UNIVERSITY OF MARYLAND
INSTITUTE FOR ADVANCED COMPUTER STUDIES
DEPARTMENT OF COMPUTER SCIENCE

EFFICIENT ITERATIVE ALGORITHMS FOR THE STOCHASTIC FINITE ELEMENT METHOD WITH APPLICATION TO ACOUSTIC SCATTERING

CS-TR #4423 / UMIACS TR #2002-102

HOWARD ELMAN[†], OLIVER G. ERNST[‡], DIANNE P. O'LEARY[§], AND MICHAEL STEWART[¶]

DECEMBER 2002

Abstract. In this study, we describe the algebraic computations required to implement the stochastic finite element method for solving problems in which uncertainty is restricted to right hand side data coming from forcing functions or boundary conditions. We show that the solution can be represented in a compact outer product form which leads to efficiencies in both work and storage, and we demonstrate that block iterative methods for algebraic systems with multiple right hand sides can be used to advantage to compute this solution. We also show how to generate a variety of statistical quantities from the computed solution. Finally, we examine the behavior of these statistical quantities in one setting derived from a model of acoustic scattering.

[†]Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742. elman@cs.umd.edu. This work was supported in part by the National Science Foundation under grants DMS9972490 and DMS0208015 and by the Office of Naval Research under grant N000140110181.

[‡]Department of Mathematics and Computer Science, TU Bergakademie Freiberg, 09596 Freiberg, Germany. ernst@math.tu-freiberg.de.

[§]Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742. oleary@cs.umd.edu. This work was supported in part by the National Science Foundation under grants CCR-97-32022 CCR-0204084 and by the Office of Naval Research under grant N000140110181.

[¶]Department of Mathematics and Statistics, Georgia State University, Atlanta, Georgia 30303-3083. mstewart@cs.umd.edu. This work was supported in part by the Office of Naval Research under grant N000140110181.

EFFICIENT ITERATIVE ALGORITHMS FOR THE STOCHASTIC FINITE ELEMENT METHOD WITH APPLICATION TO ACOUSTIC SCATTERING

HOWARD ELMAN[†], OLIVER G. ERNST[‡], DIANNE P. O'LEARY[§], AND MICHAEL STEWART[¶]

Abstract. In this study, we describe the algebraic computations required to implement the stochastic finite element method for solving problems in which uncertainty is restricted to right hand side data coming from forcing functions or boundary conditions. We show that the solution can be represented in a compact outer product form which leads to efficiencies in both work and storage, and we demonstrate that block iterative methods for algebraic systems with multiple right hand sides can be used to advantage to compute this solution. We also show how to generate a variety of statistical quantities from the computed solution. Finally, we examine the behavior of these statistical quantities in one setting derived from a model of acoustic scattering.

Key words. Stochastic, finite element, multiple right hand side, multigrid, scattering.

1. Introduction.

It is common practice for mathematical models to be studied under the assumption that data defining the models are precisely understood. In reality, however, this simplifying assumption is often not valid, and there is considerable uncertainty in specification of models. Sources of uncertainty include geological properties of transporting media, material properties of structures, and unknown aspects of boundary conditions.

One approach for addressing this issue is to treat poorly specified data as random variables having some given statistical properties such as means and higher order moments, and then to determine analogous statistical properties of solutions. For boundary value problems with uncertain data (*stochastic partial differential equations*), a methodology known as the *stochastic finite element method* has generated considerable attention in the last decade [6, 7, 11, 12, 15]. This approach starts with a boundary value problem in d-dimensional physical space. The stochastic component of the problem statement is then specified approximately using an m-dimensional auxiliary space which is derived from an underlying probability space associated with the data. The result is a $(d+m)$-dimensional model, which can be stated in a weak form on a suitable function space using a combination of standard variational constructions for the physical component of the problem together with averaging for the stochastic component. We will outline the details of this methodology in Section 2.

Once this weak formulation is specified, a numerical solution of the stochastic partial differential equation can be computed in essentially the same manner as for deterministic problems. In particular, the introduction of finite dimensional subspaces

[†]Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742. elman@cs.umd.edu. This work was supported in part by the National Science Foundation under grants DMS9972490 and DMS0208015 and by the Office of Naval Research under grant N000140110181.

[‡]Department of Mathematics and Computer Science, TU Bergakademie Freiberg, 09596 Freiberg, Germany. ernst@math.tu-freiberg.de.

[§]Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742. oleary@cs.umd.edu. This work was supported in part by the National Science Foundation under grants CCR-97-32022 CCR-0204084 and by the Office of Naval Research under grant N000140110181.

[¶]Department of Mathematics and Statistics, Georgia State University, Atlanta, Georgia 30303-3083. mstewart@cs.umd.edu. This work was supported in part by the Office of Naval Research under grant N000140110181.

leads to an algebraic system of equations whose solution can be used to approximate statistical properties of the physical solution, such as its mean, variance and covariances. Our concern in this paper is to explore the computational costs of solving the systems in question and of generating statistical analyses of the solution.

We will focus on problems where randomness only affects the right hand sides of the algebraic systems, that is, where the forcing terms or boundary data are random functions. A natural example of this arises in models of acoustic or electromagnetic scattering, where lack of information about the material properties of scatterers or the shape and structure of boundaries such as ocean bottoms leads to uncertainty in boundary conditions. We will use this model, specifically, the numerical solution of the Helmholtz equation, as a benchmark problem, and in our assessment we will explore computational issues associated with quantities such as moments and probability distributions of acoustic pressures, and how these are affected by characteristics of the problem such as wave numbers.

One of the computational tasks required is the solution of algebraic systems of equations with multiple right hand sides. In the case of uncertain boundary data, the costs of this component of the computation can be kept low using the fact that the solution has a Kronecker product structure. For our scattering example, the systems can be solved efficiently with a multigrid algorithm for the discrete Helmholtz equation [9], and we show that efficiency can be enhanced in some cases using block iterative methods for systems with multiple right hand sides [4, 10, 17, 22]. With this strategy for solving the algebraic systems, the dominant cost of the computation is that of computing statistical quantities. We also show that the Kronecker product structure of the solution allows storage costs to be kept relatively low, and moreover it enables the statistical computations to be performed using efficient matrix-oriented operations that are trivially parallelizable and amenable to implementation using Level 3 Basic Linear Algebra Subprograms (BLAS3) [8]. This means that it is possible to handle relatively fine "discretization" in the stochastic domain that would otherwise not be possible.

We note that an alternative approach for handling random right hand sides has been developed in Schwab and Todor [20], where it is shown that the mean and second moment of the solution can be computed directly, where the latter entails the solution of a fourth order equation derived for this quantity. It is shown in [20] that when the underlying differential operator is coercive, then so is the associated fourth order system, and efficient multilevel algorithms (but dependent on special sparse grids) can be developed to solve it. The approach under consideration here has the advantage that it readily yields more general statistical information such as higher order moments and probability distributions. It is also relatively straightforward to implement, essentially only requiring algorithm technology for second order problems. In particular, if, as in the example considered here, the underlying problem is *not* coercive, it is still possible to take advantage of efficient algorithms for that problem.

A summary of the contents of the paper is as follows. Section 2 contains a description of the stochastic finite element methodology and identifies the structure of the algebraic systems derived from discretization. Section 3 describes the iterative algorithms that we consider for solving the discrete Helmholtz equation and the block versions designed to handle multiple right hand sides, and then it presents some experimental results demonstrating the performance of these solvers. Section 4 then outlines the costs of computing statistical quantities associated with the solution and shows the results of these computations. Finally, Section 5 contains some concluding remarks.

2. The stochastic finite element method. We briefly describe the general methodology with an eye towards showing the structure of the algebraic systems. For our description we use the problem that we will study in experiments, the Helmholtz equation; it will be obvious that the approach is general. See [12] for complete descriptions of this methodology.

2.1. Introduction: weak formulation. A model of acoustic scattering from a bounded obstacle is given by the Helmholtz equation

$$\begin{aligned} -\Delta u - k^2 u &= f && \text{in } \mathcal{D} \\ B(u) &= g && \text{on } \Gamma \\ \tfrac{\partial u}{\partial n} &= L(u) && \text{on } \Gamma_\infty \end{aligned} \tag{2.1}$$

where the solution domain $\mathcal{D} \subset \mathbb{R}^d$ is bounded internally by the obstacle boundary $\Gamma \subset \partial\mathcal{D}$ and externally by an artificial boundary Γ_∞. The boundary differential operator B is such that Dirichlet, Neumann or Robin boundary conditions result along Γ, and L is the Dirichlet-to-Neumann operator [13] or a suitable approximation thereof.

The weak form of this problem is to find $u \in V_g$ such that

$$a(u,v) = \ell(v) \qquad \forall v \in V \tag{2.2}$$

where V and V_g denote the linear and affine subspaces of $H^1(\mathcal{D})$ of functions satisfying any homogeneous resp. inhomogeneous essential boundary conditions along Γ. In the simplest case of Dirichlet boundary data along all of Γ, the sesquilinear form $a : H^1(\mathcal{D}) \times H^1(\mathcal{D}) \to \mathbb{C}$ is

$$a(u,v) = \int_{\mathcal{D}} \left(\nabla u \cdot \nabla v - k^2 uv\right)\, dx - \int_{\Gamma_\infty} vL(u)\, ds$$

and the right hand side functional $\ell : H^1(\mathcal{D}) \to \mathbb{C}$ is

$$\ell(v) = \int_{\mathcal{D}} fv\, dx.$$

To introduce randomness into this formulation, let $(\Omega, \mathcal{A}, P)$ denote a probability space with sample space Ω, σ-algebra $\mathcal{A}$ and probability measure P. Let $\zeta : \Omega \to \mathbb{C}$ be a complex-valued random variable with $\zeta \in L^1(\Omega)$. The *mean* or *expected value* of ζ is

$$\langle \zeta \rangle = \int_\Omega \zeta(\omega)\, dP(\omega) = \int_{\mathbb{C}} z\, d\mu(z),$$

where μ is the distribution probability measure associated with ζ and defined on the Borel sets B in the complex plane by $\mu(B) = P(\zeta^{-1}(B))$. Given a bounded domain $\mathcal{D} \subset \mathbb{R}^d$ as above, a *random function*

$$u : \mathcal{D} \times \Omega \to \mathbb{C}, \qquad (x,\omega) \mapsto u(x,\omega)$$

is one that is jointly measurable with respect to Lebesgue measure on $\mathcal{D}$ and the measure P on Ω and for which

$$\left\langle \|u(\cdot,\omega)\|_{L^2(\mathcal{D})} \right\rangle < \infty.$$

The space of random functions is a Hilbert space $\tilde{L}^2(\mathcal{D} \times \Omega)$ with respect to the inner product

$$(u, v)_{\tilde{L}^2} = \left\langle (u(x,\cdot), v(x,\cdot))_{L^2(\mathcal{D})} \right\rangle .$$

The *stochastic Sobolev spaces* $\tilde{H}^k(\mathcal{D} \times \Omega)$ are defined analogously.

If any of the data in the Helmholtz equation (2.1) is random (e.g., the wave number k, forcing function f, or Dirichlet boundary data g), then the solution u will be a random function. The weak form of the stochastic problem is then to find $u \in \tilde{V}_g$ such that

$$\langle a(u, v) \rangle = \langle \ell(v) \rangle \quad \forall v \in \tilde{V}_0, \tag{2.3}$$

where $\tilde{V}_g$ and $\tilde{V}_0$ are the stochastic Sobolev spaces analogous to V_g and V_0.

2.2. The Karhunen-Loève expansion and derived weak form. We consider the development of the stochastic finite element method using the *Karhunen-Loève (KL) expansion*, a representation of random functions in series form using the eigenfunctions of the covariance operator. For the sake of concreteness, we describe its use under the assumption that the forcing function f of (2.1) is random; we will discuss other possibilities in Section 2.3.

Let the covariance function associated with f be denoted by

$$c(x, y) = \langle f(x) f(y) \rangle - \langle f(x) \rangle \langle f(y) \rangle .$$

Consider the integral equation

$$(\mathcal{C}\gamma)(x) = \lambda\, \gamma(x), \quad \text{where} \quad (\mathcal{C}\gamma)(x) = \int_{\mathcal{D}} c(x, y) \gamma(y)\, dy. \tag{2.4}$$

This is a linear integral eigenvalue problem in which, by definition, the kernel is symmetric and positive-semidefinite. It follows from the general theory of integral equations [5, Ch. 3] that $\mathcal{C}$ is a compact operator and there exists a countable sequence of eigenpairs $\{(\lambda_r, f_r)\}$ where the eigenvalues $\{\lambda_r\}$ are nonnegative and the eigenfunctions $\{f_r\}$ are orthogonal in $L^2(\mathcal{D})$. Let the eigenvalues be ordered so that $\lambda_1 \geq \lambda_2 \geq \cdots \geq 0$. The Karhunen-Loève expansion for f is

$$f(x, \boldsymbol{\xi}) = f_0(x) + \sum_{r=1}^{\infty} \sqrt{\lambda_r}\, f_r(x) \xi_r, \tag{2.5}$$

where $f_0(x) = \langle f(x) \rangle$ is the mean of f, and $\{\xi_r(\omega)\}_{r \geq 1}$ are uncorrelated random variables with mean zero and variance one [23, pp. 447ff].

For computation, the infinite series (2.5) is approximated by a finite one with, say, m terms. In general, the more localized the covariance kernel of f (the smaller the correlation length), the slower the decay of its eigenvalues and the more terms need be retained in the KL expansion to achieve good accuracy. Thus, the utility of this approach depends on the assumption that the properties of physical systems under consideration vary smoothly, i.e., there are significant correlations in the random data. In this case, it is expected that a truncated version of (2.5) with a small number m of terms in the sum is sufficient to capture the randomness in the system.

Assume now that the random function is given by a finite-term KL expansion

$$f(x, \xi) = f_0(x) + \sum_{r=1}^{m} \sqrt{\lambda_r}\, f_r(x) \xi_r. \tag{2.6}$$

Let $\mathcal{I}_r = \xi_r(\Omega)$ denote the image of ξ_r, and let $\mathcal{I} = \mathcal{I}_1 \times \cdots \times \mathcal{I}_m$. Collecting these variables into the random vector $\boldsymbol{\xi} = (\xi_1, \ldots, \xi_m)$, we have $\boldsymbol{\xi}(\Omega) \subset \mathcal{I}$. Assume that ξ_r possesses the probability density function $\rho_r : \mathcal{I}_r \to \mathbb{R}$, which gives rise to the joint density function

$$\rho(\boldsymbol{\xi}) = \rho_1(\xi_1)\rho_2(\xi_2)\cdots\rho_m(\xi_m).$$

The stochastic variational formulation of the Helmholtz equation (2.1) uses as test functions random functions in the space

$$\tilde{\mathcal{V}} = \left\{ u(x,\boldsymbol{\xi}) : \int_{\mathcal{I}} \left(\|u\|_{H^1(\mathcal{D})} \right) \rho(\boldsymbol{\xi}) d\boldsymbol{\xi} \ < \infty \right\}, \tag{2.7}$$

with trial functions in the space $\tilde{\mathcal{V}}_g$ defined analogously. The stochastic variational problem is then specified as in (2.3) with

$$\begin{aligned} \langle a(u,v) \rangle &= \int_{\mathcal{I}} \left(\int_{\mathcal{D}} \left(\nabla u \cdot \nabla v - k^2 uv \right) \, dx - \int_{\Gamma_\infty} vL(u) \, ds \right) \rho(\boldsymbol{\xi}) d\boldsymbol{\xi} \\ \langle \ell(v) \rangle &= \int_{\mathcal{I}} \left(\int_{\mathcal{D}} fv \, dx \right) \rho(\boldsymbol{\xi}) d\boldsymbol{\xi}. \end{aligned} \tag{2.8}$$

The weak solution u can be viewed as defined on a $(d+m)$-dimensional domain $\mathcal{D} \times \mathcal{I}$.

2.3. Discretization and the stochastic system. In order to establish notation, we briefly discuss the discretization of the deterministic problem (2.1), assuming Dirichlet boundary conditions $u = g$ hold on the obstacle boundary Γ. Let $\mathcal{V}^h = \text{span}\{\phi_1, \ldots, \phi_{N_x}\}$ denote a finite dimensional subspace of $H_0^1(\mathcal{D})$, and let $\mathcal{V}_g^h$ denote the affine space obtained by adding basis functions $\{\phi_{N_x+1}, \ldots, \phi_{N_x+N_E}\}$ to handle degrees of freedom on the boundary. As is well known, the discrete weak formulation entails finding

$$u_h = \sum_{j=1}^{N_x} \boldsymbol{u}_j \phi_j + \sum_{j=N_x+1}^{N_x+N_E} g(x_j)\phi_j$$

such that

$$\sum_{j=1}^{N_x} a(\phi_j, \phi_i)\boldsymbol{u}_j = \int_{\mathcal{D}} f\phi_i \, dx \ - \sum_{j=N_x+1}^{N_x+N_E} a(\phi_j,\phi_i)g(x_j) \quad \forall i = 1, \ldots, N_x \, .$$

This is a linear system of equations $A\boldsymbol{u} = \boldsymbol{f}$ where

$$\boldsymbol{f} = [(f, \phi_i)]_{i=1}^{N_x} - A_{UE}\boldsymbol{g}; \tag{2.9}$$

A_{UE} represents the coupling between degrees of freedom constrained by Dirichlet boundary conditions and other unknowns, and $\boldsymbol{g} = [g(x_j)]_{j=N_x+1}^{N_x+N_E}$ is the vector of nodal boundary values.

Now consider the stochastic problem defined by (2.3) and (2.8). For the discretization, let

$$\tilde{\mathcal{V}}^h = \text{span}\{\chi_{jq}(x,\boldsymbol{\xi}) = \phi_j(x)\psi_q(\boldsymbol{\xi}) : j = 1, \ldots, N_x, \, q = 1, \ldots, N_\xi\},$$

denote a finite-dimensional subspace of $\tilde{\mathcal{V}}$ of (2.7), where $\{\psi_1, \ldots, \psi_{N_\xi}\}$ is a basis for a finite-dimensional subspace of $L^2(\mathcal{I})$. Let $\tilde{\mathcal{V}}_g^h$ denote the affine space satisfying inhomogeneous essential boundary conditions. The discrete stochastic problem is then to find $u_h \in \tilde{\mathcal{V}}_g^h$,

$$u_h(x,\xi) = \sum_{q=1}^{N_\xi} \sum_{j=1}^{N_x} \phi_j(x)\psi_q(\xi)\boldsymbol{u}_{jq} \; + \sum_{j=N_x+1}^{N_X+N_E} \phi_j(x) g(x_j) \tag{2.10}$$

such that

$$\langle a(u_h, v_h)\rangle = \langle \ell(v_h)\rangle \quad \forall v_h \in \tilde{\mathcal{V}}^h.$$

The result is a linear system of equations, the *stochastic system*,

$$\boldsymbol{A}\boldsymbol{u} = \boldsymbol{b} \tag{2.11}$$

of order $N_x \times N_\xi$, for unknowns

$$(\boldsymbol{u}_{11}, \boldsymbol{u}_{21}, \ldots, \boldsymbol{u}_{N_x-1,N_\xi}, \boldsymbol{u}_{N_x,N_\xi})^T.$$

Once $\boldsymbol{u}$ is obtained, statistical properties of the associated random function u_h can be computed easily, see Section 4.

As we have noted, this study concerns the case where randomness only affects the right hand side of the algebraic systems generated, i.e., where the source term or boundary data is random. Let us consider the structure of the discrete problem (2.11) in this case. The entries of the finite element system matrix $\boldsymbol{A}$ are

$$\begin{aligned}
\langle a(\chi_{jq}, \chi_{ip})\rangle &= \int_{\mathcal{I}} a(\phi_j\psi_q, \phi_i\psi_p)\, \rho(\boldsymbol{\xi}) d\boldsymbol{\xi} \\
&= \left(\int_\Gamma \psi_q\psi_p\rho(\boldsymbol{\xi})\, d\boldsymbol{\xi}\right) \left(\int_{\mathcal{D}} \nabla\phi_j \cdot \nabla\phi_i - k^2\phi_j\phi_i\, dx - \int_{\Gamma_\infty} \phi_i L\phi_j\, ds\right) \\
&= \langle \psi_q\psi_p\rangle\, a(\phi_j, \phi_i),
\end{aligned}$$

for $1 \le i,j \le N_x$, $1 \le p,q \le N_\xi$. Denoting by $G \in \mathbb{R}^{N_\xi \times N_\xi}$ the Grammian matrix

$$[G]_{pq} = \langle \psi_q\psi_p\rangle\,, \quad p,\, q = 1, \ldots, N_\xi\,, \tag{2.12}$$

and by $A \in \mathbb{C}^{N_x \times N_x}$ the stiffness matrix of the deterministic equation, the coefficient matrix is seen to have the Kronecker structure

$$\boldsymbol{A} = G \otimes A.$$

Note that this implicitly determines an ordering for the rows and columns of $\boldsymbol{A}$. The rows are ordered so that for each p, indices $i = 1, \ldots, N_x$ are grouped together, and then p is ordered from 1 to N_ξ; the same grouping applies to the columns.

For the right hand side, assume as in Section 2.2 that the forcing function is random, and also assume for the moment that homogeneous Dirichlet boundary conditions $g = 0$ hold on Γ. It then follows from (2.6) and (2.8) that the entry of $\boldsymbol{f}$ corresponding to the test function $\chi_{ip} = \phi_i\psi_p$ is

$$\langle \ell(\chi_{ip})\rangle = \int_{\mathcal{I}} \ell(f, \chi_{ip})\rho(\boldsymbol{\xi}) d\boldsymbol{\xi} = \ell(f_0, \phi_i)\, \langle\psi_p\rangle + \sum_{r=1}^m \sqrt{\lambda_r}\; \ell(f_r, \phi_i)\, \langle \xi_r\psi_p\rangle\,. \tag{2.13}$$

Let us define the vectors

$$\begin{aligned} \boldsymbol{f}_r &= [(f_r,\phi_i)]_{i=1}^{N_x}\,, \quad r=0,1,\ldots,m \\ \boldsymbol{\psi}_0 &= [\langle\psi_p\rangle]_{p=1}^{N_\xi} \\ \boldsymbol{\psi}_r &= [\langle\xi_r\psi_p\rangle]_{p=1}^{N_\xi}\,, \quad r=0,1,\ldots,m \end{aligned} \tag{2.14}$$

whereupon the discrete system has the form

$$(G\otimes A)\boldsymbol{u}=\boldsymbol{f}, \qquad \boldsymbol{f}=\boldsymbol{\psi}_0\otimes\boldsymbol{f}_0+\sum_{r=1}^{m}\sqrt{\lambda_r}\,(\boldsymbol{\psi}_r\otimes\boldsymbol{f}_r)\,.$$

That is, the right hand side lies in an $(m+1)$-dimensional subspace of $\mathbb{R}^{N_\xi\times N_x}$. The solution is then

$$\begin{aligned} \boldsymbol{u} &= (G\otimes A)^{-1}\boldsymbol{f}=(G^{-1}\otimes A^{-1})\boldsymbol{f} \\ &= (G^{-1}\boldsymbol{\psi}_0)\otimes(A^{-1}\boldsymbol{f}_0)+\sum_{r=1}^{m}\sqrt{\lambda_r}\,(G^{-1}\boldsymbol{\psi}_r)\otimes(A^{-1}\boldsymbol{f}_r)\,. \end{aligned} \tag{2.15}$$

This entails the solution of $m+1$ systems of size N_ξ with coefficient matrix G, and $m+1$ systems of size N_x with coefficient matrix A. In practice, the basis functions $\{\psi_p\}$ for the stochastic component are often chosen to be orthogonal with respect to the probability measure [7, 12], in which case G is a diagonal matrix. Thus, the main computational requirement is solution of the $m+1$ systems with coefficient matrix A.

Although the derivation above is for the case of stochastic forcing function and homogeneous boundary conditions, the conclusion reached is general. For example, if a nonzero Dirichlet condition holds on Γ, then the construction is identical except $\boldsymbol{f}_0$ has the form (cf. (2.9))

$$\boldsymbol{f}_0=[(f_0,\phi_i)]_{i=1}^{N_x}-A_{UE}\boldsymbol{g}.$$

More generally, if it is Dirichlet boundary conditions that are random (we will explore this in experiments described in Section 3), then terms of the form

$$\boldsymbol{\psi}_0\otimes(A_{UE}\boldsymbol{g}_0)+\sum_r\sqrt{\lambda_r}\,(\boldsymbol{\psi}_r\otimes(A_{UE}\boldsymbol{g}_r))$$

will be included in the right hand side. Similar considerations apply for Neumann conditions on the obstacle boundary.

2.4. Implementation. The notation used in the previous section treats the unknowns $\boldsymbol{u}$ of (2.11) as a vector. In an implementation, it is in fact more convenient to treat the solution as a two-dimensional array. In particular, consider the matrices

$$F=[\boldsymbol{f}_0,\boldsymbol{f}_1,,\ldots,\boldsymbol{f}_m], \quad \Lambda=diag(1,\sqrt{\lambda_1},\ldots,\sqrt{\lambda_m}), \quad \Psi=[\boldsymbol{\psi}_0,\boldsymbol{\psi}_1,,\ldots,\boldsymbol{\psi}_m],$$

where the vectors $\{\boldsymbol{f}_r\}$ and $\{\boldsymbol{\psi}_r\}$ are defined in (2.14). Then the system (2.11) is essentially of the form

$$AU=B, \tag{2.16}$$

where $B=F\,W^T$ with $W=G^{-1}(\Psi\Lambda)$. The solution can then be represented implicitly in outer-product form as

$$U=V\,W^T, \tag{2.17}$$

where $V=A^{-1}F$ is obtained by solving the system of equations $AV=F$ with $m+1$ right hand sides.

3. Iterative solution of the stochastic system. For the problem under consideration, the coefficient matrix of (2.16) is a discrete Helmholtz operator, which is complex, symmetric and indefinite. In this section, we describe an iterative algorithm that can be used to solve this system and demonstrate its effectiveness on a set of benchmark problems.

3.1. Solution algorithm. The basic solution algorithm we use is a multigrid method designed for the Helmholtz equation, adapted to handle multiple right hand sides. As is well known, the principle behind multigrid is to combine *smoothers* to eliminate oscillatory components of the error on fine grids, together with *coarse grid corrections* to eliminate smooth components. For the Helmholtz equation, standard multigrid approaches are not effective. There are two difficulties:

1. Standard smoothers such as the Jacobi and Gauss-Seidel methods do not work because certain smooth modes are amplified by these operations.
2. The eigenvalues associated with some smooth modes change signs during the grid coarsening process, which causes the coarse grid correction to also amplify some smooth modes rather than eliminate them from the error.

These difficulties derive from the indefiniteness of the system. In [9], we developed a method that addresses them. The first difficulty is handled by replacing standard smoothers with Krylov subspace methods, i.e., GMRES iteration [19] is used as the smoother. The second one is handled by using the multigrid operation as a preconditioner for an outer Krylov subspace iteration, so that components of the error not treated correctly by the multigrid coarse grid computations are eliminated. Because the multigrid smoother is no longer a linear operator, the outer iteration must handle this via a so-called “flexible” GMRES algorithm [18]. A complete description and analysis of the preconditioning strategy is given in [9], where it is demonstrated that the algorithm exhibits “textbook multigrid” convergence behavior, that is, convergence rates that are independent of the discretization parameter; there is some dependence on the wave number k.

We also adapt this approach to handle the system (2.16) with multiple right hand sides, the number of which is denoted by m within this section. Recall that Krylov subspace methods generate an iterate at step s using a certain subspace of dimension s. Two types of Krylov subspace algorithms have been proposed for problems with multiple right hand sides:

- *Block algorithms* [2], [17] construct a subspace of dimension ms formed by the union of the s-dimensional subspaces for each right hand side. Then, for each right hand side, they find the best solution within that subspace. Deflation is used to remove vectors that become linearly dependent.
- *Seed algorithms* [3], [21] form a Krylov subspace using one of the right hand sides and then find the best solution for each of the m problems within that subspace. If the seed problem converges before the others, then a different right hand side is chosen as the seed and the algorithm is repeated.

Each of these approaches has its advantages. Seed methods tend to perform best when the right hand sides are related to each other, for example, if they arise from functions evaluated at nearby points [3]. This approach requires less storage than block methods: for systems of order N, the seed GMRES method requires storage proportional to sN, compared to smN for a block GMRES solver. On the other hand, block algorithms tend to converge more rapidly for more general right hand sides, or when a small number of eigenvalues are well-separated from the others [17]. The block algorithm also makes much better use of computer memory traffic, since each access to the coefficient matrix is used for m matrix-vector products.

In our application, the right hand side vectors (columns of F in (2.16)) derive from the orthogonal eigenvectors of the covariance matrix, and we found the seed method to be ineffective. Therefore, we restrict our attention to a block method. The idea of block iterative methods is due to O'Leary and R. Underwood. The block biconjugate gradient algorithm was described in [17], and a block quasi-minimum residual method in [10]. Algorithms for altering the block size adaptively were given in [1]. A block GMRES algorithm was presented by Vital [22].

We also need to modify the algorithm to handle the nonlinear preconditioner, as described in [4]. To present this *block flexible GMRES method* for (2.16), we use the generic notation $Ax = b$ for the linear system, and $w = M(v)$ to represent a generic *preconditioning* operation. This may be a linear operation derived from a matrix, or (as in the present setting) a nonlinear operation. Let $x_j^{(s)}$ denote an approximation of the solution to the jth equation (of m) computed at iteration s. The block flexible GMRES algorithm generates a sequence of matrices $\{V_j\}$ of dimensions $N \times m$ that together form a matrix $V = [V_1, \ldots, V_s]$, and a set of matrices $Z_j = M(V_j)$, and $Z = [Z_1, \ldots, Z_s]$. The block-Hessenberg matrix H has block entries H_{ij}. Each component $\{x_j^{(s)} : j = 1, \ldots, m\}$ of the iterative solution has the form

$$x_j^{(s)} = x_j^{(0)} + Z y_j^{(s)} \tag{3.1}$$

such that the norm of the residual $\|b_j - Ax_j^{(s)}\|$ is minimal, where $x_j^{(0)}$ is a possibly arbitrary initial value. If $M(v)$ is a linear operation (say Mv where M is a matrix), then $\mathcal{V}_s$, the span of the columns of V, is the same as

$$span\{Mr_1, (MA)Mr_1, \ldots, (MA)^{s-1}Mr_1, \ldots, Mr_m, (MA)Mr_m, \ldots, (MA)^{s-1}Mr_m\},$$

the span of the Krylov vectors. In this case, it is not necessary to store the auxiliary matrix Z. If M is a nonlinear operator, then $\mathcal{V}_s$ will not be a Krylov subspace, but the s'th iterate is still optimal among all possibilities of the form (3.1), which corresponds to an affine subspace of $\mathbb{C}^N$ of dimension ms. The algorithm is as follows:

Compute the residual $r = b - Ax$ of dimension $N \times m$.
until ($\|r_\ell\| \leq \delta$, $\ell = 1, \ldots, m$),
    % Generate a subspace of dimension ms from the residual r.
    Define V_1 to be the orthogonal factor in the QR factorization of r.
    for $i = 2, \ldots, s+1$,
        % Generate the directions defining the new basis vectors.
        $Z_{i-1} = M(V_{i-1})$
        $W = AZ_{i-1}$
        % Orthogonalize these directions against the previous ones.
        for $j = 1, \ldots, i-1$,
            $H_{j,i-1} = V_j^* W$
            $W = W - V_j H_{j,i-1}$
        end for j.
        Perform a QR factorization of W, obtaining the upper triangular factor $H_{i,i-1}$ and the orthogonal matrix V_i.
    end for i.
    % Update each of the solutions.
    for $j = 1, \ldots, m$,

FIG. 3.1. *Spatial domain and initial mesh used in spatial discretization.*

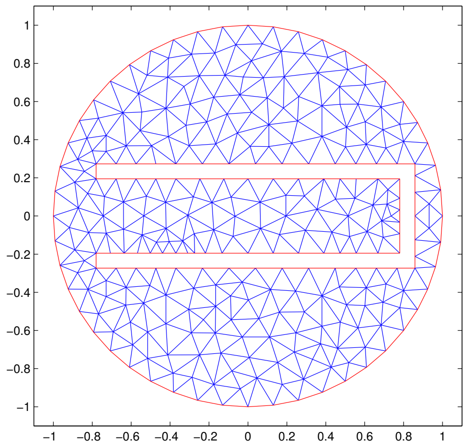

$c = V^* r_j$
Solve the least squares problem $\min_y \|c - Hy\|$.
$x_j = x_j + Zy$
end for j.
$r = b - Ax$
end until

The loop on i can break down if the matrix W becomes rank-deficient. In this case, we reduce the size of the block by dropping the dependent columns, updating the solutions, and continuing with the residuals that have not converged.

3.2. Experimental results. We tested the performance of the block flexible GMRES algorithm for solving the stochastic Helmholtz equation on the two-dimensional domain $\mathcal{D}$ consisting of the complement within the unit circle of a scatterer taken to be a semi-open cavity. Dirichlet-to-Neumann conditions are specified on the external boundary Γ_∞. The discretization in space consists of piecewise linear elements on triangles. Figure 3.1 shows the scatterer and the initial mesh used in all tests. For each wave number k, this mesh is refined until kh_{max} was on the order of $\pi/5 \approx .63$, so that there are approximately ten points per wavelength. All computations were done using MATLAB. Mesh construction was done using the MATLAB PDE TOOLBOX routines `initmesh` and `refinemesh`, which performs a uniform mesh refinement.

All uncertainty in the specification of the boundary value problem occurs in the statement of boundary conditions on Γ, the boundary of the scatterer, where Dirichlet boundary conditions $u = g$ are such that g is a random function as specified in (2.6), with mean determined by an incident plane wave $g(x_1, x_2) = -e^{ik(x_1 \cos\theta + x_2 \sin\theta)}$ at angle $\theta = \pi/4$. We assume as in [6, 7] that the random variables $\{\xi_r\}$ making up the KL expansions are uniformly distributed on an interval $\mathcal{I}_r = [-\alpha, \alpha]$, giving rise to the joint uniform distribution on $\mathcal{I} = [-\alpha, \alpha]^m$ with joint density function

$$\rho(\xi) = \left(\frac{1}{2\alpha}\right)^m . \tag{3.2}$$

The convention that $\langle \xi_i \xi_j \rangle = \delta_{ij}$ leads to the condition $\alpha = \sqrt{3}$. The stochastic

TABLE 3.1
Number of matrix-vector products required to solve $m+1$ systems arising from m-term KL-expansion, using preconditioned FGMRES. Numbers in parentheses are average iteration counts or number of block iterations.

$k = 5\pi$	$kh = .72$, $N_x = 4170$	$m = 4$	$m = 6$	$m = 8$
	Block FGMRES	35 (7)	49 (7)	54 (6)
	FGMRES	37 (7.4)	52 (7.4)	67 (7.4)
	$kh = .36$, $N_x = 16,196$	$m = 4$	$m = 6$	$m = 8$
	Block FGMRES	40 (8)	56 (8)	72 (8)
	FGMRES	45 (9.0)	63 (9.0)	81 (9.0)
$k = 10\pi$	$kh = .72$, $N_x = 16,196$	$m = 4$	$m = 6$	$m = 8$
	Block FGMRES	85 (17)	105 (15)	135 (15)
	FGMRES	153 (30.6)	214 (30.6)	276 (30.7)
	$kh = .36$, $N_x = 63,816$	$m = 4$	$m = 6$	$m = 8$
	Block FGMRES	90 (18)	119 (17)	162 (18)
	FGMRES	157 (31.4)	220(31.4)	282 (31.3)
$k = 20\pi$	$kh = .72$, $N_x = 63,816$	$m = 4$	$m = 6$	$m = 8$
	Block FGMRES	200 (40)	245 (35)	288 (32)
	FGMRES	360 (72.0)	495 (70.7)	636 (80.7)

domain $\mathcal{I} = [-\alpha, \alpha]^m$ is discretized using a uniform mesh: each of the m coordinate intervals is subdivided into n_ξ equal subintervals, resulting in $N_\xi = n_\xi^m$ elements, each of which is an m-dimensional cube with side length $h_\xi = 2\alpha/n_\xi$. Since there are no continuity requirements on the probability space, the basis functions are taken to be piecewise constants, that is, the basis function ψ_q has the value one on the cube with index q and zero elsewhere. This leads to the particularly simple diagonal structure for the Grammian matrix G of (2.12), $G = \frac{1}{N_\xi} I$. (Recall that this is always the case if the basis functions for the stochastic discretization are orthogonal with respect to the probability measure.)

To define the KL expansion, the covariance function associated with g is assumed to have the form

$$c(x_1, x_2) = \sigma^2 e^{-(|(x_1)_1 - (x_2)_1|)/2 + |(x_1)_2 - (x_2)_2|))}, \quad x_1,\, x_2 \in \Gamma.$$

A general requirement of this methodology is that the first m eigenfunctions and eigenvalues of the covariance operator, or discrete approximations to them, be available. In some circumstances, these may be obtained in closed form [12, pp. 27ff], or, alternatively, they may be approximated using a Galerkin discretization of the integral equation (2.4). For the problems we are considering, the domain of g is one-dimensional, and the Galerkin computation is inexpensive. The discrete eigenvalues and eigenvectors are computed directly from the Galerkin approximation. If the domain in question is of higher dimension, this computation can be done efficiently using sparse eigenvalue methods and fast summation techniques [14].

Table 3.1 examines the performance of the block flexible GMRES algorithm and compares it with that of the flexible GMRES algorithm (FGMRES) applied to each right hand side separately. In these tests, the stopping criterion for the solvers was for the Euclidean norm of each component of the residual to satisfy

$$\|r_j\|/\|b_j\| < 10^{-6}, \quad j = 1, \ldots, m+1.$$

For the block method, the iteration was stopped when the maximal individual residual component meets this criterion. The table shows the total number of matrix-vector products performed during the course of the computation, and in parentheses, the number of iterations required for convergence. For FGMRES, the latter number is the average for $m+1$ right hand sides; for block FGMRES, it is the number of block iterations. Note that the dimension of the spaces constructed by the block FGMRES method depends on m, the number of right hand sides, but not on the discretization parameter n_ξ associated with the stochastic domain, since F does not depend on n_ξ in (2.16).

It is clear from these results that the block methods require fewer matrix-vector products in all cases, and the difference in the number of these products becomes more dramatic as the number of right hand sides increases and also as the wave number k grows, i.e., as the problem becomes more difficult. The results provide further evidence of the mesh independent performance of the multigrid algorithm. We note, however, that as the number of steps s increases, the advantages of the block FGMRES method become less pronounced, since the overhead in generating the Krylov space grow like $m^2 s^2 N_x$, compared to $m s^2 N_x$ when the right hand sides are processed separately. A block Krylov subspace method such as QMR [10] would not suffer from this drawback, although it is not clear that this approach can be adapted to handle a nonlinear preconditioner. Because solution of the linear systems step represents a low order cost for the complete construction of statistical data (see the next section), we have not explored this issue.

4. Computation of statistical data.

Once the random function u_h of (2.10) is available, we are interested in statistical properties such as moments and probability distributions associated with it. In the case of time-harmonic wave propagation, an important quantity is the modulus $|u_h|$, which indicates the significance of the component with wave number k in the wave field. In this section, we describe the computations required to generate statistical data associated with the random function $|u_h(x,\xi)|$.

4.1. Computation of a distribution function.

Consider the construction of the probability distribution function for the maximum modulus

$$F(a) \equiv Pr(\max_x |u_h(x,\xi)| \le a). \tag{4.1}$$

Let

$$\mathcal{S}_a = \{\xi \in \mathcal{I} : \max_x |u_h(x,\xi)| \le a\}.$$

Using the definition of the joint density function (3.2), we have

$$F(a) = \int_{\mathcal{S}_a} \rho(\xi) d\xi = |\mathcal{S}_a| \left(\frac{1}{2\alpha}\right)^m.$$

FIG. 4.1. *Effect of the stochastic discretization parameter* n_ξ *on the estimate of* $Pr(\max_x |u_h(x,\xi)| \leq a)$*, for* $k = 5\pi$ *and various choices of* m*, the number of KL terms.*

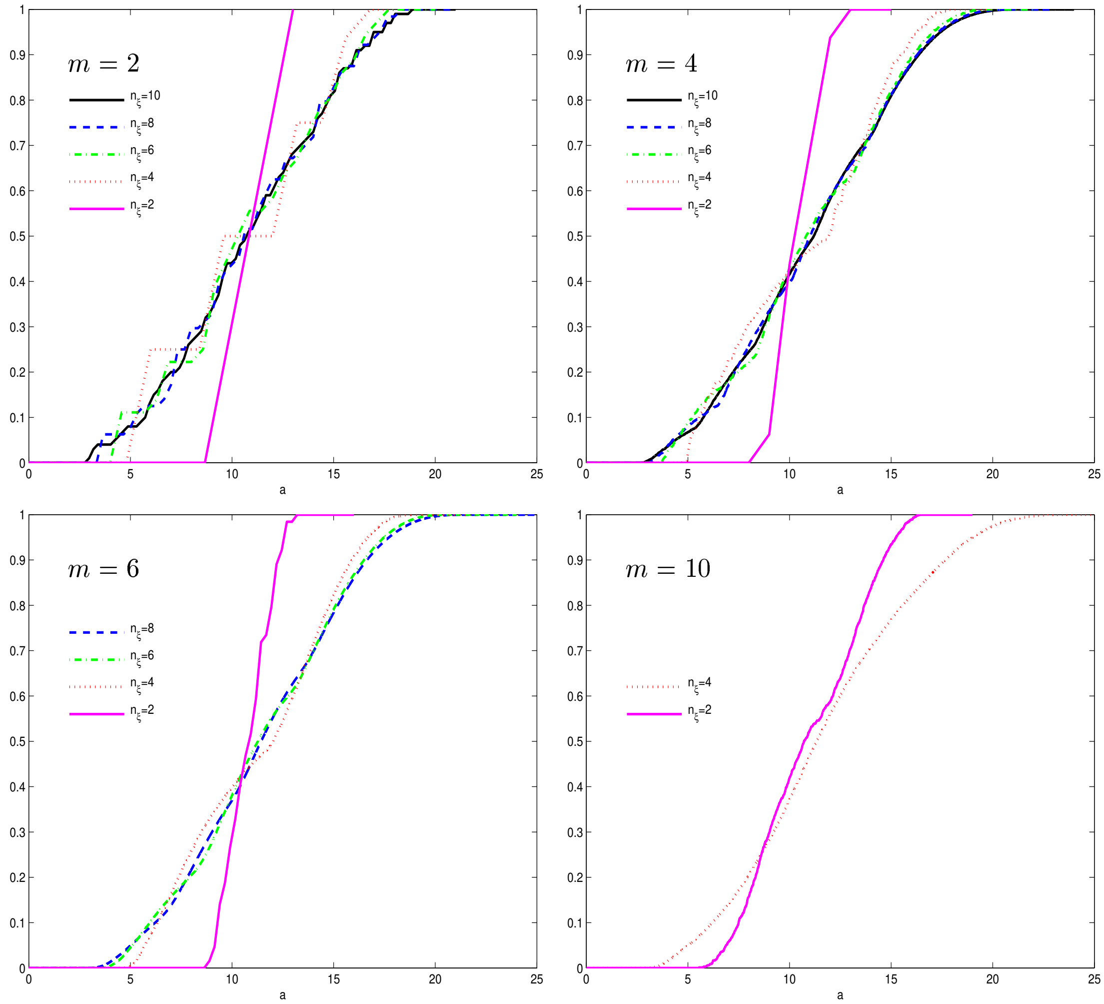

To determine the volume of $\mathcal{S}_a$, let $\xi \in \mathcal{I}$ be given, and let $q = q(\xi)$ be the index of the stochastic element $\mathcal{I}_q \subset \mathcal{I}$ containing ξ. Then

$$\max_x |u_h(x,\xi)| = \max_x \left| \sum_{j=1}^{N_x} \boldsymbol{u}_{jq} \phi_j(x) \right| = \max_j |\boldsymbol{u}_{jq}|$$

where the latter equality follows from the linearity of u_h in space. Letting

$$s_a = \left| \{ q : \max_j |\boldsymbol{u}_{jq}| \leq a \} \right|,$$

it follows that $|\mathcal{S}_a| = \frac{(2\alpha)^m}{N_\xi} s_a$, and therefore $F(a) = s_a / N_\xi$. This construction requires $\max_j |\boldsymbol{u}_{jq}|$ for each q. Once these maxima are computed, they can then be used to compute $F(a)$ for any a.

FIG. 4.2. *Effect of the number of terms m in the KL expansion on the estimate of $Pr(\max_x |u_h(x,\xi)| \le a)$, for $k = 5\pi$.*

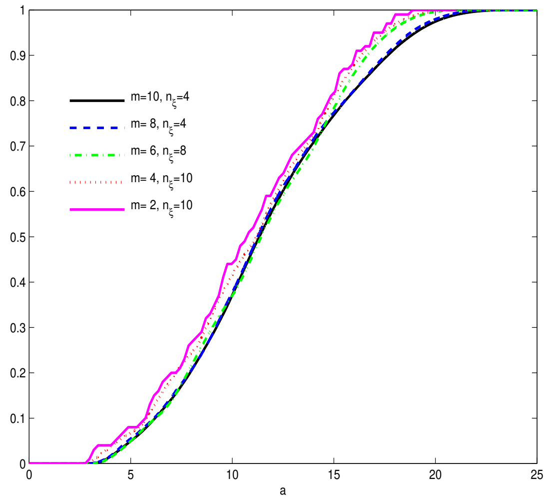

We show in Figures 4.1–4.2 the results of computing the distribution function (4.1) for various parameter values. There is no analytic expression for this quantity, so it is difficult to make a rigorous assessment of the accuracy of the computations. Nevertheless, it is possible to identify certain qualitative aspects of the results as well as to place the costs of producing them in context. First, note that discretization of the random component of the problem can be viewed from two perspectives, derived from the number of terms m used in the finite KL expansion, and, once m is fixed, from the value n_ξ of the discretization parameter in $\mathcal{I}$. Convergence of the KL expansion depends on the correlations within the process; when the finite expansion is fixed, it is shown in [7] that the error in the stochastic discretization (assuming an accurate spatial discretization) is proportional to n_ξ^{-1}. Since the number of stochastic degrees of freedom is proportional to $N_\xi = n_\xi^m$, it would be desirable to keep m as small as possible.

In Figure 4.1, we consider the impact of the two parameters m and n_ξ, for a fixed wave number $k = 5\pi$. (The spatial discretization was such that $kh = .36$.) Each subplot in this figure corresponds to a fixed value of m for which n_ξ is allowed to vary. Each plot shows convergence to a fixed curve with refinement in n_ξ, as expected. It is also noteworthy that as m is increased, the quality of the solution for fixed n_ξ appears to improve. (For example, the solution for $n_\xi = 4$ is closer to the limiting value for each successive choice $m = 2$, 4, 6.) This indicates that the constants associated with the error bounds are smaller as m increases. Figure 4.2 explores the impact of m more closely. For this example, the results suggest that $m = 8$ is an appropriate limiting value for the number of terms in the KL expansion. With $n_\xi = 4$, this yields $65,536$ spatial degrees of freedom.[1] The combination of smaller values of m together with large n_ξ (e.g., $m = 2$ and $n_\xi = 10$, yielding 1024 stochastic degrees of freedom) is able convey the qualitative structure of the distribution at significantly smaller cost.

[1]We also remark that for $m = 10$, $n_\xi = 4$ was the largest discretization parameter we could use in our MATLAB environment. This led to $N_\xi = 1,048,576$ stochastic degrees of freedom.

FIG. 4.3. *Estimated probability distribution function $Pr(\max_x |u_h(x,\xi)| \leq a)$, for various values of the wave number k.*

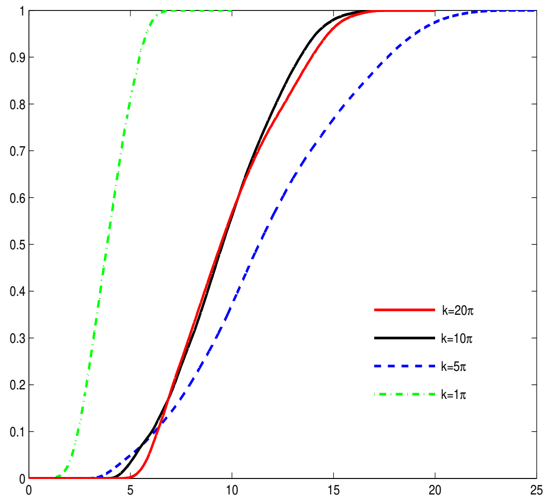

Finally, Figure 4.3 shows the estimated distribution function (4.1) for different values of k. These results suggest that this probability distribution function does not vary dramatically as the wave number increases.

4.2. Computation of higher moments. For examples of other statistical data to be computed, consider the moments of the modulus of u_h. Let $a_h(x,\xi) \equiv |u_h(x,\xi)|$, and let

$$a_h^{(\nu)}(x) \equiv \langle a_h(x,\cdot)^\nu \rangle, \quad \nu = 1, 2, 3, \ldots$$

denote the moments of a_h. We have

$$\begin{aligned}
a_h^{(\nu)}(x) &= \int_{\mathcal{I}} |u_h(x,\xi)|^\nu \rho(\xi) d\xi \\
&= \sum_{p=1}^{N\xi} \int_{\mathcal{I}_p} \left| \sum_{q=1}^{N_\xi} \left(\sum_{j=1}^{N_x} \boldsymbol{u}_{jq} \phi_j(x) \right) \psi_q(\xi) \right|^\nu \rho(\xi) d\xi \\
&= \sum_{p=1}^{N_\xi} \left(\int_{\mathcal{I}_p} \rho(\xi) d\xi \right) \left| \sum_{j=1}^{N_x} \boldsymbol{u}_{jp} \phi_j(x) \right|^\nu \\
&= \frac{1}{N_\xi} \sum_{p=1}^{N_\xi} \left| \sum_{j=1}^{N_x} \boldsymbol{u}_{jp} \phi_j(x) \right|^\nu .
\end{aligned}$$

This is straightforward to evaluate once the coefficients $\{\boldsymbol{u}_{jp}\}$ are available. In particular, the nodal values are

$$\boldsymbol{a}_i^{(\nu)} \equiv a_h^{(\nu)}(x_i) = \frac{1}{N_\xi} \sum_{p=1}^{N_\xi} |\boldsymbol{u}_{ip}|^\nu ,$$

FIG. 4.4. *Estimated mean μ_h, standard deviation σ_h, ratio σ_h/μ_h, and skewness of $|u_h|$, for $k = 5\pi$.*

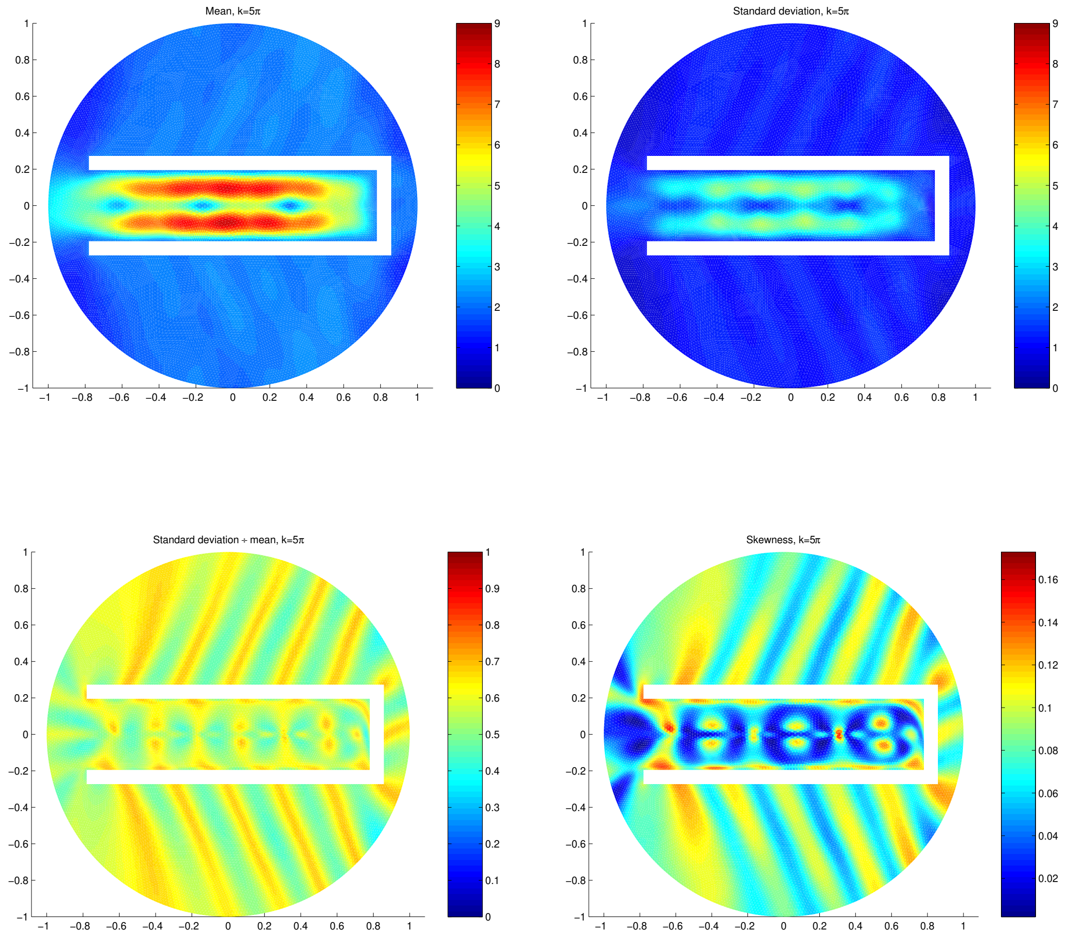

giving the piecewise linear interpolant of the ν'th moment,

$$\hat{a}_h^{(\nu)}(x) = \sum_{j=1}^{N_x} \boldsymbol{a}_j^{(\nu)} \phi_j(x).$$

The computations required for central moments $\left\langle \left(a_h - a_h^{(1)}\right)^\nu \right\rangle$ are identical in structure.

We examine some of these quantities in Figures 4.4–4.6. Four things are shown: the mean μ_h, standard deviation σ_h, ratio of standard deviation to mean, and scaled third central moment (the *coefficient of skewness* [16])

$$\frac{1}{\sigma_h^3} \left\langle \left(a_h - a_h^{(1)}\right)^3 \right\rangle$$

of a_h. The data used for these plots come from the parameter choices $m = 8$ for the truncated KL expansion, stochastic discretization parameter $n_\xi = 4$ and spatial

FIG. 4.5. *Estimated mean μ_h, standard deviation σ_h, ratio σ_h/μ_h, and skewness of $|u_h|$, for $k = 10\pi$.*

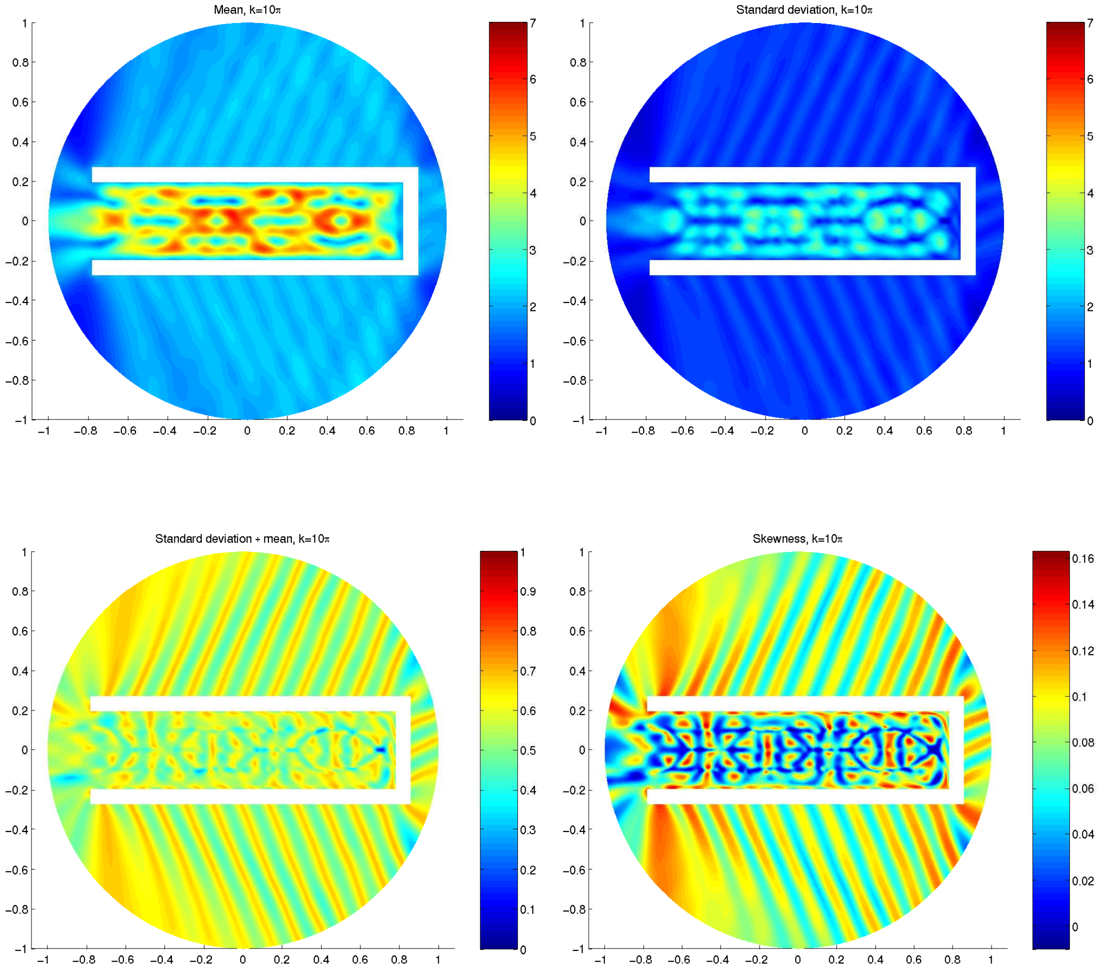

discretization satisfying $kh = .36$ for both $k = 5\pi$ and 10π and $kh = .72$ for $k = 20\pi$. Within each figure, the means and standard deviations are displayed using the same scalings. The magnitudes of the standard deviations largely mirror those of the means, and there is virtually no difference in the relative sizes of these quantities for different wave numbers. This indicates that size of the wave number k will not have a significant impact on the confidence that can be attributed to computed mean solutions. The depictions of skewness indicate that near the corner singularities, the distributions tend to be more skewed toward the right (positive direction) with respect to the mean, and inside the cavity they are skewed more toward the left; this may be of use in identifying the shape of scatterers.

Note that all these computations require the complete set of values $\{\boldsymbol{u}_{jq} : j = 1, \ldots, N_x,\ q = 1, \ldots, N_\xi\}$, which are obtained from (2.17) as

$$\boldsymbol{u}_{jq} = \sum_{r=0}^{m} v_{jr} w_{qr} \, .$$

FIG. 4.6. *Estimated mean μ_h, standard deviation σ_h, ratio σ_h/μ_h, and skewness of $|u_h|$, for $k = 20\pi$.*

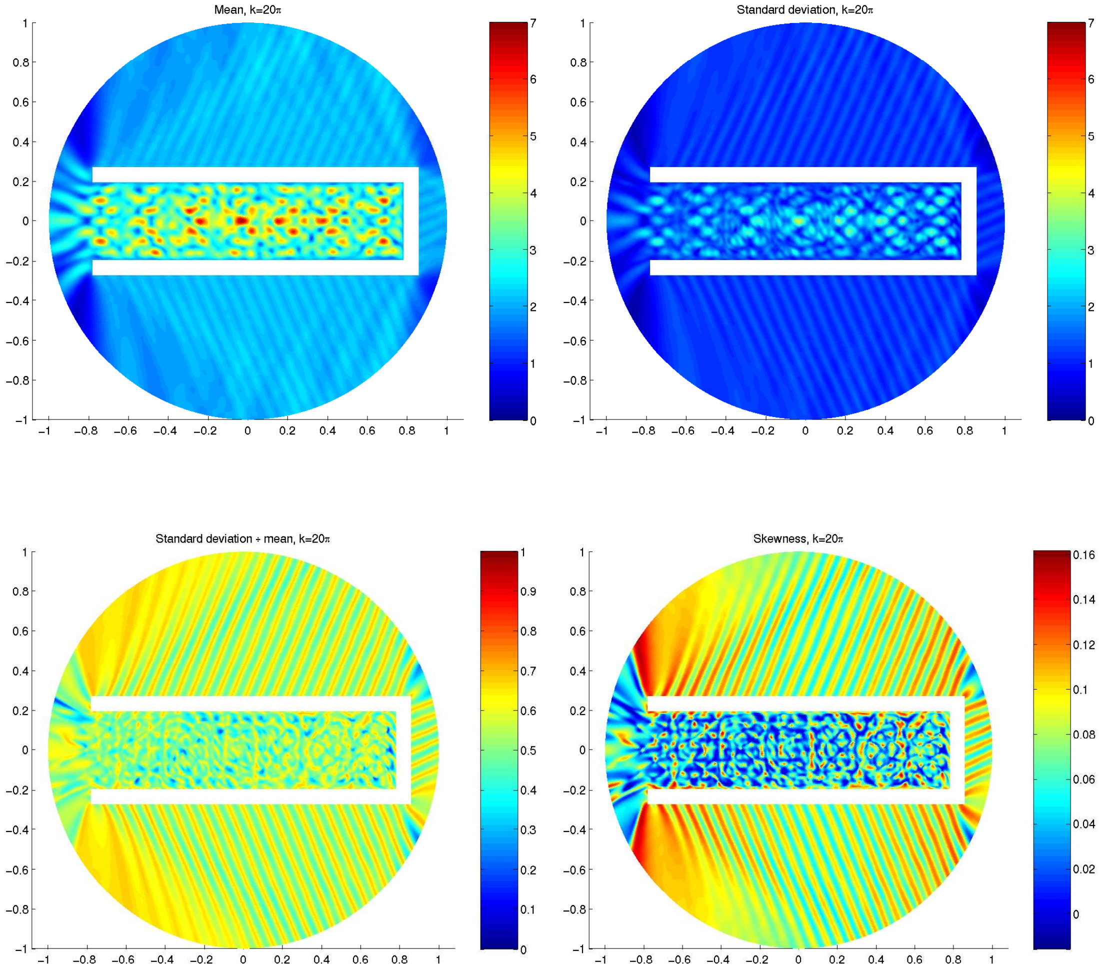

Consequently, the cost is of order $O(N_x N_\xi)$ and these computations represent the dominant expense of the process. The storage costs are also of this magnitude but can be reduced to order $m\ \max(N_x, N_\xi)$ by taking advantage of the outer-product representation (2.17) and recomputing $\boldsymbol{u}_{jq}$ whenever it is needed. The tradeoff here is a (small) additional computational expense of magnitude $O(mN_xN_\xi)$. This makes it feasible to handle large values of m or n_ξ that storage restrictions would otherwise prevent.

5. Concluding remarks. Our aim in this work was to carefully outline the computational issues associated with implementing the stochastic finite element method and processing the results for a model of acoustic scattering, where uncertainty is restricted to boundary data. We have shown that a representation of the solution in outer product form leads to significant savings in storage and also enables the relatively inexpensive computation of the random solution. The dominant cost comes from postprocessing of the solution to compute statistical data, although the outer product form in this setting reduces the storage overhead of these computations. Finally, we note that if uncertainty appears in the differential operator instead of the right hand side, then the outer product formulation of the stochastic system is not available, and this problem would be more costly to solve.

Acknowledgement: The authors wish to thank Ivo Babuška for introducing them to this topic and for many helpful discussions.

REFERENCES

[1] A. A. Nikishin an dA. Yu. Yeremin. Variable block CG algorithms for solving large symmetric positive definite linear systems on parallel computers, I: General iterative scheme. *SIAM J. Matrix Anal. and Applics.*, 16:1135–1153, 1995.

[2] W. E. Boyse and A. A. Seidl. A block QMR method for computing multiple simultaneous solutions to complex symmetric systems. *SIAM J. Sci. Comput.*, 17:263–274, 1996.

[3] T. F. Chan and W. L. Wan. Analysis of projection methods for solving linear systems with multiple right-hand sides. *SIAM J. Sci. Comput.*, 18:1698–1721, 1997.

[4] A. Chapman and Y. Saad. Deflated and augmented Krylov subspace techniques. *Numer. Linear Algebra Appl.*, 4:43–66, 1997.

[5] R. Courant and D. Hilbert. *Methods of Mathematical Physics, Volume I.* John Wiley & Sons, New York, 1989.

[6] M. K. Deb. *Solution of Stochastic Partial Differential Equations (SPDEs) Using Galerkin Method: Theory and Applications.* PhD thesis, Department of Mathematics, University of Texis at Austin, 2000.

[7] M. K. Deb, I. M. Babuška, and J. T. Oden. Solution of stochastic partial differential equations using Galerkin finite element techniques. *Comp. Meths. Appl. Mech. Engrg.*, 190:6359–6372, 2001.

[8] J. J. Dongarra, J. du Croz, I. Duff, and S. Hammarling. A set of level 3 basic linear algebra subprograms. *ACM Trans. Math. Soft.*, 16:1–17, 1990.

[9] H. C. Elman, O. G. Ernst, and D. P. O'Leary. A multigrid method enhanced by Krylov subspace iteration for discrete helmholtz equations. *SIAM J. Sci. Comput.*, 23:1291–1315, 2001.

[10] R. W. Freund and M. Malhotra. A block QMR algorithm for non-Hermitian linear systems with multiple right-hand sides. *Linear Algebra Appl.*, 254:119–157, 1997.

[11] R. Ghanem. Ingredients for a general purpose stochastic finite elements implementation. *Comp. Meths. Appl. Mech. Engrg*, 168:19–34, 1999.

[12] R. Ghanem and P. Spanos. *Stochastic Finite Elements: A Spectral Approach.* Springer-Verlag, New York, 1991.

[13] J. B. Keller and D. Givoli. Exact nonreflecting boundary conditions. *J. Comput. Phys.*, 82:172–192, 1989.

[14] R. B. Lehoucq, D. C. Sorensen, and C. Yang. *ARPACK Users' Guide.* SIAM, Philadelphia, 1998.

[15] H. G. Matthies and C. Bucher. Finite elements for stochastic media problems. *Comp. Meths. Appl. Mech. Engrg*, 168:3–17, 1999.

[16] A. M. Mood and F. A. Graybill. *Introduction to the Theory of Statistics.* McGraw-Hill, New York, 1963.

[17] D. P. O'Leary. The block conjugate gradient algorithm and related methods. *Linear Algebra Appl.*, 29:293–322, 1980.

[18] Y. Saad. A flexible inner-outer preconditioned GMRES algorithm. *SIAM J. Sci. Stat. Comput.*, 14, 1993.

[19] Y. Saad and M. H. Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Stat. Comput.*, 7:856–869, 1986.

[20] C. Schwab and R.-A. Todor. Sparse finite elements for elliptic problems with stochastic data. Technical Report Report CH-8092, Seminar f"ur Angewandte Mathematik, ETH Zurich, 2002.

[21] V. Simoncini and E. Gallopoulos. An iterative method for nonsymmetric systems with multiple right-hand sides. *SIAM J. Sci. Comput.*, 16:917–933, 1995.

[22] B. Vital. *Etude de quelques méthodes de résolution de problèmes linéaires de grande taille sur multiprosseur.* PhD thesis, Université de Rennes I, Rennes, France, Nov. 1990. cited in [21].

[23] A. M. Yaglom. *Correlation Theory of Stationary and Related Random Functions I.* Springer-Verlag, New York, 1987.