

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/320609053>

# Hybrid Fortran: High Productivity GPU Porting Framework Applied to Japanese Weather Prediction Model

Article · October 2017

CITATIONS

3

READS

41

2 authors, including:



Michel Müller

Tokyo Institute of Technology

6 PUBLICATIONS 8 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



SPPEXA [View project](#)

# Hybrid Fortran: High Productivity GPU Porting Framework Applied to Japanese Weather Prediction Model

Michel Müller and Takayuki Aoki

Tokyo Institute of Technology

**Abstract.** In this work we use the GPU porting task for the operative Japanese weather prediction model “ASUCA” as an opportunity to examine productivity issues with OpenACC when applied to structured grid problems. We then propose “Hybrid Fortran”, an approach that combines the advantages of directive based methods (no rewrite of existing code necessary) with that of stencil DSLs (memory layout is abstracted). This gives the ability to define multiple parallelizations with different granularities in the same code. Without compromising on performance, this approach enables a major reduction in the code changes required to achieve a hybrid GPU/CPU parallelization - as demonstrated with our ASUCA implementation using Hybrid Fortran.

**Keywords:** HPC, OpenACC, CUDA, GPGPU, OpenMP, Atmospheric, Weather, Parallel Programming, Granularity, Memory Layout

## 1 Introduction

With supercomputing shifting towards accelerators, the increasing need for many-core architecture support in software has created a divide between domain scientists who mainly care about modeling, and the supercomputers their applications are required to run on. The high rate of change in both hardware and software creates maintainability issues, especially with codes that are required to run on varying hardware architectures such as multi-core CPU and GPU. The aforementioned divide has widened especially in atmospheric sciences, a field with constant need for model adaptations and high demand for computing resources. Related publications suggest that productivity issues are what is holding back wider accelerator support in this field.

In this work we use the GPU porting task for a Japanese weather prediction model (“ASUCA”) as an opportunity to examine productivity issues with the current GPGPU standard “OpenACC” when it is applied to structured grid problems written in Fortran. ASUCA is the main mesoscale weather prediction model developed at the Japan Meteorological Agency. It is used in operation, generating nine-hour-forecasts every hour [18] [10].

We then propose Hybrid Fortran, a solution that is designed to increase productivity when re-targeting structured grid Fortran applications to GPU. It is an improvement over OpenACC in two major aspects:

1. Parallelization granularity is abstracted. This allows the user to have multiple granularities defined in the same codebase, depending on the targeted hardware architecture. This is a crucial advantage in order to implement ASUCA’s physical processes on GPU - a code that originally has a very coarse granularity, which is ill-matched for GPUs.
2. Memory layout is abstracted while supporting the already existing user code. More specifically, the layout is reordered at compile-time to match the target architecture, and extended with additional dimensions to match the specified parallelization granularity.

By investigating the necessary code changes with a completed implementation based on Hybrid Fortran we show that this method has enabled high productivity and performance for re-targeting ASUCA to GPU. More than 85% of the hybridized codebase is a one-to-one copy of the original CPU-only code - without counting white-space, code comments and line continuations. An equivalent OpenACC-based solution of ASUCA is estimated to require more than ten thousand additional code lines, or 12.5% of the reference codebase. The new implementation performs up to 4.9x faster when comparing one GPU to one multi-core CPU socket. On a full-scale production run with 1581 x 1301 x 58 grid size and 2km resolution, 24 Tesla P100 GPUs are shown to replace more than 50 18-core Broadwell Xeon sockets.

This paper is structured as follows: Section 1.1 introduces the application our work focuses on. In Sections 1.2 and 1.3 we introduce the main difficulties that we face when porting this application to GPU. Section 1.4 outlines the related work in terms of existing methods to solve these difficulties. We then provide a problem summary in Section 1.5. In Sections 2, 3 and 4 we discuss our solution to this problem, the underlying code transformation method, as well as productivity and performance results achieved with this solution, respectively. Finally, in Section 5 we draw conclusions and point out future work.

## 1.1 ASUCA on GPU

The regional scale weather prediction model “ASUCA” is one of the main operational forecast models in Japan [10]. It is developed by the Japan Meteorological Agency (JMA) and used in production since 2014, covering all of Japan as well as relevant ocean areas and surrounding landmasses in East Asia on a rectangular grid with two kilometer resolution, using the finite volume method for the spatial discretization [21].

ASUCA is implemented in Fortran, with multidimensional arrays stored in modules as its main data structure. It is structured as a dynamical core

interfacing with physical processes, as commonly seen in other weather models such as WRF [14] and COSMO [1].

Parallelization is applied in the horizontal domain (iterated with I and J loop indices). In order to scale to multiple nodes, ASUCA's grid is decomposed into blocks in the horizontal domain that are scattered across nodes, thus requiring halo communication. OpenMP parallel loop directives are employed as an existing intra-node CPU parallelization. Time discretization is implemented by employing a third order Runge-Kutta scheme, enabling long time steps [24]. Sound and gravity waves are treated separately using a second-order Runge-Kutta scheme to enable a higher time resolution, employing the HEVI scheme (Horizontally explicit - vertically implicit) [21].

ASUCA's dynamical core is a stencil code and thus heavily bounded by memory bandwidth [20] [3]. Since the dynamical core also dominates the runtime, GPUs are an attractive target architecture, with a memory bandwidth that is typically 5 to 7 times higher than Intel Xeon architectures of a similar generation. This work is thus motivated by the task of achieving a GPU port for ASUCA, however with the additional goal of optimizing for minimal code changes in order to achieve better acceptance from the application owner (JMA).

ASUCA is a protected asset of the Japanese government and can not be published at this time - in this paper we thus refer to code snippets to discuss our implementation.

## 1.2 Parallelization Granularity

Compared to CPUs, GPUs support a very high number of parallel threads while having a very low thread switching overhead - however with the cost of small caches available per thread and a low single-threaded performance. Furthermore, the latency experienced when off-loading code to GPU results in a very high number of threads being optimal, preferably a multiple of the available arithmetic units (i.e. tens of thousands of threads or more). This in turn leads to register pressure being a major factor for the optimization of GPU code, since the number of utilized registers scales linearly with the number of active threads. These characteristics point to an important distinction in the GPGPU programming model: A fine-grained parallelization is strongly preferred or even necessary (as subroutine calls in kernels are required to be inlined, which becomes impractical with deep call graphs).

While ASUCA's original implementation already offers a fine-grained and thus GPU-friendly dynamical core, most of its physical processes use very large kernels and thus coarse granularity. Since each vertical column ( $K$  index) can be computed independently for many of the physical processes, these computations are implemented in a single kernel in order to increase cache locality and decrease the amount of context switching and thread synchronization [12] [2].

In order to achieve GPU support, a more fine-grained parallelization is thus required for the physical processes. An automated or assisted approach for kernel

fission is desirable in order to allow for GPU acceleration while keeping cache locality and low overhead for the CPU case.

### 1.3 Memory Layout

There are two major aspects in which the memory layout on GPU differs from that on CPU:

**Stride-1 Access** On CPU, the memory layout is generally chosen such that the fastest varying dimension is mapped to stride-1 access. GPUs on the other hand require the first parallel dimension to be mapped to stride-1 in order to coalesce the memory operations, since a group of threads is executed in lockstep. In case of ASUCA these are not the same dimensions: K is the fastest varying dimension, yet it is executed sequentially in general (with a few possible exceptions in the dynamical core), while either I or J can be chosen as the first parallel dimension. The original codebase thus employs KIJ memory order. In an unpublished paper submitted for review, we show that KIJ order leads to a 7.7x slowdown on GPU (versus IJK) while IJK ordering leads to a 35% slowdown on CPU (versus KIJ) [15].

**Privatization** To create a more fine-grained parallelization for GPU, as discussed in Section 1.2, kernel fission is required. This in turn requires thread-local data structures and passed-in data slices to be extended in the parallel domain (IJ). Many scalars thus become 2D-arrays. Many 1D- become 3D-arrays.

Thus, in order to make performance portability possible, the memory layout is required to be changeable and extendable, depending on the target architecture and the parallelization granularity. To achieve this in pure Fortran however, all array specifications and accesses need to be modified and thus specialized for GPU (i.e. code duplication is necessary).

### 1.4 Related Work

The following works are related to this paper in that they describe alternative or similar methods to achieve a hybrid GPU/CPU capable port for atmospheric models. All of these works have achieved compelling speedups on GPU, this discussion therefore focuses on the productivity aspects.

**Stencil Domain-Specific Languages** Domain-specific languages applied to stencil algorithms have been one method to abstract parallelization boiler-plate and memory layout for hybrid GPU/CPU code. For this matter, direct data accesses with loop- or thread indices are abstracted in the point-wise stencil code. This generally requires a complete rewrite of existing code. We take note of the following implementations of weather models using this technique:

- Shimokawabe et al. have completed a research implementation of the ASUCA dynamical core and a portion of its physical processes using their own C++ stencil DSL library [21],
- Fuhrer et al. have implemented the dynamical core of COSMO for operational use at MeteoSwiss, using a purpose-built C++ stencil DSL library for structured grid applications (“STELLA”) [6]. The successor project currently enhancing this method is called “GridTools” [5].
- Jumah et al. have proposed a general grid definition and manipulation language (GGDML), an extension to Fortran with applicability to other languages, based on the requirements for three existing models: DYNAMICO, ICON and NICAM. This work is part of the AIMES project [11].

**Directive-Based Porting Methods** Directives are used to steer compilers on how to optimize or parallelize already existing code for a specific hardware architecture. Privatization of thread-local data is generally supported in directive-based methods, while storage ordering is not. To switch between multiple parallelization granularities, code duplication is necessary. We are aware of the following implementations of atmospheric models with this method:

- Lapillonne et al. have implemented the relevant physical processes of COSMO for operational use at MeteoSwiss using OpenACC in Fortran [13].
- Govett et al. have ported the dynamical core of the Fortran-based Non-hydrostatic Icosahedral Model (NIM) to GPU, first using their own directive-based transformation tool “F2C-ACC” and later using OpenACC (after critical performance issues and bugs were fixed by the compiler vendors). At the time of this writing we are not aware of a GPU implementation of the NIM physical processes however, an issue that has reduced the potential speedup due to the communication overhead caused by running the physical processes on CPU [7] [8].
- Norman et al. have implemented the “Accelerated Model for Climate and Energy” (ACME) for the U.S. DOE using OpenACC. As with ASUCA, ACME’s physical processes are problematic for GPU due to their coarse-grained parallelization. GPU-specific code duplication was the only solution found when using OpenACC [16].

**Granularity Optimization Methods** Kernel fusion has been the main approach to granularity optimization applied to GPGPU programming we are aware of. We take note of the following related work:

- Wahib and Maruyama have shown the effectiveness of this approach in terms of performance when applied to CUDA C kernels [23].
- Gysi and Hoefler have applied the same approach as well as loop fusion for the aforementioned STELLA stencil DSL library [9].

- Clement et al. are applying kernel fusion and an OpenACC/OpenMP targeted transformation to Fortran code in an ongoing effort as part of the C2SM project (as of yet unpublished<sup>1</sup>).

We are not aware of previous or ongoing work regarding kernel fission (and thus support for coarse-grained parallel programming) applied to GPUs.

**Memory Layout Abstraction Methods** While stencil DSLs abstract the memory layout, they also require a full rewrite of the point-wise code. The following work allows for the reuse of existing code while keeping the performance portability gains of an abstracted memory layout:

- With the C++ library “Kokkos” (part of the Trilinos project), Edwards et al. have demonstrated that existing point-wise code can be reused even when the underlying data structures are converted to an abstracted memory layout. OpenMP, Pthreads as well as CUDA are provided as backends to the user code for this library. Granularity optimizations are not supported at the time of this writing, neither is Fortran user code [4].
- With a DSL created for the climate model “ICON”, Torres et al. have shown that the Fortran syntax can be extended to allow for an abstracted memory layout. A code transformation based on the ROSE compiler [17] is employed towards that goal [22]. Sawyer et al. have subsequently built on this abstraction to port the dynamical core of ICON to GPU using OpenACC directives [19]. We are not aware of any granularity optimizations supported in the ICON DSL.

### 1.5 Problem Summary

No existing method, that we are aware of, combines memory layout abstraction and a flexible parallelization granularity with the ability to reuse existing Fortran code for GPGPU. In this work we aim at introducing such a method.

## 2 Hybrid Fortran Language Extension and Code Transformation

Hybrid Fortran has been developed as a method for porting structured grid Fortran applications to GPU. In recognition of the advantages and disadvantages of stencil DSL- and directive-based methods (outlined in Section 1.5) we have combined the advantages of both by employing the following characteristics in our approach:

1. it *does abstract* the parallel loops in order to achieve multiple parallelization granularities with the same code,

---

<sup>1</sup> These open-sourced efforts can be found at <https://github.com/C2SM-RCM/claw-compiler>.

2. it *does not abstract* the point-wise code (i.e. the loop bodies) - allowing for code reuse,
3. it *does separate* the memory layout as defined in the user code from the layout that is effectively implemented for each architecture.

Hybrid Fortran is an open-source framework and can be accessed together with a library of sample applications<sup>2</sup>.

In this Section we discuss how these characteristics have been achieved. Subsection 2.1 describes our approach to parallelization and granularity in Hybrid Fortran, while Subsection 2.2 discusses and compile-time defined memory layout and device memory handling.

## 2.1 Parallel Loop Abstraction

Consider the following kernel from JMA’s ASUCA reference implementation. As part of the dynamical core it is executed within the second-order Runge-Kutta scheme with high time resolution. It applies lateral and upper damping to ASUCA’s grid point values.

**Listing 1.1.** Lateral and upper damping kernel applied to grid point values.

```
!$OMP PARALLEL DO
do j = ny_mn, ny_mx
do i = nx_mn, nx_mx
do k = nz_mn, nz_mx
dens_ptb_damp(k,i,j) = &
& mtratio_bnd * ( dens_ref_f(k,i,j) + dens_ptb_bnd(k,i,j,1) ) &
& + tratio_bnd * ( dens_ref_f(k,i,j) + dens_ptb_bnd(k,i,j,2) ) &
& - dens_ref_f(k,i,j)
end do
end do
end do
!$OMP END PARALLEL DO
```

Using Hybrid Fortran we replace the OpenMP directives, as well as the loop instructions to be parallelized, with our parallelization DSL:

**Listing 1.2.** Lateral and upper damping kernel, modified with Hybrid Fortran.

```
@parallelRegion{ &
& domName(i,j), domSize(nx_mn:nx_mx,ny_mn:ny_mx), &
& startAt(nx_mn,ny_mn), endAt(nx_mx,ny_mx), template(TIGHT_STENCIL) &
& }
do k = nz_mn, nz_mx
dens_ptb_damp(k,i,j) = &
& mtratio_bnd * ( dens_ref_f(k,i,j) + dens_ptb_bnd(k,i,j,1) ) &
& + tratio_bnd * ( dens_ref_f(k,i,j) + dens_ptb_bnd(k,i,j,2) ) &
& - dens_ref_f(k,i,j)
end do
@end parallelRegion
```


We therefore have an explicit distinction between loops that are treated as parallelizable (and are thus restricted in their access patterns, i.e. loop carried dependencies are not supported) and loops that are always executed sequentially.

<sup>2</sup> Please refer to <https://github.com/muellermichel/Hybrid-Fortran>.



The attributes `domName` and `domSize` specify the relevant domain iterators and the domain size relevant to data objects accessed within the parallel region (this relevancy will later be discussed in more detail in Section 2.2). The attributes `startAt` and `endAt` explicitly state the region boundaries, which can be a subset of the domain size (however, if omitted, the domain size is also assumed as the region boundary).

For CPU targets, Hybrid Fortran generates an OpenMP code version very similar to the reference code shown in Listing 1.1<sup>3</sup>, with multi-core parallelization applied to the outermost loop. For GPU targets it defaults to CUDA Fortran kernels (thus generating all the necessary host- and device code boilerplate and data copy operations) with an option to use OpenACC kernels with CUDA compatible data structures (device pointers)<sup>4</sup>. The attribute `template` specifies a macro suffix used for the generated block size parameters - this allows a central configuration for the block sizes used in an application, rather than leaking this architecture-specific optimization to the user code in each kernel. If omitted, configurable default block sizes are used.



**Fig. 1.** Simplified code structure of ASUCA.

The main advantage of this parallelization DSL is the following: replacing parallelizable loops with the `@parallelRegion` construct allows the user to specify multiple granularities in the same code. Consider ASUCA’s code structure, shown in simplified form in Figure 1. It shows two selected kernels and their embedding in the call graph - the lateral and upper boundary damping already discussed

<sup>3</sup> Privatization is the main difference: Hybrid Fortran generated OpenMP code uses “firstprivate” as the default policy with an explicit “shared” clause for all arrays used in the kernel.

<sup>4</sup> OpenACC is mainly used for reduction support - Hybrid Fortran does not automatically generate reduction kernels, however it supports the “reduce” clause, which is forwarded to the generated OpenMP or OpenACC kernels.

in this section, as well as the physics kernel. Many physical processes are called within this single kernel (of which three sample processes are depicted here). This code therefore has a very coarse granularity, which is problematic on GPU as discussed in Section 1.2.

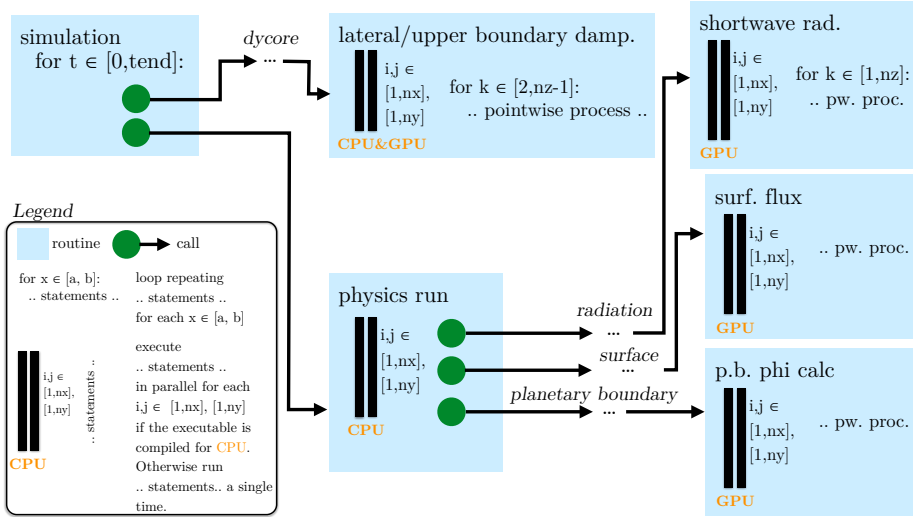


Fig. 2. Simplified code structure of ASUCA using Hybrid Fortran.

With Hybrid Fortran we can solve this problem as follows: An additional `appliesTo` attribute in the `@parallelRegion` statement allows the user to selectively apply parallel regions to either CPU or GPU. Applying the parallelization at different granularities therefore becomes possible<sup>5</sup>, by enabling user-steered kernel fission. Figure 2 shows the resulting code structure, with the physics run being split into many kernels for GPU while remaining a single coarse grained kernel for CPU. The later Listing 1.4 gives an example of how such a kernel fission works in practice.

## 2.2 Compile-time Defined Memory Layout and Device Data Region

As discussed in Section 1.3, it is necessary to consider two major aspects for implementing the memory layout: storage order on one hand and the compile-time defined granularity requiring a varying dimensionalities of data objects on the other. Part of Hybrid Fortran is an additional language extension, the `@domainDependant` construct<sup>6</sup>, as a declarative way for the user to specify

<sup>5</sup> thus obviating the need for code duplication and/or deep inlining of call trees

<sup>6</sup> Please note: While this paper follows American English, the Hybrid Fortran language extension has originally been developed following British English, which becomes

additional required information concerning data objects. This concerns memory layout as well as device memory operations, which will be discussed in this section.

**Storage Order** Revisiting the code sample from Section 2.1, the following Listing shows the specification of the routine implementing the discussed lateral and upper damping kernel:

**Listing 1.3.** Routine implementing the lateral and upper damping kernel with Hybrid Fortran.

```

subroutine lateral_and_upper_damping()
  use ref, only : dens_ref_f
  use svar, only : dens_ptb_damp
  ! ... further imports omitted
  implicit none

  @domainDependant{ &
    & attribute(autoDom, present), &
    & accPP(AT_TIGHT_STENCIL), domPP(DOM_TIGHT_STENCIL) &
  & }
  dens_ref_f, dens_ptb_damp
  @end domainDependant

  @domainDependant{ &
    & attribute(autoDom, present), &
    & accPP(AT4_TIGHT_STENCIL), domPP(DOM4_TIGHT_STENCIL) &
  & }
  dens_ptb_bnd
  @end domainDependant

  ! ... initialisation of tratio_bnd and mratio_bnd omitted
  ! ... kernel omitted (already shown in listing 1.2)
end subroutine

```

This shows the specification of the local module data object `dens_ptb_bnd` (density perturbation in the boundary layer) as well as the external module data objects `dens_ref_f` (reference density) and `dens_ptb_damp` (density perturbation in ASUCA grid).

The `autoDom` attribute is used to delegate the dimensions setup to the data object specification parser (which gathers this information in a separate pass from the source modules, here `ref` and `svar`), rather than having the user specify the dimensions explicitly again in the `@domainDependant` construct. The attributes `accPP` and `domPP` are employed to specify the macro names used to implement the dimension ordering for accesses and specification parts, respectively. These macros wrap all dimension lists in access expressions and specifications of respective data objects in the generated code. When `accPP` and `domPP` attributes are omitted, default macro names are used (for a code example after this conversion please refer to Listing 1.6). In case of Listing 1.3 we use explicit macro names for the dynamical core since the default macros

---

apparent in the spelling of “domainDependant” (<https://en.oxforddictionaries.com/definition/dependant>). We consider support for the American English spelling of this directive as part of our future efforts.

are already used with different assumptions for the physical processes (see the paragraph on “Dimensionality Changes” below).

**Device Data Region** Similar to OpenACC, in Hybrid Fortran we implement data regions by adding state attributes to data objects. The `present` attribute, shown in Listing 1.3, indicates that the respective objects are located on the device in case of GPU compilation. Analogous `transferHere` attributes are used in the main simulation routine in order to instruct Hybrid Fortran to implement the memory copy operations to- and from the device, once at the beginning and end of the simulation. For dummy variables with specified `intent`, Hybrid Fortran will use the Fortran intent information to determine the correct copy operation<sup>7</sup>, which minimizes the potential for bugs in comparison to OpenACC’s explicit `copyIn`, `copyOut` and `copy` clauses. Halo region updates, required for every timestep, are implemented explicitly in code sections guarded from CPU compilation.

**Dimensionality Changes** Due to the compile-time defined parallelization granularity, discussed in Section 2.1, it is necessary to modify the dimensionality of data objects in certain cases in the source generation. This requires hints from the framework user. Consider the following surface flux code snippet:

**Listing 1.4.** Surface flux code snippet.

```
lt = tile_land
if (tlcvr(lt) > 0.0_r_size) then
  call sf_slab_flux_land_run( &
    ! ... inputs and further tile variables omitted
    & tau_x_tile_ex(lt), tau_y_tile_ex(lt) &
    & )

  u_f(lt) = sqrt(sqrt(tau_x_tile_ex(lt) ** 2 + tau_y_tile_ex(lt) ** 2))
else
  tau_x_tile_ex(lt) = 0.0_r_size
  tau_y_tile_ex(lt) = 0.0_r_size
  ! ... further tile variables omitted
end if
! ... sea tiles code and variable summing omitted
```

Since this process is defined inside the call graph of the physics kernel, as shown in Figure 1, the relevant 2D- and 3D grid point values are already sliced and passed in as scalars or 1D-arrays, that is, data parallelism is not exposed at this level. Hybrid Fortran allows implementing this as a fine-grained kernel (as outlined in Figure 2) without modifying the computational user code, as demonstrated in the following snippet:

**Listing 1.5.** Surface flux code snippet with Hybrid Fortran.

```
@domainDependant{domName(i,j), domSize(nx,ny), attribute(autoDom, present)}
tlcvr, tau_x_tile_ex, tau_y_tile_ex, u_f
@end domainDependant
```

<sup>7</sup> Simple examples of this feature can be found in <https://github.com/muellermichel/Hybrid-Fortran/blob/v1.00rc10/examples/demo/source/example.h90>.

```

@parallelRegion{appliesTo(GPU), domName(i,j), domSize(nx,ny)}
lt = tile_land
if (tlcvr(lt) > 0.0_r_size) then
  call sf_slab_flx_land_run( &
    ! ... inputs and further tile variables omitted
    & tau_x_tile_ex(lt), tau_y_tile_ex(lt) &
    & )

  u_f(lt) = sqrt(sqrt(tau_x_tile_ex(lt) ** 2 + tau_y_tile_ex(lt) ** 2))
else
  tau_x_tile_ex(lt) = 0.0_r_size
  tau_y_tile_ex(lt) = 0.0_r_size
  ! ... further tile variables omitted
end if
! ... sea tiles code and variable summing omitted
@end parallelRegion

```

Using our parallelization DSL to provide additional dimensionality information, Hybrid Fortran is able to rewrite this code into a 2D kernel. Dimensions missing from the user code are inserted at the beginning of the dimension lists in access expressions and data object specifications. As an example, the expression `u_f(lt)` is converted to `u_f(AT(i,j,lt))`, employing the default ordering macro already mentioned in the paragraph “Storage Order”. Dimensions are extended whenever there is a match found for `domName` or `domSize` information between data objects and parallel regions within the same routine *or* in routines called within the call graph of the same routine. It is therefore necessary for Hybrid Fortran to gather global information about the application before implementing each routine.

### 2.3 Transformed Code

Revisiting Listing 1.5, the following code is generated when applying Hybrid Fortran with the OpenACC backend:

**Listing 1.6.** Surface flux code snippet after conversion with OpenACC backend.

```

!$acc kernels deviceptr(tau_x_tile_ex) deviceptr(tau_y_tile_ex) &
!$acc& deviceptr(tlcvr) deviceptr(u_f)
!$acc loop independent vector(CUDA_BLOCKSIZE_Y)
outerParallelLoop0: do j=1,ny
!$acc loop independent vector(CUDA_BLOCKSIZE_X)
  do i=1,nx
    ! *** loop body *** :
    lt = tile_land
    if (tlcvr( AT(i,j,lt) ) > 0.0_r_size) then
      call sf_slab_flx_land_run(&
        ! ... inputs and further tile variables omitted
        & tau_x_tile_ex( AT(i,j,lt) ), tau_y_tile_ex( AT(i,j,lt) ) &
        & )
      u_f( AT(i,j,lt) ) = sqrt(sqrt(tau_x_tile_ex( AT(i,j,lt) ) ** 2 + &
        & tau_y_tile_ex( AT(i,j,lt) ) ** 2))
    else
      tau_x_tile_ex( AT(i,j,lt) ) = 0.0_r_size
      tau_y_tile_ex( AT(i,j,lt) ) = 0.0_r_size
      ! ... further tile variables omitted
    end if
    ! ... sea tiles code and variable summing omitted
  end do
end do outerParallelLoop0

```

```
!$acc end kernels
```

Device data is interoperable with the CUDA Fortran backend, thus device pointers are used instead of passing the management to OpenACC. OpenACC directives together with this data type can thus be directly used in the user code as well, i.e. it remains interoperable with device code generated by Hybrid Fortran.

As noted in Section 2.2, storage ordering macros (here `AT()`) are applied to all array access statements. For the thread block setup, the configurable default sizes `CUDA_BLOCKSIZE_X/Y` are used since no template is specified for the parallel region at hand. Parallel region loops (here for indices `i` and `|verb|j`) are set up explicitly to parallelize. Other loops, such as the loop over `k`, use a `!$acc loop seq` directive to explicitly avoid parallelization and give the framework user full expressiveness over the desired granularity.

Applying CUDA Fortran backend to the same user code produces the following host code (here shown together with the routine header and footer):

**Listing 1.7.** Surface flux host code snippet after conversion with CUDA Fortran backend.

```
subroutine hfd_sf_slab_flx_tile_run( &
! ... inputs omitted
& )
use cudafor
type(dim3) :: cugrid, cublock
integer(4) :: cugridSizeX, cugridSizeY, cugridSizeZ, &
& cuerror, cuErrorMemcpy
! ... other imports and specifications omitted

cuerror = cudaFuncSetCacheConfig( &
& hfk0_sf_slab_flx_tile_run, cudaFuncCachePreferL1)
cuerror = cudaGetLastError()
if(cuerror .NE. cudaSuccess) then
! error logging omitted
stop 1
end if
cugridSizeX = ceiling(real(nx) / real(CUDA_BLOCKSIZE_X))
cugridSizeY = ceiling(real(ny) / real(CUDA_BLOCKSIZE_Y))
cugridSizeZ = 1
cugrid = dim3(cugridSizeX, cugridSizeY, cugridSizeZ)
cublock = dim3(CUDA_BLOCKSIZE_X, CUDA_BLOCKSIZE_Y, 1)
call hfk0_sf_slab_flx_tile_run <<< cugrid, cublock >>>( &
! ... inputs and further tile variables omitted
& nx, ny, tile_land, u_f, tlcvr & ! required data objects are
& tau_x_tile_ex, tau_y_tile_ex & ! automatically passed to kernel
& )
cuerror = cudaThreadSynchronize()
cuerror = cudaGetLastError()
if(cuerror .NE. cudaSuccess) then
! error logging omitted
stop 1
end if
end subroutine
```

The prefix `hfd_` is added to host routines that use device data. Hybrid Fortran also duplicates the code for a pure host version of these routines (without a name change in order to remain interoperable with code that is not passed through

Hybrid Fortran). In contexts where the data is not residing on the device, such as the setup part of an application, Hybrid Fortran automatically chooses the host version when generating the call statements at compile-time. Code residing within parallel regions is moved within a separated kernel routine (using prefix `hfki_` with `i` representing the kernel number). In case of the surface flux sample shown here, the kernel routine is generated as follows:

**Listing 1.8.** Surface flux device code snippet after conversion with CUDA Fortran backend.

```

attributes(global) subroutine hfk0_sf_slab_flux_tile_run(&
! ... inputs and further tile variables omitted
&, nx, ny, tile_land, u_f, tlcvr &
&, tau_x_tile_ex, tau_y_tile_ex &
& )
use cudafor
use pp_vardef ! defines r_size
implicit none
real(r_size), device :: u_f(DOM(nx,ny,ntlm))
real(r_size), device :: tlcvr(DOM(nx,ny,ntlm))
real(r_size), device :: tau_x_tile_ex(DOM(nx,ny,ntlm))
real(r_size), device :: tau_y_tile_ex(DOM(nx,ny,ntlm))
integer(4), value :: lt
integer(4), value :: nx
integer(4), value :: ny
integer(4), value :: tile_land
! ... other imports and specifications omitted

i = (blockidx%x - 1) * blockDim%x + threadidx%x + 1 - 1
j = (blockidx%y - 1) * blockDim%y + threadidx%y + 1 - 1
if (i .GT. nx .OR. j .GT. ny) then
return
end if
! *** loop body *** :
lt = tile_land
if (tlcvr( AT(i,j,lt) ) > 0.0_r_size) then
call hfd_sf_slab_flux_land_run( &
! ... inputs and further tile variables omitted
tau_x_tile_ex( AT(i,j,lt) ), tau_y_tile_ex( AT(i,j,lt) ) &
& )
! ... rest of loop body already shown in listing 1.6
end subroutine

```

The specification part of these kernels is automatically generated, applying device state information and converting input scalars to pass-by-value<sup>8</sup>, among other transformations.

It is notable that CUDA Fortran requires a fairly large amount of boiler plate code for grid setup, iterator setup, host- and device code separation as well as memory- and error handling - Hybrid Fortran allows the user to pass on the responsibility for that to the framework. Compared with the code generated by OpenACC however (assembly-like CUDA C or NVVM intermediate representation), the Hybrid Fortran generated CUDA Fortran code remains easily readable to programmers experienced with CUDA. Experience shows that this

<sup>8</sup> reduction kernels are thus not supported with this backend - we use the OpenACC backend selectively for this purpose, see also the discussion in the footnotes to Section 2.1.

is a productivity boost, especially in the debugging and manual performance optimization phase of a project.

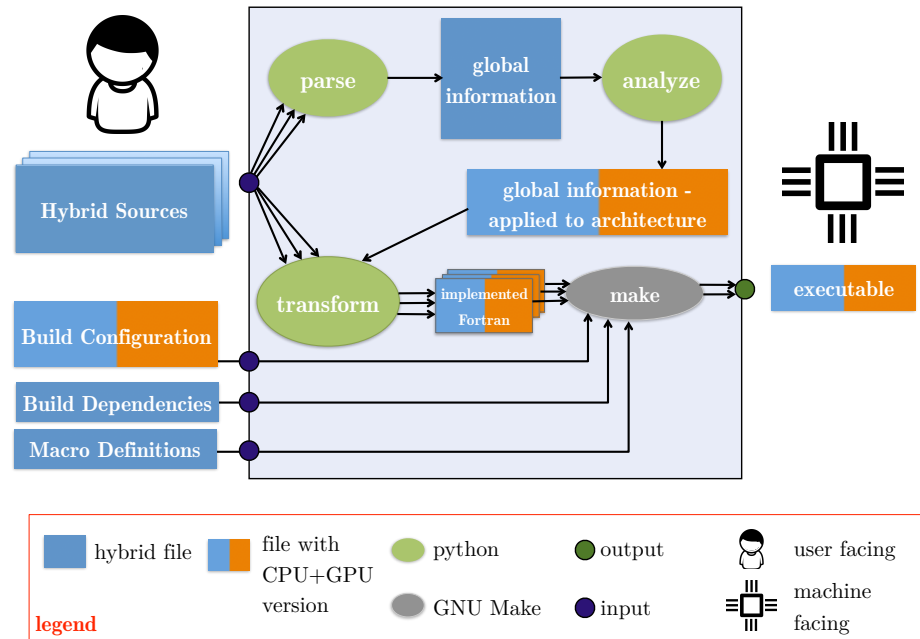
Regarding the OpenMP backend, since the surface flux example is parallelized at a much more coarse-grained level for CPU, the generated CPU code for the sample at hand is a one-to-one copy of the user code shown earlier in Listing 1.4. The parallelization is generated at a higher level in the call graph (replacing the parallel region construct with `appliesTo(CPU)` clause) as follows:

**Listing 1.9.** Surface flux device code snippet after conversion with CUDA Fortran backend.

```

!$OMP PARALLEL DO DEFAULT(firstprivate)
!$OMP& SHARED( ... inputs and outputs omitted ... )
outerParallelLoop0: do j=1,ny
  do i=1,nx
    call physics_main(i, j, &
      ! ... inputs and outputs omitted
      & )
  end do
end do outerParallelLoop0
!$OMP END PARALLEL DO
    
```

### 3 Code Transformation Method



**Fig. 3.** Hybrid Fortran software components and build workflow.



In this section we discuss code transformation method involved in implementing Hybrid Fortran’s characteristics described earlier. This process is applied transparently for the user, i.e. it is applied automatically by the means of a provided common Makefile<sup>9</sup>. Figure 3 gives an overview of the process and the components involved. We discuss this process in order of execution - each of the following enumerated items corresponds to one transformation phase:

1. To simplify the parsing in subsequent phases, Fortran continuation lines are merged.
2. Facilitating later phases, the application’s call graph and parallel region directives are parsed globally (“parse” phase in Figure 3).
3. Using the `appliesTo` information in kernels and the call graph, the position of each routine in relation to kernels is computed. Possible positions are “has kernel(s) in its caller graph”, “contains a kernel itself” and “is called inside a kernel” (“analyze” phase in Figure 3).
4. In two passes, module data object specifications are parsed and then linked against all routines with imports of such objects, together with the locally defined objects.
5. A global application model is generated, with model classes representing the modules, routines and code regions. This model can be regarded as a target hardware independent intermediate representation.
6. Each routine object is assigned an implementation class depending on the target architecture<sup>10</sup>. For each coding pattern, a separate class method of the implementation class is called by the model objects - e.g. CUDA parallelization boilerplate is generated. Using the previously gathered global kernel positioning and data object dimension information, data objects are transformed according to the behavior discussed in Section 2.2. Implementation class methods return strings that are concatenated by the model objects into source files (“transform” phase in Figure 3).
7. Code lines are split using Fortran line continuations in order to adhere to line limits imposed by Fortran compilers.
8. Macros generated by Hybrid Fortran (to implement storage reordering and configurable block sizes) are processed by using the GNU compiler toolchain. Subsequently, a user specified compiler and linker is employed in order to create the CPU and GPU executables. A common makefile is provided with the framework, however the build dependency graph is user-provided in the format of makefile rules<sup>11</sup> (“make” phase in Figure 3).

This process makes it possible to have a unified source input and create executables targeted for either multi-core CPU or many-core GPU.

<sup>9</sup> See also the “Getting Started” section in <https://github.com/muellermichel/Hybrid-Fortran/blob/v1.00rc10/doc/Documentation.pdf>.

<sup>10</sup> Hybrid Fortran allows the user to switch between varying backend implementations per routine, such as OpenACC and CUDA Fortran - the user specified information as well as the defaults given by the build system call thus steers this implementation class.

<sup>11</sup> alternatively, a dependency generator script can be configured as well

## 4 Productivity- and Performance Results

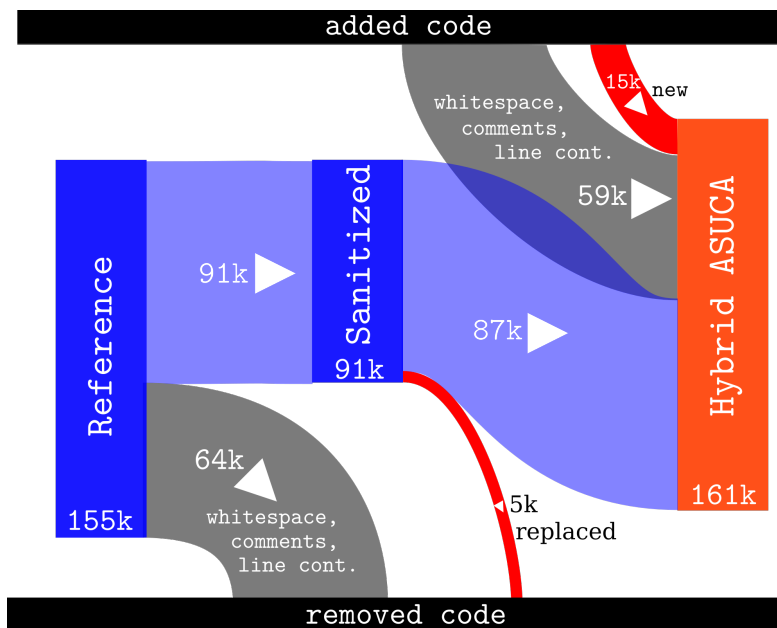
In this section we discuss the productivity- and performance results that have been achieved when applying Hybrid Fortran to weather prediction models.

At the time of this writing an additional paper has been submitted for review, in which the following results are discussed: Hybrid Fortran has been applied to both dynamical core and physical processes of ASUCA. This implementation (here referred to as “Hybrid ASUCA”) consists of 338 kernels and covers almost all modules required for an operative weather prediction, with convection being the main exception due to development time constraints. Communication between host and device has been eliminated with the exception of setup, output and halo exchange. On a 301 x 301 x 58 grid with real weather data, a 4.9x speedup has been achieved when comparing the Hybrid Fortran port on four Tesla K20x versus the JMA provided reference implementation on four 6-core Westmere Xeon X5670 (TSUBAME 2.5). The same setup executes at a speedup of 3x when comparing a single Tesla P100 versus a single 18-core Broadwell Xeon E5-2695 v4 (Reedbush-H). On a full-scale production run with 1581 x 1301 x 58 grid size and 2km resolution, 24 Tesla P100 GPUs are shown to replace more than 50 18-core Broadwell Xeon sockets [15].


In order to further examine the performance of Hybrid Fortran kernels and compare them to OpenACC and OpenMP user codes, a performance model has been constructed for a reduced weather application. The Hybrid Fortran implementation with unified code performs on par or better than OpenACC- and OpenMP- implementations separately optimized for GPU and CPU, respectively. In this application, whose runtime is dominated by the memory bandwidth used for a seven point stencil diffusion kernel, the Hybrid Fortran GPU version performs 25% better than the no-cache-model on Tesla K20x. It achieves 36% of the theoretical limit given by a perfect-cache-model. The CPU version performs 56% better than the no-cache-model and achieves 67% of the theoretical limit on Westmere Xeon X5670 [15].

To examine the productivity of our solution we have analyzed the code and compare it against the reference implementation<sup>12</sup>. The high-level results of this analysis is shown in Figure 4. In order to gain GPU support in addition to the already existing multi-core and multi-node parallelization, the code has grown by less than 4% in total, from 155k lines of code to 161k. Sanitizing the two code versions (removing white space, comments and merging continued lines), the code has grown by 12%, from 91k to 102k lines of code. 95% of the sanitized reference code is used as-is in the new implementation, while 5% or approximately 5k lines of code is replaced with approximately 15k new code lines.

<sup>12</sup> Since the input to this analysis is the closed source ASUCA codebase, full reproducibility cannot be provided in this context. However the intermediate data, the method employed to gather this data as well as a sample input is provided and documented in [https://github.com/muellermichel/hybrid-asuca-productivity-evidence/blob/master/asuca\\_productivity.xlsx](https://github.com/muellermichel/hybrid-asuca-productivity-evidence/blob/master/asuca_productivity.xlsx).



**Fig. 4.** Flow of code lines from reference implementation to Hybrid ASUCA by number of lines of code.



**Fig. 5.** New code required for Hybrid ASUCA vs. estimate of equivalent OpenACC implementation (LOC stands for “lines of code”).

Code changes and additions have the largest impact in terms of productivity. We have analyzed the additional 15k lines of code in more detail. Figure 5 shows a breakdown of these changes and compares them to an estimate of what would be required with an OpenACC-based implementation. The following methodology has been used for this analysis:

1. for the parallelization- and data layout DSL line count we have used information parsed for Hybrid ASUCA, as well as the OpenACC backend available in Hybrid Fortran, to acquire an accurate count for the directives required for an OpenACC-based implementation,
2. since OpenACC does not offer a granularity abstraction, we have used the Hybrid ASUCA’s parsed global application information to arrive at a set of routines that require a kernel positioning change (see the discussion in Section 3) - the resulting code lines require duplication for the CPU in an equivalent OpenACC implementation (shown as “CPU-only physics” in figure 5),
3. we have used the OpenACC backend in Hybrid Fortran to count the lines of code where storage order macros are introduced in order to achieve a compile-time defined data layout.

“Parallelization & data layout DSL” refers to the number of code lines for `@parallelRegion` and `@domainDependant` directives in case of Hybrid Fortran, and `!$acc` directives in case of OpenACC. Hybrid Fortran replaces the requirement for code changes to implement a varying data layout (“storage order macros”, 6098 lines of code, “LOC”) as well as a code duplication for multiple parallelizations with varying granularities (“CPU-only physics”, 7122 LOC) with a higher number of DSL code lines compared to OpenACC (4398 vs. 2521 LOC). “Modified data specifications / initializations” (3519 LOC) as well as “routine & call signature” (1381 LOC) refers to changes applied to the setup of data and call parameter lists, respectively. These changes are necessary due to device code limitations and optimizations and are largely required for both the Hybrid Fortran version as well as a potential OpenACC solution, so we use the result from Hybrid ASUCA as an estimate for what would be required with OpenACC. Finally, one physics module concerning long-wave radiation (2059 LOC), has been replaced with a version that uses less local memory per thread (a factor of 10 improvement in that regard) to make it more GPU-friendly. We again estimate that an OpenACC version would have approximately the same code size.

This result shows that an equivalent OpenACC implementation of ASUCA can be estimated to require approximately 11k LOC in additional changes compared to the Hybrid Fortran-based implementation. When comparing to the sanitized reference codebase, an OpenACC user code would require approximately 28% of code lines to be changed or added, while Hybrid Fortran currently requires 16%.

In addition to the results regarding ASUCA presented here, a library of small applications and kernel samples has been created to demonstrate Hybrid Fortran general applicability to data parallel code<sup>13</sup>.

## 5 Conclusion and Future Work

With this work we have shown that it is possible for large structured grid Fortran applications to

1. achieve a GPU implementation without rewriting major parts of the computational code,
2. abstract the memory layout and
3. allow for multiple parallelization granularities.

With our proposed method, a regional scale weather prediction model of significant importance to Japan’s national weather service has been ported to GPU, showing a speedup of up to 4.9x on single GPU compared to a single Xeon socket. When scaling up to 24 Tesla P100, less than half the number of GPUs is required compared to contemporary Xeon CPU sockets to achieve the same result. Approximately 95% of the existing codebase has been reused for this implementation and our implementation has grown by less than 4% in total, even though it is now supporting GPU as well as CPU. Through a library of results, the general applicability of our framework to data parallel Fortran code has been shown.

Considering that much of the changes still required in the user code are “mechanical” in nature, we expect additional productivity gains to be possible from further automation. We strive to achieve a solution where a Hybrid Fortran-based transformation can be applied to large structured grid applications wholesale with minimal input required by the user.

**Acknowledgments** This work has been supported by the Japan Science and Technology Agency (JST) Core Research of Evolutional Science and Technology (CREST) research program “Highly Productive, High Performance Application Frameworks for Post Peta-scale Computing”, by KAKENHI Grant-in-Aid for Scientific Research (S) 26220002 from the Ministry of Education, Culture, Sports, Science and Technology (MEXT) of Japan, by “Joint Usage/Research Center” for Interdisciplinary Large-scale Information Infrastructures (JHPCN) and “High Performance Computing Infrastructure (HPCI)” as well as by the “Advanced Computation and I/O Methods for Earth-System Simulations” (AIMES) project running under the German-Japanese priority program “Software for Exascale Computing” (SPPEXA). The authors thank the Japan Meteorological Agency for their extensive support, Tokyo University and the Global Scientific Information and Computing Center at Tokyo Institute of Technology for the use of their supercomputers Reedbush-H and TSUBAME 2.5.

<sup>13</sup> Please refer to <https://github.com/muellermichel/Hybrid-Fortran/blob/v1.00rc10/examples/Overview.md> for an overview of the available samples and their results.

## References

1. Cumming, B., Osuna, C., Gysi, T., Bianco, M., Lapillonne, X., Fuhrer, O., Schulthess, T.C.: A review of the challenges and results of refactoring the community climate code COSMO for hybrid Cray HPC systems. *Proceedings of Cray User Group* (2013)
2. Douglas, C.C., Hu, J., Kowarschik, M., Rde, U., Wei, C.: Cache optimization for structured and unstructured grid multigrid. *Electronic Transactions on Numerical Analysis* 10, 21–40 (2000)
3. Dursun, H., Nomura, K.i., Wang, W., Kunaseth, M., Peng, L., Seymour, R., Kalia, R.K., Nakano, A., Vashishta, P.: In-core optimization of high-order stencil computations. In: *PDPTA*. pp. 533–538 (2009)
4. Edwards, H.C., Trott, C.R., Sunderland, D.: Kokkos: Enabling manycore performance portability through polymorphic memory access patterns. *Journal of Parallel and Distributed Computing* 74(12), 3202 – 3216 (2014), Domain-specific languages and high-level frameworks for high-performance computing
5. Fuhrer, O.: Grid Tools: Towards a library for hardware oblivious implementation of stencil based codes. Retrieved July 13, 2017 from <http://www.pasc-ch.org/projects/2013-2016/grid-tools> (2014)
6. Fuhrer, O., Osuna, C., Lapillonne, X., Gysi, T., Cumming, B., Bianco, M., Arteaga, A., Schulthess, T.C.: Towards a performance portable, architecture agnostic implementation strategy for weather and climate models. *Supercomputing frontiers and innovations* 1(1), 45–62 (2014)
7. Govett, M., Middlecoff, J., Henderson, T.: Directive-based parallelization of the NIM weather model for GPUs. In: *Accelerator Programming using Directives (WACCPD), 2014 First Workshop on*. pp. 55–61. IEEE (2014)
8. Govett, M., Rosinski, J., Middlecoff, J., Henderson, T., Lee, J., MacDonald, A., Wang, N., Madden, P., Schramm, J., Duarte, A.: Parallelization and performance of the NIM weather model on CPU, GPU and MIC processors. *Bulletin of the American Meteorological Society* (2017)
9. Gysi, T., Hoefler, T.: Integrating STELLA & MODESTO: Definition and optimization of complex stencil programs (2017)
10. Ishida, J., Muroi, C., Kawano, K., Kitamura, Y.: Development of a new nonhydrostatic model ASUCA at JMA. *CAS/JSC WGNE Research Activities in Atmospheric and Oceanic Modelling* 40, 0511–0512 (2010)
11. Jumah, N., Kunkel, J., Zngl, G., Yashiro, H., Dubos, T., Meurdesoif, Y.: GGDML: Icosahedral models language extensions (2017)
12. Kwiatkowski, J.: Evaluation of parallel programs by measurement of its granularity. In: *International Conference on Parallel Processing and Applied Mathematics*. pp. 145–153. Springer (2001)
13. Lapillonne, X., Fuhrer, O.: Using compiler directives to port large scientific applications to GPUs: An example from atmospheric science. *Parallel Processing Letters* 24(01) (2014)
14. Mielikainen, J., Huang, B., Huang, A.: Using Intel Xeon Phi to accelerate the WRF TEMF planetary boundary layer scheme. In: *SPIE Sensing Technology+ Applications*. pp. 91240T–91240T. International Society for Optics and Photonics (2014)
15. Mller, M., Aoki, T.: New high performance GPGPU code transformation framework applied to large production weather prediction code (2017), unpublished paper submitted for review

16. Norman, M.R., Mametjanov, A., Taylor, M.: Exascale programming approaches for the accelerated model for climate and energy (2017)
17. Quinlan, D.: ROSE: Compiler support for object-oriented frameworks. *Parallel Processing Letters* 10(02n03), 215–226 (2000)
18. Sakamoto, M., Ishida, J., Kawano, K., Matsubayashi, K., Aranami, K., Hara, T., Kusabiraki, H., Muroi, C., Kitamura, Y.: Development of yin-yang grid global model using a new dynamical core ASUCA (2014)
19. Sawyer, W., Zaengl, G., Linardakis, L.: Towards a multi-node openacc implementation of the icon model. In: *EGU General Assembly Conference Abstracts*. vol. 16 (2014)
20. Shimokawabe, T., Aoki, T., Muroi, C., Ishida, J., Kawano, K., Endo, T., Nukada, A., Maruyama, N., Matsuoka, S.: An 80-fold speedup, 15.0 TFlops full GPU acceleration of non-hydrostatic weather model ASUCA production code. In: *Proceedings of the 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*. pp. 1–11. IEEE Computer Society (2010)
21. Shimokawabe, T., Aoki, T., Onodera, N.: High-productivity framework on GPU-rich supercomputers for operational weather prediction code ASUCA. In: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. pp. 251–261. SC '14, IEEE Press, Piscataway, NJ, USA (2014), <http://dx.doi.org/10.1109/SC.2014.26>
22. Torres, R., Linardakis, L., Kunkel, J., Ludwig, T.: ICON DSL: A domain-specific language for climate modeling. In: *International Conference for High Performance Computing, Networking, Storage and Analysis*, Denver, CO (2013)
23. Wahib, M., Maruyama, N.: Scalable kernel fusion for memory-bound GPU applications. In: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. pp. 191–202. SC '14, IEEE Press, Piscataway, NJ, USA (2014), <http://dx.doi.org/10.1109/SC.2014.21>
24. Wicker, L.J., Skamarock, W.C.: Time-splitting methods for elastic models using forward time schemes. *Monthly weather review* 130(8), 2088–2097 (2002)