

Multilevel Monte Carlo methods

Mike Giles

Mathematical Institute, University of Oxford

SIAM Conference on Uncertainty Quantification

April 5-8, 2016

Acknowledgements to many collaborators:

Frances Kuo, Ian Sloan (UNSW), Des Higham, Xuerong Mao (Strathclyde),
Rob Scheichl, Aretha Teckentrup (Bath), Andrew Cliffe (Nottingham),
Ruth Baker, Ben Hambly, Christoph Reisinger, Endre Süli (Oxford),
Klaus Ritter (Kaiserslautern), Lukas Szpruch (Edinburgh),
Jaime Peraire, Ferran Vidal-Codina (MIT), ...

Objectives

In presenting the multilevel Monte Carlo method, I hope to emphasise:

- the simplicity of the idea
- its flexibility – it's not prescriptive, more an approach
- there are lots of people working on a variety of applications

In doing this, I will focus on ideas rather than lots of numerical results.

Monte Carlo method

In stochastic models, we often have

$$\begin{array}{ccc} \omega & \longrightarrow & S \\ \text{random input} & & \text{intermediate variables} & \longrightarrow & P \\ & & & & \text{scalar output} \end{array}$$

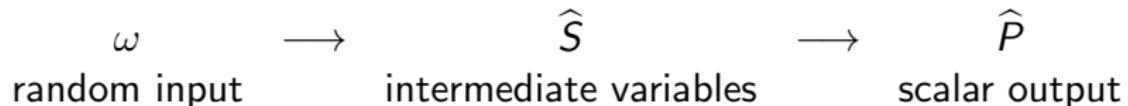
The Monte Carlo estimate for $\mathbb{E}[P]$ is an average of N independent samples $\omega^{(n)}$:

$$Y = N^{-1} \sum_{n=1}^N P(\omega^{(n)}).$$

This is unbiased, $\mathbb{E}[Y] = \mathbb{E}[P]$, and the Central Limit Theorem proves that as $N \rightarrow \infty$ the error becomes Normally distributed with variance $N^{-1} \mathbb{V}[P]$.

Monte Carlo method

In many cases, this is modified to



where \widehat{S}, \widehat{P} are approximations to S, P , in which case the MC estimate

$$\widehat{Y} = N^{-1} \sum_{n=1}^N \widehat{P}(\omega^{(n)})$$

is biased, and the Mean Square Error is

$$\mathbb{E}[(\widehat{Y} - \mathbb{E}[P])^2] = N^{-1} \mathbb{V}[\widehat{P}] + (\mathbb{E}[\widehat{P}] - \mathbb{E}[P])^2$$

Greater accuracy requires larger N and smaller weak error $\mathbb{E}[\widehat{P}] - \mathbb{E}[P]$.

SDE Path Simulation

My interest was in SDEs (stochastic differential equations) for finance, which in a simple one-dimensional case has the form

$$dS_t = a(S_t, t) dt + b(S_t, t) dW_t$$

Here dW_t is the increment of a Brownian motion – Normally distributed with variance dt .

This is usually approximated by the simple Euler-Maruyama method

$$\hat{S}_{t_{n+1}} = \hat{S}_{t_n} + a(\hat{S}_{t_n}, t_n) h + b(\hat{S}_{t_n}, t_n) \Delta W_n$$

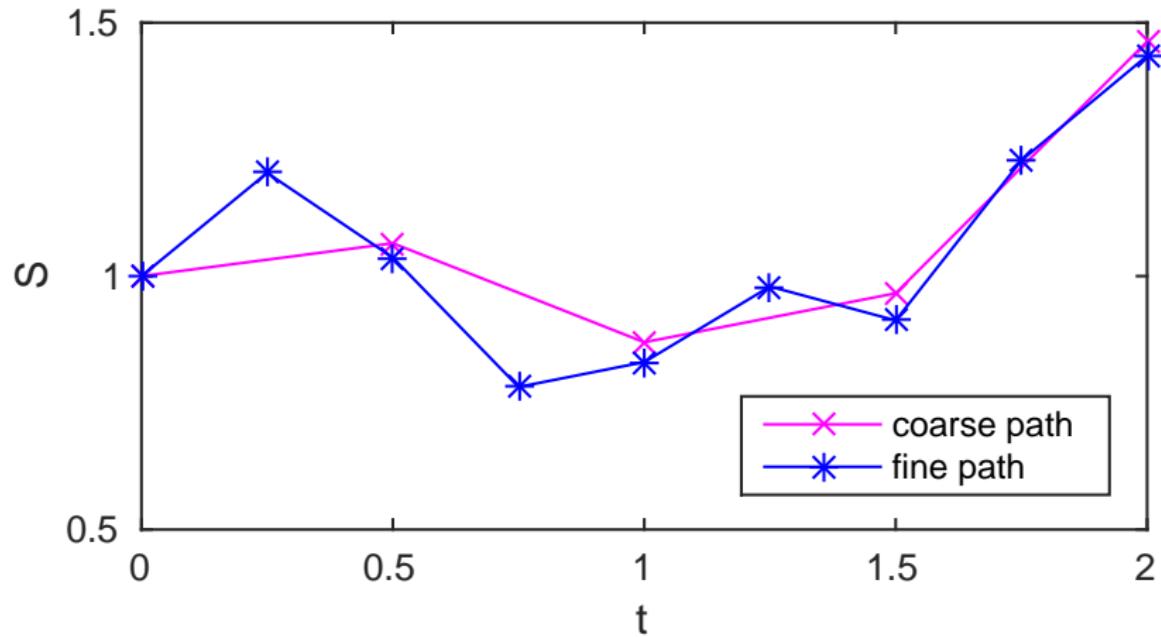
with uniform timestep h , and increments ΔW_n with variance h .

In simple applications, the output of interest is a function of the final value:

$$\hat{P} \equiv f(\hat{S}_T)$$

SDE Path Simulation

Geometric Brownian Motion: $dS_t = r S_t dt + \sigma S_t dW_t$



SDE Path Simulation

Two kinds of discretisation error:

Weak error:

$$\mathbb{E}[\hat{P}] - \mathbb{E}[P] = O(h)$$

Strong error:

$$\left(\mathbb{E} \left[\sup_{[0, T]} \left(\hat{S}_t - S_t \right)^2 \right] \right)^{1/2} = O(h^{1/2})$$

For reasons which will become clear, I prefer to use the Milstein discretisation for which the weak and strong errors are both $O(h)$.

SDE Path Simulation

The Mean Square Error is

$$N^{-1} \mathbb{V}[\hat{P}] + \left(\mathbb{E}[\hat{P}] - \mathbb{E}[P] \right)^2 \approx a N^{-1} + b h^2$$

If we want this to be ε^2 , then we need

$$N = O(\varepsilon^{-2}), \quad h = O(\varepsilon)$$

so the total computational cost is $O(\varepsilon^{-3})$.

To improve this cost we need to

- reduce N – variance reduction or Quasi-Monte Carlo methods
- reduce the cost of each path (on average) – MLMC

Two-level Monte Carlo

If we want to estimate $\mathbb{E}[\hat{P}_1]$ but it is much cheaper to simulate $\hat{P}_0 \approx \hat{P}_1$, then since

$$\mathbb{E}[\hat{P}_1] = \mathbb{E}[\hat{P}_0] + \mathbb{E}[\hat{P}_1 - \hat{P}_0]$$

we can use the estimator

$$N_0^{-1} \sum_{n=1}^{N_0} \hat{P}_0^{(0,n)} + N_1^{-1} \sum_{n=1}^{N_1} \left(\hat{P}_1^{(1,n)} - \hat{P}_0^{(1,n)} \right)$$

Benefit: if $\hat{P}_1 - \hat{P}_0$ is small, its variance will be small, so won't need many samples to accurately estimate $\mathbb{E}[\hat{P}_1 - \hat{P}_0]$, so cost will be reduced greatly.

Multilevel Monte Carlo

Natural generalisation: given a sequence $\hat{P}_0, \hat{P}_1, \dots, \hat{P}_L$

$$\mathbb{E}[\hat{P}_L] = \mathbb{E}[\hat{P}_0] + \sum_{\ell=1}^L \mathbb{E}[\hat{P}_\ell - \hat{P}_{\ell-1}]$$

we can use the estimator

$$N_0^{-1} \sum_{n=1}^{N_0} \hat{P}_0^{(0,n)} + \sum_{\ell=1}^L \left\{ N_\ell^{-1} \sum_{n=1}^{N_\ell} \left(\hat{P}_\ell^{(\ell,n)} - \hat{P}_{\ell-1}^{(\ell,n)} \right) \right\}$$

with independent estimation for each level of correction

Multilevel Monte Carlo

If we define

- C_0, V_0 to be cost and variance of \hat{P}_0
- C_ℓ, V_ℓ to be cost and variance of $\hat{P}_\ell - \hat{P}_{\ell-1}$

then the total cost is $\sum_{\ell=0}^L N_\ell C_\ell$ and the variance is $\sum_{\ell=0}^L N_\ell^{-1} V_\ell$.

Using a Lagrange multiplier μ^2 to minimise the cost for a fixed variance

$$\frac{\partial}{\partial N_\ell} \sum_{k=0}^L (N_k C_k + \mu^2 N_k^{-1} V_k) = 0$$

gives

$$N_\ell = \mu \sqrt{V_\ell / C_\ell} \implies N_\ell C_\ell = \mu \sqrt{V_\ell C_\ell}$$

Multilevel Monte Carlo

Setting the total variance equal to ε^2 gives

$$\mu = \varepsilon^{-2} \left(\sum_{\ell=0}^L \sqrt{V_\ell C_\ell} \right)$$

and hence, the total cost is

$$\sum_{\ell=0}^L N_\ell C_\ell = \varepsilon^{-2} \left(\sum_{\ell=0}^L \sqrt{V_\ell C_\ell} \right)^2$$

in contrast to the standard cost which is approximately $\varepsilon^{-2} V_0 C_L$.

The MLMC cost savings are therefore approximately:

- V_L/V_0 , if $\sqrt{V_\ell C_\ell}$ increases with level
- C_0/C_L , if $\sqrt{V_\ell C_\ell}$ decreases with level

Multilevel Path Simulation

With SDEs, level ℓ corresponds to approximation using M^ℓ timesteps, giving approximate payoff \widehat{P}_ℓ at cost $C_\ell = O(h_\ell^{-1})$.

Simplest estimator for $\mathbb{E}[\widehat{P}_\ell - \widehat{P}_{\ell-1}]$ for $\ell > 0$ is

$$\widehat{Y}_\ell = N_\ell^{-1} \sum_{n=1}^{N_\ell} \left(\widehat{P}_\ell^{(n)} - \widehat{P}_{\ell-1}^{(n)} \right)$$

using same driving Brownian path for both levels.

Analysis gives $\text{MSE} = \sum_{\ell=0}^L N_\ell^{-1} V_\ell + \left(\mathbb{E}[\widehat{P}_L] - \mathbb{E}[P] \right)^2$

To make RMS error less than ε

- choose $N_\ell \propto \sqrt{V_\ell / C_\ell}$ so total variance is less than $\frac{1}{2} \varepsilon^2$
- choose L so that $\left(\mathbb{E}[\widehat{P}_L] - \mathbb{E}[P] \right)^2 < \frac{1}{2} \varepsilon^2$

Multilevel Path Simulation

For Lipschitz payoff functions $P \equiv f(S_T)$, we have

$$\begin{aligned} V_\ell &\equiv \mathbb{V} \left[\hat{P}_\ell - \hat{P}_{\ell-1} \right] \leq \mathbb{E} \left[(\hat{P}_\ell - \hat{P}_{\ell-1})^2 \right] \\ &\leq K^2 \mathbb{E} \left[(\hat{S}_{T,\ell} - \hat{S}_{T,\ell-1})^2 \right] \\ &= \begin{cases} O(h_\ell), & \text{Euler-Maruyama} \\ O(h_\ell^2), & \text{Milstein} \end{cases} \end{aligned}$$

and hence

$$V_\ell C_\ell = \begin{cases} O(1), & \text{Euler-Maruyama} \\ O(h_\ell), & \text{Milstein} \end{cases}$$

MLMC Theorem

(Slight generalisation of version in 2008 *Operations Research* paper)

If there exist independent estimators \widehat{Y}_ℓ based on N_ℓ Monte Carlo samples, each costing C_ℓ , and positive constants $\alpha, \beta, \gamma, c_1, c_2, c_3$ such that $\alpha \geq \frac{1}{2} \min(\beta, \gamma)$ and

i) $|\mathbb{E}[\widehat{P}_\ell - P]| \leq c_1 2^{-\alpha \ell}$

ii) $\mathbb{E}[\widehat{Y}_\ell] = \begin{cases} \mathbb{E}[\widehat{P}_0], & \ell = 0 \\ \mathbb{E}[\widehat{P}_\ell - \widehat{P}_{\ell-1}], & \ell > 0 \end{cases}$

iii) $\mathbb{V}[\widehat{Y}_\ell] \leq c_2 N_\ell^{-1} 2^{-\beta \ell}$

iv) $\mathbb{E}[C_\ell] \leq c_3 2^{\gamma \ell}$

MLMC Theorem

then there exists a positive constant c_4 such that for any $\varepsilon < 1$ there exist L and N_ℓ for which the multilevel estimator

$$\hat{Y} = \sum_{\ell=0}^L \hat{Y}_\ell,$$

has a mean-square-error with bound $\mathbb{E} \left[\left(\hat{Y} - \mathbb{E}[P] \right)^2 \right] < \varepsilon^2$

with an expected computational cost C with bound

$$C \leq \begin{cases} c_4 \varepsilon^{-2}, & \beta > \gamma, \\ c_4 \varepsilon^{-2} (\log \varepsilon)^2, & \beta = \gamma, \\ c_4 \varepsilon^{-2 - (\gamma - \beta)/\alpha}, & 0 < \beta < \gamma. \end{cases}$$

MLMC Theorem

Two observations of optimality:

- MC simulation needs $O(\varepsilon^{-2})$ samples to achieve RMS accuracy ε .
When $\beta > \gamma$, the cost is optimal — $O(1)$ cost per sample on average.
(Would need multilevel QMC to further reduce costs)
- When $\beta < \gamma$, another interesting case is when $\beta = 2\alpha$, which
corresponds to $\mathbb{E}[\hat{Y}_\ell]$ and $\sqrt{\mathbb{E}[\hat{Y}_\ell^2]}$ being of the same order as $\ell \rightarrow \infty$.
In this case, the total cost is $O(\varepsilon^{-\gamma/\alpha})$, which is the cost of a single
sample on the finest level — again optimal.

MLMC generalisation

The theorem is for scalar outputs P , but it can be generalised to multi-dimensional (or infinite-dimensional) outputs with

i) $\left\| \mathbb{E}[\hat{P}_\ell - P] \right\| \leq c_1 2^{-\alpha \ell}$

ii) $\mathbb{E}[\hat{Y}_\ell] = \begin{cases} \mathbb{E}[\hat{P}_0], & \ell = 0 \\ \mathbb{E}[\hat{P}_\ell - \hat{P}_{\ell-1}], & \ell > 0 \end{cases}$

iii) $\mathbb{V}[\hat{Y}_\ell] \equiv \mathbb{E} \left[\left\| \hat{Y}_\ell - \mathbb{E}[\hat{Y}_\ell] \right\|^2 \right] \leq c_2 N_\ell^{-1} 2^{-\beta \ell}$

Original multilevel research by Heinrich in 1999 did this for parametric integration, estimating $g(\lambda) \equiv \mathbb{E}[f(x, \lambda)]$ for a finite-dimensional r.v. x .

MLMC work on SDEs

- Milstein discretisation for path-dependent options – G (2008)
- numerical analysis – G, Higham, Mao (2009), Avikainen (2009), G, Debrabant, Rößler (2012)
- financial sensitivities (“Greeks”) – Burgos (2011)
- jump-diffusion models – Xia (2011)
- Lévy processes – Dereich (2010), Marxen (2010), Dereich & Heidenreich (2011), Xia (2013), Kyprianou (2014)
- American options – Belomestny & Schoenmakers (2011)
- Milstein in higher dimensions without Lévy areas – G, Szpruch (2014)
- adaptive timesteps – Hoel, von Schwerin, Szepessy, Tempone (2012), G, Lester, Whittle (2014)

- quite natural application, with better cost savings than SDEs due to higher dimensionality
- range of applications
 - ▶ Graubner & Ritter (Darmstadt) – parabolic
 - ▶ G, Reisinger (Oxford) – parabolic
 - ▶ Cliffe, G, Scheichl, Teckentrup (Bath/Nottingham) – elliptic
 - ▶ Barth, Jenny, Lang, Meyer, Mishra, Müller, Schwab, Sukys, Zollinger (ETH Zürich) – elliptic, parabolic, hyperbolic
 - ▶ Harbrecht, Peters (Basel) – elliptic
 - ▶ Efendiev (Texas A&M) – numerical homogenization
 - ▶ Vidal-Codina, G, Peraire (MIT) – reduced basis approximation

Engineering Uncertainty Quantification

Simplest possible example:

- 3D elliptic PDE, with uncertain boundary data
- grid spacing proportional to $2^{-\ell}$ on level ℓ
- cost is $O(2^{+3\ell})$, if using an efficient multigrid solver
- 2nd order accuracy means that

$$\begin{aligned}\widehat{P}_\ell(\omega) - P(\omega) &\approx c(\omega) 2^{-2\ell} \\ \implies \widehat{P}_{\ell-1}(\omega) - \widehat{P}_\ell(\omega) &\approx 3c(\omega) 2^{-2\ell}\end{aligned}$$

- hence, $\alpha=2$, $\beta=4$, $\gamma=3$
- cost is $O(\varepsilon^{-2})$ to obtain ε RMS accuracy
- this compares to $O(\varepsilon^{-3/2})$ cost for one sample on finest level, so $O(\varepsilon^{-7/2})$ for standard Monte Carlo

PDEs with Uncertainty

I worked with Rob Scheichl (Bath) and Andrew Cliffe (Nottingham) on multilevel Monte Carlo for the modelling of oil reservoirs and groundwater contamination in nuclear waste repositories.

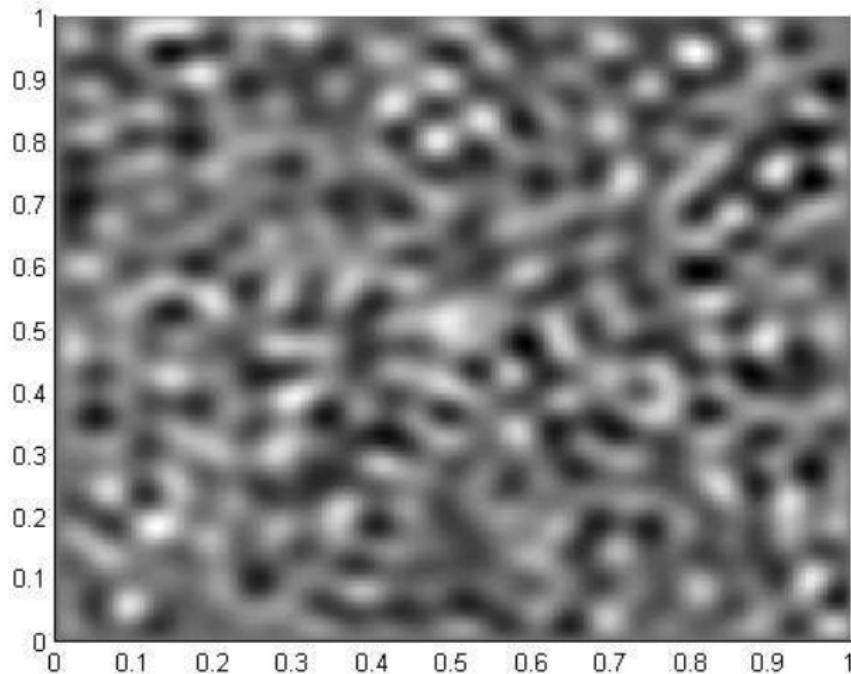
Here we have an elliptic SPDE coming from Darcy's law:

$$\nabla \cdot (\kappa(x) \nabla p) = 0$$

where the permeability $\kappa(x)$ is uncertain, and $\log \kappa(x)$ is often modelled as being Normally distributed with a spatial covariance such as

$$\text{cov}(\log \kappa(x_1), \log \kappa(x_2)) = \sigma^2 \exp(-\|x_1 - x_2\|/\lambda)$$

Elliptic SPDE



A typical realisation of κ for $\lambda = 0.01$, $\sigma = 1$.

Elliptic SPDE

Samples of $\log k$ are provided by a Karhunen-Loève expansion:

$$\log k(\mathbf{x}, \omega) = \sum_{n=0}^{\infty} \sqrt{\theta_n} \xi_n(\omega) f_n(\mathbf{x}),$$

where θ_n , f_n are eigenvalues / eigenfunctions of the correlation function:

$$\int R(\mathbf{x}, \mathbf{y}) f_n(\mathbf{y}) d\mathbf{y} = \theta_n f_n(\mathbf{x})$$

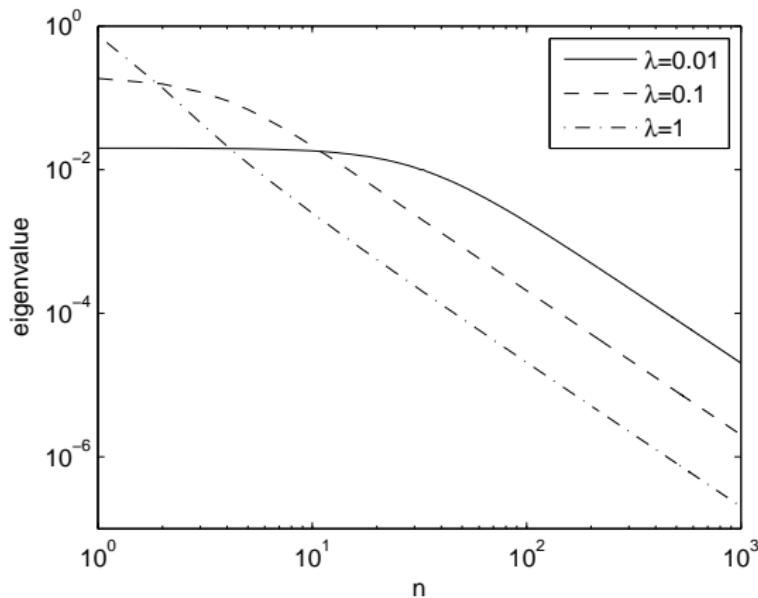
and $\xi_n(\omega)$ are standard Normal random variables.

Numerical experiments truncate the expansion.

(Latest 2D/3D work uses a more efficient FFT construction based on a circulant embedding.)

Elliptic SPDE

Decay of 1D eigenvalues



When $\lambda = 1$, can use a low-dimensional polynomial chaos approach, but it's impractical for smaller λ .

Discretisation:

- cell-centred finite volume discretisation on a uniform grid – for rough coefficients we need to make grid spacing very small on finest grid
- each level of refinement has twice as many grid points in each direction
- early numerical experiments used a direct solver for simplicity, but later work in 3D uses an efficient AMG multigrid solver with a cost roughly proportional to the total number of grid points
- later work also considers other finite element discretisations – doesn't make any substantial difference to MLMC treatment

2D Results

Boundary conditions for unit square $[0, 1]^2$:

- fixed pressure: $p(0, x_2) = 1, p(1, x_2) = 0$
- Neumann b.c.: $\partial p / \partial x_2(x_1, 0) = \partial p / \partial x_2(x_1, 1) = 0$

Output quantity – mass flux: $-\int k \frac{\partial p}{\partial x_1} dx_2$

Correlation length: $\lambda = 0.2$

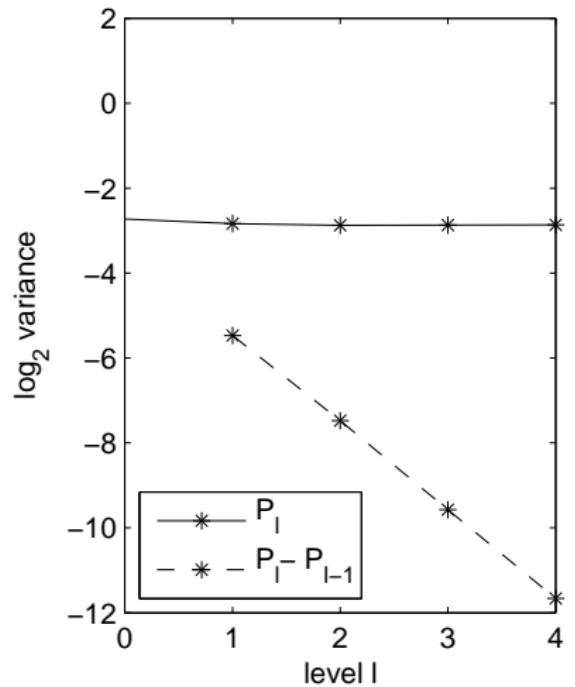
Coarsest grid: $h = 1/8$ (comparable to λ)

Finest grid: $h = 1/128$

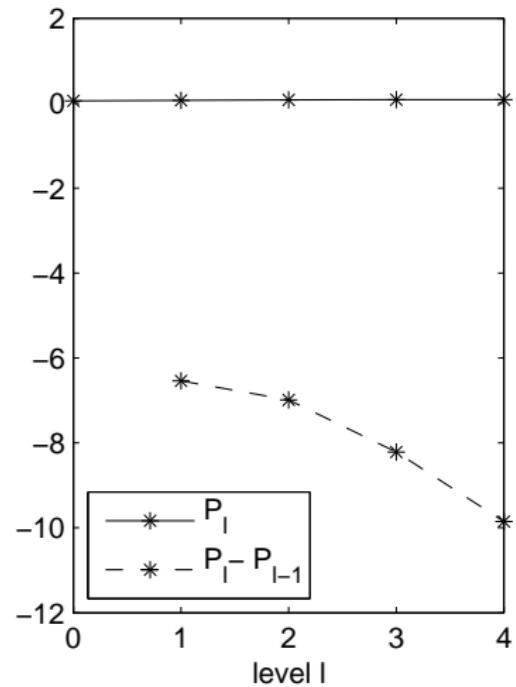
Karhunen-Loève truncation: $m_{KL} = 4000$

Cost taken to be proportional to number of nodes

2D Results

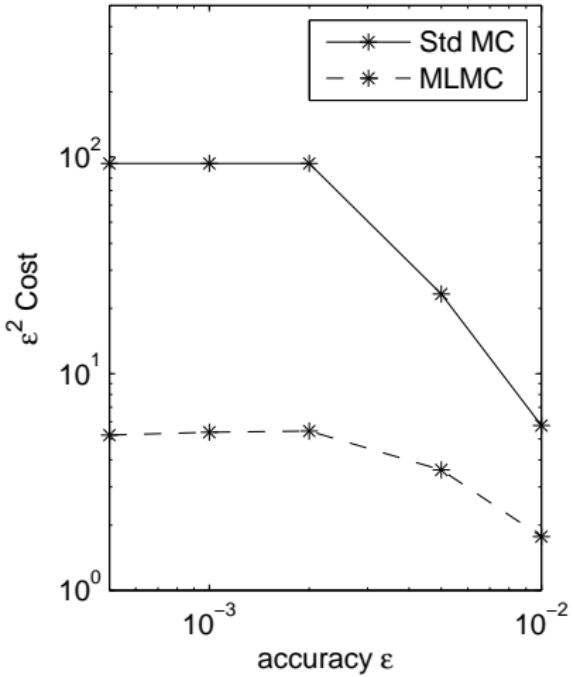
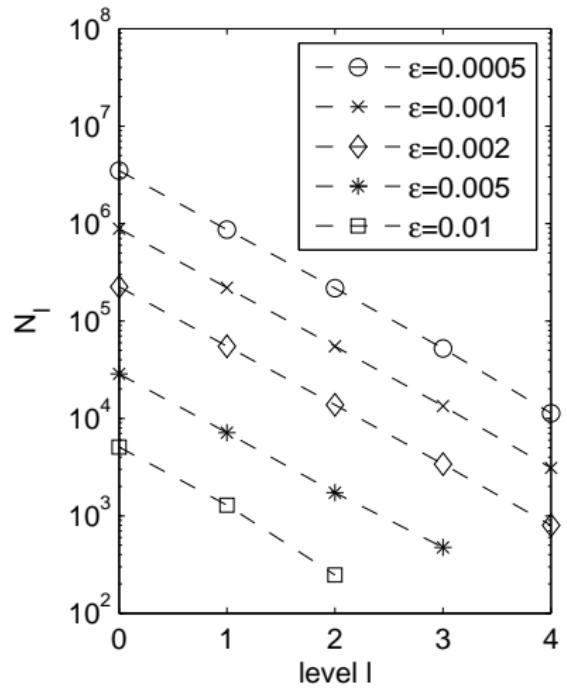


$$\mathbb{V}[\hat{P}_\ell - \hat{P}_{\ell-1}] \sim h_\ell^2$$



$$\mathbb{E}[\hat{P}_\ell - \hat{P}_{\ell-1}] \sim h_\ell^2$$

2D Results



Complexity analysis

Relating things back to the MLMC theorem:

$$\mathbb{E}[\hat{P}_\ell - P] \sim 2^{-2\ell} \implies \alpha = 2$$

$$V_\ell \sim 2^{-2\ell} \implies \beta = 2$$

$$C_\ell \sim 2^{d\ell} \implies \gamma = d \quad (\text{dimension of PDE})$$

To achieve r.m.s. accuracy ε requires finest level grid spacing $h \sim \varepsilon^{1/2}$ and hence we get the following complexity:

dim	MC	MLMC
1	$\varepsilon^{-2.5}$	ε^{-2}
2	ε^{-3}	$\varepsilon^{-2}(\log \varepsilon)^2$
3	$\varepsilon^{-3.5}$	$\varepsilon^{-2.5}$

Non-geometric multilevel

Almost all applications of multilevel in the literature so far use a geometric sequence of levels, refining the timestep (or the spatial discretisation for PDEs) by a constant factor when going from level ℓ to level $\ell + 1$.

Coming from a multigrid background, this is very natural, but it is **NOT** a requirement of the multilevel Monte Carlo approach.

All MLMC needs is a sequence of levels with

- increasing accuracy
- increasing cost
- increasingly small difference between outputs on successive levels

Reduced Basis PDE approximation

Vidal-Codina, Nguyen, G, Peraire (2014) take a fine FE discretisation:

$$A(\omega) u = f(\omega)$$

and use a reduced basis approximation

$$u \approx \sum_{k=1}^K v_k u_k$$

to obtain a low-dimensional reduced system

$$A_r(\omega) v = f_r(\omega)$$

- larger $K \implies$ greater accuracy at greater cost
- in multilevel treatment, K_ℓ varies with level
- brute force optimisation determines the optimal number of levels, and reduced basis size on each level

Other MLMC applications

- parametric integration, integral equations (Heinrich)
- multilevel QMC (Dick, G, Kuo, Scheichl, Schwab, Sloan)
- stochastic chemical reactions (Anderson & Higham, Tempone)
- mixed precision computation on FPGAs (Korn, Ritter, Wehn)
- MLMC for MCMC (Scheichl, Schwab, Stuart, Teckentrup)
- Coulomb collisions in plasma (Caflisch)
- nested simulation (Haji-Ali & Tempone, Hambly & Reisinger)
- invariant distribution of contractive Markov process (Glynn & Rhee)
- invariant distribution of contractive SDEs (G, Lester & Whittle)

Three MLMC extensions

- unbiased estimation – Rhee & Glynn (2015)
 - ▶ randomly selects the level for each sample
 - ▶ no bias, and finite expected cost and variance if $\beta > \gamma$
- Richardson-Romberg extrapolation – Lemaire & Pagès (2013)
 - ▶ reduces the weak error, and hence the number of levels required
 - ▶ particularly helpful when $\beta < \gamma$
- Multi-Index Monte Carlo – Haji-Ali, Nobile, Tempone (2015)
 - ▶ important extension to MLMC approach, combining MLMC with sparse grid methods

Multi-Index Monte Carlo

Standard “1D” MLMC truncates the telescoping sum

$$\mathbb{E}[P] = \sum_{\ell=0}^{\infty} \mathbb{E}[\Delta \hat{P}_{\ell}]$$

where $\Delta \hat{P}_{\ell} \equiv \hat{P}_{\ell} - \hat{P}_{\ell-1}$, with $\hat{P}_{-1} \equiv 0$.

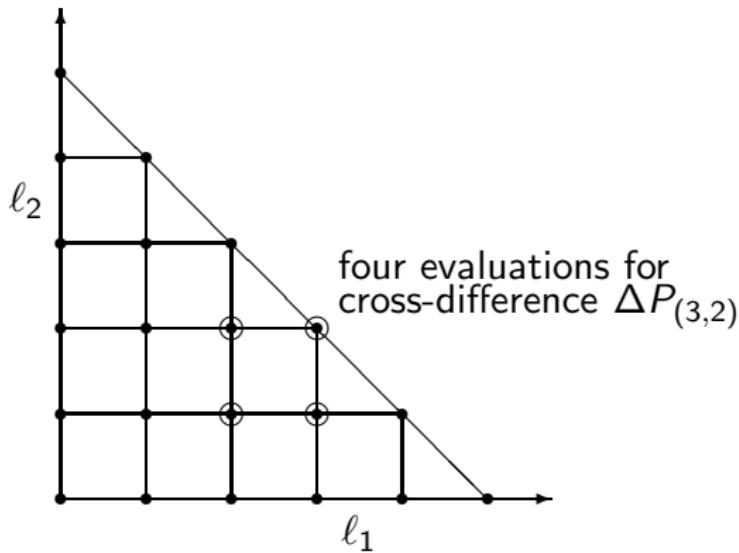
In “2D”, MIMC truncates the telescoping sum

$$\mathbb{E}[P] = \sum_{\ell_1=0}^{\infty} \sum_{\ell_2=0}^{\infty} \mathbb{E}[\Delta \hat{P}_{\ell_1, \ell_2}]$$

where $\Delta \hat{P}_{\ell_1, \ell_2} \equiv (\hat{P}_{\ell_1, \ell_2} - \hat{P}_{\ell_1-1, \ell_2}) - (\hat{P}_{\ell_1, \ell_2-1} - \hat{P}_{\ell_1-1, \ell_2-1})$

Different aspects of the discretisation vary in each “dimension” – for a 2D PDE, could use grid spacing $2^{-\ell_1}$ in direction 1, $2^{-\ell_2}$ in direction 2

Multi-Index Monte Carlo



MIMC truncates the summation in a way which minimises the cost to achieve a target MSE – quite similar to sparse grids.

Can achieve $O(\varepsilon^{-2})$ complexity for a wider range of SPDE and other applications than plain MLMC.

Conclusions

- multilevel idea is very simple; key question is how to apply it in new situations, and perform the numerical analysis
- discontinuous output functions can cause problems, but there is a lot of experience now in coping with this
- there are also “tricks” which can be used in situations with poor strong convergence
- being used for an increasingly wide range of applications; biggest computational savings when coarsest (reasonable) approximation is much cheaper than finest
- currently, getting at least $100\times$ savings for SPDEs and stochastic chemical reaction simulations

References

Webpages for my research papers and talks:

people.maths.ox.ac.uk/giles/mcmc.html

people.maths.ox.ac.uk/giles/slides.html

Webpage for new 70-page *Acta Numerica* review and MATLAB test codes:

people.maths.ox.ac.uk/giles/acta/

– contains references to almost all MLMC research

MLMC Community

Webpage: people.maths.ox.ac.uk/gilesmlmc_community.html

Abo Academi (Avikainen) – numerical analysis

Basel (Harbrecht) – elliptic SPDEs, sparse grids

Bath (Kyprianou, Scheichl, Shardlow, Yates) – elliptic SPDEs, MCMC, Lévy-driven SDEs, stochastic chemical modelling

Chalmers (Lang) – SPDEs

Duisburg (Belomestny) – Bermudan and American options

Edinburgh (Davie, Szpruch) – SDEs, numerical analysis

EPFL (Abdulle) – stiff SDEs and SPDEs

ETH Zürich (Jenny, Jentzen, Schwab) – SPDEs, multilevel QMC

Frankfurt (Gerstner, Kloeden) – numerical analysis, fractional Brownian motion

Fraunhofer ITWM (Iliev) – SPDEs in engineering

Hong Kong (Chen) – Brownian meanders, nested simulation in finance

IIT Chicago (Hickernell) – SDEs, infinite-dimensional integration, complexity analysis

Kaiserslautern (Heinrich, Korn, Ritter) – finance, SDEs, parametric integration, complexity analysis

KAUST (Tempone) – adaptive time-stepping, stochastic chemical modelling

Kiel (Gnewuch) – randomized multilevel QMC

LPMA (Frikha, Lemaire, Pagès) – numerical analysis, multilevel extrapolation, finance applications

Mannheim (Neuenkirch) – numerical analysis, fractional Brownian motion

MIT (Peraire) – uncertainty quantification, SPDEs

Munich (Hutzenthaler) – numerical analysis

Oxford (Baker, Giles, Hambly, Reisinger, Süli) – SDEs, SPDEs, nested simulation, numerical analysis, finance applications, stochastic chemical reactions, long-chain molecules

Passau (Müller-Gronbach) – infinite-dimensional integration, complexity analysis

Stanford (Glynn) – numerical analysis, randomized multilevel

Strathclyde (Higham, Mao) – numerical analysis, exit times, stochastic chemical modelling

Stuttgart (Barth) – SPDEs

Texas A&M (Efendiev) – SPDEs in engineering

UCLA (Caflisch) – Coulomb collisions in physics

UNSW (Dick, Kuo, Sloan) – multilevel QMC

Warwick (Stuart, Teckentrup) – MCMC for SPDEs

WIAS (Friz, Schoenmakers) – rough paths, fractional Brownian motion, Bermudan options

Wisconsin (Anderson) – numerical analysis, stochastic chemical modelling