

Workshop on High Throughput Heterogeneous Computing

TRACCC: the ACTS Massively Parallel Tracking Demonstrator

Stephen Nicholas Swatman¹ on behalf of the ACTS developers

Thursday, January 29, 2026

¹CERN



Introduction – TRACCC and Next Generation Triggers

- TRACCC is an ACTS subproject towards an...
 - efficient
 - massively parallel
 - track reconstruction software package
- TRACCC is supported by the **CERN NGT** project
- Goal: “remarkably increase **efficiency, sensitivity** and **modelling** of CERN experiments”
- Through the use of **novel hardware**, including **GPGPUs** (general purpose GPUs)
- **Five-year effort** to radically advance many aspects of LHC computing
- *<https://nextgentriggers.web.cern.ch/>*



Introduction – Motivation




Image adapted from ATLAS

Introduction – Motivation





Image adapted from Karl Rupp

Introduction – GPGPU computing

For around 8,000 EUR in 2026:



AMD EPYC 9555P
64 cores
360W TDP



NVIDIA RTX PRO 6000 Blackwell
24,064 cores
300W TDP

Device	Cores	\times Cycles/s	\times FLOP/cycle	= FLOP/s
AMD EPYC 9555P	64	4.40B	64	18.0T
NVIDIA RTX 6000 BW.	24,064	2.29B	2	110.2T

Embarrassingly
parallel

Invitingly
parallel

Humblingly
parallel

Terminology due to Raph Levien

Embarrassingly
parallel

Invitingly
parallel

Humblingly
parallel

Axpy
Bitcoin mining
Shaders

Terminology due to Raph Levien

Embarrassingly
parallel

Axpy
Bitcoin mining
Shaders

Invitingly
parallel

gemm
Histogramming
FFT

Humblingly
parallel

Terminology due to Raph Levien

Embarrassingly

parallel

Axpy

Bitcoin mining

Shaders

Invitingly

parallel

gemm

Histogramming

FFT

Humblingly

parallel

SpMV

(De-)compression

Sorting

Terminology due to Raph Levien

Embarrassingly

parallel

Axpy

Bitcoin mining

Shaders

Track reconstruction

Invitingly

parallel

gemm

Histogramming

FFT

Terminology due to Raph Levien

Humblingly

parallel

SpMV

(De-)compression

Sorting

Embarrassingly

parallel

Axpy

Bitcoin mining

Shaders

Invitingly

parallel

gemm

Histogramming

FFT

Humblingly

parallel

SpMV

(De-)compression

Sorting

Track reconstruction

On GPGPUs

Track reconstruction


Terminology due to Raph Levien

Challenges – GPGPU computing

- GPU threads run in **lockstep**
- One **instruction stream** is broadcast to a **group of threads** (32–64)
- Branch divergence causes idle time
- As do **unequal loop structures**
- Behaviour much like **SIMD lanes**


	m_0	m_1	m_2	m_3	t_0	t_1	t_2	t_3
int n = thread_id();	✓	✓	✓	✓	↗	↗	↗	↗
prologue();	✓	✓	✓	✓	↓	↓	↓	↓
if (0 < n < 3) {								
branch1();	✗	✓	✓	✗	↓	↗	↗	↓
} else if (n == 0) {								
branch2();	✓	✗	✗	✗	↗	↓	↓	↓
}								
epilogue();	✓	✓	✓	✓	↓	↓	↓	↓

Track Reconstruction – Clustering



Track Reconstruction – Combinatorial Kálman Filter


- The Combinatorial Kálman Filter extends seeds
- Branches frequently, contains **nested, unbound loops**
- One of the biggest **bottlenecks** and most **complicated** algorithms
- Presents *many* challenges:
 - How do we manage the **combinatorics**?
 - How do we describe our **detector** in **GPU memory**?
 - How do we keep **magnetic field** accesses fast?



Source: Paul Gessinger


Track Reconstruction – Summary

- Around **8 subproblems** with wildly different characteristics
- Map **non-trivially** to massively parallel hardware
 - Imbalance, divergence, irregular access patterns, etc.
- Requires much more than a **naive** porting exercise!



Implementation – TRACCC

- TRACCC is our **open-source** massively parallel track reconstruction pipeline
- Designed from the ground up for GPGPUs
- Algorithms often **completely** rethought
- Aim: match **physics** performance of homogeneous solutions
- See e.g. [10.5281/ZENODO.8119504](https://doi.org/10.5281/ZENODO.8119504) for more info



The screenshot shows the GitHub repository page for 'acts-project / traccc'. The repository is public and has 20 issues, 32 pull requests, and 5 discussions. It has 29 stars, 47 forks, and 5 watchers. The repository is named 'traccc' and is located at <https://github.com/acts-project/traccc>. The repository page includes sections for 'About', 'Releases', 'Packages', and 'Contributors'.

About

Demonstrator tracking chain on accelerators

Releases 16

- traccc 0.16.0 **Latest**
3 weeks ago
- + 15 releases

Packages

No packages published

[Publish your first package](#)

Contributors 19


19 recent commits

Recent commits:


- beomki-yeo Improve the ... 9a34ae0 - yesterday
- .githooks Add some docume... 3 years ago
- .github Update to detray v... last week
- .vscode Introduced CMake ... 10 months ago
- benchmarks Update to detray v... 3 weeks ago
- cmake Add -Wold-style-c... 2 months ago
- core Improve the seedi... yesterday
- data Replace hit_id of c... last month
- device Allow spacepoint r... 2 days ago
- doc/images Migrate README fl... 2 years ago
- examples Lower the value of ... last week
- extern Cell EDM Rewrite, ... last week
- extras Replace hit_id of c... last month
- io Cell EDM Rewrite, ... last week
- performance Improve the seedi... yesterday
- plugins detray update to v... 6 months ago
- simulation Update to detray v... last week
- tests Improve the seedi... yesterday
- other recent commits

[View repository](#)


- Detector descriptions are classically polymorphic, which doesn't fly in GPGPUs
- DETRAY is our heterogeneous detector geometry
 - *Crucial* component of any non-trivial reconstruction
- Tremendous amount of work by the DETRAY devs
- See [10.1088/1742-6596/2438/1/012026](https://doi.org/10.1088/1742-6596/2438/1/012026) for more info



- Reconstruction features **highly frequent, highly irregular structured grid access**
- **COVFIE** is our library for handling arbitrary vector fields incl. magnetic fields
- **Cross-platform** performance through **compile-time composition**
- Allows e.g. use of **texture memory**
- See [10.1145/3578244.3583723](https://doi.org/10.1145/3578244.3583723) for more info



Spin-Off Projects – COVIE




Source: Microsoft

- The TRACCC effort also (indirectly) produced **models** and **methods**
- Novel **derivations of Jacobian matrices**: [10.1016/j.nima.2024.169734](https://doi.org/10.1016/j.nima.2024.169734)
- Models for **thread divergence**: [10.1109/MASCOTS56607.2022.00026](https://doi.org/10.1109/MASCOTS56607.2022.00026)
- **Genetic algorithms for structured grid layouts**: [10.1145/3629526.3645034](https://doi.org/10.1145/3629526.3645034)
- Novel method for **transparent** SoA and AoS layouts
- **Throughput models** for heterogeneous task graphs


The Detector – The ODD

- ATLAS is great, but an **open-source detector** gives us some great benefits:
 - No plot approvals
 - Free code and data **sharing**
 - Ease of use for **non-ATLAS** users
 - Freedom from the **grimy real world**
- This is why we “built” the **OpenDataDetector**




The Detector – The ODD

- The ODD served as the base for the wildly successful **TrackML** Kaggle challenge
- Also serves as the main **evaluation tool** for TRACCC
- Recently released **ColliderML**: the biggest freely available **high-luminosity dataset** for e.g. ML training
- See <https://colliderml.com/>



The Good Parts – Success in Physics!

- TRACCC provides **good physics performance** on the ODD
- And we are very nearly within limits for the **ATLAS ITk**
- Given the **from-scratch** nature of TRACCC, this is an impressive result!



April 2025

Kernel	1edca0f
fit	280.68 ms
propagate_to_next_surface	118.21 ms
find_tracks	26.36 ms
count_triplets	14.16 ms
find_triplets	5.98 ms
build_tracks	1.07 ms
Total	450.89 ms

January 2026

Kernel	9bcb542
propagate_to_next_surface	7.80 ms
find_tracks	1.73 ms
ccl_kernel	825.79 μ s
count_doublets	815.01 μ s
Total	13.16 ms

The Good Parts – Success in Compute!

- We managed to increase our performance 30× in 9 months
- Current performance makes us **competitive with CPU solutions**
- **Realistic cost savings** with current solution
- But these are **percentage** savings (not orders of magnitude)
- Perhaps the benefit will increase more?

Kernel	9bcb542
propagate_to_next_surface	7.80 ms
find_tracks	1.73 ms
ccl_kernel	825.79 µs
count_doublets	815.01 µs
Total	13.16 ms

The Lessons Learned – Portability and Code Sharing


- TRACCC set out with ambitious ideas
- Share as much code **between** CPU and GPU as possible
 - In order to reduce maintenance
- Support as many **programming models** as possible
 - In order to support many devices
 - NVIDIA CUDA, AMD HIP, SYCL, etc.
- Unfortunately, neither of these approaches really worked out
 - That's R&D for you!



- Sharing code between CPUs and GPUs is tricky
- Shareable code is generally limited, watch out for:
 - Code with any **dynamic memory allocation** (incl. `std::vector`)
 - Code with **large amounts of stack usage**
 - Code with **unbound loops** (or large bound loops)
 - Early returns, **complex control flow**
- Setting out to share too much leads to issues: **start small and unify later**


The Lessons Learned – Portability

- Our approach to portability has resulted in **high maintenance and little benefit**
- “like wearing two raincoats on top of each other”
- Cross-platform support forces meeting at the **smallest common denominator**
- Recommendation, either:
 - Focus on **performance** in *one* programming model and port later; or
 - Focus on a *single* portability solution from the start




The Lessons Learned – Performance Monitoring

- HEP has a strong culture of monitoring physics performance
- For a project like TRACCC, compute performance is also a **first-class monitorable** – at kernel level
- Only last year did we get **continuous compute monitoring**
- Critical for informing **accept–reject decisions**
- Also track performance **changes** over time to find **regressions**



The Lessons Learned – Performance Monitoring




“My CPU solution runs in 10 ms and my GPU solution runs in 4 ms, so my GPU solution is 2.5 times faster”

The Lessons Learned – The Latency Myth

“Factory A makes a carpet in 10 hours and factory B takes 4 hours, so factory B produces 2.5 times more carpets”

The Lessons Learned – The Latency Myth



8 looms, 10 h. / carpet = 0.8 carpets / h

2 looms, 4 h. / carpet = 0.5 carpets / h

“Factory A makes a carpet in 10 hours and factory B takes 4 hours, so factory B produces 2.5 times more carpets”

The Lessons Learned – The Latency Myth

- For throughput-critical applications, **latency is not enough!**
- Compute **throughput** using latency: $T = \frac{N}{L}$
 - Computation of N differs for CPUs and GPUs
- If you want a latency-like metric, use **reciprocal throughput**
 - “How long does it take to produce a carpet on average?”
 - “What is the average amount of time between carpets being finished?”
- Both **measured in time**, but **semantically different!**

- **Scheduling** and **placement** remain difficult questions for us
- **Dynamic** scheduling between CPU and GPU risks hard-to-debug **runtime issues**
- Static scheduling risks **imbalance** between CPU and GPU
 - Can be alleviated with MPI/SaaS – but needs **networking**
- Requires integration of **asynchronous** execution in Gaudi


Conclusion

- Thanks to the hard work of **many**, TRACCC is a **functional, performant** track reconstruction pipeline in ACTS
- Track reconstruction is **difficult to implement** for GPGPUs due to **irregularity**
- **Solutions** to many hurdles **researched and developed**
- TRACCC currently provides **competitive performance** for **ATLAS EF tracking**
- To get **involved**: CERN Mattermost, ACTS (#traccc and friends), **bi-weekly meeting**

Backup slides

Backup – GPGPU computing

CPU architecture



GPU architecture

