

# Julia as a Unifying End-to-End Workflow Language on the Frontier Exascale System

William F. Godoy, Pedro Valero-Lara, Caira Anderson, Katrina W. Lee, Ana Gainaru, Rafael Ferreira da Silva, and Jeffrey S. Vetter

Oak Ridge National Laboratory  
Oak Ridge, TN, USA

{godoywf,valerolarap,andersonci,leekw,gainarua,silvarf,vetter}@ornl.gov

## ABSTRACT

We evaluate Julia as a single language and ecosystem paradigm powered by LLVM to develop workflow components for high-performance computing. We run a Gray-Scott, 2-variable diffusion-reaction application using a memory-bound, 7-point stencil kernel on Frontier, the US Department of Energy’s first exascale supercomputer. We evaluate the performance, scaling, and trade-offs of (i) the computational kernel on AMD’s MI250x GPUs, (ii) weak scaling up to 4,096 MPI processes/GPUs or 512 nodes, (iii) parallel I/O writes using the ADIOS2 library bindings, and (iv) Jupyter Notebooks for interactive analysis. Results suggest that although Julia generates a reasonable LLVM-IR, a nearly 50% performance difference exists vs. native AMD HIP stencil codes when running on the GPUs. As expected, we observed near-zero overhead when using MPI and parallel I/O bindings for system-wide installed implementations. Consequently, Julia emerges as a compelling high-performance and high-productivity workflow composition language, as measured on the fastest supercomputer in the world.

## KEYWORDS

Julia, end-to-end workflows, High-Performance Computing, HPC, data analysis, exascale, Frontier supercomputer, Jupyter notebooks

---

This manuscript has been authored by UT-Battelle LLC under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The publisher acknowledges the US government license to provide public access under the DOE Public Access Plan (<https://energy.gov/downloads/doe-public-access-plan>).

---

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of the United States government. As such, the Government retains a non-exclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

SC-W 2023, November 12–17, 2023, Denver, CO, USA

## ACM Reference Format:

William F. Godoy, Pedro Valero-Lara, Caira Anderson, Katrina W. Lee, Ana Gainaru, Rafael Ferreira da Silva, and Jeffrey S. Vetter. 2023. Julia as a Unifying End-to-End Workflow Language on the Frontier Exascale System. In *Workshops of The International Conference on High Performance Computing, Network, Storage, and Analysis (SC-W 2023)*, November 12–17, 2023, Denver, CO, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3624062.3624278>

## 1 INTRODUCTION

Recent emphasis on the end-to-end workflow development process for high-performance computing (HPC) applications acknowledges the increasing complexity required for achieving performance, portability, and productivity [4]. This complexity is primarily driven by two factors: (i) the evolving application requirements for experimental, observational, and computational science and (ii) the extreme heterogeneity of our computing and data generation and processing systems [11, 16, 43]. Consequently, the community vision for a research and development roadmap [14, 15] has identified key challenges posed by integrating HPC, AI, and FAIR (findable, accessible, interoperable, and reusable) [18, 41] workflows at exascale.

The Julia programming language [5] was designed to provide a powerful unifying strategy to close the gaps between scientific computing and data science. Julia unifies aspects of the scientific workflow development process: simulation, communication, visualization, parallel data I/O, AI, and interactive computing. This is achieved by providing (i) a dynamic just-in-time (JIT) compiled front end to LLVM [27], (ii) a lightweight interoperability layer for existing C and Fortran HPC codes, and (iii) a unified community ecosystem (e.g., packaging and testing). The status quo in HPC software is to write simulation code by using a compiled language (Fortran, C, or C++) and the Message Passing Interface (MPI) [22] alongside on-node standard (OpenMP [32]), vendor (CUDA [31], HIP [1]) or third-party (Kokkos [7], RAJA [3]) parallel programming models, also known as MPI+X, while data analysis and workflow orchestration tasks are programmed by using high-level interfaces and languages (e.g., Python)—not without caveats and trade-offs [42]. As

a result, Julia provides a valuable alternative in the convergence of high-productivity and high-performance that must be tested on exascale hardware.

In this work, we measure and analyze the computational performance aspects of a Gray-Scott diffusion-reaction HPC workflow application [33] written in Julia and running on Frontier, the recently deployed exascale system at Oak Ridge National Laboratory.<sup>1</sup> We focus on the computational trade-offs of solving a representative application of a 7-point stencil kernel on Frontier’s AMD’s MI250x GPUs at different scales. Simultaneously, we focus on the production and consumption of data in Jupyter Notebooks through the parallel file system. More specifically, we present the following:

- Compute measurements on AMD MI250x GPUs and comparison with a similar HIP 7-point stencil kernel
- Weak scaling on up to 4,092 GPU (or 512 nodes) using MPI showing the variability on local process time-to-solution
- Parallel I/O scaling results using the ADIOS-2 Julia bindings
- The overall feasibility of Julia as a single language and ecosystem for all above components and its use on JupyterHub for data analysis of the generated data.

The paper is organized as follows: Section 2 presents related work in the use of Julia in HPC environments. Section 3 describes the components of the tested workflow, including parallel simulations and data analysis. Section 4 describes the Julia implementation as a single unifying paradigm for GPU, MPI, parallel I/O, and data analysis using Frontier’s compute nodes and the Lustre file system with the ADIOS2 [20] library. Section 5 shows performance and weak scaling results when using up to 4,092 GPUs or 512 nodes of the Frontier supercomputer, including achieved GPU bandwidths, weak scaling, and write performance for large simulations. Conclusions and future directions are presented in Section 6. Appendix A presents the Artifact Description for the reproducibility of this study.

## 2 RELATED WORK

Recent work has reported the use of Julia in different aspects of HPC, and its use is still an area of active exploration and community engagement [8].

On the compute side, Lin and McIntosh-Smith [28] compared memory-bound (BabelStream) and compute-bound (miniBUDE) benchmark mini-app ports in Julia against several implementations on different CPU and GPU vendor systems. They concluded that Julia’s on-node performance is either on par or slightly behind the vendor compiled implementations for their use cases. Godoy

et al. [21] showed that Julia is competitive for a simple matrix multiplication kernel on modern exascale nodes when compared alongside the high-level Python’s Numba [26] JIT layer running on top of LLVM and the C++ Kokkos programming model against vendor OpenMP on CPU and vendor CUDA and HIP implementations on GPU. Faingnaert et al. [13] provide optimized GEMM kernels in Julia that are competitive with cuBLAS and CUTLASS implementations. Ranocha et al. [36] describe similar levels of performance for the `Trixi.jl` partial differential equation (PDE) solver at scale when the Julia implementation is compared against compiled HPC languages.

On the network communication side, Giordano et al. [17] measured comparable system-level scalability performance for Julia’s `MPI.jl` [6] bindings on the Arm-based Fujitsu A64FX Fugaku system running point-to-point communication benchmarks for different floating point representations. Hunold et al. [24] used a Julia port of the STREAM benchmark [30] to show negligible overheads in the collective `Bcast`, `Alltoall`, and `Allreduce` MPI operations when compared against equivalent C implementations on up to 1,152 processes. Shang et al. [40] used the Julia language ecosystem on the many-cores Sunway supercomputer to conduct quantum computational chemistry simulations. In their experiments they achieved up to 91% efficiency when measuring weak scaling on up to 21M cores. Regier et al. [37] provided Celeste, an early application of Julia on multi-threaded CPU HPC systems. Celeste achieved 1.54 petaflops using 1.3 million threads on 9,300 Knights Landing (KNL) nodes of the Cori supercomputer, while exploiting burst buffer for data I/O. More recently, Lu et al. [29] explored the integration of Julia in their framework for moving computation as data in binary or LLVM bit code representations. They found that their implementation can leverage JIT optimizations, but further exploration is needed.

To the best of our knowledge, our work represents the pioneering effort to harness the potential of Julia in modern HPC end-to-end workflows targeting a GPU-powered exascale system.


## 3 WORKFLOW DESCRIPTION

This section describes the characteristics of the computation, communication, parallel data I/O, and data analysis of the representative Gray-Scott diffusion-reaction simulation code running on Frontier.

Figure 1 shows a schematic representation of a typical parallel simulation and data analysis workflow components that apply to the current Gray-Scott use case. Julia’s value proposition lies in its capacity to utilize a single language, along with an integrated ecosystem, to write these components. This unified approach allows the same language to not only operate at scale but also to conduct data analysis in an interactive environment, thereby fostering a seamless

<sup>1</sup><https://www.olcf.ornl.gov/frontier>

connection between computational execution and analytical exploration.



**Figure 1: Gray-Scott workflow components schematic illustrating the unifying nature of Julia.**

### 3.1 Simulation

Gray-Scott is a two-variable, diffusion-reaction 3D model described by the PDEs shown in Equations (1a) and (1b):

$$\frac{\partial U}{\partial t} = D_U \nabla^2 U - UV^2 + F(1 - U) + nr \quad (1a)$$

$$\frac{\partial V}{\partial t} = D_V \nabla^2 V + UV^2 - (F + k)V \quad (1b)$$

where  $U$  and  $V$  are the output concentrations of two reacting and diffusing chemicals, while the inputs are listed as follows:

- $D_u$  and  $D_v$  are the diffusion rates for  $U$  and  $V$ ;
- $F$  is the feed rate of  $U$  into the system;
- $k$  is the kill rate of  $V$  from the system;
- $n$  is the magnitude of the noise to be added to the system; and
- $r$  is a uniformly distributed random number between  $-1$  and  $1$  for each time and spatial coordinate.

As illustrated in Equations (2a) and (2b), the set of governing equations is discretized in time,  $t$ , and space,  $i, j, k$ , on a regular normalized mesh by using simple forward and central differences, respectively:


$$U_{i,j,k}^{t+1} = U_{i,j,k}^t + \Delta t \left[ D_U \nabla^2 U_{i,j,k}^t + S_U^t \right] \quad (2a)$$

$$V_{i,j,k}^{t+1} = V_{i,j,k}^t + \Delta t \left[ D_V \nabla^2 V_{i,j,k}^t + S_V^t \right] \quad (2b)$$

where  $\Delta t$  is an input time step variable,  $S$  represents the local source terms for  $U$  and  $V$  defined in Equations (1a) and (1b), and the Laplacian operator,  $\nabla^2$ , is defined in Equation (3) for the 3D nearest-neighbor Jacobi 7-point stencil in normalized spatial units:

$$\nabla^2 U_{i,j,k}^t = -U_{i,j,k}^t + \frac{1}{6} \left[ U_{i-1,j,k}^t + U_{i+1,j,k}^t + U_{i,j-1,k}^t + U_{i,j+1,k}^t + U_{i,j,k-1}^t + U_{i,j,k+1}^t \right] \quad (3)$$

The random nature of the source term caused by  $r$  and the coupled nature of  $U$  and  $V$  provide slightly more complexity than a typical diffusion calculation (e.g., heat transfer equation), as shown in Figure 2 for the visualization of the solution of the field variable  $U$  at a center plane over a few time iterations.



**Figure 2: Gray-Scott snapshot of the solution of the 3D field  $V$  variable showing the random nature of the source terms in the PDE.**

### 3.2 On-Node Parallelization

The stencil calculation is the main computational kernel in Gray-Scott. Stencil calculations are an important class of PDE solvers on structured grids [10], and they have been successfully optimized for different hardware architectures, including GPUs [19, 23, 25].

Because each new time step value for  $U$  and  $V$  only depends on neighboring values, the problem is parallelizable on GPUs—and CPUs—thereby providing an extra memory allocation for temporary  $U_{t+1}$  and  $V_{t+1}$  field values for the next

time iteration. Figure 3 illustrates the underlying contiguous memory allocated for each variable, which is accessed from near and far away points in memory; hence, some cache misses are expected in this typical kernel. Each variable requires 7 reads (fetch) + 1 write = 8 floating point access operations per variable for a total of 16 per cell in solving Equations (2a) and (2b).

### 3.3 MPI Communication

The 3D regular domain shown in Figure 2 is decomposed using an MPI Cartesian communicator. Each subdomain source shares face cell nodes with their neighboring subdomain destination at every computational iteration. This is illustrated in Figure 4, which shows a typical MPI\_Send/MPI\_Recv pattern with *ghost cell* surfaces. Because the exchanged surfaces are not memory contiguous, we define a new strided vector type by using MPI\_Datatypes and MPI\_Type\_vector [44]. We keep the MPI exchange from CPU-allocated memory to populate the strided vector contents coming from the GPU. We did not experiment with GPU-aware MPI for this work.

### 3.4 Parallel I/O and Data Analysis

Data production and consumption is performed to enable post-processing (e.g., visualization) of the solution of the  $U$  and  $V$  field variables over time. As such, a simple visualization schema and the corresponding metadata accounting for the input provenance are stored in the resulting dataset. Due to the volume of data produced from the simulation, drastically reducing the frequency of writes to the parallel file system is often required. Listing 1 presents a sample provenance record of the data produced from the Gray-Scott simulation.

**Listing 1: Provenance of the data generated from the Gray-Scott simulation.**

```
double Du attr = 0.2
double Dv attr = 0.1
double F attr = 0.02
double U 1000*{1024, 1024, 1024}
Min/Max -0.120795 / 1.46671
double V 1000*{1024, 1024, 1024}
Min/Max 0 / 0.959875
double dt attr = 1
double k attr = 0.048
double noise attr = 0.1
int32_t step 50*scalar = 20 / 1000
Attribute visualization schemas: FIDES, VTX
```

## 4 IMPLEMENTATION

We developed GrayScott.jl as a comprehensive HPC workflow application packaged as a Julia module. The resulting Julia implementation using the AMDGPU.jl [38] package is shown in Listing 2. This implementation is a straightforward

application of a typical GPU kernel and maps each component of the kernel onto a 3D computational grid. When launching this kernel and allocating the number of threads and grid size, one must consider that Julia arrays are column-major. Therefore, the fastest index, being the first one, should be structured to avoid splitting across threads on the GPU.

Although the goal is not to increase the efficiency of the present stencil calculation, we aim to evaluate how the performance of the JIT generated Julia LLVM-IR stencil code compares with theoretical bandwidths and a simple AMD-provided Laplacian kernel. The metrics for this implementation are described and provided in Section 5.

**Listing 2: Julia AMDGPU.jl Gray-Scott kernel.**

```
using AMDGPU
using Distributions

function _laplacian(i, j, k, var)
    l = var[i - 1, j, k] + var[i + 1, j, k]
      + var[i, j - 1, k] + var[i, j + 1, k]
      + var[i, j, k - 1] + var[i, j, k + 1]
      - 6.0 * var[i, j, k]
    return l / 6.0
end

function _kernel_amdgpu!(u, v, u_temp, v_temp,
                        sizes, Du, Dv, F, K,
                        noise, dt)

    k = (workgroupIdx().x - 1) * workgroupDim().x
      + workitemIdx().x
    j = (workgroupIdx().y - 1) * workgroupDim().y
      + workitemIdx().y
    i = (workgroupIdx().z - 1) * workgroupDim().z
      + workitemIdx().z

    if k == 1 || k >= sizes[3] ||
        j == 1 || j >= sizes[2] ||
        i == 1 || i >= sizes[1]
        return
    end

    @inbounds begin
        u_ijk = u[i, j, k]
        v_ijk = v[i, j, k]

        du = Du * _laplacian(i, j, k, u)
            - u_ijk * v_ijk^2 + F * (1.0 - u_ijk)
            + noise * rand(Uniform(-1, 1))

        dv = Dv * _laplacian(i, j, k, v)
            + u_ijk * v_ijk^2 - (F + K) * v_ijk

        u_temp[i, j, k] = u_ijk + du * dt
        v_temp[i, j, k] = v_ijk + dv * dt
    end

    return nothing
end
```




Figure 3: Seven-point stencil memory access for one variable in the parallel Gray-Scott solver.




Figure 4: MPI send/receive memory patterns in a Cartesian communicator decomposition used in Gray-Scott for each variable  $U$  and  $V$ .

For managing communication, we use the MPI.jl package, which provides a high-level binding interface with the underlying MPI back end. The vector data types and MPI send/receive communication code for a single variable ghost cell surface are described in Listing 3. MPI.jl uniquely enables referencing the internal memory through `@view` to form a buffer structure prior to an actual send/receive operation. We purposely introduce variations in the amount of data shared across the  $x, y$ , and  $z$  dimensions to capture distinct communication patterns.

Listing 3: GrayScott.jl MPI.jl implementation

```
# if in GPU memory move first to CPU
u = typeof(fields.u) <: Array ? fields.u :
    Array{typeof(fields.u)}

xy_face_t = MPI.Types.create_vector(size_y + 2,
    size_x, size_x + 2, MPI.Datatype(T))
MPI.Types.commit!(xy_face_t)
...
send = MPI.Buffer(@view(u[2, 1, size_z + 1]),
    1, data_type)
recv = MPI.Buffer(@view(u[2, 1, 1]),
    1, data_type)
MPI.Sendrecv!(send, recv, comm, dest = rank1,
    source = rank2)
```

We use the ADIOS2 library via the Julia ADIOS2.jl bindings. ADIOS2 is an MPI-based library, and the variables are described by accounting for their partitioning in a parallel I/O in which each MPI process produces a block of data in the global  $U$  and  $V$  3D array variables. A typical ADIOS2 dataset is shown in Listing 1 in which the combination of all data producer blocks from the MPI subdomain form the global  $U$  and  $V$  arrays from the PDE solution. We add schema attributes to enable later visualization in ParaView [2] by using FIDES [35] and VTX readers implemented in VTK [39].

### 4.1 The Frontier Supercomputer

Frontier is equipped with AMD CPUs and GPUs and achieved 1.194 EFLOPS in the HPL benchmark [12] according to the June 2023 TOP500 list.<sup>2</sup> Table 1 lists Frontier system characteristics and the software stack used for this study. In structuring our numerical experiments, we have designed them to grow by a factor of 8, so each dimension will double ( $2^3 = 8$ ), thereby aligning with the GPU’s maximum limits of 1,024 threads per dimension in a 3D computational kernel placement (Listing 2).

## 5 RESULTS AND DISCUSSION


In this section, we present the results from our evaluation of the Gray-Scott HPC workflow application regarding (i) compute capabilities on Frontier’s MI250x AMD GPU using AMDGPU.jl, (ii) weak scalability of the simulation wall-clock times using Julia’s MPI.jl, and (iii) simulation data production and consumption rates in JupyterHub using Julia’s bindings to ADIOS2, ADIOS2.jl.<sup>3</sup>

### 5.1 Computation on MI250x GPUs

As explained in Section 3.1, the metrics for the memory-bound Gray-Scott kernel presented in Listing 2 are based on the performance of accessing the GPU allocated arrays

<sup>2</sup><https://www.top500.org/lists/top500/2023/06/>

<sup>3</sup><https://github.com/eschnett/ADIOS2.jl>



**Figure 5: Gray-Scott simulation trace obtained with rocprof showing computational load on GPU and memory transfer to CPU for communication.**

**Table 1: Summary of Frontier hardware and software used in this study.**

Frontier Characteristics	
Nodes	9,408
–CPU Architecture–	
CPU	AMD EPYC 7A53
Cores	64
Memory	DDR4 512 GB
Bandwidth	205 GB/s
–GPU Architecture–	
GPU	4×AMD MI250X
	8×GCDs
Memory	HBM2E 64 GB
Bandwidth	1,600 GB/s per GCD
–Connectivity–	
GPU-to-GPU	Infinity Fabric
Bandwidth	50–100 GB/s
GPU-to-CPU	Infinity Fabric
Bandwidth	36 GB/s
–File System–	
Type	Lustre Orion
Capacity	679 PB
Nodes	40 metadata
	450 object storage service (OSS)
Write speed	5.5 TBps
Read speed	4.5 TBps
–Software versions–	
Julia	1.9.2
AMDGPU.jl	0.4.15
ROCm	5.4.0
MPI.jl	0.20.12
Cray-MPICH	8.1.23
ADIOS2.jl	1.2.1
ADIOS2	2.8.3

$u$  and  $v$  for 7 read (fetch/load) operations and 1 write operation. We collect two types of bandwidth information:

(i) effective bandwidth, meaning the expected minimal data movement (fetch/load and write/store) according to the number of operations in a single domain, and (ii) the total measured bandwidth obtained from AMD’s rocprof profiler for `FETCH_SIZE` and `WRITE_SIZE`, which represent the total number of KB moved to and from the GPU memory (Figure 5).

The effective fetch and write sizes can be estimated from the number of floating point operations per cell by considering that edge cells execute fewer operations due to reduced stencil size. A regular grid with  $L$  number of cells per direction is provided in Equations (4a) and (4b), as reported in AMD’s training material.<sup>4</sup>

$$fetch\_size_{effective} = [L^3 - 8 - 12(L - 2)] \cdot sizeof(T) \quad (4a)$$

$$write\_size_{effective} = (L - 2)^3 \cdot sizeof(T) \quad (4b)$$

Hence, the effective and total bandwidth metrics are shown in Equations (5a) and (5b) for the effective (theoretical) and the actual (measured with rocprof) GPU data access bandwidths.

$$bandwidth_{effective} = \frac{(fetch\_size + write\_size)_{effective}}{kernel\_time} \quad (5a)$$

$$bandwidth_{total} = \frac{(FETCH\_SIZE + WRITE\_SIZE)_{rocprof}}{kernel\_time} \quad (5b)$$

Listing 4 presents the generated LLVM-IR memory accesses in the Julia implementation. These numbers are consistent with the algorithm patterns in Listing 2 of 16 loads and 2 stores. Consequently, we did not observe additional overhead typical of high-level implementations. The reasons for the bandwidth performance difference in Table 3

<sup>4</sup><https://github.com/AMD/amd-lab-notes>

**Table 2: Average bandwidth comparison of different stencil implementations on a single GPU.**

Kernel	Bandwidth (GB/s)	
	Effective	Total
Julia GrayScott.jl		
- 2-variable (application)	312	570
- 1-variable no random	312	625
HIP single variable	599	1,163
Theoretical peak MI250x	1,600	

**Table 3: The rocprow outputs for HIP 1-variable and Julia Gray-Scott.jl implementations.**

Kernel metric	HIP	GrayScott.jl	
	1-var	1-variable no random	2-variable (application)
wgr	256	512	512
lds	0	29,184	29,184
scr	0	8,192	8,192
FETCH_SIZE (GB)	25.08	25.40	50.80
WRITE_SIZE (GB)	8.35	8.38	16.78
TCC_HIT (M)	9.14	10.80	24.60
TCC_MISS (M)	8.36	8.69	17.19
Avg Duration (ms)	28.74	54.03	111.07

are beyond the IR level and might present opportunities for lower-level optimizations at the LLVM instruction set architecture level, thereby providing an opportunity to close gaps from vendor-specific optimizations as Julia’s AMDGPU.jl and the AMD HIP/ROCm stack continues to rapidly evolve.

**Listing 4: GrayScott.jl application kernel unique memory loads (14) and store (2) in LLVM-IR.**

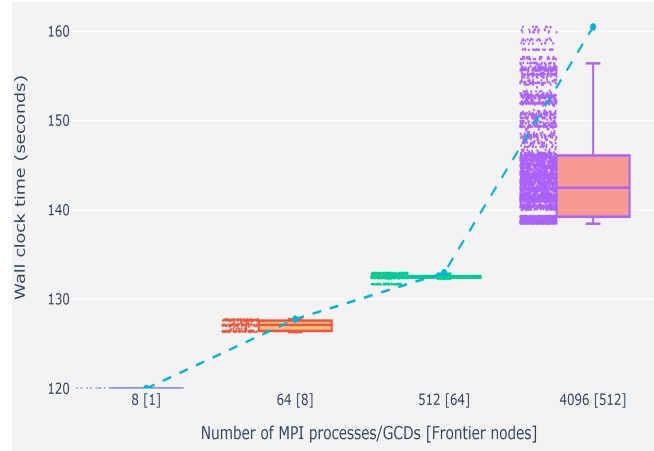
```

%94 = load double, double addrspace(1)* %93, align 8
%103 = load double, double addrspace(1)* %102, align 8
%107 = load double, double addrspace(1)* %106, align 8
%110 = load double, double addrspace(1)* %109, align 8
%114 = load double, double addrspace(1)* %113, align 8
%117 = load double, double addrspace(1)* %116, align 8
%122 = load double, double addrspace(1)* %121, align 8
%126 = load double, double addrspace(1)* %125, align 8
%312 = load double, double addrspace(1)* %311, align 8
%315 = load double, double addrspace(1)* %314, align 8
%318 = load double, double addrspace(1)* %317, align 8
%321 = load double, double addrspace(1)* %320, align 8
%325 = load double, double addrspace(1)* %324, align 8
%329 = load double, double addrspace(1)* %328, align 8
...
store double %345, double addrspace(1)* %353, align 8
store double %355, double addrspace(1)* %363, align 8
    
```

## 5.2 Weak scaling

We ran GrayScott.jl with different cell sizes using 1 GCD (GPU) per MPI process. Weak scaling results are illustrated in Figure 6 and show the total wall-clock times per MPI

process for a constant workload of  $n_x = n_y = n_z = 1,024$  per dimension, or  $1,024^3 = 1,073 \times 10^9$  cells per GPU.




**Figure 6: Weak scaling with single MPI process wall-clock time variability obtained with Gray-Scott.jl on Frontier.**

Per-process wall-clock time variability in HPC workflows is proportional to the number of MPI processes. As expected, the overall communication overhead is dictated by the slowest time-to-solution processes (dashed blue line). The Julia code could run successfully on 4,096 GPUs or 512 nodes of Frontier (5.44% of the system). A clear trend of small variability of a few percent (2%–3%) is apparent up to 512 MPI processes, whereas large variability (12%–15%) appears for the largest case. Although we do not have a baseline comparison, this initial measurement can help with further optimizations because Julia can produce realistic scenarios for MPI communication on a system such as Frontier. When attempting runs for the next size for a factor 8 at 32,768 GPUs (4,096 nodes, 50% of the Frontier system), unpredictable failures occurred at the underlying MPI layers during the ghost cell exchange stage. Nevertheless, all 32,768 GPUs showed initial runs keeping the bandwidth close to the expected value of 312 GB/s, as reported in Table 2.

In the context of HPC workflows, understanding the interplay between scale and performance is critical. To this end, we delve deeper into GPU bandwidth metrics from large-scale runs. Figure 7 shows the obtained bandwidth distributions for the initial JIT compiled run, which represents the initial overhead in running Julia kernels, and the optimized code when running on 4,096 GPUs. The initial JIT compilation represents the early overhead associated with running Julia kernels in an HPC workflow—an  $\approx 12.5\times$  cost, or 8% on average, of the bandwidth in the optimized kernel. Although JIT compilation represents an amortized cost, it should be incorporated in the estimation of total wall-clock run times.

While this overhead can be addressed in HPC workflow planning, Julia’s ahead-of-time mechanism was not explored in this study. These insights not only demonstrate Julia’s capacity for HPC workflows but also reveal areas for further research and optimization, which could enable more efficient utilization of resources in complex computational environments.




**Figure 7: Bandwidth distribution for JIT and optimized kernel code running on 4,096 Frontier GPUs for 20 simulation steps.**

### 5.3 Parallel I/O

In HPC workflows, the efficient use of parallel file systems is crucial for both saving computational results and ensuring smooth data transfer between different stages of a workflow. To this end, we evaluate using the parallel file system within the GrayScott.jl package, focusing on saving one output step of the cases presented in Figure 6. Results from using the ADIOS2.jl Julia bindings package are presented in Figure 8, revealing the performance characteristics of weak scaling in the writing of the dataset described in Section 4. The underlying ADIOS2 library uses the default BP5 format, which generates a single sub-file per node. At 512 nodes, write times are fairly flat, which represents an increase of write bandwidth up to 434 GB/s for the largest case. This represents an 8% peak file system write performance of 5.5 TB/s when using 5% of the machine (Table 1). In the context of HPC workflows, such a performance metric is significant, as parallel I/O performance is dictated by low-level POSIX write operations, the ADIOS2 buffering management, and real-time file system usage. With the Julia binding, we observed negligible overhead, indicating that Julia can serve as an effective tool in HPC workflow optimization.

Finally, we look at Julia in an end-to-end workflow development process for data analysis. We tested a Julia implementation by reading a plotting of simple 2D slices of the



**Figure 8: Weak scaling on parallel I/O showing wall-clock times and bandwidth performance when using ADIOS2.jl on Frontier.**

stored  $U$  and  $V$  3D arrays with the Makie.jl library [9]. Figure 9 shows a snapshot of the resulting JupyterHub service consuming the ADIOS2-generated data shared on Frontier’s file system. This instance demonstrates Julia’s potential as a unified language that can seamlessly connect different stages of an HPC workflow, from computational simulation to data visualization. Future work includes further performance tuning and evaluation of the trade-offs for in-memory streaming data pipelines [34].

By bridging various components in a typical HPC workflow, Julia could emerge as a compelling option for scientists and engineers working on complex, large-scale computational problems. It also sets the stage for more integrated, responsive, and adaptive HPC workflows that can fully leverage the capabilities of modern supercomputing infrastructures.

## 6 CONCLUSIONS

On up to 4,096 GPUs and MPI processes (512 nodes of the Frontier supercomputer), we used a 2-variable diffusion-reaction code, Gray-Scott, to test the performance of the Julia HPC ecosystem in the development of tightly coupled workflow components. Overall, the Julia stencil solver achieves close to a 50% bandwidth on Frontier’s MI250x AMD GPUs vs. a HIP implementation; hence, performance gaps still exist and must be closed as we look forward to future versions of the actively developed AMDGPU.jl. Meanwhile, the measured weak scaling from the MPI communication and parallel I/O components suggest that bindings available in Julia are lightweight and perfectly suitable running on top of the underlying MPI and ADIOS2 library implementations. The Julia



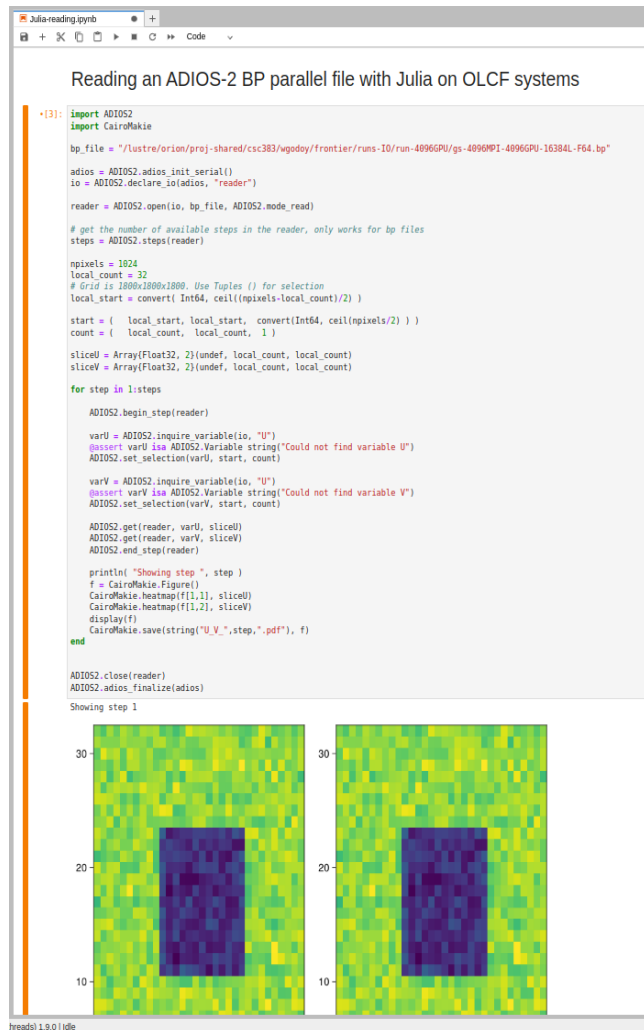


Figure 9: Snapshot of Julia data analysis implementation on the OLCF's JupyterHub server.

implementation shows similar patterns in overhead and variability typical of network and file system communication in HPC systems. We also showcased the data analysis feasibility on available JupyterHub systems. Overall, the LLVM-based Julia HPC ecosystem presents an attractive alternative for developing co-design components given the high-performance and high-productivity requirements for the end-to-end workflows that power scientific discovery (e.g., AI, FAIR).

## ACKNOWLEDGMENTS

This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the US Department of Energy Office of Science and the National Nuclear Security Administration. This research used resources

of the Oak Ridge Leadership Computing Facility and the Experimental Computing Laboratory (ExCL) at the Oak Ridge National Laboratory, which is supported by the Office of Science of the US Department of Energy under Contract No. DE-AC05-00OR22725. This work is funded, in part, by Bluestone, a X-Stack project in the DOE Advanced Scientific Computing Office with program manager Hal Finkel.

## REFERENCES

- [1] AMD. 2022. AMD ROCm v5.2 Release. [https://rocm.docs.amd.com/en/latest/Current\\_Release\\_Notes/Current-Release-Notes.html#amd-rocm-v5-2-release](https://rocm.docs.amd.com/en/latest/Current_Release_Notes/Current-Release-Notes.html#amd-rocm-v5-2-release)
- [2] Utkarsh Ayachit. 2015. *The ParaView guide: a parallel visualization application*. Kitware, Inc.
- [3] David A. Beckingsale, Jason Burmark, Rich Hornung, Holger Jones, William Killian, Adam J. Kunen, Olga Pearce, Peter Robinson, Brian S. Ryujiin, and Thomas RW Scogland. 2019. RAJA: Portable Performance for Large-Scale Scientific Applications. In *2019 IEEE/ACM International Workshop on Performance, Portability and Productivity in HPC (P3HPC)*. 71–81. <https://doi.org/10.1109/P3HPC49587.2019.00012>
- [4] Tal Ben-Nun, Todd Gamblin, D. S. Hollman, Hari Krishnan, and Chris J. Newburn. 2020. Workflows are the New Applications: Challenges in Performance, Portability, and Productivity. In *2020 IEEE/ACM International Workshop on Performance, Portability and Productivity in HPC (P3HPC)*. 57–69. <https://doi.org/10.1109/P3HPC51967.2020.00011>
- [5] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B Shah. 2017. Julia: A Fresh Approach to Numerical Computing. *SIAM Rev.* 59, 1 (Jan. 2017), 65–98. <https://doi.org/10.1137/141000671>
- [6] Simon Byrne, Lucas C. Wilcox, and Valentin Churavy. 2021. MPIjl: Julia bindings for the Message Passing Interface. *Proceedings of the JuliaCon Conferences 1*, 1 (2021), 68. <https://doi.org/10.21105/jcon.00068>
- [7] H. Carter Edwards, Christian R. Trott, and Daniel Sunderland. 2014. Kokkos: Enabling manycore performance portability through polymorphic memory access patterns. *J. Parallel and Distrib. Comput.* 74, 12 (2014), 3202–3216. <https://doi.org/10.1016/j.jpdc.2014.07.003> Domain-Specific Languages and High-Level Frameworks for High-Performance Computing.
- [8] Valentin Churavy, William F Godoy, Carsten Bauer, Hendrik Ranocha, Michael Schlotke-Lakemper, Ludovic Räss, Johannes Blaschke, Mosè Giordano, Erik Schnetter, Samuel Omlin, Jeffrey S. Vetter, and Alan Edelman. 2022. Bridging HPC Communities through the Julia Programming Language. arXiv:2211.02740 [cs.DC]
- [9] Simon Danisch and Julius Krumbiegel. 2021. Makie.jl: Flexible high-performance data visualization for Julia. *Journal of Open Source Software* 6, 65 (2021), 3349. <https://doi.org/10.21105/joss.03349>
- [10] Kaushik Datta, Mark Murphy, Vasily Volkov, Samuel Williams, Jonathan Carter, Leonid Oliker, David Patterson, John Shalf, and Katherine Yelick. 2008. Stencil computation optimization and autotuning on state-of-the-art multicore architectures. In *SC '08: Proceedings of the 2008 ACM/IEEE Conference on Supercomputing*. 1–12. <https://doi.org/10.1109/SC.2008.5222004>
- [11] Ewa Deelman, Tom Peterka, Ilkay Altintas, Christopher D Carothers, Kerstin Kleese van Dam, Kenneth Moreland, Manish Parashar, Lavanya Ramakrishnan, Michela Taufer, and Jeffrey Vetter. 2018. The future of scientific workflows. *The International Journal of High Performance Computing Applications* 32, 1 (2018), 159–175. <https://doi.org/10.1177/1094342017704893> arXiv:https://doi.org/10.1177/1094342017704893

- [12] Jack J Dongarra, Piotr Luszczek, and Antoine Petitet. 2003. The LINPACK benchmark: past, present and future. *Concurrency and Computation: practice and experience* 15, 9 (2003), 803–820.
- [13] Thomas Faingnaert, Tim Besard, and Bjorn De Sutter. 2022. Flexible Performant GEMM Kernels on GPUs. *IEEE Transactions on Parallel and Distributed Systems* 33, 9 (2022), 2230–2248. <https://doi.org/10.1109/TPDS.2021.3136457>
- [14] Rafael Ferreira da Silva, Rosa M. Badia, Venkat Bala, Debbie Bard, Timo Bremer, Ian Buckley, Silvina Caino-Lores, Kyle Chard, Carole Goble, Shantenu Jha, Daniel S. Katz, Daniel Laney, Manish Parashar, Frederic Suter, Nick Tyler, Thomas Uram, Ilkay Altintas, et al. 2023. *Workflows Community Summit 2022: A Roadmap Revolution*. Technical Report ORNL/TM-2023/2885. Oak Ridge National Laboratory. <https://doi.org/10.5281/zenodo.7750670>
- [15] Rafael Ferreira da Silva, Henri Casanova, Kyle Chard, Ilkay Altintas, Rosa M Badia, Bartosz Balis, Tainã Coleman, Frederik Coppens, Frank Di Natale, Bjoern Enders, Thomas Fahringer, Rosa Filgueira, Grigori Fursin, Daniel Garijo, Carole Goble, Dorrán Howell, Shantenu Jha, Daniel S. Katz, Daniel Laney, Ulf Leser, Maciej Malawski, Kshitij Mehta, Loïc Pottier, Jonathan Ozik, J. Luc Peterson, Lavanya Ramakrishnan, Stian Soiland-Reyes, Douglas Thain, and Matthew Wolf. 2021. A Community Roadmap for Scientific Workflows Research and Development. In *2021 IEEE Workshop on Workflows in Support of Large-Scale Science (WORKS)*. 81–90. <https://doi.org/10.1109/WORKS54523.2021.00016>
- [16] Rafael Ferreira da Silva, Rosa Filgueira, Ilia Pietri, Ming Jiang, Rizos Sakellariou, and Ewa Deelman. 2017. A Characterization of Workflow Management Systems for Extreme-Scale Applications. *Future Generation Computer Systems* 75 (2017), 228–238. <https://doi.org/10.1016/j.future.2017.02.026>
- [17] Mosè Giordano, Milan Klöwer, and Valentin Churavy. 2022. Productivity meets Performance: Julia on A64FX. In *2022 IEEE International Conference on Cluster Computing (CLUSTER)*. 549–555. <https://doi.org/10.1109/CLUSTER51413.2022.00072>
- [18] Carole Goble, Sarah Cohen-Boulakia, Stian Soiland-Reyes, Daniel Garijo, Yolanda Gil, Michael R Cruseo, Kristian Peters, and Daniel Schober. 2020. FAIR computational workflows. *Data Intelligence* 2, 1-2 (2020), 108–121.
- [19] William F. Godoy and Xu Liu. 2011. Introduction of Parallel GPGPU Acceleration Algorithms for the Solution of Radiative Transfer. *Numerical Heat Transfer, Part B: Fundamentals* 59, 1 (2011), 1–25. <https://doi.org/10.1080/10407790.2010.541359>
- [20] William F. Godoy, Norbert Podhorszki, Ruonan Wang, Chuck Atkins, Greg Eisenhauer, Junmin Gu, Philip Davis, Jong Choi, Kai Gernaschewski, Kevin Huck, Axel Huebl, Mark Kim, James Kress, Tahsin Kurc, Qing Liu, Jeremy Logan, Kshitij Mehta, George Ostrouchov, Manish Parashar, Franz Poeschel, David Pugmire, Eric Suchyta, Keichi Takahashi, Nick Thompson, Seiji Tsutsumi, Lipeng Wan, Matthew Wolf, Kesheng Wu, and Scott Klasky. 2020. Adios 2: The Adaptable Input Output System. A framework for high-performance data management. *SoftwareX* 12 (2020), 100561. <https://doi.org/10.1016/j.softx.2020.100561>
- [21] William F. Godoy, Pedro Valero-Lara, T. Elise Dettling, Christian Trefftz, Ian Jorquera, Thomas Sheehy, Ross G. Miller, Marc Gonzalez-Tallada, Jeffrey S. Vetter, and Valentin Churavy. 2023. Evaluating performance and portability of high-level programming models: Julia, Python/Numba, and Kokkos on exascale nodes. In *2023 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*. 373–382. <https://doi.org/10.1109/IPDPSW59300.2023.00068>
- [22] William Gropp, Rajeev Thakur, and Ewing Lusk. 1999. *Using MPI-2: Advanced features of the message passing interface*. MIT press.
- [23] Justin Holewinski, Louis-Noël Pouchet, and P. Sadayappan. 2012. High-Performance Code Generation for Stencil Computations on GPU Architectures. In *Proceedings of the 26th ACM International Conference on Supercomputing (San Servolo Island, Venice, Italy) (ICS '12)*. Association for Computing Machinery, New York, NY, USA, 311–320. <https://doi.org/10.1145/2304576.2304619>
- [24] Sascha Hunold and Sebastian Steiner. 2020. Benchmarking Julia’s Communication Performance: Is Julia HPC ready or Full HPC?. In *2020 IEEE/ACM Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS)*. IEEE. <https://doi.org/10.1109/pmbs51919.2020.00008>
- [25] Marcin Krotkiewski and Marcin Dabrowski. 2013. Efficient 3D stencil computations using CUDA. *Parallel Comput.* 39, 10 (2013), 533–548. <https://doi.org/10.1016/j.parco.2013.08.002>
- [26] Siu Kwan Lam, Antoine Pitrou, and Stanley Seibert. 2015. Numba: A LLVM-based Python JIT compiler. In *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*. 1–6.
- [27] Chris Lattner and Vikram Adve. 2004. LLVM: A compilation framework for lifelong program analysis & transformation. In *International Symposium on Code Generation and Optimization, 2004. CGO 2004*. IEEE, 75–86.
- [28] Wei-Chen Lin and Simon McIntosh-Smith. 2021. Comparing Julia to Performance Portable Parallel Programming Models for HPC. In *2021 International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS)*. 94–105. <https://doi.org/10.1109/PMBS54543.2021.00016>
- [29] Wenbin Lu, Luis E. Peña, Pavel Shamis, Valentin Churavy, Barbara Chapman, and Steve Poole. 2022. Bring the BitCODE-Moving Compute and Data in Distributed Heterogeneous Systems. In *2022 IEEE International Conference on Cluster Computing (CLUSTER)*. 12–22. <https://doi.org/10.1109/CLUSTER51413.2022.00017>
- [30] John D McCalpin et al. 1995. Memory bandwidth and machine balance in current high performance computers. *IEEE computer society technical committee on computer architecture (TCCA) newsletter* 2, 19-25 (1995).
- [31] NVIDIA. 2022. CUDA Toolkit Documentation - v11.7.0. <https://developer.nvidia.com/cuda-toolkit>
- [32] OpenMP Architecture Review Board. 2021. OpenMP Application Program Interface Version 5.2. <https://www.openmp.org/wp-content/uploads/OpenMP-API-Specification-5-2.pdf>
- [33] John E. Pearson. 1993. Complex Patterns in a Simple System. *Science* 261, 5118 (1993), 189–192. <https://doi.org/10.1126/science.261.5118.189> arXiv:<https://www.science.org/doi/pdf/10.1126/science.261.5118.189>
- [34] Franz Poeschel, Juncheng E, William F. Godoy, Norbert Podhorszki, Scott Klasky, Greg Eisenhauer, Philip E. Davis, Lipeng Wan, Ana Gainaru, Junmin Gu, Fabian Koller, René Widera, Michael Bussmann, and Axel Huebl. 2022. Transitioning from File-Based HPC Workflows to Streaming Data Pipelines with openPMD and ADIOS2. In *Driving Scientific and Engineering Discoveries Through the Integration of Experiment, Big Data, and Modeling and Simulation*, Jeffrey Nichols, Arthur ‘Barney’ Maccabe, James Nutaro, Swaroop Pophale, Pravalika Devineni, Theresa Ahearn, and Becky Verastegui (Eds.). Springer International Publishing, Cham, 99–118.
- [35] David Pugmire, Caitlin Ross, Nicholas Thompson, James Kress, Chuck Atkins, Scott Klasky, and Berk Geveci. 2021. Fides: a general purpose data model library for streaming data. In *High Performance Computing: ISC High Performance Digital 2021 International Workshops, Frankfurt am Main, Germany, June 24–July 2, 2021, Revised Selected Papers 36*. Springer, 495–507.
- [36] Hendrik Ranocha, Michael Schlottke-Lakemper, Andrew R. Winters, Erik Faulhaber, Jesse Chan, and Gregor J. Gassner. 2022. Adaptive numerical simulations with Trixi.jl: A case study of Julia for scientific computing. *Proceedings of the JuliaCon Conferences* 1, 1 (2022), 77.

- <https://doi.org/10.21105/jcon.00077>
- [37] Jeffrey Regier, Keno Fischer, Kiran Pamnany, Andreas Noack, Jarrett Revels, Maximilian Lam, Steve Howard, Ryan Giordano, David Schlegel, Jon McAuliffe, Rollin Thomas, and Prabhat. 2019. Cataloging the visible universe through Bayesian inference in Julia at petascale. *J. Parallel and Distrib. Comput.* 127 (2019), 89–104. <https://doi.org/10.1016/j.jpdc.2018.12.008>
- [38] Julian Samaroo, Valentin Churavy, Wiktor Phillips, Ali Ramadhan, Jason Barmpparesos, Julia TagBot, Ludovic Räss, Michel Schanen, Tim Besard, Anton Smirnov, Takafumi Arakaki, Stephan Antholzer, Alessandro, Chris Elrod, Matin Raayai, and Tom Hu. 2022. *JuliaGPU/AMDGPU.jl: v0.4.1*. <https://doi.org/10.5281/zenodo.6949520>
- [39] William J Schroeder, Lisa Sobierajski Avila, and William Hoffman. 2000. Visualizing with VTK: a tutorial. *IEEE Computer graphics and applications* 20, 5 (2000), 20–27.
- [40] Honghui Shang, Li Shen, Yi Fan, Zhiqian Xu, Chu Guo, Jie Liu, Wenhao Zhou, Huan Ma, Rongfen Lin, Yuling Yang, Fang Li, Zhuoya Wang, Yunquan Zhang, and Zhenyu Li. 2022. Large-Scale Simulation of Quantum Computational Chemistry on a New Sunway Supercomputer. In *SC22: International Conference for High Performance Computing, Networking, Storage and Analysis*. 1–14. <https://doi.org/10.1109/SC41404.2022.00019>
- [41] Raúl Sirvent, Javier Conejero, Francesc Lordan, Jorge Ejarque, Laura Rodríguez-Navas, José M. Fernández, Salvador Capella-Gutiérrez, and Rosa M. Badia. 2022. Automatic, Efficient and Scalable Provenance Registration for FAIR HPC Workflows. In *2022 IEEE/ACM Workshop on Workflows in Support of Large-Scale Science (WORKS)*. 1–9. <https://doi.org/10.1109/WORKS56498.2022.00006>
- [42] Dominik Straßel, Philipp Reusch, and Janis Keuper. 2020. Python Workflows on HPC Systems. In *2020 IEEE/ACM 9th Workshop on Python for High-Performance and Scientific Computing (PyHPC)*. 32–40. <https://doi.org/10.1109/PyHPC51966.2020.00009>
- [43] Jeffrey S. Vetter, Ron Brightwell, Maya Gokhale, Pat McCormick, Rob Ross, John Shalf, Katie Antypas, David Donofrio, Travis Humble, Catherine Schuman, Brian Van Essen, Shinjae Yoo, Alex Aiken, David Bernholdt, Suren Byna, Kirk Cameron, Frank Cappello, Barbara Chapman, Andrew Chien, Mary Hall, Rebecca Hartman-Baker, Zhiling Lan, Michael Lang, John Leidel, Sherry Li, Robert Lucas, John Mellor-Crummey, Paul Peltz Jr., Thomas Peterka, Michelle Strout, and Jeremiah Wilke. 2018. Extreme Heterogeneity 2018 - Productive Computational Science in the Era of Extreme Heterogeneity: Report for DOE ASCR Workshop on Extreme Heterogeneity. (12 2018). <https://doi.org/10.2172/1473756>
- [44] Wei Wu, George Bosilca, Rolf Vandevaraart, Sylvain Jeaugey, and Jack Dongarra. 2016. GPU-Aware Non-contiguous Data Movement In Open MPI. In *Proceedings of the 25th ACM International Symposium on High-Performance Parallel and Distributed Computing*. 231–242.

To run the codes on Frontier, we created a specific branch to avoid issues related to CUDA.jl and random number generation on AMDGPU.jl until fixes are available:  
<https://github.com/JuliaORNL/GrayScott.jl/tree/amdgpu-frontier>.

We provide a configuration file to set up the modules environment as of June/July 2023: [https://github.com/JuliaORNL/GrayScott.jl/blob/amdgpu-frontier/scripts/config\\_frontier.sh](https://github.com/JuliaORNL/GrayScott.jl/blob/amdgpu-frontier/scripts/config_frontier.sh).  
We used the latest available release of Julia, version 1.9.2.

## A ARTIFACT DESCRIPTION

The GrayScott.jl implementation is available on GitHub:  
<https://github.com/JuliaORNL/GrayScott.jl>.

The repository contains instructions for the code and how to run it locally, including setting up the JSON configuration files used in this study, which are also available:  
<https://github.com/JuliaORNL/GrayScott.jl/blob/main/examples/settings-files.json>.