





Theory of SEM discretisation and code implementation

Niclas Jansson PDC Centre for HPC Based on material by: Paul Fischer (UIUC), Aleks Obabko (ANL), Philipp Schlatter (KTH), etc...





- Spectral Element Method overview
- Solver outline
- Code implementation

Spectral Element Method







- High-order numerical methods are well suited for accurate, and efficient simulations of turbulent flows
 - Fast convergence
 - Small diffusion and dispersion errors
 - Better data volume-over-surface ratio
- Spectral elements allows for complex geometries

Exponential Convergence

- Converges exponentially fast with Nfor smooth solutions.
- Four orders of magnitude reduction in error when doubling the resolution
- For a given error,
 - Fewer grid points
 - Smaller memory footprint
 - Smaller communication volume

Exact Navier-Stokes Solution (Kovazsnay '48)

1**σ**¹

10

 $1\sigma^{1}$







Highly Scalable Methods

- High-order at low cost
 - Excellent vectorization, dense operators (Level 3 BLAS)
 - Fast, scalable coarse-grid solvers (AMG, XX^T)
 - Domain partitioned into quadrilateral elements
 - Loosely coupled elements (C0 continuity)
- Scales up to 1,000,000 processes on BG/Q
 - 2.95M elements with ninth order polynomials
 - 524 288 cores on ALCF BG/Q Mira;
 - Full scale parallel efficiency 0.6
 - two processes/core, 2000 points/rank







Poisson's equation with homogenous Dirichlet boundary conditions,

$$\begin{aligned}
-\nabla^2 u &= f & \text{in } \Omega \\
u &= 0 & \text{on } \partial \Omega.
\end{aligned}$$

Find $u \in V \subset H_0^1$ such that,

$$\int_{\Omega} \nabla u \nabla v \, d\Omega = \int_{\Omega} f \, v \, d\Omega \quad \forall v \in V.$$

Discretise $\Omega = \bigcup_{e=1}^{E} \Omega^{e}$, and let V^{N} be some approximation space where we search for a solution.



• Seek an approximate solution in V^N with "simple" basis functions $\varphi_i(x)_j^N = \{ \varphi_1(x), \varphi_2(x), \dots, \varphi_n(x) \}$

$$\Rightarrow u(x) = \sum_{i=1}^{n} \xi_i \varphi_i(x)$$

- N-th order polynomials, equidistant points
- We aim for high-order, how large can N be?
- Unstable for high order (Runge's phenomenon)



High-order spatial discretisation

• N-th order Legendre-Lagrange polynomials $h_i(\xi)$,

$$h_i(\xi) = \frac{N(N+1)^{-1} (1-\xi^2) L'_N(\xi_i)}{(\xi-\xi_i) L_N(\xi_i)} \text{ for } \xi \in [-1,1]$$

- Gauss-Lobatto-Legendre quadrature points ξ_i
- N-th Legendre polynomials L_N •
- New ansatz, tensor product polynomials $u(x, y) = \sum_{i=0}^{N} \sum_{j=0}^{N} u_{ij} h_i(x) h_j(y)$ a_{ii}
- Linear system

$$Au = f$$





High-order spatial discretisation

- Assembly of Stiffness matrix, add element stiffness matrices $A_{i,j}^k$ to the global stiffness matrix $A_{i,j}^{u_{04}, u_{14}, \dots, u_{24}, \dots, u_{34}, \dots, u_{34}}$
 - for $k \in T$ for $i = 1 \dots n$ for $j = 1 \dots n$ $A_{\mathbb{I}(i),\mathbb{I}(j)} += A_{i,j}^k$
- Sparse system with linear elements
- How large is $A_{i,j}^k$ with our ansatz?
 - Quite large with $O(N^4)$ non-zeros ($O(N^6)$ in 3D!)
- Not feasible to form *A*!









• Always work with the unassembled matrix A_L



- Data replicated in u_L (shared points)
- We only need the action $A_L u_L$ inside iterative solvers (why?)
- Can we afford to compute A_L? (memory?, compute wise?)

Matrix-Free Operator Evaluation

- Poisson's element stiffness matrix: $\int_{e} \nabla u \nabla v \, de$
- Tensor product polynomials

$$u(x, y) = \sum_{i=0}^{N} \sum_{j=0}^{N} u_{ij} h_i(x) h_j(y)$$

• Partial derivatives evaluated as matrix-matrix products

$$\frac{\partial u}{\partial r}\Big|_{\xi_i\xi_j} = \sum_{p=0}^N u_{pj} \frac{\partial u}{\partial r}\Big|_{\xi_i} = \sum_p \widehat{D}_{ip} u_{pj} = D_r u$$

- Precompute *D* for all directions
- Can we reuse the same partial derivative for all elements?







Mapped geometries

- All evaluations performed on the reference element
 - Invertible map from physical domain to reference domain
- Geometry, same tensor-product as the solution $\mathbf{x}(r,s) = \sum_{i=0}^{N} \sum_{j=0}^{N} x_{ij} h_i(r) h_j(s)$
- Partial derivatives of u using the chain rule (in \mathbb{R}^d)

$$\frac{\partial u}{\partial x_k} = \sum_{k=1}^d \frac{\partial u}{\partial r_i} \frac{\partial r_i}{\partial x_k}$$

• Group all geometric factors $\frac{\partial r_i}{\partial x_k}$

$$G_{ij}(r) = \sum_{k=1}^{d} \frac{\partial r_i}{\partial x_k} \frac{\partial r_j}{\partial x_k} J(r)$$

Jacobian from mapping







Mapped geometries

- Best to use affine (i.e., linear) transformations in order to preserve underlying GLL spacing for stability and accurate quadrature.
- Avoid singular corners ${\sim}180^{\circ}$ or ${\sim}0^{\circ}$
- Avoid high-aspect-ratio cells, if possible





Mesh anisotropy



A common refinement scenario (somewhat exaggerated):



Refinement propagation leads to

- unwanted elements in far-field
- high aspect-ratio cells that are detrimental to iterative solver performance (F. JCP'97)

Matrix-Free Operator Evaluation

EXCELLERAT

Generalized element stiffness matrix

$$\nabla u \nabla v \ de \approx \begin{pmatrix} D_R v^e \\ D_S v^e \end{pmatrix}^T \begin{pmatrix} G_{rr} & G_{rs} \\ G_{rs} & G_{ss} \end{pmatrix} \begin{pmatrix} D_r u^e \\ D_S u^e \end{pmatrix}$$
$$= (v^e) T \begin{pmatrix} D_R \\ D_S \end{pmatrix}^T \begin{pmatrix} G_{rr} & G_{rs} \\ G_{rs} & G_{ss} \end{pmatrix} \begin{pmatrix} D_r \\ D_S \end{pmatrix} u^e$$
$$= (v^e)^T D^T G^e D u^e$$

• The entire bilinear form for Poisson

$$\int_{\Omega} \nabla u \nabla v \, d\Omega = \sum_{e=1}^{E} (v^e)^T D^T G^e D u^e = \sum_{e=1}^{E} (v^e)^T A^e u^e$$

- Recap: through use of chain rule + GLL quadrature:
 - Matrix-free operator evaluation.
 - Operation count is only $O(N^4)$ not $O(N^6)$ (in 3D) using $D_R = (I \otimes I \otimes D)$
 - Memory access is 7n (G_{rr} , G_{rs} , etc., are diagonal)
 - Work is dominated by matrix-matrix products involving D_r , D_s , etc.

Direct stiffness summation

- How to ensure continuity?
 - Data replicated in $u_L = \{u^e\}_{e=1}^E$
- Boolean gather Q^T and scatter matrix Q
 - Mapping from local to global $a(u, v) = \sum_{e=1}^{E} (v^e)^T A^e u^e = (\mathbf{Q}v)^T A_L \mathbf{Q}u = v^T A u.$
- Q nor Q^T formed, only the action QQ^T is used
 - Matrix-vector product $w = Au \Rightarrow w_L = QQ^T A_L u_L$
 - Combined gather-scatter operation
 - Local u^e_{ijk} corresponds to global index $u_{\hat{\iota}}$
- Unit depth stencil for all N!



Algorithm 1 The gather-scatter kernels QQ^T .

Function $u = Q^T u_L$	Function $u_L = Qu$
for $e = 1, 2, E$ do	for $e = 1, 2, E$ do
for $i, j, k = 1, 2, n$ do	for $i, j, k = 1, 2, n$ do
$\hat{i} \leftarrow loctglb(i, j, k, e)$	$\hat{i} \leftarrow loctglb(i, j, k, e)$
$u_{\hat{i}} \leftarrow u_{\hat{i}} + u^{e}_{ijk}$	$u^e_{ijk} \leftarrow u_{\hat{i}}$
end for	end for
end for	end for

Scalable linear solvers

- Key considerations:
 - Bounded iteration counts as $N \rightarrow \infty$
 - Cost that does not scale prohibitively with number of processors, P
- Our choice:
 - CG or GMRES, preconditioned with multilevel additive Schwarz
 - Coarse-grid solve:
 - XX^T projection-based solver
 - single V-cycle of well-tuned AMG (in house or using HYPRE)
 - Projection in time extract available temporal regularity in {<u>p</u>ⁿ⁻¹, <u>p</u>ⁿ⁻², ..., <u>p</u>^{n-k}}





Time stepping

EXCELLERAT

• Navier-Stokes time advancement:

$$\frac{\partial \vec{u}}{\partial t} + \vec{u} \cdot \nabla \vec{u} = -\nabla P + \frac{1}{Re} \nabla^2 \vec{u}$$

 $-\nabla . \vec{u} = 0$

- Nonlinear term: *explicit*
 - *k* th-order backward difference formula / extrapolation
 - characteristics (Pironneau '82, MPR '90)
- Stokes problem: pressure/viscous decoupling:
 - 3 Helmholtz solves for velocity
 - (consistent) Poisson equation for pressure (computationally dominant)

("easy" w/ Jacobi-precond. CG)



$$\begin{aligned} Hu &= \left(\frac{1}{Re}A - \frac{b_0}{\Delta t}B\right)[u] = f_u \ (h_1 = 1, h_2 = -\frac{b_0}{\Delta t}) \\ Hp &= (A)[p] = f_p \ (h_1 = 1, h_2 = 0) \end{aligned}$$