

Utilizing Multicore and GPU Hardware for Multiphysics Simulation through Implicit High-order Finite Element Methods with Tensor Product Structure

Jed Brown¹, Aron Ahmadi², Matt Knepley³, Barry Smith¹

¹Mathematics and Computer Science Division, Argonne National Laboratory

²King Abdullah University of Science and Technology

³Computation Institute, University of Chicago

2012-02-15

The Roadmap

Hardware trends

- ▶ More cores (keep hearing $\mathcal{O}(1000)$ per node)
- ▶ Long vector registers (already 32 bytes for AVX and BG/Q)
- ▶ Must use SMT to hide memory latency
- ▶ Must use SMT for floating point performance (GPU, BG/Q)
- ▶ Large penalty for non-contiguous memory access

“Free flops”, but how can we use them?

- ▶ High order methods good: better accuracy per storage
- ▶ High order methods bad: work unit gets larger
- ▶ GPU threads have very little memory, must keep work unit small
- ▶ Need library composability, keep user contribution embarrassingly parallel

How to program this beast?



- ▶ Decouple physics from discretization

- ▶ Expose small, embarrassingly parallel operations to user
- ▶ Library schedules user threads for reuse between kernels
- ▶ User provides physics in kernels run at each quadrature point
- ▶ Continuous weak form: find $u \in \mathcal{V}_D$

$$v^T F(u) \sim \int_{\Omega} v \cdot f_0(u, \nabla u) + \nabla v : f_1(u, \nabla u) = 0, \quad \forall v \in \mathcal{V}_0$$

- ▶ Similar form at faces, but may involve Riemann solve
- ▶ Library manages reductions
 - ▶ Interpolation and differentiation on elements
 - ▶ Interaction with neighbors (limiting, edge stabilization)
 - ▶ Exploit tensor product structure to keep working set small
 - ▶ Assembly into solution/residual vector (sum over elements)

Nodal hp -version finite element methods



1D reference element



- ▶ Lagrange interpolants on Legendre-Gauss-Lobatto points
- ▶ Quadrature \hat{R} , weights \hat{W}
- ▶ Evaluation: \hat{B}, \hat{D}

3D reference element

$$\begin{aligned}\hat{W} &= \hat{W} \otimes \hat{W} \otimes \hat{W} & \hat{D}_0 &= \hat{D} \otimes \hat{B} \otimes \hat{B} \\ \hat{B} &= \hat{B} \otimes \hat{B} \otimes \hat{B} & \hat{D}_1 &= \hat{B} \otimes \hat{D} \otimes \hat{B} \\ & & \hat{D}_2 &= \hat{B} \otimes \hat{B} \otimes \hat{D}\end{aligned}$$

These tensor product operations are very efficient, 70% of peak flop/s

Nodal hp -version finite element methods



1D reference element

- ▶ Lagrange interpolants on Legendre-Gauss-Lobatto points
- ▶ Quadrature \hat{R} , weights \hat{W}
- ▶ Evaluation: \hat{B}, \hat{D}

3D reference element

$$\begin{aligned}\hat{W} &= \hat{W} \otimes \hat{W} \otimes \hat{W} & \hat{D}_0 &= \hat{D} \otimes \hat{B} \otimes \hat{B} \\ \hat{B} &= \hat{B} \otimes \hat{B} \otimes \hat{B} & \hat{D}_1 &= \hat{B} \otimes \hat{D} \otimes \hat{B} \\ & & \hat{D}_2 &= \hat{B} \otimes \hat{B} \otimes \hat{D}\end{aligned}$$

These tensor product operations are very efficient, 70% of peak flop/s

Operations on physical elements

Mapping to physical space

$$x^e : \hat{K} \rightarrow K^e, \quad J_{ij}^e = \partial x_i^e / \partial \hat{x}_j, \quad (J^e)^{-1} = \partial \hat{x} / \partial x^e$$

Element operations in physical space

$$B^e = \hat{B} \quad W^e = \hat{W} \Lambda(|J^e(r)|)$$

$$D_i^e = \Lambda \left(\frac{\partial \hat{x}_0}{\partial x_i} \right) \hat{D}_0 + \Lambda \left(\frac{\partial \hat{x}_1}{\partial x_i} \right) \hat{D}_1 + \Lambda \left(\frac{\partial \hat{x}_2}{\partial x_i} \right) \hat{D}_2$$

$$(D_i^e)^T = \hat{D}_0^T \Lambda \left(\frac{\partial \hat{x}_0}{\partial x_i} \right) + \hat{D}_1^T \Lambda \left(\frac{\partial \hat{x}_1}{\partial x_i} \right) + \hat{D}_2^T \Lambda \left(\frac{\partial \hat{x}_2}{\partial x_i} \right)$$

Global problem is defined by assembly

$$F(u) = \sum_e \mathcal{E}_e^T \left[(B^e)^T W^e \Lambda(f_0(u^e, \nabla u^e)) + \sum_{i=0}^d (D_i^e)^T W^e \Lambda(f_{1,i}(u^e, \nabla u^e)) \right] = 0$$

where $u^e = B^e \mathcal{E}^e u$ and $\nabla u^e = \{D_i^e \mathcal{E}^e u\}_{i=0}^2$

Representation of Jacobians, Automation

- ▶ For unassembled representations, decomposition, and assembly
- ▶ Continuous weak form: find u

$$v^T F(u) \sim \int_{\Omega} v \cdot f_0(u, \nabla u) + \nabla v : f_1(u, \nabla u) = 0, \quad \forall v \in \mathcal{V}_0$$

- ▶ Weak form of the Jacobian $J(u)$: find w

$$v^T J(u) w \sim \int_{\Omega} [v^T \quad \nabla v^T] \begin{bmatrix} f_{0,0} & f_{0,1} \\ f_{1,0} & f_{1,1} \end{bmatrix} \begin{bmatrix} w \\ \nabla w \end{bmatrix}$$
$$[f_{i,j}] = \begin{bmatrix} \frac{\partial f_0}{\partial u} & \frac{\partial f_0}{\partial \nabla u} \\ \frac{\partial f_1}{\partial u} & \frac{\partial f_1}{\partial \nabla u} \end{bmatrix} (u, \nabla u)$$

- ▶ Terms in $[f_{i,j}]$ easy to compute symbolically, AD more scalable.
- ▶ Nonlinear terms f_0, f_1 usually have the most expensive nonlinearities in the computation of scalar material parameters
 - ▶ Equations of state, effective viscosity, “star” region in Riemann solve
 - ▶ Compute gradient with reverse-mode, store at quadrature points.
 - ▶ Perturb scalars, then use forward-mode to complete the Jacobian.
 - ▶ Flip for action of the adjoint.

Conservative (non-Boussinesq) two-phase ice flow

Find momentum density ρu , pressure p , and total energy density E :

$$(\rho u)_t + \nabla \cdot (\rho u \otimes u - \eta Du_i + p1) - \rho g = 0$$

$$\rho_t + \nabla \cdot \rho u = 0$$

$$E_t + \nabla \cdot ((E + p)u - k_T \nabla T - k_\omega \nabla \omega) - \eta Du_i : Du_i - \rho u \cdot g = 0$$

- ▶ Solve for density ρ , ice velocity u_i , temperature T , and melt fraction ω using constitutive relations.
 - ▶ Simplified constitutive relations can be solved explicitly.
 - ▶ Temperature, moisture, and strain-rate dependent rheology η .
 - ▶ High order FEM, typically Q_3 momentum & energy, SUPG (yuck).
- ▶ DAEs solved implicitly after semidiscretizing in space.
- ▶ Preconditioning using nested fieldsplit

How much nesting?

$$P_1 = \begin{pmatrix} J_{uu} & J_{up} & J_{uE} \\ 0 & B_{pp} & 0 \\ 0 & 0 & J_{EE} \end{pmatrix}$$

- ▶ B_{pp} is a mass matrix in the pressure space weighted by inverse of kinematic viscosity.
- ▶ Elman, Mihajlović, Wathen, JCP 2011 for non-dimensional isoviscous Boussinesq.
- ▶ Works well for non-dimensional problems on the cube, not for realistic parameters.
 - ▶ Low-order preconditioning full-accuracy unassembled high order operator.
 - ▶ Build these on command line with PETSc PCFieldSplit.

$$P = \begin{bmatrix} \begin{pmatrix} J_{uu} & J_{up} \\ J_{pu} & 0 \end{pmatrix} & \\ \begin{pmatrix} J_{Eu} & J_{Ep} \end{pmatrix} & J_{EE} \end{bmatrix}$$

- ▶ Inexact inner solve using upper-triangular with B_{pp} for Schur.
- ▶ Another level of nesting.
- ▶ GCR tolerant of inexact inner solves.
- ▶ Outer converges in 1 or 2 iterations.

Pseudocolor

Var: Momentum Density_magnitude

1.e+06

9.e+05

6.e+05

3.e+05

0.

Max: 2.e+06

Min: 0.

Streamline

Var: Speed

1.

0.8

0.5

0.3


0.0003

Max: 1.

Min: 0.0003

Mesh

Var: dFS_E-D



DB: vht-jako-he5q2-k2em14.dhm

Contour

Var: TemperaturePotential

— 2.9

— 1.5

— 5.9

— 10.

— 15.

— 19.

— 24.

Max: 7.3

Min: -28.

Streamline

Var: Speed

— 1.1e+03

— 8.5e+02


— 5.7e+02

— 2.8e+02

— 1.5

Max: 1.1e+03

Min: 1.5



CPU traversal code

- ▶ CPU traversal computes coefficients of test functions,

<https://github.com/jedbrowndohp/>

```
while (IteratorHasPatch(iter)) {
    IteratorGetPatchApplied(iter,&Q,&jw,
        &x,&dx,NULL,NULL,
        &u,&du,&u_,&du_, &p,&dp,&p_,NULL, &e,&de,&e_,&de_);
    IteratorGetStash(iter,NULL,&stash);
    for (dInt i=0; i<Q; i++) {
        PointwiseFunction(context,x[i],dx[i],jw[i],
            u[i],du[i],p[i],dp[i],e[i],de[i],
            &stash[i], u_[i],du_[i],p_[i],e_[i],de_[i]);
    }
    IteratorCommitPatchApplied(iter,INSERT_VALUES, NULL,NULL,
        u_,du_, p_,NULL, e_,de_);
    IteratorNextPatch(iter);
}
```

- ▶ GPU version calls `PointwiseFunction()` directly.
- ▶ Unassembled Jacobian application reuses stash

```
PointwiseJacobian(context,&stash[i],dx[i],jw[i],
    u[i],du[i],p[i],dp[i],e[i],de[i],
    u_[i],du_[i],p_[i],e_[i],de_[i]);
```

Finer grained parallelism for GPUs

- ▶ One element per thread uses too much local memory.
- ▶ Would like to use *about* one quadrature point per thread.
- ▶ Tensor product requires several synchronizations

$$\begin{aligned}\tilde{u} &= (A \otimes B \otimes C)u \\ &= (A \otimes I \otimes I)(I \otimes B \otimes I)(I \otimes I \otimes C)u\end{aligned}$$

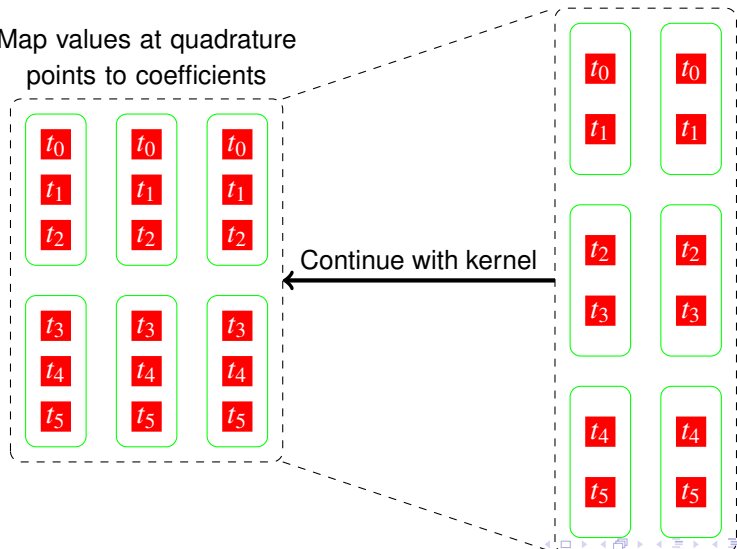
- ▶ Accumulation easy if only one thread accumulates into a location.
- ▶ Threads within a warp are implicitly synchronized, no need for `__syncthreads()`.
- ▶ Synchronization scope depends on approx order

Element	# quad pts	32T warps/element	TB size
Q_1	8	1/4	any
Q_2	27	1 (5T unused)	any
Q_3	64	2	64
Q_4	125	4 (3T unused)	128

Finer grained parallelism for GPUs, low order

Evaluate basis and process values at quadrature points

Map values at quadrature points to coefficients

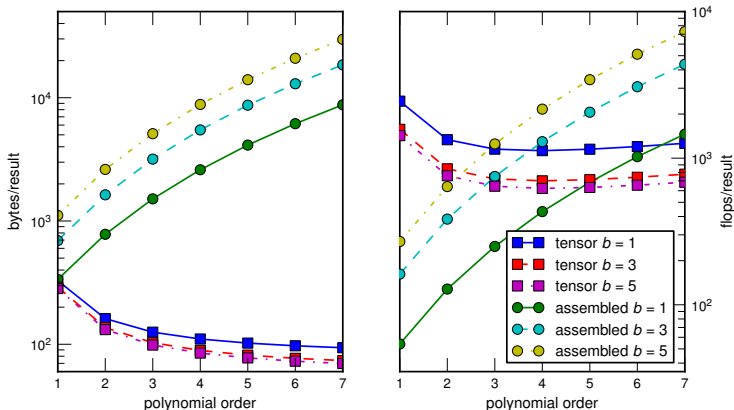


Avoiding copies

```
typedef enum {  
    PETSC_CUSP_UNALLOCATED,  
    PETSC_CUSP_GPU,  
    PETSC_CUSP_CPU,  
    PETSC_CUSP_BOTH  
} PetscCUSPFlag;
```

- ▶ Flag used for matrices and vectors.
- ▶ Data stays on GPU until it is needed on CPU (e.g. for MPI).
- ▶ Control flow for matrix and vector operations resides on CPU
 - ▶ almost all implementations run on GPU
 - ▶ can mix and match CPU-only and GPU-accelerated algorithms (but would need to pay for more copies)
- ▶ Currently always update the whole array
 - ▶ could order for low-volume updates

Performance of assembled versus unassembled



- ▶ High order Jacobian stored unassembled using coefficients at quadrature points, can use local AD
- ▶ Choose approximation order at run-time, independent for each field
- ▶ Precondition high order using assembled lowest order method
- ▶ Implementation $> 70\%$ of FPU peak, SpMV bandwidth wall $< 4\%$

Hardware Arithmetic Intensity

Operation	Arithmetic Intensity (flops/B)
Sparse matrix-vector product	1/6
Dense matrix-vector product	1/4
Unassembled matrix-vector product	≈ 8
High-order residual evaluation	> 5

Processor	BW (GB/s)	Peak (GF/s)	Balanced AI (F/B)
Sandy Bridge 6-core	21*	150	7.2
Magny Cours 16-core	42*	281	6.7
Blue Gene/Q node	43	205	4.8
GeForce 9400M	21	54	2.6
GTX 285	159	1062	6.8
Tesla M2050	144	1030	7.1

On preconditioning and multigrid

- ▶ Currently using assembled matrices for preconditioning
- ▶ Want matrix-free preconditioners for high hardware utilization
- ▶ Geometric h - and p -multigrid, could be FAS
- ▶ Smoothers build/solve with small dense matrices
 - ▶ “point” matrices: can use single threads
 - ▶ “element” matrices: need to cooperate within thread blocks
 - ▶ I want a dense linear algebra library to be called collectively within a thread block
- ▶ Multiplicative (Gauss-Seidel) is algorithmically nice
- ▶ Spectral analysis for polynomial/multi-stage smoothers
- ▶ Coarser levels better to do on CPU
 - ▶ Potential for additive correction to run concurrently

Outlook

- ▶ Sparse matrix assembly (for preconditioning)
 - ▶ > 100 GF/s for lowest order Stokes (Matt Knepley)
 - ▶ common “pointwise” physics code with CPU implementation
 - ▶ Dohp CPU version faster than libMesh and Deal.II for Q_1
 - ▶ Q_1 assembly embedded in higher order is 8% slower than hand-rolled
- ▶ Matrix-free tensor-product versions reliably get about 70% of peak flops
- ▶ Finer grained parallelism in GPU tensor product kernels
- ▶ Can't wait for OpenCL to implement indirect function calls
- ▶ Symbolic differentiation too slow, tired of hand-differentiation
 - ▶ I want source-transformation AD with indirect function calls
- ▶ Find correct amount of reuse between face and cell integration
- ▶ Riemann solves harder to vectorize
- ▶ Hide dispatch to pointwise kernels inside library
 - ▶ Easy, but scary. Library/framework becomes **F**ramework.
 - ▶ Interoperability of user-rolled, library-provided, and generated traversal code.