

# Multilevel Monte Carlo methods using approximate distributions

Mike Giles  
Oliver Sheridan-Methven

Mathematical Institute, University of Oxford

28th Biennial Conference on Numerical Analysis

June 28th, 2019

# Overview

- Multilevel Monte Carlo (MLMC)
- approximate Normal random variables
- nested MLMC
- outline of numerical analysis
- numerical results
- conclusions / future work
- (a few extra slides)

# Monte Carlo

Monte Carlo simulation uses the average

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

to estimate  $\mathbb{E}[P]$ .

To achieve a RMS error of  $\varepsilon$  requires  $N \approx \varepsilon^{-2} V$  samples, where  $V = \mathbb{V}[P]$  is the variance.

If each sample costs  $C$  then the total cost is approximately  $\varepsilon^{-2} V C$

## Two-level Monte Carlo

If  $\tilde{P} \approx P$ , then since  $\mathbb{E}[P] = \mathbb{E}[\tilde{P}] + \mathbb{E}[P - \tilde{P}]$  we can instead use

$$N_0^{-1} \sum_{n=1}^{N_0} \tilde{P}^{(n)} + N_1^{-1} \sum_{n=1}^{N_1} (P^{(n)} - \tilde{P}^{(n)}).$$

The cost of this estimator is  $N_0 C_0 + N_1 C_1$ , and the variance is  $V_0/N_0 + V_1/N_1$ , where  $V_0 \equiv \mathbb{V}[\tilde{P}]$ ,  $V_1 \equiv \mathbb{V}[P - \tilde{P}]$ .

Minimising the cost subject to the same accuracy requirement gives the total cost

$$\varepsilon^{-2} (\sqrt{V_0 C_0} + \sqrt{V_1 C_1})^2.$$

If  $C_0 = 10^{-1} C$ ,  $C_1 = C$ ,  $V_0 = V$ ,  $V_1 = 10^{-3} V$  then the total cost is  $0.121 \varepsilon^{-2} V C$ , a factor 8 savings compared to the original Monte Carlo.

## Multilevel Monte Carlo

Given a sequence of increasingly accurate (and costly) approximations  $P_0, P_1, P_2, \dots \rightarrow P$ , for example from the approximation of an SDE using  $2^\ell$  timesteps on level  $\ell$ , then

$$\mathbb{E}[P_L] = \mathbb{E}[P_0] + \sum_{\ell=1}^L \mathbb{E}[P_\ell - P_{\ell-1}]$$

and so the MLMC estimate is

$$N_0^{-1} \sum_{n=1}^{N_0} P_0^{(n)} + \sum_{\ell=1}^L N_\ell^{-1} \sum_{n=1}^{N_\ell} (P_\ell^{(\ell,n)} - P_{\ell-1}^{(\ell,n)}),$$

and the total cost ends up being approximately

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

## Approximate random variables

In some applications, generating the random numbers can be a significant cost, especially with QMC when inverting the CDF.

e.g. Poisson distribution, increments of a Lévy process, non-central chi-squared distribution (CIR model)

Even with Normal random variables, cost of conversion from uniform r.v. to Normal is non-trivial for vector implementation.

This has led to previous research by Hefter, Mayer, Ritter, amongst others.

## Approximate random variables

Simplest example: Euler-Maruyama approximation of a scalar SDE:

$$\hat{X}_{t_{m+1}} = \hat{X}_{t_m} + a(\hat{X}_{t_m})h + b(\hat{X}_{t_m})\sqrt{h}Z_m$$

where  $Z_m$  is a unit Normal r.v. generated as

$$Z_m = \Phi^{-1}(U_m)$$

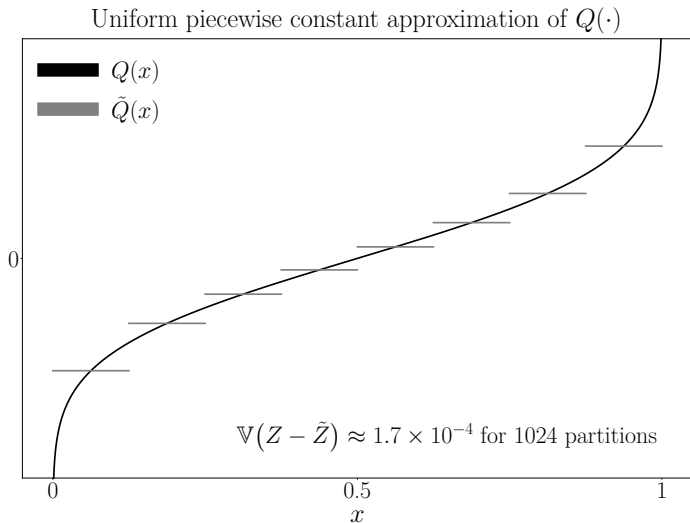
where  $U_m$  is a uniform  $(0, 1)$  r.v. and  $\Phi^{-1}$  is the inverse Normal CDF.

Suppose instead we use approximate Normals  $\tilde{Z}_m$  generated by

$$\tilde{Z}_m = \tilde{Q}(U_m)$$

where  $\tilde{Q}$  is an approximation to  $\Phi^{-1}$ .


# Approximate random variables



Good for scalar execution on CPUs – lookup in L1 cache

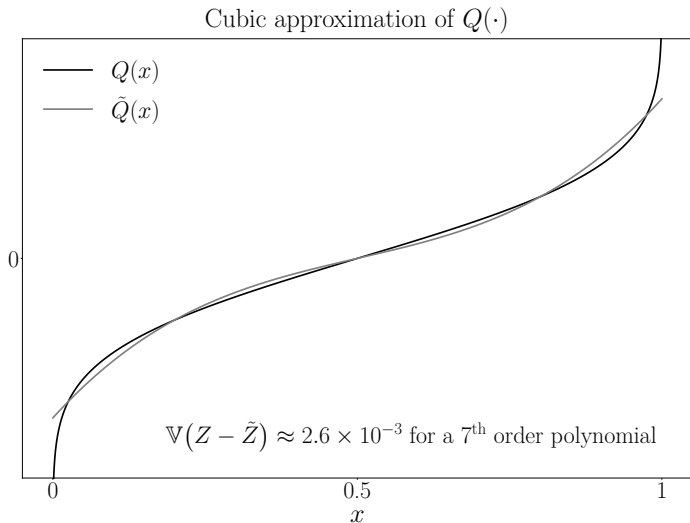


# Approximate random variables



Good for vector execution on CPUs – lookup within a vector

# Approximate random variables



Good for half-precision on GPUs – no lookup

## Approximate random variables

For an SDE approximation with a specified number of timesteps, a two-level MLMC approach gives

$$\mathbb{E}[\widehat{P}] = \mathbb{E}[\widetilde{P}] + \mathbb{E}[\widehat{P} - \widetilde{P}]$$

Further analysis proves that for  $P \equiv f(X_T)$  for a Lipschitz  $f(x)$

$$\mathbb{V}[\widehat{P} - \widetilde{P}] = O\left(\mathbb{V}[Z - \widetilde{Z}]\right)$$

so this can lead to significant savings if we have both

- $\mathbb{V}[Z - \widetilde{Z}] \ll 1$
- $\widetilde{Z}_m \equiv \widetilde{Q}(U_m)$  is much cheaper to evaluate than  $Z_m \equiv \Phi^{-1}(U_m)$

## Approximate random variables

How does this work in combination with standard timestepping MLMC?

Answer: nested MLMC

$$\begin{aligned}\mathbb{E}[P_L] &= \mathbb{E}[P_0] + \sum_{\ell=1}^L \mathbb{E}[P_\ell - P_{\ell-1}] \\ &= \mathbb{E}[\tilde{P}_0] + \mathbb{E}[P_0 - \tilde{P}_0] \\ &\quad + \sum_{\ell=1}^L \left\{ \mathbb{E}[\tilde{P}_\ell - \tilde{P}_{\ell-1}] + \mathbb{E} \left[ (P_\ell - P_{\ell-1}) - (\tilde{P}_\ell - \tilde{P}_{\ell-1}) \right] \right\}\end{aligned}$$

The pair  $(\tilde{P}_\ell, \tilde{P}_{\ell-1})$  are generated in the same way as  $(P_\ell, P_{\ell-1})$ , just replacing exact  $Z_m$  by approximate  $\tilde{Z}_m$ , for same underlying  $U_m$

# Approximate random variables

Numerical analysis – a work-in-progress

Path differences?

$$\text{(strong convergence)} \quad \widehat{X}_\ell - \widehat{X}_{\ell-1} \sim h^{1/2}$$

$$\text{(similar, but non-standard)} \quad \widetilde{X}_\ell - \widetilde{X}_{\ell-1} \sim h^{1/2}$$

$$\text{(similar, but non-standard)} \quad \widehat{X}_\ell - \widetilde{X}_\ell \sim \sqrt{\mathbb{V}[Z - \widetilde{Z}]}$$

$$\text{(tricky!)} \quad (\widehat{X}_\ell - \widehat{X}_{\ell-1}) - (\widetilde{X}_\ell - \widetilde{X}_{\ell-1}) \sim h^{1/2} \sqrt{\mathbb{V}[Z - \widetilde{Z}]}$$

# Approximate random variables

MLMC variances?

For smooth payoff functions  $f(x)$ , can prove

$$\mathbb{V}[(\hat{P}_\ell - \hat{P}_{\ell-1}) - (\tilde{P}_\ell - \tilde{P}_{\ell-1})] \sim h \mathbb{V}[Z - \tilde{Z}]$$


For continuous piecewise smooth functions (e.g. put/call functions) the analysis is much tougher. Numerical results suggest

$$\mathbb{V}[(\hat{P}_\ell - \hat{P}_{\ell-1}) - (\tilde{P}_\ell - \tilde{P}_{\ell-1})] \sim \min \left\{ h^{1/2} \mathbb{V}[Z - \tilde{Z}], h \sqrt{\mathbb{V}[Z - \tilde{Z}]} \right\}$$

but it seems tough to prove better than

$$\mathbb{V}[(\hat{P}_\ell - \hat{P}_{\ell-1}) - (\tilde{P}_\ell - \tilde{P}_{\ell-1})] \sim \min \left\{ \mathbb{V}[Z - \tilde{Z}], h \sqrt[3]{\mathbb{V}[Z - \tilde{Z}]} \right\}.$$

# Convergence between QMLMC levels



## Reduced precision arithmetic

Further computational savings can be achieved by using reduced precision arithmetic.

Previous research at TU Kaiserslautern (Korn, Ritter, Wehn and others) and Imperial College (Luk and others) has used FPGAs with complete control over the precision, but we prefer GPUs.

In the latest NVIDIA GPUs, half-precision fp16 is twice as fast as single precision fp32 (which is 2-8 times faster than double precision fp64).

Future Intel CPUs are likely to also support fp16 at twice the performance of fp32, and 4 times the performance of fp64.

In most cases, single precision is perfectly sufficient for calculating  $\hat{P}$ ; half precision can be used for  $\tilde{P}$ .

MC averaging should probably be done in double precision in both cases.



## Reduced precision arithmetic

Very important: to ensure the telescoping sum is respected, must ensure that **exactly** the same computations are performed for  $\tilde{P}$  whether on its own or in calculating  $\hat{P} - \tilde{P}$ .

The effect of half-precision arithmetic can be modelled as

$$\tilde{X}_{t_{m+1}} = \tilde{X}_{t_m} + a(\tilde{X}_{t_m})h + b(\tilde{X}_{t_m})\sqrt{h}\tilde{Z}_m + \delta\tilde{X}_{t_m}V_m$$

where  $\delta \approx 10^{-3}$  and the  $V_m$  are iid unit variance random variables.

Overall, this leads to an  $O(\delta^2/h)$  increase in the variance; if this increases it too much, the reduced precision should not be used.

## Conclusions / future work

- nested MLMC is similar to MIMC, but more general
- helpful in using approximate distributions and reduced precision
- offers significant computational savings in some situations
- in future, perform experiments with reduced precision, and extend to other distributions:
  - ▶ Poisson
  - ▶ binomial
  - ▶ non-central chi-squared (CIR process)
- also want to extend to MLQMC using quasi-uniform random numbers

## References

MBG. 'Multilevel Monte Carlo methods'. *Acta Numerica*, 24:259-328, 2015.

MBG, F.Y. Kuo, I.H. Sloan. 'Combining sparse grids, multilevel MC and QMC for elliptic PDEs with random coefficients'. *Monte Carlo and Quasi-Monte Carlo Methods 2016*, Springer, 2018.

MBG, M. Hefter, L. Mayer, K. Ritter. 'Random bit quadrature and approximation of distributions on Hilbert spaces'. *Foundations of Computational Mathematics*, 19(1):205-238, 2019.

MBG, M. Hefter, L. Mayer, K. Ritter. 'Random bit multilevel algorithms for stochastic differential equations'. *Journal of Complexity*, online January 2019.

# Generation of uniform random variables

Randomised QMC often generates points as

$$U^{(m,n)} = S^{(m)} \wedge R^{(n)}$$


where  $S^{(m)}$  is a set of Sobol points in  $[0, 1)^d$ ,  $R^{(n)}$  is a set of pseudo-random points in  $[0, 1)^d$ , and  $\wedge$  is a bitwise exclusive-OR operator.

The same approach works for  $S^{(m)}$  being pseudo-random points, and gives a point set  $\{U^{(m,n)}\}$  with pairwise independence, which is all that is needed for standard variance analysis.


This effectively eliminates the cost of uniform random number generation – two sets of 1000 random numbers give rise to  $10^6$  random numbers at an average cost of less than 2 operations each.

# MIMC compared to nested MLMC

MIMC: 4-way cancellation



nested MLMC: 2-way cancellation



## MIMC compared to nested MLMC

A key aspect of the standard MLMC for SDEs is that the Brownian increment for a timestep of  $2h$  is equal to the sum of Brownian increments for two timesteps of size  $h$ .

Another way of putting this is that if  $Z_1, Z_2$  are unit Normal r.v.'s, then

$$\sqrt{h} Z_1 + \sqrt{h} Z_2 = \sqrt{2h} Z_3$$

where  $Z_3$  is also a unit Normal r.v.

However, if  $\tilde{Z}_1, \tilde{Z}_2$  are approximate unit Normal random variables, and

$$\sqrt{h} \tilde{Z}_1 + \sqrt{h} \tilde{Z}_2 = \sqrt{2h} \tilde{Z}_3$$

then  $\tilde{Z}_3$  is an approximate Normal r.v. from a different distribution.

This is why we use the nested MLMC approach and not MIMC.