From OpenACC to OpenMP 4: Toward Automatic Translation

Nawrin Sultana
nzs0034@auburn.edu

Alexander Calvert
afc0005@auburn.edu

Jeffrey L. Overbey
joverbey@auburn.edu

Department of Computer Science and Software Engineering
Auburn University
Auburn University, AL, USA

Galen Arnold
gwarnold@illinois.edu

National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign
Urbana, IL, USA

ABSTRACT

For the past few years, OpenACC has been the primary directive-based API for programming accelerator devices like GPUs. OpenMP 4.0 is now a competitor in this space, with support from different vendors. In this paper, we describe an algorithm to convert (a subset of) OpenACC to OpenMP 4; we implemented this algorithm in a prototype tool and evaluated it by translating the EPCC Level 1 OpenACC benchmarks. We discuss some of the challenges in the conversion process and propose what parts of the process should be automated, what should be done manually by the programmer, and what future research and development is necessary in this area.

CCS Concepts

•Computing methodologies → Parallel programming languages; •Software and its engineering → Software maintenance tools;

Keywords

accelerators; GPUs; OpenACC; OpenMP; translation

1. INTRODUCTION

While CUDA and OpenCL are the canonical APIs for programming accelerator devices, scientific programmers are increasingly turning to directive-based APIs, which are more concise and declarative. Since it first appeared in 2011, OpenACC has been the standard directive-based API for scientific programmers targeting NVIDIA GPUs. OpenMP 4.0 added directives for accelerators, which were inspired by

OpenACC. Both OpenACC and OpenMP are supported in both C/C++ and Fortran.

OpenACC is supported by PGI's compilers. Intel's compilers support OpenMP 4. Cray recently announced that their support for OpenACC will be frozen at version 2.0, and future development will focus on OpenMP 4 device directives. Recent versions of GCC support OpenACC; recent versions of Clang support OpenMP 4.

So, the two APIs are rivals, at least for the time being.

Since nearly every compiler supports OpenMP for multi-core parallelism, it is reasonable to expect that the device directives added in OpenMP 4 will eventually supplant OpenACC. However, that certainly has not happened yet—and given the strong backing for OpenACC from PGI/NVIDIA, there may continue to be some rivalry between the two specifications.

In the XSEDE ecosystem, the PGI compilers are installed on Comet (at SDSC), allowing OpenACC to be used for programming its NVIDIA GPU nodes. At TACC, portions of Stampede are being upgraded to the second-generation Intel Xeon Phi (codenamed “Knight's Landing”), which can be programmed using OpenMP 4.

Since OpenACC inspired the device directives added in OpenMP 4.0, the two APIs are quite similar. In many cases, there is a 1–1 mapping between directives. However, there are some important differences as well.

In this paper, we investigate the feasibility of a tool to semi-automatically convert OpenACC code to instead use OpenMP 4 device directives. We have prototyped a tool that automates a large part of the conversion process; we have also identified parts of the conversion process that must be (or should be) left to the programmer.

We focused on converting OpenACC to OpenMP 4 *that targets the same accelerator device*—an NVIDIA GPU. We used PGI C/C++ 16.1 as our OpenACC compiler, and we used Clang 3.8 as our OpenMP compiler.

We chose this setup because it allowed us to investigate the changes in semantics and performance that result from changing the compiler and API, rather than those that result from a change in target hardware. Admittedly, for XSEDE users, a more realistic scenario would be converting Open-

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

XSEDE16, July 17 - 21, 2016, ,

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ISBN 978-1-4503-4755-6/16/07...\$15.00

DOI: <http://dx.doi.org/10.1145/2949550.2949654>

ACC on Comet (which targets NVIDIA GPUs) to OpenMP on Stampede (which targets Xeon Phi's). However, changing the target hardware in addition to the compiler and API introduces additional complexity. Keeping the target hardware fixed allowed us to focus on the differences between OpenACC and OpenMP, rather than differences in the accelerator devices themselves.

We used the EPCC Level 1 benchmarks as our test cases, along with a few miscellaneous examples, and we focused on the subset of OpenACC used by these programs. Our results show that *some* aspects of OpenACC-to-OpenMP conversion can be automated, but some manual intervention is essential. Also, both APIs rely heavily on the compiler, potentially making a change in compilers one of the biggest obstacles in converting code from OpenACC to OpenMP 4.

The remainder of this paper is organized as follows. Section 2 gives a brief history and overview of OpenACC and OpenMP. Section 3 discusses translation from OpenACC to OpenMP, focusing on complications that prohibit completely automatic translation. Section 4 describes our tool and the algorithm it uses to convert OpenACC to OpenMP device directives. Section 5 presents empirical data describing the conversion. Section 6 describes limitations of the current tool and directions for future work.

2. BACKGROUND

2.1 History of OpenACC and OpenMP

OpenACC has a short history. It was originally developed by PGI, Cray, and NVIDIA, with version 1.0 being released in 2011. Version 2.0 was released in 2013, adding support for procedure calls and nested parallelization (enabled by new CUDA capabilities), among other features. Version 2.5 was a more minor revision, made in 2015 [8].

OpenMP dates back to the late 1990s, when the OpenMP Architecture Review Board was formed to develop a specification for directive-based parallelism that could be used commonly across multiple vendors' multiprocessors. The first OpenMP specification was released for Fortran in 1997 and for C/C++ in 1998. The OpenMP specification has undergone several revisions since that time. Notably, version 4.0 added support for accelerators in 2013; an update to version 4.5 followed in 2015 [10]. The so-called *device directives* or *target directives* added in OpenMP 4 to support accelerators were based on OpenACC. However, as we will see, the two APIs are not identical.

2.2 CUDA Execution Model

OpenACC and OpenMP device directives are both designed to offload computations to accelerator devices—i.e., massively parallel devices like GPUs and Xeon Phis designed for data parallel computations. Each API has its own execution model, which is intentionally abstract to avoid coupling it to the specifics of any one device.

For the purposes of this paper, we will only focus on one type of accelerator device: NVIDIA GPUs. To understand the behavior of OpenACC and OpenMP constructs on these devices, it is helpful to understand the execution model of NVIDIA GPUs: the CUDA execution model.

CUDA-capable devices (i.e., NVIDIA GPUs) execute *kernels*, functions that execute on the device by running many GPU threads concurrently. Only one kernel executes at a time. All of the threads in a kernel are collectively called a

grid. A grid is comprised of one or more *thread blocks*, and each thread block is comprised of one or more *threads*.

In hardware, each thread block is assigned to one Streaming Multiprocessor (SM). Each SM consists of several Streaming Processors (SPs); each thread of a thread block is assigned to one of the SPs within the SM. Thread blocks are divided into 32-thread groups called *warps*. Instructions are fetched and executed on a per-warp basis, not on a per-thread basis, so the 32 threads within each warp execute instructions in single instruction, multiple data (SIMD) fashion. The SM context switches from one warp to another at every instruction issue ("zero-overhead thread scheduling"), so while threads within a warp execute in lockstep, threads in *different* warps may not. Threads within a block can coordinate using barrier synchronization, whereas threads in different blocks cannot [6].

2.3 OpenMP 4 Directives

Parallel. OpenMP's fundamental construct is the parallel construct. When a parallel construct is encountered by a thread, that thread creates a *team* of new threads, becoming that team's master. The threads are assigned numbers, with the master given number 0. The threads each execute a copy of the parallel region's code.

For. Inside a parallel region, a *for* construct may be specified. A *for* region is associated with a *for* loop, and upon encountering the region, a different thread from the current team is assigned to each iteration of the *for* loop. A parallel *for* region is also defined as a shorthand for a parallel region containing nothing but a single *for* region. OpenMP 4.0 adds a similar construct, *simd*, which divides work of the associated loop up among SIMD lanes of the current thread.

Target Data. A target data region may be specified as well. A target data region creates a data environment for execution. It may specify a device clause, which specifies a device to execute the environment on, and a map clause, which specifies a mapping of data from the host to the device.

Target Teams. A target teams construct is equivalent to a target construct containing a teams construct. The target teams construct creates a data environment and a *league* of thread teams to execute it.

Distribute. A distribute construct creates a league of teams and specifies that contained loops will be executed by those teams.

Distribute Parallel For. A distribute parallel *for* construct combines *distribute*, *parallel*, and *for* construct to specify that a loop can be executed in parallel by multiple threads in multiple teams.

The OpenMP specification [10] describes these directives in detail.

2.4 OpenACC Directives

OpenACC has fewer constructs. This is in part because OpenACC is specifically for use with massively parallel accelerator hardware, where OpenMP supports accelerators as well as multicore CPUs. However, many of OpenACC's constructs are analogous to OpenMP ones.

Parallel. OpenACC also has a parallel construct. A parallel construct, when encountered, creates *gangs of workers* to execute the region on the accelerator.

Loop. A loop construct is also available, specifying the type of parallelism for the associated loop. There is also the

<i>acc-directive</i>	→	<code>#pragma acc data <i>data-clauses</i></code>
		<code>#pragma acc parallel <i>par-clauses</i></code>
		<code>#pragma acc loop <i>type</i> <i>loop-clauses</i></code>
<i>type</i>	→	<code>gang vector</code>
		<code>gang</code>
		<code>vector</code>
		<code>seq</code>
<i>data-clause</i>	→	<code>copyin(<i>vars</i>)</code>
		<code>copyout(<i>vars</i>)</code>
		<code>copy(<i>vars</i>)</code>
		<code>create(<i>vars</i>)</code>
<i>par-clause</i>	→	<code>num_gangs(<i>n</i>)</code>
		<code>vector_length(<i>n</i>)</code>
		<code>private(<i>vars</i>)</code>
		<code>reduction(<i>reductions</i>)</code>
<i>loop-clause</i>	→	<code>independent</code>
		<code>private(<i>vars</i>)</code>
		<code>reduction(<i>reductions</i>)</code>

Figure 1: OpenACC subset grammar accepted by the translation tool. Syntax is omitted for *vars*, *reductions*, and natural numbers *n*. The nonterminals *data-clauses* and *loop-clauses* correspond to sequences of zero or more phrases derived from *data-clause* and *loop-clause*, respectively. In addition to the directives shown, `#pragma acc parallel loop` is supported, and treated as syntactic sugar for a loop directive nested directly under a parallel directive.

parallel loop abbreviation, a shorthand for a loop construct immediately inside a parallel construct.

Kernels. The kernels construct specifies a region that is to be divided into a sequence of kernels for accelerator execution. The loop and shorthand kernels loop are available for kernels directives as they are for parallel directives.

Data. A data directive is also defined. It specifies for a region what data should be copied into the accelerator upon entry to the region and copied out on exit and defines what data is to be allocated on the device for the duration of the region.

A more complete description of these directives is available in the OpenACC specification [8].

3. TRANSLATION CONSIDERATIONS

Our goal was to create a tool to convert OpenACC code to equivalent OpenMP code, automating as much of the process as possible. For the purposes of prototyping, we focused on translating C code (as opposed to C++ or Fortran).

3.1 A Subset of OpenACC

For the purposes of prototyping, we designed our tool to accept the subset of OpenACC described by the grammar in Figure 1. This includes all of the directives and clauses used by the example programs in our test suite (notably, the EPCC Level 1 benchmarks). While it omits some important features—asynchronous execution, unstructured data lifetimes, etc.—it does correspond to a useful, commonly-used subset of OpenACC.

3.2 Considerations in Translation

In theory, the goals of the present paper are modest: We are interested in translating a subset of OpenACC (Fig-

```
#pragma acc data copyin(A[0:N], B[0:N]), \
               copyout(C[0:N])
#pragma omp target data map(to:A[0:N], B[0:N]), \
               map(from:C[0:N])
{
    #pragma acc parallel loop
    #pragma omp target teams distribute parallel for
    for (int i = 0; i < N; ++i)
    {
        C[i] = A[i] + B[i];
    }
}
```

Figure 2: Vector addition kernel.

ure 1) to OpenMP 4, targeting the same accelerator devices (NVIDIA GPUs). In theory, the two APIs are very closely aligned, and the translation should be straightforward.

In practice, the situation is not so simple. The OpenACC and OpenMP specifications define an *abstract* execution model. They do not dictate how that model is mapped to specific hardware devices. This is left to the compiler. Our OpenACC programs were compiled using PGCC 16.1, the standard commercial compiler for OpenACC. We compiled OpenMP 4 using a custom build of Clang 3.8 [9] (since PGCC does not yet support OpenMP 4).¹ As we will show, directives that appear similar may be translated quite differently by the two compilers.

3.2.1 Work Distribution for Single Loops

Perhaps the simplest example of directive-based accelerator code is the vector addition kernel shown in Figure 2. It adds two vectors *A* and *B* on an accelerator device, storing the sum in *C*. While the directives direct the compiler to parallelize the loop on the accelerator, the compiler can decide (in this case) how to divide the loop iterations among threads and thread blocks on the CUDA device. When $N = 10000$:

- PGCC generates a grid of 79 thread blocks with 128 threads per block; kernel execution time is about $3\ \mu\text{s}$ on a Tesla K40.
- Clang generates a grid of 15 thread blocks with 1024 threads per block; kernel execution time is about $30\ \mu\text{s}$ on a Tesla K40.²

It is also worth noting that PGCC generates code where all threads are active in all thread blocks except the last (presumably to minimize control divergence). Clang generates code that assigns an approximately equal number of threads to each thread block, so inactive threads may be present in several blocks.

3.2.2 Work Distribution for Loop Nests

More important is the handling of loop nests, such as the simple doubly nested loops in Figure 3(a).

In OpenACC, the compiler has some freedom to decide how to parallelize the loop nest. PGCC will typically run

¹Cray’s compilers support both, but there are currently no Cray resources in the XSEDE ecosystem.

²Adding `num_teams(79) thread_limit(128)` to the OpenMP code (to match the grid and block dimensions of the OpenACC code) actually increases the kernel execution time to about $80\ \mu\text{s}$.

OpenACC (Original):

```
#pragma acc parallel loop
for (int i = 0; i < 10; i++) {
    for (int j = 0; j < 10; j++) {
        ...
    }
}
```

(a)

OpenACC (Equivalent):

```
#pragma acc parallel loop gang
for (int i = 0; i < 10; i++) {
    #pragma acc parallel loop vector
    for (int j = 0; j < 10; j++) {
        ...
    }
}
```

(b)

OpenMP 4:

```
#pragma omp target teams distribute
for (int i = 0; i < 10; i++) {
    #pragma omp parallel for
    for (int j = 0; j < 10; j++) {
        ...
    }
}
```

(c)

Figure 3: Doubly nested loops.

the outermost loop as a gang loop and the innermost loop as a vector loop. Given the OpenACC code in Figure 3(a), the compiler does exactly this: the iterations of the i -loop are divided among the thread blocks, and the iterations of the j -loop are divided among the threads within each block. The OpenACC code in Figure 3(b) is functionally equivalent but makes this distribution explicit.

In OpenMP, work distribution is not so automatic. Consider these “obvious” but incorrect translations: The `#pragma acc parallel loop` directive could be replaced by:

- `#pragma omp target teams`. In this case, the entire loop nest would be run redundantly by all thread blocks.
- `#pragma omp target teams distribute`. In this case, the iterations of the i -loop would be divided among thread blocks, but only one thread would be active within each block.
- `#pragma omp target teams distribute parallel for`. Here, the iterations of the i -loop are distributed among threads and thread blocks.

In all three cases, *the inner j -loop runs sequentially*, because the parallelization directive is applied to only one loop. OpenMP requires the programmer to explicitly specify how to run each loop in the nest.

The translation most similar to the OpenACC code is shown in Figure 3(c). The iterations of the i -loop are divided among thread blocks using the `distribute` directive, and the iterations of the j -loop are divided among threads in the block using the `parallel for` directive. OpenACC’s `gang` clause is roughly equivalent to OpenMP’s `distribute`: effectively, it distributes work among thread blocks. Likewise, OpenACC’s `vector` clause is roughly equivalent to OpenMP’s `parallel for` inside a `teams` region: it distributes work among threads within a block.

3.2.3 Data Transfer May Differ

In both OpenACC and OpenMP, if `data/target data` directives are omitted, the compiler will automatically determine what data needs to be transferred between the host and the accelerator. This is another case where compilers may differ. In the vector addition example (Figure 2), consider what happens when A , B , and C are *pointers* allocated just before the kernel (in the same function) using `malloc`. With PGCC, the `#pragma acc data` directive can be omitted; it can determine statically that 10,000 elements of A , B , and C are accessed and thus need to be copied. With Clang, if the `#pragma omp target data` directive is omitted, the kernel crashes at runtime (*device illegal address*).

4. TRANSLATION ALGORITHM AND TOOL IMPLEMENTATION

In the previous section, we showed that OpenACC and OpenMP 4 device directives handle work distribution differently, and that differences in compilers can have significant effects on the performance (and correctness) of code. This means that porting from OpenACC to OpenMP 4 will almost certainly involve some manual effort. Nevertheless, much of the translation is straightforward and tedious, which makes it worth automating.

Therefore, we designed our translation tool—and its underlying algorithm—with the following objectives:

- *The translation should be straightforward and consistent.* Each OpenACC directive should translate to the same OpenMP directive(s) every time, and the translation should mirror the structure of the OpenACC code, to the extent possible.
- *Automated translation and manual editing should be complementary.* The tool should not try to automate too much. It should not perform loop nest optimizations; it should not infer implicit data transfers; it should not make any invasive code changes. The tool should automate the tedious, most obvious parts of the translation, in a completely predictable way. More sophisticated changes—like performance tuning—should be under the control of the programmer.

The translation algorithm is as follows. The input is a C source file containing OpenACC directives. The output is the same file with OpenMP directives substituted.

1. If any of the following occur, warn the user; these structures will not be automatically converted:
 - The header file `openacc.h` is included.
 - The macro `_OPENACC` is used.
 - Functions in the OpenACC runtime library are called or referenced.
 - The `kernel`s directive is used.
2. Prepare parallel loop nests using the algorithm described in §4.1 below.
3. Remove *private* and *reduction* clauses from sequential loops as described in §4.2 below.
4. Apply the translation rules described below in §4.3 to convert OpenACC directives to OpenMP directives.

The `#pragma acc kernels` directive is effectively a request for automatic parallelization; the compiler must decide what loops to parallelize and how. There is no direct equivalent in OpenMP. Thus, our algorithm does not support the `kernels` directive. Instead, the programmer should modify the code to explicitly parallelize loops using the `parallel` directive.

4.1 Loop Nest Preparation

As noted in §3.2.2 above, OpenACC-to-OpenMP translation is complicated by the fact that work distribution is explicit in OpenMP, while it may be left to the compiler in OpenACC. We noted that the code in Figure 3(a) is equivalent to that in Figure 3(b), except the latter makes the work distribution explicit.

Before our translation tool attempts to translate any OpenACC directives, it “prepares” each parallel loop nest as follows.

- If the loop nest contains explicit gang and vector clauses, there must be exactly one gang loop and exactly one vector loop in the nest. The same loop may be both the gang and vector loop.
- If the loop nest does not contain an explicit gang loop, the outermost parallelizable loop is chosen as the gang loop.
- If the loop nest does not contain an explicit vector loop, the innermost parallelizable loop is chosen as the vector loop.
- All loops in the loop nest other than the gang and vector loop are run sequentially.
- If the loop nest does not have the form above (e.g., if it contains two gang loops, or if a vector loop appears outside a gang loop), the translation to OpenMP cannot proceed.

The last rule is necessary because OpenACC is surprisingly lax about where the nesting of gang and vector loops. For example, a **gang** loop may be nested inside a **vector** loop; there is no direct equivalent of this in OpenMP. Moreover, PGCC may ignore the **gang** and **vector** directives altogether, e.g., when a gang loop appears inside another gang loop, or a vector loop appears outside another vector loop.

The second and third rules implement a heuristic typical of parallelizing compilers (one that PGCC appears to use, in fact): the iterations of the outermost parallelizable loop are partitioned among thread blocks, and the iterations of the innermost parallelizable loop are partitioned among the CUDA threads within that block.

To determine if a loop is parallelizable (for the second and third rules), our tool uses a combination of syntactic checks and dependence analysis [12, 1]:

1. If the loop contains a **seq** clause, it is not parallelizable.
2. If the loop contains a **gang** or **vector** clause, it is parallelizable.
3. Otherwise, we perform a dependence analysis on the loop. If the loop does not carry a dependence, it is parallelizable [1]; if it does, we determine it to be non-parallelizable.

4.2 Sequential Private/Reduction Removal

After a loop nest has been prepared, there will be a gang loop, a vector loop (possibly the same loop), and zero or more sequential loops. OpenACC has a directive for sequential loops (**acc loop seq**), which permits *private* and *reduction* clauses. However, OpenMP does not have a sequential loop directive. This means that these directives—and more importantly, *private* and *reduction* clauses—must be removed from sequential loops in a way that will preserve the semantics of the *private* and *reduction* clauses.

To remove a *private* clause, we can simply re-declare the privatized variable in the loop body, as shown in Figure 4(a) (noting that this features requires C99 support, which is available in all modern C compilers). This declaration shadows the declaration in the outer scope, which gives the intended semantics of the *private* clause.

To illustrate removing a *reduction* clause, consider as a concrete example **reduction(+:x)**. There are two cases.

Case 1. This is perhaps the most commonly occurring use of the *reduction* clause: If the only accesses to **x** in the loop body are in statements of the form **x++** or **x += expression** (where **x** does not occur in *expression*), then the reduction clause can be removed with no change in semantics. This is the case in Figure 4(a), for example.

Case 2. When the above condition does not hold, it is still possible to remove a *reduction* clause. However, preserving the semantics of the *reduction* clause requires a more complex source code transformation, illustrated in Figure 4(b). (1) A “fresh” variable (**sum**) is declared to accumulate the result. (2) Since **x** is the reduction variable, a local variable **x** is declared in the loop body with the same type as the shadowed variable **x**, and it is initialized to the initializer value for a sum reduction (0). (3) At the end of the loop body, the local **x** is added to the result (**sum**). (4) After the loop, the accumulated result is added to the original **x**. (Our prototype translation tool does not yet implement this case, since it was not necessary for our evaluation.)

4.3 Directive Translation

The loop nest preparation algorithm in the previous section identifies a single gang loop and a single vector loop in each parallel loop nest, and it guarantees that the vector loop is nested under (or is the same as) the gang loop. After gang, vector, and sequential loops have been identified, it is possible to proceed with translating OpenACC directives to OpenMP directives.

Our tool implements the translation rules in Figure 5. The majority of the translation is straightforward: Each phrase in OpenACC is mapped to a corresponding phrase in OpenMP.

It is important to note that, according to the translation rules in Figure 5, *each data directive is translated as-is*. No attempt is made to infer implicit data transfers or array lengths.

Also, these rules assume that each loop in a parallel loop nest has been *explicitly* labeled with a **gang**, **vector**, or **seq** clause (or both **gang** and **vector**). Of course, this is not necessary. In fact, our tool runs the loop nest preparation algorithm implicitly; it identifies a gang and vector loop in each parallel loop nest, but it does not literally insert OpenACC **gang**, **vector**, and **seq** clauses into the source code (although it retains them if they are already present in the code). However, it *does* change the source code when remov-

```

int p, x = 100;
#pragma acc loop seq private(p) reduction(+:x)
for (int i = 0; i < 5; i++) {
    p = 100;
    x += p;
}

```

⇓

```

int p, x = 100;
for (int i = 0; i < 5; i++) {
    int p;
    p = 100;
    x += p;
}

```

(a)

```

int x = 100;
#pragma acc loop seq reduction(+:x)
for (int i = 0; i < 5; i++) {
    x = 1;
}

```

⇓

```

int x = 100;
int sum = 0; // (1)
for (int i = 0; i < 5; i++) {
    int x = 0; // (2)
    x = 1;
    sum += x; // (3)
}
x += sum; // (4)

```

(b)

Figure 4: Elimination of *private* and *reduction* clauses on a sequential loop.

ing *private* and *reduction* clauses from sequential loops, since these changes must be reflected in the resulting OpenMP code.

4.4 Translation Tool

We prototyped our OpenACC-to-OpenMP translation algorithm by extending the Eclipse C/C++ Development Tools. In addition to our translation algorithm, we added support for various data flow analyses and dependence analysis, along with a parser for OpenACC directives. Our translation can be performed from within the Eclipse user interface, but we also added a command line interface (CLI), so one can run the translation from the command line, without installing Eclipse.

5. EVALUATION

To evaluate the correctness of our tool, we used it to translate the EPCC Level 1 OpenACC benchmarks [2], which are ports of the PolyBench and PolyBench/GPU kernels [3]. We also translated several miscellaneous test programs, including naïve and tiled matrix multiplication kernels. We compiled the original OpenACC code using PGCC 16.1 and the resulting OpenMP code using Clang 3.8. Performance measurements were taken on a Tesla K40 (Comet and Stampede have Tesla K20 GPUs).

The EPCC Level 1 benchmarks contain a total of 70 OpenACC directives, three of which are *kernels* directives that our tool would not translate directly. Of the remaining 67 directives, 13 were data directives, 23 were parallel loop directives, and 31 were loop directives nested under other parallel directives. There was only one OpenACC API call: a call to `acc_init`, used to initialize the GPU and the OpenACC runtime.

This suggests that a tool like ours could be quite valuable. Only two changes needed to be made by the programmer: removing `acc_init` and converting the three *kernels* regions. The remaining 67 directives could all be converted automatically, without manual intervention.

Interestingly, while both OpenACC and Clang were targeting the same GPU device, the GPU kernels generated by Clang were almost always slower than those generated by PGCC. The naïve matrix multiplication kernel in Figure 6 can serve as a small, concrete example. With 2048×2048 matrices, the OpenACC kernel ran for about 250 ms; the

Kernel	PGI/OpenACC	Clang/OpenMP	Speedup
3MM	17241	50073	0.34
ATAH	24539	28308	0.87
BICG	25178	29561	0.85
MVT	24021	27928	0.86
SYRK	72562	213693	0.34
COV	28120	19905	1.41
COR	28640	187218	0.15
SYR2K	15231	66935	0.23
GESUMMV	37998	52765	0.72
			Median: 0.72

Table 1: Execution times for EPCC Level 1 benchmarks on a Tesla K40 (μ s).

OpenMP kernel ran for about 925 ms. Since the OpenACC and OpenMP codes were nearly identical, and they were running on the same hardware, we did not expect such massive differences in performance. Execution times for the EPCC Level 1 benchmarks on a Tesla K40 are shown in Table 1. The benchmarks were run with a data size of 1024 and 5 repetitions. Clearly, details of the compiler-generated code can have a substantial impact on performance.

6. LIMITATIONS AND FUTURE WORK

Our tool is a prototype, and our translation algorithm supports only the most commonly used OpenACC directives. Some areas for improvement are obvious. The algorithm needs to be expanded to include the entire OpenACC specification, including asynchronous kernel launches, unstructured data lifetimes, and OpenACC API calls. Our tool works only on C, while OpenACC and OpenMP both support C++ and Fortran. While the EPCC benchmarks were useful for evaluating our prototype, we need to test its effectiveness on larger code bases (i.e., application codes), ideally by working with other XSEDE collaborators.

While automating a tedious 1–1 translation is helpful, it is not the end of the conversion process. As noted earlier, a common XSEDE use case will likely be converting OpenACC code to take advantage of the forthcoming second-generation Intel Xeon Phi on Stampede. Our translation algorithm (and tool) could certainly be specialized for this process.

$$\begin{aligned}
\mathcal{T}[\text{\#pragma acc data } data\text{-clauses}] &= \text{\#pragma omp target data } \mathcal{T}[data\text{-clauses}] \\
\mathcal{T}[\text{\texttt{copyin}}(vars)] &= \text{map}(\text{to}:vars) \\
\mathcal{T}[\text{\texttt{copyout}}(vars)] &= \text{map}(\text{from}:vars) \\
\mathcal{T}[\text{\texttt{copy}}(vars)] &= \text{map}(\text{tofrom}:vars) \\
\mathcal{T}[\text{\texttt{create}}(vars)] &= \text{map}(\text{alloc}:vars) \\
\mathcal{T}[\text{\#pragma acc parallel } par\text{-clauses}] &= \text{\#pragma omp target teams } \mathcal{T}[par\text{-clauses}] \\
\mathcal{T}[\text{\texttt{num_gangs}}(n)] &= \text{num_teams}(n) \\
\mathcal{T}[\text{\texttt{vector_length}}(n)] &= \text{thread_limit}(n) \\
\mathcal{T}[\text{\texttt{private}}(vars)] &= \text{private}(vars) \\
\mathcal{T}[\text{\texttt{reduction}}(vars)] &= \text{reduction}(reductions) \\
\mathcal{T}[\text{\texttt{independent}}] &= \epsilon \\
\mathcal{T}[\text{\#pragma acc loop } type \text{ } loop\text{-clauses}] &= \begin{cases} \text{\#pragma omp distribute parallel for } \mathcal{T}[loop\text{-clauses}] & \text{if } type = \text{gang vector} \\ \text{\#pragma omp distribute } \mathcal{T}[loop\text{-clauses}] & \text{if } type = \text{gang} \\ \text{\#pragma omp parallel for } \mathcal{T}[loop\text{-clauses}] & \text{if } type = \text{vector} \\ \epsilon & \text{if } type = \text{seq} \end{cases}
\end{aligned}$$

Figure 5: Translation rules from OpenACC to OpenMP 4. The function $\mathcal{T}$ maps phrases derived from the OpenACC subset grammar to strings corresponding to one or more OpenMP directives. Sequences (e.g., *par-clauses*) are translated elementwise, that is, $\mathcal{T}[t_1 \ t_2 \ \dots \ t_n] = \mathcal{T}[t_1] \ \mathcal{T}[t_2] \ \dots \ \mathcal{T}[t_n]$. The symbol ϵ denotes the empty string.

```

#pragma omp target teams distribute
#pragma acc parallel loop gang
for (int j = 0; j < N; j++) {
    #pragma omp parallel for
    #pragma acc loop vector
    for (int i = 0; i < N; i++) {
        double t = 0.0;
        #pragma acc loop seq
        for (int k = 0; k < N; k++) {
            t += m[k][i] * n[j][k];
        }
        p[j][i] = t;
    }
}

```

Figure 6: Naïve matrix multiplication.

After the initial OpenMP code has been created, some restructuring will almost certainly be necessary. For example, we have performed some preliminary experiments with Intel’s compilers on first-generation Phi; in one case, we needed to convert global variables to local variables to ensure that the OpenMP *target data* directive would guarantee a copy of the arrays to the Phi. Cataloging (and perhaps automating) such transformations would be beneficial.

Perhaps more importantly, performance portability is not guaranteed. A direct translation of OpenACC code optimized for an NVIDIA GPU will not necessarily perform at its peak on a Xeon Phi. As noted above, we did not even obtain comparable performance targeting the *same* hardware.

7. RELATED WORK

Most related to our work is a talk by Hernandez et al. at OpenMPCon 2015 [4, 5]. This talk describes the salient differences between OpenACC 2.0 and OpenMP 4.0 and suggests a (manual) procedure for porting code from OpenACC to OpenMP. Their procedure consists of five steps: removing constructs with no OpenMP counterparts (like *kernels*), translating data regions, translating data updates, translating accelerator parallel regions, and adjusting function attribute specifiers.

Lee and Vetter [7] evaluate existing directive-based models and show that directive-based models can achieve reasonable performance compared to hand-written GPU codes. For accelerator programming, the two most promising options are OpenACC and OpenMP. Wienke et al. [11] compare both models with respect to their programmability to as-

sist developers in deciding which approach to take. Xu and colleagues [13] demonstrate the effectiveness of a hybrid model combined with OpenACC and OpenMP as a plausible solution to port scientific applications in heterogeneous architectures.

8. CONCLUSION

We have described an algorithm (§4) and a prototype tool that converts a subset of OpenACC (Figure 1) to OpenMP 4. The translation algorithm was designed to be simple and predictable, since it will necessarily be used in conjunction with manual restructuring and performance tuning. We evaluated our tool by converting the EPCC Level 1 OpenACC benchmarks from OpenACC compiled by PGI C/C++ to OpenMP 4 compiled by Clang. Both targeted an NVIDIA GPU (Tesla K40). To keep the translation process transparent, the programmer was required to manually convert OpenACC *kernels* loops, although the tool was able to automatically translate the remaining 67 of 70 directives.

9. ACKNOWLEDGMENTS

Portions of this work were supported by the United States Department of Defense under contract #DoD-H98230-15-C-0686-2.

10. REFERENCES

- [1] J. R. Allen and K. Kennedy. *Optimizing Compilers for Modern Architectures: A Dependence-based Approach*. Morgan Kaufmann, San Francisco, CA, 2002.
- [2] EPCC OpenACC benchmark suite. <https://www.epcc.ed.ac.uk/research/computing/performance-characterisation-and-benchmarking/epcc-openacc-benchmark-suite>. Accessed April 29, 2016.
- [3] S. Grauer-Gray, L. Xu, R. Searles, S. Ayalasomayajula, and J. Cavazos. Auto-tuning a high-level language targeted to GPU codes. In *Innovative Parallel Computing (InPar)*, 2012, pages 1–10, May 2012.
- [4] O. Hernandez, W. Ding, W. Joubert, D. Bernholdt, M. Eisenbach, and C. Kartsaklis. Porting OpenACC 2.0 to OpenMP 4.0: Key similarities and differences. <http://openmpcon.org/wp-content/uploads/openmpcon2015-oscar-hernandez-portingacc.pdf>. Accessed April 29, 2016.
- [5] O. Hernandez, W. Ding, W. Joubert, D. Bernholdt, M. Eisenbach, and C. Kartsaklis. YouTube: Porting OpenACC 2.0 to OpenMP 4.0: Key similarities and differences. <https://www.youtube.com/watch?v=CHMrcMUXuuY>. Accessed April 29, 2016.
- [6] D. B. Kirk and W.-m. Hwu. *Programming massively parallel processors: a hands-on approach*. Morgan-Kaufmann, 2012.
- [7] S. Lee and J. S. Vetter. Early evaluation of directive-based GPU programming models for productive exascale computing. In *Proc. SC12*, page 23. IEEE Computer Society Press, 2012.
- [8] The OpenACC application programming interface, version 2.5. http://www.openacc.org/sites/default/files/OpenACC_2pt5.pdf. Accessed June 15, 2016.
- [9] OpenMP 4.0 on NVIDIA CUDA GPUs. <https://parallel-computing.pro/index.php/9-cuda/43-openmp-4-0-on-nvidia-cuda-gpus>. Accessed April 29, 2016.
- [10] OpenMP application programming interface, version 4.5. <http://www.openmp.org/mp-documents/openmp-4.5.pdf>. Accessed June 15, 2016.
- [11] S. Wienke, C. Terboven, J. C. Beyer, and M. S. Müller. A pattern-based comparison of OpenACC and OpenMP for accelerator computing. In *Euro-Par 2014 Parallel Processing*, pages 812–823. Springer, 2014.
- [12] M. J. Wolfe. *High Performance Compilers for Parallel Computing*. Addison-Wesley, Boston, MA, 1995.
- [13] R. Xu, S. Chandrasekaran, and B. Chapman. Exploring programming multi-GPUs using OpenMP and OpenACC-based hybrid model. In *IPDPSW '13*, pages 1169–1176. IEEE, 2013.