

A study of vectorization for matrix-free finite element methods

Tianjiao Sun¹, Lawrence Mitchell², Kaushik Kulkarni³,
Andreas Klöckner³, David A Ham⁴ and Paul HJ Kelly¹

The International Journal of High Performance Computing Applications 2020, Vol. 34(6) 629–644
© The Author(s) 2020



Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/1094342020945005

journals.sagepub.com/home/hpc



Abstract

Vectorization is increasingly important to achieve high performance on modern hardware with SIMD instructions. Assembly of matrices and vectors in the finite element method, which is characterized by iterating a local assembly kernel over unstructured meshes, poses difficulties to effective vectorization. Maintaining a user-friendly high-level interface with a suitable degree of abstraction while generating efficient, vectorized code for the finite element method is a challenge for numerical software systems and libraries. In this work, we study cross-element vectorization in the finite element framework Firedrake via code transformation and demonstrate the efficacy of such an approach by evaluating a wide range of matrix-free operators spanning different polynomial degrees and discretizations on two recent CPUs using three mainstream compilers. Our experiments show that our approaches for cross-element vectorization achieve 30% of theoretical peak performance for many examples of practical significance, and exceed 50% for cases with high arithmetic intensities, with consistent speed-up over (intra-element) vectorization restricted to the local assembly kernels.

Keywords

Finite element method, vectorization, global assembly, code generation

1 Introduction

The realization of efficient solution procedures for partial differential equations (PDEs) using finite element methods on modern computer systems requires the combination of diverse skills across mathematics, programming languages and high-performance computing. Automated code generation is one of the promising approaches to manage this complexity. It has been increasingly adopted in software systems and libraries. Recent successful examples include FEniCS (Logg et al., 2012), Firedrake (Rathgeber et al., 2016) and FreeFem++ (Hecht, 2012). These software packages provide users with high-level interfaces for high productivity while relying on optimizations and transformations in the code generation pipeline to generate efficient low-level code. The challenge, as in all compilers, is to use appropriate abstraction layers that enable optimizations to be applied that achieve high performance on a broad set of programs and machines.

One particular challenge for generating high-performance code on modern hardware is vectorization. Modern CPUs increasingly rely on SIMD instructions to achieve higher throughput and better energy efficiency. Finite element computation requires the assembly of vectors and matrices which represent differential forms on discretized function spaces. This process consists of

applying a local function, often called an *element kernel*, to each mesh entity, and incrementing the global data structure with the local contribution. Typical local assembly kernels suffer from issues that can preclude effective vectorization. These issues include complicated loop structures, poor data access patterns, and short loop trip counts that are not multiples of the vector width. As we show in this paper, general purpose compilers perform poorly in generating efficient, vectorized code for such kernels. Padding and data layout transformations are required to enable the vectorization of the element kernels (Luporini et al., 2015), but the effectiveness of such approaches is not consistent across different examples. Since padding may also result in larger overheads for wider vector architectures,

¹ Department of Computer Science, Imperial College, London, UK

² Department of Computer Science, Durham University, Durham, UK

³ Department of Computer Science, University of Illinois at Urbana-Champaign, IL, USA

⁴ Department of Mathematics, Imperial College, London, UK

Corresponding author:

Tianjiao Sun, Department of Computing, Imperial College London, London SW7 2AZ, UK.

Email: tianjiao.sun14@imperial.ac.uk

new strategies are needed as vector width increases for the new generation of hardware.

Matrix-free methods avoid building large sparse matrices in applications of the finite element method and thus trade computation for storage. They have become popular for use on modern hardware due to their higher arithmetic intensity (defined as the number of floating-point operations per byte of data transfer). Vectorization is particularly important for computationally intensive high order methods, for which matrix-free methods are often applied. Previous work on improving vectorization of matrix-free operator application, or equivalently, residual evaluation, mostly focuses on exposing library interfaces to the users. Kronbichler and Kormann (2017) first perform a change of basis from nodal points to quadrature points, and provide overloaded SIMD types for users to write a quadrature-point-wise expression for residual evaluation. However, since the transformation is done manually, new operators require manual reimplementations. Knepley and Terrel (2013) also transpose to quadrature-point basis but target GPUs instead. Both works vectorize by grouping elements into batches, either to match the SIMD vector length in CPUs or the shared memory capacity on GPUs. In contrast, Müthing et al. (2017) apply an intra-kernel vectorization strategy and exploit the fact that, in 3D, evaluating both a scalar field and its three derivatives fills the four lanes of an AVX2 vector register (assuming the computation is in double precision). More recently, Kempf et al. (2018) target high order Discontinuous Galerkin (DG) methods on hexahedral meshes using automated code generation to search for vectorization strategies, while taking advantage of the specific memory layout of the data.

In this work, we present a generic and portable solution based on cross-element vectorization. Our vectorization strategy, implemented in Firedrake, is similar to that of Kronbichler and Kormann (2017) but is fully automated through code generation like that of Kempf et al. (2018). We extend the scope of code generation in Firedrake to incorporate the outer iteration over mesh entities and leverage Loopy (Klöckner, 2014), a loop code generator based loosely on the polyhedral model, to systematically apply a sequence of transformations which promote vectorization by grouping mesh entities into batches so that each SIMD lane operates on one entity independently. This automated code generation mechanism enables us to explore the effectiveness of our techniques on operators spanning a wide range of complexity and systematically evaluate our methodology. Compared with an intra-kernel vectorization strategy, this approach is conceptually well-defined, more portable, and produces more predictable performance. Our experimental evaluation demonstrates that the approach consistently achieves a high fraction of hardware peak performance while being fully transparent to end users.

The contributions of this work are as follows:

- We present the design of a code transformation pipeline that permits the generation of high-

performance, vectorized code on a broad class of finite element models.

- We demonstrate the implementability of the proposed pipeline by realizing it in the Firedrake finite element framework.
- We provide a thorough evaluation of our code generation strategy and demonstrate that it achieves a substantial fraction of theoretical peak performance across a broad range of test cases and popular C compilers.

The rest of this paper is arranged as follows. After reviewing the preliminaries of code generation for the finite element method in Section 2, we describe our implementation of cross-element vectorization in Firedrake in Section 3. In Section 4, we demonstrate the effectiveness of our approach with experimental results. Finally, we review our contributions and identify future research priorities in Section 5.

2 Preliminaries

The computation of multilinear forms using the basis functions spanning the discretized function spaces is called *finite element assembly*. When applying the matrix-free methods, one only needs to assemble linear forms, or residual forms, because matrix-vector products are essentially the assembly of linear forms which represent the *actions* of bilinear forms. Optimizing linear form assembly is therefore crucial for improving the performance of matrix-free methods. In Firedrake, one can invoke the matrix-free approach without changing the high-level problem formulation by setting solver options as detailed by Kirby and Mitchell (2018).

The general structure of a linear form L is

$$L(c_1, c_2, \dots, c_k; v) : \hat{V}_1 \times \hat{V}_2 \times \dots \times \hat{V}_k \times V \rightarrow \mathbb{R}, \quad (1)$$

where $c_i \in \hat{V}_i, i = 1 \dots k$, are arbitrary coefficient functions, and $v \in V$ is the test function. L is linear with respect to v , but possibly nonlinear with respect to the coefficient functions.

Let $\{\phi_i\}_{i=1}^N$ be the set of basis functions spanning V . Define $v_i = L(c_1, \dots, c_k; \phi_i) \in \mathbb{R}$, then the assembly of L constitutes the computation of the vector $\mathbf{v} = (v_1, \dots, v_n)$. In Firedrake, this is treated as a two-step process: *local* assembly and *global* assembly. The rest of this section highlights the computational properties of these two steps with an example, and in doing so, introduces the components and concepts in Firedrake that are relevant to the implementation of cross-element vectorization.

2.1 Local assembly

Local assembly of linear forms is the evaluation of the integrals as defined by the weak form of the differential equation on each entity (cell or facet) of the mesh. In Firedrake, the users define the problem in *Unified Form Language* (UFL) (Alnæs et al., 2014) which captures the weak

```

mesh = UnitSquareMesh(10, 10)
V = FunctionSpace(mesh, "Lagrange", 2)
v = TestFunction(V)
u = Function(V)
L = (dot(grad(u), grad(v)) + u*v) * dx
result = assemble(L)

```

Listing 1. Assembling the linear form of the Helmholtz operator in UFL.

form and the function space discretization. Then the *Two-Stage Form Compiler* (TSFC) (Homolya et al., 2018) takes this high-level, mathematical description and generates efficient C code. The intermediate representation of TSFC is a tensor algebra language called *GEM*, which supports various optimizations on the tensor operations. As an example, consider the linear form of the weak form of the positive-definite Helmholtz operator:

$$L(u; v) = \int_{\Omega} \nabla u \cdot \nabla v + uv \, dx, \quad (2)$$

Listing 1 shows the UFL syntax to assemble the linear form L as the vector `result`, on a 10×10 triangulation of a unit square. We choose to use the second-order Lagrange element, commonly known as the P_2 element, as our approximation space. Listing 2 shows a C representation

of this kernel generated by TSFC.¹ We note the following key features of this element kernel:

- The kernel takes three array arguments in this case: `coords` holds the coordinates of the current triangle, `w_0` holds u_i , the coefficients of u , and `A` stores the result.
- The first part of the kernel (lines 8–20) computes the inverse and the determinant of the Jacobian for the coordinate transformation from the reference element to the current element. This is required for *pulling back* the differential forms to the reference element. The Jacobian is constant for each triangle because the coordinate transformation is affine in this case. In the general case, the Jacobian is computed at every quadrature point.
- The constant arrays `t13`, `t14`, `t15`, `t16` are the same for all elements. `t14` represents the tabulation of basis functions at quadrature points, `t15` and `t16` represent derivatives of basis functions at quadrature points across each spatial dimension, `t13` represents the quadrature weights.
- The `ip` loop iterates over the quadrature points, evaluating the integrand in (2) and summing to approximate the integral. The `i` and `j` loops iterate

```

1  static inline void helmholtz(double *__restrict__ A, double const *__restrict__ coords,
2                               double const *__restrict__ w_0)
3  {
4      double const t13[6] = { ... };
5      double const t14[6 * 6] = { ... };
6      double const t15[6 * 6] = { ... };
7      double const t16[6 * 6] = { ... };
8      double t0 = -1.0 * coords[0];
9      double t1 = t0 + coords[2];
10     double t2 = -1.0 * coords[1];
11     double t3 = t2 + coords[5];
12     double t4 = t0 + coords[4];
13     double t5 = t2 + coords[3];
14     double t6 = t1 * t3 + -1.0 * t4 * t5;
15     double t7 = 1.0 / t6;
16     double t8 = t1 * t7;
17     double t9 = t7 * -1.0 * t4;
18     double t10 = t7 * -1.0 * t5;
19     double t11 = t3 * t7;
20     double t12 = fabs(t6);
21     for (int ip = 0; ip <= 5; ++ip) {
22         double t17 = 0.0;
23         double t18 = 0.0;
24         double t19 = 0.0;
25         for (int i = 0; i <= 5; ++i) {
26             t17 = t17 + t16[6 * ip + i] * w_0[i];
27             t18 = t18 + t15[6 * ip + i] * w_0[i];
28             t19 = t19 + t14[6 * ip + i] * w_0[i];
29         }
30         double t20 = t13[ip] * t12;
31         double t21 = t11 * t19 + t10 * t18;
32         double t22 = t9 * t19 + t8 * t18;
33         double t23 = t20 * (t21 * t10 + t22 * t8);
34         double t24 = t20 * (t21 * t11 + t22 * t9);
35         double t25 = t20 * t17;
36         for (int j = 0; j <= 5; ++j)
37             A[j] = A[j] + t15[6 * ip + j] * t23 + t16[6 * ip + j] * t25 + t14[6 * ip + j] * t24;
38     }
39 }

```

Listing 2. Local assembly kernel for the Helmholtz operator of Listing 1 in C generated by TSFC.

over the degrees of freedom performing a change of basis to values at quadrature points and then back to degrees of freedom when accumulating into the output array `A`. The extents of these loops depend on the integrals performed and the choice of function spaces respectively.

- TSFC performs a sequence of optimization passes by rewriting the tensor operations following mathematical rules before generating the loop nests (see Homolya et al. (2017) for details). This is more powerful than the C compiler's *loop-invariant code motion* optimization because, firstly, TSFC operates on the symbolic tensor expressions to explore different refactoring and reordering strategies that expose invariant sub-expressions, and secondly, non-scalar sub-expressions can also be extracted into temporary arrays to eliminate redundant computation in the loop nests. As a side effect of these transformations, the loop nests in the kernels are no longer perfectly nested, thus limiting the effectiveness of vectorization if it is only applied to the innermost loops.
- After the optimization and scheduling stage, the translation of tensor algebra to C in TSFC is a straightforward rewrite of tensor operations to loop nests. This process results in certain artifacts that are shown in Listing 2. For example, since there is no specific representation for subtractions in TSFC, negations are emitted as multiplication by `-1`. This can result in generated C code that is not as idiomatic as if written by hand. We rely on the modern C compilers to optimize such artifacts away, since readability is a secondary concern for the generated code.

2.2 Global assembly. During global assembly, the local contribution from each mesh entity, computed by the element kernel, is accumulated into the global data structure. In Firedrake, PyOP2 (Rathgeber et al., 2012) is responsible for representing and realizing the iteration over mesh entities, marshalling data in and out of the element kernels. The computation is organized as PyOP2 parallel loops, or *parloops*. A parloop specifies a computational kernel, a set of mesh entities to which the kernel is applied, and all data required for the kernel. The data objects can be directly defined on the mesh entities, or indirectly accessed through maps from the mesh entities. For instance, the signature for the global assembly of the Helmholtz operator is:

```
parloop(helmholtz, cells, L(cell2dof, INC),
        coords(cell2vert, R), u(cell2dof, R)).
```

Here `helmholtz` is the element kernel as shown in Listing 2, generated by TSFC; `cells` is the set of all triangles in the mesh; `L`, `coords`, and `u` are the global data objects that are needed to create the arguments for the element kernel, where `L` holds the result vector, `coords` holds the coordinates of the vertices of the triangles which

are needed for computing the Jacobian, and `u` holds the vector representation of function `u` (as weights of basis functions). These global data objects correspond to the kernel arguments `A`, `coords` and `w_0` respectively. The maps `cell2dof` and `cell2vert` provide indirection from mesh entities to the global data objects, and each data argument is annotated with an access descriptor (`R` for read-only, `INC` for increment access). In this example, the `L` and `u` arguments share the same map (since they are both defined on the same quadratic Lagrange space) while the `coords`, being linear, use a different map.

Listing 3 shows the C code generated by PyOP2 for the above example. The code is then JIT-compiled when the result is needed in Firedrake. In the context of vectorization, this approach, with the inlined element kernel, forms the baseline in our experimental evaluation. We note the following key features of the global assembly kernel:

- The outer loop is over mesh entities.
- For each entity, the computation can be divided into three parts: gathering the input data from global into local data (`t3` and `t4` in this case, which correspond to kernel arguments `coords` and `w_0`), calling the local assembly kernel, scattering the output data (`t2`) to the global data structure.
- The gathering and scattering of data make use of indirect addressing via base pointers (`dots`) and indices (`maps`).
- Different mesh entities might share the same degrees of freedom: parallelization of the scattering loop on line 27 must be aware of the potential for data races.
- Global assembly interacts with local assembly via a function call (line 25). While the C compiler can inline this call, it creates an artificial boundary to using loop optimization techniques that operate at the source code level. Additionally, even after inlining, outer loop vectorization over mesh entities requires that the C compiler vectorize through data-dependent array accesses. This is the software engineering challenge that has previously limited vectorization to a single local assembly kernel in Firedrake.

3 Vectorization

As one would expect, the loop nests and loop trip counts vary considerably for different integrals, meshes and function spaces that users might choose. This complexity is one of the challenges that our system specifically, and Firedrake more generally, must face in order to deliver predictable performance on modern CPUs, which have increasingly rich SIMD instruction sets.

In the prior approach to vectorization in our framework, the local assembly kernels generated by TSFC were further transformed to facilitate vectorization, as described in Luporini et al. (2015). The arrays are padded so that the trip counts of the innermost loops match multiples of the length of SIMD units. However, padding becomes less

```

1  static inline void helmholtz(double *__restrict__ A, double const *__restrict__ coords,
2                             double const *__restrict__ w_0)
3  {
4      // ... element kernel as defined previously ... //
5  }
6
7  void wrap_helmholtz(int const start, int const end, double *__restrict__ dat0, double const *__restrict__ dat1,
8                    double const *__restrict__ dat2, int const *__restrict__ map0, int const *__restrict__ map1)
9  {
10     double t2[6];
11     double t3[3 * 2];
12     double t4[6];
13
14     for (int n = start; n <= -1 + end; ++n) {
15         for (int i6 = 0; i6 <= 5; ++i6)
16             t4[i6] = dat2[map0[6 * n + i6]];
17
18         for (int i2 = 0; i2 <= 2; ++i2)
19             for (int i3 = 0; i3 <= 1; ++i3)
20                 t3[2 * i2 + i3] = dat1[2 * map1[3 * n + i2] + i3];
21
22         for (int i1 = 0; i1 <= 5; ++i1)
23             t2[i1] = 0.0;
24
25         helmholtz(t2, t3, t4);
26
27         for (int i15 = 0; i15 <= 5; ++i15)
28             dat0[map0[6 * n + i15]] += t2[i15];
29     }
30 }

```

Listing 3. Global assembly code for action of the Helmholtz operator in C generated by PyOP2.

effective for low polynomial degrees on wide SIMD units. For instance, AVX512 instructions act on 8 double-precision floats, but the loops for degree 1 polynomials on triangles only have trip counts of 3. Moreover, loop-invariant code motion is very effective in reducing the number of floating-point operations, but hoisted instructions are not easily vectorized as they are no longer in the innermost loops. This effect is more pronounced on tensor-product elements where TSFC is able to apply *sum factorization* (Homolya et al., 2017) to achieve better algorithmic complexity.

3.1 Cross-element vectorization and Loopy

Another strategy is to vectorize across several elements in the outer loop over the mesh entities, as proposed previously by Kronbichler and Kormann (2017). This approach computes the contributions from several mesh entities using SIMD instructions, where each SIMD lane handles one entity. This is always possible regardless of the complexity of the local element kernel because the computation on each entity is independent and identical. One potential downside is the increase in register and cache pressure as the working set is larger.

For a compiler, the difficulty in performing cross-element vectorization (or, more generally, outer-loop vectorization) is to automate a sequence of loop transformations and necessary data layout transformations robustly. This is further complicated by the indirect memory access in data gathering and scattering, and the need to unroll and interchange loops across the indirections, which requires significantly more semantic knowledge than that available to the C compiler.

Loopy (Klöckner, 2014) is a loop generator embedded in Python which targets both CPUs and GPUs. Loopy provides abstractions based on integer sets for loop-based computations and enables powerful transformations based on the polyhedral model (Verdoolaege, 2010). Loop-based computations in Loopy are represented as *Loopy kernels*. A Loopy kernel is a subprogram consisting of a loop domain and a partially-ordered list of scalar assignments acting on multi-dimensional arrays. The loop domain is specified as the set of integral points in the convex intersection of quasi-affine constraints, represented using the Integer Set Library (Verdoolaege, 2010). Loopy supports code generation for different environments from the same kernel by choosing different *targets*.

To integrate with Loopy, the code generation mechanisms in Firedrake were modified as illustrated in Figure 1. Instead of generating source code directly, TSFC and PyOP2 are modified to generate Loopy kernels. We have augmented the Loopy internal representation with the ability to support a generalized notion of kernel fusion through the nested composition of kernels, specifically through subprograms and inlining. This allows PyOP2 to inline the element kernel such that the global assembly Loopy kernel encapsulates the complete computation of global assembly. This holistic view of the overall computation enables robust loop transformations for vectorization across the boundary between global and local assembly. To facilitate SIMD instruction generation, we also introduced a new OpenMP target to Loopy which extends its existing C-language target to support OpenMP SIMD directives (OpenMP Architecture Review Board, 2018, §2.9.3).

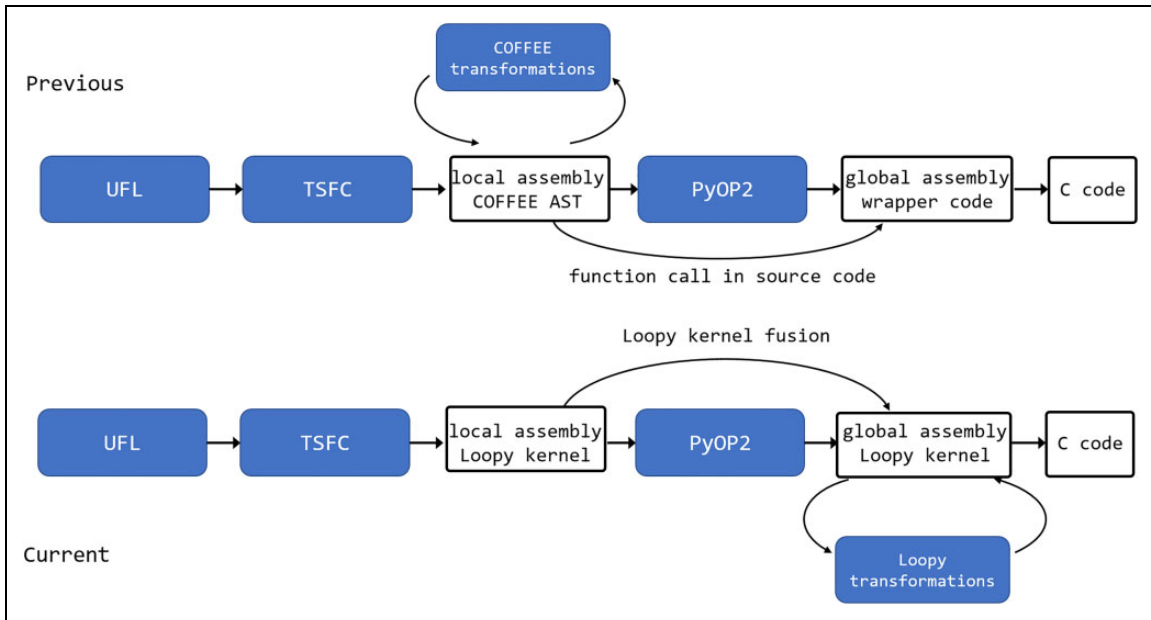


Figure 1. Integration of Loopy in Firedrake for global assembly code generation.

```

KERNEL: helmholtz
-----
ARGUMENTS:
start: type: int32
end: type: int32
dat0: type: float64, shape: (None)
// ... More arguments ... //
-----
DOMAINS:
[end, start] -> { [n] : start <= n < end }
{ [i6] : 0 <= i6 <= 5 }
// ... More domains ... //
-----
INAME_IMPLEMENTATION_TAGS:
None
-----
TEMPORARIES:
t4: type: float64, shape: (6), dim_tags: (stride:1)
// ... More temporaries ... //
-----
INSTRUCTIONS:
for n
  for i6
    t4[i6] = dat2[map0[n, i6]]
    // ... More instructions ... //
  for i15
    dat0[map0[n, i15]] += t0[0, i15]
  end n
end n

```

Listing 4. Global assembly Loopy kernel of the Helmholtz operator.

Listing 4 shows an abridged version of the global assembly Loopy kernel for the Helmholtz operator, with the element kernel fused. We highlight the following key features of Loopy kernels:

- Loop indices, such as n and $i1$, are called *inames* in Loopy, which define the iteration space. The bounds of the loops are specified by the affine constraints in *domains*.
- Loop transformations operate on kernels by rewriting the loop domain and the statements making up the kernel. In addition, each iname carries a set of

tags governing its realization in generated code, perhaps as a sequential loop, as a vector lane index, or through unrolling.

- Multi-dimensional arrays occur as *arguments* and *temporaries*. The memory layout of the data can be specified by assigning *tags* to the array dimensions.
- Dependencies between statements specify their partial order. Statement scheduling can also be controlled by assigning priorities to statements and inames.

For example, to achieve cross-element vectorization (by batching four elements into one SIMD vector in this example) we invoke the following sequence of Loopy transformations on the global assembly Loopy kernel, exploiting the domain knowledge of finite element assembly:

1. Split the outer loop n over mesh entities into n_{outer} and n_{simd} , with n_{simd} having a trip count of four. The objective is to generate SIMD instructions for the n_{simd} loops, such that each vector lane computes one iteration of the n_{simd} loops.
2. Assign the tag `SIMD` to the new iname n_{simd} . This tag informs Loopy to force the n_{simd} loop to be innermost, privatizing data by vector-expansion if necessary.

We show the change to the Loopy kernel after these transformations in Listing 5. In particular, the vector-expansion of the temporary $t4$, and the splitting (and subsequent modification of the loop domain) of the n iname.

Listing 6 shows the generated C code for the Helmholtz operator vectorized by grouping together four elements.

```

KERNEL: helmholtz_simd
-----
ARGUMENTS:
start: type: int32
end: type: int32
dat0: type: float64, shape: (None)
// ... More arguments ... //
-----
DOMAINS:
[end, start] -> { [n_outer, n_simd] :
  n_simd >= start - 4n_outer
  and 0 <= n_simd <= 3
  and n_simd < end - 4n_outer }
// ... More domains ... //
-----
INAME_IMPLEMENTATION_TAGS:
n_simd: SIMD
-----
TEMPORARIES:
t4: type: float64, shape: (6, 4),
    dim_tags: (stride:4, stride:1)
// ... More temporaries ... //
-----
INSTRUCTIONS:
for n_outer, n_simd
  for i6
    t4[i6, n_simd] = dat2[map0[n_outer*4 + n_simd, i6]]
    // ... More instructions ... //
  for i15
    dat0[map0[n_outer*4+n_simd, i15]] += t2[i15, n_simd]
end n_outer, n_simd

```

Listing 5. Changes to global assembly Loopy kernel of the Helmholtz operator after cross-element vectorization.

Apart from the previously mentioned changes, we note the following details:

- The `n_simd` loops are pushed to the innermost level. Moreover, this transformation vector-expands temporary arrays such as `t2`, `t3`, `t4` by four, with the expanded dimension labeled as varying the fastest when viewed from (linear) system memory. This ensures their accesses in the `n_simd` loops always have unit stride.
- Loopy provides a mechanism to declare arrays to be aligned to specified memory boundaries (64 bytes in this example).
- The `n_simd` loops are decorated with `#pragma omp simd` to inform C compilers to generate SIMD instructions. The exception is the writing back to the global array (lines 43–45), which is sequentialized due to potential race conditions, as different mesh entities could share the same degrees of freedom.
- The remainder loop which handles the cases where the number of elements is non-divisible by four is omitted here for simplicity.
- After cross-element vectorization, all local assembly instructions (lines 28–42) are inside `n_simd` loops, which always have trip counts of four and are unit stride. All loop-varying array accesses are unit stride in the fastest moving dimension. There are no loop-carried dependencies in `n_simd` loops. As a result, the `n_simd` loops, and therefore all local assembly instructions, are vectorizable without further consideration of dependencies. This is verified by checking the x86 assembly code and running the program with the Intel Software Development Emulator.

3.2 Vector extensions

A more direct way to inform the compiler to emit SIMD instructions without depending on OpenMP is to use *vector extensions*,² which support vector data types. These were first introduced in the GNU compiler (GCC), but are also supported in recent versions of the Intel C compiler (ICC) and Clang. Analogous mechanisms exist in various vector-type libraries, e.g. VCL (Fog 2017). To evaluate and compare with the directive-based approach from Section 3.1, we created a further code generation target in Loopy to support vector data types. When inames and corresponding array axes are jointly tagged as vector loops, Loopy generates code to compute on data in vector registers directly, instead of scalar loops over the vector lanes. It is worth noting that the initial intermediate representation of the loop is identical in each case, and that the different specializations are achieved through code transformation. Listing 7 shows the C code generated for the Helmholtz operator vectorized by batching four elements using the vector extension target. Here almost all vectorized (innermost) loops for local assembly are replaced by operations on vector variables. For instructions which do not fit the vector computation model, most noticeably the indirect data gathering (lines 18–19 24–25), or instructions containing built-in mathematics functions which are not supported on vector data types (line 37), Loopy generates scalar loops over vector lanes decorated with `#pragma omp simd`. In addition, because vector extensions do not automatically broadcast scalars, any vector instruction with a scalar rvalue is modified by adding the zero vector to the expression, as shown in lines 29 and 31.

Compared to Listing 6, using vector extensions removes most of the innermost loops, and the only remaining OpenMP SIMD directives are due to the limitations of vector extensions as explained previously.

4 Performance evaluation

We follow the performance evaluation methodology of Luporini et al. (2017) by measuring the assembly time of a range of operators of increasing complexity and polynomial degrees. Due to the large number of combinations of experimental parameters (operators, meshes, polynomial degrees, vectorization strategies, compilers, hyperthreading), we only report an illustrative portion of the results here, with the entire suite of experiments made available on the interactive online repository CodeOcean (Sun, 2019a).

4.1 Experimental setup

We performed experiments on a single node of two Intel systems, based on the Haswell and Skylake microarchitectures, as detailed in Table 1. Firedrake uses MPI for parallel execution where each MPI process handles the assembly for a subset of the domain. Hybrid MPI-OpenMP parallelization

```

1 // ... Constant array declarations ... //
2
3 void wrap_helmholtz(int const start, int const end, double *__restrict__ dat0, double const *__restrict__ dat1,
4                   double const *__restrict__ dat2, int const *__restrict__ map0, int const *__restrict__ map1)
5 {
6     double form_t1[4] __attribute__((aligned(64)));
7     double t2[6 * 4] __attribute__((aligned(64)));
8     // ... More temporary array declarations ... //
9     for (int n_outer = (start / 4); n_outer <= ((-4 + end) / 4); ++n_outer) {
10        for (int i2 = 0; i2 <= 2; ++i2) {
11            for (int i3 = 0; i3 <= 1; ++i3) {
12                #pragma omp simd
13                for (int n_simd = 0; n_simd <= 3; ++n_simd)
14                    t3[n_simd + 8 * i2 + 4 * i3] = dat1[2 * map1[12 * n_outer + 3 * n_simd + i2] + i3];
15            }
16        }
17        for (int i6 = 0; i6 <= 5; ++i6) {
18            #pragma omp simd
19            for (int n_simd = 0; n_simd <= 3; ++n_simd)
20                t4[n_simd + 4 * i6] = dat2[map0[24 * n_outer + 6 * n_simd + i6]];
21        }
22        for (int i1 = 0; i1 <= 5; ++i1) {
23            #pragma omp simd
24            for (int n_simd = 0; n_simd <= 3; ++n_simd)
25                t2[n_simd + 4 * i1] = 0.0;
26        }
27        #pragma omp simd
28        for (int n_simd = 0; n_simd <= 3; ++n_simd) {
29            form_t11[n_simd] = 0.0;
30            form_t1[n_simd] = -1.0 * t3[n_simd];
31            // ... More similar instructions ... //
32            form_t8[n_simd] = fabs(form_t7[n_simd]);
33        }
34        for (int form_ip = 0; form_ip <= 5; ++form_ip) {
35            // ... More similar loop nests ... //
36            for (int form_j = 0; form_j <= 5; ++form_j) {
37                #pragma omp simd
38                for (int n_simd = 0; n_simd <= 3; ++n_simd)
39                    t2[n_simd + 4 * form_j] += (form_t25[form_j] * form_t24[n_simd] + form_t13[n_simd + 4 * form_j] +
40                                               form_t27[form_j] * form_t26[n_simd]);
41            }
42        }
43        for (int i15 = 0; i15 <= 5; ++i15)
44            for (int n_simd = 0; n_simd <= 3; ++n_simd)
45                dat0[map0[24 * n_outer + 6 * n_simd + i15]] += t2[n_simd + 4 * i15];
46    }
47 }

```

Listing 6. Global assembly code for action of the Helmholtz operator in C vectorized by batching four elements.

is not supported, and we stress that OpenMP pragmas are only used for SIMD vectorization within a single MPI process. Because we observe that hyperthreading usually improves the performance by 5% to 10% for our applications, we set the number of MPI processes to the number of logical cores of the CPU to utilize all available computation resources. Experimental results with hyperthreading turned off are available on CodeOcean. Turbo Boost is switched off to mitigate reproducibility problems that might be caused by dynamic thermal throttling. The batch size, i.e. the number of elements grouped together for vectorization, is chosen to be consistent with the SIMD length. We use three C compilers: GCC 7.3, ICC 18.0 and Clang 5.0. The two vectorization strategies described in Section 3 are tested on all platforms. We use the listed *Base Frequency* to calculate the peak performance in Table 1. In reality, modern Intel CPUs dynamically reduce frequencies on heavy workloads with AVX2 and AVX512 instructions, which results in lower achievable performance. Running the optimized LINPACK benchmark binary provided by

Intel gives a reasonable indication of achievable peak performance for compute-bound applications.

For the benefit of reproducibility, we have archived the specific versions of Firedrake components used for the experimental evaluation on Zenodo (Sun, 2019b; Zenodo/Firedrake, 2019). An installation of Firedrake with components matching the ones used for evaluation in this paper can be obtained following the instruction at <https://www.firedrakeproject.org/download.html>, with the following command:

```
python3 firedrake-install --doi 10.5281/zenodo.2595487
```

We measure the execution time of assembling the residual for five operators: the mass matrix (“mass”), the Helmholtz equation (“helmholtz”), the vector Laplacian (“laplacian”), an elastic model (“elasticity”), and a hyperelastic model (“hyperelasticity”). The mathematical description of the operators is detailed in the supplemental material. These operators stem from real-world


```

1  typedef double double4 __attribute__ ((vector_size (32)));
2  typedef int int4 __attribute__ ((vector_size (16)));
3
4  static double4 const _zeros_double4 __attribute__ ((aligned (64))) = { 0.0 };
5
6  // ... Constant array declarations ... //
7
8  void wrap_helmholtz(int const start, int const end, double *__restrict__ dat0, double const *__restrict__ dat1,
9                    double const *__restrict__ dat2, int const *__restrict__ map0, int const *__restrict__ map1)
10 {
11     double4 form_t1 __attribute__ ((aligned (64)));
12     // ... Temporary array declarations ... //
13
14     for (int n_outer = (start / 4); n_outer <= ((-4 + end) / 4); ++n_outer) {
15         for (int i2 = 0; i2 <= 2; ++i2) {
16             for (int i3 = 0; i3 <= 1; ++i3) {
17                 #pragma omp simd
18                 for (int n_simd = 0; n_simd <= 3; ++n_simd)
19                     t3[n_simd + 8 * i2 + 4 * i3] = dat1[2 * map1[12 * n_outer + 3 * n_simd + i2] + i3];
20             }
21         }
22         for (int i6 = 0; i6 <= 5; ++i6) {
23             #pragma omp simd
24             for (int n_simd = 0; n_simd <= 3; ++n_simd)
25                 t4[n_simd + 4 * i6] = dat2[map0[24 * n_outer + 6 * n_simd + i6]];
26         }
27
28         for (int i1 = 0; i1 <= 5; ++i1)
29             t2[i1] = _zeros_double4;
30
31         form_t1l1 = _zeros_double4;
32         form_t1 = -1.0 * t3[0];
33
34         for (int form_ip = 0; form_ip <= 5; ++form_ip) {
35             // ... More similar instructions ... //
36             #pragma omp simd
37             for (int n_simd = 0; n_simd <= 3; ++n_simd)
38                 form_t8[n_simd] = fabs(form_t7[n_simd]);
39             // ... More similar instructions ... //
40             for (int form_j = 0; form_j <= 5; ++form_j)
41                 t2[form_j] += form_t25[form_j] * form_t24 + form_t13[form_j] + form_t27[form_j] * form_t26;
42         }
43
44         for (int i15 = 0; i15 <= 5; ++i15)
45             for (int n_simd = 0; n_simd <= 3; ++n_simd)
46                 dat0[map0[24 * n_outer + 6 * n_simd + i15]] += t2[i15][n_simd];
47     }
48 }

```

Listing 7. Global assembly code for action of the Helmholtz operator in C vectorized by four elements (using vector extensions).

Table I. Hardware specification for experiments.

	Haswell Xeon E5-2640 v3	Skylake Xeon Gold 6130
Base frequency	2.6 GHz	2.1 GHz
Physical cores	8	16
SIMD instruction set	AVX2	AVX512
doubles per SIMD vector	4	8
Cross-element vectorization batch size	4	8
FMA ³ units per core	2	2
FMA instruction issue per cycle	2	2
Peak performance (double-precision) ⁴	332.8 GFLOP/s	1075.2 GFLOP/s
System memory	4 × 8 GB DDR4-2133	2 × 32 GB DDR4-2666
LINPACK performance (double-precision) ⁵	262.5 GFLOP/s	678.8 GFLOP/s
Memory bandwidth ⁶	38.5 GB/s	36.6 GB/s
GCC/Clang arch flag	-march=native	-march=native
ICC SIMD flag	-xcore-avx2	-xcore-avx512 -qopt-zmm-usage=high
Other compiler flags	-O3 -ffast-math -fopenmp	-O3 -ffast-math -fopenmp
Intel Turbo Boost	OFF	OFF

applications and cover a wide range of complexity: the generated C code for the corresponding global assembly kernels exceeds hundreds of KB for the hyperelasticity operator at high polynomial degree.

We performed experiments on both 2D and 3D domains, with two types of mesh used for each case: triangles (“`tri`”) and quadrilaterals (“`quad`”) for 2D problems, tetrahedra (“`tet`”) and hexahedra (“`hex`”) for 3D problems. This large variety underscores the broad applicability of our approach. The *arithmetic intensities* and other pertinent characteristics of the operators are listed in Table 2. The memory footprint is calculated assuming perfect caching—it is thus a lower bound which results in an upper bound estimation for the arithmetic intensity. The triangular and tetrahedral meshes use an affine coordinate transformation (requiring only one Jacobian evaluation per element). The quadrilateral and hexahedral meshes use a bilinear (trilinear) coordinate transformation (requiring Jacobian evaluation at every quadrature point), which usually results in higher arithmetic intensities at low orders. In Firedrake, tensor-product elements (McRae et al., 2016) benefit from optimizations such as sum factorization to achieve lower asymptotic algorithmic complexity. They are therefore more competitive for higher order methods (Homolya et al., 2017).

We record the maximum execution time of the generated global assembly kernels on all MPI processes. This time does not include the time in synchronization and MPI data exchange for halo updates. Each experiment is run five times, and the average execution time is reported. Exclusive access to the compute nodes is ensured and threads are pinned to individual logical cores. Startup costs such as code generation time and compilation time are excluded. We use automatic vectorization by GCC without batching, compiled with the optimization flags of Table 1, as the baseline for comparison. Comparing with our cross-element strategy, the baseline represents the out-of-the-box performance of compiler auto-vectorization for the local element kernel. We note that cross-element vectorization does not alter the algorithm of local assembly except for the vector expansion, as illustrated by Listing 2 and Listing 6. Consequently, the total number of floating-point operations remains the same. The performance benefit from cross-element vectorization is therefore composable with the operation-reduction optimizations performed by the form compiler to the local assembly kernels.

4.2 Experimental results and discussion

Figures 2 to 5 show the performance of the `helmholtz` and `elasticity` operators on Haswell and Skylake, vectorized with OpenMP SIMD directives (Section 3.1), and with vector extensions (Section 3.2). We indicate the fraction of peak performance achieved on the left axis, and the fraction of the LINPACK benchmark performance on the right axis. Figures 6 and 7 compare roofline models (Williams et al., 2009) of the baseline and our approach using

vector extensions and the GCC compiler on Haswell and Skylake. The speed-up achieved is also summarized in Table 2.

4.2.1 Compiler comparison and vector extensions. When vectorizing with OpenMP SIMD directives, ICC gives the best performance for almost all test cases, followed by Clang, while GCC is significantly less competitive. The performance disparity is more pronounced on Skylake than on Haswell. However, when using vector extensions, Clang and GCC improve significantly and are able to match the performance of ICC on both Haswell and Skylake, whereas ICC performs similarly with both OpenMP SIMD directives and vector extensions.

We use the Intel Software Development Emulator⁷ to count the number of instructions executed at runtime for code generated by different compilers. The data indicate that although floating-point operations are fully vectorized by all compilers, GCC and Clang generate more load and store instructions between vector registers and memory when using OpenMP SIMD directives for vectorization. One possible reason is that GCC and Clang choose to allocate short arrays to the stack rather than the vector registers directly, causing more load on the memory subsystem.

In light of these results, we conclude that vectorization with vector extensions allows greater performance portability on different compilers and CPUs for our application. It is, therefore, our preferred strategy for implementing cross-element vectorization, and is the default option for the rest of our analysis.

4.2.2 Vectorization speed-up. Almost across the board, significant speed-up is achieved on the test cases under consideration. Slowdown occurs in two situations. On low polynomial degrees, the kernels tend to have low arithmetic intensity so that the increase in available floating point throughput through cross-element vectorization cannot compensate for the increase in the size of the working set of data. On simple operators such as `mass` on `tri` and `tetra`, the kernels have simple loop structures and the compilers can sometimes successfully apply other optimizations such as unrolling and loop interchange to achieve vectorization without batching elements in the outer loop. The pattern of speed-up is consistent across Haswell and Skylake. Higher speed-up is generally achieved on more complicated operators (e.g. `hyperelasticity`), and on tensor-product elements (`quad` and `hex`), which generally correspond to more complicated loop structure and higher arithmetic intensity due to the Jacobian recomputation at each quadrature point.

4.2.3 Achieved fraction of peak performance. We observe that the fraction of peak performance varies smoothly with polynomial degrees for cross-element vectorization in all test cases. This fulfils an important design requirement for Firedrake: small changes in problem setup by the users should not create unexpected performance degradation.

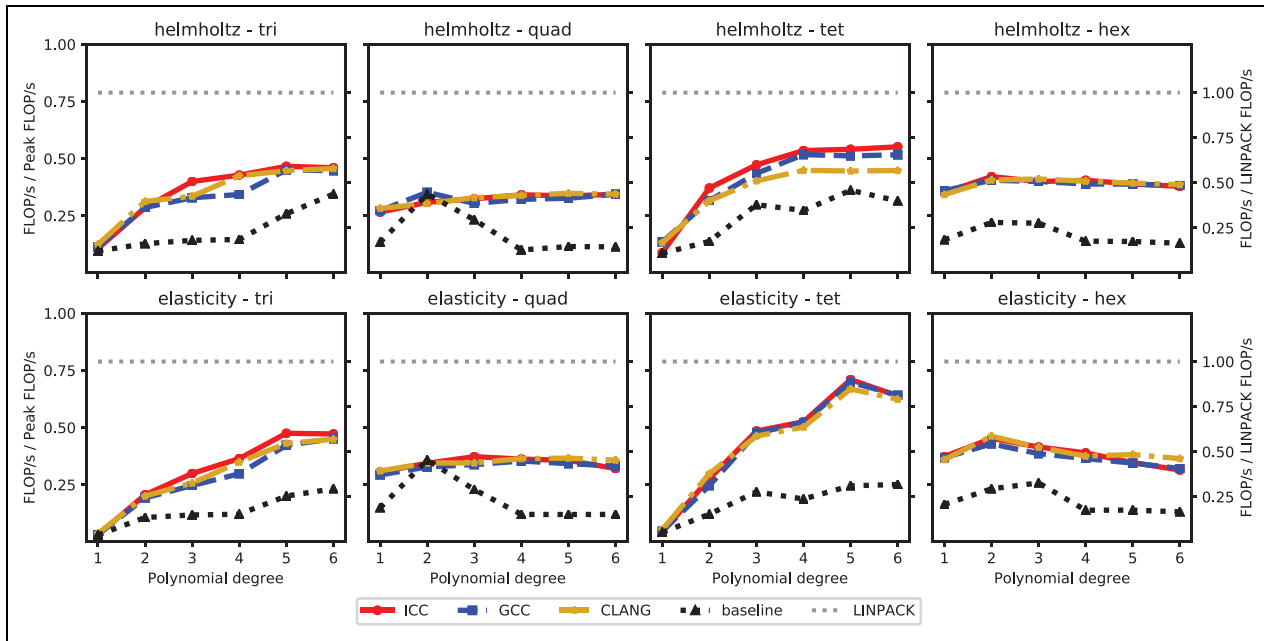


Figure 2. The fraction of peak FLOP/s (as listed in Table I) achieved by different compilers for operators {helmholtz, elasticity}, on meshes {tri, quad, tet, hex} on **Haswell** using **vector extensions** with 16 MPI processes. The dotted line indicates the fraction of peak performance achieved by LINPACK benchmark.

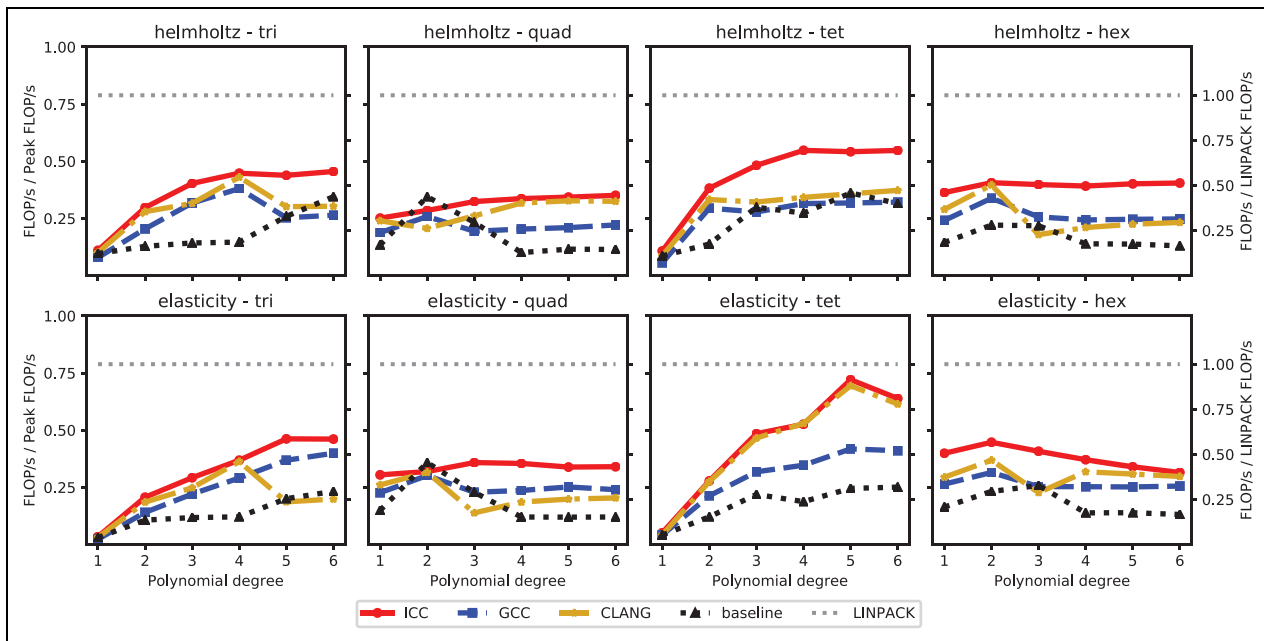


Figure 3. The fraction of peak FLOP/s (as listed in Table I) achieved by different compilers for operators {helmholtz, elasticity}, on meshes {tri, quad, tet, hex} on **Haswell** using **OpenMP SIMD directives** with 16 MPI processes. The dotted line indicates the fraction of peak performance achieved by LINPACK benchmark.

This is also shown in Figures 6 and 7 where the results are more clustered on the roofline plots after cross-element vectorization. The baseline shows performance inconsistency, especially for low polynomial degrees. For instance, for the `helmholtz` operator with degree 3 on `quad`, the quadrature loops and the basis function loops all have trip counts of 4, which fits the vector length on Haswell and results in better performance.

On simplicial meshes (`tri` and `tetra`), higher order discretization leads to kernels with very high arithmetic intensity because of the quadratic and cubic increases in the number of basis functions, and thus the loop trip counts. This is due to the current limitation that simplicial elements in Firedrake are not sum factorized. In these test cases, we observe that the baseline approaches cross-element vectorization for sufficiently high polynomial degrees. This is not

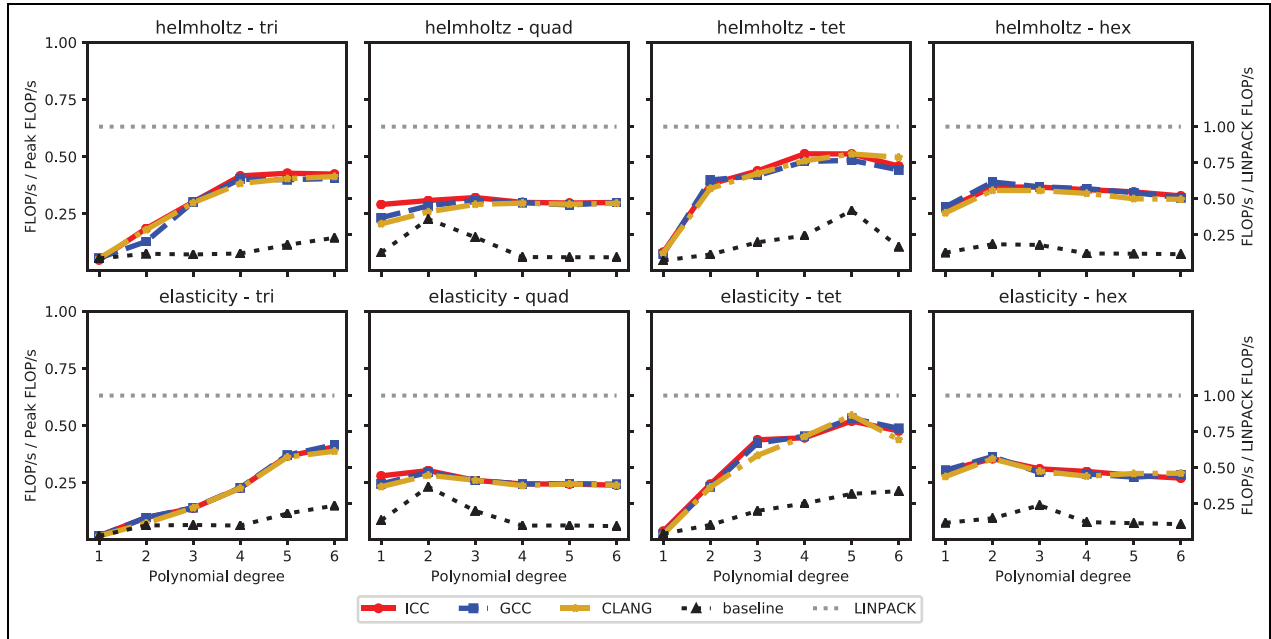


Figure 4. The fraction of peak FLOP/s (as listed in Table I) achieved by different compilers for operators $\{\text{helmholtz}, \text{elasticity}\}$, on meshes $\{\text{tri}, \text{quad}, \text{tet}, \text{hex}\}$ on **Skylake** using **vector extensions** with 32 MPI processes. The dotted line indicates the fraction of peak performance achieved by the LINPACK benchmark.

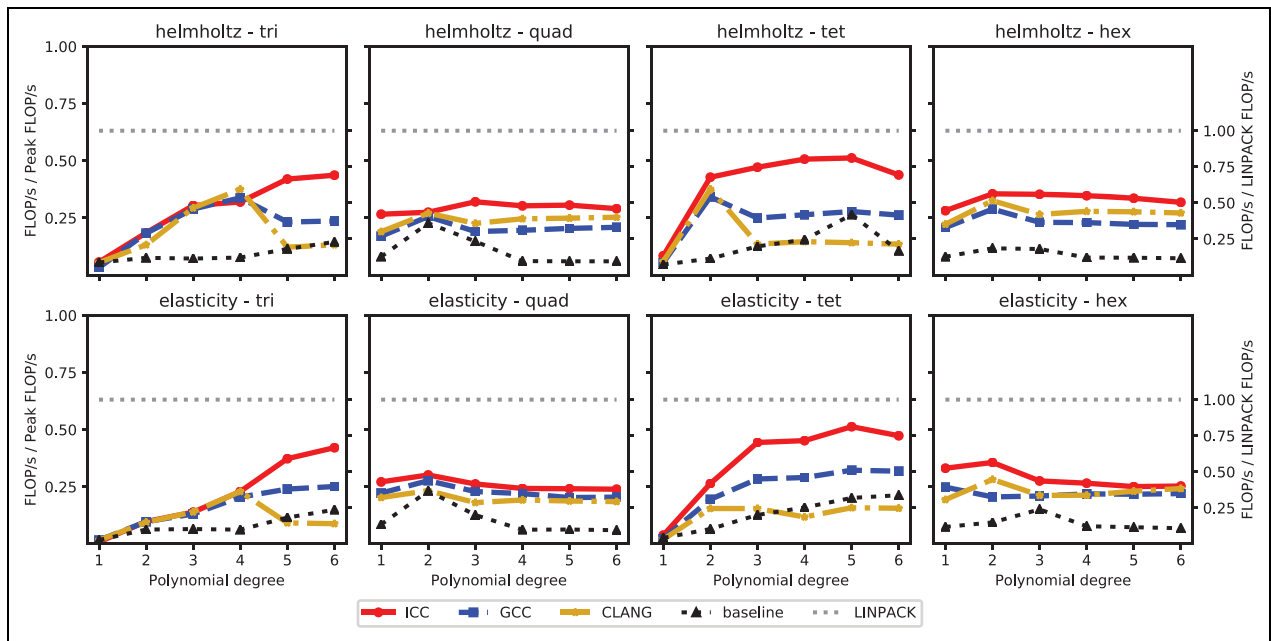


Figure 5. The fraction of peak FLOP/s (as listed in Table I) achieved by different compilers for operators $\{\text{helmholtz}, \text{elasticity}\}$, on meshes $\{\text{tri}, \text{quad}, \text{tet}, \text{hex}\}$ on **Skylake** using **OpenMP SIMD directives** with 32 MPI processes. The dotted line indicates the fraction of peak performance achieved by the LINPACK benchmark.

a serious concern for our optimization approach because the break-even degrees are very high except for simple operators such as `mass`, and ultimately tensor-product elements are more competitive for higher order methods in terms of algorithmic complexity.

We also observe that there exist a small number of test cases where the achieved peak performance is marginally higher than the LINPACK benchmark on Skylake, as

shown in Figure 7. One possible reason for this observation is thermal throttling since our test cases typically run for a shorter period of time than LINPACK. We also note that these test cases correspond to high order hyperelasticity operators on `tet` meshes, which, as noted previously, are not practically important use cases since using tensor-product elements requires many fewer floating-point operations at the same polynomial order.

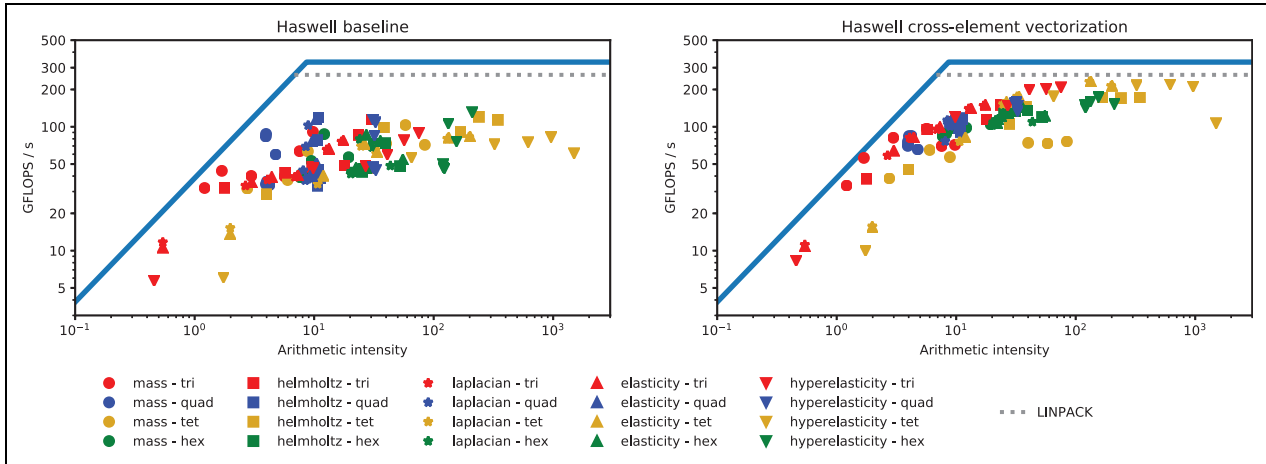


Figure 6. Roofline model of operators for baseline and cross-element vectorization using GCC on **Haswell**. The dotted lines indicate the performance of the LINPACK benchmark.

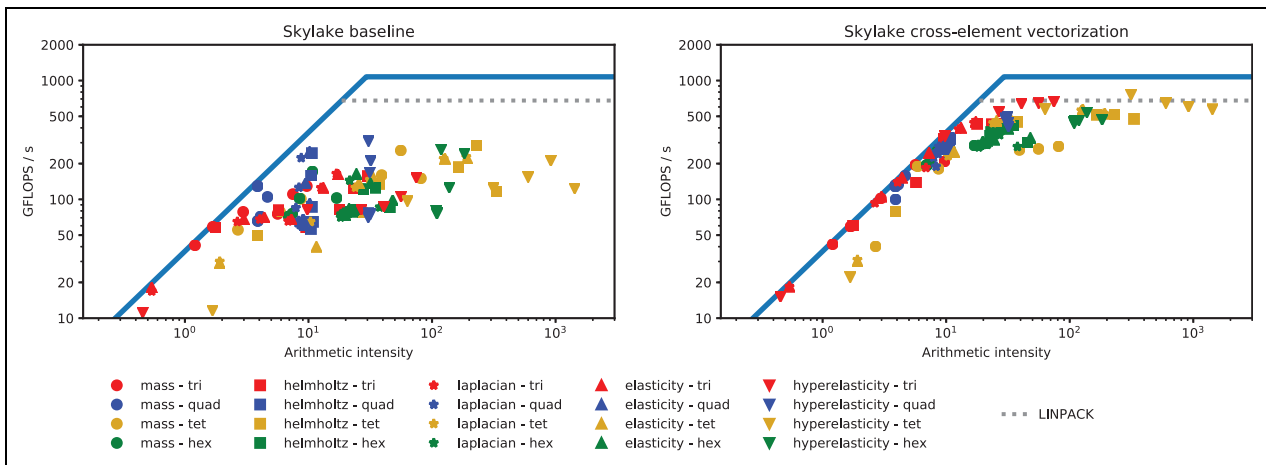


Figure 7. Roofline model of operators for baseline and cross-element vectorization using GCC on **Skylake**. The dotted lines indicate the performance of the LINPACK benchmark.

4.2.4 Tensor-product elements. We observe higher and more consistent speed-up for tensor-product elements (quad and hex) on both Haswell and Skylake. This is because, on these meshes, more computation is moved out of the innermost loop due to sum factorization, resulting in more challenging loop nests for the baseline strategy which attempts to vectorize within the element kernel. The same applies to the evaluation of the Jacobian of coordinate transformation, which is a nested loop over quadrature points after sum factorization for tensor-product elements.

The base elements of quad and hex are interval elements in 1D, thus the extents of loops over degrees of freedom increase only linearly with respect to polynomial degrees, as shown in Table 2. As a result, the baseline performance does not improve as quickly for higher polynomial degrees on quad and hex compared with tri and tet, resulting in stable speed-up for cross-element vectorization observed on tensor-product elements.

5 Conclusion and future work

We have presented a portable, general-purpose solution for delivering stable vectorization performance on modern CPUs for matrix-free finite element assembly for a very broad class of finite element operators on a large range of elements and polynomial degrees. We described the implementation of cross-element vectorization in Firedrake which is transparent to the end users. Although the technique of cross-element vectorization is conceptually simple and has been applied in hand-written kernels before, our implementation based on code generation is automatic, robust and composable with other optimization passes.

We have presented extensive experimental results on two recent Xeon processors that are commonly used in HPC applications, and compared the vectorization performance of three popular C compilers. We showed that by generating appropriate vectorizable code, and using compiler-based vector extensions, we can obtain portably high performance across all three compilers.

The write-back to global data structure is not vectorized in our approach due to possible race conditions. The newly introduced Conflict Detection instructions in the Intel AVX512 instruction set could potentially mitigate this limitation (Zhang, 2016, section 2.3). This could be achieved by informing Loopy to use the relevant intrinsics when generating code for loops with specific tags.

We have focused on the matrix-free finite element method because it is compute-intensive and more likely to benefit from vectorization. However, our methods and implementation also support matrix assembly. Firedrake relies on PETSc (Balay et al., 2017) to handle distributed sparse matrices, and PETSc requires certain data layouts for the input array when updating the global matrices. When several elements are batched together for cross-element vectorization, we need to generate code to explicitly unpack/transpose the local assembly results into individual arrays before calling PETSc functions to update the global sparse matrices for each element. Future improvement could include eliminating this overhead, possibly by extending the PETSc API.

The newly introduced abstraction layer, together with Loopy integration in the code generation and optimization pipeline, opens up multiple possibilities for future research in Firedrake. These include code generation with intrinsics instructions, loop tiling, and GPU acceleration, all of which are already supported in Loopy.

Acknowledgments

The authors would like to thank Tobias Grosser, Richard Veras, J. Ramanujam and P. Sadayappan for their valuable insights during our discussions which started at Dagstuhl Seminar 18111 on Loop Optimization. The authors are grateful to James Cownie and Andrew Mollinson at Intel Corp. as well as Koki Sagiyama at Imperial College London for providing access to the Skylake platform.



Declaration of conflicting interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was supported by the Engineering and Physical Sciences Research Council [grant numbers EP/L016796/1, EP/R029423/1], and the Natural Environment Research Council [grant number NE/K008951/1]. It was further funded by the US Navy Office of Naval Research under grant number N00014-14-1-0117 and the US National Science Foundation under grant number CCF-1524433. AK gratefully acknowledges a hardware gift from Nvidia Corporation.

ORCID iD

Tianjiao Sun  <https://orcid.org/0000-0003-4223-6700>
Paul HJ Kelly  <https://orcid.org/0000-0001-5905-1804>

Supplemental material

Supplemental material for this article is available online

Notes

1. This TSFC-generated kernel is reformatted slightly for consistency with the PyOP2 generated kernels. The kernel function names generated by TSFC and PyOP2 are long and complicated due to name mangling in Firedrake. They are shortened to operator names such as `helmholtz` in the listings for readability.
2. <https://gcc.gnu.org/onlinedocs/gcc/Vector-Extensions.html>
3. Fused multiply-add operations.
4. Calculated as $basefrequency \times \#cores \times SIMDwidth \times 2 (forFMA) \times \#issuepercycle$
5. Intel LINPACK Benchmark. <https://software.intel.com/en-us/articles/intel-mkl-benchmarks-suite>
6. STREAM triad benchmark, 2 threads per core.
7. <https://software.intel.com/en-us/articles/intel-software-development-emulator>

References

- Alnæs MS, Logg A, Ølgaard KB, et al. (2014) Unified form language: a domain-specific language for weak formulations of partial differential equations. *ACM Transactions on Mathematical Software* 40(2): 9:1–9:37.
- Balay S, Abhyankar S, Adams M, et al. (2017) *Petsc users manual revision 3.8. Technical report*. Argonne, IL: National Lab.(ANL).
- Fog A (2017) VCL—A C++ Vector Class Library. Available at: <https://www.agner.org/optimize/vectorclass.pdf> (accessed 19 March 2019).
- Hecht F (2012) New development in FreeFEM++. *Journal of Numerical Mathematics* 20(3–4): 251–266.
- Homolya M, Kirby RC and Ham DA (2017) Exposing and exploiting structure: optimal code generation for high-order finite element methods. ArXiv: 1711.02473 [cs.MS].
- Homolya M, Mitchell L, Luporini F, et al. (2018) TSFC: a structure-preserving form compiler. *SIAM Journal on Scientific Computing* 40(3): C401–C428.
- Kempf D, Heß R, Müthing S, et al. (2018) Automatic code generation for high-performance discontinuous Galerkin methods on modern architectures. *arXiv preprint arXiv:1812.08075*.
- Kirby RC and Mitchell L (2018) Solver composition across the PDE/linear algebra barrier. *SIAM Journal on Scientific Computing* 40(1): C76–C98.
- Klößner A (2014) Loo. py: transformation-based code generation for GPUs and CPUs. In: *Proceedings of ACM SIGPLAN International Workshop on Libraries, Languages, and Compilers for Array Programming*, Edinburgh, UK, 12–13 June 2014, p. 82. New York, NY: ACM.
- Knepley MG and Terrel AR (2013) Finite element integration on GPUs. *ACM Transactions on Mathematical Software (TOMS)* 39(2): 10.
- Kronbichler M and Kormann K (2017) Fast matrix-free evaluation of discontinuous Galerkin finite element operators. *arXiv preprint arXiv:1711.03590*.

- Logg A, Mardal KA and Wells GN (eds.) (2012) *Automated Solution of Differential Equations by the Finite Element Method: the FEniCS Book*. Berlin: Springer. ISBN 978-3-642-23098-1. DOI:10.1007/978-3-642-23099-8.
- Luporini F, Ham DA and Kelly PHJ (2017) An algorithm for the optimization of finite element integration loops. *ACM Transactions on Mathematical Software (TOMS)* 44(1): 3.
- Luporini F, Varbanescu AL, Rathgeber F, et al. (2015) Cross-loop optimization of arithmetic intensity for finite element local assembly. *ACM Transactions on Architecture and Code Optimization (TACO)* 11(4): 57.
- McRae ATT, Bercea GT, Mitchell L, et al. (2016) Automated generation and symbolic manipulation of tensor product finite elements. *SIAM Journal on Scientific Computing* 38(5): S25–S47.
- Müthing S, Piatkowski M and Bastian P (2017) High-performance implementation of matrix-free high-order discontinuous Galerkin methods. *arXiv preprint arXiv:1711.10885*.
- OpenMP Architecture Review Board (2018) OpenMP Application Programming Interface Version 5.0. Available: <https://www.openmp.org/wp-content/uploads/OpenMP-API-Specification-5.0.pdf> (accessed 5 November 2020).
- Rathgeber F, Ham DA, Mitchell L, et al. (2016) Firedrake: automating the finite element method by composing abstractions. *ACM Transactions on Mathematical Software* 43(3): 24:1–24:27.
- Rathgeber F, Markall GR, Mitchell L, et al. (2012) PyOP2: A High-Level Framework for Performance-Portable Simulations on Unstructured Meshes. In: *2012 SC Companion: High Performance Computing, Networking Storage and Analysis*, Salt Lake City, UT, USA, 10–16 November 2012. Piscataway: IEEE. ISBN 978-0-7695-4956-9, pp. 1116–1123. DOI:10.1109/SC.Companion.2012.134.
- Sun T (2019a) Cross-element vectorization in Firedrake. Available at: <https://www.codeocean.com/>. DOI: <https://doi.org/10.24433/CO.8386435.v2> (accessed 5 November 2020).
- Sun T (2019b) `tj-sun/firedrake-vectorization`: scripts for experimental evaluation for the manuscript on cross-element vectorization. DOI: 10.5281/zenodo.3365432 (accessed 5 November 2020).
- Verdoolaage S (2010) ISL: An integer set library for the polyhedral model. In: *International Congress on Mathematical Software*, Kobe, Japan, 13–17 September 2010, pp. 299–302. Berlin: Springer.
- Williams S, Waterman A and Patterson D (2009) Roofline: an insightful visual performance model for multicore architectures. *Communications of the ACM* 65–76. Available at: <http://dl.acm.org/citation.cfm?id=1498785> (accessed 5 November 2020).
- Zenodo/Firedrake (2019) Softwareusedin'Astudyofvectorization for matrix-free finite element methods'. DOI:10.5281/zenodo.3362177 (accessed 5 November 2020).
- Zhang B (2016) Guide to automatic vectorization with Intel AVX-512 instructions in Knights Landing processors. Available at: <https://colfaxresearch.com/knl-avx512> (accessed 19 March 2019).

Author biographies

Tianjiao Sun is a PhD student in the Department of Computing at Imperial College London, where he is a member

of the Software Performance Optimisation research group. His main research interests include domain-specific compilers and optimisations for numerical simulations.

Lawrence Mitchell is an Assistant Professor in the Department of Computer Science at Durham University. His research is in high performance computing and computational mathematics. Much of his recent focus has been in the development of compilers and software abstractions for the development of numerical models implemented using the finite element method. This research is concretely realised in the open source Firedrake project. He is particularly interested in preconditioning techniques for challenging problems in computational and atmospheric fluid dynamics.

Kaushik Kulkarni is a third year PhD student at the University of Illinois at Urbana-Champaign, where he is a member of the Scientific Computing group in the Computer Science Department. His research focuses on the area of transformation based code generation in applications motivated by the finite element method.

Andreas Klöckner is an Associate Professor in the scientific computing area within the Department of Computer Science at the University of Illinois at Urbana-Champaign. His research focuses on high-order accurate integral equation methods and fast algorithms for elliptic boundary value problems as well as code transformation for high-performance scientific computing. He is the recipient of a 2017 National Science Foundation CAREER Award. In support of his research, Dr. Klöckner has released numerous scientific software packages. Previously, he was a Courant Instructor at the Courant Institute of Mathematical Sciences at New York University with Leslie Greengard, after obtaining his PhD degree from the Division of Applied Mathematics at Brown University in 2010 with Jan Hesthaven.

David A Ham is a Senior Lecturer (Associate Professor) in the Department of Mathematics at Imperial College London. His research centres on code generation approaches to high performance finite element simulation, chiefly through the Firedrake project. He was awarded the 2015 Wilkinson Prize for Numerical Software for his contribution to the Dolfin-Adjoint automated adjoint simulation system, and is an executive editor of Geoscientific Model Development.

Paul HJ Kelly leads the Software Performance Optimisation research group in the Department of Computing at Imperial College London. His main research focus is compiler technology, specifically automated delivery of domain-specific performance optimisations.