



Calhoun: The NPS Institutional Archive
DSpace Repository

Faculty and Researchers

Faculty and Researchers' Publications

2016

Array Program Transformation with Loo.py by Example: High-Order Finite Elements

Klöckner, Andreas; Wilcox, Lucas C.; Warburton, T.

ARRAY'16 (June 14, 2016). Santa Barbara, CA, USA.
<http://hdl.handle.net/10945/48656>

This publication is a work of the U.S. Government as defined in Title 17, United States Code, Section 101. Copyright protection is not available for this work in the United States.

Downloaded from NPS Archive: Calhoun



Calhoun is the Naval Postgraduate School's public access digital repository for research materials and institutional publications created by the NPS community. Calhoun is named for Professor of Mathematics Guy K. Calhoun, NPS's first appointed -- and published -- scholarly author.

Dudley Knox Library / Naval Postgraduate School
411 Dyer Road / 1 University Circle
Monterey, California USA 93943

<http://www.nps.edu/library>

Array Program Transformation with Loo.py by Example: High-Order Finite Elements

Andreas Klöckner

University of Illinois at
Urbana-Champaign, Urbana, IL, USA
andreask@illinois.edu

Lucas C. Wilcox

Naval Postgraduate School, Monterey,
CA, USA
lwilcox@nps.edu

T. Warburton

Virginia Polytechnic Institute and State
University, Blacksburg, VA, USA
tim.warburton@vt.edu

Abstract

To concisely and effectively demonstrate the capabilities of our program transformation system Loo.py, we examine a transformation path from two real-world Fortran subroutines as found in a weather model to a single high-performance computational kernel suitable for execution on modern GPU hardware. Along the transformation path, we encounter kernel fusion, vectorization, prefetching, parallelization, and algorithmic changes achieved by mechanized conversion between imperative and functional/substitution-based code, among a number more. We conclude with performance results that demonstrate the effects and support the effectiveness of the applied transformations.

Categories and Subject Descriptors D.3.4 [Programming Languages]: Processors— Code generators; D.1.3 [Programming Languages]: Programming Techniques— Concurrent programming; G.4 [Mathematics of Computing]: Mathematical software

Keywords Code generation, high-level language, GPU, substitution rule, embedded language, high-performance, program transformation, OpenCL

1. Introduction

User-guided transformation of numerical and array-based computations is an area of sustained interest across the areas of high-performance computing, programming languages, and numerical methods. Its existence is inspired by the discrepancy between a compiler’s hypothetical ability to transform programs in ways that would ostensibly be beneficial to performance and their practical inability to do so as hampered by (1) the size of the search space of such transformations, and (2) a compiler’s ability to prove such transformations correct and/or equivalent to prior behavior.

Descriptions of such systems in the literature often restrict themselves to simple, common, “micro-benchmark” examples. It is understandable that this approach is seeing much use, since it is the least demanding in terms of the reader’s attention span, the required space, and, also, the implemented capability of the transformation system. Further, it facilitates easy comparison of performance re-

sults between systems and allows for the presentation of many examples. This article is an experiment in the opposite approach. We present a single, realistic example of program transformation drawn from the application domain of weather prediction. In taking this approach, we hope to better showcase capabilities present in our transformation system and the language mechanisms present that enable their use, all while describing the trade-offs that lead to the actual implemented design. We hope that others may follow our example and that a body of literature may arise that can serve to motivate and guide the discussion on program transformation.

Loo.py [22, 23] is a programming system for array computations that targets CPUs, GPUs, and other, potentially heterogeneous compute architectures. One salient feature of Loo.py is that programs written in it necessarily consist of two parts:

- A semi-mathematical statement of the array computation to be carried out, in terms of a *loop polyhedron* and a partially ordered *set of ‘instructions’*.
- A sequence of *kernel transformations*, driven by an ‘outer’ program in the high-level scripting language Python [41].

This strong separation is an explicit design goal, as it enables specialization of users, cleanliness of notation in either part, as well as greater flexibility in terms of transformation.

The present article demonstrates how Loo.py can function as a code generation and transformation engine for computational code originally expressed in a subset of Fortran while maintaining its full capability to transform the ingested code in a manner comprehensible and useful to the author of the original program. A number of mechanisms are described that are intended to aid the formulation of transformations on array computations in this setting. As one example of the issues that arise, the strong separation of semantics and transformation, while desirable, also poses a difficulty. For example, unlike in an annotation-based setting, where lexical proximity alone can be used to indicate what part of a program is to be transformed, this option does not exist for Loo.py, and so alternatives have to be devised.

1.1 Related work on code generation

The literature on code generation and optimization for array languages is vast, and no attempt will be made to provide a survey of the subject in any meaningful way. Instead, we will seek to highlight a few approaches that have significantly influenced the thinking behind Loo.py, are particularly similar, or provide ideas for further development. Loo.py is heavily inspired by the polyhedral model of expressing static-control programs [2, 12]. While it takes significant inspiration from this approach, the details of how a program is represented, beyond the existence of a loop domain, are quite different. High-performance compilation for GPUs, by now, is hardly a new topic, and many different approaches have

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ARRAY’16, June 14, 2016, Santa Barbara, CA, USA.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM . . . \$15.00.

<http://dx.doi.org/10.1145/>

been used, including ones using OpenMP-style directives [20, 27], ones that are fully automatic [44], ones based on functional languages [39], and ones based on the polyhedral model [42]. Other ones define an automatic, array computation middleware [15] designed as a back-end for multiple languages, including Python. Automatic, GPU-targeted compilers for languages embedded in Python also abound [5, 8, 35], most of which transform a Python AST at run-time based on various levels of annotation and operational abstraction.

User-guided program transformation based on polyhedral representation has received considerable attention over the years. Perhaps the conceptually closest prior work to the approach taken by Loo.py is CUDA-CHiLL [36], which performs source-to-source translation based on a set of user-controlled transformations [6, 19]. Loo.py and CHiLL still are not quite alike, using dissimilar intermediate representations, dissimilar levels of abstraction in the description of transformations, and a dissimilar (static vs. program-controlled) approach to transformation. Other similar projects include the AlphaZ [45] and Clay systems [1], although these projects emphasize the scheduling of a given workload rather than its algorithmic or data-based transformation. Rewriting- and substitution-focused systems such as Terra [10], like Loo.py, provide a powerful building block for DSLs, but they lack the loop transformation and parallelization capabilities afforded by polyhedral representation.

Other optimizing compilers assume a substantial amount of domain knowledge (such as what is needed for assembly of finite element matrices) and leverage this to obtain parallel, optimized code. One example of this family of code generators is COFFEE [28].

Source-to-source transformation similarly has been studied extensively, with many mature systems existing in the literature (see for instance Schordan and Quinlan [37] and Dave et al. [9]).

1.2 Related work on our example benchmark

Our benchmark detailed below is an embodiment of the ‘unstructured grid’ ‘dwarf’ of Colella’s oft-cited seven dwarfs [7]. Specifically, we consider a subclass of finite element models, the continuous [33] and discontinuous spectral element methods [3], which are well suited for simulating wave phenomena such as acoustics, elastodynamics, electromagnetics, and fluid dynamics in complex geometry. We focus on a computation from NUMA (Nonhydrostatic Unified Model of the Atmosphere) [21], the dynamical core of the U.S. Navy’s next generation nonhydrostatic atmospheric prediction system, NEPTUNE (Navy Environmental Prediction SysTem Utilizing the NUMA corE) [11, 13, 14]. Although we focus on this particular application, the computation would be similar in other application domains.

The continuous and discontinuous spectral element method has received attention from the high-performance computing community. For example, it has been used by Tufo and Fischer [40] and Komatitsch et al. [24] to win the ACM/IEEE Supercomputing Gordon Bell prize. Furthermore, it also has been ported to GPUs for seismic wave propagation using handwritten CUDA kernels [4, 17, 25]; for fluid flow using handwritten OpenCL [38] and OpenACC [18, 29] kernels, and using an in-built domain specific language targeting multiple backends [43]; for electromagnetic wave propagation using OpenACC [32]; and for elliptic problems using handwritten OCCA [30] kernels targeting multiple backends [34]. While porting part of the atmospheric climate model CAM-SE from CUDA Fortran to OpenACC its authors note that “it is highly unlikely that a literal single source code would suffice for performance portability” [31]. This provides motivation to use programmatic tools, like Loo.py, to generate the desired set of kernels needed for performance portability.

2. Structure of the Computation

To illustrate some of the transformations in Loo.py, we consider a computation performed in a simplified dynamical core of a non-hydrostatic numerical weather prediction model. This model is a discretization of Euler’s equations in a method-of-lines approach using a continuous or discontinuous Galerkin spectral element discretization in space on a curvilinear hexahedral mesh of tensor-product polynomial elements (similar to Kelly and Giraldo [21]). Focusing on explicit time integration, for our benchmark computation we consider the element volume contribution to the rate function used by the ordinary differential equation solver. This computation is common to both the continuous and discontinuous Galerkin spectral element methods, which differ by how the elements are connected. For a more detailed description of these methods than provided below, see for example Kopriva [26].

For our computation, we consider the governing equations of a dry atmosphere (without gravity and viscous terms) which are

$$\frac{\partial q_b}{\partial t} + \sum_a \frac{\partial f_{ab}(\mathbf{q})}{\partial x_a} = 0 \quad (1)$$

using density, momentum (the three components), potential temperature density and three tracer densities as prognostic variables

$$\mathbf{q} = [\rho \ U_1 \ U_2 \ U_3 \ \Theta \ Q_1 \ Q_2 \ Q_3]^\top,$$

where the a th column of the flux is

$$f_a(\mathbf{q}) = \begin{bmatrix} U_a \\ \frac{U_a U_1}{\rho} + \delta_{a1} p \\ \frac{U_a U_2}{\rho} + \delta_{a2} p \\ \frac{U_a U_3}{\rho} + \delta_{a3} p \\ U_a \Theta \\ \frac{U_a Q_1}{\rho} \\ \frac{U_a Q_2}{\rho} \\ \frac{U_a Q_3}{\rho} \end{bmatrix},$$

which is equation set 2C (without gravity and viscous terms) from Giraldo et al. [16]. Here, the spatial coordinates are $[x_1 \ x_2 \ x_3]^\top$, t is time, δ_{ab} is the Kronecker delta, and p is the pressure obtained from the equation of state $p = p_0 \left(\frac{R\Theta}{p_0}\right)^\gamma$, where p_0 is a constant reference pressure at the surface, $R = c_p - c_v$ is the gas constant given in terms of the specific heats for constant pressure and volume, c_p and c_v respectively, and $\gamma = \frac{c_p}{c_v}$ is the specific heat ratio.

The semi-discretization of Euler’s equations (1) by the continuous and discontinuous Galerkin methods can be split into volume, \mathbf{v} , and surface, \mathbf{s} , contributions as

$$\frac{\partial \mathbf{q}_b^{eijk}}{\partial t} = \mathbf{v}_b^{eijk}(\mathbf{q}) + \mathbf{s}_b^{eijk}(\mathbf{q}),$$

where e indexes the N_e elements and (i, j, k) indexes the element’s N_q^3 grid points. The surface term connects the elements together using a ‘numerical flux’ in the discontinuous Galerkin method and the direct stiffness-summation operator in the continuous Galerkin method. As the benchmark in this paper, we consider the computation of the volume term

$$\begin{aligned} \mathbf{v}_b^{eijk}(\mathbf{q}) = & \sum_{n,a} \frac{1}{J_{eijk}} D^{in} g_{a1}^{enjk} f_{ab}^{enjk}(\mathbf{q}) \\ & + \sum_{n,a} \frac{1}{J_{eijk}} D^{jn} g_{a2}^{eink} f_{ab}^{eink}(\mathbf{q}) \\ & + \sum_{n,a} \frac{1}{J_{eijk}} D^{kn} g_{a3}^{eijn} f_{ab}^{eijn}(\mathbf{q}), \end{aligned} \quad (2)$$




Figure 1. Computational layout of the volume term (2). Non-white boxes represent the intra-element grid points involved in the volume term (2) associated with the center box; the first term on the right-hand side involves the (red) boxes in a line parallel to the i -axis, the second term involves the (green) boxes in a line parallel to the j -axis, and the third term involves (blue) boxes in a line parallel to the k -axis. The group of boxes with the same value of k is called the k -slice. In addition, element reference directions r , s , and t respectively correspond to the loop indices i , j , and k .

where D is the element differentiation matrix and J and g are geometric factors related to the Jacobian of the transformation from the reference to the physical elements. The computation of the volume term is most expensive aspect of the entire solver. A visual impression of the per-element (e) loop (and data) layout for this operation is given in Figure 1.

3. Transforming the Computation

We begin our consideration with a program that calculates the flux on each element in the r -reference direction. The program enters the transformation system described as a short subroutine in Fortran [23]:

```

subroutine refFluxR(Ne, Nq, geo, D, q, rhsq, p_p0, &
  p_Gamma, p_R)
  implicit none
  integer*4 Ne, Nq, e, i, j, k, n
  real geo(Nq,Nq,Nq,11,Ne), D(Nq,Nq)
  real rhsq(Nq,Nq,Nq,8,Ne), q(Nq,Nq,Nq,8,Ne)

  real U1, U2, U3, Rh, Th, P, Q1, Q2, Q3
  real U1flx, U2flx, U3flx, Rhflx, Thflx
  real Q1flx, Q2flx, Q3flx, p_p0, p_Gamma, p_R
  real g11, g21, g31, Jinv, udotGradR, JiD

  do e = 1, Ne
    do k = 1, Nq
      do j = 1, Nq
        do i = 1, Nq
          do n = 1, Nq
            !$loopy begin tagged: local_prep
              U1 = q(n,j,k,1,e)
              U2 = q(n,j,k,2,e)
              U3 = q(n,j,k,3,e)
              Rh = q(n,j,k,4,e)
              Th = q(n,j,k,5,e)
              Q1 = q(n,j,k,6,e)
              Q2 = q(n,j,k,7,e)
              Q3 = q(n,j,k,8,e)

              g11 = geo(n,j,k,1,e)
              g21 = geo(n,j,k,2,e)
              g31 = geo(n,j,k,3,e)

              Jinv = geo(i,j,k,10,e)

```

```

  P = p_p0*(p_R*Th/p_p0) ** p_Gamma
  udotGradR = (g11*U1 + g21*U2 + g31*U3)/Rh
  !$loopy end tagged: local_prep

```

```

  JiD = Jinv*D(i,n)

```

```

  !$loopy begin tagged: compute_fluxes
  U1flx = U1*udotGradR + g11*P
  U2flx = U2*udotGradR + g21*P
  U3flx = U3*udotGradR + g31*P
  Rhflx = Rh*udotGradR
  Thflx = Th*udotGradR
  Q1flx = Q1*udotGradR
  Q2flx = Q2*udotGradR
  Q3flx = Q3*udotGradR
  !$loopy end tagged: compute_fluxes

```

```

  rhsq(i,j,k,1,e) = rhsq(i,j,k,1,e) - JiD*U1flx
  rhsq(i,j,k,2,e) = rhsq(i,j,k,2,e) - JiD*U2flx
  rhsq(i,j,k,3,e) = rhsq(i,j,k,3,e) - JiD*U3flx
  rhsq(i,j,k,4,e) = rhsq(i,j,k,4,e) - JiD*Rhflx
  rhsq(i,j,k,5,e) = rhsq(i,j,k,5,e) - JiD*Thflx
  rhsq(i,j,k,6,e) = rhsq(i,j,k,6,e) - JiD*Q1flx
  rhsq(i,j,k,7,e) = rhsq(i,j,k,7,e) - JiD*Q2flx
  rhsq(i,j,k,8,e) = rhsq(i,j,k,8,e) - JiD*Q3flx

```

```

end do
end do
end do
end do
end subroutine refFluxR

```

In many ways, the above represents a ‘typical’ Fortran code. It represents a (more or less) direct translation of the first term on the right-hand side of the volume discretization term (2), faithful to Fortran’s original spirit of ‘formula translation’. A deep nest of do loops dominates the structure. Nonetheless, a few peculiar aspects are of note: first, the code is written entirely without regard to performance. With reasonable clarity, it exhibits the computational intent. Second, a small amount of annotation is visible. Two regions, delimited by loopy begin/end tagged markers are shown. These carry no meaning of their own. They merely serve as location markers for subsequent transformation operations.

One common performance challenge that stands in direct competition with code clarity is reuse of data and intermediate results across multiple calculations. In the context of this example, a salient instance of this issue is the computation of fluxes in multiple reference directions. All such calculations refer to the same input data, compute the same (e.g. P) or related (e.g. \dots flx) intermediate results. Yet, a subroutine that merges a number of such calculations invariably more cluttered and less transparent than the code shown above.

To illustrate, we include in our consideration a further computational subroutine to compute the fluxes in the s -reference direction. This routine differs from the above merely in the data access pattern. Specifically, the input of the field variables uses indices as follows:

```

  U1 = q(i,n,k,1,e)
  U2 = q(i,n,k,2,e)
  U3 = q(i,n,k,3,e)

```

(further accesses proceed analogously) and the differentiation matrix is accessed as:

```

  JiD = Jinv*D(j,n)

```

Another minor difference is that the s -kernel uses a different set of geometric factors. All other aspects, including the final output increments, are exactly the same.

Transformations in Loo.py are expressed using the Python programming language. This code can be given either in a separate

location or, as shown here, as part of the same Fortran file in a directive-like comment block.

As we proceed through the transformations, we will occasionally denote intermediate versions of the kernel as being of a certain optimization level. Where relevant, this is indicated in parentheses as '[Opt. level N]'. These optimization levels are then later referenced in the results section to clarify the individual performance impact of each of the transformation steps.

The transformation begins by receiving the translation unit (including the comment block) in an implicitly defined variable `SOURCE` which it then translates from Fortran into Loopy's intermediate representation, resulting in two kernels which are returned:

```
!$loopy begin
! from loopy import *
! import loopy.match as m
!
! f_r, f_s = parse_fortran(SOURCE, FILENAME,
    auto_dependencies=False)
```

We omit the Fortran comment markers ('!') in the following. Next up, the r - and s -flux computation kernels are fused together. This operation joins loop axes ('inames') with matching names. The projection of the domains of each pair of kernels onto their common subset of inames must match in order for fusion to succeed. Kernel arguments with identical names are likewise merged, assuming non-conflicting declarations. Local variables are kept private to each of the fused kernels by applying a user-given suffix, here ' $_r$ ' and ' $_s$ '.

```
f = fuse_kernels([f_r, f_s], suffixes=["_r", "_s"])
```

In situations more complicated than this one, data flow between kernels may be specified as part of fusion to determine dependencies.

Next, fixing values of parameters helps determine control flow by fixing loop bounds and removing run-time data references. A non-emptiness assumption further helps reduce unnecessary conditionals.

```
f = fix_parameters(f, Nq=Nq)
f = fix_parameters(f, p_p0=1, p_Gamma=1.4, p_R=1)
f = assume(f, "Ne >= 1")
```

To further determine control flow, we set the ordering of the sequential loops making the intra-element axis k the outer loop. Furthermore, we assign the element number iname e to the abstract 'core' axis and the intra-element axes i and j to two of the 'SIMD lane' axes exposed in the OpenCL model.

```
f = set_loop_priority(f, "k")
f = tag_inames(f, "e:g.0, i:l.0, j:l.1")
```

[\setminus Opt. level 1] The following section determines the data layouts of the bulk variables q and $rhsq$. At the coarsest level, this is done by specifying a "nesting order" of axes, sorting them from fastest- to slowest-varying. This is indicated by a capital N followed by an integer. The computation of the surface term (mentioned in Section 2) amounts to scattered, indirect memory access. To better match this access to the wide bus width and relative dearth of caches on the target hardware, we use short vector types (of length four) as the basic granularity of our access. We map the 'field' axis (of length eight) of the arrays containing our degrees of freedom across the vector entries. To facilitate this, we split the axis in two as $field = 4field_outer + field_inner$, where $field_inner$ and $field_outer$ are now exposed as separate array axes arranged in "column-major" (i.e., "Fortran" order, i.e., $field_inner$ comes first. $field_inner$, being of fixed length 4, is then ready to be tagged for implementation as a short vector ("vec").

In addition, this section stipulates some identifiers for array axes which are subsequently used in automatic name generation for, e.g., inames in precomputation dealing with this array axis.

```
f = tag_data_axes(f, "D", "NO,N1")
for name in ["q", "rhsq"]:
    f = set_array_dim_names(f, name, "i,j,k,field,e")
    f = split_array_dim(
        f, (name, 3, "F"), 4, auto_split_inames=False)
    f = tag_data_axes(f, name, "NO,N1,N2,vec,N4,N3")
```

[\setminus Opt. level 2] Since the differentiation matrix D is referenced frequently, it makes sense to bring it into local scratchpad memory, which is accomplished as follows:

```
f = add_prefetch(f, "D[:,:]")
```

[\setminus Opt. level 3] The main computation in this kernel consists of two layers: first, the purely local evaluation of the flux function $f_a(\mathbf{q})$ followed by the non-local derivative operation. To allow this multi-layered computation to be transformed, we employ substitution rules as an intermediate form. These can be imagined as parametric, expression-level macros that may always be expanded. One of their main purposes is to assign a name to a certain intermediate result. The Fortran form of our kernel begins by storing field values into local variables. This can be seen in an area of the code carrying a 'loopy begin tagged: local_prep' annotation. This occurs in preparation for the pre-computation and storing of the flux values, which refer to this data. To retain the semantic content of these variable names while removing the computational inconvenience, we use an `assignment_to_subst` transformation to obtain an appropriate substitution rule. In code, the effect of this is the elimination of the temporary variable, all references to which are replaced by references to the rule, which in turn expands to the expression originally assigned to the temporary. The annotation above is used to create a list of all the variable names for which this is needed. Going forward, we will refer to this data as 'degree of freedom' data.

```
local_prep_var_names = set()
for insn in find_instructions(f, "tag:local_prep"):
    assignee = insn.assignee_name
    local_prep_var_names.add(assignee)
    f = assignment_to_subst(f, assignee)
f = assignment_to_subst(f, "JiD_r")
f = assignment_to_subst(f, "JiD_s")
```

[\setminus Opt. level 4] In a similar fashion, we make substitution rules of the flux computations, for each component of the solution vector. Given these rules, we use the `precompute` transformation to create a loop nest that computes and stores the values of the fluxes in a temporary variable. We perform this precomputation for one k -slice at a time (cf. Figure 1), yielding a family of two-dimensional calculations each determined by two pairs of two inames. The first such pair determines the access footprint, i.e., the set of inames that span the index space for which the pre-computation will replace the actual memory access and subsequent expression evaluation. The second such pair determines the inames over which the pre-computation is actually carried out. Since this transformation needs to be carried out once for the r - and once for the s -kernel, we encode a (Python) `for` loop. For the r -kernel, the first pair is (j, n) (cf. the `local_prep` section of the code). In each case, the pair of inames carrying out the pre-computation is a permutation of the new inames ii and jj . We also fix the data layout for the newly created temporary.

The two main operational aspects of the `precompute` transformation are the creation of the temporary variable and the rewriting of the input substitution rule into an assignment operation. The latter mainly involves changing subscripts to use the inames that are being used to carry out the pre-computation.

It is worth noting that, since the read operations on our degree of freedom data exist in a (sub-)substitution rule as created above, the subscript information therein tracks the changes being made. This is the case because, for Loo.py's purposes, term rewriting is applied as if the substitution rules did not exist, i.e., term rewriting, at least notionally, always operates on a 'fully expanded' version of the expression tree. Substitution rules are nonetheless preserved, however it may be the case that multiple versions of a rule are created depending on different rewrites occurring on a per-usage-site basis. Similarly, substitution rules with identical content are merged. Through this merging mechanism as well as an appropriate choice of the pre-computation inames, we are able to merge the degree-of-freedom reads from the r - and the s -kernel, ensuring that the resulting fused kernel only performs a single read access to that data.

```
comps = ["U1", "U2", "U3", "Rh", "Th", "Q1", "Q2", "Q3"]
for comp in comps:
    for ref_ax, flux_inames, flux_precomp_inames in [
        ("r", ("j", "n"), ("jj", "ii")),
        ("s", ("i", "n"), ("ii", "jj")),
    ]:
        flux_var = comp+"flx_"+ref_ax
        f = assignment_to_subst(f, flux_var)
        flux_store_name = "flux_store_" + comp + ref_ax
        f = precompute(f, flux_var+"_subst", flux_inames,
            temporary_name=flux_store_name,
            precompute_inames=flux_precomp_inames,
            default_tag=None)
        f = tag_data_axes(
            f, flux_store_name, {"ii":"N0", "jj":"N1"})
```

Next, we determine that the new inames ii and jj should be mapped to the abstract SIMD lane indices along axis 0 and 1. To conserve temporary, on-chip storage, we next request that, within each of the r - and the s -parts of the fused kernel, the temporary flux storage for each solution component be aliased, resulting in two storage areas being used, one for the r -kernel and one for the s -kernel. This aliasing automatically creates scheduling constraints that prohibit the range from the first right to the last read for each temporary to overlap with the same range for another temporary that is aliased to the same actual storage.

```
f = tag_inames(f, "ii:1.0,jj:1.1")
for ref_ax in ["r", "s"]:
    f = alias temporaries(f,
        ["flux_store_" + comp + ref_ax for comp in comps])
```

These scheduling constraints necessarily prevent that a single n (i.e., summation) loop can be used, since at any one time only two sets of precomputed flux data are available. Further, since Loo.py uses inames as unique identifiers of loops, a (kernel-level) for loop that has been opened and closed once cannot be opened again. Therefore, we use the `rename_iname` transformation to create copies of the iname n . A unique copy of the n iname is created for each pair of flux temporaries, to be used by the summation instructions that refer to them. That set of instructions is identified by the `Reads(flux_store_name)` match query.

```
for comp in comps:
    for ref_ax in ["r", "s"]:
        flux_store_name = "flux_store_" + comp + ref_ax
        f = rename_iname(f, "n", "n_"+comp,
            within=m.Reads(flux_store_name), existing_ok=True)
```

[\setminus Opt. level 5] To appropriately limit the scope of transformations, such as in the previous scenario, Loo.py provides a small, but fairly comprehensive query language to match program context, based on features of instructions and substitution rules.

Moving along, we once again use the `precompute` transformation, this time to fetch the degree of freedom data into temporary

variables. Note that $Jinv$ is excluded from this pre-computation because it follows a different access pattern.

```
for prep_vname in local_prep_var_names:
    if prep_vname.startswith("Jinv") or "s" in
        prep_vname:
            continue
    f = precompute(f,
        find_one_rule_matching(f, prep_vname+"_subst*"))
```

[\setminus Opt. level 6] Next, by using the `add_prefetch` transformation, we create a batch fetch process (which will subsequently be vectorized) to load the degree-of-freedom data across the field components on to the chip. The transformation `add_prefetch` combines (1) substitution rule creation based on an existing global (i.e., off-chip) variable, (2) replacement of appropriate references to that variable with references to that substitution rule, and (3) the application of the `precompute` transformation to that same substitution rule. The net effect of this is for off-chip data to be stored in an on-chip temporary variable and then referenced from there instead of from global memory.

```
f = add_prefetch(f, "q[ii,jj,k,:,:e]")
```

The Fortran code, as its last step of processing, performs in-place updates on the computed time derivative $rhsq$ of the degree-of-freedom data. This particular access, while easy to code, is not efficient, because it touches each of these off-chip quantities $2N_q$ times. The performance issue of this access is further exacerbated by the chosen (vectorized) memory layout, leading to non-contiguous access. Buffering this data on the chip for the duration of the kernel addresses all these performance concerns:

```
f = buffer_array(f, "rhsq", (),
    fetch_bounding_box=True, default_tag="for",
    init_expression="0",
    store_expression="base + buffer")
```

[\setminus Opt. level 7] We fix the data layouts of the on-chip versions of D , q and $rhsq$. For the latter two, we ensure that their data format matches the vectorization of their off-chip versions. To ensure that the access carrying this data onto and off of the chip proceeds using full width vectors, we also need to tag the relevant inames with the `vec` iname implementation tag. Note that the execution recipe as chosen by the iname tag is independent of the data layout, as data stored as vectors can still be accessed in a scalar fashion, hence the necessity to specify both.

```
f = tag_data_axes(f, "rhsq_buf", "vec,N0")
f = tag_data_axes(f, "q_fetch", "vec,N0")
f = tag_data_axes(f, "D_fetch", "NO,N1")
f = tag_inames(f,
    "rhsq_init_field_inner:vec,rhsq_store_field_inner:vec,"
    "rhsq_init_field_outer:unr,rhsq_store_field_outer:unr,"
    "q_dim_field_outer:unr,q_dim_field_inner:vec")
```

Our last transformation step takes advantage of the distributive law in the setting where a variable being additively updated is also buffered on the chip. In this setting, common factors can be 'pulled out' and applied only once at the time when the buffered variable is being written back to its off-chip location. This is realized by the following transformation:

```
f = collect_common_factors_on_increment(f, "rhsq_buf")
```

[\setminus Opt. level 8] This completes our walkthrough of the transformation and yields the final version of the kernel.

4. Results

Performance Throughout the Transformation. An interesting data set whose collection is enabled by the transformation-based nature

Opt. Level	Wall Time (ms)	Speedup	G FLOP/s	Bandwidth (GB s ⁻¹)
Radeon R9 FURY X				
1	25.0	1.00	88.3	263
2	13.5	1.85	164	487
3	14.2	1.77	156	448
4	14.0	1.79	373	1520
5	2.43	10.3	758	3050
6	2.13	11.7	707	2640
7	1.37	18.2	1120	1630
8	1.36	18.4	816	395
GeForce GTX TITAN X				
1	7.18	1.00	308	915
2	5.14	1.40	429	1280
3	5.06	1.42	437	1250
4	9.17	0.783	568	2320
5	1.83	3.91	1000	4040
6	2.09	3.44	722	2700
7	2.04	3.52	753	1100
8	2.06	3.48	539	261
Tesla K40c				
1	14.3	1.00	154	459
2	13.4	1.07	165	492
3	13.3	1.08	166	477
4	21.0	0.682	248	1020
5	4.44	3.22	414	1670
6	3.61	3.96	418	1560
7	3.44	4.16	447	651
8	3.28	4.35	338	164

Table 1. Performance of the kernel for the various optimization levels described in Section 3. Note, the FLOP/s and bandwidth are estimated, not measured, as described in Section 4.

of our work arises from the performance characteristics across all versions of the kernel. To this end, we have indicated various points within the preceding section as ‘optimization levels.’ Table 1 shows this data for three recent GPU architectures. These results were obtained using $N_q=8$ and $N_e=6910$. We note that even the baseline (level 1) version of the kernel makes full use of the available parallelism and concurrency on the GPUs. The speedup number indicated thus summarizes what gains in wall time can be achieved by performing additional tuning beyond such parallelization.

Two features of this data are immediately striking: first, performance is not necessarily monotonic. Some transformations may individually lose performance, but, in combination with later transformations enabled by them, make an even larger gain possible. Second, some GPU architectures are more sensitive to specific tuning efforts than others. Specifically, in our example, the AMD GPU appears to benefit to a much greater extent than the others.

For a complementary point of view, Table 1 also shows absolute performance numbers in terms of floating-point operation rate and memory bandwidth. While each computation for which timing data is displayed computes the same numerical output, it is worth while to observe that some intermediate optimization levels achieve far higher FLOP rates or memory bandwidths than the final versions. These intermediate versions are not superior despite the higher performance because they ultimately take longer to complete in terms of wall time. This does highlight that some of the transformations realized by Loo.py amount to algorithmic changes affecting the overall complexity of the computation.

Dependency on N_q . Our next dataset highlights the performance of the transformed kernel for different values of N_q , where the total number of degrees-of-freedom is kept constant. The results, given

N_q	N_e	Wall Time (ms)	G FLOP/s	Bandwidth (GB s ⁻¹)
Radeon R9 FURY X				
4	55 300	3.20	206	168
8	6910	1.36	816	395
12	2050	4.26	367	126
16	864	6.46	312	83.3
GeForce GTX TITAN X				
4	55 300	1.78	369	302
8	6910	2.07	538	260
12	2050	2.21	709	244
16	864	2.12	953	254
Tesla K40c				
4	55 300	5.54	119	97.1
8	6910	3.29	338	164
12	2050	5.35	293	101
16	864	5.54	364	97.1

Table 2. Performance of the optimization level 8 kernel for different numbers of per-element grid points. Note, the total number of grid points is equal between runs. The FLOP/s and bandwidth are estimated, not measured, as described in Section 4.

in Table 2, show that the sensitivity of the performance on N_q varies depending on the hardware, with AMD hardware tending towards greater sensitivity than Nvidia. Ultimately, and irrespective of the target hardware, it is quite likely that different tuning approaches may be needed to yield consistent performance even for each value of N_q , even on the same device—and transformation-based programming with Loo.py provides a clear path of deriving these variants from the same, ‘clean’ source code.

The results presented in this section are for Loo.py¹ generated OpenCL kernels with 32-bit integers and floating point numbers. The OpenCL build flags used with the AMD-APP (1912.5) and CUDA 7.5.0 (355.11) platforms are:

```
-cl-denorms-are-zero, -cl-fast-relaxed-math,
-cl-finite-math-only, -cl-mad-enable, and
-cl-no-signed-zeros.
```

The performance measurements, FLOP/s and bandwidth, are statically counted in the generated kernels where add, multiply, divide, fused multiply-add, and special functions count as 1 FLOP. For the bandwidth calculation, all references (read or write) to global memory are counted. Due to caching effects, this way of measuring bandwidth can exceed the peak global memory bandwidth rates for the given devices. Timings were obtained using wall clock time by performing a few untimed ‘warm-up’ rounds of kernel execution, followed by the timing loop which was run until an overall run time of at least 0.3 seconds was reached.

Acknowledgments

AK’s work on Loo.py was supported in part by US Navy ONR grant number N00014-14-1-0117, and by the National Science Foundation under grant numbers DMS-1418961 and CCF-1524433. AK also gratefully acknowledges a hardware gift from Nvidia Corporation. TW and LCW’s work was supported in part by US Navy ONR grant numbers N00014-15-WX-01603 and N00014-13-1-0873, respectively. The authors would like to thank M. Wala for helpful suggestions that led to improvements in the manuscript.

¹The kernels tested in Section 4 are generated using Loo.py from the git repository <http://git.tiker.net/trees/loopy.git> commit 7d3f70aa76d627e95940c4281088dbc644da4cf5.

References

- [1] L. Bagnères, O. Zinenko, S. Huot, and C. Bastoul. Opening Polyhedral Compiler's Black Box. Mar. 2016.
- [2] C. Bastoul. Code generation in the polyhedral model is easier than you think. In *Proceedings of the 13th International Conference on Parallel Architectures and Compilation Techniques*, page 7–16, 2004. doi:10.1109/pact.2004.1342537.
- [3] K. Black. A conservative spectral element method for the approximation of compressible fluid flow. *Kybernetika*, 35(1):133–146, 1999. doi:10.3388/dmlcz/135273.
- [4] C. Burstedde, O. Ghattas, M. Gurnis, T. Isaac, G. Stadler, T. Warburton, and L. C. Wilcox. Extreme-scale AMR. In *2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, page nil, Nov. 2010. doi:10.1109/sc.2010.25.
- [5] B. Catanzaro, M. Garland, and K. Keutzer. Copperhead: compiling an embedded data parallel language. In *ACM SIGPLAN Notices*, volume 46, pages 47–56. ACM, 2011. doi:10.1145/2038037.1941562.
- [6] C. Chen, J. Chame, and M. Hall. CHiLL: A framework for composing high-level loop transformations. Technical Report 08-897, Department of Computer Science, University of Southern California, 2008. URL <http://www.cs.usc.edu/assets/001/82866.pdf>.
- [7] P. Colella. Defining software requirements for scientific computing, 2004. Talk by Phillip Colella about DARPA's High Productivity Computing Systems.
- [8] Continuum Analytics, Inc. Numba Pro, 2014.
- [9] C. Dave, H. Bae, S.-J. Min, S. Lee, R. Eigenmann, and S. Midkiff. Cetus: A source-to-source compiler infrastructure for multicores. *Computer*, 42(12):36–42, 2009. doi:10.1109/mc.2009.385.
- [10] Z. DeVito, J. Hegarty, A. Aiken, P. Hanrahan, and J. Vitek. Terra: A Multi-Stage Language for High-Performance Computing. *Proceedings of PLDI '13*, June 2013.
- [11] J. D. Doyle. A next generation atmospheric prediction system for the Navy. Annual report, Office of Naval Research, FY15. URL <http://www.onr.navy.mil/reports/FY15/mm2doyle.pdf>.
- [12] P. Fautrier. Automatic parallelization in the polytope model. In *The Data Parallel Programming Model*, pages 79–103. Springer, 1996. doi:10.1007/3-540-61736-1_44.
- [13] S. Gaberšek. Development of the Navy's next-generation nonhydrostatic modeling system. Annual report, Office of Naval Research, FY14. URL <http://www.onr.navy.mil/reports/FY14/mmgabers.pdf>.
- [14] S. Gaberšek. Development of the Navy's next-generation nonhydrostatic modeling system. Annual report, Office of Naval Research, FY15. URL <http://www.onr.navy.mil/reports/FY15/mmgabers.pdf>.
- [15] R. Garg and L. Hendren. A compiler toolkit for array-based languages targeting CPU/GPU hybrid systems. Technical Report 2012-3, Sable Research Group, Computer Science Department, McGill University, Montreal, QC, Canada, 11 2012. URL <http://www.sable.mcgill.ca/publications/techreports/2012-3/sable-10623-jsc.pdf>.
- [16] F. X. Giraldo, M. Restelli, and M. Läuter. Semi-implicit formulations of the Navier—Stokes equations: Application to nonhydrostatic atmospheric modeling. *SIAM Journal on Scientific Computing*, 32(6):3394–3425, 2010. doi:10.1137/090775889.
- [17] D. Göldeke, D. Komatitsch, and M. Möller. *Finite and Spectral Element Methods on Unstructured Grids for Flow and Wave Propagation Problems*, pages 183–206. Springer Science + Business Media, 2014. doi:10.1007/978-3-319-06548-9_9.
- [18] J. Gong, S. Markidis, M. Schliephake, E. Laure, L. Cebamanos, A. Hart, M. Min, and P. Fischer. NekBone with optimized OpenACC directives. In *Proceedings of the Second International Workshop on Sustainable Ultrascale Computing Systems (NESUS 2015): Krakow, Poland*, pages 63–70, 2015. doi:10016/21998.
- [19] M. Hall, J. Chame, C. Chen, J. Shin, G. Rudy, and M. Khan. Loop transformation recipes for code generation and auto-tuning. In *Languages and Compilers for Parallel Computing*, pages 50–64, 2010. doi:10.1007/978-3-642-13374-9_4.
- [20] T. D. Han and T. S. Abdelrahman. hiCUDA: High-Level GPGPU Programming. *IEEE Transactions on Parallel and Distributed Systems*, 22:78–90, 2011. doi:10.1109/TPDS.2010.62.
- [21] J. F. Kelly and F. X. Giraldo. Continuous and discontinuous Galerkin methods for a scalable three-dimensional nonhydrostatic atmospheric model: Limited-area mode. *Journal of Computational Physics*, 231(24):7988–8008, 2012. doi:10.1016/j.jcp.2012.04.042.
- [22] A. Klöckner. 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*, ARRAY'14, pages 82:82–82:87, Edinburgh, Scotland, 2014. ACM. doi:10.1145/2627373.2627387.
- [23] A. Klöckner. Loo.Py: From Fortran to Performance via Transformation and Substitution Rules. In *Proceedings of the 2nd ACM SIGPLAN International Workshop on Libraries, Languages, and Compilers for Array Programming*, ARRAY 2015, pages 1–6, New York, NY, USA, 2015. ACM. doi:10.1145/2774959.2774969.
- [24] D. Komatitsch, S. Tsuboi, C. Ji, and J. Tromp. A 14.6 billion degrees of freedom, 5 teraflops, 2.5 terabyte earthquake simulation on the earth simulator. In *Proceedings of the 2003 ACM/IEEE Conference on Supercomputing*, SC '03, pages 4–, New York, NY, USA, 2003. ACM. doi:10.1145/1048935.1050155.
- [25] D. Komatitsch, G. Erlebacher, D. Göldeke, and D. Michéa. High-order finite-element seismic wave propagation modeling with MPI on a large GPU cluster. *Journal of Computational Physics*, 229(20):7692–7714, 2010. doi:10.1016/j.jcp.2010.06.024.
- [26] D. A. Kopriva. *Implementing Spectral Methods for Partial Differential Equations*. Scientific Computation. Springer Netherlands, 2009. doi:10.1007/978-90-481-2261-5.
- [27] S. Lee and R. Eigenmann. OpenMPC: Extended OpenMP programming and tuning for GPUs. In *Proceedings of the 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–11. IEEE Computer Society, 2010. doi:10.1109/sc.2010.36.
- [28] F. Luporini, A. L. Varbanescu, F. Rathgeber, G.-T. Bercea, J. Ramanujam, D. A. Ham, and P. H. J. Kelly. Cross-loop optimization of arithmetic intensity for finite element local assembly. *ACM Trans. Archit. Code Optim.*, 11(4):57:1–57:25, Jan. 2015. doi:10.1145/2687415.
- [29] S. Markidis, J. Gong, M. Schliephake, E. Laure, A. Hart, D. Henty, K. Heisey, and P. Fischer. OpenACC acceleration of the Nek5000 spectral element code. *International Journal of High Performance Computing Applications*, 29(3):311–319, 2015. doi:10.1177/1094342015576846.
- [30] D. S. Medina, A. St.-Cyr, and T. Warburton. OCCA: A unified approach to multi-threading languages. *CoRR*, abs/1403.0968, 2014. URL <http://arxiv.org/abs/1403.0968>.
- [31] M. Norman, J. Larkin, A. Vose, and K. Evans. A case study of CUDA FORTRAN and OpenACC for an atmospheric climate kernel. *Journal of Computational Science*, 9:1–6, 2015. doi:10.1062/jcs.2015.04.022.
- [32] M. Otten, J. Gong, A. Mametjanov, A. Vose, J. Levesque, P. Fischer, and M. Min. An MPI/OpenACC implementation of a high-order electromagnetics solver with GPUdirect communication. *International Journal of High Performance Computing Applications*, 2016. doi:10.1177/1094342015626584.
- [33] A. T. Patera. A spectral element method for fluid dynamics: Laminar flow in a channel expansion. *Journal of Computational Physics*, 54(3):468–488, 1984. doi:10.1016/0021-9991(84)90128-1.
- [34] J. Remacle, R. Gandham, and T. Warburton. GPU accelerated spectral finite elements on all-hex meshes. *CoRR*, abs/1506.05996, 2015. URL <http://arxiv.org/abs/1506.05996>.
- [35] A. Rubinsteyn, E. Hielscher, N. Weinman, and D. Shasha. Parakeet: A just-in-time parallel accelerator for Python. In *Presented as part of the 4th USENIX Workshop on Hot Topics in Parallelism*, Berkeley, CA, 2012. USENIX.

- [36] G. Rudy, M. Khan, M. Hall, C. Chen, and J. Chame. A programming language interface to describe transformations and code generation. In *Languages and Compilers for Parallel Computing*, pages 136–150, 2011. doi:10.1007/978-3-642-19595-2_10.
- [37] M. Schordan and D. Quinlan. A source-to-source architecture for user-defined optimizations. In *Modular Programming Languages: Joint Modular Languages Conference, Proceedings*, pages 214–223, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg. doi:10.1007/978-3-540-45213-3_27.
- [38] N. Stilwell. gNek: A GPU accelerated incompressible Navier Stokes solver. Master’s thesis, Rice University, Houston, TX, Sept. 2013. doi:1911/72043.
- [39] J. Svensson, K. Claessen, and M. Sheeran. GPGPU kernel implementation and refinement using Obsidian. *Procedia Computer Science*, 1(1):2065–2074, 2010. doi:10.1016/j.procs.2010.04.231.
- [40] H. M. Tufo and P. F. Fischer. Terascale spectral element algorithms and implementations. In *Proceedings of the 1999 ACM/IEEE Conference on Supercomputing, SC ’99*, New York, NY, USA, 1999. ACM. doi:10.1145/331532.331599.
- [41] G. van Rossum et al. The Python programming language, 1994. URL <http://python.org>.
- [42] S. Verdoolaege, J. Carlos Juega, A. Cohen, J. Ignacio Gómez, C. Tendlado, and F. Catthoor. Polyhedral parallel code generation for CUDA. *ACM Transactions on Architecture and Code Optimization*, 9(4): 54:1–54:23, Jan. 2013. doi:10.1145/2400682.2400713.
- [43] F. D. Witherden, B. C. Vermeire, and P. E. Vincent. Heterogeneous computing on mixed unstructured grids with PyFR. *Computers & Fluids*, 120:173–186, 2015. doi:10.1016/j.compfluid.2015.07.016.
- [44] Y. Yang, P. Xiang, J. Kong, and H. Zhou. A GPGPU compiler for memory optimization and parallelism management. In *ACM SIGPLAN Notices*, volume 45, page 86–97, 2010. doi:10.1145/1809028.1806606.
- [45] T. Yuki, G. Gupta, D. Kim, T. Pathan, and S. Rajopadhye. Alphaz: A system for design space exploration in the polyhedral model. In *Languages and Compilers for Parallel Computing*, pages 17–31. Springer, 2012.