Multilevel Quadrature Methods

Hans-Werner van Wyk
with Max Gunzburger
FSU

Abstract Stochastic sampling methods (such as Monte Carlo methods) are used to simulate physical systems whose model parameters are uncertain. To maintain a given level of accuracy, the spatial fidelity of the physical system being simulated and the number of samples used to estimate the stochastic quantity of interest should be sufficiently high, which could be computationally expensive. Multilevel methods aim to achieve the same overall accuracy as traditional sampling methods but at a much reduced computational cost, by making use of a series of simulation models instead of just one, each with a different level of spatial detail.

Single Level Methods

The process of computing an approximation v_h of a physical quantity of interest v from the output u_h of a numerical simulation with input parameter q is illustrated in the following diagram.

$$q \to \boxed{\text{Numerical Simulator}} \to u_h \to \boxed{\mathcal{G}} \to v_h = \mathcal{G}(u_h)$$

Here, h represents the level of spatial accuracy. In the presence of uncertainty, the parameter q is replaced by a sample of parameters giving rise to a sample of the quantity of interest. We are now interested in approximating v's expected value $\mathbb{E}[v]$ by means of a numerical quadrature scheme $\mathcal{I}_N$.

$$\left\{q^{(i)}\right\}_{i=1}^{N} \to \boxed{\text{Numerical Simulator \& } \mathcal{G}} \to \left\{v_h^{(i)}\right\}_{i=1}^{N} \to \boxed{\mathcal{I}_N} \to \mathcal{I}_N[v_h] \approx \mathbb{E}[v]$$

The overall approximation error can be decomposed as follows:

$$\|\mathbb{E}[v] - I_N[v_h]\| \leq \underbrace{\|\mathbb{E}[v - v_h]\|}_{\text{Spatial Error}} + \underbrace{\|\mathbb{E}[v_h] - \mathcal{I}_N[v_h]\|}_{\text{Sampling Error}}$$

Letting $h \downarrow 0^+$ improves the spatial fidelity of each numerical sample (and hence the spatial error), while increasing $N \uparrow \infty$ leads to a more accurate estimate $\mathcal{I}_N$ of the stochastic integral $\mathbb{E}$ (thus reducing the sampling error). Both improvements increase the computational cost.

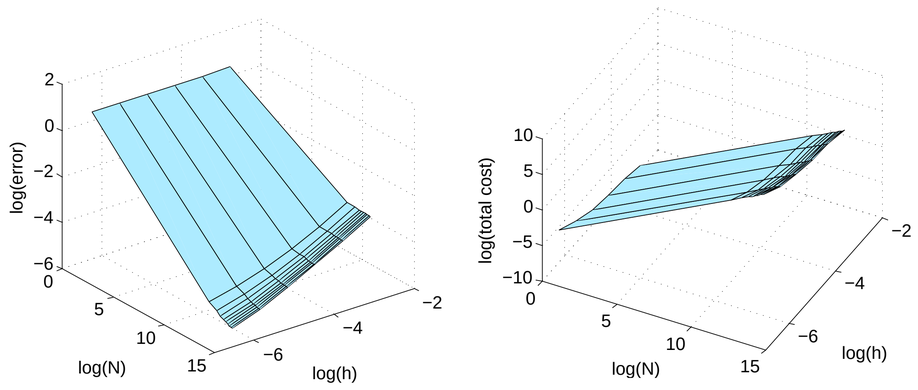

(a) Total error (b) Computational cost

Figure 1: The total error and computational cost as a function of h and N.

The Multilevel Idea

In order to achieve the same level of accuracy but at a reduced cost, we make use of a family of spatial approximations $\{v_{h_\ell}\}_{\ell=0}^{L}$ instead of a single approximation v_h. Here, h_0 is the coarsest (and cheapest) approximation level and h_L is chosen so that the spatial error is small enough. The expected value $\mathbb{E}[v]$ can now be approximated by:

$$\mathbb{E}[v] \approx \mathbb{E}[v_{h_L}] = \mathbb{E}[v_{h_0}] + \sum_{\ell=1}^{L} \mathbb{E}[v_{h_\ell} - v_{h_{\ell-1}}] \approx \mathcal{I}_{N_0}[v_{h_0}] + \sum_{\ell=1}^{L} \mathcal{I}_{N_\ell}[v_{h_\ell} - v_{h_{\ell-1}}],$$

i.e. by an initial estimate on a coarse grid h_0, together with a series of correction terms.

This formulation allows us to choose a different sample size N_ℓ to approximate the expectation of each correction term, the samples of which are computed using numerical simulations at spatial refinement level h_ℓ.

By choosing the sample sizes optimally, based on the decay rate of the correction terms and on the computational effort required to generate samples at each refinement level ℓ, we can achieve a considerable speed-up in the overall computational time.

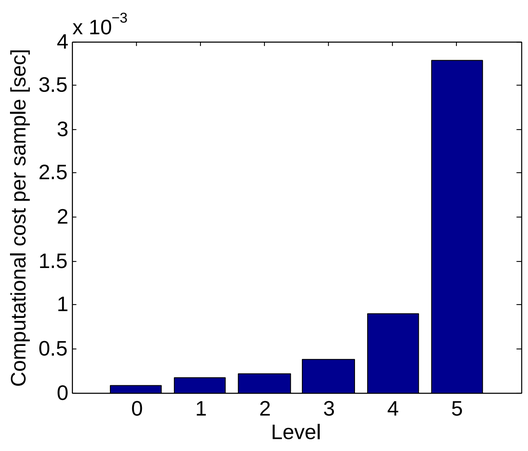

Figure 2: CPU cost per level

If C_ℓ is the computational cost of computing one sample at refinement level ℓ, then the minimization problem whose solution yields the optimal sample sizes $\{N_\ell\}$, takes the form

$$\min_{\{N_\ell\}} \sum_{\ell=0}^{L} N_\ell C_\ell \;\text{ subject to: Total Error} \leq \varepsilon,$$

where ε is the required error tolerance.

Example

As an illustrative example, consider the following simple two-point boundary value problem on $(0, 1)$

$$\begin{aligned} -\frac{d}{dx}\left(q(x,\omega)\frac{du}{dx}\right) &= 0 \;\text{ in } (0,1) \\ u(0) = 1,\; u(1) &= 0. \end{aligned} \tag{1}$$

We are interested in the expected value of the flux at the right endpoint $\mathbb{E}[v] = -\mathbb{E}[q(1)u'(1)]$. For this example we use finite volume methods to solve the underlying spatial problem and Monte Carlo averaging to estimate the expectation.

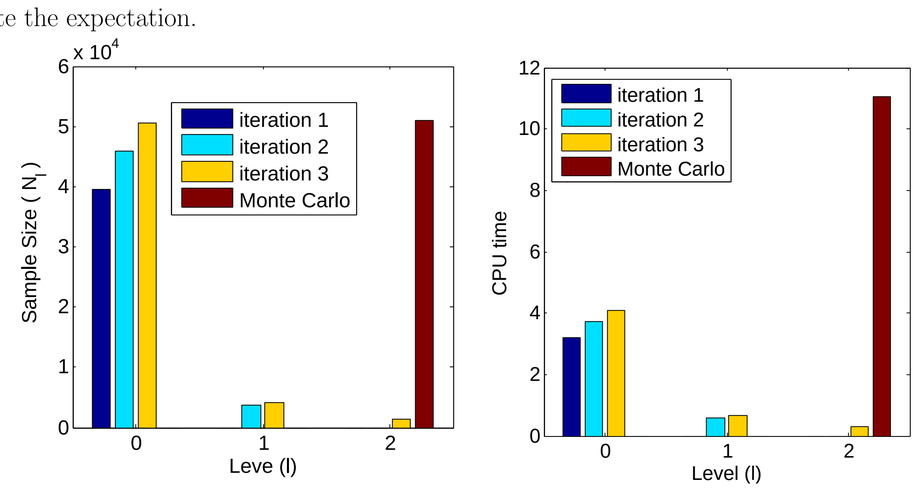

(a) Sample size per level
$N_{tot}^{MLMC} = 56030$, $N_{tot}^{MC} = 51087$

(b) Computational cost per sample
$T_{tot}^{MLMC} = 5.0623$, $T_{tot}^{MC} = 11.0718$

Figure 3: The sample sizes and computational costs associated with each level ℓ at various iterations of the multilevel algorithm.

For each iteration of the multilevel algorithm, the bulk of the samples are computed at the coarsest spatial refinement level, whereas the single level Monte Carlo method computes all samples at the highest refinement level, to guarantee that the spatial error is within tolerance. Even though the total sample size needed by the multilevel method exceeds that of the single level method, the computational time is reduced by half.

Adaptive Mesh Refinement

If the physical output of interest is spatially varying, then the correction terms $v_{h_\ell} - v_{h_{\ell-1}}$ can be used to guide adaptivity. We refine the grid in areas where the quantity of interest changes considerably from one refinement level to the next.

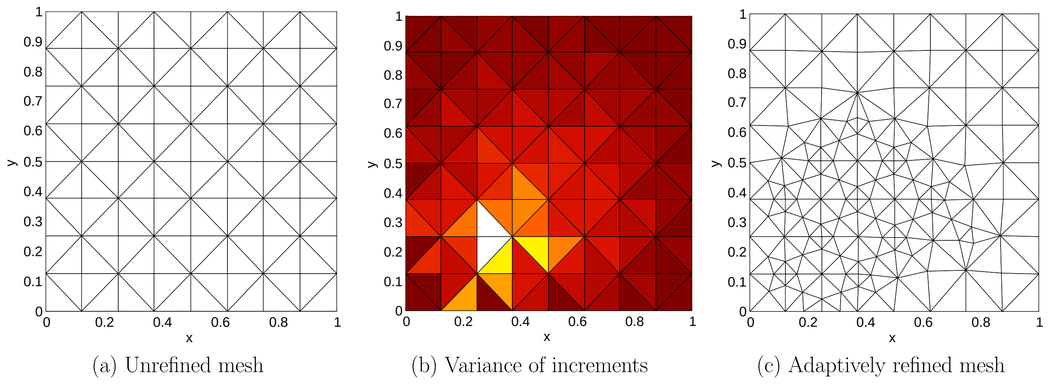

(a) Unrefined mesh (b) Variance of increments (c) Adaptively refined mesh

Figure 4: The sample sizes and computational costs associated with each level ℓ at various iterations of the multilevel algorithm.