

Achieving portability for a highly optimized GPU code for 3D Fourier Transforms at extreme problem sizes

Kiran Ravikumar¹, Oscar Hernandez², John Levesque³,
Stephen Nichols², P.K. Yeung¹

kiran.r@gatech.edu

¹Georgia Institute of Technology, ²Oak Ridge National Laboratory, ³Cray (HPE)

Performance, Portability, and Productivity in HPC Forum
September 1-2, 2020

Introduction & Motivation

Communication intensive codes on exascale heterogeneous machines?

- multidimensional Fourier transforms: fluid dynamics, signal processing, etc
- accelerators provide most of computing power, but need to move data
- communication still an issue (perhaps even more so): some new challenges

Batched asynchronous approach targeting extreme problem sizes


- problem size limited by smaller device memory compared to host
- process data in batches on GPU with entire data residing in CPU memory
- overlap operations on different batches of data, hide compute and transfer costs

CUDA Fortran (Ravikumar *et al.* SC'19) on Summit (IBM+NVIDIA)

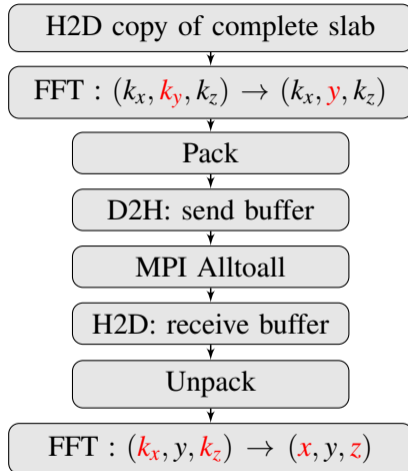
- **portability**: using OpenMP for offload, up to Version 5.0

Synchronous 3D FFT using GPUs

1D domain decomposition (Slabs)




- Slabs: GPU parallelism instead of distributed memory, fewer MPI tasks in communication
- Copy **entire slab** from CPU (host (H)) to GPU (device (D)) and back to CPU at end
- 1D FFTs using cufft or rocfft library
- MPI Alltoall among all tasks to transpose $x-y$ to $x-z$ slabs



Large problem that may not fit on GPU? Any asynchronism possible?

Batched approach

- Divide slab into np pencils and process each pencil separately ($nyp = nxp = N/np$)
- Overlap operations on different pencils to hide some data transfer and compute costs
- Overlap using one stream each (in CUDA Fortran) for data transfer and compute
- Overlap: **Compute** on ip , **HtoD** on $ip + 1$, **DtoH** on $ip - 1$ and **all-to-all** on $ip - 2$
- Non-blocking all-to-all allows overlap, `MPI_WAIT` ensures completion
- GPU-Direct can be used to avoid copies before and after all-to-all
- Repeat until all pencils (np) processed on GPU and transposed




Non-contiguous maps and strided copies

- FFTs in y : need only $a(1:nxp, 1:ny)$ on device
- In CUDA Fortran use:

```
1  cudaMemCpy2DAsync (dst, dpitch, src, spitch, &  
2  width, height, kind, stream)
```

- unique destination (dst) & source (src) buffers
- $width$ (nxp) elements out of $spitch$ (nx) to copy from first dimension, $height$ (ny) such copies



How to do it in OpenMP?

- `MAP (to:a(1:nxp, 1:ny))`: not 5.0 compliant
 - copy a to smaller $abuf$ on host then MAP
 - but this adds extra work on the host

```
1  allocate (a(nx, ny), abuf(nxp, ny))  
2  ! copy on host : abuf(:, :) = a(1:nxp, 1:ny)  
3  !$OMP TARGET DATA MAP(tofrom: abuf)  
4  ! do some work  
5  !$OMP END TARGET DATA MAP  
6  ! copy on host : a(1:nxp, 1:ny) = abuf(:, :)
```

Performance penalty due to additional operations on host

OpenMP 4.5+: `omp_target_memcpy_rect`?

- Copy rectangular subvolume from a multi-dimensional array
- Callable from C/C++, use C-Fortran interface
- *ndims*: no. of dimensions in array
- *vol*: no. of elements to copy in each dimension
- *offset*: no. of elements from *base of each dimension*, after which to copy data from/to
 - In 5.0: from *origin of dst (src)*, **need clarity**
- *dims*: no. of elements in each dimension
- Need to **account for C vs. Fortran ordering**
 - first dimension along row (*ny*) even though in Fortran it is along column

```
1  ! src on host of shape (nx, ny)
2  ! dst on device of shape (npx, ny)
3
4  ! copy src(1:npx, 1:ny) to dst(1:npx, 1:ny)
5
6  num_dims = 2
7  vol(1) = ny ; vol(2) = npx
8  dst_offset(1) = 0 ; dst_offset(2) = 0
9  src_offset(1) = 0 ; src_offset(2) = 0
10 dst_dims(1) = ny ; dst_dims(2) = npx
11 src_dims(1) = ny ; src_dims(2) = nx
12
13 omp_target_memcpy_rect(dst, src, elem_size,
    ndims, vol, dst_offset, src_offset,
    dst_dims, src_dims, dst_dev, src_dev)
```

Interoperability between OpenMP and non-blocking libraries

```
1 TARGET DATA MAP(tofrom:a)
2
3 TASK DEPEND(OUT:var) (A)
4 TARGET DATA USE_DEVICE_PTR(a)
5 FFTExecute (a, forward, stream)
6 FFTExecute (a, inverse, stream)
7 END TARGET DATA
8 END TASK
9
10 TARGET TEAMS DISTRIBUTE DEPEND(IN:var) NOWAIT
11 a(:, :, :) = a(:, :, :) / nx (B)
12 END TARGET TEAMS DISTRIBUTE
13
14 TASKWAIT
15 END TARGET DATA
```

- Task A: Non-blocking FFT call
- Task B: OpenMP kernel with dependency on previous task
- Expect kernel runs after FFT completes
- But task A is considered complete after call to library, freeing dependency
- B executes before FFT completes
- Returns incorrect results

Task dependency not sufficient to enforce required synchronization

DETACH to enforce synchronization

```
1 TARGET DATA MAP(tofrom: a)
```

```
2  
3 TASK DEPEND(out:var) DETACH(event)
```

```
4  
5 TARGET DATA USE_DEVICE_PTR(a) (A)
```

```
6 FFTExecute (a, forward, stream)
```

```
7 FFTExecute (a, inverse, stream)
```

```
8 END TARGET DATA
```

```
9  
10 hipStreamAddCallback (stream, ptr_callback, C_LOC(event), 0)
```

```
11 END TASK
```

```
12  
13 ! Copy or compute on other data (C)
```

```
14  
15 TARGET TEAMS DISTRIBUTE DEPEND(IN:var) NOWAIT
```

```
16 a(:, :, :) = a(:, :, :)/nx
```

```
17 END TARGET TEAMS DISTRIBUTE (B)
```

```
18  
19 END TARGET DATA
```

```
1 subroutine callback (stream, status, event)
```

```
2 type(c_ptr) :: event
```

```
3 integer(kind=omp_event_handle_kind) :: f_event
```

```
4 call C_F_POINTER (event, f_event)
```

```
5 call omp_fulfill_event(f_event)
```

```
6 end subroutine callback
```

1. A: launch FFT, add call to *callback* in stream where FFT is running
2. B waits as dependent on A, C executes asynchronously
3. After FFT, function *callback* is called and *event* fulfilled
4. A completes allowing B to run

Support for DETACH not yet available

Porting asynchronous CUDA Fortran to OpenMP

```
1 do ip=1,np
2   NEXT = mod(ip+1,3); CURR = mod(ip,3);
3   PREV = mod(ip-1,3); COMM = mod(ip-2,3);
4   cudaStreamWaitEvent (trans_stream, DtoH(NEXT), 0)
5   cudaMemCpy2DAsync (abuf(NEXT),a(ip+1),trans_stream)
6   cudaEventRecord (HtoD(NEXT),trans_stream)
7   cudaStreamWaitEvent (comp_stream, HtoD(CURR), 0)
8   FFTExecute (abuf(CURR), comp_stream)
9   cudaEventRecord (comp(CURR), comp_stream)
10  cudaStreamWaitEvent (trans_stream, comp(PREV), 0)
11  cudaMemCpy2DAsync (snd(ip-1), abuf(PREV), &
12    trans_stream)
13  cudaEventRecord (DtoH(PREV), trans_stream)
14  cudaEventSynchronize (DtoH(COMM))
15  MPI_IALLTOALL (snd(ip-2))
16 end do
```


```
1 do ip=1,np
2   NEXT = mod(ip+1,3); CURR = mod(ip,3);
3   PREV = mod(ip-1,3); COMM = mod(ip-2,3);
4   TASK DEPEND (IN:DtoH(NEXT), OUT:HtoD(NEXT))
5   omp_target_memcpy_rect (abuf(NEXT), a(ip+1))
6
7   TASK DEPEND (IN:HtoD(CURR), OUT:comp(CURR))
8   DETACH(event)
9   FFTExecute (abuf(CURR), comp_stream)
10  TASK DEPEND (IN:comp(PREV), OUT:DtoH(PREV))
11  omp_target_memcpy_rect (snd(ip-1), abuf(PREV))
12
13  TASK DEPEND(IN:DtoH(COMM))
14  MPI_IALLTOALL (snd(ip-2))
15
16 end do
```

- **DEPEND clause replaces** cudaEventRecord & cudaStreamWaitEvent
- **omp_target_memcpy_rect replaces** cudaMemCpy2DAsync

Performance: Non-Batched synchronous version

Summit (XL compiler) up to 1024 nodes ($\sim 22\%$ of full machine) using 1 task/GPU
Timings for 3 pairs of forward and inverse transforms

# Nodes	Prob. Size	Time (s)		
		CPU	CUDA	OMP
2	1536 ³	5.21	2.39	2.41
16	3072 ³	6.79	3.30	3.16
128	6144 ³	9.10	5.26	5.01
1024	12288 ³	10.59	4.30	4.12



- OpenMP & CUDA show similar performance ($\sim 2.6X$ speedup for $12k^3$)
- GPU: **compute** negligible but additional cost due to **copies**, 62% in **MPI**
- OpenMP data copies slower than in CUDA, but compute faster !
- OpenMP code also works with CCE compiler and AMD GPUs

Performance: Batched version

- OpenMP version: copy on host from large buffer to small buffer before UPDATE (workaround)
 - `omp_target_memcpy_rect` slow compared to workaround and `cudaMemCpy2D`
- $6k^3$ OMP is 16.1s slower than CUDA async
 - 12.4s to copy one buffer to another on host
 - 3.7s (or 20%) saving due to asynchronism?
- Work in progress: optimize OpenMP version
 - Fast rectangular copy to avoid host operations
 - DETACH will help enable asynchronism
- Both OMP codes work with CCE & AMD GPUs

6 pencils per slab

Performance on Summit using XL
OpenMP version uses workaround

# Nodes	Prob. Size	Time (s)	
		CUDA async	OMP sync
4	3072^3	10.14	26.20
32	6144^3	13.53	29.64

Production code using CUDA: $18k^3$
on 3k nodes, $\sim 3X$ speedup

Summary and Future Work

- Progress made on OpenMP offload implementation of 3D Fast Fourier Transforms
 - Porting from a successful CUDA Fortran code (SC'19) on Summit at OLCF
 - At present: batched, **synchronous**, which enable large problem sizes
- Some challenges of portability overcome, some pending full OMP 5.0 availability:
 - Strided copy b/w smaller device & larger host arrays: `omp_target_memcpy_rect`
 - Synchronizing non-blocking GPU library calls & OpenMP tasks: `DETACH`
- Future work towards 3D FFTs at massive scale, at resolution beyond $18,432^3$
 - Batched **asynchronism** algorithm (using `DETACH`) needed for optimal performance
 - Portable GPU parallelism for communication-intensive applications

Comment: Need better Fortran support (would like to see OMP 5.0 routines that are only C-callable also callable from Fortran, with good examples)