A compressed-sensing approach for closed-loop optimal control of nonlinear systems

Lionel Mathelin[a,*], Luc Pastur[a], Olivier Le Maître[a]

[a] *LIMSI-CNRS, UPR-3251, Orsay, France.*

Abstract

We present a method that seeks to combine the properties of optimal control with the robust character of closed-loop control. The method relies on the availability of a reduced model of the system to be controlled, in order to express the control problem in a low dimensional space where the system-state dependent optimal control law is subsequently approximated in a pre-processing stage. A polynomial expansion is used for the approximation, enabling fast update of the optimal control law each time a new observation of the system-state is made available. It results in a real-time compatible, efficient (optimal-like) and robust control strategy. A compressed-sensing approach is proposed to efficiently construct the approximate control law, exploiting the compressible character of the optimal control law in the retained approximation basis. The method is demonstrated for the control of the flow around a cylinder and is shown to perform as well as the much more costly receding-horizon optimal control approach, where the exact optimal control problem is actually recomputed, even in presence of large aleatoric perturbations.

Potential and remaining issues toward application to larger dimensional reduced systems are also discussed and some directions for improvement are suggested.

Key words: Nonlinear control, Closed-loop control, Compressed sensing, Model reduction

1. Introduction

With the ever more stringent search for efficiency and optimized performances in almost all aspects of everyday life, control techniques are more of interest than ever. In particular, controlling a physical system has become a widely adopted road-map towards more efficient devices. This trend does not spare the fluid mechanics community, in particular owing to strong incitations from the aeronautical industry. However, fluid systems often involve complex behaviors and dynamics, typically highly sensitive to external perturbations, and exhibiting a wide spectrum of temporal- and spatial-scales. This characteristic usually makes the control of such systems a formidable task and, despite the continuous increase in the available computational power, numerical optimization procedures based on a detailed model of the systems -including all scales- remain often out of reach. Another approach is then necessary and several techniques have been proposed to reduce the complexity of the numerical model used for the determination of the control. All these techniques aim at constructing a low dimensional numerical model (from a few tens to a few thousands

[☆] The first two authors gratefully acknowledge the financial support from the French Research National Agency (ANR) under projects CORMORED (ANR-08-BLAN-0115) and ASRMEI (JCJC08 #375619). LM is also indebted to André Veragen for numerous inspiring discussions.

[*] Corresponding author. Laboratoire d'Informatique pour la Mécanique et les Sciences de l'Ingénieur, LIMSI-CNRS, rue John von Neumann, BP 133, 91403 Orsay cedex, France.
Tel: 33-1-69-85-80-69; fax: 33-1-69-85-80-88

Email addresses: `mathelin@limsi.fr` (Lionel Mathelin), `pastur@limsi.fr` (Luc Pastur), `olm@limsi.fr` (Olivier Le Maître)

URL: `http://www.limsi.fr/Individu/mathelin` (Lionel Mathelin), `http://www.limsi.fr/Individu/olm` (Olivier Le Maître)

of degrees of freedom), yet able to reproduce the essential feature of the actual physical system. When an appropriate reduced model has been derived, it can be controlled by a variety of methods.

To achieve a good performance in controlling the low dimensional model, high order control methods, as opposed to low order methods such as Proportional-Integral-Derivative (PID) controllers, are often necessary. Among those, an appealing one is the optimal control approach, [19], which relies on an optimality principle and a model prediction over a prescribed time-horizon to derive the control command. While optimal by definition, this open-loop control is sometimes suffers a lack of robustness when applied to the actual physical system. In such a situation, the optimal control derived from the reduced model exhibits a sharp decrease in the resulting performances in case of mild drift from the assumed conditions affecting the system and in presence of perturbations. Even though perturbations may be somehow described by a probabilistic model, leading to a robust $\mathscr{H}_2$- or $\mathscr{H}_\infty$-control command, such probabilistic or worst case approaches do not account for actual perturbations and the resulting control is usually too conservative. Further, the optimal control law determined from the reduced order model is certainly non-optimal when applied to the true system, with possible performance loss calling for more robust strategies. To account for non-modeled external perturbations, a receding-horizon implementation may be used where the control command is determined over an horizon but is updated (recomputed) after only a fraction of it has elapsed. This implementation allows to improve the robustness by adapting the control to fluctuating conditions and by reseting potential drifts in the predicted dynamics of the system. It has been successfully applied to numerous problems but the method can require prohibitive computational resources, even when relying on a reduced order model, and is not suitable in particular for the real-time control of complex fluid flows involving fast time-scales has in aerodynamical problems.

Gain-scheduling is a popular alternative to achieve high performance control without heavy computations, see [30] for a review. The idea is to pre-set "envelopes" in the state-space of the system and to derive a control specific to each envelop. Several distinct controllers are then set and the state of the system determines which envelop is "active" and which, possibly closed-loop, controller is to be used. At the expense of determining as many controllers as envelopes, this approach allows for simple, yet effective, (closed-loop) global control. Numerous aeronautical applications rely on such a control scheme. While appealing, this requires several controllers to be determined, and transitions from one controller to another, when the system state drifts from one envelope to another, is a difficulty.

For linear systems, a popular approach for closed-loop, real-time, control relies on an explicit expression relating the current state of the system to the control to be applied. The control command is then simply expressed as a gain matrix applied to the current system state vector. The gain matrix is determined as the solution of a, possibly time-dependent, Riccati equation. When the control command is derived over an infinite horizon, the Differential Riccati Equation (DRE) degenerates into a time-independent gain matrix and the control then simply expresses as a matrix-vector product at each time step. In that case, an efficient closed-loop optimal control is possible, usually coupled with a state estimator (for instance, Kalman filter-based) to provide the controller with an estimate of the actual current state of the system. This approach has become very popular. However, current algorithms for solving the Algebraic Riccati Equation, [16], practically limit the size of the system state to a few thousands before the problem becomes ill-conditioned. Efforts are made in order to approximate the Riccati equation solution and allow to consider reduced order systems of hundred of thousands of state variables, see [24, 4]. Typical studies involving a Riccati equation-based approach deal with control of small perturbations affecting a linearized model such as boundary layer transition control or instability development control, see for instance [2]. One may refer to [15] for a review on the control of linear systems. While appealing, this Riccati-based approach is restricted to linear models and it precludes its use in many situations. It is obviously desirable to take advantage of the real-time feature of the Riccati-based approach without the restriction to a linear model. Unless a remarkable theoretical result, similar to that leading to the Riccati equation, is found in the non-linear context, one has little choice but to shift as much calculation as possible in a pre-processing step to be left with the very minimal amount of operations to be carried-out while on-line when the control is actually activated. This requirement is particularly crucial in the turbulent fluid mechanics context where the required control command frequencies are routinely above a kilohertz.

In the present paper, we build-up on our previous works on robust model reduction [18] and control of

reduced model [23] to propose a control strategy that combines the closed-loop character of the Riccati-based-type control allowing for performance and robustness with computational efficiency. We specifically consider complex fluid flows, for which we assume the availability of a reduced model, and requiring a sophisticated control synthesis that can not be computed in real time, even using the reduced model. The latter aspect prevents the receding time approaches used in [23] and an alternative is thus proposed to construct a cheap approximate closed-loop optimal controller of the flow. The proposed approach is two-fold. In a first, off-line, step, the optimal control law to be applied is determined for some carefully chosen initial states of the reduced model. This collection of pre-computed control laws is then used to construct an approximation of the reduced state system dependent optimal control law. When on-line, *i.e.* in the actual control situation, the current (reduced) state of the system is evaluated and the corresponding control command is reconstructed at virtually no cost, using the approximation. The sole computationally costly step is then associated to the off-line derivation of the set of pre-computed optimal laws and is directly related to the size of the necessary set to obtain the approximation. Specifically, we shall consider polynomial expansions to express the dependence of the control law with regard to the reduced system state, so the number of pre-computed control laws needed to estimate the expansion coefficients is related to the polynomial order of the approximation and the dimensionality of the reduced system state. In practice however, the approximation exhibits some degree of compressibility in the retained polynomial basis, *i.e.* many of the expansion coefficients are negligible with respect to the dominant ones. Such sparsity is exploited here making use of the compressed-sensing methodology, [5, 9]. Compressed-sensing enables to limit the number of necessary control law computations to a minimum, by focusing on the determination of the dominant coefficients in the expansion. Compressed-sensing methodology will be seen to achieve a very significant savings in the computational cost while providing a very effective approximation of the control law.

The paper is organized as follows. In section 2, we briefly describe two main ingredients of the control strategy proposed here, namely the model reduction step and the optimal control formulation. In section 3, the technique for efficiently approximate the state-dependent control law (off-line step) is introduced. In particular, the choice of the approximation basis and the compressed sensing technique are discussed. The method is demonstrated with the control of the flow around a circular cylinder in section 4. Its efficiency with respect to the receding-horizon optimal control is shown and several aspects are investigated, such as the robustness with respect to exogeneous perturbations and the impact of the approximation accuracy. Final conclusions are drawn in section 5.

2. Model Reduction and Optimal Control

We consider a complex physical system described by the time-dependent vector-valued field $\boldsymbol{u}(\boldsymbol{x}, t) \in \mathbb{R}^m$ defined on a spatial domain Ω_x. The system may be a fluid-flow where $\boldsymbol{u}(\boldsymbol{x}, t)$ is the vector containing the components of the fluid velocity, the pressure and the temperature at point $\boldsymbol{x} \in \Omega_x$ and time t. The dynamics of $\boldsymbol{u}$ is prescribed by the governing equations, for instance the Navier-Stokes equations, consisting of a set of partial differential equations. We wish to efficiently control the dynamics of the system with respect to some objectives yet to be defined, through some control actions such as modification of boundary conditions or introduction of some forcing terms. We denote $\boldsymbol{\mu}$ the control and call "control law" the time-evolution of the control: $\boldsymbol{\mu}(t)$.

The determination of the control law yielding the control objective or, more likely, approaching it, is a very computationally demanding task that can hardly be achieved directly from the system governing equations in practice. In fact, in most situations one needs first to derive a reduced order model (ROM) of the system. The purpose of the model reduction is to derive a low order (or low dimensional) system whose dynamics mimic to some extent the essential features of the original one. Owing to its low dimensionality, one expects the determination of the control law for the reduced system to be computationally much less complex, while remaining relevant when applied to the original system.

In this section, the reduced order modeling of the system is first introduced, followed by a brief description of the determination of its optimal control law.

2.1. Model reduction

We assume the system can be expanded in the separated form

$$\boldsymbol{u}(\boldsymbol{x},t) = \sum_{i=1}^{\infty} \boldsymbol{U}_i(\boldsymbol{x}) a_i(t), \tag{1}$$

where $\boldsymbol{U}_i(\boldsymbol{x})$ and $a_i(t)$ are respectively the spatial modes and temporal coefficients of the reduced model. It is assumed that the infinite family of spatial modes $\mathcal{U} = \{\boldsymbol{U}_i, i = 1, 2, \dots\}$ is linearly independent (w.r.t. some spatial inner product $\langle \cdot, \cdot \rangle_{\Omega_x}$), such that any vector field $\boldsymbol{u}$ has a unique decomposition on $\mathcal{U}$ with coefficients $(a_1, a_2, \dots)$. For reduction purpose, the expansion is truncated, retaining a finite number N of modes in the representation, leading to the N-order representation $\boldsymbol{U}^N$ of $\boldsymbol{u}$:

$$\boldsymbol{u}(\boldsymbol{x},t) \approx \boldsymbol{U}^N(\boldsymbol{x},t) = \sum_{i=1}^{N} \boldsymbol{U}_i(\boldsymbol{x}) a_i(t). \tag{2}$$

Clearly, the convergence property of $\boldsymbol{U}^N$ to $\boldsymbol{u}$ as $N \to \infty$ depends on the family $\mathcal{U}$ considered. In fact, the basis $\mathcal{U}$ should be such that it minimizes the truncature error $u - U^N$ in some sense. It also should contain essential features of the dynamics in view of constructing the reduced model of the control flow (see discussion below). Different methods are available for the determination of the set of modes $\boldsymbol{U}_i$. In the present work, we used the Proper Orthogonal Decomposition (POD) originally proposed in [20, 21] in the context of turbulent flows. The basis is constructed from a set of observations $\{\boldsymbol{u}_i := \boldsymbol{u}(\boldsymbol{x}, t_i), i = 1, \dots, m > N\}$ of the flow field at distinct times t_i (the so-called snapshots [31]). From this set, the "correlation" matrix $M_{i,j} := \langle \boldsymbol{u}_i, \boldsymbol{u}_j \rangle_{\Omega_x}$ is formed. The Singular Value Decomposition of M yields the singular values σ_i, ordered as $\sigma_1 \geq \sigma_2 \geq \dots$, and associated normalized vectors $\boldsymbol{m}^i \in \mathbb{R}^m$. Assuming rank $(M) \geq N$, we define the spatial modes as

$$\boldsymbol{U}_i(\boldsymbol{x}) := \frac{1}{\sqrt{\sigma_i}} \sum_{l=1}^{m} \boldsymbol{m}_l^i \, \boldsymbol{u}_l(\boldsymbol{x}), \quad i = 1, \dots, N. \tag{3}$$

This definition of the spatial modes has several attractive characteristics. First, the modes are orthonormal, a handy feature for computational manipulation, because (observing that M is symmetric)

$$\langle \boldsymbol{U}_i, \boldsymbol{U}_j \rangle_{\Omega_x} = \frac{1}{\sqrt{\sigma_i \sigma_j}} \sum_{l=1}^{m} \sum_{l'=1}^{m} \boldsymbol{m}_l^i \langle \boldsymbol{u}_l, \boldsymbol{u}_{l'} \rangle_{\Omega_x} \boldsymbol{m}_{l'}^j = \frac{\sigma_j}{\sqrt{\sigma_i \sigma_j}} \boldsymbol{m}^i \cdot \boldsymbol{m}^j = \delta_{i,j}. \tag{4}$$

As a consequence, the temporal coefficients can be simply retrieved by

$$a_i(t) = \langle \boldsymbol{u}(\boldsymbol{x},t), \boldsymbol{U}_i(\boldsymbol{x}) \rangle_{\Omega_x}. \tag{5}$$

Second, it can be shown that the decomposition (2) with the definition of the modes in (3) minimizes the mean square representation error, for given order N, with regard to the set of snapshots. In other words, the singular values provides an immediate way to select the relevant modes among the family.

The time-dependent spatial solution $\boldsymbol{u}(\boldsymbol{x},t)$ is now described by the N temporal coefficients $a_n(t)$ associated with their corresponding spatial modes $\boldsymbol{U}_i(\boldsymbol{x})$. A crucial issue is to ensure the reduced order model remains a good approximation of the flow when the control is activated, that is, the reduction must be robust. This critical point must be accounted for when deriving a reduced order model and the basis used in the present work includes some robustness constraints. The interested reader may refer to [23] where the derivation of the robust order reduced model used in the present work was investigated. In the rest of the paper, the ROM will then be assumed not to bring a significant source of error.

We assume that the resulting ROM is a causal autonomous dynamical system and that its dimensionality N is large enough so that the reduced order model attractor is fully developed, *i.e.*, the trajectory of the system state never crosses itself in the state space. This may be verified with the Takens procedure [33]. We further rely on the hypothesis of scale separation between modes included in the ROM ($\boldsymbol{U}_{i \leq N}$) and discarded

ones ($\boldsymbol{U}_{i>N}$). This hypothesis means that the truncated modes are somehow slaved to the dominant ones in that their dynamics is constrained by that of the modes included in the ROM. Under this hypothesis, the control command controlling the system may be assumed to be function of the first N modes solely.

To integrate the reduced order model in time, a popular choice is the Galerkin approach, see for instance [13]. Approximation (2) is introduced in the governing equations of the system and the residual is requested to be orthogonal to the space spanned by the approximation functions $\boldsymbol{U}_i$, resulting in a set of ODEs. This procedure has long been recognized as unstable, [13], and an alternative, equation-free, strategy proposed in [18] was employed instead and allows a robust and stable prediction of the system state, even over long time horizons.

2.2. Optimal control

We now focus on the derivation of the optimal control strategy. Given an initial state of the system at time t, expressed in its reduced coordinates $\boldsymbol{a}(t)$, we seek for the control law $\boldsymbol{\mu}(t)$ that drives the system toward some desirable state, or maximize some prescribed performance measure, by modifying the system dynamics. To remain as general as possible, we assume in the following that $\boldsymbol{\mu}(t)$ is a real-valued vector having N_μ components: $\boldsymbol{\mu}(t) \equiv \big(\mu_1(t) \ \dots \ \mu_{N_\mu}(t)\big) \in \mathbb{R}^{N_\mu}$. To account for the effect of the control onto the system dynamics we shall write

$$\dot{\boldsymbol{a}}(t) = \mathcal{F}_\mu(\boldsymbol{a}(t), \boldsymbol{\mu}(t)), \tag{6}$$

the dynamics of the controlled ROM. We assume a numerical procedure is available for the integration of (6) from a given control law and initial state. It is observed that (6) implicitly assumes that the effect of the control applied in the past ($t' < t$) only impacts the future dynamics (at time $\geq t$) through the resulting current state $\boldsymbol{a}(t)$.

In practice, the control is determined for a finite time interval $[t, t+T]$ as the solution of a constrained optimization problem. This problem can be conveniently reformulated as an unconstrained optimization one and reduces to the minimization of a Lagrangian $\mathscr{L}$,

$$\begin{aligned}\mathscr{L}(\boldsymbol{a}, \boldsymbol{\mu}, \boldsymbol{\lambda}) \equiv &\int_0^T \left[D^2(\boldsymbol{a}(t+\tau), \boldsymbol{\mu}(t+\tau)) + w_{\boldsymbol{\mu}}\, \boldsymbol{\mu}^T(t+\tau)\, \boldsymbol{\mu}(t+\tau)\right] d\tau \\ &+ \int_0^T \boldsymbol{\lambda}^T(\tau)\, \left[\dot{\boldsymbol{a}}(t+\tau) - \mathcal{F}_\mu(\boldsymbol{a}(t+\tau), \boldsymbol{\mu}(t+\tau))\right] d\tau,\end{aligned} \tag{7}$$

with $\boldsymbol{\lambda}(\tau) \in \mathbb{R}^N$ the time-dependent vector of Lagrange multipliers, $D^2(\boldsymbol{a}, \boldsymbol{\mu}) \in \mathbb{R}^+$ the convex objective function to minimize and $w_{\boldsymbol{\mu}}$ a user-set parameter weighting the regularization term with respect to the cost term. Upon minimization, the Lagrangian is stationary, hence satisfying a set of optimality equations,

$$\frac{d\mathscr{L}}{d\boldsymbol{\lambda}} = \dot{\boldsymbol{a}}(t+\tau) - \mathcal{F}_\mu(\boldsymbol{a}(t+\tau), \boldsymbol{\mu}(t+\tau)) = 0 \qquad \text{(State equation)}, \tag{8a}$$

$$\begin{aligned}\frac{d\mathscr{L}}{d\boldsymbol{a}} = &\, 2\, D(\boldsymbol{a}(t+\tau), \boldsymbol{\mu}(t+\tau))\, \boldsymbol{\nabla}_{\boldsymbol{a}} D(\boldsymbol{a}(t+\tau), \boldsymbol{\mu}(t+\tau)) - \dot{\boldsymbol{\lambda}}(\tau) \\ &- \boldsymbol{\lambda}^T(\tau)\, \boldsymbol{\nabla}_{\boldsymbol{a}} \mathcal{F}_\mu(\boldsymbol{a}(t+\tau), \boldsymbol{\mu}(t+\tau)) = 0 \qquad \text{(Adjoint equation)},\end{aligned} \tag{8b}$$

$$\frac{d\mathscr{L}}{d\boldsymbol{\mu}} = 2\, \boldsymbol{\mu}(t+\tau) - \boldsymbol{\lambda}^T(\tau)\, \boldsymbol{\nabla}_{\boldsymbol{\mu}} \mathcal{F}_\mu(\boldsymbol{a}(t+\tau), \boldsymbol{\mu}(t+\tau)) = 0 \qquad \text{(Optimality equation)}, \tag{8c}$$

where the derivatives are understood in the Gâteaux sense. This system is complemented with appropriate initial and terminal conditions. Specifically, the terminal condition for $\boldsymbol{\lambda}$ is $\boldsymbol{\lambda}(T) = \boldsymbol{0}$. Solving this set of equations yields the optimal control law $\boldsymbol{\mu}(t+\tau)$ and the corresponding state trajectory $\boldsymbol{a}(t+\tau)$ over the entire time-span $[t, t+T]$.

3. Approximation of the Parameterized Optimal Control

3.1. Parameterized optimal control problem

As discussed previously, one would like to recompute the control law each time a state $\boldsymbol{a}(t)$ (or its estimation) of the system is made available. In the following, we assume that updates of the system state are provided with a frequency greater that $1/T$, so that a new optimal control command $\boldsymbol{\mu}$ can be computed, using the optimal control machinery introduced above, before reaching the end of the current control time-span. However, the resolution of the optimal control problem is too costly to be performed on-line. Fortunately, inspection of the control problem shows that it can be seen as an optimization problem parameterized by the initial state $\boldsymbol{a}(t)$, specifically

$$\boldsymbol{\mu}(t+\tau) = \mathcal{M}_T\left(\tau; \boldsymbol{a}(t)\right), \quad \tau \in [0, T]\,. \tag{9}$$

What is proposed here is to construct in a pre-processing step (off-line) an approximation of $\mathcal{M}_T$, with a low evaluation cost, and to rely on this approximation for the on-line control when one can not afford sophisticated computational operations. To this end, we assume that the state vectors $\boldsymbol{a}$ of the controlled system remains in a bounded domain $\Omega_{\boldsymbol{a}} \subset \mathbb{R}^N$ at all time. For simplicity, we shall restrict ourselves to domains $\Omega_{\boldsymbol{a}}$ having product form,

$$\Omega_{\boldsymbol{a}} \equiv [a_1^-, a_1^+] \times [a_2^-, a_2^+] \times \ldots \times [a_N^-, a_N^+], \tag{10}$$

where a_n^- and a_n^+ are the lower and upper bounds of the n-th coefficient $a_n(t)$ of the system. These bounds can be estimated from observations of the actual system or from theoretical considerations. Since the dependence of the control law $\mathcal{M}_T(\tau; \boldsymbol{a})$ with regard to the initial condition $\boldsymbol{a} \in \Omega_{\boldsymbol{a}}$ is not known explicitly, one should rely on an approximation of it. It is chosen to approximate $\mathcal{M}_T(\tau; \boldsymbol{a})$ by a finite series of the form

$$\mathcal{M}_T(\tau; \boldsymbol{a}) \approx \sum_{l \in S} \mathcal{Q}_l(\tau) \Psi_l(\boldsymbol{a}), \tag{11}$$

where S is an index set.

In this work, the control laws will be discretized in time with a fixed time-step $\Delta t = T/(n_T - 1)$, such that one actually seeks for a set of n_T approximations

$$\mathcal{M}_T(\tau_i; \boldsymbol{a}) \approx \sum_{l \in S} \mathcal{Q}_{l,i} \Psi_l(\boldsymbol{a}), \quad \tau_i = (i-1)\Delta t,\, i = 1, \ldots, n_T. \tag{12}$$

However, to avoid the introduction of an additional time-index and unnecessary complex notation, the description of the procedures to construct the approximation of $\mathcal{M}_T$ will be presented in continuous time, but one should keep in mind that they apply for each $\tau_i \in [0, T]$.

Remark.
Under formulation (9), the control takes the form of a classic closed-loop where the control command directly depends on the current state of the system. As for all closed-loop configurations, one faces here the issue of knowing the state of the system and should consider the use of a state estimator. While this is a critical point and the subject of many efforts in the literature, this is not the primary focus of the present work and we will make the assumption that the state $\boldsymbol{a}$ of the system may be estimated at any time.

3.2. Polynomial approximation

To construct a representation as given in Eq. (11), a classical approach consists in selecting the Ψ_l to be polynomial functions in $\boldsymbol{a}$, and then to define the respective coefficients $\mathcal{Q}_l$ such that the approximation error is minimized in some prescribed sense. It is convenient to work on polynomial functions $\{\Psi_l\}$ forming

an orthogonal set. This calls for the definition of an inner product on $\Omega_{\boldsymbol{a}}$, which will be denoted by $\langle \cdot, \cdot \rangle_{\Omega_{\boldsymbol{a}}}$. Specifically, let $W : \mathbb{R}^N \to \mathbb{R}^+$ be a given weight function on $\Omega_{\boldsymbol{a}}$, such that

$$W(\boldsymbol{a}) = \begin{cases} 0 & \text{if } \boldsymbol{a} \notin \Omega_{\boldsymbol{a}}, \\ > 0 & \text{if } \boldsymbol{a} \in \Omega_{\boldsymbol{a}}. \end{cases} \tag{13}$$

We denote $\| \cdot \|_W$ the induced norm and we further assume that $\mathcal{M}_T(\cdot; \boldsymbol{a}) \in L_2(\Omega_{\boldsymbol{a}}, W)$, that is

$$\|\mathcal{M}_T(\tau; \cdot)\|_W^2 = \langle \mathcal{M}_T(\tau; \cdot), \mathcal{M}_T(\tau; \cdot) \rangle_{\Omega_{\boldsymbol{a}}} = \int_{\Omega_{\boldsymbol{a}}} \mathcal{M}_T(\tau; \boldsymbol{a})\, \mathcal{M}_T(\tau; \boldsymbol{a})\, W(\boldsymbol{a})\, d\boldsymbol{a} < +\infty, \quad \forall \tau \in [0, T]. \tag{14}$$

The orthonormality of the functions $\{\Psi_l\}$ then writes

$$\forall (l, l') \in S,\ \langle \Psi_l, \Psi_{l'} \rangle_{\Omega_{\boldsymbol{a}}} = \delta_{l,l'} = \begin{cases} 1 & \text{if } l = l', \\ 0 & \text{if } l \neq l'. \end{cases} \tag{15}$$

With the previous assumptions, the set $\{\Psi_l, l \in S\}$ forms an orthogonal basis which spans the approximation space $\mathcal{A}_S$:

$$\mathcal{A}_S := \operatorname{span}\{\Psi_l, l \in S\} \subseteq L_2(\Omega_{\boldsymbol{a}}, W). \tag{16}$$

Clearly, the approximation space $\mathcal{A}_S$ depends both on the selected weight W and the index set S. In this work, a constant weight $W(\boldsymbol{a} \in \Omega_{\boldsymbol{a}}) = c$ is used, allowing for great simplifications in the construction of the functions Ψ_l. Indeed, in the case of a constant weight and domains having the product form (10), the Ψ_l can be defined as the product of rescaled and translated one-dimensional Legendre polynomials, that is

$$\Psi_l(\boldsymbol{a}) = \prod_{i=1}^{N} \Lambda_{\alpha_i(l), [a_i^-, a_i^+]}(a_i), \tag{17}$$

where $\Lambda_{\alpha, [a^-, a^+]}(a)$ is the Legendre polynomial of degree α [1], rescaled and translated to the interval $[a^-, a^+]$. The approximation basis is then fully defined by prescribing the set of multi-indexes $\boldsymbol{\alpha}(l) = (\alpha_1(l), \ldots, \alpha_N(l))$ for each $l \in S$. Classically, the index set S is constructed retaining all multi-indices $\boldsymbol{\alpha}(l)$ corresponding to polynomials Ψ_l of total degree less or equal to p:

$$S = \left\{ l, |\boldsymbol{\alpha}(l)| = \sum_{i=1}^{N} \alpha_i(l) \leq p \right\}. \tag{18}$$

The degree p is called the truncation order of the polynomial approximation and the cardinality P of S is function of p and N through

$$\operatorname{card}(S) = \binom{N+p}{N} =: P. \tag{19}$$

The approximation basis being set, we now focus on the definition of the approximation and introduce a methodology for the computation of the coefficients $\mathcal{Q}_l$ in (11). Different strategies can be thought of depending on the definition of the approximation error. A sound approach is obtained by defining the approximation as the L_2-projection of $\mathcal{M}_T(\cdot; \boldsymbol{a})$ onto $\mathcal{A}_S$. Owing to the orthonormality of the basis of $\mathcal{A}_S$, this definition leads to the following expression for the coefficients $\mathcal{Q}_l(\cdot)$:

$$\mathcal{Q}_l(\tau) = \langle \mathcal{M}_T(\tau; \cdot), \Psi_l \rangle_{\Omega_{\boldsymbol{a}}} = \int_{\Omega_{\boldsymbol{a}}} \mathcal{M}_T(\tau; \boldsymbol{a})\, \Psi_l(\boldsymbol{a}) W(\boldsymbol{a})\, d\boldsymbol{a}, \quad \forall l \in S. \tag{20}$$

This definition yields the best approximation of $\mathcal{M}_T$ in $\mathcal{A}_S$ in the L_2-sense in that it minimizes the norm $\|\mathcal{M}_T - \sum_l \mathcal{Q}_l \Psi_l\|_W$.

Two difficulties arise for the computation of $\mathcal{Q}_l$ from (20). First, no closed-form is available for $\mathcal{M}_T\left(\tau;\boldsymbol{a}(t)\right)$ preventing an exact computation of the right-hand-side of (20). Second, one has to compute a high-dimensional integral when N is large. Therefore, one needs to rely on an efficient numerical strategy where the integral is approximated by a finite sum of the form

$$\mathcal{Q}_l(\tau) \simeq \sum_{q=1}^{N_q} \mathcal{M}_T\left(\tau;\boldsymbol{a}^{(q)}\right)\, \Psi_l(\boldsymbol{a}^{(q)})\,\omega^{(q)}, \quad \forall l \in S, \tag{21}$$

with $\boldsymbol{a}^{(q)} \in \Omega_{\boldsymbol{a}}$ and $\omega^{(q)}$ the integration points and weights of the numerical integration. Different methods can be used for the selection of the integration points. These range from random sampling of $\Omega_{\boldsymbol{a}}$ (Monte-Carlo integration) to the use of deterministic cubature formulas. In this work, we take advantage of the particular structure of the integration domain $\Omega_{\boldsymbol{a}}$, see (10), and of the constant weight function W to construct cubature formulas from the Smolyak's scheme, see [32, 26, 27, 29]. The construction of the cubature formulas proceeds from a sparse tensorization of a series of nested one-dimensional quadrature formulas (Konrod-Patterson delayed formulae [28] are used in this work). The procedure aims at minimizing the number N_q of integration points, and hence the number of optimal control problems to be solved, while maintaining the highest order of polynomial exactness for the resulting numerical integration.

Remark 1. *More advanced sparse grid integration methods can be used for the construction of the integration points and weights in* (21)*, with further reduction of N_q. In particular, we mention the dimension-adaptive sparse grid technique [25] and fully adaptive sparse grid [11] techniques which rely on adaptive strategies for the definition of the set of integration points. Instead of improving our numerical integration method along these directions, we focus in the following section on the compressed-sensing approach, which is expected to yield a substantial reduction in the number of optimal control law evaluations, as well as a coarser approximation (lower dimensionality of the approximation basis), and assess its efficiency compared to the original (isotropic, non-adapted) Smolyak's cubature.*

Remark 2. *An alternative to the L_2-projection is the regression approach, where one minimizes the distance between the approximation and the actual control law over a set of specific initial conditions $\{\boldsymbol{a}^{(q)}\}$. A common formulation is obtained by minimizing the sum d^2 of the squares of the residuals,*

$$d^2 = \sum_{q=1}^{N_q} \left(\mathcal{M}_T\left(\tau;\boldsymbol{a}^{(q)}\right) - \sum_{l\in S} \mathcal{Q}_l(\tau)\Psi_l\left(\boldsymbol{a}^{(q)}\right)\right)^2. \tag{22}$$

The regression and L_2-projection are equivalent in the limit $N_q \to \infty$ and when the regression points are sampled in $\Omega_{\boldsymbol{a}}$ with the density proportional to W. Although the asymptotic convergence rate toward the exact projection coefficients of the regression coefficients is generally slow with N_q compared to the sparse-grid integration approach, the regression approach is interesting when one can only afford a number of computations which is not large enough to perform a sufficiently accurate integration: the regression is generally more robust than the integration for a low N_q, which however need be at least greater than P. We shall see in the next section that the compressed-sensing approach takes advantage of this property, while also improving the quality of the coefficients through a reduction of the effective dimension of the approximation space.

3.3. A Compressed-Sensing approach

3.3.1. Basics

In this section, we focus on the approximation of the parameterized control law $\mathcal{M}_T(\tau;\boldsymbol{a})$ at a fixed time τ, so that the explicit dependence with τ will be omitted to alleviate notations. We introduce an alternative approach to the sparse grid integration discussed above that aims at further reducing the number N_q of optimal problems to solve. For sufficient regularity of the optimal law (with regard to the initial condition)

and large enough basis $\{\Psi_l\}$, one can assume that the coefficients $\mathcal{Q}_l$ exhibit -upon appropriate indexation- a decay rate bounded from above as

$$|\mathcal{Q}_l| \lesssim l^{-1/r}, \quad \forall l \in S,\ 0 < r \leq 1. \tag{23}$$

If the above condition is fulfilled, $\mathcal{M}_T$ is said compressible in the approximation basis $\{\Psi_l\}$ and a significant number of coefficients $\mathcal{Q}_l$ are then negligible, w.r.t. the dominant ones, and can be disregarded. It suggests that one should not treat identically all the coefficients but focus the available information onto the evaluation of the relevant ones, *and only them*. This would relax the constraints mentioned in remark 2 on the number of sampling points and allow for significant computational savings. Considering the problem with such objective in mind, the so-called Compressed Sensing (CS) approach is the method of choice. For a more complete introduction to the CS theory, we refer to [5, 6, 9], and restrict ourselves to the presentation of its essential steps.

The objective of CS is to construct a sparse representation minimizing in some sense the approximation error

$$\mathcal{R}_{\boldsymbol{Q}}(\boldsymbol{a}) = \mathcal{M}_T(\boldsymbol{a}) - \sum_{l \in S} \mathcal{Q}_l \Psi_l(\boldsymbol{a}) = \mathcal{M}_T(\boldsymbol{a}) - \boldsymbol{\Psi}(\boldsymbol{a})\boldsymbol{Q}. \tag{24}$$

To this end, a measurement basis $\{\phi_q(\boldsymbol{a}), 1 \leq q \leq N_q\}$ is selected together with an inner product $\langle \cdot, \cdot \rangle$ on $\Omega_{\boldsymbol{a}}$. We set for $1 \leq q \leq N_q$

$$y_q \equiv \langle \phi_q, \mathcal{M}_T \rangle, \quad \widehat{y}_q \equiv \langle \phi_q, \boldsymbol{\Psi}\boldsymbol{Q} \rangle = [S]\boldsymbol{Q}, \tag{25}$$

where the matrix $[S] \in \mathbb{R}^{N_q \times P}$ has for entries $[S]_{q,l} = \langle \phi_q, \Psi_l \rangle$. The CS solution $\boldsymbol{Q}^\star \in \mathbb{R}^P$ is defined as

$$\boldsymbol{Q}^\star = \arg\min_{\boldsymbol{Q} \in \mathbb{R}^P} \|\boldsymbol{Q}\|_0, \quad s.t. \quad \boldsymbol{Y} = [S]\,\boldsymbol{Q}, \tag{26}$$

where $\boldsymbol{Y} \equiv \left(y_1 \dots y_{N_q}\right)^T$. In words, the CS solution is the sparsest vector, that is having the minimal L_0-norm, that represents the N_q observations.

The constrained optimization problem (26) involves a non-convex cost functional and is combinatorial in nature. It results in a NP-hard problem which solution is not computable in practice for most cases. Fortunately, it was shown that the L_0-norm may be substituted with the L_1-norm, which is convex, while the solution $\boldsymbol{Q}^\star$ remains the same. Further, since the set of approximating functions $\{\Psi_l, l \in S\}$ is not a complete basis, the observations can not be exactly retrieved in general. An approximation error ε then exists which value depends on the truncation order p. The optimization problem (26) then has to be reformulated into the following inequality-constrained optimization problem:

$$\boldsymbol{Q}^\star = \arg\min_{\boldsymbol{Q} \in \mathbb{R}^P} \|\boldsymbol{Q}\|_1, \quad s.t. \quad \|\boldsymbol{Y} - [S]\,\boldsymbol{Q}\|_2 \leq \varepsilon. \tag{27}$$

In this form, the problem is called basis pursuit and may be efficiently solved by reformulating it as a linear program. The solution method relies on an optimized choice of measurement functions set $\{\phi_q\}$. In this work, the sensing elements ϕ_q are simply taken as Dirac distributions $\phi_q(\boldsymbol{a}) = \delta(\boldsymbol{a} - \boldsymbol{a}^{(q)})$, where the measurement points $\boldsymbol{a}^{(q)}$ are sampled at random in $\Omega_{\boldsymbol{a}}$. For this choice, the entries of the matrix $[S]$ are $S_{q,l} = \Psi_l\left(\boldsymbol{a}^{(q)}\right)$, $1 \leq q \leq N_q$, $1 \leq l \leq P$.

Usually, the determination of the P coefficients $\{\mathcal{Q}_l, l \in S\}$ calls for $N_q \geq P$ measurement points. However, if the signal $\mathcal{M}_T(\cdot; \boldsymbol{a})$ is compressible enough in the orthonormal trial basis $\{\Psi_l, l \in S\}$, the exact signal can be recovered with far fewer evaluation points, $N_q < P$, hence the term "compressed sensing" to refer to this technique. This remarkable result has received considerable attention in the past years.

3.3.2. Solution method

As seen above, the problem takes the form of an inequality constrained optimization problem (27), to be solved for $\boldsymbol{Q} \in \mathbb{R}^P$. A great deal of work has been devoted to this class of problems and its alternative formulations, such as the convex constrained version

$$\boldsymbol{Q}^\star \equiv \arg\min_{\boldsymbol{Q}} \|\boldsymbol{Y} - [S]\,\boldsymbol{Q}\|_2, \quad s.t. \quad \|\boldsymbol{Q}\|_1 \leq \epsilon_{\boldsymbol{Q}} \quad \text{where } \epsilon_{\boldsymbol{Q}} > 0, \tag{28}$$

and the (convex) unconstrained optimization version

$$\boldsymbol{Q}^\star \equiv \arg\min_{\boldsymbol{Q}} \|\boldsymbol{Q}\|_1 + \eta \, \|\boldsymbol{Y} - [S]\,\boldsymbol{Q}\|_2 \quad \text{where } \eta > 0. \tag{29}$$

Upon an appropriate choice of $\epsilon_{\boldsymbol{Q}}$ and η, problems (27), (28) and (29) have the same solution.

Algorithms vary depending on the formulation of the problem used. For instance, ideas from the Least Angle Regression (LARS) procedure [10] may be used to solve formulation (28), while formulation (27) may be recast as a second order cone program for which efficient algorithms exist, see for instance [3]. In the present work, the formulation (29) is used together with a memory-limited second-order quasi-Newton approach [12]. Note that this choice is in no way unique and one may consider projected gradient techniques [34], Interior Points methods (IP) [35], or iterative-shrinkage techniques [36]. For the examples shown in the next section, the memory-limited second-order quasi-Newton approach was found very efficient so that alternatives have not been investigated in the present work.

The solution to (29) depends on the value of η which is unknown. It balances the reconstructed signal norm $\|\boldsymbol{Q}\|_1$ and the approximation error $\|\boldsymbol{Y} - [S]\,\boldsymbol{Q}\|_2$. A weakly penalized observation constraints (low η) results in an approximation lying on a too low-dimensional subspace; *a contrario*, a large η may result in over-fitting on the available observations. A cross-validation technique is employed to estimate the optimal η, denoted η^*, along the Pareto front. Specifically, $\eta^\star$ is estimated through a K-fold cross-validation technique as the value of η which minimizes the mean reconstruction error over the K folds,

$$\eta^\star = \arg\min_\eta \sum_{k=1}^{K} \left\| \mathcal{M}_T^{(k)} - \boldsymbol{\Psi}^{(k)}\, \boldsymbol{Q}^\star(\eta) \right\|_2, \tag{30}$$

where $\mathcal{M}_T^{(k)}$ and $\boldsymbol{\Psi}^{(k)}$ are evaluated on the k-th set of points $\left\{\boldsymbol{a}^{(q,k)}, q = 1, \ldots, N_q/K, k = 1, \ldots, K\right\}$, independent of the set of measurement points used to construct $\boldsymbol{Y}$. The procedure is repeated K times so that each set of N_q/K points alternatively serves for the approximation and the residual norm estimation. $K = 3$ is used in this work. More sophisticated cross-validation techniques such as the Leave-One-Out or the .632+ bootstrap method may lead to lower variance error estimation but are deemed too costly.

Remark 3. *As mentioned at the beginning of the section, the compressed sensing approach is used to approximate $\mathcal{M}_T(\tau; \boldsymbol{a})$ at fixed times τ_i, $i = 1, \ldots, n_T$. Therefore, the complete approximation of the control law involves the solution of n_T independent problems of type* (27)*. Provided the time-extend between two successive discrete times τ_i is low compared to the control law intrinsic time-scale, it is expected that $\boldsymbol{Q}^\star$ does not drastically change from τ_i to τ_{i+1} so that the solution $\boldsymbol{Q}^\star(\tau_i)$ may be efficiently used as an initial condition for the optimization problem at τ_{i+1}.*

Alternatively, a global optimization problem involving an extended solution vector $\widetilde{\boldsymbol{Q}}^T = \left(\boldsymbol{Q}^T(\tau_1) \ldots \boldsymbol{Q}^T(\tau_{n_T})\right) \in \mathbb{R}^{n_T \times P}$ may be considered. The approximation is then described over the whole discretized time domain by the components of $\widetilde{\boldsymbol{Q}}$. This approach involves a high-dimensional ($n_T \times P$) optimization problem and an accordingly large matrix $[S]$. When the time horizon requires a significant number of time steps, the size of the optimization problem grows to intractable level and this later approach is not thought to be a tractable route.

3.4. Algorithm

To complete the exposition of the method, and before showing example results, we provide a sketch of the overall control procedure, distinguishing particularly the off- and in-line components. We assume that a reduced basis $\{\boldsymbol{U}_i,\ i = 1, \ldots, N\}$ has been provided, together with a relevant control domain $\Omega_{\boldsymbol{a}}$ and associated reduced model $\mathcal{F}_\mu$. The procedure is described for the case of the Compressed-Sensing approach only, as it is similar in structure for the integration approach.

- **Off-line step.**
  - 1 Define the set of Compressed -Sensing points $\{\boldsymbol{a}^{(q)},\ q = 1, \ldots, N_q\}$.

– 2 **For** each point $\boldsymbol{a}^{(q)}$, $q = 1, \ldots, N_q$ **do**
Compute and record $\mathcal{M}_T(\tau; \boldsymbol{a}^{(q)})$, $\tau \in [0, T]$, the solution to the optimal control problem (8) with the initial condition $\boldsymbol{a}(t = 0) = \boldsymbol{a}^{(q)}$.

– **End**

– 3 **For** each time τ_i, $i = 1, \ldots, n_T$ **do**
Solve the Compressed-Sensing problem (27) for the set of data $\{\mathcal{M}_T(\tau_i, \boldsymbol{a}^{(q)}),\ q = 1, \ldots, N_q\}$ and record solution $\mathcal{Q}_l(\tau_i)$, $l = 1, \ldots, P$.

– **End**

- **On-line step.**

  – 4 Get/estimate the current reduced state $\boldsymbol{a}$ of the system and set $\tau = 0$.

  – 5 **Repeat**

    * Find i such that $\tau_{i-1} < \tau \leq \tau_i$.
    * For $l = 1, \ldots, P$, interpolate $\mathcal{Q}_l(\tau)$ between $\mathcal{Q}_l(\tau_{i-1})$ and $\mathcal{Q}_l(\tau_i)$.
    * Reconstruct $\boldsymbol{\mu}(\tau) = \sum_{l=1}^{P} \mathcal{Q}_l(\tau)\Psi_l(\boldsymbol{a})$.
    * Apply the optimal control and integrate the system state $\boldsymbol{a}$ over a time step $\delta\tau$ using the reduced model equations (6).
    * $\tau \leftarrow \tau + \delta\tau$.

  – **Until** a new estimate of the state is made available.

  – 6 Reset $\boldsymbol{a}$ to the new state, reset $\tau \leftarrow 0$, and continue from step 4.

For the **off-line** component, we observe that both steps 2 and 3 are computationally demanding. We remark however that if they need be treated sequentially, step 2 first, the two can be parallelized in a straightforward manner: over the set of CS points for step 2, and over the set of recorded times for step 3. In addition, although N_q may be large, we recall that the optimal problems are based on the reduced model which dimensionality is not that large. In fact, the detailed model is not used at any step. We further stress that the overall cost of the off-line component primarily scales with N_q, which is the number of control problems and the dimension of the CS problem that has to be solved, and, to a lower extend, on the control time-range T since it only impacts step 2 (provided n_T remains constant, see discussion below). Note that, in view of optimal control efficiency, one can consider a long time-range T, and record/construct approximations of the control law only for some early times τ_i. This is particularly sensible as one expects new state observations at a sufficient high rate, compared to $1/T$, as to avoid the emergence of significant errors in the time-integration of the controlled ROM, due to unmodeled effects and external perturbations.

Focusing now on the **on-line** component, we see that it essentially amounts to the time-integration of the controlled reduced model using time-interpolation of the optimal law reconstructed for the current state estimate. In this work, we used the simple piecewise constant time interpolation for the approximation coefficients $\mathcal{Q}_l(\tau)$, namely

$$\mathcal{Q}_l(\tau) = \mathcal{Q}_l(\tau_i), \quad \tau_{i-1} < \tau \leq \tau_i. \tag{31}$$

Once the coefficients $\mathcal{Q}_l(\tau)$ have been obtained, the reconstruction of $\boldsymbol{\mu}(\tau)$ only involves the evaluation of P N-variate polynomials at an effectively low computational cost. It thus remains to integrate the reduced model equation $\dot{\boldsymbol{a}} = \mathcal{F}_\mu(\boldsymbol{a}, \boldsymbol{\mu})$ to obtain the new state estimate at time $\tau + \delta\tau$. The sequence of time-interpolation, polynomial reconstruction, time-integration is then repeated until a new observation of the state is made available. In fact, re-initialization of the optimal control strategy at a high frequency rate is here possible: in the examples shown in the next section, the simulations use a new state estimate available at a frequency $(T/n_T)^{-1}$ so we actually only have to record and construct the approximation of the control law at $\tau = \tau_1$.

4. An Application Example

In this Section we illustrate the effectiveness of the proposed methods on a simple flow control problem introduced in section 4.1. We then assess and compare the efficiency and numerical complexity of the two procedures (compressed sensing and integration approaches) for the approximation of the parameterized optimal control law, in section 4.2. The impact on the resulting control performance of both the polynomial order in the control law approximation and model reduction errors are subsequently studied in section 4.3. Finally, the control robustness induced by the capability of restarting the control strategy on the fly is demonstrated in section 4.4.

4.1. Description of the configuration

The methodology presented in the previous section is now demonstrated on a simple example. The considered the control of the two-dimensional cross-flow around a circular cylinder in laminar incompressible regime corresponding to a moderate Reynolds number $Re = 180$ (based on the cylinder radius). The control aims at minimizing the total drag of the cylinder in cross-flow. While rather academic, this simple configuration is thought to involve many essential features to demonstrate the proposed methodology efficiency. In particular, the flow involves a spectrum of spatial and temporal scales requiring a fine spatial and time discretization for its simulation, thus calling for a reduced order model approximation to derive optimal control strategies in a suitable low dimensional space. Further, the control needs be time-dependent, due to the vortex-shedding dynamics, and should be updated with the phase of the flow. Finally, the control has be robust as the drag is sensitive to external perturbations which may contaminate the incident flow. The general system configuration is sketched in Fig. 1.

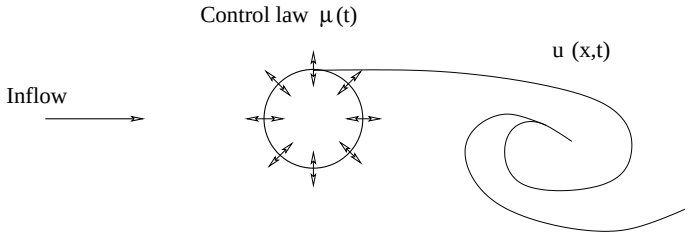

Figure 1: Sketch of the control configuration. The control is achieved by uniformly imposing the normal velocity at the cylinder surface.

The control consists in imposing of a normal flow at the surface of the cylinder, which may then be thought of as porous. Without loss of generality, the normal flow is considered uniform over the surface, leading to a unique time-dependent control parameter, $N_\mu = 1$.

The flow is numerically simulated with a Ψ–ω formulation (stream-function – vorticity). Details about the numerical flow solver are provided in [17, 18]. We rely on a $\mathcal{O}$-type 180×180-mesh having 32400 nodes in space. This detailed discrete flow field is both too fine to be compatible with the real-time control constraint and not observable in practice, preventing a state estimator to be available. An observable reduced model is then to be derived and choice is made here to employ a classic Proper Orthogonal Decomposition. Details on the construction of the reduced model basis can be found in [18]. In the following we use a reduced model of order 9.

As said above, the objective of the control is to decrease the total drag of the cylinder. This objective is supplemented with an upper bound constraint on the control intensity (normal flow velocity) which is justified by several reasons. First, for sake of realism, the control intensity has to be bounded since it is physically achieved by actuators with technologically limited capabilities. Second, the unconstrained optimization problem has unbounded solutions resulting in a physically non-realistic configurations. Finally, and somehow related to the second point, the reduced order model, while robust, can not remain accurate if the flow is too distorted from its non-controlled counter-part.

4.2. Compressed-sensing approximation

The relevance of a compressed-sensing technique for identifying the polynomial coefficients approximating the parameterized control law, as described in Section 3, is first investigated. The accuracy of the CS approximation is assessed and the complexity of its determination, in terms of number of measurement points, is also contrasted with that of the sparse grid integration approach based on Smolyak cubature. Let $\epsilon^\star$ be the relative L_2-norm of the approximation residual:

$$\epsilon^\star = \frac{1}{\|\mathcal{M}_T(\tau;\boldsymbol{a})\|_W} \left\| \mathcal{M}_T(\tau;\boldsymbol{a}) - \sum_{l\in S} \mathcal{Q}_l(\tau)\ \Psi_l(\boldsymbol{a}) \right\|_W , \tag{32}$$

where coefficients $\mathcal{Q}_l$ are evaluated either using (29) (CS approach) or Eq. (21) (Smolyak cubature approach). Since the exact control law $\mathcal{M}_T$ is unknown, the $\epsilon^\star$ has to be estimated, and we again rely on a Smolyak cubature formula with high order of accuracy to obtain well converged error estimates.

Residual norms are plotted in Fig. 2 for approximations involving a variable number N_q of sampling points and different polynomial total orders p. For the L_2-projection approximation, the Smolyak scheme level is varied from 2 ($N_q{}^{(2)} = 163$) to 7 ($N_q{}^{(7)} = 87823$). For the CS-based approach, different sizes of the set of samples were considered, from $N_q = 20$ to $N_q = 3000$. The samples used in the CS-flavor are selected at random within $\Omega_{\boldsymbol{a}}$.

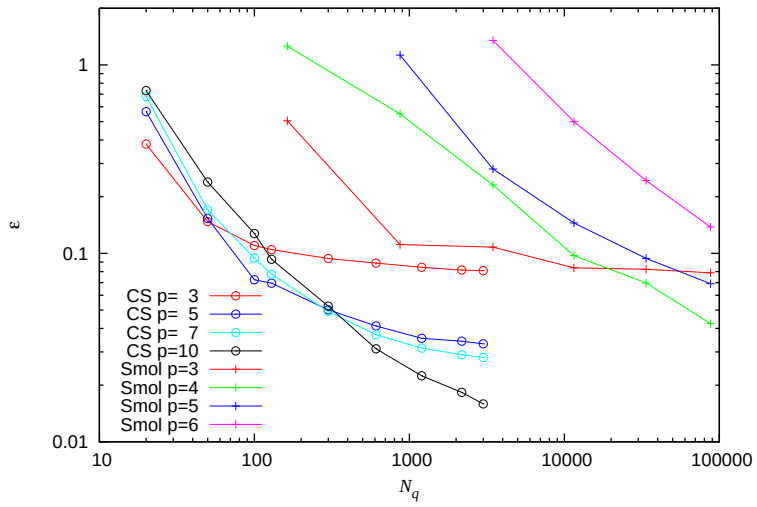

Figure 2: L_2-norm of the residual of the approximation at different polynomial orders p for the Smolyak cubature-based projection (labeled 'Smol') and the compressed-sensing approach (labeled 'CS').

Figure 2 shows that for every approximation order p, the CS approximation achieved an error comparable to the Smolyak approach but with a much lower number of sampling points. For instance, to reach an error $\epsilon^\star \simeq 7\ 10^{-2}$ using a $p = 5$ approximation, the L_2-projection using a Smolyak cubature requires about $N_q \simeq 10^5$ sample points while the CS-based approach only necessitates $N_q \simeq 100$ measurement points. The directly impact the overall cost of obtaining the approximation as it is directly related to the number of sample points: CS approach allows to reduce by about three orders of magnitude the construction cost of the parameterized optimal law as compared to the Smolyak integration approach in that case. Similar observations hold for other polynomial orders and strongly support the CS approach proposed in this work.

Further, for a fixed polynomial order p, the L_2-error decays as we increase the number of measurement and integration points, although after a certain value of N_q the errors level-off, reflecting the limited approximation properties of the truncated basis. Note also that as the dimension of the polynomial basis

increases, larger N_q are needed to achieved similar accuracy in the approximation, while the CS approach is systematically doing a much better job than the Smolyak integration.

The behavior of the CS approach with varying number of measurement points and expansion orders can be better appreciated in Figure 3 where plotted are the L_2-errors $\epsilon^\star$ corresponding to different polynomial orders p obtained using two sets of measurement points with respective dimensions $N_q = 100$ and $N_q = 609$. In the $N_q = 100$ case, the approximation accuracy improves when the polynomial order increases from $p = 0$ to $p = 4$ but deteriorates as p grows beyond. A $p = 10$-approximation then achieves a resulting error about four times larger than using $p = 4$. When $N_q = 609$, while the number of measurements points is sufficiently large to reasonably estimate the approximation coefficients even for the most detailed basis ($p = 10$), the accuracy of the approximation hardly improves when p is larger than 5, again as a symptom of the p / N_q compromise, then letting $p \simeq 5$ appear as a reasonable choice. The impact of the polynomial order onto the controller performance will be specifically investigated in section 4.3.1.

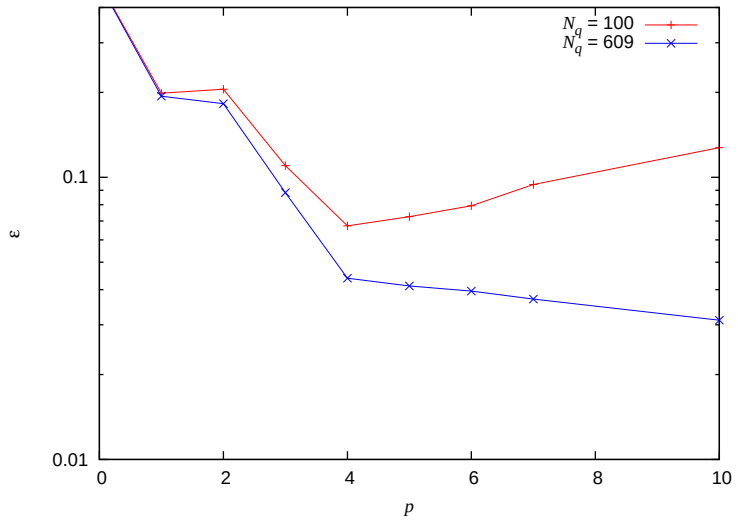

Figure 3: L_2-norm of the approximation residual for different polynomial orders p and two measurement points having size $N_q = 100$ and $N_q = 609$. CS-approach.

The ability of the CS-based approach to accurately approximate the control law can be appreciated from Fig. 4 where the exact parameterized control law ($\mathcal{M}_T(\tau_1)$) is plotted together with its polynomial approximation obtained with the CS approach using $N_q = 3000$ points (left) and from the Smolyak integration approach using $N_q = 87823$ points (right). To ease the visualization, the control law is only shown in the $\{a_1 - a_4\}$-plane (other coefficients being set to value $a_i = 0.5$, their mid-value).

The approximations are seen to be fairly close to the exact control law over the prescribed state vector coefficients domain. The plateau near the edge ($a_1 = 0, a_4 = 1$) is due to the bounds on the control intensity. While both approaches lead to a reasonable approximation, the CS-approximation relying on $N_q = 3000$ and a $p = 10$ polynomial order achieves a better fit. It was seen above in Fig. 2 that this small set of observations is sufficiently large for the polynomials coefficients to be accurately estimated. Conversely, the Smolyak integration approach does not achieve that level of performance as it relies on a $p = 4$ polynomial order, the best achievable with the present $N_q = 87823$ observations set. To highlight the better and much less computationally intensive approximation achieved by the CS approach, the difference between the exact $\{a_1 - a_4\}$ control law and the two approximations are plotted in Fig. 5. Again, the CS technique is clearly seen to achieve a much better approximation while needing roughly 30 times less control law computations. As it clearly outperforms the integration approach, the compressed sensing methodology will be used in all of the computations shown in the remainder of the paper. While not critical for the particular problem treated here, the dimension N of the ROM space being moderate, but in the perspective of dealing with

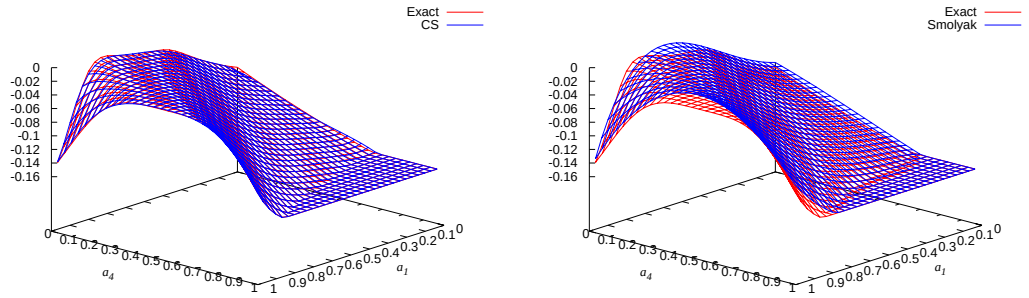

Figure 4: Comparison of the exact control law and its approximations using CS (left) and Smolyak integration (right). Only the plane $\{a_1 - a_4\}$ is shown for clarity. The CS approximation is constructed using $N_q = 3000$ measurements for $p = 10$, while the Smolyak approach uses $N_q = 87823$ integration points for $p = 4$.

more complex systems, a coarse measurement set having only $N_q = 300$ points is used in the following to limit the number of off-line optimal control law resolutions.

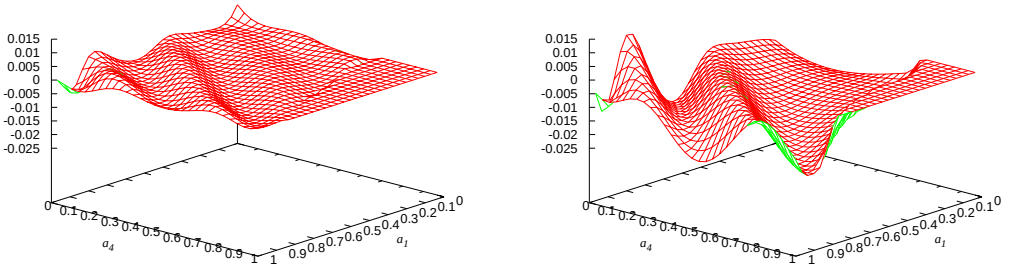

Figure 5: Local error in the control command approximation using CS (left) and Smolyak integration (right). Only the plane $\{a_1 - a_4\}$ is shown for clarity. The CS approximation (left) is constructed using $N_q = 3000$ measurements for $p = 10$, while the Smolyak approach (right) uses $N_q = 87823$ integration points for $p = 4$.

4.3. Approximate controller

4.3.1. Impact of the polynomial order

Having analyzed above the errors in the polynomial approximates of the parametrized optimal law, we now focus on the impact of these approximation errors on the resulting control efficiency. To this end, we perform simulations of the reduced model, without control up to time $t = 500$, when the closed-loop control is eventually turned on. The simulations are performed using approximations with $p = 3$, 4 and 5 in order to assess the impact of the polynomial order. Taking advantage of the moderate dimensionality of the reduced model, we also provide, for comparison, the case of the simulation using a receding-horizon optimal control where the control problem is recomputed at each time step. The later approach, while obviously much more involved from the computational point of view, should be considered as the exact optimal control solution of the reduced system. The simulation results are presented in Figure 6 as the time histories of the control intensities ($\mu(t)$) and the resulting drag coefficients ($C_D(t)$).

Before discussing the convergence of the approximation with increasing polynomial order p, we comment briefly the control on the control law and its performance. The control intensity plots show that

after switching the control on, its intensity oscillates at roughly twice the natural flow frequency (Strouhal number) and remains negative at all times. In fact, a negative intensity corresponds to a suction flow at the cylinder surface. While increasing the viscous dissipation, suction tends to re-attach the boundary layers and postpones flow separation from the surface, with lower pressure drag as a result. The net result of these two trends is a decrease of the total drag. On the contrary, a positive intensity would correspond to a surface blowing, promoting boundary layer separations and widening the near wake albeit with a slight decrease of the viscous drag. The computed intensities are then physically consistent. For the drag coefficient signals, which without control oscillate in the range $1 \leq C_D \leq 1.75$ at twice the Strouhal frequency, we observe that, owing to the control, both the mean drag and its amplitude of fluctuation are reduced. Note however, that the absolute performance of the control, in terms of drag reduction, obviously depends on the compromise one makes between decrease of the drag and control intensity through the choice of the weighting parameter $\omega_{\boldsymbol{\mu}}$. This aspect will not be investigated, as the focus of the paper is on the construction of approximate parameterized optimal control commands for closed-loop type strategy, and not on control performance.

Returning to the analysis of the plots at different orders, the agreement with the exact receding-horizon optimal control is seen to improve from $p = 3$ to $p = 5$, as one might have already anticipated from the result of the previous section. In particular, for $p = 5$, the agreement is excellent, and both the control intensity and drag coefficients are virtually the same as for the receding-horizon optimal strategy. This observation supports the choice of using $p = 5$ in the remainder of the paper.

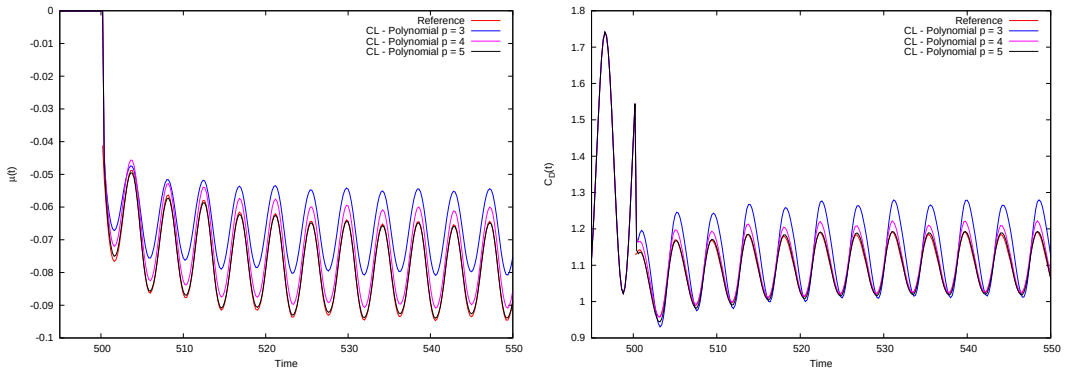

Figure 6: Time evolutions of the control intensities (left) and drag coefficients (right), for the closed-loop control based on the polynomial approximation (CL-Polynomial) with $p = 3 - 5$ and the time receding open-loop optimal control (Reference).

Remark.
While a 5-th order polynomial appears good enough in the present problem, this value can not be generalized. In fact, depending on the reduced model, its dynamics and the optimal control problem formulation (*e.g.*, the constraints considered), there is no guarantee that the optimal control law remains smooth with regard to the initial state of the system. In case where poor regularity is expected, more general approximation bases should be considered and the compressed-sensing approach appears again better suited from this perspective, as it can even handle and work with redundant dictionaries. More specifically, it is no problem adding-up various approximation bases while still being able to retrieve the components which most efficiently approximate the actual "signal". For instance, in addition to the present polynomial basis, it may be beneficial to include compactly supported elements, such as wavelets, and/or enrichment functions. Investigations along these lines are the subject of on-going efforts.

4.3.2. Phase-portraits

To gain further insight into the impact of the control law approximation, we provide in Figure 7 a comparison between the time evolution of the controlled system reduced state for both the approximate control law and the receding-horizon optimal control strategy. Plotted are the projection of the phase

portraits in the plane of dominant modes $\{a_1 - a_2\}$ (left plot) and in the plane of highest order modes $\{a_8 - a_9\}$. Also shown are the corresponding trajectory of the uncontrolled system to appreciate the effect of the control on the dynamics. For the dominant modes, it is remarkable that the approximate and reference controllers lead to indistinguishable phase-portraits, again providing a clear indication of the effectiveness of the approximation strategy. May be even more remarkable is the close agreement between the controlled trajectories for the higher order modes a_8 and a_9: although some differences are visible at the plot scale, the reference and approximate control laws yield trajectories in close agreement considering the significant changes in dynamics, with or without control.

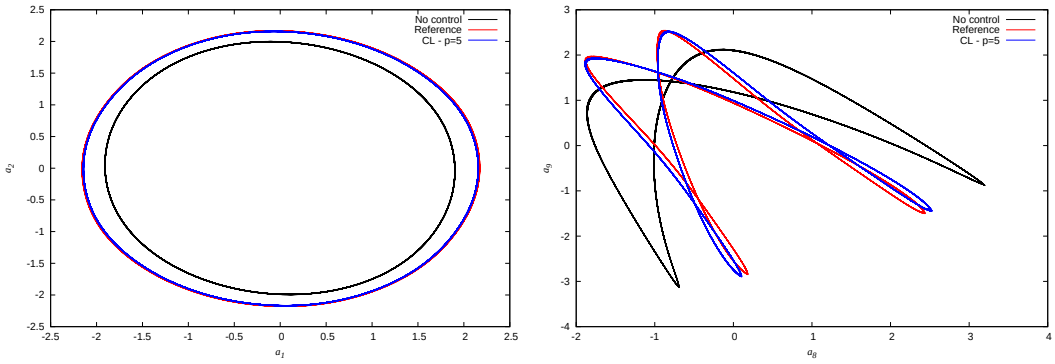

Figure 7: Projection of phase portraits in planes $\{a_1 - a_2\}$ (left) and $\{a_8 - a_9\}$ (left), using approximate closed-loop (CL - $p = 5$) and time receding optimal laws (Reference). Also shown are the phase portraits without control.

4.4. Robustness of the control

4.4.1. Stability analysis

A crucial issue is to ensure that the control strategy is stable, in the sense that a small deviation from the trajectory of the controlled system in the phase space will not result in a diverging regime where the system would divert in time from its nominal trajectory. Instead, the control must be such that any small deviation of the system state induced by some perturbation is damped and that the system eventually returns to the nominal trajectory. Different strategies can be followed to study the stability of a periodic non-linear dynamical system as considered in this work. Here, the stability is investigated in terms of the Floquet matrix [14] which represents the Jacobian, with respect to perturbation, of the trajectory deviation after a period of the system. The deviation of the system trajectory is considered in a Poincaré section. If the largest module of the Floquet matrix eigenvalues is larger than unity, some perturbations to the system will not be damped and the system may destabilize. Several Poincaré sections were considered along the trajectory of the controlled system and the corresponding Floquet matrices were determined. The resulting eigenvalues were found to remain the same along all the Poincaré sections considered, with eigenvalues modules in the range 0 to 0.76, thus providing strong arguments that the controlled periodic system is linearly stable around its optimally controlled trajectory. The upper spectrum of the Floquet matrix at a point of the trajectory is plotted in Figure 8.

4.4.2. Real-life control

The performance of the controller has so far been evaluated in ideal conditions where the future-predicting model accounts exactly for the operating conditions. To validate the whole control approach, it is necessary to consider more realistic situations, where limitations or flaws exist in the model and exogenous perturbations occur. This solicits the robust character of the controller. Without loss of generality, random perturbations were added to the $1st$ reduced order model mode $a_1(t)$ solely (recall that the first mode is the dominant one so its perturbation is likely to have strong impact on the system future). Specifically, a random noise is added to the first component of each new system state. The noise is Gaussian, with zero mean, and time-correlated

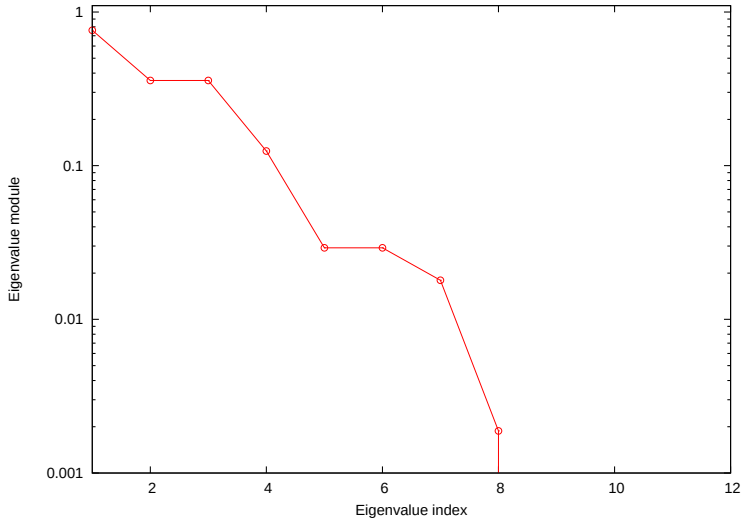

Figure 8: Upper spectrum of the Floquet matrix determined at a point along the controlled system trajectory.

with a characteristic time scale equal to one period of the natural frequency of the flow. The noise level is selected so as to significantly impact the system dynamics while ensuring a trajectory remaining in the domain of the control law approximation. The impact of the perturbation can be appreciated in Figure 9 where the resulting applied control and cylinder drag are plotted as a function of time. These plots have to be compared with their unperturbed counterparts shown in Figure 6. We observe that the control intensity (left plot) has now a wide spectrum of time-scales, contrary to the unperturbed case where, asymptotically, a periodic control was obtained. The drag coefficient (right plot) exhibits similar shaky time-evolution and is now fluctuating over a much broader range than for the unperturbed and uncontrolled problem. These large fluctuations underline the large perturbation introduced for the purpose of robustness assessment. The robustness can in turn be appreciated from Figure 9 where the two cases of approximate closed-loop control and reference receding-horizon optimal control are reported. It is seen that despite a weakly predictable environment, the approximate closed-loop control is able to almost exactly reproduce the optimal control command determined from a receding-horizon optimal procedure. Recalling that this latter open-loop procedure recomputes the control to be applied in the future whenever new information about the state of the system is available (here, at each time-step) and thus constitutes the optimum which can be achieved in a causal framework, the efficiency of the proposed closed-loop approach is impressive: it achieves nearly the same control performance for an in-line cost amounting to a fraction (few orders of magnitude less in the present case) of the receding-horizon approach.

5. Concluding remarks

In this work, an efficient real-time control method was presented. The efficiency comes from the underlying optimality criteria the control command synthesis relies on, while the real-time character is made possible thanks to the very low necessary number of computational operations to reconstruct and update the optimal control when plugged-in. The essence of the proposed method is that most computations are carried-out off-line in a preprocessing step, under mild hypotheses about the behavior of the system to be controlled. In particular, the system is assumed a) to be amenable to an accurate enough reduced order approximation and b) that its state vector representation remains in a given compact domain. The optimal control over a given time-horizon, for a given initial state and control objective, is approximated using an ad-hoc basis (orthogonal polynomials were used but other choices are possible as discussed below). This approximation allows to reconstruct the control command achieving optimally a prescribed objective to

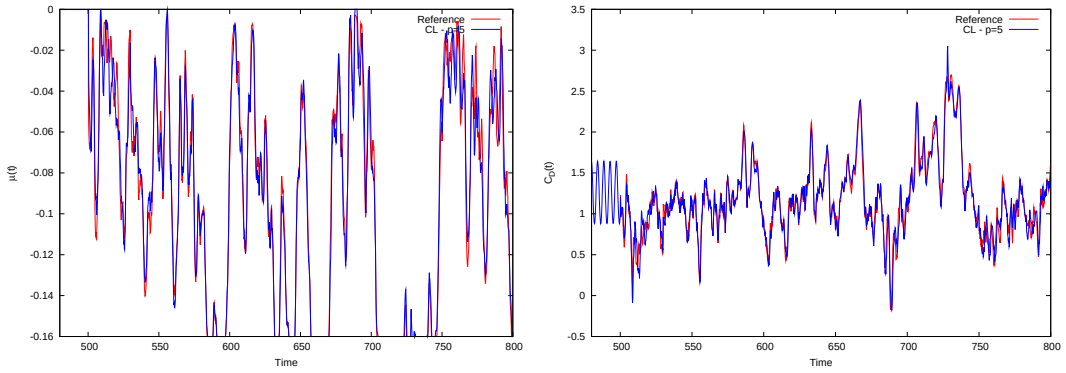

Figure 9: Time-histories of the control intensity (left) and drag coefficient (right) for random perturbations of the dynamics added to dominant reduced mode a_1.Compared are the simulations with the proposed approximate closed-loop strategy (CL - $p = 5$) and the receding-horizon optimal control strategy (Reference).

be applied to the system knowing its current state. Further, since the control is determined based on the current system state, it is updated whenever a new state of the system is made available. Thus, the method accounts for the external / unmodeled perturbations affecting the system. While potentially CPU involved, the preprocessing step, where the parametrized control law approximation is determined, is embarrassingly straightforward to parallelize since it relies on independent computations. In addition, we have shown that compressed-sensing is a promising approach to obtain accurate approximations of the parametrized optimal control law from a limited number calculations. It strongly suggests that compressed-sensing would be a key ingredient in applying the proposed method to high dimensional problems involving hundreds of modes.

The method was applied to the control of the total drag of the flow around a circular cylinder at a low Reynolds number. Results have shown that the proposed methodology achieves performance almost as good as a receding-horizon optimal control approach while necessitating orders of magnitude less computational power when on-line. The quality of the approximation and some related issues were discussed. In the example considered, $5th$-order Legendre polynomials were found sufficient to provide results virtually as good as with a receding-horizon optimal control approach. However, other situations may require higher order polynomial approximations or even the introduction of enrichment functions in the approximation basis to present steep variations or discontinuity of the optimal law. In that case, enriched approximation bases would need be considered, and compressed-sensing seems again well-suited for that purpose (contrary to the integration approach).

In addition, the robust character of the method was demonstrated by considering large aleatory perturbations. The control command was seen to adapt well to fluctuating conditions. This proves the effectiveness of the proposed method which bridges closed-loop capability (the control depends on the current state of the system) and non-linear optimal control nature, thus improving upon the linear framework of Riccati approaches. Supported by its performance, reasonable computational complexity and manageable issues, the method is thought to be of interest for many areas related to systems control. Investigation of some of the specific issues discussed above (case of optimal control laws lacking regularity with regard to the system state, high-dimensional constrained-optimization problem, ...) is the subject of on-going efforts.

References

[1] Abramowitz, M., Stegun, I.: Handbook of mathematical functions. Dover (1970).

[2] Åkervik, E, Hœpffner, J, Ehrenstein, U, Henningson, D.S.: Optimal growth, model reduction and control in a separated boundary layer flow using global eigenmodes. J. Fluid Mech. **579**, 305-314 (2007).

[3] Becker, S, Bobin, J, Candès, E.: NESTA: a fast and accurate first-order method for sparse recovery. Caltech Institute of Technology, Tech. Rep. (2009).

[4] Benner, P, Jing-Rebecca, L, Pentzl, T.: Numerical solution of large-scale Lyapunov equation, Riccati equations, and linear-quadratic optimal control problems. Num. Lin. Alg. Appl. **15**, 755-777 (2008).

[5] Candès, E.J, Tao, T.: Near-optimal signal recovery from random projections: universal encoding strategies. IEEE Trans. Inform. Theory **52**, 5406-5425 (2006).
[6] Candès, E.J, Romberg, J, Tao, T.: Stable signal recovery from incomplete measurements. Comm. Pure Appl. Math. **59**, 1207-1223 (2005).
[7] Candès, E.J, Romberg, J.: Quantitative robust uncertainty principles and optimally sparse decompositions. Found. Comput. Math. **6**(2), 227-254 (2006).
[8] Clenshaw, C, Curtis, A.: A method for numerical integration on a automatic computer. Numer. Math. **2**, 197-205 (1960).
[9] Donoho, D.L.: Compressed sensing. IEEE Trans. Inform. Theory **52**(4), 1289-1306 (2006).
[10] Efron, B, Hastie, T, Johnstone, I, Tibshirani, R.: Least Angle Regression. Ann. Phys. **32**, 407-499 (2004).
[11] Gerstner, T, Griebel, M.: Numerical integration using sparse grids. Numer. Algorithms **18**, 209-232 (1998).
[12] Gilbert, J.C, Lemaréchal, C.: Some numerical experiments with variable-storage quasi-Newton algorithms. Math. Program. **45**, 407-435 (1989).
[13] Iollo, A, Lanteri, S, Désidéri, J.-A.: Stability properties of POD-Galerkin approximations for the compressible Navier-Stokes equations. Theo. Comput. Fluid Dyn. **13**, 377-396 (2000).
[14] Iooss, G, Joseph, D.D..: Elementary stability and bifurcation theory, Springer-Verlag, second edition (1990).
[15] Kim, J, Bewley, T.R.: A Linear Systems Approach to Flow Control. Annual Review of Fluid Mechanics. **39**, 383-417 (2007).
[16] Laub, A.J.: A Schur method for solving Algebraic Riccati equations. IEEE Trans. Auto. Control. **24**, 913-921 (1979).
[17] Le Maître, O, Scanlan, R. H., and Knio, O. M.: Estimation of the flutter derivatives of an NACA airfoil by means of Navier-Stokes simulation. J. Fluids Struct. **17**(1), 1-28 (2003).
[18] Le Maître, O, Mathelin, L.: Equation-free model reduction for complex dynamical systems. Int. J. Numer. Meth. Fluids. **63**(2), 163-184 (2009).
[19] Lewis, F, Syrmos, V: Optimal control, John Wiley & Sons, Inc, second edition (1995).
[20] Lumley J.L.: The structure of inhomogeneous turbulent flows. In A.M. Iaglom and V.I. tatarski editors, Atmospheric Turbulence and Radio Wave Propagation, 221-227 (1967).
[21] Lumley J.L.: Coherent structures in turbulence. In R.E. Meyer editor, Transition in Turbulence, 315-342 (1981).
[22] Mathelin, L, Gallivan, K.: A compressed sensing approach for partial differential equations with random input data. Tech. Rep. FSU10-12, 42 p. (2010).
[23] Mathelin, L, Le Maître, O.: Robust control of uncertain cylinder wake flows based on robust reduced order models. Comput. Fluids **38**, 1168-1182, (2009).
[24] Morris, K, Navasca, C.: Solution of Algebraic Riccati Equations Arising in Control of Partial Differential Equations. In: J.P. Zolesio and J. Cagnol, eds, Control and Boundary Analysis, Lecture Notes in Pure Appl. Math., vol. 240, CRC Press, Boca Raton (2005).
[25] Nobile, F, Tempone, R, Webster, C.G.: An Anisotropic Sparse Grid Stochastic Collocation Method for Partial Differential Equations with Random Input Data. SIAM J. Num. Anal. **46**, 2411-2442 (2008).
[26] Novak, E, Ritter, K.: High-dimensional integration of smooth functions over cubes. Numer. Math. **75**, 79-97 (1996).
[27] Novak, E, Ritter, K.: Simple cubature formulas with high polynomial exactness. Constructive Approximation **15**, 499-522 (1999).
[28] Petras, K.: On the Smolyak cubature error for analytic functions. Advances in Computational Mathematics **12**, 71-93 (2000).
[29] Petras, K.: Fast calculation in the Smolyak algorithm. Numer. Algorithms **26**, 93-109 (2001).
[30] Rugh, W.J, Shamma, J.S.: Research on gain scheduling. Automatica. **36**, 1401-1425 (2000).
[31] Sirovitch, L.: Turbulence and the dynamics of coherent structures part i: Coherent structures. Quart. Appl. Math. **45**:3 561-571 (1987).
[32] Smolyak, S.A.: Quadrature and interpolation formulas for tensor products of certain classes of functions. Dokl. Akad. Nauk. SSSR **4**, 240-243 (1963).
[33] Takens, F.: Detecting strange attractors in turbulence. Lecture Notes in Mathematics **898**, 366-381 (1981).
[34] Van den Berg, E, Friedlander, M.P.: Probing the Pareto frontier for basis pursuit solutions. SIAM J. Sci. Comput. **31**(2), 890-912 (2008).
[35] Wright, S.J.: Primal-Dual Interior-Point Methods. SIAM Publications (1997).
[36] Zibulevsky, M, Elad, M.: L1-L2 optimization in signal and image processing. IEEE Signal Proc. Mag.. **27**(3), 76-88 (2010).