Improving Performance of Uncertainty Quantification Methods on Advanced Computing Architectures

Eric Phipps, H. Carter Edwards, Jonathan Hu
Sandia National Laboratories

Problem

Uncertainty quantification (UQ) is an important scientific driver for pushing to the exascale, potentially enabling rigorous and accurate predictive simulation for many problems that are intractable today.

Nearly all UQ approaches repeatedly sample deterministic simulation codes at different realizations of the input data, resulting in performance limited to that of each realization.

Many PDE simulations do not achieve high performance on multicore (CPU/GPU/Accelerator) architectures due to:

- Random, uncoalesced memory accesses
- Inability to exploit consistent vectorization

Approach

In many cases, the code path, processor instructions, and memory access patterns are very similar from realization to realization.

Idea: Propagate a UQ information together through the forward simulation leveraging reuse and fine-grained parallelism.

FEM Residual Equations $f(u,y)=0 \implies F(U,Y)=0$

Embedded Ensemble Propagation

Ensemble PDE Solution ***Input Data Ensemble*** ***Ensemble Residual***

$$U=\sum_{i=1}^m u_i \otimes e_i, \quad Y=\sum_{i=1}^m y_i\otimes e_i, \quad F=\sum_{i=1}^m f(u_i,y_i)\otimes e_i$$

Ensemble Jacobian

$$\frac{\partial F}{\partial U}=\sum_{i=1}^m \frac{\partial f}{\partial u_i}\otimes e_i e_i^T$$

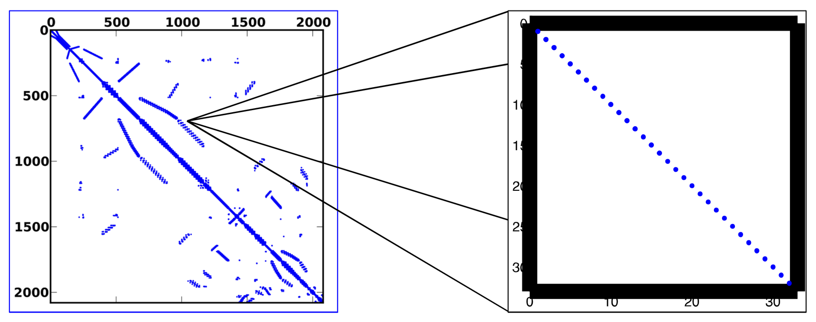

Embedded Stochastic Galerkin (SG)

Polynomial chaos approximation

$$u(y)\approx \sum_{i=0}^P u_i\psi_i(y), \quad u_i=\langle u\psi_i\rangle \equiv \int_\Gamma u(y)\psi_i(y)\rho(y)dy, \quad \langle \psi_i\psi_j\rangle=\delta_{ij}$$

SG Solution ***SG Parameter*** ***SG Residual***

$$U=\sum_{i=0}^P u_i\otimes e_i, \quad Y=\sum_{i=0}^P y_i\otimes e_i, \quad F=\sum_{i=0}^P f_i\otimes e_i, \quad f_i=\langle f\psi_i\rangle$$

SG Jacobian

$$\frac{\partial F}{\partial U}\approx \sum_{k=0}^P A_k\otimes G_k, \quad A_k=\langle \frac{\partial f}{\partial u}\psi_k\rangle, \quad G_k(i,j)=\langle \psi_i\psi_j\psi_k\rangle$$

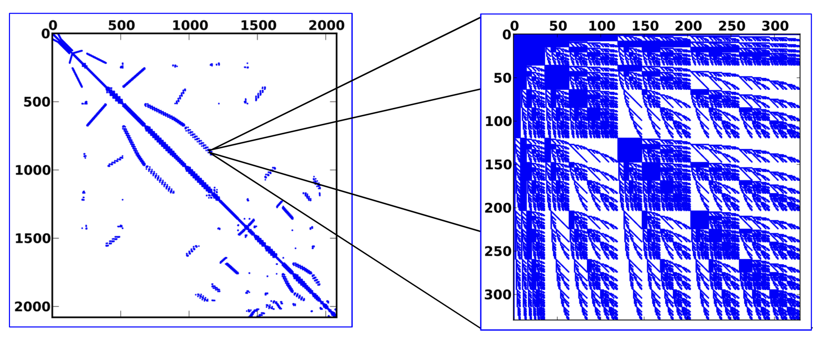

Implementation

- Propgate UQ information (samples, SG coefficients) at scalar level of simulation
- Each UQ-dependent datum becomes a small array
- Map coarse-grained UQ parallelism to fine-grained hardware parallelism (threads, SIMD, SIMT) operating on array

Apply to C++ PDE codes via template-based generic programming:

- Template PDE assembly/solver code on scalar type
- Instantiate template code on UQ scalar type
- UQ scalar type implements all relevant operations using fine-grained parallelism
- Use *thread team interface* for kernel launch and functor

Approach implemented by **Stokhos** embedded uncertainty quantification library on top of **Kokkos** portable manycore performance library (H.C. Edwards, D. Sunderland, C. Trott).

Potential speed-up in PDE residual/Jacobian evaluation:

- Amortize MPI latency in halo exchange
- Reuse mesh data info
- Contiguous loads/stores in gather/scatter
- Vector arithmetic

Potential speed-up in sparse solvers:

- Amortize MPI latency in halo exchange, dot products
- Reuse matrix graph info
- Contiguous loads in matrix-vector products
- Vector arithmetic
- Increased work/thread for smaller matrices in algebraic-multigrid preconditioners

Proxy Application

Techniques prototyped in FENL proxy application:

$$-\nabla\cdot(\kappa(x,y)\nabla u)+u^2=0$$

- 3-D, linear FEM discretization on unit cube, unstructured mesh
- "KL-like" random field model for uncertain diffusion coefficient
- Hybrid MPI+X parallelism via **Tpetra + Kokkos**
- Thread-scalable PDE assembly using Kokkos
- CG linear solver via **Belos**
- Smoothed aggregation algebraic multi-grid using **MueLu**
- Embedded ensemble propagation via Smolyak sparse-grids
- Embedded stochastic Galerkin
- Released with **Trilinos-Couplings** package

Performance of embedded approaches compared to traditional non-intrusive polynomial chaos sampling via Smolyka sparse grids:

$$u_i=\int_\Gamma u(y)\psi_i(y)\rho(y)dy\approx \sum_{k=0}^Q w^k u^k\psi_i(y^k), \quad f(u^k,y^k)=0$$

Trilinos Equinox

Computational Results

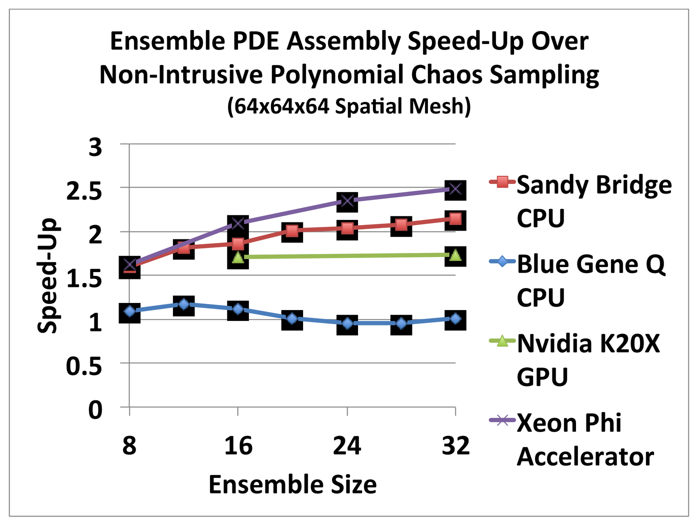

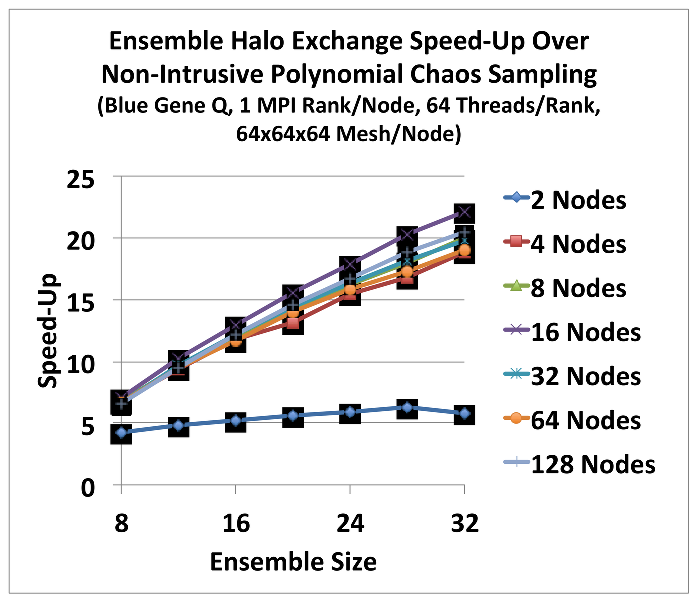

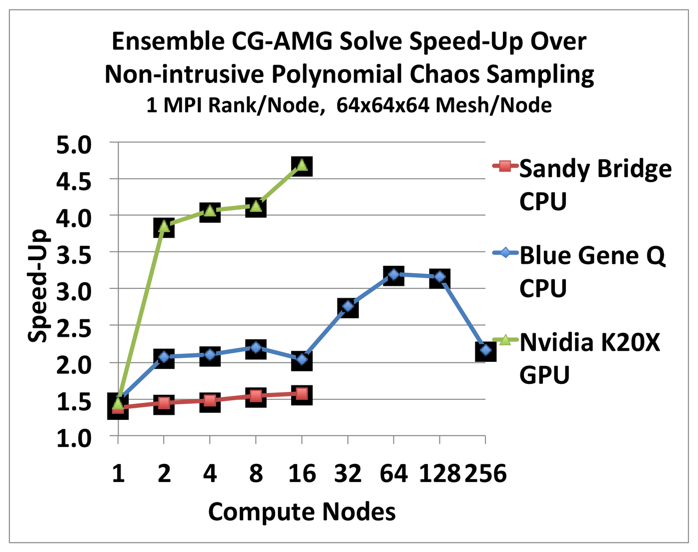

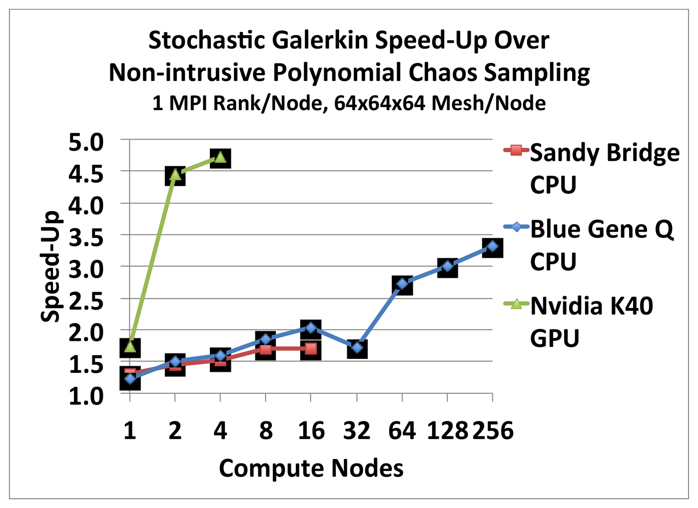

Significance

- Substantially improved aggregate UQ performance on emerging multicore and manycore architectures
- Strategy for ensuring high performance for UQ calculations on next-generation platforms
- Mitigates need for application to fully utilize fine-grained parallelism resources
- Path for impacting applications that doesn't require developers to explicitly manage uncertainty propagation

Support provided by:

U.S. DEPARTMENT OF ENERGY Office of Science

LDRD LABORATORY DIRECTED RESEARCH & DEVELOPMENT

UNCLASSIFIED

SAND 2014-xxxx

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Sandia National Laboratories

NNSA National Nuclear Security Administration