

MLMC techniques for discontinuous functions

Michael B. Giles

Abstract The Multilevel Monte Carlo (MLMC) approach usually works well when estimating the expected value of a quantity which is a Lipschitz function of intermediate quantities, but if it is a discontinuous function it can lead to a much slower decay in the variance of the MLMC correction. This article reviews the literature on techniques which can be used to overcome this challenge in a variety of different contexts, and discusses recent developments using either a branching diffusion or adaptive sampling.

1 Introduction

The Multilevel Monte Carlo (MLMC) method is based on the telescoping sum

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

where \widehat{P}_ℓ represents an approximation to an output quantity of interest P on level ℓ , with the weak error $|\mathbb{E}[\widehat{P}_\ell - P]|$ and MLMC variance $\mathbb{V}[\widehat{P}_\ell - \widehat{P}_{\ell-1}]$, both decreasing as the level ℓ increases, but with the corresponding computational cost per sample increasing.

If \widehat{Y}_ℓ has expected value $\mathbb{E}[\widehat{P}_\ell - \widehat{P}_{\ell-1}]$, with variance V_ℓ and cost C_ℓ , then for a given target RMS error ε , the number of levels L and the number of independent samples on each level can be optimised [13, 14] to give an overall cost which is

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

Michael B. Giles

University of Oxford Mathematical Institute, Woodstock Rd, Oxford OX2 6GG, UK
e-mail: mike.giles@maths.ox.ac.uk

If $C_\ell V_\ell \rightarrow 0$ as $\ell \rightarrow \infty$, then the cost is dominated by the first term from level 0, and the cost is approximately $\varepsilon^{-2} C_0 V_0$, so proportional to ε^{-2} .

If $C_\ell V_\ell \rightarrow \text{const}$ as $\ell \rightarrow \infty$, then the contributions to the cost are spread almost equally across all levels and the cost is approximately $\varepsilon^{-2} L^2 C_L V_L$, proportional to $\varepsilon^{-2} |\log \varepsilon|^2$ if $\mathbb{E}[\widehat{P}_\ell - P]$ decays exponentially with ℓ .

Even worse, if $C_\ell V_\ell \rightarrow \infty$ then the cost is dominated by the contribution from the finest level and so is approximately $\varepsilon^{-2} C_L V_L$ which is $O(\varepsilon^{-2-(\gamma-\beta)/\alpha})$ if $\mathbb{E}[\widehat{P}_\ell - P] \sim 2^{-\alpha\ell}$, $V_\ell \sim 2^{-\beta\ell}$ and $C_\ell \sim 2^{\gamma\ell}$.

In most MLMC applications, P is a smooth function of some intermediate solution quantities, such as the solution of an SDE, a PDE with stochastic coefficients or initial/boundary data, or an estimate of an inner conditional expectation. Under these circumstances we usually have $\beta \geq \gamma$ and so the MLMC complexity is $O(\varepsilon^{-2})$ or $O(\varepsilon^{-2} |\log \varepsilon|^2)$.

This article is concerned with the small but important class of applications where P is a discontinuous function of the intermediate quantities, and because of this the MLMC variance V_ℓ can decay much more slowly, leading to the complexity falling into the third category of being $O(\varepsilon^{-2-(\gamma-\beta)/\alpha})$.

The good news is that there has been considerable research within the MLMC community to address this challenge. This article surveys the variety of methods which have been developed in the hope that this can aid researchers meeting similar challenges in future applications.

To illustrate things, we begin by detailing two specific application challenges which have motivated much of this research. We then discuss the many approaches which have been developed, several of which have borrowed ideas from the literature on computing sensitivities (the ‘‘greeks’’ in mathematical finance literature) of the form

$$\frac{\partial}{\partial \alpha} \mathbb{E} [f(\omega, \alpha)]$$

using the pathwise sensitivity approach [24] (also known as Infinitesimal Perturbation Analysis, IPA for short [30]) or Likelihood Ratio Method [31], or from methods from improving integrand smoothness to improve the rate of convergence for QMC integration [1, 6].

1.1 Challenge 1: nested expectation

Suppose f is a scalar function and we want to estimate the nested expectation $\mathbb{E} [f(\mathbb{E}[Z|X])]$, where the outer expectation is with respect to a random variable X and we will assume that the inner conditional expectation $\mathbb{E}[Z|X]$ has a bounded density near zero.

A very simple MLMC treatment ¹ uses $M_\ell = 2^\ell M_0$ inner samples on level ℓ , so estimators on level 0 and the higher levels are simply

¹ Note that if f is smooth, or at least Lipschitz, then it is better to use an ‘‘antithetic’’ estimator [8, 14, 15, 18], but this does not give a better order of convergence when f is discontinuous.

$$\widehat{Y}_0 = f(\overline{Z}^{(0, M_0)}), \quad \widehat{Y}_\ell = f(\overline{Z}^{(\ell, M_\ell)}) - f(\overline{Z}^{(\ell, M_{\ell-1})}),$$

where $\overline{Z}^{(\ell, M_\ell)}$ and $\overline{Z}^{(\ell, M_{\ell-1})}$ represent independent averages of M_ℓ and $M_{\ell-1}$ independent samples of Z , all conditional on the same value of X [14, 15].

If $\mathbb{V}[Z|X]$ is finite, and f is Lipschitz with constant L_f , then

$$\begin{aligned} \mathbb{E} \left[\left(f(\overline{Z}^{(\ell, M_\ell)}) - f(\mathbb{E}[Z|X]) \right)^2 \mid X \right] &\leq L_f^2 \mathbb{E} \left[\left(\overline{Z}^{(\ell, M_\ell)} - \mathbb{E}[Z|X] \right)^2 \mid X \right] \\ &= L_f^2 M_\ell^{-1} \mathbb{V}[Z|X], \end{aligned}$$

and hence $\mathbb{E}[\widehat{Y}_\ell^2|X] \leq 4L_f^2(M_\ell^{-1} + M_{\ell-1}^{-1})\mathbb{V}[Z|X]$ for $\ell > 0$. If $\mathbb{V}[Z|X]$ is uniformly bounded it follows that $V_\ell = O(M_\ell^{-1})$. If the cost of each conditional sample of Z is $O(1)$ then $C_\ell = O(M_\ell)$ and hence the complexity is $O(\varepsilon^{-2} |\log \varepsilon|^2)$.

Unfortunately, the situation is significantly poorer when f is the Heaviside step function H defined by $H(x)=0$ if $x < 0$, and $H(x)=1$ if $x \geq 0$. This occurs in many applications because $\mathbb{P}[\mathbb{E}[Z|X] > K] = \mathbb{E}[H(\mathbb{E}[Z|X] - K)]$, so it corresponds to the probability of a conditional expectation exceeding some threshold K , which is a very important quantity in risk calculations.

If $K=0$ and $E[Z|X]$ has a bounded density near zero then there is an $O(M_\ell^{-1/2})$ probability that $|E[Z|X]| = O(M_\ell^{-1/2})$, which is the circumstance under which there is an $O(1)$ probability that $\widehat{Y}_\ell = \pm 1$ due to $\overline{Z}^{(\ell, M_\ell)}$ being positive and $\overline{Z}^{(\ell, M_{\ell-1})}$ being negative, or vice versa. Hence $V_\ell \approx O(M_\ell^{-1/2})$ and the complexity is approximately $O(\varepsilon^{-5/2})$ [19].

This challenge is the primary motivation for Section 7, and also arises in the context of Section 3.

1.2 Challenge 2: discontinuous payoff function

In the case of a scalar SDE

$$dS_t = a(S_t) dt + b(S_t) dW_t, \quad (1)$$

with an output quantity of interest $P \equiv f(S_T)$, the standard estimator is

$$\widehat{Y}_\ell = \widehat{P}_\ell - \widehat{P}_{\ell-1}$$

where the same Brownian motion W_t is used to calculate both \widehat{P}_ℓ and $\widehat{P}_{\ell-1}$, but with different uniform timesteps h_ℓ and $h_{\ell-1}$.

If f is Lipschitz with constant L_f , then

$$V_\ell \leq \mathbb{E} \left[(\widehat{P}_\ell - \widehat{P}_{\ell-1})^2 \right] \leq L_f^2 \mathbb{E} \left[(\widehat{S}_\ell - \widehat{S}_{\ell-1})^2 \right]$$

where \widehat{S}_ℓ is the level ℓ numerical approximation to S_T . Hence, based on standard strong convergence results [28] we have $V_\ell = O(h_\ell)$ for an Euler-Maruyama discretisation of the SDE, and $V_\ell = O(h_\ell^2)$ for the first order Milstein discretisation. The cost C_ℓ is $O(h_\ell^{-1})$ in both cases, giving MLMC complexities of $O(\varepsilon^{-2} |\log \varepsilon|^2)$ and $O(\varepsilon^{-2})$, respectively,

In mathematical finance, a digital call option payoff is 0 or 1, depending on whether S_T is below or above the strike K , so the payoff function can be written as $f(S_T) = H(S_T - K)$. The MLMC problem is that a small difference between the coarse and fine paths can give a payoff difference of ± 1 if the two paths straddle the strike, i.e. are on different sides of the strike.

When using the Euler-Maruyama approximation of the SDE, $\widehat{S}_\ell - \widehat{S}_{\ell-1} = O(h_\ell^{1/2})$. Speaking loosely (see [4, 21] for the rigorous analysis) an $O(h_\ell^{1/2})$ fraction of fine/coarse pairs straddle the strike, so $V_\ell = O(h_\ell^{1/2})$ and hence the complexity is $O(\varepsilon^{-5/2})$.

Similarly, using the Milstein approximation gives $\widehat{S}_\ell - \widehat{S}_{\ell-1} = O(h_\ell)$ so $V_\ell = O(h_\ell)$. This is clearly better, and gives a complexity which is $O(\varepsilon^{-2} |\log \varepsilon|^2)$, but there is still the problem that most MLMC samples Y_ℓ are zero on the finer levels, so the kurtosis is $O(h_\ell^{-1})$ which causes problems in practice in estimating V_ℓ accurately to determine the number of samples N_ℓ to use on level ℓ . In addition, there is the difficulty that the Milstein discretisation of multi-dimensional SDEs often requires the simulation of Lévy areas, though this problem can be addressed through the use of an antithetic estimator [23].

This challenge is the primary motivation for Sections 2, 4, 5 and 6, also also arises in Sections 3 and 7.

2 Explicit smoothing

The pathwise sensitivity analysis (or IPA) approach to compute the parameter sensitivities known as Greeks in mathematical finance [24] requires that the payoff function f is continuous and piecewise smooth. This is clearly a problem with digital options, and one standard approach is to smooth the payoff function by replacing the Heaviside step function H with a smoothed approximation $H_\delta(x) \equiv g(x/\delta)$, with $g(x) \rightarrow 0$ as $x \rightarrow -\infty$ and $g(x) \rightarrow 1$ as $x \rightarrow +\infty$, so the discontinuity is smoothed over a width of $O(\delta)$.

For financial reasons, the preference is often to use a one-sided smoothing, such as the piecewise linear approximation shown in yellow in Figure 1. This one-sided approximation introduces a weak error, or bias, which is $O(\delta)$. If it is used for MLMC, then $H'_\delta(S_T) = \delta^{-1}$ for the $O(\delta)$ fraction of the paths which end up in the ramp region, and therefore $V_\ell = O(\delta \times (\delta^{-1})^2) = O(\delta^{-1})$. Hence the optimal choice of δ involves a tradeoff between bias and variance.

The bias can be reduced by making the smoothing anti-symmetric about $x = 0$ so that $H_\delta(x) - H(x) = -(H_\delta(-x) - H(-x))$, for example by choosing $g(x) \equiv \Phi(x)$

as illustrated in orange in Figure 1. If S_T has the smooth probability density $\rho(S)$ then the weak error is

$$\int_{-\infty}^{\infty} (H_{\delta}(S-K) - H(S-K)) \rho(S) dS = \delta \int_{-\infty}^{\infty} (g(x) - H(x)) \rho(K+x\delta) dx$$

and a Taylor series expansion of $\rho(K+x\delta)$ results in the asymptotic error expansion

$$a_1\rho(K)\delta + a_2\rho'(K)\delta^2 + a_3\rho''(K)\delta^3 + a_4\rho'''(K)\delta^4 + O(\delta^5)$$

where

$$a_k = \int_{-\infty}^{\infty} x^{k-1} (g(x) - H(x)) dx.$$

If $g(x) - H(x) = -(g(-x) - H(-x))$ then $a_1 = a_3 = 0$, and

$$a_2 = 2 \int_0^{\infty} x(g(x) - 1) dx, \quad a_4 = 2 \int_0^{\infty} x^3(g(x) - 1) dx.$$

If $g(x)$ is monotonic, then $a_2 \neq 0$, but by considering non-monotonic functions such as $g(x) = (4/3)\Phi(x) - (1/3)\Phi(2x)$ it is possible to set $a_2 = 0$ making the weak error $O(\delta^4)$. Hence, to achieve $O(\varepsilon)$ accuracy overall we need $\delta = O(\varepsilon^{1/4})$, and then on the coarsest levels $V_{\ell} = O(\delta^{-1}) = O(\varepsilon^{-1/4})$ so the overall complexity is approximately $O(\varepsilon^{-2-1/4})$ in the best cases where the overall cost is dominated by the cost on the coarsest levels.

Giles, Nagapetyan & Ritter [22] used explicit smoothing for estimating cumulative distribution functions (CDFs). For a scalar random variable X , to estimate $C(x) = \mathbb{P}(X < x) = \mathbb{E}[H(x-X)]$, the approach they adopted was to use MLMC to estimate $C(x_j)$ for a set of spline points x_j and then interpolate these values with a cubic spline. Overall, their method balanced three weak errors, the SDE discretisation error on the finest level, the smoothing error due to H_{δ} , and the cubic spline interpolation error, in addition to the MLMC sampling error.




Fig. 1 Two explicitly smoothed versions of the Heaviside step function for a digital call option

3 Integration/differentiation and Malliavin calculus

Krumscheid & Nobile [29] used a slightly different approach for estimating CDFs, particularly in the context of risk estimation. Starting from the identity

$$\frac{d}{dx} \mathbb{E} [\max(0, x - S_T)] = \mathbb{E} [H(x - S_T)]$$

they used MLMC to estimate $\mathbb{E} [\max(0, x_j - S_T)]$ for a set of spline points x_j , interpolated these with a cubic spline, and then differentiated the spline to obtain an approximation to the CDF $C(x)$. This avoids the extra weak error due to smoothing the Heaviside function, but differentiating the cubic spline amplifies the noise in the spline data.

On a similar note, Altmayer & Neuenkirch [2] used Malliavin calculus integration by parts to treat discontinuous payoffs based on solutions of the Heston stochastic volatility SDE. They observed that asymptotically this improves the MLMC variance on the finer levels, but it increases the variance on coarse levels. To address this, they split the payoff into a smooth part which they treated with the standard MLMC approach, and a compact-support discontinuous part for which they used the Malliavin MLMC.

Malliavin calculus was originally developed for computing sensitivities, so this is another example of the literature on sensitivity calculations being exploited to develop improved MLMC algorithms.

4 Conditional expectation

When using the first order Milstein discretisation for an SDE, one way to improve the MLMC variance for digital options is to switch to the Euler-Maruyama approximation for the final timestep, and then take the conditional expectation with respect to the final fine path Brownian increment ΔW [12, 17].

For the fine path approximation of the scalar SDE (1) with N timesteps of size h_ℓ , the path value S_T at the final time T is given by

$$\widehat{S}_T^f = \widehat{S}_{T-h_\ell}^f + a(\widehat{S}_{T-h_\ell}^f) h_\ell + b(\widehat{S}_{T-h_\ell}^f) \Delta W_N,$$

and therefore the conditional expected value for the digital call option is

$$\widehat{P}_\ell^f = \mathbb{E} \left[H(\widehat{S}_T^f - K) \mid \widehat{S}_{T-h_\ell}^f \right] = \Phi \left(\frac{\widehat{S}_{T-h_\ell}^f + a(\widehat{S}_{T-h_\ell}^f) h_\ell - K}{b(\widehat{S}_{T-h_\ell}^f) \sqrt{h_\ell}} \right).$$

Similarly, for the coarse path with coarse timestep $h_{\ell-1} = 2 h_\ell$, the Brownian increment for the final coarse timestep is the sum of the last two Brownian increments for the fine path, $\Delta W_{N-1} + \Delta W_N$, and therefore

$$\widehat{S}_T^c = \widehat{S}_{T-h_{\ell-1}}^c + a(\widehat{S}_{T-h_{\ell-1}}^c) h_{\ell-1} + b(\widehat{S}_{T-h_{\ell-1}}^c) (\Delta W_{N-1} + \Delta W_N),$$

from which we obtain

$$\begin{aligned} \widehat{P}_{\ell-1}^c &= \mathbb{E} \left[H(\widehat{S}_T^c - K) \mid \widehat{S}_{T-h_{\ell-1}}^c, \Delta W_{N-1} \right] \\ &= \Phi \left(\frac{\widehat{S}_{T-h_{\ell-1}}^c + a(\widehat{S}_{T-h_{\ell-1}}^c) h_{\ell-1} + b(\widehat{S}_{T-h_{\ell-1}}^c) \Delta W_{N-1} - K}{b(\widehat{S}_{T-h_{\ell-1}}^c) \sqrt{h_{\ell}}} \right). \end{aligned}$$

With $\widehat{Y}_{\ell} \equiv \widehat{P}_{\ell} - \widehat{P}_{\ell-1}$, numerical analysis [17] proves that $V_{\ell} \approx O(h_{\ell}^{3/2})$ so the MLMC complexity is $O(\varepsilon^{-2})$. Heuristically, this is because there is an $O(h_{\ell}^{1/2})$ probability of paths being within $O(h_{\ell}^{1/2})$ of the strike K , and for these

$$\widehat{S}_{T-h_{\ell-1}}^f - \widehat{S}_{T-h_{\ell-1}}^c = O(h_{\ell}), \quad \frac{\partial \widehat{P}}{\partial \widehat{S}} = O(h_{\ell}^{-1/2}) \implies \widehat{P}_{\ell} - \widehat{P}_{\ell-1} = O(h_{\ell}^{1/2}),$$

so $V_{\ell} \approx O(h_{\ell}^{1/2} \times (h_{\ell}^{1/2})^2) = O(h_{\ell}^{3/2})$. In addition, the kurtosis is improved to $O(h_{\ell}^{-1/2})$.

Unfortunately, the conditional expectation approach does not help when the Euler-Maruyama discretisation is used for the entire path since $\widehat{S}_{T-h_{\ell-1}}^f - \widehat{S}_{T-h_{\ell-1}}^c = O(h_{\ell}^{1/2})$ and so $\widehat{P}_{\ell} - \widehat{P}_{\ell-1} = O(1)$

The use of this kind of conditional expectation is a standard technique for smoothing the payoff to enable IPA/pathwise sensitivity calculations [24]. Another example is a down-and-out barrier option, where the option is knocked out if the path drops below a certain value. In this case the payoff can be smoothed by computing the probability of this happening, conditional on the computed path approximations at discrete timesteps [24]. Again, this works well for MLMC when using the first order Milstein discretisation [12, 17], but it does not help with the Euler-Maruyama discretisation.

A different kind of conditional expectation smoothing was introduced by Achtsis, Cools & Nuyens [1] and Bayer, Siebenmorgen & Tempone [6] to improve the convergence of QMC computations, and then used by Bayer, Ben Hammouda & Tempone [5] to improve the MLMC variance for digital options.

In its simplest form, they split the random inputs for the numerical simulation into a scalar Z and the remainder Z_r , and express the desired MLMC level ℓ expectation as

$$\mathbb{E}[\widehat{P}_{\ell} - \widehat{P}_{\ell-1}] = \mathbb{E} \left[\mathbb{E}[\widehat{P}_{\ell} - \widehat{P}_{\ell-1} \mid Z_r] \right]$$

and observe that in many financial applications it is possible to perform this split in a way such that the conditional expectations $\mathbb{E}[\widehat{P}_{\ell} \mid Z_r]$, $\mathbb{E}[\widehat{P}_{\ell-1} \mid Z_r]$ are smooth functions of Z_r , and can be evaluated analytically or very accurately by 1D numerical quadrature when there is just a single discontinuity with respect to changes in Z .

For a scalar SDE, Z could be the terminal value of the driving Brownian motion, in which case Z_r would represent the other Normal random variables required for a Brownian Bridge construction of the Brownian increments.

5 Change of measure

Another approach to treating digital options using the Milstein discretisation is to use a change of measure [9, 14], which has connections to the Likelihood Ratio Method (LRM) that is used for sensitivity analysis [31].

For both the fine and coarse paths, we have conditional Gaussian distributions for \widehat{S}_T , with slightly different means and variances, as illustrated in Figure 2. We can therefore perform a change of measure to the same Gaussian distribution with mean μ and variance σ^2 , also illustrated in Figure 2, and then pick the same sample \widehat{S}_T for both paths from this common Gaussian distribution.

The resulting MLMC estimator is

$$\widehat{Y}_\ell = f(\widehat{S}_T) (R_\ell - R_{\ell-1})$$

where $R_\ell, R_{\ell-1}$ are the respective Radon-Nikodym derivatives for the fine and coarse paths. For the scalar SDE (1), R_ℓ is

$$R_\ell = \frac{\sigma}{b(\widehat{S}_{T-h_\ell}^f) \sqrt{h_\ell}} \frac{\exp\left(-(\widehat{S}_T - \widehat{S}_{T-h_\ell}^f - a(\widehat{S}_{T-h_\ell}^f) h_\ell)^2 / (2 b^2(\widehat{S}_{T-h_\ell}^f) h_\ell)\right)}{\exp\left(-(\widehat{S}_T - \mu)^2 / (2\sigma^2)\right)}$$

and $R_{\ell-1}$ is defined similarly. It can be shown that the difference $R_\ell - R_{\ell-1}$ is approximately $O(h_\ell^{1/2})$, which implies that $V_\ell \approx O(h_\ell)$. To improve the variance we note that the conditional expected value of Radon-Nikodym derivatives is always 1,




Fig. 2 Coarse and fine path conditional Gaussian distributions, plus third common distribution




Fig. 3 Illustration of coarse and fine jump-diffusion paths with jumps before and after $T = 1$.

i.e. $\mathbb{E}[R_\ell | \widehat{S}_{T-h_\ell}^f] = \mathbb{E}[R_{\ell-1} | \widehat{S}_{T-h_{\ell-1}}^c, \Delta W_{N-1}] = 1$, and therefore we can change the definition of \widehat{Y}_ℓ to

$$\widehat{Y}_\ell = \left(f(\widehat{S}_T) - f(\mu) \right) (R_\ell - R_{\ell-1})$$

without changing its expected value. This estimator is now non-zero only when \widehat{S}_T and μ are on opposite sides of the strike K , which occurs for an $O(h_\ell^{1/2})$ fraction of coarse/fine paths. Hence the new MLMC variance V_ℓ is approximately $O(h_\ell^{3/2})$, as with the use of the analytic conditional expectation.

The benefit of this approach is that it works well in multiple dimensions when it is often not possible to evaluate the analytic conditional expectation [9, 14]. However, again it does not help with the full path Euler-Maruyama discretisation because that gives $R_\ell - R_{\ell-1} = O(1)$.

An earlier use of a change of measure in an MLMC computation was by Xia [33, 34] for a Merton-style jump-diffusion SDE with a path-dependent jump rate $\lambda(S, t)$. The challenge in this application, as illustrated in Figure 3, is that the coarse and fine paths will jump at different times, and one might jump just before the final time T , and the other just after, leading to a large jump in the computed value of $f(S_T)$. The path-dependent jump rate was treated by using the thinning technique of Glasserman & Merener [25], over-sampling possible jump times using a uniform rate $\lambda_{sup} > \lambda(S, t)$ and then using an acceptance/rejection step to select the real jump times. Xia modified this with a change of measure to ensure the same acceptance/rejection decision for both the fine and coarse paths so that they both jump at the same time. This leads to an estimator of the form

$$\widehat{Y}_\ell = \widehat{P}_\ell R_\ell - \widehat{P}_{\ell-1} R_{\ell-1}.$$

When combined with a first order Milstein discretisation of the SDE between the jump times, this gives $V_\ell = O(h_\ell^2)$ for Lipschitz payoff functions such as a standard put or call option [33, 34].

6 Splitting

Returning to the challenge of digital options arising from the solution of an SDE, a third approach is to use path-splitting to generate an unbiased estimate of the conditional expectation introduced in Section 4 [14].

This is a variant of the general splitting technique [3]. As illustrated in Figure 4, it involves performing a standard fine path simulation up until one timestep before the final time T , and then performing multiple independent simulations of the final timestep, averaging the payoff for each of these to get an approximation of the conditional expectation. The same is done for the coarse path except that each of the splits uses the same ΔW_{N-1} that was used for the second to last fine path timestep.

Since the computational cost of the path up to the splitting time is $O(h_\ell^{-1})$, it means that up to $O(h_\ell^{-1})$ splits can be used without increasing the path cost significantly. If M_ℓ splits are used, then the standard splitting variance analysis gives

$$\mathbb{V}[\widehat{Y}_\ell] = \mathbb{V} \left[\mathbb{E}[\widehat{P}_\ell - \widehat{P}_{\ell-1} \mid \{\Delta W_n\}_{n < N}] \right] + M_\ell^{-1} \mathbb{E} \left[\mathbb{V}[\widehat{P}_\ell - \widehat{P}_{\ell-1} \mid \{\Delta W_n\}_{n < N}] \right].$$

As discussed previously $\mathbb{V} \left[\mathbb{E}[\widehat{P}_\ell - \widehat{P}_{\ell-1} \mid \{\Delta W_n\}_{n < N}] \right] = O(h_\ell^{3/2})$, and similarly it can be argued that $\mathbb{E} \left[\mathbb{V}[\widehat{P}_\ell - \widehat{P}_{\ell-1} \mid \{\Delta W_n\}_{n < N}] \right] = O(h_\ell)$. Therefore choosing M_ℓ to lie between $O(h_\ell^{-1})$ and $O(h_\ell^{-1/2})$ ensures the benefits of the splitting are obtained without significantly increasing the computational cost per sample.




Fig. 4 Path splitting in final timestep to estimate conditional expectation

As an additional bonus, one can use the more accurate Milstein discretisation for the final timestep, instead of switching to the Euler-Maruyama discretisation. Burgos [9, 10] gives more details of the analysis, and also used the same approach for pathwise sensitivity analysis for a variety of financial options.

Giles & Bernal [16] also used splitting for Feynman-Kac functionals arising from stopped diffusions, SDE simulations which terminate when the solution path leaves a prescribed domain. The issue here is that when a fine path exits, there is an $O(h_\ell^{1/2})$ probability that the corresponding coarse path does not leave until much later. This is addressed by estimating a conditional expectation by splitting the coarse path into $O(h_\ell^{-1/2})$ independent sub-simulations which continue until each of them leaves the domain. V_ℓ is improved from $O(h_\ell^{1/2})$ to approximately $O(h_\ell)$ without a significant increase in the cost per sample, and finally the MLMC complexity achieved is $O(\varepsilon^{-2} |\log \varepsilon|^3)$.

None of the three methods introduced so far (conditional expectation, change of measure, splitting) helps when using the Euler-Maruyama discretisation. For this, a new method has recently been developed by Giles & Haji-Ali [20].

It again uses splitting, but inspired by the simulation of branching diffusions, it considers splits at multiple deterministic times, as illustrated in Figure 5 which shows the logical structure of a set of split paths. Here we are considering a simulation on the unit time interval. A single pair of fine/coarse paths is calculated up to time $t = 1/2$, with the number of fine timesteps being $\frac{1}{2}h_\ell^{-1}$. This simulation is then split into two separate independent simulations up to time $t = 3/4$, with the two simulations between them accounting for an additional $\frac{1}{2}h_\ell^{-1}$ fine timesteps. There are further splits at $t = 3/4$, then at $t = 7/8$, and so on, with the final split when there is just one coarse timestep left.

The total number of fine timesteps simulated is $O(h_\ell^{-1} |\log h_\ell|)$ so the computational cost is only slightly increased compared to the original method with a single pair of fine/coarse paths. \tilde{Y}_ℓ is defined to be the average of the values $\hat{P}_\ell - \hat{P}_{\ell-1}$ for each of the final paths, and it can be proved that its variance is $O(h_\ell)$, the same




Fig. 5 Repeated path splitting to estimate conditional expectation

asymptotic order of convergence as for Lipschitz payoff functions [20]. The kurtosis is also improved, so this technique fully addresses the challenge of using MLMC with the Euler-Maruyama discretisation to estimate digital option values.

7 Adaptive sampling

We return now to the challenge of estimating the nested expectation $\mathbb{E} [H (\mathbb{E}[Z|X])]$ and we note that we only need an accurate estimate of the inner conditional expectation $\mathbb{E}[Z|X]$ when it is near zero. This observation is the basis for the development of adaptive sampling by Broadie, Du & Moallemi [7] within a standard Monte Carlo procedure. This was then extended to adaptive sampling combined with MLMC by Giles & Haji-Ali [19] by defining the number of inner samples M_ℓ on level ℓ to be

- $M_\ell = 2^\ell M_0$ inner samples when $|\mathbb{E}[Z|X]| \gg \sqrt{\mathbb{V}[Z|X]/(2^\ell M_0)}$

This is the smallest number of samples used on level ℓ . $\sqrt{\mathbb{V}[Z|X]/(2^\ell M_0)}$ is the standard deviation of the Monte Carlo estimate for $\mathbb{E}[Z|X]$, so the inequality means that this number of samples is sufficient to be very sure that the estimate has the correct sign.

- $M_\ell = 4^\ell M_0$ inner samples when $|\mathbb{E}[Z|X]| = O(\sqrt{\mathbb{V}[Z|X]/(4^\ell M_0)})$

This is the maximum number of samples used on level ℓ . In this case, the estimate of $\mathbb{E}[Z|X]$ may have the incorrect sign, but this will only happen when $|\mathbb{E}[Z|X]| = O(2^{-\ell})$ which occurs with probability $O(2^{-\ell})$. Likewise, the total cost of the higher number of samples in this region is $O(2^{-\ell} \times 4^\ell) = O(2^\ell)$, so it does not significantly increase the overall average cost.

- $2^\ell M_0 < M_\ell < 4^\ell M_0$ for intermediate values

In this region the number of samples is chosen to be very sure that the estimate of $\mathbb{E}[Z|X]$ has the correct sign, and at the same time the total cost is $O(2^\ell)$.

Overall, this adaptive sampling approach leads to $C_\ell \sim 2^\ell$, $V_\ell \sim 2^{-\ell}$ and hence a complexity of roughly $O(\varepsilon^{-2})$ [19]. However, the kurtosis is $O(2^\ell)$ since only an $O(2^{-\ell})$ fraction of the outer samples give non-zero values for \widehat{Y}_ℓ .

Haji-Ali, Spence & Teckentrup [27] have further extended this to estimate quantities of the form

$$\mathbb{P}[G \in \Omega] \equiv \mathbb{E}[1_{G \in \Omega}]$$

where G is a d -dimensional random variable which cannot be sampled directly. In their paper they consider in particular the two challenges in this article. In the context of the digital option with the Euler-Maruyama discretisation on the unit time interval, the adaptive sampling varies the timestep used on level ℓ so that

- $h_\ell = 2^{-\ell}$ when $|\widehat{S}_\ell - K|$ is large compared to the strong error in the path approximation
- $h_\ell = 4^{-\ell}$ when $|\widehat{S}_\ell - K|$ is of the same order as the strong error

- $2^{-\ell} < h_\ell < 4^{-\ell}$ for intermediate values

A Brownian bridge construction is used when the timestep needs to be refined as part of the adaptation procedure from its initial value $h_\ell = 2^{-\ell}$. The adaptation again leads to $C_\ell \sim 2^\ell$, $V_\ell \sim 2^{-\ell}$ and hence a complexity of roughly $O(\varepsilon^{-2})$, but there is again a high kurtosis [27].

In earlier research, Elfverson, Hellman & Malqvist [11] considered estimation of $\mathbb{E}[H(X)]$ where X cannot be sampled exactly but there is a sequence of approximations $X'_0, X'_1, X'_2, \dots, X$ of increasing accuracy and increasing cost. Motivated by PDE applications with a well-behaved truncation error so that there are uniform geometric bounds on $|X'_j - X|$, level ℓ in their method uses

$$\widehat{X}_\ell = X'_j, \quad j = \min\{\ell, \min j : |\widehat{X}'_j - X| < |X|\}$$

and achieves similarly good MLMC benefits. This idea is essentially the same as in the work of Haji-Ali *et al* but requiring a uniform bound on $|X'_j - X|$ is significantly more restrictive than the bounds on $\mathbb{E}[|X'_j - X|^q]$ for some $q > 2$ required by Haji-Ali *et al*.

A final comment is that the analysis of Haji-Ali, Spence & Teckentrup can be generalised to a product of an indicator function and a Lipschitz function, $\mathbb{E}[1_{G \in \Omega} f(S)]$, and so can handle barrier options. Furthermore, Haji-Ali & Spence have extended the adaptive sampling methodology to an extremely challenging triply-nested expectation which arises in mathematical finance [26]. By incorporating the randomised MLMC treatment of Rhee & Glynn [32] to handle the time discretisation of the underlying SDEs as well as the sampling for the inner conditional expectations, they achieve an overall complexity of approximately $O(\varepsilon^{-2})$ which is very impressive for such a difficult application.




Fig. 6 Error distributions for two conditional expectations with i) few samples being needed to ensure the correct sign (left), and ii) many samples being insufficient to ensure the correct sign (centre). The blue line represents the Heaviside step function.

8 Conclusions

It is worth repeating that in most MLMC applications the output quantity of interest is a Lipschitz function of the intermediate simulation quantities, so good strong convergence for the intermediate quantities leads automatically to a good rate of convergence of the MLMC variance V_ℓ .

For those applications in which the function is discontinuous, this article shows there is an extensive literature with a variety of different approaches to improve the MLMC variance and try to recover the optimal $O(\varepsilon^{-2})$ complexity. It is notable that many of these methods have adapted ideas from Monte Carlo sensitivity analysis which also has problems with discontinuous functionals. It is hoped that this survey will assist future researchers facing similar challenges in other new application areas.

Acknowledgements This paper is based on research with many students, postdocs and other collaborators and I am grateful to all of them. Funding for the research has been provided primarily by the UK Engineering and Physical Sciences Research Council through grants EP/E031455/1, EP/H05183X/1 and EP/P020720/2. The paper was written while visiting the Oden Institute at UT Austin, and I thank my hosts for their warm hospitality.

References

1. N. Achtsis, R. Cools, and D. Nuyens. Conditional sampling for barrier option pricing under the LT method. *SIAM Journal on Financial Mathematics*, 4:327–352, 2013.
2. M. Altmayer and A. Neuenkirch. Multilevel Monte Carlo quadrature of discontinuous payoffs in the generalized Heston model using Malliavin integration by parts. *SIAM Journal on Financial Mathematics*, 6(1):22–52, 2015.
3. S. Asmussen and P.W. Glynn. *Stochastic Simulation*. Springer, New York, 2007.
4. R. Avikainen. On irregular functionals of SDEs and the Euler scheme. *Finance and Stochastics*, 13(3):381–401, 2009.
5. C. Bayer, C. Ben Hammouda, and R. Tempone. Multilevel Monte Carlo combined with numerical smoothing for robust and efficient option pricing and density estimation. *ArXiv preprint: 2003.05708*, 2020.
6. C. Bayer, M. Siebenmorgen, and R. Tempone. Smoothing the payoff for efficient computation of basket option prices. *Quantitative Finance*, 18(3):491–505, 2018.
7. M. Broadie, Y. Du, and C.C. Moallemi. Efficient risk estimation via nested sequential simulation. *Management Science*, 57(6):1172–1194, 2011.
8. K. Bujok, B. Hambly, and C. Reisinger. Multilevel simulation of functionals of Bernoulli random variables with application to basket credit derivatives. *Methodology and Computing in Applied Probability*, 17(3):579–604, 2015.
9. S. Burgos. *The computation of Greeks with multilevel Monte Carlo*. DPhil thesis, University of Oxford, 2014.
10. S. Burgos and M.B. Giles. Computing Greeks using multilevel path simulation. In L. Plaskota and H. Woźniakowski, editors, *Monte Carlo and Quasi-Monte Carlo Methods 2010*, pages 281–296. Springer, 2012.
11. D. Elfverson, F. Hellman, and A. Målqvist. A multilevel Monte Carlo method for computing failure probabilities. *SIAM/ASA Journal on Uncertainty Quantification*, 4(1):312–330, 2016.
12. M.B. Giles. Improved multilevel Monte Carlo convergence using the Milstein scheme. In A. Keller, S. Heinrich, and H. Niederreiter, editors, *Monte Carlo and Quasi-Monte Carlo Methods 2006*, pages 343–358. Springer, 2008.

13. M.B. Giles. Multilevel Monte Carlo path simulation. *Operations Research*, 56(3):607–617, 2008.
14. M.B. Giles. Multilevel Monte Carlo methods. *Acta Numerica*, 24:259–328, 2015.
15. M.B. Giles. MLMC for nested expectations. In *Contemporary Computational Mathematics - A Celebration of the 80th Birthday of Ian Sloan*. Springer, 2018.
16. M.B. Giles and F. Bernal. Multilevel estimation of expected exit times and other functionals of stopped diffusions. *SIAM/ASA Journal on Uncertainty Quantification*, 6(4):1454–1474, 2018.
17. M.B. Giles, K. Debrabant, and A. Röbler. Analysis of multilevel Monte Carlo path simulation using the Milstein discretisation. *Discrete and Continuous Dynamical Systems - series B*, 24(8):3881–3903, 2019.
18. M.B. Giles and T. Goda. Decision-making under uncertainty: using MLMC for efficient estimation of EVPP. *Statistics and Computing*, 29(4):739–751, 2019.
19. M.B. Giles and A.-L. Haji-Ali. Multilevel nested simulation for efficient risk estimation. *SIAM/ASA Journal on Uncertainty Quantification*, 7(2):497–525, 2019.
20. M.B. Giles and A.-L. Haji-Ali. Multilevel path branching for digital options. *arXiv pre-print: 2209.03017*, 2022.
21. M.B. Giles, D.J. Higham, and X. Mao. Analysing multilevel Monte Carlo for options with non-globally Lipschitz payoff. *Finance and Stochastics*, 13(3):403–413, 2009.
22. M.B. Giles, T. Nagapetyan, and K. Ritter. Multilevel Monte Carlo approximation of distribution functions and densities. *SIAM/ASA Journal on Uncertainty Quantification*, 3(1):267–295, 2015.
23. M.B. Giles and L. Szpruch. Antithetic multilevel Monte Carlo estimation for multi-dimensional SDEs without Lévy area simulation. *Annals of Applied Probability*, 24(4):1585–1620, 2014.
24. P. Glasserman. *Monte Carlo Methods in Financial Engineering*. Springer, New York, 2004.
25. P. Glasserman and N. Merener. Convergence of a discretization scheme for jump-diffusion processes with state-dependent intensities. *Proc. Royal Soc. London A*, 460:111–127, 2004.
26. A.-L. Haji-Ali and J. Spence. Efficient risk estimation for the credit valuation adjustment. *in preparation*, 2022.
27. A.-L. Haji-Ali, J. Spence, and A. Teckentrup. Adaptive multilevel Monte Carlo for probabilities. *SIAM Journal on Numerical Analysis*, 60(4):2125–2149, 2022.
28. P.E. Kloeden and E. Platen. *Numerical Solution of Stochastic Differential Equations*. Springer, Berlin, 1992.
29. S. Krumscheid and F. Nobile. Multilevel Monte Carlo approximation of functions. *SIAM/ASA Journal on Uncertainty Quantification*, 6(3):1256–1293, 2018.
30. P. L’Ecuyer. A unified view of the IPA, SF and LR gradient estimation techniques. *Management Science*, 36(11):1364–1383, 1990.
31. P. L’Ecuyer. On the interchange of derivative and expectation for likelihood ratio derivative estimators. *Management Science*, 41(4):738–748, 1995.
32. C.-H. Rhee and P.W. Glynn. Unbiased estimation with square root convergence for SDE models. *Operations Research*, 63(5):1026–1043, 2015.
33. Y. Xia. *Multilevel Monte Carlo for jump processes*. DPhil thesis, University of Oxford, 2014.
34. Y. Xia and M.B. Giles. Multilevel path simulation for jump-diffusion SDEs. In L. Plaskota and H. Woźniakowski, editors, *Monte Carlo and Quasi-Monte Carlo Methods 2010*, pages 695–708. Springer, 2012.