Preconditioners based on Voronoi quantizers of random coefficient fields for the iterative solves of stochastic elliptic partial differential equations

Nicolas Venkovic[1], Paul Mycek[2], Luc Giraud[3] and Olivier Le Maître[4]

[1]NXP Semiconductors
[2]Cerfacs
[3]Inria, Bordeaux Sud-Ouest
[4]CNRS, École Polytechnique

20 June 2023

SECURE CONNECTIONS
FOR A SMARTER WORLD

Stochastic elliptic partial differential equations (PDEs)

- Let us consider the domain $\Omega := [0,1]^2$ with boundary $\partial\Omega$, and the set Θ of all possible outcomes. We search for $u : \overline{\Omega} \times \Theta \to \mathbb{R}$ such that

$$\nabla \cdot [\kappa(x,\theta)\nabla u(x,\theta)] = f(x) \; \forall \, x \in \Omega$$
$$u(x,\theta) = 0 \qquad \forall \, x \in \partial\Omega$$

is almost surely satisfied.

- We assume $f : \Omega \to \mathbb{R}$ is square integrable, i.e., $f \in L^2(\Omega)$ where

$$\|f\|_\Omega^2 := \int_\Omega |f(x)|^2 \mathrm{d}x < \infty \iff f \in L^2(\Omega).$$

- We further assume the random coefficient field $\kappa : \Omega \times \Theta \to \mathbb{R}$ is such that $P[\theta \in \Theta : \kappa(\cdot,\theta) \in \mathcal{A}] = 1$ where $\mathcal{A} := \{\kappa \in L^2(\Omega), \kappa(x) > 0 \; \forall \, x \in \Omega\}$.
- We define the set $L^2(\Omega,\Theta)$ of 2nd order stochastic processes such that

$$\mathbb{E}[\|\kappa(\cdot,\theta)\|_\Omega^2] < \infty \iff \kappa \in L^2(\Omega,\Theta).$$

- See Babuska al. (2004) for more details on existence and uniqueness of u.

Babuska, Ivo, Raúl Tempone, and Georgios E. Zouraris. "Galerkin finite element approximations of stochastic elliptic partial differential equations." SIAM Journal on Numerical Analysis 42.2 (2004): 800-825.

Karhunen-Loève (KL) representation of coefficient fields

- Let κ be a real-valued 2$^{\text{nd}}$ order stochastic processes, i.e., $\kappa \in L^2(\Omega, \Theta)$, with zero mean and known covariance $C(x, x') = \mathbb{E}[\kappa(x, \cdot)\kappa(x', \cdot)]$.
- Then, the truncated KL expansion κ_N of κ minimizes the representation error $\mathbb{E}[\|\kappa - \kappa_N\|_\Omega^2]$ over N-dimensional function spaces. It is given by

$$\kappa_N(x, \theta) := \sum_{\alpha=1}^{N} \sqrt{\lambda_\alpha}\xi_\alpha(\theta)\Phi_\alpha(x)$$

where $(\lambda_\alpha, \Phi_\alpha) \in \mathbb{R}^+ \times L^2(\Omega)$ is the α-th dominant eigen-pair of the covariance function and Φ_α is a normalized eigenfunction. That is, $(\lambda_\alpha, \Phi_\alpha)$ is solution of the Fredholm integral equation:

$$\int_\Omega C(x, x')\Phi(x')\mathrm{d}x' = \lambda\Phi(x), \quad \|\Phi\|_\Omega^2 = 1.$$

- The random variables (RVs) ξ_α, a.k.a. the "stochastic coordinates" of κ_N, are uncorrelated with zero mean and unit variance, i.e., $\mathbb{E}[\xi_\alpha\xi_\beta] = \delta_{\alpha\beta}$.
- In case κ is a Gaussian process, the RVs are independent.

State-of-the-art

Two main approaches have been used to characterize the uncertainty of u.

- ▶ Approaches based on **polynomial chaos (PC)** expansions:

Leveraging the KL expansion, an approximate functional representation of the random solution $u : \Omega \times \Theta \to \mathbb{R}$ is built in the form of a spectral expansion

$$u_M(x, \boldsymbol{\xi}(\theta)) := \sum_{\alpha=0}^{M} u_\alpha(x) \Psi_\alpha(\boldsymbol{\xi}(\theta))$$

where $\Psi_\alpha : \Theta \to \mathbb{R}$ is set a priori from a finite polynomial basis. The computation of the coefficients $u_\alpha : \Omega \to \mathbb{R}$ is done by stochastic Galerkin, regression or collocation. Once equipped with a spectral expansion, statistics can be computed on the basis of approximate solution realizations.

- ▶ Approaches based on **Monte Carlo (MC)** sampling:

Statistics are computed on the basis of solution realizations $u(\cdot, \theta)$ which are obtained by solving a deterministic equations with the corresponding realizations $\kappa(\cdot, \theta) \in \mathcal{A}$ of the coefficient field. The spatial discretization of the deterministic equation for a given event θ leads to an SPD linear system

$$\mathbf{A}(\theta)\mathbf{u}(\theta) = \mathbf{b}(\theta).$$

Limit preconditioning strategies

- $\mathbf{A}(\theta)$ being SPD, $\mathbf{A}(\theta)\mathbf{u}(\theta) = \mathbf{b}(\theta)$ can be solved by conjugate gradient (CG), i.e., we search for iterates $\mathbf{u}^{(j)}(\theta)$ such that:

$$\begin{aligned} \mathbf{u}^{(j)}(\theta) - \mathbf{u}^{(0)} &\in \mathcal{K}^{(j)}(\mathbf{A}(\theta), \mathbf{r}^{(0)}(\theta)) \\ \mathbf{r}^{(j)}(\theta) &\perp \mathcal{K}^{(j)}(\mathbf{A}(\theta), \mathbf{r}^{(0)}(\theta)) \end{aligned}$$

where $\mathcal{K}^{(j)}(\mathbf{A}(\theta), \mathbf{r}^{(0)}(\theta)) := \mathrm{Span}\{\mathbf{r}^{(0)}(\theta), \mathbf{A}(\theta)\mathbf{r}^{(0)}(\theta), \ldots, \mathbf{A}^{j-1}(\theta)\mathbf{r}^{(0)}(\theta)\}$ is the Krylov subspace of $\mathbf{A}(\theta)$ generated by $\mathbf{r}^{(0)}(\theta)$.

- J is the (random) number of solver iterations to reach a backward error of 10^{-6}.

Median realization preconditioner

A single SPD preconditioner $\mathbf{M}_{0,\bullet}^{-1}$ is defined based on $\mathbf{A}(\boldsymbol{\xi} = 0)$. We then search for iterates $\mathbf{u}^{(i)}(\theta)$ such that:

$$\begin{aligned} \mathbf{u}^{(j)}(\theta) - \mathbf{u}^{(0)} &\in \mathcal{K}^{(j)}(\mathbf{M}_{0,\bullet}^{-1}\mathbf{A}(\theta), \mathbf{M}_{0,\bullet}^{-1}\mathbf{r}^{(0)}(\theta)) \\ \mathbf{r}^{(j)}(\theta) &\perp \mathcal{K}^{(j)}(\mathbf{M}_{0,\bullet}^{-1}\mathbf{A}(\theta), \mathbf{M}_{0,\bullet}^{-1}\mathbf{r}^{(0)}(\theta)) \end{aligned}$$

Limit preconditioning strategies

- $\mathbf{A}(\theta)$ being SPD, $\mathbf{A}(\theta)\mathbf{u}(\theta) = \mathbf{b}(\theta)$ can be solved by conjugate gradient (CG), i.e., we search for iterates $\mathbf{u}^{(j)}(\theta)$ such that:

$$\begin{aligned} \mathbf{u}^{(j)}(\theta) - \mathbf{u}^{(0)} &\in \mathcal{K}^{(j)}(\mathbf{A}(\theta), \mathbf{r}^{(0)}(\theta)) \\ \mathbf{r}^{(j)}(\theta) &\perp \mathcal{K}^{(j)}(\mathbf{A}(\theta), \mathbf{r}^{(0)}(\theta)) \end{aligned}$$

where $\mathcal{K}^{(j)}(\mathbf{A}(\theta), \mathbf{r}^{(0)}(\theta)) := \mathrm{Span}\{\mathbf{r}^{(0)}(\theta), \mathbf{A}(\theta)\mathbf{r}^{(0)}(\theta), \ldots, \mathbf{A}^{j-1}(\theta)\mathbf{r}^{(0)}(\theta)\}$ is the Krylov subspace of $\mathbf{A}(\theta)$ generated by $\mathbf{r}^{(0)}(\theta)$.

- J is the (random) number of solver iterations to reach a backward error of 10^{-6}.

Realization-dependent ideal preconditioner

For every single realization θ, a preconditioner $\mathbf{M}_{\bullet}^{-1}(\theta)$ is defined based on $\mathbf{A}(\theta)$. We then search for iterates $\mathbf{u}^{(i)}(\theta)$ such that:

$$\begin{aligned} \mathbf{u}^{(j)}(\theta) - \mathbf{u}^{(0)} &\in \mathcal{K}^{(j)}(\mathbf{M}_{\bullet}^{-1}(\theta)\mathbf{A}(\theta), \mathbf{M}_{\bullet}^{-1}(\theta)\mathbf{r}^{(0)}(\theta)) \\ \mathbf{r}^{(j)}(\theta) &\perp \mathcal{K}^{(j)}(\mathbf{M}_{\bullet}^{-1}(\theta)\mathbf{A}(\theta), \mathbf{M}_{\bullet}^{-1}(\theta)\mathbf{r}^{(0)}(\theta)) \end{aligned}$$

Results of limit preconditioning strategies

- We consider a covariance $C(x, x') = \exp(-\|x - x'\|^2/0.1^2)$.

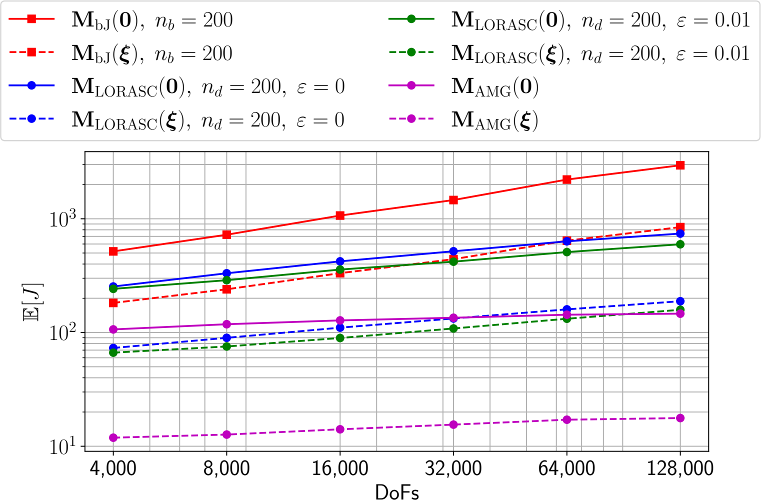

- Speedup of $\mathbf{M}_{\bullet}^{-1}(\boldsymbol{\xi})$: 2-3X (bJ), 3X (LORASC), 10X (AMG).

Alternative preconditioning strategies

- Using $\mathbf{M}_\bullet(\mathbf{0})$ leads to large numbers of solver iterations, while $\mathbf{M}_\bullet(\boldsymbol{\xi})$ entails significant preconditioner setup times, e.g., computing factorizations for every single realization $\boldsymbol{\xi}(\theta)$.
- Alternative preconditioning strategies are needed.
- We consider preconditioning strategies which consist of both:
  1. A P-quantizer $q : \kappa \in \mathcal{A} \mapsto \hat{\kappa} \in \hat{\mathcal{A}}$ with centroidal coefficient fields in the codebook $\hat{\mathcal{A}} := \{\hat{\kappa}_p \in \mathcal{A},\ p = 1, \ldots, P\}$. The quantizer q serves as a compact representation of the random coefficient field.
  2. A preconditioner $\mathbf{M} : \kappa \in \mathcal{A} \mapsto \mathbf{M}(\kappa) \in \mathrm{Sym}^+_{n\times n}(\mathbb{R})$.
     We are interested in the composition $\mathbf{M} \circ q : \kappa \in \mathcal{A} \mapsto \{\hat{\mathbf{M}}_1, \ldots, \hat{\mathbf{M}}_P\}$.
     The preconditioners $\hat{\mathbf{M}}_1, \ldots, \hat{\mathbf{M}}_P$ are not known explicitly.
     Instead, as we consider solving cycles of algebraic multigrid solvers for $\mathbf{A}(\hat{\kappa}_1), \ldots, \mathbf{A}(\hat{\kappa}_P)$, we simply know how to efficiently compute the mapping $\mathbf{x} \mapsto \hat{\mathbf{M}}_p^{-1}\mathbf{x}$ for $p = 1, \ldots, P$.

Alternative preconditioning strategies

- We assume there exists an invertible map $T_1 : L^2(\Omega) \to L^2(\Omega)$ such that $T_1^{-1}\kappa$ is a Gaussian process with zero mean.
- We know the truncated KL expansion $\hat{T}_1^{-1}\kappa$ which approximates $T_1^{-1}\kappa$ with m dominant eigen-pairs $(\lambda_\alpha, \Phi_\alpha) \in \mathbb{R}^+ \times L^2(\Omega)$.
- Let us introduce the following projection:

$$\hat{P}_1^{-1} : f \in L^2(\Omega) \mapsto \begin{bmatrix} \lambda_1^{-1/2}\langle \Phi_1, f\rangle_\Omega \\ \vdots \\ \lambda_m^{-1/2}\langle \Phi_m, f\rangle_\Omega \end{bmatrix}, \langle f, g\rangle_\Omega := \int_\Omega f(x)g(x)\mathrm{d}x \;\forall\; f, g \in L^2(\Omega)$$

$$\text{and}\;\; \hat{P}_1 : \boldsymbol{\xi} \in \mathbb{R}^m \mapsto \sum_{k=1}^{m} \lambda_k^{1/2}\xi_k\Phi_k(\cdot) \in L^2(\Omega) \;\; \text{s.t.} \;\; \hat{T}_1^{-1} = \hat{P}_1 \circ \hat{P}_1^{-1} \circ T_1^{-1}.$$

- We introduce a quantizer q_2 of $\mathbb{R}^m$ and an invertible map $T_2 : \mathbb{R}^m \to \mathbb{R}^m$ which we use as follows to define q:

$$q : \kappa(\cdot) \in \mathcal{A} \mapsto \widetilde{T}(q_2(\widetilde{T}^{-1}\kappa(\cdot))) \in \hat{\mathcal{A}} \subset \mathcal{A}$$

where $\widetilde{T} := T_1 \circ \hat{P}_1 \circ T_2$ and $\hat{\mathcal{A}}$ is the codebook induced by q.

Optimal preconditioning strategies

- The representation error of $\kappa(\cdot,\Theta)$ by $q(\kappa(\cdot,\Theta))$ is the distortion

$$w(q,d) := \mathbb{E}[d(\kappa,q(\kappa))] = \int_\Theta d(\kappa(\cdot,\theta),q(\kappa(\cdot,\theta)))\mathrm{d}\mu(\theta)$$

where the distortion functional, a.k.a. divergence, $d:\mathcal{A}\times\mathcal{A}\to[0,\infty)$ measures proximity between realizations of the coefficient field.

- Every Voronoi quantizer q has a codebook $\hat{\mathcal{A}} := \{\hat{\kappa}_1,\dots,\hat{\kappa}_P\}\subset\mathcal{A}$ s.t.

$$q:\kappa\in\mathcal{A}\mapsto\sum_{p=1}^{P}\hat{\kappa}_p\mathbf{1}[\kappa\in\mathcal{A}_p],\ \mathcal{A}_p\subset\{\kappa\in\mathcal{A},\ d(\kappa,\hat{\kappa}_p)\leq d(\kappa,\hat{\kappa}_q),\ q=[1,P]\}$$

and so that $\mathcal{A}_1,\dots,\mathcal{A}_P$ form a Borel partition of $\mathcal{A}$.

- The local distortions $w_p(q,d) := \mathbb{E}[d(\kappa,q(\kappa))\,|\,\kappa\in\mathcal{A}_p] = w_p(\hat{\kappa}_p,d)$ and attribution frequencies $f_p := \mu(\mathcal{A}_p)$ form a decomposition of distortion

$$w(\hat{\mathcal{A}},d) = \sum_{p=1}^{P} f_p w_p(\hat{\kappa}_p,d).$$

- For a partition $\mathcal{A}_1,\dots,\mathcal{A}_P$, the distortion is minimized by selecting centroidal fields $\hat{\kappa}_p$ which minimize the local distortions $w_p(\hat{\kappa}_p,d)$.

Computation of stationary quantizers

- Remember that we let $q(\kappa(\cdot)) = \widetilde{T}(q_2(\widetilde{T}^{-1}\kappa(\cdot)))$ be induced by a vector quantizer q_2 of $T_2^{-1}(\boldsymbol{\xi})$.
- We are interested in L^2 quantizers with distortions given by

$$w_2(q_2) := \mathbb{E}[\|T_2^{-1}(\boldsymbol{\xi}) - q_2(T_2^{-1}(\boldsymbol{\xi}))\|^2] = \int_\Theta \|T_2^{-1}(\boldsymbol{\xi}) - q_2(T_2^{-1}(\boldsymbol{\xi}))\|^2 \mathrm{d}\mu_{\boldsymbol{\xi}}(\theta).$$

- We let q_2 be a Voronoi quantizer, and we denote the partition of $T_2^{-1}(\mathbb{R}^m)$ induced by q_2 as $\mathcal{H}_1, \dots, \mathcal{H}_P$ so that

$$q_2(T_2^{-1}(\boldsymbol{\xi})) := \sum_{p=1}^{P} \hat{\boldsymbol{\eta}}_p \mathbf{1}[T_2^{-1}(\boldsymbol{\xi}) \in \mathcal{H}_p]$$

where $\hat{\boldsymbol{\eta}}_p := T_2^{-1}(\hat{\boldsymbol{\xi}}_p)$.

- The distortion of q_2 admits the following decomposition:

$$w_2(q_2) = \sum_{p=1}^{P} w_{2,p}(q_2)\mu_{\boldsymbol{\xi}}(T_2^{-1}(\mathcal{H}_p))$$

where $w_{2,p}(q_2) := \mathbb{E}[\|T_2^{-1}(\boldsymbol{\xi}) - q_2(T_2^{-1}(\boldsymbol{\xi}))\|^2 \mid T_2^{-1}(\boldsymbol{\xi}) \in \mathcal{H}_p]$.

Computation of stationary quantizers

- The computation of $\mu_{\boldsymbol{\xi}}(T_2^{-1}(\mathcal{H}_p))$ and $w_{2,p}(q_2)$ is intractable.
- In practice, we use an empirical measure of distortion. Given an n_s-sample $\kappa_1, \dots, \kappa_{n_s}$ of i.i.d. realizations of the coefficient field, we compute $\boldsymbol{\xi}_s := \hat{P}_1^{-1}(T_1^{-1}\kappa_s)$ for $s = 1, \dots, n_s$ and approximate $w_2(q_2)$ with

$$w_2^{(n_s)}(q_2) := \frac{1}{n_s}\sum_{s=1}^{n_s} \|T_2^{-1}(\boldsymbol{\xi}_s) - q_2(T_2^{-1}(\boldsymbol{\xi}_s))\|^2$$

which is also given by

$$w_2^{(n_s)}(q_2) = \sum_{p=1}^{P} f_{2,p}^{(n_s)} w_{2,p}^{(n_s)}(q_2) \;\text{ where }\; f_{2,p}^{(n_s)} := \frac{1}{n_s}\sum_{s=1}^{n_s} \mathbf{1}[T_2^{-1}(\boldsymbol{\xi}_s) \in \mathcal{H}_p]$$

is the empirical measure of $\mathcal{H}_p$ associated with $\boldsymbol{\xi}_1, \dots, \boldsymbol{\xi}_{n_s}$, and

$$w_{2,p}^{(n_s)}(q_2) := \frac{1}{f_{2,p}^{(n_s)} n_s}\sum_{s=1}^{n_s} \|T_2^{-1}(\boldsymbol{\xi}_s) - q_2(T_2^{-1}(\boldsymbol{\xi}_s))\|^2 \mathbf{1}[T_2^{-1}(\boldsymbol{\xi}_s) \in \mathcal{H}_p].$$

- Several algorithms compute stationary quantizers q_2 on the basis of these empirical measures, e.g., k-means, competitive learning vector quantization (CLVQ), ...

Choices of the map $T_2(\boldsymbol{\xi})$

- Upon defining the map $T_2 : \mathbb{R}^m \to \mathbb{R}^m$, we introduce some control over the design of q_2 and its underlying codebook, as well as of q.

- First, we aim to define T_2 so as to minimize the $L^2(\Omega)$-distortion of $\hat{T}_1^{-1}\kappa$. By orthonormality of the eigenfunctions of the KL expansion, we have:

$$\|\hat{T}_1^{-1}\kappa(\cdot,\theta)\|_\Omega^2 = \left\| \sum_{k=1}^{m} \lambda_k^{1/2}\Phi_k(x)\xi_k(\theta) \right\|_\Omega^2 = \sum_{k=1}^{m} \lambda_k \xi_k(\theta)^2 = \boldsymbol{\xi}(\theta)^T \boldsymbol{\Lambda} \boldsymbol{\xi}(\theta)$$

which can be recast as $\|\boldsymbol{\Lambda}^{1/2}\boldsymbol{\xi}(\theta)\|^2$ for all $\theta \in \Theta$. Consequently, the map $T_2^{-1} : \boldsymbol{\chi} \mapsto \boldsymbol{\Lambda}^{1/2}\boldsymbol{\chi}$ is such that $\|T_2^{-1}\boldsymbol{\xi}(\theta)\|^2 = \|\hat{T}_1^{-1}\kappa(\cdot,\theta)\|_\Omega^2$.

- Second, we aim at designing stationary quantizers with constant frequencies. We consider

$$T_2^{-1} : \boldsymbol{\chi} \mapsto \boldsymbol{\Lambda}^{1/2} F_\xi \circ \boldsymbol{\chi} \;\; \text{where} \;\; F_\xi \circ \boldsymbol{\chi} = \begin{bmatrix} F_\xi(\chi_1) \\ \vdots \\ F_\xi(\chi_m) \end{bmatrix}$$

where $F_\xi(\chi) = \Pr[\xi \leq \chi]$. Our experiments show that this choice of T_2^{-1} yields stationary quantizers q_2 with $f_1 \approx \cdots \approx f_P$.

Quantizations based on deterministic grids

- We want a quantizer for which the number m of KL modes considered depends on the number P of preconditioners.
- To indicate the number m of KL modes in the quantization, we write $q_2^{(m)}$.
- For $m = 1$, we use

$$q_2^{(1)}(\xi) = T_2^{-1}(0)\mathbf{1}[-s/2 \leq \xi < s/2] + T_2^{-1}(-s)\mathbf{1}[\xi < s/2] + T_2^{-1}(s)\mathbf{1}[s/2 \leq \xi]$$

so as to provide symmetric design. Moreover, in order to have constant attribution frequencies, we let $s = 2F_\xi^{-1}(2/3) \approx 0.8614$.
- For higher numbers m of KL modes, we have

$$q_2^{(m)}(\boldsymbol{\xi}) = \sum_{p=0}^{2^m} T_2^{-1}(\hat{\boldsymbol{\xi}}_p)\mathbf{1}[T_2^{-1}(\boldsymbol{\xi}) \in \mathcal{H}_p]$$

where $\mathcal{H}_0, \ldots, \mathcal{H}_{2^m}$ form a Voronoi partition of $T_2^{-1}(\mathbb{R}^m)$ and are given such that

$$\mathcal{H}_p \subset \left\{T_2^{-1}(\boldsymbol{\xi}),\ \boldsymbol{\xi} \in \mathbb{R}^m,\ \|\boldsymbol{\xi} - \hat{\boldsymbol{\xi}}_p\| \leq \|\boldsymbol{\xi} - \hat{\boldsymbol{\xi}}_q\|,\ q = 0, \ldots, 2^m\right\}$$

with centers $\hat{\boldsymbol{\xi}}_0, \ldots, \hat{\boldsymbol{\xi}}_{2^m}$.

Quantizations based on deterministic grids

Algorithm 1 GetGridCoordinates(s, m)

Require: Grid parameter s,
Number of KL modes m

Ensure: Centroids $T_2^{-1}(\hat{\boldsymbol{\xi}}_0), \ldots, T_2^{-1}(\hat{\boldsymbol{\xi}}_{2^m})$ of quantizer $q_2^{(m)}$.

$\hat{\boldsymbol{\xi}}_0 := \mathbf{0}$
$p := 1$
for $\hat{\xi}_{p,1} \in (-s, s)$ **do**
  for $\hat{\xi}_{p,2} \in (-s, s)$ **do**
    $\vdots$
    for $\hat{\xi}_{p,m} \in (-s, s)$ **do**
      $\hat{\boldsymbol{\xi}}_p := [\hat{\xi}_{p,1}, \ldots, \hat{\xi}_{p,m}]^T$
      $p := p + 1$
    end for
  end for
end for
return $T_2^{-1}(\hat{\boldsymbol{\xi}}_0), \ldots, T_2^{-1}(\hat{\boldsymbol{\xi}}_{2^m})$

Attribution frequencies for different stationary quantizers

Let $m = 2$ and $\boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_2)$. We compute stationary quantizers q_2 of $\mathbb{R}^2$ with different T_2.

$$T_2^{-1}(\boldsymbol{\xi}) = \boldsymbol{\Lambda}^{1/2}\boldsymbol{\xi}$$

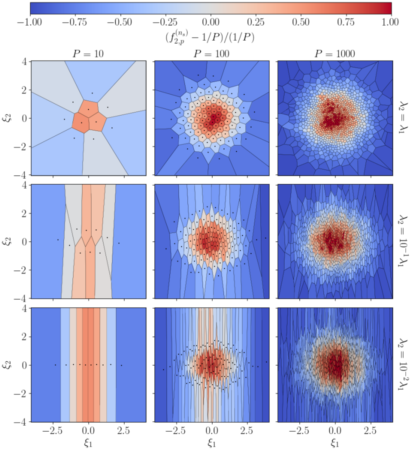

$$T_2^{-1}(\boldsymbol{\xi}) = \boldsymbol{\Lambda}^{1/2} F_\xi \circ \boldsymbol{\xi}$$

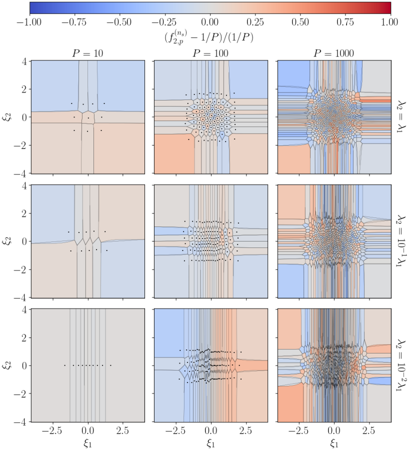

Effect of preconditioning strategies from different stationary quantizers and deterministic grids

- Let $T_1^{-1}\kappa := \log \kappa$ be a Gaussian process with $C(x, x') = \exp\left(\frac{-\|x-x'\|^2}{0.1^2}\right)$.
- $\mathbb{E}[J]$ is estimated with 100,000 realizations and $\mathbf{A}$ is 100,000-dimensional.

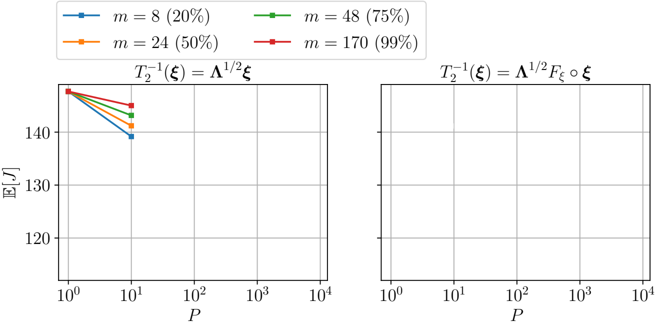

Effect of preconditioning strategies from different stationary quantizers and deterministic grids

- Let $T_1^{-1}\kappa := \log \kappa$ be a Gaussian process with $C(x,x') = \exp\left(\frac{-\|x-x'\|^2}{0.1^2}\right)$.
- $\mathbb{E}[J]$ is estimated with 100,000 realizations and $\mathbf{A}$ is 100,000-dimensional.

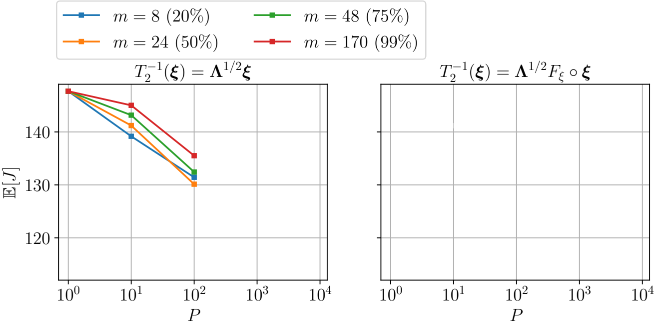

Effect of preconditioning strategies from different stationary quantizers and deterministic grids

- Let $T_1^{-1}\kappa := \log \kappa$ be a Gaussian process with $C(x, x') = \exp\left(\frac{-\|x-x'\|^2}{0.1^2}\right)$.
- $\mathbb{E}[J]$ is estimated with 100,000 realizations and $\mathbf{A}$ is 100,000-dimensional.

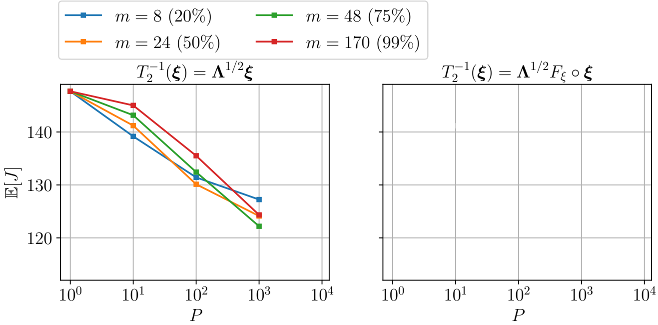

- For some m, $\mathbb{E}[J]$ stagnates passed some value of P. The smaller m the faster it happens.

Effect of preconditioning strategies from different stationary quantizers and deterministic grids

- Let $T_1^{-1}\kappa := \log \kappa$ be a Gaussian process with $C(x, x') = \exp\left(\frac{-\|x-x'\|^2}{0.1^2}\right)$.
- $\mathbb{E}[J]$ is estimated with 100,000 realizations and $\mathbf{A}$ is 100,000-dimensional.

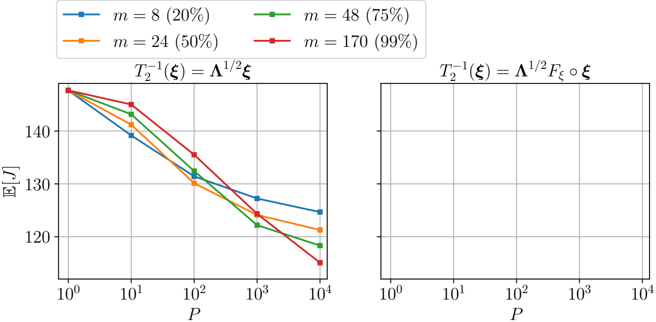

- For some m, $\mathbb{E}[J]$ stagnates passed some value of P. The smaller m the faster it happens.

Effect of preconditioning strategies from different stationary quantizers and deterministic grids

- Let $T_1^{-1}\kappa := \log\kappa$ be a Gaussian process with $C(x,x') = \exp\left(\frac{-\|x-x'\|^2}{0.1^2}\right)$.
- $\mathbb{E}[J]$ is estimated with 100,000 realizations and $\mathbf{A}$ is 100,000-dimensional.

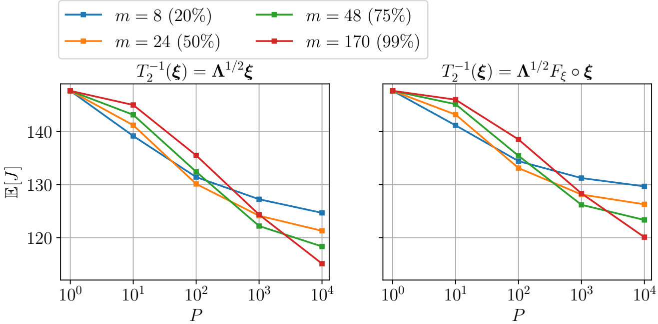

- For some m, $\mathbb{E}[J]$ stagnates passed some value of P. The smaller m the faster it happens.
- A similar behavior is observed for $T_2^{-1}(\boldsymbol{\xi}) = \boldsymbol{\Lambda}^{1/2}F_\xi \circ \boldsymbol{\xi}$, but with larger values of $\mathbb{E}[J]$.

Effect of preconditioning strategies from different stationary quantizers and deterministic grids

- Let $T_1^{-1}\kappa := \log\kappa$ be a Gaussian process with $C(x,x') = \exp\left(\frac{-\|x-x'\|^2}{0.1^2}\right)$.
- $\mathbb{E}[J]$ is estimated with 100,000 realizations and $\mathbf{A}$ is 100,000-dimensional.

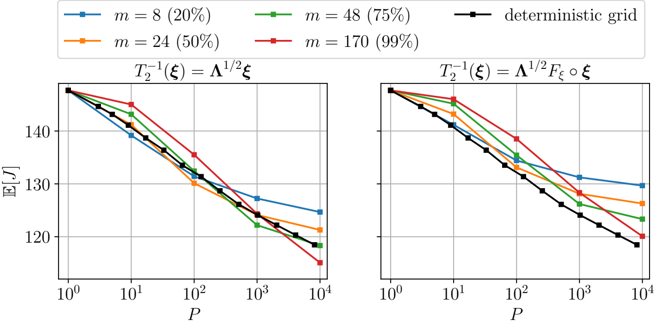

- For some m, $\mathbb{E}[J]$ stagnates passed some value of P. The smaller m the faster it happens.
- A similar behavior is observed for $T_2^{-1}(\boldsymbol{\xi}) = \boldsymbol{\Lambda}^{1/2}F_\xi \circ \boldsymbol{\xi}$, but with larger values of $\mathbb{E}[J]$.
- Using the deterministic grid prevents the stagnation of $\mathbb{E}[J]$.

Distribution of cumulated number of solver iterations for different stationary quantizers

- Let $\hat{T}_1^{-1}\kappa$ have $m = 8$ KL modes (for 20% energy) and P=1,000.
- The realizations of the simulation are denoted by $\hat{\Theta} \in \Theta$ with the partition $\hat{\Theta}_1, \ldots, \hat{\Theta}_P$.
- We are interested by number of linear solves n_p per preconditioner, and the cumulated number of solver iterations $\sum_{\theta\in\hat{\Theta}_p} J(\theta)$.

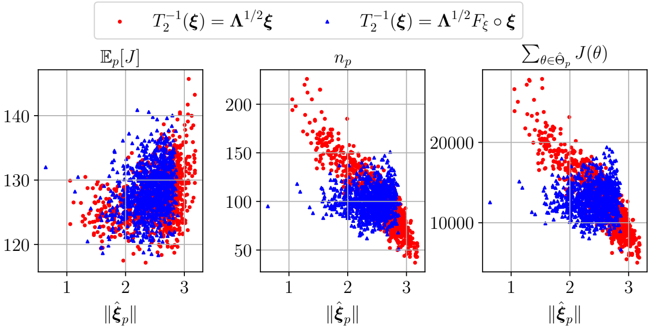

- The strategy with $T_2^{-1} = \boldsymbol{\Lambda}^{1/2}F_\xi \circ \boldsymbol{\xi}$ is more load balanced.

Conclusions and perspectives

Conclusions:

- A lot of improvement can be made compared to using a single constant preconditioner.
- The optimal energy level of the approximating coefficient field of a preconditioning strategy based on stationary quantizers depends on P. Selecting an optimal m requires some preprocessing.
- For sequential simulations, the average number of solver iterations is minimized by letting $T_2^{-1} := \mathbf{\Lambda}^{1/2}\boldsymbol{\xi}$.
- For distributed simulations, the most balanced distribution of the cumulated number of solver iterations is obtained by letting $T_2^{-1} := \mathbf{\Lambda}^{1/2}F_{\xi} \circ \boldsymbol{\xi}$.
- Using a deterministic grid such that the approximating coefficient field has an energy which increases with P prevents stagnation of $\mathbb{E}[J]$ and does not require preprocessing.

Future endeavor:

- Speeding-up the setup of AMG preconditioners for random coefficient fields when using fixed meshes and discretization.

Effect of truncation of the approximating KL expansion on realization-dependent ideal preconditioning

- Let $T_1^{-1}\kappa := \log\kappa$ be a Gaussian process with $C(x,x') = \exp\left(\frac{-\|x-x'\|^2}{0.1^2}\right)$.
- Relative energy is the variance of $\hat{T}_1^{-1}\kappa$ given by $\sum_{k=1}^{m}\lambda_k$.
- Discretization with 100,000 DoFs.

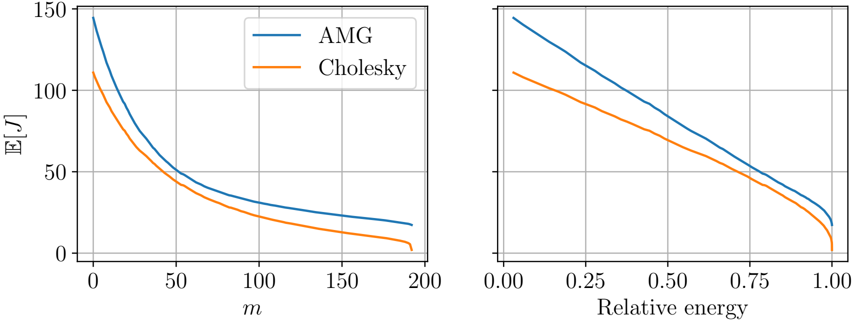

- AMG preconditioners are nearly as effective as Cholesky factorizations, especially for large values of relative energy.
- Nearly linear dependence of $\mathbb{E}[J]$ on relative energy.

Local interpolation of preconditioners

- Let us consider the case of distributed simulations in which the m-th node (out of $M \leq P$) stores P_m preconditioners $\mathbf{M}^{-1}(\hat{\boldsymbol{\xi}}_1^{(m)}), \dots, \mathbf{M}^{-1}(\hat{\boldsymbol{\xi}}_{P_m}^{(m)})$ in memory.
- When $\boldsymbol{\xi}$ is drawn close to a centroid of the m-th node, we follow the work of Zahm and Nouy (2016) and leverage the availability of local preconditioners to approximate $\mathbf{M}^{-1}(\boldsymbol{\xi})$ with an interpolation of the form

$$\hat{\mathbf{M}}_m^{-1}(\boldsymbol{\xi}) = \sum_{p=1}^{P_m} \alpha_p^{(m)}(\boldsymbol{\xi}) \mathbf{M}^{-1}(\hat{\boldsymbol{\xi}}_p^{(m)}).$$

- An ideal choice for $\alpha_1^{(m)}, \dots, \alpha_{P_m}^{(m)} \in \mathbb{R}$ is to minimize the condition number of $\hat{\mathbf{M}}_m^{-1}(\boldsymbol{\xi})\mathbf{A}(\boldsymbol{\xi})$. This, however, is a Clarke regular pseudoconvex optimization problem which is not worth solving for every realization of $\boldsymbol{\xi}$.

Zahm, Olivier, and Anthony Nouy. "Interpolation of inverse operators for preconditioning parameter-dependent equations." SIAM Journal on Scientific Computing 38.2 (2016): A1044-A1074.

Local interpolation of precondtioners

- Another more computationally feasible alternative is to minimize the Frobenius norm $\|\mathbf{I}-\hat{\mathbf{M}}_m^{-1}(\boldsymbol{\xi})\mathbf{A}(\boldsymbol{\xi})\|_F$. This leads to solving

$$\mathbf{B}(\boldsymbol{\xi})\begin{bmatrix}\alpha_1^{(m)}(\boldsymbol{\xi})\\ \vdots \\ \alpha_{P_m}^{(m)}(\boldsymbol{\xi})\end{bmatrix}=\begin{bmatrix}\mathrm{tr}\left(\mathbf{M}^{-1}(\hat{\boldsymbol{\xi}}_1^{(m)})\mathbf{A}(\boldsymbol{\xi})\right)\\ \vdots \\ \mathrm{tr}\left(\mathbf{M}^{-1}(\hat{\boldsymbol{\xi}}_{P_m}^{(m)})\mathbf{A}(\boldsymbol{\xi})\right)\end{bmatrix}$$

where $\mathbf{B}(\boldsymbol{\xi})$ has components

$$B_{pq}(\boldsymbol{\xi})=\mathrm{tr}\left(\left(\mathbf{M}^{-1}(\hat{\boldsymbol{\xi}}_p^{(m)})\mathbf{A}(\boldsymbol{\xi})\right)^T\mathbf{M}^{-1}(\hat{\boldsymbol{\xi}}_q^{(m)})\mathbf{A}(\boldsymbol{\xi})\right)\ ,\quad (p,q)\in[1,P_m]^2.$$

- The computation of the matrices $\mathbf{M}^{-1}(\hat{\boldsymbol{\xi}}_p)\mathbf{A}(\boldsymbol{\xi})$ for $p=1,\dots,P_m$ requires large numbers of preconditioner applications whose cost may surpass the gain obtained by local interpolation. Random sketching is used to significantly reduce the number of necessary preconditioner applications.

- **Results**: we applied Shepard interpolation and the minimizer presented here with and without random sketching to build local interpolations of preconditioners. All of our attempts failed to improve convergence.