections present the details of the GPA and how it is used to propagate sets in dynamical systems. A non-intrusive technique is then introduced followed by a theoretical analysis of the computational complexity of both the GPA and the non-intrusive counterpart. The methods considered are then applied to the propagation of uncertainties in a scalable dynamical system, in order to test their limits and compare their performance. Experimental tests are used to derive an empirical convergence analysis. Finally, application to space dynamics is illustrated by four instances of the two-body problem in low-Earth orbit, considering uncertainties both on states and model parameters. Some final remarks on the future development and applications of GPA methodology are given in the conclusions.

Implementation by the authors of all methods discussed in the paper can be accessed at https://github.com/strath-ace/smart-uq within the C++ library SMART-UQ (the Strathclyde Mechanical and Aerospace Research Toolbox for Uncertainty Quantification).

2. Set Propagation in Dynamical Systems

Consider the dynamical system:

$$\begin{cases} \dot{\mathbf{x}} & = f(\mathbf{x}, \mathbf{b}, \mathbf{t}) \\ \mathbf{x}(t_0) & = \mathbf{x_0} \end{cases}, \tag{1}$$

where $\mathbf{b} \in \Upsilon \subseteq \mathbb{R}^q$ is a vector of model parameters. $t \in T \subseteq \mathbb{R}$ and the initial conditions have value $\mathbf{x_0} \in \Sigma_0 \subseteq \mathbb{R}^c$ so that $d = q + c$.

The goal is to propagate set $\Omega = \Upsilon \times \Sigma_0$ through dynamical system (1) from time $t = t_0$ to $t = t_f$ and obtain the set Σ_t, at time t_f, defined as:

$$\Sigma_t = \left\{ \mathbf{x}_f | \mathbf{x}_f = \mathbf{x}_0 + \int_{t_0}^{t_f} f(\mathbf{x}, \mathbf{b}, \tau) d\tau, \forall \mathbf{b} \in \Upsilon \wedge \mathbf{x_0} \in \Sigma_0 \right\} \tag{2}$$

under the assumption that the solution of system (1) exists and is unique. Since an exact representation of Σ_t would require the propagation of an infinite number of sample vectors $(\mathbf{x}_0, \mathbf{b})$, if Σ_t is a compact set and $\mathbf{x}_f$ is continuous in $\mathbf{x}_0$ and $\mathbf{b}$, we can approximate Σ_t with:

$$\tilde{\Sigma}_t = \left\{ \mathbf{x}_f | \mathbf{x}_f = P_{n,d}(\mathbf{x}_0, \mathbf{b}), \forall \mathbf{b} \in \Upsilon \wedge \mathbf{x_0} \in \Sigma_0 \right\} \tag{3}$$

where $P_{n,d}$ is a polynomial of degree (order) n in d dimensions (variables) (see Eq.(7)).

The polynomial $P_{n,d}$ can be constructed in a number of ways. In this paper we present two computationally different approaches: a propagation of a polynomials representation of Ω via generalised polynomial algebra, and an interpolation of a limited number of propagate samples $\mathbf{x}_{f,k}$ with Chebyshev multivariate polynomials.

3. Generalised Polynomial Algebra

Given a continuous, piece-wise differentiable, function $f(\mathbf{x}) : \Omega \subset \mathbb{R}^d \to \mathbb{R}$, we consider the approximation

$$f(\mathbf{x}) = P(\mathbf{x}) + r(\varepsilon) = \sum_{\mathbf{i},|\mathbf{i}|\leq n} p_{\mathbf{i}}\alpha_{\mathbf{i}}(\mathbf{x}) + r(\varepsilon)\,, \tag{4}$$

where, without loss of generality, $\Omega = [-1,1]^d$, $\mathbf{x} \in \Omega$, $\mathbf{i} \in [0,n]^d \subset \mathbb{N}^d$, $|\mathbf{i}| = \sum_{r=1}^{d} i_r$, $r(\varepsilon)$ is a remainder (with $\varepsilon \in \Omega$), and $\alpha_{\mathbf{i}}(\mathbf{x})$ is a polynomial basis of choice, up to order n. The number of coefficients for a complete expansion is given by

$$\mathcal{N}_{d,n} = \binom{n+d}{d} = \frac{(n+d)!}{n!d!}\,, \tag{5}$$

The polynomial $P(\mathbf{x})$ belongs to the function space $\mathcal{P}_{n,d}(\alpha_i)$ of polynomials of order n in d dimensions, in the α_i basis. The definition of the polynomials can be extended to a generic hyper-rectangle $\overline{\Omega} = [\mathbf{a},\mathbf{b}] \subseteq \mathbb{R}^d$; being $\tau : \overline{\Omega} \to \Omega$ the linear mapping between the two regions, the generalised expansions are defined over $\overline{\Omega}$ by

$$\alpha_{\mathbf{i}}(\mathbf{x}) = \alpha_{\mathbf{i}}(\tau(\mathbf{x}'))\,, \tag{6}$$

where $\mathbf{x}' \in \overline{\Omega}$. So without loss of generality the domain Ω is considered hereafter.

3.1. The monomial basis

A polynomial of degree n in d dimensions with coefficients $c_{\mathbf{i}}$ in a field $\mathcal{K}$ can be written as the sum of multi-degree monomials as:

$$P_{n,d}(\mathbf{x}) = \sum_{\mathbf{i},|\mathbf{i}|\leq n} c_{\mathbf{i}}\mathbf{x}^{\mathbf{i}} \tag{7}$$

with the monomial $\mathbf{x^i}$ defined as:

$$\mathbf{x^i} = \prod_j^d x_j^{i_j} \tag{8}$$

and $\mathbf{i} = (i_1, i_2, ..., i_j, ..., i_d)$ a multi-index. When it is possible to represent $P_{n,d}(\mathbf{x})$ in monomial form, one can represent Ω in the preferred polynomial basis and then rewrite the resulting $P_{n,d}(\mathbf{x})$ with an equivalent monomial expansion. Although, in the literature, the translation into monomials was found to be a ill-conditioned operation, see Trefethen (2013), and can lead to large coefficients, it has two main advantages: it significantly reduces the computational cost and allows one to maintain the core algebraic operations unchanged regardless of the polynomial representation of f. As it will be explained in the remainder of the paper, since the expansion is truncated after the transformation in monomials, one has to account for an additional approximation error that derives from the truncation. Nonetheless we showed that the resulting methods provides a fast and accurate propagation in the cases investigated in this paper. Thus in the remainder of this paper all polynomials are assumed to be written in monomial form and the resulting algebra is defined on the space of the monomial basis. We define this change of basis with the notation:

$$\nu \; : \; \alpha_i \longrightarrow \phi_i \tag{9}$$

where ϕ is the monomial basis.

The idea is to expand in the desired polynomial basis, convert the resulting polynomial in monomial basis and then use an algebra over the space of monomial basis to propagate the polynomial through dynamical system (1).

3.2. Polynomial Algebra Over $\mathcal{P}_{n,d}(\phi_i)$

The function space $\mathcal{P}_{n,d}(\phi_i)$ can be equipped with a set of elementary arithmetic operations, generating an algebra on the space of polynomials such that, given two functions f_A and f_B, and two polynomials A and B, which approximate f_A, and respectively f_B,

$$f_A(\mathbf{x}) \oplus f_B(\mathbf{x}) \sim A(\mathbf{x}) \otimes B(\mathbf{x}), \tag{10}$$

where $\oplus \in \{+, -, \cdot, /\}$ and $\otimes$ is the corresponding operation in $\mathcal{P}_{n,d}(\phi_i)$. This allows one to define the algebra $(\mathcal{P}_{n,d}(\phi_i), \otimes)$, of dimension $\dim(\mathcal{P}_{n,d}(\phi_i), \otimes) = \mathcal{N}_{d,n}$, the elements of which belong to the polynomial ring in d indeterminates $\mathbb{R}[\mathbf{x}]$ and have degree up to n. Each element $P(\mathbf{x})$ of the algebra, is uniquely identified by the set of its coefficients $\mathbf{c} = \{c_{\mathbf{i}} \; : \; |\mathbf{i}| \le n\} \in \mathbb{R}^{\mathcal{N}_{d,n}}$.

3.2.1. Addition and Subtraction

The operations of addition and subtraction are defined as follows: being $A(\mathbf{x})$ and $B(\mathbf{x})$ two elements of $(\mathcal{P}_{n,d}(\phi_i), \otimes)$, identified by the set of coefficients $\mathbf{a}, \mathbf{b} \in \mathbb{R}^{\mathcal{N}_{d,n}}$, respectively, the result of their sum or difference is

$$C(\mathbf{x}) = A(\mathbf{x}) \pm B(\mathbf{x}), \tag{11}$$

identified by the set of coefficients $\mathbf{c} \in \mathbb{R}^{\mathcal{N}_{d,n}}$ such that

$$\mathbf{c} = \mathbf{a} \pm \mathbf{b}. \tag{12}$$

Since the operations of addition and subtraction do not change the degree of the polynomials if one applies a truncation to a given order the result is $C_n(\alpha(\mathbf{x})) = C_n(\phi(\mathbf{x})) = C(\mathbf{x})$, with $C_n(\alpha(\cdot))$ and $C_n(\phi(\cdot))$ the polynomials in the original polynomial basis and in the monomial basis, respectively, truncated to order n.

3.2.2. Multiplication

The product is defined as

$$A(\mathbf{x}) \cdot B(\mathbf{x}) = \left(\sum_{\mathbf{i}, |\mathbf{i}| \leq n} a_{\mathbf{i}} \mathbf{x}^{\mathbf{i}} \right) \left(\sum_{\mathbf{i}, |\mathbf{i}| \leq n} b_{\mathbf{i}} \mathbf{x}^{\mathbf{i}} \right). \tag{13}$$

where the product of two monomials is:

$$\mathbf{x}^{\mathbf{i}} \cdot \mathbf{x}^{\mathbf{j}} = \begin{cases} \mathbf{x}^{\mathbf{i}+\mathbf{j}} & \text{if } |\mathbf{i}+\mathbf{j}| \leq n \\ 0 & \text{otherwise} \end{cases}, \tag{14}$$

so that the product polynomial is truncated to order n.

3.2.3. Composition

Any polynomial $G(\mathbf{z})$ in the functional space $\mathcal{P}_{\nu,\delta}(\phi_i)$ can be interpreted as a finite set of additions and multiplications. Hence, with these operators as defined for $(\mathcal{P}_{n,d}(\phi_i), \otimes)$, it is straightforward to define a composition rule

$$\circ : \mathcal{P}_{\nu,\delta}(\phi_i) \times [\mathcal{P}_{n,d}(\phi_i)]^{\delta} \longrightarrow \mathcal{P}_{n,d}(\phi_i) \tag{15}$$

such that, if $G(\mathbf{z})$ approximates a $\delta-$variate function $g(\mathbf{y}(\mathbf{x})) = g(\mathbf{z})$ in $\mathcal{P}_{\nu,\delta}(\phi_i)$, then

$$g(\mathbf{y}(\mathbf{x})) \sim G(\mathbf{z}) \circ \mathbf{Y}(\mathbf{x}), \tag{16}$$

7

where $\mathbf{y}(\mathbf{x})$ is a $\delta-$array of $d-$dimensional sets and $\mathbf{Y}(\mathbf{x})$ their polynomial expansions in $(\mathcal{P}_{n,d}(\phi_i), \otimes)$.

Furthermore, the composition rule is used to define in $(\mathcal{P}_{n,d}(\phi_i), \otimes)$ a division operator, as well as the counterpart of any elementary function, as follows: being $h(z)$ any of the functions $\{1/z, \sin(z), \cos(z), \exp(z), \log(z), ...\}$, $H(z)$ an univariate polynomial of degree ν approximating $h(z)$, $y(\mathbf{x})$ a multivariate set and $Y(\mathbf{x})$ its polynomial expansion in the algebra, the $h-$mapping of y is approximated by

$$h(y(\mathbf{x})) \sim H(z) \circ Y(\mathbf{x}) \,. \tag{17}$$

In this case

$$\circ \,:\, \mathcal{P}_{n,1}(\phi_i) \times \mathcal{P}_{n,d}(\phi_i) \longrightarrow \mathcal{P}_{n,d}(\phi_i) \,. \tag{18}$$

3.3. Range Estimation and Truncation Error

The means by which $H(z) \sim h(z)$ is computed will affect the capability to propagate a set with given properties, leading to different polynomial algebras. For instance, in Taylor Algebra, $H(z)$ is the order-n Maclaurin expansion of $h(z - Y(\mathbf{x} = 0))$, so the $h-$mapping preserves the local approximation around the central point. If a polynomial algebra implements a global approximation over an interval, or is based on finite-support polynomials, the generation of $H(z)$ requires information on $h(z)$, not only in the vicinity of a point, but on a set of values taken by $Y(\mathbf{x})$ in $\bar{\Omega}$. In other words, it is necessary to bound the range of $Y(\mathbf{x})$ by computing

$$\underline{y} \leq \min_{\mathbf{x} \in \bar{\Omega}} Y(\mathbf{x}) \quad, \quad \bar{y} \geq \max_{\mathbf{x} \in \bar{\Omega}} Y(\mathbf{x}) \tag{19}$$

prior to approximating $h(z)$ over the interval $I = \left[\underline{y}, \bar{y}\right]$. To be noted that a rigorous use of Taylor expansions would require the estimation of the radius of convergence.

In the GPA approach implemented in this work, the authors approximated the infinite Chebyshev series of $h(z)$ with a Chebyshev interpolation $H_T(z)$ of order 100 over I. In the rest of the discussion, this approximation will be obviated for the sake of simplicity, as if $H_T(z)$ were precisely the infinite series. Then $H(z) \in \mathcal{P}_{n,1}(\phi_i)$ is obtained by truncation of $H_T(z)$, in the Chebyshev basis, up to order $\bar{n} \geq n$. The choice of $\bar{n}$ depends on the range-bounding method. Note that the order-$\bar{n}$ Chebyshev interpolation and the Chebyshev series truncated to order $\bar{n}$ constitute distinct operators, and

thus lead to different polynomial expansions of $h(z)$ over I. Both are near-minimax order-$\bar{n}$ polynomial approximations, with an identical expression for the upper bound on their remainder Rivlin (1974). However, the latter is preferred for small values of $\bar{n}$, since it has been experimentally found to lead to smoother approximations in the applications considered in this paper.

Computing the range of a multivariate polynomial in a hyperrectangle is a central problem in numerical analysis. If the polynomial is not linear, finding tight bounds in (19) is an NP-hard problem because it requires finding the global maximum and minimum of the polynomial over $\hat{\Omega}$, see Nesterov (2000). Hence a range estimate is sought based on known upper (respectively lower) bounds on the maximum (respectively minimum) value of the polynomial. This implies that

$$\left[\underline{y}^*, \bar{y}^*\right] = Range(Y(\mathbf{x})) \subseteq I = \left[\underline{y}, \bar{y}\right] \tag{20}$$

often with proper inclusion. This phenomenon is referred to as *range overestimation*, and can constitute an important source of approximation error in the propagation, especially when $(\bar{y}^* - \underline{y}^*)/(\bar{y} - \underline{y}) \ll 1$. Indeed $H(z)$, approximating $h(z)$ over I, might account for values of the elementary function that are actually out of domain. These superfluous variations generate noise, thus lowering the approximation quality within the true domain.

In the GPA approach proposed in this paper we employed the following range estimation:

$$\bar{y} = c_{\mathbf{0}} + \sum_{\mathbf{i} \neq \mathbf{0}} |c_{\mathbf{i}}| \tag{21}$$

$$\underline{y} = c_{\mathbf{0}} - \sum_{\mathbf{i} \notin Q} |c_{\mathbf{i}}|$$

where $Y(\mathbf{x}) = \sum_{\mathbf{i}} c_{\mathbf{i}} \mathbf{x}^{\mathbf{i}}$ and $Q = \{\mathbf{i} \,|\, i_j \equiv 0 \,(\mathrm{mod}\,2)\,,\ j = 1, \ldots, d\}$. More precise polynomial range-bounding methods exist, e.g. Garloff (1985) presents two approaches for the multivariate case. These have not been explored due to their higher computational complexity. The bounds in (21) are of fast evaluation, but might lead to some range overestimation. In order to mitigate this effect, the approximation $H(z)$ is obtained by Chebyshev-basis truncation of $H_T(z)$ up to an order of expansion

$$\bar{n} := \lfloor 1.5n + 1 \rfloor . \tag{22}$$

The resulting approximation is then converted into monomial basis, and truncated to order n for its expression in $(\mathcal{P}_{n,1}(\phi_i), \otimes)$. This choice of $\bar{n}$ has been found to yield a decrease in maximum absolute error of the final state expansions of up to one order of magnitude in the numerical experiments presented in Section 6 , with respect to direct usage of the order-n Chebyshev expansion. This is counter-intuitive, as the order-n Chebyshev expansion is a near-minimax approximation in $\mathcal{P}_{n,1}(\phi_i)$ and is generally more precise over $[\underline{y}^*, \bar{y}^*]$ than the approximation obtained as suggested above. Nonetheless, it is shown in the subsequent of this section that this procedure can improve the approximation over I and thus mitigate the error induced by range over-estimation by (21), at a lower computational cost than more precise range bounding. In consequence, A lower $\bar{n} \geq n$ is recommended if a more precise range-bounding approach is used, or small deviations from linearity are expected. On the other hand, progressively increasing $\bar{n}$ within the propagation can enhance the approximation quality for long propagations through non-linear systems.

It is now interesting to study the remainder of the expansion $H(z)$. Let us assume $H(z) \in \mathcal{P}_{n,1}(\phi_i)$ is an expansion truncated at order n in the monomial basis. We call $H(z; \bar{n} = n) \in \mathcal{P}_{n,1}(\phi_i)$ the near- minimax order-n Chebyshev expansion, whereas $H(z; \bar{n} > n) \in \mathcal{P}_{n,1}(\phi_i)$ is the monomial-basis truncation, up to order n, of an order-$\bar{n}$ Chebyshev expansion, obtained as suggested above. The respective remainders in a point of the domain are bounded by

$$r_n = |h(z) - H(z; \bar{n} = n)| \leq \frac{1}{2^n (n+1)!} |h^{(n+1)}(\xi_n)| \tag{23}$$

$$r_{\bar{n}} = |h(z) - H(z; \bar{n} > n)| \leq \underbrace{\frac{1}{2^{\bar{n}} (\bar{n}+1)!} |h^{(\bar{n}+1)}(\xi_{\bar{n}})|}_{r_{\bar{n}}^{approx.}} + \underbrace{|\hat{z}^{n+1}| |\Xi(\hat{z})|}_{r_{\bar{n}}^{trunc.}}$$

where $\hat{z} = 2(z - \underline{y})/(\bar{y} - \underline{y}) - 1$; $\xi_n, \xi_{\bar{n}} \in [\underline{y}, \bar{y}]$; $\Xi(\hat{z})$ is a polynomial of degree $\bar{n} - n - 1$; and bracketed superscripts indicate order of derivative with respect to $\hat{z}$. Furthermore, the coefficients of $\Xi(\hat{z})$ are the coefficients that have been truncated to express $H(\hat{z}; \bar{n} > n)$ in $(\mathcal{P}_{n,1}(\phi_i), \otimes)$. Since $\hat{z} \in [\hat{\underline{z}}, \hat{\bar{z}}] \subseteq [-1, 1]$, one can in turn use these coefficients to bound $|\Xi(\hat{z})|$, e.g. by (21), and obtain an upper bound on the truncation-related term $r_{\bar{n}}^{trunc}$ that is independent on $\hat{z}$. With this in mind, it is possible to state that:

- In the vicinity of the midpoint $z_0 = (\underline{y} + \bar{y})/2$, the remainder $r_{\bar{n}}$ equals
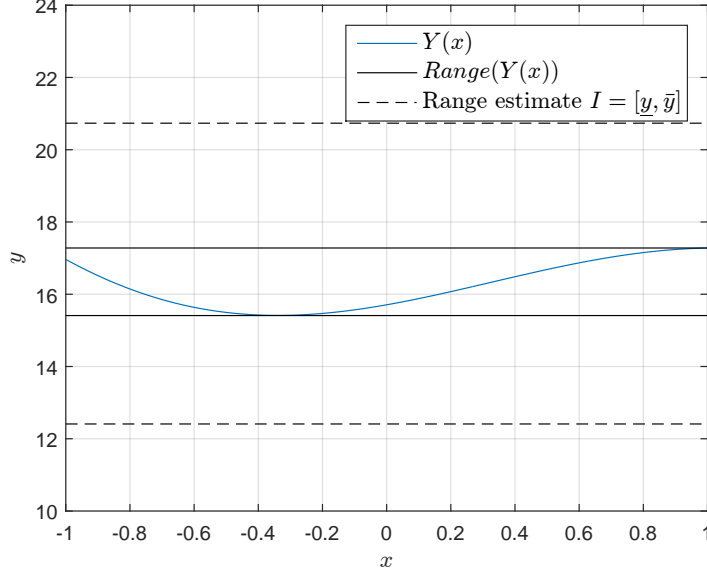
10

Figure 1: Range overestimation on $Y(x) = 5\pi(1 + 0.1x + 0.1x^2 - 0.1x^3 - 0.01x^4 + 0.01x^5)$.

the remainder of an order-$\bar{n}$ Chebyshev approximation, $r_{\bar{n}}^{approx.}$. With $\bar{n} > n$, it is expected that $r_{\bar{n}}^{approx} \leq r_n$.

- The bounds on $r_{\bar{n}}$ will grow towards the edges of the interval, as $r_{\bar{n}}^{trunc}$ is affected by a factor $|\hat{z}^{n+1}|$.

These effects are exemplified in Figures 1 to 3 with elementary function $\sin(z)$ as $h(z)$, and a degree-5 polynomial as $Y(x)$. The approximations $H(z; \bar{n} = 5) \circ Y(x)$ and $H(z; \bar{n} = 8) \circ Y(x)$ are evaluated in $(\mathcal{P}_{5,1}(\phi_i), \otimes)$. Figure 1 shows the magnitude of the range overestimation by (21) in this case. Figure 2 illustrates how the truncation order (22) affects the error distribution in the approximation of $h(z)$; the error increases at the extrema of I, but decreases in the central region, notably within the true range of $Y(x)$. Figure 3 compares the final approximations. Note that $H(z; \bar{n} = 5)$ presents small error towards the extrema of I, whereas $H(z; \bar{n} = 5) \circ Y(x)$ does not. This is due to the fact that yet an additional truncation error is created during the composition, when evaluating the powers of $Y(x)$ in $(\mathcal{P}_{5,1}(\phi_i), \otimes)$ by (14).
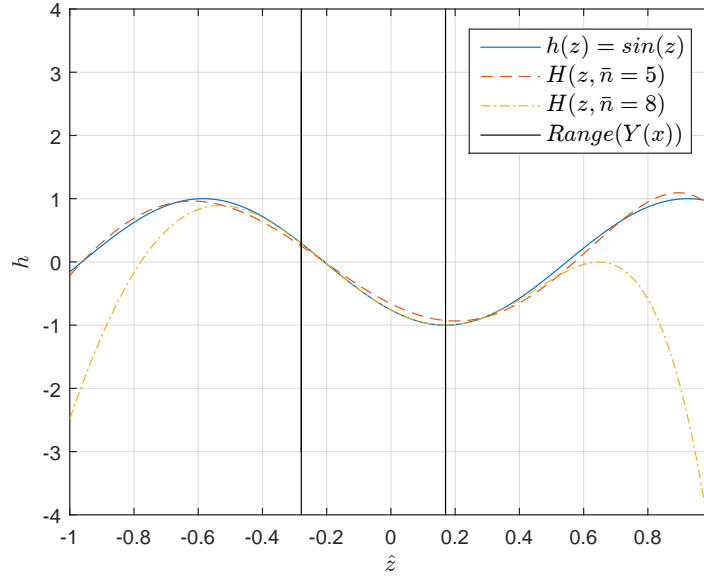
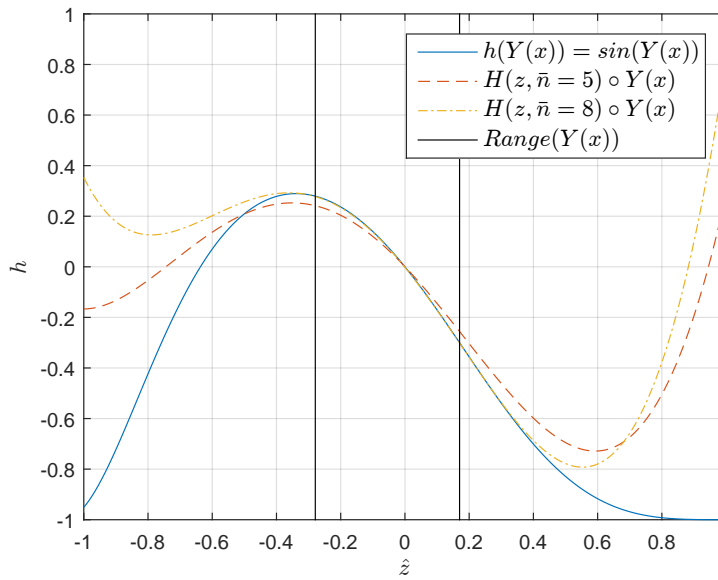Figure 2: Comparison of order-5 approximations of $sin(z)$ over $I = [\underline{y}, \bar{y}]$.

Figure 3: Comparison of order-5 approximations of $sin(Y(x))$ over $I = [\underline{y}, \bar{y}]$.

3.4. Propagation Through Dynamical Systems

Consider now a numerical scheme ψ used to integrate system (1) and the map:

$$\mathbf{x}_k \leftarrow \psi(\mathbf{x}_{k-1}, \mathbf{b}) \tag{24}$$

The idea is to represent $\mathbf{x}_{k-1}$ and $\mathbf{b}$ as elements of the algebra $(\mathcal{P}_{n,d}(\phi_i), \otimes)$ and applying all algebraic operation defined in $(\mathcal{P}_{n,d}(\phi_i), \otimes)$ to the integration scheme ψ. Thus we replace real numbers with polynomials of degree n, replace all elementary functions in (1) with their polynomial representations up to degree n and apply all algebraic operations in map (24) to the elements of $\mathcal{P}_{n,d}(\phi_i)$.

As an example, consider the simple Euler scheme:

$$\mathbf{x}_k = \mathbf{x}_{k-1} + f(\mathbf{x}_{k-1}, \mathbf{b})\Delta t \tag{25}$$

and represent $\mathbf{x}_0 \simeq \mathbf{X}_0 := (X_1(\mathbf{x}), ..., X_c(\mathbf{x}))$ and $\mathbf{b} \simeq \mathbf{B} := (B_1(\mathbf{x}), .., B_q(\mathbf{x}))$ as elements of the algebra $(\mathcal{P}_{n,d}(\phi_i), \otimes)$. With this scheme, at each integration step one has:

$$\mathbf{X}_k = \mathbf{X}_{k-1} + \Delta t \, \mathbf{F}_{k-1}, \tag{26}$$

where $\mathbf{F}_{k-1}$ is the polynomial approximation of $f(\mathbf{x}_{k-1}, \mathbf{b})$, obtained by expanding all elementary functions in f in Chebyshev series of $\mathbf{x}_{k-1}$ and $\mathbf{b}$ and evaluating the resulting expansion by replacing $\mathbf{x}_{k-1}$ and $\mathbf{b}$ with $\mathbf{X}_{k-1}$ and $\mathbf{B}$ and applying all the algebraic operations defined in previous sections. Hence $\mathbf{X}_k$ is the polynomial representation of the system flow at the k^{th} time-step.

To be noted that the nature of the set to be propagate or the elementary functions to be represented can suggest a different polynomial representation. For example, for unbounded sets, Hermite or Taylor polynomials might be more appropriate. Nonetheless, the scheme proposed in this paper remains unchanged, as well as the main contributions to the associated computational complexity. Hence the use of the term generalised polynomial algebra (or GPA).

4. Interpolation via Multivariate Chebyshev Polynomials

The non-intrusive counterpart of a GPA is the approximation of a generic multivariate function $f(\mathbf{x})$ by means of interpolation or regression techniques

13

in the chosen basis. Hence, as before, given a function $f(\mathbf{x}) : \Omega \subset \mathbb{R}^d \to \mathbb{R}$, its polynomial approximation on the functional space $\mathcal{P}_{n,d}(\alpha_i)$ is defined as

$$f(\mathbf{x}) = P(\mathbf{x}) + r(\varepsilon) = \sum_{\mathbf{i},|\mathbf{i}|\leq n} p_{\mathbf{i}}\alpha_{\mathbf{i}}(\mathbf{x}) + r(\varepsilon), \tag{27}$$

where Ω is a generic hypercube and $r(\varepsilon)$ is the reminder (with $\varepsilon \in \Omega$). The coefficients $p_{\mathbf{i}}$ are computed via interpolation or regression. This means that the space Ω is sampled and the function f evaluated on each of the sample points. The coefficients are then computed by solving the linear system

$$HP = F, \tag{28}$$

where

$$H = \begin{bmatrix} \alpha_1(\mathbf{x}_1) & \dots & \alpha_{N_K}(\mathbf{x}_1) \\ \vdots & \ddots & \vdots \\ \alpha_1(\mathbf{x}_{N_p}) & \dots & \alpha_{N_K}(\mathbf{x}_{N_p}) \end{bmatrix}, \quad P = \begin{bmatrix} p_1 \\ \vdots \\ p_{N_K} \end{bmatrix}, \quad F = \begin{bmatrix} f_1 \\ \vdots \\ f_{N_p} \end{bmatrix},$$

$N_p = |\mathcal{S}|$ is the cardinality of the set of sample points $\{\mathbf{x}_1, \dots, \mathbf{x}_{N_p}\}$ and the components of F are the evaluations of the model at the i-th sample point, $f_i = f(\mathbf{x}_i)$. If a sample is used such that $N_p = N_K$ and H is non-singular, it is possible interpolate in the sample. A regression approach can be used with $N_p > N_K$. In this paper we restricted to interpolation, using two sampling schemes. One we called *full basis*, for it solves for the coefficients of the full Chebyshev basis of order n and dimension d. As regards the sampling, it takes $N_P = N_K = \mathcal{N}_{d,n}$ random samples in a Latin Hypersquare pattern McKay et al. (1979) , Iman et al. (1980). The other, that we called *reduced basis*, uses isotropic Smolyak sparse grids Smolyak (1963). The sparse grids are generated following the Clenshaw-Curtis rule Gerstner and Griebel (1998). The extrema of the univariate Chebyshev polynomials are used as suggested and described in Judd et al. (2014). The sample consists of a subset of the tensor product of these extrema. The grids are fully nested, i.e. the grid of level l contains all points of the grid of level $l-1$. The corresponding reduced-basis polynomial is of degree 2^l, but with $N_P = N_K < \mathcal{N}_{d,2^l}$. A number of the coefficients is excluded from the basis and assumed equal to zero. The reduced basis tends to exclude cross-product terms under the assumption that higher-order correlations are negligible.

Consider now dynamical system (1), the approximation of the expansion at time t_N can be computed by means of non-intrusive technique by solving the

linear system (28) where H is the evaluation of the multivariate polynomial base at the points sampled in $\Omega = \Upsilon \times \Sigma_0$ and the right hand side of the equation contains the values of the terminal states obtained integrating the dynamics with initial conditions and model parameters taken from the set of sample points. If the state vector to be approximated in the dynamical system has dimension $c = \dim(\Sigma_0)$, then the right hand side of equation (28) is a matrix of dimension $N_p \times c$. Hence the approximation of the states is equivalent to the solution of c linear systems. Note that it is not always necessary to obtain polynomial approximations of and only of the states in the state vector. Only a subset of the propagated states might be of interest. Or it might be useful to approximate a quantity that is not a state of (1), but a derived expression. The set of propagated quantities requiring approximation by (28) is referred to as the *quantities of interest*.

5. Computational complexity analysis

The goal of this section is to derive an estimation of the computational cost of both the intrusive and the non-intrusive approach to decide whether, for a given accuracy, one is preferable to the other (assuming that both are applicable). The computational cost measured in terms of floating-point operations, will be estimated by means of complexity analysis of the subroutines involved in each methodology.

The computational complexity of propagation with GPA depends on the size of the algebra, $\mathcal{N}_{d,n}$, the computational complexity of the operations necessary for the evaluation of the right hand side of the dynamical system and the computational complexity of the numerical integrator. On the other hand, the complexity of the non-intrusive method is proportional to the number of points in the sample, N_p, the computational complexity of the numerical integrator and the solution of a $N_p \times N_K$ linear system whenever a polynomial representation needs to be obtained.

A more complete and rigorous analysis of the computational complexity of each method is given in the following subsections.

5.1. Computational Complexity of GPA

The computational complexity of the GPA can be derived as the complexity of each elementary algebraic operation times the total number of operations.

5.1.1. Complexity of addition and multiplication

Addition or subtraction, of two polynomials, as well as multiplication or division by a scalar value, require exactly $\mathcal{N}_{d,n}$ floating-point operations, where $\mathcal{N}_{d,n}$ is given by the expression in Eq. (5). Their computational complexity is, therefore, $O(\mathcal{N}_{d,n})$. The multiplication of two polynomials is a more expensive operation and an upper limit on the number of floating-point operations is given by the following lemma.

Lemma 1. *Given the multiplication defined in (13), the upper limit on the number of floating-point operations is given by:*

$$F_{d,n}^{*} = 2\mathcal{N}_{2d,n}, \tag{29}$$

Proof. A constructive proof of this lemma can be obtained by considering the order truncation rule applied during multiplication:

$$\mathbf{x^{i}} \cdot \mathbf{x^{j}} = \begin{cases} \mathbf{x^{i+j}} & \text{if } |\mathbf{i+j}| \leq n \\ 0 & \text{otherwise} \end{cases} \tag{30}$$

and rewriting the truncation condition $|\mathbf{i} + \mathbf{j}| \leq n$ as

$$i_1 + i_2 + \dots + i_d + j_1 + j_2 + \dots + j_d \leq n$$

which, introducing the variable $n' = n - |\mathbf{i+j}|$, $0 \leq n' \leq n$, can be expressed without loss of generality as a Diophantine equation in $2d + 1$ variables, namely

$$n' + i_1 + i_2 + \dots + i_d + j_1 + j_2 + \dots + j_d = n$$

whose number of non-negative solutions N_c equals the number of n-multisubsets of a set of cardinality $2d + 1$, i.e. Brualdi (1977)

$$N_c = \left(\!\!\binom{2d+1}{n}\!\!\right) = \binom{n+2d}{n} = \binom{n+2d}{2d} = \mathcal{N}_{2d,n}.$$

Each solution corresponds to a term of polynomial p_1 being multiplied by a term of polynomial p_2 and actually contributing by addition to a term of the product polynomial $p_1 p_2$, therefore

$$F_{d,n}^{*} = 2N_c = 2\mathcal{N}_{2d,n} \tag{31}$$
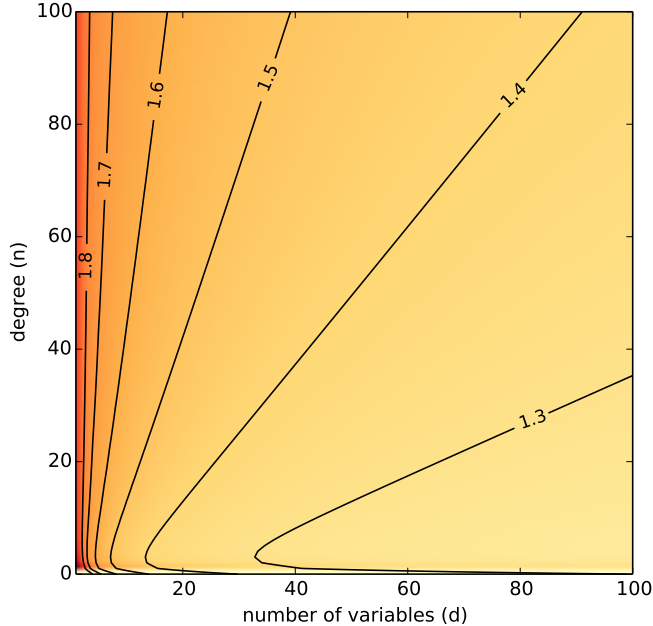
$$\square$$

Figure 4: Contour plot of $log_{\mathcal{N}_{d,n}}(F^*_{d,n})$ for different combinations of d and n

Note this lemma is obtained assuming precomputed indices for the multiplication operations, i.e. omitting the cost of the operation $\mathbf{i} + \mathbf{j}$. A map $(\mathbf{i}, \mathbf{j}) \rightarrow \mathbf{i} + \mathbf{j}$ is obtained and stored upon instantiation of the polynomial algebra.

The result implies that it is not possible to obtain a generally-valid expression for the complexity as a function of $\mathcal{N}_{d,n}$ uniquely; due to the dependence on $\mathcal{N}_{2d,n}$, equivalent-size algebras with different combinations of d and n will present different multiplication cost, as is shown in Figure 4. It can be derived from the result in Eq. (31) that best-case complexity, as measured with respect to $\mathcal{N}_{d,n}$, occurs for low degree and high dimension. In the extreme case of $n = 1$ (linear algebra in $\mathcal{N}_{d,n} - 1$ variables), the cost of the multiplication is linear in $\mathcal{N}_{d,n}$:

$$F^*_{d,n=1} = F^*_{\mathcal{N}_{d,n}-1,1} = 4\mathcal{N}_{d,n} - 2 = O(\mathcal{N}_{d,n}) \tag{32}$$

while worst-case complexity occurs when $d = 1$ (univariate algebra of degree

$\mathcal{N}_{d,n} - 1$), yielding

$$F^*_{d=1,n} = F^*_{1,\mathcal{N}_{d,n}-1} = \mathcal{N}_{d,n}(\mathcal{N}_{d,n} + 1) = O(\mathcal{N}^2_{d,n}) \tag{33}$$

Hence we can say that the multiplication algorithm has overall complexity $F^*_{d,n} = O(\mathcal{N}^p_{d,n})$, $1 \le p \le 2$..

Additionally the following result can be demonstrated.

Lemma 2. *Given the multiplication defined in (13), an upper limit on the number of floating-point operations is given by:*

$$F^*_{d,n} \le 2^{n+1}\mathcal{N}_{d,n}, \quad \forall d, n \in \mathbb{N}, \tag{34}$$

Proof. This lemma can be proven by induction on n. In fact, for $n = 0$ one has:

$$F^*_{d,0} = 2\mathcal{N}_{2d,0} = 2\frac{(2d + 0)!}{(2d)!0!} = 2 = 2\frac{(d + 0)!}{d!0!} = 2^{0+1}\mathcal{N}_{d,0}$$

While for a generic n one has that:

$$F^*_{d,n+1} = 2\frac{(2d + n + 1)!}{(2d)!(n + 1)!} = 2\frac{2d + n + 1}{n + 1}\frac{(2d + n)!}{(2d)!n!} = \frac{2d + n + 1}{n + 1}F^*_{d,n};$$

$$2^{n+2}\mathcal{N}_{d,n+1} = 2^{n+2}\frac{(d + n + 1)!}{d!(n + 1)!} = 2\frac{d + n + 1}{n + 1}\cdot2^{n+1}\frac{(d + n)!}{d!n!} = \frac{2d + 2n + 2}{n + 1}2^{n+1}\mathcal{N}_{d,n};$$

Since

$$\frac{2d + n + 1}{n + 1} \le \frac{2d + 2n + 2}{n + 1} \quad \forall d, n \in \mathbb{N}$$

It follows that

$$F^*_{d,n} \le 2^{n+1}\mathcal{N}_{d,n} \Rightarrow F^*_{d,n+1} \le 2^{(n+1)+1}\mathcal{N}_{d,n+1}$$

Which, together with the proof for $n = 0$, gives that

$$F^*_{d,n} \le 2^{n+1}\mathcal{N}_{d,n}, \quad \forall d, n \in \mathbb{N}.$$

$\square$

Lemma 2 allows one to express the overall complexity alternatively as

$$F^*_{d,n} = O(2^{n+1}\mathcal{N}_{d,n}),$$

18

a result that depicts very accurately the behaviour for low n and $d \gg n$, becoming more conservative as n grows.

As a final remark, one might wish to avoid precomputing the indices, for low-memory applications, and perform the operation $\mathbf{i} + \mathbf{j}$ for each contribution to $p_1 p_2$. In this case, the complexity of the multiplication is affected by an additional factor $O(d)$, in exchange of a memory save of $\mathcal{N}_{2d,n}$ integers. Computing indices has a complexity of $O(2d\mathcal{N}_{2d,n})$ too, but is performed only once per instantiation of the algebra.

5.1.2. Complexity of Composition and Elementary Functions

The composition $H(z) \circ p(\mathbf{x})$ of $p \in (\mathcal{P}_{d,n}(\alpha_i), \otimes)$ with an unidimensional expansion $H(z)$ has complexity $O(n \cdot F_{d,n}^*)$, where $F_{d,n}^*$ is the complexity of multiplying two polynomials $p_1, p_2 \in (\mathcal{P}_{d,n}(\alpha_i), \otimes)$.

Note that the application $h(p(\mathbf{x}))$ of an elementary function $h(z)$ to an expansion $p \in (\mathcal{P}_{d,n}(\alpha_i), \otimes)$ is a two-step process involving the approximation $H(z) \sim h(z)$ and the composition $H(z) \circ p(\mathbf{x})$. The complexity associated to the approximation step depends on the approximation method of choice. In the GPA implementation proposed hereby, where a Chebyshev expansion of $h(z)$ is obtained over the range of $p(\mathbf{x})$, the approximation complexity is $O(\mathcal{N}_{d,n} + n^3)$. The main contribution $O(\mathcal{N}_{d,n})$ corresponds to the estimation of the range of $p(\mathbf{x})$ by (21). The $O(n^3)$ contribution corresponds to the change of basis applied to $H(z)$ in order to manipulate in monomials; this requires evaluating a monomial basis the expression

$$H(z) = \sum_{i=0}^{n} c_i^H \cdot T_i(z) \,,$$

where c_i^H are the coefficients of the Chebyshev expansion of $H(z)$ and $T_i(z)$ is the univariate Chebyshev basis evaluated in z, i.e. over the range of $p(\mathbf{x})$. This evaluation is linear in terms of polynomial multiplications. Thus the change of basis can be done in complexity $O(n \cdot F_{1,n}) = O(n^3)$. In Taylor Algebra, the approximation complexity is lower, in general $O(n)$ for elementary functions. Since all of these are dominated by the complexity of the composition, one can state that the whole process $h(p(\mathbf{x}))$ has asymptotic complexity $O(n \cdot F_{d,n}^*)$ in both algebras. Nonetheless, the cost of approximation is found to cause a small but perceptible difference in run-time in moderately high-dimensional problems in Ortega et al. (2016). Note also that, should composition be defined between Chebyshev and monomial expansions, using, for instance,

19

Clenshaw's algorithm Clenshaw (1955), it would not be necessary to apply a change of basis to $H(z)$, so the contribution $O(n^3)$ could be avoided. This option has not been explored in this implementation. However, the expected gain in performance is small in high-dimensional problems where $\mathcal{N}_{d,n} \gg n^3$.

Division of two polynomials is treated as the application of the elementary function $inv(z) = z^{-1}$, followed by multiplication.

The composition of a polynomial $p \in (\mathcal{P}_{d,n}(\alpha_i), \otimes)$ with an array of polynomials $\mathbf{q} \in (\mathcal{P}_{d,n}(\alpha_i), \otimes)^d$ has complexity $O(d\mathcal{N}_{d,n}F_{d,n}^*)$. It is, therefore, the asymptotically most expensive operation in an algebra of size $\mathcal{N}_{d,n}$.

5.1.3. Total Complexity of GPA

Consider η the propagation of a set of states of dimension d through a dynamical system and assume the use of a numerical integration scheme that performs m integration steps. Every time the numerical integrator evaluates the right hand side of the dynamical system it performs F^{RHS} operations. An estimation of the total computational complexity of GPA applied to the propagation of a set of states is derived, assuming that all the F^{RHS} operations are multiplications among elementary functions.

Theorem 3. *Given the Truncated Polynomial Algebra defined in Section (3) the total complexity for m steps of integration is:*

$$F_I^\eta \sim (d + mF^{RHS})F_{d,n}^* = 2(d + mF^{RHS})\mathcal{N}_{2d,n} \qquad (35)$$

where, the additional d is the cost of the precomputation of indices.

Proof. Direct from the preceding results. $\qquad \square$

Note that a more accurate cost estimation for the intrusive approach is obtained, if the operation count F^{RHS} is conducted so that it ignores additions, subtractions and multiplications by scalars, and accounts for n operations for each division and elementary function application. Nevertheless, for the sake of simplicity in the comparison with the non-intrusive approach, equation 35 will be considered an acceptable estimation hereafter.

5.2. Computational Complexity of Chebyshev Multivariate Interpolation

The non-intrusive Chebyshev multivariate interpolation is a polynomial representation containing a subset of N_K non-null coefficients of the truncated Chebyshev series of order n in d variables. This representation is built

taking $N_p = N_K$ samples belonging to Ω, and the process of sampling has complexity $O(d \cdot N_K)$. Each sample is then propagated forward in time, thus, the total cost of the propagation of all the samples is

$$F^{prop.} = m N_K F^{RHS}$$

where F^{RHS}, as for GPA, is the cost of the evaluation of the right hand side of the dynamical system and m is the number of integration steps.

In order to obtain a polynomial representation of the propagated states, s linear systems need to be solved, with s the number of states that need representation, i.e. the quantities of interest. Each linear system has N_K equations in N_K unknowns and is in the form:

$$\mathrm{H} \mathbf{c}_{j,k} = \mathbf{f}_{j,k} \tag{36}$$

where $\mathbf{c}_{j,k}$ is the vector of coefficients of the polynomial representation of state f_j at the k^{th} time-step, $\mathbf{f}_{j,k}$ is the vector containing all N_K values propagated for state f_j at the k^{th} time-step, and H is the $N_K \times N_K$ matrix obtained from evaluation of the polynomial basis in the sample, not depending on j or k. Assuming solution by LU decomposition, the total cost is that of:

- Computing H, with complexity $O(N_K^2 d)$.

- Decomposing H, with complexity $O(\frac{2}{3} N_K^3)$.

- Solving all the equation systems using the LU decomposition of H, with complexity $O(N_K^2 s)$.

Therefore, ignoring lesser contributions, the overall complexity of the non-intrusive approach is given by:

$$F_{NI}(d, n) \sim m N_K F^{RHS} + \frac{2}{3} N_K^3 \tag{37}$$

Note that the computational complexity is not driven by the sampling method or polynomial basis, but only by the cardinality of the sample and number of coefficients of the basis, these being equal in the case of interpolation. The choice of one or another sampling method and polynomial basis will however impact the accuracy obtained for a given cost.

5.3. Theoretical Comparison

This subsection will provide a theoretical comparison between the computational cost of the intrusive and the non-intrusive approaches, in order to decide when one approach has to be preferred to the other. The assumption is that for the same size of the algebra, the two methods provide roughly the same accuracy of the final polynomial representation. Hence this comparison will assume $N_K^{NI} \sim \mathcal{N}_{d,n}^I$, where N_K^{NI} is the number of coefficients of the non-intrusive interpolation, and $\mathcal{N}_{d,n}^I$ is the size of the intrusive algebra. In section 6 we will present empirical convergence analysis that will verify the hypothesis on the accuracy for two examples of dynamical systems.
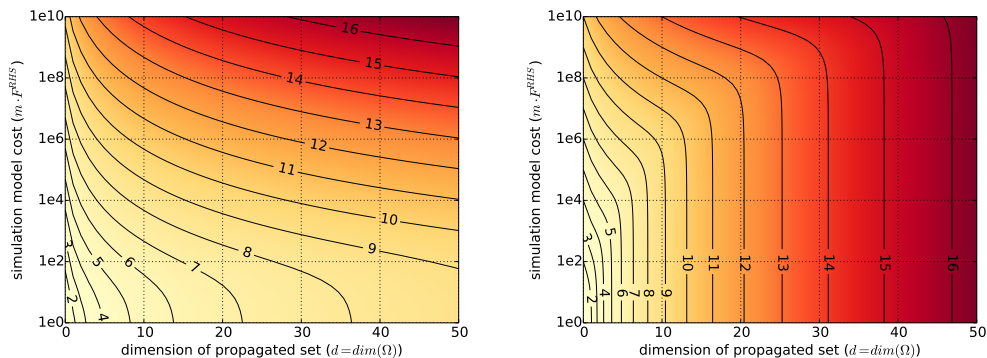
Figure 5: Contour plots of $log_{10}(F_I^\eta)$ (left) and $log_{10}(F_{NI}^\eta)$ (right) for an entire uncertainty propagation application. The degree of the intrusive algebra is set to $n = 4$. For the non-intrusive case, it is assumed that $N_K \sim \mathcal{N}_{d,n=4}$.

Under this assumption, the orders of magnitude of F_I^η and F_{NI}^η are shown in Figure 5 for an intrusive algebra of degree $n = 4$. Note that in the non-intrusive case the contour lines have a vertical trend, indicating that the cost presents very fast growth with the number of uncertain variables, but is barely sensitive to an increase in the cost of the simulation model, mF^{RHS}. This trend changes where mF^{RHS} is large enough as to overcome the effect of the term of complexity $O(N_K^3)$, associated with the solution of the equation systems and independent of the simulation cost. On the other hand, with GPA the sensitivity to an increase in d is remarkably smaller due to the lower complexity with respect to the size of the algebra that derives from the lack of model-training. Nevertheless, for a fixed number of variables, the overall cost grows faster with the cost of the simulation model.
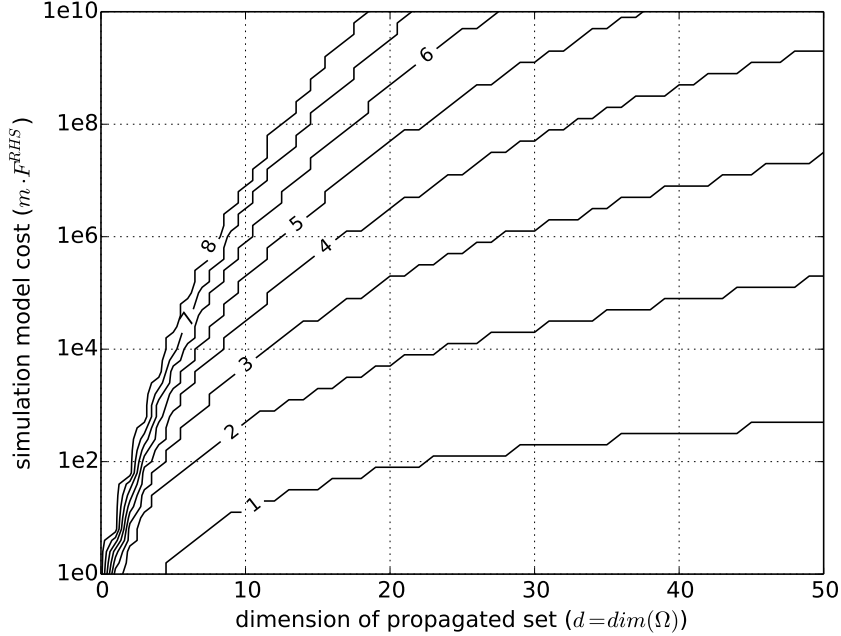
22

Figure 6: Computational cost cross-over lines between intrusive and non-intrusive approaches discussed. Each line corresponds to a degree of the polynomial algebra n. Under the line, the intrusive approach of degree n is less expensive than any non-intrusive approach with similar size of the algebra ($N_K \sim \mathcal{N}_{d,n}$). Note that no information on absolute cost is extracted from this figure.

Therefore, assuming that the same size of the algebra is required to both approaches to reach the desired level of accuracy, it will always be less expensive to use the intrusive method on high-dimensional problems with low simulation model cost. On the other hand, it will always be preferable to use interpolation on a propagated sample for low-dimensional problems with high simulation model cost, provided both methods are applicable. Figure 6 depicts the cross-over lines delimiting the regions of the plane $< d, mF^{RHS} >$ where it is cheaper to use one or another approach, for degrees 1 to 8 of the intrusive approach. For instance, the cross-over line of $n = 3$ intersects $d = 30$ at $mF^{RHS} \sim 10^6$. This indicates that, if the cost of a single simulation mF^{RHS} is lower than 10^6 floating-point operations, then intrusive propagation in 30 variables and degree 3 will be less expensive than any non-intrusive interpolation with equivalent size of the algebra, $N_K^{NI} \sim \mathcal{N}_{d=30,n=3}^I = 5456$.

23

Furthermore, this stands $\forall d \geq 30$.

6. Numerical Experiments

In this section the intrusive and non-intrusive approaches are put to the test on two problems: a multi-spring-mass system and an orbital dynamics problem. These problems are used to test the scalability of both the intrusive and the non-intrusive methods with the number of dimensions.

The goal is to test the ability of both methods to provide an accurate representation of a given set propagate over a given length of time. In all numerical tests, the initial set is defined as follows: given a generic parameter vector $\mathbf{p}$ with nominal value $\bar{\mathbf{p}}$ and range of variability of its components defined by $\mathbf{u}_p$, $\mathbf{p}$ is assumed to take any value in the box $[\bar{\mathbf{p}} - \mathbf{u}_p, \bar{\mathbf{p}} + \mathbf{u}_p]$.

The metrics we use to compare the two methods are the CPU time and the accuracy of the polynomial approximations. In all cases, the experiments have been conducted with double precision on a 2.8GHz machine. Unless stated otherwise, the following test conditions are used:

- For GPA, multiplication with precomputed indices has been used. Runtime measured includes all precomputations necessary for instantiation of the algebra, allocation of the polynomials, propagation and storage of the expansions.

- For the non-intrusive method, a sample of size equal to the size of the algebra, $N_p = N_K$, has been used for training. Runtime measured includes sampling, construction and decomposition of the interpolation matrix, propagation of the sample and interpolation of the expansions. Sampling technique is detailed in the experiments section.

- Intrusive and non-intrusive methods are compared against direct propagation of a validation sample. The Root Mean Square Error (RMSE)

$$ RMSE = \sqrt{\frac{1}{N_s} \sum_{i=1}^{N_s} (\hat{x}_i - x_i)^2} \quad , $$

  is used as error metric to compare the overall accuracy of the polynomial representations, where x_i is the true value of a state obtained by

24

forward integration of a sample point, $\hat{x}_i$ is the approximate value computed evaluating the polynomial expansion obtained with the approach under discussion, and N_s is the cardinality of the validation sample $\mathcal{V}$.

For both the GPA and Chebyshev interpolation, convergence will be analysed empirically, as a function of the size of the algebra for a given dimensionality of Ω. Besides $RMSE$, accuracy results regarding convergence studies are also presented in terms of the maximum absolute error,

$$\text{Max. abs. err.} = \max_{\mathcal{V}}(|\hat{x}_i - x_i|) \quad .$$

The validation sample $\mathcal{V}$ is generated with a Latin Hypersquare. A cardinality of $N_s = 10^4$ is used for all experiments, after confirming stability of the error metric with respect to the sample.

The results are presented in the following subsections.

6.1. Linear Dynamical Systems

For this numerical experiment a non-damped spring-mass system with L masses is considered. The L masses can move in a one-dimensional space subject to the force of multiple springs. The second-order dynamical equations governing the motion of the L masses can be rewritten as a first-order system with $2L$ states as follows:

$$\mathbf{M}\frac{d\mathbf{y}}{dt} + \mathbf{K}\mathbf{y} = 0 \tag{38}$$

with initial conditions at time t_0:

$$\begin{aligned} x_1(t_0) &= x_{1,0} \\ x_i(t_0) &= x_{i,0} \\ x_L(t_0) &= x_{L,0} \\ v_i(t_0) &= v_{i,0} \end{aligned} \quad , \tag{39}$$

for $i = 1, ..., L$ and

$$\mathbf{y} = \begin{bmatrix} x_1 \\ \vdots \\ x_L \\ v_1 \\ \vdots \\ v_L \end{bmatrix} , \quad \mathbf{M} = \left[\begin{array}{c|c} \begin{matrix} 1 & & \\ & \ddots & \\ & & 1 \end{matrix} & \\ \hline & \begin{matrix} m_1 & & \\ & \ddots & \\ & & m_L \end{matrix} \end{array}\right] ,$$

$$
\mathbf{K} = \left[
\begin{array}{cccc|ccc}
 & & & & -1 & & \\
 & & & & & \ddots & \\
 & & & & & & -1 \\
\hline
k_0 + k_1 & -k_1 & & & & & \\
-k_1 & k_1 + k_2 & \ddots & & & & \\
 & \ddots & \ddots & -k_{L-1} & & & \\
 & & -k_{L-1} & k_{L-1} + k_L & & &
\end{array}
\right],
$$

where, for $1 \leq i \leq L$, the mass, position and velocity of particle i are noted m_i, x_i, and v_i respectively, and k_{i-1}, k_i are the stiffness constants of the springs before and after particle i.

This test case is selected to analyse how the performance of a method scales with dimensionality, as isolated from other effects as possible. The equations are linear with respect to the state vector y, but the test case requires non-linear analysis as uncertainty in the model parameters is considered.

6.1.1. Experimental Setup

Simulations were run for multiple values of L, namely $1 \leq L \leq 13$. In all simulations, the following set of nominal parameters and initial conditions were used:

$$
\bar{\mathbf{p}} = [\bar{x}_{1,0}, \bar{x}_{L,0}, \bar{x}_{i,0}, \bar{v}_{i,0}, \bar{m}_i, \bar{k}_i]^T = [-0.25, 0.25.0.0, 0.0, 1.0, 1.0]^T
$$

with the exception that for $L = 1$, $\bar{x}_{1,0} = \bar{x}_{L,0} = 0.25$. Note that here and in the following test case we use an overbar to indicate quantities within the set $\hat{\Omega}$ around which one would expand the dynamics in Taylor series if one was using a Taylor algebra. We call these values, nominal, because in the context of uncertainty propagation one would consider these values as not affected by any uncertainty. In order to consider nonlinearities in the dynamics, the initial set includes the model parameter m_i as well as in the position component x_i. As a consequence, the vector $\mathbf{u}_p$ has components:

$$
\begin{cases}
u_{x_{i,0}} & = \quad 0.10 \quad\quad 1 \leq i \leq L \\
u_{m_i} & = \quad 0.05 \quad\quad 1 \leq i \leq L
\end{cases}, \tag{40}
$$

This yields $2L$ polynomial expansions to be propagated, for x_i and v_i, each in $d = 2L$ variables, $x_{i,0}$ and m_i. The initial set was propagated for $m = 10^4$ time-steps with $\Delta t = 5 \cdot 10^{-3}$ using a forward Euler scheme.

The intrusive propagation is conducted with an algebra of degree $n = 4$. The non-intrusive interpolation is built on a fully-nested Clenshaw-Curtis sparse grid of level 3 using a reduced polynomial basis of degree 8. These parameters are chosen following precision and run-time criteria after an empirical convergence analysis. Figure 7 shows the evolution of the size of the polynomial representations with respect to the dimensionality.
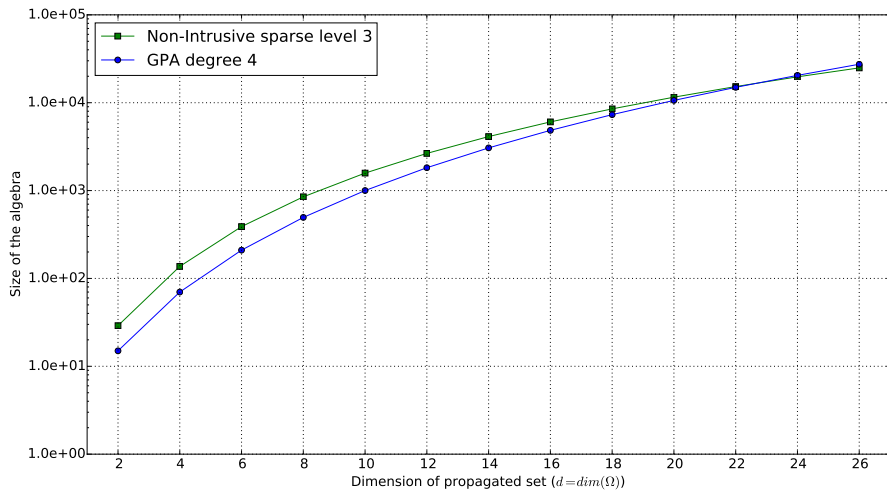

Figure 7: Size of the polynomial representations vs. dimensionality.

6.1.2. Empirical Convergence Analysis

This subsection empirically verifies the assumption on the accuracy of intrusive and non-intrusive approaches for the case of the multi-spring problem. For Chebyshev interpolation we use a fully-nested Clenshaw-Curtis sparse grid with reduced Chebyshev basis as described in Section 4. Both the degree of the intrusive algebra and the level of the sparse grid are increased. The accuracy results are shown in Figures 8, 9 (RMSE) and Figures 10, 11 (Max. abs. err.). The abscissa is the size of the polynomial expansions obtained, $\mathcal{N}_{d,n}^{I}$ and N_{K}^{NI} for the intrusive and non-intrusive approaches respectively. The experiment is repeated for several dimensionality values of the uncertainty space d, namely 2, 4, 8 and 16. Note all results obtained are in line with the assumption that imposing $N_{K}^{NI} \sim \mathcal{N}_{d,n}^{I}$ leads to similar accuracy of the final polynomial expansions for the intrusive and non-intrusive approaches considered.
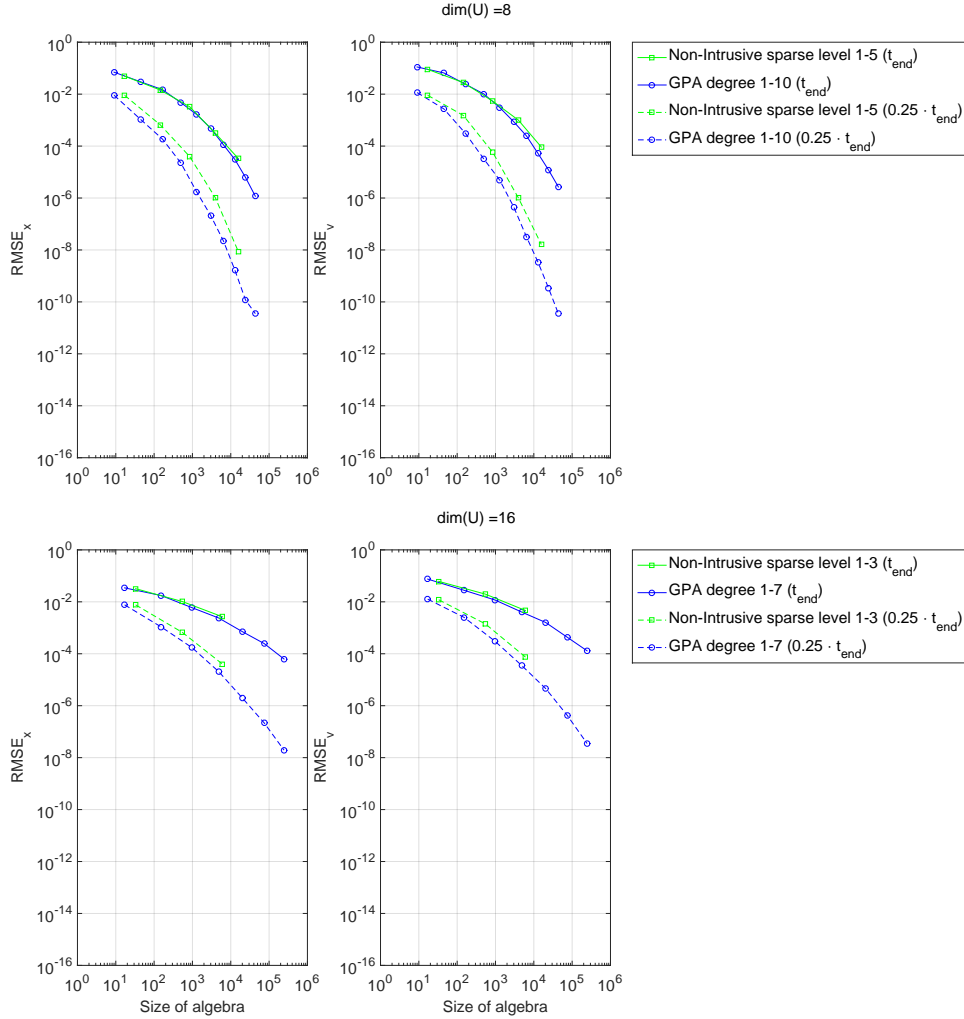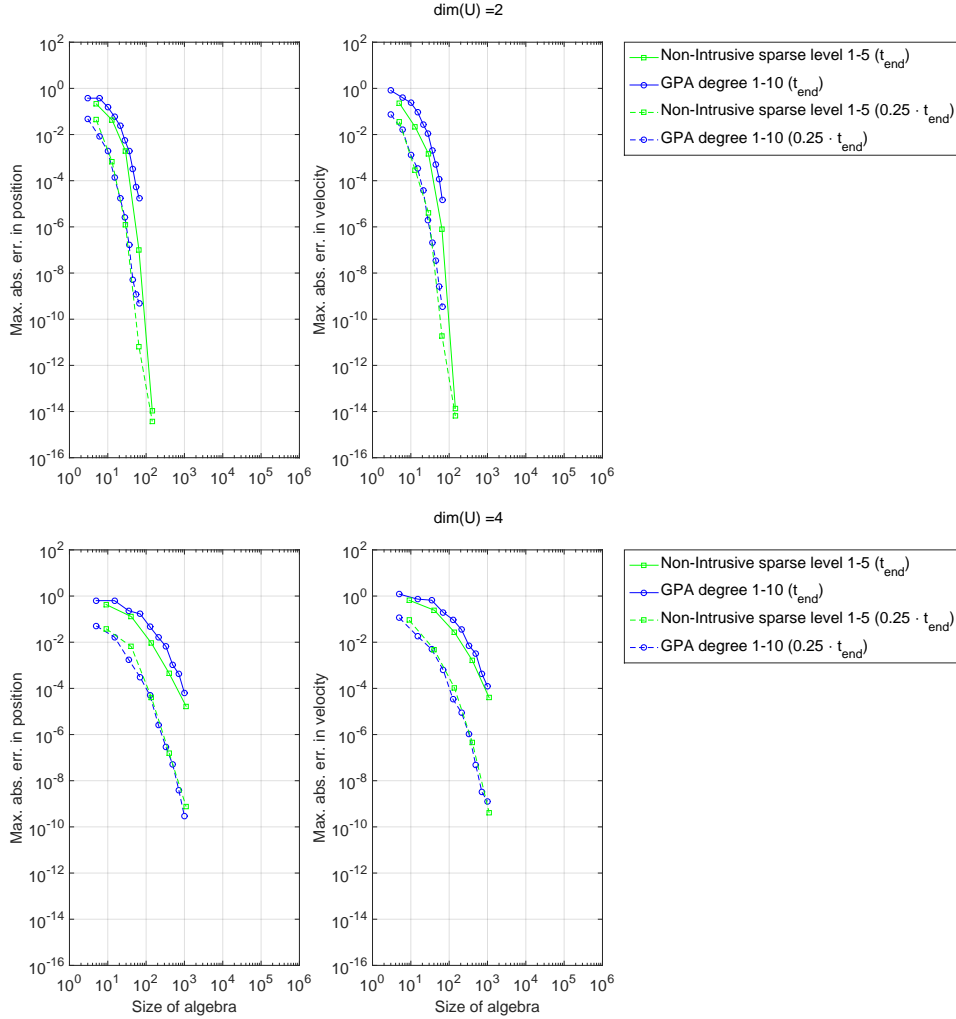
27
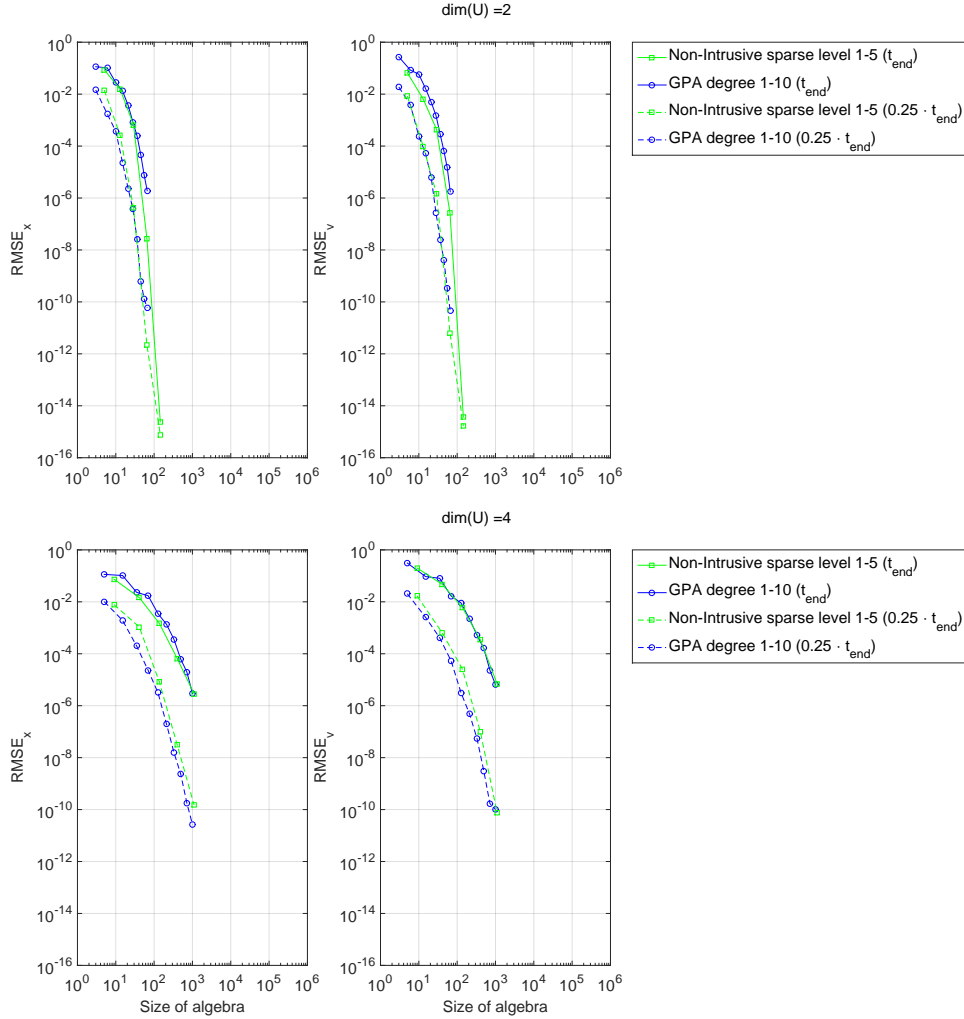
Figure 8: Accuracy of the final polynomial expansions with respect to the size of the algebra for the intrusive and non-intrusive approaches on a multi-DOF spring-mass system. The maximum RMSE among all position (left) or velocity (right) expansions is displayed for 2 (above) and 4 dimensions (below).

6.1.3. Runtime Analysis

Figure 12 shows the runtime necessary to propagate the initial set until t_{end} versus the dimensionality of the set. Only one snapshot of the final set was generated. The result in Figure 12 is in line with the theory and shows a crossing point at which the intrusive method is computationally faster than its non-intrusive counterpart. In this case, the cross-over point

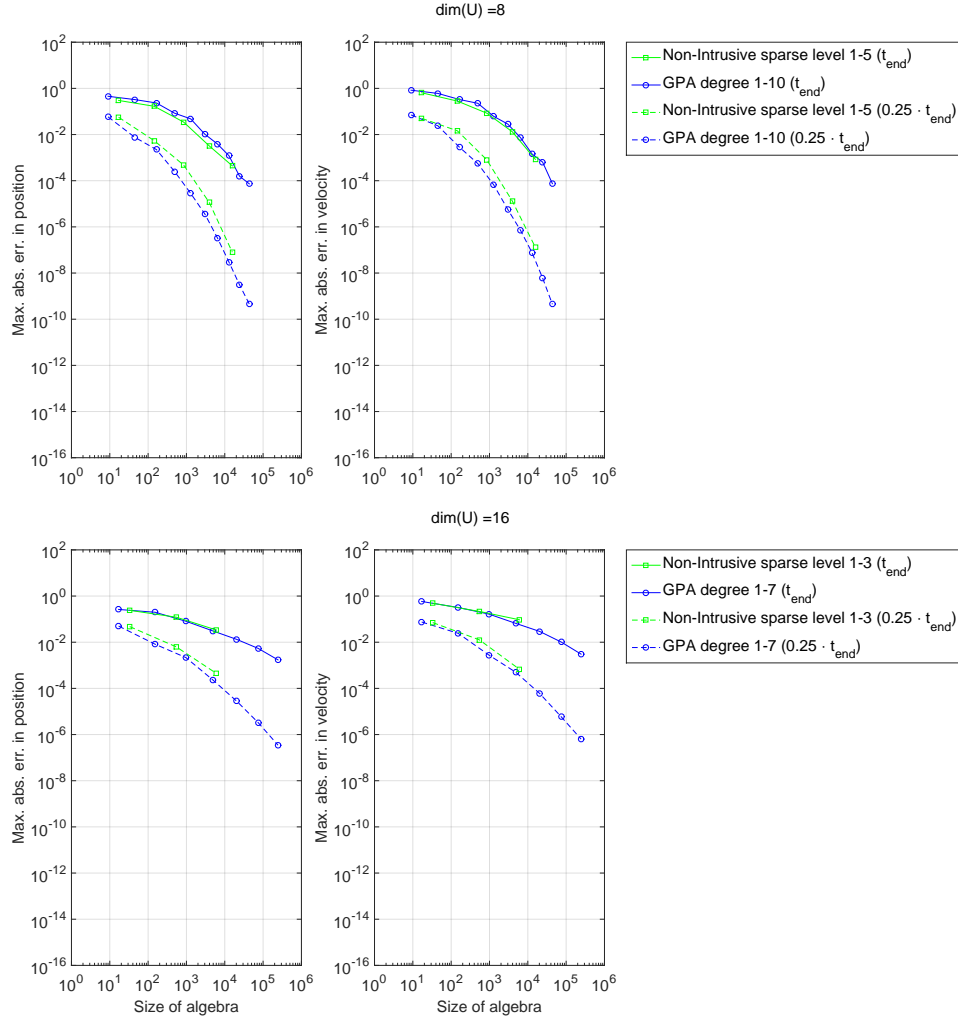Figure 9: Accuracy of the final polynomial expansions with respect to the size of the algebra for the intrusive and non-intrusive approaches on a multi-DOF spring-mass system. The maximum RMSE among all position (left) or velocity (right) state expansions is displayed for 8 (above) and 16 dimensions (below).

is at $d = 16$. The problem presents, for $d = 16$, a simulation model cost of $mF^{RHS} \approx 10^{5.6}$ (where $m = 10^4$ and $F^{RHS} \approx 40$). The cross-over point is thus predicted with acceptable accuracy to be around 15 uncertain variables by the corresponding theoretical cross-over line in Figure 6, reproduced in Figure 14. In Figure 12 runtime has to be intended as the computational time of F^η floating-point operations. Figure 13 shows the cost breakdown of

Figure 10: Accuracy of the final polynomial expansions with respect to the size of the algebra for the intrusive and non-intrusive approaches on a multi-DOF spring-mass system. The maximum absolute error in the sample and among all position (left) or velocity (right) state expansions is displayed for 2 (above) and 4 dimensions (below).

the non-intrusive approach.

Figure 15 shows that, except for very low dimensions, the approximations obtained with the intrusive approach are of comparable or superior accuracy to those obtained by the non-intrusive approach. The differences in accuracy in the lowest dimensionality cases are a consequence of the larger size of the algebra of the non-intrusive approach – see Figure 7. These are in line with

Figure 11: Accuracy of the final polynomial expansions with respect to the size of the algebra for the intrusive and non-intrusive approaches on a multi-DOF spring-mass system. The maximum absolute error in the sample and among all position (left) or velocity (right) state expansions is displayed for 8 (above) and 16 dimensions (below). For the latter, some points are omitted due to computational limitations.

the results in Figures 8 and 9.

6.2. Propagation of Sets in Space Dynamics

The suitability of the proposed methodology for space applications is tested on a three-dimensional two-body dynamical system representing a
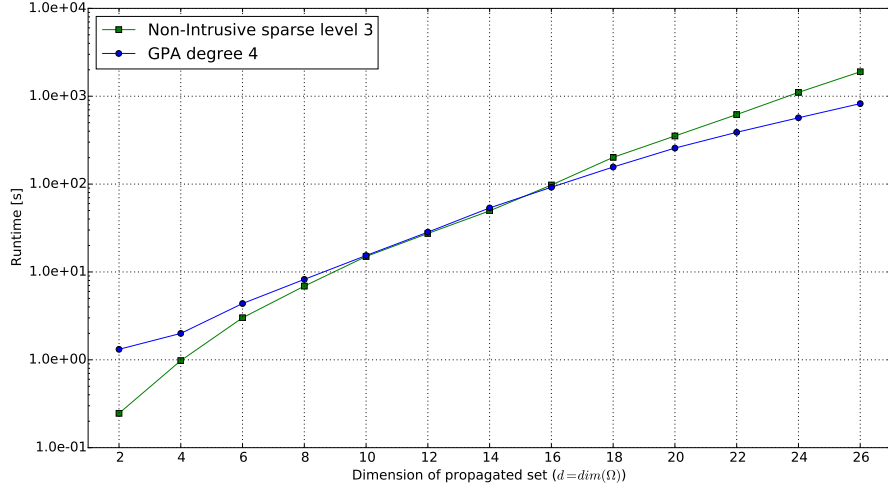
Figure 12: Runtime vs. dimension of the propagated set for the multi-DOF spring-mass system.
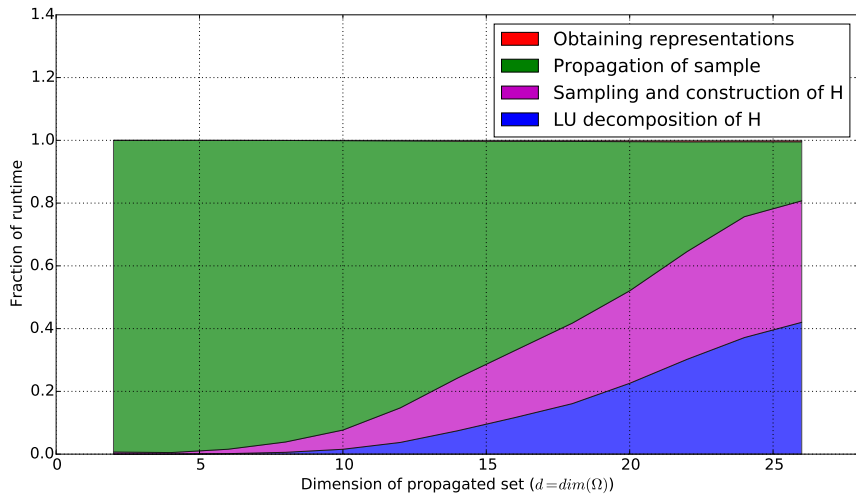
Figure 13: Cost breakdown for the non-intrusive approach applied to the multi-DOF spring-mass system.

spacecraft orbiting in Low Earth Orbit. A constant thrust $\mathbf{T}$, the atmospheric drag and a constant generic perturbing acceleration ϵ are included in the dynamics in order to increase the number of uncertain variables and the
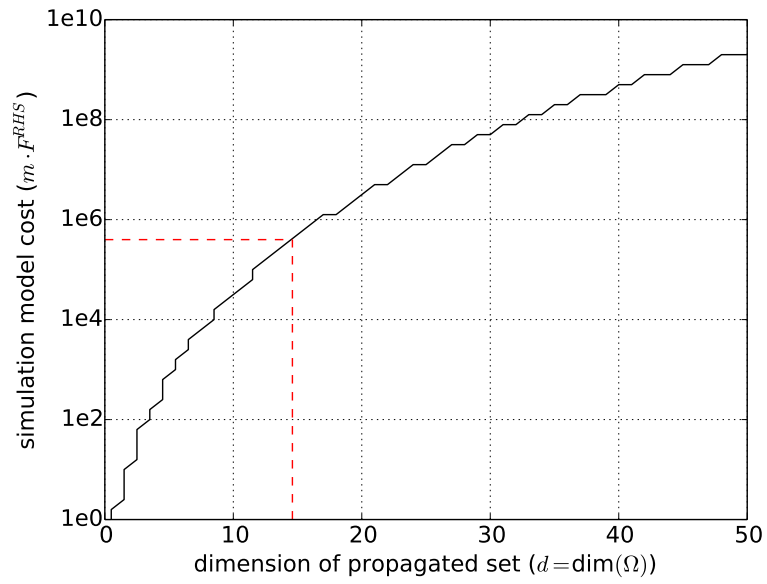
Figure 14: Theoretical cross-over point for the multi-spring problem ($mF^{RHS} \approx 10^{5.6}$), following the assumptions in Section 5.
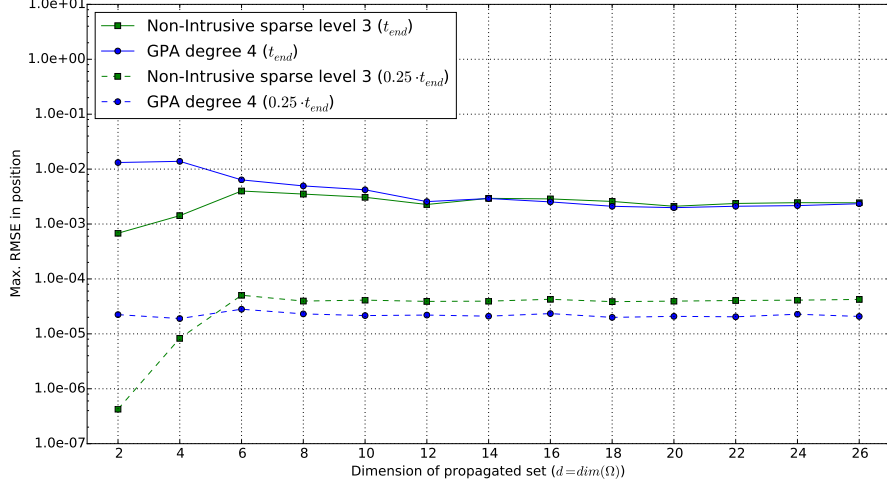
Figure 15: Maximum RMSE in uncertain states vs. dimension of propagated set for $t = t_{end}$ and $t = 0.25 \cdot t_{end}$, multi-DOF spring-mass system test case.

complexity of the problem. In an Earth-centered inertial reference frame the equation of motion, in vector form, governing the dynamics of the spacecraft is:

$$\ddot{\mathbf{x}} = -\frac{\mu}{\|\mathbf{x}\|^3}\mathbf{x} + \frac{\mathbf{T}}{m} + \frac{1}{2}\rho\frac{C_D A}{m}\|\mathbf{v}_{rel}\|\mathbf{v}_{rel} + \epsilon \tag{41}$$

where μ is the gravitational parameter of the Earth, $\mathbf{T}$ the thrust vector, ρ the atmospheric density, $\frac{C_D A}{m}$ the ballistic coefficient of the spacecraft and $\mathbf{v}_{rel}$ its velocity relative to the atmosphere. The mass of the spacecraft m is considered as a seventh state with governing equation:

$$\dot{m} = -\alpha\|\mathbf{T}\|, \tag{42}$$

The initial conditions at time t_0 are:

$$\begin{aligned}
x(t_0) &= x_0 \\
y(t_0) &= y_0 \\
z(t_0) &= z_0 \\
v_x(t_0) &= v_{x,0} \quad , \\
v_y(t_0) &= v_{y,0} \\
v_z(t_0) &= v_{z,0} \\
m(t_0) &= m_0
\end{aligned} \tag{43}$$

34

All quantities are normalised with the following scale factors:

$$\text{DU} = 6378136\,[\text{m}]\,, \quad \text{TU} = 806.78\,[\text{s}]\,, \quad \bar{m}_0 = 2000\,[\text{kg}]\,.$$

corresponding to the planetary canonical units of the Earth and to the nominal initial mass of the spacecraft considered.

6.2.1. Experimental Setup

The non-dimensional set of nominal initial conditions selected is:

$$\bar{\mathbf{p}} = [\bar{x}_0, \bar{y}_0, \bar{z}_0, \bar{v}_{x,0}, \bar{v}_{y,0}, \bar{v}_{z,0}, \bar{m}_0]^T = [1.1505, 0, 0, 0, 0.9297, 0, 1]^T$$

Nominally the thrust magnitude is $2.5513 \cdot 10^{-5}$ and the vector points in the y direction, with $\bar{\alpha} = 0.2635$. The density ρ follows an exponential atmospheric model

$$\rho = \rho_b exp\left(-\frac{r - r_b}{H_b}\right)\,,$$

with $\bar{\rho}_b = 680.451$, $r_b = 1.1411$ and $\bar{H}_b = 0.0284$ corresponding to the initial altitude Vallado (1997), and $\bar{C}_D A = 1.081610^{-13}$. The nominal value of the magnitude of the perturbation ϵ is set to zero.

Four instances of the problem have been investigated with identical nominal dynamics and progressively increasing dimension of the propagated set. Case 1 considers a set defined only by the initial conditions whereas cases 2 to 4 progressively include up to 10 model parameters. The definition of the initial set for each of the four cases can be found in Table 1. In all cases the degree of the algebra was set to $n = 4$ for GPA, whereas different non-intrusive configurations were tested, namely Chebyshev interpolation on a Latin Hypercube Sample with degree 4 and 5 and interpolation with reduced Chebyshev basis on fully-nested Clenshaw-Curtis sparse grids of level 2 and 3. For the sake of a fair comparison, extended results are presented only for those configurations whose the accuracy on the final states is comparable to that provided by GPA. The equations are integrated until $t_{end} = 7.8175$ with a fixed-step-size Runge-Kutta 4th-order scheme with $\Delta t = 0.0124$, resulting in $m = 630$ time-steps.

6.2.2. Empirical Convergence Analysis

This subsection empirically verifies the assumption on the accuracy of intrusive and non-intrusive approaches for the propagation in space dynamics. Test case 4 of dimension 17 is used for the comparison. Cases 1 to 4

Table 1: Initial set for the four cases of propagation in space dynamics.

Test case	1	2	3	4
$u_{\mathbf{x}(0)}$	$1.5679 \cdot 10^{-4}$	$1.5679 \cdot 10^{-4}$	$1.5679 \cdot 10^{-4}$	$1.5679 \cdot 10^{-4}$
$u_{\mathbf{v}(0)}$	$6.3246 \cdot 10^{-4}$	$6.3246 \cdot 10^{-4}$	$6.3246 \cdot 10^{-4}$	$6.3246 \cdot 10^{-4}$
$u_{m(0)}$	$0.05\,\%$	$0.05\,\%$	$0.05\,\%$	$0.05\,\%$
$u_{\mathbf{T}},\ u_{\alpha}$	–	$5.00\,\%$	$5.00\,\%$	5.00%
$u_{\rho_b},\ u_H,\ u_{C_D}$	–	–	$1.00\,\%$	$1.00\,\%$
u_{ϵ}	–	–	–	$1.0205 \cdot 10^{-5}$
d	7	11	14	17

present very similar overall accuracy results due to the propagation being driven by the variability of the initial states, which are identical in all cases and have much larger impact on the dynamics than the variability of model parameters.

For Chebyshev interpolation we use a fully-nested Clenshaw-Curtis sparse grid with reduced Chebyshev basis as in section 6.1. Besides, we include in the comparison a non-intrusive interpolation with a full Chebyshev basis on a Latin Hypersquare Sample, which is a stochastic non-structured sampling method. The degree of the intrusive algebra and the level of the sparse grid are increased as well as the degree of the non-intrusive full basis. The accuracy results are shown in Figure 16 (RMSE) and Figure 17 (maximum absolute error). The abscissa is the size of the polynomial expansions obtained, $\mathcal{N}_{d,n}^{I}$ and N_{K}^{NI} for the intrusive and non-intrusive approaches respectively. Error measures are presented for each of the final position and velocity state expansions. Figures 16 and 17 show that, when comparing the intrusive approach to interpolation on a sparse grid, imposing $N_{K}^{NI} \sim \mathcal{N}_{d,n}^{I}$ leads to similar accuracy of the final approximations. Both methods dominate the non-intrusive approach with full basis. This difference in performance between the non-intrusive approaches derives from the benefits of using a structured sample such as in the fully-nested Clenshaw-Curtis sparse grids, with respect to interpolation on a non-structured Latin Hypersquare Sample. The GPA approach matches the accuracy of the better-performing non-intrusive approach.
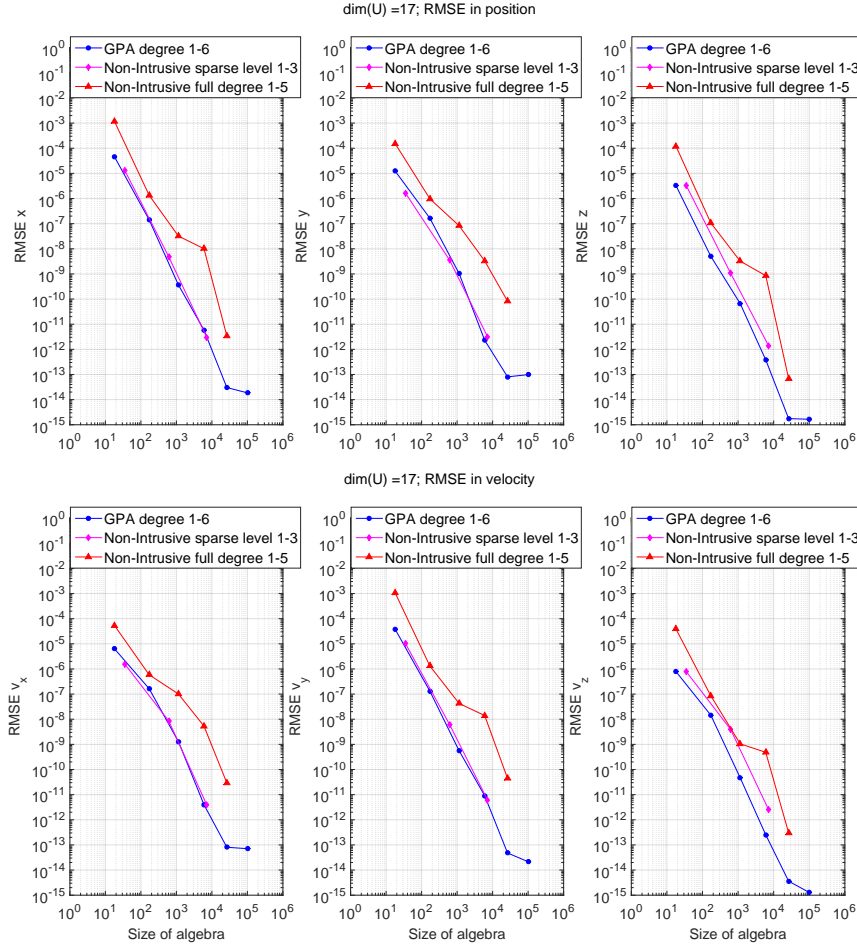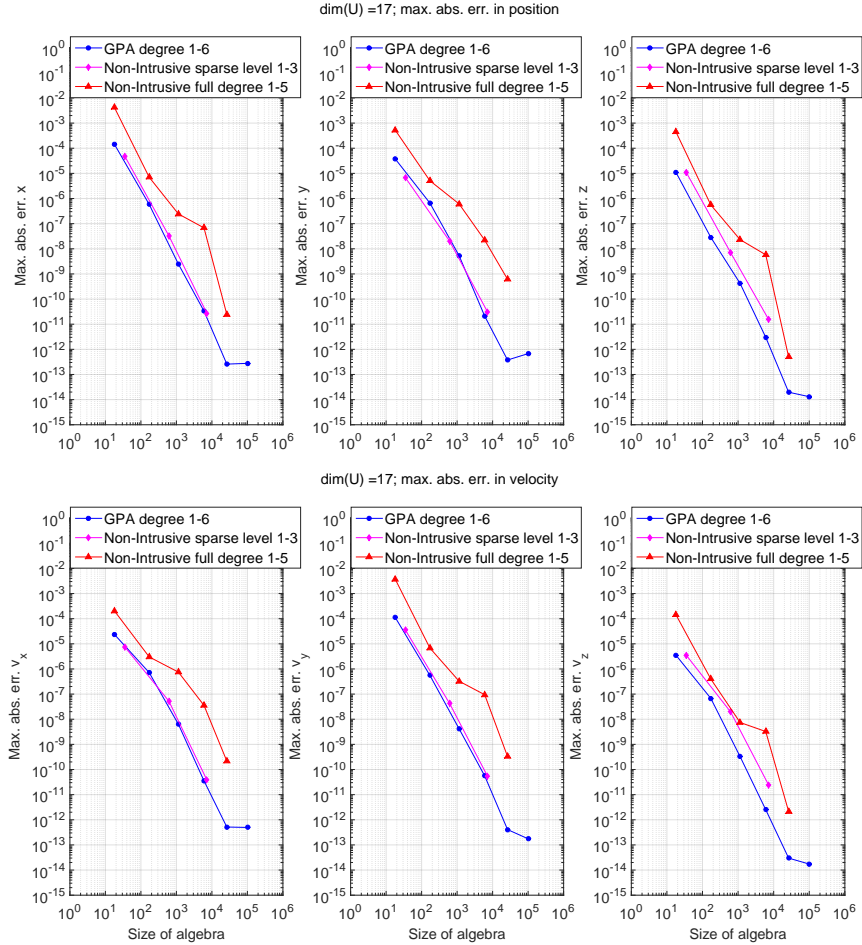
Figure 16: Accuracy of the final polynomial expansions with respect to the size of the algebra for the intrusive and non-intrusive approaches on the propagation of space dynamics case 4. The final RMSE is displayed for each of the position and velocity state expansions.

6.2.3. Runtime Analysis

The degree of the algebra was set to $n = 4$ for GPA, whereas various non-intrusive configurations have been tested. Figure 18 shows that interpolation with a full basis of degree 4 and interpolation with reduced basis on a sparse grid of level 2 present RMSE values that are completely dominated by those of GPA degree 4. On the other hand, the results obtained with full basis of degree 5 and sparse grid of level 3 are considered comparable in accuracy to those of GPA degree 4 and will, therefore, be the only included in the rest of this section.

Figure 17: Accuracy of the final polynomial expansions with respect to the size of the algebra for the intrusive and non-intrusive approaches on the propagation of space dynamics case 4. The final maximum absolute error in the sample is displayed for each of the position and velocity state expansions.

Depictions of the uncertain regions and RMSE values in **x** and **v** are presented in Figures 19 to 21. All three methods provide very precise representation of the uncertain region, proving their worth for uncertainty propagation in orbital mechanics. The mean final approximation errors obtained are in the order of only 10^{-1} mm and 10^{-4} mm/s in the plane of nominal motion $< x, y >$.

As hinted by the convergence analysis, it is necessary to increase the degree of the non-intrusive method with full basis to 5 for its accuracy to match the

38
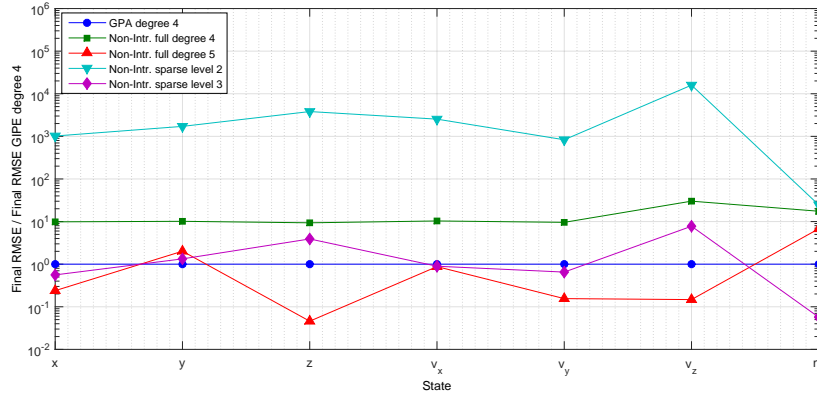
Figure 18: RMSE in states at the final snapshot for several non-intrusive configurations, propagation in space dynamics case 4. Each RMSE value is scaled with the corresponding RMSE value of GPA degree 4.

one provided by GPA degree 4. Whereas for the sparse-grid version a level 3 was necessary, which approximately matches the size of the full-basis algebra of degree 4 for this range of dimensions – see Figure 7.

Figure 22 shows the CPU time necessary for each case and method. The reduced-basis method (sparse grids) is faster than the full-basis one for all cases. However, cross-over points exist between the intrusive and each of the non-intrusive methods due to the lower complexity of GPA to an increase in dimensionality. GPA is thus the fastest approach among those with comparable accuracy for the cases with the propagated set with dimensions 14 and 17. The operation count of this case amounts to $m \cdot F^{RHS} \approx 10^{5.2}$; Figure 25 shows the cross-over points as predicted by the formulae in Section 5. Comparing it with Figure 22, one observes a slight underestimation in the exact location of the experimental cross-over points, which derive from all the simplifications made in the cost computation. Nonetheless, the deviations are small and the theory predicts accurately which approach will be most efficient for each case considered.

Figure 23 shows the cost breakdown of the non-intrusive method with full basis, split between (by asymptotical cost): decomposition LU of matrix H, sampling and evaluation of the matrix H, propagation of sample sets and reconstruction of the final uncertainty region. Figure 24 shows the same breakdown for the non-intrusive sparse method of level 3.
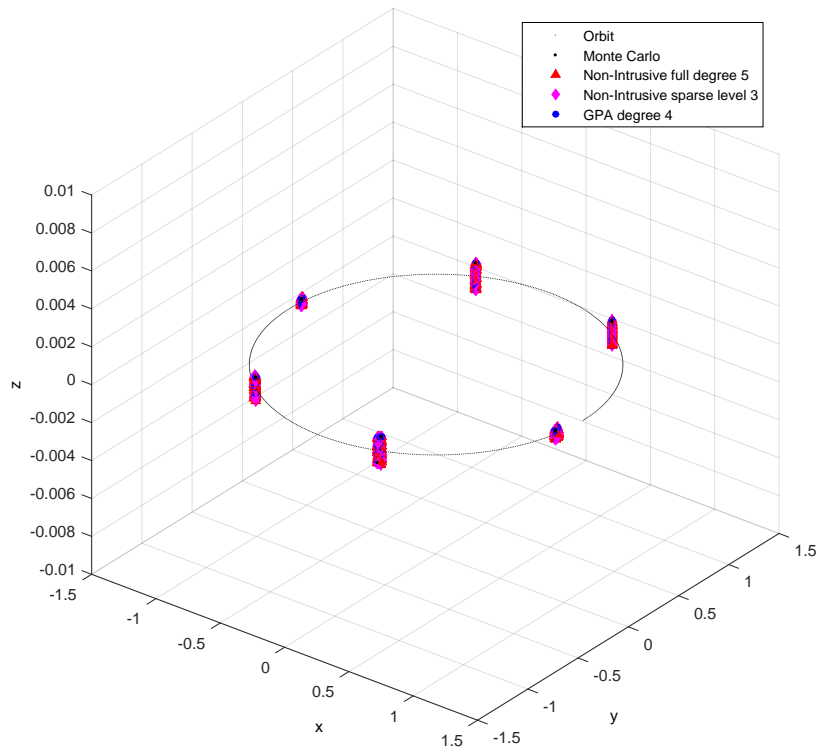
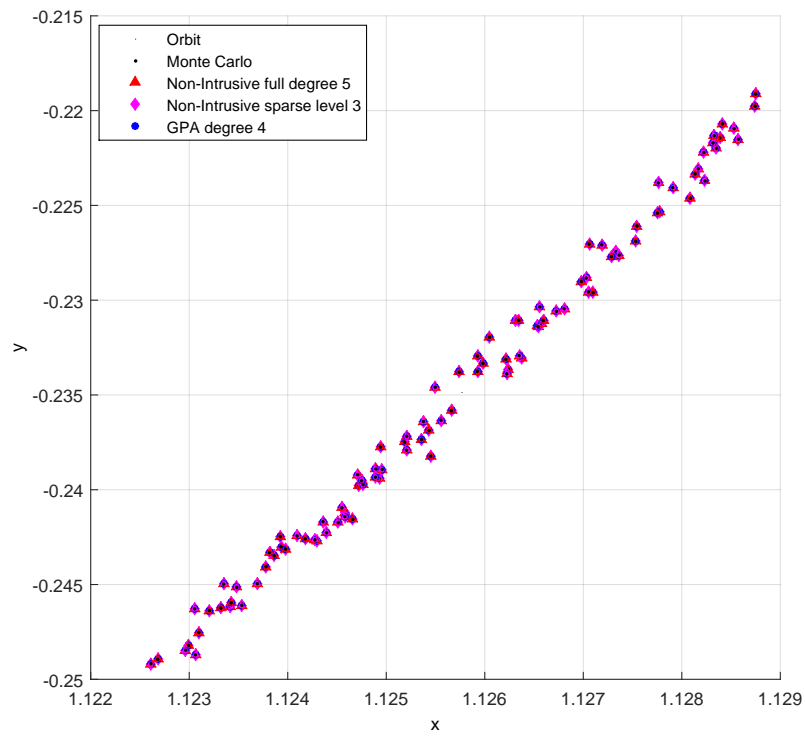Figure 19: Propagated sets and nominal orbit, propagation in space dynamics case 4.

Figure 20: Detail of the final set, propagation in space dynamics case 4.
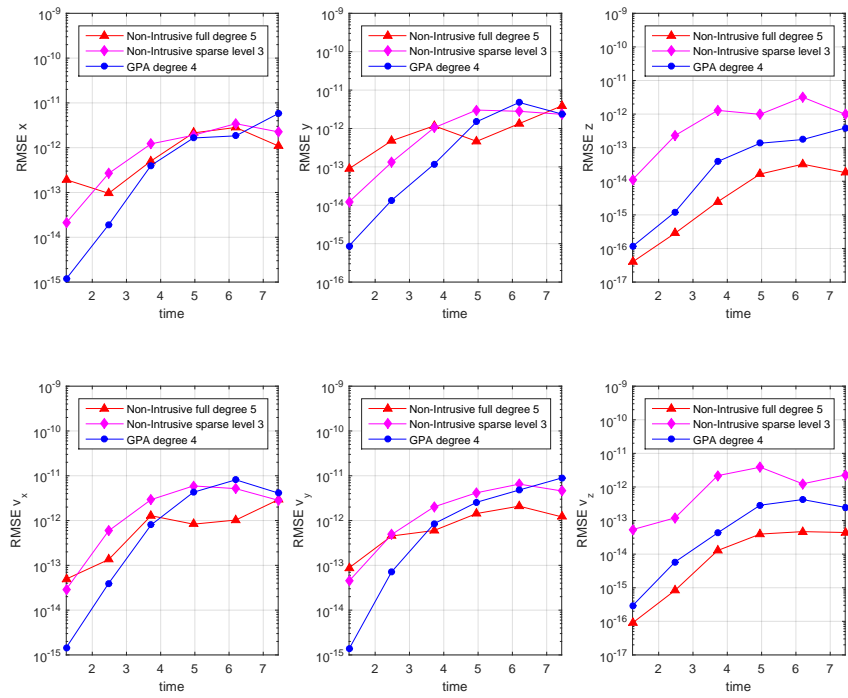
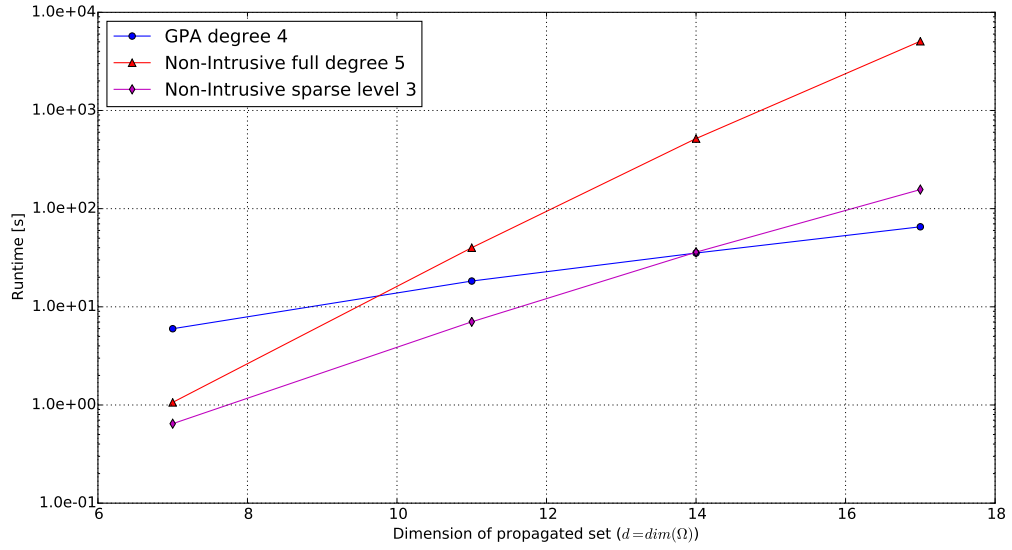Figure 21: RMSE in **x** and **v** states, propagation in space dynamics case 4.

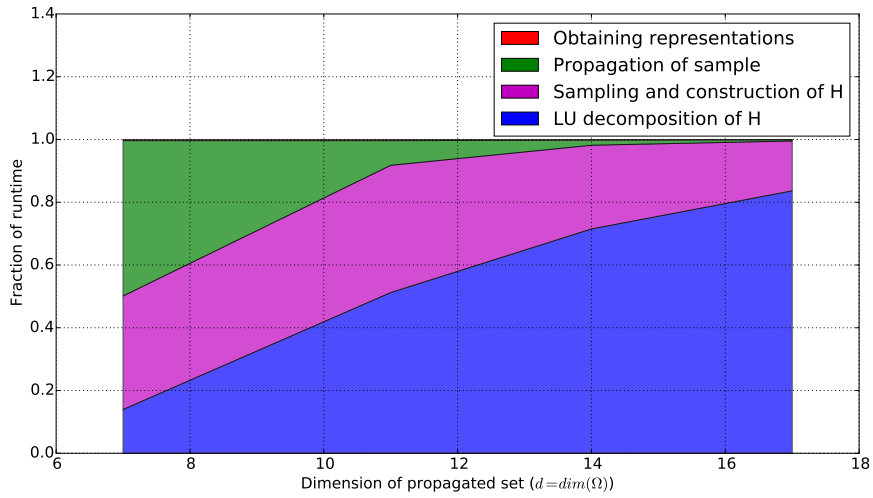Figure 22: Runtime vs. dimension of set, propagation in space dynamics.

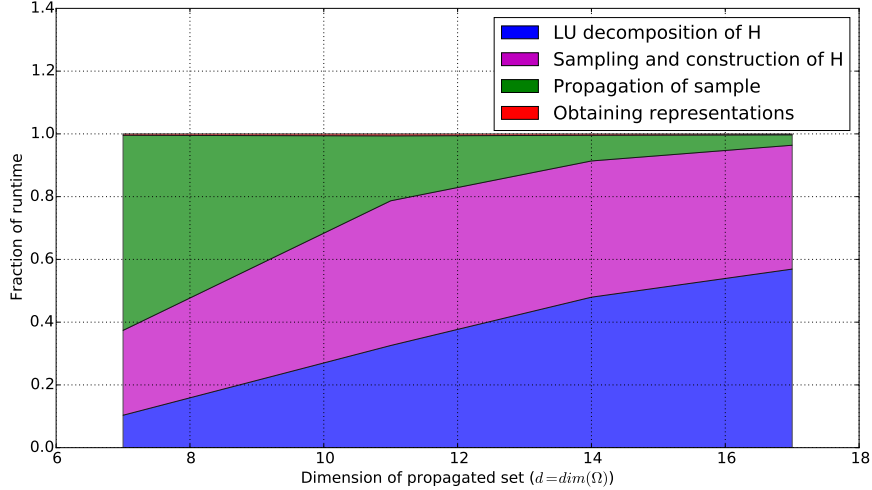Figure 23: Cost breakdown of non-intrusive method, full basis degree 5

43

Figure 24: Cost breakdown of non-intrusive sparse method, reduced basis level 3

7. Conclusions

The paper presented a Generalised Polynomial Algebra (GPA) approach to propagate generic sets through dynamical systems and compared, theoretically and empirically, its time complexity against a non-intrusive counterpart.

The complexity analysis found that there is a well defined cross-over line that predicts when the intrusive approach is computationally more efficient than its non-intrusive counterpart for the same accuracy. The same crossing point was found also in the numerical experiments suggesting that for low-dimensional problems a non-intrusive approach is advantageous. Likewise for long propagation times when the deformation of the propagated set is limited a non-intrusive approach is advisable.

For higher-dimensional problems the computational cost associated to the generation of the polynomial representation exceeds the cost of the sampling process and renders the non-intrusive approach more expensive than the intrusive one. This was shown to be true for both a full and a sparse basis.

It was also shown that in the cases considered in this paper the intrusive approach could achieve higher accuracy for a lower order of the polynomials than non-intrusive interpolation with a full Chebyshev basis on a Latin Hypersquare Sample, and not lower than interpolation on fully-nested
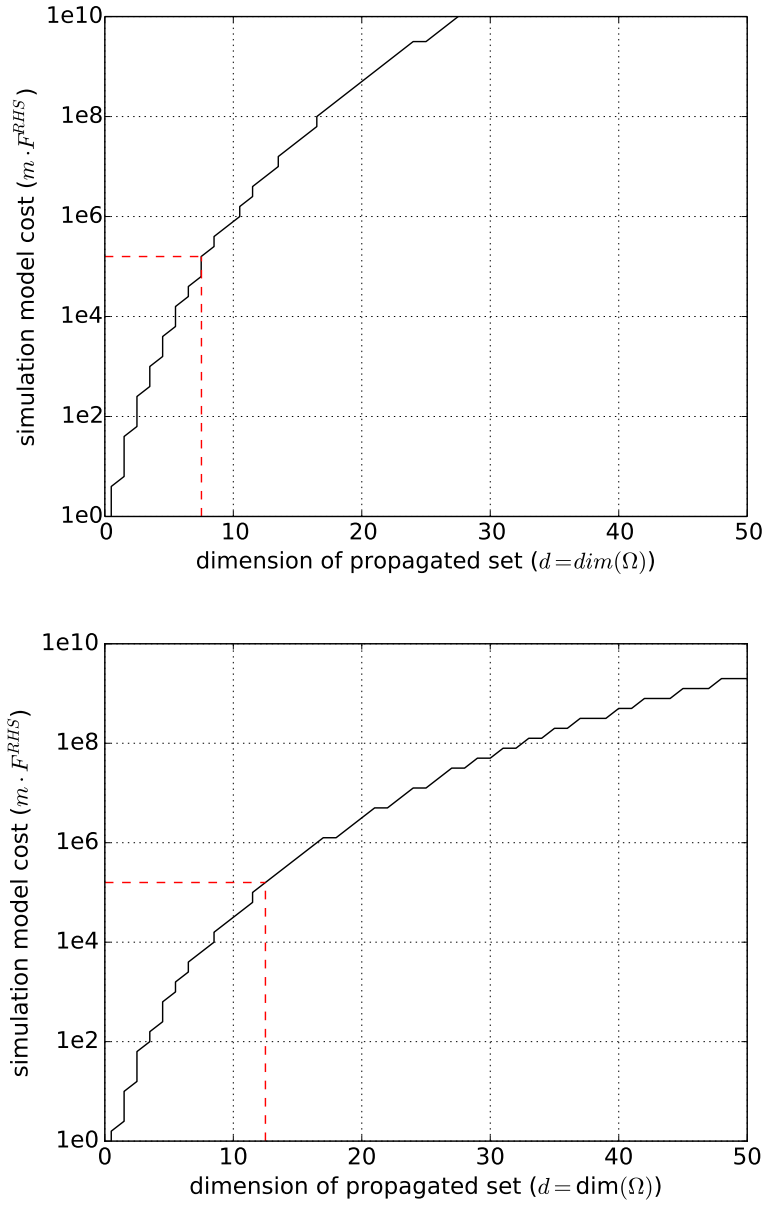
44

Figure 25: Theoretical cross-over points for the propagation in space dynamics problem ($mF^{RHS} \approx 10^{5.2}$). Above, tailored for the comparison with non-intrusive full basis of degree 5. Below, following the assumptions in Section 5.

Clenshaw-Curtis sparse grids with a reduced Chebyshev basis. Thus one can argue that GPA presents overall good performance and scalability compared to its non-intrusive counterpart.

Further research directions include the use of reduced-bases also in the intrusive approach, as well as the adaptation of the expansion to maintain the required level of accuracy. This is particularly relevant when one or more singularities exist in the domain of interest. Future work will present a direct comparison with other intrusive approaches.

8. Acknowledgment

References

Amparo Gil, Javier Segura, N. M. T., 2008. Numerical Methods for Special Functions. Society for Industrial and Applied Mathematics.

Armellin, R., Di Lizia, P., Bernelli-Zazzera, F., Berz, M., 2010. Asteroid close encounters characterization using differential algebra: the case of apophis. Celestial Mechanics and Dynamical Astronomy 107 (4), 451–470.

Berz, M., 1986. The new method of TPSA algebra for the description of beam dynamics to high orders. Tech. Report AT–6: TN–86–16, Los Alamos National Laboratory, Los Alamos, N.M.

Berz, M., 1987. The method of power series tracking for the mathematical description of beam dynamics. Nuclear Instruments and Methods A258, 431–436.

Berz, M., 1997. From Taylor series to Taylor models. In: in AIP Conference Proceedings 405. pp. 1–23.

Brisebarre, N., Joldes, M., 2010. Chebyshev interpolation polynomial-based tools for rigorous computing. In: Proceedings of the 2010 International Symposium on Symbolic and Algebraic Computation. ISSAC '10. ACM, New York, NY, USA, pp. 147–154.
URL http://doi.acm.org/10.1145/1837934.1837966

Brualdi, R., 1977. Introductory Combinatorics. North-Holland.

Clenshaw, C. W., 07 1955. A note on the summation of chebyshev series. Mathematical Tables and Other Aids to Computation 9, 118–120.

De Mars, K. J., Jah, M. K., 2013. Probabilistic initial orbit determination using gaussian mixture models. Journal of Guidance, Control, and Dynamics 36 (5), 1324–1335.

Di Lizia, P., Armellin, R., Lavagna, M., 2008. Application of high order expansions of two-point boundary value problems to astrodynamics. Celestial Mechanics and Dynamical Astronomy 102 (4), 355–375.

Di Lizia, P., Armellin, R., Morselli, A., Bernelli-Zazzera, F., 2014. High order optimal feedback control of space trajectories with bounded control. Acta Astronautica 94 (1), 383–394.
URL http://eprints.soton.ac.uk/360515/

Epstein, C., Miranker, W., Rivlin, T., 1982. Ultra-arithmetic i: Function data types. Mathematics and Computers in Simulation 24 (1), 1 – 18.
URL http://www.sciencedirect.com/science/article/pii/0378475482900453

Garloff, J., 1985. Convergent bounds for range of multivariate polynomials. In: Proceedings of the International Symposium on Interval Mathemantics. pp. 36–57.

George B. Arfken, F. E. H., 2013. Mathematical Methods for Physicists (Seventh Edition). Elsevier.

Gerstner, T., Griebel, M., Jan 1998. Numerical integration using sparse grids. Numerical Algorithms 18 (3), 209.
URL https://doi.org/10.1023/A:1019129717644

Iman, R., Davenport, J., Zeigler, D., 1 1980. Latin hypercube sampling (program user's guide). [lhc, in fortran].

J.C. Mason, D. C. H., 2002. Chebyshev Polynomials. Chapman and Hall/CRC.

Jones, B. A., Doostan, A., Born, G. H., March 2013. Nonlinear propagation of orbit uncertainty using non-intrusive polynomial chaos. Journal of Guidance Control Dynamics 36, 430–444.

Jorba, A., Zou, M., 2005. A software package for the numerical integration of odes by means of high-order Taylor methods. Experimental Mathematics 14, 99–117.

Judd, K., Maliar, L., Maliar, S., Valero, R., 07 2014. Smolyak method for solving dynamic economic models: Lagrange interpolation, anisotropic grid and adaptive domain. Journal of Economic Dynamics and Control 44.

Makino, K., Berz, M., 2003. Taylor models and other validated functional inclusion methods. International Journal of Pure and Applied Mathematics 4 (4), 379–456.

McKay, M., Beckman, R., Conover, W., 1979. A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. Technometrics 21 (2), 239–245.

Nesterov, Y., 2000. Squared functional systems and optimization problems. In: In: Frenk H., Roos K., Terlaky T., Zhang S. (eds) High Performance Optimization. Applied Optimization, vol.33. Springer, Boston, MA, pp. 405–440.

Ortega, C., Serra, R., Riccardi, A., Vasile, M., September 2016. De-orbiting and re-entry analysis with generalised intrusive polynomial expansions. 67th International Astronautical Congress, Guadalajara, Mexico.

Park, R. S., Scheeres, D. J., 2006. Nonlinear mapping of gaussian statistics: Theory and applications to spacecraft trajectory design. Journal of Guidance, Control, and Dynamics 29 (6).

Rivlin, T. J., 1974. The Chebyshev Polynomials. Wiley.

Smolyak, S., 1963. Quadrature and interpolation formulas for tensor products of certain classes of functions. Dofl. Akad. Nauk 158, 1042–1045.

Tardioli, C., Kubicek, M., Vasile, M., Minisci, E., Riccardi, A., August 9–13 2015. Comparison of non-intrusive approaches to uncertainty propagation in orbital mechanics. AAS Astrodynamics Specialists Conference, Vail, Colorado, USA AAS 15-545.

Trefethen, L. N., 2013. Approximation Theory and Approximation Practice. Society for Industrial and Applied Mathematics.

Vallado, D., 1997. Fundamentals of astrodynamics and applications. Springer.