Domain Decomposition methods for the Karhunen-Loève decomposition and stochastic elliptic PDEs

Olivier Le Maître[1], Omar Knio[2,3]
Paul Mycek[3,4] and Andres Contreras[3]

[1] LIMSI CNRS, Orsay, France
[2,3] KAUST Saudi-Arabia
[3] Duke University
[3] Cerfacs

Séminaire MaNu, LJLL UPMC

High-Performance Computing - Parallel Computing

Context :

- Development of massively parallel machines
- Solution of large scale PDEs' based problem
- New architectures $\longrightarrow$ evolution of strategies and algorithms
- Based on Divide to Conquer paradigms
- Specificities of UQ problems (forward problem).

Objectives :

- Acceleration and implementation of (UQ) solvers on massively parallel machines
- **Domain decomposition** method
- Incorporate resilience properties into algorithms (exascale machine)

Focus on two contributions

1. Parallel DD method for stochastic fields **decomposition**
2. Parallel DD method for **Monte-Carlo sampling** in elliptic problems.

Stochastic Elliptic Equation

Generic model

$$\nabla \cdot (K(\theta)\nabla U(\theta)) = -f \quad (+\mathrm{BCs})$$

- Coefficient K is uncertain
- Model problem appearing in multiple domains : porous media flow, elasticity, thermal sciences, electromagnetism, ...
- Extensively analyzed and used for benchmark

The (now ?) classics : [Ghanem & Spanos, 1989]

- Parametrization of $K(\theta)$ using a **finite** number of RVs **KL expansion**
- Exploit the smoothness of U w.r.t. the RVs to build **spectral expansions**
- Possibly very high-dimensional problem (number of RVs)
- Sparse grid, PGD, low rank, adaptive constructions, ...

The current questions :

- Parametrization
- New architectures $\longrightarrow$ evolution of strategies and algorithms

Stochastic field $K(\boldsymbol{x}, \theta)$

Consider a bounded spatial domain Ω and an probability space $(\Theta, \Sigma_\Theta, d\mu)$.

- Let $\mathcal{V}(\Omega)$ is an inner product space,

$$\forall u \in \mathcal{V}, \|u\|_\Omega = \langle u, u \rangle_\Omega^{1/2} < \infty$$

- and $L_2(\Theta)$ the space of 2nd order random variables

$$\forall u \in L_2(\Theta), \|u\|_\Theta = \langle u, u \rangle_\Theta^{1/2} = \mathbb{E}\left[u^2\right]^{1/2} < \infty$$

Let $G \in L_2(\mathcal{V}, \Theta)$, *i.e.* **2nd order stochastic process G**

Separated representation of $G(\boldsymbol{x}, \theta)$:

$$G(\boldsymbol{x}, \theta) \approx G^{\mathrm{N}}(\boldsymbol{x}, \theta) = \sum_{l=1}^{\mathrm{N}} \phi_l(\boldsymbol{x}) \eta_l(\theta),$$

minimizing the truncation error

$$\epsilon_N^2 = \min_{\{\phi_l, \eta_l\}} \mathbb{E}\left[\left\| G(\boldsymbol{x}, \theta) - \sum_{l=1}^{\mathrm{N}} \phi_l(\boldsymbol{x}) \eta_l(\theta) \right\|_\Omega^2\right]$$

Karhunen-Loeve expansion

The optimal representation is given by

$$G^{\mathrm{N}}(\boldsymbol{x},\theta) = \mathbb{E}\left[G(\boldsymbol{x},\cdot)\right] + \sum_{l=1}^{\mathrm{N}} \sqrt{\lambda_l}\Phi_l(\boldsymbol{x})\xi_l(\theta).$$

where

- (λ_l, Φ_l) are the **eigenpairs** of the covariance function

$$C(\boldsymbol{x},\boldsymbol{x}') := \mathbb{E}\left[G(\boldsymbol{x},\cdot)G(\boldsymbol{x}',\cdot)\right] - \mathbb{E}\left[G(\boldsymbol{x},\cdot)\right]\mathbb{E}\left[G(\boldsymbol{x}',\cdot)\right],$$

- satisfying

$$\langle C(\cdot,\boldsymbol{x}'),\Phi_l(\boldsymbol{x}')\rangle_\Omega = \lambda_l\Phi_l(\boldsymbol{x}), \quad \|\Phi_l\|_\Omega = 1.$$

- **C is symmetric and non-negative** such that $\lambda_l \geq 0$ and the KL error is

$$\epsilon_N^2 = \mathbb{E}\left[\left\|G(\boldsymbol{x},\theta) - \mathbb{E}\left[G(\boldsymbol{x},\theta)\right] - \sum_{l=1}^{\mathrm{N}} \sqrt{\lambda_l}\Phi_l(\boldsymbol{x})\eta_l(\theta)\right\|_\Omega^2\right] = \sum_{l=\mathrm{N}+1}^{\infty} \lambda_l.$$

- Bi-orthonormal decomposition : $\langle \Phi_l, \Phi_{l'}\rangle_\Omega = \delta_{l,l'}$ and $\mathbb{E}\left[\xi_l, \xi_{l'}\right] = \delta_{l,l'}$

G Gaussian $\Rightarrow \xi_i$ *iid* $N(0,1)$

Illustration of the spectral decay

$C(x, x') = \exp(-(x - x')^2/2L^2)$, Gaussian field

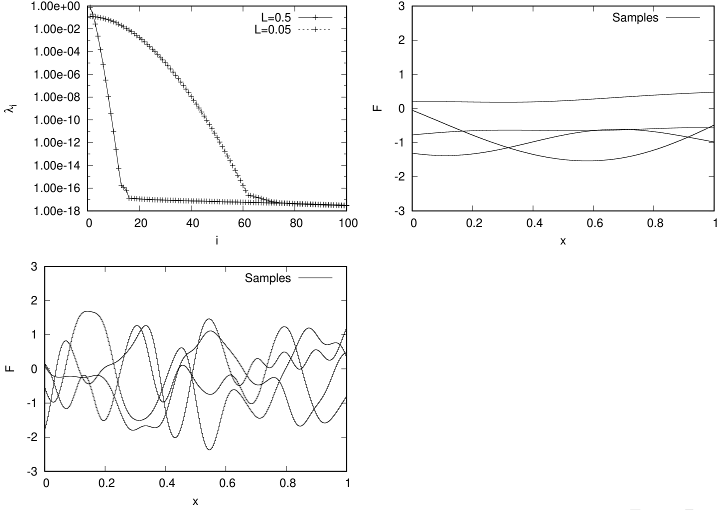

Stochastic Elliptic Equation

Log-normal field $K(\theta)$ approximated by truncated KL expansion of its log :

$$K^{\mathrm{N}}(\boldsymbol{x},\boldsymbol{\xi}(\theta)) := \exp\left[G^{N}(\boldsymbol{x},\boldsymbol{\xi})\right], \quad \boldsymbol{\xi} = (\xi_1 \cdots \xi_{\mathrm{N}}) \sim N(0, I_{\mathrm{N}}).$$

The problem becomes (homogeneous essential BCs)

$$\begin{aligned} &-\boldsymbol{\nabla}\cdot\left(K^{\mathrm{N}}(\boldsymbol{x},\boldsymbol{\xi})\boldsymbol{\nabla}U(\boldsymbol{x},\boldsymbol{\xi})\right) = -f, && \boldsymbol{x}\in\Omega \\ &U(\boldsymbol{x},\boldsymbol{\xi}) = 0, && \boldsymbol{x}\in\partial\Omega \end{aligned}$$

The solution is sought in $\mathcal{V}\otimes L_2(\Xi) := L_2(\mathcal{V},\Xi)$.

Consider an **Hilbertian basis** $\{\Psi_\alpha, \alpha = 1, 2, \dots\}$ **of** $L_2(\Xi)$ then

$$L_2(\mathcal{V},\Xi) \ni U(\boldsymbol{x},\boldsymbol{\xi}) = \sum_{\alpha=1}^{\infty} u_i(\boldsymbol{x})\Psi_\alpha(\boldsymbol{\xi}), \quad u_i \in \mathcal{V}, \Psi_\alpha \in L_2(\Xi).$$

PC expansion : Gaussian ξ corresponds to Hermite polynomials
The truncated PC expansion of $U(\boldsymbol{x},\boldsymbol{\xi})$:

$$U(\boldsymbol{x},\boldsymbol{\xi}) \approx \sum_{\boldsymbol{\alpha}\in\mathcal{A}} u_i(\boldsymbol{x})\Psi_{\boldsymbol{\alpha}}(\boldsymbol{\xi}), \quad \mathcal{A} = \{\boldsymbol{\alpha}\in\mathbb{N}^{\mathrm{N}}, |\boldsymbol{\alpha}| = \sum_{i=1}^{\mathrm{N}}\alpha_i \le \mathrm{N_o}\}, \quad |\mathcal{A}| = \frac{(\mathrm{N}+\mathrm{N_o})!}{\mathrm{N}!\mathrm{N_o}!}.$$

Limitations of Spectral Methods

Solution Methods :

- Galerkin methods : large set of $|\mathcal{A}|$ coupled problems
- Non-intrusive (projection) methods : large set of deterministic simulations ($k \times |\mathcal{A}|$)
- Cost increases with N and N^o .
- Non-isotropic truncature strategies `Nobile, Tamellini,...`
- Require smooth covariance with characteristic length scale comparable to the domain size `J. Charrier`

For many practical problems, **rough parameter fields**

- N must be large (with not so slow decay)
- High-dimensional problem : low-rank and separated approximations `[Nouy, Cohen, Schwab,...]`
- **KL decomposition is costly to compute for large meshes**
- Monte-Carlo is a viable alternative for this problems

Table of content

Karhunen-Loeve (KL) Eexpansion

Let $G \in L_2(\Omega, \Theta)$ be a **centered second order stochastic process** in Ω with covariance function $C : \Omega \times \Omega \mapsto \mathbb{R}$:

$$C(\boldsymbol{x}, \boldsymbol{x}') = \mathbb{E}\left[G(\boldsymbol{x}, \cdot), G(\boldsymbol{x}', \cdot)\right].$$

The Karhunen-Loeve expansion of G writes as

$$G(\boldsymbol{x}, \theta) = \sum_{\alpha=1}^{\infty} \sqrt{\lambda_\alpha}\eta_\alpha(\theta)\Phi_\alpha(\boldsymbol{x}),$$

where $\lambda_1 \geq \lambda_2 \geq \cdots$ are the leading eigenvalues of

$$\int_\Omega C(\boldsymbol{x}, \boldsymbol{x}')\Phi_\alpha(\boldsymbol{x}')d\boldsymbol{x}' = \lambda_\alpha \Phi_\alpha(\boldsymbol{x}).$$

Upon truncation with normalized Φ_α,

$$G(\boldsymbol{x}, \theta) \approx G_N(\boldsymbol{x}, \theta) \equiv \sum_{\alpha=1}^{N} \sqrt{\lambda_\alpha}\eta_\alpha(\theta)\Phi_\alpha(\boldsymbol{x}), \quad \text{where} \quad \eta_\alpha(\theta) = \int_\Omega G(\boldsymbol{x}, \theta)\Phi_\alpha(\boldsymbol{x})d\boldsymbol{x}.$$

Optimal approximation : $\|G - G_N\|^2_{L_2(\Omega,\Theta)} = \sum_{\alpha>N} \lambda_\alpha.$

Numerical KL Expansion

Galerkin approximation. **Finite-element** approximation of the KL modes Φ^h_α in a **finite dimensional subspace** $\mathcal{V}_h$, through

$$\Phi_\alpha(\boldsymbol{x}) \approx \Phi^h_\alpha(\boldsymbol{x}) \equiv \sum_{k=1}^{Q} c_{\alpha,k} v_k(\boldsymbol{x}).$$

The weak form of the eigenproblem is obtained by requiring

$$\langle \lambda \Phi^h(\boldsymbol{x}) - \int_\Omega C(\boldsymbol{x}, \boldsymbol{x}')\Phi^h(\boldsymbol{x}')d\boldsymbol{x}', V(\boldsymbol{x})\rangle_\Omega = 0 \quad \forall V \in \mathcal{V}_h.$$

Finite dimensional **generalized eigenproblem** for the Q coordinates in $\mathcal{V}_h$,

$$[S]\boldsymbol{c}_\alpha = \lambda_\alpha [M]\boldsymbol{c}_\alpha,$$

$$[S]_{k,k'} = \int_\Omega C(\boldsymbol{x}, \boldsymbol{x}')v_k(\boldsymbol{x})v_{k'}(\boldsymbol{x}')d\boldsymbol{x}d\boldsymbol{x}', \quad [M]_{k,k'} = \int_\Omega v_k(\boldsymbol{x})v_{k'}(\boldsymbol{x}')d\boldsymbol{x}d\boldsymbol{x}'$$

Numerical complexity when $\dim \mathcal{V}_h$ is large (full matrix).

Domain Decomposition

Domain partitioning : partition the domain Ω into D **non-overlapping** subdomains

$$\bar{\Omega} = \overline{\bigcup_{d=1}^{D} \Omega_d}, \quad \Omega_i \cap \Omega_{j \neq i} = \emptyset.$$

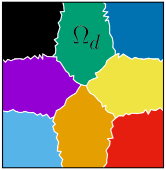

Compute local eigenmodes of subdomain d solving **the local eigenproblem**

$$\int_{\Omega_d} C(\boldsymbol{x}, \boldsymbol{x}')\tilde{\phi}_\beta^{(d)}(\boldsymbol{x}')d\boldsymbol{x}' = \lambda_\beta^{(d)}\tilde{\phi}_\beta^{(d)}(\boldsymbol{x}), \quad \left\| \tilde{\phi}_\beta^{(d)} \right\|_{\Omega_d} = 1.$$

and extend outside of Ω_d through

$$\forall \boldsymbol{x} \in \bar{\Omega}, \ \phi_\beta^{(d)}(\boldsymbol{x}) = \begin{cases} \tilde{\phi}_\beta^{(d)}(\boldsymbol{x}), & \boldsymbol{x} \in \Omega_d, \\ 0, & \boldsymbol{x} \notin \Omega_d. \end{cases}$$

Domain Decomposition

For each subdomain retain the $m_d > 0$ dominant eigenfunctions to form a reduced basis

$$\mathcal{B} = \bigcup_{d=1}^{D} \mathcal{B}_d, \quad \mathcal{B}_d = \left\{ \phi_\beta^{(d)}, \beta = 1, \ldots, m_d \right\}.$$

Denote $\mathcal{V}_\mathcal{B}$ the linear span of $\mathcal{B}$, and **approximate the global modes** as

$$\mathcal{V}_\mathcal{B} \ni \hat{\Phi}(\boldsymbol{x}) = \sum_{d=1}^{D} \sum_{\beta=1}^{m_d} a_\beta^{(d)} \phi_\beta^{(d)}(\boldsymbol{x}) \approx \Phi(\boldsymbol{x}).$$

Apply Galerkin method to cast the **reduced eigenproblem**,

$$\begin{bmatrix} [\hat{K}_{11}] & \cdots & [\hat{K}_{1D}] \\ \vdots & \ddots & \vdots \\ [\hat{K}_{D1}] & \cdots & [\hat{K}_{DD}] \end{bmatrix} \begin{Bmatrix} \boldsymbol{a}^{(1)} \\ \vdots \\ \boldsymbol{a}^{(D)} \end{Bmatrix} = \Lambda \begin{Bmatrix} \boldsymbol{a}^{(1)} \\ \vdots \\ \boldsymbol{a}^{(D)} \end{Bmatrix},$$

where

$$[\hat{K}_{i,j}]_{\alpha,\beta} = \int_{\Omega_i} \int_{\Omega_j} C(\boldsymbol{x}, \boldsymbol{x}') \phi_\alpha^{(i)}(\boldsymbol{x}) \phi_\beta^{(j)}(\boldsymbol{x}') d\boldsymbol{x} d\boldsymbol{x}', \quad 1 \le \alpha \le m_i, \ 1 \le \beta \le m_j.$$

Reduced Problem

Dimension : denote $n_t = \sum_{d=1}^{D} m_d = \dim \mathcal{B}$.
Reduced operator : $[\hat{K}] \in \mathbb{R}^{n_t \times n_t}$ is symmetric and positive definite.
The n_t eigenvalues Λ_α can be ordered as

$$\Lambda_1 \geq \Lambda_2 \geq \cdots \geq \Lambda_{n_t} \geq 0.$$

Truncature : let $1 \leq \hat{N} \leq n_t$, the approximation of U is finally

$$G(\boldsymbol{x}, \theta) \approx \hat{G}_{\hat{N}}(\boldsymbol{x}, \theta) \equiv \sum_{\alpha=1}^{\hat{N}} \sqrt{\Lambda_\alpha}\, \hat{\eta}_\alpha(\theta) \hat{\Phi}_\alpha(\boldsymbol{x}), \quad \hat{\Phi}_\alpha(\boldsymbol{x}) = \sum_{d=1}^{D} \sum_{\beta=1}^{m_d} a_{\alpha,\beta}^{(d)} \phi_\beta^{(d)}(\boldsymbol{x}).$$

Remarks :

- The approach is **suitable for parallel implementation**
- The local problems can eventually enable direct solvers
- Could use different discretizations over distinct subdomains
- $\mathcal{V}_{\mathcal{B}} \not\subseteq \mathcal{V}_h$
- n_t **is fixed by the targeted error** on G and is not $\dim \mathcal{V}_h$.

Example

Consider $\Omega = [0, 1]^2$ and the squared exponential covariance

$$C(\boldsymbol{x}, \boldsymbol{x}') = \exp\left[-\frac{\|\boldsymbol{x} - \boldsymbol{x}'\|_2^2}{2L^2}\right].$$

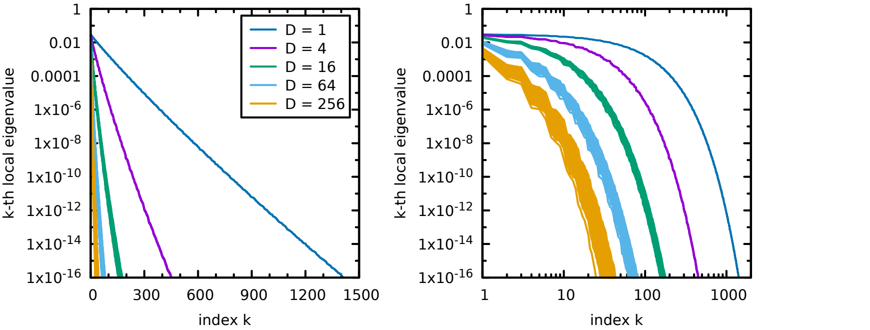

FIGURE – Spectra of local decompositions for $L = 0.1$, and different D as indicated.

Example

Differences between $\hat{G}_N$ and G_N for given N, using two metrics :

$$\epsilon_{spec} = \frac{\sum_{k=1}^{N} |\lambda_k - \Lambda_k|}{\sum_{k=1}^{N} |\lambda_k|}.$$

and $\epsilon^2_{sub}(G_N)$ defined as

$$\epsilon^2_{sub}(V) = \frac{\mathbb{E}\left[\left\| V(\boldsymbol{x},\theta) - \sum_{\alpha=1}^{N} \left\langle V(\boldsymbol{x},\theta), \hat{\Phi}_\alpha(\boldsymbol{x}) \right\rangle_\Omega \hat{\Phi}_\alpha(\boldsymbol{x}) \right\|^2_\Omega\right]}{\mathbb{E}\left[\| V(\boldsymbol{x},\theta)\|^2_\Omega\right]}.$$

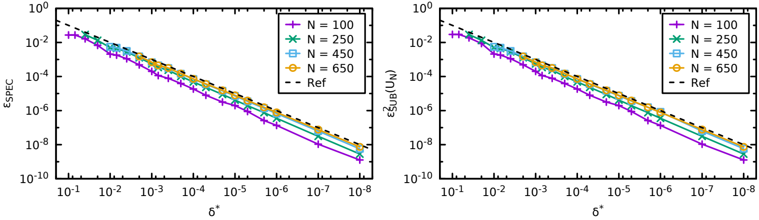

FIGURE – Computations use $D = 80$ and $L = 0.1$.

Size of reduced problem

Increasing the number of subdomains : effect on the reduced problem dimension.

D	n_t	$\bar{m} \pm \sigma_{m_d}$
20	431	21.55 ± 1.43
40	542	13.55 ± 0.59
80	741	9.26 ± 0.56
160	983	6.14 ± 0.35
320	1682	5.26 ± 0.44
640	2,306	3.60 ± 0.53
1280	3,840	3.00 ± 0.00

TABLE – Progression of n_t for different values of D with $\delta^2 = 2 \times 10^{-3}$ and $L = 0.1$.

Anisotropy and size distribution effects

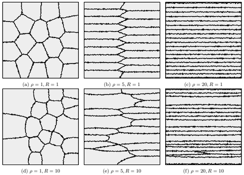

FIGURE – Controlling the aspect ratio (ρ) and area dispersion (R).

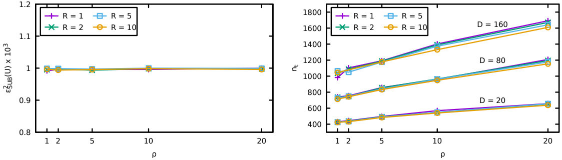

FIGURE – $\epsilon^2_{sub}(G)$ and reduced basis dimension versus ρ for a target accuracy $\delta^2 = 2 \times 10^{-3}$.

Computational Efficiency

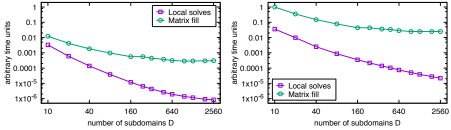

(a) P_0 finite elements. (b) P_2 finite elements.

FIGURE – Sequential computation : timing local decompositions and reduced problem assembly (fixed target accuracy).

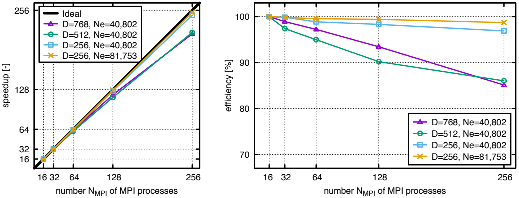

FIGURE – Parallel speedup (left) and efficiency (right) versus the number of MPI processes.

Stochastic Elliptic Problem

Consider

$$\nabla \cdot (K(\boldsymbol{x}, \theta)\nabla u(\boldsymbol{x}, \theta)) = -f(\boldsymbol{x}), \quad u(\boldsymbol{x} \in \Gamma) = 0,$$

in particular for **log-normal random field** :

$$\log(K - \kappa_{\text{min}}) = G(\boldsymbol{x}, \theta) \sim N(\mu, C).$$

Upon deterministic FE discretization, it comes

$$A[\theta]\boldsymbol{u}(\theta) = \boldsymbol{b}(\theta),$$

where $\boldsymbol{u}$ is a random vector of $\mathbb{R}^{\mathrm{N_e}}$ (e.g. nodal values on a FE mesh).
Solved by

- Monte Carlo : proceed by sampling K to compute samples of $\boldsymbol{u}(\theta)$ and estimate averages
- Stochastic Spectral Method (PC, PGD, low-rank) : **parametrization** of $K(\boldsymbol{x}, \boldsymbol{\xi})$ to construct a functional representation of $\boldsymbol{u}(\boldsymbol{\xi}) \approx \sum_\alpha \boldsymbol{u}_\alpha \Psi_\alpha(\boldsymbol{\xi})$.

Domain Decomposition method

Consider again a non overlapping partition of the FE mesh. The stochastic vector $\boldsymbol{u}$ can be slip into $\boldsymbol{u}_{\text{in}}$ and $\boldsymbol{u}_{\text{bd}}$ containing the internal unknowns and inner boundary values. The discrete system can be recast as

$$\begin{bmatrix} [\boldsymbol{A}_{\text{bd,bd}}] & [\boldsymbol{A}_{\text{bd,in}}] \\ [\boldsymbol{A}_{\text{in,bd}}] & [\boldsymbol{A}_{\text{in,in}}] \end{bmatrix} \begin{pmatrix} \boldsymbol{u}_{\text{bd}} \\ \boldsymbol{u}_{\text{in}} \end{pmatrix} = \begin{pmatrix} \boldsymbol{b}_{\text{bd}} \\ \boldsymbol{b}_{\text{in}} \end{pmatrix},$$

Or $\widehat{[\boldsymbol{A}]}\boldsymbol{u}_{\text{bd}} = \widehat{\boldsymbol{b}}$, where

$$\widehat{[\boldsymbol{A}]} \doteq [\boldsymbol{A}_{\text{bd,bd}}] - [\boldsymbol{A}_{\text{bd,in}}][\boldsymbol{A}_{\text{in,in}}]^{-1}[\boldsymbol{A}_{\text{in,bd}}], \quad \widehat{\boldsymbol{b}} \doteq \boldsymbol{b}_{\text{bd}} - [\boldsymbol{A}_{\text{bd,in}}][\boldsymbol{A}_{\text{in,in}}]^{-1}\boldsymbol{b}_{\text{in}}.$$

$[\boldsymbol{A}_{\text{in,in}}]$ has **diagonal block structure** :

- Applying $[\boldsymbol{A}_{\text{in,in}}]^{-1}$ amounts to solve D **local problems** for the inner nodes of the subdomains $\Omega^{(d)}$
- Can be carried out in parallel
- Suggest solving the condensed problem in a **matrix free** approach with parallel computation of $\widehat{[\boldsymbol{A}]}\boldsymbol{v}$.

Condensed Problem

The condensed problem

$$\widehat{[\boldsymbol{A}]}(\theta)\boldsymbol{u}_{\mathrm{bd}}(\theta) = \widehat{\boldsymbol{b}}(\theta),$$

can be expressed as subdomains contributions :

$$\widehat{[\boldsymbol{A}]}(\theta) = \sum_{d=1}^{\mathrm{D}} \widehat{[\boldsymbol{A}]}^{(d)}(\theta), \quad \widehat{\boldsymbol{b}}(\theta) = \sum_{d=1}^{\mathrm{D}} \widehat{\boldsymbol{b}}^{(d)}(\theta).$$

Elementary contributions can be determined solving a sequence of local problems with deterministic boundary conditions.

Their solutions then depends on (log of) $\kappa(\boldsymbol{x}, \theta)$ for $\boldsymbol{x} \in \Omega^{(d)}$. In other words, $\widehat{[\boldsymbol{A}]}^{(d)}(\theta)$ and $\widehat{\boldsymbol{b}}^{(d)}(\theta)$ can be expanded in terms of **local KL coefficients** :

$$(\widehat{[\boldsymbol{A}]}, \widehat{\boldsymbol{b}})^{(d)}(\theta) = (\widehat{[\boldsymbol{A}]}, \widehat{\boldsymbol{b}})^{(d)}(\boldsymbol{\xi}^{(d)}(\theta)) \approx \sum_{\alpha} (\widehat{[\boldsymbol{A}]}, \widehat{\boldsymbol{b}})^{(d)}_{\alpha} \Psi_{\alpha}(\boldsymbol{\xi}^{(d)}(\theta)),$$

where $\boldsymbol{\xi}^{(d)} \sim N(0, I_{m_d})$.
Constructing approximations is manageable provided the subdomains are small enough so m_d is small.

MC Sampling

To sample the condensed problem, we have to solve

$$\widehat{[\boldsymbol{A}]}(\theta)\boldsymbol{u}_{\mathrm{bd}}(\theta) = \widehat{\boldsymbol{b}}(\theta)$$

with

$$\widehat{[\boldsymbol{A}]}(\theta) = \sum_{d=1}^{\mathrm{D}} \widehat{[\boldsymbol{A}]}^{(d)}(\theta) \approx \sum_{d=1}^{\mathrm{D}} \sum_{\alpha} \widehat{[\boldsymbol{A}]}_{\alpha}^{(d)} \Psi_{\alpha}(\boldsymbol{\xi}^{(d)}(\theta)).$$

Amounts to **sample jointly** the local KL coefficients.

- **Directly generate sample of the condensed problem**, without solving any local problems
- Leads to significant computational saving
- Still need to solve (**once**) local problems to yield full solution sample.

MC Sampling : realizations

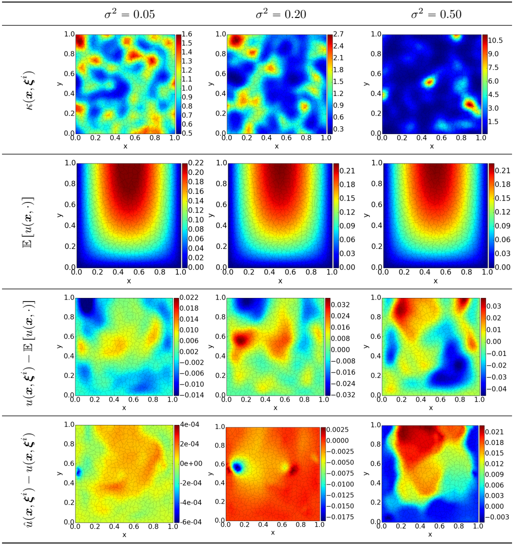

MC Sampling

Expectation of $u(\boldsymbol{x}, \theta)$

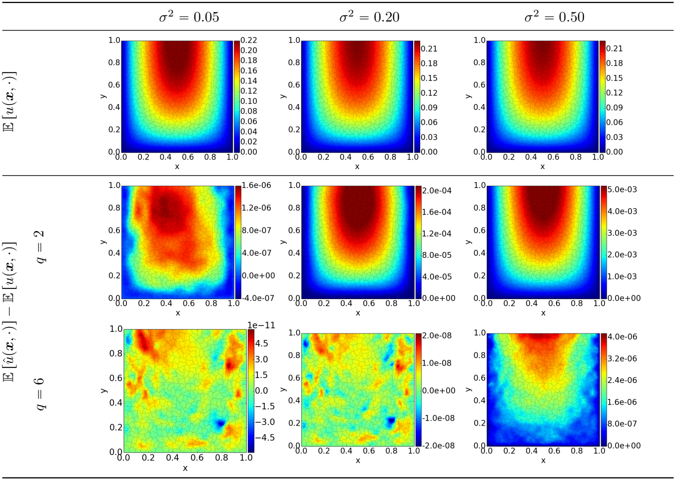

Figure 2: Top row shows $\mathbb{E}\left[u(\boldsymbol{x}, \cdot)\right]$, estimated via Monte Carlo using 500,000 samples, for three different values of σ. Rows two and three show $\mathbb{E}\left[u\right](\boldsymbol{x}, \cdot) - \mathbb{E}\left[u(\boldsymbol{x}, \cdot)\right]$ for two different values of the polynomial order q. For each i, $\hat{u}(\boldsymbol{x}, \boldsymbol{\xi}^i)$ was computed using $D = 480$ and $L = 0.1$.

MC Sampling

Standard deviation of $u(\boldsymbol{x}, \theta)$

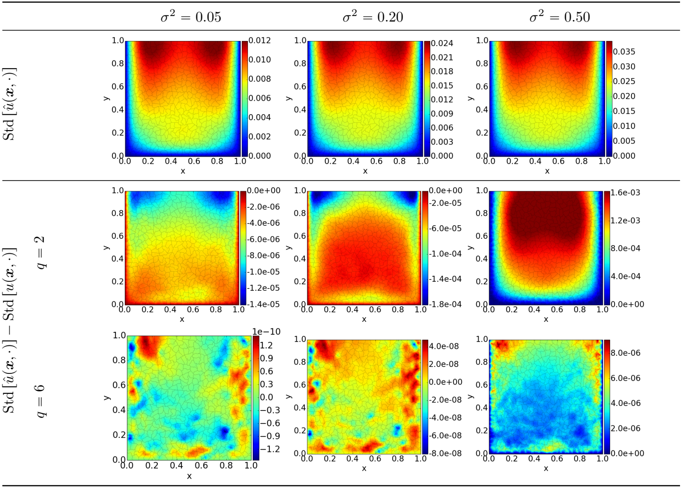

Figure 3: Top row shows $\mathrm{Std}\,[u(\boldsymbol{x}, \cdot)]$, estimated via Monte Carlo using 500,000 samples, for three different values of σ. Rows two and three show $\mathrm{Std}\,[u(\boldsymbol{x}, \cdot)] - \mathrm{Std}\,[u(\boldsymbol{x}, \cdot)]$ for two different values of the polynomial order q. For each i, $\hat{u}(\boldsymbol{x}, \boldsymbol{\xi}^i)$ was computed using $D = 480$ and $L = 0.1$.

MC Sampling

L-2 norm of error in mean :

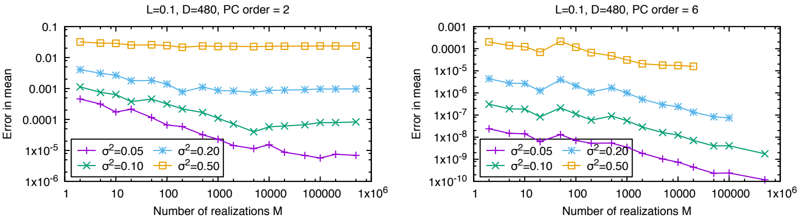

(a) Error in mean, $q = 2$ (b) Error in mean, $q = 6$

Figure 4: Norm of the error in the mean solution as a function of the number of realizations.

MC Sampling

Error on mean and standard deviation of solution

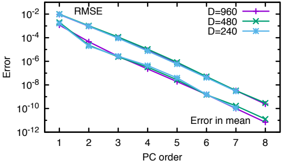

(a) $\sigma^2 = 0.05$

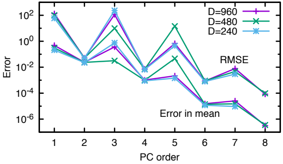

(b) $\sigma^2 = 0.5$

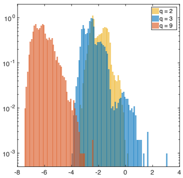

(a) $\log_{10} \|u - \hat{u}\|_{L_2(\Omega)}$

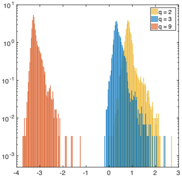

(b) $\log_{10} \|\widehat{[\mathbf{A}]} - \widetilde{[\mathbf{A}]}\|_F$

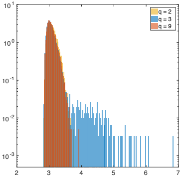

(c) $\log_{10} \operatorname{cond}\left(\widetilde{[\mathbf{A}]}\right)$

Figure 7: Log-Histograms of the error norm $\|u - \hat{u}\|_{L2(\Omega)}$ (left), of the approximation error on condensed operator $\|\widehat{[\mathbf{A}]} - \widetilde{[\mathbf{A}]}\|_F$ (center), and of the condition number of the approximate system $\operatorname{cond}\left(\widetilde{[\mathbf{A}]}\right)$ (right) for PC orders $\mathrm{N_o} = 2, 3, 9$. Case of G with $\sigma^2 = 0.5$ and $L = 1$.

MC Sampling

Non positive approximation of the condensed operator

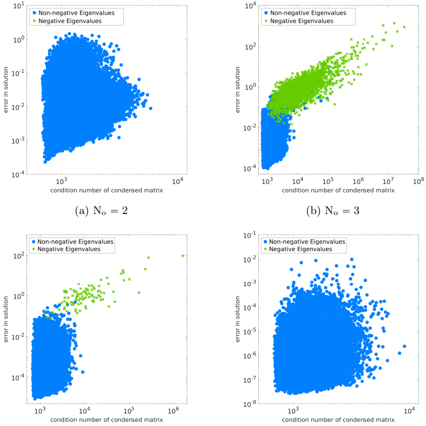

(a) $N_o = 2$ (b) $N_o = 3$ (c) $N_o = 5$ (d) $N_o = 9$

Figure 8: Samples of the error in the solution $\|u - \hat{u}\|_{L2(\Omega)}$ as a function of the condition number cond $[\widetilde{\mathbf{A}}]$. The samples are colored according to the sign of the smallest eigenvalue of $[\widetilde{\mathbf{A}}]$. Different PC orders as indicated.

PC approximation

Complexity analysis

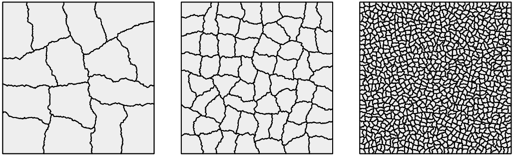

(a) D = 15. (b) D = 60. (c) D = 960.

Figure 9: Partitions of the computational mesh into different numbers of subdomains D as indicated.

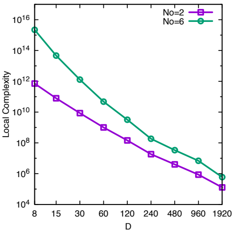

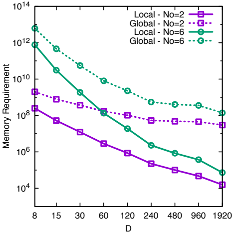

Figure 10: Local complexity (left plot) and local and global memory requirements (right plot), as a function of the number of subdomains D and for two PC degree $N_o = 2$ and $N_o = 6$. Note that both plots use a log-log scale.

Parallel efficiency

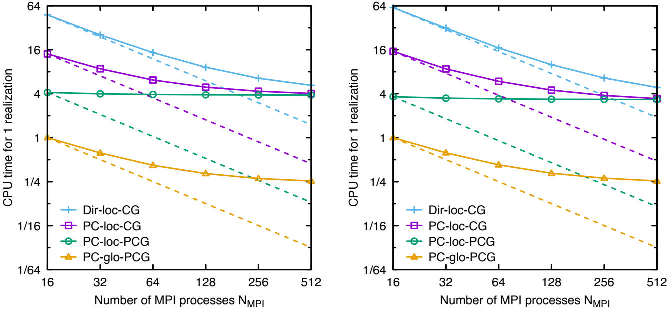

(a) $N_e = 163{,}272$. (b) $N_e = 327{,}334$.

Figure 12: Scaled CPU times, to generate one sample, as a function of the number of MPI processes N_{MPI}. The dashed lines represent ideal parallel scaling.

Conclusions and Remarks

- Complexity Reduction by means of local solves and local parametrization
- PC expansions of local operators is effective and avoid costly online assembly
- Computational efficiency and scalability to be improved
- **Resilient aspects** must be added

Next / ongoing :

- Accelerated Schwartz method using local PC approximations as preconditionners (Thèse João Reis, CMAP)
- Exploit spectral information from previous solves to precondition new problem -recycling Krylov- (thèse Nicolas Venkovic, Cerfacs)
- Extension to a **multi-level framework** : MC convergence & preconditioning
- "Condense / compress" the local / global problems : H-matrix, Low rank approximation, . . .

- A. Contreras, P. Mycek, O. Le Maître, F. Rizzi, B. Debusschere and O. Knio, Parallel Domain Decomposition Strategies for Stochastic Elliptic Equations. Part B : Accelerated Monte-Carlo Sampling with Local PC Expansions, *SIAM J. Sci. Comp.*, **40** :4, C547-C580,pp. (2018).
- A. Contreras, P. Mycek, O. Le Maître, F. Rizzi, B. Debusschere and O. Knio, Parallel Domain Decomposition Strategies for Stochastic Elliptic Equations. Part A : Local KL Representations, *SIAM J. Sci. Comp.*, **40** :4, pp. C520-546, (2018).

References

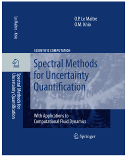

O.P. Le Maître and O.M. Knio, Spectral Methods for Uncertainty Quantification, with application to CFD, Scientific Computing Series, Springer, (2010).
`https://perso.limsi.fr/olm/`

- A. Contreras, P. Mycek, O. Le Maître, F. Rizzi, B. Debusschere and O. Knio, Parallel Domain Decomposition Strategies for Stochastic Elliptic Equations. Part B : Accelerated Monte-Carlo Sampling with Local PC Expansions, *SIAM J. Sci. Comp.*, **40** :4, C547-C580,pp. (2018).
- A. Contreras, P. Mycek, O. Le Maître, F. Rizzi, B. Debusschere and O. Knio, Parallel Domain Decomposition Strategies for Stochastic Elliptic Equations. Part A : Local KL Representations, *SIAM J. Sci. Comp.*, **40** :4, pp. C520-546, (2018).
- P. Mycek *et al.* A resilient domain decomposition polynomial chaos solver for uncertain elliptic PDEs, *Computer Phycis Communications*, **216**, pp. 18-34, (2017).
- P. Mycek *et al.* Discrete a priori bounds for the detection of corrupted PDE solutions in exascale computations, *SIAM J. Sci. Comp.*, **39** :1, pp. 1-28, (2017).
- F. Rizzi *et al.* Exploring the Interplay of Resilience and Energy Consumption for a Task-Based Partial Differential Equations Preconditioner, *Parallel Computing*, (in press).
- F. Rizzi *et al.* Partial differential equations preconditioner resilient to soft and hard faults, *The International Journal of High Performance Computing Applications*, (in press).