

# ATLAS's Accelerator Language Choice... *is* C++

Attila Krasznahorkay

# The Setup



A	CS	14	HL-LHC Technology decision: CUDA or one of its less-proprietary competitors	3/31/2024			3/31/2024	Q1 2024	Will be concluded when developer documentation has been provided.
A	CS	14.1	Full parallelization pattern recommendation to collaboration	3/31/2024			3/31/2024	Q1 2024	
A	CS	14.2	Design patterns/tutorial on GPU migration	3/31/2024			3/31/2024	Q1 2024	

1.3.4 EF Tracking	Milestone	Choice of FPGA family	2024-03-21	2024-03-21
1.3.4 EF Tracking	Milestone	Choice of GPU language	2024-03-21	2024-03-21
1.3.4 EF Tracking	Milestone	Choice of host interface platform	2024-03-21	2024-03-21

- We have milestones defined for choosing a language to program GPUs with, in both the S&C R2R4 planning document and in TDAQ's EF Tracking schedule
  - Today's talk is meant to address both of these milestones, and put them both to rest (for the time being...)


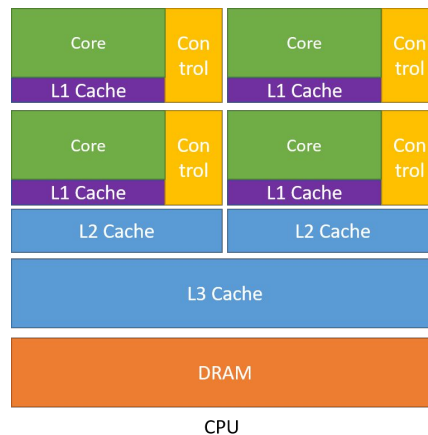
- What this talk is meant to address is
  - How custom, hand-written algorithms, meant to run on GPUs, are to be written for the offline and trigger software
  - How people should organise their code
  - Just a little bit of how people are expected to build/test their code
- What this talk is **not** meant to address is
  - What sort of GPUs ATLAS would buy/use in the short and long terms
  - How we would deal with machine learning training and inference
  - How we would (efficiently) schedule the execution of hand-written GPU algorithms in our offline and trigger software

# GPU Programming Basics

# The Architecture



- GPUs have a lot more processing cores than CPUs do
  - And they can have orders of magnitudes more threads in flight than a CPU
- However all those cores are not nearly as independent as CPU cores are
  - Task based multithreading, like what we do in the offline/trigger code, does not fit to them
  - We must use a SIMT “approach” for our code
- Generally, we write functions (kernels) that would be executed on tens / hundreds of thousands of threads at the same time



# Memory Management

- All “primary” languages provide low-level ways of (de-)allocating and copying memory
- Which APIs are fairly easy to write GPU SDK independent abstractions on top of
  - One such abstraction ([vecmem](#)) is available in Athena since December
- Generally, memory handling is not the biggest issue when choosing a GPU SDK / language

```

__host__ __device__ cudaError_t cudaMalloc ( void** devPtr , size_t size )

```

Allocate memory on the device.

**Parameters**

`devPtr`  
- Pointer to allocated device memory

`size`  
- Requested allocation size in bytes

```

void * malloc_device (size_t size, const queue &q, const property_list &propList, const detail::code_location
&CodeLoc=detail::code_location::current())
void * aligned_alloc_device (size_t alignment, size_t size, const device &dev, const context &ctx, const
detail::code_location &CodeLoc=detail::code_location::current())
void * aligned_alloc_device (size_t alignment, size_t size, const device &dev, const context &ctx, const
property_list &propList, const detail::code_location &CodeLoc=detail::code_location::current())
void * aligned_alloc_device (size_t alignment, size_t size, const queue &q, const detail::code_location
&CodeLoc=detail::code_location::current())

```

```

◆ jagged_vector_buffer() [3/3]
template<typename TYPE >
template<typename SIZE_TYPE , std::enable_if_t< !std::is_integral< SIZE_TYPE >::value &&std::is_unsigned< SIZE_TYPE >::value, bool >>
vecmem::data::jagged_vector_buffer< TYPE >::jagged_vector_buffer ( const std::vector< SIZE_TYPE > & capacities,
memory_resource & resource,
memory_resource * host_access_resource = nullptr,
buffer_type type = buffer_type::fixed_size
)

```

Constructor from a vector of (“inner vector”) sizes.

**Parameters**

`capacities` Simple vector holding the capacities/sizes of the “inner vectors” for the jagged vector buffer.

`resource` The device accessible memory resource, which may also be host accessible.

`host_access_resource` An optional host accessible memory resource. Needed if resource is not host accessible.

`type` The type (resizable or not) of the buffer

# Writing / Launching Kernels



```
277 // Kernel filling a jagged vector to its capacity
278 __global__ void fillTransformKernel(
279     vecmem::data::jagged_vector_view<int> vec_data) {
280
281     // Find the current index.
282     const std::size_t i = hipBlockIdx_x * hipBlockDim_x + hipThreadIdx_x;
283     if (i >= vec_data.size()) {
284         return;
285     }
286
287     // Create a device vector on top of the view.
288     vecmem::jagged_device_vector<int> vec(vec_data);
289
290     // Fill the vectors to their capacity.
291     while (vec[i].size() < vec[i].capacity()) {
292         vec[i].push_back(1);
293     }
294 }
295
296 void fillTransform(vecmem::data::jagged_vector_view<int> vec) {
297
298     // Launch the kernel
299     hipLaunchKernelGGL(fillTransformKernel, vec.size(), 1, 0, nullptr, vec);
300
301     // Check whether it succeeded to run.
302     VECMEM_HIP_ERROR_CHECK(hipGetLastError());
303     VECMEM_HIP_ERROR_CHECK(hipDeviceSynchronize());
304 }
```

```
108 // Now define the Alpaka work division
109 auto const deviceProperties = ::alpaka::getAccDevProps<Acc>(devAcc);
110 auto const maxThreadsPerBlock = deviceProperties.m_blockThreadExtentMax[0];
111 auto const threadsPerBlock = maxThreadsPerBlock;
112 auto const blocksPerGrid =
113     (sp_size + threadsPerBlock - 1) / threadsPerBlock;
114 auto const elementsPerThread = 1u;
115 auto workDiv = WorkDiv{blocksPerGrid, threadsPerBlock, elementsPerThread};
116 auto bufAcc = ::alpaka::allocBuf<float, uint32_t>(devAcc, sp_size);
117
118 ::alpaka::exec<Acc>(queue, workDiv, CountGridCapacityKernel{}, m_config,
119     &m_axes.first, &m_axes.second, spacepoints_view,
120     &grid_capacities_view);
121 ::alpaka::wait(queue);
```

- In all cases we write “some function”
  - Depending on the language this may need to be a standalone function, or could be even something like a functor or lambda
  - The compiler needs to recognize it as a “kernel” function, to generate the appropriate binaries for it
- We tell “some runtime API” to launch this function, with a set of arguments, on a selected number of GPU threads

# Device Code

- Is mostly pretty standard C++
  - “Kernel” functions can call as many “device” functions as they wish, just like in regular C++
- With (mostly) the following extensions:
  - Atomic operations on memory shared by all / some of the threads
  - Cooperative usage of memory dedicated to a block of threads
  - Synchronization points between the threads
- Other, language specific features also exist, but were not needed in the tracking R&D so far
  - The calo clusterization R&D code uses one CUDA specific feature right now that we’ll need to see about...

```

19 namespace tracc::device {
20
21     TRACCC_HOST_DEVICE
22     inline void find_doublets(
23         const std::size_t globalIndex, const seedfinder_config& config,
24         const sp_grid_const_view& sp_view,
25         const doublet_counter_collection_types::const_view& dc_view,
26         device_doublet_collection_types::view mb_doublets_view,
27         device_doublet_collection_types::view mt_doublets_view) {
28
29         // Check if anything needs to be done.
30         const doublet_counter_collection_types::const_device doublet_counts(
31             dc_view);
32         if (globalIndex >= doublet_counts.size()) {
33             return;
34         }
35
36         // Get the middle spacepoint that we need to be looking at.
37         const doublet_counter_middle_sp_counter 73 static __global__
38         void signalToNoiseKernel(Helpers::CUDA_kernel_object<CellStateArr> cell_state_arr,
39                                 Helpers::CUDA_kernel_object<ClusterInfoArr> clusters_arr,
40                                 Helpers::CUDA_kernel_object<TopoAutomatonBrowsingTemporaries> temporaries,
41                                 const Helpers::CUDA_kernel_object<CellInfoArr> cell_info_arr,
42                                 const Helpers::CUDA_kernel_object<CellNoiseArr> noise_arr,
43                                 const Helpers::CUDA_kernel_object<GeometryArr> geometry,
44                                 const Helpers::CUDA_kernel_object<TopoAutomatonOptions> opts)
45         // Set up the device containers.
46         // Set up the device containers.
47         const const_sp_grid_device sp_grid(sp_v 77
48         device_doublet_collection_types::device 78
49         device_doublet_collection_types::device 80
50
51         // Get the spacepoint that we're evalu 82
52         // as the "middle" spacepoint.
53         const internal_spacepoint(spacepoint) m 84
54         sp_grid.bin(middle_sp_counter.m_sp_m 85
55         .at(middle_sp_counter.m_sp_m 87
56
57         // Find the reference (start) index of 88
58         // where the doublets are recorded.
59         const unsigned int mid_bot_start_idx = 92
60         const unsigned int mid_top_start_idx = 94
61
62         for (int cell = index; cell < NCaloCells; cell += grid_size)
63         {
64             const int cell_sampling = geometry->sampling(cell);
65             const float cellEnergy = cell_info_arr->energy[cell];
66
67             if (!(cell_info_arr->is_valid(cell) || opts->uses_calorimeter_by_sampling(cell_sampling))
68                 {
69                 cell_state_arr->clusterTag[cell] = TACTag::make_invalid_tag();
70                 temporaries->secondary_array[cell] = TACTag::make_invalid_tag();
71                 continue;
72             }
73
74             float sigNoiseRatio = 0.00001f;
75             //It's what's done in the CPU implementation...
76             if (cell_info_arr->is_bad(cell, opts->treat_L1_predicted_as_good)
77                 {
78                 const int gain = cell_info_arr->gain[cell];
79
80                 float cellNoise = 0.f;
81
82                 if (opts->use_two_gaussian && geometry->is_tile(cell))
83                 {
84                     cellNoise = noise_arr->get_double_gaussian_noise(cell, gain, cellEnergy);
85                 }
86                 else
87             }
88         }
89     }
90 }

```



# GPU R&D Takeaways

# Considered / Tested Languages

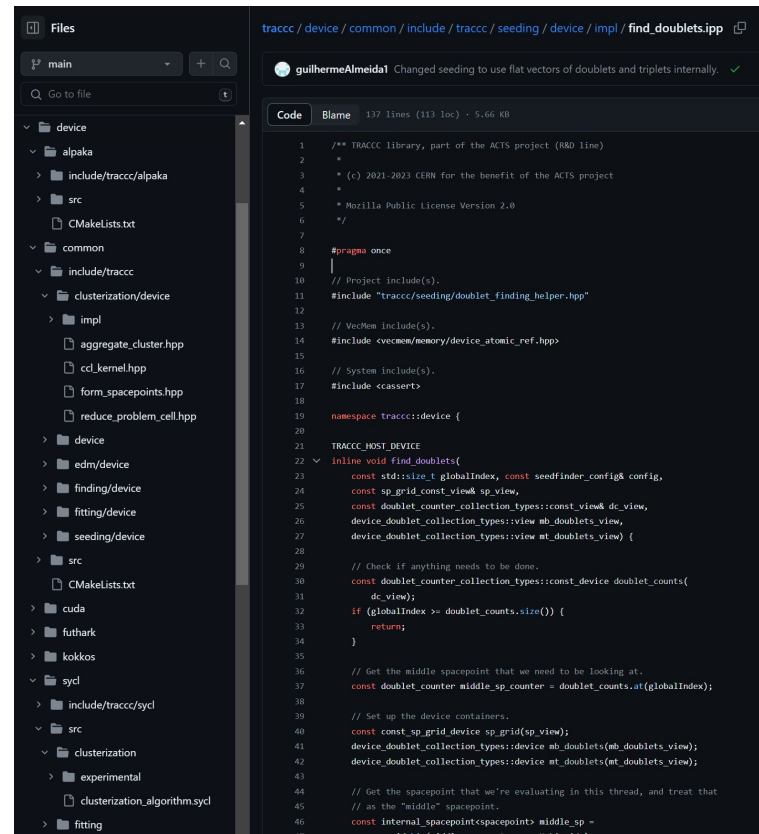


	<u>CUDA</u>	<u>HIP</u>	<u>SYCL</u>	<u>Kokkos</u>	<u>Alpaka</u>	<u>std::par</u>
NVIDIA	✓	✓	✓	✓	✓	✓
AMD	✗	✓	✓	✓	✓	✓
Intel	✗	✗	✓	✓	✓	✓

- The “support matrix” is a bit misleading
  - To produce NVIDIA, AMD and Intel binaries, you must have CUDA, HIP and oneAPI available respectively. Something like Alpaka, or even oneAPI, needs CUDA in the background for producing NVIDIA binaries!
- I.e. the “abstraction layers” don’t help a lot with “platform support” or licensing 😞

# The tracc Lesson(s)

- In the GPU Tracking R&D ([tracc](#)) we've set up a very large amount of code sharing between the different GPU languages
  - Just by carefully thinking it over how we should organise the source code
- The amount of language specific code is not negligible, but can be written/translated **very** automatically
  - Didn't try it myself with tracc, but [HIPIFY](#), [DPCT](#), [SYCLomatic](#), etc. could even do >90% of this automated work when/if needed.



The screenshot shows a code editor with a file explorer on the left and a code editor on the right. The file explorer shows a directory structure for 'device' with sub-directories like 'alpa', 'common', 'edm/device', 'finding/device', 'fitting/device', 'seeding/device', 'src', and 'sycl'. The code editor shows the file 'tracc/device/common/include/tracc/seeding/device/impl/find\_doublets.hpp' with the following content:

```

1  /** TRACCC Library, part of the ACTS project (R&D line)
2  *
3  * (c) 2021-2023 CERN for the benefit of the ACTS project
4  *
5  * Mozilla Public License Version 2.0
6  */
7
8  #pragma once
9
10 // Project include(s).
11 #include "tracc/seeding/doublet_finding_helper.hpp"
12
13 // VecMem include(s).
14 #include <vecmem/memory/device_atomic_ref.hpp>
15
16 // System include(s).
17 #include <cassert>
18
19 namespace tracc::device {
20
21 TRACCC_HOST_DEVICE
22 inline void find_doublets(
23     const std::size_t globalIndex, const seedfinder_config config,
24     const sp_grid const_view sp_view,
25     const doublet_counter_collection_types::const_view dc_view,
26     device_doublet_collection_types::view mb_doublets_view,
27     device_doublet_collection_types::view mt_doublets_view) {
28
29     // Check if anything needs to be done.
30     const doublet_counter_collection_types::const_device doublet_counts(
31         dc_view);
32     if (globalIndex >= doublet_counts.size()) {
33         return;
34     }
35
36     // Get the middle spacepoint that we need to be looking at.
37     const doublet_counter middle_sp_counter = doublet_counts.at(globalIndex);
38
39     // Set up the device containers.
40     const const_sp_grid device_sp_grid(sp_view);
41     device_doublet_collection_types::device mb_doublets(mb_doublets_view);
42     device_doublet_collection_types::device mt_doublets(mt_doublets_view);
43
44     // Get the spacepoint that we're evaluating in this thread, and treat that
45     // as the "middle" spacepoint.
46     const internal_spacepoint<spacepoint> middle_sp =

```

# The traccs Lesson(s)



CUDA

```
10 // Project include(s).
11 #include "traccs/definitions/qualifiers.hpp"
12 #include "traccs/device/fill_prefix_sum.hpp"
13 #include "traccs/edm/device/doublet_counter.hpp"
14 #include "traccs/seeding/detail/seeding_config.hpp"
15 #include "traccs/seeding/detail/spacepoint_grid.hpp"
16
17 // System include(s).
18 #include <cstring>
19
20 namespace traccs::device {
21
22     /// Function used for calculating the number of spacepoint doublets
23     ///
24     /// The count is necessary for allocating the appropriate amount of memory
25     /// for storing the information of the candidates in a next step.
26     ///
27     /// @param[in] globalIndex The index of the current thread
28     /// @param[in] config Seedfinder configuration
29     /// @param[in] sp_view The spacepoint grid to count doublets on
30     /// @param[in] sp_ps_view Prefix sum for iterating over the spacepoint grid
31     /// @param[out] doublet_view Collection storing the number of doublets for each
32     /// spacepoint
33     /// @param[out] nMidBot Total number of middle-bottom doublets
34     /// @param[out] nMidTop Total number of middle-top doublets
35     ///
36     TRACCC_HOST_DEVICE
37     inline void count_doublets(
38         std::size_t globalIndex, const seedfinder_config& config,
39         const sp_grid_const_view& sp_view,
40         const vecmem::data::vector_view<const prefix_sum_element_t>& sp_ps_view,
41         doublet_counter_collection_types::view doublet_view, unsigned int& nMidBot,
42         unsigned int& nMidTop);
43
44 } // namespace traccs::device
45
46 // Include the implementation.
47 #include "traccs/seeding/device/impl/count_doublets.hpp"
```

Generic C++

```
37 namespace traccs::cuda {
38 namespace kernels {
39
40     /// CUDA kernel for running @c traccs::device::count_doublets
41     __global__ void count_doublets(
42         seedfinder_config config, sp_grid_const_view sp_grid,
43         vecmem::data::vector_view<const device::prefix_sum_element_t> sp_prefix_sum,
44         device::doublet_counter_collection_types::view doublet_counter,
45         unsigned int& nMidBot, unsigned int& nMidTop) {
46
47         device::count_doublets(threadIdx.x + blockIdx.x * blockDim.x, config,
48                               sp_gr 192 // Count the number of doublets that we need to produce.
49                               nMidT 193 kernels::count_doublets<<<nDoubletCountBlocks, nDoubletCountThreads, 0,
50                               194 stream>>>)(
51                               195 m_seedfinder_config, g2_view, sp_grid_prefix_sum_buff,
52                               196 doublet_counter_buffer, (*globalCounter_device).m_nMidBot,
53                               197 (*globalCounter_device).m_nMidTop);
54                               198 CUDA_ERROR_CHECK(cudaGetLastError());
55 }
```

```
119 details::get_queue(m_queue)
120 .submit([& (::sycl::handler& h) {
121     h.parallel_for<kernels::count_doublets>(
122         doubletCountRange,
123         [config = m_seedfinder_config, g2_view, sp_grid_prefix_sum_view,
124          doublet_counter_view,
125          aux_globalCounter] (::sycl::nd_item<1> item) {
126             device::count_doublets(item.get_global_linear_id(), config,
127                                   g2_view, sp_grid_prefix_sum_view,
128                                   doublet_counter_view,
129                                   (*aux_globalCounter).m_nMidBot,
130                                   (*aux_globalCounter).m_nMidTop);
131         });
132     });
133     .wait_and_throw();
```

SYCL

# Conclusions

# The Choice of No Choice



- Instead of buying into a specific SDK, we have to structure all our new GPU code such as to make it easy/trivial to use different SDKs with the same “core” code
  - Performance penalties for using a “non-native” SDK are pretty minimal at the moment, but we will continue monitoring this
- Depending on how licensing and technical developments go, we may very well come out with a recommended SDK in the end
  - But even at that point, code will be structures so that it would still be easy to use from other SDKs as well at a later date
- For now, CUDA will be the easiest to use with Athena nightlies, inside the CERN firewall
  - oneAPI can be used already today from CVMFS, with a bit of manual environment setup
    - Even with GCC 13! Today!
  - HIP can not be installed on CVMFS just yet, but hopefully soon...

- If you're just now starting out, have a look at:
  - [Control/AthenaExamples/AthExCUDA](#)
  - [Control/AthenaExamples/AthExSYCL](#)
  - Will add slightly more elaborate examples, with code sharing between CUDA and SYCL, in not too long
- If you have a working setup already, just continue using it
  - Though if you're not in contact with people from HCAF, please get in touch with us! To make sure that your code would be future-proof.

# Backup



# Code Snippets



```
10 namespace tracco::device {
11
12     TRACCC_HOST_DEVICE
13     inline void form_spacepoints(
14         const std::size_t globalIndex,
15         measurement_collection_types::const_view measurements_view,
16         cell_module_collection_types::const_view modules_view,
17         const unsigned int measurement_count,
18         spacepoint_collection_types::view spacepoints_view) {
19
20         // Get device copy of input parameters
21         const measurement_collection_types::const_device measurements_device(
22             measurements_view);
23
24         // Check if anything needs to be done
25         if (globalIndex >= measurement_count) {
26             return;
27         }
28
29         // Get device copy of input parameters
30         const cell_module_collection_types::const_device modules_device(
31             modules_view);
32
33         spacepoint_collection_types::device spacepoints_device(spacepoints_view);
34
35         // Get the measurement for this index
36         const measurement& meas = measurements_device.at(globalIndex);
37         // Get the current cell module
38         const cell_module& mod = modules_device.at(meas.module_link);
39         // Form a spacepoint based on this measurement
40         point3 local_3d = {meas.local[0], meas.local[1], 0.};
41         point3 global = mod.placement.point_to_global(local_3d);
42
43         // Fill the result object with this spacepoint
44         spacepoints_device[globalIndex] = {global, meas};
45     }
46 } // namespace tracco::device
```

```
19 namespace tracco::device {
20
21     TRACCC_HOST_DEVICE
22     inline void count_doublets(
23         const std::size_t globalIndex, const seedfinder_config& config,
24         const sp_grid_const_view& sp_view,
25         const vecmem::data::vector_view<const prefix_sum_element_t>& sp_ps_view,
26         doublet_counter_collection_types::view doublet_view, unsigned int& nMidBot,
27         unsigned int& nMidTop) {
28
29         // Check if anything needs to be done.
30         vecmem::device_vector<const prefix_sum_element_t> sp_prefix_sum(sp_ps_view);
31         if (globalIndex >= sp_prefix_sum.size()) {
32             return;
33         }
34
35         // Get the middle spacepoint that we need to be looking at.
36         const prefix_sum_element_t middle_sp_idx = sp_prefix_sum[globalIndex];
37
38         // Set up the device containers.
39         const const_sp_grid_device sp_grid(sp_view);
40         doublet_counter_collection_types::device doublet_counter(doublet_view);
41
42         // Get the spacepoint that we're evaluating in this thread, and treat that
43         // as the "middle" spacepoint.
44         const internal_spacepoint<spacepoint> middle_sp =
45             sp_grid.bin(middle_sp_idx.first).at(middle_sp_idx.second);
46
47         // The IDs of the neighbouring bins along the phi and Z axes of the
48         // grid.
49         const detray::dindex_range phi_bins =
50             sp_grid.axis_p0().range(middle_sp.phi(), config.neighbor_scope);
51         const detray::dindex_range z_bins =
52             sp_grid.axis_p1().range(middle_sp.z(), config.neighbor_scope);
53         assert(z_bins[0] <= z_bins[1]);
54
55         // The number of middle-bottom candidates found for this thread's middle
56         // spacepoint.
57         unsigned int n_mb_cand = 0;
58         // The number of middle-top candidates found for this thread's middle
59         // spacepoint.
60         unsigned int n_mt_cand = 0;
```

```
123     template <typename barrier_t>
124     TRACCC_DEVICE inline void cell_kernel(
125         const index_t threadId, const index_t blkDim, const unsigned int blockId,
126         const cell_collection_types::const_view cells_view,
127         const cell_module_collection_types::const_view modules_view,
128         const index_t max_cells_per_partition,
129         const index_t target_cells_per_partition, unsigned int& partition_start,
130         unsigned int& partition_end, unsigned int& outI, index_t* f, index_t* gf,
131         barrier_t& barrier, measurement_collection_types::view measurements_view,
132         unsigned int& measurement_count,
133         vecmem::data::vector_view<unsigned int> cell_links) {
134
135         // Get device copy of input parameters
136         const cell_collection_types::const_device cells_device(cells_view);
137         const cell_module_collection_types::const_device modules_device(
138             modules_view);
139         measurement_collection_types::device measurements_device(measurements_view);
140
141         const unsigned int num_cells = cells_device.size();
142
143         /*
144          * First, we determine the exact range of cells that is to be examined
145          * by this block of threads. We start from an initial range determined
146          * by the block index multiplied by the target number of cells per
147          * block. We then shift both the start and the end of the block forward
148          * (to a later point in the array); start and end may be moved different
149          * amounts.
150          */
151         if (threadId == 0) {
152             unsigned int start = blockId * target_cells_per_partition;
153             assert(start < num_cells);
154             unsigned int end =
155                 std::min(num_cells, start + target_cells_per_partition);
156             outI = 0;
157
158             /*
159              * Next, shift the starting point to a position further in the
160              * array; the purpose of this is to ensure that we are not operating
161              * on any cells that have been claimed by the previous block (if
162              * any).
163              */
164             while (start != 0 &&
165                  cells_device[start - 1].module_link ==
166                  cells_device[start].module_link &&
167                  cells_device[start].channel1 <=
168                  cells_device[start - 1].channel1 + 1) {
169                 ++start;
```

# Code Snippets



```
41 // CUDA kernel for running @c tracc::device::ccl_kernel
42 __global__ void ccl_kernel(
43     const cell_collection_types::const_view cells_view,
44     const cell_module_collection_types::const_view modules_view,
45     const index_t max_cells_per_partition,
46     const index_t target_cells_per_partition,
47     measurement_collection_types::view measurements_view,
48     unsigned int& measurement_count,
49     vecmem::data::vector_view<unsigned int> cell_links) {
50     __shared__ unsigned int partition_start, partition_end;
51     __shared__ unsigned int outi;
52     extern __shared__ index_t shared_v[];
53     index_t* f = &shared_v[0];
54     index_t* f_next = &shared_v[max_cells_per_partition];
55     tracc::cuda::barrier barry_r;
56
57     device::ccl_kernel(threadIdx.x, blockDim.x, blockIdx.x, cells_view,
58                       modules_view, max_cells_per_partition,
59                       target_cells_per_partition, partition_start,
60                       partition_end, outi, f, f_next, barry_r,
61                       measurements_view, measurement_count, cell_links);
62 }
```

```
127 // Launch ccl kernel. Each thread will handle a single cell.
128 kernels::
129     ccl_kernel<<<(num_partitions, threads_per_partition,
130                 2 * max_cells_per_partition * sizeof(index_t), stream)>>(
131         cells, modules, max_cells_per_partition,
132         m_target_cells_per_partition, measurements_buffer,
133         *num_measurements_device, cell_links);
```

```
117 // Run ccl kernel
118 details::get_queue(m_queue)
119     .submit([&](::sycl::handler& h) {
120         vecmem::sycl::local_accessor<unsigned int> shared_uint(3, h);
121         vecmem::sycl::local_accessor<index_t> shared_idx(
122             2 * max_cells_per_partition, h);
123
124         h.parallel_for<kernels::ccl_kernel>(
125             cclKernelRange, [=](::sycl::nd_item<1> item) {
126                 index_t* f = &shared_idx[0];
127                 index_t* f_next = &shared_idx[max_cells_per_partition];
128                 unsigned int& partition_start = shared_uint[0];
129                 unsigned int& partition_end = shared_uint[1];
130                 unsigned int& outi = shared_uint[2];
131                 tracc::sycl::barrier barry_r(item);
132
133                 device::ccl_kernel(
134                     item.get_local_linear_id(), item.get_local_range(0),
135                     item.get_group_linear_id(), cells, modules,
136                     max_cells_per_partition, target_cells_per_partition,
137                     partition_start, partition_end, outi, f, f_next,
138                     barry_r, measurements_view,
139                     *aux_num_measurements_device, cell_links_view);
140             });
141         });
142     .wait_and_throw();
```



<http://home.cern>