IAC–18–C1.8.5x46264

AN INTRUSIVE POLYNOMIAL ALGEBRA MULTIPLE SHOOTING APPROACH TO THE SOLUTION OF OPTIMAL CONTROL PROBLEMS

Cristian Greco[a*], Marilena Di Carlo[a], Massimiliano Vasile[a], Richard Epenoy[b]

[a] *Mechanical and Aerospace Department, University of Strathclyde, 75 Montrose Street, Glasgow, United Kingdom G1 1XJ*, {c.greco , marilena.dicarlo , massimiliano.vasile}@strath.ac.uk
[b] *Centre National d'Etudes Spatiales (CNES), 18 Avenue Edouard Belin, Toulouse, France 31401*, richard.epenoy@cnes.fr
[*] Corresponding author

Abstract

This paper proposes an approach to the solution of optimal control problems under uncertainty, that extends the classical direct multiple shooting transcription to account for random variables defined on extended sets. The proposed approach employs a Generalised Intrusive Polynomial Expansion to model and propagate uncertainty. The development of a generalised framework for a direct multiple shooting transcription of the optimal control problem starts with the discretisation of the time domain in sub-segments. At the beginning of each segment, the state spatial distribution is modelled with a multivariate polynomial and then propagated to the sub-interval final time. Continuity conditions are implicitly imposed at the boundary of two adjacent segments, a critical operation because it requires the continuity of two extended sets. The Intrusive Polynomial Algebra aNd Multiple shooting Approach (IPANeMA) developed in this paper can handle optimal control problems under a wide range of uncertainty models, e.g. nonparametric, expensive to sample, and imprecise probability distributions. In this paper, the approach is applied to the design of a low-thrust trajectory to a Near-Earth Object with uncertain initial conditions.

Keywords: optimal control under uncertainty; robust control; generalised multiple shooting; intrusive polynomial algebra; low-thrust trajectory optimisation.

1. Introduction

Optimal control problems aim at finding the optimal control law for a single trajectory evolving in a deterministic nonlinear system. Hence, the resulting control is valid only for the computed reference trajectory. However, in real-life applications, perfect compliance to the reference trajectory is impossible to achieve as uncertainty always affects the system; uncertainty is due to both imperfect state knowledge and to unknown model parameters. Furthermore, for nonlinear systems and large time-scales, even small deviations from a pointwise trajectory can lead to significant differences as the system evolves over time. In space applications, low-thrust missions are rather sensitive to trajectory deviations. Indeed, due to the limited control authority, long periods of maximum thrust are required to build up the nominal orbital changes. If possible uncertainty is not taken into account during the stages of trajectory design, this may leave no room for compensation maneuvers. One common cause of trajectory deviation in low-thrust trajectories is missed-thrust due to sub-systems partial failure or external causes, like experienced by Dawn and Hayabusa missions.

To compensate for possible deviations, currently the practical solution is to consider propellant margins and enforced coasting arcs in the reference trajectory design [9]. Several research works developed methods to deal with an optimal control formulation which models uncertainty directly. Methods based on model predictive control or closed-loop formulations takes into account directly correction terms based on possible deviations from the desired trajectory [11][6]. A method based on Taylor polynomials algebra has been developed to deal with uncertain boundary conditions around a reference trajectory [5]. In addition, stochastic differential dynamic programming has been applied to space trajectory optimisation with uncertainty with an expected value formulation [8].

The common baseline of these works is the presence of a desired reference trajectory and undesired deviations from it. This problem statement can be either formulated implicitly, by working with expected valued objective and constraints, or explicitly, by trying to compensate the trajectory deviations. Furthermore, often these techniques can deal only with simple families of probability density distributions to represent uncertainty.

This paper proposes a tool for the transcription and solution of optimal control problem with uncertainty phrased under a more general probabilistic framework. Specifically, the premise of a reference trajectory is abandoned in favor of an extended uncertainty set representation. Each sample within the uncertainty set is a fully admissible pointwise trajectory with associated probability density. Aiming at a probabilistic framework, the objective function and constraint formulation are framed accordingly, since an expected value formulation would result too limited. From here, the goal is to compute control profiles able to optimally steer the uncertain region to a final target set, while minimising the modified objective function. The developed tool is a generalised multiple shooting for the optimal control problem transcription, coupled with intrusive polynomial algebra for the uncertainty propagation.

The paper is structured as follows. Section 2 introduces the formulation of the addressed optimal control problem under uncertainty. Section 3 presents the main development, first introducing the intrusive polynomial algebra propagation, and then integrating it with the novel generalised multiple shooting framework using an expectation formulation for the objective function and constraints. Within this framework, a specific approach is proposed based on polynomial reinitialisation and successive sampling. The developed tool is then applied to the optimisation of a low-thrust rendezvous trajectory to a Near-Earth Object in Section 4. Finally, Section 5 concludes the paper with the final remarks.

2. Optimal Control under Uncertainty

Generally, the deterministic optimal control problem statement is formulated as follows:

$$\begin{aligned} &\min_{\mathbf{u}(t)\in U} J \\ &\text{s.t.} \quad \dot{\mathbf{x}} = f(t,\mathbf{x},\mathbf{u},\mathbf{d}) \\ &\qquad \mathbf{g}(t,\mathbf{x},\mathbf{u},\mathbf{d}) \in \boldsymbol{G} \\ &\qquad \boldsymbol{\psi}(t_0,\mathbf{x}_0,t_f,\mathbf{x}_f) \in \boldsymbol{\Psi} \end{aligned} \qquad [1]$$

The objective function J can be in Bolza form in the most general case, i.e. with both end-cost and integral terms, while both path $\mathbf{g}$ and boundary $\boldsymbol{\psi}$ constraints could be imposed. A set inclusion formulation has been used to describe simultaneously both the admissible cases of equality and inequality constraints. This framework is suitable for a single trajectory optimisation.

When uncertainties and random factors come into play, a set of admissible trajectories is associated to a single control. For this situation, the framework above results too limited. The main challenge addressed in this section is the formulation of a general optimal control problem of a dynamical system under uncertainty.

The general uncertainty vector is denoted as $\mathbf{Z}$ with probability density distribution $p(\boldsymbol{\xi})$. Usually, it models possible uncertainty affecting the initial state and model parameters. As a result, the state and model parameters become random variables themselves $\mathbf{X} = \mathbf{X}(\mathbf{Z})$ and $\mathbf{D} = \mathbf{D}(\mathbf{Z})$. The lower case letters $\boldsymbol{\xi}$, $\mathbf{x}$ and $\mathbf{d}$ denote an admissible realisation. On the other hand, the control variables will be treated as completely deterministic input since possible disturbances on the control can be modelled in the dynamics as a multiplicative noise incorporated in $\mathbf{Z}$.

The dynamical equations induce the state density distribution to evolve over time $p(\mathbf{x})$. In the general nonlinear case, computing directly its evolution is an ambitious, and when possible laborious, task. Therefore, this problem is often tackled with sampling techniques. Indeed, the dynamical equations can be used directly as a map from the state and parameter sample space at a given time to the state sample space at another time. The distribution at the time of interest is then reconstructed according to the sample responses, usually by fitting a parametric (possibly discrete) distribution.

In the transition from the deterministic to the uncertain setting, the main divergence lies in the definition of objective functions and constraints depending on random variables. Specifically, since the random variable state affects their value, generally they turn out to be random variables themselves. How to formulate an objective or constraint on a random variable is a design choice that directly affects the interpretation and result of the optimisation process. Common choices in stochastic programming are to impose so-called objective or constraints *in expected value* or *in probability* [2]. In order to have a single notation, we will write both the possible formulations *in expectation* form with the auxiliary function ϕ. The expectation formulation is flexible as it allows to define a variety of different quantities just selecting the appropriate function. From here, depending on the quantity of interest, the expectation can be minimised or constrained. Specifically for the constraints, the set of acceptable constraint values $\boldsymbol{\Phi}$ for the inclusion relationship should be defined accordingly.

It is worth describing in detail how the expectation formulation encloses common cases of objective or constraints forms:

- *in expected value*, for which the function ϕ is the mapping between the trajectory realisation and the quantity of interest. As an example, the expected value of the final state $\mathbf{X}_f$ can be constrained to be equal to a target state $\mathbf{x}_f$. In this case, the auxiliary function is the identity mapping of the final state

$$\phi_{\boldsymbol{\psi}}(t_f, \mathbf{X}_f) = \mathbf{X}_f \ , \qquad [2]$$

resulting in the constraint formulation

$$\mathbb{E}[\mathbf{X}_f] \in \boldsymbol{\Phi}_{\boldsymbol{\psi}} = \{\mathbf{x}_f\} \ ; \qquad [3]$$

- *in probability*, for which the indicator function of a particular event should be employed. For example, we can ask the final state to reach a target region A with probability larger than or equal to $1-\alpha$. The auxiliary function is defined as

$$\phi_{\psi}(t_f, \mathbf{X}_f) = \mathbb{I}_A(\mathbf{X}_f) \; , \qquad [4]$$

where

$$\mathbb{I}_A(\mathbf{X}_f = \mathbf{x}_f) = \begin{cases} 1 & \text{if } \mathbf{x}_f \in A \\ 0 & \text{if } \mathbf{x}_f \notin A \; . \end{cases} \qquad [5]$$

From here, the constraint is formulated as

$$P(\mathbf{X}_f \in A) = \mathbb{E}[\mathbb{I}_A(\mathbf{X}_f)] \in \mathbf{\Phi}_{\psi} = [1-\alpha, 1] \; . \qquad [6]$$

- according to higher order moments, for which a specific ϕ and acceptable set $\mathbf{\Phi}$ are selected accorginly.

Among the possible choices, the specific form to minimise or constraint is a design decision, which translates the question *"what do we want to optimise?"*. Generally, different design choices lead to different optimisation results.

Given these premises, the optimal control under uncertainty to be tackled in this paper is formulated as

$$\begin{aligned} \min_{\mathbf{u}(t)\in U} \; & \mathbb{E}[\phi_J] \\ \text{s.t.} \; & \dot{\mathbf{X}} = f(t, \mathbf{X}, \mathbf{u}, \mathbf{D}) \\ & \mathbb{E}[\phi_{\mathbf{g}}(t, \mathbf{X}, \mathbf{u}, \mathbf{D})] \in \mathbf{\Phi}_{\mathbf{g}} \\ & \mathbb{E}[\phi_{\psi}(t_0, \mathbf{X}_0, t_f, \mathbf{X}_f)] \in \mathbf{\Phi}_{\psi} \; , \end{aligned} \qquad [7]$$

in a comparable way to the classical deterministic optimal control problem. The objective auxiliary function ϕ_J may depend on all the random variables when in the Bolza form. Nonetheless, its dependencies have not been written explicitly for conciseness of notation. Clearly, fully deterministic objective and constraints are still allowed. In that case, the corresponding function would fall back to Equation (1) notation. One common example is the case of the objective function only depending on the deterministic control.

This paper considers the case of deterministic dynamics affected by epistemic uncertainty, but with no intrinsic stochastic terms. Hence, for a fixed uncertain sample $\boldsymbol{\xi} \in \Omega_{\boldsymbol{\xi}}$, the resulting trajectory $\mathbf{x}(t)$ is deterministic.

In general, an optimal control problem is infinite dimensional with no closed form solution. This implies that a finite dimensional approximation is required to practically compute a solution with a numerical solver. Transcription methods are schemes which convert a dynamical optimal control problem into static constrained optimisation one, which is possible to solve with well-established numerical routines, e.g. NLP solvers. The next section will introduce a general transcription framework for optimal control problems with uncertainty in the form of Equation (7).

3. Generalised Direct Multiple Shooting

In the family of transcription techniques, direct methods aim at finding a sequence of control profiles which progressively decrease both the objective function and the constraint's violation. Within direct methods, the deterministic *multiple shooting* works by discretising the independent variable interval into n_i sub-segments $[t_i, t_{i+1}]$, which are then handled as independent. As a consequence, the state $\mathbf{x}_i$ at the beginning of each segment has to be treated as free variable. Within a segment, the control is parameterised using a functional form with free parameters $\boldsymbol{\beta}_i$, such that the control profile has the finite-dimensional representation $\mathbf{u}_i(t) = \mathcal{U}_i(t, \boldsymbol{\beta}_i)$. Once these free variable are set by the optimisation solver, each state $\mathbf{x}_i$ is integrated from t_i to t_{i+1}. Continuity constraints are added at the boundary of two adjacent segments to ensure continuity of the final solution.

However, when uncertainties are introduced, this pointwise method is not sufficient anymore. This section introduces a generalised shooting framework to deal with potential uncertainty in the initial conditions and dynamical model under general form objective and constraints as in Equation (7). A generalised intrusive polynomial expansion approach is used to represent the state variable evolution as function of the uncertain variables in a finite-dimensional space [7]. The resulting approach is named IPANeMA (Intrusive Polynomial Algebra aNd Multiple shooting Approach) for short.

3.1 Intrusive polynomial algebra propagation

The initial uncertain state sample domain $\Omega_{\mathbf{x}_0}$, induced by the random variable $\mathbf{Z}$, is bounded by a q-degree n_ξ-variables polynomial representation $\mathbf{P}_{\mathbf{X}_0} \in T_{q,n_\xi}$, where T_{q,n_ξ} is the resulting polynomial space. Since we are interested in the time evolution of this region, the polynomials approximating the state vector are function of all the n_ξ random variables involved:

$$\begin{aligned} \mathbf{P}_{\mathbf{X}}(t) &= \sum_{j=1}^{\mathcal{N}} \alpha_j(t) \mathcal{P}_j(\boldsymbol{\xi}) \\ \mathbf{P}_{\mathbf{X}}(t_0) &= \mathbf{P}_{\mathbf{X}_0} \; , \end{aligned} \qquad [8]$$

where $\mathcal{N} = \binom{n_\xi + q}{q}$ is the algebra dimension of the resulting functional space T_{q,n_ξ}, and $\mathcal{P}_j$ one of its multivariate polynomial basis. Keeping the ordering constant, each element in the polynomial space is uniquely identified by a specific vector of coefficients.

In this functional space, a set of algebraic operations between polynomials can be defined. Denoting with V and Z the multivariate polynomial approximations in T_{q,n_ξ} of v and z, the algebraic operation $\oplus = \{+, -, \cdot, /\}$ between real-valued functions has its correspondent $\otimes$ in the polynomial space:

$$v(\boldsymbol{\xi}) \oplus w(\boldsymbol{\xi}) \sim V(\boldsymbol{\xi}) \otimes W(\boldsymbol{\xi}) \in T_{q,n_\xi} \, . \qquad [9]$$

The result of the addition (or equivalently subtraction) of two elements of T_{q,n_ξ} is still an element of the same functional space. On the other hand, the result of multiplication needs to be truncated to restore the order q. Multiplication of two polynomials is an expensive operation, hence the approach used in this analysis will rely on a monomial basis, which guarantees lower computational costs. A composition rule is defined to handle division and other elementary functions such as trigonometric functions, exponents, logarithms, etc. These polynomial operations are implemented using the C++ overloading operator within the Strathclyde Mechanical and Aerospace Research Toolbox for Uncertainty Quantification (SMART-UQ) [7].

Given this set of operations, any integrator for the propagation of ordinary differential equations can be easily templated to work with generalised polynomial expansions. This feature enables to propagate the initial hyper-region $\mathbf{P}_{\mathbf{X}_0}$ through the dynamical system constraints in Equation (7).

The uncertain model parameters are handled equivalently. The parameter sample domain $\Omega_{\mathbf{d}}$, induced by the random variable $\mathbf{Z}$, is bounded by a constant multivariate polynomial $\mathbf{P}_{\mathbf{D}} \in T_{q,n_\xi}$, which is composed in $\mathbf{P}_{\mathbf{X}}(t)$ through the dynamical operations.

3.2 Multiple shooting framework

Intrusive polynomial algebra could be used directly to propagate $\mathbf{P}_{\mathbf{X}_0}$, the initial uncertain region, through the dynamics to obtain the final region $\mathbf{P}_{\mathbf{X}_f}$. The latter polynomial mapping could then be used to compute the objective function and the constraints (see section 3.3) to complete a loop of the optimisation process. This scheme can be seen as a generalised single shooting transcription method.

Despite its simplicity, this polynomial algebra-assisted single-shooting suffers severely from the renowned curse of dimensionality. Indeed, intrusive polynomial algebra scales badly with increasing number of uncertain variables, precisely as $(n_\xi + q)!\,/(n_\xi!\,q!\,)$. Hence, the cost of each algebraic operation involved in the numerical propagation grows dramatically.

When the uncertainties affect the system evolution sequentially (e.g. multi-phase trajectories, discretised control with disturbances, etc.), this issue can be mitigated adopting a multiple shooting scheme. In this development, we will consider the uncertain vector to be composed of uncertain initial conditions and model parameters $\mathbf{Z} = [\mathbf{X}_0, \mathbf{D}]$, and consequently an admissible realisation is $\boldsymbol{\xi} = [\mathbf{x}_0, \mathbf{d}]$. In particular, the uncertain parameter vector is partitioned as $\mathbf{D} = [\mathbf{D}_0, \mathbf{D}_1, \ldots, \mathbf{D}_i, \ldots]$, such that the parameters $\mathbf{D}_i$ affect the system only in the discretised interval $[t_i, t_{i+1}]$. From here, the goal is to develop a transcription method such that each subsegment can be treated independently, and consequently the algebra dimension in the i-th segment reduces to $n_{\xi_i} = n_s + d_i$, namely the number of the uncertain state variables $\mathbf{X}(t_i)$ at the beginning of the segment and the number of uncertain parameters $\mathbf{D}_i$ affecting the system evolution for $t \in [t_i, t_{i+1}]$. With this partition, the accumulation of uncertainties is avoided.

If this goal is achieved, each segment can be treated as a single shooting where the polynomial representation of the initial condition $\mathbf{P}^{(g)}_{\mathbf{X}_i}$ at t_i is propagated to $\mathbf{P}^{(p)}_{\mathbf{X}_{i+1}}$ at t_{i+1} under the effect of uncertain parameters $\mathbf{D}_i$ only.

3.2.1 Reinitialisation Approach

The main difficulty of the proposed discretisation arises from the necessity to impose continuity conditions between two hyper-dimensional sets at the boundary of two adjacent segments. For polynomial algebra, this continuity requirement could be translated into a reinitialisation approach: the propagated polynomial representation $\mathbf{P}^{(p)}_{\mathbf{X}_{i+1}}$, function of $\mathbf{X}_i$ and $\mathbf{D}_i$, is reinitialised to the polynomial $\mathbf{P}^{(g)}_{\mathbf{X}_{i+1}}$, initially function of $\mathbf{X}_{i+1}$ only. Indeed, it is worth stressing that the terms interacting with $\mathbf{D}_{i+1}$ arise only during the propagation, i.e. for $t \in (t_{i+1}, t_{i+2}]$.

However, in general, it is not possible to fully describe a multivariate polynomial with another polynomial of smaller number of (initial) variables and same degree. To overcome this intrinsic issue, the reinitalised polynomial will be constructed to bound the propagated one. In particular, the propagation phase is carried out as follows:

1. Initialise $i = 0$, $\mathbf{P}^{(g)}_{\mathbf{X}_i} = \mathbf{P}_{\mathbf{X}_0}$;
2. Propagate uncertainty region $\mathbf{P}^{(g)}_{\mathbf{X}_i}$ at t_i to $\mathbf{P}^{(p)}_{\mathbf{X}_{i+1}}$ at t_{i+1} through intrusive polynomial algebra;
3. Compute lower $\mathbf{XL}_{i+1}$ and upper $\mathbf{XU}_{i+1}$ polynomial ranges of $\mathbf{P}^{(p)}_{\mathbf{X}_{i+1}}$;
4. Reinitialise uncertainty region $\mathbf{P}^{(g)}_{\mathbf{X}_{i+1}}$ as hyper-box with range $\mathbf{XL}_{i+1}$ and $\mathbf{XU}_{i+1}$;
5. Update $i = i + 1$ and repeat steps 2-5 while $i < n_i$

In this way, all the possible state realisations are included, therefore granting any pointwise trajectory continuity. This simple approach comes at the expense of propagating larger regions than strictly needed. Graphically, this procedure can be visualised as in Figure 1 for a two-dimensional example.

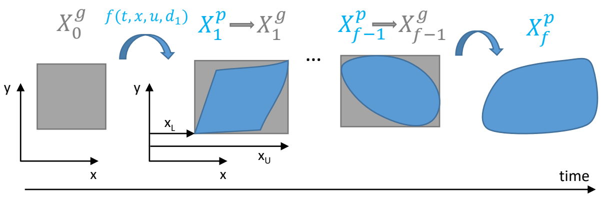

Fig. 1: Graphical sketch of intrusive polynomial propagation approach for the generalised multiple shooting. The gray boxes represent the reinitalisation hyper-boxes, whereas the blue regions depict the propagated polynomials.

The result of this propagation approach is a chain of polynomial surrogates describing the state at time t_{i+1} as function of the state at time t_i and uncertain parameters within the corresponding interval. Therefore, a recursive polynomial surrogate of the final state $\mathbf{X}_f$ is available as a function of the initial conditions $\mathbf{X}_0$ and all the uncertain parameters $\mathbf{D}$. At this step however, the hyper-box reinitialisation caused the final state surrogate to be an over-estimation of the true final uncertain space in general.

The routine to recover the actual terminal region is achieved by successive sampling. In the simplest form, the final hyper-region computation algorithm is described as follows:

1. Initialise $i = 0$
2. Sample the initial uncertain space:
   $\mathbf{x}_i^{(s)} \in \Omega_{\mathbf{x}_0}$
3. Sample the $i-$uncertain parameter space:
   $\mathbf{d}_i \in \Omega_{\mathbf{d}_i}$
4. Propagate each particle from t_i to t_{i+1} with polynomial surrogate $\mathbf{P}^{(p)}_{\mathbf{X}_{i+1}}$:
   $(\mathbf{x}_i^{(s)}, \mathbf{d}_i) \rightarrow \mathbf{x}_{i+1}$
5. Each response state is scaled within the polynomial input domain using the same ranges $\mathbf{XL}_{i+1}$ and $\mathbf{XU}_{i+1}$ used for polynomial reinitialisation:
   $\mathbf{x}_{i+1} \rightarrow \mathbf{x}_{i+1}^{(s)}$
6. Update $i = i + 1$ and repeat steps 3-5 while $i < n_i$ (skip step 5 for last iteration)

A graphical depiction of the recovery strategy is plotted in Figure 2.

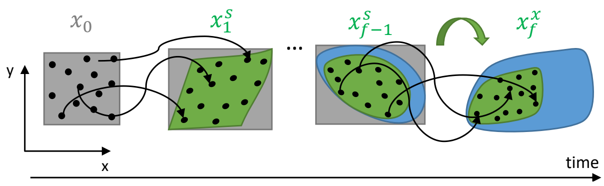

Fig. 2: Graphical sketch of the recovery approach for the generalised multiple shooting. The gray boxes represent the reinitialisation hyper-boxes, the blue regions depict the propagated polynomials, while the grey areas symbolise the true uncertainty regions reconstructed by the black samples.

It is worth noting that the samples can be propagated at any intermediate time of interest $\bar{t} \in (t_i, t_{i+1})$ without the need of further discretisation. Trivially, an intermediate polynomial can be saved during the propagation phase, and the samples $(\mathbf{x}_i^{(s)}, \mathbf{d}_i)$ propagated to $\bar{t}$ through it. Hence, the general result of this approach is a surrogate model $\widetilde{F}_{\bar{t}} : \Omega_{\mathbf{x}_0} \times \Omega_{\mathbf{d}_{0:i}} \rightarrow \mathbb{R}^{n_s}$ that maps the uncertain initial conditions and parameters to the state vector at any time $\bar{t}$. The uncertain parameter space $\Omega_{\mathbf{d}_{0:i}}$ takes into account only the model uncertain parameters $\mathbf{D}_{0:i} = [\mathbf{D}_0, \ldots, \mathbf{D}_i]$ which entered the system not later than the time of interest.

With the developed generalised multiple shooting approach, the uncertain space dimensionality is kept as low as possible in each discretisation interval. Furthermore, the outer reinitialisation strategy intrinsically implies pointwise trajectory continuity. This property removes the need of explicit defect constraints and intermediate free variables in the transcription, hence reducing the dimensionality of the associated constrained optimisation problem. The only free variables to be optimised are the control parameters in each sub-segment. Another powerful upside of this method is that the sampling in steps 2-3 is agnostic to the probability distribution nature. Therefore, the method is equally suitable for any probability distribution.

The missing bit for a complete transcription scheme for uncertain optimal control problems is the computation of expectations of generic functions as introduced in Section 2.

3.3 *Objective and constraint computation*

In the problem statement development, the general expectation formulation was chosen to represent a wide class of possible constraints and objective functions. Indeed, the generic function ϕ is considered a design choice to be selected according to the sought quantity of interest. Furthermore, the expectation can be computed at any fixed-time $\bar{t}$, or even on a time-span of interest.

For the case considered in this paper, the expectation over the uncertain distribution can be computed by means of the surrogate model $\widetilde{F}_{\bar{t}}$, approximating the true mapping, as defined in the previous section:

$$\begin{aligned}\mathbb{E}[\phi(\mathbf{X}(\bar{t}))] &\approx \mathbb{E}[\phi(\widetilde{F}_{\bar{t}}(\mathbf{Z}_{0:i}))] \\ &= \int_{\Omega_{\boldsymbol{\xi}_{0:i}}} \phi(\widetilde{F}_{\bar{t}}(\boldsymbol{\xi}_{0:i}))p(\boldsymbol{\xi}_{0:i})d\boldsymbol{\xi}_{0:i} \ ,\end{aligned} \quad [10]$$

where the random variable is defined as $\mathbf{Z}_{0:i} = [\mathbf{X}_0, \mathbf{D}_{0:i}]$, its realisation as $\boldsymbol{\xi}_{0:i} = [\mathbf{x}_0, \mathbf{d}_{0:i}]$, and the uncertain domain $\Omega_{\boldsymbol{\xi}_{0:i}} = \Omega_{\mathbf{x}_0} \times \Omega_{\mathbf{d}_{0:i}}$.

In the general case, this integral has no closed-form solution and numerical techniques shall be applied. Exploiting the inexpensive surrogate map, two main sample-based alternatives can be considered:

- Monte Carlo methods for the estimation of the expectation: the samples in steps 2-3 of the recovery strategy shall be drawn according to the uncertain variables propability distributions $p(\mathbf{x}_0)$ and $p(\mathbf{d}_{0:i})$. Then, the expected value can be computed as:

$$\mathbb{E}[\phi(\widetilde{F}_{\bar{t}}(\mathbf{Z}_{0:i}))] \approx \frac{1}{N}\sum_{j=1}^{N} \phi(\widetilde{F}_{\bar{t}}(\boldsymbol{\xi}^{(j)}_{0:i})) \ ; \quad [11]$$

- Quadrature schemes for the computation of the integral: the samples are chosen according to a quadrature scheme with corresponding weights w_j, resulting in the integral approximation

$$\int_{\Omega_{\boldsymbol{\xi}_{0:i}}} \phi(\widetilde{F}_{\bar{t}}(\boldsymbol{\xi}_{0:i}))p(\boldsymbol{\xi}_{0:i})d\boldsymbol{\xi}_{0:i} \approx \sum_{j=1}^{N} w_j \phi(\widetilde{F}_{\bar{t}}(\boldsymbol{\xi}^{(j)}_{0:i}))p(\boldsymbol{\xi}^{(j)}_{0:i}) \ . \quad [12]$$

The latter scheme shall be preferred when the probability distribution is complex to sample but rather easy to evaluate, or when the expectation should be evaluated for a set of different density distributions.

It is worth suggesting that when this approximation is included in a NLP local optimisation solver with finite-difference derivative computation, sampling grids should be kept constant within a major NLP step. Indeed, if the grids are varied between the reference and perturbed propagations, the derivative values would result highly inaccurate, leading the optimiser to compute unreliable descent directions.

Although probability constraints (or equivalently objectives) are an intuitive and general tool to impose conditions on random variables, the indicator function discontinuity introduces important numerical challenges when coupled with derivative-based optimisers. Indeed, although in theory the expectation operator should smooth the discontinuity, the final constraint is usually computed by sample-based numerical approximations (as in eq. (11) or (12)), which cause the constraint response to be piecewise constant with discontinuous jumps. Local derivative-based optimisers cannot cope with such functions.

To overcome this numerical issue, the developed tool substitutes the indicator function by a smoother approximation obtained by convolution [2], a general technique to modify the shape of a function according to a smoothing function h. To simplify the convolution application to a scalar function of scalar variable, the membership condition of a sample belonging to a region A will be expressed by an auxiliary continuous scalar function $\eta_A : \mathbb{R}^{n_s} \rightarrow \mathbb{R}$ such that:

$$\begin{cases} |\eta_A(\mathbf{X} = \mathbf{x})| \leq 1 & \text{if } \mathbf{x} \in A \\ |\eta_A(\mathbf{X} = \mathbf{x})| > 1 & \text{if } \mathbf{x} \notin A \ . \end{cases} \quad [13]$$

Hence, the indicator function previously defined is equivalent to $\mathbb{I}_A(\mathbf{X}) = \mathbb{I}_{[-1,+1]}(\eta_A(\mathbf{X}))$. Now, for a state realisation $\mathbf{x}$, the convolution of the indicator function with a smoothing function h results in the function:

$$\begin{aligned}\mathbb{I}^{(r)}_{[-1,+1]}(\eta_A(\mathbf{X} = \mathbf{x})) &= \int_{-\infty}^{+\infty} \mathbb{I}_{[-1,+1]}(y)\frac{1}{r}h\Big(\frac{\eta_A(\mathbf{x}) - y}{r}\Big)dy \\ &= \int_{-1}^{+1} \mathbb{I}_{[-1,+1]}(y)\frac{1}{r}h\Big(\frac{\eta_A(\mathbf{x}) - y}{r}\Big)dy \ ,\end{aligned} \quad [14]$$

with $r > 0$ a small positive scaling parameter. The integration interval is restricted to the interval $[-1, +1]$ because of the function η_A definition. The function h is chosen to result in a proper approximation of the original function. Specifically, $h : \mathbb{R} \rightarrow \mathbb{R}$ shall be non-negative, symmetric, with an unique maximum in 0, and it shall integrate to 1. These properties imply $\lim_{r\rightarrow 0} h(\cdot/r)/r = \delta$, the Dirac delta. Hence, for $r \rightarrow 0$ the convolution result tends to the original indicator function [2].

It is worth mentioning that objective and constraints not falling under the expectation formulation are possible. As an example, if we are interested in the final state ending in the target region A, one alternative is to constraint the maximum deviation to be under a set threshold, e.g. $\max(\{|\eta_A(\mathbf{x}^{(j)})|: j = 1, \ldots, N\}) \leq \rho$. Similar objective and constraint functions are rather test case specific and hence not explicitly accounted for, but nonetheless possible in the developed framework.

3.4 Transcribed problem

IPANeMA is meant to transcribe the optimal control under uncertainty into a finite-dimensional constrained optimisation problem. In the current implementation, the constrained optimisation is solved using WORHP as nonlinear programming solver [12]. Differently from a deterministic multiple shooting, the resulting transcribed problem is dense and low-dimensional, as no intermediate state guesses or explicit continuity constraints have been introduced. Hence, the control parameters $\boldsymbol{\beta}_i$ per each sub-segment are the only free variables.

To avoid a new expensive intrusive polynomial propagation each time a free variable vector is set within the optimisation routine, the deterministic control can be expanded in polynomial representation as well. Specifically, the control parameters domain $\Omega_{\boldsymbol{\beta}_i}$ can be bounded by a time-static multivariate polynomial $\mathbf{B}_i \in T_{q,n_{\xi_i}}$, where the number of uncertain variables n_{ξ_i} should be increased accordingly. Then, the polynomial control profile in each interval follows according to the parameter-control relationship $\mathbf{U}_i(t) = \mathcal{U}_i^{(p)}(t, \mathbf{B}_i)$, where by $\mathcal{U}_i^{(p)}$ it is intended the corresponding polynomial operator of $\mathcal{U}_i$. For a fixed value $\boldsymbol{\beta}_i \in \Omega_{\boldsymbol{\beta}_i}$, the control polynomial representation reduces to the deterministic control $\mathbf{u}_i(t) = \mathcal{U}_i(t, \boldsymbol{\beta}_i)$. With this procedure, only one uncertainty polynomial propagation is needed, and it can be precomputed before the optimisation cycle.

4. Test case

The developed method is applied to the optimisation of a space trajectory. The goal of the set-up mission is to compute the optimal-fuel rendezvous to the near-Earth asteroid 99942 Apophis (2004 MN_4) with a low-thrust spacecraft departing from the Earth sphere of influence. The initial date of the interplanetary leg is 22/10/2026 for a total time of flight of 628 days. The engine has maximum thrust of $T_{max} = 53$ mN, for a spacecraft initial mass of $m_0 = 644.3$ kg. The reference mission employs an initial excess velocity of magnitude $v_\infty^{ref} = 3.34$ km/s and azimuth angle $\alpha_\infty^{ref} = 35.17$ deg deviation from the $x-$axis in the Earth-centered inertial reference frame.

As this case is meant at assessing the suitability and performance of the developed method to preliminary design of robust space trajectories, a few simplifying assumptions will be used. Namely, only the Sun pull is considered as gravitational force, and only the planar trajectory is studied.

The spacecraft planar motion is described in equinoctial coordinate system for the in-plane coordinates [3]:

$$\begin{aligned} &a \\ &P_1 = e \sin(\Omega + \omega) \\ &P_2 = e \cos(\Omega + \omega) \, . \end{aligned} \qquad [15]$$

The governing equations are expressed in the Gauss' planetary form in a radial-transverse reference frame. The fast angular variable L, i.e. the true longitude, can be used as independent variable to replace the time evolution. Under the enforced assumption of a low-thrust control magnitude significantly smaller than the local gravitational force, the resulting system of equations is [13]:

$$\begin{aligned} \frac{da}{dL} &= \frac{2a^3B^2}{\mu}\left[\frac{P_2 \sin L - P_1 \cos L}{\Phi^2(L)} f_R + \frac{1}{\Phi(L)} f_T\right] \\ \frac{dP_1}{dL} &= \frac{B^4 a^2}{\mu}\left[-\frac{\cos L}{\Phi^2(L)} f_R + \left(\frac{P_1 + \sin L}{\Phi^3(L)} + \frac{\sin L}{\Phi^2(L)}\right) f_T\right] \\ \frac{dP_2}{dL} &= \frac{B^4 a^2}{\mu}\left[+\frac{\sin L}{\Phi^2(L)} f_R + \left(\frac{P_2 + \cos L}{\Phi^3(L)} + \frac{\cos L}{\Phi^2(L)}\right) f_T\right] , \end{aligned} \qquad [16]$$

where $B = \sqrt{(1 - P_1^2 - P_2^2)}$ and $\Phi(L) = 1 + P_1 \sin L + P_2 \cos L$. The radial and transverse control components are controlled in terms of acceleration magnitude and azimuth angle:

$$\mathbf{f} = \begin{bmatrix} f_R \\ f_T \end{bmatrix} = \begin{bmatrix} \epsilon \sin \alpha \\ \epsilon \cos \alpha \end{bmatrix} . \qquad [17]$$

For the deterministic reference case, the terminal constraint is imposed by requiring the matching of the spacecraft final state with Apophis in equinoctial elements at the time of arrival. The optimal-fuel objective to minimise is the trajectory ΔV. The first guess for the reference trajectory has been generated by the deterministic single-shooting tool FABLE (Fast Analytical Boundary-value Low-thrust Estimator) [4], which transcribes the optimal control problem into a sequence of coast and constant thrust arcs. In FABLE, the dynamics in Equation (16) is analytically propagated using a first-order expansion in the perturbing control acceleration [13]. The resulting objective value is $\Delta V = 2.0318$ km/s.

In the following, a fictitious scenario is considered to introduce uncertainty. Telemetry has reported a partial failure in the interplanetary orbit injection phase, but accurate information about the new spacecraft state after failure is not available yet. From the partial data, the uncertainty has been traced back to the injection velocity. Specifically, the azimuth angle error $\Delta\alpha_\infty^{ref}$ is modelled with zero-mean normal distribution about the reference value with $1\sigma = 0.25$ deg, while the magnitude velocity error Δv_∞^{ref} is modelled with a zero-mean reversed Gaussian tail distribution with $1\sigma = 25$ m/s, where only negatives values are admissible. The uncertain vector, composed as $\mathbf{Z} = [\Delta\alpha_\infty^{ref}, \Delta v_\infty^{ref}]$, induces uncertainty in the initial conditions in planar equinoctial elements. The resulting initial set of uncertainty in equinoctial elements is displayed in Figure 3. The sample colour indicates the associated probability density value according to the distribution definition above. Superimposed there are the expected value of the distribution and 1σ ellipsoid projections according to the samples covariance. By construction, the distribution is asymmetric, and the expected value significantly deviates from the mode of the distribution in the dark red coloured area. Hence, the first two moments are not fully representative of the real distribution.

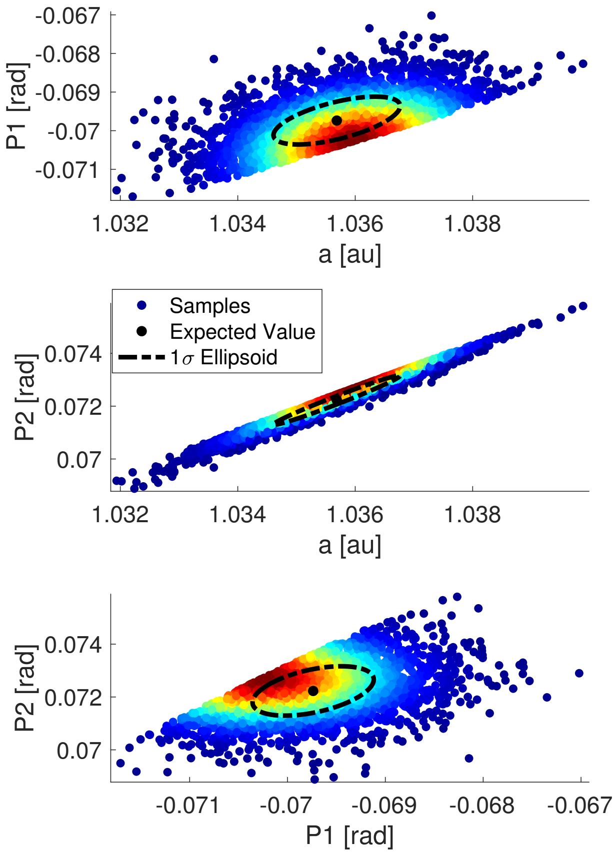

Fig. 3: Initial set of uncertainty in equinoctial elements as a result of uncertainty in the excess velocity with superimposed the expected value and 1σ ellipsoid resulting from to the sample distribution.

To counteract this partial failure, it has been decided to compute a modified control profile able to steer the initial uncertain region into a target zone around the asteroid state. With this approach, the goal is to limit the required correction manoeuvres, either in-flight or at arrival, when accurate measurements will be available.

This optimal control problem under uncertainty is formulated by substituting the final boundary condition with a probability constraint. Specifically, the probability of the final uncertain state to belong to a target ellipsoid T is required to be above a given threshold. The probability constraint is defined thanks to the auxiliary positive continuous function that, for a given state realisation $\mathbf{x}$, is defined as

$$\eta_T(\mathbf{X}=\mathbf{x}) = (\mathbf{x}-\boldsymbol{\mu})^T M(\mathbf{x}-\boldsymbol{\mu}) = \begin{cases} \leq 1 & \text{if } \mathbf{x} \in T \\ > 1 & \text{if } \mathbf{x} \notin T \end{cases}, \quad [18]$$

where $\boldsymbol{\mu}$ is the target ellipsoid center, i.e. the asteroid state at the time at arrival, while M is defined such that its eigenvectors are the ellipsoid principal axes and its eigenvalues are the reciprocals of the semi-axes squared. In this test case, M is defined as symmetric, while its eigenvalues follow from the set semi-axes $10^{-3} \cdot [2.2\,\text{au}, 2.0\,\text{rad}, 3.7\,\text{rad}]$. The probability threshold has been set to 95%.

To solve this optimal control under uncertainty, 5-degree Chebyshev polynomials are employed for the intrusive propagation, which have been already used in aerospace applications because of their superior global convergence properties [1, 10]. As regards the transcription, the following settings are used: 6 discretisation intervals; piecewise constant control; 200 samples for Monte Carlo approximation of Equation (10); biquadratic function

$$h(z) = 15(1-z^2)^2 \mathbb{I}_{[-1,+1]}/16$$

for the convolution operator [2]. The resulting control profile and the first guess are shown in Figure 4.

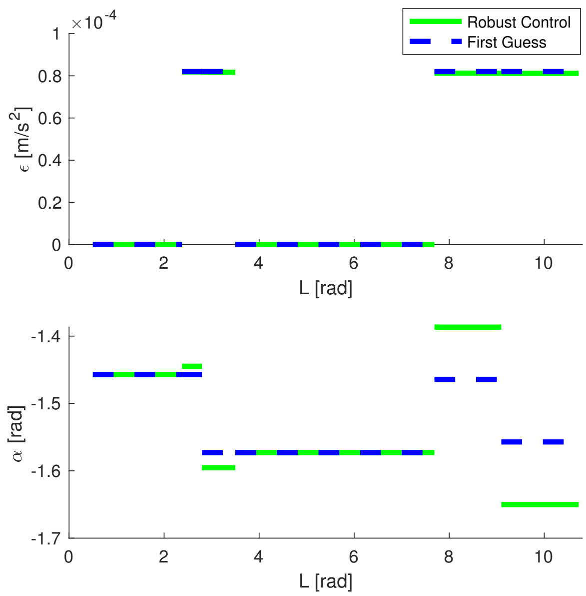

Fig. 4: Optimised robust control profile components versus first guess control.

While the thrust magnitude is essentially unaltered, implying a robust ΔV objective value, the thrust angle changed significantly to steer the final region within the required target ellipsoid. The optimisation routine finished with a probability of 95.3% associated to the final state arriving within the target region, improving the trajectory reliability from the value of 22.7% associated to the first guess control.

The resulting L-evolution of the uncertain region is shown in Figure 5, where the trajectory of the initial expected value is highlighted with a black dotted line.

The optimised final uncertain area projected in two dimensional planes is displayed in Figure 6 together with the target ellipsoid. It is worth noting how the shape and

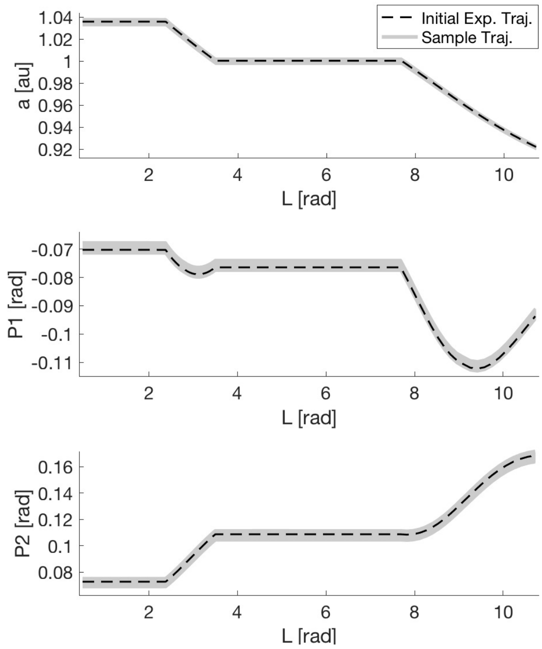

Fig. 5: Uncertain set evolution as function of independent variable L. The trajectory of the initial expected value is reported with a black dotted line.

extension of the uncertain region have remained essentially unaffected during the dynamics propagation. The reason behind this effect is twofold. First, the equinoctial state variables are integrals of motion of the two body problem, hence only partially affected by the small perturbative force in a limited time-span. Technically, the dynamical system is not fully controllable because of the limited low-thrust control authority. Second, for the given dynamics and initial uncertain set, it is not possible to make all the possible trajectories converge within an arbitrary final region with a single open-loop control profile.

For validation, two key approximations employed in the optimisation routine are checked, namely the surrogate propagation by intrusive polynomial algebra, and the probability approximation by convolution on a rather small set of samples. For the former, 10^5 samples drawn from the initial distribution are reintegrated with the refined control profile, with both the polynomial surrogate and a numerical fourth order Runge-Kutta integrator for comparison. The resulting root-mean-square error is in the order of 10^{-5} per state component, which confirms that intrusive polynomial algebra produces a satisfactory propagation approximation. As for the latter, the final probability is then computed with the indicator function directly, i.e. without the convolution approximation, on the extended validation set of RK4 propagated samples. The resulting probability is 94.0%, slightly lower than the value obtained in the optimisation loop. This discrepancy results partly from the convolution approximation, but mainly from the different orders of magnitude of uncertainty samples employed. More uncertain samples can be used in the optimisation loop to improve the solution accuracy. Nonetheless, for the current test case, the obtained results are considered highly satisfactory.

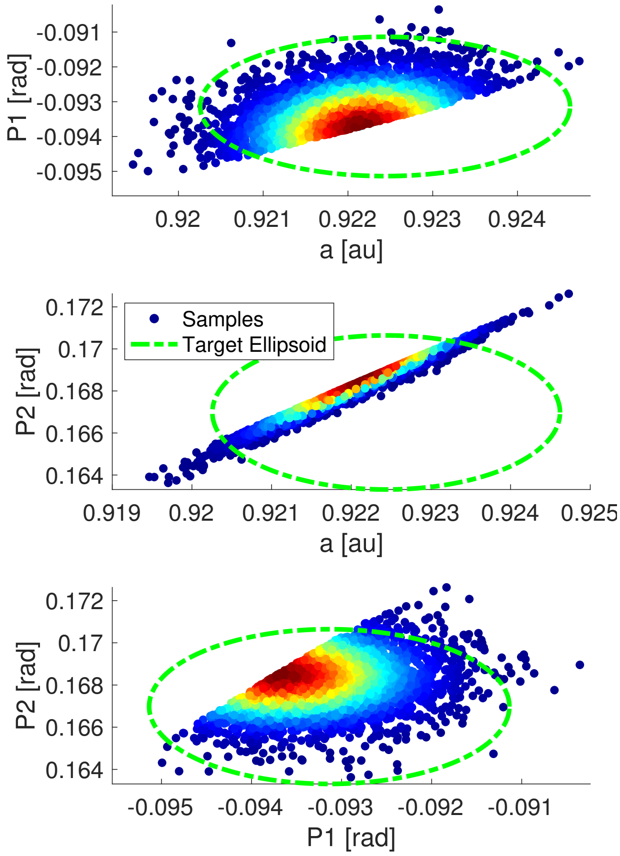

Fig. 6: Final set of uncertainty in equinoctial elements resulting from optimised control profile and final target region.

The reliability percentages associated to different settings are summarised in the following table.

Table 1: Probability of final target matching for different settings.

Setting	Samples	Convolution	$P(\mathbf{X}_f \in T)$
First guess	$2 \cdot 10^2$	Yes	22.7%
Guess validation	10^5	No	23.3%
Robust solution	$2 \cdot 10^2$	Yes	95.3%
Robust validation	10^5	No	94.0%

5. Conclusions

This paper presented the development of an intrusive polynomial assisted multiple shooting transcription for the solution of optimal control problems affected by uncertainty.

First, a general problem statement is introduced, which reformulates in expectation the constraints and objective functions affected by uncertainty. Thanks to the intermediate auxiliary function, this general notation is shown to be flexible in describing a variety of formulations, e.g. expected value, probability, different statistics, etc.

Then, IPANeMA is presented as tool for the transcription of the infinite-dimensional optimal control problem under uncertainty into a constrained optimisation possible to solve with a NLP solver. IPANeMA integrates a novel multiple shooting framework with a generalised intrusive polynomial expansion to represent and propagate uncertainty regions. One approach based on reinitialisation by bounding hyper-boxes is proposed, which reduces the intrusive algebra dimension, and intrinsically handles the continuity conditions between two adjacent segments with no need of additional constraints. A sample-based recovery strategy is employed as the reinitialisation approach requires propagating wider regions than the actual one.

The developed framework is capable of handling both uncertainty in the initial state and in the model parameters. Furthermore, IPANeMA is suitable to work with a large variety of uncertainty models, hence it is not restricted to purely Gaussian, uniform or other basic probability distribution families. Indeed, the sample-based strategy developed naturally suits the approximate computation of constraints and objective functions formulated in expectation. A specific convolution approach is integrated to deal with numerical complexities in the optimisation loop introduced by the probability formulation.

Finally, IPANeMA is successfully applied to the robust optimisation of a low-thrust rendezvous trajectory to the near-Earth asteroid 99942 Apophis. In particular, the found control law steers the initial uncertain region to a target ellipsoid around the asteroid, with a probabilistic constraints satisfied within the required threshold.

As for future developments, the main challenge is to integrate observations within IPANeMA for the computation of an updated control law which takes into account the measurement information. The developed framework seems to suit naturally an interface with a particle filter.

Acknowledgements

This work was partially funded by the European Commission's H2020 programme, through the H2020-MSCA-ITN-2016 UTOPIAE Marie Curie Innovative Training Network, grant agreement 722734.

References

[1] C. O. Absil, R. Serra, A. Riccardi, and M. Vasile. De-orbiting and re-entry analysis with generalised intrusive polynomial expansions. *67th International Astronautical Congress*, 2016.

[2] L. Andrieu, G. Cohen, and F. Vzquez-Abad. Stochastic programming with probability constraints. *arXiv preprint arXiv:0708.0281*, 2007.

[3] R. Broucke and P. J. Cefola. On the equinoctial orbit elements. *Celestial Mechanics*, 5(3):303–310, 1972.

[4] M. Di Carlo, J. M. Romero Martin, and M. Vasile. CAMELOT: Computational-Analytical Multi-fidElity Low-thrust Optimisation Toolbox. *CEAS Space Journal*, 10:25–36, 2018.

[5] P. Di Lizia, R. Armellin, F. Bernelli-Zazzera, and M. Berz. High order optimal control of space trajectories with uncertain boundary conditions. *Acta Astronautica*, 93:217–229, 2014.

[6] E.D. Gustafson. *Stochastic Optimal Control of Spacecraft*. PhD thesis, The University of Michigan, 2010.

[7] C. Ortega, A. Riccardi, M. Vasile, and C. Tardioli. SMART-UQ: Uncertainty Quantification Toolbox for Generalised Intrusive and non Intrusive Polynomial Algebra. In *6th International Conference on Astrodynamics Tools and Techniques*, 2016.

[8] N. Ozaki, S. Campagnola, C.H. Yam, and R. Funase. Differential dynamic programming approach for robust-optimal low-thrust trajectory design considering uncertainty. In *25th International Symposium on Space Flight Dynamics*, 2015.

[9] M. D. Rayman and S.N. Williams. Design of the first interplanetary solar electric propulsion mission. *Journal of Spacecraft and Rockets*, 3:589–595, 2002.

[10] A. Riccardi, C. Tardioli, and M. Vasile. An intrusive approach to uncertainty propagation in orbital mechanics based on Tchebycheff polynomial algebra. *Advances in Astronautical Sciences*, pages 707–722, 2015.

[11] A. Richards and J. How. Robust stable model predictive control with constraint tightening. In *American Control Conference 2006 IEEE*, 2006.

[12] D. Wassel, F. Wolff, J. Vogelsang, and C. Buskens. The ESA NLP-Solver WORHP - Recent Developments and Applications. In *Modeling and Optimization in Space Engineering*, pages 85–110. Springer, New York, 2013.

[13] F. Zuiani and M. Vasile. Extended analytical formulas for the perturbed Keplerian motion under a constant control acceleration. *Celestial Mechanics and Dynamical Astronomy*, 121:275–300, 2015.