CEMRACS 2016: A-HA! Accelerated Higher-order finite element Assembly through the use of a duality-based algorithm

Lead: Thomas Dubos (LMD, Ecole Polytechnique, dubos@lmd.polytechnique.fr) Participant: Chris Eldred (LAGA, University of Paris 13, chris.eldred@gmail.com) Participant: Evaggelos Kritsikis (LAGA, University of Paris 13, wkrits@gmail.com) Participant: Manel Tayachi (Indiana University, Bloomington, manel.tayachi@gmail.com) Participant: Fabrice Voitus (CNRM, Meteo France, fabrice.voitus@meteo.fr) Participant: (CNRM, Meteo France, fabrice.voitus@meteo.fr) Funding: Météo-France, ANR HEAT (ANR-14-CE23-0010, AIMES projet (ANR-15-SPPE-001), Labex MME-DII

As part of the ANR HEAT project, an investigation into the suitability of higher-order mixed finite element type methods for use in an atmospheric dynamical core has been started. In particular, a new type of element (histopolation elements) that combines conservation, higher-order accuracy and excellent wave dispersion properties was developed ([?]). However, a significant bottleneck to the use of this method in an operational model is the efficiency of the resulting solver. This project aims to investigate the use a duality-based algorithm ([?]) in accelerating the both the assembly and matrix-free action computation for variational forms built using the new elements.

1 Introduction

Over the last 10 years, significant interest has been growing in the use of structure preserving methods for atmospheric dynamical cores. A very important class of these methods are represented by the finite element exterior calculus ([?]), which constructs a discrete de Rham cohomology using sets of carefully chosen mixed finite element spaces. This allows the construction of models that conserve important invariants such as total energy (at least under spatial discretization). A pioneer in the application of these methods to atmospheric dynamical cores has been the GungHo project ([?]), which is working towards a non-hydrostatic model in Eulerian coordinates. As a part of the HEAT project, a similar effort has begun, with the aim of producing a version of the hydrostatic primitive equation model Dynamico ([?]) based on the newly developed histopolation elements ([?]). These new elements allow a model that combines excellent wave dispersion properties with conservation and higher-order accuracy.

Central to the operational use of such a scheme in an atmospheric dynamical core, however, is the resulting computational performance. A major component of this is the time required to either assemble matrices or compute the matrix-free action of a variational form. There has been a great deal of work on accelerated finite element assembly (such as the FEniCS and Firedrake projects, [?] and [?]), but most of it is specialized towards affine transformations and simplicial elements. In contrast, we have tensor product elements (we also have a tensor-product grid) and non-affine transformations (in fact, our Jacobians are specified analytically and vary with quadrature point). The current approach uses a standard loop nest finite element assembly algorithm, optimized with sum factorization (exploiting the tensor product structure of the elements) and direct indexing (using the tensor product structure of the grid). Recently, [?] has shown that use of a dualitybased algorithm combined with a dense linear algebra (BLAS) library can offer greatly improved performance over a standard loop nest algorithm for both assembly and matrix-free action for quadrilateral mixed finite elements with arbitrary Jacobians (especially at higher-order). The central aim of this project is to determine if this is also true for the histopolation elements.

2 **Project Description**

Following KIRBY REF, we will investigate the use of a duality based assembly algorithm combined with optimized dense library libraries (BLAS) to accelerate the matrix assembly process for the newly developed histopolation elements. This algorithm is described below:

2.1 Duality-based Assembly Algorithm

Consider the variational form associated with the Laplacian term of the standard Poisson equation:

$$\int_{\Omega} \vec{\nabla} v \cdot \vec{\nabla} v d\Omega \tag{1}$$

where $u, v \in H^1$ are the basis functions and Ω represents a physical grid cell. In order to actually calculate this integral, it is first transformed to a reference cell Ω_K as

$$\int_{\Omega_K} J^{-T} \vec{\nabla} u \cdot J^{-T} \vec{\nabla} v |J| d\Omega$$
⁽²⁾

where J is the Jacobian of the transform from the reference grid cell to the physical grid cell and |J| is its determinant. This integral is then evaluated using numerical quadrature as

$$\sum_{n} \sum_{m} \sum_{q} w_{q} (J_{q}^{-T} D u_{q,n})^{T} (J_{q}^{-T} D v_{q,m}) |J|_{q}$$
(3)

where the sum over q is a sum over quadrature points (w_q are the quadrature weights and J_q is the Jacobian transform evaluated at quadrature points), the sums over n and m are sums over the basis functions and $Du_{q,n}$ are the derivatives of the basis functions evaluated at quadrature points. This can be rewritten (using the fact that $(Ab)^T = b^T A^T$ (this is the duality part) as

$$\sum_{n} \sum_{m} \sum_{q} Du_{q,n}^{T} |J|_{q} J_{q}^{-1} J_{q}^{-T} Dv_{q,m} w_{q}$$
(4)

Letting $G_q = J_q^{-1} J_q^{-T} |J|_q$, $Dv_{q,m}^w = w_q Dv_{q,m}$ and $\Gamma_{q,m} = GDv_{q,m}^w$ we have

$$\sum_{n} \sum_{m} \sum_{q} Du_{q,n}^{T} \Gamma_{q,m}$$
(5)

Therefore, assuming pre-tabulated basis functions/derivatives, quadrature-weighted basis functions/derivatives variants (ie $Dv_{q,m}^w$) and Jacobians, the assembly of the local element tensor for the Laplacian operator for a batch of N elements can be split into three steps:

- 1. Compute $G_q = J_q^{-1} J_q^{-T} |J|_q$ for the entire batch of N elements
- 2. Compute $\Gamma_{q,m} = GDv_{q,m}^w$ for the entire batch of N elements, using sum factorization to accelerate the process
- 3. Compute $\sum_{n} \sum_{m} \sum_{q} Du_{q,n}^{T} \Gamma_{q,m}$ for the entire batch of N elements

A significant fraction of the computational work (especially at higher-order) is done in step 3, which can be cast as a matrix-matrix multiplication and therefore can be accelerated through the use of an optimized BLAS library. This is the essence of the technique.

2.2 Matrix Free Action Variant and Relation to Kirby 2012

As shown in [?], both the standard loop nest and BLAS based variants of this duality algorithm provide increased assembly speed for traditional finite elements relative to a standard assembly algorithm, at a very modest increase in local storage required. The same basic algorithm can also be used to compute the matrix-free action of the variational form, again with the majority of the work occurring as a matrix-matrix multiplication. These two algorithms (assembly and matrixfree action) are almost identical to the ones outlined in [?], with the exception of the use of sum factorization and direct indexing to accelerate step 2. Both of these additional optimizations are facilitated by the use of a tensor product basis and tensor product grid. Although this algorithm has been illustrated only for H^1 finite elements and a simple constant coefficient variational form, it can be trivially extended to other spaces (including H(div) and H(curl)), other discretizations (such as the histopolation elements) and to variational forms with non-constant coefficients (See [?] for details).

2.3 Software Framework

In order to facilitate the development of the histpolation elements, a software framework was developed on top of the PETSc library ([?]). It is designed around the use of both tensor-product elements and tensor product grids. This framework currently has the following capabilities (among others):

- 1. Support for tensor product elements (including both histpolation and mixed finite elements) in an arbitrary number of dimensions.
- 2. Automated generation of the assembly code for a variational form (the user must provide the local element assembly kernel code, however).
- 3. MPI based parallelism through the PETSc library.

3 Objectives

The primary objective of this project is to assess the suitability of this new algorithm for accelerating the assembly and matrix-free action of the new histopolation elements. A secondary objective is to see what, if any, additional performance gains come from the use of sum factorization and direct indexing in step 2. The main milestones of this project are:

- 1. Implement the duality based assembly algorithm, both standard loop nest and BLAS variants that additionally exploit the tensor product structure of the basis functions and grid through sum factorization and direct indexing.
- 2. Implement the duality based matrix action algorithm, both standard loop nest and BLAS variants that additionally exploit the tensor product structure of the basis functions and grid through sum factorization and direct indexing.
- 3. Compare the performance of the new algorithms with the existing ones on a set of variational forms representative of those encountered in discretization of geophysical fluid flow problems relevant to the HEAT project.

In addition, if time permits, the overall solver performance for a relevant model of geophysical fluid flow (such as the rotating shallow water equations discretized using a semi-implicit method) will be investigated through a comparative study of preconditioners and iterative solvers (including matrix-free versions such as geometric multigrid).

4 Participants

The following members of the HEAT project will be present for at least part of CEMRACS 2016:

- 1. Thomas Dubos: 4th and 5th weeks (August 8th-19th)
- 2. Evaggelos Kritsikis: Last week (August 22nd-26th)
- 3. Chris Eldred: All 6 weeks
- 4. Fabrice Voitus: 3rd week (August 1st-6th)

Manel Tayachi will be present all 6 weeks.

5 Funding

The presence of the participants above at CEMRACS 2016 will be funded through the ANR HEAT project, as part of grant ANR-14-CE23-0010-01, and through the AIMES projet (ANR-15-SPPE-001), Météo-France and the Labex MME-DII.