


**Multidisciplinary
Simulation, Estimation, and Assimilation Systems
Reports in Ocean Science and Engineering**

MSEAS-11

**Numerical Schemes and Computational
Studies for Dynamically
Orthogonal Equations**

by

**Mattheus P. Ueckermann
Pierre F. J. Lermusiaux
Themis P. Sapsis**



**Department of Mechanical Engineering
Massachusetts Institute of Technology**

Cambridge, Massachusetts

August 2011

Numerical Schemes and Computational Studies for Dynamically Orthogonal Equations

M. P. Ueckermann^a, P. F. J. Lermusiaux^a, T. P. Sapsis^a

^a*Department of Mechanical Engineering, Massachusetts Institute of Technology, 77 Mass. Avenue,
Cambridge, MA 02139. Tel.: +1-617-324-5172,
email: mpuecker@mit.edu, email: pierrel@mit.edu, email: sapsis@mit.edu*

Abstract

The quantification of uncertainties is critical when systems are nonlinear and have uncertain terms in their governing equations or are constrained by limited knowledge of initial and boundary conditions. Such situations are common in multiscale, intermittent and non-homogeneous fluid and ocean flows, and other non-linear dynamical systems. The Dynamically Orthogonal (DO) field equations provide an efficient time-dependent adaptive methodology to predict the probability density functions of such dynamics. The present work derives efficient computational schemes for the DO methodology applied to unsteady stochastic Navier-Stokes and Boussinesq equations, and illustrates and studies the numerical aspects of these schemes. Semi-implicit projection methods are developed for the mean and for the orthonormal modes that define a basis for the evolving DO subspace, and time-marching schemes of first to fourth order are used for the stochastic coefficients. Conservative second-order finite-volumes are employed in physical space with new advection schemes based on Total Variation Diminishing methods. Other results specific to the DO equations include: (i) the definition of pseudo-stochastic pressures to obtain a number of pressure equations that is linear in the subspace size instead of quadratic; (ii) symmetric advection schemes for the stochastic velocities; (iii) the use of generalized inversion to deal with singular subspace covariances or deterministic modes; and (iv) schemes to maintain orthonormal modes at the numerical level. While (i) and (ii) are specific to fluid flows, (iii) and (iv) are important for any system of equations discretized using the DO methodology. To verify the correctness of our implementation and study the properties of our schemes and their variations, a set of stochastic flow benchmarks are defined including asymmetric Dirac and symmetric lock-exchange flows, lid-driven cavity flows, and flows past objects in a confined channel. Different Reynolds number and Grashof number regimes are employed to illustrate robustness. Optimal convergence under both time and space refinements is shown as well as the convergence of the probability density functions with the number of stochastic realizations.

Keywords: Uncertainty Quantification, Navier-Stokes, Boussinesq, Dynamical Orthogonality, Projection Methods, Total Variation Diminishing, Error Subspace Statistical Estimation, Ocean Modeling, Data Assimilation

1. Introduction

Quantifying uncertainty is becoming increasingly important in many scientific and engineering applications. This is in part because the accuracy of an answer is now often as critical as the answer itself. Our present motivation is uncertainty prediction for computational fluid dynamics (CFD) applications, specifically in the context of realistic ocean predictions. In ocean dynamics, it is challenging to model multi-scale, intermittent, non-stationary and non-homogeneous uncertainties. Already a single evaluation of an ocean model is costly and straightforward stochastic modeling methods are prohibitively expensive [36, 46], particularly

when dealing with longer-term unsteady nonlinear dynamics. Fortunately, the recently developed Dynamically Orthogonal (DO) field equations [50, 49, 51] provide efficient, tractable equations for uncertainty prediction in large-scale CFD and ocean applications. While these DO equations have been solved numerically using a simple finite-difference scheme, the specific properties of the DO equations warrant novel integration and discretization schemes. Hence, our present goals are to derive efficient computational schemes for the DO methodology applied to unsteady stochastic Navier-Stokes and Boussinesq dynamics, and to illustrate and study the numerical aspects of these schemes. While we are specifically focused on incompressible fluid flows, the use of generalized inversion to deal with singular subspace covariances or deterministic modes, and schemes to maintain orthonormal modes at the numerical level are relevant to any system solved using the DO method.

Stochastic modeling approaches can be categorized as either non-intrusive or intrusive. Non-intrusive approaches have the advantage that the deterministic version of a model can be used to generate an ensemble of sample solutions from which the statistics can be calculated. The disadvantage is that a large number of samples are often required, leading to large computational costs. Intrusive approaches require the update of existing codes or the development of a new code, where the resulting system of equations can be larger than the original, deterministic system. However, intrusive methods are usually computationally less expensive than non-intrusive methods, and the statistics are explicitly available. Below we give a brief review of existing methods, focusing on their computational aspects. For more complete references and reviews we refer to e.g.: Ghanem and Spanos [16], Kloeden and Platen [24], Doucet et al. [9], Mathelin et al. [43], Eldred et al. [10], Jakeman and Roberts [22], Najm [47], Xiu [66, 67], Le Maître and Knio [27].

The non-intrusive Monte-Carlo method provides access to the full statistics of the problem. Its computational cost does not strictly depend on the size of the system, but more on the number of truly independent random variables, and convergence rates are often proportional to the square root of the number of samples. The efficiency can be improved for example by using more elaborate Monte-Carlo schemes [e.g. 9], including particle filters or mixtures of weighted kernels, e.g. Gaussian kernels [5, 45]. Nonetheless, a large number of function evaluations are needed due to the slow (square root) convergence, which can limit accuracy in large-scale applications.

The Polynomial chaos expansion (PCE), pioneered by Ghanem and Spanos [16] and based on the theory by Wiener [64, 2, 65], has become popular because it can represent and propagate large uncertainties through complex models. Both non-intrusive [e.g. 63, 25, 21, 10] and intrusive [e.g. 7, 26, 61, 42, 47] versions have been employed, but both can suffer from the curse of dimensionality. That is, the PCE of a dynamical model scales as $\frac{(p+s)!}{p!s!}$, where p is the largest degree of the polynomials used, and s is the number of independent random variables. For problems that require large p (e.g. non-Gaussian) and large s (e.g. large ocean or fluid simulations), the number of function evaluations (non-intrusive), or the size of the system of equations (intrusive), can quickly become prohibitive. For s larger than p , the storage of a PCE scales as $\mathcal{O}(s^p)$ while its computational cost for Navier-Stokes flows scales as $\mathcal{O}(s^{2p})$ due to the quadratic nonlinearity of advection terms. The large cost of these methods have prompted the use of non-Gaussian random variables [48], the development of generalized PCE [68] to speed up convergence in the polynomial degree (i.e. reduce p), and

the development of adaptive schemes that only evaluate the necessary terms in the PCE (Li and Ghanem [38]).

PCEs have been successful in many CFD applications. In the case of unsteady incompressible fluid dynamics, Le Maître et al. [28] used a PCE scheme to study mixing in a two-dimensional (2D) microchannel and improved the efficiency of their solution scheme by decoupling the velocity-pressure equations using a projection method. Wan and Karniadakis [62] studied the long term behavior of a generalized PCE first using the 1D advection equation and then a 2D noisy flow past a circular cylinder. The authors showed that multi-element generalized PCEs can significantly improve the accuracy for long time integration, but caution that the cost for large s remain high. Other applications include fluid-structure interactions, [e.g. 69], turbulence [e.g. 41], and aerodynamics [e.g. 54]. Other examples are also studied in the above references.

Motivated by the multi-scale, intermittent and non-homogeneous uncertain ocean fields, the Error Subspace Statistical Estimation (ESSE) method was developed [35, 29, 30]. It uses a Karhunen-Loève (KL) expansion [23, 39] but with time varying and adaptive basis functions. This generalized KL expansion is initialized by a multi-scale scheme [31] and evolved using stochastic, data-assimilative and adaptive, Monte-Carlo ensemble schemes. The computational cost of ESSE predictions scales as the size q of the Monte-Carlo ensemble required, i.e. $\mathcal{O}(q)$. This ensemble size q varies with time and is linked to, but a bit larger than, the evolving size of the error subspace itself s (which gives the storage cost scaling in $\mathcal{O}(s)$).

The DO equations for dynamically evolving stochastic fields [50, 49] were derived to approximate the Fokker-Planck equation (or Liouville equation if no stochastic forcing is used) and capture the dominant stochastic subspace while being computationally tractable. The DO methodology also starts from a truncated generalized Karhunen-Loève expansion but derives the governing equations for the mean, the modes and their coefficients. In this derivation, a key condition is imposed: the rate-of-change of the stochastic subspace is dynamically orthogonal to the subspace itself. The DO subspace basis, i.e. the DO modes, as well as the probability density functions (pdfs), i.e. the stochastic coefficients, thus evolve only according to the dynamics of the system. This renders the DO decomposition efficient and limits divergence issues. Consequently, the computational scaling is only dependent on the number of random variables, s , even when dealing with non-Gaussian processes: specifically, the storage scales as $\mathcal{O}(s)$ and computational cost as $\mathcal{O}(s^2)$ for Navier-Stokes equations. The size s is in general a function of time so as to adapt to the dynamically evolving uncertainties and boundary conditions [51]. The DO methodology has been applied to several Navier-Stokes flows and their stochastic dynamics has been studied, including mean-mode and mode-mode energy transfers for 2D flows and heat transfers [49, 52]. However, the DO method has not yet been applied to Boussinesq flows, and the numerical challenges of the DO decomposition for the stochastic Navier-Stokes and Boussinesq equations have not been thoroughly examined nor resolved. This explains the need for the present study.

In what follows, the equations for incompressible stochastic Navier-Stokes and Boussinesq dynamics are given (§2) and their DO decomposition is outlined in Appendix A. In §3, the discretization in time is developed, discussing explicit and implicit schemes. Importantly, we first combine DO pressure terms and define new “pseudo-stochastic pressures” in the original continuous DO equations. With this new definition, given that the pressure Poisson equa-

tions and the corresponding matrix inversions often dominate the computational cost, the cost of DO integration schemes is substantially reduced: instead of scaling as $O(s^2)$ [51], it scales as $O(s)$ (as long as the solution of the pressure dominates the cost of the scheme). The time integration schemes are then derived. For the mean and the modes, we employ projection methods [17], outlining schemes of first and second order. For the stochastic coefficients, we obtain several time-marching schemes of first to fourth order, and briefly discuss their extension to the cases of additive and multiplicative random forcing. The discretizations of the physical space and stochastic subspace are given in §4. For the former, the discretization of diffusion operators is straightforward. That of advection operators requires special attention: since the deterministic DO modes have arbitrary signs, how to apply upwinding based on total variation diminishing properties is a key question we investigate. For the stochastic subspace, a number of possible discretizations are outlined, including the direct Monte-Carlo scheme. In §5, the questions of how to deal with singular covariances and how to maintain orthonormal modes in the presence of numerical round-off and truncation errors are discussed. For the applications in §6, a set of benchmarks are defined and utilized to illustrate the properties of our DO numerics scheme. Specifically, a verification benchmark based on an asymmetric Dirac-stochastic lock-exchange flow is defined and used to test the implementation. A symmetric stochastic lock-exchange is then employed to evaluate the new advection schemes for DO modes. The spatial and temporal convergence is studied with a stochastic lid-driven cavity flow. The discretization of the stochastic coefficients is examined using a flow over a square cylinder in a confined channel. Each flow benchmark is purposely chosen to be different in part to illustrate the robustness of our DO numerics. We also expect that such benchmarks can serve as standard tests for future schemes and implementations. Lastly, conclusions and discussions are in §7.

We note that a condensed version of this report was submitted to the Journal of Computational Physics, and is currently under peer review.

2. Stochastic Dynamically Orthogonal Boussinesq Equations

This section defines the differential equations that we solve and then briefly outlines the stochastic DO methodology. The temporal and spatial discretizations are derived in the subsequent section.

The deterministic components of the partial differential equations (PDEs) that we solve on a domain \mathcal{D} are non-dimensional Boussinesq equations¹, in the same form as in Härtel

¹The dimensional variables, denoted with a hat, have been non-dimensionalized using: $\hat{t} = t\sqrt{\frac{\hat{h}}{\hat{g}'}}$; $\hat{\mathbf{x}} = \mathbf{x}\hat{h}$; $\hat{\mathbf{u}} = \mathbf{u}\sqrt{\hat{g}'\hat{h}}$; $\hat{\rho} = \hat{\rho}_{\min} + \rho(\hat{\rho}_{\max} - \hat{\rho}_{\min})$

et al. [19],

$$\begin{aligned}\nabla \cdot \mathbf{u} &= 0, \quad \mathbf{x} \in \mathcal{D}, \\ \frac{\partial \mathbf{u}}{\partial t} - \frac{1}{\sqrt{Gr}} \nabla^2 \mathbf{u} &= -\nabla \cdot (\mathbf{u}\mathbf{u}) - \nabla p + \rho \mathbf{e}^g, \quad \mathbf{x} \in \mathcal{D}, \\ \frac{\partial \rho}{\partial t} - \frac{1}{Sc\sqrt{Gr}} \nabla^2 \rho &= -\nabla \cdot (\mathbf{u}\rho), \quad \mathbf{x} \in \mathcal{D}.\end{aligned}\tag{1}$$

The non-dimensional variables are: $\mathbf{u}(\mathbf{x}, t) = [u, v, w]$, the velocity in 3D; $\rho(\mathbf{x}, t)$, the density; and, $p(\mathbf{x}, t)$, the pressure. The vector \mathbf{e}^g is a unit-vector in the direction of gravity, (\mathbf{x}, t) are the non-dimensional space and time variables, $Gr = \frac{\hat{g}'\hat{h}^3}{\hat{\nu}^2}$ is the Grashof number which is the ratio of buoyancy forces to viscous forces, $Sc = \hat{\nu}/\hat{K}$ is the Schmidt number which is the ratio of kinematic viscosity $\hat{\nu}$ to molecular diffusivity \hat{K} for the density field, $\hat{g}' = \hat{g} \frac{(\hat{\rho}_{\max} - \hat{\rho}_{\min})}{\hat{\rho}_{\text{avg}}}$ is the reduced gravity, and \hat{h} is the vertical length-scale. In what follows, we denote the total dynamical rate-of-change in the prognostic eqns. (1) for velocity and density by $\mathcal{L}^{\mathbf{u}}$ and \mathcal{L}^{ρ} , respectively, i.e. $\frac{\partial \mathbf{u}}{\partial t} = \mathcal{L}^{\mathbf{u}}$ and $\frac{\partial \rho}{\partial t} = \mathcal{L}^{\rho}$.

The latter prognostic equation for density originates from the thermodynamic energy equation and an equation of state (it arises from another form of the Boussinesq approximation frequently used in ocean modeling which retains the temperature and salinity fields as state variables, e.g. [6, 18]). We emphasize that for problems without density-driven flows, $\sqrt{Gr} \equiv Re$, that is, the square root of the Grashof number is the Reynolds number. The approach and numerical schemes that we derive in this manuscript are directly applicable to the Navier-Stokes equations.

We are interested in solving eqns. (1) in their stochastic form. We thus introduce the set of random events ω belonging to a measurable sample space Ω and consider the stochastic velocity, density and pressure fields: $\mathbf{u}(\mathbf{x}, t; \omega)$; $\rho(\mathbf{x}, t; \omega)$ and $p(\mathbf{x}, t; \omega)$. This leads to stochastic dynamical rates-of-change $\mathcal{L}^{\mathbf{u}}$ and \mathcal{L}^{ρ} . If these rate-of-changes are themselves uncertain, for example due to parameter or model uncertainties, then they also depend explicitly on ω . In this study, we mostly focus on uncertainties arising due to uncertain initial conditions. We define general stochastic initial conditions as

$$\begin{aligned}\mathbf{u}(\mathbf{x}, 0; \omega) &= \mathbf{u}_0(\mathbf{x}; \omega), \quad \mathbf{x} \in \mathcal{D}, \quad \omega \in \Omega, \\ \rho(\mathbf{x}, 0; \omega) &= \rho_0(\mathbf{x}; \omega), \quad \mathbf{x} \in \mathcal{D}, \quad \omega \in \Omega,\end{aligned}\tag{2}$$

and stochastic boundary conditions as

$$\begin{aligned}\mathbf{u} &= \mathbf{g}_D(\mathbf{x}, t; \omega), \quad \mathbf{x} \in \partial\mathcal{D}_D, \quad \omega \in \Omega, \\ \frac{\partial \mathbf{u}}{\partial n} &= \mathbf{g}_N(\mathbf{x}, t; \omega), \quad \mathbf{x} \in \partial\mathcal{D}_N, \quad \omega \in \Omega, \\ \rho &= g_{D_\rho}(\mathbf{x}, t; \omega), \quad \mathbf{x} \in \partial\mathcal{D}_{D_\rho}, \quad \omega \in \Omega, \\ \frac{\partial \rho}{\partial n} &= g_{N_\rho}(\mathbf{x}, t; \omega), \quad \mathbf{x} \in \partial\mathcal{D}_{N_\rho}, \quad \omega \in \Omega,\end{aligned}\tag{3}$$

where the boundary conditions are separated into Dirichlet and Neumann conditions for the

velocity and density fields (pressure boundary conditions are considered later). The resulting multivariate stochastic eqns. (1)-(3) define the problem to be solved. As in the deterministic case, specifics of the solution depend on the initial and boundary conditions chosen.

The DO decomposition of these equations can be obtained from [49, 50] and a summary is provided in Appendix A. In short, the DO methodology begins with a generalized Karhunen-Loève expansion truncated to $s(t)$ terms [51]. The vector of prognostic state variables $\Phi(\mathbf{x}, t; \omega) = [\mathbf{u}, \rho]^T$ is decomposed into the sum of a deterministic mean component $\bar{\Phi}(\mathbf{x}, t)$, with s deterministic modes $\Phi_i(\mathbf{x}, t)$, each mode multiplied by a stochastic coefficient $Y_i(t; \omega)$. This decomposition is first substituted into eqns. (1)-(3). The DO condition, the rate-of-change of the stochastic subspace is dynamically orthogonal to the subspace itself, is then utilized. Orthogonality is defined by the spatial inner-product $\langle \mathbf{a}, \mathbf{b} \rangle_{\mathcal{D}} = \int_{\mathcal{D}} \sum_i (a^i b^i) d\mathcal{D}$ for arbitrary vectors of spatial functions $\mathbf{a} = [a^1, a^2, \dots]^T$ and $\mathbf{b} = [b^1, b^2, \dots]^T$. In general, we note that this definition of the inner product assumes that the different components of the state vector have been properly normalized [33, 52]. This is not guaranteed from the simple deterministic non-dimensionalization used in eqn. (1). In fact, an additional stochastic normalization is usually needed, reflecting the stochastic initial and boundary conditions. After some manipulation (see Appendix A and [50, 49]) time-evolution equations for the mean, modes, and stochastic coefficients, which are completely determined by the dynamics, are obtained. A major contribution of this manuscript is to derive efficient discretizations in time and space for these equations and to evaluate the resulting computational schemes through a set of new benchmarks for stochastic Boussinesq dynamics.

3. Semi-implicit Time Discretization

Solving the deterministic version of the system of equations (1) implicitly in time often requires not only a large matrix inversion at each time-step, but also iterations at each time-step to deal with the non-linear advection terms, e.g. [13]. Discretizing their stochastic version (1)-(3) using a brute-force Monte-Carlo scheme would have similar costs per realizations, hence a total cost equal to that of the deterministic version but multiplied by the size of the ensemble. If a DO decomposition is used, solving the DO system (A.5)-(A.12) implicitly would require a matrix inversion $(s^2 + s + 1)$ times larger than for (1) since the mean and the modes are coupled through the pressure and non-linear advection terms (the number of pressure equations are: s^2 for p_{ij} 's, s for p_i 's and 1 for \bar{p}). While it is possible to solve such systems, our goal here is to discretize (A.5)-(A.12) such that the equations decouple, resulting in an efficient solution scheme. This section describes how this decoupling is achieved. First we explain why we treat some terms explicitly and others implicitly. We then define new pseudo-stochastic pressures that substantially reduce computational costs, develop Projection methods for DO equations so as to split the velocity and pressure terms, and present time marching schemes for the stochastic coefficients. The complete time discretizations, combining these time marching schemes with the projection schemes, are summarized at the end. The spatial discretizations of physical space and of the stochastic subspace are given in Sect. §4.

3.1. Explicitly and Implicitly Treated Terms

Much of the decoupling is achieved by treating some terms explicitly, resulting in a semi-implicit scheme. First, we choose to advance the stochastic coefficients explicitly, because then $\mathbf{C}_{Y_i Y_j}$ and $\mathbf{M}_{Y_j Y_m Y_n}$ can be treated as constants when evolving the mean and the modes, and no iteration is required to solve (A.5), (A.8) and (A.11). Somewhat similarly, we treat the inner product terms $\langle \mathbf{Q}_i, \mathbf{\Phi}_j \rangle_{\mathcal{D}} \mathbf{\Phi}_j$ in (A.8) explicitly to avoid iterations. Next, we treat the non-linear advection terms explicitly, which is often done in the Projection method community (e.g. [17]). This does impose a stability constraint on the time-step size, a Courant - Friedrichs - Lewy (CFL) condition. Third, we treat the linear diffusion terms implicitly because they do not couple the equations and the resulting diagonal-dominant matrices can be inverted efficiently. While they could also be treated explicitly, this imposes a much harsher stability constraint that could result in very small timesteps. Thus, to partially decouple the evolution equations, we advance the stochastic coefficients, inner product terms and non-linear advection explicitly. However, these equations are still coupled through the pressure.

3.2. The Pseudo-Stochastic Pressures

In this section we first review the explicit treatment of pressure which results in a scheme requiring $s^2 + s + 1$ solutions of stochastic Pressure Poisson equations (PPEs) per timestep. Then, we discuss our new definition of pseudo-stochastic pressures that reduces the expense to $s + 1$ and show that it is a valid definition that does not change the order of accuracy.

One approach for handling the pressure, which was adopted in Sapsis and Lermusiaux [50], is to treat it explicitly. This approach takes advantage of the fact that the full stochastic pressure can be recovered at any time instant by taking the divergence of (A.5) and (A.8), inserting the decomposition (A.2), and using the divergence-free form on continuity, noting that $\nabla \cdot \frac{\partial \mathbf{u}}{\partial t} = 0$. The result gives the stochastic Pressure Poisson equations (PPEs) for our system

$$\begin{aligned}\nabla^2 \bar{p} &= -\nabla \cdot [\nabla \cdot (\bar{\mathbf{u}}\bar{\mathbf{u}}) - \bar{\rho}\mathbf{e}^g], \\ \nabla^2 p_i &= -\nabla \cdot [\nabla \cdot (\bar{\mathbf{u}}\mathbf{u}_i) + \nabla \cdot (\mathbf{u}_i\bar{\mathbf{u}}) - \rho_i\mathbf{e}^g], \\ \nabla^2 p_{ij} &= -\nabla \cdot [\nabla \cdot (\mathbf{u}_j\mathbf{u}_i) + \nabla \cdot (\mathbf{u}_i\mathbf{u}_j)].\end{aligned}\tag{4}$$

A disadvantage of this explicit pressure approach is that it is expensive; to recover the full stochastic pressure, $1 + s + s^2$ Poisson equations need to be inverted, and this would often be the dominating cost of the scheme. Oceanic applications are expected to require $s \sim \mathcal{O}(10^2 - 10^3)$ [33], which would be very expensive. Another disadvantage is that the velocity computed with an explicit scheme will not be divergence-free after each timestep.

We can reduce the number of PPEs to $s + 1$ by defining new pseudo-stochastic pressures. The purpose of the pressure in divergence-free flows is to enforce continuity. In our stochastic equations, continuity needs to be satisfied by the mean and each modal velocity field independently (we assume that the divergence-free continuity equation is exact, without any errors in its form). Also, each of these velocity fields only needs a single scalar field in order to satisfy the continuity constraint. By inspection of equations (A.5) and (A.8), we therefore define new pseudo-stochastic pressures, which are a combination of the mean, linear-, and

quadratic-modal pressures:

$$\begin{aligned}\check{\check{p}} &= \bar{p} + \mathbf{C}_{Y_i Y_j} p_{ij}, \\ \check{p}_i &= p_i + \mathbf{C}_{Y_i Y_j}^{-1} \mathbf{M}_{Y_j Y_m Y_n} p_{mm}.\end{aligned}\tag{5}$$

With this definition, the quadratic modal pressures are eliminated from (A.5) and (A.8). Thus, to evolve the mean and modes, we no longer need to solve for the quadratic pressures. However, substituting (5) into the equation for the evolution of the stochastic coefficients (A.11), we find that the second term on the right-hand-side of (A.11),

$$\langle \nabla p_{mm} + \nabla \cdot (\mathbf{u}_n \mathbf{u}_m), \mathbf{u}_i \rangle_{\mathcal{D}} (Y_m Y_n - \mathbf{C}_{Y_m Y_n})$$

retains the projection of the quadratic stochastic pressure terms in the subspace. At first, this would indicate that the quadratic modal pressures are still needed, but for commonly used boundary conditions, the projection cancels, i.e. the inner product $\langle \nabla p_{mm}, \mathbf{u}_i \rangle_{\mathcal{D}}$ is zero (see [52]). The quadratic stochastic pressure term in (A.11) can be dropped without any penalty. Thus, by defining new pseudo-stochastic pressures (5), we have shown that we reduced the number of PPEs from $s^2 + s + 1$ to the expected $s + 1$.

3.3. Projection Methods for the Mean and Modes

To obtain a numerically divergence-free velocity, we use a Projection method. A large number of different Projection methods exist; for a recent review, see [17]. Projection methods are known for excellent efficiency, but the proper specification of boundary conditions remains a long-standing issue. While advances in this area are still being made [15, 53], we have chosen to use the “incremental pressure-correction scheme in rotational form” proposed by Timmermans et al. [58], which has a proven temporal accuracy [17]. We first summarize the classic versions of the scheme, then adapt them for the mean and modes.

Classic Projection Scheme. Intermediate velocities are first solved for using a first or second order time-accurate discretization. In both cases, the contributions from known velocities (at previous or intermediate timesteps) are written as $\mathbf{F}(\mathbf{u}^{k^*}, \rho^{k^*})$, where the time instant t^{k^*} determines the order of the scheme in time. Specifically, for first order in time, $k^* = k - 1$, and

$$\begin{aligned}\frac{\tilde{\mathbf{u}}^k}{\Delta t} - \frac{1}{\sqrt{Gr}} \nabla^2 \tilde{\mathbf{u}}^k &= -\nabla p^{k-1} + \mathbf{F}(\mathbf{u}^{k-1}, \rho^{k-1}), \quad \mathbf{x} \in \mathcal{D}, \\ \tilde{\mathbf{u}}^k &= \mathbf{g}_D, \quad \mathbf{x} \in \partial \mathcal{D}_D, \\ \frac{\partial \tilde{\mathbf{u}}^k}{\partial n} &= \mathbf{g}_N, \quad \mathbf{x} \in \partial \mathcal{D}_N.\end{aligned}$$

This is followed by the computation of the pressure correction, θ , so as to satisfy continuity,

$$\begin{aligned}\nabla^2 \theta &= \frac{\nabla \cdot \tilde{\mathbf{u}}^k}{\Delta t}, \\ \frac{\partial \theta}{\partial n} &= 0, \quad \mathbf{x} \in \partial \mathcal{D}_D, \\ \theta &= g_P, \quad \mathbf{x} \in \partial \mathcal{D}_N,\end{aligned}$$

where g_P is the prescribed pressure difference at the boundary. With this correction, the full pressure and velocity fields can be recovered at the next timestep, i.e.

$$\begin{aligned}\mathbf{u}^k &= \tilde{\mathbf{u}}^k - \nabla \theta^k, \\ p^k &= p^{k-1} + \theta^k - \nu \nabla \cdot \tilde{\mathbf{u}}^k.\end{aligned}$$

The first-order time-accurate discretization, $k^* = k-1$, uses the contributions from known velocities (at old timesteps) as $\mathbf{F}(\mathbf{u}^{k-1}, \rho^{k-1})$. The second-order time integration scheme is obtained by appropriately choosing the guess values $(\bullet)^{k^*}$ with $t_{k-1} \leq t_{k^*} \leq t_k$, and using a higher order backwards differencing method. For additional properties of such schemes, we refer to Timmermans et al. [58], Guermond et al. [17].

Projection Scheme for the Mean. We evolve the mean fields modifying the classic projection method to account for the moments of the stochastic coefficients and of the chosen explicit and implicit terms (Sect. §3.1). Starting from the PDEs (A.5) for the mean, we obtain:

$$\begin{aligned}\frac{\tilde{\mathbf{u}}^k}{\Delta t} - \frac{1}{\sqrt{Gr}} \nabla^2 \tilde{\mathbf{u}}^k &= \frac{\bar{\mathbf{u}}^{k-1}}{\Delta t} - \{\nabla \cdot (\bar{\mathbf{u}}\bar{\mathbf{u}})\}^{k^*} - \nabla \check{p}^{k^*} + \check{\rho}^{k^*} \mathbf{e}^g \\ &\quad - \mathbf{C}_{Y_i Y_j}^{k^*} \{\nabla \cdot (\mathbf{u}_j \mathbf{u}_i)\}^{k^*},\end{aligned}\tag{6a}$$

$$\nabla^2 \bar{\theta}^k = \frac{1}{\Delta t} \nabla \cdot \tilde{\mathbf{u}}^k,\tag{6b}$$

$$\bar{\mathbf{u}}^k = \tilde{\mathbf{u}}^k - \Delta t \nabla \bar{\theta}^k,\tag{6c}$$

$$\check{p}^k = \check{p}^{k-1} + \bar{\theta}^k - \nu \nabla \cdot \tilde{\mathbf{u}}^k,\tag{6d}$$

$$\frac{\bar{\rho}^k}{\Delta t} - \frac{1}{Sc\sqrt{Gr}} \nabla^2 \bar{\rho}^k = \frac{\bar{\rho}^{k-1}}{\Delta t} - \{\nabla \cdot (\bar{\mathbf{u}}\bar{\rho})\}^{k^*} - \mathbf{C}_{Y_i Y_j}^{k^*} \{\nabla \cdot (\mathbf{u}_j \rho_i)\}^{k^*},\tag{6e}$$

with deterministic boundary conditions:

$$\begin{aligned}\tilde{\mathbf{u}}^k &= \bar{\mathbf{g}}_D, \quad \frac{\partial \bar{\theta}}{\partial n} = 0, \quad \mathbf{x} \in \partial \mathcal{D}_D, \\ \frac{\partial \tilde{\mathbf{u}}^k}{\partial n} &= \bar{\mathbf{g}}_N, \quad \bar{\theta} = \bar{g}_P, \quad \mathbf{x} \in \partial \mathcal{D}_N, \\ \bar{\rho}^k &= \bar{\mathbf{g}}_{D_\rho}, \quad \mathbf{x} \in \partial \mathcal{D}_{D_\rho}, \\ \frac{\partial \bar{\rho}^k}{\partial n} &= \bar{\mathbf{g}}_{N_\rho}, \quad \mathbf{x} \in \partial \mathcal{D}_{N_\rho}.\end{aligned}\tag{7}$$

In the above, the time instant t^{k*} , either previous or intermediate, determines the order of the scheme in time. Specifically, for first order in time, $k^* = k - 1$; for second order, we refer to [58, 17]. We note that a difference between a classic projection scheme and the above DO mean scheme is the presence of the covariances and third moments of the coefficients Y_i 's (see Sect. §3.4). A related one is the coupling between the differential equations for the mean, mode and coefficients (see Sect. §3.5).

Projection Scheme for the Modes. As for the mean, the modes are evolved by modifying the classic projection method for the eqns. (A.8). We obtain:

$$\begin{aligned} \frac{\tilde{\mathbf{u}}_i^k}{\Delta t} - \frac{1}{\sqrt{Gr}} \nabla^2 \tilde{\mathbf{u}}_i^k &= \frac{\mathbf{u}_i^{k-1}}{\Delta t} - \{\nabla \cdot (\mathbf{u}_i \bar{\mathbf{u}})\}^{k*} - \{\nabla \cdot (\bar{\mathbf{u}} \mathbf{u}_i)\}^{k*} - \nabla \tilde{p}_i^{k*} + \rho_i^{k*} \mathbf{e}^g \\ &\quad - \mathbf{C}_{Y_i Y_j}^{-1, k*} \mathbf{M}_{Y_j Y_m Y_n}^{k*} \{\nabla \cdot (\mathbf{u}_n \mathbf{u}_m)\}^{k*} \\ &\quad - \langle \mathbf{Q}_i, \mathbf{\Phi}_j \rangle_{\mathcal{D}}^{k*} \mathbf{u}_j^{k*}, \end{aligned} \quad (8a)$$

$$\nabla^2 \theta_i^k = \frac{1}{\Delta t} \nabla \cdot \tilde{\mathbf{u}}_i^k, \quad (8b)$$

$$\mathbf{u}_i^k = \tilde{\mathbf{u}}_i^k - \Delta t \nabla \theta_i^k, \quad (8c)$$

$$\tilde{p}_i^k = \tilde{p}_i^{k-1} + \theta_i^k - \nu \nabla \cdot \tilde{\mathbf{u}}_i^k, \quad (8d)$$

$$\begin{aligned} \frac{\rho_i^k}{\Delta t} - \frac{1}{Sc\sqrt{Gr}} \nabla^2 \rho_i^k &= \frac{\rho_i^{k-1}}{\Delta t} - \{\nabla \cdot (\mathbf{u}_i \bar{\rho})\}^{k*} - \{\nabla \cdot (\bar{\rho} \mathbf{u}_i)\}^{k*} \\ &\quad - \mathbf{C}_{Y_i Y_j}^{-1, k*} \mathbf{M}_{Y_j Y_m Y_n}^{k*} \{\nabla \cdot (\mathbf{u}_n \rho_m)\}^{k*} \\ &\quad - \langle \mathbf{Q}_i, \mathbf{\Phi}_j \rangle_{\mathcal{D}}^{k*} \rho_j^{k*}. \end{aligned} \quad (8e)$$

with boundary conditions:

$$\begin{aligned} \tilde{\mathbf{u}}_i^k &= \mathbf{g}_{i,D}(=0), \quad \frac{\partial \theta_i}{\partial n} = 0, \quad \mathbf{x} \in \partial \mathcal{D}_D, \\ \frac{\partial \tilde{\mathbf{u}}_i^k}{\partial n} &= \mathbf{g}_{i,N}(=0), \quad \theta_i = g_{i,P}(=0), \quad \mathbf{x} \in \partial \mathcal{D}_N, \\ \rho_i^k &= \mathbf{g}_{i,D_\rho}(=0), \quad \mathbf{x} \in \partial \mathcal{D}_{D_\rho}, \\ \frac{\partial \rho_i^k}{\partial n} &= \mathbf{g}_{i,N_\rho}(=0), \quad \mathbf{x} \in \partial \mathcal{D}_{N_\rho}, \end{aligned} \quad (9)$$

where, again, the scheme is first order for $k^* = k - 1$. Since we focus on the numerics of the DO differential equations, we assume that the stochastic boundary forcings are null, i.e. g_i 's are null in eqns. (9). In the present study, we do not exemplify problems with stochastic forcing, but we touch on the numerical modifications required in such cases (see §3.4 and [32] for an example of stochastic ocean models). The time evolution of the Y_i 's is discussed next.

3.4. Time Integration Scheme for the Stochastic Coefficients

To integrate the Ordinary Differential Equations (ODEs) (A.11) for the stochastic coefficients, we assume that all variables are available at time t_{k-1} and that we integrate forward to time t_k .

First, we consider the case where the original governing differential equations only contain uncertain initial conditions, and no stochastic forcing. This case corresponds to the examples and benchmarks we consider later. It allows the use of classic time-marching integration algorithms or, in other words, appropriate ODE solvers. We consider a given realization of a coefficient Y_i at time t_{k-1} . To integrate to time t_k , we first define an approximation to the time-rate of change of this coefficient at a given instant $t_{k-1} \leq t \leq t_k$ as follows:

$$\begin{aligned} \left. \frac{dY_i}{dt} \right|_{Y(t)} &= \left\langle \frac{1}{\sqrt{Gr}} \nabla^2 \mathbf{u}_m - \nabla \cdot (\mathbf{u}_m \bar{\mathbf{u}}) - \nabla \cdot (\bar{\mathbf{u}} \mathbf{u}_m) - \nabla \check{p}_m + \rho_m \mathbf{e}^g, \mathbf{u}_i \right\rangle_{\mathcal{D}}^{k-} Y_m(t) \\ &+ \left\langle \frac{1}{Sc\sqrt{Gr}} \nabla^2 \rho_m - \nabla \cdot (\mathbf{u}_m \bar{\rho}) - \nabla \cdot (\bar{\rho} \mathbf{u}_m), \rho_i \right\rangle_{\mathcal{D}}^{k-} Y_m(t) \\ &- \langle \nabla \cdot (\mathbf{u}_n \mathbf{u}_m), \mathbf{u}_i \rangle_{\mathcal{D}}^{k-} \left(Y_m(t) Y_n(t) - \mathbf{C}_{Y_m Y_n}^{k-} \right) \\ &- \langle \nabla \cdot (\mathbf{u}_n \rho u_m), \rho_i \rangle_{\mathcal{D}}^{k-} \left(Y_m(t) Y_n(t) - \mathbf{C}_{Y_m Y_n}^{k-} \right), \end{aligned} \quad (10)$$

where $(.)^{k-}$ indicates that the modal quantities are estimates of their values at time t . The numerically exact option is to choose $t_{k-} = t$, while the cheapest is to take $t_{k-} = t_{k-1}$ since at that time, all modal quantities are available from the previous time-step. Using this derivative definition, we have employed several time-marching schemes of varying order to advance the Y_i in time, including: a low-storage 4th-order-accurate explicit Runge-Kutta integrator of the form

$$\begin{aligned} Y^{(0)} &= Y(t), \quad y^{(0)} = 0, \\ y^{(m)} &= a_m y^{(m-1)} + \Delta t \left. \frac{dY}{dt} \right|_{Y^{(m-1)}, t+c_m \Delta t} \quad \text{for } m \in [1 \dots 5], \\ Y^{(m)} &= Y^{(m-1)} + b_m y^{(m)} \quad \text{for } m \in [1 \dots 5], \\ Y(t + \Delta t) &= Y^{(5)}, \end{aligned}$$

where the coefficients a_m, b_m, c_m are given in Carpenter and Kennedy [3]; a 2nd-accurate explicit Runge-Kutta scheme (Heun's version)

$$\begin{aligned} Y^{(0)} &= Y(t) + \Delta t \left. \frac{dY}{dt} \right|_{Y(t)}, \\ Y(t + \Delta t) &= Y(t) + \frac{\Delta t}{2} \left\{ \left. \frac{dY}{dt} \right|_{Y(t)} + \left. \frac{dY}{dt} \right|_{Y^{(0)}} \right\}, \end{aligned}$$

and the first-order-accurate explicit Euler scheme

$$Y(t + \Delta t) = Y(t) + \Delta t \left. \frac{dY}{dt} \right|_{Y^t}.$$

For the Euler scheme, $k^- = k - 1$. For each stage of the two RK schemes, $\frac{dY}{dt}$ is evaluated

using (10). If $t_{k^-} = t$, the modal quantities (the rates for the Y_i 's) are advanced to intermediate times using the modes and mean PDEs above (which is expensive). If $k^- = k - 1$, the mean and modal quantities are not advanced, but the $Y_m(t)$'s are still updated at intermediate times. Of course, the formal order of accuracy of that scheme is limited by these mean/modal terms kept constant as will be shown later (§6). While $\mathbf{C}_{Y_m Y_n}^{k^-}$ could be recalculated at each time level, our simulations showed better results if they were kept at the same time k^- as modal quantities (in that case, coefficients and subspace remain consistent).

Second, we consider the case where the original governing equations contain stochastic forcing in the form of zero-mean Wiener processes, added linearly to the deterministic form of the PDEs (1). These stochastic forcings then appear in the governing equation (A.11) for the stochastic coefficients. Without entering details, the above marching schemes are then augmented with stochastic integrals over $t_k - t_{k-1}$ for the contribution of these forcings [24, 20]. If the stochastic forcings are of multiplicative form in the original governing equations, then expectations and coupled integrals of stochastic terms are in general also in the mean and mode equations of Sect. §3.3. This is not considered in the present study.

3.5. Complete Time Integration Scheme

We now summarize the complete time-discretization scheme from t_{k-1} to t_k . Since we have decoupled the equations (6), (8), and (10), the order in which they are solved is not important. In fact, (6), (8), and (10) could be solved in parallel. Presently, we employ the following, serial approach:

1. Calculate/extrapolate the statistics ($\mathbf{C}_{Y_m Y_n}, \mathbf{M}_{Y_j Y_m Y_n}$) to the approximated times k^*, k^- and store for later use
2. Calculate/extrapolate the advection terms ($\nabla \bar{\mathbf{u}} \bar{\mathbf{u}}, \nabla \bar{\mathbf{u}} \mathbf{u}_i + \nabla \mathbf{u}_i \bar{\mathbf{u}}, \nabla \mathbf{u}_i \mathbf{u}_j$) to the approximated times k^*, k^- and store for later use
3. Advance the Y_i 's using (10), and one of the ODE solvers in §3.4
4. Advance the mean $\bar{\mathbf{u}}$ using (6)
5. Advance the modes \mathbf{u}_i using (8).

For the modes and mean, we choose $k^* = k^- = k - 1$, resulting in a first order accurate scheme for the present applications (§6). For second order accuracy, we would use k^* as defined in Timmermans et al. [58], Guermond et al. [17]. For the stochastic coefficients, a higher order ODE solver may be used in step 3 which may reduce the magnitude of the integration error. However, since DO equations are coupled, the order of accuracy of that step is influenced by the choice of k^* and k^- used for the time-integration of the modes and mean. Hence, if $k^* = k^- = k - 1$, the overall accuracy of the time-integration is in general expected to be first order.

4. Spatial Physical and Stochastic Subspace Discretizations

In this section, we start by describing the 2D spatial discretization of (6) and (8) on a structured grid (§4.1). We employ a standard conservative finite volume discretization of second-order. Special treatment is needed for the advection by the modes since they are basis vectors in the stochastic subspace and thus do not have a preferential direction. We finally discuss the discretization of the s -dimensional probability subspace in §4.2.

4.1. Spatial Discretization of the Physical Space

The domain \mathcal{D} is discretized into a finite number of non-overlapping control volumes. Presently, we use rectangular control volumes, which form a structured Cartesian grid that has uniform spacing in the x and y directions. Several choices exist for the relative placement of velocity and pressure control volumes. Here we choose to use a standard staggered C-grid [13, 40], where the u - and v -velocity control volumes are displaced half a grid-cell in the x - and y -directions relative to the pressure and density control volumes, respectively.

4.1.1. Diffusion Operator

The diffusion operator ∇^2 is discretized by using central boundary fluxes. That is, for the field ϕ at the center (x, y) of a control volume boundary, we have the following boundary flux in the x -direction (similar for y -direction)

$$\left. \frac{\partial \phi}{\partial x} \right|_{(x,y)} = \frac{\phi(x + \frac{\Delta x}{2}, y) - \phi(x - \frac{\Delta x}{2}, y)}{\Delta x},$$

which, on a structured grid, is exactly equivalent to the second-order accurate central finite-difference scheme. For the advection operator, the simple second-order central flux is well-known to be unstable (e.g. [4]), and needs more careful treatment, as described next.

4.1.2. Advection Operator

For advection by a velocity component u , we use a standard Total Variation Diminishing (TVD) scheme, with an monotonized central (MC) symmetric flux limiter [60]. The scheme can be written for a variable η as:

$$F(\eta_{i-\frac{1}{2}}) = u_{i-\frac{1}{2}} \frac{\eta_i + \eta_{i-1}}{2} - \left| u_{i-\frac{1}{2}} \right| \frac{\eta_i - \eta_{i-1}}{2} \left[1 - \left(1 - \left| u_{i-\frac{1}{2}} \frac{\Delta t}{\Delta x} \right| \right) \Psi(r_{i-\frac{1}{2}}) \right], \quad (11)$$

where the MC slope limiter $\Psi(r)$ is defined as

$$\Psi(r) = \max \left\{ 0, \min \left[\min \left(\frac{1+r}{2}, 2 \right), 2r \right] \right\},$$

and the variable r as

$$r_{i-\frac{1}{2}} = \frac{\left[\frac{1}{2} \left(u_{i-\frac{1}{2}} + \left| u_{i-\frac{1}{2}} \right| \right) (\eta_{i-1} + \eta_{i-2}) + \frac{1}{2} \left(u_{i-\frac{1}{2}} - \left| u_{i-\frac{1}{2}} \right| \right) (\eta_{i+1} + \eta_i) \right]}{u_{i-\frac{1}{2}} (\eta_i - \eta_{i-1})},$$

where $u_{i-\frac{1}{2}}$ is used without interpolation for the density advection while a second-order central scheme is used for the non-linear u and v advection. For more on TVD schemes, we refer to the excellent text by [37].

A possible issue with using this scheme for DO equations arises from the realization that the absolute value of the velocity $|\mathbf{u}|$, is a function of the full velocity $\bar{\mathbf{u}} + Y_i u_i$ which, depending on the specific realization, may be either positive or negative; in other words, $|\mathbf{u}|$

is positive, but stochastic. Fortunately, we never need the full velocity to evolve the mean and modes in (A.5) and (A.8) (see also (6) and (8)). In fact, in the case of the mean velocity $\bar{\mathbf{u}}$, its absolute value is deterministic. Therefore, the advection of the mean velocity by the mean velocity $\bar{\mathbf{u}} \cdot \nabla \bar{\mathbf{u}}$, and the advection of the velocity modes by the mean velocity $\bar{\mathbf{u}} \cdot \nabla \mathbf{u}_i$ can use the classic TVD method without modification. Advection of the mean by the modes $\mathbf{u}_i \cdot \nabla \bar{\mathbf{u}}$ and of the modes by the modes $\mathbf{u}_i \cdot \nabla \mathbf{u}_j$, however, need additional consideration. Similar statements apply for the advection of the mean density and of the density modes by either the mean velocity (classic scheme is fine) or by the modes (additional considerations are needed).

Here we propose three arguments for three different advection schemes that can be used for these “advection by the modes” terms. First, if we examine the equations from the perspective of the numerical scheme only, a preferential advection direction will be present. In this case we simply use the TVD scheme unmodified. Next, we argue that, since the stochastic coefficients are zero mean, then the probability of $Y_i < 0$ is equal to the probability that $Y_i > 0$. This suggests that \mathbf{u}_i should not have a preferential direction of propagation, in which case a central differencing advection scheme (CDS) could be used. Last, recognizing that the CDS scheme may cause oscillations, we still wish to limit the flux in some way. A direct approach, then, is to use the TVD scheme in both directions, and average the results. That is, the present sign of the modal velocity is used first to calculate the advective terms, then the negative of the modal velocity is used, and the two results are averaged. We call this the symmetric TVD or TVD* scheme. Note that the TVD* scheme is not a true TVD scheme and is thus not guaranteed to be oscillation-free for all realizations.

To better understand the behavior of the TVD* scheme, consider three scenarios: when calculating the flux using the negative and positive modal velocities i) neither requires limiting, ii) both require full limiting, iii) only one requires full limiting. For i), both use a CDS flux and the average of the two is a CDS flux. For ii), both use their respective UW fluxes, which, when averaged together, again gives a CDS flux. This may appear undesirable but it is sensible since the velocity should be fully up-winded from both sides with equal probability, resulting in no preferential direction of up-winding, which gives CDS. For case iii) we depart from a CDS scheme since here only one side uses an UW flux and the other uses CDS. When averaged together, 75% of the total flux will come from the up-winded side, and only 25% from the CDS side. Since any combination of these three scenarios can happen in practice, the TVD* scheme is a significant departure from the CDS scheme.

The three proposed schemes are tested in §6.2, where we find that the TVD* scheme performs best. Improving this TVD* is still possible, since minor oscillations can remain. A proper flux-limited advection scheme may be derived by looking at the characteristics of the hyperbolic parts of the full system. However, since the system has $(s + 1) \times d$ equations, where d is the dimension of the problem, this analysis is left for future research.

4.2. Discretization of Stochastic Subspace

The stochastic coefficients exist in an s -dimensional space, which could become large, from $O(10)$ to $O(10^3)$ based on our experience. In most cases, there is also no strict bound on the value that a stochastic coefficient can take. Thus, evenly dividing the s -dimensional space is not feasible. To discretize the uncertainty subspace, other schemes are thus used. They include: (a) non-uniform discretizations of the subspace, either using structured or

unstructured grids, possibly using schemes based on finite-volumes or finite-elements [16, 41]; (b) solve a discretized version of the PDEs for the probability densities of the coupled s coefficients, e.g. solve Fokker-Planck equations [50]; (c) parameterize the probability space, either using Polynomial Chaos [16, 68, 42, 14], in our case extended to time-dependent polynomials, or other parameterizations such as Gaussian mixtures [44, 56] or particle filters [9], and (d), use a Monte-Carlo approach [35, 30].

We have employed a few of these schemes. Here, we only illustrate a Monte-Carlo scheme: for each realization the time-integration of Sect. §3.4 is used. In general, the expected error for the mean and covariance is of $\mathcal{O}\left(\frac{1}{\sqrt{q}}\right)$, where q is the number of samples. For efficient results, it is important that samples are generated in regions where the probability is relatively high, based on importance sampling [9]. At initial time t_0 , the distribution of the Y_i is generated using the specified initial probability density given by eqns. (2). Here, our focus is on numerical schemes and we restrict ourselves to simple distributions such as Gaussians or Dirac functions.

Thus, we discretize Y_i by generating q samples, and forming a $q \times s$ dimensional matrix $Y_{r,i}$. Each row in $Y_{r,i}$ corresponds to one of the samples, and each column corresponds to one of the modes. During the time-integration step, each sample of $Y_{r,i}$ is then advanced using (10), which is done efficiently. In all cases computed so far, q can be large, e.g. $\sim \mathcal{O}(10^4 - 10^5)$, but still sufficiently small such that advancing (10) for every sample is not the dominating cost of the whole scheme.

A drawback to this Monte-Carlo approach is that rare events will not be captured unless a very large number of samples are used. Alternative methods, such as Gaussian mixtures [55, 56] or other approaches mentioned above can then be used as an alternative.

5. Implementation Details

In this section we describe selected implementation details that arise due to numerical issues. In particular, we discuss how to deal with possibly poorly conditioned covariance matrices in the stochastic subspace, as well as the orthonormalization of the modes and decorrelation of the stochastic coefficients.

5.1. Dealing with a Singular Covariance Matrix

The covariance matrix may be singular or poorly conditioned if one or more of the stochastic coefficients have zero or very small variance compared to other modes. This situation for example arises if a system has deterministic initial conditions, but becomes uncertain through forcing, boundaries, parameters, numerical uncertainties, or other causes. The initial covariance matrix is then simply zero: its inverse is not defined. Special treatment is thus needed for such cases since the inverse of the covariance matrix is required in (8).

Fortunately, this problem is also common in data assimilation where it is resolved using generalized Moore-Penrose inversions [1, 35]. In the particular case of the DO equations, the inverse of the covariance matrix is multiplied by the third moments in (8),

$$\mathbf{C}_{Y_i Y_j}^{-1,k*} \mathbf{M}_{Y_j Y_m Y_n}^{k*}, \quad (12)$$

which, for most physical processes, goes to zero for the eigenvalues of $\mathbf{C}_{Y_i Y_j}^{k*}$ that go to zero. However, to ensure a numerically stable estimate of the inverse of the covariance matrix, we employ a generalized Moore-Penrose inverse, which amounts to truncate the singular values less than a defined tolerance or to set them to that tolerance (with the former, the inverse of a zero covariance, i.e. deterministic initial conditions, is zero). This results in a stable numerical simulation, as exemplified for the Lock Exchange validation benchmark in §6.1.

5.2. Orthonormalization

The DO equations enforce orthonormal modes and it is important to maintain this property numerically when integrating over time. Analytically, if modes are orthonormal initially, orthonormality is maintained because

$$\frac{\partial \langle \Phi_i, \Phi_j \rangle_{\mathcal{D}}}{\partial t} = \left\langle \frac{\partial \Phi_i}{\partial t}, \Phi_j \right\rangle_{\mathcal{D}} + \left\langle \Phi_i, \frac{\partial \Phi_j}{\partial t} \right\rangle_{\mathcal{D}} = 0,$$

where the DO condition $\left\langle \Phi_i, \frac{\partial \Phi_j}{\partial t} \right\rangle_{\mathcal{D}} = 0, \quad \forall i, j \in [1, 2, \dots, s]$ was used. At the discrete level, this property is maintained, up to truncation and round-off errors. Even if the modes are orthonormal at a given time-step, not all integration schemes over the next time-step will conserve the discrete orthonormality.

Let's first consider the analytical integration from time t_{k-1} to t_k . For the modes, we have $\Phi_i^k = \Phi_i^{k-1} + \Delta t \frac{\partial \Phi_i}{\partial t}$ where the exact time integral is denoted as $\Delta t \frac{\partial \Phi_i}{\partial t} = \int_{t_{k-1}}^{t_k} \frac{\partial \Phi_i}{\partial t} dt$. For the inner product at t_k , we then have:

$$\begin{aligned} \langle \Phi_i^k, \Phi_j^k \rangle_{\mathcal{D}} &= \left\langle \Phi_i^{k-1} + \Delta t \frac{\partial \Phi_i}{\partial t}, \Phi_j^{k-1} + \Delta t \frac{\partial \Phi_j}{\partial t} \right\rangle_{\mathcal{D}}, \\ &= \langle \Phi_i^{k-1}, \Phi_j^{k-1} \rangle_{\mathcal{D}} + \Delta t \left\langle \Phi_j^{k-1}, \frac{\partial \Phi_i}{\partial t} \right\rangle_{\mathcal{D}} \\ &\quad + \Delta t \left\langle \Phi_i^{k-1}, \frac{\partial \Phi_j}{\partial t} \right\rangle_{\mathcal{D}} + \Delta t^2 \left\langle \frac{\partial \Phi_i}{\partial t}, \frac{\partial \Phi_j}{\partial t} \right\rangle_{\mathcal{D}}. \end{aligned}$$

Since $\langle \Phi_i^k, \Phi_j^k \rangle_{\mathcal{D}} = \langle \Phi_i^{k-1}, \Phi_j^{k-1} \rangle_{\mathcal{D}} = \delta_{ij}$, the remaining terms sum to zero,

$$\Delta t \left\langle \Phi_j^{k-1}, \frac{\partial \Phi_i}{\partial t} \right\rangle_{\mathcal{D}} + \Delta t \left\langle \Phi_i^{k-1}, \frac{\partial \Phi_j}{\partial t} \right\rangle_{\mathcal{D}} + \Delta t^2 \left\langle \frac{\partial \Phi_i}{\partial t}, \frac{\partial \Phi_j}{\partial t} \right\rangle_{\mathcal{D}} = 0 \quad (13)$$

Considering now a p^{th} -order-accurate discrete approximation of $\frac{\partial \Phi_i}{\partial t}$, its error, ε_i , is of $\mathcal{O}(\Delta t^p)$. Assuming that the spatial inner product is computed exactly, this error in the time integration leads to an error, \mathcal{E}_{ij} , in the inner product at time k , i.e. $\langle \Phi_i, \Phi_j \rangle_{\mathcal{D}}^{k, \text{discrete}} = \langle \Phi_i, \Phi_j \rangle_{\mathcal{D}}^k + \mathcal{E}_{ij}$. To estimate the magnitude of this \mathcal{E}_{ij} made over one time step, we assume

an exactly orthonormal inner product at t_{k-1} . By discrete integration to t_k , we then have:

$$\begin{aligned} \langle \Phi_i, \Phi_j \rangle_{\mathcal{D}}^k + \mathcal{E}_{ij} &= \langle \Phi_i^{k-1}, \Phi_j^{k-1} \rangle_{\mathcal{D}} + \Delta t \left\langle \Phi_j^{k-1}, \frac{\partial \overline{\Phi_i}}{\partial t} + \varepsilon_i \right\rangle_{\mathcal{D}} \\ &\quad + \Delta t \left\langle \Phi_i^{k-1}, \frac{\partial \overline{\Phi_j}}{\partial t} + \varepsilon_j \right\rangle_{\mathcal{D}} \\ &\quad + \Delta t^2 \left\langle \frac{\partial \overline{\Phi_i}}{\partial t} + \varepsilon_i, \frac{\partial \overline{\Phi_j}}{\partial t} + \varepsilon_j \right\rangle_{\mathcal{D}}. \end{aligned}$$

Using (13),

$$\begin{aligned} \mathcal{E}_{ij} &= \Delta t \langle \Phi_j^{k-1}, \varepsilon_i \rangle_{\mathcal{D}} + \Delta t \langle \Phi_i^{k-1}, \varepsilon_j \rangle_{\mathcal{D}} \\ &\quad + \Delta t^2 \left\langle \varepsilon_i, \frac{\partial \overline{\Phi_j}}{\partial t} \right\rangle_{\mathcal{D}} + \Delta t^2 \left\langle \frac{\partial \overline{\Phi_i}}{\partial t}, \varepsilon_j \right\rangle_{\mathcal{D}} + \Delta t^2 \langle \varepsilon_i, \varepsilon_j \rangle_{\mathcal{D}}, \\ &\sim 2\Delta t \mathcal{O}(1) \mathcal{O}(\Delta t^p) + 2\Delta t^2 \mathcal{O}(\Delta t^p) \mathcal{O}(\Delta t^{-1}) + \Delta t^2 \mathcal{O}(\Delta t^{2p}), \\ &\sim \mathcal{O}(\Delta t^{p+1}). \end{aligned}$$

Therefore, the error in the orthonormality will always be Δt smaller than that of the numerical scheme.

The orthonormality can be corrected indirectly by enforcing that the solution and the error are numerically orthonormal, $\langle \Phi_i, \Phi_j \rangle_{\mathcal{D}}^k + \mathcal{E}_{ij} = \delta_{ij}$, at the end of a time step. This has to be done with care: the summation $Y_{r,i} \Phi_i$ produces specific realizations, and changing the basis without modifying the coefficients will change the specific realizations. Because Φ_i and $Y_{r,i}$ are linked, various schemes for performing the orthonormalization exist. They are described in Appendix B.

6. Numerical Applications

In this section we present four benchmarks used to verify and study the schemes described above. To ensure that the implementation is solving the desired equations, we compare the stochastic code to a deterministic code (§6.1) for a version of the lock-exchange dynamical problem [19]. Next, we examine three advection schemes proposed for DO equations in §4.1.2, using a symmetric version of the lock-exchange problem (§6.2). Then, we evaluate the spatial and temporal convergence using the lid-driven cavity flow (§6.3). Finally, we study the discretization of the stochastic coefficients using the flow over a square cylinder in a confined channel (§6.4).

For evaluating errors, we use the L^2 norm. At a single time instance, the norm is $\|\Phi\|_2 = \sqrt{\langle \Phi, \Phi \rangle_{\mathcal{D}}}$ for the whole state vector, or $\|\phi\|_2 = \sqrt{\int_{\mathcal{D}} \phi^2 d\mathcal{D}}$ for a single component. For the convergence studies (§6.3) we also integrate over time, using $\|\Phi\|_2^T = \sqrt{\int_0^T \langle \Phi, \Phi \rangle_{\mathcal{D}} dt}$ for the state vector, and $\|Y_i\|_2 = \sqrt{\int_0^T \mathbb{E}[Y_i Y_i] dt}$ for the stochastic coefficients.




Figure 1: Lock-exchange problem: an initial barrier separating light and heavy fluid is removed, and the flow is allowed to evolve. Uncertainty in our studies originates from not knowing the initial density differences between the fluids. This benchmark is used to verify the correctness of the implementation, not the DO methodology.

6.1. Lock-Exchange Verification Benchmark

The purpose of this benchmark is to verify the numerical implementation. For example, our use of stochastic pseudo-pressures does not in theory introduce additional errors. However, one needs to verify that their numerical implementation is accurate, solving the correct equations. Ideally, problems with analytical solutions should be used to verify a code, however, constructing a valid analytical solution of (6), (8) and (10) with multiple stochastic modes and coefficients is neither trivial, nor does it lend itself to compact expressions. In this section we address this problem by defining a numerical benchmark and then using it to verify the present DO code.

6.1.1. Lock-Exchange Verification Benchmark: Setup

We verify our stochastic DO code by comparing it to a deterministic NS code which *has* been thoroughly verified [34, 59]. This deterministic code uses the same second-order Finite volume scheme and first order backwards difference Projection method as the stochastic code. While the DO code is inherently more complicated with coupled equations, the major differences are in the advection scheme for the stochastic modes (see §4.1.2 and §6.2), and in the need to also evolve the stochastic coefficients.

The deterministic code is used non-intrusively with a Monte-Carlo method to generate an ensemble of independent realizations. These references are then compared to realizations from the DO code which solves the coupled DO equations. The benchmark (Fig. [1]) is based on the lock-exchange problem [19], where uncertainty is introduced by prescribing four possible initial density differences. In other words, while the exact difference between the densities is unknown, it is known that only four possibilities of equal probability exists: i.e. the pdf is initialized as four discrete Dirac delta function. Essentially, we employ a single run of the DO code to try to replicate four independent deterministic runs. While not a practical use of the DO method, it is a challenging benchmark for verifying its numerical implementation.

To capture all of the uncertainty, three DO modes ($s = 3$) are needed for the stochastic simulation. Four density difference were chosen because three DO modes is the minimum




Figure 2: Initialization of the lock-exchange problem (density field: mean, mode 1 and marginal pdf, modes 2 and 3). The initial velocity is zero, and stochasticity is introduced through the first (of three) orthonormal modes for the density. The vertical length scale used for non-dimensionalization is the half-height of the channel ($\hat{h} = 1$). Initializing with four stochastic samples for the first mode, $Y_{r,1} = [-0.18, -0.06, 0.04, 0.20]^T$, we have four possible initial conditions corresponding to $\Delta\rho = [1, 0.84, 0.74, 0.62]$. The two remaining modes are initialized as described in the text, and do not introduce additional stochasticity.

number required to have full energy interactions between the mean and modes of the DO simulation [50, 49]. Also, to fully verify the implementation, we need to ensure that the initial pdf is non-symmetric so that the third moments are non-zero (in (10) for example).

General setup: The Schmidt number is kept constant, $Sc = 1$, and we present results for $Gr = 1.25 \times 10^6$ and $Gr = 4 \times 10^4$, (although other Gr were studied). The four density differences of equal probability are $\Delta\rho = [1, 0.84, 0.74, 0.62]$, with the initial density profile prescribed by

$$\rho(x, y, t=0) = \frac{\Delta\rho}{2} \tanh(2x/l_\rho),$$

where we take $l_\rho = 1/64$. Initially the velocity is zero everywhere. Free-slip boundary conditions are used at the domain boundaries (Fig. [1]). The domain is discretized using $\Delta x = \Delta y = 1/256$, and $\Delta t = 1/512$, which is sufficient resolution for these Grashoff numbers [19].

Mean initialization: The mean density profile (Fig. [2]-a) uses the hyperbolic tan function specified above with a mean density difference $\Delta\rho = 0.8$. The mean pressure and velocity are zero everywhere, initially.

Mode initialization: The density profile for the first mode is the hyperbolic tan profile above (Fig. [2]-b), but normalized. The two remaining modes are arbitrary, since they do not introduce initial uncertainty (see below). They are set to:

$$\rho_2(\mathbf{x}, t=0) = \begin{cases} (\Delta\rho - |\bar{\rho}|)\text{sign}(\bar{\rho})|\sin(\pi y)| & \text{if } (\Delta\rho - |\bar{\rho}|)\text{sign}(\bar{\rho})\sin(\pi y) > 0 \\ 0 & \text{otherwise,} \end{cases}$$

$$\rho_3(\mathbf{x}, t=0) = \begin{cases} (\Delta\rho - |\bar{\rho}|)\text{sign}(\bar{\rho})|\sin(\pi y)| & \text{if } (\Delta\rho - |\bar{\rho}|)\text{sign}(\bar{\rho})\sin(\pi y) < 0 \\ 0 & \text{otherwise,} \end{cases}$$

where these are orthonormalized numerically as described in §5.2. The pressure and velocity for all modes are zero everywhere, initially.

Stochastic coefficient initialization: The discrete pdf for the first stochastic coefficient is

specified with four samples $Y_{r,1} = [-0.18, -0.06, 0.04, 0.20]^T$, (Fig. [2]-c). The next two coefficients do not introduce additional uncertainty because we specify perfectly correlated samples, $Y_{r,2} = Y_{r,3} = \epsilon \cdot [-0.18, -0.06, 0.04, 0.20]^T$, where ϵ is a small constant chosen such that $\sum_i \text{Var}(Y_{r,i}) = \text{Var}(Y_{r,1})$ numerically. This means that the inverse of the covariance matrix is very ill-conditioned (or numerically singular), so the pseudo-inverse is required during the initial stages of the simulation (see §5.1). Also, because the pdf is discrete, we calculate moments using the biased estimator $\mathbf{C}_{Y_i Y_j} \approx \frac{1}{q} \sum_r Y_{r,i} Y_{r,j}$, instead of the usual unbiased estimator $\mathbf{C}_{Y_i Y_j} \approx \frac{1}{q-1} \sum_r Y_{r,i} Y_{r,j}$. The initial fields and pdf are shown in Fig. [2].

6.1.2. Lock-Exchange Verification Benchmark: Results and Discussion

The outputs from the stochastic run are reported in Fig. [3] for both Grashoff numbers. Comparisons with the deterministic runs are shown in Fig. [4] and Fig. [5] for $Gr = 4 \times 10^4$ and $Gr = 1.25 \times 10^6$, respectively. Finally, the evolution of the differences between the stochastic and deterministic runs for both Grashoff numbers are shown in Fig. [6].

We see excellent agreement between the stochastic realizations and the deterministic runs (Fig. [5]-Fig. [6]) for this challenging benchmark. Particularly, for the lower Grashoff number flow, the local error is less than 0.2% everywhere. It is non-trivial that the complex DO implementation is capable of reproducing multiple deterministic runs in a single simulation. Thus, based on these results and many other tests (not shown), we conclude that our implementation is correct, that is, we are solving the intended equations.

The growth of differences between the deterministic and stochastic simulations (Fig. [6]) should be explained by the main differences between the stochastic and deterministic solvers, in particular the advection schemes, and the evolution of the stochastic coefficients. We found that the magnitude of the differences over time is larger for coarser space and time resolution runs. This suggests the error is due to spatial and/or temporal truncation error. The advection scheme does not contribute significantly to the error at the reported resolution, since using a CDS advection scheme instead of the TVD* scheme for the stochastic modes did not change the reported results significantly. Reducing the time-step to $\Delta t = 1/1024$ reduced the error at the final time from $\sim 2.1\%$ to $\sim 1.05\%$ for the higher Grashoff number flow. This indicates that the error is dominated by a temporal truncation compounded error of approximately $\mathcal{O}(\Delta t)$ (as expected, see §6.3). The primary source of this compounded error may be from the evolution of the stochastic coefficients and/or the modes.

We examine the errors more closely to determine the primary source of the small numerical errors. Note that the density differences are either a bit too small (first realization in Fig. [5]) or a bit too large (last three realizations in Fig. [5]), causing phase errors as observed around the density interface. These phase errors can be caused by small errors in the relative magnitudes of the stochastic coefficients. We found that different orthonormalization strategies (§5.2), which affect the relative magnitudes of the stochastic coefficients, can also impact the magnitude of the phase errors. In particular, when using a Gram-Schmidt orthonormalization, the average difference was greater than 3% (or 50% worse) for the high Gr case at the final time. Also, considering that the larger Gr number flow has larger stochastic coefficients (Fig. [3]) than the lower Gr flow, the same relative error in the stochastic coefficients would lead to a larger phase error (observed when comparing Fig. [4] to Fig. [5]). Thus, this benchmark's results indicate that the primary source of error originates from the time evolution of the stochastic coefficients.

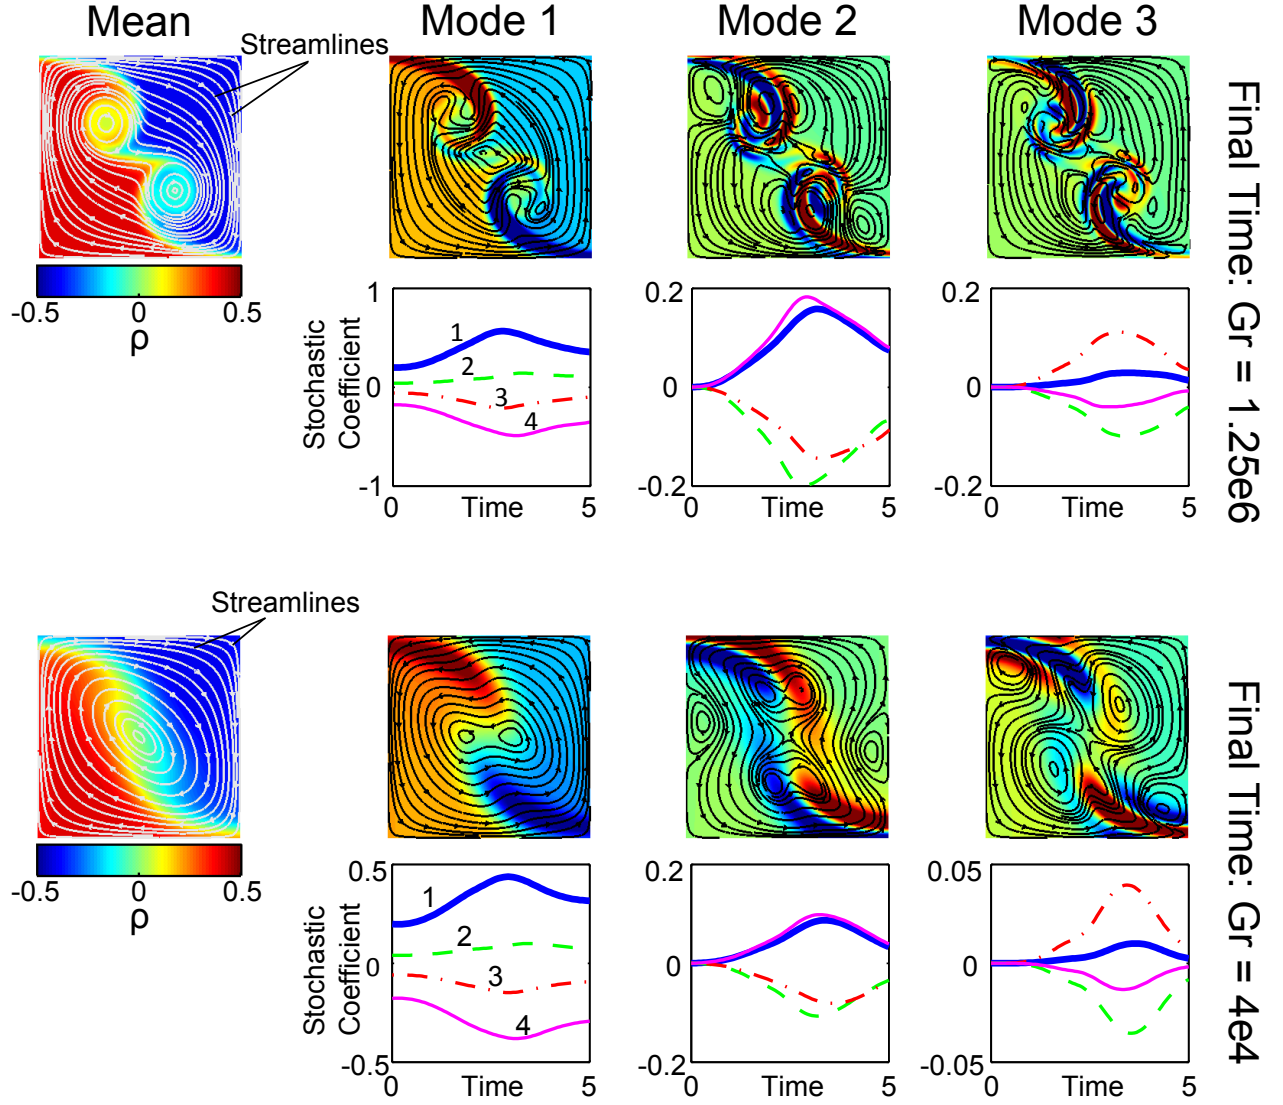


Figure 3: The DO mean, modes at non-dimensional time $t = 5$, and evolution of stochastic coefficients, for $Gr = 1.25 \times 10^6$ (top) and $Gr = 4 \times 10^4$ (bottom). The higher-Gr flow has sharper gradients, and its coefficients are larger than the lower-Gr flow. Streamlines shown over density in color. Note that the sign of the contribution from a mode depends on the sign of the stochastic coefficients.

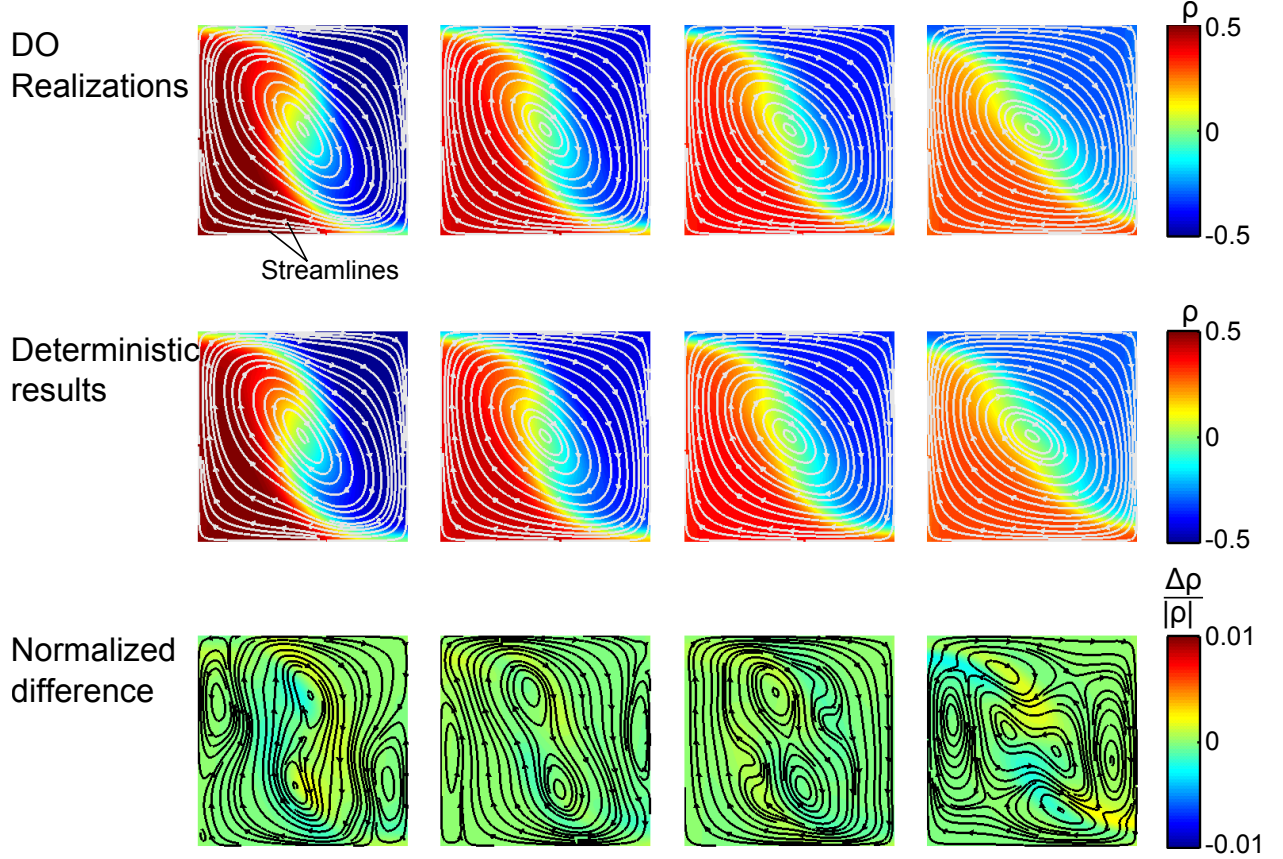


Figure 4: The four DO realizations (top row), the deterministic runs (middle row), and the DO realizations minus the deterministic runs normalized by $\|\Phi_{\text{deterministic}}\|_2$ (bottom row) for $Gr = 4 \times 10^4$ (resolution 256×256 with 512×5 time-steps). The stochastic solver result agrees with the deterministic solver results, which validates our DO numerical schemes. Streamlines shown over density in color.




Figure 5: As Fig. [4] but with $Gr = 1.25 \times 10^6$. The higher-Gr flow with sharper gradients has larger errors than the lower-Gr flow, but they are still small.




Figure 6: The relative errors for both Gr tend to grow over time but can decrease. In both cases, the DO mean field has a smaller error than the average error of the realizations.

In summary, using the proposed benchmark, we have verified the correct implementation of the stochastic code. Also, based on the results, we suggest that the accuracy of the scheme will benefit most from improving the temporal discretization for problems that require long time integration.

6.2. Effect of Advection Scheme

The purpose of this benchmark is to test the three different advection schemes proposed (§4.1.2). We again use a version of the lock-exchange problem because the sharp density interface will highlight numerical oscillations. We modify the problem by introducing symmetry, which should be maintained numerically. Finally, for simplicity we only consider a single stochastic mode with a bimodal continuous pdf.

General setup: The Schmidt number is kept constant, $Sc = 1$, and we present results for $Gr = 1.25 \times 10^6$. Initially the velocity is zero everywhere. Free-slip boundary conditions are used at the boundaries of the domain (Fig. [1]). The domain is discretized using $\Delta x = \Delta y = 1/64$, and $\Delta t = 1/256$. A lower resolution is used here compared to the cases in §6.1 in order to highlight the symmetry errors and numerical oscillations

Mean initialization: The mean density, pressure, and velocity are all zero everywhere, initially.

Mode initialization: The density profile for the mode is the same normalized hyperbolic tan profile used in §6.1. The pressure and velocity for this mode are zero everywhere, initially.

Stochastic coefficient initialization: The bimodal Gaussian continuous pdf is represented by 10,000 samples. To ensure that any asymmetry comes from the numerical errors only, we exactly enforce symmetry in the initial conditions as follows. We first generate 2,500 samples ($Y_{r_{2500}}$) from a zero mean Gaussian distribution with standard deviation $\sigma = 0.01e^1 \approx 0.027$. Next, these samples are duplicated to obtain a representation of the bimodal pdf that is exactly symmetric at the numerical level:

$$Y_{r,1} = \left[Y_{r_{2500}} - \frac{1}{2}, \quad -Y_{r_{2500}} - \frac{1}{2}, \quad Y_{r_{2500}} + \frac{1}{2}, \quad -Y_{r_{2500}} + \frac{1}{2} \right]^T.$$

These are the 10,000 samples that we evolve. They are illustrated in the first row of Fig. [7]. Of course, these samples are correlated. The procedure should not be used in general; it is used here solely to evaluate how good are numerical schemes at maintaining symmetry.

6.2.1. Effect of Advection Scheme: Results

The result of this simulation is shown in Fig. [7] for the three advection schemes. While the CDS scheme maintains symmetry of the mean, modes, pdf, and realizations, some clear numerical oscillations are present, particularly evident in the realizations. While there are no oscillations in the TVD scheme, it clearly loses symmetry in the mean, modes, pdf, and realizations, as can be seen (with aid of the dashed guide lines) in Fig. [7]. The TVD* scheme only develops minor oscillations, which can be barely detected when examining the realizations, and completely retains symmetry. Thus, the new TVD* scheme is the preferred scheme among the three.

Initially, we can represent all density realizations exactly with one mode. However, as different densities evolve at different rates, one mode becomes insufficient to represent the

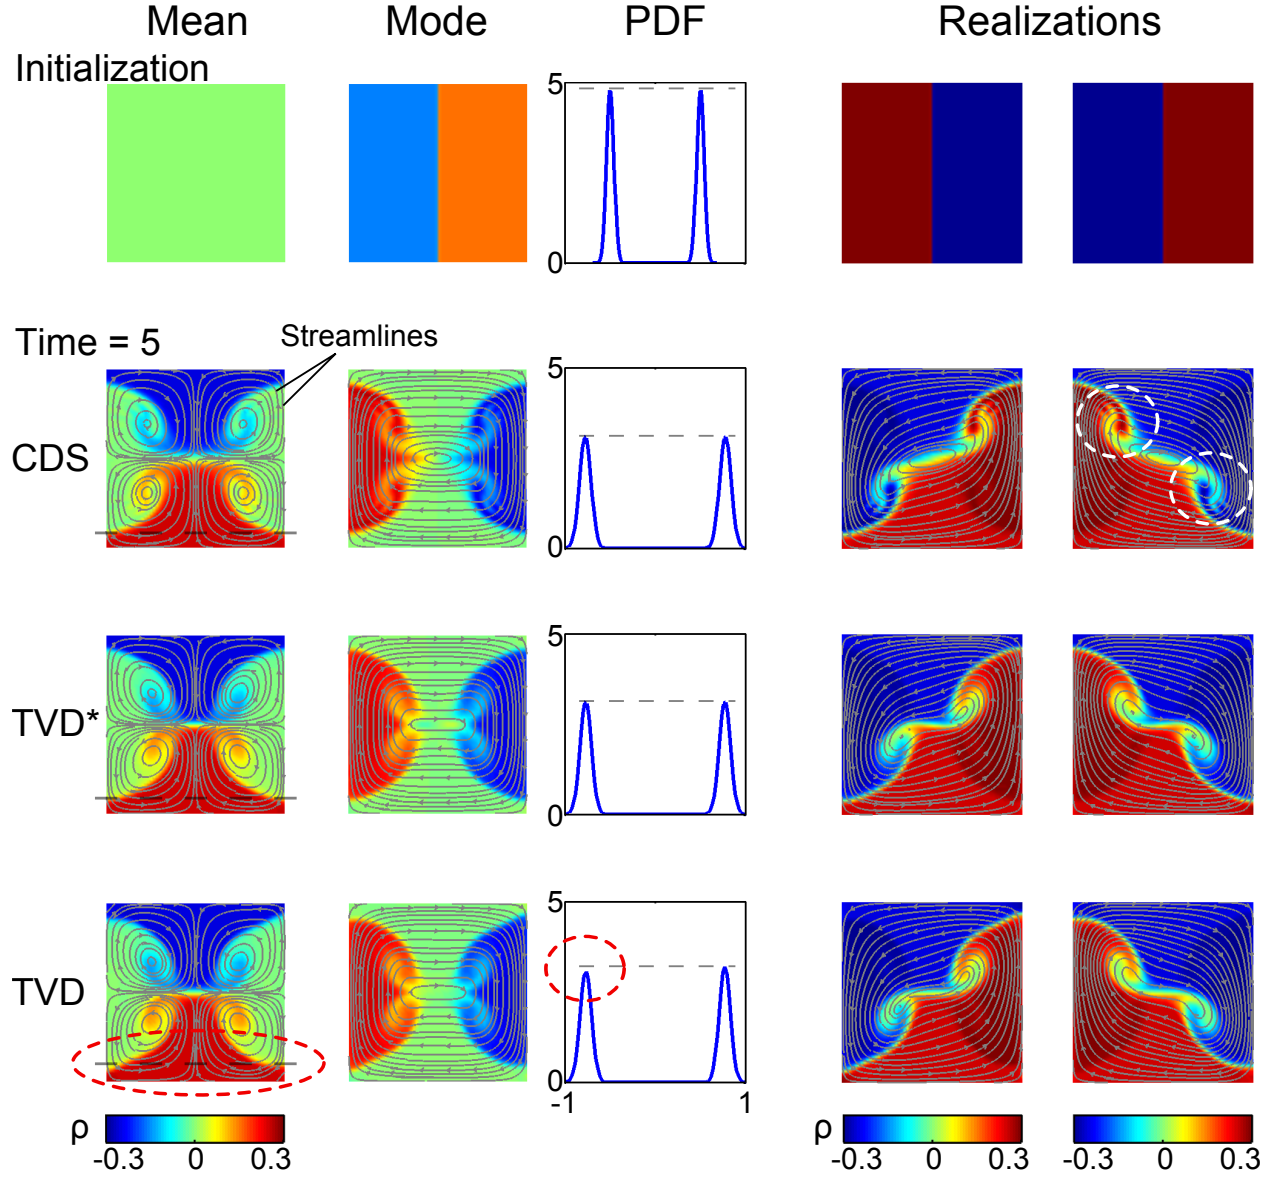


Figure 7: Symmetric lock-exchange problem ($Gr = 1.25 \times 10^6$) with grid resolution 64×64 and $\Delta t = 1/256$ using various advection schemes for the modes. Only the stochastic density is non-zero initially, with a bi-modal pdf(top row). The two most extreme realizations (largest and smallest stochastic coefficients) are plotted in each case (right column). Our new averaged TVD* scheme only has minor oscillations and retains symmetry (third row), while the CDS advection scheme suffers from large oscillations (white dashed circles, second row), and the one-sided TVD scheme loses symmetry (red dashed circles, last row). Streamlines shown over density in color.

uncertainty. Hence, spurious gradients can appear in the reconstructed realizations. We purposely chose to use only one mode in this benchmark to also illustrate that if the number of modes is fixed, errors occur. In general, we do not keep the number of modes fixed [51].

In summary, we found that the TVD* scheme performs adequately and that our DO implementation can reproduce vastly different realizations of a given problem.

6.3. Numerical Convergence Analysis

The purpose of this benchmark is mainly to show that the implemented scheme is converging. Here we use the classical lid-driven cavity flow, and examine the numerical convergence under spatial and temporal refinement of each component separately. This benchmark does not have a variable density, and so we report the Reynolds number instead of the Grashoff number. We also completed convergence tests with density, with results analogous to those presented below.

6.3.1. Numerical Convergence Analysis: Setup




Figure 8: The lid-driven cavity flow is a classical benchmark used to verify convergence of numerical implementations. Uncertainty for this case will be introduced through the initial conditions.

General setup: We present results for a Reynolds number of $Re = 500$, although other Reynolds number cases ($Re \in [100, 1000]$) were also studied, giving similar results. The flow is driven by a deterministic boundary condition at the top of the enclosed box Fig. [8], with no-slip velocity boundary conditions, and uncertain initial conditions. The finest resolution uses $\Delta x = \Delta y = 1/512$, and $\Delta t = 1/4096$, which is sufficient at this Reynolds number [12, 11], and the order of convergence is approximated as $\mathcal{O} \approx \log \left(\frac{\|\phi_{2N_t} - \phi_{4N_t}\|_2}{\|\phi_{2N_t} - \phi_{N_t}\|_2} \right) / \log(2)$. We examine the difference between the fine and coarse space resolutions by interpolating (using splines) the fine solution onto the coarse resolution grid, and taking the L^2 norm over the interior of the domain, $\mathcal{D}_I \in [0.25, 0.75] \times [0.25, 0.75]$, to avoid the boundary condition singularities at the top two corners.

Mean initialization: For a challenging case, the mean velocity and pressure are initially zero, everywhere.

Mode initialization: The velocity modes are initialized by specifying the stream function

$$\psi_{M,N}(x, y) = C_{M,N} \sin(\pi x) \sin(\pi M x) \sin(\pi y) \sin(\pi N y),$$

where $C_{M,N} = \frac{1}{\pi} \sqrt{\frac{N^2+1}{16} \left(\frac{3}{2}\right)^{\delta(N-1)} + \frac{N^2+1}{16} \left(\frac{3}{2}\right)^{\delta(M-1)}}$ is the normalization constant; the delta function $\delta(x)$ takes the value 1 if $x = 0$ and 0 otherwise. The velocity modes are then specified as

$$u_i = -\frac{\partial}{\partial y} \psi_{(M,N)_i}, \quad v_i = \frac{\partial}{\partial x} \psi_{(M,N)_i},$$

where $(M, N) = \{(1, 1), (1, 2), (1, 3)\}$. The initial pressure for the modes is specified as zero everywhere.

Stochastic coefficient initialization: The pdf is created using 5,000 samples of zero mean Gaussian distributions with variances $\text{Var}(Y_{r5000,i}) = e^{(1-M_i-N_i)}$. The number of samples is not critical in this case, since the samples are the same from one run to the next. That is, we only test spatial and temporal convergence, not stochastic convergence. Since this is, again, a numerical test, the Y_i samples are purposely created using a procedure as in §6.2,

$$Y_{r,i}^* = \begin{bmatrix} Y_{r5000,i} \\ -Y_{r5000,i} \end{bmatrix}.$$

To ensure that the final generated samples have numerical variances exactly as specified, we correct the samples using the numerically calculated variance

$$Y_{r,i} = Y_{r,i}^* \frac{\sqrt{\text{Var}(Y_{r5000,i})}}{\sqrt{\text{Var}^q(Y_{r,i})}},$$

where $\text{Var}^q(a_{r,i}) = \frac{1}{q-1} \sum_{r=1}^q a_{r,i}$ is the calculated sample variance. The initialization for this problem can be seen in the first row of Fig. [9].

6.3.2. Numerical Convergence Analysis: Results and Discussion

Three time snapshots of the reference solution are shown in Fig. [9]. While the marginal seems to remain approximately Gaussian, the sample-scatter diagram clearly shows the very non-Gaussian behavior for this benchmark. Also, we confirm from Fig. [10] that the variances of the modes are decreasing, which is expected since this problem has a deterministic steady-state solution. Thus, the stochastic solution is behaving as expected.

Next, examining the convergence, we see that the numerical error for all components decrease with temporal and spatial refinement. The convergence is near optimal at large grid sizes for all variables. These results with the velocity components separated are tabulated in Table [1] and Table [2], which also shows that the stochastic coefficients are converging optimally. Even though a fourth-order RK method is used to advance the stochastic coefficients, the total-DO convergence is first order (based on choices made in §3.5). Thus, we observe near-optimal convergence for all variables, which suggests that the implementation is correct.

6.4. Stochastic Convergence

The purpose of the fourth benchmark is to further assess numerical performance, demonstrate that a significant number of modes can be used, and study the effect of the stochastic discretization. Specifically, we quantify the effects on accuracy of the number of stochastic




Figure 9: Evolution of reference solution for the lid-driven cavity flow ($Re = 500$) with stochastic initial conditions over time (resolution 512×512 with 4096×5 time-steps). This case demonstrates that the DO representation with our implementation is able to capture and evolve non-Gaussian statistics (see sample scatter plot) for a continuous pdf. The marginal pdfs are normalized on the plot, and the variance can be read off Fig. [10]. Streamlines shown over pressure in color. The sample scatter is colored by the $s = 3$ stochastic coefficient.




Figure 10: Evolution of the stochastic energy ($\text{Var}(Y_i)$) for the lid-driven cavity reference solution. The tick-marks on the time-axis corresponds to the time snapshots in Fig. [9].




Figure 11: The control volume size is held fixed at $\Delta x = 1/512$ for the time convergence (left), and the time step size is held fixed at $\Delta t = 1/4096$ for the spatial convergence (right). The error ($\|\phi_{2N} - \phi_N\|_2$) decreases with both temporal- and spatial-refinement for each component, and convergence is near-optimal (order 1 in time and 2 in space).




Figure 12: Laminar vortex shedding over a square cylinder in a channel. Here uncertainty originates from the initial conditions and uncertain vortex shedding: depending on the perturbation, the first vortex could either be shed above or below the cylinder. Within a stochastic framework, however, if uncertainties are initially symmetric, the mean and modes should remain symmetric, and this is evaluated with our DO numerics.

.	N_t	Mean		Mode 1		Mode 2		Mode 3	
		$\ e\ _2$	\mathcal{O}	$\ e\ _2$	\mathcal{O}	$\ e\ _2$	\mathcal{O}	$\ e\ _2$	\mathcal{O}
P	2048	2.9e-04	1.5	5.8e-04	1.5	3.0e-03	1.5	2.3e-03	1.5
	1024	5.7e-04	1.0	1.2e-03	1.0	5.9e-03	0.99	4.7e-03	1.0
u	2048	2.7e-04	1.5	3.7e-04	1.5	1.6e-03	1.5	2.2e-03	1.5
	1024	5.4e-04	1.0	7.3e-04	1.0	3.1e-03	.99	4.3e-03	1.0
v	2048	2.5e-04	1.5	4.7e-04	1.5	1.8e-03	1.5	2.4e-03	1.5
	1024	5.0e-04	1.0	9.5e-04	1.0	3.5e-03	0.99	4.8e-03	1.0
Y_i	2048	.	.	3.1e-04	1.5	5.1e-04	1.5	9.9e-04	1.5
	1024	.	.	6.2e-04	1.0	1.0e-03	1.0	2.0e-03	1.0

Table 1: Temporal convergence of lid-driven cavity flow. Tabulated is the error ($e = \|\phi_{2N_t} - \phi_{N_t}\|_2$) between the solutions using $\Delta t = 1/(2N_t)$ and using $\Delta t = 1/N_t$, and the approximate order of convergence \mathcal{O} . The grid size is fixed at $\Delta x = 1/512$.

samples and of the time-order of integration. To do so, we extend the classic shedding of vortices by a uniform flow as it encounters a symmetric obstacle to stochastic DO computations. Since this problem is symmetric, the stochastic solution can only lose symmetry if numerical or external perturbations initiate the non-symmetric laminar shedding of vortices. However, if perturbations are symmetric, there should be no preferential direction for vortex shedding. Thus, a carefully initialized simulation should be able to capture symmetric directions: this provides an excellent test to assess numerical performance.

6.4.1. Stochastic Convergence: Setup

The benchmark is an open flow in a frictionless pipe with a square cylindrical obstacle (Fig. [12]), which is a classic test for deterministic flow solvers.

General setup: We present results for a Reynolds number of $Re = 100$, although other cases were also studied. The flow is driven by a deterministic uniform inlet boundary condition (left of domain), with slip velocity boundary conditions at the top and bottom, open boundary conditions at the outlet, and symmetric uncertain initial conditions. All simulations use a resolution of 336×63 in space, and 63×40 in time. We choose to integrate until $t = 40$ because this allows the statistics to reach steady values. At $t = 40$ the mean velocity has traveled through the domain 2.5 times.

Mean initialization: The mean velocity and pressure are initially zero, everywhere.

Mode initialization: The exact shape of the initial stochastic perturbations are not important since they are advected out of the domain. However, to maintain symmetry, perturbations have to be symmetric. We initialize the velocity modes by specifying the stream function

$$\psi_{M,N}(x, y) = C_{M,N} \sin(\pi x/a) \sin(\pi M x/a) \sin(\pi y/b) \sin(\pi N y/b),$$

.	N_x	Mean		Mode 1		Mode 2		Mode 3	
		$\ e\ _2$	\mathcal{O}	$\ e\ _2$	\mathcal{O}	$\ e\ _2$	\mathcal{O}	$\ e\ _2$	\mathcal{O}
P	256	1.8e-04	2.2	5.4e-04	1.8	1.9e-03	2.3	2.9e-03	2.0
	128	1.0e-03	2.5	2.1e-03	2.0	1.0e-02	2.4	1.1e-02	2.0
	64	4.5e-03	2.2	9.6e-03	2.2	4.7e-02	2.2	3.0e-02	1.5
u	256	2.0e-04	2.1	5.2e-04	2.1	1.2e-03	1.9	3.6e-03	2.1
	128	1.0e-03	2.4	2.7e-03	2.4	4.9e-03	2.1	1.7e-02	2.2
	64	4.5e-03	2.1	1.3e-02	2.3	2.5e-02	2.4	6.7e-02	2.0
v	256	1.9e-04	2.0	5.6e-04	2.1	1.2e-03	1.9	4.0e-03	2.1
	128	8.6e-04	2.2	2.8e-03	2.3	5.1e-03	2.0	1.9e-02	2.3
	64	3.6e-03	2.1	1.4e-02	2.4	2.6e-02	2.4	7.8e-02	2.0
Y_i	256	.	.	4.8e-04	2.1	7.2e-04	2.1	1.3e-03	2.2
	128	.	.	2.2e-03	2.2	3.4e-03	2.2	6.4e-03	2.3
	64	.	.	7.5e-03	1.8	1.5e-02	2.1	2.1e-02	1.7

Table 2: Spatial convergence of lid-driven cavity flow. Tabulated is the error ($e = \|\phi_{2N_x} - \phi_{N_x}\|_2$) between the refined ($2N_x \times 2N_x$) and present ($N_x \times N_x$) grid, and the approximate order of convergence \mathcal{O} . The time step is fixed at its smallest value $\Delta t = 1/4096$.

where $C_{M,N}$ is the normalization constant (as in §6.3), and $a = 16$, $b = 3$ are the width and height of the domain respectively. The velocity modes are then specified as

$$u_i = -\frac{\partial}{\partial y} \psi_{(M,N)_i} B_M, \quad v_i = \frac{\partial}{\partial x} \psi_{(M,N)_i} B_M,$$

where

$$(M, N) = \{(1, 1), (2, 1), (1, 2), (3, 1), (1, 3), (2, 2), (4, 1), (1, 4), (3, 2), (2, 3)\},$$

and B_M is a smoothing function created numerically from the domain mask. B_M is created from the mask by iteratively averaging each control volume by its own value and its four neighbors for $\frac{2}{\Delta y}$ iterations. That is, at iteration k

$$B_M^k(i, j) = \frac{1}{5} (B_M^{k-1}(i, j) + B_M^{k-1}(i-1, j) + B_M^{k-1}(i, j-1) + B_M^{k-1}(i+1, j) + B_M^{k-1}(i, j+1)).$$

The initial pressure for the modes is specified as zero everywhere.

Stochastic coefficient initialization: To ensure initial symmetry, the samples for the stochastic coefficients are created using the same procedure as in §6.3, using the variances $Var(Y_{r,i}) = e^{(2-M_i-N_i)}$ (note the difference of +1 in the exponential from §6.3). The reference solution uses 10^5 samples for the stochastic coefficients, and a 4th order RK time integration scheme (§3.5).

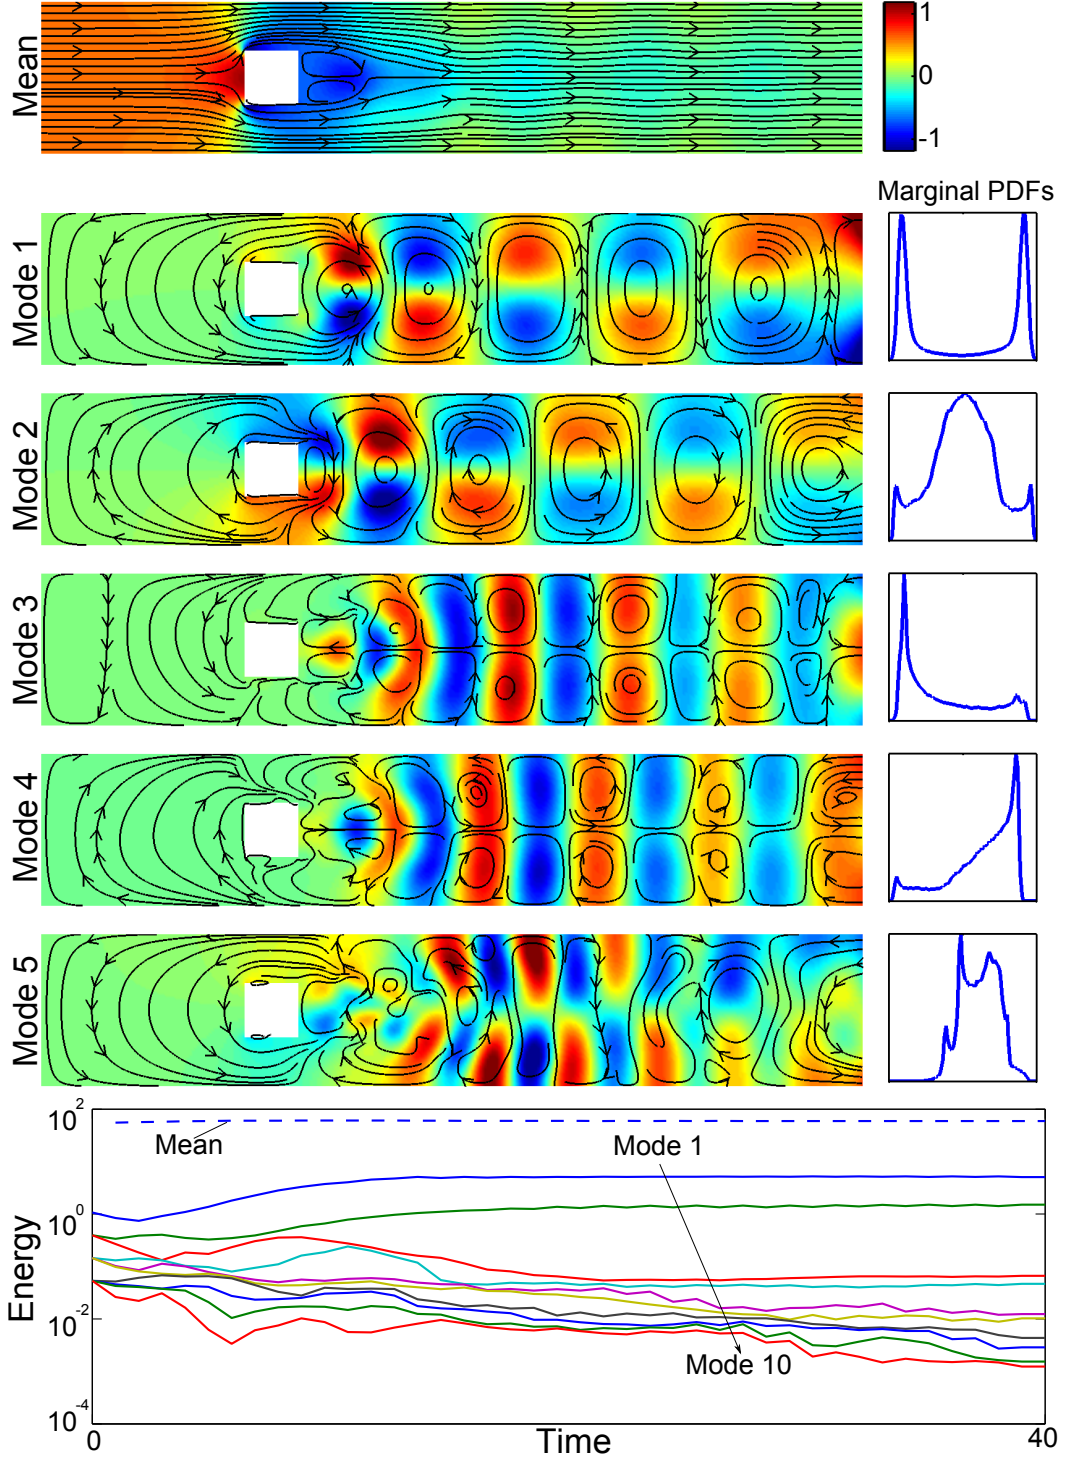


Figure 13: Mean field and first 5 modes with marginal pdfs at final non-dimensional time $T=40$ for $Re=100$, and the evolution of the mean, $\langle \mathbf{u}, \mathbf{u} \rangle_{\mathcal{D}}$ and stochastic energy, $\text{Var}(Y_i)$, for the reference solution (resolution 63×336 with 63×40 time-steps). Streamlines shown over pressure in color. Our scheme and implementation retains (anti)-symmetry for the most important first four modes.




Figure 14: As in Fig. [13], but showing two realizations where the vortex is shed in opposite directions. The colorbar for the pressure is as on Fig. [13].




Figure 15: Energy of the stochastic coefficients, $\text{Var}(Y_i)$, over time for: (top) different number of samples ($\mathcal{O}(4)$ time integration); and (bottom), different time integration schemes (10,000 samples). Trends are well captured in all cases, but there are noticeable errors for less energetic modes after long integration times when a small number of samples and a low order time discretization scheme is used. A relatively small number of samples (10,000) can be used for this benchmark since nearly identical results compared to 100,000 were found for the most energetic four modes.

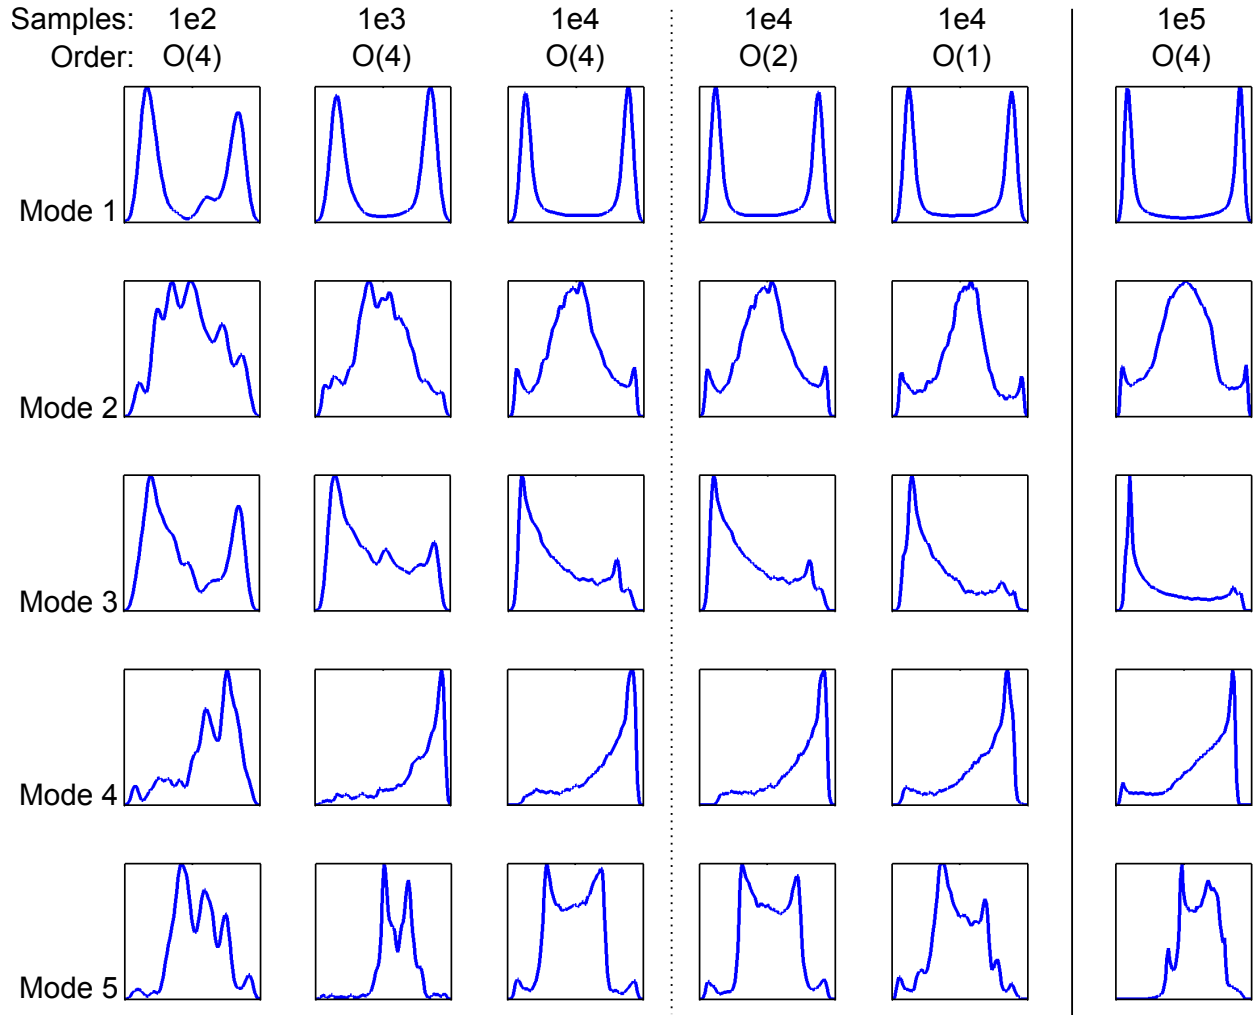


Figure 16: Marginal probability density functions of first five modes for the flow over a square cylinder benchmark at final time (columns 1-3, increasing sample sizes; columns 4-5, decreasing order in time; column 6, reference). With 10,000 samples for the stochastic coefficients, the continuous marginal pdfs are well-represented, although the bimodal peaks lose some symmetry. The marginals have similar shapes for all sample sizes, but the best representations have larger sizes. Overall, the order of the time-integration scheme has less effect than the sample size (the overall order in time is one).

6.4.2. Stochastic Convergence: Results and Discussion

For the reference solution, we find that excellent symmetry is maintained for the first four stochastic modes (Fig. [13]). From the realizations (Fig. [14]), the scheme clearly captures both shedding directions. We find that for fewer samples and lower time integration accuracy, the symmetry is not as well maintained (not shown). However, for a sufficient number of samples, symmetry is maintained for all Reynolds numbers we tested. From Fig. [13], we see that variances seem to reach a steady value after an initial transient period. The smaller variances take longer to reach a steady value; the highest modes still evolve after the final time step.

For our sequential implementation, the reference DO simulation (Fig. [13]) was computed in about 7.5 hours on a 2.4Ghz computer. Each time-step required the inversion of 11 pressure equations, as opposed to the 111 required for the scheme without the stochastic pseudo pressure. This roughly translates to a 1000 % increase in efficiency, or about 3 days saving in terms of computational time for this computer. Since then, we have used as many as $s = 25$ modes for data assimilation applications [56, 57].

Sample sizes: Examining the evolution of the stochastic coefficients for different sample sizes (top of Fig. [15]), we see that the differences are larger for higher modes. For the most energetic first four coefficients, the results agree well with the reference solution when using $> 1,000$ samples. After the 5th coefficient, the differences become larger, and in the 10th coefficient, the differences begin sooner and have larger relative magnitudes. Note that the logarithmic scale magnifies smaller errors, and the differences for the 10th are fully insignificant when compared to the variance of the first mode. Thus, with a relatively small number of samples, it is possible to capture the stochastic coefficient's variances accurately. Also, it appears as though the solution converges when the number of samples increases, which increases our confidence in this approach.

Not only are the evolution of the variances well-represented by a smaller number of samples, the shape of the marginals are also well-reproduced (Fig. [16]). We observe that the marginals are well-reproduced using $> 1,000$ samples, and for all time integration schemes (Fig. [16]). The magnitudes of the bimodal peaks in the first mode are more symmetric with increased sample sizes, which suggests using upwards of 10,000 samples is, perhaps, advised. In fact, using such a number of samples does not affect the overall cost of the solution scheme. That the marginals are well-represented using smaller sizes is encouraging, since it suggests that a small number of samples (that won't negatively impact the efficiency of the scheme) could be used for problems with large s : this would not substantially affect accuracy.

Order in time: We are here only varying the order of the ODE solver used to evolve the stochastic coefficients: the overall order of the DO scheme is kept fixed at first order (see §3.4). Examining results (bottom of Fig. [15]), we see that the second and fourth order ODE solver agree well for all coefficients. A second order solver thus seems sufficient.

In summary, we simulated a stochastic version of the classical flow over a square cylinder benchmark. We found that our DO solver maintains excellent symmetry, and the use of pseudo-pressures reduces the cost by 1000% when using 10 stochastic modes. We also found that the statistics is converging with the number of samples and that accuracy benefits from using a second-order solver for the coefficients, even though the overall DO scheme is

limited to first order. Finally, smaller sample sizes for the coefficients still produced accurate time-evolving statistic and marginal pdfs, which is encouraging.

7. Conclusions and Future Research

We have derived efficient computational schemes for the DO methodology applied to unsteady stochastic Navier-Stokes and Boussinesq equations, and illustrated and studied the numerical aspects of these schemes. For the discretizations in time, we developed semi-implicit projection methods for the mean and the modes, and we employed time-marching schemes of first to fourth order for the stochastic coefficients. For the discretizations of the physical space, we employ conservative second-order finite-volumes, with a special treatment for the advection terms based on a TVD scheme with monotonized central symmetric flux limiter. Several alternate discretizations in time and space were outlined and illustrated by comparisons. We also addressed several numerical issues specific to the DO method for fluid and ocean flows. In particular we have shown: how to define pseudo-stochastic pressures to reduce the number of matrix inversions for the pressure from $\mathcal{O}(s^2)$ to $\mathcal{O}(s)$; how to treat advection by the stochastic modes using symmetric approximate TVD schemes; how to deal with singular subspace covariances by generalized inversion; and how to maintain orthonormal modes at the numerical level and so account for truncation and round-off errors during time integration. Finally, we evaluated our schemes using a varied set of stochastic flows, hence illustrating robustness but also providing benchmarks for future schemes and implementations.

Using an asymmetric Dirac-stochastic lock-exchange benchmark, we found excellent agreement among multiple realizations from a deterministic code and the realizations generated from a single stochastic simulation. This validated our numerical implementation, and confirmed that the pseudo-stochastic pressure approach works in practice. Using a symmetric version of the lock-exchange benchmark, we showed that the symmetric TVD advection scheme (TVD*) works well, while the CDS scheme suffers from numerical oscillations, and the non-symmetric TVD scheme loses symmetry.

Using a lid-driven cavity flow with uncertain initial conditions, we showed that each component converged near-optimally under both time and space refinement. Additionally, we showed that even when a higher-order time integrator is used for the stochastic coefficients, their accuracy is still limited to approximately first order, as discussed in §3.4. This benchmark also demonstrated an expected decay in the variance of the stochastic coefficients, which were shown to be very non-Gaussian.

Finally, using a stochastic flow past a square cylinder in a confined channel, we showed that the stochastic coefficients converged with increased samples sizes and orders of time integration. We found that using $> 10,000$ samples and a second order time-integrator yielded adequate performance. This benchmark also demonstrated that our discretized DO methodology successfully captures both vortex shedding directions, resulting in a fully symmetric mean field. Also, we noted that using our newly defined pseudo-stochastic pressure, we were 1000% more efficient computationally when using 10 modes.

Possible future studies include the efficient extension of our present framework to flows with uncertain parameters in the governing equations and with stochastic forcing at the boundary and in the interior. Also, our results suggest that simulations with long time

integration will benefit from more accurate time-integration scheme. Then, while our proposed symmetric TVD-based advection and orthonormalization schemes for the stochastic modes were shown to perform adequately, more optimal treatments are possible. Alternative means of discretizing the stochastic coefficients are worth exploring for applications where rare events are important. Finally, specific multi-resolution DO schemes can be derived for multiscale stochastic fluid and ocean flows. Unstructured grids as well as nested approaches [8, 18] will then be useful. The utilization of the present schemes as well as these future advances will allow the prediction of uncertainty in realistic simulations of multi-physics fluid systems, over a wide range of applications, from micro-nano fluid engineering to multiscale ocean and climate studies.

8. Acknowledgements

We are very thankful to the MSEAS group members, in particular Mr. Sondergaard and Mr. Lolla for helpful discussions. We thank Dr. Shirokoff for the many discussions about Projection methods. We are grateful to the Office of Naval Research for support under grants N00014-08-1-1097 (ONR6.1), N00014-09-1-0676 (QPE) and N00014-09-1-0676 (Science of Autonomy – A-MISSION) to the Massachusetts Institute of Technology (MIT). MPU and PFJL also thank the Natural Sciences and Engineering Research Council (NSERC) of Canada for the Postgraduate Scholarship partially supporting the graduate studies and research of MPU at MIT.

Appendix A. DO Navier-Stokes Equations

In this Appendix we present the stochastic DO equations for Boussinesq dynamics, and briefly summarize the required steps (from [50, 49, 52]) to obtain these equations.

Starting with a generalized Karhunen-Loève expansion [32, 49], one decomposes the solution of eqns. (1)-(3) into a Dynamically Orthogonal (DO) field expansion [50, 49, 51] for the velocity, density and pressure² fields

$$\mathbf{\Phi}(\mathbf{x}, t; \omega) = \bar{\mathbf{\Phi}}(\mathbf{x}, t) + \sum_{i=1}^s Y_i(t; \omega) \mathbf{\Phi}_i(\mathbf{x}, t), \quad (\text{A.1})$$

$$\mathbf{\Phi} \equiv \bar{\mathbf{\Phi}} + Y_i \mathbf{\Phi}_i,$$

$$p(\mathbf{x}, t; \omega) = \bar{p}(\mathbf{x}, t) + \sum_{i=1}^s Y_i(t; \omega) p_i(\mathbf{x}, t) + \sum_{i=1}^s \sum_{j=1}^s Y_i(t; \omega) Y_j(t; \omega) p_{ij}(\mathbf{x}, t), \quad (\text{A.2})$$

$$p \equiv \bar{p} + Y_i p_i + Y_i Y_j p_{ij},$$

where $\mathbf{\Phi} = [\Phi^1, \Phi^2, \dots]^T = [\mathbf{u}, \rho]^T$ is the vector of prognostic state variables. The scalar $s = s(t)$ defines the time-dependent dimension of the stochastic subspace, i.e. s is a discrete number of stochastic terms retained from a complete expansion. The field functions $\mathbf{\Phi}_i(\mathbf{x}, t)$ are the s orthonormal deterministic modes and the $Y_i(t; \omega)$ are their s zero-mean stochastic

²For convenience, we changed the sign of p_{ij} from the original definition in the Sapsis and Lermusiaux [50]

coefficients, in general non-Gaussian. In our notation, we use the Einstein summation exclusively for summations related to the stochastic expansion. The decomposition of pressure into a mean, linear modal and quadratic modal component follows from the Pressure Poisson Equations (see §3.2 or Sapsis and Lermusiaux [50]). Since both the modes and stochastic coefficients are functions of time, a redundancy arises, which is resolved by the DO condition,

$$\left\langle \Phi_i, \frac{\partial \Phi_j}{\partial t} \right\rangle_{\mathcal{D}} = 0, \quad \forall i, j \in [1, 2, \dots, s], \quad (\text{A.3})$$

where the inner product is defined as $\langle \mathbf{a}, \mathbf{b} \rangle_{\mathcal{D}} = \int_{\mathcal{D}} \sum_i (a^i b^i) d\mathcal{D}$ for arbitrary vectors of spatial functions $\mathbf{a} = [a^1, a^2, \dots]^T$ and $\mathbf{b} = [b^1, b^2, \dots]^T$.

Using the DO condition, an exact set of equations can be obtained that governs the evolution of the mean, modes, and stochastic coefficients of the generalized Karhunen-Loève expansion. The only approximation arises from the truncation of the DO expansion to $s(t)$ terms. First substitute (A.1) into eqns. (1)-(3) to obtain (using a Langevin notation):

$$\begin{aligned} \frac{\partial \bar{\mathbf{u}}}{\partial t} + \frac{dY_i}{dt} \mathbf{u}_i + Y_i \frac{\partial \mathbf{u}_i}{\partial t} &= \mathcal{L}^{\mathbf{u}} (\bar{\mathbf{u}} + Y_i \mathbf{u}_i, \bar{\rho} + Y_i \rho_i, p, \mathbf{x}, t; \omega), \\ \frac{\partial \bar{\rho}}{\partial t} + \frac{dY_i}{dt} \rho_i + Y_i \frac{\partial \rho_i}{\partial t} &= \mathcal{L}^{\rho} (\bar{\rho} + Y_i \rho_i, \bar{\mathbf{u}} + Y_i \mathbf{u}_i, \mathbf{x}, t; \omega), \end{aligned} \quad (\text{A.4})$$

and DO decomposed versions of (2)-(3). It is from these equations, within which the DO decomposition was inserted, that the equations for the mean, modes and their coefficients are obtained, using the expectation operator, the spatial inner product, and eqns. (A.2)-(A.3).

Mean. To obtain a rate of change for the mean fields, the idea is to eliminate the random components in the left-hand-sides of eqns. (A.4). Hence, taking the expectation of eqns. (A.4) it can be found that the evolution of the mean fields are governed by

$$\begin{aligned} \nabla \cdot \bar{\mathbf{u}} &= 0, \quad \mathbf{x} \in \mathcal{D}, \\ \frac{\partial \bar{\mathbf{u}}}{\partial t} - \frac{1}{\sqrt{Gr}} \nabla^2 \bar{\mathbf{u}} &= -\nabla \cdot (\bar{\mathbf{u}} \bar{\mathbf{u}}) - \nabla \bar{p} + \bar{\rho} \mathbf{e}^g \\ &\quad - \mathbf{C}_{Y_i Y_j} (\nabla p_{ij} + \nabla \cdot (\mathbf{u}_j \mathbf{u}_i)), \quad \mathbf{x} \in \mathcal{D}, \\ \frac{\partial \bar{\rho}}{\partial t} - \frac{1}{Sc \sqrt{Gr}} \nabla^2 \bar{\rho} &= -\nabla \cdot (\bar{\mathbf{u}} \bar{\rho}) - \mathbf{C}_{Y_i Y_j} \nabla \cdot (\mathbf{u}_j \rho_i), \quad \mathbf{x} \in \mathcal{D}, \end{aligned} \quad (\text{A.5})$$

with the initial and boundary conditions given by

$$\begin{aligned} \bar{\mathbf{u}}(\mathbf{x}, 0) &= \bar{\mathbf{u}}_0, \quad \mathbf{x} \in \mathcal{D}, \\ \bar{\rho}(\mathbf{x}, 0) &= \bar{\rho}_0, \quad \mathbf{x} \in \mathcal{D}, \end{aligned} \quad (\text{A.6})$$

$$\begin{aligned} \bar{\mathbf{u}} &= \bar{\mathbf{g}}_D, \quad \mathbf{x} \in \partial \mathcal{D}_D, \\ \frac{\partial \bar{\mathbf{u}}}{\partial n} &= \bar{\mathbf{g}}_N, \quad \mathbf{x} \in \partial \mathcal{D}_N, \\ \bar{\rho} &= \bar{g}_{D_\rho}, \quad \mathbf{x} \in \partial \mathcal{D}_{D_\rho}, \\ \frac{\partial \bar{\rho}}{\partial n} &= \bar{g}_{N_\rho}, \quad \mathbf{x} \in \partial \mathcal{D}_{N_\rho}, \end{aligned} \quad (\text{A.7})$$

where $\mathbf{C}_{Y_i Y_j} = E^\omega [Y_i Y_j]$ is an element of the covariance matrix in the stochastic/error subspace. Deterministic initial and boundary conditions (\bullet quantities) are assigned to the mean. Note that in general the vector $\nabla \cdot (\mathbf{u}_j \mathbf{u}_i)$ differs from $\nabla \cdot (\mathbf{u}_i \mathbf{u}_j)$ (recall that $\nabla \cdot (\mathbf{u}_j \mathbf{u}_i)_q = \frac{\partial \mathbf{u}_{j_r} \mathbf{u}_{i_q}}{\partial x_r} \neq \frac{\partial \mathbf{u}_{i_r} \mathbf{u}_{j_q}}{\partial x_r}$; e.g., $\frac{\partial v_j u_i}{\partial y} \neq \frac{\partial v_i u_j}{\partial y}$).

Modes. The evolution of the modes is also obtained from eqns. (A.4). To do so, the idea is to eliminate the random coefficients in front of the time derivatives of the modes. The essential steps are to multiply these equations with a stochastic coefficient Y_j , apply the expectation operator, and substitute an expression for $E^\omega [\frac{\partial Y_i}{\partial t} Y_j]$, which is obtained by projecting the equation unto Φ_k and imposing the DO condition. From this the following governing evolution equations for the modes can be found:

$$\begin{aligned} \nabla \cdot \mathbf{u}_i &= 0, \quad \mathbf{x} \in \mathcal{D}, \\ \frac{\partial \mathbf{u}_i}{\partial t} &= \mathbf{Q}_i^{\mathbf{u}} - \langle \mathbf{Q}_i, \Phi_j \rangle_{\mathcal{D}} \mathbf{u}_j, \quad \mathbf{x} \in \mathcal{D}, \\ \frac{\partial \rho_i}{\partial t} &= \mathbf{Q}_i^{\rho} - \langle \mathbf{Q}_i, \Phi_j \rangle_{\mathcal{D}} \rho_j, \quad \mathbf{x} \in \mathcal{D}, \end{aligned} \tag{A.8}$$

where

$$\begin{aligned} \mathbf{Q}_i &= [\mathbf{Q}_i^{\mathbf{u}}, \mathbf{Q}_i^{\rho}]^T = \left[\mathbf{C}_{Y_i Y_j}^{-1} E^\omega [\mathcal{L}^{\mathbf{u}} Y_j], \mathbf{C}_{Y_i Y_j}^{-1} E^\omega [\mathcal{L}^{\rho} Y_j] \right]^T, \\ \mathbf{Q}_i^{\mathbf{u}} &= \frac{1}{\sqrt{Gr}} \nabla^2 \mathbf{u}_i - \nabla \cdot (\mathbf{u}_i \bar{\mathbf{u}}) - \nabla \cdot (\bar{\mathbf{u}} \mathbf{u}_i) - \nabla p_i + \rho_i \mathbf{e}^g \\ &\quad - \mathbf{C}_{Y_i Y_j}^{-1} \mathbf{M}_{Y_j Y_m Y_n} (\nabla p_{mm} + \nabla \cdot (\mathbf{u}_n \mathbf{u}_m)), \\ \mathbf{Q}_i^{\rho} &= \frac{1}{Sc \sqrt{Gr}} \nabla^2 \rho_i - \nabla \cdot (\mathbf{u}_i \bar{\rho}) - \nabla \cdot (\bar{\mathbf{u}} \rho_i) \\ &\quad - \mathbf{C}_{Y_i Y_j}^{-1} \mathbf{M}_{Y_j Y_m Y_n} (\nabla \cdot (\mathbf{u}_n \rho_m)), \end{aligned}$$

and $\mathbf{M}_{Y_j Y_m Y_n} = E^\omega [Y_j Y_m Y_n]$ is a third moment. The right-hand-sides in eqns. (A.8) correspond the total rate of change of the subspace (without a DO condition) minus the projection of this rate of change on the subspace itself (which is subtracted to ensure the DO condition). We note that in general $\nabla \cdot (\bar{\mathbf{u}} \mathbf{u}_i) \neq \nabla \cdot (\mathbf{u}_i \bar{\mathbf{u}})$ since $\frac{\partial \bar{v} u_i}{\partial y} \neq \frac{\partial v_i \bar{u}}{\partial y}$ for example.

The initial and boundary conditions for the modes are obtained from those of the full stochastic fields, eqns. (2) and (3), but reduced to their dominant initial error subspace of size s_0 ,

$$\begin{aligned} \mathbf{u}_i(\mathbf{x}, 0) &= \mathbf{u}_{i,0}(\mathbf{x}), \quad \mathbf{x} \in \mathcal{D}, \\ \rho_i(\mathbf{x}, 0) &= \rho_{i,0}(\mathbf{x}), \quad \mathbf{x} \in \mathcal{D}, \end{aligned} \tag{A.9}$$

and

$$\begin{aligned}
\mathbf{u}_i &= \mathbf{g}_{i,D}, & \mathbf{x} &\in \partial\mathcal{D}_D, \\
\frac{\partial \mathbf{u}_i}{\partial n} &= \mathbf{g}_{i,N}, & \mathbf{x} &\in \partial\mathcal{D}_N, \\
\rho_i &= g_{i,D_\rho}, & \mathbf{x} &\in \partial\mathcal{D}_{D_\rho}, \\
\frac{\partial \rho_i}{\partial n} &= g_{i,N_\rho}, & \mathbf{x} &\in \partial\mathcal{D}_{N_\rho}.
\end{aligned} \tag{A.10}$$

Coefficients. Finally, to obtain the evolution of the stochastic coefficients from eqns. (A.4), the idea is to eliminate the modes in the term containing the time derivatives of the random coefficients. To do so, project the evolution eqns. (A.4) onto each mode i , apply the DO conditions, and essentially impose that each coefficient is of zero mean. The resulting governing stochastic ODEs are

$$\begin{aligned}
\frac{dY_i}{dt} &= \langle \mathcal{L} - E^\omega [\mathcal{L}], \Phi_i \rangle_{\mathcal{D}}, \\
&= \langle \mathbf{F}_m, \Phi_i \rangle_{\mathcal{D}} Y_m - \langle \nabla p_m + \nabla \cdot (\mathbf{u}_m \mathbf{u}_m), \mathbf{u}_i \rangle_{\mathcal{D}} (Y_m Y_n - \mathbf{C}_{Y_m Y_n}) \\
&\quad - \langle \nabla \cdot (\mathbf{u}_m \rho_m), \rho_i \rangle_{\mathcal{D}} (Y_m Y_n - \mathbf{C}_{Y_m Y_n}),
\end{aligned} \tag{A.11}$$

where $\mathcal{L} = [\mathcal{L}^u, \mathcal{L}^\rho]^T$ and $\mathbf{F}_m = [\mathbf{F}_m^u, \mathbf{F}_m^\rho]^T$ with

$$\begin{aligned}
\mathbf{F}_m^u &= \frac{1}{\sqrt{Gr}} \nabla^2 \mathbf{u}_m - \nabla \cdot (\mathbf{u}_m \bar{\mathbf{u}}) - \nabla \cdot (\bar{\mathbf{u}} \mathbf{u}_m) - \nabla p_m + \rho_m \mathbf{e}^g, \\
\mathbf{F}_m^\rho &= \frac{1}{Sc\sqrt{Gr}} \nabla^2 \rho_m - \nabla \cdot (\mathbf{u}_m \bar{\rho}) - \nabla \cdot (\bar{\rho} \mathbf{u}_m),
\end{aligned}$$

The initial conditions for the coefficients are obtained from those of the full stochastic fields, eqn. (2), by projection onto each initial mode i , $\Phi_{i,0}$, and removal of the mean. This leads to:

$$\begin{aligned}
Y_i(t_0; \omega) &= \langle \Phi_0 - \bar{\Phi}_0, \Phi_{i,0} \rangle_{\mathcal{D}}, \\
&= \langle \mathbf{u}_0 - \bar{\mathbf{u}}_0, \mathbf{u}_{i,0} \rangle_{\mathcal{D}} + \langle \rho_0 - \bar{\rho}_0, \rho_{i,0} \rangle_{\mathcal{D}}, \quad \omega \in \Omega.
\end{aligned} \tag{A.12}$$

In this Appendix we stated the DO equations for the mean, modes, and stochastic coefficients of the stochastic Boussinesq equations.

Appendix B. Numerical orthonormalization procedure

In this Appendix, we present our newly developed, numerically efficient, orthonormalization procedure based on the discussion in §5.2.

Consider the pictorial analogy showing different strategies when orthonormalizing the basis (Fig. [B.17]). In case (a), the argument is that the coefficients were evolved correctly, and all the error lies in the basis evolution. This is not true because we evolve the coefficients by taking an inner product of the equations using the bases at the old time-step. In case (b), the argument is that the energy and ‘direction’ of the realizations were correctly evolved. When




Figure B.17: There are different ways to treat the stochastic coefficients when orthonormalizing the modes. In this pictorial analogy, the standard 2D cartesian bases and coordinates represent the modes and stochastic coefficients, respectively.

correcting the bases, the direction and energy are maintained, but the actual realization is modified. In case (c), the argument is that the realizations were evolved correctly. However, it is unlikely that the errors in the modes and coefficients cancel during the time evolution. Also, while not evident from the pictorial analogy, the energy of the coefficients change in case (c). This is because the length of the basis changes when orthonormalized, and this length change is then accounted for in the coefficients in order to maintain the coordinate of the realization. The case (b) strategy is a compromise between cases (a) and (c), and it results in the lowest error for the lock-exchange cases in §6.1. The errors were the largest for case (a), and the errors for case (c) were marginally larger compared to case (b).

To be clear, case (b) is shown to perform best for the specific cases in §6.1, and these results are not proven to be universally applicable. Also, we made no attempt to provide an exhaustive list of all the orthonormalization strategies. For example, arguing that the coefficients still live in the bases of the old timestep, one can consider projecting the coefficients from this old basis onto the new basis before orthonormalization. Therefore, the issue of the optimal orthonormalization strategy is left open to future research. Nonetheless, we believe our current approach is sufficient for most applications since no strategy will influence the scheme's overall order of accuracy. Any improved strategy will only reduce the size of the error, which will be of order Δt smaller than the truncation error. Therefore, any accuracy issues due to orthonormalization can be resolved by reducing the size of the timestep when using the current scheme. Our orthonormalization procedure for case (b) is described next.

For case (b), we want the new orthonormalized basis to: approximately recover the specific realizations before the orthonormalization; and exactly maintain the stochastic energy. Mathematically we, ideally, want both:

$$\Phi_{h,i}(Y_{r,i})^T \approx \Phi_{h,i}^O(Y_{r,i}^O)^T, \quad (\text{B.1})$$

$$\text{Tr}(\mathbf{C}_{Y_{r,i}Y_{r,j}}^{k\star}) = \text{Tr}(\mathbf{C}_{Y_{r,i}Y_{r,j}}^{k\star,O}), \quad (\text{B.2})$$

where the orthonormalized basis is defined such that $\langle \Phi_{h,i}^O, \Phi_{h,j}^O \rangle_{\mathcal{D}} = \delta_{ij}$, and the trace operator Tr simply sums the diagonal entries of a matrix. In order to preserve the second

property, we first rotate the stochastic system such that the covariance matrix is diagonal:

$$\mathbf{C}_{Y_i Y_j}^k = \mathbb{E}[Y_{r,i} Y_{r,j}] = \mathbf{V}_{DC} \mathbf{D}_{DC} \mathbf{V}_{DC}^T, \quad (\text{B.3})$$

$$Y_{r,i}^{DC} = Y_{r,i} \mathbf{V}_{DC}, \quad (\text{B.4})$$

$$\Phi_{h,i}^{DC} = \Phi_{h,i} \mathbf{V}_{DC}, \quad (\text{B.5})$$

which gives the decorrelated (or DC) stochastic coefficients. By starting with a decorrelated system, we have a unique and convenient reference frame which we can use to enforce the $\text{Tr}(\mathbf{C}_{Y_i Y_j}^{k\star})$ property.

Next we perform the orthonormalization as follows

$$\mathbf{M}_{ij} = \langle \Phi_{h,i}^{DC}, \Phi_{h,j}^{DC} \rangle_{\mathcal{D}} = \mathbf{V}_M \mathbf{D}_M \mathbf{V}_M^T, \quad (\text{B.6})$$

$$Y_{r,i}^{OC} = Y_{r,i}^{DC} \mathbf{V}_M \sqrt{\mathbf{D}_M}, \quad (\text{B.7})$$

$$\Phi_{h,i}^{OC} = \Phi_{h,i}^{DC} \mathbf{V}_M \sqrt{\mathbf{D}_M^{-1}}. \quad (\text{B.8})$$

The above could be accomplished by directly calculating the Singular Value Decomposition of the $\Phi_{h,i}^{DC}$ matrix, but the above procedure has the advantage of enabling the user to specify any desired (possibly non-linear) inner product. In either case, the two approaches are similar in efficiency and give the same results. To verify that (B.1) is satisfied, we substitute the above expression in (B.1)

$$\begin{aligned} \Phi_{h,i}^{DC} (Y_{r,i}^{DC})^T &= \Phi_{h,i}^{OC} (Y_{r,i}^{OC})^T, \\ &= \Phi_{h,i}^{DC} \mathbf{V}_M \sqrt{\mathbf{D}_M^{-1}} \left(Y_{r,i}^{DC} \mathbf{V}_M \sqrt{\mathbf{D}_M} \right)^T, \\ &= \Phi_{h,i}^{DC} \mathbf{V}_M \sqrt{\mathbf{D}_M^{-1}} \sqrt{\mathbf{D}_M} \mathbf{V}_M^T (Y_{r,i}^{DC})^T, \\ &= \Phi_{h,i}^{DC} (Y_{r,i}^{DC})^T, \end{aligned}$$

since the eigen-vectors are orthonormal $\mathbf{V}_M \mathbf{V}_M^T = \mathbf{I}$. Therefore the first property, (B.1), is exactly satisfied at this stage. Stopping here would give the case (c) strategy.

Finally, we decorrelate the samples again, while ensuring that the stochastic energy is preserved:

$$\mathbf{C}_{Y_i Y_j}^{k,OC} = \mathbb{E}[Y_{r,i}^{OC} Y_{r,j}^{OC}] = \mathbf{V}_O \mathbf{D}_O \mathbf{V}_O^T, \quad (\text{B.9})$$

$$Y_{r,i}^O = Y_{r,i}^{OC} \mathbf{V}_O \sqrt{\frac{\text{Tr}(\mathbf{D}_{DC})}{\text{Tr}(\mathbf{D}_O)}}, \quad (\text{B.10})$$

$$\Phi_{h,i}^O = \Phi_{h,i}^{OC} \mathbf{V}_O. \quad (\text{B.11})$$

Note that $\text{Tr}(\mathbf{D}_O) = \text{Tr}(\mathbf{D}_{DC} \mathbf{D}_M)$, but that $\mathbf{D}_O \neq \mathbf{D}_{DC} \mathbf{D}_M$. This ensures that (B.2) is maintained, but now we can only approximately satisfy (B.1). The over-all correction is

then as follows:

$$Y_{r,i}^O = Y_{r,i} \mathbf{V}_{DC} \mathbf{V}_M \sqrt{\mathbf{D}_M} \mathbf{V}_O \sqrt{\frac{\text{Tr}(\mathbf{D}_{DC})}{\text{Tr}(\mathbf{D}_O)}}, \quad (\text{B.12})$$

$$\Phi_{h,i}^O = \Phi_{h,i} \mathbf{V}_{DC} \mathbf{V}_M \sqrt{\mathbf{D}_M^{-1}} \mathbf{V}_O. \quad (\text{B.13})$$

This procedure is robust even if $\mathbf{C}_{YY}^{k,OC}$ is singular, however it does require a non-singular \mathbf{M}_{ij} which will normally be the case if modes are properly initialized. Also, the number of floating point operations is dominated by the calculation of \mathbf{M}_{ij} if the number of spatial points exceeds the number of realizations. Nonetheless, the overall cost of the orthonormalization is inexpensive compared to the inversion of the pressure Poisson equations.

References

- [1] Bennett AF (1992) Inverse Methods in Physical Oceanography., 1st edn. Cambridge University Press, New York
- [2] Cameron RH, Martin WT (1947) The orthogonal development of nonlinear functionals in series of Fourier-Hermite functionals. *Ann of Math* 48:385–392
- [3] Carpenter MH, Kennedy C (1994) Fourth-order 2N-storage Runge-Kutta schemes. NASA Report TM 109112, NASA Langley Research Center, Hampton, VA
- [4] Chapra SC, Canale RP (2006) Numerical Methods for Engineers, 5th edn. McGraw-Hill Higher Education, Boston, MA
- [5] Chen D, Liu J (2000) Mixture Kalman filters. *J Roy Statist Soc Ser A* 62:493–508
- [6] Cushman-Roisin B, Beckers JM (2010) Introduction to geophysical fluid dynamics: Physical and Numerical Aspects. Academic Press.
- [7] Deb MK, Babuška I, Oden J (2001) Solution of stochastic partial differential equations using Galerkin finite element techniques. *Comput Methods Appl Mech Eng* 190:6359–6372
- [8] Deleersnijder E, Lermusiaux PFJ (2008) Multi-scale modeling: nested grid and unstructured grid approaches. *Ocean Dynamics* 58:335–336
- [9] Doucet A, de Freitas N, Gordon N (2001) Sequential Monte-Carlo Methods in Practice. Springer-Verlag
- [10] Eldred MS, Webster CG, Constantine P (2008) Evaluation of Non-Intrusive Approaches for Wiener-Askey Generalized Polynomial Chaos. In: Proceedings of the 49th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference (10th AIAA Non-Deterministic Approaches Conference), paper AIAA-2008-1892, Schaumburg, IL
- [11] Erturk E (2009) Discussions on driven cavity flow. *Int J Numer Meth Fluids* 60:275–294
- [12] Erturk E, Corke TC, Gökçöl C (2005) Numerical Solutions of 2-D Steady Incompressible Driven Cavity Flow at High Reynolds Numbers. *Int J Numer Meth Fluids* 48:747–774
- [13] Ferziger JH, Peric M (2002) Computational Methods for Fluid Dynamics, 3rd edn. Springer, New York, NY
- [14] Foo J, Karniadakis GE (2010) Multi-element probabilistic collocation method in high dimensions. *J Comput Phys* 229:1536–1557
- [15] Frochte J, Heinrichs W (2009) A splitting technique of higher order for the Navier-Stokes equations. *J Comput Appl Math* 228:373–390
- [16] Ghanem R, Spanos P (1991) Stochastic Finite Elements: A Spectral Approach. Springer-Verlag
- [17] Guermond J, Mineev P, Shen J (2006) An overview of projection methods for incompressible flows. *Comput Methods Appl Mech Engrg* 195:6011–6045
- [18] Haley PJ Jr, Lermusiaux PFJ (2010) Multiscale two-way embedding schemes for free-surface primitive-equations in the Multidisciplinary Simulation, Estimation and Assimilation System (MSEAS). *Ocean Dynamics* 60:1497–1537
- [19] Härtel C, Meiburg E, Necker F (2000) Analysis and direct numerical simulation of the flow at a gravity-current head. Part 1. Flow topology and front speed for slip and no-slip boundaries. *J Fluid Mech* 418:189–212
- [20] Higham DJ (2001) An Algorithmic Introduction to Numerical Simulation of Stochastic Differential Equations. *SIAM Review* 43(3):525–546
- [21] Hosder S, Walters R, Perez R (2006) A nonintrusive polynomial chaos method for uncertainty propagation in CFD simulations. In: AIAA Aerosp. Sci. Meet. Exhib, 44th, AIAA-2006-0891, Reno, Nevada

- [22] Jakeman JD, Roberts SG (2008) Stochastic Galerkin and collocation methods for quantifying uncertainty in differential equations: a review. *Aust NZ Ind Appl Math J* 50:C815–C830
- [23] Karhunen K (1946) Zur spektraltheorie stochastischer prozesse. *Ann Acad Sci Fennicae* 34:17
- [24] Kloeden PE, Platen E (1999) *Numerical Solution of Stochastic Differential Equations*. Springer-Verlag, Berlin
- [25] Le Maître O, Reagan M, Najm H, Ghanem R, Knio O (2002) A stochastic projection method for fluid flow II. Random process. *J Comput Phys* 181:9–44
- [26] Le Maître O, Najm H, Ghanem R, Knio O (2004) Multi-resolution analysis of Wiener-type uncertainty propagation schemes. *J Comput Phys* 197:502–531
- [27] Le Maître OP, Knio OM (2010) *Spectral Methods for Uncertainty Quantification*, 1st edn. Scientific Computation, Springer, New York
- [28] Le Maître OP, Knio OM, Najm HN, Ghanem RG (2001) A Stochastic Projection Method for Fluid Flow I. Basic Formulation. *J Comput Phys* 173:481–511
- [29] Lermusiaux PFJ (1999) Data assimilation via error subspace statistical estimation. Part II: Middle atlantic bight shelfbreak front simulations and ESSE validation. *Mon Weather Rev* 127:1408–1432
- [30] Lermusiaux PFJ (2001) Evolving the subspace of the three-dimensional multiscale ocean variability: Massachusetts Bay. *J Mar Syst* 29:385–422
- [31] Lermusiaux PFJ (2002) On the mapping of multivariate geophysical fields: Sensitivities to size, scales, and dynamics. *J Atmos Ocean Technol* 19(10):1602–1637
- [32] Lermusiaux PFJ (2006) Uncertainty estimation and prediction for interdisciplinary ocean dynamics. *J Comput Phys* 217:176–199
- [33] Lermusiaux PFJ (2007) Adaptive Modeling, Adaptive Data Assimilation and Adaptive Sampling. *Physica D* 230:172–196, Special issue on “Mathematical Issues and Challenges in Data Assimilation for Geophysical Systems: Interdisciplinary Perspectives.” C.K.R.T. Jones and K. Ide, Eds.
- [34] Lermusiaux PFJ (2009) 2.29: Numerical Fluid Mechanics. *Lectures Notes*, MIT, Cambridge, MA
- [35] Lermusiaux PFJ, Robinson AR (1999) Data assimilation via error subspace statistical estimation. Part I: Theory and schemes. *Mon Weather Rev* 127:1385–1407
- [36] Lermusiaux PFJ, Chiu CS, Gawarkiewicz G, Abbot P, Robinson A, Miller R, Haley P, Leslie W, Majumdar S, Pang A, Lekien F (2006) Quantifying uncertainties in ocean predictions. *Oceanography* 19:92–105
- [37] Leveque RJ (2002) *Finite Volume Methods for Hyperbolic Problems*. Cambridge University Press
- [38] Li R, Ghanem R (1998) Adaptive polynomial chaos expansions applied to statistics of extremes in nonlinear random vibration. *Prob Eng Mech* 13:125–136
- [39] Loève M (1977) *Probability Theory*. Springer-Verlag, Berlin/New York
- [40] Logutov OG, Lermusiaux PFJ (2008) Inverse Barotropic Tidal Estimation for Regional Ocean Applications . *Ocean Modelling* 25:17–34
- [41] Lucor D, Meyers J, Sagaut P (2007) Sensitivity analysis of large-eddy simulations to subgrid-scale-model parametric uncertainty using polynomial chaos. *J Fluid Mech* 585:255–279
- [42] Marzouk YM, Najm HN (2009) Dimensionality reduction and polynomial chaos acceleration of Bayesian inference in inverse problems. *J Comput Phys* 228(6):1862–1902
- [43] Mathelin L, Hussaini MY, Zang TA (2005) Stochastic approaches to uncertainty quantification in CFD simulations. *Numer Algor* 38(1-3):209–236
- [44] McLachlan G, Peel D (2000) *Finite mixture models*. Wiley series in probability and statistics: Applied probability and statistics, Wiley
- [45] van der Merwe R, Wan E (2003) Gaussian mixture sigma-point particle filters for sequential probabilistic inference in dynamic state-space models. In: *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing, ICASSP, Hong-Kong*
- [46] Miller R (2007) Topics in data assimilation: Stochastic processes. *Physica D* 230(1-2):17 – 26
- [47] Najm HN (2009) Uncertainty Quantification and Polynomial Chaos Techniques in Computational Fluid Dynamics. *Annu Rev Fluid Mech* 41:35–52
- [48] Nouy A (2007) A generalized spectral decomposition technique to solve a class of linear stochastic partial differential equations. *Comput Methods Appl Mech Engrg* 196:4521–4537
- [49] Sapsis TP (2010) Dynamically orthogonal field equations for stochastic fluid flows and particle dynamics. Ph.D.Thesis, Massachusetts Institute of Technology, Dept. of Mech. Eng.
- [50] Sapsis TP, Lermusiaux PFJ (2009) Dynamically orthogonal field equations for continuous stochastic dynamical systems.

- [51] Sapsis TP, Lermusiaux PFJ (2011) Dynamical criteria for the evolution of the stochastic dimensionality in flows with uncertainty. *Physica D* Under Review
- [52] Sapsis TP, Lermusiaux PFJ, Ueckermann MP (2011) Statistical model of Navier-Stokes equations using dynamical orthogonality. *J Fluid Mech* In preparation.
- [53] Shirokoff D, Rosales R (2011) An efficient method for the incompressible Navier-Stokes equations on irregular domains with no-slip boundary conditions, high order up to the boundary. *J Comput Phys* 230(23):8619 – 8646
- [54] Simon F, Guillen P, Sagaut P, Lucor D (2010) A gPC-based approach to uncertain transonic aerodynamics. *Comput Methods Appl Mech Engrg* 199(17-20):1091 – 1099
- [55] Sondergaard T (2011) Data assimilation with gaussian mixture models using the dynamically orthogonal field equations. S.M.Thesis, Massachusetts Institute of Technology, Dept. of Mech. Eng.
- [56] Sondergaard T, Lermusiaux PFJ (2011) Data Assimilation with Gaussian Mixture Models using the Dynamically Orthogonal Field Equations. Part I: Theory and Scheme. Submitted to Mon Weather Rev
- [57] Sondergaard T, Lermusiaux PFJ (2011) Data Assimilation with Gaussian Mixture Models using the Dynamically Orthogonal Field Equations. Part II: Applications. Submitted to Mon Weather Rev
- [58] Timmermans L, Mineev P, Van De Vosse F (1996) An approximate projection scheme for incompressible flow using spectral elements. *Int J Numer Methods Fluids* 22:673–688
- [59] Ueckermann MP, Lermusiaux PFJL (In Prep. 2011) 2.29 Finite Volume MATLAB Framework Documentation. Reports in Ocean Science and Engineering 14, Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts
- [60] Van Leer B (1977) Towards the ultimate conservative difference scheme. IV. A new approach to numerical convection. *J Comput Phys* 23(3):276 – 299
- [61] Wan X, Karniadakis GE (2005) An adaptive multi-element generalized polynomial chaos method for stochastic differential equations. *J Comput Phys* 209:617–642
- [62] Wan X, Karniadakis GE (2006) Long-term behavior of polynomial chaos in stochastic flow simulations. *Comput Methods Appl Mech Engrg* 195:5582–5596
- [63] Webster M, Sokolov A (2000) A methodology for quantifying uncertainty in climate projections. *Clim Change* 46:417–446
- [64] Wiener N (1938) The homogeneous chaos. *Amer J Math* 60:897–936
- [65] Wiener N (1958) *Nonlinear Problems in Random Theory*. MIT Technology Press and John Wiley and Sons
- [66] Xiu D (2009) Fast numerical methods for stochastic computations: A review. *Commun Comput Phys* 5(2-4):242–272
- [67] Xiu D (2010) *Numerical Methods for Stochastic Computations: A Spectral Method Approach*. Princeton University Press
- [68] Xiu D, Karniadakis GE (2002) The Wiener-Askey polynomial chaos for stochastic differential equations. *SIAM J Sci Comput* 24:619–644
- [69] Xiu D, Lucor D, Su CH, Karniadakis GE (2002) Stochastic Modeling of Flow-Structure Interactions Using Generalized Polynomial Chaos. *J Fluids Eng* 124(1):51–60