A Distributed Robust Optimal Control Framework Based on Polynomial Chaos

Patrick Piprek, Sébastien Gros, and Florian Holzapfel

Abstract This study is concerned with the development of a robust open-loop optimal control (ROC) framework that distributes different generalized polynomial chaos (gPC) sub-problems from the non-intrusive stochastic collocation (SC) method. This distributed open-loop optimal control (DOC) approach yields a number of smaller open-loop optimal control problems (OCPs) that can be solved independently of each other and are only connected by a small number of connection variables. These connection variables are introduced based on the specifics of the used cost and constraint functions and describe the coupling in the gPC expansion when e.g., calculating the variance. Overall, the definition as a DOC problem yields a faster and more reliable way to solve the ROC problem than by a full, connected problem. Here, the study shows the applicability of the proposed method in an air race example with the optimization of mean values and variances.

1 Introduction

With increasing computational power, the calculation of robust optimal trajectories became a popular research topic [2, 6, 8, 9]. Especially the introduction of gPC by XIU AND KARNIADAKIS in 2002 [14] is widely applied as it provides an efficient method to introduce parametric uncertainties in OCPs. Overall, the gPC method

Patrick Piprek
Technical University of Munich, Institute of Flight System Dynamics, Boltzmannstraße 15, 85748 Garching bei Muenchen, e-mail: patrick.piprek@tum.de

Sébastien Gros
Norwegian University of Science and Technology, Department of Engineering Cybernetics, O. S. Bragstads Plass 2D, 7034 Trondheim, e-mail: sebastien.gros@ntnu.no

Florian Holzapfel
Technical University of Munich, Institute of Flight System Dynamics, Boltzmannstraße 15, 85748 Garching bei Muenchen, e-mail: florian.holzapfel@tum.de

is a spectral representation of uncertainties. This means that the approach utilizes an expansion with deterministic expansion coefficients and orthogonal polynomials that depend on the uncertainties [13].

This study concentrates on the development of a bi-level approach using a DOC technique: In this context, gPC was already applied within a bi-level, non-distributed open-loop optimal control (OC) framework for the optimization of noise minimal approach trajectories with wind uncertainties [10]. Therefore, it is suited for the algorithm developed within this research project.

DOC is a powerful tool in OC: It defines a methodology to split up a large OCP into smaller OCPs that are coupled by connection variables. These smaller problems are then easier to solve because they can be solved independently of each other. On the downside, the problems must be solved multiple times, instead of only once as for the large scale problem. Still, it is generally easier regarding both time and complexity, to solve multiple smaller, parallelized problems instead of a single large problem (which might not even be possible due to e.g., the size). Thus, this study applies the methods of DOC to ROC with gPC. These problems can be distributed fairly well due to the nature of the gPC expansion.

An overview on already developed gPC DOC frameworks, mainly for model predictive control (MPC) and with intrusive changes to the dynamic model by gPC, can be found in [3, 7]: At first, study [3] shows a combination of DOC in the context of MPC with gPC. The authors apply there method to linear systems and use an intrusive gPC formulation to rewrite the deterministic equations into a stochastic form. Their goal is to solve a linear probabilistic constraint. It should be noted that this is opposite to the formulation of this study that conserves the deterministic baseline formulation and thus does not require and alteration of the original problem. Study [7] introduces a similar idea in the context of stochastic DOC for non-convex problems. The authors again use the intrusive reformulation of the gPC method to change the problem formulation from the deterministic to the stochastic domain.

To show the implementation of a DOC framework with gPC, this study is organized as follows: In Section 2 an overview on different used methodologies to implement the DOC framework is given. The optimization model is introduced in Section 3. The results for different DOC reference scenarios are given in Section 4. Within Section 5 conclusive remarks on the implementation aspects of the proposed algorithm and an outlook are stated.

2 Methodology

This section gives an overview on the methods used within this study as well as some characteristics of their implementation: Therefore, Subsection 2.1 introduces the general OCP formulation. Then, Subsection 2.2 gives a review of the gPC method and how to calculate the statistics for the DOC formulation and Subsection 2.3 concludes with an overview of the DOC methodology developed in this study.

2.1 Direct Open-Loop Optimal Control

We define the OCP (it should be reminded here that we consider open-loop OC here) for this study as follows [1]:

$$\begin{aligned} \min_{\mathbf{x},\mathbf{u},\mathbf{p},t_f} \quad & J\left(\mathbf{x},\mathbf{u},\mathbf{p},t_f\right) \\ \text{s.t.} \quad & \mathbf{f}(\mathbf{x},\mathbf{u},\mathbf{p};\theta) = \dot{\mathbf{x}}, \\ & \mathbf{c}(\mathbf{x},\mathbf{u},\mathbf{p}) \leq \mathbf{0}, \\ & \boldsymbol{\psi}(\mathbf{x},\mathbf{u},\mathbf{p}) = \mathbf{0} \end{aligned} \tag{1}$$

The OCP is depending on the states of the system $\mathbf{x} \in \mathbb{R}^{n_x \times 1}$, the controls $\mathbf{u} \in \mathbb{R}^{n_u \times 1}$, optimizable time-invariant parameters $\mathbf{p} \in \mathbb{R}^{n_p \times 1}$, and the final time t_f. Further on, these optimization variables are also combined in the vector $\mathbf{z} = \left[t_f, \mathbf{p}^T, \mathbf{x}^T, \mathbf{u}^T\right]^T$. The scalar valued cost function J is expressed as a BOLZA COST FUNCTIONAL as follows:

$$J\left(\mathbf{x},\mathbf{u},\mathbf{p},t_f\right) = e\left(\mathbf{x}\left(t_f\right),\mathbf{u}\left(t_f\right),\mathbf{p},t_f\right) + \int_0^{t_f} L\left(\mathbf{x},\mathbf{u},\mathbf{p}\right) \mathrm{d}t \tag{2}$$

The objective is to minimize the cost functional consisting of the Mayer term e that depicts the cost index at the final point in time and the Lagrange term L that describes the cost index over the optimization time interval.

The OCP in Eq. 1 is subject to the following constraints: First of all, the state dynamics $\dot{\mathbf{x}}$ ensuring a feasible trajectory must be fulfilled. Additionally, inequality path and point constraints $\mathbf{c}$ as well equality path and point constraints $\boldsymbol{\psi}$ are required (e.g., boundary conditions or load factor limits).

Furthermore, compared to the normally used OCP formulation, we introduce within this study that the state dynamics are subject to uncertainties θ for which we know the probability density function (pdf). These uncertainties are symbolized by θ in the state dynamics equations.

Generally, the OCP stated in Eq. 1 is solved using direct methods available from the OC toolbox FALCON.M [11]. Direct methods first of discretize the problem into a nonlinear programming problem (NLP) and afterwards optimize this discretized problem [1]. Here, a transcription by trapezoidal collocation is used [1]. To solve the discretized OCP, we use the gradient-based NLP solver IPOPT [12].

2.2 Generalized Polynomial Chaos

This section gives an overview of the gPC method: At first, Subsection 2.2.1 introduces the basics of the gPC method. Subsection 2.2.2 then describes the SC method that is used to calculate the expansion coefficients. The general calculation of the

statistical moments is presented in Subsection 2.2.3, while Subsection 2.2.4 extends these results to the DOC framework of this study.

2.2.1 Theoretical Background

The gPC method was originally developed by XIU AND KARNIADAKIS in 2002 [14]. Essentially, it is a spectral representation of the uncertain response of a system as follows [14]:

$$\mathbf{y}(\mathbf{z};\theta) \approx \sum_{m=0}^{M-1} \hat{\mathbf{y}}^{(m)}(\mathbf{z})\, \Phi^{(m)}(\theta), \qquad M-1 = \binom{N+D}{N} \tag{3}$$

In Eq. 3, the output variables, whose response should be approximated by the gPC expansion, are given by $\mathbf{y}$. The dimension of the gPC expansion is given by M, the number of uncertain parameters by N, and the highest order of the expansion polynomials by D. The multivariate expansion polynomials are depicted by Φ and are generally orthogonal polynomials. For a scalar polynomial ϕ the orthogonality relation is defined as follows [14]:

$$\int_{\Omega} \phi^{(m)}(\theta)\, \phi^{(n)}(\theta)\, \rho(\theta)\, \mathrm{d}\theta = \left(h^{(m)}\right)^2 \delta_{mn} = \begin{cases} \left(h^{(m)}\right)^2 & \text{if } m = n \\ 0 & \text{else} \end{cases}, \quad m,n \in \mathbb{N}_0 \tag{4}$$

The scalar orthogonal polynomial is symbolized by ϕ in Eq. 4, while Ω is the random space, i.e., the support of the pdf $\rho(\theta)$. The KRONECKER DELTA is given by δ_{mn} and $h^{(m)}$ is a normalization factor. For convenience, we assume that all orthogonal polynomials in this article are appropriately normalized. It should be noted that the relation between orthogonal polynomials, pdf, and its support is summarized in Table 1.

Now, the expansion coefficients, $\hat{\mathbf{y}}^{(m)}$, in Eq. 3 are obtained by the following integral equation [13]:

$$\hat{\mathbf{y}}^{(m)}(\mathbf{z}) = \int_{\Omega} \mathbf{y}(\mathbf{z};\theta)\, \Phi^{(m)}(\theta)\, \rho(\theta) \mathrm{d}\theta \tag{5}$$

For getting the gPC expansion in Eq. 3, the efficient calculation of the integral in Eq. 5 is important. Within this study, we use the SC approach as introduced in Subsection 2.2.2.

2.2.2 Stochastic Collocation

The SC approach is a Gaussian quadrature approach to get an approximation of the integral for the expansion coefficients in Eq. 5. This is a particularly viable choice as Gaussian quadrature schemes are designed with respect to the pdf of the orthogonal polynomials as introduced in Table 1 [2, 13].

Table 1: *Continuous Probability Distribution - Orthogonal Polynomial Connection For Standard gPC according to Wiener-Askey scheme [13].*

Probability Distribution	Probability Density Function	Support	Symbol	Orthogonal Polynomial
Gaussian/Normal	$\frac{1}{\sqrt{2\pi}}\exp\left(\frac{\theta^2}{2}\right)$	$]-\infty,\infty[$	$\mathscr{N}(\mu,\sigma)$	Hermite
Gamma	$\frac{\theta^{\alpha}\exp(-\theta)}{\Gamma(\alpha+1)}$	$[0,\infty[$	$\gamma(\mu,\sigma,\alpha)$	Laguerre
Beta	$\frac{\Gamma(\alpha+\beta+2)}{2^{\alpha+\beta+1}\Gamma(\alpha+1)\Gamma(\beta+1)}(1-\theta)^{\alpha}(1+\theta)^{\beta}$	$]-1,1[$	$\mathscr{B}(a,b,\alpha,\beta)$	Jacobi
Uniform	$\frac{1}{2}$	$]-1,1[$	$\mathscr{U}(a,b)$	Legendre

Generally, the Gaussian quadrature approximates the integral in Eq. 5 using a discrete expansion at a set of nodes $\theta^{(j)}$ with corresponding integration weights $\alpha^{(j)}$. This yields an approximation for Eq. 5 as follows [13]:

$$\hat{\mathbf{y}}^{(m)}(\mathbf{z}) = \int_{\Omega} \mathbf{y}(\mathbf{z};\theta)\,\Phi^{(m)}(\theta)\,\rho(\theta)\mathrm{d}\theta \approx \sum_{j=1}^{Q} \mathbf{y}\left(\mathbf{z};\theta^{(j)}\right)\Phi^{(m)}\left(\theta^{(j)}\right)\alpha^{(j)} \tag{6}$$

Here, Q is the number of chosen nodes that directly relates to the accuracy of the approximation. The nodes $\theta^{(j)}$ are the zeros of the orthogonal polynomial of order Q, while the weights $\alpha^{(j)}$ are calculated based on the pdf applying an integration of LAGRANGE polynomials as follows [13, p. 40]:

$$\alpha^{(j)} = \int_{\Omega} \rho(\theta) \prod_{\substack{i=1 \\ i\neq j}}^{Q} \frac{\theta-\theta^{(i)}}{\theta^{(j)}-\theta^{(i)}}\mathrm{d}\theta \tag{7}$$

Then, Eq. 6 reduces the problem of calculating the integral of the expansion coefficients in Eq. 5 to a deterministic sampling problem at the nodes $\theta^{(j)}$ and a subsequent evaluation of the discrete sum in Eq. 6. Take into account that for multiple uncertainties the CURSE OF DIMENSIONALITY becomes a problem. To overcome this issue sparse grid implementations can be used in engineering applications [13].

2.2.3 Statistical Moments

Statistical moments, such as mean or variance, can be calculated directly from the gPC expansion in Eq. 3. The mean is given as follows [13]:

$$\mathbb{E}[\mathbf{y}(\mathbf{z};\theta)] \approx \int_{\Omega} \left(\sum_{m=0}^{M-1} \hat{\mathbf{y}}^{(m)}(\mathbf{z})\,\Phi^{(m)}(\theta)\right)\rho(\theta)\mathrm{d}\theta = \hat{\mathbf{y}}^{(0)}(\mathbf{z}) \tag{8}$$

Take into account that Eq. 8 depicts that the mean is only depending on the first expansion coefficient.

Then, the variance (squared standard deviation σ) is calculated as follows [13]:

$$\sigma^2\left[\mathbf{y}\left(\mathbf{z};\theta\right)\right] = \mathbb{E}\left[\left(\mathbf{y}\left(\mathbf{z};\theta\right) - \mathbb{E}\left[\mathbf{y}\left(\mathbf{z};\theta\right)\right]\right)^2\right] \approx \sum_{m=1}^{M-1}\left[\hat{\mathbf{y}}^{(m)}\left(\mathbf{z}\right)\right]^2 \tag{9}$$

Again, the variance in Eq. 9 only depends on the expansion coefficients.

Eqs. 8–9 are now a representation of the uncertain system using mean and variance and can therefore be used to approximate the robust system response. Higher order statistical moments, such as skewness and kurtosis, can be calculated from the expansion formula in Eq. 3 as well, but are not required in this study.

2.2.4 Statistics in Distributed Optimization

Evaluating the statistical moments from gPC in the DOC framework requires rewriting Eqs. 8–9. This is due to the fact that the DOC framework relies on a distributed evaluation of the physical trajectories at the SC nodes, $\theta^{(j)}$, while the statistical moments are evaluated using the gPC expansion coefficients. To achieve this projection, the SC expansion formula in Eq. 6 must be inserted directly in the statistical moment calculation of Eqs. 8–9.

At first, the mean value in Eq. 8 can be rewritten as follows:

$$\mathbb{E}\left[\mathbf{y}\left(\mathbf{z};\theta\right)\right] \approx \hat{\mathbf{y}}^{(0)}\left(\mathbf{z}\right) \approx \sum_{j=1}^{Q}\mathbf{y}\left(\mathbf{z};\theta^{(j)}\right)\Phi^{(0)}\left(\theta^{(j)}\right)\alpha^{(j)} \tag{10}$$

It should be noted that Eq. 10 is already the desired distributed solution of the mean value even without introducing connection variables. This is due to the fact that the summands in Eq. 10 can be solved for each SC node individually. Thus, a DOC setup that e.g., only tries to optimize a mean value is indeed perfectly decoupled and can be solved independently in one step, i.e., without a connection problem.

Now, the variance in Eq. 9 can be written in distributed form starting from:

$$\sigma^2\left[\mathbf{y}\left(\mathbf{z};\theta\right)\right] \approx \sum_{m=1}^{M-1}\left[\hat{\mathbf{y}}^{(m)}\left(\mathbf{z}\right)\right]^2 \approx \sum_{m=1}^{M-1}\left[\sum_{j=1}^{Q}\mathbf{y}\left(\mathbf{z};\theta^{(j)}\right)\Phi^{(m)}\left(\theta^{(j)}\right)\alpha^{(j)}\right]^2 \tag{11}$$

To get a distributed representation of the variance, we need the binomial formula

$$\left[\sum_{j=1}^{Q}x^{(j)}\right]^2 = \sum_{j=1}^{Q}\left[x^{(j)}\right]^2 + \sum_{j\neq i}x^{(i)}x^{(j)}. \tag{12}$$

We can then rewrite the squared sum in Eq. 11 based on Eq. 12 as follows:

$$\begin{aligned}\sigma^2\left[\mathbf{y}\left(\mathbf{z};\theta\right)\right] \approx &\sum_{m=1}^{M-1}\sum_{j=1}^{Q}\left[\mathbf{y}\left(\mathbf{z};\theta^{(j)}\right)\right]^2\left[\Phi^{(m)}\left(\theta^{(j)}\right)\right]^2\left[\alpha^{(j)}\right]^2 \\ &+\sum_{m=1}^{M-1}\left[\sum_{j\neq i}\mathbf{y}\left(\mathbf{z};\theta^{(j)}\right)\Phi^{(m)}\left(\theta^{(j)}\right)\alpha^{(j)}\mathbf{y}\left(\mathbf{z};\theta^{(i)}\right)\Phi^{(m)}\left(\theta^{(i)}\right)\alpha^{(i)}\right]\end{aligned} \tag{13}$$

Now, we can split Eq. 13 in parts that are only dependent on the expansion coefficients and the SC nodes respectively as follows:

$$\begin{aligned}\sigma^2\left[\mathbf{y}\left(\mathbf{z};\theta\right)\right] \approx & \sum_{j=1}^{Q}\left[\mathbf{y}\left(\mathbf{z};\theta^{(j)}\right)\right]^2\left[\alpha^{(j)}\right]^2\sum_{m=1}^{M-1}\left[\Phi^{(m)}\left(\theta^{(j)}\right)\right]^2 \\ & +\sum_{j\neq i}\left[\mathbf{y}\left(\mathbf{z};\theta^{(j)}\right)\underbrace{\mathbf{y}\left(\mathbf{z};\theta^{(i)}\right)}_{\nu^{(i)}}\alpha^{(j)}\alpha^{(i)}\sum_{m=1}^{M-1}\Phi^{(m)}\left(\theta^{(j)}\right)\Phi^{(m)}\left(\theta^{(i)}\right)\right]\end{aligned} \tag{14}$$

Take into account that the first line in Eq. 14 is again decoupled from other SC nodes and can be handled distributed (as we know the orthogonal polynomials). On the other hand, the second line in Eq. 14 does not exhibit this preferable behavior: Here, a connection between the current distributed problem j and all other distributed problems i is imminent. Thus, when optimizing the variance the introduction of connection variables in the DOC framework is required.

It should be noted in the context of Eq. 14 that the connection problem still remains easy to solve. To show this we look at the second line of Eq. 14 and introduce the connection variables $\nu^{(i)}$ as follows:

$$\widetilde{\sigma}^2\left[\mathbf{y}\left(\mathbf{z};\theta\right)\right] = \sum_{j\neq i}\left[\mathbf{y}\left(\mathbf{z};\theta^{(j)}\right)\nu^{(i)}\alpha^{(j)}\alpha^{(i)}\sum_{m=1}^{M-1}\Phi^{(m)}\left(\theta^{(j)}\right)\Phi^{(m)}\left(\theta^{(i)}\right)\right] \tag{15}$$

Now, the connection variables for each distributed gPC problem j that are summed in Eq. 15, must fulfill the following condition enforced within the connection problem:

$$\xi^{(j)} = \nu^{(j)} - \mathbf{y}\left(\mathbf{z};\theta^{(j)}\right) \stackrel{!}{=} \mathbf{0} \tag{16}$$

Here, $\mathbf{y}\left(\mathbf{z};\theta^{(j)}\right)$ are known values in the connection problem after solving the DOC problems. Thus, Eq. 16 can be solved by the NEWTON method.

Then, the derivative of Eq. 16, required for the NEWTON method, is given as follows:

$$\frac{\mathrm{d}\xi^{(j)}}{\mathrm{d}\nu^{(j)}} = \mathbf{I} - \underbrace{\frac{\mathrm{d}\mathbf{y}\left(\mathbf{z};\theta^{(j)}\right)}{\mathrm{d}\nu^{(j)}}}_{\mathbf{0}} \tag{17}$$

It should be noted that the second addend equates to zero as a direct consequence of Eqs. 14 and 15: Here, it is imminent that the connection variables only introduce a cross-coupling, while no direct coupling is introduced. The sum with the connection variables is here only evaluated for elements that are not equal to the current distributed gPC problem j. Thus, a differentiation with respect to the same exponent j yields zero.

The structure in Eq. 17 is beneficial as we do not need to provide any sensitivities to the upper level connection problem and also justifies the seemingly complex rewriting of the original Eq. 11 to Eq. 14. Finally, the update of Eq. 16 using a NEWTON step with the derivative in Eq. 17 is given as follows:

$$\nu_{new}^{(j)} = \nu^{(j)} - \tau \left[\nu^{(j)} - \mathbf{y}\left(\mathbf{z}; \theta^{(j)}\right) \right] \tag{18}$$

Here, $\tau \in]0;1]$ is a step size (line search) parameter in order to find the best possible update to get as close to the current problem solution as possible.

2.3 Distributed Optimal Control

In a DOC framework, the OCP in Eq. 1 is distributed, i.e., divided into multiple smaller OCPs that can be solved independently. This independence must normally be secured by introducing connection variables that enforce constraints or mitigate an interaction between now distributed problems to ensure that still the original problem is solved (as shown in Eqs. 14 and 15). The overall strategy is also referred to as primal decomposition. The connection problem is solved separately from the distributed OCP in an upper level connection problem.

Overall, the $j-th$ DOC problem, i.e., the OCP solved at the gPC node $\theta^{(j)}$, can be stated according to Eq. 1 as follows:

$$\begin{aligned} \min_{\mathbf{x}^{(j)}, \mathbf{u}^{(j)}, \mathbf{p}^{(j)}} \quad & J^{(j)}\left(\mathbf{x}^{(j)}, \mathbf{u}^{(j)}, \mathbf{p}^{(j)}\right) + \widetilde{J}^{(j)}\left(\nu^{(i)}\right) \\ \text{s.t.} \quad & \mathbf{f}^{(j)}\left(\mathbf{x}^{(j)}, \mathbf{u}^{(j)}, \mathbf{p}^{(j)}; \theta^{(j)}\right) + \widetilde{\mathbf{f}}^{(j)}\left(\nu^{(i)}\right) = \dot{\mathbf{x}}^{(j)}, \\ & \mathbf{c}^{(j)}\left(\mathbf{x}^{(j)}, \mathbf{u}^{(j)}, \mathbf{p}^{(j)}\right) + \widetilde{\mathbf{c}}^{(j)}\left(\nu^{(i)}\right) \leq \mathbf{0}, \\ & \psi^{(j)}\left(\mathbf{x}^{(j)}, \mathbf{u}^{(j)}, \mathbf{p}^{(j)}\right) + \widetilde{\psi}^{(j)}\left(\nu^{(i)}\right) = \mathbf{0} \end{aligned} \tag{19}$$

Here, $\nu^{(i)}$ symbolizes the connection variables, while $\widetilde{.}$ over a function depicts functions with dependence on this connection variables. Take into account that this vector is no optimization variable in the distributed, single OCP, but is adapted according to a connection problem (e.g., Eq. 18) to ensure that the original solution of the connected problem is achieved. Further note that the connection variables, as shown in e.g., Eq. 15, depend on the other distributed problems ($i \neq j$). Thus, they are distinguished by another superscript. Overall, the connection variables thus merely provides an additional gradient shaping in the DOC context for the OCPs.

It should be noted that the distributed problem must still fulfill the original OCP formulation. Thus, e.g., the cost function must fulfill the following equality:

$$\sum_j \left[J^{(j)}\left(\mathbf{x}^{(j)}, \mathbf{u}^{(j)}, \mathbf{p}^{(j)}\right) + \widetilde{J}^{(j)}\left(\nu^{(i)}\right) \right] \equiv J \tag{20}$$

The same also applies for the constraints in Eq. 19. Take into account that Eq. 20 also directly implies a strategy to update the connection variables in the connection problem.

As an example for the definition of the DOC with statistical moments, we look at the following cost functional:

$$J = \mathbb{E}\left[y\left(\mathbf{z};\theta\right)\right] + k \cdot \sigma^2\left[y\left(\mathbf{z};\theta\right)\right] \tag{21}$$

Here, k is a scaling factor and y the variable whose mean and variance should be optimized (here in the scalar context for simplicity).

Now, we can use Eqs. 10 and 14 to rewrite Eq. 21 as follows:

$$\begin{aligned} J = &\sum_{j=1}^{Q} y\left(\mathbf{z};\theta^{(j)}\right) \Phi^{(0)}\left(\theta^{(j)}\right) \alpha^{(j)} + k \left\{ \sum_{j=1}^{Q} \left[y\left(\mathbf{z};\theta^{(j)}\right)\right]^2 \left[\alpha^{(j)}\right]^2 \sum_{m=1}^{M-1} \left[\Phi^{(m)}\left(\theta^{(j)}\right)\right]^2 \right. \\ &\left. + \sum_{j\neq i} \left[y\left(\mathbf{z};\theta^{(j)}\right) \underbrace{\nu^{(i)}}_{y\left(\mathbf{z};\theta^{(i)}\right)} \alpha^{(j)}\alpha^{(i)} \sum_{m=1}^{M-1} \Phi^{(m)}\left(\theta^{(j)}\right) \Phi^{(m)}\left(\theta^{(i)}\right) \right] \right\} \end{aligned} \tag{22}$$

Here, Eq. 22 directly provides us with the rule to split up the cost functional in Eq. 20 as the first line is only dependent on the $j-th$ OCP, while the second line has the additional influence of the connection variables:

$$\begin{aligned} J^{(j)} &= y\left(\mathbf{z};\theta^{(j)}\right) \Phi^{(0)}\left(\theta^{(j)}\right) \alpha^{(j)} + k \cdot \left[y\left(\mathbf{z};\theta^{(j)}\right)\right]^2 \left[\alpha^{(j)}\right]^2 \sum_{m=1}^{M-1} \left[\Phi^{(m)}\left(\theta^{(j)}\right)\right]^2 \\ \widetilde{J}^{(j)}\left(\nu^{(i)}\right) &= k \cdot y\left(\mathbf{z};\theta^{(j)}\right) \nu^{(i)} \alpha^{(j)}\alpha^{(i)} \sum_{m=1}^{M-1} \Phi^{(m)}\left(\theta^{(j)}\right) \Phi^{(m)}\left(\theta^{(i)}\right) \end{aligned} \tag{23}$$

The general structure of a DOC formulation in the context of OC is also introduced in Figure 1. It is evident that the upper level gives the connection variables to the lower level. The lower level then provides the updated trajectories that the upper level uses to improve on the connection variables. The procedure stops when the magnitude in the update of the connection variables becomes small.

Furthermore, the algorithm that solves the DOC, and is used within this study, can also be schematically written as given in Algorithm 1.

Algorithm 1 Implemented DOC framework algorithm with gPC.

Initialize connection variables $\nu^{(i)}$
while $||$Eq. 16$|| > \varepsilon$ (e.g., $\varepsilon = 10^{-5}$) **do**
  for all $\theta^{(j)}$ **do**
    Solve OCP in Eq. 19 (e.g., using [11]) at each $\theta^{(j)}$ in distributed manner, i.e., independent of each other, with current connection variables $\nu^{(i)}$
  end for
  Update the connection variables $\nu^{(i)}$ using Eq. 18 (and a line search if required)
end while

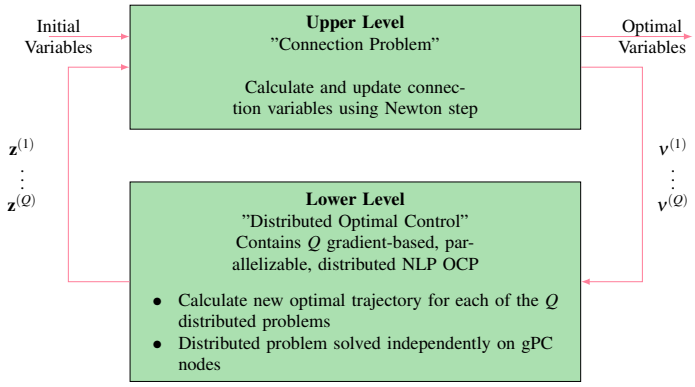

Fig. 1: *General structure of the distributed optimal control problem formulation with connection problem in upper level and the standard distributed optimal control problems in lower level.*

3 Dynamic Model and Optimization Problem

This section summarizes the equations of motion (EoM) (Section 3.1) and the OCP setup (Section 3.2) for the air race optimization model used in the case studies in Section 4.

3.1 Aircraft Dynamic Equations

The following subsection summarizes the EoM for a rigid-body aerobatic aircraft. This aircraft model has already been introduced and successfully used in other OC related applications [4, 5].

The EoMs for the x, y, and z position in a local coordinate system is defined as follows:

$$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{z} \end{bmatrix} = \begin{bmatrix} V \cdot \cos(\chi) \cdot \cos(\gamma) \\ V \cdot \sin(\chi) \cdot \cos(\gamma) \\ -V \cdot \sin(\gamma) \end{bmatrix} \tag{24}$$

The translational EoMs are described using the kinematic velocity V, the kinematic course angle χ, and the kinematic climb angle γ. From NEWTONs second law, the translational dynamics are propagated as follows:

$$\dot{V} = \frac{1}{m} \cdot F_{T,x}, \quad \dot{\chi} = \frac{1}{m \cdot V \cdot \cos(\gamma)} \cdot F_{T,y}, \quad \dot{\gamma} = -\frac{1}{m \cdot V} \cdot F_{T,z} \tag{25}$$

Here, F_T symbolizes the total force (i.e., thrust, aerodynamics, and gravitation) in each axis direction.

The attitude EoMs are described by the kinematic angle of attack α, the kinematic bank angle μ, and the kinematic sideslip angle β:

$$\dot{\mu} = \omega_x^{KB} - \tan(\beta)\left(\omega_y^{KB}\cos(\mu) + \omega_z^{KB}\sin(\mu)\right) \tag{26a}$$

$$\dot{\alpha} = \frac{1}{\cos(\beta)}\left(\omega_y^{KB}\cos(\mu) + \omega_z^{KB}\sin(\mu)\right) \tag{26b}$$

$$\dot{\beta} = \omega_y^{KB}\sin(\mu) - \omega_z^{KB}\cos(\mu) \tag{26c}$$

In Eqs. 26a–26c, the components of $\left(\omega^{KB}\right)_K$ denote the kinematic rotational velocities between the Kinematic frame K and the Body-fixed frame B, denoted in the Kinematic frame K:

$$\left(\omega^{KB}\right)_K = \begin{bmatrix}\omega_x^{KB}\\ \omega_y^{KB}\\ \omega_z^{KB}\end{bmatrix}_K = -\begin{bmatrix}-\dot{\chi}\cdot\sin(\gamma)\\ \dot{\gamma}\\ \dot{\chi}\cdot\cos(\gamma)\end{bmatrix}_K + M_{KB}\cdot\begin{bmatrix}p\\ q\\ r\end{bmatrix}_B \tag{27}$$

The last set of equations for the rigid body dynamics are the rotational EoMs obtained from momentum conservation. They describe the evolution of the kinematic roll rate p, pitch rate q, and yaw rate r. Using the total moments $(\mathbf{M}_T)_B$, the inertia tensor $\mathbf{I}$ with respect to the center of gravity and the Body-fixed frame, and the kinematic angular body rate vector $\left(\omega^{OB}\right)_B$ between the Orientation (NED) frame (O) and the Body-fixed frame (B) the angular accelerations are computed as follows:

$$\left(\dot{\omega}^{OB}\right)_B = \left[\dot{p}\ \dot{q}\ \dot{r}\right]^T = \mathbf{I}^{-1}\left[(\mathbf{M}_T)_B - \left(\omega^{OB}\right)_B \times \mathbf{I}\cdot\left(\omega^{OB}\right)_B\right] \tag{28}$$

The thrust lever dynamics are modeled as a first order lag:

$$\dot{\delta}_T = \frac{1}{T_{Engine}}\left(\delta_{T,c} - \delta_T\right) \tag{29}$$

The thrust lever directly provides the thrust force T for the EoM using the following relation:

$$F_{T,x} = T = \delta_T \cdot T_{ref} \tag{30}$$

3.2 Optimal Control Problem Setup

The OCP consists of two phases that model three race gate positions to be passed in wings-level position ($\mu = 0$) and for a defined direction χ. The following initial boundary conditions (*IBC*) for the states $[x, y, z, \chi, \gamma, \mu]$ define these in the beginning of each phase (the other states are free):

Phase	IBC_{lb}	IBC_{ub}
1	$[0,0,0,0,0,0]$	$[0,0,0,0,0,0]$
2	$[10^3,0,0,0,0,0]$	$[10^3,0,0,0,0,0]$

Furthermore, the final boundary conditions (FBC) for the same states in the last phase are defined as:

Phase	FBC_{lb}	FBC_{ub}
2	$[10^3, 140, 0, \pi, 0, 0]$	$[10^3, 140, 0, \pi, 0, 0]$

The final times of the three phases with the respective lower and upper bounds t_f^{lb}, t_f^{ub} and scaling t_f^S are defined as:

Phase	$t_{f,i}$	$t_{f,i}^{lb}$	$t_{f,i}^{ub}$	$t_{f,i}^S$
$i=1$	$10s$	$0s$	$40s$	10^0
$i=2$	$20s$	$0s$	$40s$	10^0

The states with the respective lower and upper bounds x^{lb}, x^{ub}, and scalings x^S are as follows:

State	Name	x^{lb}	x^{ub}	x^S
x	x-Position	$-5 \cdot 10^3$	$5 \cdot 10^3$	10^{-2}
y	y-Position	$-5 \cdot 10^3$	$5 \cdot 10^3$	10^{-2}
z	z-Position	$-5 \cdot 10^3$	$0 \cdot 10^1$	10^{-2}
χ	Course angle	-3π	3π	10^0
γ	Climb angle	$-\pi$	π	10^0
μ	Bank angle	$-\pi$	π	10^0
V	Velocity	25	100	10^{-1}
α	Angle of attack	-0.1π	0.1π	10^0
β	Angle of sideslip	-0.1π	0.1π	10^0
p	Roll rate	-2π	2π	10^0
q	Pitch rate	$-\pi$	π	10^0
r	Yaw rate	$-\pi$	π	10^0
δ_T	Thrust lever	0	1	10^0

The controls with the respective lower and upper bounds u^{lb}, u^{ub}, and scalings u^S are as follows:

Control	Name	u^{lb}	u^{ub}	u^S
ξ	Aileron deflection	$-\frac{\pi}{8}$	$\frac{\pi}{7}$	10^0
η	Elevator deflection	$-\frac{\pi}{7}$	$\frac{\pi}{7}$	10^0
ζ	Rudder deflection	$-\frac{\pi}{6}$	$\frac{\pi}{6}$	10^0
$\delta_{T,c}$	Thrust lever position	0	1	10^0

Furthermore, the load factor $(n_{T,z})_B$ in the z_B-direction of the Body-fixed frame (B) is constrained for all three phases:

$$-10 \leq (n_{T,z})_B = \frac{(F_{T,z})_B}{mg} \leq 2 \tag{31}$$

The normalized state and control discretization for each of the two phases is 0.005 and we solve the OCPs as well as the connection problem by IPOPT to a tolerance of 10^{-5}. Finally, the uncertain system response is approximated using a 3^{rd} order gPC expansion.

4 Distributed Optimal Control Test Cases

This section covers the test cases of the DOC framework: At first, a minimization of the mean final time is described in Subsection 4.1. Then, Subsection 4.2 looks at the additional optimization of the load factor variance that yields a more complex DOC problem including connection variables.

4.1 Minimization of Mean Final Time

This test case features a fairly simple DOC problem in order to show the capabilities of the approach as well as to verify the results. For this optimization, we consider the reference thrust in Eq. 30 to be normally distributed as follows:

$$T_{ref} \in \mathcal{N}(\mu = 5500N, \sigma = 100N) \tag{32}$$

This uncertainty in the reference thrust is a consequence of different used engines as well as environmental conditions (e.g., air density). The order of the gPC expansion is three (i.e., $M = D = 3, N = 1$), as this is sufficient to approximate the mean value.

Overall, we look at the minimization of the mean final time, for which the overall cost function is as follows:

$$J = \hat{t_f}_{,2}^{(0)} \approx \sum_{j=1}^{Q=3} {t_f}_{,2}^{(j)} \Phi^{(0)}\left(\theta^{(j)}\right) \alpha^{(j)} \tag{33}$$

From Eq. 33, the DOC cost function can be written down like in Eq. 23 ($k = 0$):

$$J^{(j)} = {t_f}_{,2}^{(j)} \Phi^{(0)}\left(\theta^{(j)}\right) \alpha^{(j)}, \qquad j = 1,2,3 \tag{34}$$

Take into account that we do not need any connection variables as also already stated in Subsection 2.2.4. This consequently means that we converge in one iteration of the DOC problem. Furthermore, this implies that the added cost terms in Eq. 19 equate to zero:

$$\widetilde{J}^{(j)} = 0, \qquad j = 1,2,3 \tag{35}$$

The results of the DOC are visualized in Figure 2: Here, a 3D comparison of the trajectory from DOC (solid blue) and from the reference gPC solution of the OCP, i.e., a non-robust uncertain representation (dashed red). This representation is calculated by solving standard time-minimal OCPs at the nodes $\theta^{(j)}$ and afterwards creating the gPC expansion from it. Overall, the results are matching fairly good and both methods are converging to the same solution.

It should be noted that small errors still occur as also depicted in Figure 3. These errors are based on the fact that we are not solving the same numerical problem although the analytical problem is the same. This is a consequence of the fact that

we are scaling the cost function differently. In the reference case the cost function is just the numeric value while we scale this numeric value by the expansion weights in the DOC case (Eq. 10). This also shows in the slightly different cost function for DOC case of $J_{distr} = 16.6623642599$ compared to the value of the reference case of $J_{ref} = 16.6519073935$. This shows that the slight variations in the scaling do have an influence in the cost function and which local minima is found (i.e., which result is optimal compared to the solver tolerances). In the end, this also leads to the variations in the trajectory that can be seen in the results of Figure 3. Nonetheless, this does not nullify the results of the proposed method as a good and correct scaling is a general issue for OC and always must be chosen carefully.

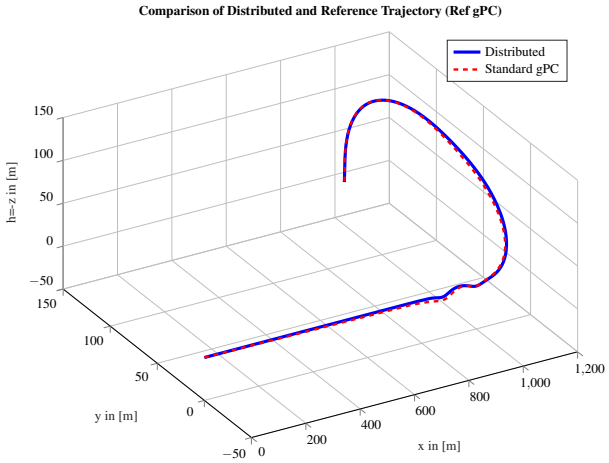

Fig. 2: *Comparison of optimal trajectories for distributed optimal control (solid blue) and reference, non-robust (dashed red) gPC optimization.*

4.2 Minimization of Load Factor Variance

Within this section, we additionally look at the incorporation of a Lagrange cost on the load factor variance in Eq. 33. This cost implements the pilot's desire to always fly with the same "feeling". This means that the pilot wants to fly the trajectory similar even with an uncertainty as this gives him the best, i.e., the well-known, feeling for the aircraft and his flight. The uncertainty remains as defined in Eq. 32. The expansion order also remains three (i.e., $M = D = 3, N = 1$).

Here, the DOC compared to Subsection 4.1 is changing as we need an interconnection between the different phases to calculate the variance. The distributed cost related to the load factor variance is defined as follows:

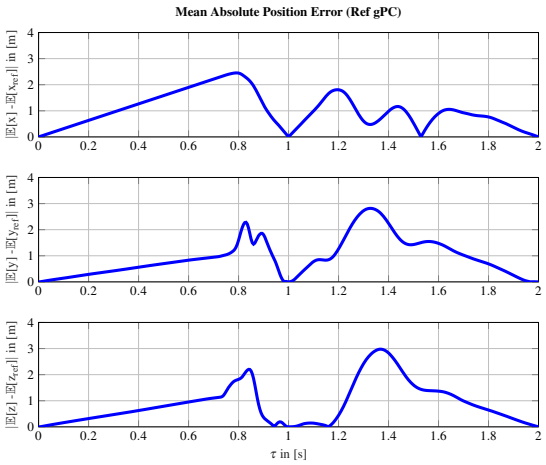

Fig. 3: *Absolute mean position error over time between distributed optimal control and reference, non-robust gPC optimization.*

$$\begin{aligned}\sigma^2\left[\left(n_z^{(j)}\right)_B\right] = &\left[\left(n_z^{(j)}\right)_B\right]^2\left[\alpha^{(j)}\right]^2\sum_{m=1}^{M-1}\left[\Phi^{(m)}\left(\theta^{(j)}\right)\right]^2 \\ &+\sum_{j\neq i}\left[\left(n_z^{(j)}\right)_B\underbrace{\left(n_z^{(i)}\right)_B}_{\nu^{(i)}}\alpha^{(j)}\alpha^{(i)}\sum_{m=1}^{M-1}\Phi^{(m)}\left(\theta^{(j)}\right)\Phi^{(m)}\left(\theta^{(i)}\right)\right]\end{aligned} \tag{36}$$

It should be noted that the connection variables $\nu^{(i)} = \left(n_z^{(i)}\right)_B$ are introduced in Eq. 36 to connect the load factor histories of the other SC phases to be able to calculate the variance. The introduced connection variables are updated in a bi-level coordination problem such that the original problem formulation is solved (Figure 1).

Looking at the DOC problem statement in Eq. 19 and the general distributed cost in Eq. 23, the following cost function influences can be identified:

$$\begin{aligned}J^{(j)} &= \left[\left(n_z^{(j)}\right)_B\right]^2\left[\alpha^{(j)}\right]^2\sum_{m=1}^{M-1}\left[\Phi^{(m)}\left(\theta^{(j)}\right)\right]^2, \quad j=1,2,3 \\ \widetilde{J}^{(j)} &= \sum_{j\neq i}\left[\left(n_z^{(j)}\right)_B\nu^{(i)}\alpha^{(j)}\alpha^{(i)}\sum_{m=1}^{M-1}\Phi^{(m)}\left(\theta^{(j)}\right)\Phi^{(m)}\left(\theta^{(i)}\right)\right], \quad j=1,2,3\end{aligned} \tag{37}$$

Thus, we have indeed found a distributed description of the original coupled problem. It should be noted that the cost in Eq. 37 is combined with the already distributed cost in Eq. 33 to form the overall cost function (similar to Eq. 22).

Generally, the connection variables in Eq. 36 must fulfill the condition as follows:

$$\nu^{(j)} - \left(n_z^{(j)}\right)_B \equiv 0, \quad j = 1,2,3 \tag{38}$$

Using a NEWTON step to fulfill the condition in the coordination problem iteratively, as proposed in Eq. 18, we get the following update equation:

$$\nu_{new}^{(j)} = \nu^{(j)} - \tau \left[\nu^{(j)} - \left(n_z^{(j)}\right)_B\right], \quad j = 1,2,3 \tag{39}$$

In this study, the line search parameter is chosen according to the criterion that minimizes the standard deviation difference of the updated and the current problem.

Take into account that the simple update law in Eq. 39 requires no sensitivity or inversion of any matrix. This makes the update fairly easy and fast to calculate. It should also be noted that this simple update law form is based on the transformation done to the original definition of the variance in Subsection 2.2.4. Originally, this rewriting resulted in a generally more complex formula, but a formula that is easier to handle in the coordination problem as seen here (e.g., when no sensitivities are readily available).

Now, we look at the results for the robust distributed optimization: At first, Figure 4 depicts a comparison of the DOC results of this section (in the following: solid blue line) and the DOC reference results from Section 4.1 (in the following: dashed red line). Once more, the reference results are obtained as a non-robust solution by solving standard time-minimal OCPs at the nodes $\theta^{(j)}$ and afterwards creating the gPC expansion from it. It is imminent that the DOC solution takes a larger turn radius within the steep turn. This consequently yields a minimization of the mean load factor as seen in Figure 5. It is obvious that the load factor is not as aggressive and specifically not as often on its boundary condition.

In combination with the also plotted standard deviation that is also significantly reduced over the integration interval, we achieve an increased robustness. The overall area is reduced from the value of ≈ 4.38 for the reference case, to a value of ≈ 2.96 in the robust case, which equates to roughly 33%. This increased robustness is achieved at the cost of a reduced optimality: Specifically, the optimal trajectory time is increased from $\approx 16.66s$ in the reference case to $\approx 18.28s$ for the robust case. This is an increase of approximately 10%.

Overall, the results show that a robustness increase can be achieved by the proposed DOC framework. This robustness increase comes at the cost of a less optimal solution. The user can choose the weighting between optimality and robustness by appropriately scaling the two terms. This yields to the Pareto problem introduced in Eq. 22 that is subject to further investigations.

In order to better understand how the DOC problem achieves a more robust trajectory description, Figures 6 and 7 show the mean values and standard deviations of the three longitudinal plane states height, climb angle, and velocity respectively. The height state shows the behavior that is already known from Figure 4: The robust DOC result does climb more gradually to an overall smaller height. This is also seen from the climb angle results. The climb is more gradual and does not reach the same maximal value as in the reference case. Another interesting maneuver that the aircraft takes in the first flight phase is to reduce height in order to climb before the

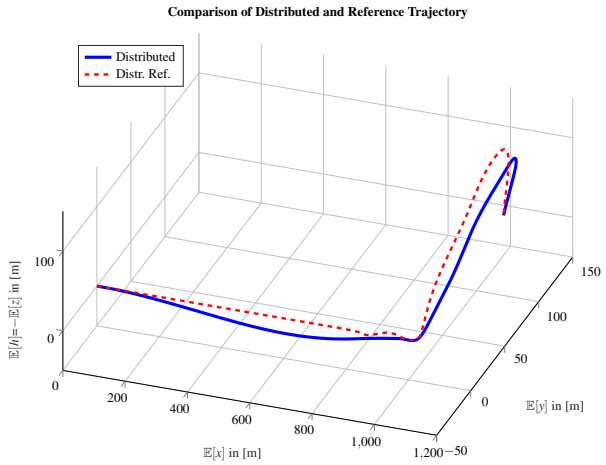

Fig. 4: *Comparison of Optimal Trajectories for distributed (solid blue) and reference, non-robust (dashed red) optimization.*

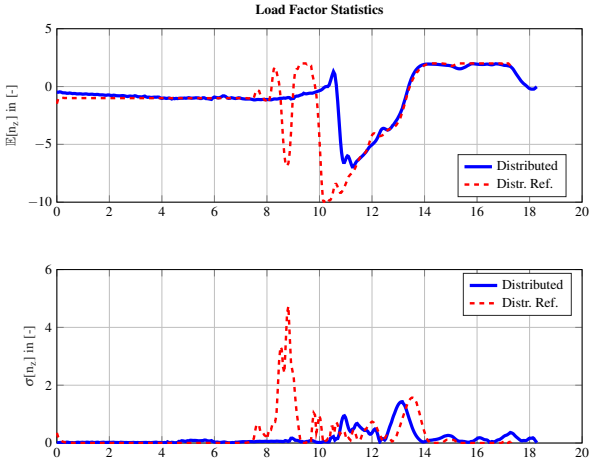

Fig. 5: *Comparison of load factor statistics for distributed (solid blue) and reference, non-robust (dashed red) optimization.*

turn. This yields a reduction of the velocity before the turn is entered, which again reduces the load factor.

The standard deviations in Figure 7 depict that these values are smaller for the robust DOC case in the height and the climb angle from a global perspective. Contrary, the velocity has a very large standard deviation, especially in the first phase, with the already mentioned descent-climb maneuver. This is a consequence of the uncertainty in the maximum thrust that influences how the maneuver must be exe-

cuted. Overall, the standard deviations of the robust DOC case show a more gradual behavior than in the reference case.

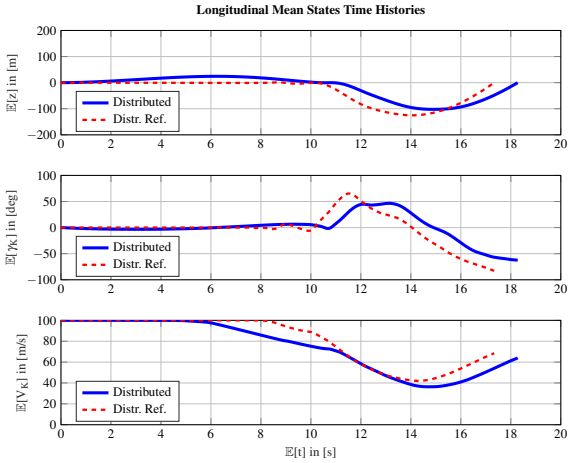

Fig. 6: *Comparison of specific mean longitudinal states for distributed (solid blue) and reference, non-robust (dashed red) optimization.*

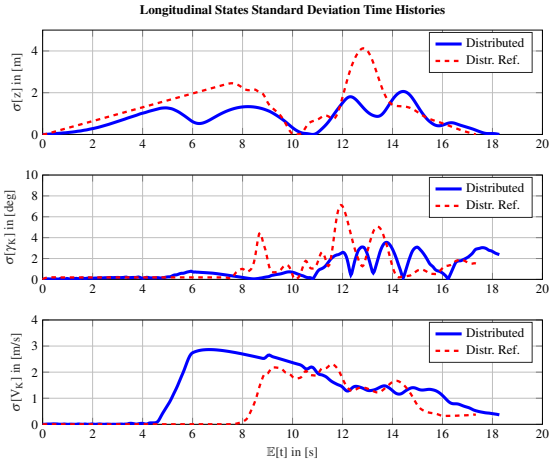

Fig. 7: *Comparison of specific longitudinal states' standard deviation for distributed (solid blue) and reference, non-robust (dashed red) optimization.*

5 Conclusions

This study presented an efficient method for ROC that relies on a distributed solution of the gPC problem. Therefore, a specifically tailored DOC framework was developed that could efficiently handle the resulting gPC connection problem efficiently. Here, the connection variables are part of the statistical moments, e.g., the variance, that can only be calculated with the knowledge of the optimal solution of all distributed problems. With the formulation as introduced in this study the connection problem becomes a simple zero search that does not require the knowledge of sensitivities. This makes the solution fairly easy and efficient.

One important additional finding was that the distributed problem reduces to a single solution of the unconnected problems in case only mean values are optimized or constrained. This makes the framework very efficient for these kinds of cases. For these cases the framework showed a very good match with the reference results.

On the other hand the optimization of higher order moments (e.g., variance) requires the introduction of additional connection variables. But with a reformulation of the gPC formulas of e.g., the variance using binomial theorems, an easy form of the connection problem can be achieved.

Future research should be concerned with adaptive and efficient line search procedure for the connection problem as it has shown itself that always taking a full step is not suitable in order to reach a good update and converge fast.

Furthermore, chance constraints can be treated with the proposed robust DOC framework. Here, the chance constraint must be distributed as well and the results should give a probability for the trajectory to be within a specified constraint tolerance.

Finally, an important topic in the context of ROC in general and robust DOC in specific is the already mentioned Pareto problem that is encountered. This problem is natural and based on the trade-off of optimality and robustness. Here, fast methods for Pareto optimization in the context of the DOC framework should be developed in order to find the Pareto optimum efficient, fast, and reliable.

Acknowledgements This work was supported by the Deutsche Forschungsgemeinschaft (DFG) through the TUM International Graduate School of Science and Engineering (IGSSE).

References

[1] Betts JT (2010) Practical methods for optimal control and estimation using nonlinear programming, 2nd edn. Advances in design and control, Society for Industrial and Applied Mathematics (SIAM 3600 Market Street Floor 6 Philadelphia PA 19104), Philadelphia, Pa., URL `http://epubs.siam.org/ebooks/siam/advances_in_design_control/dc19`

[2] Cottrill G (2012) Hybrid Solution of Stochastic Optimal Control Problems Using Gauss Pseudospectral Method and Generalized Polynomial Chaos Algorithms. Dissertation, Air Force Institute of Technology, Ohio

[3] Dai L, Xia Y, Gao Y (2015) Distributed Model Predictive Control of Linear Systems with Stochastic Parametric Uncertainties and Coupled Probabilistic Constraints. SIAM Journal on Control and Optimization 53(6):3411–3431, doi: 10.1137/140994290

[4] Diepolder J, Bittner M, Piprek P, Grüter B, Holzapfel F (2017) Facilitating aircraft optimal control based on numerical nonlinear dynamic inversion. In: 25th Mediterranean Conference on Control and Automation (MED), pp 141–146, doi: 10.1109/MED.2017.7984108

[5] Fisch F (2011) Development of a Framework for the Solution of High-Fidelity Trajectory Optimization Problems and Bilevel Optimal Control Problems. Dissertation, Technische Universität München, München, URL http://nbn-resolving.de/urn/resolver.pl?urn:nbn:de:bvb:91-diss-20110221-1001868-1-4

[6] Harmon FG (2017) Hybrid solution of nonlinear stochastic optimal control problems using Legendre Pseudospectral and generalized Polynomial Chaos algorithms. In: 2017 American Control Conference (ACC), IEEE, Piscataway, NJ, pp 2642–2647, doi: 10.23919/ACC.2017.7963351

[7] Jiang Y, Nimmegeers P, Telen D, van Impe J, Houska B (2017) A Distributed Optimization Algorithm for Stochastic Optimal Control. IFAC-PapersOnLine 50(1):11,263–11,268, doi: 10.1016/j.ifacol.2017.08.1618

[8] Li X, Nair PB, Zhang Z, Gao L, Gao C (2014) Aircraft Robust Trajectory Optimization Using Nonintrusive Polynomial Chaos. Journal of Aircraft 51(5):1592–1603, doi: 10.2514/1.C032474

[9] Matsuno Y, Tsuchiya T (2014) Stochastic 4D trajectory optimization for aircraft conflict resolution. In: IEEE Aerospace Conference, 2014, IEEE, Piscataway, NJ, pp 1–10, doi: 10.1109/AERO.2014.6836275

[10] Richter M, Holzapfel F (2013) Robust Noise Optimal Approach Trajectories. AIAA Guidance, Navigation, and Control (GNC) Conference, Boston doi: 10.2514/6.2013-4556

[11] Rieck M, Bittner M, Grüter B, Diepolder J, Piprek P (2018) FALCON.m User Guide. URL www.falcon-m.com

[12] Wächter A, Biegler LT (2006) On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. Mathematical Programming 106(1):25–57, doi: 10.1007/s10107-004-0559-y

[13] Xiu D (2010) Numerical methods for stochastic computations: A spectral method approach. Princeton University Press, Princeton, N.J, URL http://lib.myilibrary.com/detail.asp?ID=293631

[14] Xiu D, Karniadakis GE (2002) The Wiener-Askey Polynomial Chaos for Stochastic Differential Equations. SIAM Journal on Scientific Computing 24(2):619–644, doi: 10.1137/S1064827501387826