



SYCL Support for Continental-Scale Ecological Observations: Scalable and Portable Blending of Massive Image Mosaics

Dr. Steve Petruzza, Dr. Valerio Pascucci

The University of Utah & Lawrence Livermore National Laboratory Intel oneAPI Center of Excellence
Utah State University



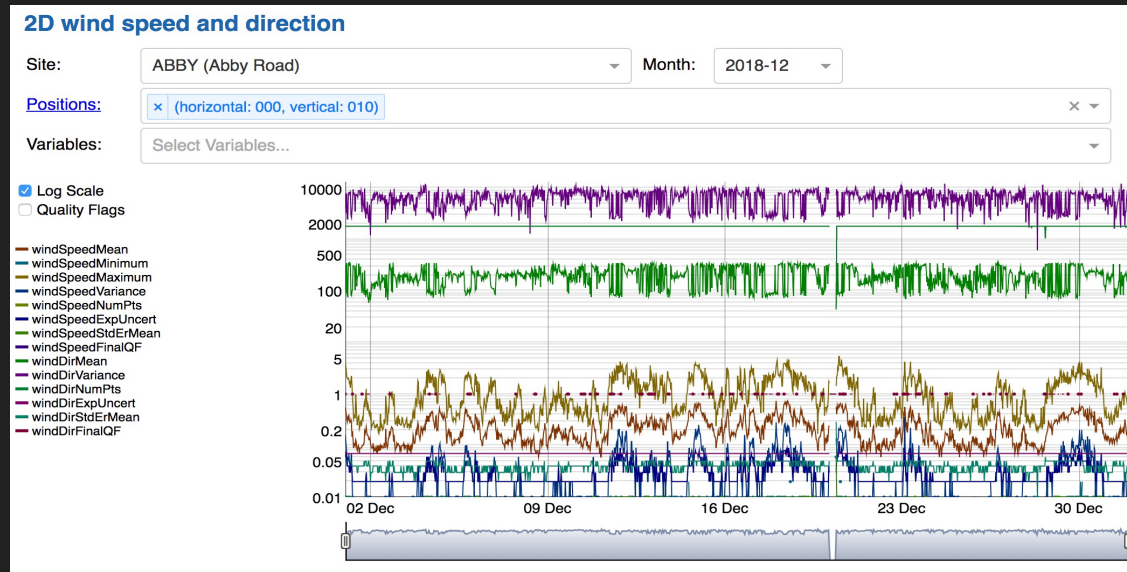
National Ecological Observatory

Network



NEON data

- NEON has a large amount of data that is shared with the community through their **data portal**
- For some data, such as sensor measurements, the portal provides an **interactive** navigation system
- For others, like **Airborne Observation Platforms data**, there is a long list of image files...
- There is a need to present all AOP data interactively, where the users can preview, navigate, and select/access/download the data they need



Include	Filename
<input checked="" type="checkbox"/>	2017_ABBY_1_546000_5060000_image.tif
<input checked="" type="checkbox"/>	2017_ABBY_1_546000_5061000_image.tif
<input checked="" type="checkbox"/>	2017_ABBY_1_546000_5062000_image.tif
<input checked="" type="checkbox"/>	2017_ABBY_1_546000_5063000_image.tif
<input checked="" type="checkbox"/>	2017_ABBY_1_546000_5064000_image.tif
<input checked="" type="checkbox"/>	2017_ABBY_1_546000_5065000_image.tif
<input checked="" type="checkbox"/>	2017_ABBY_1_546000_5066000_image.tif
<input checked="" type="checkbox"/>	2017_ABBY_1_546000_5067000_image.tif
<input checked="" type="checkbox"/>	2017_ABBY_1_546000_5068000_image.tif
<input checked="" type="checkbox"/>	2017_ABBY_1_546000_5069000_image.tif
<input checked="" type="checkbox"/>	2017_ABBY_1_546000_5070000_image.tif
<input checked="" type="checkbox"/>	2017_ABBY_1_546000_5071000_image.tif
<input checked="" type="checkbox"/>	2017_ABBY_1_546000_5072000_image.tif
<input checked="" type="checkbox"/>	2017_ABBY_1_546000_5073000_image.tif
<input checked="" type="checkbox"/>	2017_ABBY_1_546000_5074000_image.tif
<input checked="" type="checkbox"/>	2017_ABBY_1_547000_5059000_image.tif
<input checked="" type="checkbox"/>	2017_ABBY_1_547000_5060000_image.tif

Showing 1 to 100 of 20,850 entries

NEON Airborne Operational Platform Data

- NEON – National Ecological Observatory Network
- 59 AOP sites
- 1-2 datasets per site/year
- Dataset size range from 20 to 300GB
- Total data collection 100s TBs



Access to Airborne Remote Sensing Data

AOP Data to Hard Drive Request

There are several ways, users can access airborne data:

- Download the data from the [NEON data portal](#) (recommended for smaller amounts of data)
- Programmatically access the data with the [NEON Data API](#) or using the [NEON Utilities](#) GitHub repo (>1 GB downloads)
- Mail in a hard drive to receive your data

Please fill out the form below if you are interested in receiving a hard drive of AOP data, and we will respond with a recommended hard drive size as well as mailing instructions.

Please note that requesters must mail us a SATA drive at least the size of the data they are requesting. The drive must be blank.

First Name

Last Name

Email

[Privacy](#) [Terms](#)



NEON | High-resolution orthore...

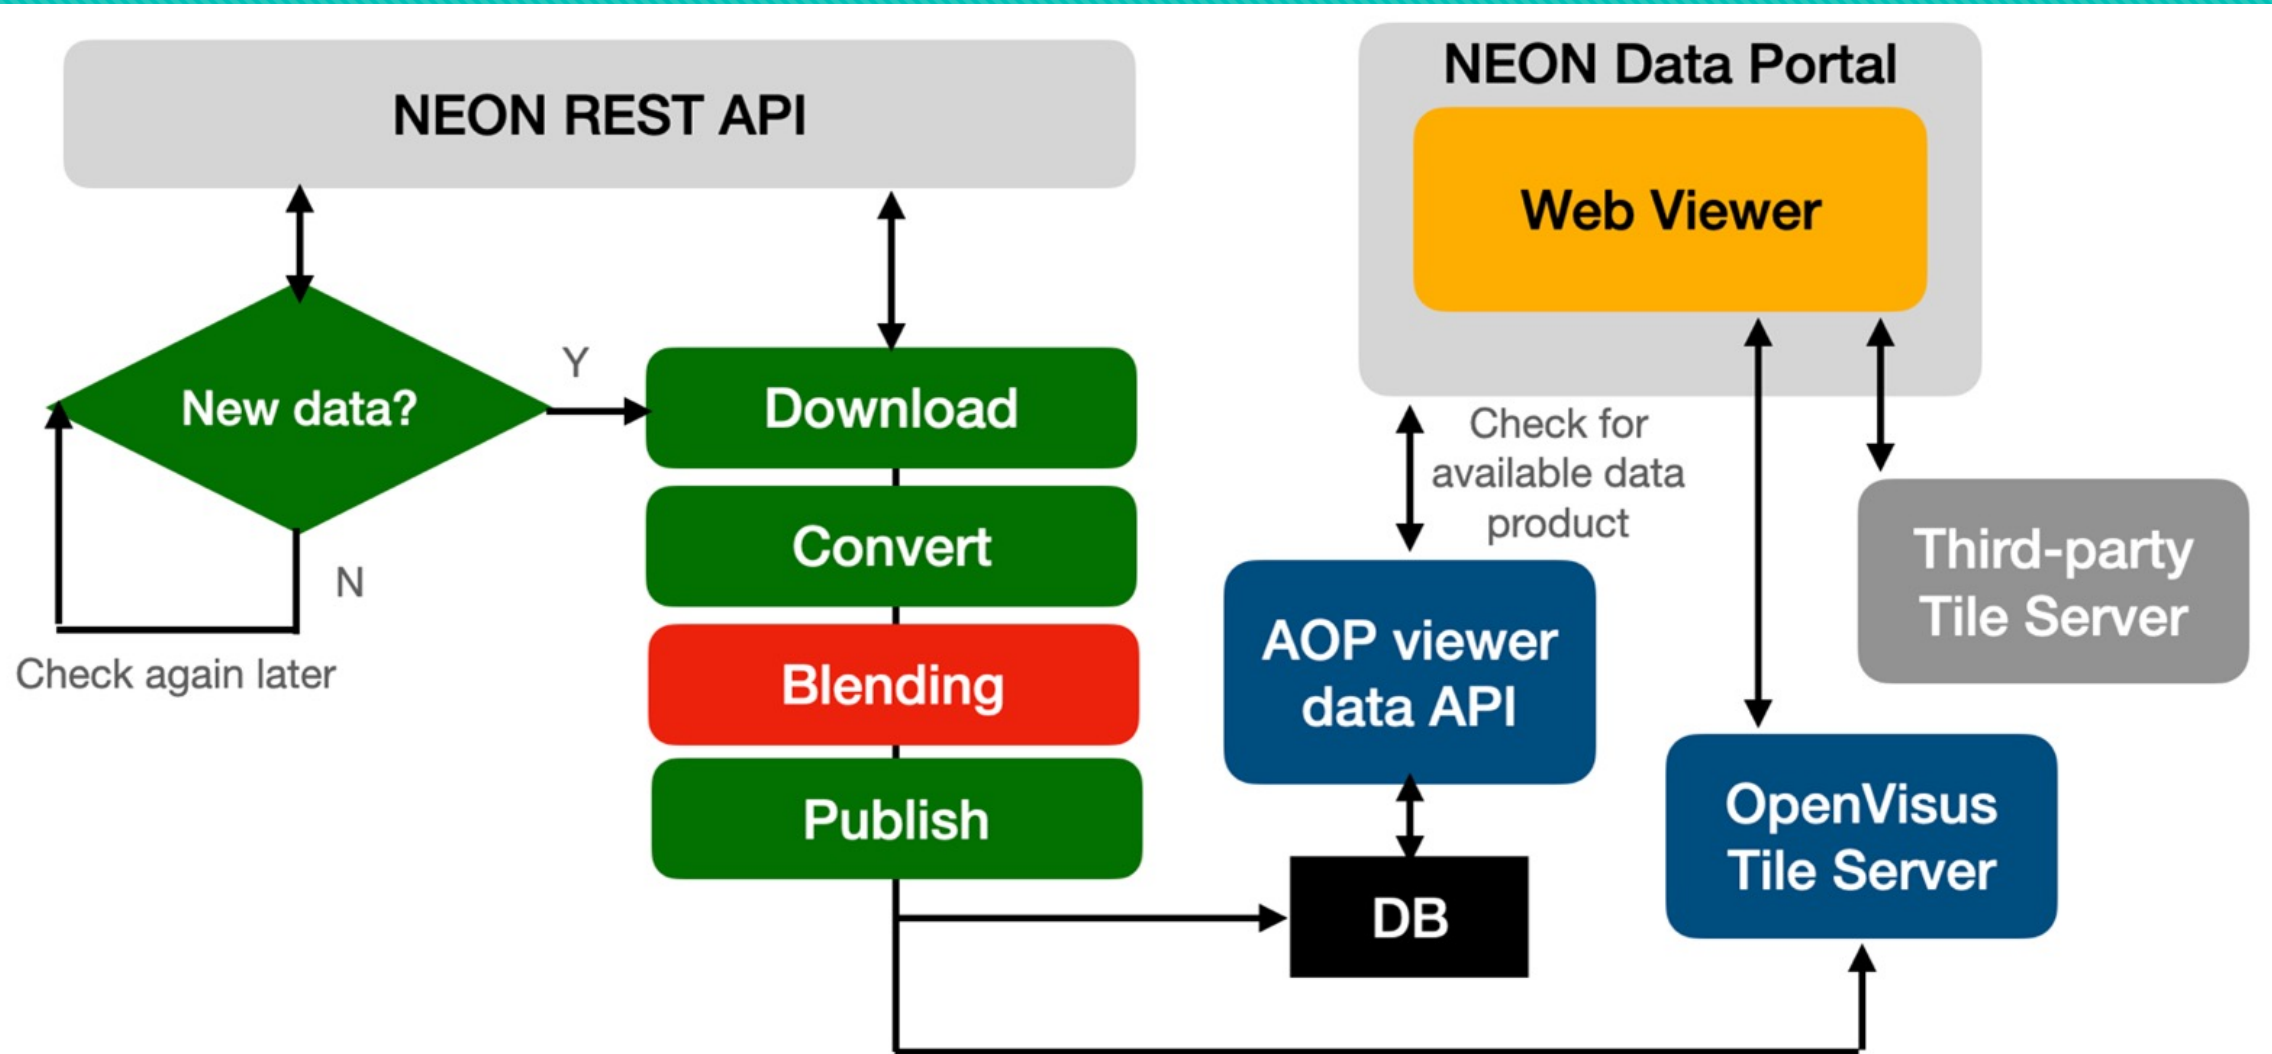
https://data.neonscience.org/data-products/DP3.30010.0

This viewer allows for interactive exploration of remotely sensed data from the Airborne Observation Platform (AOP). Change the field site and flight for this data product using the tools below to stream different data into view. Pan and zoom in the view to stream higher resolution imagery. This pilot data viewer is provided through a collaboration with the [Visus Project at the University of Utah](#) and more updates are planned for the future.

Site Year Flight

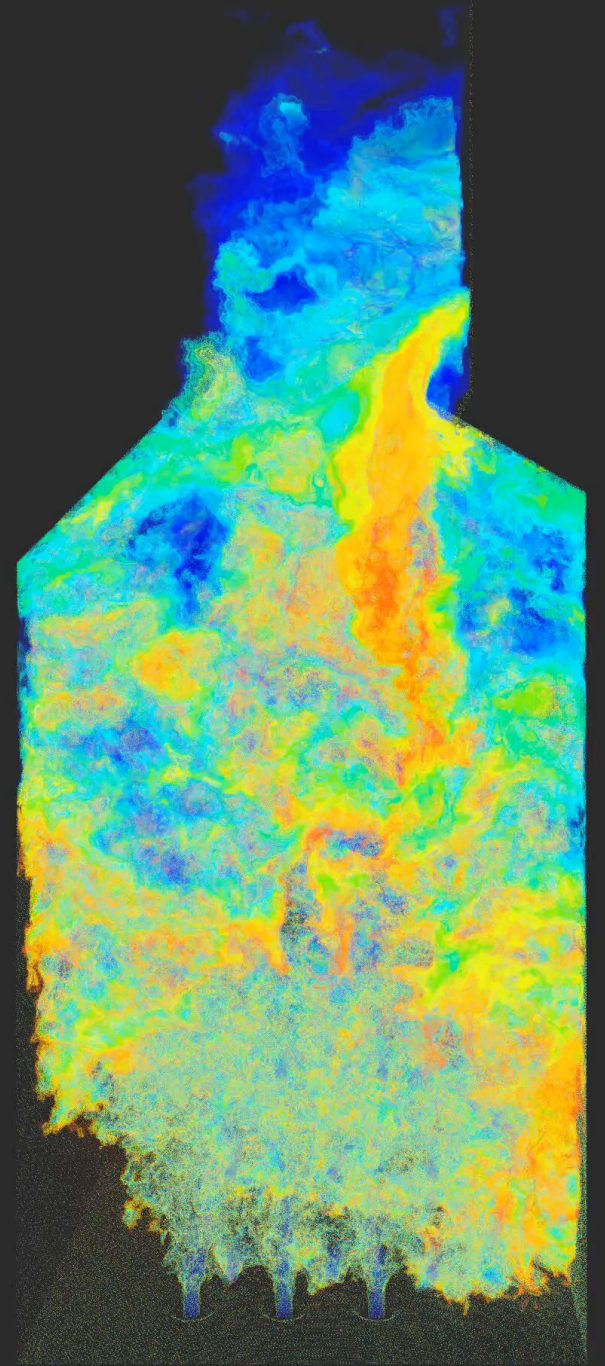
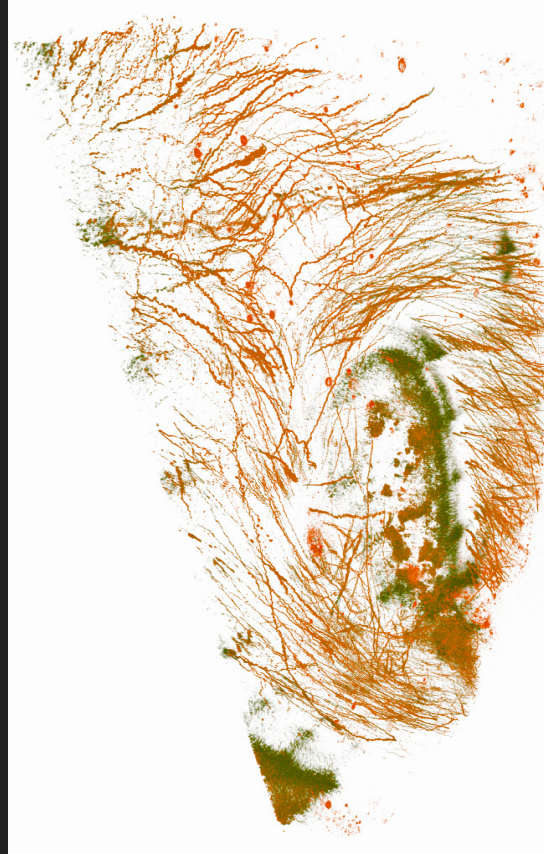
powered by OpenVisus [Download](#) Abby Road, WA -- June 2017 -- Flight 1/1

Data analysis and visualization pipeline



OpenVisus streaming visualization powered by Intel OSPRay

- Large scale combustion simulations
- Neuroscience data (600GB microscopy scan of Macaque monkey visual cortex)
- OpenVisus data streaming for large scale simulation
- High quality rendering at high framerate on local workstations using Intel OSPRay



Gradient domain image processing



Gradient domain image processing

- images processed in the gradient space (not pixel)
- find the closest smooth image to the guiding gradient with a minimal least squared error
- equivalent to solving a 2D Poisson equation
 - direct solutions can be fast using FFT, but with low accuracy
 - or by discretizing the equations into a large linear system that can be solved iteratively (e.g., Conjugate Gradient)
- Requires to build and solve the system in memory

Challenges with NEON AOP data

Remove major artifacts that prevent a scientific investigation

Input



with iterative solver



Challenges with NEON AOP data

Remove major artifacts that prevent a scientific investigation

Input

with iterative solver

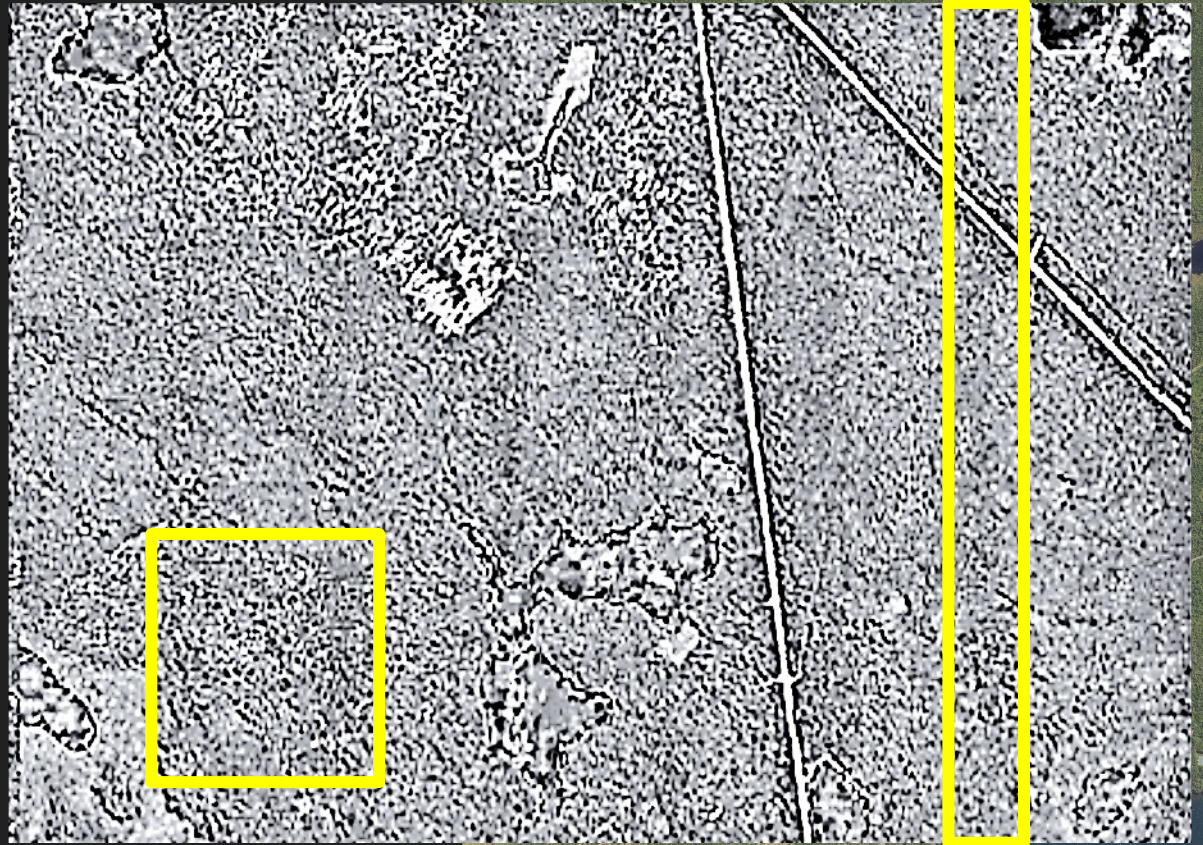


Challenges with NEON AOP data

Remove major artifacts that prevent a scientific investigation

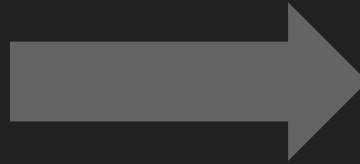
Input

with iterative solver



Code conversion from CUDA to SYCL

Name
..
CMakeLists.txt
SOLVER.cu
common.h
helper_cuda.h
helper_string.h
main.cpp
poisson_solver.cpp
poisson_solver.h



Name
..
MainSourceFiles.yaml
SOLVER.dp.cpp
common.h
helper_cuda.h
helper_cuda.h.yaml
helper_string.h
main.cpp
poisson_solver.cpp.dp.cpp
poisson_solver.h
poisson_solver.h.yaml

Code conversion from CUDA to SYCL

poisson_solver / cg-solver / src / SOLVER.cu 

 spetruzza more testing clean up

Code Blame Executable File · 269 lines (210 loc) · 8.47 KB

```
1 #include <cuda.h>
2 #include <cuda_runtime.h>
3 #include <cuda_runtime_api.h>
4 #include "helper_cuda.h"
5 #include <stdio.h>
6 #include <cusblas.h>
7 #include <cusparse.h>
8 #include <time.h>
9 #include "common.h"
10 #include <iostream>
11
12 #ifndef clamp
13 #define clamp(value,a,b) (((value)<(a))?a:(((value)>(b))?b:(value)))
14 #endif
15
16 __global__ void mult_noshared(int dimx, int dimy, const float3 *x_old, float3 *x, uchar3 *map_data){
17     int i = blockIdx.x * blockDim.x + threadIdx.x;
18     int j = blockIdx.y * blockDim.y + threadIdx.y;
19     int idx = j*dimx+i;
20
21     if(i < dimx && j < dimy && map_data[idx].z != 255){
22         int x0 = idx-1;
23         int x1 = idx+1;
24         int y0 = idx-dimx;
25         int y1 = idx+dimx;
26
27         float3 temp;
```

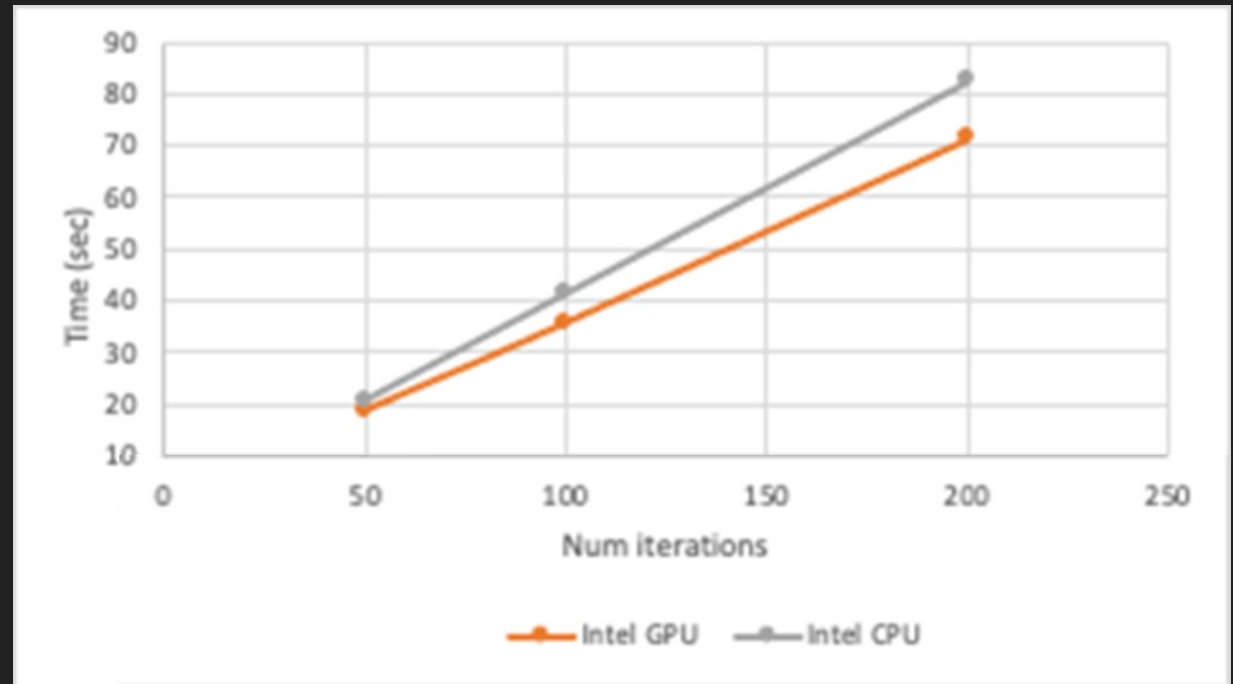
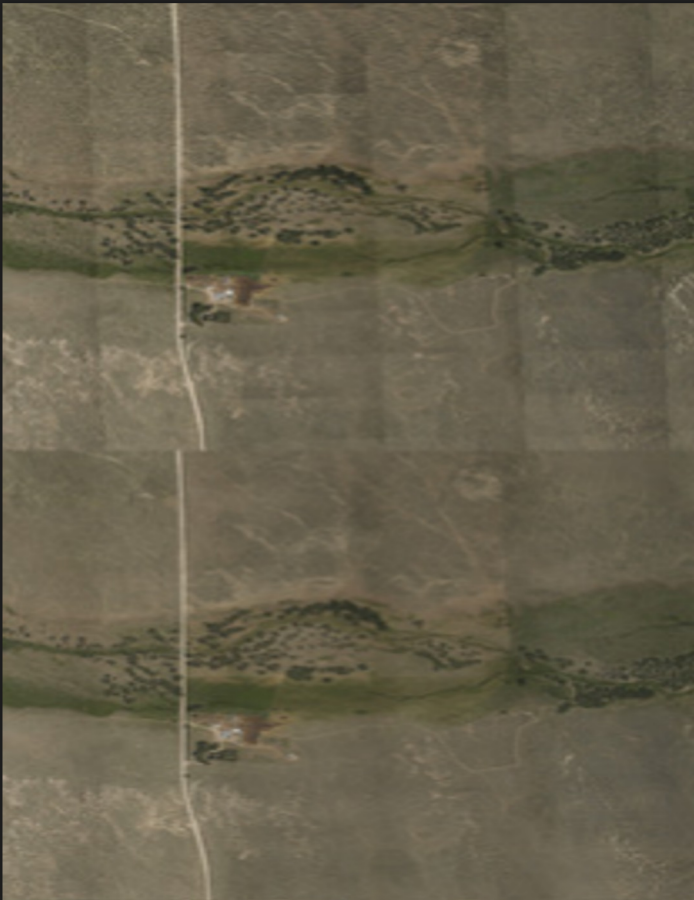
poisson_solver / cg-solver / sycl_converted / SOLVER.dp.cpp 

 spetruzza add converted sycl code

Code Blame 617 lines (561 loc) · 29.9 KB

```
1 #include <sycl/sycl.hpp>
2 #include <dpct/dpct.hpp>
3 #include "helper_cuda.h"
4 #include <stdio.h>
5 #include <dpct/blas_utils.hpp>
6
7 #include <dpct/sparse_utils.hpp>
8
9 #include <time.h>
10 #include "common.h"
11 #include <iostream>
12
13 #ifndef clamp
14 #define clamp(value,a,b) (((value)<(a))?a:(((value)>(b))?b:(value)))
15 #endif
16
17 void mult_noshared(int dimx, int dimy, const sycl::float3 *x_old,
18                 sycl::float3 *x, sycl::uchar3 *map_data,
19                 const sycl::nd_item<3> &item_ct1) {
20     int i = item_ct1.get_group(2) * item_ct1.get_local_range(2) +
21         item_ct1.get_local_id(2);
22     int j = item_ct1.get_group(1) * item_ct1.get_local_range(1) +
23         item_ct1.get_local_id(1);
24     int idx = j*dimx+i;
25
26     if (i < dimx && j < dimy && map_data[idx].z() != 255) {
27         int x0 = idx-1;
```

Experimental study: multiple solver iterations



Prior Intel® DevCloud for oneAPI

Intel CPU: 11th Gen Intel® Core™ i9-11900KB @ 3.30GHz

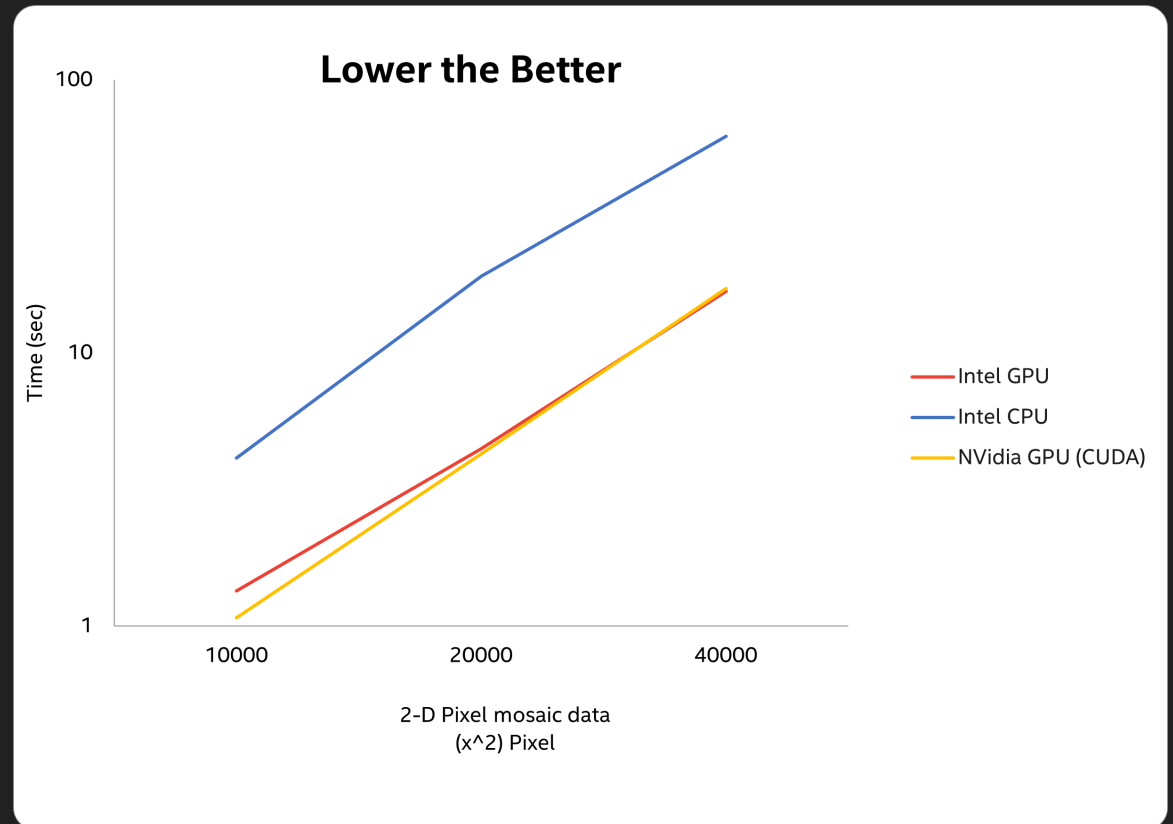
Intel GPU: Intel® Xeon® E-2176 P630 processors (2018)

Experimental Study: Increasing problem size on different computing platforms



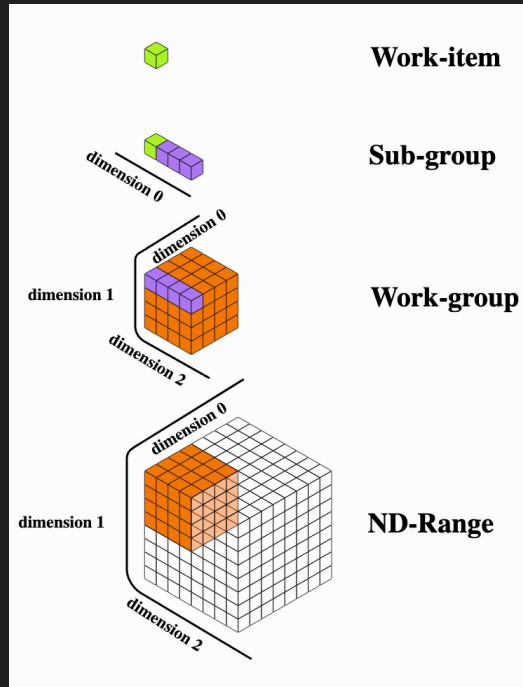
Data:

- [Results_v1](#)
- [Results_v2](#)



- CPU Intel® Xeon® Platinum 8480
- GPU Intel® Data Center GPU Max 1100
- GPU Nvidia RTX A6000

Execution Model



CUDA	SYCL
thread	work-item
warp	sub-group
block	work-group
grid	ND-range

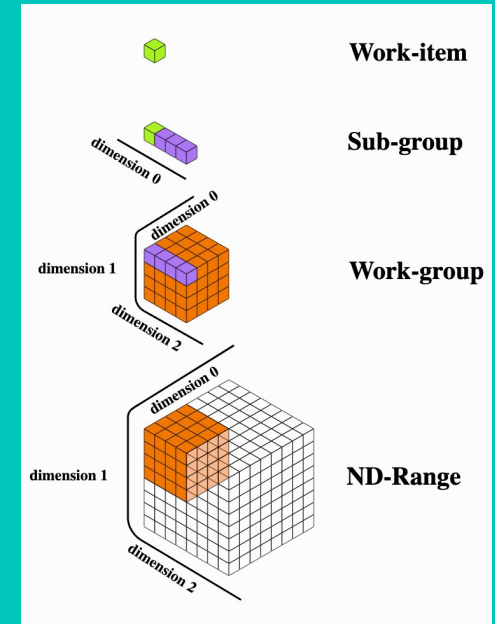
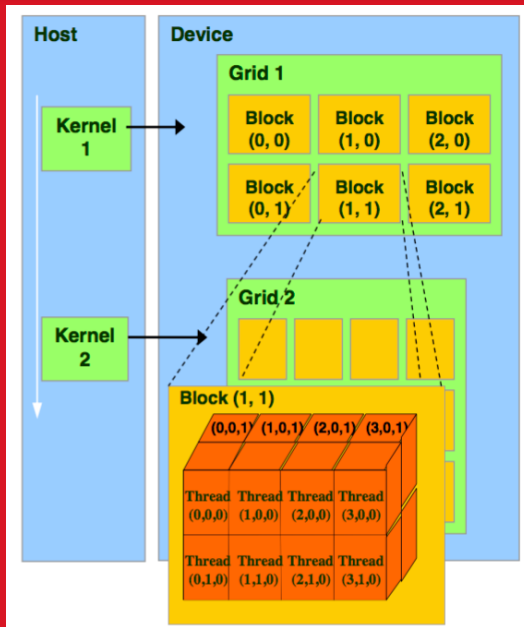
Original CUDA Code	Migrated SYCL Code
<pre>__global__ void foo() { int a = threadIdx.x; } int main() { dim3 size_1(100, 200, 300); dim3 size_2(5, 10, 20); foo<<<size_1, size_2>>>(); }</pre>	<pre>void foo(sycl::nd_item<3> item) { int a = item.get_local_id(2); } int main() { sycl::queue q; sycl::range<3> size_1(300, 200, 100); sycl::range<3> size_2(20, 10, 5); q.parallel_for(sycl::nd_range<3>(size_1 * size_2, size_2), [=](sycl::nd_item<3> item) [[sycl::reqd_sub_group_size(32)]] { foo(item); }); }</pre>

Migration CUDA->SYCL

Kernel functions

```
float_to_char<<<numBlocks_mult, threadsPerBlock>>>(dimx, dimy, _d_xvec, _d_sol_data);
```

CUDA



```
q.parallel_for(  
    sycl::nd_range<3>(numBlocks_mult * threadsPerBlock, threadsPerBlock),  
    [=](sycl::nd_item<3> item_ct1) {  
        float_to_char(dimx, dimy, _d_xvec, _d_sol_data, item_ct1);  
    });
```

SYCL

Memory Model

Original CUDA Code

```
__global__ void foo() {  
    __shared__ int shm[16];  
    shm[0] = 2;  
}  
  
int main() {  
    foo<<<1, 1>>>();  
}
```

Migrated SYCL Code

```
void foo(int *shm) {  
    shm[0] = 2;  
}  
  
int main() {  
    sycl::queue q;  
    q.submit([&](sycl::handler &cgh) {  
  
        sycl::local_accessor<int>  
shm_acc(sycl::range<1>(16), cgh);  
        cgh.parallel_for(  
            sycl::nd_range<3>(sycl::range<3>(1, 1,  
1), sycl::range<3>(1, 1, 1)), [=]  
(sycl::nd_item<3> item_ct1) {  
                foo(shm_acc.get_pointer());  
            });  
    });  
};
```

Shared memory

Global, constant and unified memory

Original CUDA Code

```
void foo() {  
    int *mem1, *mem2;  
  
    cudaMalloc(&mem1, 10);  
    cudaMallocManaged(&mem2, 10);  
}
```

Migrated SYCL Code

```
void foo() {  
    sycl::queue q;  
    int *mem1, *mem2;  
  
    mem1 = sycl::malloc_device<int>(10, q);  
    mem2 = sycl::malloc_shared<int>(10, q);  
}
```

Migrate CUDA->SYCL

CUBLAS -> Intel OneMKL (Math Kernel Library)

```
rho = cublasSdot(N, (float *)_d_p, 1, (float *)_d_p, 1);
```

CUDA

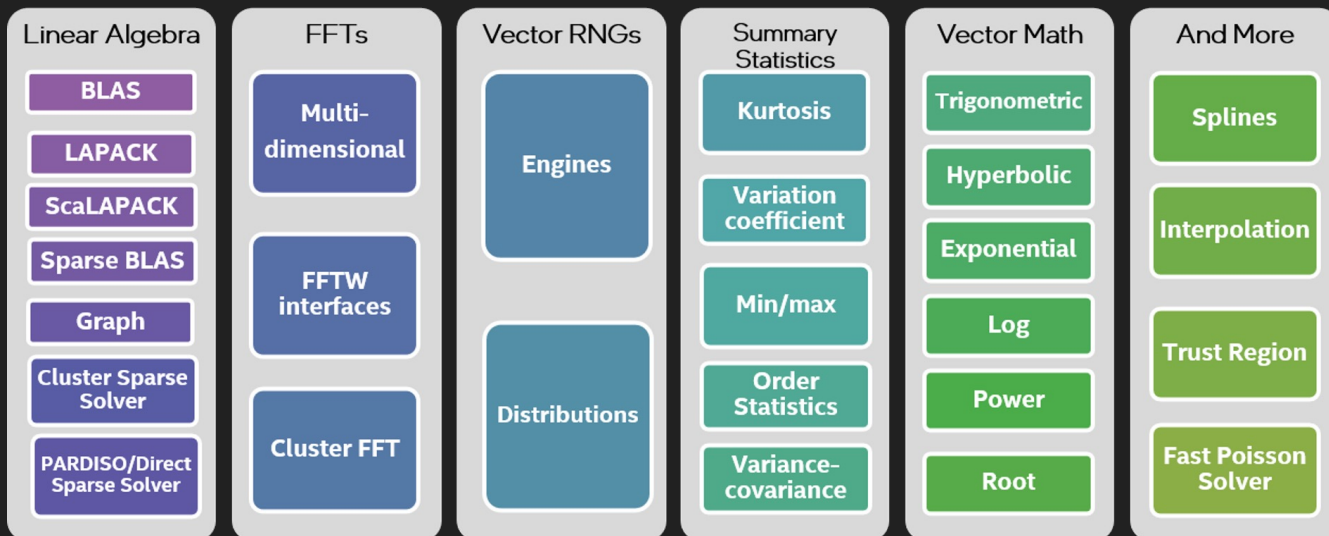
```
float *res_temp_ptr_ct1 =
```

```
    sycl::malloc_shared<float>(1, dpct::get_default_queue());
```

Intel OneMKL

```
rho = *res_temp_ptr_ct1;
```

```
oneapi::mkl::blas::column_major::dot(q, N, (float *)_d_p, 1, (float *)_d_p, 1, &rho).wait();
```



Conclusions

- **SYCL enabled portability of software for use in the Cloud on different computing devices**
- **The Intel OneAPI compatibility tools allowed easy transition from vendor specific implementation of image blending algorithm to SYCL**
- **Experimental results demonstrated great performance portability**
- **Next Steps: SYCL implementation of ZFP compression library**

Reference: "Interactive Visualization and Portable Image Blending of Massive Aerial Image Mosaics". Steve Petruzza, Brian Summa, Amy Gooch, Christine Laney, Tristan Goulden, John Schreiner, Steven Callahan, Valerio Pascucci.
In press at *International Workshop on Big Data Analytics for Climate Change, IEEE International Conference on Big Data, 2023*