Stochastic Simulation: Lecture 12

Christoph Reisinger

Oxford University Mathematical Institute

Modified from earlier slides by Prof. Mike Giles.

PDEs with Uncertainty

Looking at the history of numerical methods for PDEs, the first steps were about improving the modelling:

- 1D $\rightarrow$ 2D $\rightarrow$ 3D
- steady $\rightarrow$ unsteady
- laminar flow $\rightarrow$ turbulence modelling $\rightarrow$ large eddy simulation $\rightarrow$ direct Navier-Stokes
- simple geometries (e.g. a wing) $\rightarrow$ complex geometries (e.g. an aircraft in landing configuration)
- adding new features such as combustion, coupling to structural / thermal analyses, etc.

. . . and then engineering switched from analysis to design.

PDEs with Uncertainty

The big move now is towards handling uncertainty:

- uncertainty in modelling parameters
- uncertainty in geometry
- uncertainty in initial conditions
- uncertainty in spatially-varying material properties
- inclusion of stochastic source terms

Engineering wants to move to "robust design" taking into account the effects of uncertainty.

Other areas want to move into Bayesian inference, starting with an *a priori* distribution for the uncertainty, and then using data to derive an improved *a posteriori* distribution.

PDEs with Uncertainty

Examples:

- Long-term climate modelling:

  Lots of sources of uncertainty including the effects of aerosols, clouds, carbon cycle, ocean circulation (`http://climate.nasa.gov/uncertainties`)

- Short-range weather prediction

  Considerable uncertainty in the initial data due to limited measurements

PDEs with Uncertainty

- Engineering analysis

  Perhaps the biggest uncertainty is geometric due to manufacturing tolerances

- Nuclear waste repository and oil reservoir modelling

  Considerable uncertainty about porosity of rock

- Finance

  Stochastic forcing due to market behaviour

PDEs with Uncertainty

Motivated by the modelling of oil reservoirs and groundwater contamination, there has been a lot of MLMC/MLQMC research on an elliptic SPDE coming from Darcy's law:

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

where the permeability $\kappa(x)$ is uncertain due to lack of knowledge.

A log-Normal model uses

$$\log \kappa(x) = u_0(x) + u(x)$$

where $u_0(x)$ is the mean, and $u(x)$ is Normally distributed pointwise, with zero mean and spatial covariance

$$\mathbb{E}[u(x_1) \, u(x_2)] = K(x_1, x_2).$$

Karhunen-Loève expansion

Under "suitable" conditions, the linear operator

$$u(x) \longrightarrow \int_D K(x,y)\, u(y)\, \mathrm{d}y$$

has a complete set of orthonormal eigenfunctions $e_j(x)$, and associated eigenvalues $\lambda_j > 0$ such that

$$\int_D K(x,y)\, e_j(y)\, \mathrm{d}y = \lambda_j\, e_j(x)$$

Hence, $u(x)$ can be expressed as

$$u(x) = \sum_j Z_j\, e_j(x)$$

where

$$Z_j = \int_D e_j(x)\, u(x)\, \mathrm{d}x$$

Karhunen-Loève expansion

It follows that Z_j is Normally distributed with zero mean, and

$$\begin{aligned}
\mathbb{E}[Z_j\, Z_k] &= \int_D \int_D e_j(x)\; \mathbb{E}[u(x)\, u(y)]\, e_k(y)\; \mathrm{d}x\, \mathrm{d}y \\
&= \int_D \int_D e_j(x)\; K(x,y)\, e_k(y)\; \mathrm{d}x\, \mathrm{d}y \\
&= \lambda_k \int_D e_j(x)\, e_k(x)\; \mathrm{d}x \\
&= \lambda_k \delta_{j,k}
\end{aligned}$$

Thus the Z_j are independent, and have variance λ_j.

Re-scaling, we can make the Z_j have unit variance and define $u(x)$ as

$$u(x) = \sum_j \sqrt{\lambda_j}\, Z_j\, e_j(x)$$

which is the Karhunen-Loève expansion.

Karhunen-Loève expansion

Note that this is simply the continuous generalisation of the PCA factorisation in which we ended up with

$$u = U \Lambda^{1/2} Z \equiv \sum_j \sqrt{\lambda_j}\, Z_j\, U_j$$

which U_j, λ_j being the eigenvectors and eigenvalues of the covariance matrix Σ.

They both have a very similar computational drawback: if there are M grid points in the domain D, then PCA has M eigenmodes so the cost per sample is $O(M^2)$, in addition to an $O(M^3)$ setup cost and an $O(M^2)$ memory requirement.

The K-L expansion is even worse as there are an infinite number of eigenmodes, but in both cases the cost can be reduced by truncating the expansion – keeping only the leading terms. Even so, it can still be the dominant cost in a calculation, much more than the solution of the elliptic PDE!

Elliptic SPDE

Often the covariance function is assumed to be stationary – i.e. a function of the separation $x_1 - x_2$.

"Exponential" covariance:

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

"Gaussian" covariance:

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

Typically in real applications σ is large, and λ is small – both of these mean that PDE methods such as stochastic collocation are very expensive

Elliptic SPDE

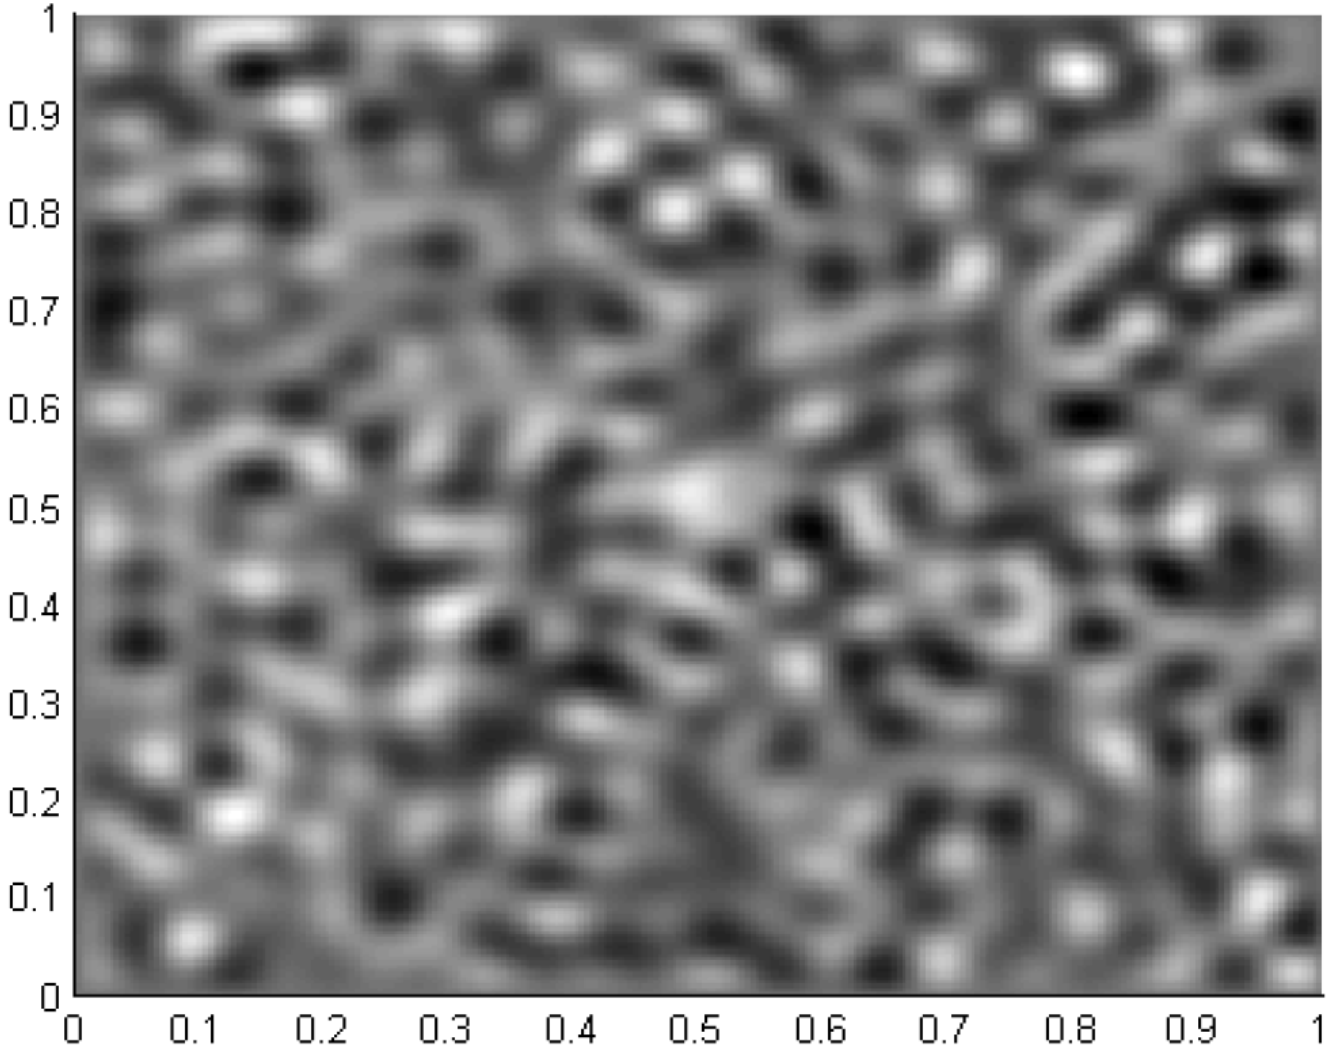

A typical realisation of κ for exponential covariance with $\lambda = 0.01,\ \sigma = 1$.

Elliptic SPDE

Decay of 1D eigenvalues

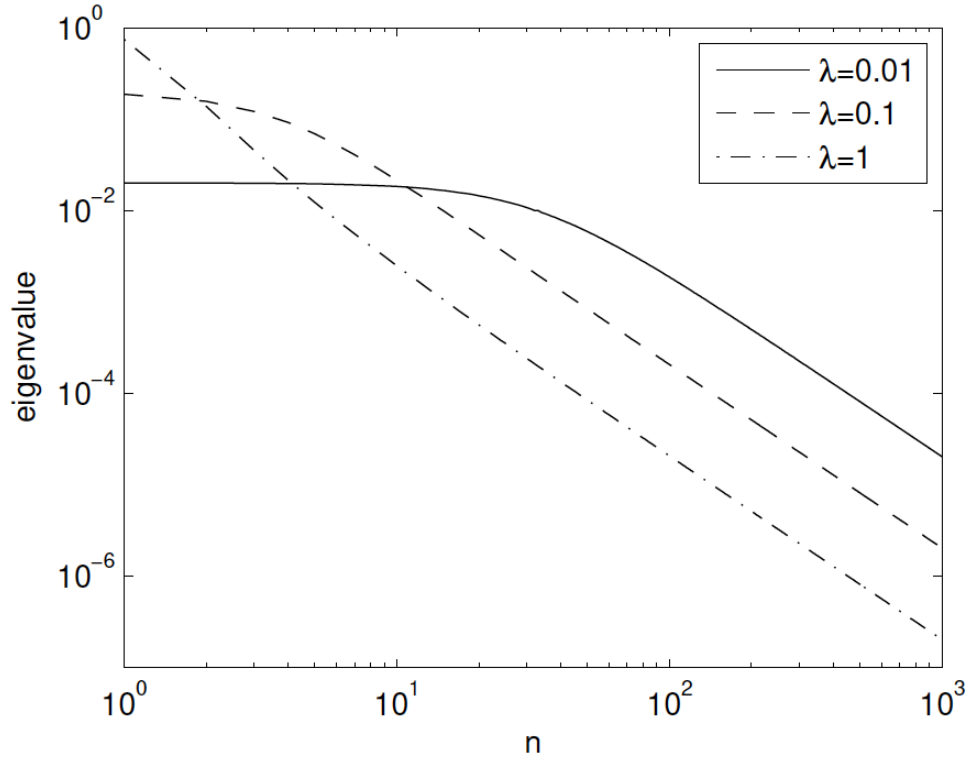

When $\lambda = 1$, can use a low-dimensional approximation, but it's impractical for smaller λ.

SPDEs with multilevel

The approach is very simple, in principle:

- use a sequence of grids of increasing resolution in space (and time)
- as with SDEs, determine the optimal allocation of computational effort on the different levels
- the savings can be much greater because the cost goes up more rapidly with level

MLMC Theorem

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 a computational cost C with bound

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

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

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

- hence, $\alpha = 2,\ \beta = 4,\ \gamma = 3$
- cost is $O(\varepsilon^{-2})$ to obtain ε RMS accuracy
- in comparison, cost is $O(\varepsilon^{-3/2})$ for a single calculation with ε accuracy

Elliptic SPDE

Some early 2D experiments performed by Cliffe, Giles, Scheichl and Teckentrup (2013):

- 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
- these experiments used a direct solver for simplicity, but later work used an efficient AMG multigrid solver with a cost roughly proportional to the total number of grid points

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: $\left[-\int k\,\frac{\partial p}{\partial x_1}\,\mathrm{d}x_2\right]_{x_1=1}$

Correlation length: $\lambda = 0.2$

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

Finest grid: $h = 1/128$

Karhunen-Loève truncation: used the leading 4000 modes

Cost taken to be proportional to number of nodes

2D Results

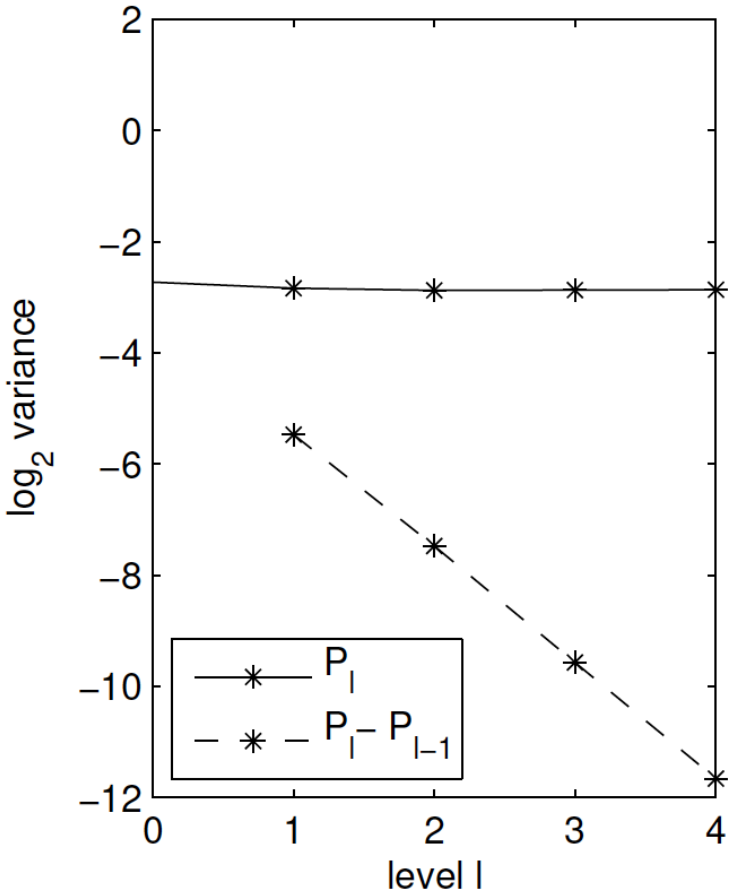

$$\mathbb{V}[\widehat{P}_\ell - \widehat{P}_{\ell-1}] \sim h_\ell^2$$

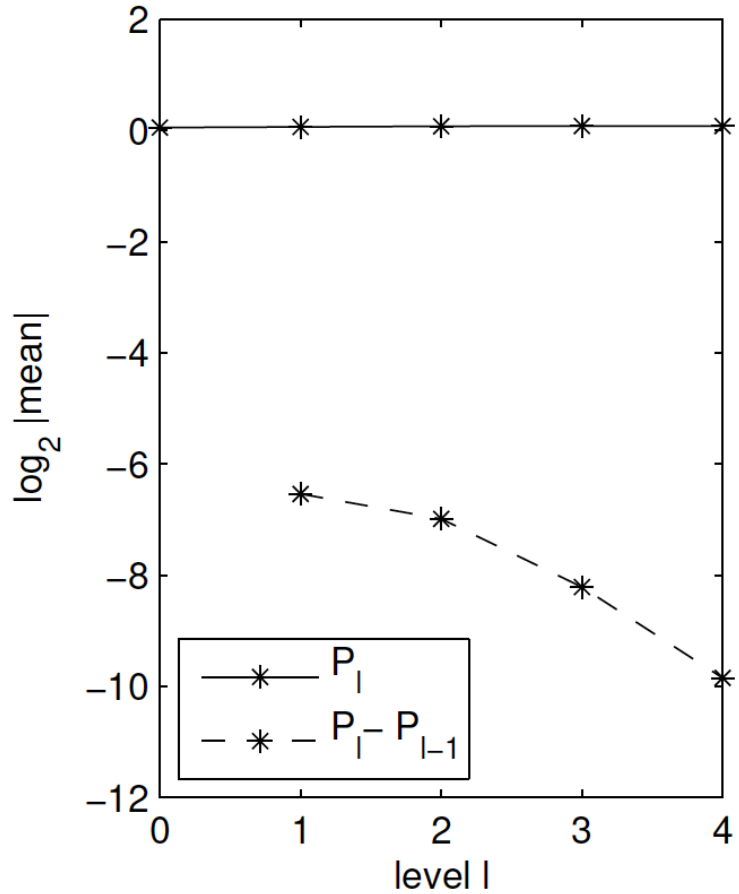

$$\mathbb{E}[\widehat{P}_\ell - \widehat{P}_{\ell-1}] \sim h_\ell^2$$

2D Results

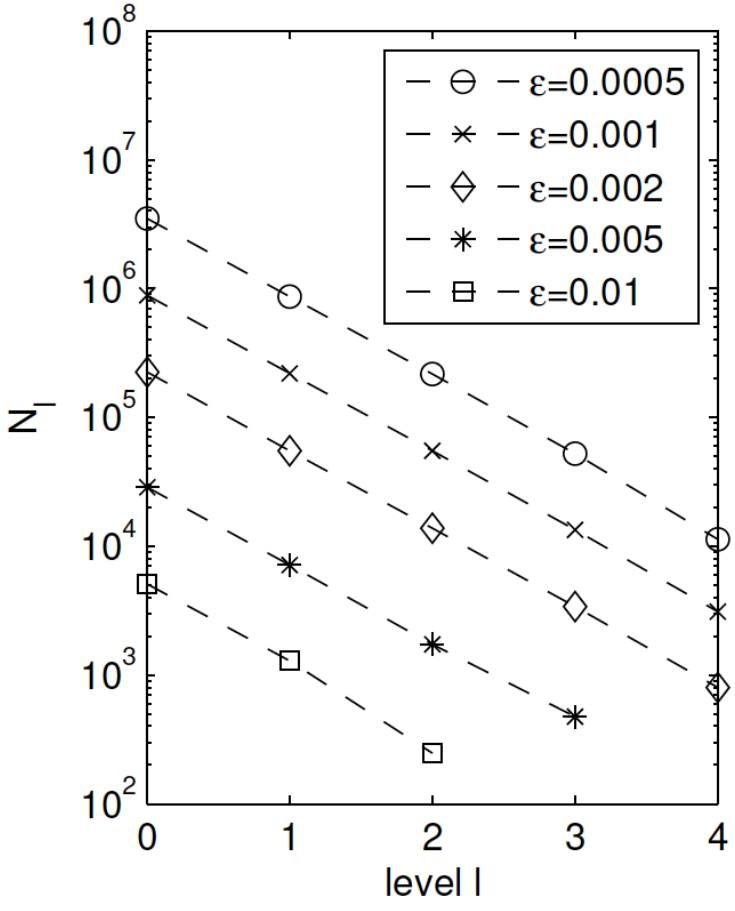
10^8
10^7
10^6
10^5
10^4
10^3
10^2
N_l
0
1
2
3
4
level l
ε=0.0005
ε=0.001
ε=0.002
ε=0.005
ε=0.01

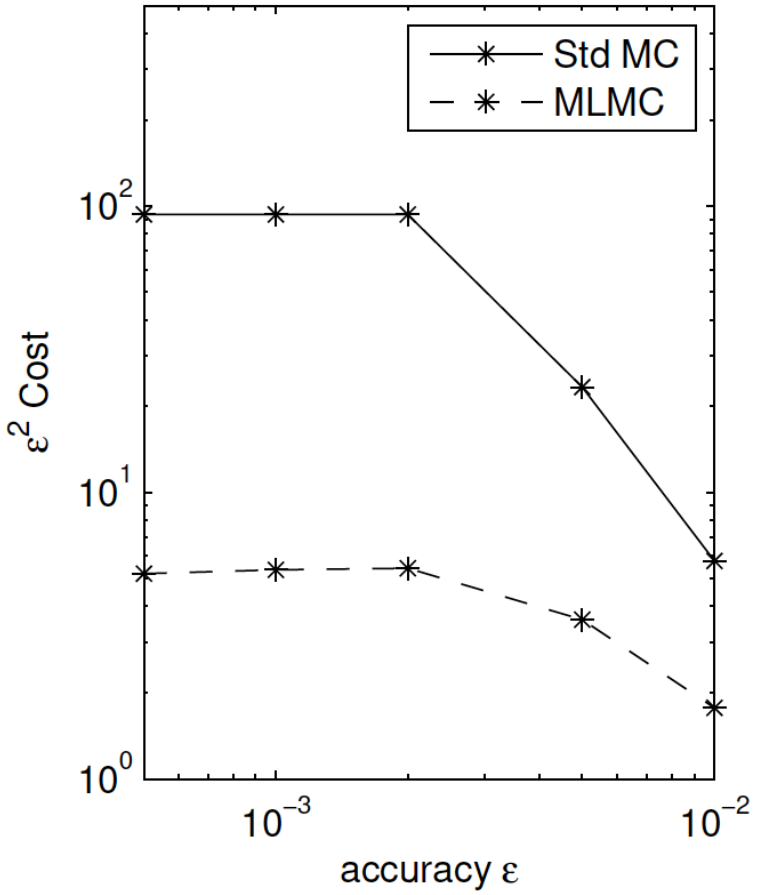
10^2
10^1
10^0
ε^2 Cost
10^{-3}
10^{-2}
accuracy ε
Std MC
MLMC

Complexity analysis

Relating things back to the MLMC theorem:

$$\begin{aligned} \mathbb{E}[\widehat{P}_\ell - P] &\sim 2^{-2\ell} \implies \alpha = 2 \\ V_\ell &\sim 2^{-2\ell} \implies \beta = 2 \\ C_\ell &\sim 2^{d\ell} \implies \gamma = d \quad \text{(dimension of PDE)} \end{aligned}$$

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	$\varepsilon^{-3.5}$	$\varepsilon^{-2.5}$

SPDEs

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

Parabolic SPDE

Parabolic SPDE arises in [filtering and] credit modelling (Giles & Reisinger, 2012)

$$\mathrm{d}p = -\mu \frac{\partial p}{\partial x}\,\mathrm{d}t + \frac{1}{2}\frac{\partial^2 p}{\partial x^2}\,\mathrm{d}t + \sqrt{\rho}\frac{\partial p}{\partial x}\,\mathrm{d}W$$

with absorbing boundary $p(0,t) = 0$

- derived in limit as number of firms $\longrightarrow \infty$
- x is distance to default
- $p(x,t)$ is probability density function
- $\mathrm{d}W$ term corresponds to systemic risk
- $\partial^2 p/\partial x^2$ comes from idiosyncratic risk

Parabolic SPDE

- numerical discretisation combines Milstein time-marching with central difference approximations
- coarsest level of approximation uses 1 timestep per quarter, and 10 spatial points
- each finer level uses four times as many timesteps, and twice as many spatial points – ratio is due to numerical stability constraints
- mean-square stability theory, with and without absorbing boundary
- computational cost $C_\ell \propto 8^\ell$
- numerical results suggest variance $V_\ell \propto 8^{-\ell}$
- can prove $V_\ell \propto 16^{-\ell}$ when no absorbing boundary

Parabolic SPDE

Fractional loss on equity tranche of a 5-year CDO:

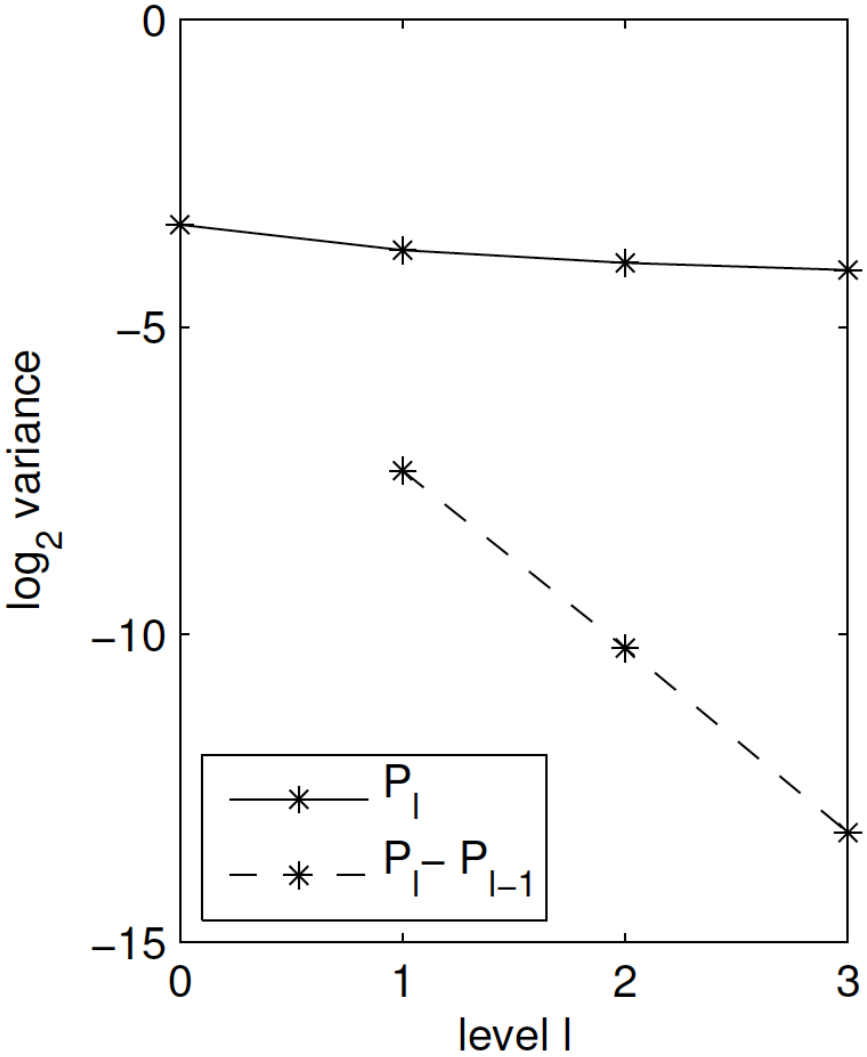

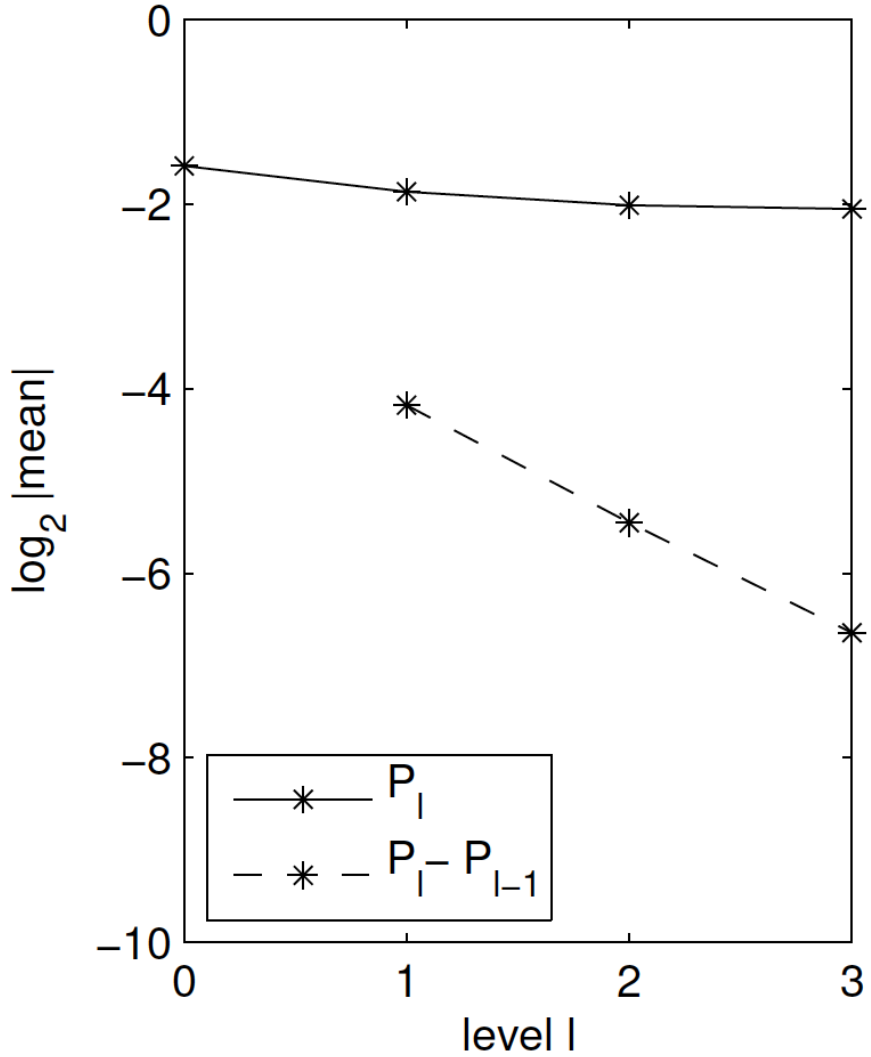

Parabolic SPDE

Fractional loss on equity tranche of a 5-year CDO:

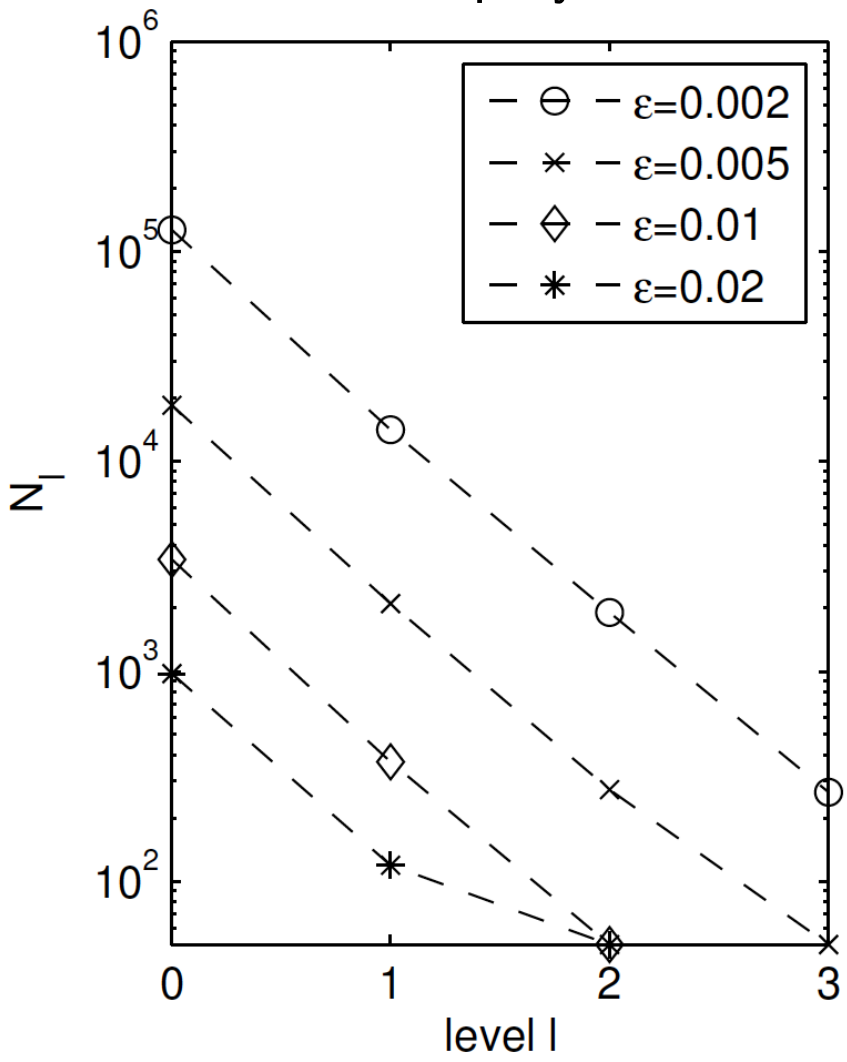

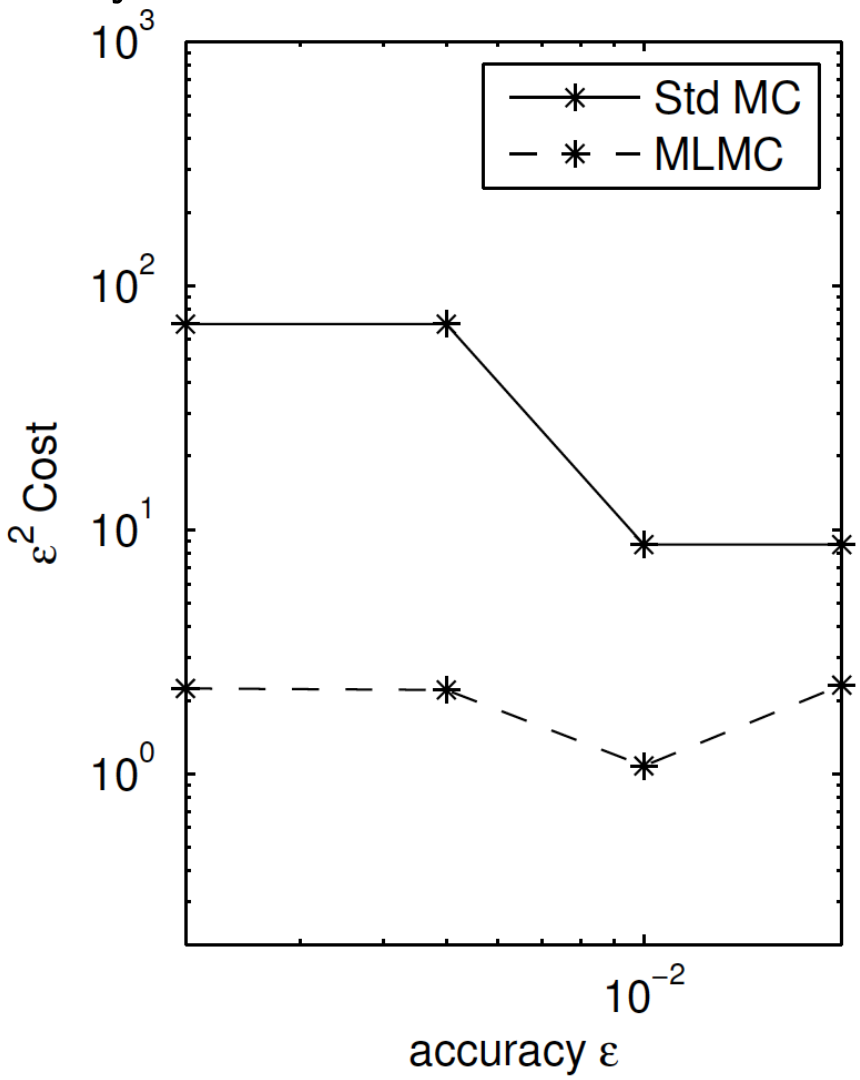

Parabolic SPDE

Milstein and central difference discretisation leads to

$$\begin{aligned} v_j^{n+1} &= v_j^n - \frac{\mu\, k + \sqrt{\rho\, k}\, Z_n}{2h}\left(v_{j+1}^n - v_{j-1}^n\right) \\ &\quad + \frac{(1-\rho)\, k + \rho\, k\, Z_n^2}{2h^2}\left(v_{j+1}^n - 2v_j^n + v_{j-1}^n\right) \end{aligned}$$

where $Z_n \sim N(0,1)$.

Considering a Fourier mode

$$v_j^n = g_n \exp(ij\theta), \quad |\theta| \leq \pi$$

leads to ...

Parabolic SPDE

$$g_{n+1} = \left(a(\theta) + b(\theta)\, Z_n + c(\theta)\, Z_n^2\right) g_n,$$

where

$$\begin{aligned}
a(\theta) &= 1 - \frac{i\,\mu\,k}{h}\,\sin\theta - \frac{2\,(1-\rho)\,k}{h^2}\,\sin^2\tfrac{\theta}{2}, \\
b(\theta) &= -\frac{i\sqrt{\rho\,k}}{h}\,\sin\theta, \\
c(\theta) &= -\frac{2\,\rho\,k}{h^2}\,\sin^2\tfrac{\theta}{2}.
\end{aligned}$$

Parabolic SPDE

Following the approach of mean-square stability analysis (e.g. see Higham)

$$\begin{aligned}\mathbb{E}[\,|g_{n+1}|^2] &= \mathbb{E}\left[(a + b\,Z_n + c\,Z_n^2)(a^* + b^* Z_n + c^* Z_n^2)\;|g_n|^2\right] \\ &= \left(\,|a+c|^2 + |b|^2 + 2|c|^2\,\right)\;\mathbb{E}\left[\,|g_n|^2\right]\end{aligned}$$

so stability requires $|a+c|^2 + |b|^2 + 2|c|^2 \;\leq\; 1$ for all θ, which leads to a timestep stability limit:

$$\begin{aligned}\mu^2 k &\leq 1 - \rho, \\ \frac{k}{h^2} &\leq (1 + 2\rho^2)^{-1}.\end{aligned}$$

Additional analysis extends this to include the effect of boundary conditions.

Parabolic SPDE

This can be extended to finite domains using matrix stability analysis, writing the discrete equations as

$$V_{n+1} = (A + B\,Z_n + C\,Z_n^2)\;V_n, \quad \text{where}$$

$$A = I - \frac{\mu\,k}{2h}\,D_1 + \frac{(1-\rho)\,k}{2h^2}\,D_2, \quad B = -\frac{\sqrt{\rho\,k}}{2h}\,D_1, \quad C = \frac{\rho\,k}{2h^2}\,D_2,$$

and D_1 and D_2 look like

$$D_1 = \begin{pmatrix} 0 & 1 & & \\ -1 & 0 & 1 & \\ & -1 & 0 & 1 \\ & & -1 & 0 \end{pmatrix}, \quad D_2 = \begin{pmatrix} -2 & 1 & & \\ 1 & -2 & 1 & \\ & 1 & -2 & 1 \\ & & 1 & -2 \end{pmatrix}.$$

Parabolic SPDE

$$\begin{aligned}\mathbb{E}[\,V_{n+1}^T V_{n+1}] &= \mathbb{E}\left[V_n^T (A^T + B^T Z_n + C^T Z_n^2)(A + B\,Z_n + C\,Z_n^2)\;V_n\right] \\ &= \mathbb{E}\left[V_n^T \left((A+C)^T(A+C) + B^T B + 2\,C^T C\right) V_n\right]\end{aligned}$$

D_1 is anti-symmetric and D_2 is symmetric, and

$$D_1 D_2 - D_2 D_1 = E_1 - E_2, \quad D_1^2 = D_3 + E_1 + E_2$$

where D_3 looks like

$$D_3 = \begin{pmatrix} -3 & 0 & 1 & \\ 0 & -2 & 0 & 1 \\ 1 & 0 & -2 & 0 \\ & 1 & 0 & -3 \end{pmatrix},$$

Parabolic SPDE

and E_1 and E_2 are zero apart from one corner element,

$$E_1 = \begin{pmatrix} 2 & & \\ & & \\ & & \end{pmatrix}, \qquad E_2 = \begin{pmatrix} & & \\ & & \\ & & 2 \end{pmatrix}$$

This leads to

$$\begin{aligned} &\mathbb{E}\left[V_n^T \left((A+C)^T(A+C) + B^T B + 2\,C^T C\right) V_n\right] \\ &\quad = \quad \mathbb{E}\left[V_n^T M V_n\right] - (e_1+e_2)\,\mathbb{E}[(v_1^n)^2] - (e_1-e_2)\,\mathbb{E}[(v_{J-1}^n)^2], \end{aligned}$$

where e_1 and e_2 are scalars and

$$M = I - \frac{k}{h^2} D_2 + \frac{k^2}{4\,h^4} D_2^2 - \left(\frac{\rho k}{4\,h^2} + \frac{\mu^2 k^2}{4\,h^2}\right) D_3.$$

Parabolic SPDE

It can be verified that the m^{th} eigenvector of M is a Fourier mode and the associated eigenvalue is

$$|a(\theta_m)+c(\theta_m)|^2 + |b(\theta_m)|^2 + 2|c(\theta_m)|^2$$

where $a(\theta), b(\theta), c(\theta)$ are the same functions as before.

In the limit $h, k/h \to 0$, $e_1 \pm e_2 > 0$, and therefore the Fourier stability condition

$$\sup_{\theta} \left\{ |a(\theta)+c(\theta)|^2 + |b(\theta)|^2 + 2|c(\theta)|^2 \right\} \leq 1$$

is also a sufficient condition for mean-square matrix stability.

References

K.A. Cliffe, MBG, R. Scheichl, A.L. Teckentrup. "Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients". Computing and Visualization in Science, 14(1):3-15, 2011.

A. Barth, Ch. Schwab, N. Zollinger. "Multi-level Monte Carlo finite element method for elliptic PDEs with stochastic coefficients". Numerische Mathematik, 119(1):123-161, 2011.

M. Croci, MBG, M.E. Rognes, P.E. Farrell. "Efficient white noise sampling and coupling for multilevel Monte Carlo with non-nested meshes". SIAM/ASA Journal on Uncertainty Quantification, 6(4):1454-1474, 2018.

N. Hale, N.J. Higham, L.N. Trefethen. "Computing $A^\alpha, \log(A)$, and related matrix functions by contour integrals". SIAM Journal on Numerical Analysis, 46(5):2505-2523, 2008.

Key references

D.J. Higham, "Mean-square and asymptotic stability of the stochastic theta method", SIAM Journal of Numerical Analysis, 38(3):753-769, 2000.

MBG, C. Reisinger. "Stochastic finite differences and multilevel Monte Carlo for a class of SPDEs in finance". SIAM Journal of Financial Mathematics, 3(1):572-592, 2012.

F. Vidal-Codina, N.-C. Nguyen, MBG, J. Peraire. 'A model and variance reduction method for computing statistical outputs of stochastic elliptic partial differential equations'. Journal of Computational Physics, 297:700-720, 2015.

F. Vidal-Codina, N.-C. Nguyen, MBG, J. Peraire. "An empirical interpolation and model-variance reduction method for computing statistical outputs of parametrized stochastic partial differential equations". SIAM/ASA Journal on Uncertainty Quantification, 4(1):244-265, 2016

Appendix: Alternative ways to generate Gaussian fields

a) Cholesky factorisation of covariance matrix

Same bad points as PCA factorisation: $O(M^3)$ setup cost, $O(M^2)$ memory requirement, $O(M^2)$ cost per sample

b) use of H-matrices (H stands for Hierarchical)

Promising new research direction:

`https://en.wikipedia.org/wiki/Hierarchical_matrix`

J. Dölz, H. Harbrecht, and C. Schwab, Covariance regularity and H-matrix approximation for rough random fields, Numerische Mathematik, 135:1045-1071, 2017.

Alternative ways to generate Gaussian fields

c) circulant embedding

In 1D, for a uniform grid and a stationary covariance function, the covariance matrix looks like:

$$\begin{pmatrix} a_0 & a_1 & a_2 & a_3 & a_4 & a_5 \\ a_1 & a_0 & a_1 & a_2 & a_3 & a_4 \\ a_2 & a_1 & a_0 & a_1 & a_2 & a_3 \\ a_3 & a_2 & a_1 & a_0 & a_1 & a_2 \\ a_4 & a_3 & a_2 & a_1 & a_0 & a_1 \\ a_5 & a_4 & a_3 & a_2 & a_1 & a_0 \end{pmatrix}$$

which can be inserted into a larger circulant matrix ...

Alternative ways to generate Gaussian fields

$$A = \begin{pmatrix} a_0 & a_1 & a_2 & a_3 & a_4 & a_5 & a_4 & a_3 & a_2 & a_1 \\ a_1 & a_0 & a_1 & a_2 & a_3 & a_4 & a_5 & a_4 & a_3 & a_2 \\ a_2 & a_1 & a_0 & a_1 & a_2 & a_3 & a_4 & a_5 & a_4 & a_3 \\ a_3 & a_2 & a_1 & a_0 & a_1 & a_2 & a_3 & a_4 & a_5 & a_4 \\ a_4 & a_3 & a_2 & a_1 & a_0 & a_1 & a_2 & a_3 & a_4 & a_5 \\ a_5 & a_4 & a_3 & a_2 & a_1 & a_0 & a_1 & a_2 & a_3 & a_4 \\ a_4 & a_5 & a_4 & a_3 & a_2 & a_1 & a_0 & a_1 & a_2 & a_3 \\ a_3 & a_4 & a_5 & a_4 & a_3 & a_2 & a_1 & a_0 & a_1 & a_2 \\ a_2 & a_3 & a_4 & a_5 & a_4 & a_3 & a_2 & a_1 & a_0 & a_1 \\ a_1 & a_2 & a_3 & a_4 & a_5 & a_4 & a_3 & a_2 & a_1 & a_0 \end{pmatrix}$$

in which each row (and column) is the same as the previous, but rotated by 1 position

Hence the eigenvectors are Fourier modes, and the corresponding eigenvalues come from a Fourier transform of the first row.

Alternative ways to generate Gaussian fields

This give

$$A = F \Lambda F^T = L L^T$$

where

$$L = F \Lambda^{1/2}$$

and therefore

$$L Z = \sum_j \sqrt{\lambda_j} Z_j F_j$$

which can be computed in $O(M \log M)$ cost.

There are some minor technical difficulties (making sure the embedded matrix has positive eigenvalues), but it extends to 2D and 3D and works well when doing MLMC on regular grids.

Alternative ways to generate Gaussian fields

d) spatial white noise

If $\dot{W}$ is spatial white noise. then the solution of

$$(I - \kappa^{-2}\nabla^2)^k u = \dot{W}$$

is a Matérn field with covariance of the form

$$K(x, y) = \frac{\sigma^2}{2^{\nu-1}\Gamma(\nu)}(\kappa r)^\nu \mathcal{K}_\nu(\kappa r), \quad r = \|x - y\|_2,$$

where $\nu = 2k - d/2$, and $\mathcal{K}_\nu$ is a Bessel function of the second kind – the Matérn class includes the exponential and Gaussian cases mentioned previously.

Alternative ways to generate Gaussian fields

What is white noise $\dot{W}$?

It is a generalised stochastic field defined by its effect on L_2-integrable test functions ϕ_j so that

$$\langle \dot{W}, \phi_j \rangle \sim N\left(0, \int_D \phi_j^2 \, \mathrm{d}x\right)$$

and

$$\mathbb{E}\left[\langle \dot{W}, \phi_j \rangle \, \langle \dot{W}, \phi_k \rangle\right] = \int_D \phi_j \, \phi_k \, \mathrm{d}x$$

If a domain D is split up into a number of disjoint pieces D_1, D_2, D_3, ..., then $\dot{W}$ can be decomposed into the sum of its restrictions onto each of those pieces, and the effects of each are independent.

Hence, can independently simulate the effect of each, and then sum them up.

Alternative ways to generate Gaussian fields

In a MLMC setting, working with coarse and fine grids composed of triangles (2D) or tetrahedra (3D), can create a finer supermesh of triangles/tetrahedra such that each new cell has a non-zero intersection with one and only one coarse and fine cell.

If $\dot{W}\Big|_{\Delta}$ is the restriction of $\dot{W}$ to this cell, then we can create a small covariance matrix for the test functions which are non-zero on this cell:

$$A_{\Delta} \equiv \mathbb{E}\left[\langle \dot{W}\Big|_{\Delta} \phi_j \rangle \, \langle \dot{W}\Big|_{\Delta} \phi_k \rangle\right] = \int_{\Delta} \phi_j \, \phi_k \, \mathrm{d}x$$

A_{Δ} is small, so can use Cholesky factorisation to generate required samples of $\langle \dot{W}\Big|_{\Delta} \phi_j \rangle$ on both grids.

Alternative ways to generate Gaussian fields

Additional complication: for some values of the Matérn parameter ν, need to solve

$$(I - \kappa^{-2}\nabla^2)^k u = \dot{W}$$

for non-integer values of k.

How do we do that? What does it even mean?

We rely on a complex contour integral representation of generalised matrix functions:

$$f(A) = \frac{1}{2\pi \mathrm{i}} \int_\Gamma f(z)\,(zI - A)^{-1}\,\mathrm{d}z$$

and approximate the integral.

Final comments

- PDEs with random inputs / boundary data have been well explored
- PDEs with random coefficients have also been well explored – growing body of literature on numerical analysis too
- Stochastic PDEs with white noise or Brownian noise inputs have received much less attention, so still more to be done?

  I'm not even sure how much is proven concerning wellposedness and numerical analysis