MLMC for multi-dimensional reflected Brownian diffusions

Mike Giles [1] Kavita Ramanan [2]

[1]Mathematical Institute, University of Oxford

[2]Division of Applied Mathematics, Brown University

MCQMC'16, Stanford University

August 15, 2016

Outline

- Multilevel Monte Carlo
- Multi-dimensional reflected diffusions
  - numerical discretisations
  - adaptive timestepping
  - numerical analysis
  - numerical results
- Conclusions

Multilevel Monte Carlo

MLMC 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 $\widehat{P}_\ell$ represents an approximation of some output P on level ℓ.

In simple SDE applications with uniform timestep $h_\ell = 2^{-\ell}\, h_0$, if the weak convergence is

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

and $\widehat{Y}_\ell$ is an unbiased estimator for $\mathbb{E}[\widehat{P}_\ell - \widehat{P}_{\ell-1}]$, based on N_ℓ samples, with variance

$$\mathbb{V}[\widehat{Y}_\ell] = O(N_\ell^{-1}\, 2^{-\beta\,\ell}),$$

and expected cost

$$\mathbb{E}[C_\ell] = O(N_\ell\; 2^{\gamma\,\ell}), \quad \ldots$$

Multilevel Monte Carlo

... then the finest level L and the number of samples N_ℓ on each level can be chosen to achieve an RMS error of ε at an expected cost

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

I always try to get $\beta > \gamma$, so the main cost comes from the coarsest levels – use of QMC can then give substantial additional benefits.

With $\beta > \gamma$, can also randomise levels to eliminate bias (Rhee & Glynn, Operations Research, 2015).

Multilevel Monte Carlo

The standard estimator for SDE applications is

$$\widehat{Y}_\ell = N_\ell^{-1} \sum_{n=0}^{N_\ell} \left(\widehat{P}_\ell(W^{(n)}) - \widehat{P}_{\ell-1}(W^{(n)}) \right)$$

using the same Brownian motion $W^{(n)}$ for the n^{th} sample on the fine and coarse levels.

However, there is some freedom in how we construct the coupling provided $\widehat{Y}_\ell$ is an unbiased estimator for $\mathbb{E}[\widehat{P}_\ell - \widehat{P}_{\ell-1}]$.

Have exploited this with an antithetic estimator for multi-dimensional SDEs which don't satisfy the commutativity condition.

(G, Szpruch: AAP 2014)

Multilevel Monte Carlo

Also, uniform timestepping is not required – it is fairly straightforward to implement MLMC using non-nested adaptive timestepping.

(G, Lester, Whittle: MCQMC14 proceedings)

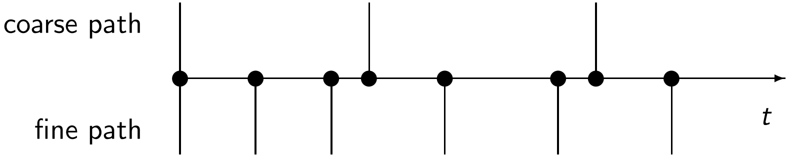

Also, interesting possibilities for applications with discontinuous output functionals.

Reflected diffusions

Motivating application comes from modelling of network queues

Reflected Brownian diffusion with constant volatility in a domain D has SDE

$$\mathrm{d}x_t = a(x_t)\,\mathrm{d}t + b\,\mathrm{d}W_t + \nu(x_t)\,\mathrm{d}L_t,$$

where L_t is a local time which increases when x_t is on the boundary ∂D.

$\nu(x)$ can be normal to the boundary (pointing inwards), but in some cases it is not and reflection from the boundary includes a tangential motion.

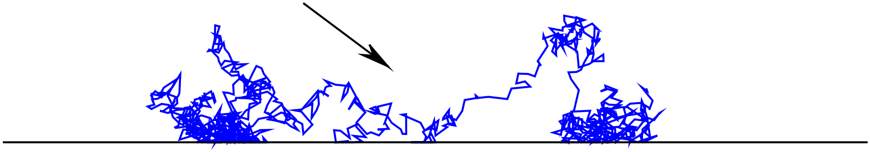

Reflected diffusions

A penalised version is

$$\begin{aligned} \mathrm{d}x_t &= a(x_t)\,\mathrm{d}t + b\,\mathrm{d}W_t + \nu(x_t)\,\mathrm{d}L_t, \\ \mathrm{d}L_t &= \lambda \max(0, -d(x_t))\,\mathrm{d}t, \qquad \lambda \gg 1 \end{aligned}$$

where $d(x)$ is signed distance to the boundary (negative means outside) and $\nu(x)$ is a smooth extension from the boundary into the exterior.

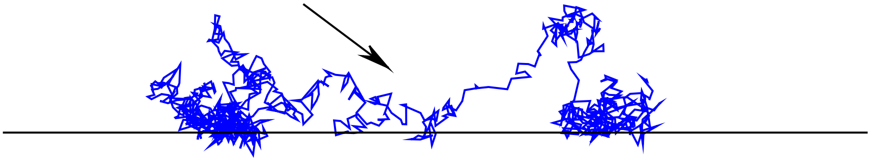

Reflected diffusions

When D is a polygonal domain, this generalises to

$$\mathrm{d}x_t = a(x_t)\,\mathrm{d}t + b\,\mathrm{d}W_t + \sum_{k=1}^{K} \nu_k(x_t)\,\mathrm{d}L_{k,t},$$

with a different ν_k and local time $L_{k,t}$ for each boundary face.

The corresponding penalised version is

$$\begin{aligned}
\mathrm{d}x_t &= a(x_t)\,\mathrm{d}t + b\,\mathrm{d}W_t + \sum_{k=1}^{K} \nu_k(x_t)\,\mathrm{d}L_{k,t}, \\
\mathrm{d}L_{k,t} &= \lambda\, \max(0, -d_k(x_t))\,\mathrm{d}t, \qquad \lambda \gg 1
\end{aligned}$$

where $d_k(x_t)$ is signed distance to the boundary face with a suitable extension.

Numerical approximations

3 different numerical treatments in literature:

- projection (Gobet, Słomiński): predictor step

$$\widehat{X}^{(p)} = \widehat{X}_{t_n} + a(\widehat{X}_{t_n}, t_n)\, h_n + b\, \Delta W_n,$$

  followed by correction step

$$\widehat{X}_{t_{n+1}} = \widehat{X}^{(p)} + \nu(\widehat{X}^{(p)})\ \Delta \widehat{L}_n,$$

  with $\Delta \widehat{L}_n > 0$ if needed to put $\widehat{X}_{t_{n+1}}$ on boundary

- reflection (Gobet): similar but with double the value for $\Delta \widehat{L}_n$
  – can give improved $O(h)$ weak convergence

- penalised (Słomiński): Euler-Maruyama approximation of penalised SDE with $\lambda = O(h^{-1})$, giving convergence as $h \to 0$

Numerical approximations

Concern:

- because b is constant, Euler-Maruyama method corresponds to first order Milstein scheme, suggesting an $O(h)$ strong error
- however, all three treatments of boundary reflection lead to a strong error which is $O(h^{1/2})$ – this is based primarily on empirical evidence, with only limited supporting theory
- if the output quantity of interest is Lipschitz with respect to the path then

$$\mathbb{V}\left[\widehat{P}-P\right] \;\leq\; \mathbb{E}\left[(\widehat{P}-P)^2\right] \;\leq\; c^2\,\mathbb{E}\left[\sup_{[0,T]}(\widehat{X}_t-X_t)^2\right]$$

  so the variance is $O(h)$

- OK, but not great – would like $O(h^\beta)$ with $\beta>1$ for $O(\varepsilon^{-2})$ MLMC complexity

Adaptive timesteps

Simple idea: use adaptive timestep based on distance from the boundary

- far away, use uniform timestep $h_\ell = 2^{-\ell}\, h_0$
- near the boundary, use uniform timestep $h_\ell = 2^{-2\ell}\, h_0$
- in between, define $h_\ell(x)$ to vary smoothly based on distance $d(x)$

What do we hope to achieve?

- strong error $O(2^{-\ell}) \implies$ MLMC variance is $O(2^{-2\ell})$
- computational cost per path $O(2^{\ell})$
- $\beta\!=\!2,\ \gamma\!=\!1$ in MLMC theorem $\implies$ complexity is $O(\varepsilon^{-2})$

Adaptive timesteps

In intermediate zone, want negligible probability of taking a single step and crossing the boundary.

Stochastic increment in Euler timestep is $b\,\Delta W$, so define h_ℓ so that

$$(\ell{+}3)\,\|b\|_2\,\sqrt{h_\ell} = d$$

Final 3-zone max-min definition of h_ℓ is

$$h_\ell = \max\left(2^{-2\ell} h_0, \min\left(2^{-\ell} h_0, (d/((\ell{+}3)\,\|b\|_2)^2\right)\right)$$

Balancing terms, gives

- boundary zone up to $d = O(2^{-\ell})$
- intermediate zone up to $d = O(2^{-\ell/2})$

Adaptive timesteps

Balancing terms, gives

- boundary zone up to $d \approx O(2^{-\ell})$
- intermediate zone up to $d \approx O(2^{-\ell/2})$

If $\rho(y,t)$, the density of paths at distance y from the boundary at time t, is uniformly bounded then the computational cost per unit time is approximately

$$\int_0^\infty \frac{\rho(y,t)\,\mathrm{d}y}{h_\ell(y)} \;\sim\; \underbrace{2^{2\ell}\times 2^{-\ell}}_{\textit{boundary}} + \underbrace{\int_{O(2^{-\ell})}^{O(2^{-\ell/2})} \frac{\mathrm{d}y}{y^2}}_{\textit{intermediate}} + \underbrace{2^{\ell}\times 1}_{\textit{interior}} \;\approx\; O(2^\ell)$$

so we get similar cost contributions from all 3 zones.

Numerical analysis

Theorem (Computational cost)

If

- *the density $\rho(y,t)$ for the SDE paths at distance y from the boundary is uniformly bounded*
- *the numerical discretisation with the adaptive timestep has strong convergence $O(2^{-\ell})$*

then the computational cost is $o(2^{(1+\delta)\ell})$ for any $0<\delta\ll 1$.

The second condition is needed to bound the difference between the distributions of the paths and their numerical approximations.

Numerical analysis

Theorem (Strong convergence)

If

- *the drift a is constant*
- *a uniform timestep discretisation has $O(h^{1/2})$ strong convergence*
- *the adaptive timestep h_ℓ is rounded to the nearest multiple of the boundary zone timestep*

then the strong convergence is $O(2^{-\ell})$

The proof is based on a comparison with a discretisation using the uniform boundary zone timestep:

- adaptive numerical discretisation is exact when boundary not crossed
- almost zero probability of crossing the boundary unless in the boundary zone using the uniform timestep

Numerical results

Simple test case:

- 3D Brownian motion in a unit ball
- normal reflection at the boundary
- $x_0 = 0$
- aim is to estimate $\mathbb{E}[\|x\|_2^2]$ at time $t\!=\!1$.
- implemented with both projection and penalisation schemes

Numerical results

Projection method:

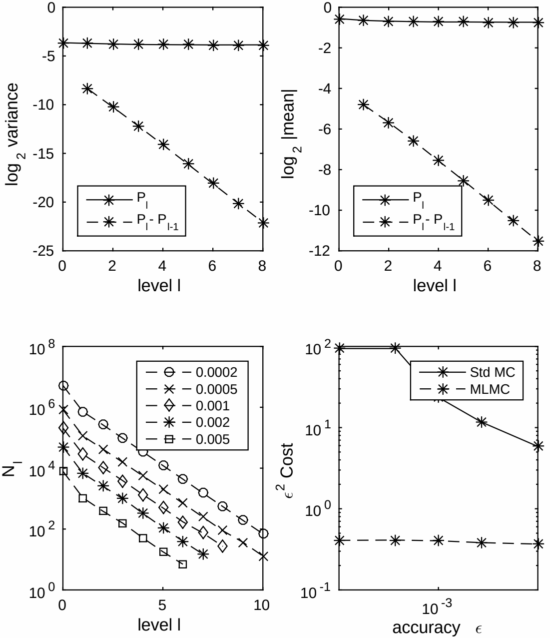

Numerical results

Penalisation method:

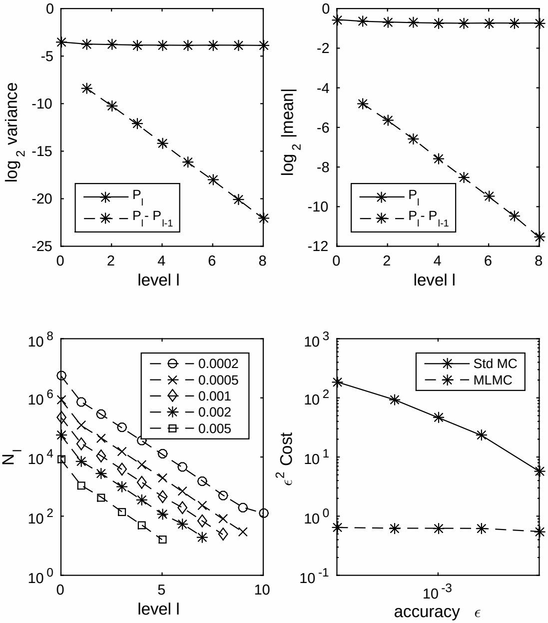

Conclusions

Initial research is promising:

- natural use of localised adaptive timestepping to reduce errors
- $O(\varepsilon^{-2})$ complexity for ε RMS error
- significant progress with numerical analysis
- numerical results are also encouraging

Future challenges:

- prove that for constant drift a and timestep h, the strong error is $O(h^{1/2})$ for reflected diffusions with oblique reflections, preferably for generalised penalisation method for polygonal boundaries
- extend analysis to include errors in local time
- extend analysis to general drift and adaptive timesteps

Webpages:

http://people.maths.ox.ac.uk/gilesm/

http://people.maths.ox.ac.uk/gilesm/mlmc_community.html