

Essentially Local Approaches to Exascale UQ

F. Battaglia¹, C. Beattie¹, D. Engel³, S. Gugercin¹, C.J. Roy¹, A. Sandu^{*}, M. Shahnam², E. de Sturler¹, X. Sun³, M. Syamlal², L. T. Watson¹

¹Virginia Tech, ²NETL, ³PNNL

*Computer Science Department, Virginia Tech. E-mail: sandu@cs.vt.edu. Phone: 540-231-2193.

Many applications of interest to DOE involve mathematically rigorous uncertainty quantification (UQ) methodologies for the estimation and representation of numerical, modeling, and physical uncertainties, the propagation of aleatory and epistemic uncertainties through a model, the solution of inverse problems, and decision making. For example, consider applications involving multiphase flows, such as combustion, pyrolysis, gasification, carbon capture, and chemical vapor deposition processes [14, 25]. The solutions of multiple equations for reacting flows are both computationally time and memory intensive [6, 11, 15, 24, 25]. The complex physical and chemical phenomena are not well understood, and are subject to considerable uncertainties. Estimating and modeling these uncertainties is essential for predicting quantities of interest such as the percentage of captured carbon, and for solving inverse problems such as parameter calibration using data, optimal design, and optimal control. However, performing UQ further increases the computational and memory demands of already very large simulations.

Current UQ approaches face considerable challenges at extreme scales. Approaches based on discretizing the probabilistic dimension, e.g., polynomial chaos, suffer from the curse of dimensionality. Markov chain Monte Carlo (MCMC) methods in large dimensional state spaces typically require extensive sampling of system response. The mix of aleatory and epistemic uncertainties is compounded by inherent (and frequent) system errors for various high performance computing architectures.

At the same time, the nature of UQ methodologies presents unique mathematical opportunities for achieving high end scalability and resiliency. Specifically, the effect of software and hardware errors can be handled algorithmically in the same statistical framework as other uncertainties, thus becoming part of a comprehensive and resilient approach to UQ. The key idea to achieve scalability is to develop essentially local algorithms that reveal and exploit extreme levels of concurrency between different subdomains, physical processes, independent model samples, scales, reduced order models, and levels of complexity. Essentially local methods reduce the need for expensive distant communication and can efficiently leverage complex memory hierarchies and heterogeneous nodes. New algorithms must be asynchronous, using delayed remote information, and resilient to incomplete information and hardware failures.

We discuss several areas where future mathematical innovation is the key to enabling successful UQ at the exascale: sampling strategies, linear and nonlinear solvers, local model reduction, and model validation.

Sampling strategies at the extreme scale. A promising uncertainty propagation approach is based on probability bounds analysis [7–9], which involves segregated propagation of probabilistic (i.e., aleatory) and interval (i.e., epistemic) parameters. In this approach, the outer loop must employ sampling, while the inner loop can employ various uncertainty propagation approaches [19]. Other methods for propagating both aleatory and epistemic uncertainty include evidence theory [21] and fuzzy set theory [17]. The widely used MCMC methods to sample the posterior distribution of uncertain parameters generally converge slowly, and the outlook worsens considerably if output response statistics are needed for a wide variety of distinct dynamic inputs. New MCMC strategies are needed that use only local information for the proposal probability densities and use delayed and incomplete remote information from other subsystems in order to decide acceptance/rejection.

Extreme-scale solvers. For the numerical algorithms and mathematical software underpinning UQ, nowhere are the exascale challenges more daunting than for solvers of linear and nonlinear systems of equations. New multilevel preconditioners, Krylov subspace, and multigrid methods are needed that maintain the essential components for fast convergence, but drastically reduce the communication costs. The new algorithms need to exploit smoothness, randomized algorithms, compression, model reduction, and asynchronous computations with multilevel communication to combine increased levels of local computations with carefully balanced, asynchronous global computations and minimized, multilevel global communication [13]. Promising scalable nonlinear solvers include homotopy methods (e.g., [23]), nonlinear Gauss-Seidel variants and asynchronous fixed point iteration, nonlinear multigrid, and quasi-Newton-like approximations. Important aspects will be new theory for the use of basis functions of almost locally supported translates, reduced accuracy/power subcomputations, and stochastic solver components.

Local model reduction. For extremely large models, the computational burden of MCMC can be addressed in part by strategically using approaches that utilize efficient, high-fidelity, reduced-order models (ROMs) during the creation of statistical ensembles. This use of ROMs as low cost surrogates for sampling of the system response has been well-studied[10], although even for modestly large problems the naive use of model surrogates may distort the derived posterior QoI distributions [16]. For extremely large models, the usual model reduction approaches inevitably require heavy memory traffic, moving widely distributed data across several processors. Deploying the final (global) reduced model as a surrogate for a model having a widely distributed interface with other systems also requires significant data motion and so is likely to create performance limitations in exascale settings. In order to retain the advantages that model reduction can provide for computational problems of this scale, one must rethink the way reduced models are both *constructed* and *employed*. Instead of trying to capture the underlying large-scale complex dynamics with a single (global) reduced model as is typically done, we are exploring the construction of *local* reduced models that model local dynamics, exchange information only with neighbors, yet in aggregate are able to encode the fine grain structure of the dynamic response across multiple resolution scales. Construction and adjustment of local models utilize observed system response at the various observed scales, combining POD methods [5] with response information derived from known or surmised system structure [3, 12]. Interpolatory model reduction methods are well developed and may be flexibly deployed in such settings [1, 4]. These methods have the flexibility to allow creation of ROMs that function as unbiased estimators of QoIs [2].

Model validation. Model validation is the assessment of a model response relative to experimental data [18]. The recently developed method of manufactured universes (MMU) provides a framework for the rigorous assessment of model form uncertainty [22]. MMU involves the generation of a manufactured reality, or universe, from which experimental observations can be made in the presence of experimental measurements errors in both the model inputs and the outputs. A simplified model of the system is then applied to this universe, including the propagation of input uncertainty and possibly numerical solution errors. Since the true behavior of reality is known in this manufactured universe, the exact evaluation of modeling errors can be performed. This approach could be used to compare and contrast different approaches for both estimating structural model uncertainty and calibrating (i.e., updating) the model in the presence of random experimental measurement error, experimental bias errors, modeling errors, and uncertainties (both random and those due to lack of knowledge) in the model inputs. This approach could also be used to quantify additional uncertainties that are introduced by extrapolating the model uncertainty to conditions where no experimental data are available [20], for example, from laboratory scale to industrial scale.

References

- [1] A.C. Antoulas, C.A. Beattie, and S. Gugercin. Interpolatory model reduction of large-scale dynamical systems. In J. Mohammadpour and K. Grigoriadis, editors, *Efficient Modeling and Control of Large-Scale Systems*. Springer, 2010.
- [2] Ulrike Baur, Christopher Beattie, Peter Benner, and Serkan Gugercin. Interpolatory projection methods for parameterized model reduction. *SIAM Journal on Scientific Computing*, 33(5):2489–2518, 2011.
- [3] C. Beattie and S. Gugercin. Structure-preserving model reduction for nonlinear port-hamiltonian systems. In *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, pages 6564–6569, Dec.
- [4] Christopher Beattie and Serkan Gugercin. Realization-independent \mathcal{H}_2 -approximation. In *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, pages 4953–4958. IEEE, 2012.
- [5] Christopher A Beattie, Jeff Borggaard, Serkan Gugercin, and Traian Iliescu. A domain decomposition approach to pod. In *Decision and Control, 2006 45th IEEE Conference on*, pages 6750–6756. IEEE, 2006.
- [6] M. Deza, N. P. Franka, T. J. Heindel, and F. Battaglia. CFD modeling and X-ray imaging of biomass in a fluidized bed. *ASME Journal of Fluids Engineering*, 11:111303, 131.
- [7] S. Ferson and L.R. Ginzburg. Different methods are needed to propagate ignorance and variability. *Reliability Engineering and System Safety*, 54:133–144, 1996.
- [8] S. Ferson and J.G. Hajagos. Arithmetic with uncertain numbers: rigorous and (often) best possible answers. *Reliability Engineering and System Safety*, 85:135–152, 2004.
- [9] S. Ferson, W.L. Oberkampf, and A. Ginzburg. Model validation and predictive capability for the thermal challenge problem. *Computer Methods in Applied Mechanics and Engineering*, 2008.
- [10] M Frangos, Y Marzouk, K Willcox, and B van Bloemen Waanders. Surrogate and reduced-order modeling: A comparison of approaches for large-scale statistical inverse problems. *Large-Scale Inverse Problems and Quantification of Uncertainty*, pages 123–149, 2010.
- [11] Y. Ge, M. S. Gordon, F. Battaglia, and R. O. Fox. heoretical study of the pyrolysis of methyl-trichlorosilane in the gas phase. 3. Reaction rate constant calculations. *Journal of Physical Chemistry A*, 114(6):2384–2392, 2010.
- [12] Serkan Gugercin, Rostyslav V. Polyuga, Christopher Beattie, and Arjan van der Schaft. Structure-preserving tangential interpolation for model reduction of port-hamiltonian systems. *Automatica*, 48(9):1963 – 1974, 2012.
- [13] J. He, L. T. Watson, and M. Sosenkina. Algorithm 897: VTDIRECT95: serial and parallel codes for the global optimization algorithm DIRECT. *ACM Trans. Math. Software*, 36:Article 17, 1–17, 2009.
- [14] S. K. Kanholly, J. Chodak, B. Y. Lattimer, and F. Battaglia. Modeling and predicting gas-solid fluidized bed dynamics to capture nonuniform inlet conditions. *ASME Journal of Fluids Engineering*, 134(11):111303, 2012.
- [15] D. Law, S. J. Jones, T. J. Heindel, and F. Battaglia. A combined numerical and experimental study of hydrodynamics for an air-water external loop airlift reactor. *ASME Journal of Fluids Engineering*, 133(2):021301, 2011.

- [16] Jing Li and Dongbin Xiu. Evaluation of failure probability via surrogate models. *Journal of Computational Physics*, 229(23):8966–8980, 2010.
- [17] Baoding Liu. *Uncertain Theory: An Introduction to its Axiomatic Foundations*. Springer-Verlag, 2004.
- [18] W. L. Oberkampf and C. J. Roy. *Verification and Validation in Scientific Computing*. Cambridge University Press, 2010.
- [19] C. J. Roy and M. S. Balch. A holistic approach to uncertainty quantification with application to supersonic nozzle thrust. *International Journal for Uncertainty Quantification*, 2012.
- [20] C. J. Roy and W. L. Oberkampf. A comprehensive framework for verification, validation, and uncertainty quantification in scientific computing. *Computer Methods in Applied Mechanics and Engineering*, 200:2131–2144, 2011.
- [21] G. Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, 1976.
- [22] H. F. Stripling, M. L. Adams, R. G. McClarren, and B. K. Mallick. The method of manufactured universes for validating uncertainty quantification methods. *Reliability Engineering and System Safety*, 96:1242–1256, 2011.
- [23] H.-J. Su, J. M. McCarthy, M. Sosonkina, and L. T. Watson. Algorithm 857: POLSYS_GLP—a parallel general linear product homotopy code for solving polynomial systems of equations. *ACM Trans. Math. Software*, 32:561–579, 2006.
- [24] J. Sun, F. Battaglia, and S. Subramaniam. Hybrid two-fluid dem simulation of gas-solid fluidized beds. *ASME Journal of Fluids Engineering*, 129(11):1394–1403, 2007.
- [25] N. Xie, F. Battaglia, , and S. Pannala. Effects of using two- versus three-dimensional computational modeling of fluidized beds: Part II, analysis. *Powder Technology*, 182(1):14–24, 2008.