

Towards GPU-based tracking in ACTS and ATLAS

Stephen Nicholas Swatman
May 3, 2024

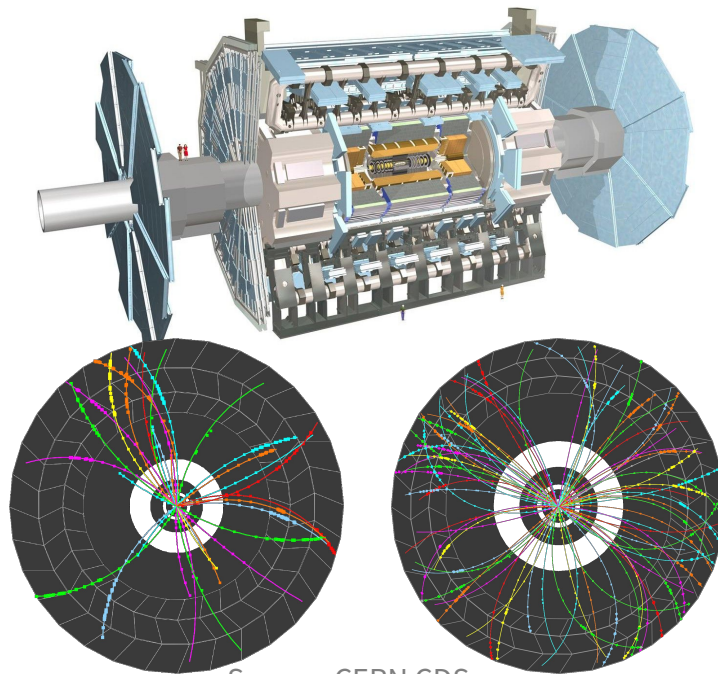


Introduction – About Me & My PhD

- Finished a MSc at **UvA + VU** in 2019, graduation project at **Nikhef**
- Started a PhD in **CS** at **UvA (PCS) + CERN** in 2020
- Now writing up – aim to defend in **autumn**
- *Golden staple*-like PhD thesis
- In this talk: very brief overview of my work, and some insights into **GPU tracking**

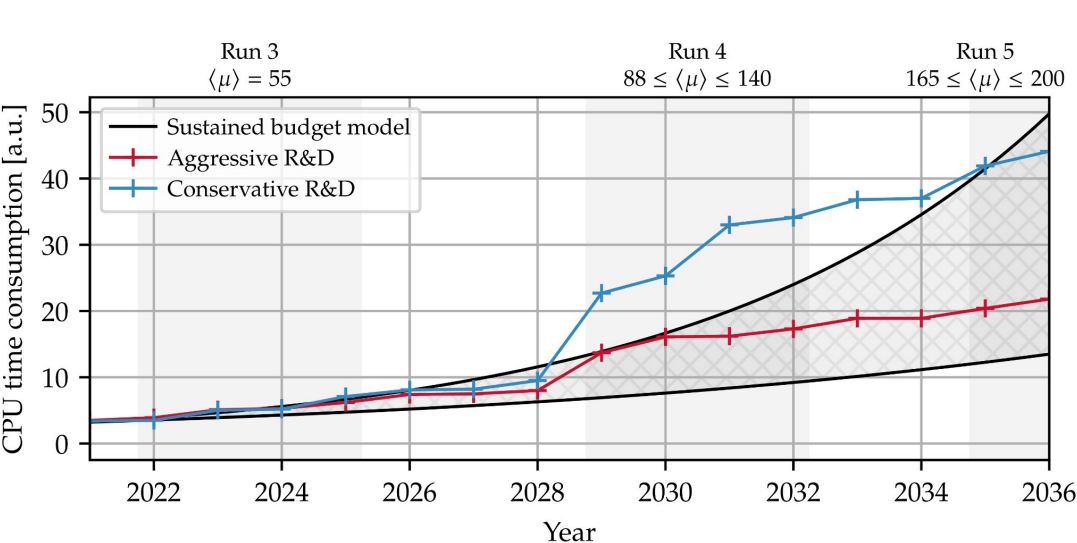
Problem Statement – Track Reconstruction

- Increased pile-up threatens to make track reconstruction infeasible for HL-LHC, FCC
- More **efficient compute** → more **available resources** → more interesting **physics**
- Compute complexity scales $\sim O(\mu^2)$
- This will require research into novel **algorithms**, novel **hardware**, etc.
- In my case: how can this run on **GPUs**?



Source: CERN CDS

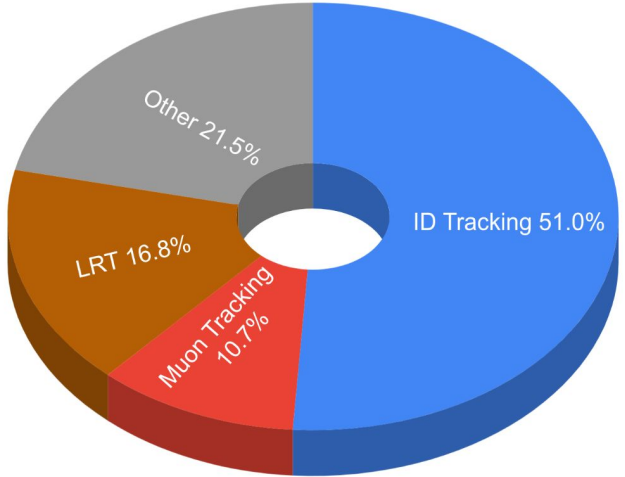
Problem Statement – Track Reconstruction



Numbers by ATLAS, plot by me

ATLAS
Updated Reconstruction
Fraction of Total Time

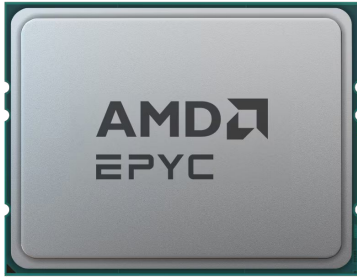
LHC Fill 6291
 $\langle \mu \rangle = 50$
4.8 s per event



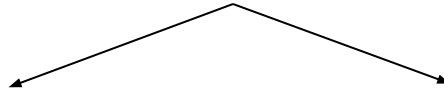
10.1007/s41781-023-00111-y

Massively Parallel Computing

€10,000

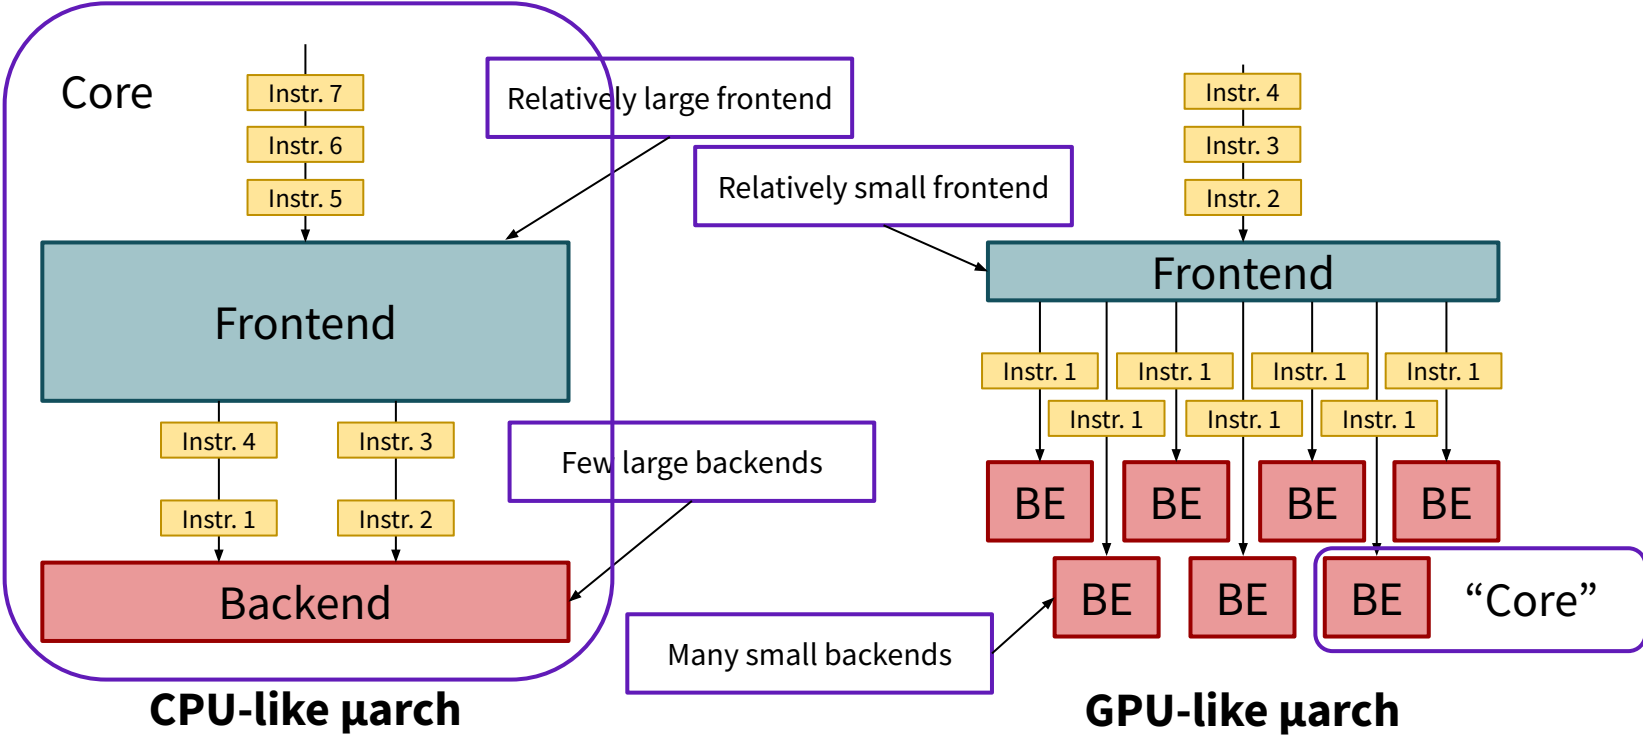


AMD EPYC 9554
64 cores
360W TDP



NVIDIA RTX 6000
18,176 cores
300W TDP

Massively Parallel Computing



Massively Parallel Computing

Can be programmed independently

Cores × **Cycles/s** × **FLOP/cycle** = **FLOP/s**



64

3.75B

48

11.5 TFLOP/s

↓ **8×**



18,176

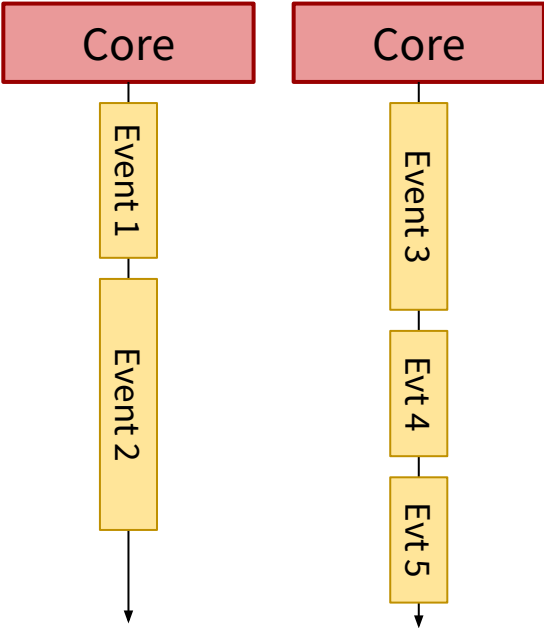
2.51B

2

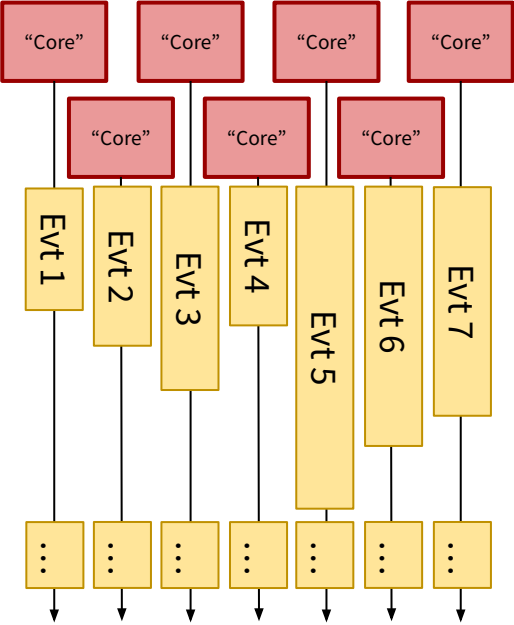
91.1 TFLOP/s

Must be carefully programmed in lockstep

Massively Parallel Computing



CPU-like march

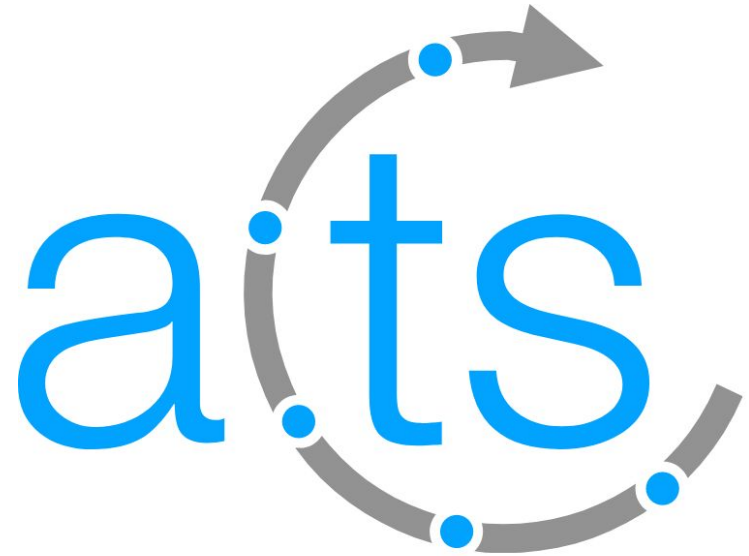


GPU-like march

**Can track reconstruction be
implemented efficiently on
massively parallel systems?**

ACTS – A Common Tracking Software

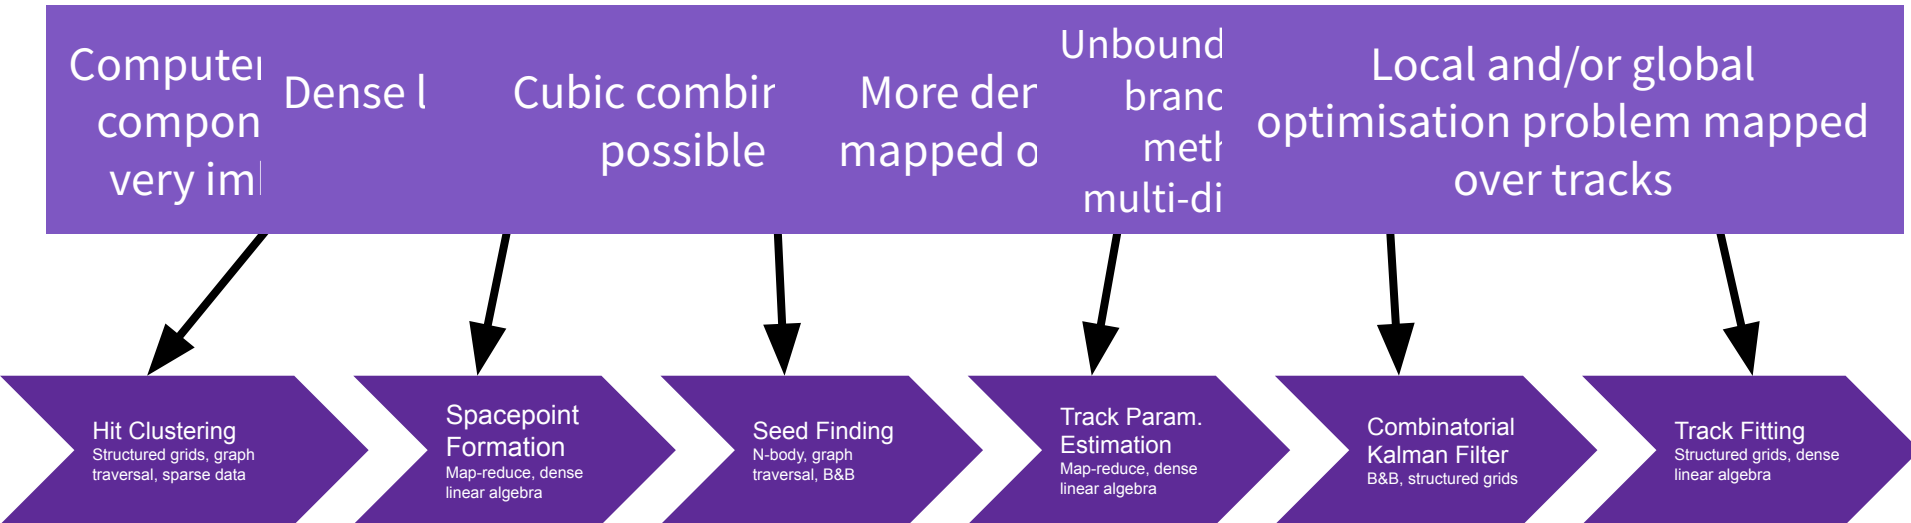
- **ACTS is A Common Tracking Software**
- Goals are to be...
 - **Feature complete**
 - **Well-written**
 - **Extensible**
 - **Experiment-agnostic** (SPHENIX, ePIC, etc.)
 - an **R&D platform**
- Originally based on **ATLAS tracking**
- Now being migrated **back into Athena**



**What are the challenges in
developing track
reconstruction algorithms to
massively parallel
architectures?**

Problem Statement – State-of-the-Art

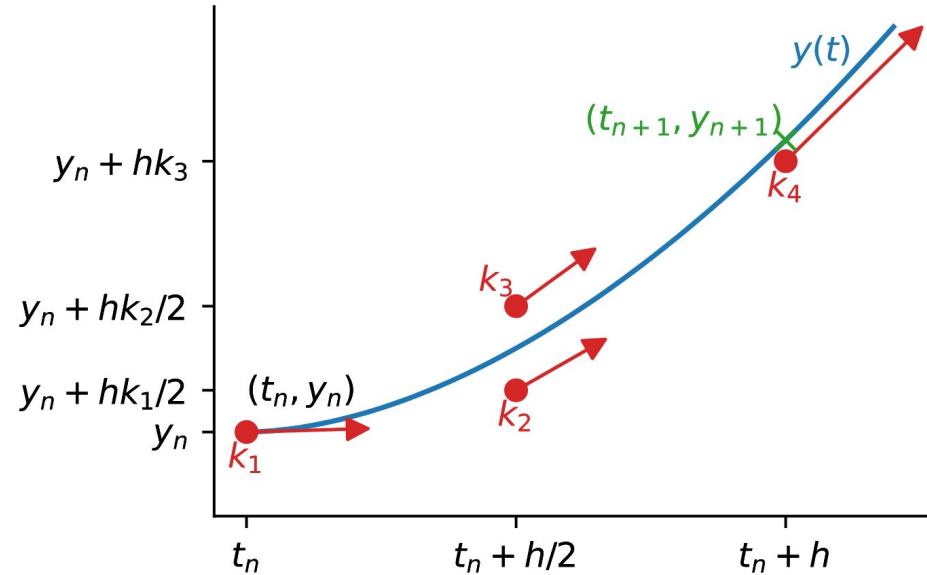
- Track reconstruction consists of a **task graph** of structurally different algorithms
- Here denoted according to their “13 dwarves” classification
- Varying degrees of **complexity** and **challenge**...



**How can structured grid data
be represented in order to
maximize the efficiency of
arbitrary computations?**

Research – Vector Fields – Amuse Bouche

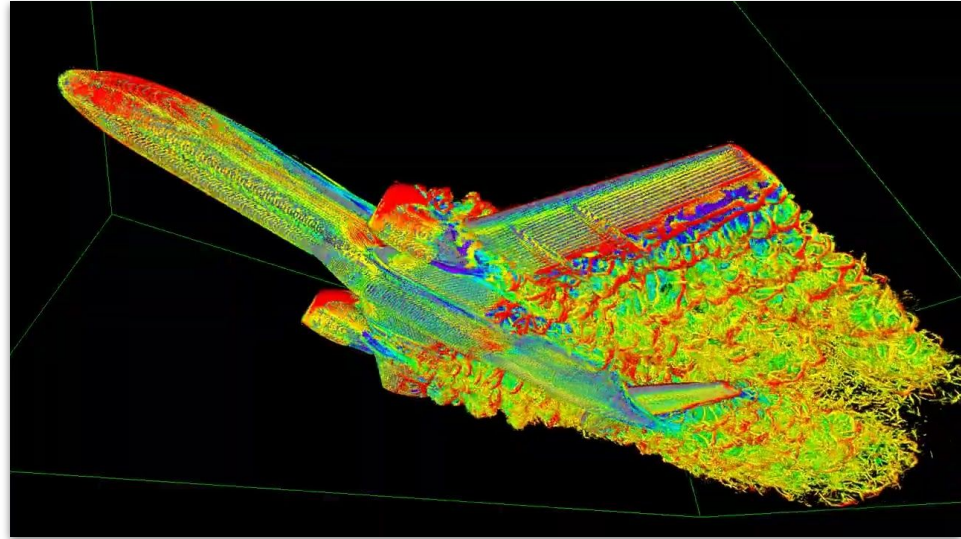
- Most experiment **magnetic fields** are fairly homogeneous, but not quite
- Those irregularities matter when e.g. **propagating particle motion**
- Storing ATLAS B-field: **~200 MB**
- **Accessed** millions of times per second!
- How do we do that **quickly**?
- We don't even know how to do that on CPUs, and if we did it would not translate!



Source: ACTS Project

Research – Vector Fields

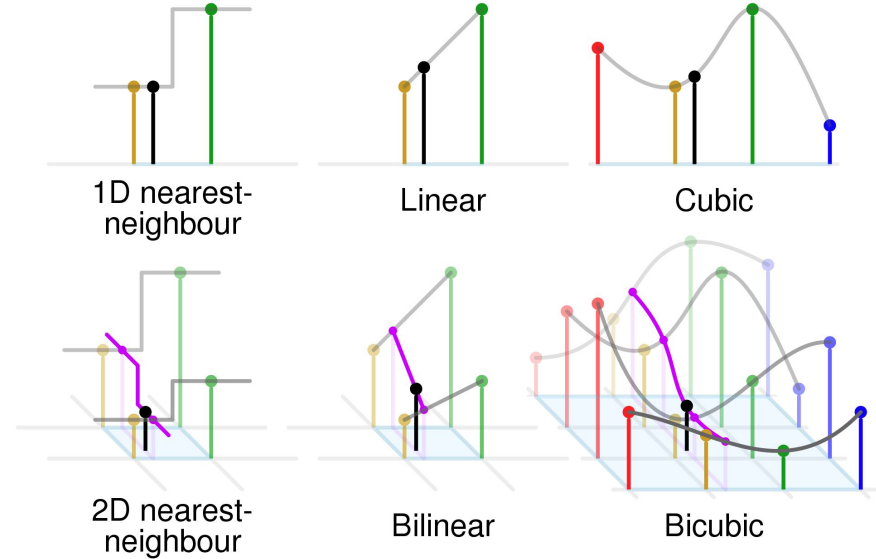
- **Multi-dimensional data** is everywhere HPC
 - Magnetic fields
 - But also CFD, lattice QCD, etc.
- Must be accessed very frequently in **hard-to-predict patterns** for iterative numerical methods
- Increasing **cache efficiency** can improve **performance** for these kernels
- **Design space is large**
 - Many **functional** and **non-functional** properties come into play



Source: Moritz Lehmann

Research – Vector Fields

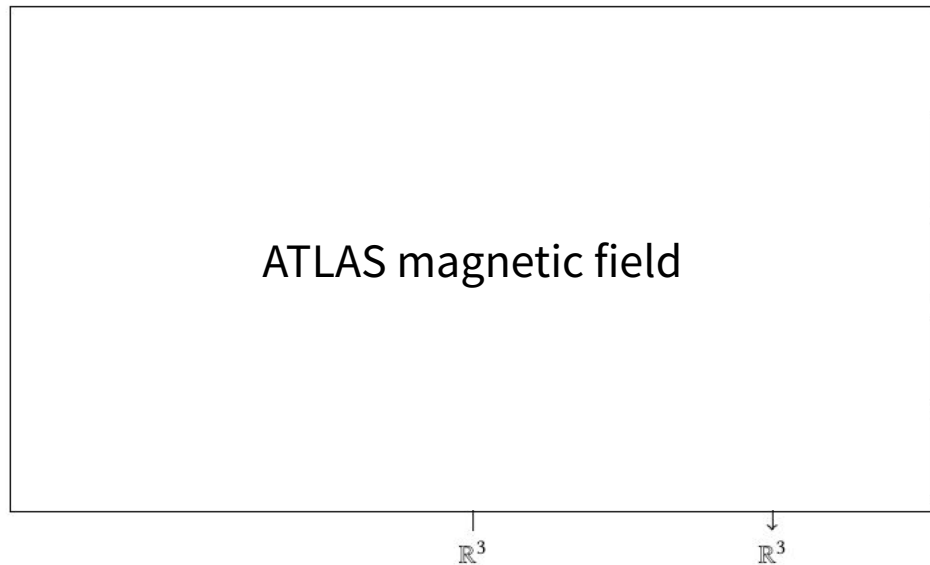
- Interpolation methods (F & EF)
 - NN, linear, cubic, etc.
- Boundary checking (F & EF)
 - Wrap, mirror, zero, etc.
- Array layout (EF)
 - Row-major, column-major, *Morton*, etc.
- Coordinate transformation (F & EF)
 - Affine, polar, etc.
- Storage location (EF)
 - Main memory, CUDA memory, texture memory, etc.
- Potentially many more!
- Across many **devices** and with many access patterns



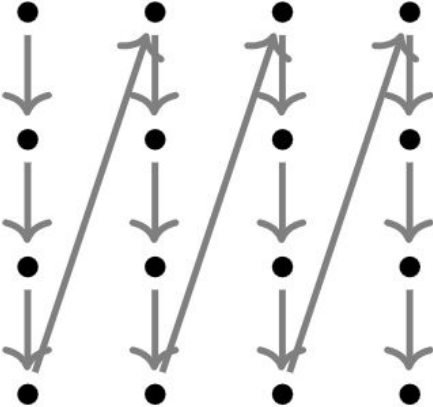
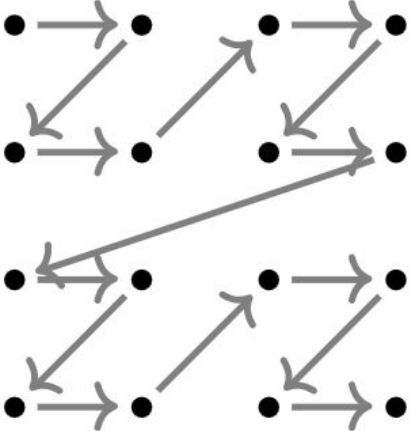
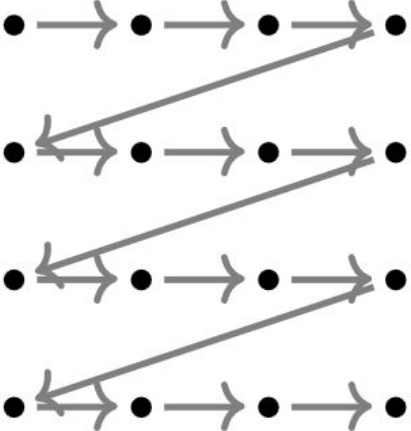
Source: Wikipedia

Research – Vector Fields

- We developed a **category-theoretical** method for **decomposing** this design space (like on the previous slide)
- Can be re-composed at **compile time** with zero run-time overhead
- Allows our method to serve as both a **benchmarking suite** and **design space exploration tool...**
- ...as well as a **library** for implementing heterogeneous multi-dimensional arrays

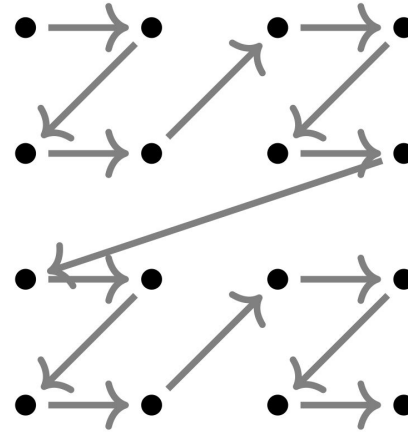


Research – Morton Layouts – Amuse Bouche



Research – Morton Layouts

- The Morton curve provides **balanced locality** compared to row-major and column-major layouts
- Calculated by **interleaving** the bits of the **binary expansions** of the input coordinates!
 - Can be done efficiently on modern commodity hardware
- But what if you were to interleave the bits in **arbitrary patterns?**



$$f(5, 3, 4) = f(101_2, 011_2, 100_2) = 101010110_2 = 342_{10}$$

Research – Morton Layouts

$$f(1011_2, 0101_2) = \frac{\begin{array}{r} 01000101_2 \\ \vee 00100010_2 \\ \hline 01100111_2 \end{array}}{=} = 103_{10}$$

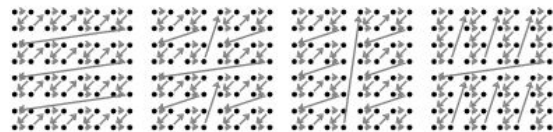
$$f(1011_2, 0101_2) = \frac{\begin{array}{r} 01000110_2 \\ \vee 00010001_2 \\ \hline 01010111_2 \end{array}}{=} = 87_{10}$$

Research – Morton Layouts

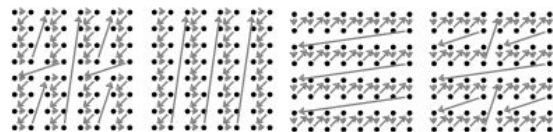
- Turns out this gives you *very* large families of array layouts, all of which have **different cache properties!**
- We propose that **evolutionary algorithms** can be used to efficiently explore this design space
- We show that we can significantly **improve performance** – up to 10 times in extreme cases
- Automated, **problem-agnostic** optimisation method!



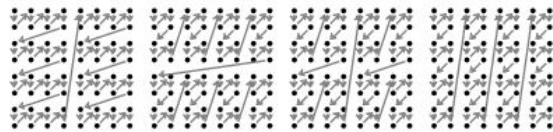
(a) [0,0,0,1,1,1] (b) [0,0,1,0,1,1] (c) [0,0,1,1,0,1] (d) [0,0,1,1,1,0]



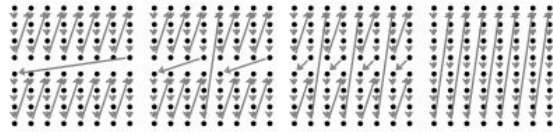
(e) [0,1,0,0,1,1] (f) [0,1,0,1,0,1] (g) [0,1,0,1,1,0] (h) [0,1,1,0,0,1]



(i) [0,1,1,0,1,0] (j) [0,1,1,1,0,0] (k) [1,0,0,0,1,1] (l) [1,0,0,1,0,1]



(m) [1,0,0,1,1,0] (n) [1,0,1,0,0,1] (o) [1,0,1,0,1,0] (p) [1,0,1,1,0,0]



(q) [1,1,0,0,0,1] (r) [1,1,0,0,1,0] (s) [1,1,0,1,0,0] (t) [1,1,1,0,0,0]

Research – Morton Layouts

- This was accepted to **ICPE'24**
- To be presented in **London** next week!
- <https://doi.org/10.1145/3629526.3645034> (not yet active)



Using Evolutionary Algorithms to Find Cache-Friendly Generalized Morton Layouts for Arrays

Stephen Nicholas Swatman^{*}
University of Amsterdam
Amsterdam, The Netherlands
s.n.swatman@uva.nl

Ana-Lucia Varbanescu
University of Twente
Enschede, The Netherlands
a.l.varbanescu@utwente.nl

Andy D. Pimentel
University of Amsterdam
Amsterdam, The Netherlands
a.d.pimentel@uva.nl

Andreas Salzburger
CERN
Geneva, Switzerland
andreas.salzburger@cern.ch

Attila Krasznahorkay
CERN
Geneva, Switzerland
attila.krasznahorkay@cern.ch

ABSTRACT

The layout of multi-dimensional data can have a significant impact on the efficacy of hardware caches and, by extension, the performance of applications. Common multi-dimensional layouts include the canonical row-major and column-major layouts as well as the Morton curve layout. In this paper, we describe how the Morton layout can be generalized to a very large family of multi-dimensional data layouts with widely varying performance characteristics. We posit that this design space can be efficiently explored using a combinatorial evolutionary methodology based on genetic algorithms. To this end, we propose a chromosomal representation for such layouts as well as a methodology for estimating the fitness of array layouts using cache simulation. We show that our fitness function correlates to kernel running time in real hardware, and that our evolutionary strategy allows us to find candidates with favorable simulated cache properties in four out of the eight real-world applications under consideration in a small number of generations. Finally, we demonstrate that the array layouts found using our evolutionary method perform well not only in simulated environments but that they can effect significant performance gains—up to a factor ten in extreme cases—in real hardware.

CCS CONCEPTS

• Software and its engineering → Software performance.
• Mathematics of computing → Combinatorial optimization.
• Information systems → Data layout.

KEYWORDS

Morton curve, Z-order curve, space-filling curves, array layout, multi-dimensional data, evolutionary algorithms, caching, locality

1 INTRODUCTION

Structured multi-dimensional data are ubiquitous in high-performance computing [5]: three-dimensional fluid simulations, dense linear algebra operations, and stencil kernels are just a few examples of applications which rely fundamentally on multi-dimensional arrays. In spite of the importance of such applications, however, most modern computer systems have one-dimensional memories: from the perspective of the programmer, memory is nothing more than a very large one-dimensional array of bytes. This discrepancy

^{*}Also with CERN.

between application requirements and hardware design requires programmers to carefully consider array layout in injective functions which translate multi-dimensional indices into one-dimensional memory addresses.

Although array layouts do not impact the functional properties of programs, choosing a suitable layout can significantly impact application performance in modern processors with complex cache hierarchies [4]. Exploiting these caches is of critical importance to achieving high performance in all but purely compute-bound applications, but doing so requires locality of access—both temporal and spatial—in memory. Kernels often exhibit locality in multiple dimensions, and a well-chosen array layout maximizes the degree to which this application-level locality is translated to the address-level locality that caches are designed to exploit; as a result, that layout increases the efficacy of hardware caching and—by extension—the performance of an application.

Data in two-dimensions is commonly laid out in row-major order (shown in Figure 2a for an 8×8 array) or column-major order (Figure 2b) which provide good locality of access in a single dimension, but poor locality in all others. Thankfully, the design space for data orderings—in two dimensions or more—extends far beyond these canonical layouts: the Morton layout (Figure 2c), for example, is a layout based on a space-filling curve which provides balanced locality between multiple dimensions [46, 42]. Our work explores a family of data layouts which generalize the Morton order, and allow us to carefully tune the cache behavior in multiple dimensions to match a given application.

The design space of the aforementioned family of data layouts is dauntingly large; indeed, the number of possible layout for arrays at scales applicable to real-world problems is so large that it renders exhaustive search infeasible. In order to find suitable array layouts in tractable amounts of time, we propose to employ genetic algorithms—heuristics known to be able to efficiently find high-quality solutions in large search spaces [35]. To this end, we design a chromosomal representation of Morton-like array layouts, as well as a fitness function that uses cache simulation to estimate the performance of individual array layouts. Finally, we evaluate our evolutionary strategy and the array layouts it discovers.

In short, our paper makes the following contributions:

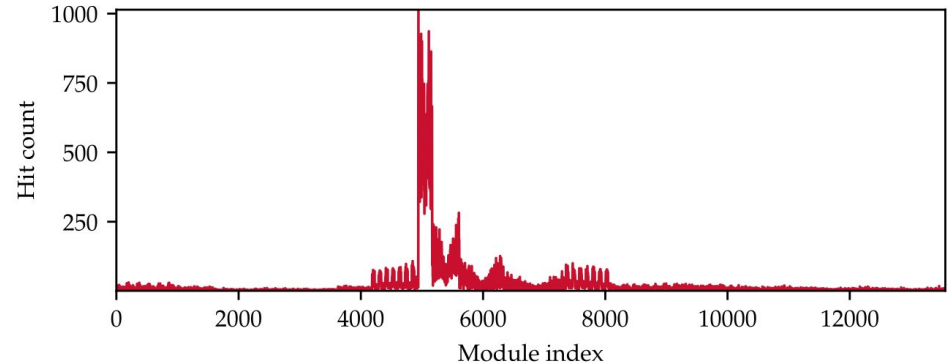
- We characterize the design space given by a generalization of the Morton array layout, and we show that that the size

arXiv:2309.07002v2 [cs.LG] 5 Feb 2024

How can the effects of thread imbalance in SIMT workloads be modelled and how can they be mitigated?

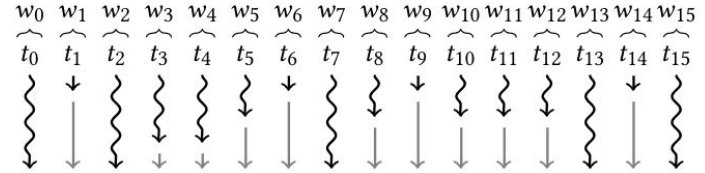
Research – Thread Imbalance – Amuse Bouche

- We know we cannot assign one event to one thread: **lockstep execution**
- Turns out we also cannot assign one module to one thread: **too imbalanced**
- We will need **fundamentally** different algorithms for **clustering**
- But can we **predict** what will work?
 - And can we use the graph on the right?

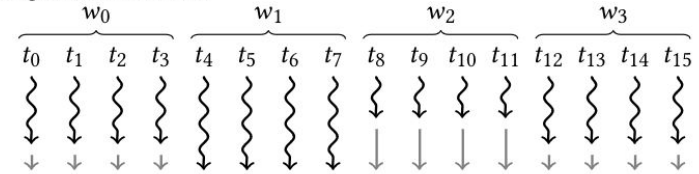


Research – Thread Imbalance

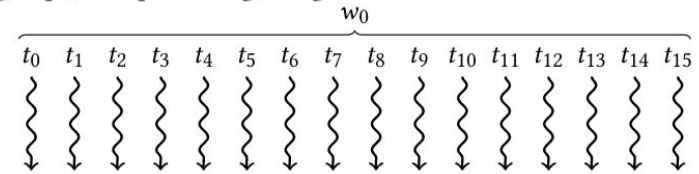
- A lot of our kernels are **imbalanced**: different threads execute **different amounts of work**
 - Different cluster sizes
 - Different branching factors
- This is not a problem on CPUs, but **strongly impacts GPU performance**
- Can be mitigated using **thread coarsening** or **thread refinement**
- But what is the **performance impact** of these techniques?



(a) Per-thread granularity: each thread in a thread group processes a single unit of work.



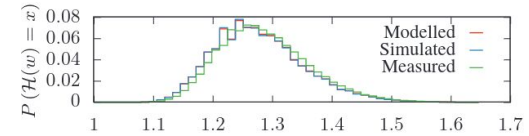
(b) Intermediate granularity: threads in a thread group form sub-groups, each processing a single unit of work.



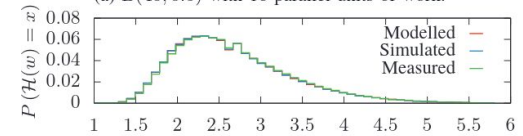
(c) Per-group granularity: all threads in a thread group process a single unit of work.

Research – Thread Imbalance

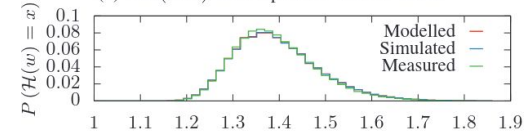
- We propose a **statistical model** that can infer the overhead of SIMT execution
- Metric for “*how suitable for GPU execution is this workload*”
- Uses **only** distribution of **thread load**: no hardware details, etc.
- Allows **early evaluation** of feasibility of different **parallelism strategies!**



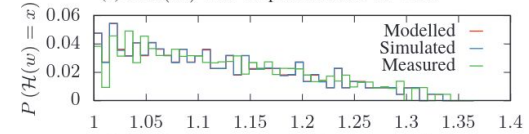
(a) $B(40, 0.5)$ with 16 parallel units of work.



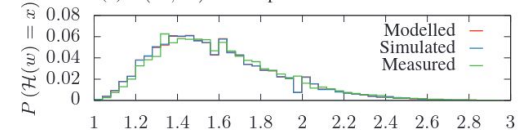
(b) $Geo(0.05)$ with 8 parallel units of work.



(c) $Pois(30)$ with 32 parallel units of work.



(d) $U(20, 40)$ with 2 parallel units of work.



(e) $NB(5, 0.3)$ with 4 parallel units of work.

Research – Thread Imbalance

- This work was published in MASCOTS'22
- Which was nice and also in Nice
- <https://doi.org/10.1109/MASCOTS56607.2022.00026>



Modelling Performance Loss due to Thread Imbalance in Stochastic Variable-Length SIMT Workloads

Stephen Nicholas Swatman¹, Ana-Lucia Varbanescu¹, Attila Krisznanorkay¹, Andy Pimentel²

¹ University of Amsterdam, Amsterdam, The Netherlands

² European Organisation for Nuclear Research, Geneva, Switzerland

³ University of Twente, Enschede, The Netherlands

Abstract—When designing algorithms for single-instruction multiple-thread (SIMT) devices such as general purpose graphics processing units (GPGPUs), thread imbalance is an important performance consideration. Thread imbalance can emerge in iterative applications where workloads are of variable length, because threads processing larger amounts of work will cause threads with less work to idle. This form of thread imbalance influences the design space of algorithms—particularly in terms of processing granularity—but we lack models to quantify its impact on application performance. In this paper, we present a statistical model for quantifying the performance loss due to thread imbalance for iterative SIMT applications with stochastic, variable-length workloads. Our model is designed to operate with minimal knowledge of the implementation details of the algorithm, relying solely on an understanding of the probability distribution of the lengths of the workloads. We validate our model against a synthetic benchmark based on a Monte Carlo simulation of matrix exponentiation, and show that our model achieves nearly perfect accuracy. Compared to empirical data extracted from real hardware, our model maintains a high degree of accuracy, predicting mean performance loss within a margin of 2%.

Index Terms—SIMT, imbalance, performance modelling

I. INTRODUCTION

As the landscape of high-performance computing has evolved over recent years, single-instruction multiple-thread (SIMT) processors—usually in the form of general purpose graphics processing units (GPGPUs)—have become popular for high-performance computation in many domains [1]. By sacrificing the independence of individual processing cores, SIMT processors are able to use a significantly more processing cores, and thus provide much more raw processing power, compared to their traditional multiple-instruction multiple-data (MIMD) counterparts [2].

However, not every conceivable computational workload can be efficiently handed off to an SIMT device. The increased raw processing power of these devices comes at the cost of reduced flexibility, and algorithms must be carefully designed to run efficiently on SIMT devices, lest their computational process goes to waste. One important consideration when programming SIMT devices is the concept of *thread divergence*. In an SIMT device, a group of threads can—by definition—perform only a single, common instruction at a time; colloquially, these threads run in *lockstep*. Thus, cases

where the execution paths of threads diverge will cause some of the threads to be idle. If care is not taken to minimize thread divergence in algorithms designed to run on SIMT devices, it can severely degrade performance [3].

Thread divergence emerges not only in situations with conditional branches in the common *if-else* sense, but it can also arise in iterative processes in the form of *thread imbalance*. When the number of iterations of a loop varies between threads, the result is divergence: threads will be idle until the thread with the largest amount of work has performed the necessary number of iterations. Throughout this paper, we refer to workloads where the number of iterations is not fixed and may differ between threads as *variable-length workloads*. When the number of iterations is described by some probabilistic process, we refer to them as *stochastic workloads*. While it is well understood that thread imbalance in variable-length workloads is detrimental to the performance of SIMT devices [3], [4], we are unaware of any quantitative models that predict exactly how much performance is lost.

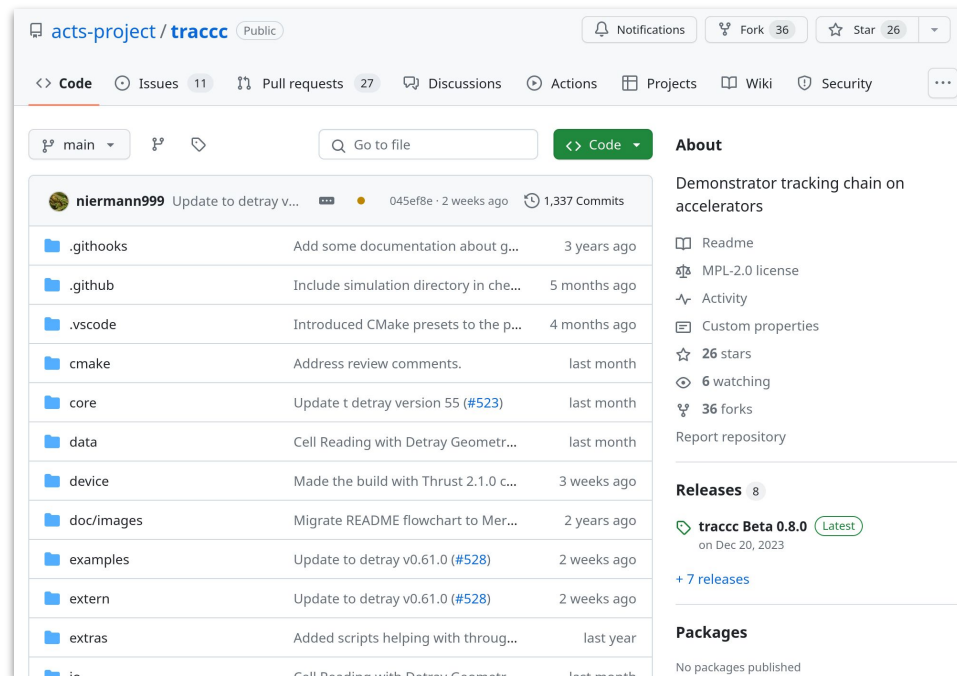
The question how we can model the impact of thread imbalance in stochastic variable-length workloads is the core focus of this paper. With this work, we are the first to design and implement an accurate statistical model for the expected performance loss of a given application, given only that it is an iterative process, that it is executed on an SIMT device, and that the number of iterations required to complete the process follows a known (albeit arbitrarily complex) distribution. We validate our model using empirical measurements gathered using a dedicated benchmark running on an NVIDIA GPU. The results of this validation show that our model agrees with simulated data with a relative error of less than 0.1%, and that it agrees with measurements taken on a real device within 2%.

Our accurate model can help motivate more precisely the design process of (future) SIMT applications—in particular in terms of processing granularity—in domains where stochastic iterative processes are common, such as machine learning [5], cryptography [6], graph processing [7] and scientific computing [8]. The importance of thread imbalance and granularity is further supported by our own results, which show (in Table I) that thread imbalance in SIMT devices can lead to execution that is nearly four times slower if thread granularity is not chosen carefully.

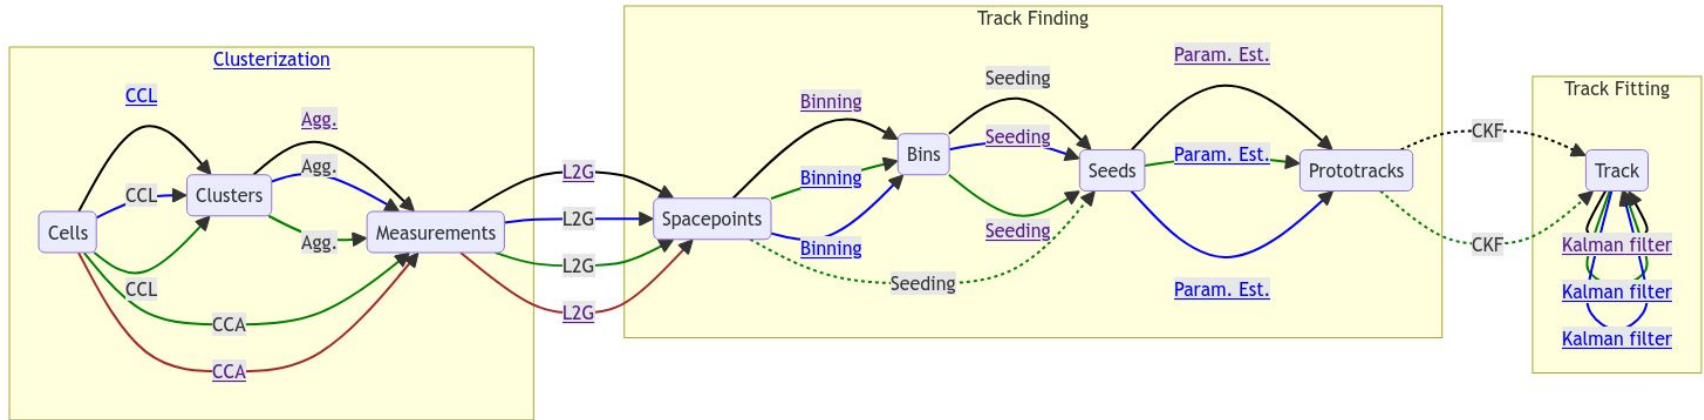
**How do the extra-functional
properties of novel track
reconstruction algorithms
compare to
state-of-the-art solutions?**

tracc - GPU Tracking Demonstrator

- Now: back to the application under study
- We have developed a **track reconstruction chain** for massively parallel devices
- Integrates **novel algorithms**
- Tracking on **TrackML** and **ODD**-like detectors

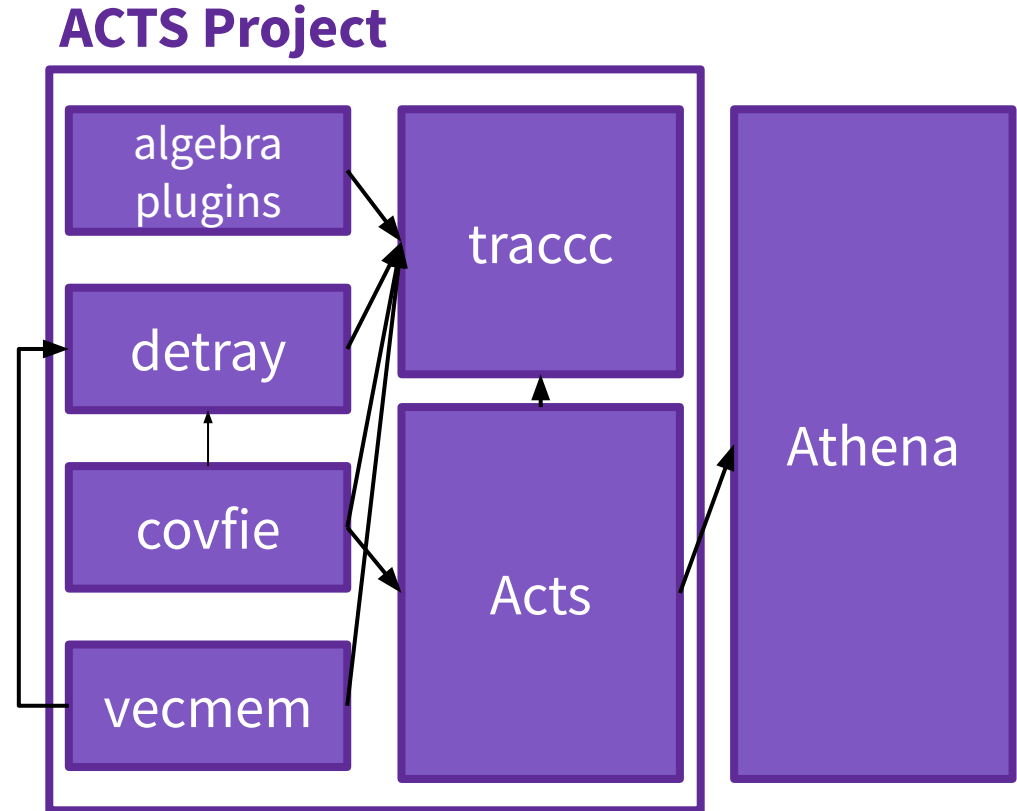


traccc – GPU Tracking Demonstrator



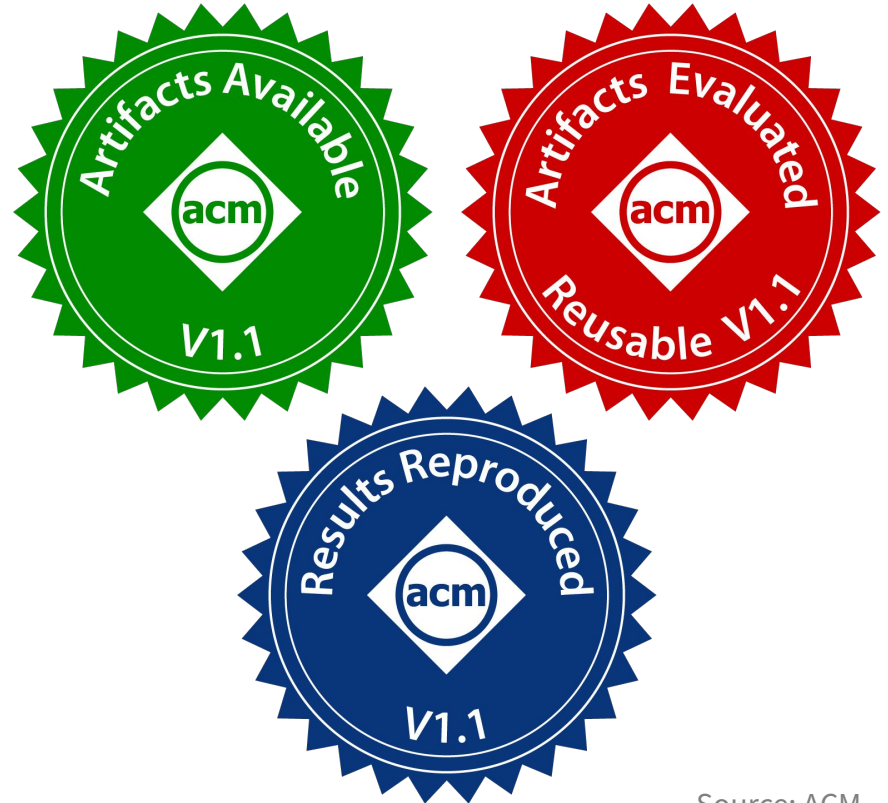
ACTS Project – Subprojects

- R&D consists of many **subprojects** for **HEP** and **HPC** in general
- **traccc**: tracking demonstrator
- **algebra-plugins**: linear algebra
- **detray**: detector description
- **vecmem**: memory management
- **covfie**: vector field storage
 - Plasma physics at HZDR



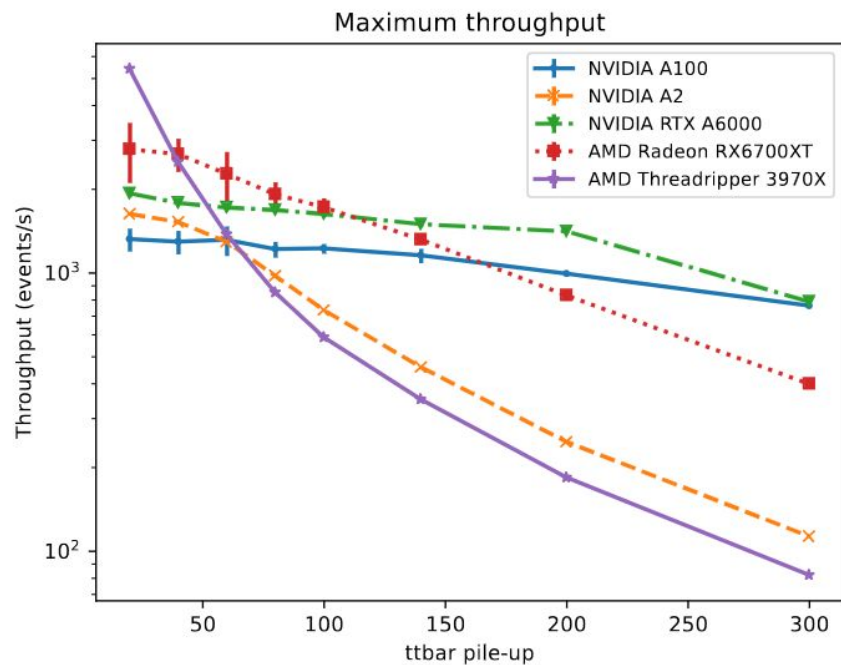
Key Point – Reproducibility

- Aim: to write software and develop methods that can be **used and improved on**
 - LHC lifetime: 17 more years
- Try to avoid “PhDware”: software that becomes unusable after the end of the PhD
- **Artifact** evaluation tracks take extra effort but reward handsomely
 - Nice **stickers**
 - In some cases taken into account for reviews, rebuttals, etc.
- Also served on AE committees for SC, CGO, ICPE, and ICPP



Results

- Preliminary results on **TrackML** data show that our GPU-based solutions work well
- **Outperform** similarly priced CPUs at **higher pile-up values**
- **Factor ~10** gain in throughput



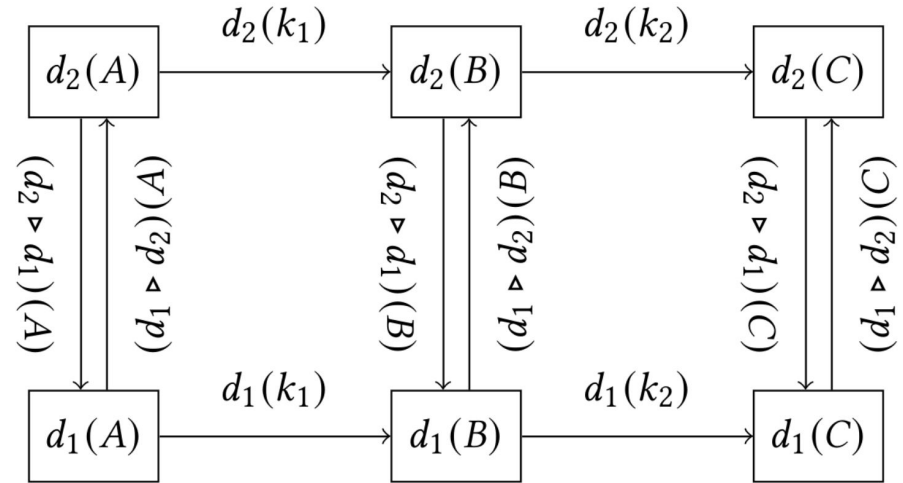
Source: Guilherme Metelo Rita de Almeida

Outstanding Challenges

- So far tested only on **simple geometries**: how do we integrate e.g. **ITk**
- **Combinatorial Kálmán Filter**: important step with high combinatorics
 - How do we **distribute branches** over threads?
- Integration into **Athena** is ongoing work
- How to **schedule** and **place** algorithms?
 - GPUs have **separate memories**: transfers are not free

Research – Throughput Models

- Our scheduling is an open problem, but can we somehow estimate the **throughput** of our **task graph on heterogeneous systems...**
- ...using only the throughputs of the **individual kernels?**
- We propose that we can create an **optimistic upper bound** for this based on work in the data flow community
- Using **linear programming** we can solve a resource-constrained **maximum flow problem!**



$$\begin{aligned}
 &\text{maximise} && \sum_{e \in E^-(t)} x_e f(e) \\
 &\text{subject to} && \forall v \in T' \setminus \{s, t\} : \sum_{e \in E^+(v)} x_e f(e) = \sum_{e \in E^-(v)} x_e f(e) \\
 &&& \forall e \in E : 0 \leq x_e \leq 1 \\
 &&& \forall r \in D \cup I : 0 \leq \sum_{e \in Q(r)} x_e \leq 1
 \end{aligned}$$

Conclusions

- ATLAS needs **aggressive R&D** to tackle **high- μ compute challenges**
- **Massively parallel** track reconstruction under development in **ACTS**
- Despite **irregular workloads**, we can exploit **GPUs** well in **TrackML** and **ODD**
- **Performance** is very **competitive** at $\mu \geq 100$
- Complex **geometries**, **scheduling**, and **placement** remain open questions



The ACTS project
<https://github.com/acts-project/>



Bi-weekly R&D meeting
<https://indico.cern.ch/category/16958/>