



# Vecpar - A portability parallelization library

---

Georgiana Mania<sup>1,2</sup>   Nicholas Styles<sup>1</sup>   Michael Kuhn<sup>3</sup>   Andreas Salzburger<sup>4</sup>   Beomki Yeo<sup>5</sup>  
Thomas Ludwig<sup>2</sup>

<sup>1</sup>DESY   <sup>2</sup>University of Hamburg   <sup>3</sup>Otto von Guericke University   <sup>4</sup>CERN   <sup>5</sup>University of California, Berkeley

Motivation

Vecpar library

Code example

Evaluation

Conclusion

- 7-10x increase in data acquisition rate
- 10x more events (both real and simulated)
- higher degree of event complexity (mostly due to an increase in track multiplicity)
- higher physics fidelity is required
- increased efficiency needed in exploiting the available hardware resources




Figure 1: Computing model for ATLAS

<sup>1</sup>ATLAS Software and Computing HL-LHC Roadmap [Col22]

World largest supercomputers use accelerators to boost performance and reduce electricity consumption per computation, with NVIDIA being the most popular vendor [SDSM]

- Summit (US#1) has 27,648 NVIDIA V100s
- JUWELS Booster (EU#1) has 2,744 NVIDIA A100

Some of the next exascale supercomputers will employ other vendors like

- AMD for Frontier Supercomputer at ORNL and El Capitan Supercomputer at LLNL
- Intel for Aurora Supercomputer at ANL

*Code portability is now more important than ever!*

Many general solutions to target accelerators with different levels of complexity, portability and performance through

- language extensions and proprietary frameworks: NVIDIA CUDA, AMD HIP, Intel oneAPI  
*(limited portability & code rewrite needed using new language)*
- compiler directives (vendor-agnostic): OpenMP, OpenACC  
*(limited portability when using single source code)*
- portability libraries: Kokkos, Alpaka  
*(code rewrite needed to fit the abstractions)*

*Choosing the right solution depends on the application and the accepted compromise.*

- Easy to use in plain C++ & limited code changes to adopt the API
- Abstract the notion of architecture (as much as possible)
- Support single-source code & generated code for different targets and/or optimizations
- Compilable by main stream C++ compilers (e.g. clang)
- Portable (target shared-memory CPU and NVIDIA GPUs) but also easily extendable to new architectures like AMD/Intel GPUs and potentially Big Data platforms (Google Cloud/AWS)
- Ensure improved wall-clock performance over the initial (sequential) implementation

*We designed vecpar having all these in mind.*

Motivation

Vecpar library

Code example

Evaluation

Conclusion

Regardless of the platform, data parallelism

- means executing the **same task** for several elements of a collection in parallel (e.g. OpenMP parallel-for loop)
- requires the data to be partitioned and distributed among the workers (e.g. threads of any kind: OpenMP, TBB, CUDA, etc)

Moreover, to **offload** computations to a GPU, some preconditions have to be met

- the function pointer to the actual parallel task has to be copied to the GPU
- the data used for the computations needs to be accessible from the device

*Vecpar handles both data distribution and host↔device transfers out of the box.*



Vecpar builds on the **map-filter-reduce** paradigm from functional programming

- (map **f** coll) - apply function **f** to each element of the **collection** to produce a different collection of the same size, with the same or different type of elements
- (filter **p** coll) - keep in the output collection only the elements from the input **collection** which satisfy the predicate **p**
- (reduce **f** coll **init**) - reduce the elements of the **collection** using function **f** and store it in the result initialized with **init**

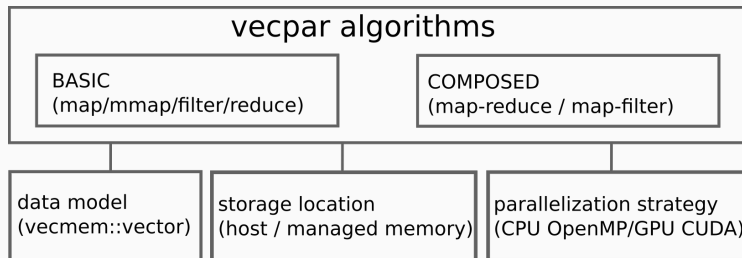
```

1 ;; transform a collection of measurements into space points
2 (map toSpacePoints `(meas1 meas2.. measN)) ;; (sp1 sp2.. spN)
3 ;; keep only the space points (sp) with activation above a threshold
4 (filter isAboveThreshold? `(sp1 sp2 sp3)) ;; (sp1 sp2)
5 ;; gather volumes during propagation; ps = propagation state
6 (reduce addVolume `(ps1 ps2.. psN) []) ;; `(volume1 volume2.. volumeM)

```

**Listing 1:** Code samples in Clojure programming language [Hic20]

vecpar algorithms offer an abstraction layer over low-level specializations enabled by compile-time polymorphism for the two backends: CPU OpenMP and GPU CUDA



**Figure 2:** Vecpar architecture

C++ code which extends vecpar algorithms can be compiled (with clang) for different architectures. cmake examples:

```
1 find_package(vecpar REQUIRED)
2
3 target_link_libraries(cpu_exe
4     vecpar::all
5     vecmem::core vecmem::cuda
6     OpenMP::OpenMP_CXX
```

**Listing 2:** cmake compile for CPU configuration

```
1 find_package(vecpar REQUIRED)
2
3 target_link_libraries(gpu_exe
4     vecpar::all
5     vecmem::core vecmem::cuda
6     CUDA::cudart)
7
8 target_compile_options(gpu_exe PUBLIC
9     $<$<COMPILE_LANGUAGE:CXX>:-x cuda
10                                     --offload -arch=sm_XY>)
```

**Listing 3:** cmake compile for GPU configuration

Motivation

Vecpar library

**Code example**

Evaluation

Conclusion

- Runge-Kutta-Nyström (RKN) stepper is used to estimate the position and momentum of charged particles by integrating the equation of motion
- It is a compute intense algorithm due to the operations of the numerical integrator
- RKN Stepper is ideal for parallelization because the estimated position can be computed in parallel for all the tracks since the data is independent
- The implementation used for this talk is available in detrax project <sup>2</sup>, which is part of the ACTS R&D line

---

<sup>2</sup><https://github.com/acts-project/detrax>

```
1  __global__ void rk_stepper_test_kernel (vecmem::data::vector_view<free_track_parameters> tracks_data, const
    ↪ vector3 B) {
2      int gid = threadIdx.x + blockIdx.x * blockDim.x;
3      vecmem::device_vector<free_track_parameters> tracks(tracks_data);
4      if (gid >= tracks.size()) { // Prevent overflow
5          return;
6      }
7      auto& traj = tracks.at(gid); // Get a track
8      // Define RK stepper
9      rk_stepper_type rk(B);
10     // Forward direction
11     rk_stepper_type::state forward_state(traj);
12     for (unsigned int i_s = 0; i_s < rk_steps; i_s++) {
13         rk.step(forward_state);
14     }
15     // Backward direction
16     traj.flip();
17     rk_stepper_type::state backward_state(traj);
18     for (unsigned int i_s = 0; i_s < rk_steps; i_s++) {
19         rk.step(backward_state);
20     }
21 }
```

**Listing 4:** CUDA implementation of the RKN stepper in detray project (Feb 2022)

```
1 struct rk_stepper_algorithm :  
2     public vecpar::algorithm::parallelizable_mmap<free_track_parameters , vector3 >{  
3  
4     TARGET free_track_parameters& map(free_track_parameters& traj , vector3 B) override {  
5  
6         // Define RK stepper  
7         rk_stepper_type rk(B);  
8         // Forward direction  
9         rk_stepper_type::state forward_state(traj);  
10        for (unsigned int i_s = 0; i_s < rk_steps; i_s++) {  
11            rk.step(forward_state);  
12        }  
13        // Backward direction  
14        traj.flip();  
15        rk_stepper_type::state backward_state(traj);  
16        for (unsigned int i_s = 0; i_s < rk_steps; i_s++) {  
17            rk.step(backward_state);  
18        }  
19        return traj;  
20    }  
21 };
```

*No index-based calculation needed → portable code*

**Listing 5:** C++ implementation of the RKN stepper using vecpar library

Motivation

Vecpar library

Code example

**Evaluation**

Conclusion




**Physics setup** 10,000 simulated tracks with origin position (0,0,0) and charge -1, uniform magnetic field of  $B=(0,0,2T)$ , RKN computes 100 integration steps for each track


Config	Environment 1	Environment 2 <sup>3</sup>
Arch	1 socket x 8 cores x 2 threads	2 sockets x 10 cores x 2 threads
CPU	Intel(R) Core(TM) i7-10870H @ 2.20GHz	Intel(R) Xeon(R) Gold 5115 @ 2.40GHz
GPU	NVIDIA GeForce RTX 3060	NVIDIA Tesla V100
CUDA Driver	510.47.03	510.47.03
CUDA Version	11.6	11.6
C++ compiler	clang 14	clang 14

**Table 1:** Test environments used for performance evaluation

<sup>3</sup>ATLAS-GPU01 node at the National Analysis Facility (NAF) at DESY



**(a)** Comparison between Env1 (16 OpenMP threads) and Env2 (40 OpenMP threads), double precision, cmath backend



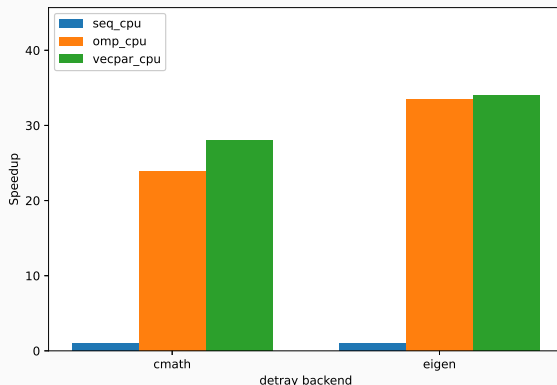
**(b)** Strong-scaling evaluation: Multi-threading implementations in simple precision for cmath/detray backends, on Env2

**Figure 3:** Mean and standard deviation for Runge-Kutta-Nyström stepper


Vecpar implementation speedups over the sequential implementation for different data storage backends

- 28x - cmath
- 34x - eigen


Vecpar implementation adds no overhead in comparison to a hard-coded OpenMP pragma implementation.



**Figure 4:** Speed-up factors over initial sequential version (seq\_cpu) for the multi-threading hard-coded OpenMP (omp\_cpu) and vecpar (vecpar\_cpu) implementations using detray cmath/eigen backends, double precision on Env2




(a) Fastmath disabled




(b) Fastmath enabled

**Figure 5:** Mean and standard deviation for Runge-Kutta-Nyström stepper with grid configuration: 157x64 CUDA threads, clang as host compiler for nvcc



(a) Fastmath disabled



(b) Fastmath enabled

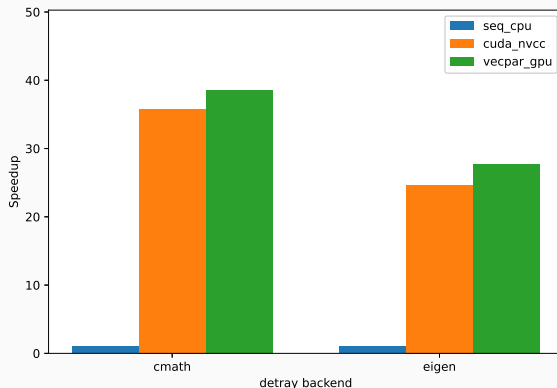
**Figure 6:** Mean and standard deviation for Runge-Kutta-Nyström stepper, cmath backend, simple /double precision, Env2

Similar results were obtained with the eigen backend.


Vecpar implementation speedups over the sequential implementation for different data storage backends

- 38x - cmath
- 27x - eigen


Vecpar implementation compiled with clang is faster than the CUDA implementation compiled with nvcc.



**Figure 7:** Speed-up factors over initial sequential version (seq\_cpu) for GPU CUDA (cuda\_nvcc) and vecpar (vecpar\_gpu), detray cmath/eigen backends in double precision, with fastmath disabled, on Env2



(a) Simple precision




(b) Double precision

**Figure 8:** Speedup for vecpar implementation over the base (sequential) CPU implementation, using cmath/eigen storage backends, in simple/double precision, on Env2

Statistics for GPU provided by  
Nsight Compute Tool

- Kernel's theoretical occupancy: 66.67%
- Kernel's achieved occupancy: 65.62%
- Kernel execution time: 128ms
- Memory throughput: 130GB/s
- L1 cache hit rate: 84.91%
- L2 cache hit rate: 92.63%



**Figure 9:** Speedup diagram for RKN stepper, cmath/eigen backends, in simple precision, 1 million tracks, fastmath disabled, on Env1



- Speedups of 27-70x over sequential implementation using the same code base compiled for different platforms
- Vecpar shows comparable wall-clock times to the hand-tuned OpenMP and CUDA implementations while vecpar is using one single-source file.

Motivation

Vecpar library

Code example

Evaluation

Conclusion

### vecpar library

- is open-source on github <sup>4</sup>
- supports x86 and NVIDIA GPUs
- uses vecmem library <sup>5</sup> to handle data structures and memory allocations.
  - currently only vecmem::vectors are supported as iterable collections for parallel loops
  - managed memory is fully supported by both backends while host memory is supported for CPU backend only (GPU work in progress)
- is covered by automated tests using googletest infrastructure <sup>6</sup>

*Disclaimer: vecpar library is in early development phase!*

---

<sup>4</sup><https://github.com/wr-hamburg/vecpar>

<sup>5</sup><https://github.com/acts-project/vecmem>

<sup>6</sup><https://github.com/google/googletest>

- Offloading of an algorithm chain on a GPU
- Use of more complex data types (e.g. 2D vectors)
- Automatic performance optimizations
- A new (HIP) backend to target AMD GPUs

*Suggestions and contributions are highly welcomed! Email me at [georgiana.mania@desy.de](mailto:georgiana.mania@desy.de)*

## References

- [Col22] ATLAS Collaboration. **ATLAS Software and Computing HL-LHC Roadmap**. Technical report, CERN, Geneva, Mar 2022.
- [Hic20] Rich Hickey. **A history of clojure**. *Proc. ACM Program. Lang.*, 4(HOPL), jun 2020.
- [SDSM] Erich Strohmaier, Jack Dongarra, Horst Simon, and Martin Meuer. **Top500 list**. <https://www.top500.org/lists/top500/2021/11/>. Accessed: 2022-05-02.