

---

Christian Ketelsen  
**An Element-Free Adaptive Algebraic Multigrid Method  
for Markov Chain Monte Carlo Simulations**

Center for Applied Scientific Computing  
Lawrence Livermore National Laboratory  
Box 808  
L-560  
Livermore  
CA 94551-0808  
`ketelsen1@llnl.gov`  
Panayot S. Vassilevski

We are interested in solving a 2<sup>nd</sup> -order elliptic equation where there is some uncertainty in the background field,  $a$ :

$$\mathcal{L}(a(x, \omega), u(x, \omega)) = f, \tag{1}$$

subject to some boundary conditions. Monte Carlo methods have, for some time, been the standard solution method for such problems. In this setting, background field  $a(x, \omega)$  is assumed to come from a prescribed probability distribution. A great many samples of  $a$  are drawn and the forward problem is solved to obtain an ensemble of solutions  $\{u(x, \omega_i)\}_{i=1}^N$ . Moments of the ensemble can then be computed to quantify uncertainty in the solution.

Often times, background field  $a$  belongs to a probability distribution that is difficult to sample from. Such is the case in simulations of subsurface flow and quantum field theory (QFT). In such cases, sampling is achieved using the computationally expensive Markov Chain Monte Carlo (MCMC) method. In MCMC, samples are obtained via an accept-reject process. Given the current realization  $a(x, \omega_k)$ , proposed realization  $a(x, \omega'_k)$  is generated via some simple transition rule (perturbing the value in each grid cell with a normal random variable, for example). The forward solution of (1) is then computed and the proposal is accepted or rejected based on some functional of the proposed solution. The vast majority of proposals are rejected, making the solution of the discrete forward problems the major computational bottleneck in the sampling procedure. At each step in the MCMC method, a discretized version of (1) must be solved. This leads to a sequence of linear systems with changing matrix coefficients:

$$A(\omega)x = b. \tag{2}$$

If the correlation lengths in the stochastic process defining  $a$  are small compared to the size of the domain, the solution of even a single realization of (2) is nontrivial and may require adaptivity. Our goal is: assuming that we have constructed an efficient multilevel solver for (2) based on some realization of  $A(\omega)$ , we would like to efficiently construct a multilevel solver for (2) based on some new realization,  $A(\hat{\omega})$ . The nature of MCMC makes this possible because consecutive realizations of the background  $a$  are often similar.

In this talk, we will describe one such adaptive two-level solver in the setting of smoothed aggregation (or SA) algebraic multigrid (AMG) which is outlined as follows. Based on a given hierarchy for  $A(\omega)$ , we build a respective two-level solver for the matrix  $A(\hat{\omega})$  corresponding to the new parameter  $\hat{\omega}$ . We test the old hierarchy by iterating on the homogeneous problem  $A(\hat{\omega})e = 0$ . If convergence is acceptable we retain the old hierarchy. If convergence stalls, then after several iterations on the homogeneous problem an algebraically smooth error mode,  $x_{\text{bad}}$ , has been exposed. This mode is added to the coarse space to improve performance. This is achieved by first localizing the problem to every aggregate exploiting an “element-free” approach. For the localized quadratic forms, we solve generalized eigenvalue problem. The cost of this computation is greatly reduced by searching only in the subspace spanned by  $x_{\text{bad}}$  and the basis functions from the original hierarchy. The choice of this subspace is motivated by the fact that  $A(\omega)$  and  $A(\hat{\omega})$ , obtained from MCMC, are in some sense nearby, hence some of the old basis functions may be adequate whereas the missing part is provided by  $x_{\text{bad}}$ . The new (unsmoothed) interpolant is then built based on the first few of the low-lying modes. The adaptive procedure is repeated as needed.