Multilevel Monte Carlo methods

Mike Giles

Mathematical Institute, University of Oxford

Seminar at Reading University

March 3, 2015

Outline

- introduction to key ideas
- some example applications
- challenges and generalisations
- my current research
  - dynamics of long-chain molecules
  - high-dimensional PDEs

Objectives

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

- the simplicity of the idea
- its flexibility
- that it's not prescriptive, more an approach
- scope for improved performance through being creative
- lots of people working on a variety of applications

I will focus on ideas rather than lots of numerical results.

Monte Carlo method

Given a function f of a random input ω, to estimate the value of $\mathbb{E}[f]$ we can use the Monte Carlo estimate

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

based on N independent samples $\omega^{(n)}$.

By the Central Limit Theorem, as $N \to \infty$, the error in this estimate becomes Normally distributed, with variance $N^{-1}\mathbb{V}[f]$.

The error lies within 3 s.d. with probability 99.7%, giving us a confidence interval.

Control variate

Classic approach to variance reduction: approximate $\mathbb{E}[f]$ using

$$N^{-1}\sum_{n=1}^{N}\left\{f(\omega^{(n)}) - \lambda\left(g(\omega^{(n)}) - \mathbb{E}[g]\right)\right\}$$

where

- control variate g has known expectation $\mathbb{E}[g]$
- g is well correlated with f, and optimal value for λ can be estimated by a few samples

For the optimal value of λ, the variance is reduced by factor $(1-\rho^2)$, where ρ is the correlation between f and g.

Two-level Monte Carlo

If we want to estimate $\mathbb{E}[f_1]$ but it is much cheaper to simulate $f_0 \approx f_1$, then since

$$\mathbb{E}[f_1] = \mathbb{E}[f_0] + \mathbb{E}[f_1 - f_0]$$

we can use the estimator

$$N_0^{-1} \sum_{n=1}^{N_0} f_0^{(0,n)} \;+\; N_1^{-1} \sum_{n=1}^{N_1} \left(f_1^{(1,n)} - f_0^{(1,n)} \right)$$

Two differences from standard control variate method:

- $\mathbb{E}[f_0]$ is not known, so has to be estimated
- $\lambda = 1$

Benefit: if $f_1 - f_0$ is small, won't need many samples to accurately estimate $\mathbb{E}[f_1 - f_0]$, so cost will be reduced greatly.

Multilevel Monte Carlo

Natural generalisation: given a sequence $f_0, f_1, \ldots, f_L$

$$\mathbb{E}[f_L] = \mathbb{E}[f_0] + \sum_{\ell=1}^{L} \mathbb{E}[f_\ell - f_{\ell-1}]$$

we can use the estimator

$$N_0^{-1} \sum_{n=1}^{N_0} f_0^{(0,n)} \; + \; \sum_{\ell=1}^{L} \left\{ N_\ell^{-1} \sum_{n=1}^{N_\ell} \left(f_\ell^{(\ell,n)} - f_{\ell-1}^{(\ell,n)} \right) \right\}$$

with independent estimation for each level

Multilevel Monte Carlo

If we define

- C_0, V_0 to be cost and variance of f_0
- C_ℓ, V_ℓ to be cost and variance of $f_\ell - f_{\ell-1}$

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

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

gives

$$N_\ell = \mu \sqrt{V_\ell / C_\ell} \quad \Longrightarrow \quad N_\ell\, C_\ell = \mu \sqrt{V_\ell\, C_\ell}$$

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:

- 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 Monte Carlo

This analysis treated the N_ℓ as real variables. Rounding them up to the nearest integer gives the following result:

Theorem: With V_ℓ and C_ℓ as defined previously, an estimate $\widehat{Y}$ with RMS accuracy ε,

$$\text{MSE} \equiv \mathbb{E}\left[(\widehat{Y} - \mathbb{E}[f_L])^2\right] \leq \varepsilon^2$$

can be obtained at computational cost

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

Note: this assumes perfect knowledge of V_ℓ and C_ℓ. In practice V_ℓ at least usually needs to be estimated.

Multilevel Path Simulation

Motivated by computational finance applications, in 2006 I introduced MLMC for SDEs (stochastic differential equations).

$$\mathrm{d}S_t = a(S_t, t)\,\mathrm{d}t + b(S_t, t)\,\mathrm{d}W_t$$

Level ℓ corresponds to approximation using 2^ℓ timesteps, giving approximate payoff $\widehat{P}_\ell$.

Choice of finest level L depends on weak error (bias).

Multilevel decomposition gives

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

Multilevel Monte Carlo

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

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

To make RMS error less than ε

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

MLMC Theorem

(Slight generalisation of original version)

If there exist independent estimators $\widehat{Y}_\ell$ 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) $\left|\mathbb{E}[\widehat{P}_\ell - P]\right| \leq c_1\, 2^{-\alpha\,\ell}$

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

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

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

with an expected computational cost C with bound

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

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 $\mathbb{E}[\widehat{Y}_\ell]$ and $\sqrt{\mathbb{E}[\widehat{Y}_\ell^2]}$ being of the same order as $\ell \to \infty$.

  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) $\left\| \mathbb{E}[\widehat{P}_\ell - P] \right\| \le c_1 \, 2^{-\alpha \ell}$

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

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 Challenges

- not always obvious how to couple coarse and fine levels
  i.e. what does $\widehat{P}_\ell(\omega^{(n)}) - \widehat{P}_{\ell-1}(\omega^{(n)})$ mean?

- some creativity required to handle discontinuous functionals, where a small difference between the underlying coarse and fine simulations can produce an $O(1)$ difference in the output

- numerical analysis to determine the decay rate of V_ℓ can be tough

Brownian Diffusion SDEs

Brownian increments for coarse path obtained by summing increments for fine path – very simple and natural

I like the Milstein discretisation which gives first order strong convergence

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

so for payoffs which are Lipschitz functions of the final state we get

$$\widehat{P}_\ell - \widehat{P}_{\ell-1} = O(h_\ell)$$

and hence $V_\ell = O(h_\ell^2)$.

However, not so easy for lookback, digital and barrier options. Also, in multiple dimensions sometimes requires Lévy areas, but can be avoided by an antithetic treatment, (G & Szpruch, 2013).

Financial application

- basket of 5 underlying assets, modelled by Geometric Brownian Motion

$$\mathrm{d}S_i = r\, S_i\, \mathrm{d}t + \sigma_i\, S_i\, \mathrm{d}W_i$$

  with correlation between 5 driving Brownian motions

- Milstein numerical approximation

- standard call option is piecewise linear function of average at final time T

- digital call option is discontinuous function of average

Financial application

Standard call option:

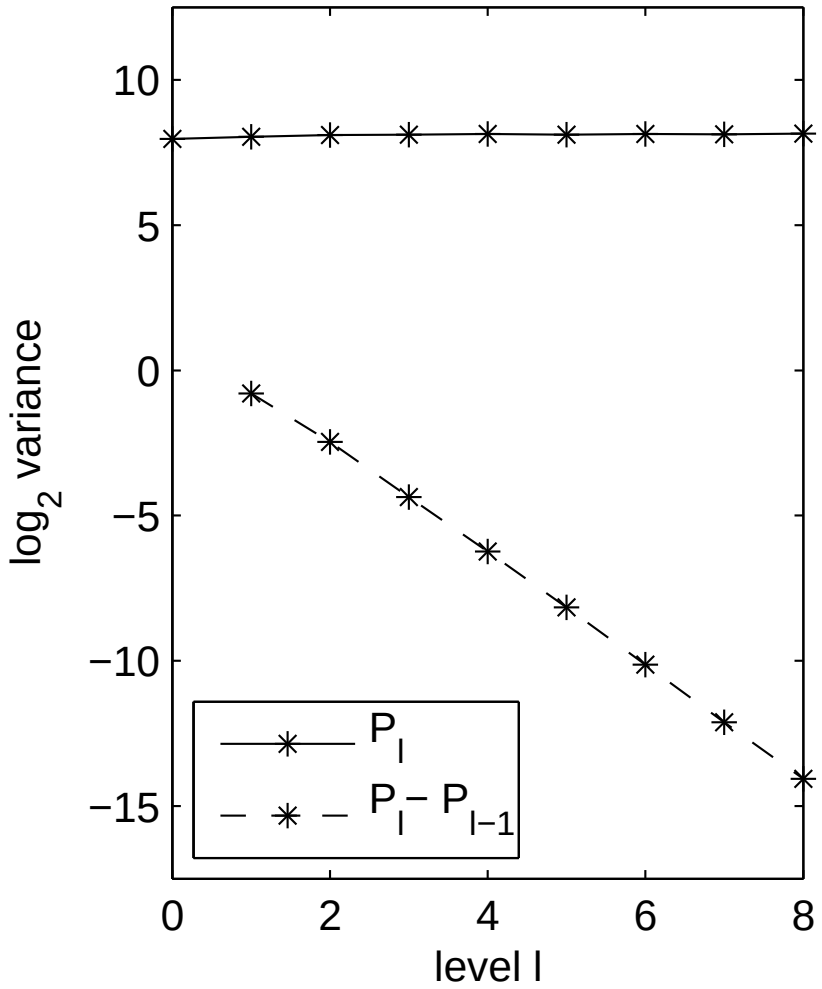

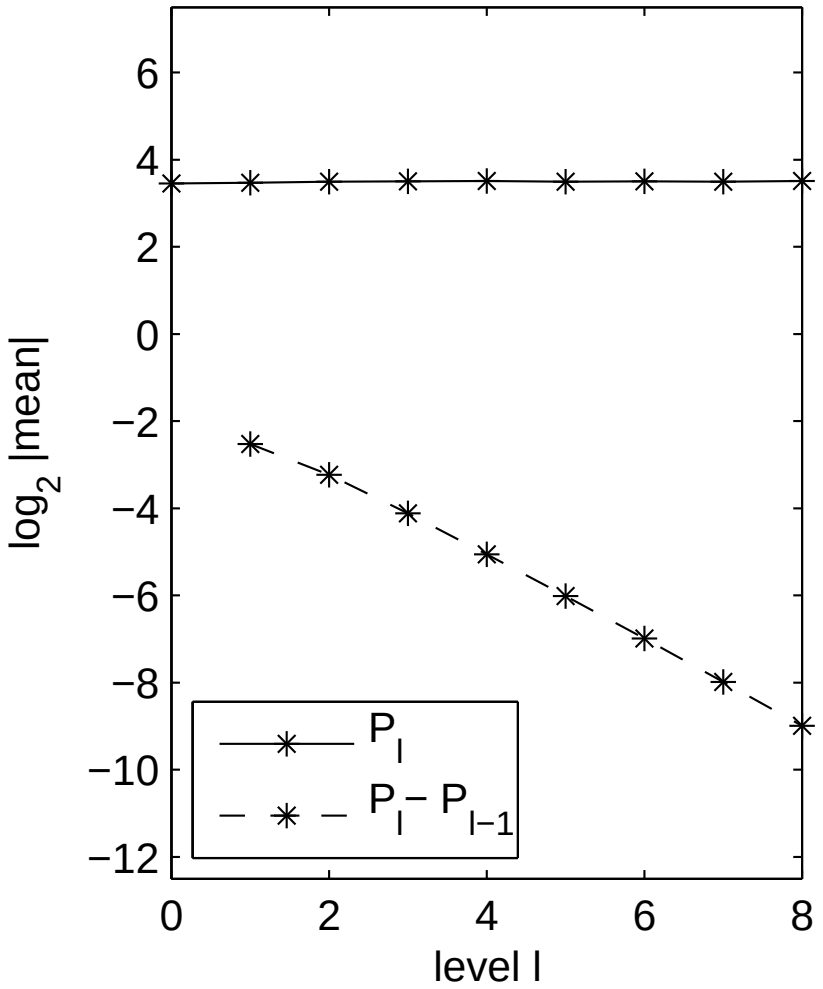

Financial application

Standard call option:

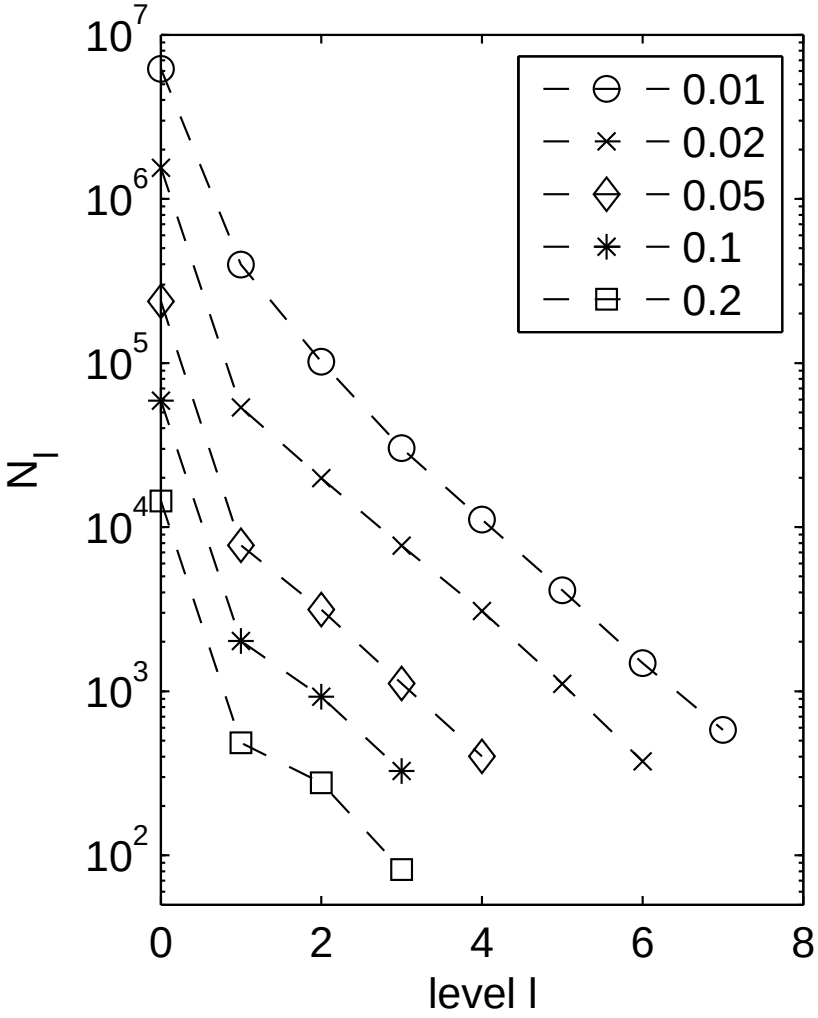

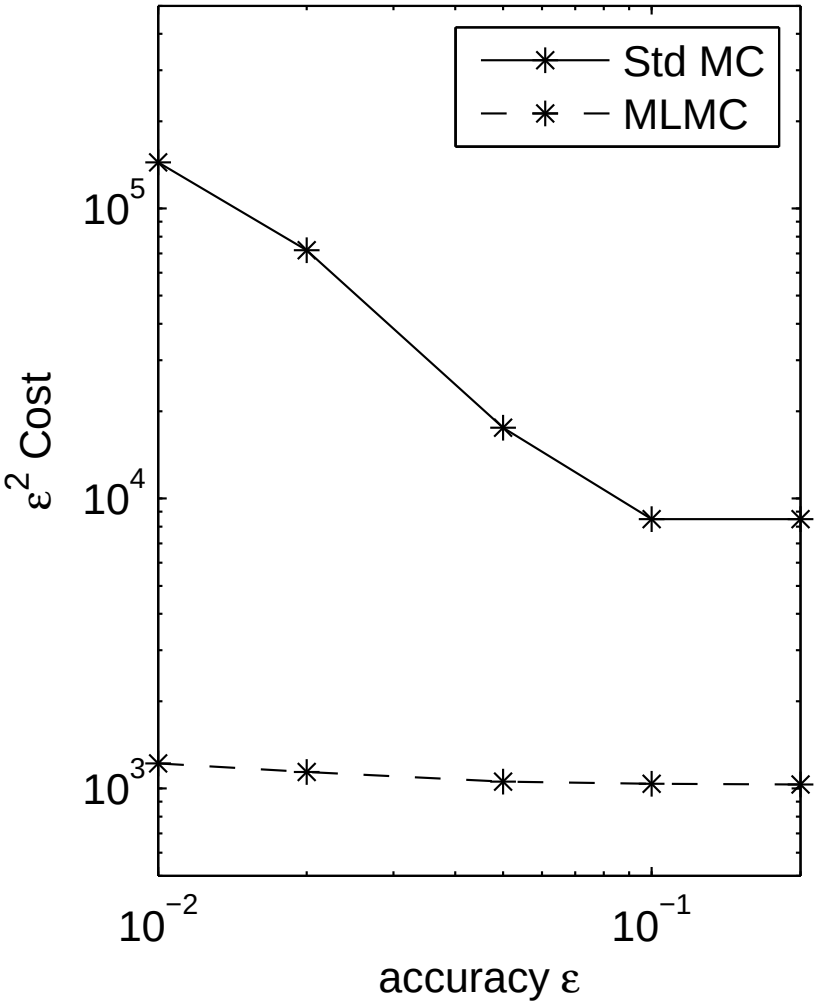

Digital options

In a digital option, the payoff is a discontinuous function of the final state.

Using the Milstein approximation, first order strong convergence means that $O(h_\ell)$ of the simulations have coarse and fine paths on opposite sides of a discontinuity.

Hence,

$$\widehat{P}_\ell - \widehat{P}_{\ell-1} = \begin{cases} O(1), & \text{with probability } O(h_\ell) \\ O(h_\ell), & \text{with probability } O(1) \end{cases}$$

so

$$\mathbb{E}[\widehat{P}_\ell - \widehat{P}_{\ell-1}] = O(h_\ell), \quad \mathbb{E}[(\widehat{P}_\ell - \widehat{P}_{\ell-1})^2] = O(h_\ell),$$

and hence $V_\ell = O(h_\ell)$, not $O(h_\ell^2)$

Digital options

Three fixes:

- Conditional expectation: using the Euler discretisation instead of Milstein for the final timestep, conditional on all but the final Brownian increment, the final state has a Gaussian distribution, with a known analytic conditional expectation in simple cases
- Splitting: split each path simulation into M paths by trying M different values for the Brownian increment for the last fine path timestep
- Change of measure: when the expectation is not known, can use a change of measure so the coarse path takes the same final state as the fine path — difference in the "payoff" now comes from the Radon-Nikodym derivative

These all effectively smooth the payoff – end up with $V_\ell = O(h_\ell^{3/2})$.

Financial application

Digital call option:

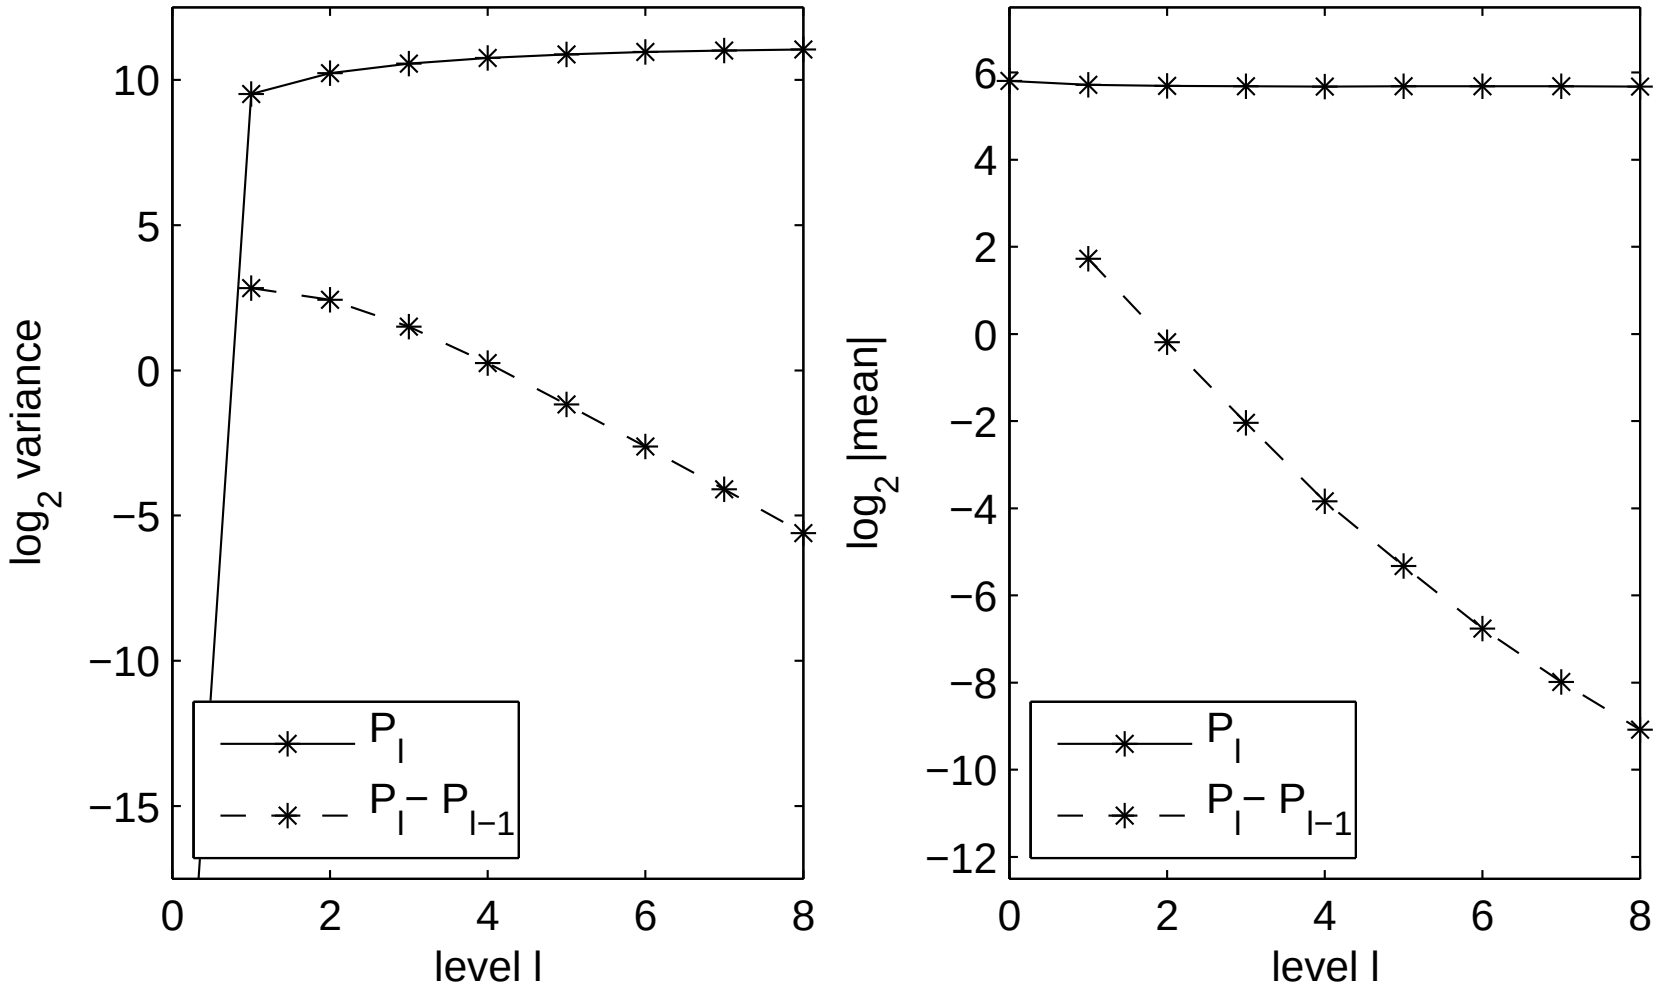

Financial application

Digital call option:

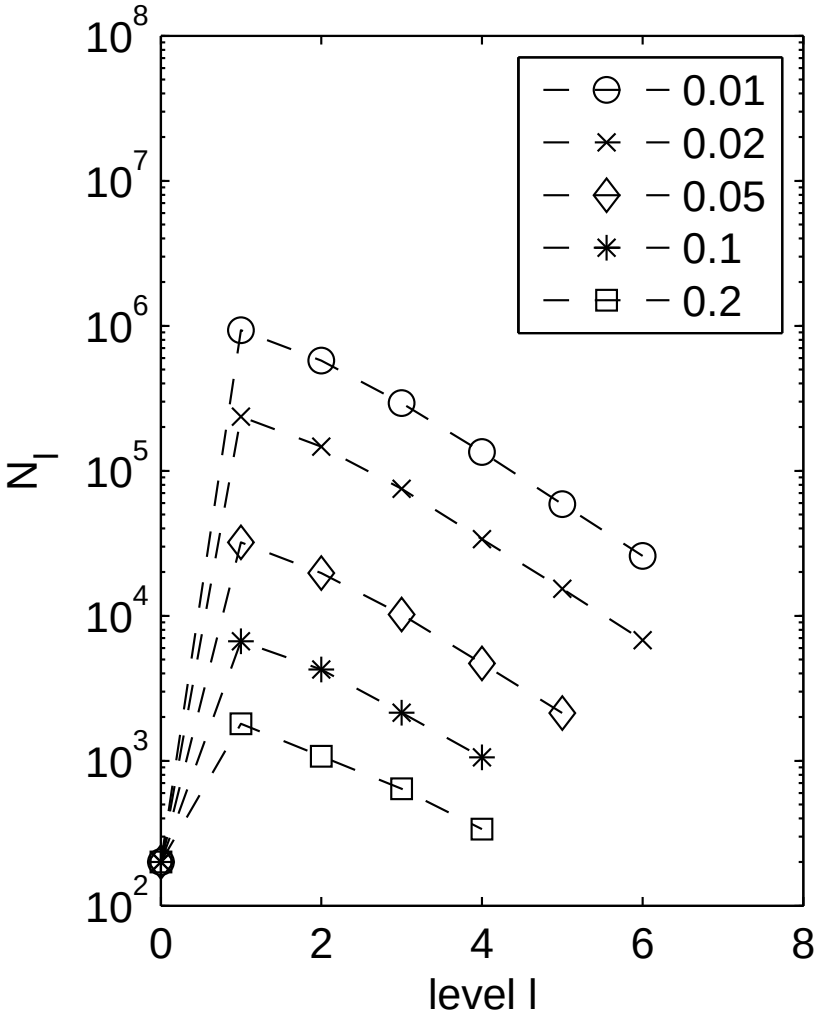

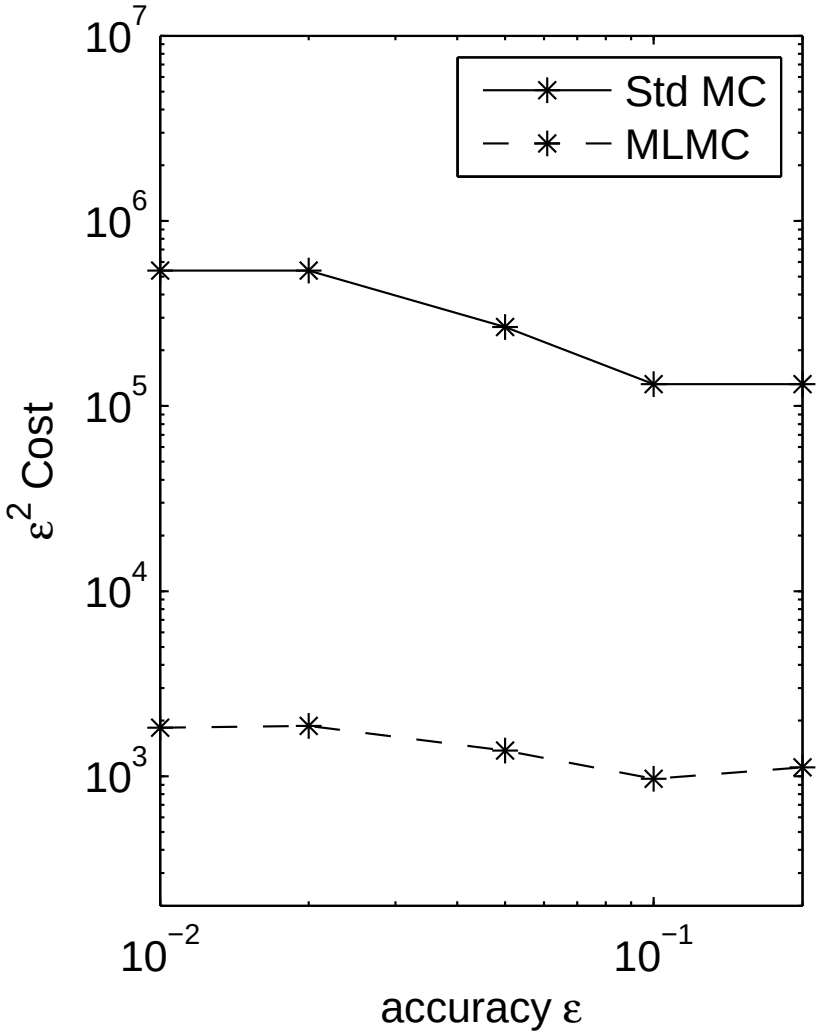

SPDEs

- quite natural application, with better cost savings than SDEs due to higher dimensionality

- range of applications
  - Graubner & Ritter (Darmstadt $\rightarrow$ Kaiserslautern) – 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
  - G, Hou, Zhang (Caltech) – numerical homogenization

Engineering Uncertainty Quantification

- consider 3D elliptic PDE, with uncertain boundary data
- use 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

$$\widehat{P}_\ell(\omega) - \widehat{P}(\omega) \approx c(\omega)\, 2^{-2\ell}$$

$$\Longrightarrow \quad \widehat{P}_{\ell-1}(\omega) - \widehat{P}_\ell(\omega) \approx 3\, c(\omega)\, 2^{-2\ell}$$

- hence, $\alpha = 2,\ \beta = 4,\ \gamma = 3$
- cost is $O(\varepsilon^{-2})$ to obtain ε RMS accuracy

Other MLMC Applications

- parametric integration, integral equations (Heinrich)
- multilevel QMC (Dick, G, Kuo, Scheichl, Schwab, Sloan)
- Lévy-driven SDEs (Dereich, Heidenreich)
- 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)

Recent MLMC Extensions

- unbiased estimation through randomisation of levels (Glynn, Rhee)
  - good for $\beta > \gamma$
- Richardson/Romberg extrapolation (Lemaire, Pagès)
  - good for $\beta < \gamma$
- Multi-Index Monte Carlo (Haji-Ali, Nobile, Tempone)
  - combines MLMC with sparse grid methods
  - potentially very important for SPDE applications

New project 1: FENE molecules in a fluid (Süli, Ye)

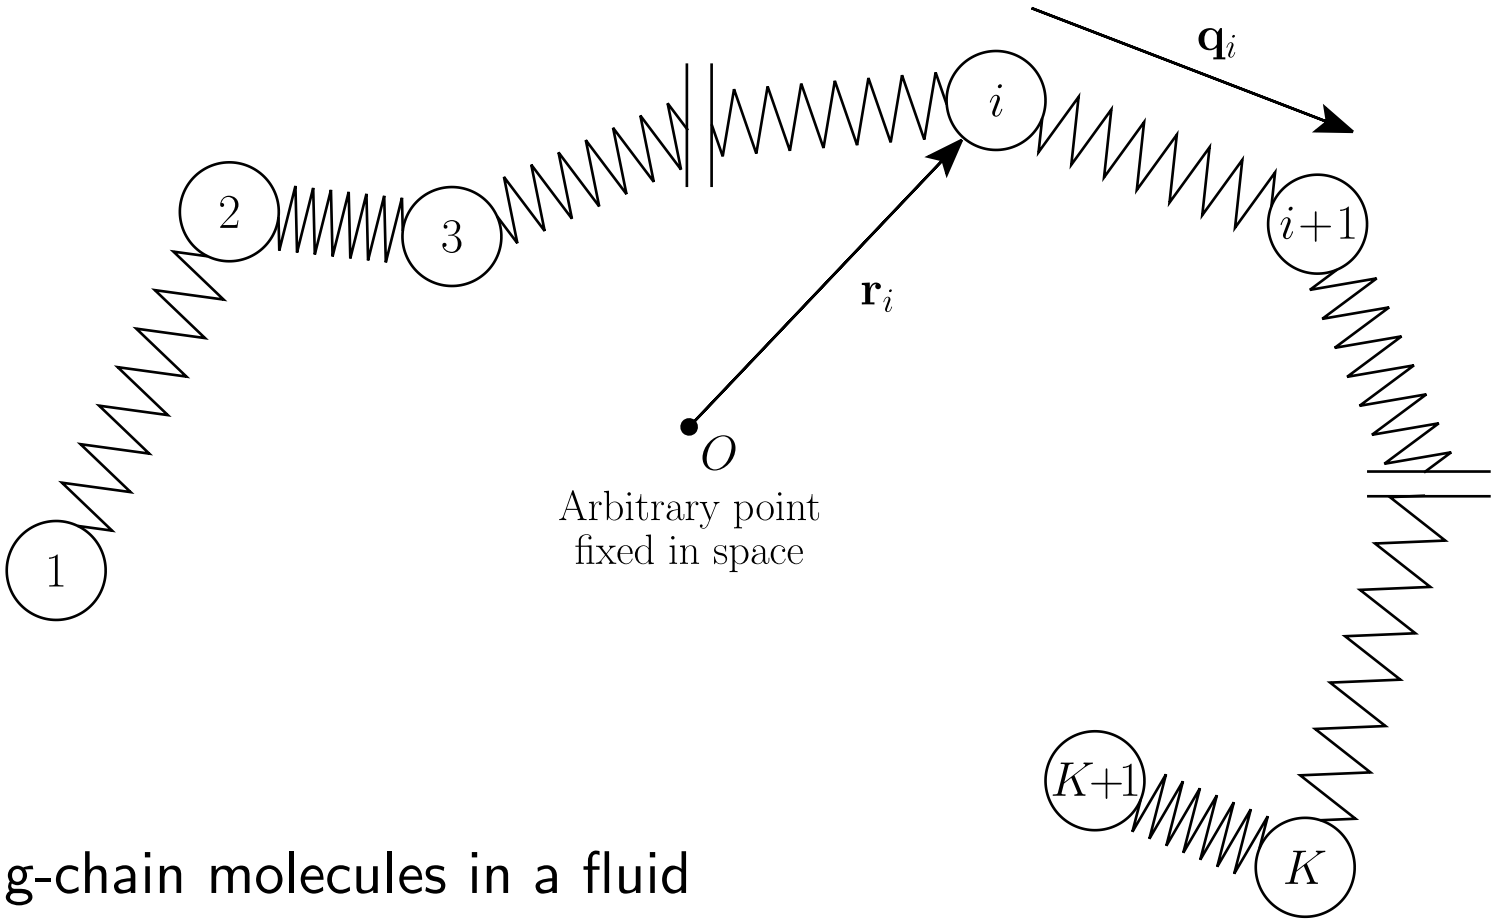

- long-chain molecules in a fluid
- modelled as ball-and-spring systems, subject to
  - force due to Finitely Extensible Nonlinear Elastic bond energy
  - force due to local rate-of-strain tensor $\partial v/\partial x$
  - random forcing due to fluid fluctuations

Modelling

The coupled system of SDEs can be written collectively as

$$\mathrm{d}q = (K\,q - D\,\nabla V)\,\mathrm{d}t + \sqrt{2}\;L\;\mathrm{d}W$$

where

- $V(q) \equiv \sum_i U_i(\|q_i\|^2/2)$ is the total bond energy, with $U_i(\|q_i\|^2/2) \to \infty$ as $\|q_i\|^2 \to 1$
- K is block diagonal, due to the fluid strain-rate tensor $\partial v/\partial x$
- L and D are of the form

$$L = \begin{pmatrix} -I & I & & \\ & -I & I & \\ & & -I & I \end{pmatrix}, \quad D = \begin{pmatrix} 2I & -I & \\ -I & 2I & -I \\ & -I & 2I \end{pmatrix} = L\,L^T.$$

Numerical approximation

The SDE is approximated as

$$q_{n+1} = q_n + (K\, q_n - D\, \nabla V(q_n))\, h_n + \sqrt{2}\; L\; \Delta W_n$$

using an adaptive timestep h_n.

No bond length should exceed 1 – try to ensure this through the restrictions:

$$\begin{aligned} h_n\, U_i'(\|q_{i,n}\|^2/2)\; \|q_{i,n}\| &\leq 1 - \|q_{i,n}\| \\ 5\sqrt{2\, h_n} &\leq 1 - \|q_{i,n}\| \end{aligned}$$

where $q_{i,n}$ is the i^{th} bond vector at timestep n (and then use clamping if this fails).

This sets an upper bound on the timestep – smaller values need to be chosen for accuracy.

Multilevel Monte Carlo simulation

First challenge: how does MLMC work with adaptive time-stepping?

Actually, surprisingly easy — on level ℓ use

$$h_n = 2^{-\ell}\ \frac{\min_i(1-\|q_{i,n}\|)^2}{\max(2\beta, 50)}$$

Coarse and fine paths each compute their own adaptive timesteps independently – this ensures the telescoping sum works correctly

But what is involved in coarse and fine paths using same driving Brownian motion?

Multilevel Monte Carlo simulation

As time proceeds, Brownian increments are generated as needed at discretisation times which are a union of coarse and fine path times:

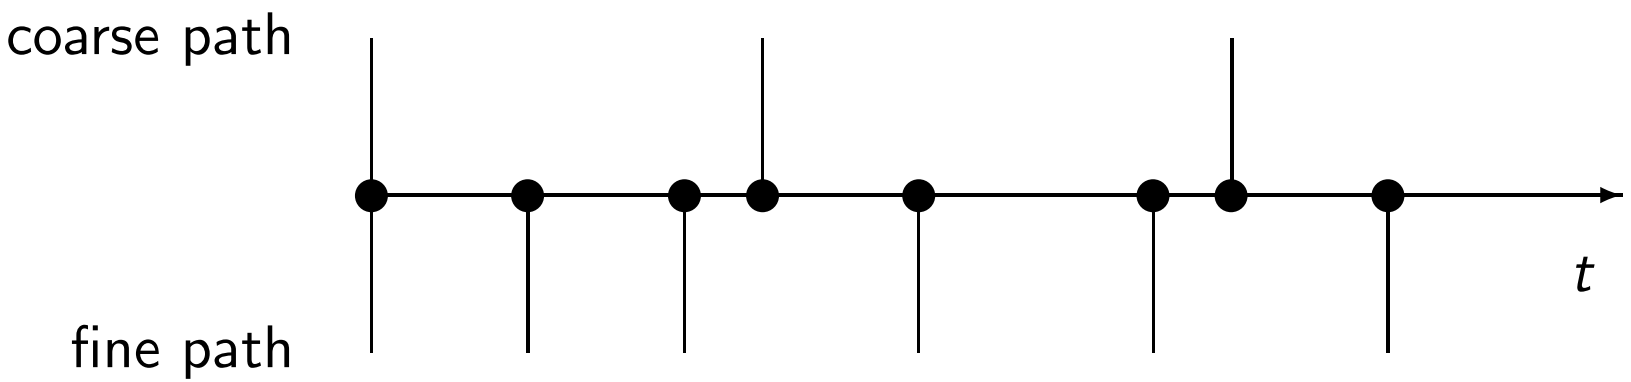

The fact that the timesteps are not nested is not a problem – strong convergence still ensures a strong coupling between the coarse and fine paths, because both approximate the true path.

Multilevel Monte Carlo simulation

Second challenge: we want to approximate a functional of the equilibrium distribution, the limit as time $T \to \infty$.

Key idea here comes from research by Rhee & Glynn (2014) on contracting Markov chains.

- on level ℓ we perform simulations for period $[-T_\ell, 0]$, evaluating the output at time $t=0$, and let $T_\ell \to \infty$ as ℓ increases.
- when doing the paths on levels ℓ and $\ell-1$, use the same Brownian motion for overlapping period $[-T_{\ell-1}, 0]$
- due to contraction property, effect of different starting points decays exponentially as $T_\ell \to \infty$

Numerical results: all works well, but numerical analysis looks very difficult because drift is not Lipschitz

New project 2: Feynman-Kac (Francisco Bernal)

Suppose that X_t satisfies the SDE

$$\mathrm{d}X_t = a(X_t, t)\,\mathrm{d}t + b(X_t, t)\,\mathrm{d}W_t$$

in bounded domain D, where W_t is an uncorrelated Brownian motion, and let

$$u(x,t) = \mathbb{E}\left[\int_t^\tau E(t,s)\,f(X_s,s)\,\mathrm{d}s + E(t,\tau)\,g(X_\tau,\tau) \mid X_t = x\right]$$

where τ is the first time at which X_t leaves D and

$$E(t_0,t_1) = \exp\left(-\int_{t_0}^{t_1} V(X_t,t)\,\mathrm{d}t\right).$$

Feynman-Kac theorem

If $f(x,t), g(x,t), V(x,t), a(x,t), b(x,t)$ are all Lipschitz continuous, then the Feynman-Kac theorem states that $u(x,t)$ satisfies the PDE

$$\frac{\partial u}{\partial t} + \sum_j a_j \frac{\partial u}{\partial x_j} + \tfrac{1}{2} \sum_{j,k,l} b_{j,k} b_{k,l} \frac{\partial^2 u}{\partial x_j \partial x_l} - V(x,t)\, u(x,t) + f(x,t) = 0$$

in domain D, subject to $u(x,t) = g(x,t)$, on the boundary ∂D.

Hence, can estimate $u(x,t)$ solution to a high-dimensional PDE at particular points (x,t), by Monte Carlo simulation of SDE.

This also extends to linear and nonlinear functionals of the PDE solution.

Numerical approximation

Let $\widehat{X}_t$ be the piecewise-constant Euler-Maruyama approximation and define

$$\widehat{E}(t_0, t_1) = \exp\left(-\int_{t_0}^{t_1} V(\widehat{X}_t, t)\,\mathrm{d}t\right),$$

and let

$$\widehat{u}(x,t) = \mathbb{E}\left[\int_t^{\widehat{\tau}} \widehat{E}(t,s)\, f(\widehat{X}_s, s)\,\mathrm{d}s + \widehat{E}(t,\widehat{\tau})\, g(\widehat{X}_{\widehat{\tau}}, \widehat{\tau}) \mid \widehat{X}_t = x\right].$$

with the Euler-Maruyama discretisation beginning at time t, and with $\widehat{\tau}$ being the exit time.

Numerical approximation

The Euler-Maruyama method has strong accuracy $O(h^{1/2})$, and the natural definition of τ gives an $O(h^{1/2})$ weak error too.

For standard Monte Carlo method, ε RMS accuracy needs $O(\varepsilon^{-2})$ paths, each with $h = O(\varepsilon^2)$, so total cost is $O(\varepsilon^{-4})$

Gobet reduced this to $O(\varepsilon^{-3})$ by shifting the boundary by $O(h^{1/2})$ to improve the weak accuracy to $O(h)$.

Alternatively, Higham *et al* use MLMC to achieve $O(\varepsilon^{-3}|\log\varepsilon|^3)$ complexity without shifting the boundary.

Numerical approximation

MLMC challenge:

When coarse or fine path exits the domain, the other one is within $O(h^{1/2})$ of boundary. However, there is a $O(h^{1/2})$ probability that it will not exit the domain until much later $\Longrightarrow V_\ell = O(h^{1/2})$.

How can we do better?

Similar to digital options, using splitting to give multiple instances of the second path.

$O(h^{1/2})$ time to expected exit of second path, so can afford to use $O(h^{-1/2})$ copies of second path.

This gives an approximation to the conditional expectation resulting in $\widehat{P}_\ell - \widehat{P}_{\ell-1} \approx O(h^{1/2})$, so $V_\ell \approx O(h)$.

Numerical results confirm this – numerical analysis is underway.

Conclusions

- multilevel idea is very simple; key is how to apply it in new situations
- 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 (helpful) approximation is much cheaper than finest
- currently, getting at least 100$\times$ savings for SPDEs and stochastic chemical reaction simulations

References

Webpage for my research/papers:

`people.maths.ox.ac.uk/gilesm/mlmc.html`

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

`people.maths.ox.ac.uk/gilesm/acta/`

– contains references to almost all MLMC research, including some very early related work by Achi Brandt

MLMC Community

Webpage: people.maths.ox.ac.uk/gilesm/mlmc_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, von Schwerin) – 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) – SDEs, SPDEs, numerical analysis, finance applications, stochastic chemical modelling
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
UTS (Baldeaux) – 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