REAL-TIME FLOW ESTIMATION FROM REDUCED ORDER MODELS AND SPARSE MEASUREMENTS

A Picard, M Ladvig, Valentin Resseguier, D Heitz, E Mémin, B Chapron

REAL-TIME FLOW ESTIMATION FROM REDUCED ORDER MODELS AND SPARSE MEASUREMENTS

A. M. Picard[(1)], M. Ladvig[(1)], V. Resseguier[(1)], D. Heitz[(2)], E. Mémin[(3)] and B. Chapron[(4)]

[(1)] Lab, SCALIAN DS, Espace Nobel, 2 Alle de Becquerel, 35700 Rennes, FRANCE,

agustinmartin.picard@scalian.com, valentin.resseguier@scalian.com

[(2)] INRAE, UR OPAALE, 17 Avenue de Cucill, CS 64427, 35044 Rennes Cedex, FRANCE

[(3)] Fluminance team, Inria, Campus de Beaulieu, 263 Avenue Gnral Leclerc, 35042 Rennes, FRANCE

[(4)] LOPS, Ifremer, 1625 Route de Sainte-Anne, 29280 Plouzan, France

ABSTRACT

To successfully monitor and actively control hydrodynamic and aerodynamic systems (e.g. aircraft wings), it can be critical to estimate and predict the unsteady flow around them in real-time. Thus, we introduce a new algorithm to couple on-board measurements with fluid dynamics simulations and prior data in real-time without the need to rely on large computational infrastructure. This is achieved through a combination of a Proper Orthogonal Decomposition Galerkin method, stochastic closure – model under location uncertainty – and a particle filtering scheme. Impressive numerical results have been obtained for a 3-dimensional wake flows at moderate Reynolds for up to 14 vortex shedding cycles after the learning window, using a single measurement point.

1. INTRODUCTION

Accurate aeroelastic and aerodynamics active control (e.g. active flutter suppression, flight stability augmentation, gust and loads alleviation) can require state observers. However, estimating – or even predicting – in real-time an unsteady turbulent flow state from sparse measurements can be challenging. Through statistical estimation techniques, sensor observations can be assimilated to flow dynamical models' predictions, but for this to be a viable strategy, some difficulties must be surmounted first.

Firstly, the system must be able to resolve sufficient spatio-temporal scales for the system to be stable, and this in real-time. Though tempting, tackling this problem with purely data-driven fluid dynamics models (typically learned through machine learning techniques) may not be accurate and/or robust enough. On the other hand, pure physics-based model simulations (e.g. LES, RANS) are often too slow for real-time applications and a speed-up would most likely push it to miss important spatio-temporal scales. Hence, Reduced Order Models (ROM) represent a nice compromise (see e.g. [6] for some aeroelastic applications). In particular, the Proper Orthogonal Decomposition (POD) -Galerkin method is a model that relies on physical equations while constraining the solution to live inside a small subspace learned from data. Nevertheless, the unsteady CFD ROM state of art limits itself mostly to deterministic ROMs (often linear and/or with purely data-driven calibration) with limited prediction capabilities, especially due to the chaotic and intermittent nature of turbulence, and closure problems. The simulation-measurement coupling, known in the literature as data assimilation, can be in itself a challenging task too, but thanks to the weather forecast community, there have been many advances in the field, with promising research and many operational perspectives thriving lately. Notwithstanding, algorithms fully addressing non-linear dynamics – like fluid mechanics – are limited by the available computational power and the dynamical model's accuracy quantification. The former could easily be overcome by the use of the use of our proposed ROM, whereas the second would require a method that can combine two distinct information sources, each with its own accuracy: predictions and measurements. As this not currently possible with the ROM alone, we will rely on so-called dynamics under location uncertainty – a random fluid mechanics framework designed specifically for this purpose [3, 4]. This paper will be organized as follows: section 2 will recall the main aspects of POD-Galerkin ROM, section 3 will present our main contribution: a randomized version of POD-Galerkin ROM, and explain the data assimilation procedure, and finally, section 5 will showcase its potential through some of our numerical results.

2. POD-ROM

The goal of Reduced Order Models (ROM) is to reduce the computational cost of simulations by drastically constricting the solution's degrees of freedom. This gain is typically enabled by a combination of simulation data and modeling based on physical equations. In CFD, solutions i.e. the velocity field have as many degrees of freedom as there are grid points in the spatial domain (typically in the order of 10^6). Thus, to achieve such a dimensionality reduction using ROMs, velocity fields are traditionally

decomposed as follows:

$$v(x,t) = \underbrace{w(x,t)}_{\substack{\text{Resolved}\\ \text{by the ROM}}} + \underbrace{v'(x)}_{\substack{\text{Unresolved}\\ \text{by the ROM}}}, \quad (1)$$

with

$$w(x,t) = \underbrace{\overline{v}(x)}_{\substack{\text{Time}\\ \text{averaged}}} + \underbrace{\sum_{1=1}^{n} b_i(t)\phi_i(x)}_{\substack{\text{Unsteady}\\ \text{component}}}, \quad (2)$$

with $n \in [1, 10^2]$. Proper orthogonal decomposition (POD) learns the time averaged $\overline{v}(x)$ and the spatial modes $\phi(x)$ through principal component analysis (PCA) applied to a set of high-resolution simulation solutions (learning period). Then, one can project the physical equations (e.g. Navier-Stokes equations) onto these spatial modes, thus providing a system of n coupled ordinary differential equations that describe the evolution of the temporal modes $b_i(t)$. Through the time integration of this low-dimensional system and from equation (2), we can predict an a priori estimation of the velocity field at any given time. Consequently, this ROM construction scheme can be seen as a midpoint between fully data-driven methods and pure physics-based model, as it takes advantage of both, the available simulation data and physical modeling in order to infer reliable predictions more efficiently.

3. MODEL UNCERTAINTY QUANTIFICATION AND DATA ASSIMILATION

Models under location uncertainty are a type of random CFD model [3, 4] that provides both an efficient closure (i.e. an efficient way to model the effect of the neglected dynamics degrees of freedom v') and an efficient quantification of the error induced by this closure. Though for it to be tractable, two assumptions must be made: i) the time decorrelation of the unresolved velocity component v' (see eq. (1)) and the stochastic transport – up to some forces F – of the resolved velocity component w. With Ito stochastic calculus notations, this reads:

$$\begin{aligned} \frac{Dw_k}{Dt} &= \partial_t w_k + \left(w - \tfrac{1}{2}(\nabla\cdot a)^T + v'\right)\cdot\nabla w_k - \tfrac{1}{2}\nabla\cdot(a\nabla w_k), \\ &= F_k, \end{aligned} \quad (3)$$

with

$$a_{pq} = \overline{v'_p v'_q}\ \tau_{v'}, \quad (4)$$

being the unresolved velocity (Eulerian) absolute diffusivity, and $\tau_{v'}$ the unresolved velocity correlation time. In order to express that uncertainty (inherent to any closure in CFD), multiple simulations are run in parallel using the stochastic model so as to efficiently realize the most probable future states of the fluid system. Since that stochastic closure is based on physics, its robustness is proved and calibrations can be performed from available physical quantities, hence almost no tuning nor fitting is required. As per chapter 8 of [5], we have implemented the POD-Galerkin of the Navier-Stokes model under location uncertainty, along with the corresponding technical statistical estimations based on stochastic calculus presented there. Interested readers can refer to [5] for more details. The ensuing low-dimensional system incorporates noise terms, thus enabling model uncertainty quantification and strengthening its forecasting capabilities to allow for more stable forecasts beyond the learning period.

4. MEASUREMENTS AND DATA ASSIMILATION

The final block in our stochastic POD-ROM model's pipeline is a particle filter [1] to integrate the real-time measurements coming from all the different sensors. In particular, this scheme was tested with a two dimensional, two component particle image velocimetry (2D2C PIV) experimental arrangement with a linear observation model as per 6.

$$\begin{aligned} y &= Hv + \varepsilon_y, & (5)\\ &= \sum_i (H\phi_i) b_i + \left(Hv' + \varepsilon_y\right), & (6) \end{aligned}$$

where, in order to mimic the PIV measurement process, the linear operator H incorporates a 3-dimensional spatial smoothing operation, the occlusion of the horizontal plane and its corresponding component in the velocity field. Additionally, to make the data assimilation task more challenging, the information relating to a subset of points in the grid was obscured though this operator. Indeed, estimating a vector b of $n \sim 10$ components from a noisy linearly-dependent observation vector y of $M_{PIV} \sim 10^4 (\gg n)$ components could otherwise be solved using a simple least-squares procedure. The parameters inside H and ε_y are estimated using experimental data, comparing the hot wire and PIV measurements' spectrum (using a Taylor assumption). Note that the unresolved velocity v''s strong influence on the final observation model's uncertainty is taken into account through a noise term in (6).

For the first tests presented here, we consider only synthetically generated measurements produced using the aforementioned observation model (5), as this allow us to know the exact values of the velocity field everywhere on the 3-dimensional space.

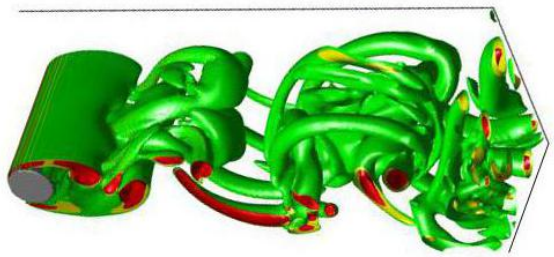

Figure 1: Q-criterion – 13 vortex shedding cycles after the learning period – from the 3D DNS at Reynolds 300.

5. NUMERICAL RESULTS

Our data assimilation procedure has been applied to a 3-dimensional cylinder wake flows at Reynolds number 300. In wake flows, vortices are shed from the cylinder in a pseudo-periodic way. As shown by figure 1, even at this moderate Reynolds, the flow is already three-dimensional and very complex. Before our data assimilation tests, we performed direct numerical simulations (DNS) using Incompact3d, a high-order flow solver based on the discretization of the incompressible Navier-Stokes equations [2]. The simulation's spatial grids correspond to state space dimensions of about 10^7. 80 vortex shedding cycles are used for the ROM constructions (learning period). The remaining vortex shedding cycles are used to test and to build synthetic measurements.

Figures 2, 3, 4 and 5 show the Q-criterion[1] isosurfaces of the data assimilation results, with $n = 2$, 4, 6 and 8 modes respectively . Here, a single spatial resolution point of the synthetic PIV data is assimilated ten times for each of the vortex shedding cycles. The observation point is outside but close to the recirculation zone. Specifically, its coordinates are $x = 1.31D$ (streamwise direction), $z = 0$ (spanwise direction) and $y = -1.48D$ (orthogonal direction), where the cylinder is centered on $(0,0,-)$ and has a diameter D. Despite the little amount of information contained in the measurements, the ROM shows a very good predictive power even outside the learning period. The proposed observers are almost identical to the references, positively demonstrating the huge potential of our approach. Those references (top plots in figures 2, 3, 4 and 5) are the theoretical performance limits: they correspond to the case of exactly known temporal modes $b_i(t)$ in the equation (2). The differences between those theoretical optimums and the DNS reference (see figure 1) is the unresolved velocity v', this difference is restricted to small scales and 3 dimensional effects at 300. In order to better appreciate the potential of our method, figures 2, 3, 4 and 5 display predictions from a benchmark POD-ROMs – POD-Galerkin with optimally fitted

[1] Q-criterion is a classical visual tool of CFD to visualize vortices.

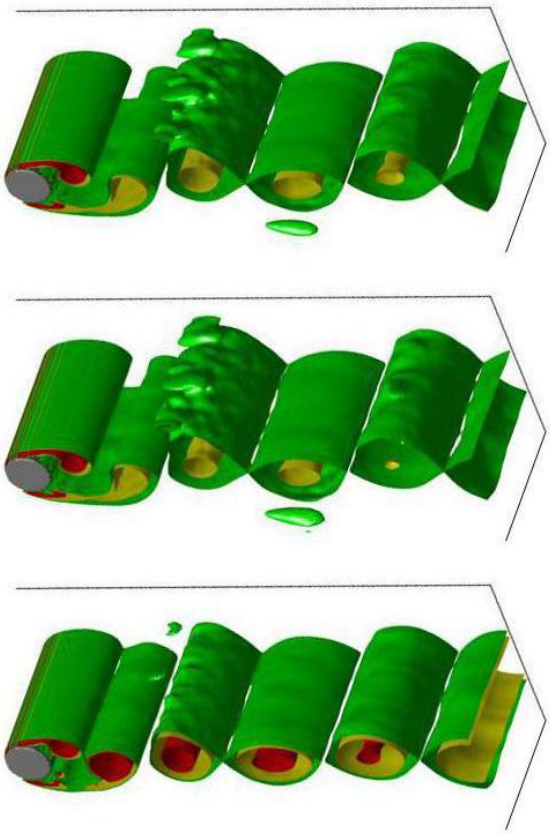

Figure 2: Q-criterion – 13 vortex shedding cycles after the learning period – from the reference 2-dimensional representation (projection of the 3D DNS at Reynolds 300 onto the POD modes) (top), our reduced data assimilation (middle) and a reduced data assimilation benchmark (bottom).

eddy viscosity and additive noise. The same data assimilation algorithm is used there. Nevertheless, the benchmark method strongly differs from the reference whereas our method sticks to it.

Q-criterion iso-surfaces provide a good qualitative analysis. For a quantitative one, figures 6, 7, 8 and 9 plot the global velocity prediction normalized error from the end of the learning period until 14 vortex shedding cycles later. Again, our method performs extremely well whereas the benchmark perform extremely bad.

To generate the simulation dataset, the offline high-resolution highly-optimized CFD code needs to run for several hours on a supercomputer, whereas the off-line ROM construction just takes a few hours on a laptop for the current non-parallelized MATLAB code. Meanwhile, the online ROM should allow for measurements to be assimilated on the fly (i.e. at one vortex shedding cycle every 5 seconds) on a non-parallelized Python implementation, also on a laptop computer. Consequently, a considerable speed-up should be expected for their conversion

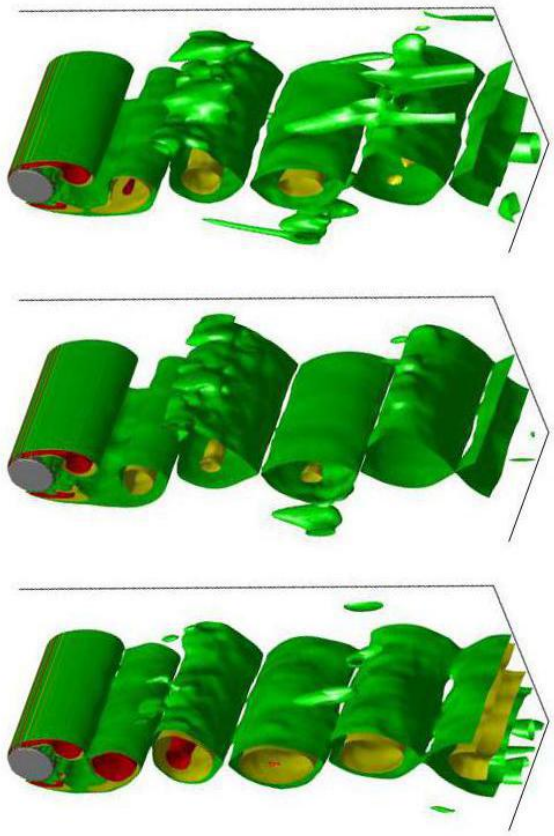

Figure 3: Q-criterion – 13 vortex shedding cycles after the learning period – from the reference 4-dimensional representation (projection of the 3D DNS at Reynolds 300 onto the POD modes) (top), our reduced data assimilation (middle) and a reduced data assimilation benchmark (bottom).

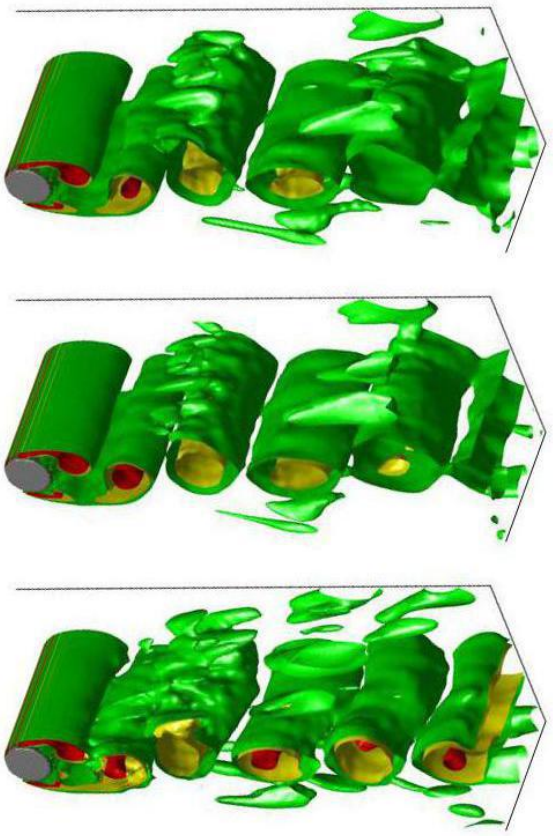

Figure 4: Q-criterion – 13 vortex shedding cycles after the learning period – from the reference 6-dimensional representation (projection of the 3D DNS at Reynolds 300 onto the POD modes) (top), our reduced data assimilation (middle) and a reduced data assimilation benchmark (bottom).

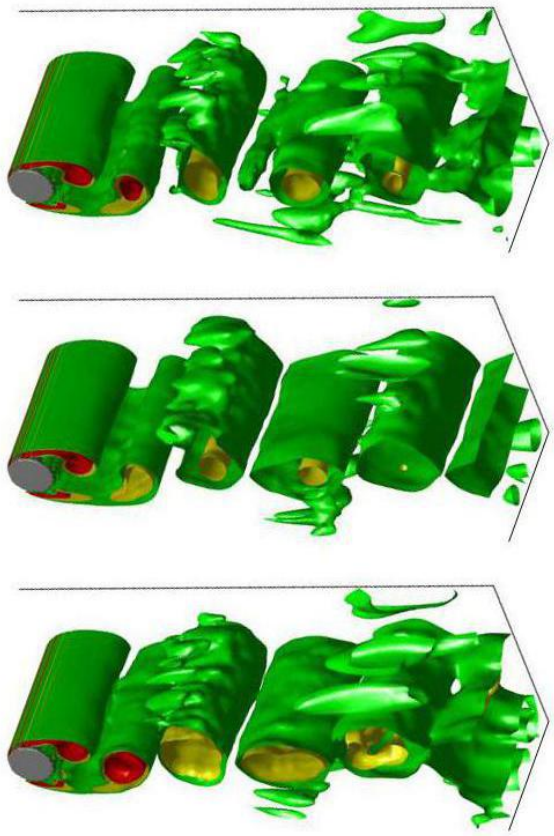

Figure 5: Q-criterion – 13 vortex shedding cycles after the learning period – from the reference 8-dimensional representation (projection of the 3D DNS at Reynolds 300 onto the POD modes) (top), our reduced data assimilation (middle) and a reduced data assimilation benchmark (bottom).

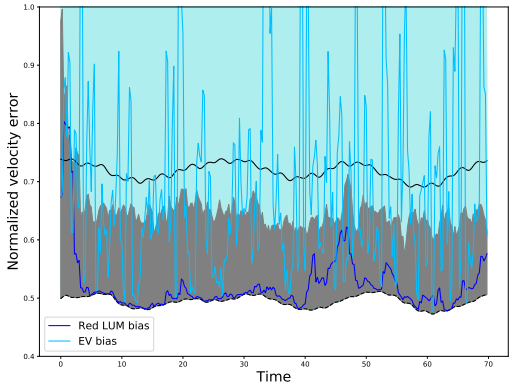

Figure 6: Global velocity prediction normalized error (after the learning period) with $n = 2$ modes for our reduced data assimilation (dark blue) and a reduced data assimilation benchmark (light blue). The shade colors (grey and light blue) correspond to the respective estimated *a posteriori* standard deviations. The dashed black line at the bottom is the POD truncation error. The solid black line at the top is the error obtained by setting all temporal modes b_i to 0, i.e. keeping only the time averaged velocity $\overline{v}$.

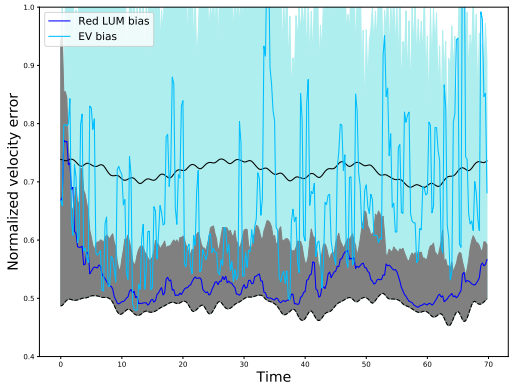

Figure 7: Global velocity prediction normalized error (after the learning period) with $n = 4$ modes for our reduced data assimilation (dark blue) and a reduced data assimilation benchmark (light blue). The shade colors (grey and light blue) correspond to the respective estimated *a posteriori* standard deviations. The dashed black line at the bottom is the POD truncation error. The solid black line at the top is the error obtained by setting all temporal modes b_i to 0, i.e. keeping only the time averaged velocity $\overline{v}$.

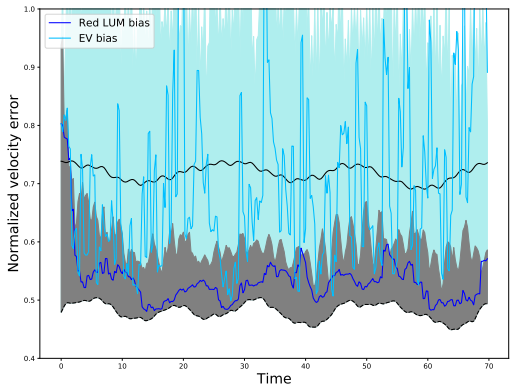

Figure 8: Global velocity prediction normalized error (after the learning period) with $n = 6$ modes for our reduced data assimilation (dark blue) and a reduced data assimilation benchmark (light blue). The shade colors (grey and light blue) correspond to the respective estimated *a posteriori* standard deviations. The dashed black line at the bottom is the POD truncation error. The solid black line at the top is the error obtained by setting all temporal modes b_i to 0, i.e. keeping only the time averaged velocity $\bar{v}$.

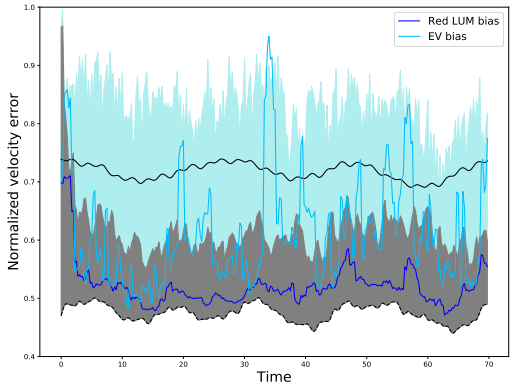

Figure 9: Global velocity prediction normalized error (after the learning period) with $n = 8$ modes for our reduced data assimilation (dark blue) and a reduced data assimilation benchmark (light blue). The shade colors (grey and light blue) correspond to the respective estimated *a posteriori* standard deviations. The dashed black line at the bottom is the POD truncation error. The solid black line at the top is the error obtained by setting all temporal modes b_i to 0, i.e. keeping only the time averaged velocity $\bar{v}$.

to C++ and code parallelization.

6. CONCLUSION

In this paper, we have presented a novel accurate yet efficient method to estimate and predict a flow velocity field globally in space from sparse measurements. Our algorithm is based on a stochastic, low-dimensional system built using the knowledge of physics equations and simulated DNS data, and enabling the coupling with real-time measurements. Comparisons with a benchmark approach have shown a remarkable improvement. With our method, a single measurement point was enough to accurately predict the unsteady velocity field in the whole 3-dimensional space.

Naturally, as a continuation to this work we will perform some tests with actual particle velocimetry (PIV) measurements. Some enhancements to address more turbulent flows are also underway. Finally, we envisage the introduction of parametric ROMs and ROM construction from noisy and/or incomplete data.

REFERENCES

[1] A. Doucet and A. Johansen. A tutorial on particle filtering and smoothing: Fifteen years later. *Handbook of Nonlinear Filtering*, 12:656–704, 2009.

[2] S. Laizet and E. Lamballais. High-order compact schemes for incompressible flows: a simple and efficient method with the quasi-spectral accuracy. *J. Comp. Phys.*, 228(15):5989–6015, 2009.

[3] E. Mémin. Fluid flow dynamics under location uncertainty. *Geophysical & Astrophysical Fluid Dynamics*, 108(2):119–146, 2014.

[4] V. Resseguier, E. Mémin, and B. Chapron. Geophysical flows under location uncertainty, part II: Quasi-geostrophic models and efficient ensemble spreading. Manuscript submitted for publication in Geophysical & Astrophysical Fluid Dynamics, 2017.

[5] Valentin Resseguier. *Mixing and fluid dynamics under location uncertainty*. PhD thesis, 2017.

[6] Jeffrey P Thomas, Earl H Dowell, and Kenneth C Hall. Three-dimensional transonic aeroelasticity using proper orthogonal decomposition-based reduced-order models. *Journal of Aircraft*, 40(3):544–551, 2003.