The Compadre Toolkit for Native Degrees-of-Freedom

Paul Kuberry, Mauro Perego, Nat Trask, Pavel Bochev
Center for Computing Research, Sandia National Laboratories

Coupling Approaches for Next Generation Architectures (CANGA)

Native DoFs

Different codes may employ different discretizations of the same PDE due to different designs, e.g., stabilized vs. compatible

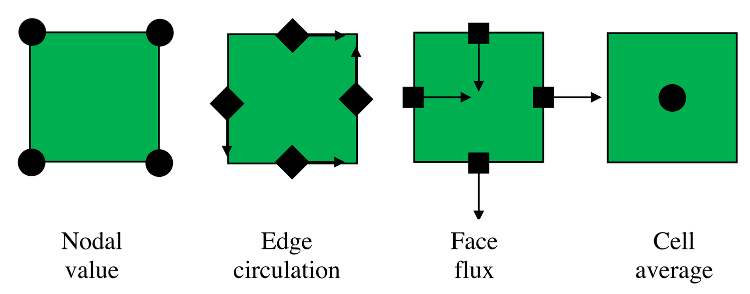

The same field may be represented differently in a coupled multi-physics simulation, e.g.,

- Raviart-Thomas (H(div)) velocity vs. nodal (H1) velocity in Darcy-Stokes coupling

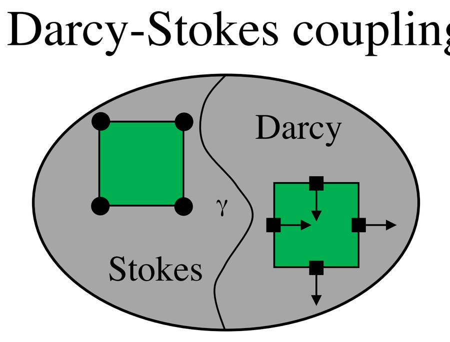

The field may be represented by the same type of discretization but on a different cell shape:

- Raviart-Thomas on tets, Raviart-Thomas on hexes and mimetic difference on polyhedrons

Research Problem

Provide a field reconstruction capability that:

- Accepts any reasonable native DoF type
- Reconstructs accurately a field from a type without having to know the internal reconstruction procedure for each code utilizing this DoF type

Generalized Moving Least Squares

Given a set of sampling functionals λ_j and a vector space V, we seek the element of V that best matches a set of scattered degrees of freedom, and use that to approximate a target functional τ over V

1. (GMLS Approximate) Constrained Optimization formulation:

$$\tau^h_{\bar{\mathbf{x}}}(u) := \sum_{i \in I_{\bar{\mathbf{x}}}} a^i_{\tau_{\bar{\mathbf{x}}}} \lambda_i(u), \quad \{a^i_{\tau_{\mathbf{x}}}\} = \arg\min_{a^i} \sum_{i \in I_{\bar{\mathbf{x}}}} \frac{|a^i|^2}{W(\bar{\mathbf{x}}, \mathbf{x}_i)},$$

$$\left(\tau^h_{\bar{\mathbf{x}}}(p) = \tau_{\bar{\mathbf{x}}}(p)\right) \text{ s. t. } \tau_{\bar{\mathbf{x}}}(p) = \sum_{i \in I_{\bar{\mathbf{x}}}} a^i \lambda_i(p), \ \forall p \in P$$

2. (Practical Recipe) Least Square formulation:

$$\tau^h_{\bar{\mathbf{x}}}(u) := \tau_{\bar{\mathbf{x}}}(p_{\bar{\mathbf{x}},u}), \quad p_{\bar{\mathbf{x}},u} = \arg\min_{p \in P} \sum_{i \in I_{\bar{\mathbf{x}}}} (\lambda_i(u) - \lambda_i(p))^2 \, W(\bar{\mathbf{x}}, \mathbf{x}_i)$$

$$\mathrm{I}_{\bar{\mathbf{x}}} := \{i : W(\bar{\mathbf{x}}, \mathbf{x}_i) > 0\}$$

Example: Point evaluation reconstruction of derivatives from nodal data in 1D.

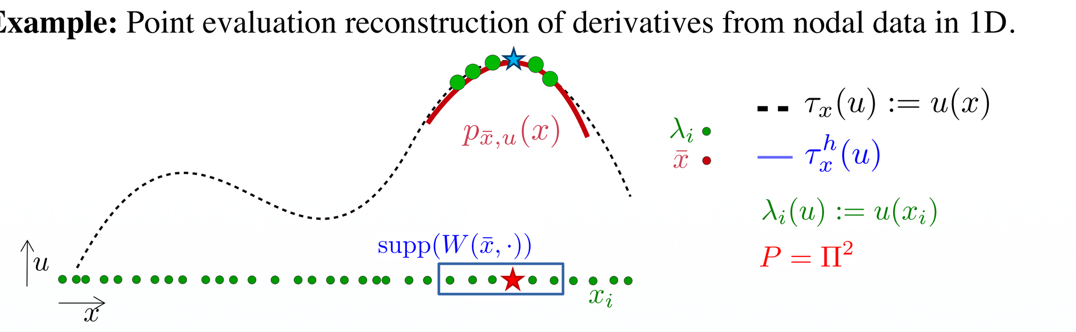

Sampling Functional Choices

GMLS allows for the reconstruction of a vector field from Raviart-Thomas H(div) conforming scalar representations or from Nedelec H(curl) conforming representations by choosing $\boldsymbol{\delta}_x$ (point evaluation) as τ, and λ_j as:

$$u(x_i) \cdot n(x_i) \text{ or } \int_{f_i} u(x) \cdot n(x) \, df_i$$

and

$$u(x_i) \cdot t(x_i) \text{ or } \int_{e_i} u(x) \cdot t(x) \, de_i,$$

respectively.

Compadre Toolkit

The **Compadre Toolkit** provides a massively parallel solution framework for setting up and solving the quadratic programs (QP) defined for meshless discretizations such as Generalized Moving Least Squares (GMLS).

The coefficients generated by the toolkit are amenable to **meshless remap and PDE solution,** allowing users to harness meshless discretizations such as GMLS while executing these parallel **communication-sparse, computationally-dense** kernels on modern architectures.

The only third party library (TPL) needed by the toolkit is Kokkos, a performance portability library produced by Sandia, which allows code to be written once that will target multiple architectures and thread-parallel frameworks such as OpenMP, Pthreads, and Cuda.

Kokkos source code is included and automatically built by the toolkit, effectively reducing the requirements of Compadre to zero TPLs.

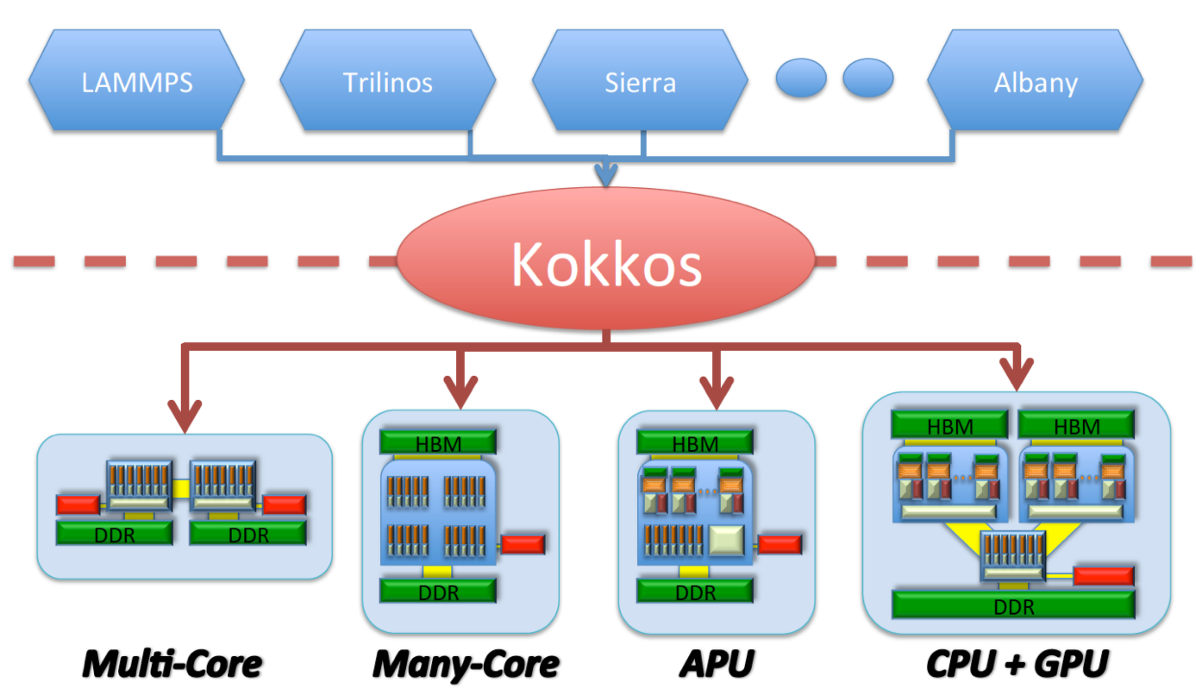

https://github.com/kokkos

Caption: Target heterogeneous computational architectures via Kokkos.

Coupling Climate Codes

The remap capability supports many different discretizations, easily allowing users to couple codes using discretizations including finite elements, finite volumes, finite differences, spectral elements, as well as fields generated by other meshless methods.

Disparate mesh discretizations can be coupled as is required by the Canga SciDAC-BER to demonstrate remap between cell-averaged finite volumes and spectral elements.

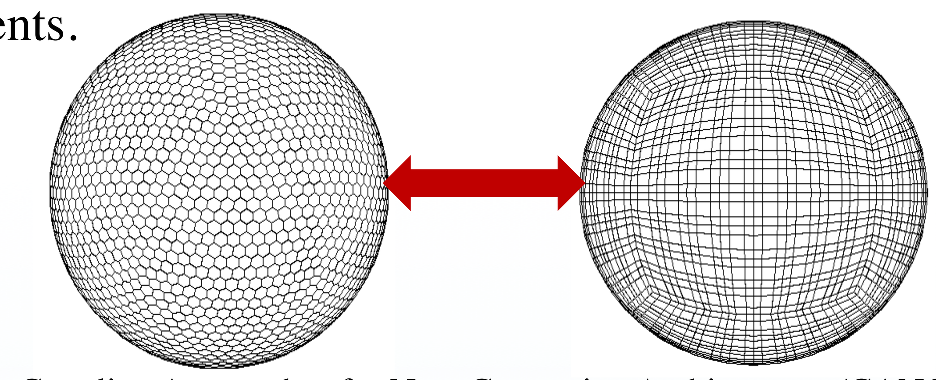

Caption: Coupling Approaches for Next Generation Architectures (CANGA)

Additionally, the Compadre Toolkit comes with an optional Python module generation that allows user to call a limited subset of the GMLS functionality, using their GPU (if selected). The Python package is now available through the Pypi repository and can be installed by running:

```
>> pip install compadre
```

Results

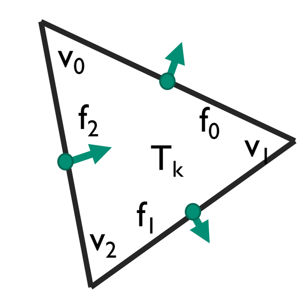

Raviart-Thomas (Face) elements

Low order Raviart-Thomas, data is contained at midpoints of faces $\{f_i\}$ and represents either $u(x_i) \cdot n(x_i)$ or $\int_{f_i} u(x) \cdot n(x) \, df_i$.

Nedelec (Edge) elements

Low order Nedelec, data is contained at midpoints of edges $\{e_i\}$ and represents either $u(x_i) \cdot t(x_i)$ or $\int_{e_i} u(x) \cdot t(x) \, de_i$.

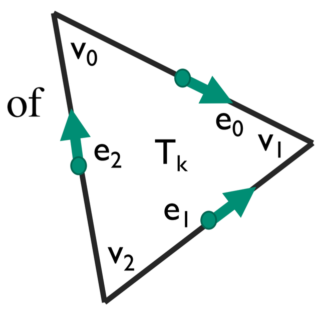

Visual Representation of RT Data

Face normal (Raviart-Thomas) field data

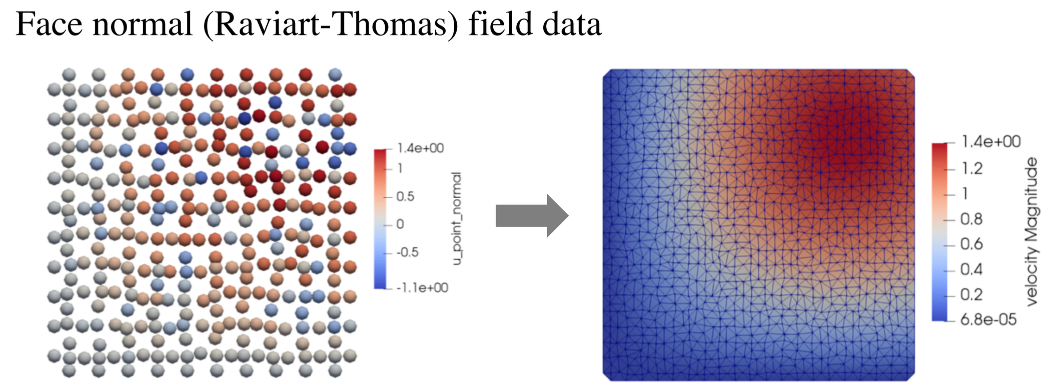

P1 Lagrange field data

Remap of H(div) and H(curl) DoFs

Reconstruction of a vector field represented as Raviart-Thomas and Nedelec DoFs

Exact solution:

$$\vec{v}(x,y) = \begin{pmatrix} \sin(x)\sin(y) \\ -\sin(x)\sin(y) \end{pmatrix}$$

Degree of basis for reconstruction: 4

Case 1
Edges from a quasiuniform mesh

Refinement sequence

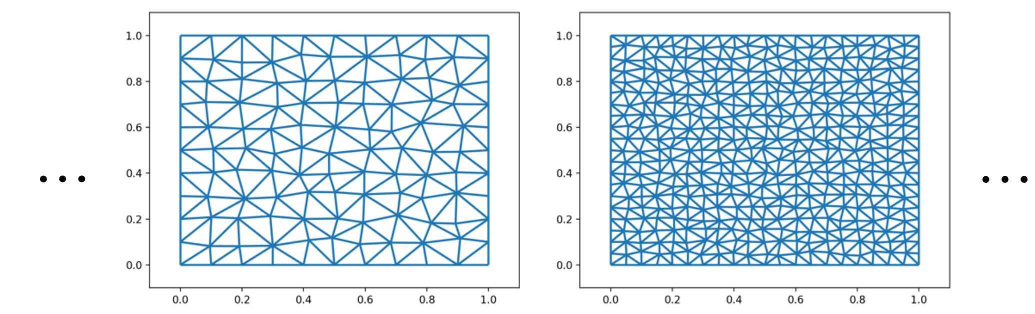

Regular Mesh								
	Face Elements (Raviart-Thomas)				Edge Elements (Nedelec)			
	Integral		Point		Integral		Point	
h	l2 Error	Rate	l2 Error	Rate	l2 Error	Rate	l2 Error	Rate
0.02	6.64559E-05	-	7.28634E-05	-	4.94356E-05	-	6.77058E-05	-
0.01	4.78883E-06	3.79	4.58658E-06	3.99	2.34348E-06	4.40	2.5238E-06	4.75
0.005	8.4105E-08	5.83	8.22289E-08	5.80	7.40259E-08	4.98	6.62357E-08	5.25
0.0025	2.10075E-09	5.32	1.95143E-09	5.40	2.03705E-09	5.18	1.71913E-09	5.27
0.00125	5.11811E-11	5.36	4.7762E-11	5.35	4.57059E-11	5.48	4.04744E-11	5.41

Case 2
Point cloud of edge data, randomly rotated and dilated

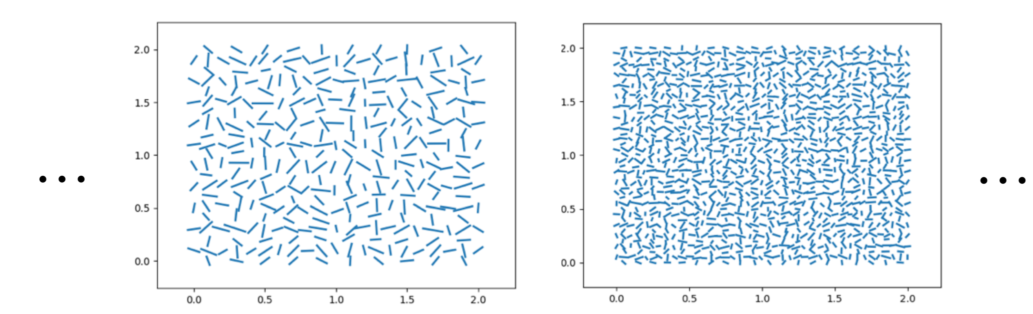

Randomly Rotated Edges from Dilated Mesh								
	Face Elements (Raviart-Thomas)				Edge Elements (Nedelec)			
	Integral		Point		Integral		Point	
h	l2 Error	Rate	l2 Error	Rate	l2 Error	Rate	l2 Error	Rate
0.02	0.0016478	-	0.0017719	-	0.0016324	-	0.0016149	-
0.01	5.777E-05	4.83	5.941E-05	4.90	3.549E-05	5.52	3.466E-05	5.54
0.005	1.075E-06	5.75	1.039E-06	5.84	1.03E-06	5.11	1.045E-06	5.05
0.0025	2.647E-08	5.34	2.583E-08	5.33	2.458E-08	5.39	2.388E-08	5.45
0.00125	7.168E-10	5.21	6.746E-10	5.26	7.279E-10	5.08	6.875E-10	5.12

Performance for Standard Point Data Remap

Scaling study of **a single CPU vs threaded CPU** (Intel Power8) vs **GPU** (NVIDIA Tesla P100)

QR Basis Order	# Targets	Time (s)					Speedup	Time (s)	Speedup
		1	8	16	32	64	CPU	GPU	GPU
2	300,000	196.76	32.5777	12.6224	8.1136	6.21444	32x	1.88265	104x
3	150,000	296.269	37.5049	18.8854	12.3385	9.66613	31x	2.77242	107x
4	50,000	293.762	36.9651	18.9772	12.4823	9.80607	30x	7.17664	41x
5	30,000	459.164	58.8537	29.6574	19.4458	15.3464	30x	10.3913	44x
6	10,000	377.407	48.3725	24.3342	16.0421	12.4498	30x	10.9297	35x
9	1,000	409.659	52.5378	26.4115	17.1199	13.6088	30x	20.6386	20x

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.
SAND No. 2019-7513 C