

STOCHASTIC GALERKIN METHODS FOR THE STEADY-STATE NAVIER-STOKES EQUATIONS *

BEDŘICH SOUSEDÍK[†] AND HOWARD C. ELMAN[‡]

Abstract. We study the steady-state Navier-Stokes equations in the context of stochastic finite element discretizations. Specifically, we assume that the viscosity is a random field given in the form of a generalized polynomial chaos expansion. For the resulting stochastic problem, we formulate the model and linearization schemes using Picard and Newton iterations in the framework of the stochastic Galerkin method, and we explore properties of the resulting stochastic solutions. We also propose a preconditioner for solving the linear systems of equations arising at each step of the stochastic (Galerkin) nonlinear iteration and demonstrate its effectiveness for solving a set of benchmark problems.

1. Introduction. Models of mathematical physics are typically based on partial differential equations (PDEs) that use parameters as input data. In many situations, the values of parameters are not known precisely and are modeled as random fields, giving rise to *stochastic partial differential equations*. In this study we focus on models from fluid dynamics, in particular the stochastic Stokes and the Navier-Stokes equations. We consider the viscosity as a random field modeled as colored noise, and we use numerical methods based on spectral methods, specifically, the generalized polynomial chaos (gPC) framework [10, 14, 26, 27]. That is, the viscosity is given by a gPC expansion, and we seek gPC expansions of the velocity and pressure solutions.

There is a number of reasons to motivate our interest in Navier-Stokes equations with stochastic viscosity. For example, the exact value of viscosity may not be known, due to measurement error, the presence of contaminants with uncertain concentrations, or of multiple phases with uncertain ratios. Alternatively, the fluid properties might be influenced by an external field, with applications for example in magnetohydrodynamics. Specifically, we assume that the viscosity ν depends on a set of random variables ξ . This means that the Reynolds number,

$$\text{Re}(\xi) = \frac{UL}{\nu(\xi)},$$

where U is the characteristic velocity and L is the characteristic length, is also stochastic. Consequently, the solution variables are random fields, and different realizations of the viscosity give rise to realizations of the velocities and pressures. As observed in [19], there are other possible formulations and interpretations of fluid flow with stochastic Reynolds number for example, where the velocity is fixed but the volume of fluid moving into a channel is uncertain so the uncertainty derives from the Dirichlet inflow boundary condition.

We consider models of steady-state stochastic motion of an incompressible fluid moving in a domain $D \subset \mathbb{R}^2$. Extension to three-dimensional models is straightforward. We formulate the stochastic Stokes and Navier-Stokes equations using the

*This work is based upon work supported by the U. S. Department of Energy Office of Advanced Scientific Computing Research, Applied Mathematics program under Award Number DE-SC0009301, and by the U. S. National Science Foundation under grants DMS1418754 and DMS1521563.

[†]Department of Mathematics and Statistics, University of Maryland, Baltimore County, 1000 Hilltop Circle, Baltimore, MD 21250 (sousedik@umbc.edu).

[‡]Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742 (elman@cs.umd.edu)

stochastic finite element method, assuming that the viscosity has a general probability distribution parametrized by a gPC expansion. We describe linearization schemes based on Picard and Newton iteration for the stochastic Galerkin method, and we explore properties of the solutions obtained, including a comparison of the stochastic Galerkin solutions with those obtained using other approaches, such as Monte Carlo and stochastic collocation methods [26]. Finally, we propose efficient hierarchical preconditioners for iterative solution of the linear systems solved at each step of the nonlinear iteration in the context of the stochastic Galerkin method. Our approach is related to recent work by Powell and Sylvester [19]. However, besides using a general parametrization of the viscosity, our formulation of the stochastic Galerkin system allows straightforward application of state-of-the-art deterministic preconditioners by wrapping them in the hierarchical preconditioner developed in [22]. For alternative preconditioners see, e.g., [2, 11, 17, 18, 20, 23, 25]. Finally, we note that there exist related approaches based on stochastic perturbation methods [13], important developments also include reduced-order models such as [4, 24], and an overview of existing methods for stochastic computational fluid dynamics can be found in the monograph [14].

The paper is organized as follows. In Section 2, we recall the deterministic steady-state Navier-Stokes equations and their discrete form. In Section 3, we formulate the model with stochastic viscosity, derive linearization schemes for the stochastic Galerkin formulation of the model, and explore properties of the resulting solutions for a set of benchmark problems that model the flow over an obstacle. In Section 4 we introduce a preconditioner for the stochastic Galerkin linear systems solved at each step of the nonlinear iteration, and in Section 5 we summarize our work.

2. Deterministic Navier-Stokes equations. We begin by defining the model and notation, following [6]. For the deterministic Navier-Stokes equations, we wish to find velocity \vec{u} and pressure p such that

$$-\nu \nabla^2 \vec{u} + (\vec{u} \cdot \nabla) \vec{u} + \nabla p = \vec{f}, \quad (2.1)$$

$$\nabla \cdot \vec{u} = 0, \quad (2.2)$$

in a spatial domain D , satisfying boundary conditions

$$\vec{u} = \vec{g}, \quad \text{on } \Gamma_{\text{Dir}}, \quad (2.3)$$

$$\nu \nabla \vec{u} \cdot \vec{n} - p \vec{n} = \vec{0}, \quad \text{on } \Gamma_{\text{Neu}}, \quad (2.4)$$

where $\partial D = \bar{\Gamma}_{\text{Dir}} \cup \bar{\Gamma}_{\text{Neu}}$, and assuming sufficient regularity of the data. Dropping the convective term $(\vec{u} \cdot \nabla) \vec{u}$ from (2.1) yields the Stokes problem

$$-\nu \nabla^2 \vec{u} + \nabla p = \vec{f}, \quad (2.5)$$

$$\nabla \cdot \vec{u} = 0. \quad (2.6)$$

The mixed variational formulation of (2.1)–(2.2) is to find $(\vec{u}, p) \in (V_E, Q_D)$ such that

$$\nu \int_D \nabla \vec{u} : \nabla \vec{v} + \int_D (\vec{u} \cdot \nabla \vec{u}) \vec{v} - \int_D p (\nabla \cdot \vec{v}) = \int_D \vec{f} \cdot \vec{v}, \quad \forall \vec{v} \in V_D, \quad (2.7)$$

$$\int_D q (\nabla \cdot \vec{u}) = 0, \quad \forall q \in Q_D, \quad (2.8)$$

where (V_D, Q_D) is a pair of spaces satisfying the inf-sup condition and V_E is an extension of V_D containing velocity vectors that satisfy the Dirichlet boundary conditions [3, 6, 12].

Let $c(\vec{z}; \vec{u}, \vec{v}) \equiv \int_{\Omega} (\vec{z} \cdot \nabla \vec{u}) \cdot \vec{v}$. Because the problem (2.7)–(2.8) is nonlinear, it is solved using a linearization scheme in the form of Newton or Picard iteration, derived as follows. Consider the solution (\vec{u}, p) of (2.7)–(2.8) to be given as $\vec{u} = \vec{u}^n + \delta \vec{u}^n$ and $p = p^n + \delta p^n$. Substituting into (2.7)–(2.8) and neglecting the quadratic term $c(\delta \vec{u}^n; \delta \vec{u}^n, \vec{v})$ gives

$$\nu \int_D \nabla \delta \vec{u}^n : \nabla \vec{v} + c(\delta \vec{u}^n; \vec{u}^n, \vec{v}) + c(\vec{u}^n; \delta \vec{u}^n, \vec{v}) - \int_D \delta p^n (\nabla \cdot \vec{v}) = R^n(\vec{v}), \quad (2.9)$$

$$\int_D q (\nabla \cdot \delta \vec{u}^n) = r^n(q), \quad (2.10)$$

where

$$R^n(\vec{v}) = \int_D \vec{f} \cdot \vec{v} - \nu \int_D \nabla \vec{u}^n : \nabla \vec{v} - c(\vec{u}^n; \vec{u}^n, \vec{v}) + \int_D p^n (\nabla \cdot \vec{v}), \quad (2.11)$$

$$r^n(q) = - \int_D q (\nabla \cdot \vec{u}^n). \quad (2.12)$$

Step n of the *Newton iteration* obtains $(\delta \vec{u}^n, \delta p^n)$ from (2.9)–(2.10) and updates the solution as

$$\vec{u}^{n+1} = \vec{u}^n + \delta \vec{u}^n, \quad (2.13)$$

$$p^{n+1} = p^n + \delta p^n. \quad (2.14)$$

Step n of the *Picard iteration* omits the term $c(\delta \vec{u}^n; \vec{u}^n, \vec{v})$ in (2.9), giving

$$\nu \int_D \nabla \delta \vec{u}^n : \nabla \vec{v} + c(\vec{u}^n; \delta \vec{u}^n, \vec{v}) - \int_D \delta p^n (\nabla \cdot \vec{v}) = R^n(\vec{v}), \quad (2.15)$$

$$\int_D q (\nabla \cdot \delta \vec{u}^n) = r^n(q). \quad (2.16)$$

Consider the discretization of (2.1)–(2.2) by a div-stable mixed finite element method; for experiments discussed below, we used Taylor-Hood elements [6]. Let the bases for the velocity and pressure spaces be denoted $\{\phi_i\}_{i=1}^{N_u}$ and $\{\varphi_i\}_{i=1}^{N_p}$, respectively. In matrix terminology, each nonlinear iteration entails solving a linear system

$$\begin{bmatrix} \mathbf{F}^n & \mathbf{B}^T \\ \mathbf{B} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \delta \mathbf{u}^n \\ \delta \mathbf{p}^n \end{bmatrix} = \begin{bmatrix} \mathbf{R}^n \\ \mathbf{r}^n \end{bmatrix}, \quad (2.17)$$

followed by an update of the solution

$$\mathbf{u}^{n+1} = \mathbf{u}^n + \delta \mathbf{u}^n, \quad (2.18)$$

$$\mathbf{p}^{n+1} = \mathbf{p}^n + \delta \mathbf{p}^n. \quad (2.19)$$

For Newton's method, \mathbf{F}^n is the Jacobian matrix, a sum of the vector-Laplacian matrix \mathbf{A} , the vector-convection matrix \mathbf{N}^n , and the Newton derivative matrix \mathbf{W}^n ,

$$\mathbf{F}^n = \mathbf{A} + \mathbf{N}^n + \mathbf{W}^n, \quad (2.20)$$

where

$$\begin{aligned}\mathbf{A} &= [a_{ab}], \quad a_{ab} = \nu \int_D \nabla \phi_b : \nabla \phi_a, \\ \mathbf{N}^n &= [n_{ab}^n], \quad n_{ab}^n = \int_D (u^n \cdot \nabla \phi_b) \cdot \phi_a, \\ \mathbf{W}^n &= [w_{ab}^n], \quad w_{ab}^n = \int_D (\phi_b \cdot \nabla u^n) \cdot \phi_a.\end{aligned}$$

For Picard iteration, the Newton derivative matrix \mathbf{W}^n is dropped, and $\mathbf{F}^n = \mathbf{A} + \mathbf{N}^n$. The divergence matrix B is defined as

$$\mathbf{B} = [b_{cd}], \quad b_{cd} = \int_D \phi_d (\nabla \cdot \varphi_c). \quad (2.21)$$

The residuals at step n of both nonlinear iterations are computed as

$$\begin{bmatrix} \mathbf{R}^n \\ \mathbf{r}^n \end{bmatrix} = \begin{bmatrix} \mathbf{f} \\ \mathbf{g} \end{bmatrix} - \begin{bmatrix} \mathbf{P}^n & \mathbf{B}^T \\ \mathbf{B} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{u}^n \\ \mathbf{p}^n \end{bmatrix}, \quad (2.22)$$

where $\mathbf{P}^n = \mathbf{A} + \mathbf{N}^n$ and \mathbf{f} is a discrete version of the forcing function of (2.1).¹

3. The Navier-Stokes equations with stochastic viscosity. Let $(\Omega, \mathcal{F}, \mathcal{P})$ represent a complete probability space, where Ω is the sample space, \mathcal{F} is a σ -algebra on Ω and \mathcal{P} is a probability measure. We will assume that the randomness in the model is induced by a vector of independent, identically distributed (i.i.d.) random variables $\xi = (\xi_1, \dots, \xi_N)^T$ such that $\xi : \Omega \rightarrow \Gamma \subset \mathbb{R}^N$. Let $\mathcal{F}_\xi \subset \mathcal{F}$ denote the σ -algebra generated by ξ , and let $\mu(\xi)$ denote the joint probability density measure for ξ . The expected value of the product of random variables u and v that depend on ξ determines a Hilbert space $T_\Gamma \equiv L^2(\Omega, \mathcal{F}_\xi, \mu)$ with inner product

$$\langle u, v \rangle = \mathbb{E}[uv] = \int_\Gamma u(\xi) v(\xi) d\mu(\xi), \quad (3.1)$$

where the symbol \mathbb{E} denotes mathematical expectation.

3.1. The stochastic Galerkin formulation. The counterpart of the variational formulation (2.7)–(2.8) consists of performing a Galerkin projection on the space T_Γ using mathematical expectation in the sense of (3.1). That is, we seek the velocity \vec{u} , a random field in $V_E \otimes T_\Gamma$, and the pressure $p \in Q_D \otimes T_\Gamma$, such that

$$\mathbb{E} \left[\int_D \nu \nabla \vec{u} : \nabla \vec{v} + \int_D (\vec{u} \cdot \nabla \vec{u}) \vec{v} - \int_D p (\nabla \cdot \vec{v}) \right] = \mathbb{E} \left[\int_D \vec{f} \cdot \vec{v} \right] \quad \forall \vec{v} \in V_D \otimes T_\Gamma, \quad (3.2)$$

$$\mathbb{E} \left[\int_D q (\nabla \cdot \vec{u}) \right] = 0 \quad \forall q \in Q_D \otimes T_\Gamma. \quad (3.3)$$

The stochastic counterpart of the Newton iteration (2.9)–(2.10) is

$$\mathbb{E} \left[\int_D \nu \nabla \delta \vec{u}^n : \nabla \vec{v} + c(\vec{u}^n; \delta \vec{u}^n, \vec{v}) + c(\delta \vec{u}^n; \vec{u}^n, \vec{v}) - \int_D \delta p^n (\nabla \cdot \vec{v}) \right] = R^n, \quad (3.4)$$

$$\mathbb{E} \left[\int_D q (\nabla \cdot \delta \vec{u}^n) \right] = r^n, \quad (3.5)$$

¹Throughout this study, we use the convention that the right-hand sides of discrete systems incorporate Dirichlet boundary data for velocities.

where

$$R^n(\vec{v}) = \mathbb{E} \left[\int_D \vec{f} \cdot \vec{v} - \int_D \nu \nabla \vec{u}^n : \nabla \vec{v} - c(\vec{u}^n; \vec{u}^n, \vec{v}) + \int_D p^n (\nabla \cdot \vec{v}) \right], \quad (3.6)$$

$$r^n(q) = -\mathbb{E} \left[\int_D q (\nabla \cdot \vec{u}^n) \right]. \quad (3.7)$$

The analogue for Picard iteration omits $c(\delta \vec{u}^n; \vec{u}^n, \vec{v})$ from (3.4):

$$\mathbb{E} \left[\int_D \nu \nabla \delta \vec{u}^n : \nabla \vec{v} + c(\vec{u}^n; \delta \vec{u}^n, \vec{v}) - \int_D \delta p^n (\nabla \cdot \vec{v}) \right] = R^n. \quad (3.8)$$

In computations, we will use a finite-dimensional subspace $T_P \subset T_\Gamma$ spanned by a set of polynomials $\{\psi_\ell(\xi)\}$ that are orthogonal with respect to the density function μ , that is $\langle \psi_k, \psi_\ell \rangle = \delta_{k\ell}$. This is referred to as the gPC basis; see [10, 27] for details and discussion. For T_P , we will use the space spanned by multivariate polynomials in $\{\xi_j\}_{j=1}^N$ of total degree P , which has dimension $M = \binom{N+P}{P}$. We will also assume that the viscosity is given by a gPC expansion

$$\nu = \sum_{\ell=0}^{M_\nu-1} \nu_\ell(x) \psi_\ell(\xi), \quad (3.9)$$

where $\{\nu_\ell(x)\}$ is a set of given deterministic spatial functions.

3.2. Stochastic Galerkin finite element formulation. We discretize (3.4) (or (3.8)) and (3.5) using div-stable finite elements as in Section 2 together with the gPC basis for T_P . For simplicity, we assume that the right-hand side $\vec{f}(x)$ and the Dirichlet boundary conditions (2.3) are deterministic. This means in particular that, as in the deterministic case, the boundary conditions can be incorporated into right-hand side vectors (specified as \mathbf{y} below). Thus, we seek a discrete approximation of the solution of the form

$$\vec{u}(x, \xi) \approx \sum_{k=0}^{M-1} \sum_{i=1}^{N_u} u_{ik} \phi_i(x) \psi_k(\xi) = \sum_{k=0}^{M-1} \vec{u}_k(x) \psi_k(\xi), \quad (3.10)$$

$$p(x, \xi) \approx \sum_{k=0}^{M-1} \sum_{j=1}^{N_p} p_{jk} \varphi_j(x) \psi_k(\xi) = \sum_{k=0}^{M-1} p_k(x) \psi_k(\xi), \quad (3.11)$$

The structure of the discrete operators depends on the ordering of the unknown coefficients $\{u_{ik}\}$, $\{p_{jk}\}$. We will group velocity-pressure pairs for each k , the index of stochastic basis functions (and order equations in the same way), giving the ordered list of coefficients

$$u_{1:N_u,0}, p_{1:N_p,0}, u_{1:N_u,1}, p_{1:N_p,1}, \dots, u_{1:N_u,M-1}, p_{1:N_p,M-1}. \quad (3.12)$$

To describe the discrete structure, we first consider the stochastic version of the Stokes problem (2.5)–(2.6), where the convection term $c(\cdot, \cdot, \cdot)$ is not present in (3.4) and (3.8). The discrete stochastic Stokes operator is built from the discrete components of the vector-Laplacian

$$\mathbf{A}_\ell = [a_{\ell,ab}], \quad a_{\ell,ab} = \left(\int_D \nu_\ell(x) \nabla \phi_b : \nabla \phi_a \right), \quad \ell = 1, \dots, M_\nu - 1, \quad (3.13)$$

which are incorporated into the block matrices

$$\mathcal{S}_0 = \begin{bmatrix} \mathbf{A}_0 & \mathbf{B}^T \\ \mathbf{B} & \mathbf{0} \end{bmatrix}, \quad \mathcal{S}_\ell = \begin{bmatrix} \mathbf{A}_\ell & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}, \quad \ell = 1, \dots, M_\nu - 1. \quad (3.14)$$

These operators will be coupled with matrices arising from terms in T_P ,

$$\mathbf{H}_\ell = [h_{\ell,jk}], \quad h_{\ell,jk} \equiv \mathbb{E}[\psi_\ell \psi_j \psi_k], \quad \ell = 0, \dots, M_\nu - 1, \quad j, k = 0, \dots, M - 1. \quad (3.15)$$

Combining the expressions from (3.13), (3.14) and (3.15) and using the ordering (3.12) gives the discrete stochastic Stokes system

$$\left(\sum_{\ell=0}^{M_\nu-1} \mathbf{H}_\ell \otimes \mathcal{S}_\ell \right) \mathbf{v} = \mathbf{y}, \quad (3.16)$$

where \otimes corresponds to the matrix Kronecker product. The unknown vector \mathbf{v} corresponds to the ordered list of coefficients in (3.12) and the right-hand side is ordered in an analogous way. Note that \mathbf{H}_0 is the identity matrix of order M .

REMARK 3.1. *With this ordering, the coefficient matrix contains a set of M block 2×2 matrices of saddle-point structure along its block diagonal, given by*

$$\mathcal{S}_0 + \sum_{\ell=1}^{M_\nu-1} h_{\ell,jj} \mathcal{S}_\ell, \quad j = 0, \dots, M - 1.$$

This enables the use of existing deterministic solvers for the individual diagonal blocks. An alternative ordering based on the blocking of all velocity coefficients followed by all pressure coefficients, considered in [19], produces a matrix of global saddle-point structure.

The matrices arising from the linearized stochastic Navier-Stokes equations augment the Stokes systems with stochastic variants of the vector-convection matrix and Newton derivative matrix appearing in (2.20). In particular, at step n of the nonlinear iteration, let $\bar{u}_\ell^n(x)$ be the ℓ th term of the velocity iterate (as in the expression on the right in (3.10) for $k = \ell$), and let

$$\begin{aligned} \mathbf{N}_\ell^n &= [n_{\ell,ab}^n], \quad n_{\ell,ab}^n = \int_D (\bar{u}_\ell^n \cdot \nabla \phi_b) \cdot \phi_a, \\ \mathbf{W}_\ell^n &= [w_{\ell,ab}^n], \quad w_{\ell,ab}^n = \int_D (\phi_b \cdot \nabla \bar{u}_\ell^n) \cdot \phi_a. \end{aligned}$$

Then the analogues of (3.13)–(3.14) are

$$\mathbf{F}_\ell^n = \mathbf{A}_\ell + \mathbf{N}_\ell^n + \mathbf{W}_\ell^n, \quad \text{for the stochastic Newton method} \quad (3.17)$$

$$\mathbf{F}_\ell^n = \mathbf{A}_\ell + \mathbf{N}_\ell^n, \quad \text{for stochastic Picard iteration,} \quad (3.18)$$

so for Newton's method

$$\mathcal{F}_0^n = \begin{bmatrix} \mathbf{F}_0^n & \mathbf{B}^T \\ \mathbf{B} & \mathbf{0} \end{bmatrix}, \quad \mathcal{F}_\ell^n = \begin{bmatrix} \mathbf{F}_\ell & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}, \quad (3.19)$$

and as above, for Picard iteration the Newton derivative matrices $\{\mathbf{W}^n\}$ are dropped. Note that $\ell = 0, \dots, \widehat{M} - 1$ here, where $\widehat{M} = \max(M, M_\nu)$. (In particular, if $M_\nu > M$,

we set $\mathbf{N}_\ell^n = \mathbf{W}_\ell^n = 0$ for $\ell = M + 1, \dots, M_\nu - 1$.) Step n of the stochastic nonlinear iteration entails solving a linear system and updating,

$$\left[\sum_{\ell=0}^{\widehat{M}-1} \mathbf{H}_\ell \otimes \mathcal{F}_\ell^n \right] \delta \mathbf{v}^n = \mathcal{R}^n, \quad \mathbf{v}^{n+1} = \mathbf{v}^n + \delta \mathbf{v}^n, \quad (3.20)$$

where

$$\mathcal{R}^n = \mathbf{y} - \left[\sum_{\ell=0}^{\widehat{M}-1} \mathbf{H}_\ell \otimes \mathcal{P}_\ell^n \right] \mathbf{v}^n, \quad (3.21)$$

\mathbf{v}^n and $\delta \mathbf{v}^n$ are vectors of current velocity and pressure coefficients and updates, respectively, ordered as in (3.12), \mathbf{y} is the similarly ordered right-hand side determined from the forcing function and Dirichlet boundary data, and

$$\mathcal{P}_0^n = \begin{bmatrix} \mathbf{A}_0 + \mathbf{N}_0^n & \mathbf{B}^T \\ \mathbf{B} & \mathbf{0} \end{bmatrix}, \quad \mathcal{P}_\ell^n = \begin{bmatrix} \mathbf{A}_\ell + \mathbf{N}_\ell^n & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix};$$

note that the $(1, 1)$ -blocks here are as in (3.18).

3.3. Sampling methods. In experiments described below, we compare some results obtained using stochastic Galerkin methods to those obtained from Monte Carlo and stochastic collocation. We briefly describe these approaches here.

Both Monte Carlo and stochastic collocation methods are based on sampling. This entails the solution of a number of mutually independent deterministic problems at a set of sample points $\{\xi^{(q)}\}$, which give realizations of the viscosity (3.9). That is, a realization of viscosity $\nu(\xi^{(q)})$ gives rise to deterministic functions $\vec{u}(\cdot, \xi^{(q)})$ and $p(\cdot, \xi^{(q)})$ on D that satisfy the standard deterministic Navier-Stokes equations, and to finite-element approximations $\vec{u}^{(q)}(x)$, $p^{(q)}(x)$.

In the Monte Carlo method, the N_{MC} sample points are generated randomly, following the distribution of the random variables ξ , and moments of the solution are obtained from ensemble averaging. For stochastic collocation, the sample points consist of a set of predetermined *collocation points*. This approach derives from a methodology for performing quadrature or interpolation in multidimensional space using a small number of points, a so-called sparse grid [8, 16]. There are two ways to implement stochastic collocation to obtain the coefficients in (3.10)–(3.11), either by constructing a Lagrange interpolating polynomial, or, in the so-called pseudospectral approach, by performing a discrete projection into T_P [26]. We use the second approach because it facilitates a direct comparison with the stochastic Galerkin method. In particular, the coefficients are determined using a quadrature

$$u_{ik} = \sum_{q=1}^{N_q} \vec{u}^{(q)}(x_i) \psi_k(\xi^{(q)}) w^{(q)}, \quad p_{ik} = \sum_{q=1}^{N_q} p^{(q)}(x_i) \psi_k(\xi^{(q)}) w^{(q)},$$

where $\xi^{(q)}$ are collocation (quadrature) points, and $w^{(q)}$ are quadrature weights. We refer, e.g., to [14] for an overview and discussion of integration rules.

3.4. Example: flow around an obstacle. In this section, we present results of numerical experiments for a model problem given by a flow around an obstacle

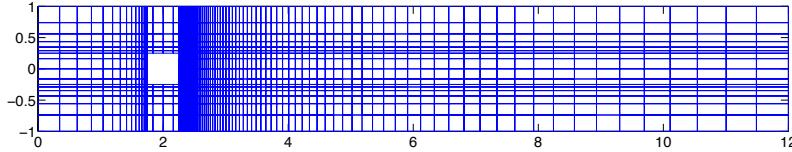


FIG. 3.1. Finite element mesh for the flow around an obstacle problem.

in a channel of length 12 and height 2. The viscosity (3.9) was taken to be a truncated lognormal process with mean values $\nu_0 = 1/50$ or $1/150$, which corresponds to mean Reynolds numbers $\text{Re}_0 = 100$ or 300 , respectively, and its representation was computed from an underlying Gaussian random process using the transformation described in [9]. That is, for $\ell = 0, \dots, M_\nu - 1$, $\psi_\ell(\xi)$ is the product of N univariate Hermite polynomials, and denoting the coefficients of the Karhunen-Loëve expansion of the Gaussian process by $g_j(x)$ and $\eta_j = \xi_j - g_j$, $j = 1, \dots, N$, the coefficients in the expansion (3.9) are computed as

$$\nu_\ell(x) = \mathbb{E}[\psi_\ell(\eta)] \exp \left[g_0(x) + \frac{1}{2} \sum_{j=1}^N (g_j(x))^2 \right].$$

The covariance function of the Gaussian field, for points $X_1 = (x_1, y_1)$, $X_2 = (x_2, y_2) \in D$, was chosen to be

$$C(X_1, X_2) = \sigma_g^2 \exp \left(-\frac{|x_2 - x_1|}{L_x} - \frac{|y_2 - y_1|}{L_y} \right),$$

where L_x and L_y are the correlation lengths of the random variables ξ_i , $i = 1, \dots, N$, in the x and y directions, respectively, and σ_g is the standard deviation of the Gaussian random field. The correlation lengths were set to be equal to 25% of the width and height of the domain, i.e. $L_x = 3$ and $L_y = 0.5$. The coefficient of variation of the lognormal field, defined as $\text{CoV} = \sigma_\nu/\nu_0$ where σ_ν is the standard deviation, was 10% or 30%. The stochastic dimension was $N = 2$. The lognormal formulation was chosen to enable exploration of a general random field in which the viscosity guaranteed to be positive. See [1] for an example of the use of this formulation for diffusion problems and [28] for its use in models of porous media.

We implemented the methods in **Matlab** using **IFISS 3.3** [5]. The spatial discretization uses a stretched grid, discretized by 1520 Taylor-Hood finite elements; the domain and grid are shown in Figure 3.1. There are 12,640 velocity and 1640 pressure degrees of freedom. the degree used for the polynomial expansion of the solution was $P = 3$, and the degree used for the expansion of the lognormal process was $2P$, which ensures a complete representation of the process in the discrete problem [15]. With these settings, $M = 10$ and $M_\nu = \widehat{M} = 28$, and \mathbf{H}_ℓ is of order 10 in (3.20).

Consider first the case of $\text{Re}_0 = 100$ and $\text{CoV} = 10\%$. Figure 3.2 shows the mean horizontal and vertical components of the velocity and the mean pressure (top), and the variances of the same quantities (bottom). It can be seen that there is symmetry in all the quantities, the mean values are essentially the same as we would expect in the deterministic case, and the variance of the horizontal velocity component is concentrated in two “eddies” and is larger than the variance of the vertical velocity

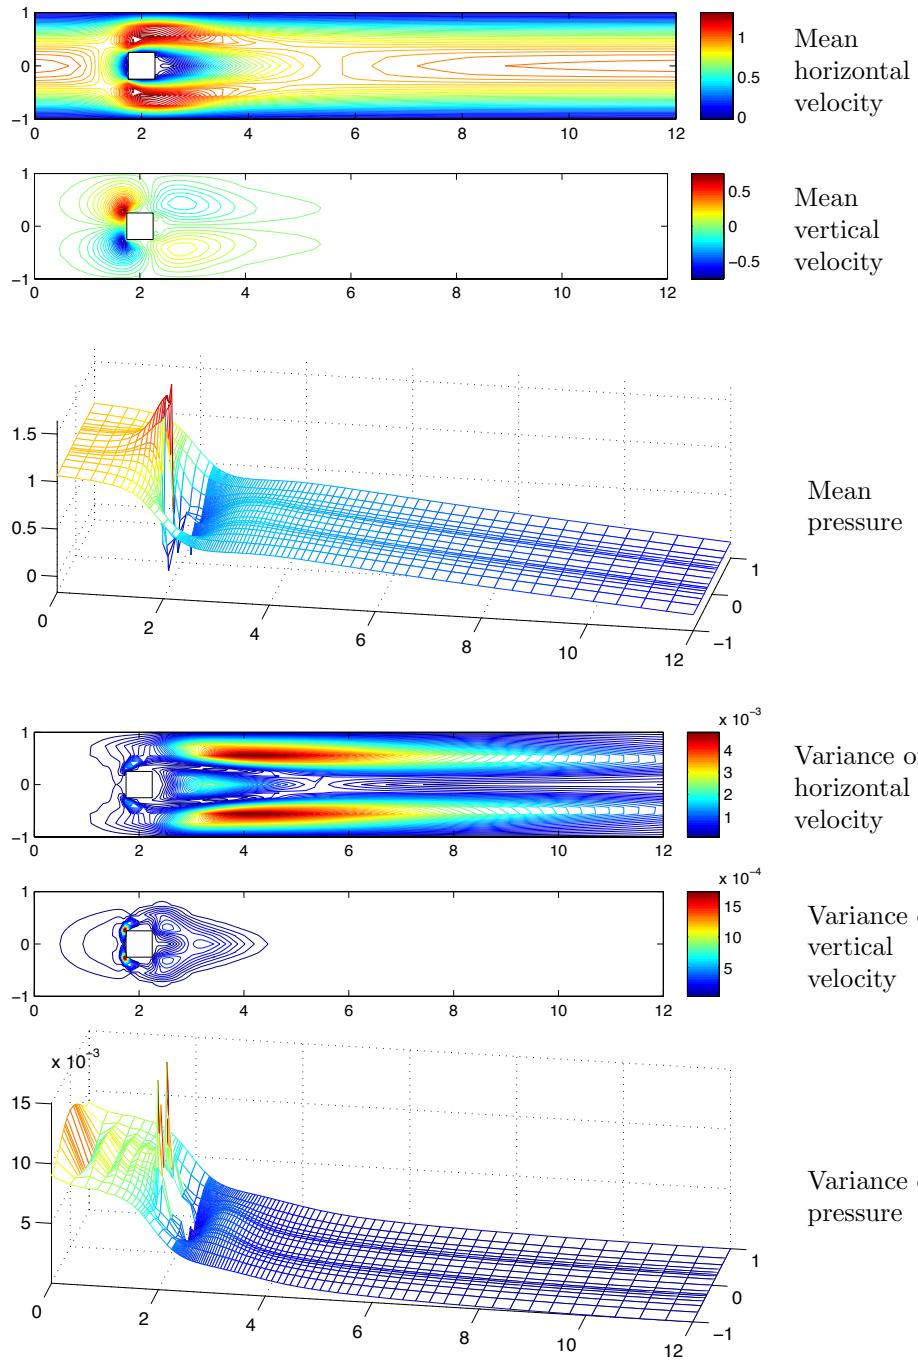


FIG. 3.2. Mean horizontal and vertical velocities and pressure (top) and variances of the same quantities (bottom), for $Re = 100$ and $CoV = 10\%$.

component. Figure 3.3 illustrates values of several coefficients of expansion (3.10) of the horizontal velocity. All the coefficients are symmetric, and as the index increases they become more oscillatory and their values decay. We found the same trends for the coefficients of the vertical velocity component and of the pressure. Our observations are qualitatively consistent with numerical experiments of Powell and Silvester [19].

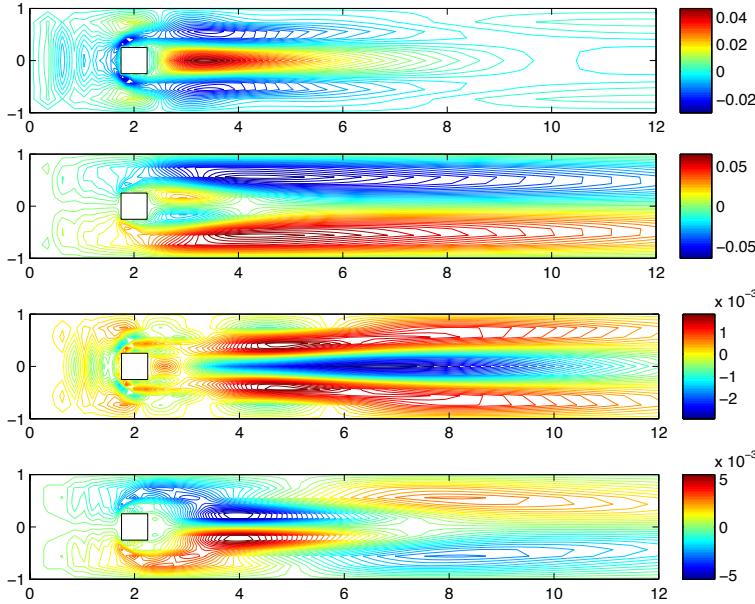


FIG. 3.3. Coefficients 1 – 4 of the gPC expansion of the horizontal velocity, $Re_0 = 100$ and $CoV = 10\%$.

We also tested (the same) $Re_0 = 100$ with increased $CoV = 30\%$. We found that the mean values are essentially the same as in the previous case; Figure 3.4 shows the variances, which display the same qualitative behavior but have values that are approximately 10 times larger than for the case $CoV = 10\%$.

A different perspective on the results is given in Figure 3.5, which shows estimates of the probability density function (pdf) for the horizontal velocity at two points in the domain, $(4.0100, -0.4339)$ and $(4.0100, 0.4339)$. These are locations at which large variances of the solution were seen, see Figures 3.2 and 3.4. The results were obtained using Matlab's `ksdensity` function. It can be seen that with the larger value of CoV , the support of the velocity pdf is wider, and except for the peak values, for fixed CoV the shapes of the pdfs at the two points are similar, indicating a possible symmetry of the stochastic solution. For this benchmark, we also obtained analogous data using the Monte Carlo and collocation sampling methods; it can be seen from the figure that these methods produced similar results. ²

Next, we consider a larger value of the mean Reynolds number, $Re_0 = 300$. Figure 3.6 shows the means and variances for the velocities and pressure for $CoV = 10\%$. It is evident that increased Re_0 results in increased values of the mean quantities, but they are again similar to what would be expected in the deterministic case. The

²The results for Monte Carlo were obtained using 10^3 samples, and those for collocation were found using a Smolyak sparse grid with Gauss-Hermite quadrature and grid level $n_q = 4$.

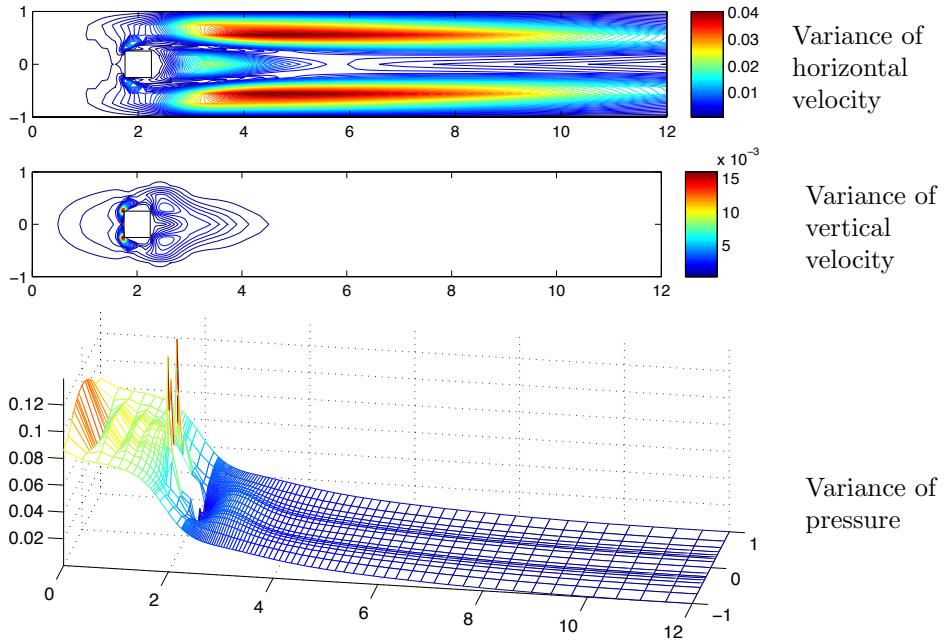


FIG. 3.4. Variances of velocity components and pressure for $Re_0 = 100$ and $CoV = 30\%$.

variances exhibit wider eddies than for $Re_0 = 100$, and in this case there is only one region of the largest variance in the horizontal velocity, located just to the right of the obstacle; this is also a region with increased variance of the pressure.

In similar tests for the larger value $CoV = 30\%$, we found that the mean values are essentially the same as for $CoV = 10\%$, and Figure 3.7 shows the variances of velocities and pressures. From the figure it can be seen that the variances are qualitatively the same but approximately 10 times larger than for $CoV = 10\%$, results similar to those found for $Re_0 = 100$.

As above, we also examined estimated probability density functions for the velocities and pressures at a specified point in the domain, in this case, at the point $(3.6436, 0)$, taken again from the region in which the solution has large variance. Figure 3.8 shows these pdf estimates, from which it can be seen that the three methods for handling uncertainty are in close agreement for each of the two values of CoV .

Finally, we show in Figure 3.9 the results of one additional experiment with $Re_0 = 300$ and $CoV = 10\%$, where estimated pdfs of the first velocity component u_x were computed at three points near the inflow boundary, $(x, y) = (0.5, 0)$, $(1, 0)$ and $(1.5, 0)$. These plots show some effects of spatial accuracy as well as differences between the stochastic Galerkin method and sampling methods, as represented by collocation (results for Monte Carlo were identical). The images on the top and bottom left show results for a uniform mesh of width $h = 1/8$ and for two refined meshes in which the horizontal mesh width to the left of the obstacle is reduced to $h/2$ and $h/4$. The image in the bottom right provides a more detailed view of the fine-grid results; in this image, the width of the horizontal window is the same (0.01) for the three subplots but the vertical heights are different. Several trends are evident:

- The differences between the Galerkin and sampling results become less pronounced as the mesh is refined, and at points further from the inflow.

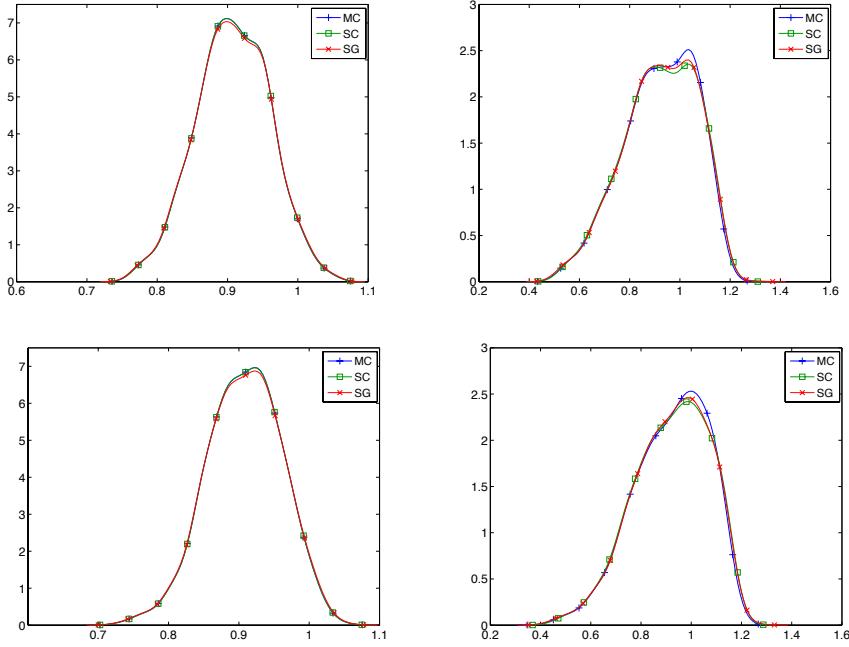


FIG. 3.5. Estimated pdfs of the velocities u_x with $Re_0 = 100$, $CoV = 10\%$ (left) and 30% (right) at the points with coordinates $(4.0100, -0.4339)$ (top) and $(4.0100, 0.4339)$ (bottom).

- The pdfs at points nearer the boundary are much narrower in general, and this is more dramatically so for collocation.
- When the spatial accuracy is low, both methods miss some features of the pdf. The Galerkin method produces a pdf that is overly smooth. The sampling methods produce a pdf that captures its “spiky” nature near the inflow, but the location of the spike is not correct.

We believe these effects stem from the fact that the inflow boundary is deterministic (with $u_y = 1$ at $y = 0$), and the effects of variability in the viscosity are felt less strongly at points near the inflow boundary. Moreover, the coupled nature of the Galerkin system tends to have a smoothing effect on the pdf near the inflow, especially when spatial accuracy is low. At points further from the deterministic inflow boundary, these effects and differences become less dramatic.

3.5. Nonlinear solvers. We briefly comment on the nonlinear solution algorithm used to generate the results of the previous section. The nonlinear solver was implemented by modifying the analogue for deterministic systems in IFISS. It uses a hybrid strategy in which an initial approximation is obtained from solution of the stochastic Stokes problem (3.16), after which several steps of Picard iteration (equation (3.20) with \mathcal{F}_ℓ specified using (3.19) and (3.18)) are used to improve the solution, followed by Newton iteration (\mathcal{F}_ℓ from (3.17)). A convergent iteration stopped when the Euclidian norm of the algebraic residual (3.21) satisfied $\|\mathcal{R}^n\|_2 \leq \epsilon \|y\|_2$ where $\epsilon = 10^{-8}$ and y is as in (3.16).

In the experiments described in Section 3.4, we used values of the Reynolds number, $Re = 100$ and 300 , and for each of these, two values of the coefficient of variation,

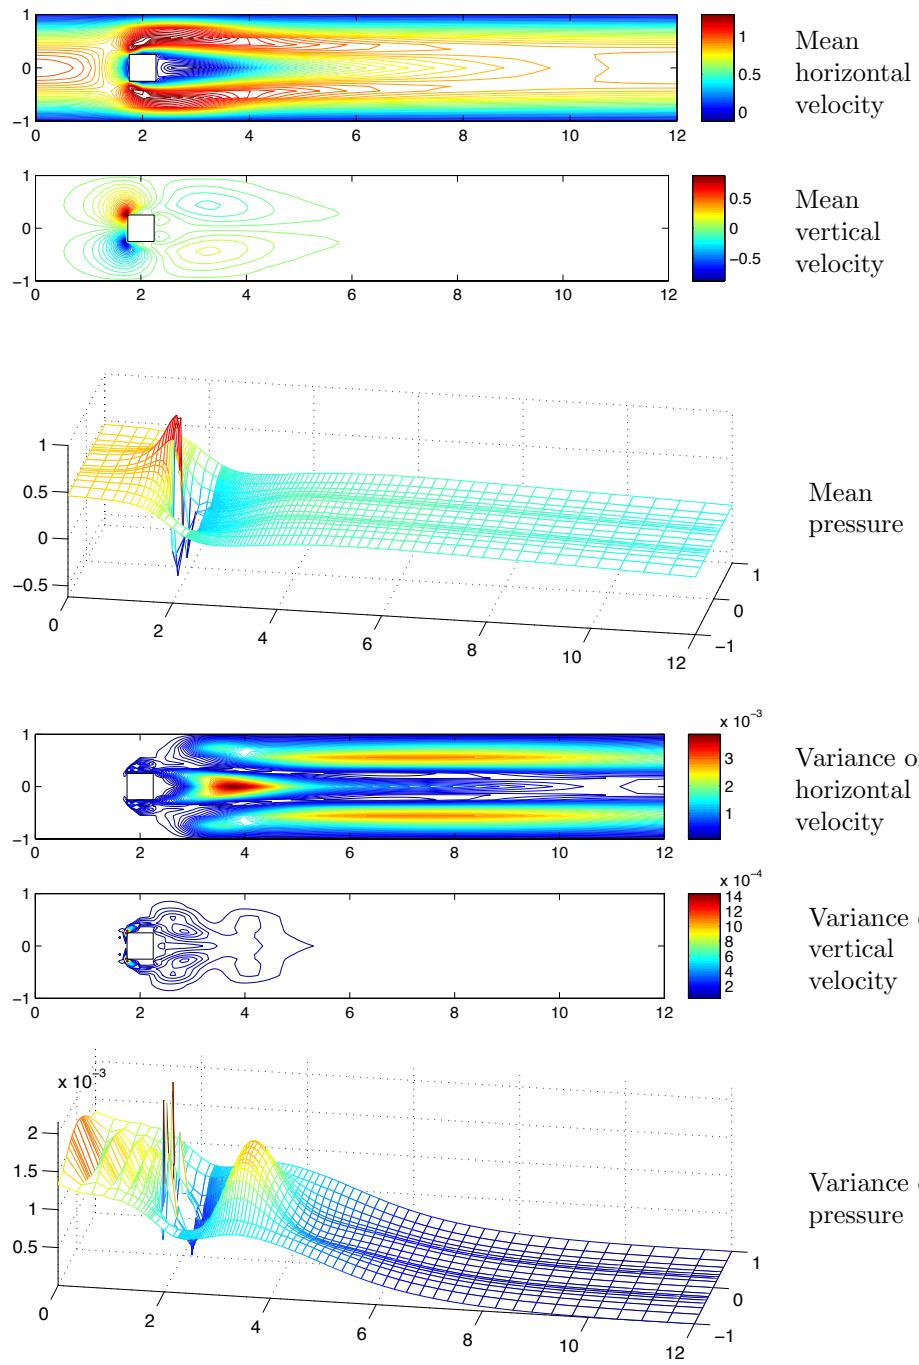


FIG. 3.6. Mean horizontal and vertical velocities and pressure (top) and variances of the same quantities (bottom), for $Re_0 = 300$ and $CoV = 10\%$.

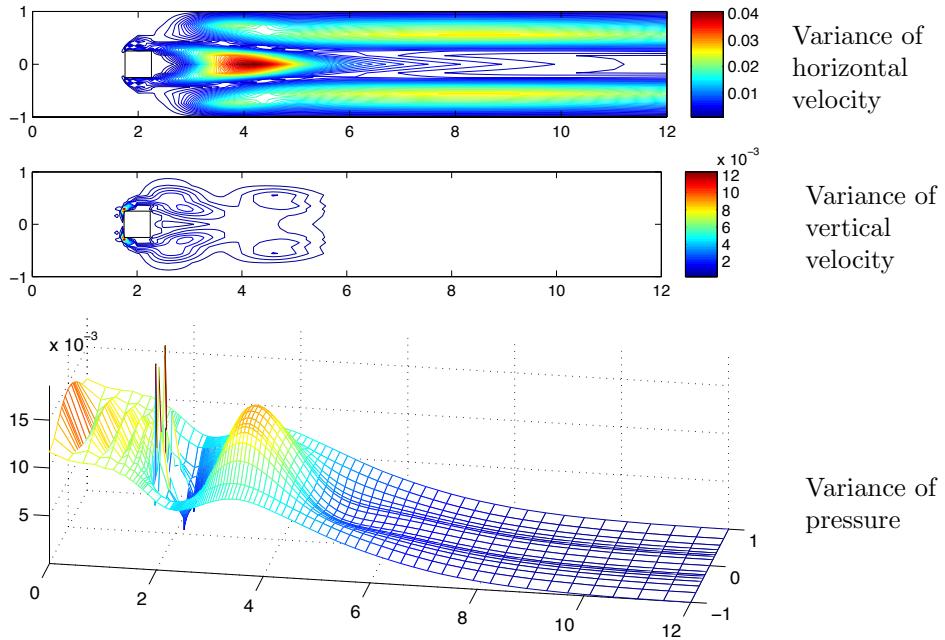


FIG. 3.7. Variances of velocity components and pressure for $Re_0 = 300$ and $CoV = 30\%$.

$CoV = 10\%$ and 30% . We list here the numbers of steps leading to convergence of the nonlinear algorithms that were used to generate the solutions discussed above. Direct solvers were used for the linear systems; we discuss a preconditioned iterative algorithm in Section 4 below.

$Re = 100, CoV = 10\%$:	6 Picard steps	1 Newton step(s)
$Re = 100, CoV = 30\%$:	6	3
$Re = 300, CoV = 10\%$:	20	1
$Re = 300, CoV = 30\%$:	20	2

Thus, a larger CoV (larger standard deviation of the random field determining uncertainty in the process) leads to somewhat larger computational costs. For $Re = 300$, the nonlinear iteration was not robust with 6 initial Picard steps (for the stochastic Galerkin method as well as the sampling methods); 20 steps was sufficient.

We also explored an *inexact* variant of these methods, in which the coefficient matrix of (3.20) for the Picard iteration was replaced by the block diagonal matrix $\mathbf{H}_0 \otimes \mathcal{F}_0^n$ obtained from the mean coefficient. For $CoV = 10\%$, with the same number of (now inexact) Picard steps as above (6 for $Re = 100$ and 20 for $Re = 300$), this led to just one extra (exact) Newton step for $Re = 100$ and no additional steps for $Re = 300$. On the other hand, for $CoV = 30\%$, this inexact method failed to converge.

4. Preconditioner for the linearized systems. The solution of the linear systems required during the course of the nonlinear iteration is a computationally intensive task, and use of direct solvers may be prohibitive for large problems. In this section, we present a preconditioning strategy for use with Krylov subspace methods to solve these systems, and we compare its performance with that of several other techniques. The new method is a variant of the hierarchical Gauss-Seidel preconditioner developed in [22].

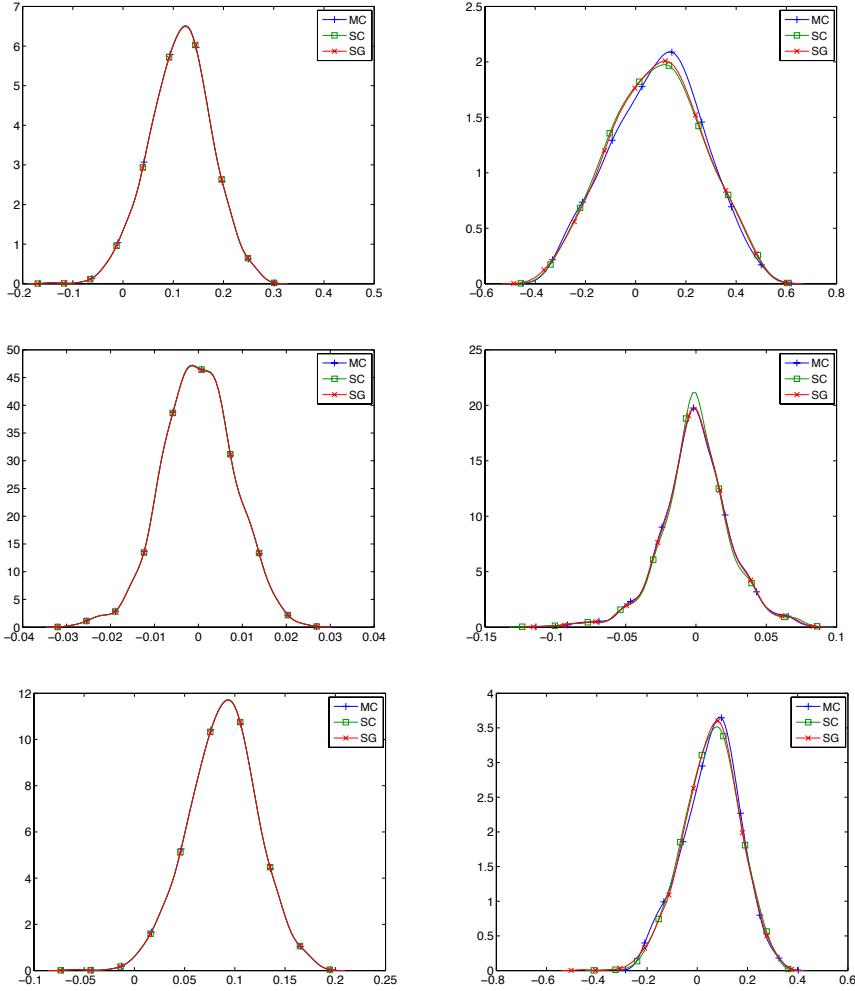


FIG. 3.8. Estimated pdfs of the velocities u_x (top), u_y (middle), and pressure p (bottom) with $Re_0 = 300$ and $CoV = 10\%$ (left) and 30% (right) at the point with coordinates $(3.6436, 0)$.

4.1. Structure of the matrices and the preconditioner. We first recall the structure of the matrices $\{\mathbf{H}_\ell\}$ of (3.15). More comprehensive overviews of these matrices can be found in [7, 15]. The matrix structure can be understood through knowledge of the coefficient matrix $c_P \equiv \sum_{\ell=0}^{M_\nu-1} h_{\ell,jk}$ where $j, k = 0, \dots, M - 1$. The block sparsity structure depends on the type of coefficient expansion in (3.9). If only linear terms are included, that is $\psi_\ell = \xi_\ell$, $\ell = 1, \dots, N$, then the coefficients $h_{\ell,jk} = \mathbb{E}[\xi_\ell \psi_j \psi_k]$ yield a Galerkin matrix with a block sparse structure. In the more general case, $h_{\ell,jk} = \mathbb{E}[\psi_\ell \psi_j \psi_k]$ and the stochastic Galerkin matrix becomes fully block dense. In either case, for fixed ℓ and a set of degree \mathcal{P} polynomial expansions, with $1 \leq \mathcal{P} \leq P$, the corresponding coefficient matrix $c_{\mathcal{P}}$ has a hierarchical structure

$$c_{\mathcal{P}} = \begin{bmatrix} c_{\mathcal{P}-1} & b_{\mathcal{P}}^T \\ b_{\mathcal{P}} & d_{\mathcal{P}} \end{bmatrix}, \quad \mathcal{P} = 1, \dots, P.$$

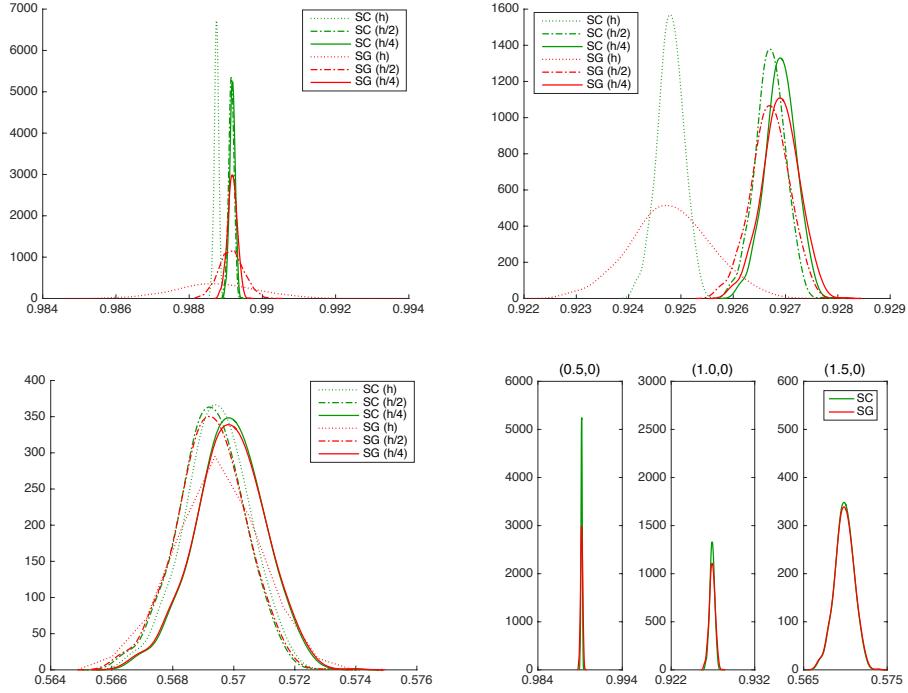


FIG. 3.9. Estimated pdfs of the velocity u_x at points $(0.5, 0)$ (top left) $(1, 0)$ (top right) and $(1.5, 0)$ (bottom left), with $\text{Re}_0 = 300$ and $\text{CoV} = 10\%$ and three meshes. Bottom right: detailed depiction of the pdfs for each grid point obtained from the finest mesh.

Now, let \mathcal{A}_P denote the global stochastic Galerkin matrix corresponding to either a Stokes problem (3.16) or a linearized system (3.20); we will focus on the latter system in the discussion below. The matrix \mathcal{A}_P also has a hierarchical structure

$$\mathcal{A}_P = \begin{bmatrix} \mathcal{A}_{P-1} & \mathcal{B}_P \\ \mathcal{C}_P & \mathcal{D}_P \end{bmatrix}, \quad P = 1, \dots, P, \quad (4.1)$$

where \mathcal{A}_0 is the matrix of the mean, derived from ν_0 in (3.9). This hierarchical structure is shown in the left side of Figure 4.1.

We will write vectors with respect to this hierarchy as

$$x_{(0:P)} = \begin{bmatrix} x_{(0)} \\ x_{(1)} \\ \vdots \\ x_{(P)} \end{bmatrix},$$

where $x_{(q)}$ includes all indices corresponding to polynomial degree q , blocked by spatial ordering determined by (3.12). With this notation, the global stochastic Galerkin linear system has the form

$$\mathcal{A}_P x_{(0:P)} = f_{(0:P)}. \quad (4.2)$$

To formulate the preconditioner for (4.2), we let $\tilde{\mathcal{A}}_0$ represent an approximation

of \mathcal{A}_0 and $\tilde{\mathcal{D}}_{\mathcal{P}}$ represent an approximation of $\mathcal{D}_{\mathcal{P}}$. In particular, let

$$\tilde{\mathcal{D}}_{\mathcal{P}} = \begin{bmatrix} \tilde{\mathcal{A}}_0 & & \\ & \ddots & \\ & & \tilde{\mathcal{A}}_0 \end{bmatrix}, \quad (4.3)$$

where the number of diagonal blocks is given by \mathcal{P} . We will need the action of the inverse of $\tilde{\mathcal{D}}_{\mathcal{P}}$, or an approximation to it, which can be obtained using an LU-factorization of $\tilde{\mathcal{A}}_0$, or using some preconditioner for $\tilde{\mathcal{A}}_0$, or using a Krylov subspace solver. A preconditioner $\mathcal{P} : w_{(0:\mathcal{P})} \rightarrow v_{(0:\mathcal{P})}$ for (4.2) is then defined as follows:

ALGORITHM 4.1. [Approximate hierarchical Gauss-Seidel preconditioner (ahGS)]
Solve (or solve approximately)

$$\tilde{\mathcal{A}}_0 v_{(0)} = w_{(0)}, \quad (4.4)$$

and, for $\mathcal{P} = 1, \dots, P$, solve (or solve approximately)

$$\tilde{\mathcal{D}}_{\mathcal{P}} v_{(\mathcal{P})} = (w_{(\mathcal{P})} - \mathcal{C}_{\mathcal{P}} v_{(0:\mathcal{P}-1)}). \quad (4.5)$$

The cost of preconditioning can be reduced further by truncating the matrix-vector (MATVEC) operations used for the multiplications by the submatrices $\mathcal{C}_{\mathcal{P}}$ in (4.5). The idea is as follows. The system (4.2) can be written as

$$\sum_{j=0}^{M-1} \sum_{\ell=0}^{M_{\nu}-1} h_{\ell,jk} \mathcal{F}_{\ell} x_j = f_k, \quad k = 0, \dots, M-1, \quad (4.6)$$

and the MATVEC with \mathcal{A}_P is given by

$$v_j = \sum_{k=0}^{M-1} \sum_{\ell=0}^{M_{\nu}-1} h_{\ell,jk} \mathcal{F}_{\ell} u_k, \quad (4.7)$$

where the indices $j, k \in \{0, \dots, M-1\}$ correspond to nonzero blocks in \mathcal{A}_P . The truncated MATVEC is an inexact evaluation of (4.7) proposed in [22, Algorithm 1], in which the summation over $\ell = 0, \dots, M_{\nu}-1$ is replaced by summation over a subset $\mathcal{M}_t \subseteq \{0, \dots, M_{\nu}-1\}$. Figure 4.1 shows the hierarchical structure of the matrix and of the ahGS preconditioning operator. Both images in the figure correspond to the choice $P = 3$, so that the hierarchical preconditioning operation (4.4)–(4.5) requires four steps. Because $N = 4$, the matrix block size is $M = \binom{N+P}{P} = 35$. The block-lower-triangular component of the image on the right in Figure 4.1 shows the hierarchical structure of the ahGS preconditioning operator with truncation. For the matrix in the left panel, $M_{\nu} = \binom{N+2P}{2P} = 210$, but the index set \mathcal{M}_t includes terms with indices at most $M-1$ in the accumulation of sums used for $\mathcal{C}_{\mathcal{P}}$. These two heuristics, approximation by (4.3) and the truncation of MATVECs, significantly improve the sparsity structure of the preconditioner, on both the block diagonal (through the first technique) and the block lower triangle (through the second). In the next section, we will also consider truncated MATVECs with smaller maximal indices.

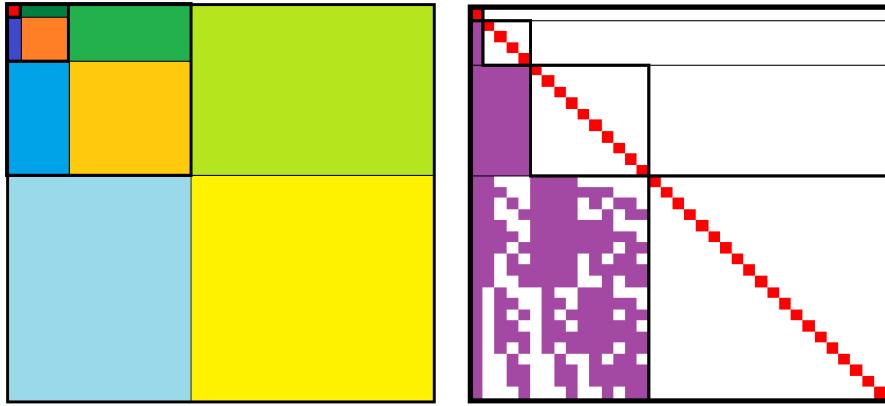


FIG. 4.1. *Hierarchical structure of the stochastic Galerkin matrix (4.1) (left) and splitting operator $L + \text{diag}(D)$ for the approximate hierarchical Gauss-Seidel preconditioner (ahGS) with the truncation of the MATVEC (right).*

4.2. Numerical experiments. In this section, we describe the results of experiments in which the ahGS preconditioner is used with GMRES to solve the systems arising from both Picard and Newton iteration.³ Unless otherwise specified, in the experiments discussed below, the block-diagonal matrices of (4.3), (4.4), (4.5) are taken to be $\tilde{\mathcal{A}}_0 = \mathcal{A}_0$, which corresponds to the first summand $\mathbf{H}_0 \otimes \mathcal{F}_0^n$ in (3.20) and represents an approximation to the diagonal blocks of \mathcal{D}_P ; see also Remark 3.1. We also compare the performance of ahGS preconditioning with several alternatives:

- Mean-based preconditioner (MB) [17, 18], where the preconditioning operator is the block-diagonal matrix $\mathbf{H}_0 \otimes \mathcal{F}_0^n = I \otimes \mathcal{F}_0^n$ derived from the mean of ν .
- The Kronecker-product preconditioner [25] (denoted K), given by $\widehat{\mathbf{H}}_0 \otimes \mathcal{F}_0^n$, where $\widehat{\mathbf{H}}_0$ is chosen to minimize a measure of the difference $\mathcal{A}_P - \widehat{\mathbf{H}}_0 \otimes \mathcal{F}_0^n$.
- Block Gauss-Seidel preconditioner (bGS), in which the preconditioning operation entails applying $(\tilde{\mathcal{D}}_P + L)^{-1}$ determined using the block lower triangle of \mathcal{A}_P , that is with the approximation of the diagonal blocks by \mathcal{A}_0 but without MATVEC truncation. This is an expensive operator but enables an understanding of the impact of various efforts to make the preconditioning operator more sparse.
- ahGS(PCD), a modification of ahGS in which the block-diagonal matrices of (4.4)–(4.5) are replaced by the *pressure convection-diffusion* (PCD) approximation [6] to the mean matrix \mathcal{F}_0 .
- ahGS(PCD-it), in which the block-diagonal solve of (4.5) is replaced by an approximate solve determined by twenty steps of a PCD-preconditioned GMRES iteration. For this strategy, $\tilde{\mathcal{A}}_0$ in (4.4) and the blocks of $\tilde{\mathcal{D}}_P$ in (4.5) are taken to be the complete sum of (3.20), but the approximate solutions to these systems are obtained using a preconditioned inner iteration.

Four of the strategies, the ahGS, mean-based, Kronecker-product and bGS preconditioners, require the solution of a set of block-diagonal systems with the structure

³The preconditioning operator may be nonlinear, for example, if the block solves in (4.4)–(4.5) are performed approximately using Krylov subspace methods, so that care must be made in the choice of the Krylov subspace iterative method. For this reason, we used a flexible variant of the GMRES method [21].

of a linearized Navier-Stokes operator of the form given in (2.17). We used direct methods for these computations. All results presented are for $\text{Re} = 100$; performance for $\text{Re} = 300$ were essentially the same. We believe this is because of the exact solves performed for the mean operators.

The results for the first step of Picard iteration are in Tables 4.1–4.4. All tests started with a zero initial iterate and stopped when the residual $r^{(k)} = f_{(0:P)} - \mathcal{A}_P x_{(0:P)}^{(k)}$ for the k 'th iterate satisfied $\|r^{(k)}\|_2 \leq 10^{-8} \|f_{(0:P)}\|_2$ in the Euclidian norm. With other parameters fixed and no truncation of the MATVEC, Table 4.1 shows the dependence of GMRES iterations on the stochastic dimension N , Table 4.2 shows the dependence on the degree of polynomial expansion P , and Table 4.3 shows the dependence on the coefficient of variation CoV . It can be seen that the numbers of iterations with the ahGS preconditioner are essentially the same as with the block Gauss-Seidel (bGS) preconditioner, and they are much smaller compared to the mean-based (MB) and the Kronecker product preconditioners. When the exact solves with the mean matrix are replaced by the mean-based modified pressure-convection-diffusion (PCD) preconditioner for the diagonal block solves, the iterations grow rapidly. **On the other hand, if the PCD preconditioner is used as part of an inner iteration as in ahGS(PCD-it), then good performance is recovered.** This indicates that a good preconditioner for the mean matrix is an essential component of the global preconditioner for the stochastic Galerkin matrix.

Table 4.4 shows the iteration counts when the MATVEC operation is truncated in the action of the preconditioner. Truncation decreases the cost per iteration of the computation, and it can be also seen that performance can actually be improved. For example, with $\ell_t = 2$, the number of iterations is the smallest. Moreover there are only 21 nonzeros in the lower triangular part of the sum of coefficient matrices $\{\mathbf{H}_\ell\}$ (each of which has order 10) used in the MATVEC with $\ell_t \leq 2$, compared to 63 nonzeros when no truncation is used; there are 203 nonzeros in the set of 28 full matrices $\{\mathbf{H}_\ell\}$.

TABLE 4.1

For Picard step: dependence on stochastic dimension N of GMRES iteration counts, for various preconditioners, with polynomial degree $P = 3$ and coefficient of variation $\text{CoV} = 30\%$. M is the block size of the stochastic Galerkin matrix, M_ν the number of terms in (3.9) and ngdof the size of the stochastic Galerkin matrix.

N	M	M_ν	ngdof	MB	K	bGS	ahGS	ahGS(PCD)	ahGS(PCD-it)
1	4	7	57,120	63	36	30	30	201	29
2	10	28	142,800	102	66	58	54	357	48
3	20	84	285,600	145	109	88	82	553	73

Results for the first step of Newton iteration, which comes after six steps of Picard iteration, are summarized in Tables 4.5–4.8. As above, the first three tables show the dependence on stochastic dimension N (Table 4.5), polynomial degree P (Table 4.6), and coefficient of variation CoV (Table 4.7). It can be seen that all iteration counts are higher compared to corresponding results for the Picard iteration, but the other trends are very similar. In particular, the performances of the ahGS and bGS preconditioners are comparable except for the case when $N = 3$, and $P = 3$ (the last rows in Tables 4.5 and 4.6). Nevertheless, checking results in Table 4.8, which shows the effect of truncation, it can be seen that with the truncation of the MATVEC the iteration counts of the ahGS and bGS preconditioners can be further

TABLE 4.2

For Picard step: dependence on polynomial degree P of GMRES iteration counts, for various preconditioners, with stochastic dimension $N = 3$ and coefficient of variation $CoV = 30\%$. Other headings are as in Table 4.1.

P	M	M_ν	$ngdof$	MB	K	bGS	ahGS	ahGS(PCD)	ahGS(PCD-it)
1	4	10	57,120	26	22	12	11	144	13
2	10	35	142,800	63	49	29	26	300	29
3	20	84	285,600	145	109	88	82	553	73

TABLE 4.3

For Picard step: dependence on coefficient of variation CoV of GMRES iteration counts, for various preconditioners, with stochastic dimension $N = 2$ and polynomial degree $P = 3$. Other headings are as in Table 4.1.

$CoV(\%)$	MB	K	bGS	ahGS	ahGS(PCD)	ahGS(PCD-it)
10	16	14	7	7	186	10
20	36	27	17	16	259	17
30	102	66	58	54	357	48

improved. Indeed, running one more experiment for the aforementioned case when $N = 3$, and $P = 3$, it turns out that with $\ell_t = 2$ the number of iterations with the ahGS and bGS preconditioners are 118 and 120, respectively and with $\ell_t = 3$ they are 101 and 104, respectively. That is, the truncation leads to fewer iterations also in this case, and the performance of ahGS and bGS preconditioners is again comparable.

Finally, we briefly discuss computational costs. For any preconditioner, each GMRES step entails a matrix-vector product by the coefficient matrix. For viscosity given by a general probability distribution, this will typically involve a block-dense matrix, and, ignoring any overhead associated with increasing the number of GMRES steps, this will be the dominant cost per step. The mean-based preconditioner requires the action of the inverse of the block-diagonal matrix $I \otimes \mathcal{F}_0^n$. This has relatively small amount of overhead once the factors of \mathcal{F}_0^n are computed. The ahGS preconditioner without truncation effectively entails a matrix-vector product by the block lower-triangular part of the coefficient matrix, so its overhead is bounded by 50% of the cost of a multiplication by the coefficient matrix. This is an overestimate because it ignores the approximation of the block diagonal and the effect of truncation. For example, consider the case with stochastic dimension $N = 2$ and polynomial expansions of degree $P = 3$ for the solution and $2P = 6$ for the viscosity; this gives $M = 10$ and $M_\nu = 28$. In Tables 4.4 and 4.8, M_t indicates how many matrices are used in the MATVEC operations and $nnz(c_\ell)$ is the number of nonzeros in the sum of the lower triangular parts of the coefficient matrices $\{\mathbf{H}_\ell \mid \ell = 0, \dots, M_t - 1\}$. With complete truncation, $\ell_t = 0$, and ahGS reduces to the mean-based preconditioner. With no truncation, $\ell_t = 6$, and because the number of nonzeros in $\{\mathbf{H}_\ell\}$ is 203, the overhead of ahGS is $63/203$, less than 30% of the cost of multiplication by the coefficient matrix. If truncation is used, in particular when the iteration count is the lowest ($\ell_t = 2$), the overhead is only $21/203 \doteq 10\%$. Note that with increasing stochastic dimension and degree of polynomial expansion, the savings will be higher because the ratio of the sizes of the blocks $\mathcal{C}_P/\mathcal{D}_P$ decreases as P increases, see (4.1). Last, the mean-based preconditioner is embarrassingly parallel; the ahGS preconditioner requires $P+1$ sequential steps, although each of these steps is also highly parallelizable.

TABLE 4.4

For Picard step: number of GMRES iterations when the preconditioners use truncated MATVEC. The matrices corresponding to higher degree expansion of the coefficient than ℓ_t are dropped from the action of the preconditioner. Here $N = 2$, $P = 3$, M_t is the number of terms used in the inexact (truncated) evaluation of the MATVEC (4.7), $nnz(c_\ell)$ is the number of nonzeros in the lower triangular parts of the coefficient matrices (3.15) after the truncation. Other headings are as in Table 4.1.

	ℓ_t	0	1	2	3	6
setup	M_t	1	3	6	10	28
	$nnz(c_\ell)$	0	12	21	43	63
bGS						
CoV (%)	10	16	9	7	7	7
	20	36	19	15	17	17
	30	102	55	43	57	58
ahGS						
CoV (%)	10	16	9	7	7	7
	20	36	19	14	16	16
	30	102	55	35	55	54

TABLE 4.5

For Newton step: dependence on stochastic dimension N of GMRES iteration counts, for various preconditioners, with polynomial degree $P = 3$ and coefficient of variation $CoV = 30\%$. Other headings are as in Table 4.1.

N	M	M_ν	$ngdof$	MB	K	bGS	ahGS	ahGS(PCD)	ahGS(PCD-it)
1	4	7	57,120	73	42	32	32	309	40
2	10	28	142,800	126	80	77	69	546	63
3	20	84	285,600	235	170	128	151	1011	117

The Kronecker preconditioner is more difficult to assess because it does not have block-diagonal structure, and we do not discuss it here.

5. Conclusion. We studied the Navier-Stokes equations with stochastic viscosity given in terms of polynomial chaos expansion. We formulated the stochastic Galerkin method and proposed its numerical solution using a stochastic versions of Picard and Newton iteration, and we also compared its performance in terms of accuracy with that of stochastic collocation and Monte Carlo method. Finally, we presented a methodology of Gauss-Seidel hierarchical preconditioning with approximation using the mean-based diagonal block solves and a truncation of the MATVEC operations. The advantage of this approach is that neither the matrix nor the preconditioner need to be formed explicitly, and the ingredients include only the matrices from the polynomial chaos expansion and a good preconditioner for the mean-value deterministic problem, it allows an obvious parallel implementation, and it can be written as a “wrapper” around existing deterministic code.

REFERENCES

- [1] I. Babuška, F. Nobile, and R. Tempone. A stochastic collocation method for elliptic partial differential equations with random input data. *SIAM Journal on Numerical Analysis*, 45:1005–1034, 2007.
- [2] M. Brezina, A. Doostan, T. Manteuffel, S. McCormick, and J. Ruge. Smoothed aggregation

TABLE 4.6

For Newton step: dependence on polynomial degree P of GMRES iteration counts, for various preconditioners, with stochastic dimension $N = 3$ and coefficient of variation $CoV = 30\%$. Other headings are as in Table 4.1.

P	M	M_ν	$ngdof$	MB	K	bGS	ahGS	ahGS(PCD)	ahGS(PCD-it)
1	4	10	57,120	27	23	12	11	219	32
2	10	35	142,800	68	55	33	32	450	42
3	20	84	285,600	235	170	128	151	1011	117

TABLE 4.7

For Newton step: dependence on coefficient of variation CoV of GMRES iteration counts, for various preconditioners, with stochastic dimension $N = 2$ and polynomial degree $P = 3$. Other headings are as in Table 4.1.

$CoV(\%)$	MB	K	bGS	ahGS	ahGS(PCD)	ahGS(PCD-it)
10	17	15	8	8	322	28
20	40	30	20	19	379	35
30	126	80	77	69	546	63

algebraic multigrid for stochastic PDE problems with layered materials. *Numerical Linear Algebra with Applications*, 21(2):239–255, 2014.

- [3] F. Brezzi and M. Fortin. *Mixed and Hybrid Finite Element Methods*. Springer-Verlag, New York – Berlin – Heidelberg, 1991.
- [4] H. C. Elman and Q. Liao. Reduced basis collocation methods for partial differential equations with random coefficients. *SIAM/ASA Journal on Uncertainty Quantification*, 1:192–217, 2013.
- [5] H. C. Elman, A. Ramage, and D. J. Silvester. IFIGS: A computational laboratory for investigating incompressible flow problems. *SIAM Review*, 56:261–273, 2014.
- [6] H. C. Elman, D. J. Silvester, and A. J. Wathen. *Finite elements and fast iterative solvers: with applications in incompressible fluid dynamics*. Numerical Mathematics and Scientific Computation. Oxford University Press, New York, second edition, 2014.
- [7] O. G. Ernst and E. Ullmann. Stochastic Galerkin matrices. *SIAM Journal on Matrix Analysis and Applications*, 31(4):1848–1872, 2010.
- [8] T. Gerstner and M. Griebel. Numerical integration using sparse grids. *Numerical Algorithms*, 18:209–232, 1998.
- [9] R. Ghanem. The nonlinear Gaussian spectrum of log-normal stochastic processes and variables. *Journal of Applied Mechanics*, 66(4):964–973, 1999.
- [10] R. G. Ghanem and Pol D. Spanos. *Stochastic Finite Elements: A Spectral Approach*. Springer-Verlag New York, Inc., New York, NY, USA, 1991. (Revised edition by Dover Publications, 2003).
- [11] L. Giraldi, A. Nouy, and G. Legrain. Low-rank approximate inverse for preconditioning tensor-structured linear systems. *SIAM Journal on Scientific Computing*, 36(4):A1850–A1870, 2014.
- [12] V. Girault and P.-A. Raviart. *Finite element methods for Navier-Stokes equations*. Springer-Verlag, Berlin, 1986.
- [13] M. Kamiński. *The Stochastic Perturbation Method for Computational Mechanics*. John Wiley & Sons, 2013.
- [14] O. Le Maître and O. M. Knio. *Spectral Methods for Uncertainty Quantification: With Applications to Computational Fluid Dynamics*. Scientific Computation. Springer, 2010.
- [15] H. G. Matthies and A. Keese. Galerkin methods for linear and nonlinear elliptic stochastic partial differential equations. *Computer Methods in Applied Mechanics and Engineering*, 194(1216):1295–1331, 2005.
- [16] E. Novak and K. Ritter. High dimensional integration of smooth functions over cubes. *Numerische Mathematik*, 75(1):79–97, 1996.
- [17] M. F. Pellissetti and R. G. Ghanem. Iterative solution of systems of linear equations arising in the context of stochastic finite elements. *Advances in Engineering Software*, 31(8-9):607–616, 2000.

TABLE 4.8

For Newton step: number of GMRES iterations when the preconditioners use the truncation of the MATVEC. Headings are as in Table 4.4.

	ℓ_t	0	1	2	3	6
setup	M_t	1	3	6	10	28
	$nnz(c_\ell)$	0	12	21	43	63
	bGS					
CoV (%)	10	17	10	8	8	8
	20	40	22	18	20	20
	30	126	68	57	75	77
ahGS						
CoV (%)	10	17	10	8	8	8
	20	40	22	17	19	19
	30	126	68	45	70	69

- [18] C. E. Powell and H C. Elman. Block-diagonal preconditioning for spectral stochastic finite-element systems. *IMA Journal of Numerical Analysis*, 29(2):350–375, 2009.
- [19] C. E. Powell and D. J. Silvester. Preconditioning steady-state Navier–Stokes equations with random data. *SIAM Journal on Scientific Computing*, 34(5):A2482–A2506, 2012.
- [20] E. Rosseel and S. Vandewalle. Iterative solvers for the stochastic finite element method. *SIAM Journal on Scientific Computing*, 32(1):372–397, 2010.
- [21] Y. Saad. A flexible inner-outer preconditioned GMRES algorithm. *SIAM Journal on Scientific Computing*, 14(2):461–469, 1993.
- [22] B. Sousedík and R. G. Ghanem. Truncated hierarchical preconditioning for the stochastic Galerkin FEM. *International Journal for Uncertainty Quantification*, 4(4):333–348, 2014.
- [23] B. Sousedík, R. G. Ghanem, and E. T. Phipps. Hierarchical Schur complement preconditioner for the stochastic Galerkin finite element methods. *Numerical Linear Algebra with Applications*, 21(1):136–151, 2014.
- [24] L. Tamellini, O. Le Maître, and A. Nouy. Model reduction based on Proper Generalized Decomposition for the stochastic steady incompressible Navier–Stokes equations. *SIAM Journal on Scientific Computing*, 36(3):A1089–A1117, 2014.
- [25] E. Ullmann. A Kronecker product preconditioner for stochastic Galerkin finite element discretizations. *SIAM Journal on Scientific Computing*, 32(2):923–946, 2010.
- [26] D. Xiu. *Numerical Methods for Stochastic Computations: A Spectral Method Approach*. Princeton University Press, 2010.
- [27] D. Xiu and G. E. Karniadakis. The Wiener-Askey polynomial chaos for stochastic differential equations. *SIAM Journal on Scientific Computing*, 24(2):619–644, 2002.
- [28] D. Zhang. *Stochastic Methods for Flow in Porous Media. Coping with Uncertainties*. Academic Press, San Diego, CA, 2002.