

TOWARDS EXASCALE SOLVERS FOR STOCHASTIC FEM

ITERATIVE SUBSTRUCTURING IN THE FRAMEWORK OF STOCHASTIC FINITE ELEMENTS

Howard C. Elman¹ and Bedřich Sousedík

Precise values of coefficients in the physical models using partial differential equations (PDEs) are often not known. In such situations, the coefficients are typically treated as random variables in an attempt to quantify uncertainty in the underlying problem. The most popular approach for solution of such problems is the Monte-Carlo method [21]. This method is robust and versatile but tends to exhibit slow convergence. In the last two decades, significant effort has been devoted to development of methods that leverage regularity of the solution and outperform Monte-Carlo methods, at least for problems with stochastic dimensions that are not too large. The most promising developments, known generically as *spectral methods* [11, 28], include stochastic Galerkin methods [1, 3, 11] and stochastic collocation methods [2, 20, 29]. The first approach uses a global Galerkin projection that translates the stochastic PDE into one large coupled algebraic system. The second one samples the stochastic PDE at a predetermined set of collocation points, which yields a set of uncoupled deterministic systems that can then be used to interpolate the solution in the entire random input domain. Because extending legacy software to support the collocation points is relatively simple, collocation methods are often regarded as non-intrusive. In contrast, Galerkin methods are intrusive and require the development of new solvers, although when this is done effectively, Galerkin methods may be more efficient than collocation methods [9].

It is not unusual today to see deterministic computations discretized using millions of degrees of freedom. The introduction of parameters into this mix will further increase the complexity by multiple orders of magnitude. Thus, this problem class is ripe for treatment on exascale machines. Our aim here is to explore, test and refine iterative methods for solving in an exascale environment the algebraic systems that arise from spectral treatment of stochastic PDEs. Specific directions of research include:

1. Block structure in stochastic Galerkin methods. Matrices arising from stochastic Galerkin methods have a block structure derived from the tensor product form of the discretization. Component blocks have a sparsity structure given by the underlying spatial discretization, and the full block structure depends on the expansion used for random field that represents uncertainty in the model. One such expansion, the Karhunen-Loève (KL) expansion [16], leads to block sparse structure of global matrices. Another is the generalized polynomial chaos (gPC) expansion [30], which is suitable for expansions of random fields with general probability distributions whose dependence on parameters is nonlinear, for example, where the random coefficients follow a lognormal distribution. This generality comes at a cost, however, since it leads to global matrices with a *block dense* structure. Although this is viewed as a liability of this approach, in an exascale setting this may be less of a concern. For example, it may be possible to load and process component dense blocks arising from gPC discretizations while simultaneously performing matrix-vector operations (MATVECS) associated with sparse spatial discretization in ways so that data movement associated with MATVECS is masked. These ideas can be tested in combination with recently developed [22, 23] preconditioners that take advantage of the recursive hierarchy in the (“outer”) block structure, which is due to the stochastic Galerkin projection. These methods do not require either the global matrix or the matrix of the preconditioner to be formed explicitly, and there is significant opportunity to mask communication with dense matrix operations.

¹Corresponding author: Department of Computer Science, University of Maryland College Park, MD 20742, USA, phone: (301)405-2694, email: elman@cs.umd.edu

2. Reduced-order models for stochastic PDEs. Reduced-order models [4, 5, 13, 14] are designed to make large-scale simulations more amenable to efficient solution when resources are scarce. However, many aspects of this methodology can benefit from use of large-scale computing resources. A brief statement of the approach follows. For a discrete parameter-dependent PDE $\mathcal{F}_\xi(\mathbf{u}) = 0$ of (large) order N , a set of solutions $\{\mathbf{u}_j = \mathbf{u}_{\xi_j}\}_{j=1}^n$, $n \ll N$, is obtained with the goal that $\mathcal{S}_n = \text{span}\{\mathbf{u}_j\}$ capture the dynamics of the model. Methods of uncertainty quantification [8] (and general simulation) can be performed on the smaller space using, for example, a Galerkin condition $U^{(n)T} \mathcal{F}_\xi(\mathbf{u}_R) = U^{(n)T} \mathcal{F}_\xi(U^{(n)}\boldsymbol{\eta}) = 0$ where $U^{(n)}$ is a basis for \mathcal{S}_n . Components of this process that lend themselves to large-scale computation include:

(a) *Construction of the reduced basis.* This can be done using direct search methods [4] or by optimizing an error functional [5] to identify parameters to use for the construction of \mathcal{S}^n . In a computing environment where communication overhead decreases efficiency of sparse computations, local optimizations within regions (of the parameter space) can be performed in order to produce (perhaps) sub-optimal reduced bases much more rapidly than optimal ones.

(b) *Solving the reduced-order model.* In conventional computing environments, the size of the reduced basis may be large enough to make the reduced model more expensive to solve by direct methods than applying fast (e.g., multigrid) methods to the underlying sparse problem. (We have seen [8] bases on the order of 1000 for problems with many parameters.) This difficulty may be turned around in an exascale setting, where sparse solvers could be communication bound but the processing of dense sets for reduced models are used to advantage to reduce the impact of data movement. This idea also complements the previous discussion, whereby we may deliberately *allow* the reduced basis to grow.

(c) *Overlapping construction and use of reduced bases.* When reduced basis methods are used for uncertainty quantification, it may be that a partial reduced basis is capable of resolving some parts of the parameter space. Therefore, rather than waiting for the full construction of the basis, we can allow the reduced computations to begin earlier and use error estimates to assess the utility of reduced solutions obtained from partially computed reduced bases.

3. Adaptive multilevel domain decomposition methods. We can combine these ideas with domain decomposition methods for the spatial components of problems. We have developed the *Adaptive-Multilevel BDDC* method [24], which is currently the most advanced variant of the Balancing Domain Decomposition by Constraints (BDDC) method introduced in [6, 7, 10]. It is also relatively straightforward to extend the BDDC method to a *multilevel* method [18, 26, 27]. This effectively removes the bottleneck associated with solution of coarse grid problems, which represents the primary limitation on efficiency of domain decompositions on large-scale parallel computers. The method also allows for *adaptive* selection of the components of the method [17, 19] on every decomposition level, which enables detection of the troublesome parts of the problem and improve convergence rates at modest computational cost.

Thus, our aim is to combine the natural parallelism of state-of-the-art domain decomposition methods with new approaches for treating uncertainty using block methods and reduced basis methods. Our expectation is that communication overhead can be reduced through use of dense matrix methods and overlapping of dense and sparse computation. To the best of our knowledge, use of domain decomposition solvers for stochastic computations has been studied only in a handful of publications, e.g., [12, 15, 25, 31], and the state of the art does not reflect their full potential. We expect the large-dimensional algebraic problems that arise from PDEs with uncertain coefficients to present an ideal testbed for exascale computing.

References

- [1] I. BABUŠKA, R. TEMPONE, AND G. E. ZOURARIS, *Galerkin finite element approximations of stochastic elliptic partial differential equations*, SIAM J. Numer. Anal., 42 (2005), pp. 800–825.
- [2] I. BABUŠKA, F. NOBILE, AND R. TEMPONE, *A stochastic collocation method for elliptic partial differential equations with random input data*, SIAM Review, 52 (2010), pp. 317–355.
- [3] I. BABUŠKA, R. TEMPONE, AND G. ZOURARIS, *Solving elliptic boundary value problems with uncertain coefficients by the finite element method: the stochastic formulation*, Comput. Meth. Appl. Mech. Eng., 194 (2005), pp. 1251–1294.
- [4] S. BOYAVAL, C. L. BRIS, T. LELIÈVRE, Y. MADAY, N. C. NGUYEN, AND A. T. PATERA, *Reduced basis techniques for stochastic problems*, Arch. Comput. Meth. Engin., 17 (2010), pp. 435–454.
- [5] T. BUI-THANH, K. WILLCOX, AND O. GHATTAS, *Model reduction for large-scale systems with high-dimensional parametric input space*, SIAM J. Sci. Comput., 30 (2008), pp. 3270–3288.
- [6] J.-M. CROS, *A preconditioner for the Schur complement domain decomposition method*, in Domain Decomposition Methods in Science and Engineering, I. Herrera, D. E. Keyes, and O. B. Widlund, eds., National Autonomous University of Mexico (UNAM), México, 2003, pp. 373–380. 14th International Conference on Domain Decomposition Methods, Cocoyoc, Mexico, January 6–12, 2002.
- [7] C. R. DOHRMANN, *A preconditioner for substructuring based on constrained energy minimization*, SIAM J. Sci. Comput., 25 (2003), pp. 246–258.
- [8] H. C. ELMAN AND Q. LIAO, *Reduced basis collocation methods for partial differential equations with random coefficients*, Tech. Rep. TR-2012-09, University of Maryland Institute for Advanced Computer Studies, 2012. To appear in SIAM/ASA J. Uncertainty Quantification.
- [9] H. C. ELMAN, C. W. MILLER, E. T. PHIPPS, AND R. S. TUMINARO, *Assessment of collocation and galerkin approaches to linear diffusion equations with random data*, International Journal for Uncertainty Quantifications, 1 (2011), pp. 19–33.
- [10] Y. FRAGAKIS AND M. PAPADRAKAKIS, *The mosaic of high performance domain decomposition methods for structural mechanics: Formulation, interrelation and numerical efficiency of primal and dual methods*, Comput. Methods Appl. Mech. Engrg., 192 (2003), pp. 3799–3830.
- [11] R. G. GHANEM AND P. D. SPANOS, *Stochastic Finite Elements: A Spectral Approach*, Springer-Verlag New York, Inc., New York, NY, USA, 1991.
- [12] D. GHOSH, P. AVERY, AND C. FARHAT, *A FETI-preconditioned conjugate gradient method for large-scale stochastic finite element problems*, Comput. Methods Appl. Mech. Engrg., 80 (2009), pp. 914–931.
- [13] M. A. GREPL, Y. MADAY, N. C. NGUYEN, AND A. T. PATERA, *Efficient reduced-basis treatment of nonaffine and nonlinear partial differential equations*, M2AN, 41 (2007), pp. 575–605.
- [14] B. HAASDONK AND M. OHLBERGER, *Reduced basis method for finite volume approximations of parametrized linear evolution equations*, M2AN, 42 (2008), pp. 277–301.

- [15] C. JIN, X.-C. CAI, AND C. LI, *Parallel domain decomposition methods for stochastic elliptic equations*, SIAM J. Sci. Comput., 29 (2007), pp. 2096–2114.
- [16] M. LOÈVE, *Probability Theory. II*, Graduate Texts in Mathematics, Vol. 46, Springer-Verlag, New York, 4th ed., 1978.
- [17] J. MANDEL AND B. SOUSEDÍK, *Adaptive selection of face coarse degrees of freedom in the BDDC and the FETI-DP iterative substructuring methods*, Comput. Methods Appl. Mech. Engrg., 196 (2007), pp. 1389–1399.
- [18] J. MANDEL, B. SOUSEDÍK, AND C. R. DOHRMANN, *Multispace and multilevel BDDC*, Computing, 83 (2008), pp. 55–85.
- [19] J. MANDEL, B. SOUSEDÍK, AND J. ŠÍSTEK, *Adaptive BDDC in three dimensions*, Math. Comput. Simulation, 82 (2012), pp. 1812–1831.
- [20] F. NOBILE, R. TEMPONE, AND C. G. WEBSTER, *A sparse grid stochastic collocation method for partial differential equations with random input data*, SIAM J. Numer. Anal., 46 (2008), pp. 2309–2345.
- [21] C. P. ROBERT AND G. CASELLA, *Monte Carlo Statistical Methods*, Springer-Verlag New York, Inc., Secaucus, NJ, USA, second ed., 2005. (Springer Texts in Statistics).
- [22] B. SOUSEDÍK AND R. G. GHANEM, *Truncated hierarchical preconditioning for the stochastic Galerkin FEM*. Submitted to International Journal for Uncertainty Quantification, 2013.
- [23] B. SOUSEDÍK, R. G. GHANEM, AND E. T. PHIPPS, *Hierarchical Schur complement preconditioner for the stochastic Galerkin finite element methods*, Numer. Linear Algebra Appl., (2013).
- [24] B. SOUSEDÍK, J. ŠÍSTEK, AND J. MANDEL, *Adaptive-Multilevel BDDC and its parallel implementation*, Computing, (2013).
- [25] W. SUBBER AND A. SARKAR, *Domain decomposition method of stochastic PDEs: a two-level scalable preconditioner*, Journal of Physics: Conference Series, 341 (2012), p. 012033.
- [26] X. TU, *Three-level BDDC in three dimensions*, SIAM J. Sci. Comput., 29 (2007), pp. 1759–1780.
- [27] —, *Three-level BDDC in two dimensions*, Internat. J. Numer. Methods Engrg., 69 (2007), pp. 33–59.
- [28] D. XIU, *Numerical Methods for Stochastic Computations: A Spectral Method Approach*, Princeton University Press, 2010.
- [29] D. XIU AND J. S. HESTHAVEN, *High-order collocation methods for differential equations with random inputs*, SIAM J. Sci. Comput., 27 (2005), pp. 1118–1139.
- [30] D. XIU AND G. E. KARNIADAKIS, *The Wiener-Askey polynomial chaos for stochastic differential equations*, SIAM J. Sci. Comput., 24 (2002), pp. 619–644.
- [31] K. ZHANG, R. ZHANG, Y. YIN, AND S. YU, *Domain decomposition methods for linear and semilinear elliptic stochastic partial differential equations*, Applied Mathematics and Computation, 195 (2008), pp. 630–640.