A Compiler for Variational Forms

ROBERT C. KIRBY
The University of Chicago
and
ANDERS LOGG
Toyota Technological Institute at Chicago

As a key step towards a complete automation of the finite element method, we present a new algorithm for automatic and efficient evaluation of multilinear variational forms. The algorithm has been implemented in the form of a compiler, the FEniCS Form Compiler (FFC). We present benchmark results for a series of standard variational forms, including the incompressible Navier–Stokes equations and linear elasticity. The speedup compared to the standard quadrature-based approach is impressive; in some cases the speedup is as large as a factor of 1000.

Categories and Subject Descriptors: G.4 [**Mathematics of Computing**]: Mathematical Software—*Algorithm design and analysis; efficiency*; G.1.8 [**Numerical Analysis**]: Partial Differential Equations—*Finite element methods*

General Terms: Algorithms, Performance

Additional Key Words and Phrases: Variational form, compiler, finite element, automation

1. INTRODUCTION

The finite element method provides a general mathematical framework for the solution of differential equations and can be viewed as a machine that automates their discretization; given the variational formulation of a differential equation, the finite element method generates a discrete system of equations for the approximate solution.

This generality of the finite element method is seldom reflected in codes, which are often very specialized and can only solve one particular or a small set of differential equations.

This work was supported by the United States Department of Energy under grant DE-FG02-04ER25650.
Authors' addresses: R. C. Kirby, Department of Computer Science, University of Chicago, 1100 East 58th Street, Chicago, IL 60637; email: kirby@cs.uchicago.edu; A. Logg, Toyota Technological Institute at Chicago, University Press Building, 1427 East 60th Street, Chicago, IL 60637; email: logg@simula.no.
Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage, and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, or to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permission may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax: +1(212) 869-0481, or permission@acm.org.
© 2006 ACM 0098-3500/06/0900-0417 $5.00

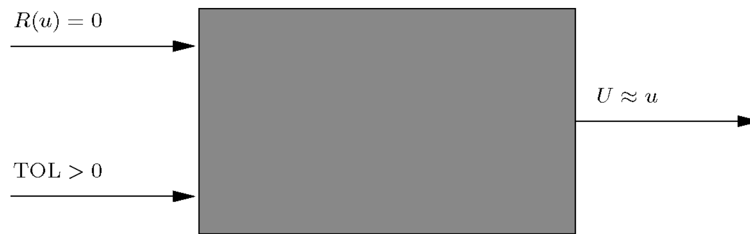

Fig. 1. The Automation of Computational Mathematical Modeling (ACMM).

There are two major reasons that the finite element method has yet to be fully automated; the first is the complexity of the task itself, and the second is that specialized codes often outperform general codes. We address both these concerns in this article.

A basic task of the finite element method is the computation of the element stiffness matrix from a bilinear form on a local element. In many applications, the computation of element stiffness matrices accounts for a substantial part of the total run-time for the code [Kirby et al. 2005a]. This routine constitutes a small amount of code, but can be tedious to get both correctly and efficiently. While the standard quadrature-based approach to computing the element stiffness matrix works on very general variational forms, it is well-known that precomputing certain quantities in multilinear forms can improve the efficiency of building finite element matrices.

The methods discussed in this for efficient computation of element stiffness matrices are based on ideas previously presented in Kirby et al. [2005a; 2005b], where the basic idea is to represent the element stiffness matrix as a tensor product. A similar approach had been implemented earlier in the finite element library DOLFIN [Hoffman et al. 2005; Hoffman and Logg 2002], but only for linear elements. The current article generalizes and formalizes these ideas and presents an algorithm for the generation of the tensor representation of element stiffness matrices and for evaluation of the tensor product. This algorithm has been implemented in the form of the compiler FFC for variational forms; the compiler takes as input a variational form in mathematical notation and automatically generates efficient code (C or C++) for the computation of element stiffness matrices and their insertion into a global sparse matrix. This includes the generation of code both for computation of element stiffness matrices and local-to-global mappings of degrees of freedom.

1.1 FEniCS and the Automation of CMM

FFC, the FEniCS Form Compiler, is a central component of FEniCS [Hoffman et al. 2005], a project for the Automation of Computational Mathematical Modeling (ACMM). The central task of ACMM, as formulated in Logg [2004], is to create a machine that takes as input a model $R(u) = A(u) - f$, a tolerance TOL > 0, and a norm $\|\cdot\|$ (or some other measure of quality), and produces as output an approximate solution $U \approx u$ that satisfies the accuracy requirement $\|U - u\| <$ TOL using a minimal amount of work (see Figure 1). This includes an aspect of *reliability* (the produced solution should satisfy the accuracy

requirement) and an aspect of *efficiency* (the solution should be obtained with minimal work).

In many applications, several competing models are under consideration, and we would like to computationally compare them. Developing separate, special-purpose codes for each model is prohibitive. Hence, a key step of ACMM is the *automation of discretization*, that is, the automatic translation of a differential equation into a discrete system of equations, and as noted previously, this key step is automated by the finite element method. The FEniCS Form Compiler FFC may then be viewed as an important step towards the automation of the finite element method, and thus an important step towards the automation of CMM.

FEniCS software is free software. In particular, FFC is licensed under the GNU General Public License [Free Software Foundation 1991]. All FEniCS software is available for download on the FEniCS website [Hoffman et al. 2005]. In Section 5.6, we return to a discussion of the different components of FEniCS and their relation to FFC.

1.2 Current Finite Element Software

Several emerging projects seek to automate important aspects of the finite element method. By developing libraries in existing languages or new domain-specific languages, software tools may be built that allow programmers to define variational forms and other parts of a finite element method with succinct, mathematical syntax. Existing C++ libraries for finite elements include DOLFIN [Hoffman et al. 2005; Hoffman and Logg 2002], Sundance [Long 2003], deal.II [Bangerth et al. 2005], Diffpack [Langtangen 1999], and FEMSTER [Castillo et al. 2004; 2005]. Projects developing domain-specific languages for finite element computation include FreeFEM [Pironneau et al. 2005] and GetDP [Dular and Geuzaine 2005]. A precursor to the FEniCS project, Analysa [Bagheri and Scott 2003], was a Scheme-like language for finite element methods. Earlier work on object-oriented frameworks for finite element computation include that of Mackie [1992] and Masters et al. [1997].

While these tools are effective at exploiting modern software engineering to produce workable systems, we believe that additional mathematical insight will lead to even more powerful codes with more general approximating spaces and more powerful algorithms. The FEniCS project is more ambitious than to just collect and implement existing ideas.

1.3 Design Goals

The primary design goal for FFC is to accept as input "any" multilinear variational form and *any* finite element, and to generate code that will run with close to optimal performance.

We will make precise in Section 3.2 forms and elements the compiler can currently handle (general multilinear variational forms with coefficients over affine simplices).

A secondary goal for FFC is to create a new standard in form evaluation; hopefully, FFC can become a standard tool for practitioners solving partial

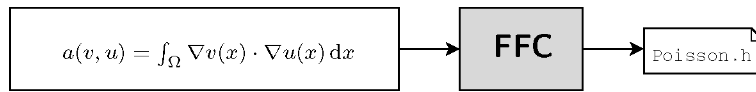

Fig. 2. The form compiler FFC takes as input a variational form and generates code for evaluation of the form.

differential equations using the finite element method. In addition to generating very efficient code for the evaluation of the element stiffness matrix, FFC thus removes the burden of having to implement the code from the developer. Furthermore, if the code for the element stiffness matrix is generated by a trusted compiler that has gone through rigorous testing, it is easier to achieve correctness of a simulation code.

The primary output target of FFC is the C++ library DOLFIN. By default, FFC accepts as input a variational form and generates code its evaluation in DOLFIN, as illustrated in Figure 2. Although FFC works closely with other FEniCS components, such as DOLFIN, it has an abstraction layer that allows it to be hooked-up to multiple backends. One example of this is the newly added ASE (ANL SIDL environment, [Balay et al. 2005a]) format added to FFC, which allows forms to be compiled for the next generation of PETSc [Balay et al. 2005b].

1.4 The Compiler Approach

It is widely accepted that in developing software for scientific computing there is a tradeoff between generality and efficiency; a software component that is general in nature, that is, it accepts a wide range of inputs, is often less efficient than one which performs the same job on a more limited set of inputs. As a result, most codes used by practitioners for the solution of differential equations are very specific.

However, by using a compiler approach, it is possible to combine generality and efficiency, without loss of either generality or efficiency. This is possible because our compiler works on a very small family of inputs (multilinear variational forms) with sharply defined mathematical properties. Our domain-specific knowledge allows us to generate much better code than if we were to use general-purpose compiler techniques.

1.5 Outline of This Article

Before presenting the main algorithm, we give a short background on the implementation of the finite element method and the evaluation of variational forms in Section 2. The main algorithm is then outlined in Section 3. In Section 4, we compare the complexity of form evaluation for the algorithm used by FFC with the standard quadrature-based approach. We then discuss the implementation of the form compiler in some detail in Section 5.

In Section 6, we compare the CPU-time for evaluating a series of standard variational forms using code automatically generated by FFC, as well as hand-coded quadrature-based implementations. The speedup is in all cases

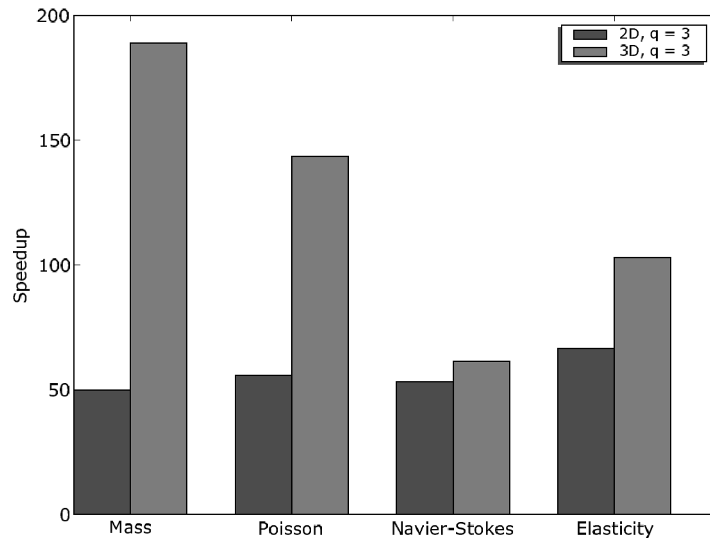

Fig. 3. Speedup for a series of standard variational forms (here compiled for cubic Lagrange elements on triangles and tetrahedra, respectively).

significant (in the case of cubic Lagrange elements on tetrahedra, a factor of 100, see Figure 3).

Finally, in Section 7, we summarize the current features and shortcomings of FFC and give directions for future development and research.

2. BACKGROUND

In this section, we present a brief background on the finite element method. The material is standard [Ciarlet 1976; Hughes 1987; Brenner and Scott 1994; Eriksson et al. 1996], but is included here to give a context for the presentation of the form compiler and to summarize the notation used throughout the remainder of this article.

For simplicity, we here consider only linear partial differential equations and note that these play an important role in the discretization of nonlinear partial differential equations (in Newton or fixed-point iterations).

2.1 Variational Forms

We work with the standard variational formulation of a partial differential equation: Find $u \in V$ such that

$$a(v, u) = L(v) \quad \forall v \in \hat{V}, \tag{1}$$

where $a : \hat{V} \times V \to \mathbb{R}$ is a bilinear form, $L : \hat{V} \to \mathbb{R}$ a linear form, and $(\hat{V}, V)$ a pair of suitable function spaces. For the standard example, Poisson's equation $-\Delta u(x) = f(x)$ with homogeneous Dirichlet conditions on a domain Ω, the bilinear form a is given by $a(v, u) = \int_\Omega \nabla v(x) \cdot \nabla u(x)\,\mathrm{d}x$, and the linear form L is given by $L(v) = \int_\Omega v(x) f(x)\,\mathrm{d}x$, and $\hat{V} = V = H_0^1(\Omega)$.

The finite element method discretizes Equation (1) by replacing $(\hat{V}, V)$ with a pair of (piecewise polynomial) discrete function spaces. With $\{\hat{\phi}_i\}_{i=1}^M$ as a basis for the test space $\hat{V}$ and $\{\phi_i\}_{i=1}^M$ a basis for the trial space V, we can expand the approximate solution U of Equation (1) in the basis functions of V, $U = \sum_{j=1}^M \xi_j \phi_j$, and obtain a linear system $A\xi = b$ for the *degrees of freedom* $\{\xi_j\}_{j=1}^M$ of the approximate solution U. The entries of the matrix A and the vector b defining the linear system are given by

$$\begin{aligned} A_{ij} &= a(\hat{\phi}_i, \phi_j), \quad i, j = 1, \ldots, M, \\ b_i &= L(\hat{\phi}_i), \quad i = 1, \ldots, M. \end{aligned} \tag{2}$$

2.2 Assembly

The standard algorithm for computing the matrix A (or vector b) is *assembly*; the matrix is computed by iteration over the elements K of a triangulation $\mathcal{T}$ of Ω, and the contribution from each local element is added to the global matrix A.

To see this, we note that if the bilinear form a is expressed as an integral over the domain Ω, we can write the bilinear form as a sum of element bilinear forms, $a(v, u) = \sum_{K \in \mathcal{T}} a_K(v, u)$, and thus

$$A_{ij} = \sum_{K \in \mathcal{T}} a_K(\hat{\phi}_i, \phi_j), \quad i, j = 1, \ldots, M. \tag{3}$$

In the case of Poisson's equation, the element bilinear form a_K is defined by $a_K(v, u) = \int_K \nabla v(x) \cdot \nabla u(x)\,\mathrm{d}x$.

Let now $\{\hat{\phi}_i^K\}_{i=1}^n$ be the restriction to K of the subset of $\{\hat{\phi}_i\}_{i=1}^M$ supported on K and let $\{\phi_i^K\}_{i=1}^n$ be the corresponding local basis for V. Furthermore, let $\hat{\iota}(\cdot, \cdot)$ be a mapping from the local to the global numbering scheme (local-to-global mapping) for the basis functions of $\hat{V}$, so that $\hat{\phi}_i^K$ is the restriction to K of $\phi_{\hat{\iota}(K,i)}$, and let $\iota(K, \cdot)$ be the corresponding mapping for V. We may now express an algorithm for the computation of the matrix A (Algorithm 1).

Algorithm 1 A = Assemble $(a, \mathcal{T}, \hat{V}, V)$

$A = 0$
for $K \in \mathcal{T}$
  for $i = 1, \ldots, n$
    for $j = 1, \ldots, n$
      $A_{\hat{\iota}(K,i)\iota(K,j)} = A_{\hat{\iota}(K,i)\iota(K,j)} + a_K\left(\hat{\phi}_i^K, \phi_j^K\right)$
    end for
  end for
end for

Alternatively, we may define the *element matrix* A^K by

$$A_{ij}^K = a_K\left(\hat{\phi}_i^K, \phi_j^K\right) \quad i, j = 1, \ldots, n, \tag{4}$$

and separate the computation on each element K into two steps: the computation of the element matrix A^K and the insertion of A^K into A (Algorithm 2).

Algorithm 2 A = Assemble $(a, \mathcal{T}, \hat{V}, V)$

$A = 0$
for $K \in \mathcal{T}$
  Compute A^K according to (4)
  Add A^K to A using the local-to-global mappings $(\hat{\iota}(K, \cdot), \iota(K, \cdot))$
end for

Separating the two concerns of computing the element matrix A^K and adding it to the global matrix A as in Algorithm 2 has the advantage that we may use an optimized library routine for adding the element matrix A^K to the global matrix A. Sparse matrix libraries such as PETSc [Balay et al. 1997; 2004; 2005b] often provide optimized routines for this type of operation. Note that the cost of adding A^K to A may be substantial, even with an efficient implementation of the sparse data structure for A [Kirby et al. 2005a].

As we shall next see, we may also take advantage of the separation of concerns of Algorithm 2 to optimize the computation of the element matrix A^K. This step is automated by the form compiler FFC. Given a bilinear (or multilinear) form a, FFC automatically generates code for the run-time computation of element matrix A^K.

3. EVALUATION OF MULTILINEAR FORMS

In this section, we present the algorithm used by FFC to automatically generate efficient code for run-time computation of the element matrix A^K.

3.1 Multilinear Forms

Let $\{V_i\}_{i=1}^r$ be a given set of discrete function spaces defined on a triangulation $\mathcal{T}$ of $\Omega \subset \mathbb{R}^d$. We consider a general multilinear form a defined on the product space $V_1 \times V_2 \times \cdots \times V_r$:

$$a : V_1 \times V_2 \times \cdots \times V_r \rightarrow \mathbb{R}. \tag{5}$$

Typically, $r = 1$ (linear form) or $r = 2$ (bilinear form), but the form compiler FFC can handle multilinear forms of arbitrary *arity* r. Forms of higher arity appear frequently in applications and include variable coefficient diffusion and advection of momentum in the incompressible Navier–Stokes equations.

Let now $\{\phi_i^1\}_{i=1}^{M_1}, \{\phi_i^2\}_{i=1}^{M_2}, \ldots, \{\phi_i^r\}_{i=1}^{M_r}$ be bases of $V_1, V_2, \ldots, V_r$ and let $i = (i_1, i_2, \ldots, i_r)$ be a multiindex. The multilinear form a then defines a rank r tensor given by

$$A_i = a\left(\phi_{i_1}^1, \phi_{i_2}^2, \ldots, \phi_{i_r}^r\right). \tag{6}$$

In the case of a bilinear form, the tensor A is a matrix (the stiffness matrix), and in the case of a linear form, the tensor A is a vector (the load vector).

As discussed in the previous section, to compute the tensor A by assembly, we need to compute the *element tensor* A^K on each element K of the triangulation $\mathcal{T}$ of Ω. Let $\{\phi_i^{K,1}\}_{i=1}^{n_1}$ be the restriction to K of the subset of $\{\phi_i^1\}_{i=1}^{M_1}$ supported on K and let us define the local bases on K for $V_2, \ldots, V_r$ similarly. The rank r

element tensor A^K is then defined by

$$A_i^K = a_K(\phi_{i_1}^{K,1}, \phi_{i_2}^{K,2}, \dots, \phi_{i_r}^{K,r}). \tag{7}$$

3.2 Evaluation by Tensor Representation

The element tensor A^K can be efficiently computed by representing A^K as a special tensor product. Under some mild assumptions which we shall make precise to follow, the element tensor A^K can be represented as the tensor product of a *reference tensor* A^0 and *geometry tensor* G_K:

$$A_i^K = A_{i\alpha}^0 G_K^\alpha, \tag{8}$$

or more generally, a sum $A_i^K = A_{i\alpha}^{0,k} G_{K,k}^\alpha$ of such tensor products, where i and α are multiindices and we use the convention that repetition of an index means summation over this index. The rank of the reference tensor is the sum of rank $r = |i|$ of the element tensor and rank $|\alpha|$ of the geometry tensor G_K. As we shall see, the rank of the geometry tensor depends on the specific form and is typically a function of the Jacobian of the mapping from the reference element and any variable coefficients appearing in the form.

Our goal is to develop an algorithm that converts an abstract representation of a multilinear form into: (i) the values of the reference tensor A^0 and (ii) an expression for evaluating the geometry tensor G_K for any given element K. Note that A^0 is fixed and independent of the element K and may thus be precomputed. Only G_K has to be computed on each element. As we shall see in Section 4, for a wide range of multilinear forms, this allows for computation of the element tensor A^K using far fewer floating-point operations than if the element tensor A^K were computed by quadrature on each element K.

To see how to obtain the tensor representation of Equation (8), we fix a small set of operations, allowing only those multilinear forms that can be expressed through these operations, and observe how the tensor representation of Equation (8) transforms under these operations. As we shall see in the following, this covers a wide range of multilinear forms (but not all). We first develop the tensor representation in abstract form and then present a number of test cases that exemplify the general notation in Section 3.3.

As basic elements, we take the local basis functions $\{\phi_\gamma\}_\gamma = \cup_i \{\phi_j^{K,i}\}_{j=1}^{n_i}$ for a set of finite element spaces V_i, $i = 1, 2, \dots$, including the finite element spaces $V_1, V_2, \dots, V_r$ on which the multilinear form is defined. Allowing addition $\phi_1 + \phi_2$ and multiplication with scalars $\alpha\phi$, we obtain a vector space $\mathcal{A}$ of linear combinations of basis functions. Since $\phi_1 - \phi_2 = \phi_1 + (-1)\phi_2$ and $\phi/\alpha = (1/\alpha)\phi$, we can easily equip the vector space with subtraction and division by scalars.

We next equip our vector space with multiplication between elements of the vector space. We thus obtain an algebra $\mathcal{A}$ of linear combinations of products of basis functions. Finally, we extend $\mathcal{A}$ by adding differentiation $\partial/\partial x_i$ with

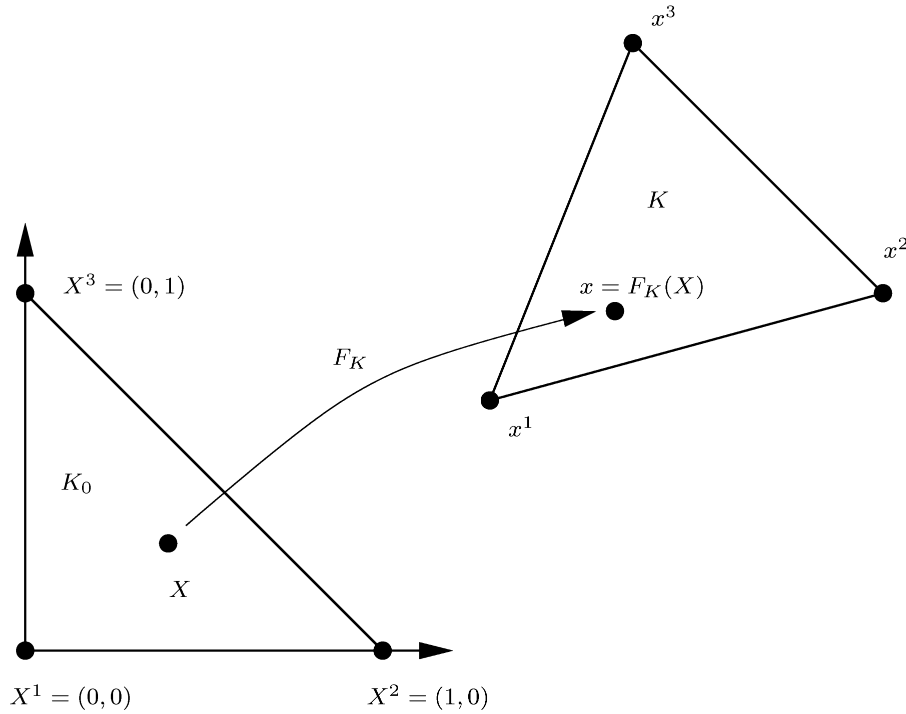

Fig. 4. The affine mapping F_K from the reference element K_0 to the current element K.

respect to a coordinate direction x_i, $i = 1, \ldots, d$, on K, to obtain

$$\mathcal{A} = \left\{ v : v = \sum c_{(\cdot)} \prod \frac{\partial^{|(\cdot)|} \phi_{(\cdot)}}{\partial x_{(\cdot)}} \right\}, \tag{9}$$

where $(\cdot)$ represents some multiindex.

To summarize, $\mathcal{A}$ is the algebra of linear combinations of products of basis functions or derivatives of basis functions that is generated from the set of basis functions through addition $(+)$, subtraction $(-)$, and multiplication $(\cdot)$, including multiplication with scalars, division by scalars $(/)$, and differentiation $\partial/\partial x_i$. Note that if the basis functions are vector-valued (or tensor-valued), the algebra is generated from the set of scalar components of the basis functions.

We may now state precisely the multilinear forms that the form compiler FFC can handle, namely, those multilinear forms that can be expressed as integrals over K (or the boundary of K) of elements of the algebra $\mathcal{A}$. Note that not all integrals over K of elements of $\mathcal{A}$ are multilinear forms; in particular, each product needs to be linear in each argument of the form.

The tensor representation of Equation (8) now follows by a standard change of variables using an affine mapping $F_K : K_0 \to K$ from a reference element K_0 to the current element K (see Figure 4). With $\{\Phi_\gamma\}_\gamma$ as the basis functions on the reference element corresponding to $\{\phi_\gamma\}_\gamma$, defined by $\Phi_\gamma = \phi_\gamma \circ F_K$, we obtain the following representation of the element tensor A^K corresponding to

$v_i = \left(\sum c_{(\cdot)} \prod \frac{\partial^{|(\cdot)|}\phi_{(\cdot)}}{\partial x_{(\cdot)}}\right)_i$:

$$\begin{aligned} A_i^K &= a_K\left(\phi_{i_1}^{K,1}, \phi_{i_2}^{K,2}, \ldots, \phi_{i_r}^{K,r}\right) = \int_K v_i \,\mathrm{d}x \\ &= \left(\int_K \sum c_{(\cdot)} \prod \frac{\partial^{|(\cdot)|}\phi_{(\cdot)}}{\partial x_{(\cdot)}} \,\mathrm{d}x\right)_i = \sum \left(c_{(\cdot)} \int_K \prod \frac{\partial^{|(\cdot)|}\phi_{(\cdot)}}{\partial x_{(\cdot)}} \,\mathrm{d}x\right)_i \\ &= \sum \left(c_{(\cdot)} \det F_K' \prod \frac{\partial X_{(\cdot)}}{\partial x_{(\cdot)}}\right)_\alpha \left(\int_{K_0} \prod \frac{\partial^{|(\cdot)|}\Phi_{(\cdot)}}{\partial X_{(\cdot)}} \,\mathrm{d}X\right)_{i\alpha} \\ &= A_{i\alpha}^{0,k} G_{K,k}^\alpha, \end{aligned} \tag{10}$$

where

$$A_{i\alpha}^{0,k} = \left(\int_{K_0} \prod \frac{\partial^{|(\cdot)|}\Phi_{(\cdot)}}{\partial X_{(\cdot)}} \,\mathrm{d}X\right)_{i\alpha}, \tag{11}$$

and

$$G_{K,k}^\alpha = \left(c_{(\cdot)} \det F_K' \prod \frac{\partial X_{(\cdot)}}{\partial x_{(\cdot)}}\right)_\alpha. \tag{12}$$

We have here used the fact that the mapping is affine to pull the determinant and transforms of the derivatives outside of the integral. For a discussion of nonaffine mappings, including the Piola transform and isoparametric mapping, see Section 3.4.

Note that the expression for the geometry tensor $G_{K,k}$ implicitly contains a summation if an index is repeated twice. Also note that the geometry tensor contains any variable coefficients appearing in the form.

As we shall see in Section 5, the representation of a multilinear form as an integral over K of an element of $\mathcal{A}$ is automatically available to the form compiler FFC, since a multilinear form must be specified using the basic operations that generate $\mathcal{A}$.

3.3 Test Cases

To make these ideas concrete, we write down the explicit tensor representation of Equation (8) of the element tensor A^K for a series of standard forms. We return to these test cases in Section 6 when we present benchmark results for each Test Case.

Test Case 1: The Mass Matrix. As a first simple example, we consider the computation of the mass matrix M with $M_{i_1 i_2} = a(\phi_{i_1}^1, \phi_{i_2}^2)$ and the bilinear form a given by

$$a(v, u) = \int_\Omega v(x)u(x) \,\mathrm{d}x. \tag{13}$$

By a change of variables given by the affine mapping $F_K : K_0 \to K$, we obtain

$$A_i^K = \int_K \phi_{i_1}^{K,1}(x)\phi_{i_2}^{K,2}(x) \,\mathrm{d}x = \det F_K' \int_{K_0} \Phi_{i_1}^1(X)\Phi_{i_2}^2(X) \,\mathrm{d}X = A_i^0 G_K, \tag{14}$$

where $A^0_i = \int_{K_0} \Phi^1_{i_1}(X)\Phi^2_{i_2}(X)\,\mathrm{d}X$ and $G_K = \det F'_K$. In this case, the reference tensor A^0 is a rank-two tensor (a matrix) and the geometry tensor G_K is a rank-zero tensor (a scalar). By precomputing the reference tensor A^0, we may thus compute the element tensor A^K on each element K by merely multiplying the precomputed reference tensor with the determinant of (the derivative of) the affine mapping F_K.

Test Case 2: Poisson's Equation. As a second example, consider the bilinear form for Poisson's equation,

$$a(v, u) = \int_\Omega \nabla v(x) \cdot \nabla u(x)\,\mathrm{d}x. \tag{15}$$

By a change of variables as previously, we obtain the following representation of the element tensor A^K:

$$\begin{aligned} A^K_i &= \int_K \nabla\phi^{K,1}_{i_1}(x) \cdot \nabla\phi^{K,2}_{i_2}(x)\,\mathrm{d}x \\ &= \det F'_K \frac{\partial X_{\alpha_1}}{\partial x_\beta}\frac{\partial X_{\alpha_2}}{\partial x_\beta}\int_{K_0} \frac{\partial\Phi^1_{i_1}(X)}{\partial X_{\alpha_1}}\frac{\partial\Phi^2_{i_2}(X)}{\partial X_{\alpha_2}}\,\mathrm{d}X = A^0_{i\alpha}G^\alpha_K, \end{aligned} \tag{16}$$

where $A^0_{i\alpha} = \int_{K_0} \frac{\partial\Phi^1_{i_1}(X)}{\partial X_{\alpha_1}}\frac{\partial\Phi^2_{i_2}(X)}{\partial X_{\alpha_2}}\,\mathrm{d}X$ and $G^\alpha_K = \det F'_K \frac{\partial X_{\alpha_1}}{\partial x_\beta}\frac{\partial X_{\alpha_2}}{\partial x_\beta}$. We see that the reference tensor A^0 is here a rank-four tensor and the geometry tensor G_K is a rank-two tensor (one index for each derivative appearing in the form).

Test Case 3: Navier–Stokes. Consider now the nonlinear term $u \cdot \nabla u$ of the incompressible Navier–Stokes equations,

$$\begin{aligned} \dot{u} + u \cdot \nabla u - \nu\Delta u + \nabla p &= f, \\ \nabla \cdot u &= 0. \end{aligned} \tag{17}$$

To solve the Navier–Stokes equations by fixed-point iteration (see, for example, Eriksson et al. [2003]), we write the nonlinear term in the form $u \cdot \nabla u = w \cdot \nabla u$ with $w = u$ and consider w as fixed. We then obtain the following bilinear form:

$$a(v, u) = a_w(v, u) = \int_\Omega v(x) \cdot (w(x) \cdot \nabla)u(x)\,\mathrm{d}x. \tag{18}$$

Note that we may alternatively think of this as a trilinear form, $a = a(v, u, w)$.

Let now $\{w^K_\alpha\}_\alpha$ be the expansion coefficients for w in a finite element basis on the current element K, and let $v[i]$ denote component i of a vector-valued function v. Furthermore, assume that u and w are discretized using the same discrete space $V = V_2$. We then obtain the following representation of the element tensor A^K:

$$\begin{aligned} A^K_i &= \int_K \phi^{K,1}_{i_1}(x) \cdot (w(x) \cdot \nabla)\phi^{K,2}_{i_2}(x)\,\mathrm{d}x \\ &= \det F'_K \frac{\partial X_{\alpha_3}}{\partial x_{\alpha_1}} w^K_{\alpha_2} \int_{K_0} \Phi^1_{i_1}[\beta](X)\Phi^2_{\alpha_2}[\alpha_1](X)\frac{\partial\Phi^2_{i_2}[\beta](X)}{\partial X_{\alpha_3}}\,\mathrm{d}X = A^0_{i\alpha}G^\alpha_K, \end{aligned} \tag{19}$$

where $A^0_{i\alpha} = \int_{K_0} \Phi^1_{i_1}[\beta](X)\Phi^2_{\alpha_2}[\alpha_1](X)\frac{\partial\Phi^2_{i_2}[\beta](X)}{\partial X_{\alpha_3}}\,\mathrm{d}X$ and $G^\alpha_K = \det F'_K \frac{\partial X_{\alpha_3}}{\partial x_{\alpha_1}} w^K_{\alpha_2}$. In this case, the reference tensor A^0 is a rank-five tensor and the geometry tensor

G_K is a rank-three tensor (one index for the derivative, one for the function w, and one for the scalar product).

Test Case 4: Linear Elasticity. Finally, consider the strain-strain term of linear elasticity,

$$\begin{aligned} a(v,u) &= \int_\Omega \frac{1}{4}(\nabla v + (\nabla v)^\top) : (\nabla u + (\nabla u)^\top)\,\mathrm{d}x \\ &= \int_\Omega \frac{1}{4}\frac{\partial v_i}{\partial x_j}\frac{\partial u_i}{\partial x_j}\,\mathrm{d}x + \frac{1}{4}\frac{\partial v_i}{\partial x_j}\frac{\partial u_j}{\partial x_i}\,\mathrm{d}x + \frac{1}{4}\frac{\partial v_j}{\partial x_i}\frac{\partial u_i}{\partial x_j}\,\mathrm{d}x + \frac{1}{4}\frac{\partial v_j}{\partial x_i}\frac{\partial u_j}{\partial x_i}\,\mathrm{d}x. \quad (20)\end{aligned}$$

Considering here only the first term, a change of variables leads to the following representation of the element tensor $A^{K,1}$:

$$\begin{aligned} A_i^{K,1} &= \int_K \frac{1}{4}\frac{\partial \phi_{i_1}^{K,1}{}_{[\beta_1]}(x)}{\partial x_{\beta_2}}\frac{\partial \phi_{i_2}^{K,2}{}_{[\beta_1]}(x)}{\partial x_{\beta_2}}\,\mathrm{d}x \\ &= \frac{1}{4}\det F_K'\frac{\partial X_{\alpha_1}}{\partial x_{\beta_2}}\frac{\partial X_{\alpha_2}}{\partial x_{\beta_2}}\int_{K_0}\frac{\partial \Phi_{i_1}^1{}_{[\beta_1]}(X)}{\partial X_{\alpha_1}}\frac{\partial \Phi_{i_2}^2{}_{[\beta_1]}(X)}{\partial X_{\alpha_2}}\,\mathrm{d}X = A_{i\alpha}^{0,1}G_{K,1}^{\alpha}, \quad (21)\end{aligned}$$

where $A_{i\alpha}^{0,1} = \int_{K_0}\frac{\partial \Phi_{i_1}^1{}_{[\beta_1]}(X)}{\partial X_{\alpha_1}}\frac{\partial \Phi_{i_2}^2{}_{[\beta_1]}(X)}{\partial X_{\alpha_2}}\,\mathrm{d}X$ and $G_{K,1}^{\alpha} = \frac{1}{4}\det F_K'\frac{\partial X_{\alpha_1}}{\partial x_{\beta_2}}\frac{\partial X_{\alpha_2}}{\partial x_{\beta_2}}$. Here, the reference tensor $A^{0,1}$ is a rank-four tensor and the geometry tensor $G_{K,1}$ is a rank-two tensor (one index for each derivative).

3.4 Extensions

The current implementation of FFC supports only affinely mapped Lagrange elements and linear problems, but it is interesting to consider the generalization of our approach to other kinds of function spaces, such as Raviart-Thomas [Raviart and Thomas 1977] elements for $H(\mathrm{div})$, and curvilinear mappings such as arise with isoparametric elements, as well as how nonlinear problems may also be automated.

3.4.1 *$H(div)$ and $H(curl)$ Conforming Elements.* The implemention of $H(\mathrm{div})$ or $H(\mathrm{curl})$ elements requires two kinds of generalizations to FFC. First of all, the basis functions are mapped from the reference element by the Piola transform [Brezzi and Fortin 1991], rather than the standard change of coordinates. With $F_K : K_0 \to K$ as the standard affine mapping for an element K, F_K' the Fréchet derivative of the mapping, and $\det F_K'$ its determinant, the Piola mapping is defined by $\mathcal{F}_K(\Psi) = \frac{1}{\det F_K'}F_K'(\Psi \circ (F_K)^{-1})$. Since our tools already track Jacobians and determinants for differentiating through affine mappings, it should be straightforward to support the Piola mapping. Second, defining the mapping between local and global degrees of freedom becomes more complicated, as we must keep track of the directions of vector components, as done in FEMSTER [Castillo et al. 2004; 2005].

As an example of using the Piola transform, we consider Raviart-Thomas [Raviart and Thomas 1977] elements with the standard (vector-valued) nodal basis $\{\Psi_i\}_{i=1}^d$ on the reference element. We compute the mass matrix M with

$M_{i_1 i_2} = a(\psi_{i_1}, \psi_{i_2})$ and the bilinear form a given by

$$a(v, u) = \int_\Omega v(x) \cdot u(x) \, \mathrm{d}x. \tag{22}$$

On K, the basis functions are given by $\psi_i^K = \mathcal{F}_K(\Psi_i)$ and computing the element tensor A^K, we obtain

$$\begin{aligned} A_i^K &= \int_K \psi_{i_1}^K(x) \cdot \psi_{i_2}^K(x) \, \mathrm{d}x \\ &= \frac{1}{\det F_K'} \frac{\partial x_\beta}{\partial X_{\alpha_1}} \frac{\partial x_\beta}{\partial X_{\alpha_2}} \int_{K_0} \Psi_{i_1[\alpha_1]} \Psi_{i_2[\alpha_2]} \, \mathrm{d}X = A_{i\alpha}^0 G_K^\alpha, \end{aligned} \tag{23}$$

where $A_{i\alpha}^0 = \int_{K_0} \Psi_{i_1[\alpha_1]} \Psi_{i_2[\alpha_2]} \, \mathrm{d}X$ and $G_K^\alpha = \frac{1}{\det F_K'} \frac{\partial x_\beta}{\partial X_{\alpha_1}} \frac{\partial x_\beta}{\partial X_{\alpha_2}}$. We see that, just as for Poisson, the reference tensor A^0 is rank-four and the geometry tensor G_K is rank-two.

3.4.2 *Curvilinear Elements.* Our techniques may be generalized to cases in which the Jacobian varies spatially within elements, such as when curvilinear elements general quadrilaterals, or hexahedra are used. In cases such, we can replace integration with summation over quadrature points and obtain a formulation based on tensor contraction, albeit with higher run-time complexity than for affine elements.

To illustrate this, we consider the case of the very simple bilinear form

$$a(v, u) = \int_\Omega v \frac{\partial u}{\partial x_1} \, \mathrm{d}x. \tag{24}$$

If the mapping from the reference element K_0 to an element K is curvilinear, then we will be unable to pull the Jacobian and derivatives out of the integral. However, we will still be able to phrase a run-time tensor contraction with an extra index for the quadrature points. The element matrix for Equation (24) is

$$A_i^K = \int_K \phi_{i_1} \frac{\partial \phi_{i_2}}{\partial x_1} \, \mathrm{d}x, \tag{25}$$

and changing coordinates, we obtain

$$A_i^K = \int_{K_0} \Phi_{i_1} \frac{\partial \Phi_{i_2}}{\partial X_{\alpha_1}} \frac{\partial X_{\alpha_1}}{\partial x_1} \det F_K' \, \mathrm{d}X. \tag{26}$$

Approximating the integral by quadrature, we let $\{X_k\}_{k=1}^N$ be a set of quadrature points on the reference element K_0, with $\{w_k\}_{k=1}^N$ as the corresponding weights. We thus obtain the representation

$$A_i^K \approx \tilde{A}_i^K = \sum_{k=1}^N w_k \Phi_{i_1}(X_k) \frac{\partial \Phi_{i_2}}{\partial X_{\alpha_1}}(X_k) \frac{\partial X_{\alpha_1}}{\partial x_1}(X_k) \det F_K'(X_k) = \tilde{A}_{i\alpha}^0 G_K^\alpha, \tag{27}$$

where

$$\tilde{A}_{i\alpha}^0 = w_{\alpha_2} \Phi_{i_1}(X_{\alpha_2}) \frac{\partial \Phi_{i_2}}{\partial X_{\alpha_1}}(X_{\alpha_2}) \tag{28}$$

and

$$G_\alpha = \frac{\partial X_{\alpha_1}}{\partial x_1}(X_{\alpha_2}) \det F_K'(X_{\alpha_2}). \tag{29}$$

Note that $\tilde{A}^0_\alpha$ may be computed entirely at compile-time, as can an expression for G_α, whereas the values of G_α depend on the geometry given at run-time. The tensors to be contracted at run-time have one extra dimension compared to a situation where the mapping is affine. Still, this computation (once the geometry tensor is computed on an element) is readily phrased as a matrix-vector product. Hence, we have given an example of how the more traditional way of expressing quadrature and our precomputation are both instances of a high-rank tensor contraction. We could interpret this formulation as saying that quadrature (expressed as we have here) gives a general model for computation and that precomputation is possible as a compile-time optimization in the case where the mapping is affine.

An open interesting problem would be to study under which conditions we could specialize a system such as FFC to use bases with tensor-product decompositions (available on unstructured, as well as structured, shapes [Karniadakis and Sherwin 1999]) and automatically generate efficient matrix-vector products, as are manually implemented in spectral element methods.

3.4.3 *Nonlinear Forms.* For a nonlinear variational problem,

$$a(v, u) = L(v) \quad \forall v \in \hat{V}, \tag{30}$$

where a is nonlinear in u, we can solve by direct fixed-point iteration on the unknown u, as discussed earlier for Test Case 3, or we can compute the (Fréchet) derivative a'_u of the nonlinear form a and solve by Newton's method. For multilinear forms, the nonlinear residual can be defined and the Jacobian can be constructed using the current capabilities of FFC.

To solve the variational problem of Equation (30) by Newton's method, we differentiate with respect to u to obtain a variational problem for the *increment* δu:

$$a'_u(v, \delta u) = -(a(v, u) - L(v)) \quad \forall v \in \hat{V}. \tag{31}$$

As an example, consider the nonlinear Poisson's equation $-c(u)\Delta u = f$ with $c(u) = u$. For more general c, considering the projection of $c(u)$ into the finite element space leads to a multilinear form. Differentiating the form with respect to u, we obtain the following variational problem: Find $\delta u \in V$ such that

$$\int_\Omega \delta u \nabla v \cdot \nabla u + u \nabla v \cdot \nabla \delta u \, \mathrm{d}x = \int_\Omega v f \, \mathrm{d}x - \int_\Omega u \nabla v \cdot \nabla u \, \mathrm{d}x \quad \forall v \in \hat{V}. \tag{32}$$

Our current capabilities would allow us to define two forms: one to evaluate the nonlinear residual and another to construct the Jacobian matrix. This is sufficient to set up a nonlinear solver. Obviously, extending FFC to symbolically differentiate the nonlinear form would be more satisfying. We remark that the code Sundance [Long 2003] currently has such capabilities. It should be possible to leverage such tools in the future to combine our precomputation techniques with automatic differentiation.

3.5 Optimization

We consider here three different kinds of optimization that could be built into FFC in the future. For one of them, the current code is generated entirely straightline as a sequence of arithmetic and assignment. It should be possible to store the tensor A^0 in a contiguous array. Moreover, each G_K may be considered a tensor or flattened into a vector. In the latter case, the action of forming the element matrix for one element may be written as a matrix-vector multiply using the level-2 BLAS. Once this observation is made, it is straightforward to see that we could form several G_K vectors and make better use of cache by computing several element matrices at once by a matrix-matrix multiply and the level-3 BLAS.

This corresponds to a coarse-grained optimization. In other work [Kirby et al. 2005a, 2005b], we have seen that for many forms, the entries in A^0 are related such that various entries of the element matrices may be formed in fewer operations. For example, if two entries of A^0 are close together in Hamming distance, then the contraction of one entry with G_K can be computed efficiently from the other. As our code for optimization, FErari, evolves, we will integrate it with FFC as an optimizing backend. It will be simple to compare the output of FErari to the best performance using the BLAS, and let FFC output the best of the two (which may be highly problem-dependent).

Finally, optimizations that arise from the variational form itself will be fruitful to explore in the future. For example, it should be possible to detect when a variational form is symmetric within FFC, as this leads to fewer operations to form the associated matrix. Moreover, for forms over vector-valued elements that have a Cartesian product basis (each basis function has support in only one component), other kinds of optimizations are appropriate. For example, the viscosity operator for Navier-Stokes is the vector Laplacian, which can be written as a block-diagonal matrix with one axis for each spatial dimension. By "taking apart" the basis functions, we hope to uncover this block structure, which will lead to more efficient compilation and hopefully, to more efficient code.

4. COMPLEXITY OF FORM EVALUATION

We now compare the proposed algorithm based on tensor representation to the standard quadrature-based approach. As we shall see, tensor representation can be much more efficient than quadrature for a wide range of forms.

4.1 Basic Assumptions and Notation

To analyze the complexity of form evaluation, we make the following simplifying assumptions:

- the form is bilinear, that is, $r = |i| = 2$;
- the form can be represented as one tensor product, that is, $A_i^K = A_{i\alpha}^0 G_K^\alpha$;
- the basis functions are scalar; and
- integrals are computed exactly, that is, the order of the quadrature rule must match the polynomial order of the integrand.

We shall use the following notation: q is the polynomial order of the basis functions on every element, p is the total polynomial order of the integrand of the form, d is the dimension of Ω, n is the dimension of the function space on an element, and N is the number of quadrature points needed to integrate polynomials of degree p exactly.

Furthermore, let n_f be the number of functions appearing in the form. For the aforementioned Test Cases 1–4, $n_f = 0$ in all cases, except Test Case 3 (Navier–Stokes) where $n_f = 1$. We use n_D to denote the number of differential operators. For Test Cases 1–4, we have $n_D = 0$ in Case 1 (the mass matrix), $n_D = 2$ in Case 2 (Poisson), $n_D = 1$ in Case 3 (Navier–Stokes), and $n_D = 2$ in Case 4 (linear elasticity).

4.2 Complexity of Tensor Representation

The element tensor A^K has n^2 entries. The number of basis functions n for polynomials of degree q in d dimensions is $\sim q^d$. To compute each entry A_i^K of the element tensor A^K using tensor representation, we need to compute the tensor product between $A_{i\cdot}^0$ and G_K. The geometry tensor G_K has rank $n_f + n_D$ and the number of entries of G_K is $n^{n_f} d^{n_D}$. The cost for computing the n^2 entries of the element tensor A^K using tensor representation is thus

$$T_T \sim n^2 n^{n_f} d^{n_D} \sim (q^d)^2 (q^d)^{n_f} d^{n_D} \sim q^{2d+n_f d} d^{n_D}. \tag{33}$$

Note that there is no run-time cost associated with computing the tensor representation of Equation (8), since this is computed at compile-time. Also note that we have not taken into account any of the optimizations discussed in Section 3.5. These optimizations can in some cases significantly reduce the operation count.

4.3 Complexity of Quadrature

To compute each entry A_i^K of the element tensor A^K using quadrature, we need to evaluate an integrand of total order p at N quadrature points. The number of quadrature points needed to integrate polynomials of order p exactly in d dimensions is $N \sim p^d$. Since the form is bilinear with basis functions of order q, the total order is $p = 2q + n_f q - n_D$. It is difficult to estimate precisely the cost of evaluating the integrand at each quadrature point, but a reasonable estimate is $n_f + n_D d + 1$. Note that we assume all basis functions and their derivatives have been pretabulated at all quadrature points on the reference element.

We thus obtain the following estimate of the total cost for computing the n^2 entries of the element tensor A^K using quadrature:

$$\begin{aligned} T_Q &\sim n^2 N (n_f + n_D d + 1) \sim (q^d)^2 p^d (n_f + n_D d + 1) \\ &\sim q^{2d} (2q + n_f q - n_D)^d (n_f + n_D d + 1). \end{aligned} \tag{34}$$

4.4 Tensor Representation vs. Quadrature

Comparing tensor representation with quadrature, the speedup of tensor representation is

$$T_Q/T_T \sim \frac{(2q + n_f q - n_D)^d (n_f + n_D d + 1)}{q^{n_f d} d^{n_D}}. \tag{35}$$

We immediately note that there can be a significant speedup for $n_f = 0$, since T_T/n^2 is now independent of the polynomial degree q. In particular, we note that for the mass matrix ($n_f = n_D = 0$) the speedup is $T_Q/T_T \sim (2q)^d$, and for Poisson's equation ($n_f = 0$, $n_D = 2$) the speedup is $T_Q/T_T \sim (2q - 2)^d$ $(2d + 1)/d^2$. As we shall next see, the speedup for Test Cases 1–4 is significant, even for $q = 0$.

On the other hand, we note that quadrature may be more efficient if n_f is large. Also, if we take into account that underintegration is possible (choosing a smaller N than given by the polynomial order p), it is less clear which approach is most efficient for any given case. It is known [Ciarlet 1976] that second-order elliptic variational problems using polynomials of degree q require an integration rule that is exact only on polynomials of degree $2q - 2$ to ensure the proper convergence rate, regardless of the arity of the form. However, our overall flop count is lower for bilinear and likely, trilinear forms, and at any rate, our code is simpler for compilers to optimize than quadrature loops.

Ultimately, we may thus imagine an intelligent system that automatically makes the choice between tensor representation and quadrature in each specific situation.

5. IMPLEMENTATION

We now discuss a number of important aspects of the implementation of the form compiler FFC. We also write down the forms for the Test Cases discussed earlier in Section 3.3 in the language of the form compiler FFC. Basically, we can consider FFC's functionality as broken into three phases. First, it takes an expression for a multilinear form and generates the tensor A^0. While doing this, it derives an expression for the evaluation of the element tensor G_K from the affine mapping and coefficients of the form. Finally, it generates code for evaluating G_K and contracting it with A^0, and for constructing the local-to-global mapping.

5.1 Parsing of Forms

The form compiler FFC implements a domain-specific language (DSL) for variational forms, using Python as the host language. A language of variational forms is obtained by overloading the appropriate operators, including addition `+`, subtraction `-`, multiplication (`*`), and differentiation `.dx(·)` for a hierarchy of classes corresponding to the algebra $\mathcal{A}$ discussed previously in Section 3.2. FFC thus uses the built-in parser of Python to process variational forms.

Table I. Unary Operators and Their Results for the Classes `BasisFunction` (B), `Function` (F), `Product` (P), and `Sum` (S)

op	B	F	P	S
-	P	S	P	S
.dx(·)	P	S	S	S

Table II. Binary Operators and Their Results for the Classes `BasisFunction` (B), `Function` (F), `Product` (P), and `Sum` (S)

+/-	B	F	P	S
B	S	S	S	S
F	S	S	S	S
P	S	S	S	S
S	S	S	S	S

*	B	F	P	S
B	P	S	P	S
F	S	S	S	S
P	P	S	P	S
S	S	S	S	S

5.2 Generation of the Tensor Representation

The basic elements of the algebra are objects of the class `BasisFunction`, representing (derivatives of) basis functions of some given finite element space. Each `BasisFunction` is associated with a particular finite element space and different `BasisFunctions` may be associated with different finite element spaces. Products of scalars and (derivatives of) basis functions are represented by the class `Product`, and sums of such products are represented by the class `Sum`. In addition, we include a class `Function`, representing linear combinations of basis functions (coefficients). In Tables I and II, we summarize the basic unary and binary operators, respectively, implemented for the hierarchy of classes.

Note that by declaring a common base class for `BasisFunction`, `Product`, `Sum`, and `Function`, some of the operations can be grouped together to simplify the implementation. As a result, most operators will directly yield a `Sum`. Also note that the algebra of `Sums` is closed under the preceding operations.

By associating with each object one or more *indices*, implemented by the class `Index`, an object of type `Sum` automatically represents a tensor, and by differentiating between different types of indices, an object of type `Sum` automatically encodes the tensor representation of Equation (8). FFC differentiates between four different types of indices: *primary*, *secondary*, *auxiliary*, and *fixed*. A primary index (i) is associated with the multiindex of the element tensor A^K, a secondary index (α) is associated with the multiindex of the geometry tensor G_K, and thus, secondary indices indicate along which dimensions to compute the tensor product $A^0_{i\alpha} G^\alpha_K$. Auxiliary indices (β) are internal within the reference tensor A^0 or geometry tensor G_K and must be repeated exactly twice; summation is performed over each auxiliary index β before the tensor product is computed by summation over secondary indices α. Finally, a fixed index is a given constant index that cannot be evaluated. Fixed indices are used to represent, for example, a derivative in a fixed coordinate direction.

Implicit in our algebra is a *grammar* for multilinear forms. We could explicitly write an EBNF grammar and use tools such as `lex` and `yacc` to create a compiler for a domain-specific language. However, by limiting ourselves to

overloaded operators, we successfully embed our language as a high-level library in Python.

To make this concrete, consider Test Case 2 of Section 3.3, Poisson's equation. The tensor representation $A_i^K = A_{i\alpha}^0 G_K^\alpha$ is given by

$$\begin{aligned} A_{i\alpha}^0 &= \int_{K_0} \frac{\partial \Phi_{i_1}^1(X)}{\partial X_{\alpha_1}} \frac{\partial \Phi_{i_2}^2(X)}{\partial X_{\alpha_2}} \, \mathrm{d}X, \\ G_K^\alpha &= \det F_K' \frac{\partial X_{\alpha_1}}{\partial x_\beta} \frac{\partial X_{\alpha_2}}{\partial x_\beta}. \end{aligned} \tag{36}$$

There are in this instance two primary indices (i_1 and i_2), two secondary indices (α_1 and α_2), and one auxiliary index (β).

5.3 Evaluation of Integrals

Once the tensor representation of Equation (8) has been generated, FFC computes all entries of the reference tensor(s) by quadrature on the reference element. The quadrature rule is automatically chosen to match the polynomial order of each integrand. FFC uses FIAT [Kirby 2004; Kirby 2005] as the finite element backend; FIAT generates the set of basis functions as well as the quadrature rule, and evaluates both the basis functions and their derivatives at the quadrature points.

Although FIAT supports many families of finite elements, the current version of FFC only supports general-order continuous/discontinuous Lagrange finite elements and first-order Crouzeix–Raviart finite elements on triangles and tetrahedra (or any other finite element with nodes given by pointwise evaluation). Support for other families of finite elements will be added in future versions.

Computing integrals is the most expensive step in the compilation of a form. The typical run-time (of the compiler) ranges between 0.1 and 30 seconds, depending on the type of form and finite element.

5.4 Generation of Code

When a form has been parsed, the tensor representation generated, and all integrals computed, code is generated for the evaluation of geometry tensors and tensor products. The form compiler FFC has been designed to allow for generation of code in multiple different languages. Code is generated according to a specific *format* (which is essentially a Python dictionary) that controls the output code being generated (see Figure 5). The current version of FFC supports four output formats: C++ code for DOLFIN [Hoffman et al. 2005; Hoffman and Logg 2002], LaTeX code (for verification and presentation purposes), a raw format that merely lists the values of reference tensors, and the recently added ASE format [Balay et al. 2005a] for the compilation of forms for the next generation of PETSc. FFC can be easily extended with new output formats, including, for example, Python, C, or Fortran.

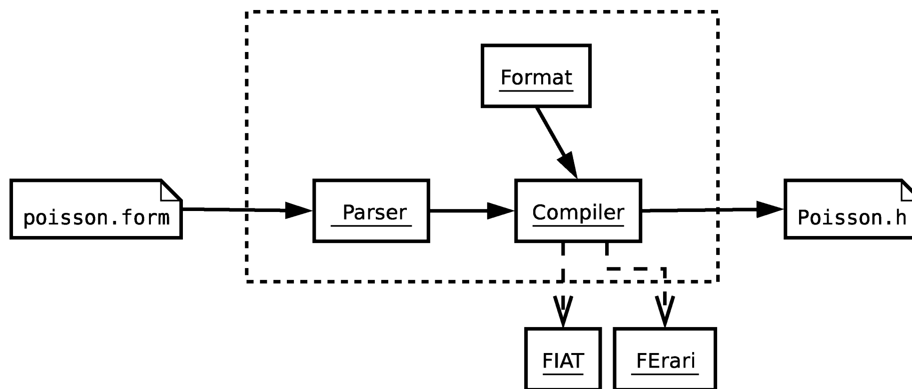

Fig. 5. Diagram of the components of the form compiler FFC.

Table III. Complete Code for Specification of Test Case 2, Poisson's equation, with piecewise cubics on tetrahedra in the language of FFC

```
element = FiniteElement("Lagrange", "tetrahedron", 3)

v = BasisFunction(element)
u = BasisFunction(element)
f = Function(element)
i = Index()

a = v.dx(i)*u.dx(i)*dx
L = v*f*dx
```

Alternatively, the form can be specified in terms of standard operators: `a = dot(grad(v), grad(u))*dx`.

5.5 Input/Output

FFC can be used either as a Python package or from the command-line. We here give a brief description of how FFC can be called from the command-line to generate C++ code for DOLFIN. To use FFC from the command-line, we specify the form in a text file in a special language for variational forms, which is simply Python equipped with the hierarchy of classes and operators discussed previously in Section 5.1. In Table III we give the complete code for the specification of Test Case 2, Poisson's equation. Note that FFC uses tensor notation, and thus the summation over the index `i` is implicit. Also note that the integral over an element K is denoted by `*dx`.

Assuming that the form has been specified in the file `Poisson.form`, it can be compiled using the command `ffc Poisson.form`. This generates the C++ file `Poisson.h`, to be included in a DOLFIN program. In Table IV, we include part of the output generated by FFC with input given by the code from Table III. In addition to this code, FFC generates code for the mapping $\iota(\cdot, \cdot)$ from local-to-global degrees of freedom for each finite element space associated with the form. Note that the values of the 10×10 element tensor A^K (in the case of cubics on triangles) are stored as one contiguous array (`block`),

Table IV. Part of the Code Generated by FFC for the Input Code from Table III

```
void eval(double block[], const AffineMap& map) const
{
 // Compute geometry tensors
double G0_0_0 = map.det*(map.g00*map.g00 + map.g01*map.g01);
double G0_0_1 = map.det*(map.g00*map.g10 + map.g01*map.g11);
double G0_1_0 = map.det*(map.g10*map.g00 + map.g11*map.g01);
double G0_1_1 = map.det*(map.g10*map.g10 + map.g11*map.g11);

 // Compute element tensor
block[0] = 4.249999999999996e-01*G0_0_0 + 4.249999999999995e-01*G0_0_1 +
           4.249999999999995e-01*G0_1_0 + 4.249999999999995e-01*G0_1_1;
block[1] = -8.749999999999993e-02*G0_0_0 - 8.749999999999995e-02*G0_0_1;
block[2] = -8.750000000000005e-02*G0_1_0 - 8.750000000000013e-02*G0_1_1;
...
block[99] = 4.049999999999997e+00*G0_0_0 + 2.024999999999998e+00*G0_0_1 +
            2.024999999999998e+00*G0_1_0 + 4.049999999999995e+00*G0_1_1;
}
```

since this is what the linear algebra backend of DOLFIN (PETSc) requires for assembly.

5.6 Completing the Toolchain

With the FEniCS project [Hoffman et al. 2005], we have the beginnings of a working system realizing (in part) the Automation of Computational Mathematical Modeling (ACMM), and the form compiler FFC is just one of several components needed to complete the toolchain. FIAT automates the generation of finite element spaces and FFC automates the evaluation of variational forms. Furthermore, PETSc [Balay et al. 1997; 2004; 2005b], automating the solution of discrete systems, is used as the solver backend of FEniCS. A common C++ interface for different FEniCS components is provided by DOLFIN.

A complete automation of CMM, as outlined in Logg [2004], is a major task and we hope that by a modular approach, we can contribute to this automation.

6. BENCHMARK RESULTS

As noted earlier, the speedup for the code generated by the form compiler FFC can, in many cases, be significant. We now present a comparison with the standard quadrature-based approach for the test cases discussed in Section 3.3.

The forms were compiled for a range of polynomial degrees using FFC version 0.1.6. This version of FFC does not take into account any of the optimizations discussed in Section 3.5, other than not generating code for multiplication with zero entries of the reference tensor.

For quadrature-based code, all basis functions and their derivatives were pretabulated at the quadrature points using FIAT. Loops for all scalar products were completely unrolled.

In all cases, we have used the *collapsed-coordinate* Gauss-Jacobi rules described by Karniadakis and Sherwin [1999]. These take tensor-product Gaussian integration rules over the square and both cube as well as map them

Table V. Speedups T_Q/T_T for Test Cases 1–4 in 2D and 3D

Form	$q = 1$	$q = 2$	$q = 3$	$q = 4$	$q = 5$	$q = 6$	$q = 7$	$q = 8$
Mass 2D	12	31	50	78	108	147	183	232
Mass 3D	21	81	189	355	616	881	1442	1475
Poisson 2D	8	29	56	86	129	144	189	236
Poisson 3D	9	56	143	259	427	341	285	356
Navier–Stokes 2D	32	33	53	37	—	—	—	—
Navier–Stokes 3D	77	100	61	42	—	—	—	—
Elasticity 2D	10	43	67	97	—	—	—	—
Elasticity 3D	14	87	103	134	—	—	—	—

to the reference simplex. These rules are not the best-known (see, for example, Dunavant [1985]), but they are fairly simple to generate for an arbitrary degree. Eventually, these rules will be integrated with FIAT, but even if we reduce the number of quadrature points by a factor of five, FFC still outperforms quadrature.

The codes were compiled with gcc (g++) version 3.3.6 and the benchmark results we present to follow were obtained on an Intel Pentium 4 (CPU 3.0 GHz, 2GB RAM) running Debian GNU/Linux. The times reported are for the computation of each entry of the element tensor on one million elements (scaled). The total time can be obtained by multiplying with n^2 the number of entries of the element tensor. The complete source-code for the benchmarks can be obtained from the FEniCS website [Hoffman et al. 2005].

6.1 Summary of Results

In Table V, we summarize the results for Test Cases 1–4. In all cases, the speedup T_Q/T_T is significant, ranging between a factor of 10–1500.

From Section 4, we know that the speedup for the mass matrix should grow as q^d, but from Table V it is clear that the speedup is not quadratic for $d = 2$ and $d = 3$; an optimum seems to be reached at around $q = 8$.

The reason that the predicted speedups are not obtained in practice is that the complexity estimates presented in Section 4 only account for the number of floating-point operations. When the polynomial degree q grows, the number of lines of code generated by the form compiler increases. FFC unrolls all loops and generates one line of code for each entry of the element tensor to be computed. For a bilinear form, the number of entries is $n^2 \sim q^{2d}$. With $q = 8$, the number of lines of code generated is about 25,000 for the mass matrix and Poisson in 3D (see Figure 6). As the number of lines of code grows, memory access becomes more important and dominates run-time. Using BLAS to compute tensor products, as discussed earlier, might lead to more efficient memory traffic.

Note, however, that although the optimal speedup is not obtained, the speedup is in all cases significant, even at $q = 1$.

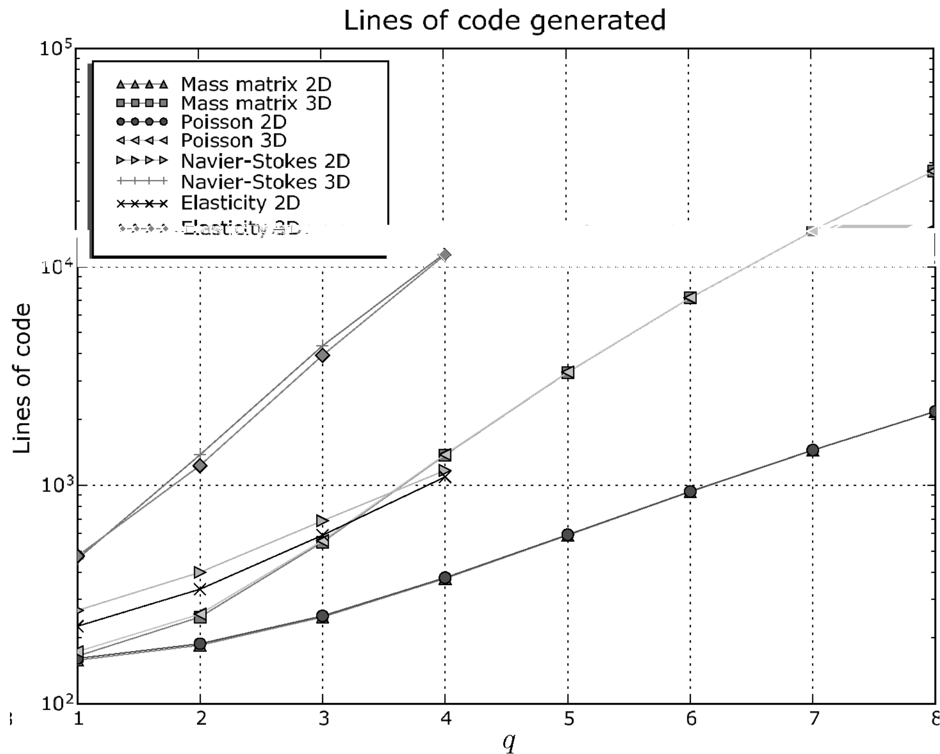

Fig. 6. Lines of code generated by the form compiler FFC as function of the polynomial degree q.

6.2 Results for Test Cases

In Figures 7–10, we present the results for Test Cases 1–4, discussed in Section 3.3. In connection to each of the results, we include the specification of the form in the language used by the form compiler FFC. Because of limitations in the current implementation of FFC, the comparison is made for polynomial order $q \leq 8$ in Test Cases 1 and 2, and $q \leq 4$ in Test Cases 3 and 4. Higher polynomial order is possible, but is very costly to compile (compare Figure 6).

7. CONCLUDING REMARKS AND FUTURE DIRECTIONS

We have demonstrated a proof-of-concept form compiler that for a wide range of variational forms can generate code that gives significant speedups compared to the standard quadrature-based approach.

The form compiler FFC is still in its early stages of development, but is already in production use. A number of basic modules based on FFC have been implemented in DOLFIN and others are currently being developed (Navier–Stokes and updated elasticity). This will serve as a testbed for the future development of FFC.

Future plans for FFC include adding support for integrals over the boundary (adding the operator `*ds` to the language), support for the automatic

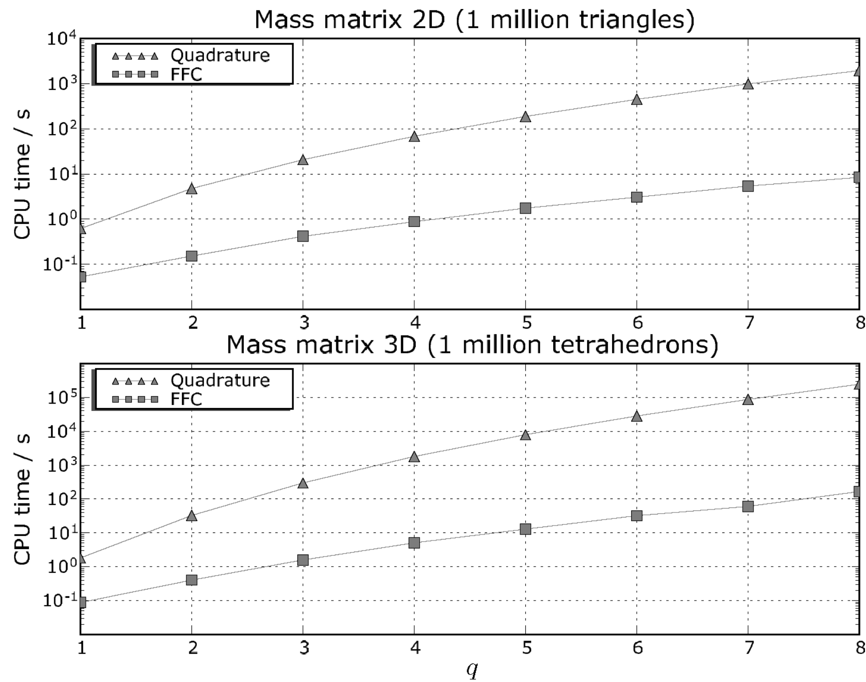

Fig. 7. Benchmark results for Test Case 1, the mass matrix, specified in FFC by `a = v*u*dx`.

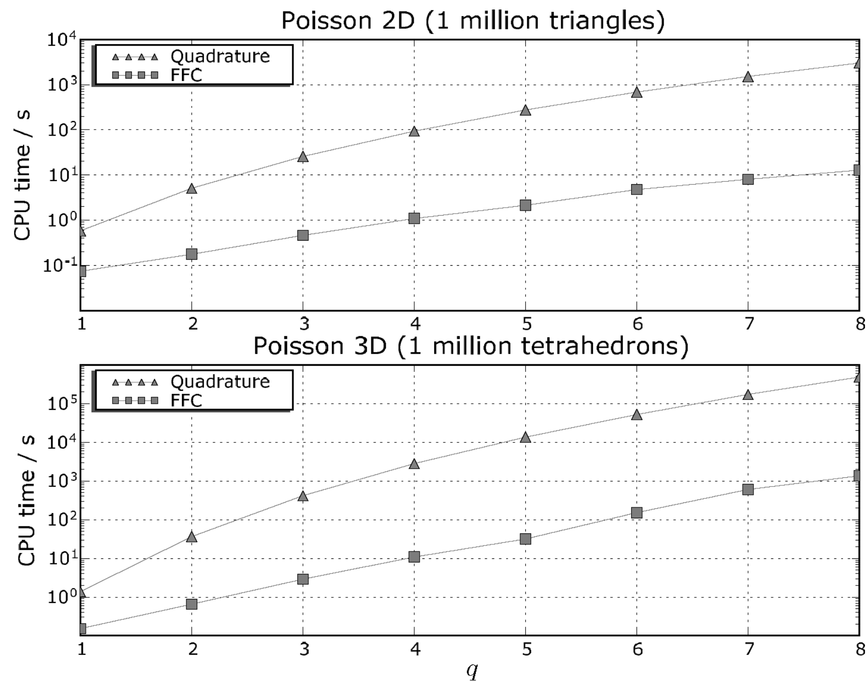

Fig. 8. Benchmark results for Test Case 2, Poisson's equation, specified in FFC by `a = v.dx(i)*u.dx(i)*dx`.

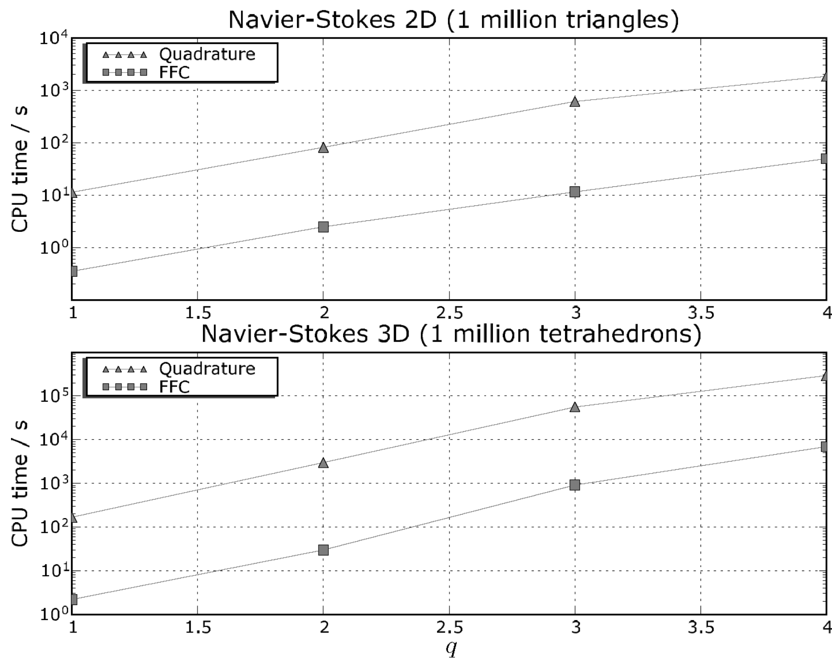

Fig. 9. Benchmark results for Test Case 3, the nonlinear term of the incompressible Navier–Stokes equations, specified in FFC by `a = v[i]*w[j]*u[i].dx(j)*dx`.

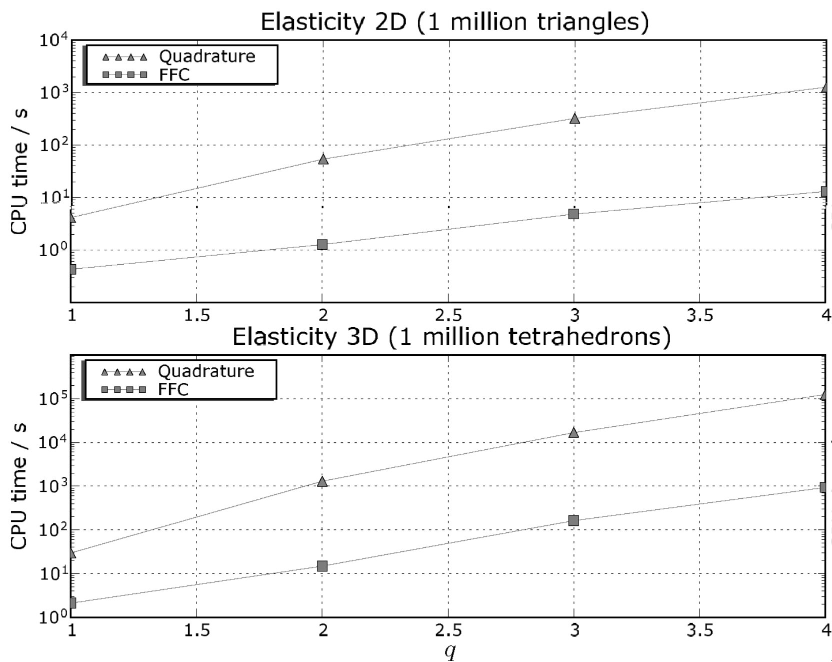

Fig. 10. Benchmark results for Test Case 4, the strain-strain term of linear elasticity, specified in FFC by `a = 0.25*(v[i].dx(j) + v[j].dx(i)) * (u[i].dx(j) + u[j].dx(i)) * dx`.

differentiation of nonlinear forms and automatic generation of dual problems and *a posteriori* error estimators [Eriksson et al. 1995; Becker and Rannacher 2001], optimization through FErari [Kirby et al. 2005a; 2006b], and adding support for new families of finite elements, including elements that require nonaffine mappings from the reference element. In addition to general-order continuous/discontinuous Lagrange finite elements and Crouzeix-Raviart [Crouzeix and Raviart 1973] finite elements, the plan is to add support for Raviart-Thomas [Raviart and Thomas 1977], Nedelec [Nédélec 1980], Brezzi-Douglas-Marini [Brezzi et al. 1985], Brezzi-Douglas-Fortin-Marini [Brezzi and Fortin 1991], Taylor-Hood [Boffi 1997; Brenner and Scott 1994], and Arnold-Winther [Arnold and Winther 2002] elements.

Furthermore, the fact that Python is an interpreted language does impose a penalty on the performance of the compiler (but not on the generated code). However, this can be overcome by more extensive use of BLAS in numerically intensive parts of the compiler, such as the precomputation of integrals on the reference element. We also plan to investigate the use of BLAS for evaluation of tensor products as an alternative to generating explicit unrolled code. Other topics of interest include automatic verification of the correctness of the code generated by the form compiler [Kirby et al. 2005c].

ACKNOWLEDGMENTS

We wish to thank the FEniCS team, particularly Johan Hoffman, Johan Jansson, Claes Johnson, Matthew Knepley, and Ridgway Scott, for substantial suggestions and comments regarding this article. We also want to thank one of the referees for pointing out how tensor representation can be generalized to a quadrature-based evaluation scheme for nonaffine mappings.

REFERENCES

ARNOLD, D. N. AND WINTHER, R. 2002. Mixed finite elements for elasticity. *Numer. Math. 92*, 3, 401–419.

BAGHERI, B. AND SCOTT, R. 2003. Analysa. http://people.cs.uchicago.edu/~ridg/al/aa.html.

BALAY, S., BUSCHELMAN, K., EIJKHOUT, V., GROPP, W. D., KAUSHIK, D., KNEPLEY, M. G., MCINNES, L. C., SMITH, B. F., AND ZHANG, H. 2004. PETSc users manual. Tech. Rep. ANL-95/11 - Revision 2.1.5, Argonne National Laboratory.

BALAY, S., BUSCHELMAN, K., GROPP, W. D., KAUSHIK, D., KNEPLEY, M. G., MCINNES, L. C., SMITH, B. F., AND ZHANG, H. 2005a. Anl anl sidl environmentsidl environment. http://www-unix.mcs.anl.gov/ase/.

BALAY, S., BUSCHELMAN, K., GROPP, W. D., KAUSHIK, D., KNEPLEY, M. G., MCINNES, L. C., SMITH, B. F., AND ZHANG, H. 2005b. PETSc. http://www.mcs.anl.gov/petsc/.

BALAY, S., EIJKHOUT, V., GROPP, W. D., MCINNES, L. C., AND SMITH, B. F. 1997. Efficient management of parallelism in object-oriented numerical software libraries. In *Modern Software Tools in Scientific Computing*, E. Arge et al., Eds. Birkhäuser Press, Cambridge, MA, 163–202.

BANGERTH, W., HARTMANN, R., AND KANSCHAT, G. 2005. `deal.II` differential equations analysis library. http://www.dealii.org.

BECKER, R. AND RANNACHER, R. 2001. An optimal control approach to a posteriori error estimation in finite element methods. *Acta Numerica 10*, 1–102.

BOFFI, D. 1997. Three-Dimensional finite element methods for the Stokes problem. *SIAM J. Numer. Anal. 34*, 2, 664–670.

BRENNER, S. C. AND SCOTT, L. R. 1994. *The Mathematical Theory of Finite Element Methods*. Springer Verlag.

BREZZI, F., DOUGLAS, JR., J., AND MARINI, L. D. 1985. Two families of mixed finite elements for second-order elliptic problems. *Numer. Math. 47*, 2, 217–235.

BREZZI, F. AND FORTIN, M. 1991. *Mixed and hybrid finite element methods*. Springer series in Computational Mathematics, vol. 15. Springer Verlag, New York.

CASTILLO, P., KONING, J., RIEBEN, R., AND WHITE, D. 2004. A discrete differential forms framework for computational electromagnetics. *Comput. Modeling Eng. Sci. 5*, 4, 331–346.

CASTILLO, P., RIEBEN, R., AND WHITE, D. 2005. Femster: An object-oriented class library of discrete differential forms. To appear in *ACM Trans. Math. Softw*.

CIARLET, P. G. 1976. *Numerical Analysis of the Finite Element Method*. Les Presses de l'Universite de Montreal.

CROUZEIX, M. AND RAVIART, P. A. 1973. Conforming and nonconforming finite element methods for solving the stationary stokes equations. *RAIRO Anal. Numer. 7*, 33–76.

DULAR, P. AND GEUZAINE, C. 2005. GetDP: A general environment for the treatment of discrete problems. http://www.geuz.org/getdp/.

DUNAVANT, D. A. 1985. High-Degree efficient symmetrical Gaussian quadrature rules for the triangle. *Int. J. Numer. Methods Eng. 21*, 6, 1129–1148.

ERIKSSON, K., ESTEP, D., HANSBO, P., AND JOHNSON, C. 1995. Introduction to adaptive methods for differential equations. *Acta Numerica 4*, 105–158.

ERIKSSON, K., ESTEP, D., HANSBO, P., AND JOHNSON, C. 1996. *Computational Differential Equations*. Cambridge University Press, New York.

ERIKSSON, K., ESTEP, D., AND JOHNSON, C. 2003. *Applied Mathematics: Body and Soul*. Vol. III. Springer Verlag.

FREE SOFTWARE FOUNDATION. 1991. GNU gpl. http://www.gnu.org/copyleft/gpl.html.

HOFFMAN, J., JANSSON, J., JOHNSON, C., KNEPLEY, M., KIRBY, R. C., LOGG, A., AND SCOTT, L. R. 2005. FEniCS. http://www.fenics.org/.

HOFFMAN, J., JANSSON, J., AND LOGG, A. 2005. DOLFIN. http://www.fenics.org/dolfin/.

HOFFMAN, J. AND LOGG, A. 2002. DOLFIN: Dynamic object-oriented library for FINite element computation. Tech. Rep. 2002–06, Chalmers Finite Element Center Preprint series.

HUGHES, T. J. R. 1987. *The Finite Element Method: Linear Static and Dynamic Finite Element Analysis*. Prentice-Hall, Upper Saddle River, NJ.

KARNIADAKIS, G. E. AND SHERWIN, S. J. 1999. Spectral/hp element methods for CFD. In *Numerical Mathematics and Scientific Computation*. Oxford University Press, New York.

KIRBY, R. C. 2004. FIAT: A new paradigm for computing finite element basis functions. *ACM Trans. Math. Softw. 30*, 502–516.

KIRBY, R. C. 2005. Optimizing FIAT with the level-3 BLAS. Submitted to *ACM Trans. Math. Softw*.

KIRBY, R. C., KNEPLEY, M., LOGG, A., AND SCOTT, L. R. 2005a. Optimizing the evaluation of finite element matrices. To appear in *SIAM J. Sci. Comput*.

KIRBY, R. C., KNEPLEY, M., AND SCOTT, L. R. 2005b. Evaluation of the action of finite element operators. Submitted to *BIT*.

KIRBY, R. C., STROUT, M. M., HOVLAND, P., AND SCOTT, L. R. 2005c. Verification of scientific code using rationality analysis. In preparation.

LANGTANGEN, H. P. 1999. *Computational Partial Differential Equations – Numerical Methods and Diffpack Programming*. Lecture Notes in Computational Science and Engineering. Springer Verlag.

LOGG, A. 2004. Automation of computational mathematical modeling. Ph.D. thesis, Chalmers University of Technology, Sweden.

LONG, K. 2003. Sundance, a rapid prototyping tool for parallel PDE-constrained optimization. In *Large-Scale PDE-Constrained Optimization*. Lecture Notes in Computational Science and Engineering. Springer Verlag.

MACKIE, R. I. 1992. Object-Oriented programming of the finite element method. *Int. J. Num. Meth. Eng. 35*, 425–436.

MASTERS, I., USMANI, A. S., CROSS, J. T., AND LEWIS, R. W. 1997. Finite element analysis of solidification using obect-oriented and parallel techniques. *Int. J. Numer. Meth. Eng. 40*, 2891–2909.

Nédélec, J.-C. 1980. Mixed finite elements in $\mathbf{R}^3$. *Numer. Math. 35*, 3, 315–341.
Pironneau, O., Hecht, F., and Hyaric, A. L. 2005. FreeFEM. http://www.freefem.org/.
Raviart, P.-A. and Thomas, J. M. 1977. A mixed finite element method for second-order elliptic problems. In *Mathematical Aspects of Finite Element Methods; Proceeding of the Conference Consiglio Naz. delle Ricerche (CNR)*, Rome, held in 1975. Lecture Notes in Mathematics, vol. 606 Springer Verlag, 292–315.

Received June 2005; revised October 2005; accepted November 2005