Performance

CFD Examples

# Performance-portable interfaces for applications in Scientific Computing

Valeria Barra<sup>1</sup>

collaborators: Jeremy Thompson<sup>1,2</sup>, Leila Ghaffari<sup>1</sup>, Jed Brown<sup>1</sup>,

Yohann Dudouit<sup>3</sup>

<sup>1</sup> Department of Computer Science, CU Boulder

<sup>2</sup> Department of Applied Math, CU Boulder

<sup>3</sup> Lawrence Livermore National Laboratory

(remotely for) SIAM-AN20

July 14, 2020







Performance

CFD Examples

University of Colorado

Boulder

#### Motivation: Why matrix-free? And why high-order?



Memory bandwidth (left) and FLOPs per dof (right) to apply a Jacobian matrix, obtained from discretizations of a b-variable PDE system. Assembled matrix vs matrix-free (exploits the tensor product structure by either storing at q-points or computing on the fly)

[Courtesy: Jed Brown]

Introduction	Performance	CFD Examples
0000000	0000	000000000000

#### Overview

- For decades, high-order numerical methods have been considered too expensive
- A sparse matrix is no longer a good representation for high-order operators. In particular, the Jacobian of a nonlinear operator is known to rapidly lose sparsity as the order is increased
- libCEED uses a matrix-free operator description, based on a purely algebraic interface, where user only specifies action of weak form operators
- libCEED operator representation is optimal with respect to the FLOPs needed for its evaluation, as well as the memory transfer needed for operator evaluations (matvec)
  - Matrix-free operators that exploit tensor-product structures reduce the work load from  $O(p^6)$  (for sparse matrix) to  $O(p^4)$ , and memory storage from  $O(p^6)$  to  $O(p^3)$
- We demonstrate the usage of libCEED, its integration with other packages, and some PETSc application examples

# libCEED: the library within CEED (Center for Efficient Exascale Discretizations)

- Primary target: high-order finite/spectral element methods (FEM/SEM) exploiting tensor-product structure
- Open source (BSD-2 license) C library with Fortran and Python interfaces
- Releases: v0.1 (January 2018), v0.2 (March 2018), v.0.3 (September 2018), v0.4 (March 2019), v0.5 (September 2019), v0.6 (March 2020)



For latest release:

Kolev T., Fischer P., Abdelfattah A., Ananthan S., **Barra V.**, Beams N., Brown J. et al., *CEED ECP Milestone Report: Improve performance and capabilities of CEED-enabled ECP applications on Summit/Sierra* (2020, March 31<sup>st</sup>) DOI: http://doi.org/10.5281/zenodo.3860804

University of Colorado Boulder

Performance

CFD Examples

#### libCEED backends



University of Colorado Boulder

## libCEED decomposition





#### $A = \mathcal{P}^T \mathcal{E}^T \mathbf{B}^T D \mathbf{B} \mathcal{E} \mathcal{P}$





Performance

CFD Examples

#### Point-wise QFunctions

User-defined QFunctions:

 $-\nabla\cdot(\kappa(x)\nabla u)$ 



Introduction
000000000

CFD Examples

#### Point-wise QFunctions

User-defined QFunctions:

 $-\nabla\cdot (\kappa(x)\nabla u)$ 

or from libCEED's Gallery:

 $\nabla\cdot(\nabla\mathfrak{u})$ 

are point-wise functions that do not depend on element resolution, topology, or basis order



Performance

CFD Examples

#### Point-wise QFunctions

User-defined QFunctions:

 $-\nabla\cdot (\kappa(x)\nabla u)$ 

or from libCEED's Gallery:

$$\nabla \cdot (\nabla \mathfrak{u})$$



are point-wise functions that do not depend on element resolution, topology, or basis order



Introduction
00000000

CFD Examples

### Point-wise QFunctions

User-defined QFunctions:

 $-\nabla \cdot (\kappa(\mathbf{x}) \nabla \mathbf{u})$ 

or from libCEED's Gallery:

 $\nabla \cdot (\nabla \mathfrak{u})$ 

are point-wise functions that do not depend on element resolution, topology, or basis order



Boulder

CFD Examples

#### libCEED's Python interface





Performance

CFD Examples

#### libCEED's Python interface





7/24

Performance

CFD Examples

#### libCEED's Python interface



#### More details:

Barra V., Brown J., Thompson J., Dudouit Y., *High-performance* operator evaluations with ease of use: *libCEED's Python in-*terface, Proceedings of the SciPy 2020 conference (2020, July 12) url: http://conference.scipy.org/proceedings/scipy2020/libceed-paper.html



CFD Examples

#### Docs and tutorials

More info on our Python interface and interactive Jupyter notebook tutorials can be found at:

https://hub.gke.mybinder.org/ user/ceed-libceed-tyuu81m6/lab



Our (very first!) user manual can be found at:

https://libceed.readthedocs.io





8/24

Skylake

Performance ●○○○ CFD Examples

# Performance on Skylake: AVX



Introduction 00000000 Skylake Performance ○●○○ CFD Examples

## Performance on Skylake: libXSMM



Performance

CFD Examples

#### Noether

#### Performance on an AMD EPYC: libXSMM

#### Noether (2x EPYC 7452), gcc-10



Figure: 2x AMD EPYC 7452 (32-core) with gcc-10 compiler. LIBXSMM blocked backend (q = P + 2, P = p + 1) with respect to time (left) and problem size (right)



Performance ○○○● CFD Examples

Noether

#### Preliminary GPU results: MFEM + libCEED



Results by Yohann Dudouit on Lassen (LLNL): CUDA-ref (left) and CUDA-gen (right) backends performance for BP3 on a NVIDIA V100 GPU.

University of Colorado Boulder

# A miniapp: a compressible Navier-Stokes solver

Compressible Navier-Stokes equations in conservative form:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot \mathbf{U} = 0$$
, (1a)

$$\frac{\partial \mathbf{U}}{\partial t} + \nabla \cdot \left( \frac{\mathbf{U} \otimes \mathbf{U}}{\rho} + P \mathbf{I}_3 \right) + \rho g \mathbf{k} = \nabla \cdot \boldsymbol{\sigma} \,, \tag{1b}$$

$$\frac{\partial E}{\partial t} + \nabla \cdot \left( \frac{(E+P)\mathbf{U}}{\rho} \right) = \nabla \cdot (\mathbf{u} \cdot \boldsymbol{\sigma} + k \nabla T) , \qquad (1c)$$



Performance

CFD Examples

# A miniapp: a compressible Navier-Stokes solver

Compressible Navier-Stokes equations in conservative form:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot \mathbf{U} = 0$$
, (1a)

$$\frac{\partial \boldsymbol{U}}{\partial t} + \nabla \cdot \left( \frac{\boldsymbol{U} \otimes \boldsymbol{U}}{\rho} + P \boldsymbol{I}_3 \right) + \rho g \boldsymbol{k} = \nabla \cdot \boldsymbol{\sigma} \,, \tag{1b} \label{eq:eq:eq_alpha}$$

$$\frac{\partial E}{\partial t} + \nabla \cdot \left( \frac{(E+P)\mathbf{U}}{\rho} \right) = \nabla \cdot \left( \mathbf{u} \cdot \boldsymbol{\sigma} + k \nabla T \right) \,, \tag{1c}$$

where 
$$\boldsymbol{\sigma} = \boldsymbol{\mu} (\nabla \boldsymbol{u} + (\nabla \boldsymbol{u})^\mathsf{T} + \boldsymbol{\lambda} (\nabla \cdot \boldsymbol{u}) \mathbf{I}_3) \text{, and}$$



Introduction

## A miniapp: a compressible Navier-Stokes solver

Compressible Navier-Stokes equations in conservative form:

$$rac{\partial 
ho}{\partial t} + 
abla \cdot \mathbf{U} = \mathbf{0}$$
 , (1a)

$$\frac{\partial \boldsymbol{U}}{\partial t} + \nabla \cdot \left( \frac{\boldsymbol{U} \otimes \boldsymbol{U}}{\rho} + P \boldsymbol{I}_3 \right) + \rho g \boldsymbol{k} = \nabla \cdot \boldsymbol{\sigma} \,, \tag{1b} \label{eq:eq:eq_alpha}$$

$$\frac{\partial E}{\partial t} + \nabla \cdot \left( \frac{(E+P)\mathbf{U}}{\rho} \right) = \nabla \cdot \left( \mathbf{u} \cdot \boldsymbol{\sigma} + k \nabla T \right) \,, \tag{1c}$$

where 
$$\boldsymbol{\sigma} = \boldsymbol{\mu} (\nabla \boldsymbol{u} + (\nabla \boldsymbol{u})^\mathsf{T} + \boldsymbol{\lambda} (\nabla \cdot \boldsymbol{u}) \mathbf{I}_3) \text{, and}$$

 $(c_{p}/c_{\nu}-1)\left(\mathsf{E}-\mathbf{U}\cdot\mathbf{U}/(2\rho)-\rho gz\right)=\mathsf{P}\quad\leftarrow \ \mathrm{pressure}$ 

- $\mu \leftarrow dynamic viscosity$
- $g \hspace{.1in} \leftarrow \hspace{.1in} \operatorname{gravitational} \hspace{.1in} \operatorname{acceleration}$
- $k \ \leftarrow \ {\rm thermal \ conductivity}$
- $\lambda \ \ \leftarrow \ \, {\rm Stokes \ hypothesis \ constant}$
- $c_p \hspace{0.1in} \leftarrow \hspace{0.1in} \mathrm{specific \ heat, \ constant \ pressure}$
- $c_\nu \quad \leftarrow \ {\rm specific \ heat, \ constant \ volume}$



ntroduction	Performance	CFD Examples
0000000	0000	00000000000

### Vector form

The system (1) can be rewritten in vector form

$$\frac{\partial \mathbf{q}}{\partial t} + \nabla \cdot \mathbf{F}(\mathbf{q}) = \mathbf{S}(\mathbf{q}), \qquad (2)$$

for the state variables

$$\begin{array}{ll} q = & \left( \begin{array}{c} \rho \\ \mathbf{U} \equiv \rho \mathbf{u} \\ \mathbf{E} \equiv \rho e \end{array} \right) \begin{array}{l} \leftarrow \mbox{ volume mass density} \\ \leftarrow \mbox{ momentum density} \\ \leftarrow \mbox{ energy density} \end{array}$$



(3)

Introduction	Performance	CFD Examples
0000000	0000	00000000000

#### Vector form

The system (1) can be rewritten in vector form

$$\frac{\partial \mathbf{q}}{\partial t} + \nabla \cdot \mathbf{F}(\mathbf{q}) = \mathbf{S}(\mathbf{q}), \qquad (2)$$

 $\begin{array}{ll} \mbox{for the state variables} \\ q = & \left( \begin{array}{c} \rho \\ u \equiv \rho u \\ E \equiv \rho e \end{array} \right) \begin{array}{l} \leftarrow \mbox{ volume mass density} \\ \leftarrow \mbox{ momentum density} \\ \leftarrow \mbox{ energy density} \end{array} \end{array}$ 

(3)

where

$$\begin{split} F(q) = & \left( \begin{array}{c} U \\ (U \otimes U)/\rho + PI_3 - \sigma \\ (E+P)U/\rho - (u \cdot \sigma + k \nabla T) \end{array} \right) \\ S(q) = & - \left( \begin{array}{c} 0 \\ \rho g \hat{k} \\ 0 \end{array} \right) \end{split}$$

University of Colorado Boulder

,

ntroduction	
0000000	

CFD Examples

#### Space discretization

We use high-order finite/spectral elements: high-order Lagrange polynomials over non-uniformly spaced nodes,  $\{x_i\}_{i=0}^p$ , the Legendre-Gauss-Lobatto (LGL) points (roots of the p<sup>th</sup>-order Legendre polynomial P<sub>p</sub>). We let  $\mathbb{R}^3 \supset \Omega = \bigcup_{e=1}^{N_e} \Omega_e$ , with N<sub>e</sub> disjoint hexaedral elements.

The physical coordinates are  $\mathbf{x}=(x,y,z)\in\Omega_e,$  while the reference coords are  $\mathbf{X}=(X,Y,Z)\in\mathbf{I}=[-1,1]^3.$ 



Performance

CFD Examples

#### Space discretization

We use high-order finite/spectral elements: high-order Lagrange polynomials over non-uniformly spaced nodes,  $\{x_i\}_{i=0}^p$ , the Legendre-Gauss-Lobatto (LGL) points (roots of the p<sup>th</sup>-order Legendre polynomial P<sub>p</sub>). We let  $\mathbb{R}^3 \supset \Omega = \bigcup_{e=1}^{N_e} \Omega_e$ , with N<sub>e</sub> disjoint hexaedral elements.

The physical coordinates are  $x=(x,y,z)\in\Omega_e,$  while the reference coords are  $X=(X,Y,Z)\in I=[-1,1]^3.$ 

Define the discrete solution

$$\mathbf{q}_{N}(\mathbf{x},t)^{(e)} = \sum_{k=1}^{P} \psi_{k}(\mathbf{x}) \mathbf{q}_{k}^{(e)}$$
 (4)

with P the number of nodes in the element e.

We use tensor-product bases  $\psi_{kji} = h_i(X)h_j(Y)h_k(Z)$ .



## Strong and weak formulations

The strong form of (3):

$$\int_{\Omega} \nu \left( \frac{\partial \mathbf{q}_{N}}{\partial t} + \nabla \cdot \mathbf{F}(\mathbf{q}_{N}) \right) \, d\Omega = \int_{\Omega} \nu \mathbf{S}(\mathbf{q}_{N}) \, d\Omega \, , \, \forall \nu \in \mathcal{V}_{p} \tag{5}$$

with  $\mathcal{V}_p=\{\nu\in H^1(\Omega_e)\,|\,\nu\in P_p(I), e=1,\ldots,N_e\}.$  Weak form:

$$\int_{\Omega} \nu \frac{\partial \mathbf{q}_{N}}{\partial t} \, d\Omega + \int_{\Gamma} \nu \widehat{\mathbf{n}} \cdot \mathbf{F}(\mathbf{q}_{N}) \, d\Omega - \int_{\Omega} \nabla \nu \cdot \mathbf{F}(\mathbf{q}_{N}) \, d\Omega =$$

$$\int_{\Omega} \nu \mathbf{S}(\mathbf{q}_{N}) \, d\Omega, \, \forall \nu \in \mathcal{V}_{p}$$
(6)



#### Strong and weak formulations

The strong form of (3):

$$\int_{\Omega} \nu \left( \frac{\partial \mathbf{q}_{N}}{\partial t} + \nabla \cdot \mathbf{F}(\mathbf{q}_{N}) \right) \, d\Omega = \int_{\Omega} \nu \mathbf{S}(\mathbf{q}_{N}) \, d\Omega \, , \, \forall \nu \in \mathcal{V}_{p} \tag{5}$$

with  $\mathcal{V}_p=\{\nu\in H^1(\Omega_e)\,|\,\nu\in P_p(I), e=1,\ldots,N_e\}.$  Weak form:

$$\int_{\Omega} \nu \frac{\partial \mathbf{q}_{N}}{\partial t} \, d\Omega + \int_{\Gamma} \nu \widehat{\mathbf{n}} \cdot \mathbf{F}(\mathbf{q}_{N}) \, d\Omega - \int_{\Omega} \nabla \nu \cdot \mathbf{F}(\mathbf{q}_{N}) \, d\Omega = \int_{\Omega} \nu \mathbf{S}(\mathbf{q}_{N}) \, d\Omega, \, \forall \nu \in \mathcal{V}_{p}$$
(6)

Explicit time discretization:

$$q_N^{n+1} = q_N^n + \Delta t \sum_{i=1}^s b_i k_i$$
, (7)

adaptive Runge-Kutta-Fehlberg (RKF4-5) method

Implicit time discretization:

$$\begin{split} f(\boldsymbol{q}_{N}) &\equiv g(t^{n+1}, \boldsymbol{q}_{N}, \dot{\boldsymbol{q}}_{N}) = 0, \\ \dot{\boldsymbol{q}}_{N}(\boldsymbol{q}_{N}) &= \alpha \boldsymbol{q}_{N} + \boldsymbol{z}_{N} \end{split} \tag{8}$$

 $\alpha$ -method

University of Colorado Boulder

CFD Examples

#### Application example: Density current

# A cold air bubble drops by convection in a neutrally stratified atmosphere.

Its initial condition is defined in terms of the Exner pressure,  $\pi(x, t)$ , and potential temperature,  $\theta(x, t)$ , that relate to the state variables via

$$\rho = \frac{P_0}{(c_p - c_\nu)\theta(\mathbf{x}, t)} \pi(\mathbf{x}, t)^{\frac{c_\nu}{c_p - c_\nu}}, \qquad (9a)$$

$$e = c_{\nu}\theta(\mathbf{x}, t)\pi(\mathbf{x}, t) + \mathbf{u} \cdot \mathbf{u}/2 + gz, \qquad (9b)$$

where  $P_0$  is the atmospheric pressure.

BCs: free slip for u, no-flux for mass and energy densities.



Intro	duc	tion
000	000	0000

CFD Examples

#### Density current

# order: p= 10, $\Omega=[0,6000]^2~m\times[0,3000]~m,$ elem. resolution: 500 m, FEM nodes: 893101



University of Colorado Boulder

Performance

CFD Examples

#### Recent Developments: Implicit time-stepping





CFD Examples

## Recent Developments: PHASTA Integration

In collaboration with PHASTA (FastMath) we have worked on libCEED's integration.



[Ref: phasta.scigap.org]



20 / 24

CFD Examples

#### Recent Developments: Stabilization methods

We have added Streamline Upwind (SU) and Streamline Upwind/Petrov-Galerkin (SUPG) stabilization methods to our Navier-Stokes example.

For the advection case:

Not stabilized version.

Stabilized version.



Performance

CFD Examples

#### Recent Developments: BPs on the cubed-sphere

Converted BP1 (Mass operator) & BP3 (Poisson's equation) on the cubed-sphere as a prototype for shallow-water equations solver

$$\frac{\partial \mathbf{u}}{\partial t} = -(\omega + f)\hat{\mathbf{k}} \times \mathbf{u} - \nabla \left(\frac{1}{2}|\mathbf{u}|^2 + g(\mathbf{h} + \mathbf{h}_s)\right)$$
(10a)  
$$\frac{\partial \mathbf{h}}{\partial t} = -\nabla \cdot (\mathbf{h}_0 + \mathbf{h})\mathbf{u}$$
(10b)  
$$(10b)$$

CFD Examples

### Conclusions

- We have showed libCEED's performance portability on several architectures, when integrated with PETSc and MFEM
- We have demonstrated the use of libCEED with PETSc for the numerical high-order solutions of
  - Full compressible Navier-Stokes equations
- We have included implicit time-stepping and SU/SUPG stabilization methods







### Outlook

Ongoing and future work:

- Algorithmic differentiation of Q-Functions
- Ongoing work on CUDA and HIP optimizations
- Complete SWE solver on the cubed-sphere
- We always welcome contributors and users https://github.com/CEED/libCEED





## Outlook

Ongoing and future work:

- Algorithmic differentiation of Q-Functions
- Ongoing work on CUDA and HIP optimizations
- Complete SWE solver on the cubed-sphere
- We always welcome contributors and users https://github.com/CEED/libCEED



Acknowledgements: Exascale Computing Project (17-SC-20-SC)

# Thank you!



CFD Examples

# libCEED backends

/cpu/self/ref/*:	with * reference serial and blocked implementations
/cpu/self/avx/*:	AVX (Advanced Vector Extensions instruction sets)
	with * reference serial and blocked implementations
/cpu/self/xsmm/*:	LIBXSMM (Intel library for small dense/sparse mat-multiply)
	with * reference serial and blocked implementations
/*/occa:	OCCA (just-in-time compilation)
	with *: CPU, GPU, OpenMP (Open Multi-Processing: API),
	OpenCL (framework for CPUs, GPUs, etc.)
/gpu/magma:	CUDA MAGMA (dense Linear Algebra library for GPUs and
	multicore architectures) kernels
/gpu/cuda/*:	CUDA with *: ref (reference pure CUDA kernels),
	reg (CUDA kernels using one thread per element),
	shared, optimized CUDA kernels using shared memory
	gen, optimized CUDA kernels using code generation

Same source code can call multiple CEEDs with different backends. On-device operator implementation with unique interface



24 / 24

Boulder

#### Tensor contractions

Let  $\{x_i\}_{i=0}^p$  denote the LGL nodes with the corresponding interpolants  $\{\psi_i^p\}_{i=0}^p$ . Choose a quadrature rule with nodes  $\{q_i^Q\}_{i=0}^Q$  and weights  $\{w_i^Q\}$ . The basis evaluation, derivative, and integration matrices are  $B_{ij}^{Qp} = \psi_j^p(q_i^Q)$ ,  $D_{ij}^{Qp} = \vartheta_x \psi_j^p(q_i^Q)$ , and  $W_{ij}^Q = w_i^Q \delta_{ij}$ . In 3D:

$\mathbf{D} = \mathbf{D} \lor \lor \mathbf{D} \lor \lor \mathbf{D}$	$\mathbf{B} = \mathbf{B} \otimes$	$B \otimes B$	(11)
---	-----------------------------------	---------------	------

$$\mathbf{D}_0 = \mathbf{D} \otimes \mathbf{B} \otimes \mathbf{B} \tag{12}$$

$$\mathbf{D}_1 = \mathbf{B} \otimes \mathbf{D} \otimes \mathbf{B} \tag{13}$$

$$\mathbf{D}_2 = \mathbf{B} \otimes \mathbf{B} \otimes \mathbf{D} \tag{14}$$

$$\mathbf{W} = \mathcal{W} \otimes \mathcal{W} \otimes \mathcal{W} \tag{15}$$

These tensor-product operations cost  $2(p^3Q + p^2Q^2 + pQ^3)$  and touch only  $O(p^3 + Q^3)$  memory. In the spectral element method, when the same LGL points are reused for quadrature (i.e., a collocated method with Q = p + 1), then  $\mathbf{B} = \mathbf{I}$  and  $\mathbf{D}$  reduces to  $O(p^4)$ .

Performance

CFD Examples

#### Geometry on the sphere



Transform  $\overset{\circ}{\mathbf{x}} = (\overset{\circ}{\mathbf{x}}, \overset{\circ}{\mathbf{y}}, \overset{\circ}{z})$  on the sphere  $\hookrightarrow$  $\mathbf{x} = (x, y, z)$  on the discrete surface  $\hookrightarrow$  $\mathbf{X} = (X, Y) \in \mathbf{I} = [-1, 1]^2$ 

$$\frac{\partial \overset{\circ}{\mathbf{x}}}{\partial \mathbf{X}}_{(3\times 2)} = \frac{\partial \overset{\circ}{\mathbf{x}}}{\partial \mathbf{x}}_{(3\times 3)} \frac{\partial \mathbf{x}}{\partial \mathbf{X}}_{(3\times 2)}$$

$$|\mathbf{J}| = \left| \operatorname{col}_1 \left( \frac{\partial \overset{\circ}{\mathbf{x}}}{\partial \mathbf{X}} \right) \times \operatorname{col}_2 \left( \frac{\partial \overset{\circ}{\mathbf{x}}}{\partial \mathbf{X}} \right) \right|$$

