

Abstractions and Directives for Adapting Wavefront Algorithms to Future Architectures*

Robert Searles
University of Delaware
Newark, Delaware
rsearles@udel.edu

Sunita
Chandrasekaran
University of Delaware
Newark, Delaware
schandra@udel.edu

Wayne Joubert
Oak Ridge National Lab
Oak Ridge, Tennessee
joubert@ornl.gov

Oscar Hernandez
Oak Ridge National Lab
Oak Ridge, Tennessee
oscar@ornl.gov

ABSTRACT

Architectures are rapidly evolving, and exascale machines are expected to offer billion-way concurrency. We need to rethink algorithms, languages and programming models among other components in order to migrate large scale applications and explore parallelism on these machines. Although directive-based programming models allow programmers to worry less about programming and more about science, expressing complex parallel patterns in these models can be a daunting task especially when the goal is to match the performance that the hardware platforms can offer. One such pattern is wavefront. This paper extensively studies a wavefront-based miniapplication for Denovo, a production code for nuclear reactor modeling. We parallelize the Koch-Baker-Alcouffe (KBA) parallel-wavefront sweep algorithm in the main kernel of Minisweep (the miniapplication) using CUDA, OpenMP and OpenACC. Our OpenACC implementation running on NVIDIA's next-generation Volta GPU boasts an 85.06x speedup over serial code, which is larger than CUDA's 83.72x speedup over the same serial implementation. Our experimental platform includes SummitDev, an ORNL representative architecture of the upcoming Summit supercomputer. Our parallelization effort across platforms also motivated us to define an abstract parallelism model that is architecture independent, with a goal of creating software abstractions that can be used by applications employing the wavefront sweep motif.

ACM Reference Format:

Robert Searles, Sunita Chandrasekaran, Wayne Joubert, and Oscar Hernandez. 2018. Abstractions and Directives for Adapting Wavefront Algorithms to Future Architectures. In *Proceedings of ACM PASC Conference (PASC'18)*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

*Notice: This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The US government retains and the publisher, by accepting the article for publication, acknowledges that the US government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for US government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>).

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

PASC'18, July 2018, Congress Center Basel, Switzerland

© 2018 Copyright held by the owner/author(s).

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM.

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Hardware architectures are rapidly evolving. High performance computing nodes are becoming increasingly heterogeneous. The current and anticipated exascale accelerated node architectures are heterogeneous [9]. They are expected to contain a mix of throughput and latency optimized cores [42]. Such a balanced mixture of cores is expected to manage different types of parallelism available in an algorithm. Memory has advanced as well. 3-D memory stacking with memory moving on-socket provides increased bandwidth and faster communication.

Such diverse architectures require their own code optimization strategies, while the application developers prefer a “write-once” code development strategy in which a single code will execute efficiently on all targeted architectures. In addition to considering the underlying hardware, a programming model is also expected to address requirements of applications and their algorithms. The programming language that implements the model should provide the right abstractions to improve the productivity of scientific developers. Programmers often resort to a trade-off between achieving portability and high performance. Why? The issue is two-fold. Adequate application parallelism will not be exposed to the hardware architecture if the algorithm is structured in a way that limits the level of concurrency that a programming model can benefit from. Secondly, such a single code representation is possible only if the programming abstractions are carefully crafted for the programming models to provide informative hints to the compilers to generate optimized code across platforms.

Directives allow us to abstract the rich feature set of hardware architectures and incrementally improve, port, and maintain the codebase across platforms. However, there still remains a gap in the way that they do not adequately expose and parallelize some of the complex algorithms often found in applications. One such case is a wavefront-based algorithm that is of critical importance to solving scientific problems in multiple science domains. Wavefront algorithms are useful for problems for which the computed result values have dependencies, requiring that results be computed in stages (wavefronts) for which each stage's results depends on results computed in previous stages.

In this paper, our objective is to study this type of algorithm and identify challenges in exposing parallelism using high-level abstractions that can be lowered to de facto parallel programming languages. We study the Minisweep proxy application [32], that is illustrative of the complexities in a wavefront-based algorithm. Parallelizing Minisweep using current directive-based APIs revealed the gaps in their expressivity and features. We address this issue by designing and envisioning an abstract parallelism model that

highlights these gaps with a combination of notations. Having these proposed notations in a programming language are key to exposing and mapping wavefront-based parallelism to multiple architectures.

For Minisweep, we use OpenACC and OpenMP along with CUDA. OpenACC, since its inception in 2012, is being widely used to port large scale applications spanning several domains such as ANSYS [40], GAUSSIAN [21], and Icosahedral non-hydrostatic (ICON) [41] to massively parallel architectures. Similarly, OpenMP 4.5 is being deployed to applications such as Pseudo-Spectral Direct Numerical Simulation-Combined Compact Difference (PSDNS-CCD3D) [17] and a CFD code for turbulent flow simulation, Quicksilver [37], a Monte Carlo Transport code.

1.1 Application Under Study

The Minisweep proxy application [32] is part of the Profugus radiation transport miniapp project [5] that reproduces the computational pattern of the sweep kernel of the Denovo S_n radiation transport code [20]. The sweep kernel is responsible for most of the computational expense (80-99%) of Denovo. Denovo, a production code for nuclear reactor neutronics modeling, is in use by a current DOE INCITE project to model the International Thermonuclear Experimental Reactor (ITER) fusion reactor [6]. The many runs of this code required to perform reactor simulations at high node counts makes it an important target for efficient mapping to accelerated architectures.

This study involves S_n radiation transport algorithms for solving the linear Boltzmann equation [30]. Here, a continuum model is used to simulate the density of particles of a given energy and direction of motion within a 3-D volume. The approach yields a six dimensional problem (3-D in space, 2-D in angular particle direction and 1-D in particle energy) that is appropriately discretized in each dimension. Minisweep includes neutronics calculations for nuclear reactor [3] and fusion reactor [6] design, radiation shielding, nuclear forensics and radiation detection. The large number of problem dimensions available in the S_n transport algorithm affords significant opportunities for parallelism on manycore parallel systems. However, the recursive nature of the wavefront calculation in the spatial dimensions is a challenge to efficient parallelization.

Denovo was one of six applications selected for early application readiness on ORNL's Titan system under the Center for Accelerated Application Readiness (CAAR) project [12] and is part of the Exnihilo code suite which received an R&D 100 award for modeling the Westinghouse AP1000 reactor [2]. Minisweep can be considered a successor to the well-known Sweep3D benchmark [1] and is similar to other S_n wavefront codes including Kripke [4], SN (Discrete Ordinates) Application Proxy (SNAP) [7] and PARTISN [13].

Minisweep is used as a vehicle to examine parallelization of wavefront algorithms in general. However, it has multiple computational motifs (dense and sparse linear algebra, structured grids) and parallelism requirements (halo communications, hierarchical synchronizations, atomic updates) which make the study of this algorithm relevant to a much broader spectrum of codes.

1.2 Contributions

Our work presents the following contributions:

- An abstract representation elucidating architectural, memory and threading challenges to programming models for

such complex wavefront algorithms as used in Minisweep that can be broadly applicable to applications with a similar computational motif. More details in Section 1.1 and Section 8.

- A performance-portable implementation of these abstractions using OpenACC to offload portions of an application to a variety of parallel architectures.
- A description of the challenges in existing programming models, and extensions that will allow programmers to overcome the obstacle of recursivity in the spatial dimensions of wavefront algorithms without requiring large modifications to the code base.

2 OVERVIEW OF SWEEP ALGORITHM

The S_n transport sweep algorithm possesses features common to wavefront algorithms in general yet has structure specific to the requirements of S_n transport. It can be considered in two parts: first, a wavefront algorithm relating the computations between gridcells of a 3-D grid, and second the computations performed on a single gridcell within this wavefront sweep across a grid.

2.1 Grid-level computations

We consider here a 3-D structured grid, with locally connected gridcells; see Figure 1. Importantly, the result computed at a gridcell is dependent on the results computed at the three neighbor gridcells in the upstream x , y and z directions; thus the computation is described by a four-point stencil. This dependency puts a restriction on the order in which results can be computed. One possible ordering is a series of wavefronts described by a sequence of planes of gridcells starting at a corner of the grid and sweeping through the whole grid (Figure 1). Other orderings are allowed as well, as long as the dependencies are satisfied.

Modeling the physical problem requires modeling particle flux in all directions. To accomplish this, an execution instance of the algorithm performs a total of eight sweeps, one starting at each corner of the domain. These directions are referred to as "octants." The results of all eight octant sweeps are added together form the final result.




Figure 1: Wavefront computational pattern

2.2 Gridcell-level computations

To further describe the algorithm, we define array v_s with dimensions $v_s(n_x, n_y, n_z, n_u, n_e, n_a, n_o)$. The n_x, n_y, n_z dimensions refer to the spatial grid size. The dimension $n_o = 8$ is the octants axis, across which results are summed for the final result. The value n_a is a set of angular directions, and n_e is the number of energy groups, each representing a decoupled instance of the problem. Finally, n_u represents a set of unknowns for each gridcell based on the spatial discretization, e.g., finite element coefficients for a gridcell. Though v_s is relevant to key computations of the algorithm, the arrays that actually hold the input and output of the algorithm have the form

$v(n_x, n_y, n_z, n_u, n_e, n_m)$. Here v_s and v are related by the fact that the n_a, n_o axes are compressed into n_m moments to form v from v_s . The computation at a gridcell then is composed of the following steps:

- (1) Moments-to-angles conversion. For a given octant and energy group the input $v_{in}(i_x, i_y, i_z, *, i_e, *)$ is transformed into array $v_{in,s}(i_x, i_y, i_z, *, i_e, *, i_o)$ by a small matrix-vector product that relates the n_m moments to the n_a angles. The matrix depends on the octant but is independent of spatial location, energy group and unknown.
- (2) Face contribution. The upstream components from the sweep are added to $v_{in,s}(i_x, i_y, i_z, *, i_e, *, i_o)$.
- (3) Solve. An operation is performed on the array values $v_{in,s}(i_x, i_y, i_z, *, i_e, *, i_o)$ which is coupled between the n_u unknowns but decoupled in all other dimensions. The result is $v_{out,s}(i_x, i_y, i_z, *, i_e, *, i_o)$.
- (4) Face update. The values of $v_{out,s}(i_x, i_y, i_z, *, i_e, *, i_o)$ are stored for use downstream by the sweep.
- (5) Angles-to-moments conversion. Array $v_{out,s}(i_x, i_y, i_z, *, i_e, *, i_o)$ is transformed and added to $v_{out}(i_x, i_y, i_z, *, i_e, *)$ with another matrix-vector product depending on octant only.

The computation can be represented mathematically as follows. To simplify we assume a single octant direction (+x,+y,+z) and a single energy group. Given matrix $M_{AM} \in R^{n_a \times n_m}$ we have

$$(v_{in,s;i_x,i_y,i_z})_u = M_{AM} \cdot (v_{in;i_x,i_y,i_z})_u,$$

where i_x, i_y, i_z is the gridcell spatial coordinate and $(\cdot)_u$ denotes the restriction of the vector to gridcell unknown u . Then

$$(v_{out,s;i_x,i_y,i_z})_a = S((v_{in,s;i_x,i_y,i_z} - v_{out,s;i_x-1,i_y,i_z} - v_{out,s;i_x,i_y-1,i_z} - v_{out,s;i_x,i_y,i_z-1})_a)$$

where $(\cdot)_a$ is the restriction of the vector to angle a . Notice this represents the wavefront recursion proper. Here the function $S : R^{n_u} \rightarrow R^{n_u}$ is a solve process which for actual transport solvers is related to the spatial discretization used but for Minisweep is a synthetic function chosen to give a known analytic solution for correctness checking; in either case, the computational cost is minor, thus the actual choice is not material to algorithm performance. Finally, for matrix $M_{MA} \in R^{n_m \times n_a}$,

$$(v_{out;i_x,i_y,i_z})_u = M_{MA} \cdot (v_{out,s;i_x,i_y,i_z})_u.$$

2.3 Summary of problem axes

The problem dimensions and their respective couplings are thus summarized as follows:

- **Space:** in the x, y and z dimensions, each gridcell depends on results from three upstream cells, based on octant direction, resulting in a wavefront problem.
- **Octant:** results for different octants can be calculated independently, the only dependency being that results from different octants at the same entry of v'_{out} are added together and thus are a race hazard, depending on how the computation is done.
- **Energy:** computations for different energy groups have no coupling and can be considered separate problem instances, enabling flexible parallelization.

- **Moment, angle:** for fixed energy, octant and spatial location, the moments and angles are interrelated by small (dense) matrix-vector products.
- **Unknown:** when the problem is represented as angles, a computation is performed which may couple (only) the unknowns within the gridcell; for all other parts of the computation, elements on this axis are fully independent.

3 PARALLELIZING THE SWEEP ALGORITHM

To map the sweep algorithm to a parallel system, it is of paramount importance to minimize data motion as well as maximize parallelism, these being increasingly critical for high performance on exascale systems. As a result, the general guiding principle is that spatial dimensions must be the outermost loops, due to their sparse coupling, whereas moment, angle and unknown loops must be innermost due to the strong all-to-all couplings. The specific approach to parallelizing each axis is as follows:

Space: Spatial parallelism is based on the Koch-Baker-Alcouffe (KBA) algorithm [25]. Here the 3-D structured grid is decomposed to processors with a 2-D tiling in x and y (Figure 2). Each processor's part of the grid is decomposed into blocks along the z axis. Then a block wavefront process is applied starting at the corner block of the domain. Processors proceed in a series of parallel steps, with one block wavefront computed at each step and block face information communicated between consecutive steps. For a GPU or other accelerated processor, the KBA block described above is further decomposed into subblocks, and the computation is arranged into a series of subblock wavefronts, which are then mapped to parallel threads.




Figure 2: KBA parallel wavefront algorithm

Octant: Minisweep assigns compute threads to the eight wavefronts corresponding to the eight octant directions. The wavefronts are independent; however there is a potential race condition when two or more threads are updating the same KBA block on the same KBA block wavefront step. The solution used in Minisweep is a grid coloring approach. The KBA block is split in half along each dimension resulting in eight "semiblocks." Eight semiblock steps are taken, and for each step every one of the (up to) eight active wavefronts is assigned to a different semiblock. This is arranged so that the wavefront dependencies are satisfied. A synchronization and thread fence are required between consecutive semiblock steps.

Energy: Since the algorithm is embarrassingly parallel along the energy axis, this problem axis can be decomposed in any way: across nodes, across threads in a node, in a core or vector unit, or any combination.

Moment, Angle: The moment and angle axes are coupled to each other in an all-to-all fashion via two small matrix-vector products. The moment, angle, and unknown dimensions of the relevant arrays are ordered to be most rapidly varying, enabling efficient stride-1 memory access. When possible, the matrix-vector products

are arranged to fit entirely within a vector unit. If n_a or n_m exceed the vector unit size, a blocking strategy is used with the computation fitting within the vector unit. Importantly, the moments-to-angles transform is threaded in angle and the angles-to-moments transform is threaded in moment (otherwise a reduction across threads and/or vector lanes is required). Because threads must be reassigned between moment and angle threads, a synchronization and memory fence is required between these operations.

Unknown: No couplings exist along the unknowns axis when the small matrix-vector products are performed, therefore this axis permits some opportunity for threading here. However, for the inner solve computation in angle, each unknown may require different kinds of computations. To prevent poor use of vector units, these computations are kept serial.

4 ABSTRACT PARALLELISM MODEL

Minisweep defines and implements a set of high level abstractions to describe the parallelism of its algorithm such that these abstractions can be used by any similar application implementing the sweep motif. The goal of these abstractions is to achieve productivity and performance portability across different architectures (GPUs, Xeon Phi, CPUs, etc). These abstractions can be instantiated using general purpose parallel programming languages like CUDA, OpenMP, and OpenACC. The need for these abstractions is also suggestive of possible shortcomings in parallel programming languages and suggests the need for extensions to support applications of this type.

The large number of problem dimensions inherent in S_n transport solves makes the need for managing thread parallelism axes and hierarchical memories via use of abstraction layers acute. However, the techniques described here are applicable to many other problems requiring multidimensional parallelism, for example, batched dense linear algebra, block sparse linear solvers, and others.

Abstract machine model: Modern compute node hardware has an execution hierarchy. For example, a compute node may be composed of multiple GPUs, each with multiple cores possessing hardware threads and employing vector units composed of vector lanes. Some of these have co-located memories, for example node main memory, GPU high bandwidth memory or GPU shared memory associated with a streaming multiprocessor (SM) core. Execution threads are also associated with each level: for NVIDIA GPUs, in-warp threads execute in lock-step within a warp, in-threadblock threads are associated with an SM, and the thread grid is associated with the GPU. One can thus view a node as a hierarchy of execution units, with local memory and compute threads, and in particular hardware threads can be thought of as indexed as a tuple depending on the location in the hierarchy. Threads also have characteristics based on location, e.g., thread synchronization across different cores of a node may be impossible or much slower compared to on-core synchronization. Likewise memories at different levels have different speeds, and thread access to memories may have NUMA effects depending on the level. Note that these concepts readily apply to heterogeneous node as well as homogeneous systems.

To abstract the characteristics of heterogeneous / homogeneous architectures, we use the concept of “place,” borrowing ideas from X10 and Chapel (“locales”). A place is an abstract location (in our case, a node of a computer) where work can execute using local

executions units with local memories where threads can possibly synchronize with each other (e.g., barriers, memory fences) and access its memory. What we want to explore is how to generalize a flat model of places (e.g., that originally X10 and Chapel used to abstract an entire system) to more local hierarchical abstractions with concepts such as Hierarchical Tree Places [48] or Chapel (Sub-locales) [14] to abstract the new trends of complex memory hierarchies and accelerators. An architecture can be described as a set of “hierarchical places,” where at the higher level places communicate with each other via message passing or remote puts/gets. In the node, places can be nested in order to abstract the memory hierarchy of an architecture and its local execution units. In our abstract machine model, a nested place can access the memory of its parent place, but sometimes cannot synchronize with sibling places. This restriction is due to the way a child place can be mapped to GPU SMs, across which it is not possible to synchronize.

We use the term “place threads” to refer to execution threads associated with each of the specific places. In Minisweep, place threads are created during execution via an adaptor function which instantiates or destroys the threads for the requested places in the underlying architecture; in practice, this is implemented for example by launching a CUDA kernel or entering an OpenMP parallel region, depending on programming language implementation on the given architecture.

Abstract arrays: In Minisweep, an “abstract (multidimensional) array” is defined as an object that consists of a list of dimensions and a base pointer associated with a memory place in the hierarchy. Array elements are accessed using a multiindex through an indexing function. In this way the memory layout is controlled by an abstraction layer that can be easily modified based on the architecture. An abstract array thus has a local view within the place (and its children) at which it is allocated.

Abstract threads: Each independent variable of the science problem is assigned an abstract “threading axis” of abstract thread indices assigned to the corresponding problem axis. For example, the axis of n_e energy group values is assigned a set of (n_e or fewer) abstract compute thread indices used to compute those values. The collection of these abstract thread indices (which can be used to describe threads applied to the problem dimensions of energy, octant, y location, z location, etc.) form a tuple or thread multiindex, which will be later bound to a place thread.

Instantiation of the abstract threaded region: A fundamental operation for multithreaded or accelerated codes is entry into a fork-join parallel region. In Minisweep this is abstracted as a multithreaded region that instantiates the abstract threads. These regions can be nested and mapped to the same or different places.

Binding of abstract threads to places: A mapping of the abstract thread multiindex to a place thread multiindex is made based on the type of parallelism required. For example, since spatial wavefronts of the sweep algorithm require barriers between wavefronts, the spatial y and z dimensions must necessarily be mapped to threads within a place (e.g., threadblock on a GPU) that allows synchronization. By comparison, energy groups are fully decoupled, thus no restrictions are placed on where the abstract energy thread axis is mapped.

Parallel worksharing construct: This construct schedules work along problem axes to a set of abstract threads and executes the

work in parallel. The array index values for a given problem axis are distributed to the abstract thread indices via a block decomposition. For example, the full set of n_e energy groups is partitioned into blocks which are in turn assigned to abstract energy threads. Table 1 describes the mapping of problem axes and associated abstract threads to the place thread hierarchy for the algorithm discussed in Section 2. It is evident that the ability to synchronize a subset of threads, akin to a barrier within an MPI sub-communicator, would be of benefit, since synchronization does not scale well to large thread counts. Figure 3 is a simplified version of S_n sweep par-

```

// Abstract Arrays Allocation
AbstractArrayAllocation(vin(nx,ny,nz,ne,mm,mu): place_main)
AbstractArrayAllocation(vout(nx,ny,nz,ne,mm,mu): place_main)
AbstractArrayAllocation(neighbors(num_neighbors,ne,na,nu): place_main)
// Multithreaded Regions for Abstract Threads
AbstractMultithreadedRegion(abstract_threads_e: place_main) {
AbstractMultithreadedRegion(abstract_threads_a {
    abstracts_thread_xy: place_local) {
// Do_All Parallel Worksharing
Do_All(e in range(0,ne): abstract_thread_e) {
AbstractArrayAllocation(vs(na,nu): place_local)
do(w in range(0,w_max)) {
// Do_All Parallel Worksharing
Do_All((x,y) in wavefront(w): abstract_thread_xy) {
z = z_coord(x,y,w)
//Do_All Parallel Worksharing Matrix-Vector Product
Do_All(a in range(0,na): abstract_thread_a) {
do(u in range(0,mu)) {
vs(a,u) = 0
do(m in range(0,mm)) {
vs(a,u) += a_from_m(a,m) * vin(x,y,z,e,m,u)
}
} // End of Do_All(a)
// Do_All Parallel Worksharing
Do_All(a in range(0,na): abstract_thread_a) {
// Apply upstream wavefront dependencies
do(i neighbor of (x,y,z) in wavefront(w-1)) {
do(u in range(0,mu)) {
vs(a,u) -= neighbors(i,e,a,u)
}
}
// Computation based on unknowns
solve(vs,a)
// Save downstream wavefront dependencies
do(i neighbor of (x,y,z) in wavefront(w+1)) {
do(u in range(0,mu)) {
neighbors(i,e,a,u) = vs(a,u)
}
}
} // End of Do_All(a)
//Do_All Parallel Worksharing Matrix-Vector Product
Do_All(m in range(0,mm): abstract_thread_m) {
do(u in range(0,mu)) {
vout(x,y,z,e,m,u) = 0
do(a in range(0,na)) {
vout(x,y,z,e,m,u) += m_from_a(a,m) * vs(a,u)
}
}
} // End of wavefront loop w
AbstractArrayFree(vs)
} // End of Do_All(e)
} // End of AbstractMultithreadedRegions
AbstractArrayFree(vin, vout, neighbors)
    
```

Figure 3: Abstract representation of wavefront algorithm

problem dimension	dependency type	GPU threading	Intel Phi threading
energy	(none)	grid	OpenMP thread
octant	coloring	threadblock	CPU thread
spatial y	wavefront	threadblock	CPU thread
spatial z	wavefront	threadblock	CPU thread
moment	all-to-all	warp, serial	vector, serial
angle	all-to-all	warp, serial	vector, serial
unknown	all-to-all	warp, serial	vector, serial

Table 1: Problem dimensions mapping to thread hierarchy. Parallelization using the abstract parallelism model. The pseudocode shows the allocation of the required arrays and definition of the hierarchical thread regions, followed by nested parallel loop over energy groups, serial loop over wavefronts, parallel loop over grid-cells in the wavefront, and then the three threaded operations of moment-to-angles, solve and angles-to-moments.

5 TRANSLATION OF ABSTRACT PARALLELISM MODEL

This section shows flavors of how different models, CUDA, OpenMP and OpenACC parallelize Minisweep. The narrative also discusses what we need (referring to Section 4) and what the models lack discussed in Section 6.

5.1 CUDA

The main sweep function, Sweeper_sweep(), of Minisweep has a KBA pipeline loop to support the KBA block sweep calculations and related asynchronous face communication between nodes using MPI. For the CUDA case, faces and KBA blocks are also transferred to and from the GPU asynchronously; for systems not requiring offload, these calls do nothing. A mirrored array datatype, resembling the underlying mechanisms of OpenACC, maintains copies of an array on CPU and GPU and manages transfers; an accessor function returns the CPU or GPU pointer depending on where the computation takes place. For details, see [24]. The key kernel operation of Minisweep is the block sweep operation, in function Sweeper_sweep_block(). This is launched as a CUDA kernel or alternatively is initiated as a parallel region in OpenMP, as described above in the abstract model. Since energy groups are independent, the energy thread axis is mapped off-threadblock into the GPU thread grid; all other axes are mapped in-threadblock due to coupling requirements as described earlier. As discussed in Section 2, an execution instance of the algorithm performs a total of eight sweeps and the CUDA port of minisweep performs these sweeps.

5.2 OpenMP

OpenMP is used to run Minisweep natively on a multicore or many-core processor with OpenMP 3.1 parallel directives and the OpenMP 4.0 simd directive. The model also implements KBA. OpenMP runs with a single thread of execution until the block sweep function is encountered, at which point threads are spawned in energy, octant and the y and z spatial dimensions. The temporary arrays placed in GPU shared memory for the CUDA case are now CPU arrays, with one part of the array reserved for each compute thread. Since the OpenMP model uses a SIMD loop rather than thread numbers to access vector lanes, loops are placed in the code for the angle, moment and unknown dimensions, and each of these dimensions is assigned only a single thread; for the CUDA case, however, these axes receive multiple threads and the SIMD for loop is removed. The OpenMP port models the particle flux in all directions, i.e. performs a total of eight sweeps using the tasks concept.

5.3 OpenACC

Our OpenACC proof-of-concept for our abstract parallelism model consists of two parts: parallelizing the initialization of faces at the beginning of the sweep and parallelizing the in-gridcell computations. Also something to note is that this proof-of-concept only sweeps in one direction eight times thus computing the same number of computations one would compute using CUDA or OpenMP that is already sweeping in eight directions. OpenACC does not provide a vocabulary for tasks to follow a similar approach to that of OpenMP to sweep in eight directions, hence further study and exploration is required.

Parallelization of the face initializations involves five nested loops (spatial decomposition of gridcells, unknowns within gridcell, energy groups and angles). However, it is worth noting that PGI's OpenACC compiler only provides us with two levels of parallelism currently: gang (block) and vector (thread). The compiler is yet to thoroughly exploit the worker level of parallelism. So, in order to map a loop nest of five loops onto the accelerator to achieve full parallelization, we utilized OpenACC's collapse clause to collapse a specified number of nested loops into one large loop, which we can then map at either the gang or vector level. For Minisweep's face initializations, we collapse the outer three loops (corresponding to the unknowns and two spatial dimensions of the gridcells) and execute at the gang level. We also collapse the inner two loops (corresponding to the energy groups and angles) and execute at the vector level.

Parallelization of the in-gridcell computations in Minisweep is not as trivial, as there are data dependencies between gridcells, as mentioned in Section 2.1. To that end, we utilize the KBA parallel sweep algorithm (discussed in Section 3) in order to exploit gang-level parallelism across the x , y , and z gridcell loops. Since there is currently no existing high-level language that provides functionality for implementing this type of parallel sweep, the programmer must modify the loop nest manually in order to achieve the desired behavior. This involves creating an outer *wavefront* loop that iterates over the wavefront decomposition, as discussed in Section 2.1 and shown in Figure 1. The computations within these wavefronts can be parallelized, albeit not trivially. First we must parallelize across the inner two dimensions: y and x . This spawns a number of threads on the GPU. Within each of these threads, we calculate our z value based off of the thread's y and x values and the wavefront iteration number. Then, we can perform a bounds check to determine whether that z value is within the bounds of the wavefront being examined (denoted by the current wavefront iteration number). This allows us to exploit parallelism across gridcells, while still accounting for data-dependencies between wavefront iterations.

The embarrassingly parallel in-gridcell computations are performed for each energy group within each gridcell. We mark these computations for execution at the *vector* level. A representation of the result is shown in Figure 4. Note that this code snippet is also the serial code if one were to simply remove all the directives.

6 PROGRAMMING MODEL LIMITATIONS

6.1 General

In all cases, inadequacies of current compilers required that some code be rewritten in an unnecessarily low-level fashion to obtain correctness and/or performance. This seems to be a systemic challenge, insofar as it is difficult for compiler teams to develop mature and performant compilers for frequently changing complex processor hardware. Programming models support vectorization in different ways, leading to portability challenges. CUDA treats vector lanes as threads, whereas OpenMP uses SIMD loops and OpenACC has a vector clause for parallel loops. Such differences can lead to increased use of undesirable `ifdefs` if it is required to support these multiple programming models. Developers would prefer a single highly performant programming model with a high level of

```

/*--- Loop over wavefronts ---*/
for ( wavefront = 0; wavefront < num_wavefronts; wavefront+=1) {

    /*---KBA threading---*/
    #pragma acc loop independent gang, collapse(2)
    for( iy=0; iy<dim_y; ++iy )
        for( ix=0; ix<dim_x; ++ix ) {

        int iz = wavefront - (ix + iy);
        if ( iz >= 0 && iz <= wavefront && iz < dim_z ) {

            /*---moments to angles---*/
            #pragma acc loop independent vector, collapse(3)
            for( ie=0; ie<dim_ne; ++ie )
                for( iu=0; iu<NU; ++iu )
                    for( ia=0; ia<dim_na; ++ia ) {
                    P result = (P)0;
            #pragma acc loop seq
                for( im=0; im < dim_nm; ++im )
                    { /*---moments to angles conversion---*/ }

            /*---solve---*/
            #pragma acc loop independent vector, collapse(2)
            for( ie=0; ie<dim_ne; ++ie )
                for( ia=0; ia<dim_na; ++ia )
                    { /*---solve calculation---*/ }

            /*---angles to moments---*/
            #pragma acc loop independent vector, collapse(3)
            for( ie=0; ie<dim_ne; ++ie )
                for( iu=0; iu<NU; ++iu )
                    for( im=0; im<dim_nm; ++im ) {
            #pragma acc loop seq
                P result = (P)0;
                for( ia=0; ia<dim_na; ++ia )
                    { /*---angles to moments conversion---*/ }
            }
        }
    }
}

```

Figure 4: Sweep loop nest with OpenACC annotations

abstraction targeting all architectures rather than the need to use multiple programming models within a code.

The Minisweep code requires in several places a thread synchronization or barrier over only a subset of threads. A barrier across fewer threads could potentially run much faster in current hardware. This feature is not currently supplied by any of the programming models, though in principle a barrier across a subset of OpenMP threads could be written, and the new CUDA 9 Cooperative Groups feature may be useful here.

The Minisweep design makes it easy to change the mapping of machine threads to abstract problem threads and problem dimensions. A more challenging goal is to allow easy modification of the execution hierarchy. Such a design would allow easy loop order permutation and other loop restructuring operations, loop blocking to optimize cache use or reduce loop overheads, and on-demand reassignment of loop axes either to parallel threads or alternatively serial execution. Such changes generally require motion of significant portions of code, e.g., to optimize for loop invariant quantities. Presently this must be done by hand, and is not directly supported by programming models or imperative programming languages as currently conceived. Likewise, the use of accessor functions in Minisweep permits easy modification of memory locale and layout for an array. One must still however schedule memory transfers across the memory hierarchy manually for peak performance. Automatic transfers via paging/caching such as the CUDA Unified Memory feature and similar functionality for Intel Phi on-package memory will simplify programming for this, however past experience has shown that manual prefetching of data across the hierarchy is sometimes necessary to attain high performance. As memory layers proliferate, e.g., with inclusion of NVRAM, managing this will become more challenging.

6.2 CUDA

CUDA by nature provides a lower level programming model compared to directives-based methods. Though the CUDA runtime API provides a slightly higher abstraction level than the CUDA driver API, both cases require `ifdefs` to make a code portable between CUDA for GPUs and standard C/C++ for conventional architectures. CUDA has the advantage that vector lanes are addressed explicitly as threads, resulting in reliable vectorization. However, certain coding constructs can lead to losses in performance in unexpected ways. For example, in the course of developing the Denovo sweeper and Minisweep, it was observed that when loop bounds were passed into a CUDA kernel within a `struct`, performance was noticeably degraded compared to when passed in as scalars. Furthermore, in some cases a `for` loop that was provably one-trip at compile time ran slower than when the loop was altogether removed, necessitating use of an `ifdef` to make a single CUDA / OpenMP-SIMD code. CUDA additionally has limitations with respect to deep copy of structs and classes—since pointers in a host struct are invalid on the device—though this is improving with the support of GPU Unified Memory. In short, limitations of this nature can make it challenging to raise the abstraction level in CUDA codes and maintain performance portability with other platforms.

6.3 OpenMP

Intel Phi performance typically depends on the effective vectorization of loops, using the native `simd` directive or alternatively the OpenMP `simd` directive. In the process of porting Minisweep to the Intel Phi using the Intel compiler, challenges to loop vectorization were encountered. In one case an array accessor function needed to be flattened by removing its use of a `struct` in order to enable the loop to vectorize. In another case the compiler failed to remove a provably loop invariant quantity from a loop, inhibiting vectorization. Also, the compiler would not vectorize the outermost loop of a deep loop nest, though CUDA had no problem threading this loop. The differing treatment of vector lanes as threads by CUDA and by SIMD loops in OpenMP required the undesirable use of special case code to handle the differences. Also, CUDA generally favors larger kernels to minimize kernel launch overhead and maximize data reuse, whereas with the Intel compiler it is difficult or impossible to vectorize large, complex loops in one piece. These differences made it challenging to support the different platforms without special case programming. Overall, the difficulty of predicting a priori when a complex loop would vectorize and the need at times to rewrite code at a lower abstraction level was detrimental to writing maintainable, performance portable code. We also explored converting current the OpenMP 3.1 code to 4.5. Adapting `doacross` for this type of wavefront problem would have been a potential direction to take. However `doacross` assumes a flat memory hierarchy (shared memory) but what we need for our type of case study is to map data objects to a memory hierarchy (e.g. place and child place) that would allow the wavefront computations to be more data-centric and be scheduled where the data is.

6.4 OpenACC

Similarly to OpenMP, we faced a number of challenges when implementing our parallelization strategy discussed in Section 5.3 in OpenACC. Our parallelization efforts also identified compiler bugs

that we reported to the PGI team. The first issue was handling array accesses in the original Minisweep implementation. Accessor functions are used to calculate the address of the flattened 5-dimensional array accesses that occur throughout the `Sweeper_sweep` function, as described in Section 4. These functions returned the address of the array access in question, which was then dereferenced by the `Sweeper_sweep` function in order to perform the manipulation on the array element. OpenACC requires that we use the `routine` directive to convert these function calls to routines. However, the compiler was unable to properly generate routine code for functions utilizing external variables like these array accesses do. We had to eliminate the use of these functions and inline the calculation of the array address into the array accesses within the given loops, resulting in a more traditional array access. Unfortunately this is detrimental to efforts to raise the abstraction level of the code.

Another issue was related to loop bounds. In Minisweep, input parameters are stored in a globally defined `struct`. Since these values are representing the sizes of each dimension of the application, they are used later as loop bounds. However, while parallelizing, OpenACC does not assume that no aliasing is being done since this `struct` is defined globally (out of scope). There are two simple solutions to this issue. First, the compiler flag `-Msafept` can be used to specify that there is no aliasing. However, this would not be the best option for this application as there is some sort of pointer aliasing present elsewhere. Instead, we simply extract the value of the dimension being used for a given loop bound and store it in an integer variable prior to the start of the accelerated loop.

The final issue we faced was identified as a compiler bug in PGI OpenACC 17.10. OpenACC can use `kernels` or `parallel` to generate code from an accelerated region. With `kernels`, the onus is on the compiler to check for dependencies and generate code, whereas with `parallel`, the onus is on the programmer; failure to do so will result in inaccurate results. In our case, we observed that even though we used the `parallel` directive, the compiler still performed dependency checks as if we had used `kernels` directive. We confirmed this behavior by parallelizing a loop with a known dependency. The compiler generated parallel code but showed incorrect results. We then added a `collapse` clause to the end of this loop directive, and we specified that it should be collapsed with the next loop in the loop nest, which contained no such dependencies. The result here was that the compiler reported that it parallelized this collapsed loop nest, but the results of the computation were accurate. Due to the increased runtime, we were able to conclude that the inner loop was indeed executing in parallel, but the outer loop was executing in serial despite what the compiler had reported. All of these issues were easily overcome in practice, but identifying them presented significant challenges along the way.

7 EVALUATION & RESULTS

As a validation of portability, Table 3 shows Minisweep results for one GPU of the Titan Cray XK7 system (CUDA), one GPU of the Summitdev IBM Minsky system (CUDA) and one node of the Percival Cray XC40 KNL system (self-hosted OpenMP 4.0). The problem solved has $n_e = 64$, $n_a = 32$, and $n_u = 4$, with $n_x, n_y, n_z = 32$. The codes are not fully optimized, in particular one of the inner loops for the OpenMP-KNL case did not vectorize. However, all cases across different hardware and software environments attained

Machine	CPU	GPU
NVIDIA PSG (V100)	Intel Xeon E5-2698 v3 (16 cores)	NVIDIA Tesla V100 (16GB HBM2)
NVIDIA PSG (P100)	Intel Xeon E5-2698 v3 (16 cores)	NVIDIA Tesla P100 (16GB HBM2)
NVIDIA PSG (K40)	Intel Xeon E5-2690 v2 (10 cores)	NVIDIA Tesla K40 (12GB GDDR5)
ORNL Titan	AMD Opteron 6274 (16 cores)	NVIDIA Tesla K20X (6GB GDDR5)
ORNL Summitdev	IBM POWER8 (10 cores)	NVIDIA Tesla P100 (16GB HBM2)
ORNL Percival	Intel Xeon Phi 7230 (64 cores)	N/A

Table 2: Specifications of the nodes in the systems we used to test different configurations of Minisweep.

a similar 4-5% of peak flop rate, a typical figure for this algorithm which has significant memory accesses, register usage and integer index calculations. This result suggests that the code is in fact performance portable, since reasonable performance is reached for all systems.

System	Cores (SMs)	GF/s peak	GF/s	% peak GF/s
Titan(K20X)	14	1311	55.9	4.26
Summitdev(P100)	56	5312	244.8	4.61
Percival(Phi7230)	64	2662	124.9	4.69

Table 3: Comparative performance on several platforms.

We evaluate the effectiveness of our abstract wavefront parallelism model by comparing the runtimes of our parallel implementations of Minisweep (described in Section 5) to the runtime of a serial version of the code on multiple HPC systems. Table 2 describes the hardware available on nodes of each system. Note that the NVIDIA Professional Service Group (PSG) machines and the ORNL Titan machine are existing state-of-the-art HPC systems, while ORNL Summitdev is a development cluster representative of what hardware will be present on nodes in ORNL’s next-gen supercomputer *Summit* [27]. We also utilized the PSG cluster’s V100 node, which houses NVIDIA’s next-generation GPU, which will be present on nodes in *Summit*. We used PGI’s 17.10 compiler to compile our OpenACC and OpenMP. We have also used GCC 6.3.0 for and ICC 17.0 for OpenMP codes. Compiling the code using Intel’s OpenMP compiler was not successful and required code restructuring to take advantage of SIMD in minisweep.

Our experimental configuration is a representative example of what a real run of Minisweep within the Denovo radiation transport code looks like. Our problem dimensions are designed to be as large as we can fit on a single GPU: $n_e = 64$, $n_a = 32$, and $n_u = 4$, with $n_x, n_y, n_z = 32$, on K20x/K40 and $n_x, n_y, n_z = 64$ on P100/V100.

Figures 5 and 6 present the results when running different implementations of Minisweep using this configuration in the form of speedups over the baseline serial implementation on existing HPC systems. Note that the speedup results presented were obtained by calculating the average of a series of runs for each implementation. There are a few notable results. First, our multicore CPU GCC’s OpenMP (3.1) and OpenACC implementations yield favorable speedups. Note that GCC’s OpenMP performed better than PGI’s OpenMP. As mentioned in Section 5.3, we have currently parallelized the in-cell computations, as well as the spatial decomposition utilizing the KBA parallel sweep algorithm to resolve data

dependencies, as discussed in Section 2.1. This implementation boasts a larger speedup than our OpenMP GCC version, as well as our CUDA configuration when parallelized over the same problem dimensions. Our OpenACC KBA configurations yields an additional layer of parallelism across spatial dimensions and shows a much larger speedup compared to configurations which only execute in-cell computations in parallel. This leads us to conclude that there is additional performance to be gained, albeit not trivial to implement. It is also worth noting that our OpenACC implementation running on NVIDIA’s next-generation Volta GPU boasts an 85.06x speedup over serial code, which is larger than the 83.72x speedup over the same serial implementation achieved by CUDA. This supports our claim that our proposed extension to existing high-level programming models is worthwhile, both from a performance standpoint, as well as a programming productivity standpoint. Currently, without major code modification, this challenge cannot be overcome.

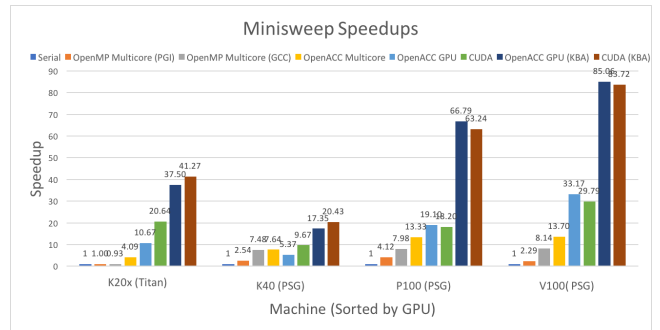


Figure 5: Minisweep’s speedups over serial using different runtime configurations. CUDA version is parallelized along the same dimensions as the OpenACC GPU configuration. The corresponding KBA configurations utilize the KBA blocking method for additional parallelism across spatial dimensions.

Absolute runtimes for GPU configurations utilizing the KBA parallel sweep algorithm are presented in Figure 7. As shown, our OpenACC GPU implementation performs well compared to its CUDA counterpart in all cases. In addition to its excellent GPU performance, it is worth noting that this same OpenACC implementation was used to obtain results on our multicore CPU platforms by simply recompiling and specifying a different target. No additional code modifications were necessary to achieve this demonstration of portability. We contend that this provides additional evidence for the importance of an extension that would allow us to parallelize the outer spatial dimensions, yielding additional parallelism across




Figure 6: Minisweep’s speedups over serial using different runtime configurations: ORNL’s next-gen Summitdev cluster.

gridcells without requiring a major coding effort on the part of the programmer. As stated in Section 4, this type of abstraction will benefit any wavefront-type code that performs some type of spatial dimension sweep. Our OpenACC proof-of-concept of such an abstraction demonstrates this on a real-world wavefront-type application used at a major national laboratory. Since these types of codes are very common in computational scientific applications, we contend that this contribution has far-reaching implications for modern-day HPC applications.




Figure 7: Absolute runtimes (sec) of OpenACC & CUDA experiments on all GPUs used. Note that the V100/P100 problem size is an order of magnitude > K40/K20x configuration, as mentioned earlier.

Scientifically, neutron flow simulation can take a considerable amount of computing time. If not for speeding up these simulation runs using accelerators, scientists will have to resort to simulating a model that is potentially less accurate leading to questionable results. This can be quite a risky task to rely on for scientists working in nuclear reactor facilities.

8 RELATED WORK

Considered as far back as 1974 by Lamport [28], wavefront computations are found in linear equation solvers [31, 35], gene sequence alignment [43] and radiation transport [11, 25], iterative solution methods [38], particle physics simulations [26], and parallel solution of triangular systems of linear equations [22]. Smith Waterman, a local sequence alignment algorithm, has mapped wavefront algorithm to GPUs [39], on Cell BE [45] and on FPGAs [15, 49]. ASCI Sweep3D wavefront application undergoes rapid succession of

wavefront and solves a 1-group neutron transport problem on IBM Blue Gene/P machine [23] by using blocking techniques. Preliminary studies to use TBB, Cilk, CnC, and OpenMP 3.0 for wavefront in [19] indicate that a higher-level template is required for less experienced users. AWE Chimaera [34], NAS-LU [10] use a variation called 'hyperplane' algorithm [29] and [36] discusses acceleration of generalized pipeline wavefront applications.. Proxy apps such as Kripke [4], SNAP [7] (mimicking communication patterns of PARTISN [13] transport code) are wavefront codes investigating different data layout patterns and parallelism. Coarray Fortran-based Sweep3D’s comparable performance to that of the MPI is discussed in [18]. Approaches using loop skewing [46] to derive the wavefront method of execution of nested loops is discussed in CHiLL [16], a polyhedral compiler transformation framework. Other related work include [44, 47]. Other approaches are High Productive Computing System (HPCS) languages: Chapel [14], X10 [33] and Fortress [8] but do not offer enough abstractions or vocabulary for heterogeneous platforms. HTP [48] proposed a hierarchical tree place to map to an architecture and schedules task to different different nodes in the tree.

However most of the above discussed strategies are solutions to specific problem types, or incurs steep learning curve thus making them not quite so favorable to easily adapt for large scale applications. Taking this up as a challenge, our work proposes solution using directives and demonstrate the parallelization of a multi-dimensional Minisweep using OpenACC on different platforms. The parallelization process revealed shortcomings in current directive-based model that we address by proposing an abstract model for expressing wavefront parallelisms in programming models.

9 CONCLUSION & FUTURE WORK

This paper examines the challenges faced when porting a wavefront application to state-of-the-art HPC systems using directives using Minisweep as the case study. We present a performance-portable OpenACC implementation of Minisweep, as well as an analysis of this implementation’s performance compared to optimized multicore CPU and GPU implementations of the same code using OpenMP and CUDA, respectively. Results demonstrate that utilizing high-level parallel programming abstractions, such as OpenACC, can achieve comparable performance to low-level, optimized parallel implementations. All of these designs and implementations are reflections of an abstract parallelism model that we propose for wavefront algorithms. This model motivates enhancing programming models with software abstractions to parallelize wavefront problems on multiple platforms, while minimizing programmer overhead.

As part of our ongoing work, we plan on applying our findings outlined by this work on additional applications to further demonstrate the efficacy of our abstract parallelism model, formulating and automating the process of transforming loop nests into the KBA method discussed in Section 3, and investigating techniques for performing wavefront sweeps across multiple GPUs and/or nodes in HPC systems.

ACKNOWLEDGMENT

This material is based upon work supported by the U.S. Department of Energy, Office of science, and this research used resources of the

Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

The authors would also like to thank NVIDIA for donating us a Titan Xp and providing us access to their Professional Service Group (PSG) machine that we have used for this work.

REFERENCES

- [1] 1995. The ASCI SWEEP3D README File. (1995). <http://www.ccs3.lanl.gov/PAL/software.shtml> [Online; accessed 24-June-2014].
- [2] 2014. Scientists successfully test code that models neutrons in reactor core. *R&D Magazine* (2014). <https://www.rdmag.com/news/2014/02/scientists-successfully-test-code-models-neutrons-reactor-core>.
- [3] 2014. Supercomputer team wins award for core work. *World Nuclear News* (2014). <http://www.world-nuclear-news.org/NN-Supercomputer-team-wins-award-for-core-work-0707147.html> [Online; accessed 15-August-2017].
- [4] 2017. Kripke. (2017). <https://codesign.llnl.gov/kripke.php> [Online; accessed 15-August-2017].
- [5] 2017. ORNL-CEES. <https://github.com/ORNLCES/Profugus>. (2017).
- [6] 2017. Quadrillions of calculations per second for fusion. *ITER Newline* (2017). https://www.iter.org/ajax/www/pop/wd_700/lang/_urldepth_0/id_newsline_ofinterest-670 [Online; accessed 15-August-2017].
- [7] 2017. SNAP: SN (Discrete Ordinates) Application Proxy. (2017). <https://github.com/lanl/SNAP> [Online; accessed 15-August-2017].
- [8] Eric Allen, David Chase, Joe Hallett, Victor Luchangco, Jan-Willem Maessen, Sukyoung Ryu, Guy L Steele Jr, Sam Tobin-Hochstadt, Joao Dias, Carl Eastlund, et al. 2005. The Fortress language specification. *Sun Microsystems* 139, 140 (2005), 116.
- [9] James A Ang, Richard F Barrett, Robert E Benner, D Burke, C Chan, J Cook, David Donofrio, Simon D Hammond, Karl S Hemmert, SM Kelly, et al. 2014. Abstract machine models and proxy architectures for exascale computing. In *Hardware-Software Co-Design for High Performance Computing (Co-HPC), 2014*. IEEE, 25–32.
- [10] David H Bailey, Eric Barszcz, John T Barton, David S Browning, Robert L Carter, Leonardo Dagum, Rod A Fatoohi, Paul O Frederickson, Thomas A Lasinski, Rob S Schreiber, et al. 1991. The NAS parallel benchmarks. *The International Journal of Supercomputing Applications* 5, 3 (1991), 63–73.
- [11] Christopher Baker, Gregory Davidson, Thomas M Evans, Steven Hamilton, Joshua Jarrell, and Wayne Joubert. 2012. High performance radiation transport simulations: preparing for Titan. In *SC 2012 International Conference for. IEEE*, 1–10.
- [12] Christopher G. Baker, Gregory G. Davidson, Thomas M. Evans, Steven P. Hamilton, Joshua J. Jarrell, and Wayne Joubert. 2012. High Performance Radiation Transport Simulations: Preparing for TITAN. In *Proceedings of Supercomputing Conference SC12*.
- [13] Randal S. Baker. 2014. PARTISN on Advanced/Heterogeneous Processing Systems. (2014). <http://permlink.lanl.gov/object/tr?what=info:lanl-repo/lareport/LA-UR-13-20948> [Online; accessed 24-June-2014].
- [14] Bradford L Chamberlain, David Callahan, and Hans P Zima. 2007. Parallel programmability and the chapel language. *The International Journal of High Performance Computing Applications* 21, 3 (2007), 291–312.
- [15] Sunita Chandrasekaran, Shilpa Shanbagh, Ramkumar Jayaraman, Douglas L Maskell, and Hui Yan Cheah. 2013. C2FPGA?A dependency-timing graph design methodology. *JPDC* 73, 11 (2013), 1417–1429.
- [16] Chun Chen, Jacqueline Chame, and Mary Hall. 2008. *CHILL: A framework for composing high-level loop transformations*. Technical Report. Technical Report 08-897, U. of Southern California.
- [17] M. P. Clay, D. Buaria, and P. K. Yeung. 2017. Improving Scalability and Accelerating Petascale Turbulence Simulations Using OpenMP. <http://openmpcon.org/conf2017/program/>. (2017). To Appear.
- [18] Cristian Coarfa, Yuri Dotsenko, and John Mellor-Crummey. 2006. Experiences with Sweep3D implementations in Co-array Fortran. *The Journal of Supercomputing* 36, 2 (2006), 101–121.
- [19] Antonio J Dios, Angeles G Navarro, Rafael Asenjo, Francisco Corbera, and Emilio L Zapata. 2011. A case study of the task-based parallel wavefront pattern. In *PARCO*. 65–72.
- [20] T. M. Evans, W. Joubert, S. P. Hamilton, S. R. Johnson, J. A. Turner, G. G. Davidson, and T. M. Pandya. 2015. Three-Dimensional Discrete Ordinates Reactor Assembly Calculations on GPUs. In *ANS 2015*.
- [21] Roberto Gomperts. 2016. Quantum Chemistry on GPUs. <http://images.nvidia.com/content/tesla/pdf/quantum-chemistry-may-2016-mb-slides.pdf>. (2016).
- [22] Michael T Heath and Charles H Romine. 1988. Parallel solution of triangular systems on distributed-memory multiprocessors. *SIAM J. Sci. Statist. Comput.* 9, 3 (1988), 558–588.
- [23] Accelerated Strategic Computing Initiative. 1995. The ASCI sweep3d benchmark code. (1995).
- [24] Wayne Joubert. 2017. Minisweep. <https://github.com/wdj/minisweep>. (2017).
- [25] K.R. Koch, R.S. Baker, and R.E. Alcouffe. 1992. Solution of the first-order form of the 3-D discrete ordinates equation on a massively parallel processor. *Transactions of the American Nuclear Society* 65 (1992), 198–199.
- [26] Kenneth R Koch, Randal S Baker, and Raymond E Alcouffe. 1992. Solution of the first-order form of the 3-D discrete ordinates equation on a massively parallel processor. *Transactions of the American Nuclear Society* 65, 108 (1992), 198–199.
- [27] Oak Ridge National Lab. 2017. SUMMIT. <https://www.olcf.ornl.gov/summit/>. (2017).
- [28] Leslie Lamport. 1974. The Parallel Execution of DO Loops. *Commun. ACM* 17(2) (1974), 83–93.
- [29] Leslie Lamport. 1974. The parallel execution of DO loops. *Commun. ACM* 17, 2 (1974), 83–93.
- [30] Edward W. Larsen and Jim E. Morel. 2010. Advances in Discrete-Ordinates Methodology. In *Nuclear Computational Science: A Century in Review*, Youssry Azmy and Enrico Sartori (Eds.), Springer, New York, Chapter 1, 1–84.
- [31] Ruipeng Li and Yousef Saad. 2013. GPU-accelerated preconditioned iterative linear solvers. *Journal of Supercomputing* 63(2) (February 2013), 443–466. <https://doi.org/10.1007/s11227-012-0825-3> <http://link.springer.com/article/10.1007/978-1-1227-012-0825-3> [Online; accessed 24-June-2014].
- [32] Bronson Messer, Ed D’Azevedo, Judy Hill, Wayne Joubert, Mark Berrill, and Christopher Zimmer. 2016. MiniApps derived from production HPC applications using multiple programming models. *The International Journal of High Performance Computing Applications* 0, 0 (2016), 1094342016668241. <https://doi.org/10.1177/1094342016668241> arXiv:<http://dx.doi.org/10.1177/1094342016668241>
- [33] Josh Milthorpe, V Ganesh, Alistair P Rendell, and David Grove. 2011. X10 as a parallel language for scientific computation: Practice and experience. In *IPDPS, 2011 IEEE International. IEEE*, 1080–1088.
- [34] Gihan R Mudalige, Mary K Vernon, and Stephen A Jarvis. 2008. A plug-and-play model for evaluating wavefront computations on parallel architectures. In *IPDPS, IEEE International Symposium on. IEEE*, 1–14.
- [35] S.J. Pennycook, S.D. Hammond, G.R. Mudalige, S.A. Wright, and S.A. Jarvis. 2012. On the Acceleration of Wavefront Applications using Distributed Many-Core Architectures. *Comput. J.* 55, 2 (2012), 138–153. <https://doi.org/10.1093/comjnl/bxr073> <http://comjnl.oxfordjournals.org/content/55/2/138> [Online; accessed 24-June-2014].
- [36] Simon J Pennycook, Simon D Hammond, Gihan R Mudalige, Steven A Wright, and Stephen A Jarvis. 2011. On the acceleration of wavefront applications using distributed many-core architectures. *Comput. J.* 55, 2 (2011), 138–153.
- [37] David F Richards, Ryan C Bleile, Patrick S Brantley, Shawn A Dawson, Michael Scott McKinley, and Matthew J O’Brien. 2017. Quicksilver: A Proxy App for the Monte Carlo Transport Code Mercury. In *Cluster Computing (CLUSTER), 2017 IEEE International Conference on. IEEE*, 866–873.
- [38] François Roddier and Claude Roddier. 1991. Wavefront reconstruction using iterative Fourier transforms. *Applied Optics* 30, 11 (1991), 1325–1327.
- [39] Edans Flavius de O Sandes and Alba Cristina MA de Melo. 2011. Smith-waterman alignment of huge sequences with gpu in linear space. In *IPDPS, 2011 IEEE International. IEEE*, 1199–1211.
- [40] Sunil Sathe. 2016. Accelerating the ANSYS Fluent R18.0 Radiation Solver with OpenACC. <https://tinyurl.com/yacxh5g7>. (2016).
- [41] Will Sawyer, Guenther Zaengl, and Leonidas Linardakis. 2014. Towards a multi-node OpenACC Implementation of the ICON Model. In *EGU General Assembly Conference Abstracts*, Vol. 16.
- [42] John Shalf. 2013. Computer Architecture for the Next Decade: Adjusting to the new normal for computing. *EEHPC Workshop* (2013).
- [43] T. F. Smith and M. S. Waterman. 1981. Identification of Common Molecular Subsequences. *Journal of Molecular Biology* 147(1) (1981), 195–197.
- [44] Anand Venkat, Mahdi Soltan Mohammadi, Jongsoo Park, Hongbo Rong, Rajkishore Barik, Michelle Mills Strout, and Mary Hall. 2016. Automating wavefront parallelization for sparse matrix computations. In *Proc. of SC16*. IEEE Press, 41.
- [45] Adrianto Wirawan, Kwoh Chee Keong, and Bertil Schmidt. 2007. Parallel DNA sequence alignment on the cell broadband engine. In *International Conference on PPAI*. Springer, 1249–1256.
- [46] Michael Wolfe. 1986. Loops skewing: The wavefront method revisited. *International Journal of Parallel Programming* 15, 4 (1986), 279–293.
- [47] Justin M Wozniak, Timothy G Armstrong, Michael Wilde, Daniel S Katz, Ewing Lusk, and Ian T Foster. 2013. Swift/t: Large-scale application composition via distributed-memory dataflow processing. In *Cluster, Cloud and Grid Computing (CCGrid), 2013 13th IEEE/ACM International Symposium on. IEEE*, 95–102.
- [48] Yonghong Yan, Jisheng Zhao, Yi Guo, and Vivek Sarkar. 2009. Hierarchical Place Trees: A Portable Abstraction for Task Parallelism and Data Movement.. In *LCPC*, Vol. 10. Springer, 172–187.
- [49] Peiheng Zhang, Guangming Tan, and Guang R Gao. 2007. Implementation of the Smith-Waterman algorithm on a reconfigurable supercomputing platform. In *Proc. of HPRCTA: held in conjunction with SC07*. ACM, 39–48.