

Exceptional service in the national interest

Developing a GPU-enabled DPG Toolkit

Experiences with *Intrepid2*

J. Plews, G. Bunting, J. Rouse, C. Dohrmann

Trilinos User Group (TUG) Meeting 2021 December 01



SAND2021-15074 C

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and 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.



Motivation: High-frequency wave propagation





Example: Semi-infinite duct



Motivation: High-frequency wave propagation

• Excitation frequency: <u>20 Hz</u>



• Analytical solution: $p(z) = \rho_o c U(\omega) e^{-ikz}$



Motivation: High-frequency wave propagation

• Excitation frequency: <u>100 Hz</u>



Same number of elements per wavelength as 20 Hz case, 4.7x larger L_2 error!

Rouse, Walsh, ASA 2019

Motivation: Address pollution in high-frequency problems

• Relative error for classical Galerkin finite element solution to the Helmholtz equation:



Interpolation error: how well the shape functions approximate the solution

F

- <u>'Pollution effect'</u>: the accumulation of interpolation error over the length of the structure
- For constant relative error the second term must remain constant with increasing L



Motivation: Address pollution in high-frequency problems

- Brute force: Domain decomposition and massively parallel FE
 - Pushed as far as possible with Galerkin FE and traditional element formulations
 - Sierra Mechanics implementation is the state of the art
- Higher-order FE methods
 - Going beyond quadratic approximations with *p*-elements drastically improves efficiency
 - Enriched methods

n.

- Discontinuous Petrov-Galerkin
 - Developed by Prof. Demkowicz at UT Austin
- Others in the literature ...

Discontinuous Petrov-Galerkin (DPG) discretizations

- Trial functions are continuous, test functions are discontinuous
 - Compute optimal test functions on the fly
 - Built-in error indicator

61

- Ultraweak formulation
 - Positive definite matrices
- Pollution is diffusive rather than dispersive
 - Pollution free in 1D
- Unconditionally stable
 - Allows for *hp*-adaptivity



L. Demkowicz and J. Gopalakrishnan. *ICES REPORT 15-20, The University of Texas at Austin*, 2015. J. Zitelli and I. Muga and L. Demkowicz and J. Gopalakrishnan and D. Pardo and V. M. Calo. *Journal of Computational Physics*, 230(7): 2406-2432, 2011.

S. Petrides. PhD Dissertation, The University of Texas at Austin, 2019.

mini_dpg: A Trilinos-based toolkit for FE and DPG discretizations

- mini_dpg mimics Sierra apps like Structural Dynamics and Solid Mechanics
 - Three-dimensional, MPI-parallel, heterogeneous (GPU-accelerated)
 - ...but with <u>advanced discretizations</u> in mind
- Intrepid2: finite element spaces for interpolation and integration
 - High-order approximations: nodal (Lagrange) and *p*-hierarchical bases
 - *H-div, H-grad*, L² polynomial spaces
- <u>STK</u> mesh and IO

- Parallel data structure for large FE meshes
- Cell topologies, computing and defining field data on nodes, edges, faces, and cells
- <u>Kokkos</u> parallelism, batch dense linear algebra, data management
- <u>Tpetra</u>, <u>Amesos2</u>, and <u>Belos</u> sparse, parallel linear solves
 - Also, p-preconditioning strategies, work in progress by Clark Dohrmann

Numerical results: Verification

- p-convergence of DPG vs. CG (Galerkin) formulation of <u>Poisson's equation</u>
 - 4x4x1 hex grid

Method of manufactured solutions for testing and verification:

 $\nabla^2 \phi = f$ $\phi = \sin(\pi x) \sin(\pi y) \sin(\pi z)$ $\to f = -3\pi^2 \sin(\pi x) \sin(\pi y) \sin(\pi z)$





Numerical results: Verification problem performance

- GPU acceleration has benefits (even at low *p*-orders for DPG)
- Memory usage in DPG is a concern







<u>Memory</u>

<u>MPI + OpenMP</u>: 2 MPI ranks x 4 threads per rank <u>MPI + CUDA</u>: 2 MPI ranks + 1 Nvidia Volta GPU

Numerical results: High-order adaptive visualization

- High-order visualization/output
- <u>Strategy</u>: *h*-adapt "viz" mesh and interpolate to points based on desired error
- <u>Example</u>: 2D acoustic duct (Helmholtz), 1 element per wavelength

p=6 CGFE solution (Re)





p=7 DPG solution (Re)





Example: Discretizing Poisson's equation

3

Intrepid2-based assembly of PDE terms is <u>simple for continuous Galerkin</u>...

intrepid2_basis::Basis hex = create_p_hierarchical_hex_basis("hierarchical", "hgrad", 1); auto fs = intrepid2_basis::FunctionSpace::create(iface, hex, hexNodes, hexNodeGids); auto transformDShape = fs->compute_transformed_derivatives(); // stiffness fs->compute weighted integral(stiffness, transformDShape, transformDShape);

• <u>Example</u>: Discretizing Poisson's equation

- DPG discretizations introduce more complexity and bookkeeping
- Introduce <u>collections</u> of function spaces and bases

Multiple bases	<pre>BasisCollection b; short pInteg = 2 * porderTest; b.add_basis(Basis("hex8", "dg_hierarchical", "hgrad", porderTest, pInteg), false); b.add_basis(Basis("hex8", "dg_hierarchical", "hdiv", porderTest, pInteg), false); b.add_basis(Basis("hex8", "hierarchical", "l2", porder - 1, pInteg), true); b.add_basis(Basis("hex8", "hierarchical", "hdiv", porder, pInteg), false); b.add_basis(Basis("hex8", "hierarchical", "hdiv", porder, pInteg), false); b.add_basis(Basis("hex8", "hierarchical", "hgrad", porder, pInteg), false);</pre>
Function spaces corresponding to trace and element spaces	<pre>FunctionSpaceCollection fc(cellCoords, cellVertexGids, bases); fc.add_function_space_instance("v", "dg_hierarchical", "hgrad", 1); fc.add_function_space_instance("q", "dg_hierarchical", "hdiv", 1); fc.add_function_space_instance("phi", "hierarchical", "l2", 1); fc.add_function_space_instance("psi", "hierarchical", "l2", 3); fc.add_trace_function_space_instance("phiHat", "hierarchical", "hgrad", 1); fc.add_trace_function_space_instance("psiHatN", "hierarchical", "hdiv", 1);</pre>

- <u>Example</u>: Discretizing Poisson's equation
 - Intrepid2-based implementation exploits integrals over <u>sub-matrices</u> to accommodate diverse range of PDE terms

<pre>auto stiffPhiQ = Kokkos::subview(stiffness, Kokkos::ALL(), phiDofs, qDofs);</pre>
<pre>auto stiffPsiQ = Kokkos::subview(stiffness, Kokkos::ALL(), psiDofs, qDofs);</pre>
<pre>auto stiffPsiV = Kokkos::subview(stiffness, Kokkos::ALL(), psiDofs, vDofs);</pre>
// Phat
// DHat auto stiffPhiHatO = Kokkos··subview(stiffness, Kokkos··ALL(), phiHatDofs, gDofs);
auto stiffPsiHatNV = Kokkos::subview(stiffness, Kokkos::ALL(), psiHatNDofs, vDofs);
// в
fsHvol->compute_weighted_integral(stiffPhiQ, transformShapeHvol, transformDerivHdiv, -1.0);
isHvol->compute_weighted_integral(stiffPsiV, transformShapeHvolvec, transformShapeHdiv, -1.0);
ishvor >compute_weighted_integrar(stillisiv, transformshapenvorvet, transformberivngrad, 1.0),
<pre>for (unsigned iface = 0; iface < hex.num faces(); ++iface)</pre>
auto qDotN = compute_values_dotted_with_normals(transformShapeHdivOnFace, hexFaceNormals);
auto psihacbotn = compute_values_dotted_with_normals(transformshapeFluxcghdiv, nexFaceNormals);
// Bhat
fsTraceCgHgrad->compute weighted integral(stiffPhiHatQ, transformShapeTraceCgHgrad, qDotN);
fsFluxCgHdiv->compute_weighted_integral(stiffPsiHatNV, psiHatDotN, transformShapeHgradOnFace);

Element contributions

E

Side contributions

- Example: Discretizing Poisson's equation
 - Aside: mini_dpg is also leveraging
 - Kokkos batched BLAS/LAPACK operations
 - and Intrepid2 projection utilities

```
// BTGinvl <- B^H * Ginv * l
// BTGinvB <- B^H * Ginv * B
dpgk::dense_kokkos_gemm(false, false, alpha, BT, get_ginv_b_subview(), beta, BTGinvB);
dpgk::dense_kokkos_gemv(false, alpha, BT, get_ginv_l_subview(), beta, BTGinvl);</pre>
```

 <u>Room for improvement</u>: Return values, handling array sizes, avoiding excess memory use/repeat allocations

auto transformShapeFluxCgHdiv = fsFluxCgHdiv->compute_transformed_values();
Each allocates a
Kokkos::DynRankView<> -convenient sizing
// q dot n = transformShapeHDiv . cellSideNormals
// q dot n = transformShapeHDiv . cellSideNormals
auto psiHatDotN = compute values dotted with normals(transformShapeFluxCgHdiv, hexFaceNormals);

- Memory overhead due to arrays
- Strategies for array allocation and reuse?
- Opportunities for more expression templates?

• <u>Room for improvement</u>: Opportunities to exploit parallel asynchrony

Operations on <u>non-</u> <u>overlapping</u> arrays could be <u>asynchronous and non-</u> <u>blocking</u>

fsHvol->compute_weighted_integral(stiffPhiQ, transformShapeHvol, transformDerivHdiv, -1.0); fsHvol->compute_weighted_integral(stiffPsiQ, transformShapeHvolVec, transformShapeHdiv, -1.0); fsHvol->compute_weighted_integral(stiffPsiV, transformShapeHvolVec, transformDerivHgrad, -1.0);

```
for (unsigned iface = 0; iface < hex.num_faces(); ++iface)</pre>
```

```
•••
```

// Bhat

fsTraceCgHgrad->compute_weighted_integral(stiffPhiHatQ, transformShapeTraceCgHgrad, qDotN);
fsFluxCgHdiv->compute_weighted_integral(stiffPsiHatNV, psiHatDotN, transformShapeHgradOnFace);

Takeaways and topics for future work

- mini_dpg is leveraging Trilinos packages to enable high-order discretizations on threedimensional PDE discretizations of interest
 - More general than expected thanks to flexibility in Intrepid2
 - Elastodynamics implementation pending

- GPU acceleration in Intrepid2 and Kokkos batched linear algebra prove valuable at higher polynomial orders
 - Work to do to prove performance on realistic problems
- Memory high-water mark, usability, and performance
 - Matrix-free computations and efficient preconditioners
 - Avoid repeated memory allocations...especially on GPUs
 - Asynchronous, nonblocking kernels?