Journal of Computational Physics 233 (2013) 272–294

Contents lists available at SciVerse ScienceDirect

Journal of Computational Physics

journal homepage: www.elsevier.com/locate/jcp

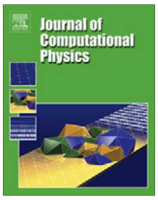

Numerical schemes for dynamically orthogonal equations of stochastic fluid and ocean flows

M.P. Ueckermann, P.F.J. Lermusiaux *, T.P. Sapsis [1]

Department of Mechanical Engineering, Massachusetts Institute of Technology, 77 Mass. Avenue, Cambridge, MA 02139, United States

ARTICLE INFO

Article history:
Received 3 November 2011
Received in revised form 25 August 2012
Accepted 27 August 2012
Available online 21 September 2012

Keywords:
Uncertainty quantification
Navier–Stokes
Boussinesq
Dynamical orthogonality
Projection methods
Total variation diminishing
Error subspace statistical estimation
Ocean modeling
Data assimilation

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. The dynamically orthogonal (DO) field equations provide an adaptive methodology to predict the probability density functions of such flows. 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 DO modes, 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 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. To verify 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.

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 [35,45], particularly when dealing with longer-term unsteady nonlinear dynamics. Fortunately, the recently developed dynamically orthogonal (DO) field equations [49,48,50] provide efficient, tractable equations for uncertainty

---

* Corresponding author. Tel.: +1 617 324 5172.
E-mail addresses: mpuecker@mit.edu (M.P. Ueckermann), pierrel@mit.edu (P.F.J. Lermusiaux), sapsis@mit.edu, sapsis@cims.nyu.edu (T.P. Sapsis).

[1] New York University, Courant Institute of Mathematical Sciences, United States.

http://dx.doi.org/10.1016/j.jcp.2012.08.041

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.

Stochastic modeling approaches can be categorized as either non-intrusive or intrusive (for reviews, we refer to e.g. [16,23,9,42,10,21,46,66,67,26,59]). 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 [5,44]. Nonetheless, a large number of function evaluations are needed due to the slow convergence, which can limit accuracy in large-scale applications.

The polynomial chaos expansion (PCE), pioneered by [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. [63,24,20,10]) and intrusive (e.g. [7,25,61,41,46]) versions have been employed, but both can suffer from the curse of dimensionality (see [59] for a brief cost scaling analysis). This has prompted the use of non-Gaussian random variables [47], the development of generalized PCE [68] to speed up convergence in the polynomial degree, and the development of adaptive schemes that only evaluate the necessary terms in the PCE [37].

PCEs have been successful in many CFD applications. In the case of unsteady incompressible fluid dynamics, [27] 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. [62] showed that multi-element generalized PCEs can significantly improve the accuracy for long time integration in the case of 2D noisy flows past a circular cylinder. Other applications include fluid–structure interactions, (e.g. [69]), turbulence (e.g. [40, e.g.]), and aerodynamics (e.g. [53]).

Motivated by the multi-scale, intermittent and non-homogeneous uncertain ocean fields, the error subspace statistical estimation (ESSE) method was developed [34,28,29]. It uses a generalized Karhunen–Loève (KL) expansion [22,38] with time varying and adaptive basis functions that is initialized by a multi-scale scheme [30] and evolved using a stochastic, data-assimilative and adaptive ensemble. The computational cost of ESSE is outlined in [59].

The DO equations for dynamically evolving stochastic fields [49,48] 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. The DO computational scaling is only dependent on the number of random variables, s: specifically, the storage scales as $\mathcal{O}(s)$ and computational cost as $\mathcal{O}(s^2)$ for Navier–Stokes equations. The size s in general adapts to the dynamically evolving uncertainties and boundary conditions [50]. 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 [48,51]. However, it has not yet been applied to Boussinesq flows, and efficient numerical schemes for the DO decomposition of such flows have not been yet obtained nor evaluated. This explains the need for the present study.

In what follows, the incompressible stochastic Navier–Stokes and Boussinesq equations are given (Section 2). In Section 3, the discretization in time is developed, discussing explicit and implicit schemes. New "pseudo-stochastic pressures" are defined. With this definition, the cost of DO integration schemes is substantially reduced: instead of scaling as $\mathcal{O}(s^2)$ [50], it scales as $\mathcal{O}(s)$ (as long as solving the pressure Poisson equations 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. The discretizations of the physical space and stochastic subspace are given in Section 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 Section 5, the questions of how to deal with singular covariances and how to maintain orthonormal modes numerically are discussed. For the applications in Section 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 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 test robustness. Lastly, conclusions and discussions are in Section 7.

2. Stochastic dynamically orthogonal boussinesq equations

The deterministic components of the partial differential equations (PDEs) that we solve on a domain $\mathcal{D}$ are non-dimensional Boussinesq equations,[2] in the same form as in [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: $\mathbf{u}(\mathbf{x}, \mathbf{t}) = [\mathbf{u}, \mathbf{v}, \mathbf{w}]$, 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, $Sc = \hat{\nu}/\widehat{K}$ is the Schmidt number which is the ratio of kinematic viscosity $\hat{\nu}$ to molecular diffusivity $\widehat{K}$ 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 Eqs. (1) for velocity and density by $\mathcal{L}^{\mathbf{u}}$ and $\mathcal{L}^{\rho}$, 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 Eq. (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 $\mathcal{L}^{\rho}$. 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 Eqs. (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 [48,49] and a summary is provided in A. In short, the DO methodology begins with a generalized Karhunen–Loève expansion truncated to $s(t)$ terms [50]. The vector of prognostic state variables $\mathbf{\Phi}(\mathbf{x}, t; \omega) = [\mathbf{u}, \rho]^T$ is decomposed into the sum of a deterministic mean component $\bar{\mathbf{\Phi}}(\mathbf{x}, t)$, with s deterministic modes $\mathbf{\Phi}_i(\mathbf{x}, t)$, each mode multiplied by a stochastic coefficient $Y_i(t; \omega)$. This decomposition is first substituted into Eqs. (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 $\mathbf{a} = [a^1, a^2, \ldots]^T$ and $\mathbf{b} = [b^1, b^2, \ldots]^T$. In general, we note that this definition of the inner product assumes that the different components of the state vector have been properly normalized [32,51]. This is not guaranteed from the simple deterministic non-dimensionalization used in Eq. (1). In fact, an additional stochastic normalization is usually needed, reflecting the stochastic initial and boundary conditions. After some manipulation (see A and [49,48]) 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.

---

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

3. Semi-implicit time discretization

Solving the deterministic version of the system of Eq. (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 efficient solution schemes. 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 are summarized at the end. The spatial discretizations of physical space and of the stochastic subspace are given in Section 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 $\mathbf{C}_{Y_iY_j}$ and $\mathbf{M}_{Y_jY_mY_n}$ 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

We now first overview 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 define new pseudo-stochastic pressures that reduce the expense to $s+1$ while still providing the valid solution.

The direct explicit handling of pressure was used in [49]. 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 is the stochastic PPEs:

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

Hence, to recover the full stochastic pressure, $1+s+s^2$ Poisson equations need to be inverted, which is expensive and can often dominate the cost of the scheme. Oceanic applications expected to require $s \sim \mathcal{O}(10^2 - 10^3)$ [32] 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 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 Eqs. (A.5) and (A.8), we therefore define new pseudo-stochastic pressures, which are a combination of the mean, linear-, and quadratic-modal pressures:

$$\begin{aligned} \breve{\bar{p}} &= \bar{p} + \mathbf{C}_{Y_iY_j} p_{ij}, \\ \breve{p}_i &= p_i + \mathbf{C}^{-1}_{Y_iY_j} \mathbf{M}_{Y_jY_mY_n} p_{mn}. \end{aligned} \tag{5}$$

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),

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

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 $\langle \nabla p_{mn}, \mathbf{u}_i \rangle_{\mathcal{D}}$ is zero (see [51]). 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 [15,52]. Here, we chose to use the "incremental pressure-correction scheme in rotational form" [57], which has a proven temporal accuracy [17]. The classic versions of the scheme for our deterministic equations are outlined in [59], next we adapt them for the mean and modes.

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 (Section 3.1). Starting from the PDEs (A.5) for the mean, we obtain:

$$\frac{\tilde{\bar{\mathbf{u}}}^k}{\Delta t} - \frac{1}{\sqrt{Gr}} \nabla^2 \tilde{\bar{\mathbf{u}}}^k = \frac{\bar{\mathbf{u}}^{k-1}}{\Delta t} - \{\nabla \cdot (\bar{\mathbf{u}}\bar{\mathbf{u}})\}^{k^\star} - \nabla \breve{\bar{p}}^{k^\star} + \bar{\rho}^{k^\star} \mathbf{e}^g - \mathbf{C}^{k^\star}_{Y_iY_j} \{\nabla \cdot (\mathbf{u}_j \mathbf{u}_i)\}^{k^\star}, \tag{6a}$$

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

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

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

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

with deterministic boundary conditions:

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

In the above, the time instant $t^{k^\star}$, either previous or intermediate, determines the order of the scheme in time. Specifically, for first order in time, $k^\star = k - 1$; for second order, we refer to [57,17]. 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 Section 3.4). A related one is the coupling between the differential equations for the mean, mode and coefficients (see Section 3.5).

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

$$\frac{\tilde{\mathbf{u}}_i^k}{\Delta t} - \frac{1}{\sqrt{Gr}} \nabla^2 \tilde{\mathbf{u}}_i^k = \frac{\mathbf{u}_i^{k-1}}{\Delta t} - \{\nabla \cdot (\mathbf{u}_i \bar{\mathbf{u}})\}^{k^\star} - \{\nabla \cdot (\bar{\mathbf{u}} \mathbf{u}_i)\}^{k^\star} - \nabla \breve{p}_i^{k^\star} + \rho_i^{k^\star} \mathbf{e}^g - \mathbf{C}^{-1,k^\star}_{Y_iY_j} \mathbf{M}^{k^\star}_{Y_jY_mY_n} \{\nabla \cdot (\mathbf{u}_n \mathbf{u}_m)\}^{k^\star} - \langle \mathbf{Q}_i, \mathbf{\Phi}_j \rangle^{k^\star}_{\mathcal{D}} \mathbf{u}_j^{k^\star}, \tag{8a}$$

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

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

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

$$\frac{\rho_i^k}{\Delta t} - \frac{1}{Sc\sqrt{Gr}} \nabla^2 \rho_i^k = \frac{\rho_i^{k-1}}{\Delta t} - \{\nabla \cdot (\mathbf{u}_i \bar{\rho})\}^{k^\star} - \{\nabla \cdot (\bar{\mathbf{u}} \rho_i)\}^{k^\star} - \mathbf{C}^{-1,k^\star}_{Y_iY_j} \mathbf{M}^{k^\star}_{Y_jY_mY_n} \{\nabla \cdot (\mathbf{u}_n \rho_m)\}^{k^\star} - \langle \mathbf{Q}_i, \mathbf{\Phi}_j \rangle^{k^\star}_{\mathcal{D}} \rho_j^{k^\star} \tag{8e}$$

with boundary conditions:

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

where, again, the scheme is first order for $k^\star = k-1$. Since we focus on the numerics of the DO equations, we assume that the stochastic boundary forcings are null, i.e. g_i's are null in Eqs. (9). 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.

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 employed later. 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 $t_{k-1} \leqslant t \leqslant t_k$ as follows:

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

where $(\cdot)^{k^-}$ 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 (10), we have compared several time-marching schemes of varying order to advance the Y_i in time (see [59], including: a low-storage 4th-order-accurate explicit Runge–Kutta integrator [3]; a 2nd-accurate explicit Runge–Kutta scheme (Heun's version); and the first-order-accurate explicit Euler scheme. 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 (Section 6). While $\mathbf{C}^{k^-}_{Y_mY_n}$ 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). The extension to the case where governing equations contain stochastic forcing in the form of zero-mean Wiener processes is discussed in [59].

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 Eqs. (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 ($\mathbf{C}_{Y_mY_n}, \mathbf{M}_{Y_jY_mY_n}$) to the approximated times $k^\star, k^-$ 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 $k^\star, k^-$ and store for later use
3. Advance the Y_i's using (10), and one of the ODE solvers in Section 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 $k^\star = k^- = k-1$, resulting in a first order accurate scheme for the present applications (Section 6). For second order accuracy, we would use $k^\star$ as defined in [57,17]. 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 $k^\star$ and k^- used for the time-integration of the modes and mean. Hence, if $k^\star = k^- = k-1$, 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 (Section 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 Section 4.2.

4.1. Spatial discretization of the physical space

The domain $\mathcal{D}$ is discretized into non-overlapping control volumes, presently forming structured Cartesian grid with rectangular uniform spacing in the x and y directions. Several choices exist for the relative placement of velocity and pressure control volumes. Here we employ a standard staggered C-grid [13,39], 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, see [59]. For the advection operator, the simple second-order central flux is well-known to be unstable (e.g. [4]), and it needs more careful treatment.

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:

$$F(\eta_{i-\frac{1}{2}}) = u_{i-\frac{1}{2}}\frac{\eta_i + \eta_{i-1}}{2} - \left|u_{i-\frac{1}{2}}\right|\frac{\eta_i - \eta_{i-1}}{2}C\left(u_{i-\frac{1}{2}}, \Delta t, \Delta x, \eta_{i-2}, \eta_{i-1}, \eta_i, \eta_{i+1}\right), \tag{11}$$

where the function C is defined in [59] and $u_{i-\frac{1}{2}}$ 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 [36].

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. Nonetheless, we have three potential schemes for advection by the modal velocities.

The three proposed schemes are tested in Section 6.2, where we find that the TVD* scheme performs best. Improving this TVD* is still possible, since minor oscillations can remain. We have shown [59] that at most 25% of the flux in the TVD* scheme will be upwinded and the method reduces to a CDS scheme if both directions require no slope limiting or the same amount of slope limiting, which explains why oscillations can remain. A proper flux-limited advection scheme may be derived using the characteristics of the hyperbolic parts of the full system. This system has $(s+1) \times d$ equations, where d is the problem's dimension, and the 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 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 [16,40]; (b) solve a discretized version of the PDEs for the probability densities of the coupled s coefficients, e.g. solve Fokker–Planck equation [49]; (c) parameterize the probability space, either using polynomial chaos [16,68,41,14], in our case extended to time-dependent polynomials, or other parameterizations such as Gaussian mixtures [43,55] or particle filters [9], and (d), use a Monte–Carlo approach [34,29].

We have employed a few of these schemes. Here, we only illustrate a Monte–Carlo scheme: for each realization the time-integration of Section 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 Eq. (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 of elements $Y_{r,i}$. Each row in this matrix corresponds to one of the samples and each column to one of the modes. During the time-integration step, each sample $Y_{r,i}$ 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 does not at all dominate the 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 mixture models [54,55] or other approaches mentioned above can then be used.

5. Implementation details

In this section we describe selected implementation details. 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 [1,34]. In the particular case of the DO equations, the inverse of the covariance matrix is multiplied by the third moments in (8),

$$\mathbf{C}^{-1,k^\star}_{Y_iY_j}\mathbf{M}^{k^\star}_{Y_jY_mY_n}, \tag{12}$$

which, for most physical processes, goes to zero for the eigenvalues of $\mathbf{C}^{k^\star}_{Y_iY_j}$ 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 benchmark in Section 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

$$\frac{\partial\langle\mathbf{\Phi}_i,\mathbf{\Phi}_j\rangle_{\mathcal{D}}}{\partial t} = \left\langle\frac{\partial\mathbf{\Phi}_i}{\partial t},\mathbf{\Phi}_j\right\rangle_{\mathcal{D}} + \left\langle\mathbf{\Phi}_i,\frac{\partial\mathbf{\Phi}_j}{\partial t}\right\rangle_{\mathcal{D}} = 0,$$

where the DO condition $\left\langle\mathbf{\Phi}_i,\frac{\partial\mathbf{\Phi}_j}{\partial t}\right\rangle_{\mathcal{D}} = 0,\ \forall i,j \in [1,2,\ldots,s]$ 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 us first consider the analytical integration from time t_{k-1} to t_k. For the modes, we have $\mathbf{\Phi}_i^k = \mathbf{\Phi}_i^{k-1} + \Delta t\overline{\frac{\partial\mathbf{\Phi}_i}{\partial t}}$ where the exact time integral is denoted as $\Delta t\overline{\frac{\partial\mathbf{\Phi}_i}{\partial t}} = \int_{t_{k-1}}^{t_k}\frac{\partial\mathbf{\Phi}_i}{\partial t}dt$. For the inner product at t_k, we then have:

$$\begin{aligned}\left\langle\mathbf{\Phi}_i^k,\mathbf{\Phi}_j^k\right\rangle_{\mathcal{D}} &= \left\langle\mathbf{\Phi}_i^{k-1} + \Delta t\overline{\frac{\partial\mathbf{\Phi}_i}{\partial t}},\mathbf{\Phi}_j^{k-1} + \Delta t\overline{\frac{\partial\mathbf{\Phi}_j}{\partial t}}\right\rangle_{\mathcal{D}},\\ &= \left\langle\mathbf{\Phi}_i^{k-1},\mathbf{\Phi}_j^{k-1}\right\rangle_{\mathcal{D}} + \Delta t\left\langle\mathbf{\Phi}_j^{k-1},\overline{\frac{\partial\mathbf{\Phi}_i}{\partial t}}\right\rangle_{\mathcal{D}} + \Delta t\left\langle\mathbf{\Phi}_i^{k-1},\overline{\frac{\partial\mathbf{\Phi}_j}{\partial t}}\right\rangle_{\mathcal{D}} + \Delta t^2\left\langle\overline{\frac{\partial\mathbf{\Phi}_i}{\partial t}},\overline{\frac{\partial\mathbf{\Phi}_j}{\partial t}}\right\rangle_{\mathcal{D}}.\end{aligned}$$

Since $\left\langle\mathbf{\Phi}_i^k,\mathbf{\Phi}_j^k\right\rangle_{\mathcal{D}} = \left\langle\mathbf{\Phi}_i^{k-1},\mathbf{\Phi}_j^{k-1}\right\rangle_{\mathcal{D}} = \delta_{ij}$, the remaining terms sum to zero,

$$\Delta t\left\langle\mathbf{\Phi}_j^{k-1},\overline{\frac{\partial\mathbf{\Phi}_i}{\partial t}}\right\rangle_{\mathcal{D}} + \Delta t\left\langle\mathbf{\Phi}_i^{k-1},\overline{\frac{\partial\mathbf{\Phi}_j}{\partial t}}\right\rangle_{\mathcal{D}} + \Delta t^2\left\langle\overline{\frac{\partial\mathbf{\Phi}_i}{\partial t}},\overline{\frac{\partial\mathbf{\Phi}_j}{\partial t}}\right\rangle_{\mathcal{D}} = 0. \tag{13}$$

Considering now a pth-order-accurate discrete approximation of $\overline{\frac{\partial\mathbf{\Phi}_i}{\partial t}}$, 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.

$\langle \boldsymbol{\Phi}_i, \boldsymbol{\Phi}_j \rangle_{\mathcal{D}}^{k,\text{discrete}} = \langle \boldsymbol{\Phi}_i, \boldsymbol{\Phi}_j \rangle_{\mathcal{D}}^{k} + \mathcal{E}_{ij}$. 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:

$$\langle \boldsymbol{\Phi}_i, \boldsymbol{\Phi}_j \rangle_{\mathcal{D}}^{k} + \mathcal{E}_{ij} = \left\langle \boldsymbol{\Phi}_i^{k-1}, \boldsymbol{\Phi}_j^{k-1} \right\rangle_{\mathcal{D}} + \Delta t \left\langle \boldsymbol{\Phi}_j^{k-1}, \overline{\frac{\partial \boldsymbol{\Phi}_i}{\partial t}} + \varepsilon_i \right\rangle_{\mathcal{D}} + \Delta t \left\langle \boldsymbol{\Phi}_i^{k-1}, \overline{\frac{\partial \boldsymbol{\Phi}_j}{\partial t}} + \varepsilon_j \right\rangle_{\mathcal{D}} + \Delta t^2 \left\langle \overline{\frac{\partial \boldsymbol{\Phi}_i}{\partial t}} + \varepsilon_i, \overline{\frac{\partial \boldsymbol{\Phi}_j}{\partial t}} + \varepsilon_j \right\rangle_{\mathcal{D}}.$$

Using (13),

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

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

The orthornomality can be corrected indirectly by enforcing that the solution and the error are numerically orthonormal, $\langle \boldsymbol{\Phi}_i, \boldsymbol{\Phi}_j \rangle_{\mathcal{D}}^{k} + \mathcal{E}_{ij} = \delta_{ij}$, at the end of a time step. This has to be done with care: the summation $Y_{r,i}\boldsymbol{\Phi}_i$ 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 [59].

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 (Section 6.1) for a version of the lock-exchange dynamical problem [19]. Next, we examine three advection schemes proposed for DO equations in Section 4.1.2, using a symmetric version of the lock-exchange problem (Section 6.2). Then, we evaluate the spatial and temporal convergence using the lid-driven cavity flow (Section 6.3). Finally, we study the discretization of the stochastic coefficients using the flow over a square cylinder in a confined channel (Section 6.4).

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

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 [33,58]. 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 Section 4.1.2 and Section 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 number required to have full energy interactions between the mean and modes of the DO simulation [49,48]. 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 $\text{Gr} = 1.25 \times 10^6$ and $\text{Gr} = 4 \times 10^4$, (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

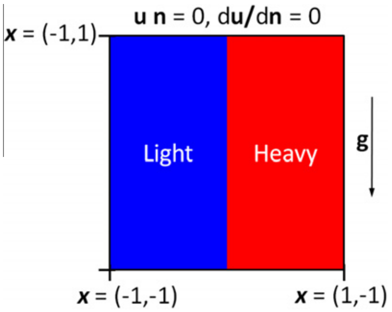

Fig. 1. Lock-exchange problem: an initial barrier separating light and heavy fluid is removed, and the flow is allowed to evolve. Uncertainty in our studies originates from not knowing the initial density differences between the fluids. This benchmark is used to verify the correctness of the implementation,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

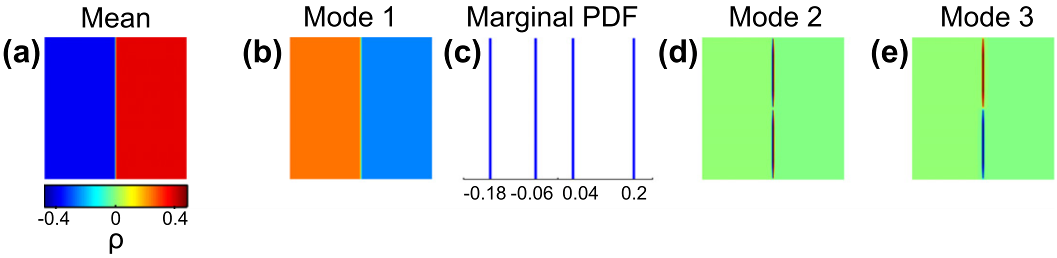

Fig. 2. Initialization of the lock-exchange problem (density field: mean, mode 1 and marginal pdf, modes 2 and 3). The initial velocity is zero. Stochasticity is introduced through the first (of three) orthonormal modes for the density. The vertical length scale used for non-dimensionalization is the half-height of the channel ($\hat{h} = 1$). Initializing with four 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 uncertainty.

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

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 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:

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

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

where these are orthonormalized numerically as described in [59]. The pressure and velocity for all modes are zero everywhere, initially.

Stochastic coefficient initialization: The discrete pdf for the first stochastic coefficient is specified using 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 $\sum_i \mathrm{Var}(Y_{r,i}) = \mathrm{Var}(Y_{r,1})$ 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 Section 5.1). Also, because the pdf is discrete, we calculate moments using the biased estimator $\mathbf{C}_{Y_iY_j} \approx \frac{1}{q}\sum_r Y_{r,i}Y_{r,j}$, instead of the usual unbiased estimator $\mathbf{C}_{Y_iY_j} \approx \frac{1}{q-1}\sum_r Y_{r,i}Y_{r,j}$. 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 shown in Fig. 3 for both Gr. Comparisons with the deterministic runs are in Fig. 4 and Fig. 5 for $\mathrm{Gr} = 4 \times 10^4$ and 1.25×10^6, respectively. Finally, the evolution of the differences between the stochastic and deterministic runs for both Gr are in Fig. 6. We obtain excellent agreement among the stochastic and deterministic runs.

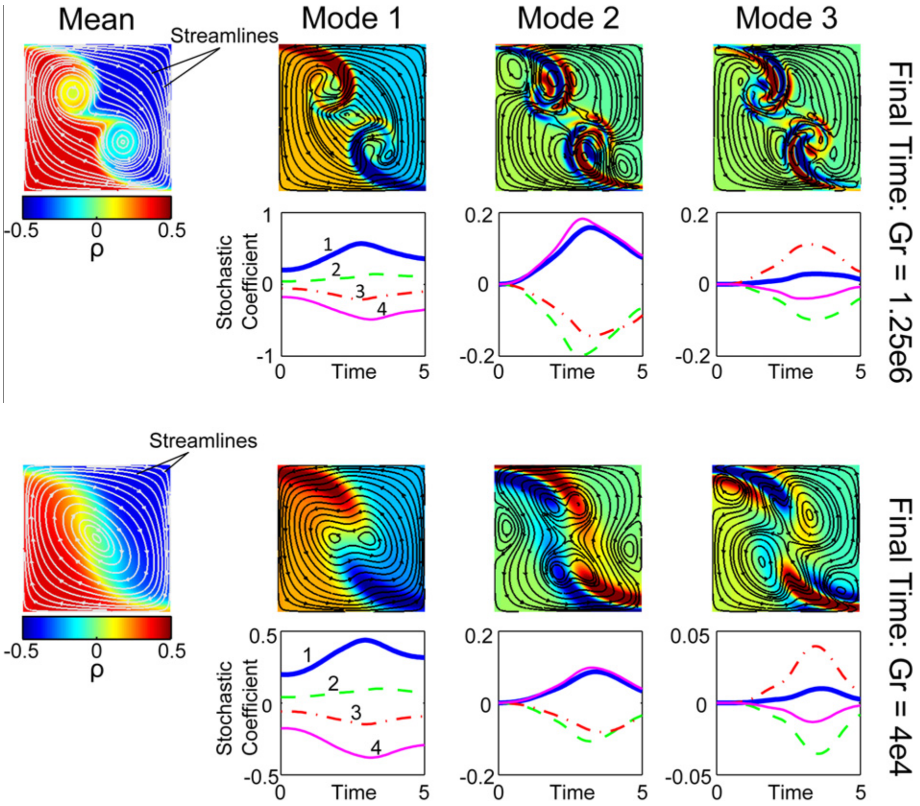

Fig. 3. The DO mean, modes at non-dimensional time $t = 5$, and evolution of stochastic coefficients, for $\mathrm{Gr} = 1.25 \times 10^6$ (top) and $\mathrm{Gr} = 4 \times 10^4$ (bottom). The higher-Gr flow has sharper gradients, and its coefficients are larger than the lower-Gr flow. Streamlines shown over density in color. Note that the sign of the contribution from a mode depends on the sign of the stochastic coefficients. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Particularly, for the lower-Gr 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. Based on these results and many other tests (not shown), our implementation is adequate, that is, we are solving the intended equations.

The growth of differences between the deterministic and stochastic simulations (Fig. 6) is due to differences between the stochastic and deterministic solvers, in particular the advection schemes and 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 modes did not change results significantly. Dividing the time-step by two to $\Delta t = 1/1024$ reduced the error at the final time from ~2.1% to ~1.05% for the higher-Gr flow. This indicates that the error is dominated by a temporal truncation compounded error of approximately $\mathcal{O}(\Delta t)$ (as expected, see Section 6.3). Hence, the primary source of this compounded error is the evolution of the stochastic coefficients and/or the modes. In [59], we show that it is the former, i.e. the stochastic coefficients time integration, that dominates this error.

In summary, using a discrete Dirac pdf benchmark, we showed that the DO implementation is correct. Also, based on the results, we suggest that the accuracy of the scheme will benefit most from improving the temporal discretization, especially 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 (Section 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.

6.2.1. Effect of advection scheme: setup

General setup: The Schmidt number is kept constant, $Sc = 1$, and we present results for $\mathrm{Gr} = 1.25 \times 10^6$. Initially the velocity is zero everywhere. Free-slip boundary conditions are used at the boundaries of the domain (Fig. 1). The domain is

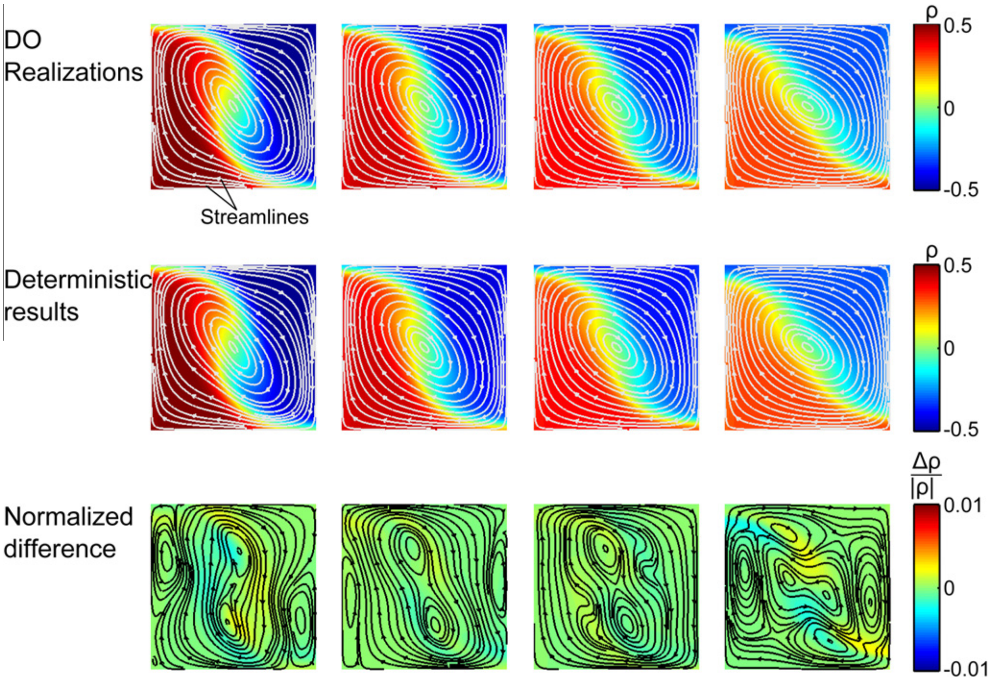

Fig. 4. The four DO realizations (top row), the deterministic runs (middle row), and the DO realizations minus the deterministic runs normalized by $\|\mathbf{\Phi}_{\text{deterministic}}\|_2$ (bottom row) for Gr = 4×10^4 (resolution 256×256 with 512 × 5 time-steps). The stochastic solver result agrees with the deterministic solver results, which validates our DO numerical schemes. Streamlines shown over density in color. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

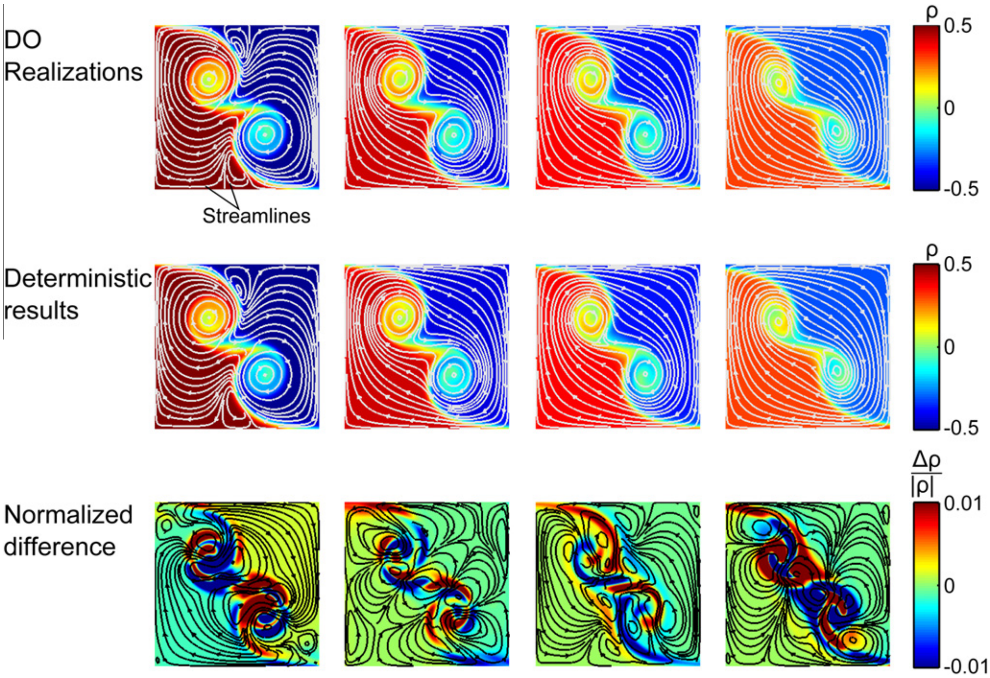

Fig. 5. As Fig. 4 but with Gr = 1.25×10^6. The higher-Gr flow with sharper gradients has larger errors than the lower-Gr flow, but they are still small.

discretized using $\Delta x = \Delta y = 1/64$, and $\Delta t = 1/256$. A lower resolution is used here compared to the cases in Section 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.

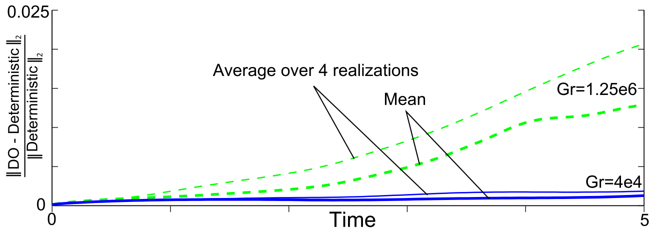

Fig. 6. The relative errors for both Gr tend to grow over time but can decrease. In both cases, the DO mean field has a smaller error than the average error of the realizations.

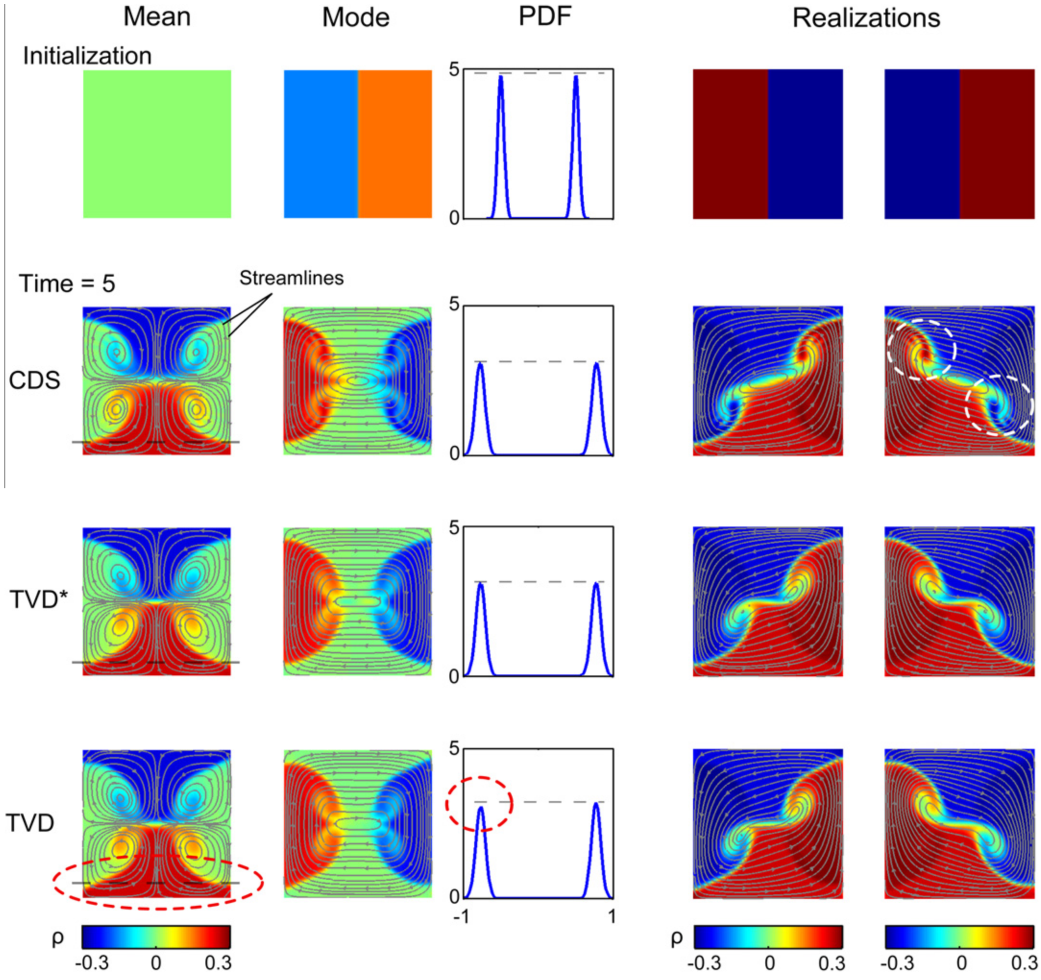

Fig. 7. Symmetric lock-exchange problem (higher $\mathrm{Gr} = 1.25 \times 10^6$) with grid resolution 64×64 and $\Delta t = 1/256$ using various advection schemes for the modes. Only the stochastic density is non-zero initially, with a bimodal pdf (top row). The two most extreme realizations (largest 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. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Mode initialization: The density profile for the mode is the same normalized hyperbolic tan profile used in Section 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 2500 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:

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

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.

6.2.2. 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 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 [50].

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 now 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 has a fixed density, and so we report the Re number instead of the Gr. Comparisons of the lid-driven cavity flows at different resolutions are provided in [59]; in what follows, we only illustrate the convergence results. We also completed convergence tests with variable densities: results (not shown here) are analogous.

6.3.1. Numerical convergence analysis: setup

General setup: We present results for $Re = 500$, although other cases ($Re \in [100, 1000]$) were also studied, giving similar results. The flow is driven by a deterministic boundary condition at the top of an enclosed cavity (see [59]), 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 Re (see [12,11]). Considering time, the order of convergence is approximated as $\mathcal{O} \approx \log\left(\frac{\|\phi_{2N_t} - \phi_{4N_t}\|_2}{\|\phi_{2N_t} - \phi_{N_t}\|_2}\right) / \log(2)$ where N_t is the number of time-steps (the same approach is used for space). 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 5000 samples of zero mean Gaussian distributions with variances $\mathrm{Var}(Y_{r_{5000},i}) = e^{(1-M_i-N_i)}$. 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 Section 6.2,

$$Y^*_{r,i} = \begin{bmatrix} Y_{r_{5000},i} \\ -Y_{r_{5000},i} \end{bmatrix}.$$

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

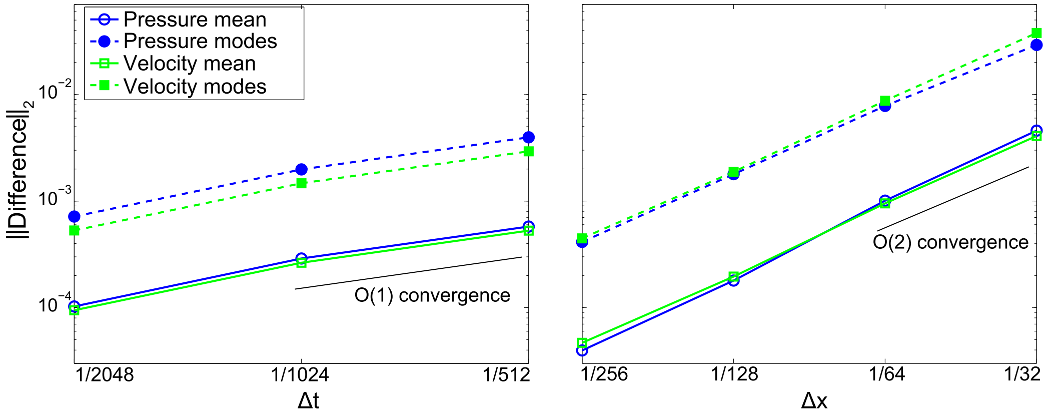

Fig. 8. The control volume size is held fixed at $\Delta x = 1/512$ for the time convergence (left), and the time step size is held fixed at $\Delta t = 1/4096$ for the spatial convergence (right). The error ($\|\phi_{2N} - \phi_N\|_2$) decreases with both temporal- and spatial-refinement for each component, and convergence is near-optimal (order 1 in time and 2 in space).

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 $\Delta x = 1/512$.

.	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

$$Y_{r,i} = Y^*_{r,i} \frac{\sqrt{\mathrm{Var}(Y_{r_{5000},i})}}{\sqrt{\mathrm{Var}^q(Y_{r,i})}},$$

where $\mathrm{Var}^q(a_{r,i}) = \frac{1}{q-1}\sum_{r=1}^{q} a_{r,i}$ is the calculated sample variance. For illustrations, we refer to [59].

6.3.2. Numerical convergence analysis: results and discussion

The convergence results are illustrated in Fig. 8. The reference high-resolution non-Gaussian DO simulation and its comparisons with simulations at different resolutions are shown in [59]. We find that the numerical errors for all components decrease with temporal and spatial refinement as expected. The convergence is near optimal at large grid sizes for all variables. These results with the velocity components separated are tabulated in Tables 1 and 2: 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 in time (based on choices made in Section 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 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.

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 $\Delta t = 1/4096$.

.	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

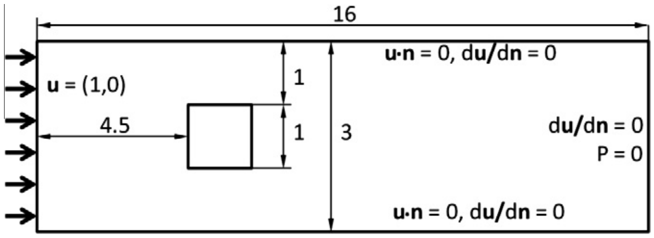

Fig. 9. Laminar vortex shedding over a square cylinder in a channel. Here uncertainty originates from the initial conditions and uncertain vortex shedding: depending on the perturbation, the first vortex could either be shed above or below the cylinder. Within a stochastic framework, however, if uncertainties are initially symmetric, the mean and modes should remain symmetric, and this is evaluated with our DO numerics.

6.4.1. Stochastic convergence: setup

The benchmark is based on an open flow in a frictionless pipe with a square cylindrical obstacle (Fig. 9), a classic test for deterministic solvers.

General setup: We present results for $Re = 100$, although other Re were 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

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

where $C_{M,N}$ is the normalization constant (as in Section 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

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

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

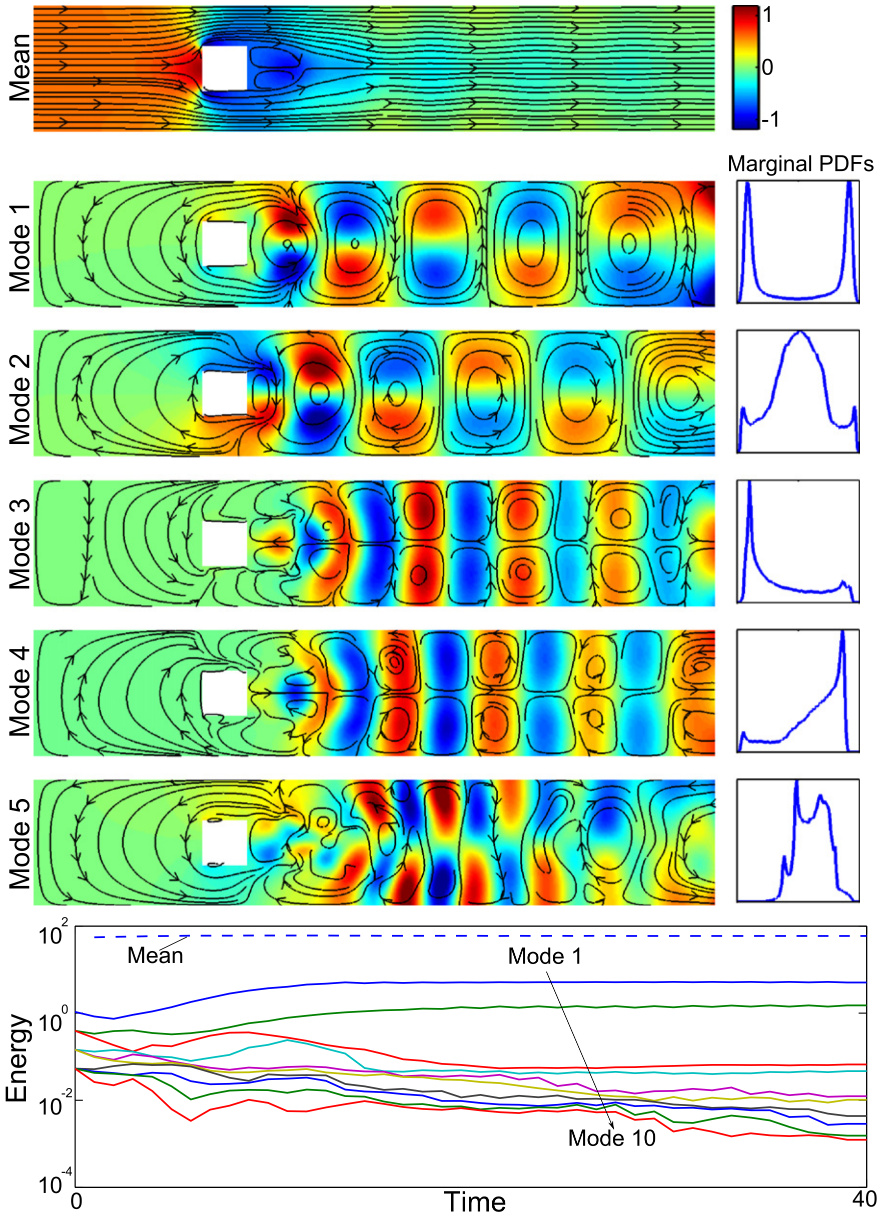

Fig. 10. Mean field and first 5 modes with marginal pdfs at final non-dimensional time $T = 40$ for Re = 100, and the evolution of the mean, $\langle \mathbf{u}, \mathbf{u} \rangle_{\mathcal{D}}$ and stochastic energy, $\mathrm{Var}(Y_i)$, for the reference solution (resolution 63×336 with 63×40 time-steps). Streamlines shown over pressure in color. Our scheme and implementation retains (anti)-symmetry for the most important first four modes. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

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

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. 10). From the realizations (Fig. 11), 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. 10, 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.

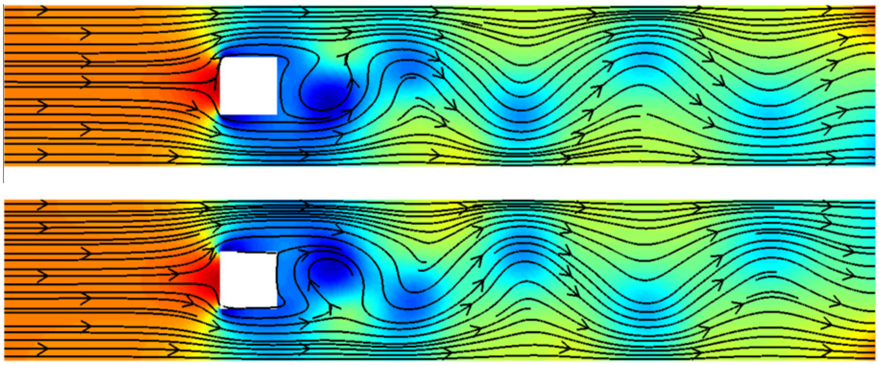

Fig. 11. As in Fig. 10, but showing two realizations where the vortex is shed in opposite directions. The colorbar for the pressure is as on Fig. 10. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

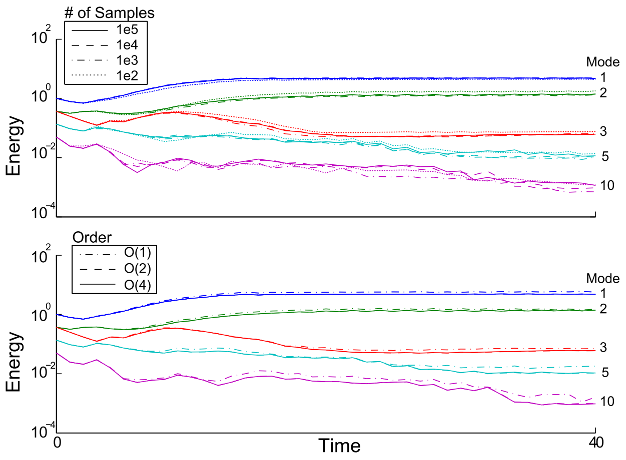

Fig. 12. Energy of the stochastic coefficients, Var(Y_i), over time for: (top) different number of samples ($\mathcal{O}(4)$ time integration); and (bottom), different time integration schemes (10,000 samples). Trends are well captured in all cases, but there are noticeable errors for less energetic modes after long integration times when a small number of samples and a low order time discretization scheme is used.

For our sequential implementation, the reference DO simulation (Fig. 10) was computed in about 7.5 h on a 2.4 Ghz 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 [55,56].

Sample sizes: Examining the evolution of the variance of the stochastic coefficients for a few sample sizes (top of Fig. 12), differences are larger for higher modes. For the first four, results agree well with the reference when using >1000 samples. After the 5th coefficient, differences are relatively larger, and in the 10th, they begin sooner. However, these differences are insignificant when compared to the variance of the first mode (note the logarithmic scale), indicative of convergence.

The shapes of marginals are also well-reproduced for 1000 samples or more, for all time integration schemes (Fig. 13). 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. Overall, the solution is converging with the number of samples: the cases with 10,000 and 100,000 samples give nearly identical results. We note that using such numbers of samples does not affect the overall cost. On the other hand, since marginals are still well-represented using smaller sizes, a small number of samples could be used for problems with larger s.

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 Section 3.4). Examining results (bottom of Fig. 12), 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.

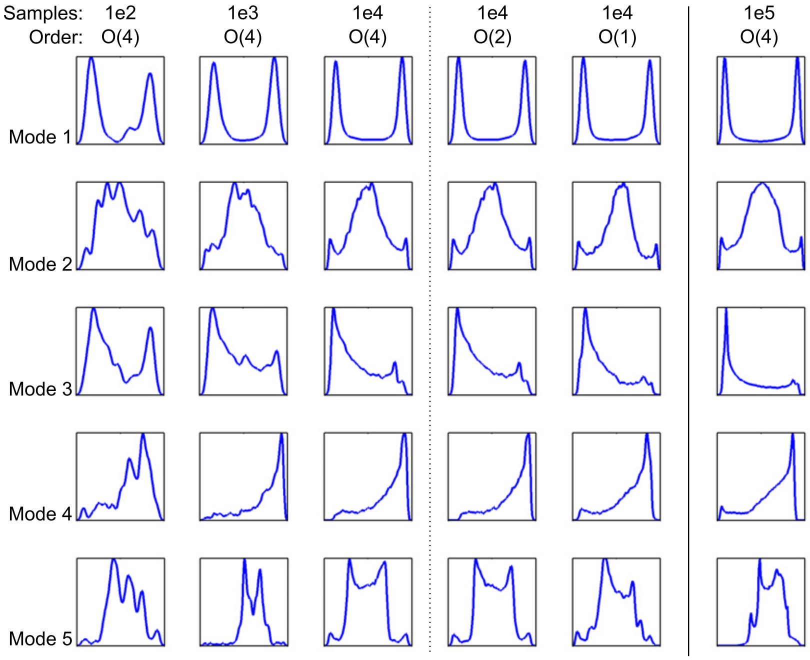

Fig. 13. Marginal probability density functions of first five modes for the flow over a square cylinder benchmark at final time (columns 1–3, increasing sample sizes; columns 4–5, decreasing order in time; column 6, reference). With 10,000 samples for the stochastic coefficients, the continuous marginal 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).

7. Conclusions and future research

We derived efficient computational schemes for the DO methodology applied to unsteady stochastic Navier–Stokes and Boussinesq simulations, and illustrated and studied the numerical aspects of these schemes. For the discretizations in time, we develop semi-implicit projection methods for the mean and modes, and we employ 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. We 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, accounting 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 DO simulation. This validated our numerical implementation, and confirmed that the pseudo-stochastic pressure approach works in practice. Using a symmetric version of this lock-exchange benchmark, we showed that the symmetric TVD-based advection scheme (TVD*) works well, while the CDS scheme suffers from numerical oscillations, and the classic TVD scheme loses symmetry.

Using a lid-driven cavity flow with uncertain initial conditions, we showed that state variables converged near-optimally under both time and space refinement, even though pdfs were non-Gaussian. We also showed that even when a higher-order time integrator is used for the stochastic coefficients, accuracy is still limited to the overall order of the DO scheme (Section 3.4).

Finally, using a stochastic flow past a square cylinder in a confined channel, we showed that the stochastic coefficients converge 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 numerical DO scheme successfully captures both vortex shedding directions, maintaining a fully symmetric mean field. Also, with the 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. 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 can allow the prediction of uncertainty in realistic simulations of multiphysics fluid systems, over a wide range of applications, from micro-nano fluid engineering to multiscale ocean and climate studies.

Acknowledgements

We are very thankful to the MSEAS group, 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), N00014-09-1-0676 (Science of Autonomy – A-MISSION) and N00014-12-1-0944 (ONR6.2) 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 [49,48,51]) to obtain these equations.

Starting with a generalized Karhunen-Loève expansion [31,48], one decomposes the solution of Eqs. (1)–(3) into a Dynamically Orthogonal (DO) field expansion [49,48,50] for the velocity, density and pressure[3] fields

$$\begin{aligned} &\mathbf{\Phi}(\mathbf{x},t;\omega) = \bar{\mathbf{\Phi}}(\mathbf{x},t) + \sum_{i=1}^{s} Y_i(t;\omega)\mathbf{\Phi}_i(\mathbf{x},t), \\ &\mathbf{\Phi} \equiv \bar{\mathbf{\Phi}} + Y_i\mathbf{\Phi}_i, \end{aligned} \tag{A.1}$$

$$\begin{aligned} &p(\mathbf{x},t;\omega) = \bar{p}(\mathbf{x},t) + \sum_{i=1}^{s} Y_i(t,\omega)p_i(\mathbf{x},t) + \sum_{i=1}^{s}\sum_{j=1}^{s} Y_i(t;\omega)Y_j(t;\omega)p_{ij}(\mathbf{x},t), \\ &p \equiv \bar{p} + Y_ip_i + Y_iY_jp_{ij}, \end{aligned} \tag{A.2}$$

where $\mathbf{\Phi} = [\Phi^1, \Phi^2, \ldots]^T = [\mathbf{u}, \rho]^T$ is the vector of prognostic state variables. The scalar $s = s(t)$ defines the time-dependent dimension of the stochastic subspace [50], 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 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 Section 3.2 or [49]). Since both the modes and stochastic coefficients are functions of time, a redundancy arises, which is resolved by the DO condition,

$$\left\langle \mathbf{\Phi}_i, \frac{\partial \mathbf{\Phi}_j}{\partial t} \right\rangle_{\mathcal{D}} = 0, \quad \forall i,j \in [1,2,\ldots,s], \tag{A.3}$$

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

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 Eqs. (1)–(3) to obtain (using a Langevin notation):

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

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 Eqs. (A.2) and (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 Eq. (A.4). Hence, taking the expectation of Eq. (A.4) it can be found that the evolution of the mean fields are governed by

---

[3] For convenience, we changed the sign of p_{ij} from the original definition in [49].

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

with the initial and boundary conditions given by

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

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

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

Modes. The evolution of the modes is also obtained from Eq. (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 $E^\omega\left[\frac{\partial Y_i}{\partial t}Y_j\right]$, which is obtained by projecting the equation unto $\mathbf{\Phi}_k$ and imposing the DO condition. From this the following governing evolution equations for the modes can be found:

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

where

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

and $\mathbf{M}_{Y_jY_mY_n} = E^\omega[Y_jY_mY_n]$ is a third moment. The right-hand-sides in Eq. (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, Eqs. (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

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

Coefficients. Finally, to obtain the evolution of the stochastic coefficients from Eq. (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 Eq. (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

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

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

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

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

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

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

References

[1] A.F. Bennett, Inverse Methods in Physical Oceanography, 1st ed., Cambridge University Press, New York, 1992.

[2] R.H. Cameron, W.T. Martin, The orthogonal development of nonlinear functionals in series of Fourier–Hermite functionals, Ann. Math. 48 (1947) 385–392.

[3] M.H. Carpenter, C. Kennedy, Fourth-order 2N-storage Runge–Kutta schemes. NASA Report TM 109112, NASA Langley Research Center, Hampton, VA, 1994.

[4] S.C. Chapra, R.P. Canale, Numerical Methods for Engineers, 5th edn. McGraw-Hill Higher Education, Boston, MA, 2006.

[5] D. Chen, J. Liu, Mixture Kalman filters, J. Roy. Statist. Soc. Ser. A 62 (2000) 493–508.

[6] B. Cushman-Roisin, J.M. Beckers, Introduction to Geophysical Fluid Dynamics: Physical and Numerical Aspects, Academic Press, 2010.

[7] M.K. Deb, I. Babuška, J. Oden, Solution of stochastic partial differential equations using Galerkin finite element techniques, Comput Methods Appl. Mech. Eng. 190 (2001) 6359–6372.

[8] E. Deleersnijder, P.F.J. Lermusiaux, Multi-scale modeling: nested grid and unstructured grid approaches, Ocean Dynam. 58 (2008) 335–336.

[9] A. Doucet, N. de Freitas, N. Gordon, Sequential Monte–Carlo Methods in Practice, Springer-Verlag, 2001.

[10] M.S. Eldred, C.G. Webster, P. Constantine, 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, 2008.

[11] E. Erturk, Discussions on driven cavity flow, Int. J. Numer. Methods Fluids 60 (2009) 275–294.

[12] E. Erturk, T.C. Corke, C. Gökçöl, Numerical solutions of 2-D steady incompressible driven cavity flow at high Reynolds numbers, Int. J. Numer. Methods Fluids 48 (2005) 747–774.

[13] J.H. Ferziger, M. Peric, Computational methods for fluid dynamics, 3rd ed., Springer, New York, NY, 2002.

[14] J. Foo, G.E. Karniadakis, Multi-element probabilistic collocation method in high dimensions, J. Comput. Phys. 229 (2010) 1536–1557.

[15] J. Frochte, W. Heinrichs, A splitting technique of higher order for the Navier–Stokes equations, J. Comput. Appl. Math. 228 (2009) 373–390.

[16] R. Ghanem, P. Spanos, Stochastic Finite Elements: A Spectral Approach, Springer-Verlag, 1991.

[17] J. Guermond, P. Minev, J. Shen, An overview of projection methods for incompressible flows, Comput. Methods Appl. Mech. Eng. 195 (2006) 6011–6045.

[18] P.J. Haley Jr., P.F.J. Lermusiaux, Multiscale two-way embedding schemes for free-surface primitive-equations in the multidisciplinary simulation, estimation and assimilation system (MSEAS), Ocean Dynam. 60 (2010) 1497–1537.

[19] C. Härtel, E. Meiburg, F. Necker, 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 (2000) 189–212.

[20] S. Hosder, R. Walters, R. Perez. A nonintrusive polynomial chaos method for uncertainty propagation in CFD simulations, in: AIAA Aerospace Science Meeting Exhibition, 44th, AIAA-2006-0891, Reno, Nevada, 2006.

[21] J.D. Jakeman, S.G. Roberts, Stochastic Galerkin and collocation methods for quantifying uncertainty in differential equations: a review, Aust. NZ Ind. Appl. Math. J. 50 (2008) C815–C830.

[22] K. Karhunen, Zur spektraltheorie stochastischer prozesse, Ann. Acad. Sci. Fennicae 34 (1946) 17.

[23] P.E. Kloeden, E. Platen, Numerical Solution of Stochastic Differential Equations, Springer-Verlag, Berlin, 1999.

[24] O. Le Maître, M. Reagan, H. Najm, R. Ghanem, O. Knio, A stochastic projection method for fluid flow II. Random process, J. Comput. Phys. 181 (2002) 9–44.

[25] O. Le Maître, H. Najm, R. Ghanem, O. Knio, Multi-resolution analysis of Wiener-type uncertainty propagation schemes, J. Comput. Phys. 197 (2004) 502–531.

[26] O.P. Le Maître, O.M. Knio, Spectral methods for uncertainty quantification, Scientific Computation, 1st ed., Springer, New York, 2010.

[27] O.P. Le Maître, O.M. Knio, H.N. Najm, R.G. Ghanem, A stochastic projection method for fluid flow I. Basic formulation, J. Comput. Phys. 173 (2001) 481–511.

[28] P.F.J. Lermusiaux, Data assimilation via error subspace statistical estimation. Part II: Middle atlantic bight shelfbreak front simulations and ESSE validation, Mon. Weather Rev. 127 (1999) 1408–1432.

[29] P.F.J. Lermusiaux, Evolving the subspace of the three-dimensional multiscale ocean variability: Massachusetts Bay, J. Mar. Syst. 29 (2001) 385–422.

[30] P.F.J. Lermusiaux, On the mapping of multivariate geophysical fields: sensitivities to size, scales, and dynamics, J. Atmos. Ocean Technol. 19 (10) (2002) 1602–1637.

[31] P.F.J. Lermusiaux, Uncertainty estimation and prediction for interdisciplinary ocean dynamics, J. Comput. Phys. 217 (2006) 176–199.

[32] P.F.J. Lermusiaux, Adaptive modeling, adaptive data assimilation and adaptive sampling, Physica D 230 (2007) 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..

[33] P.F.J. Lermusiaux, 2.29: Numerical Fluid Mechanics. Lectures Notes, MIT, Cambridge, MA, 2009.

[34] P.F.J. Lermusiaux, A.R. Robinson, Data assimilation via error subspace statistical estimation. Part I: Theory and schemes, Mon. Weather Rev. 127 (1999) 1385–1407.

[35] P.F.J. Lermusiaux, C.S. Chiu, G. Gawarkiewicz, P. Abbot, A. Robinson, R. Miller, P. Haley, W. Leslie, S. Majumdar, A. Pang, F. Lekien, Quantifying uncertainties in ocean predictions, Oceanography 19 (2006) 92–105.

[36] R.J. Leveque, Finite Volume Methods for Hyperbolic Problems, Cambridge University Press, 2002.

[37] R. Li, R. Ghanem, Adaptive polynomial chaos expansions applied to statistics of extremes in nonlinear random vibration, Prob. Eng. Mech. 13 (1998) 125–136.

[38] M. Loève, Probability Theory, Springer-Verlag, Berlin/New York, 1977.

[39] O.G. Logutov, P.F.J. Lermusiaux, Inverse Barotropic tidal estimation for regional ocean applications, Ocean Model. 25 (2008) 17–34.

[40] D. Lucor, J. Meyers, P. Sagaut, Sensitivity analysis of large-eddy simulations to subgrid-scale-model parametric uncertainty using polynomial chaos, J. Fluid Mech. 585 (2007) 255–279.

[41] Y.M. Marzouk, H.N. Najm, Dimensionality reduction and polynomial chaos acceleration of Bayesian inference in inverse problems, J. Comput. Phys. 228 (6) (2009) 1862–1902.

[42] L. Mathelin, M.Y. Hussaini, T.A. Zang, Stochastic approaches to uncertainty quantification in CFD simulations, Numer. Algor. 38 (1–3) (2005) 209–236.

[43] G. McLachlan, D. Peel, Finite mixture models, Wiley Series in Probability and Statistics: Applied Probability and Statistics, Wiley, 2000.

[44] R. van der Merwe, E. Wan, 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, 2003.

[45] R. Miller, Topics in data assimilation: stochastic processes, Physica D 230 (1,2) (2007) 17–26.

[46] H.N. Najm, Uncertainty quantification and polynomial chaos techniques in computational fluid dynamics, Annu. Rev. Fluid Mech. 41 (2009) 35–52.

[47] A. Nouy, A generalized spectral decomposition technique to solve a class of linear stochastic partial differential equations, Comput. Methods Appl. Mech. Eng. 196 (2007) 4521–4537.

[48] T.P. Sapsis, Dynamically orthogonal field equations for stochastic fluid flows and particle dynamics, Ph.D. Thesis, Massachusetts Institute of Technology, Dept. of Mechanical Engineering, 2010.

[49] T.P. Sapsis, P.F.J. Lermusiaux, Dynamically orthogonal field equations for continuous stochastic dynamical systems, Physica D 238 (2009) 2347–2360.

[50] T.P. Sapsis, P.F.J. Lermusiaux, Dynamical criteria for the evolution of the stochastic dimensionality in flows with uncertainty, Physica D 241 (1) (2012) 60–76.

[51] T.P. Sapsis, P.F.J. Lermusiaux, M.P. Ueckermann, Statistical model of Navier–Stokes equations using dynamical orthogonality, J. Fluid. Mech, submitted for publication.

[52] D. Shirokoff, R. Rosales, 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) (2011) 8619–8646.

[53] F. Simon, P. Guillen, P. Sagaut, D. Lucor, A gPC-based approach to uncertain transonic aerodynamics, Comput. Methods Appl. Mech. Eng. 199 (17-20) (2010) 1091–1099.

[54] T. Sondergaard, Data assimilation with gaussian mixture models using the dynamically orthogonal field equations, S.M. Thesis, Massachusetts Institute of Technology, Dept. of Mechanical Engineering, 2011.

[55] T. Sondergaard, P.F.J. Lermusiaux, Data assimilation with gaussian mixture models using the dynamically orthogonal field equations. Part I: Theory and scheme. Mon Weather Rev. in press.

[56] T. Sondergaard, P.F.J. Lermusiaux, Data assimilation with gaussian mixture models using the dynamically orthogonal field equations. Part II: Applications. Mon Weather Rev. in press.

[57] L. Timmermans, P. Minev, F. Van De Vosse, An approximate projection scheme for incompressible flow using spectral elements, Int. J. Numer. Methods Fluids 22 (1996) 673–688.

[58] M.P. Ueckermann, P.F.J.L. Lermusiaux, 2.29 Finite volume MATLAB framework documentation. Reports in Ocean Science and Engineering, 14, Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts, 2012.

[59] M.P. Ueckermann, P.F.J.L. Lermusiaux, T. Sapsis, Numerical schemes and computational studies for dynamically orthogonal equations. Reports in Ocean Science and Engineering 10, Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts, 2011.

[60] B. Van Leer, Towards the ultimate conservative difference scheme. IV. A new approach to numerical convection, J. Comput. Phys. 23 (3) (1977) 276–299.

[61] X. Wan, G.E. Karniadakis, An adaptive multi-element generalized polynomial chaos method for stochastic differential equations, J. Comput. Phys. 209 (2005) 617–642.

[62] X. Wan, G.E. Karniadakis, Long-term behavior of polynomial chaos in stochastic flow simulations, Comput. Methods Appl. Mech. Eng. 195 (2006) 5582–5596.

[63] M. Webster, A. Sokolov, A methodology for quantifying uncertainty in climate projections, Clim. Change 46 (2000) 417–446.

[64] N. Wiener, The homogeneous chaos, Amer. J. Math. 60 (1938) 897–936.

[65] N. Wiener Nonlinear Problems in Random Theory, MIT Technology Press and John Wiley and Sons; 1958.

[66] D. Xiu, Fast numerical methods for stochastic computations: a review, Commun. Comput. Phys. 5 (2–4) (2009) 242–272.

[67] D. Xiu, Numerical Methods for Stochastic Computations: A Spectral Method Approach, Princeton University Press, 2010.

[68] D. Xiu, G.E. Karniadakis, The Wiener–Askey polynomial chaos for stochastic differential equations, SIAM J. Sci. Comput. 24 (2002) 619–644.

[69] D. Xiu, D. Lucor, C.H. Su, G.E. Karniadakis, Stochastic modeling of flow-structure interactions using generalized polynomial chaos, J. Fluids Eng. 124 (1) (2002) 51–60.