Eric Phipps

Scalable Uncertainty Quantification Through Simultaneous Propagation, Krylov Subspace Recycling and Adjoint Enhancement

Sandia National Laboratories
PO Box 5800 MS-1318
Albuquerque
NM 87185
etphipp@sandia.gov
Eric Cyr
Mike Eldred
Mike Parks

Quantifying uncertainties in computational simulations is a key challenge for predictive computational science. An important task within uncertainty quantification is the propagation of input data uncertainty to the corresponding simulation output quantities of interest. While many uncertainty propagation approaches have been investigated throughout the literature, many of these approaches involve sampling the simulation at a prescribed set of values for the input data. For example, Monte Carlo methods require evaluation of the simulation at random realizations of the input data, whereas stochastic collocation and non-intrusive polynomial chaos methods require evaluations on either structured or unstructured grids. In each of these cases, a large number of samples can be required when high accuracy is required, the space of uncertain input data has high dimension, or the simulation quantities of interest lack regularity with respect to the uncertain input data. If the simulations themselves are computationally expensive, the large number of samples can lead to intractable uncertainty quantification problems. Thus significant improvements can be made by reducing the number of samples required and reducing the cost of evaluating the simulation on an ensemble of input data realizations.

To this end, we pursue two approaches to reducing the cost of sampling-based uncertainty propagation methods. The first is adjoint enhancement whereby gradients of quantities of interest are computed alongside response values. Often the gradient with respect to an arbitrarily large number of input parameters can be computed in a small multiple of the time required for a single simulation, if the simulation code has an intrusive adjoint propagation capability. The challenge is designing uncertainty propagation approaches that incorporate gradient information effectively to reduce the overall number of samples. We present several techniques based on gradient enhancement of global orthogonal polynomial expansions and local and global interpolants, where the gradients are computed by automatic differentiation techniques. These adjoint-enhanced approaches can then be embedded within adaptive collocation procedures that estimate anisotropy or sparsity.

Second, we investigate reducing the cost of computing multiple simulation samples by propagating batches of them together. For large-scale simulation codes involving Newton-type nonlinear solvers for both steady-state and implicit time integration, this involves computing a sequence of residual vectors and Jacobian matrices for each sample, and then solving each linear system. For these methods, we leverage template-based generic programming techniques to intrusively evaluate multiple residuals/Jacobians together. This can significantly improve vectorization and data locality making the cost of propagating N samples significantly less than N times the cost of one sample. Furthermore, because the Jacobian matrices are relatively small perturbations from one another, we explore reusing a single preconditioner over a range of samples. This amortizes the cost of preconditioner construction which can be a significant contribution for large-scale computations. Finally, we apply Krylov basis recycling to accelerate convergence for families of linear systems. Recycling solvers target the solution of slowly changing sequences of linear systems, and reuse information generated from solving previous systems in the sequence to accelerate convergence for the next system in the sequence.

To demonstrate the effectiveness of these techniques, they are applied to large-scale uncertainty propagation problems in fluid dynamics that are representative of a class of challenging problems in multi-physics simulation. For solving the linear systems arising from discretizing the fluid dynamics problem, we pursue application of block-preconditioners which have good parallel scalability properties but are expensive to construct.