



# Systematically Exploring High-Performance Representations of Vector Fields Through Compile-Time Composition

---

Stephen Nicholas Swatman<sup>1,2</sup>, Ana-Lucia Varbanescu<sup>3</sup>, Andy Pimentel<sup>1</sup>, Andreas Salzburger<sup>2</sup>, Attila Krasznahorkay<sup>2</sup>

Monday, April 17, 2023

<sup>1</sup>University of Amsterdam <sup>2</sup>CERN <sup>3</sup>University of Twente



# Introduction: Vector Fields

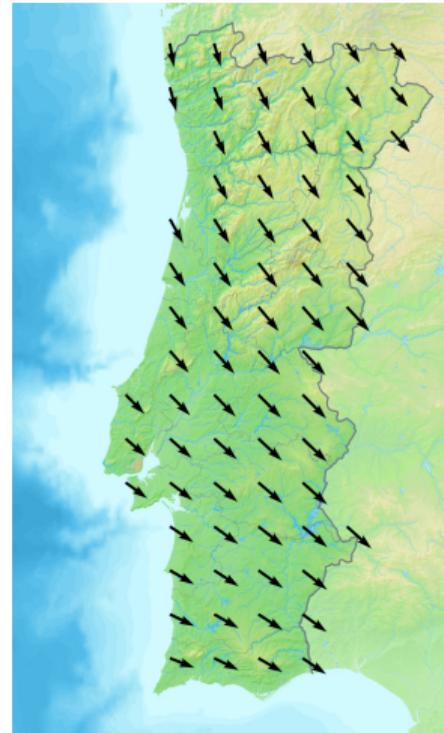
- Vector fields map vectors onto elements of a set
- Common case: mapping vectors onto structured grids



Credit: Wikimedia Commons user Waldyrious, CC BY-SA 4.0

# Introduction: Vector Fields

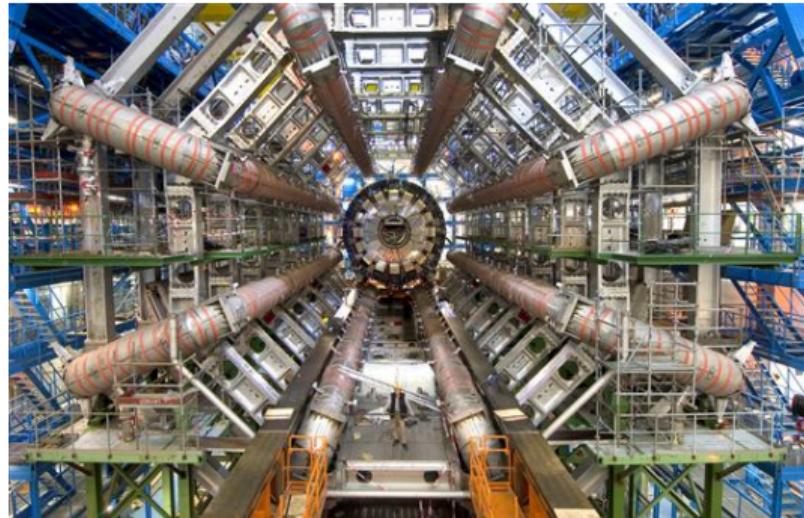
- Vector fields map vectors onto elements of a set
- Common case: mapping vectors onto structured grids
- Useful for modelling flows like wind...



Credit: Wikimedia Commons user Waldyrious, CC BY-SA 4.0

# Introduction: Vector Fields

- Vector fields map vectors onto elements of a set
- Common case: mapping vectors onto structured grids
- Useful for modelling flows like wind...



Credit: ATLAS Collaboration

# Introduction: Vector Fields

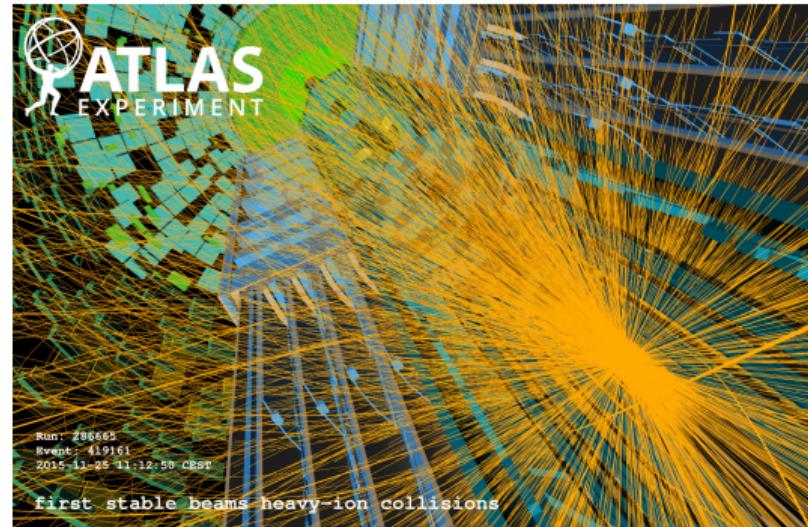
- Vector fields map vectors onto elements of a set
- Common case: mapping vectors onto structured grids
- Useful for modelling flows like wind...
- ...or magnetic fields in high-energy physics



Credit: ATLAS Collaboration

# Introduction: Vector Fields

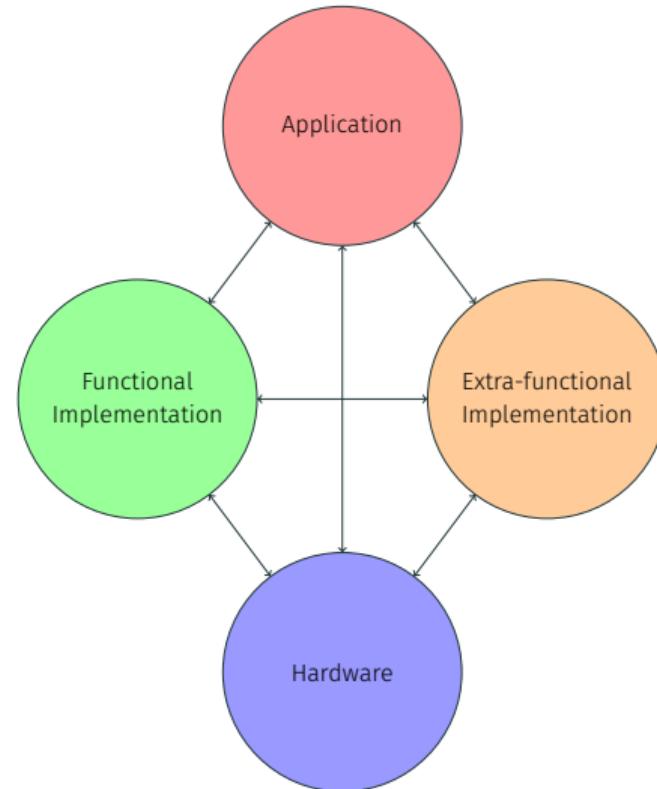
- Vector fields map vectors onto elements of a set
- Common case: mapping vectors onto structured grids
- Useful for modelling flows like wind...
- ...or magnetic fields in high-energy physics
- Performance-relevant data structures, but which representation is best?



Credit: ATLAS Collaboration

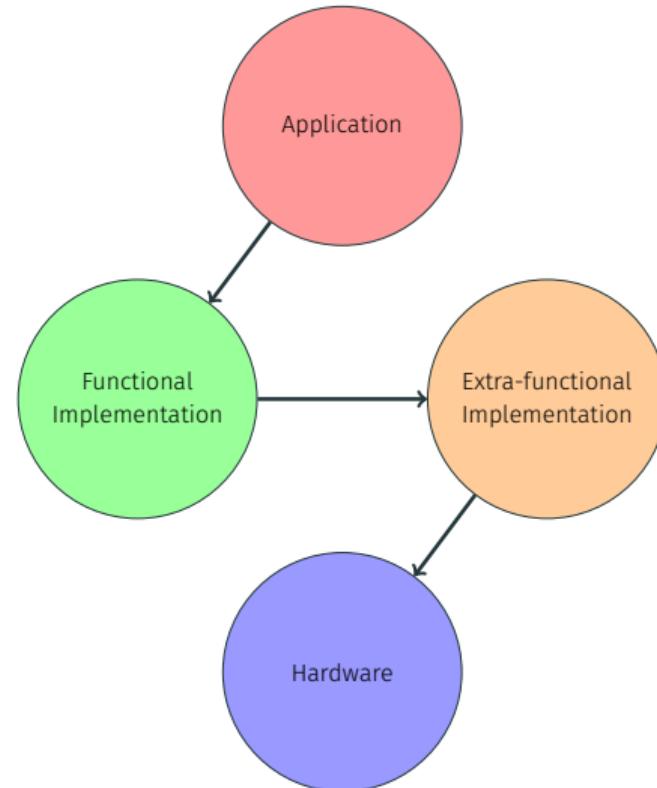
# Design Space Definition

- Large design space: performance depends heavily on the *interaction between*:
  - Application
  - Functional implementation
  - Extra-functional implementation
  - Hardware



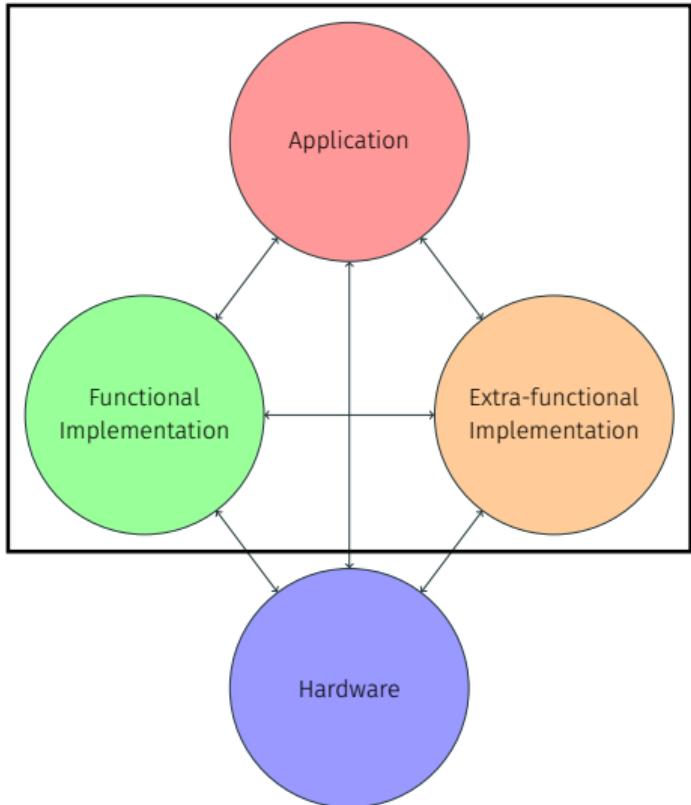
# Design Space Definition

- Example: application migrates from two-dimensional approximation to three-dimensional field:
  - Functional implementation of the field changes...
  - ...the storage size changes (extra-functional behaviour)...
  - ...and the hardware cache behaviour changes



# Design Space Exploration

- Our goal: allow systematic exploration of representations of vector fields
  - *Define* the design space
  - *Compose* representations from simple components
  - Facilitate exploration through *benchmarking*



# Application: Access Patterns

- In our suite, *applications* are modelled by their *access patterns*, capturing:
  - Locality of access
  - Computational side-effects
  - Functional requirements
- Variants:
  - CPU and GPGPU implementations
  - For patterns based on propagation through fields: order of execution
  - Depth-first or breadth-first propagation of agents

Name	Vr.	Field type
SCAN	2	$\prod_{d,d':\mathbb{N}} S^{d+d'} \rightarrow T$
RANDOM	2	$\prod_{d:\mathbb{N}} S^d \rightarrow T$
EULER	3	$\prod_{d:\mathbb{N}} \mathbb{R}^d \rightarrow \mathbb{R}^d$
RK4	3	$\prod_{d:\mathbb{N}} \mathbb{R}^d \rightarrow \mathbb{R}^d$
LORENTZ	6	$\mathbb{R}^3 \rightarrow \mathbb{R}^3$

## Implementation: Storage Backends

---

- Implementation of field data structure remains dauntingly complex
- Many choices of implementation
  - Functional: interpolation, boundary checking, geometric transformations, etc.
  - Extra-functional: data layouts, hardware acceleration, etc.
- This behaviour can (must!) be further decomposed into *primitive* behaviour and *transformers*

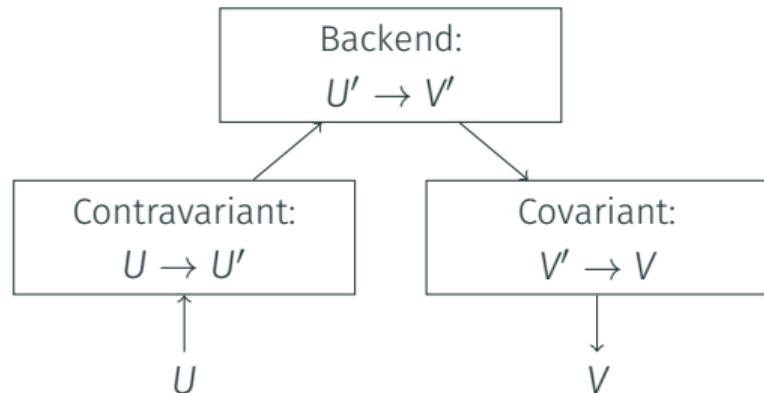
# Implementation: Primitive Storage

- Behaviour that cannot be meaningfully deconstructed
  - Basic array access (ARRAY, CUDAARRAY)
  - More complex storage like CUDA textures (CUDATEX) or opaque arrays (CUDAPITCH)
  - Constant-valued and analytical fields (CONSTANT, ANALYTICAL)

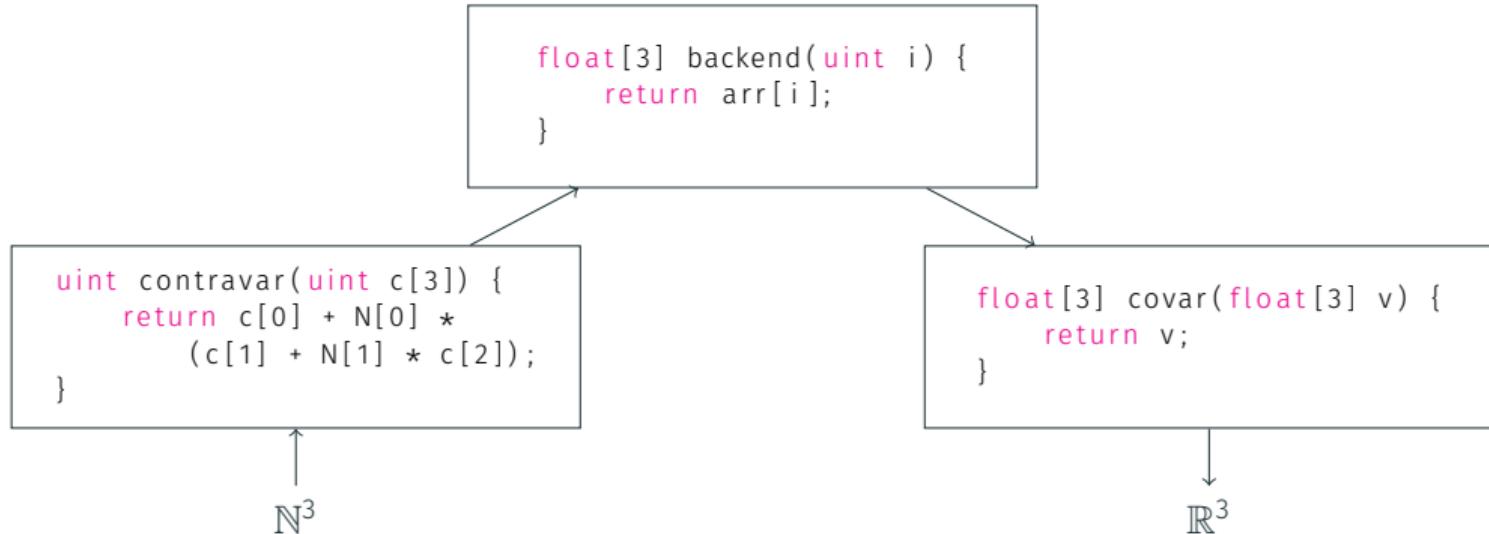
Name	Pltf.	Field type
ARRAY	CPU	$\mathbb{N} \rightarrow T$
CUDAARRAY	CUDA	$\mathbb{N} \rightarrow T$
CUDAPITCH	CUDA	$\prod_{d:\llbracket 1,3 \rrbracket} \mathbb{N}^d \rightarrow T$
CUDATEX	CUDA	$\prod_{d:\llbracket 1,3 \rrbracket} \prod_{d':\llbracket 1,4 \rrbracket} \mathbb{R}^d \rightarrow \mathbb{R}^{d'}$
ANALYTIC	CPU	$S \rightarrow T$
CONSTANT	Any	$S \rightarrow T$

# Implementation: Transformers

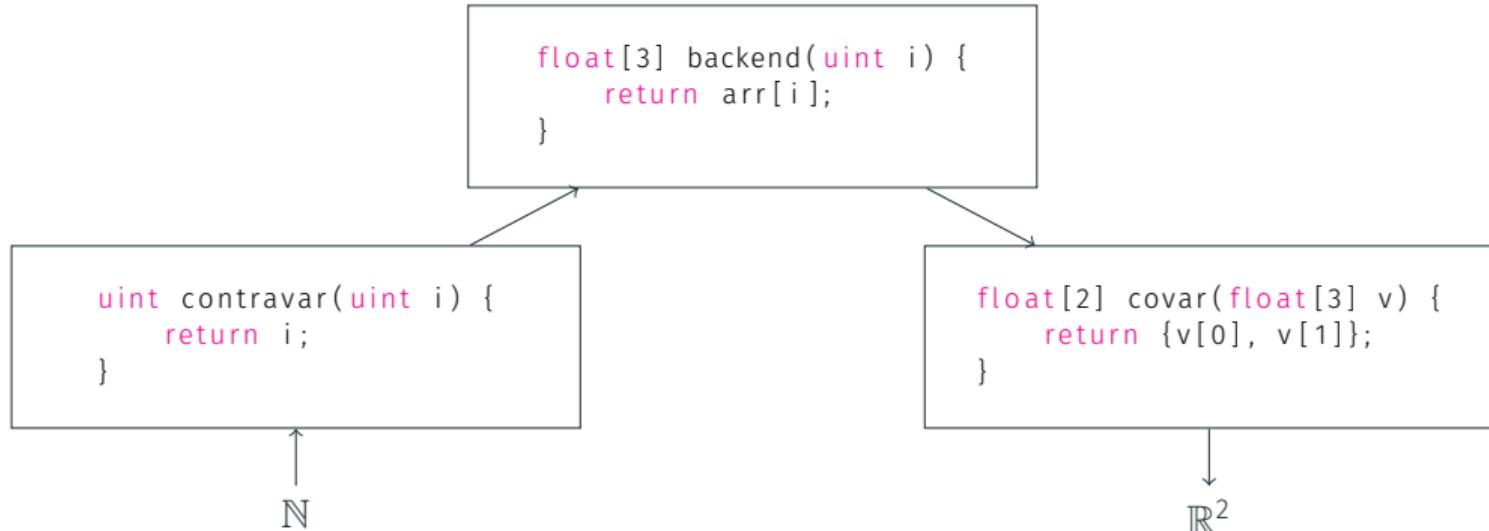
- Additional behaviour encoded as *transformers*, built of two parts:
  - Contravariant component allows to *modify input*
  - Covariant component allows to *modify output*



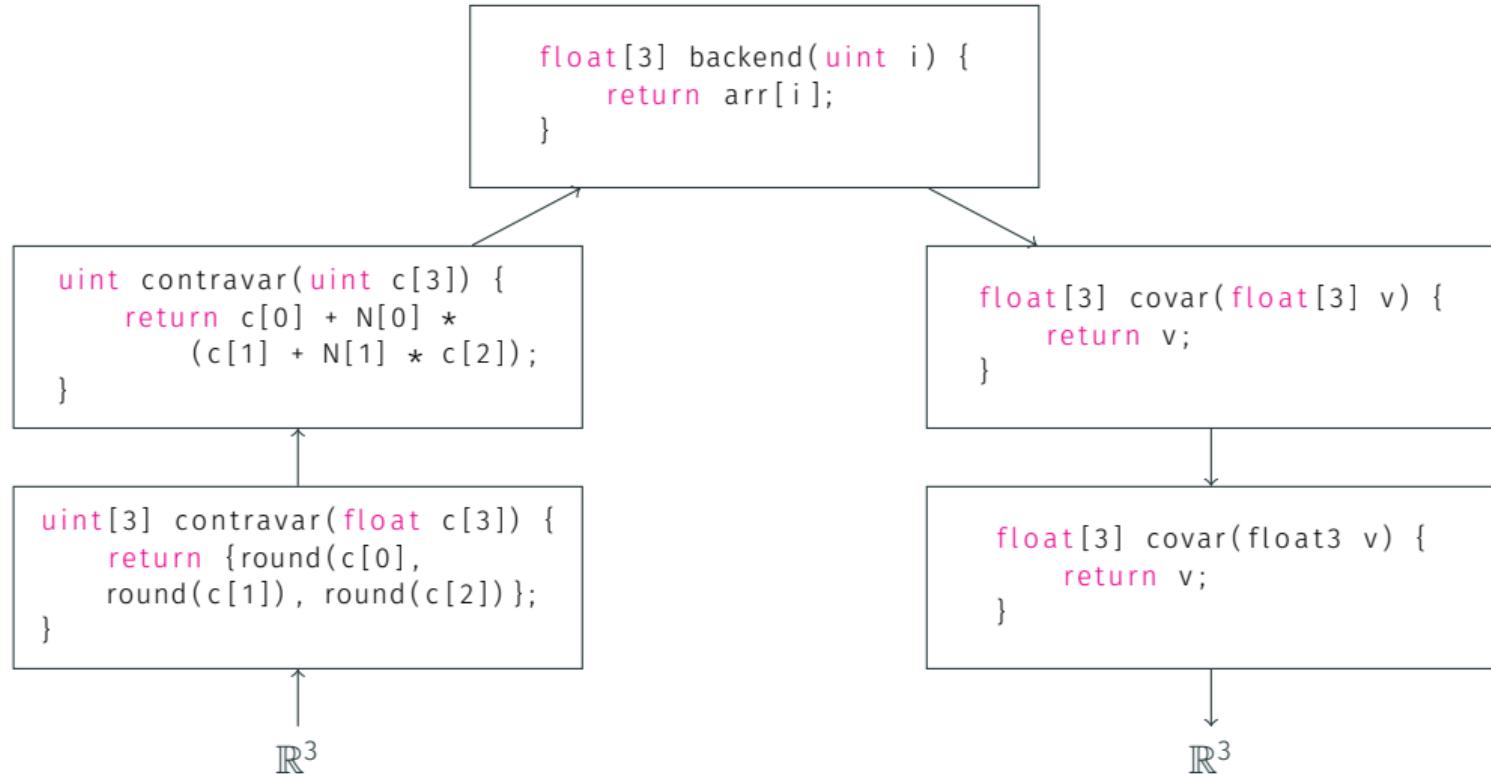
# Implementation: Transformers



# Implementation: Transformers



# Implementation: Transformers



# Implementation: Transformers

- We provide a variety of transformers for common implementation details
- These can be composed as long as the types line up
- Many transformers have additional variants (e.g. WRAP)

Name	Type
LEXIC.	$\prod_{n:\mathbb{N}} (\mathbb{N}^n \rightarrow \mathbb{N}) \times (T \rightarrow T)$
MORTON	$\prod_{n:\mathbb{N}} (\mathbb{N}^n \rightarrow \mathbb{N}) \times (T \rightarrow T)$
HILBERT	$(\mathbb{N}^2 \rightarrow \mathbb{N}) \times (T \rightarrow T)$
SHUFFLE	$\prod_{n:\mathbb{N}} \prod_{p:\mathfrak{S}_n} (S^n \rightarrow S^n) \times (T \rightarrow T)$
OOB	$(S \rightarrow S + \mathbb{1}) \times (T + \mathbb{1} \rightarrow T)$
WRAP	$(S \rightarrow S) \times (T \rightarrow T)$
NEAREST	$\prod_{n:\mathbb{N}} (\mathbb{R}^n \rightarrow \mathbb{N}^n) \times (T \rightarrow T)$
LINEAR	$\prod_{n:\mathbb{N}} (\mathbb{R}^n \rightarrow \mathbb{N}^{2^n n} \times \mathbb{R}^n) \times (\mathbb{R}^{2^n n} \times \mathbb{R}^n \rightarrow \mathbb{R}^n)$
AFFINE	$\prod_{n:\mathbb{N}} (\mathbb{R}^n \rightarrow \mathbb{R}^n) \times (T \rightarrow T)$

## Benefit: Extensibility

---

- New access pattern has “free” benchmarks with hundreds of backends
- New primitive storage composes with many transformers
- New transformer is compatible with many access patterns

# Composition in High-Performance Kernels

- Composition is usually achieved through run-time polymorphism
- Incurs overhead from dynamic function dispatch
- Manageable for most applications...
- ...but deal-breaking in HPC kernels

```
1 backend f1;
2 layer3 f2(f1);
3 layer2 f3(f2);
4 layer1 f4(f3);
5
6 vector3 v = f4.at(5.f, 2.f, -2.f);
```

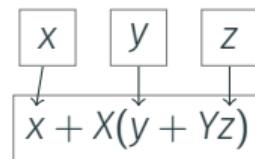
# Compile-Time Composition

- Instead of composing at *run-time*, we compose behaviour at *compile-time*
- Use of C++ templating fixes behaviour at compile-time
  - Eliminates overhead due to dispatching
  - Increases optimisation space of the compiler
- We achieve native performance *and* preserve flexibility

```
1 using storage = layer1<
2   layer2<
3     layer3<
4       backend <... >
5     >,
6   ...
7 >,
8 ...
9 >;
```

# Compile-Time Composition

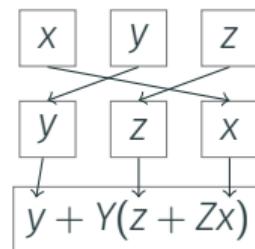
```
1 using storage = lexicographic<ulong3,  
ulong3>;
```



```
1 ; g++ 12.2.0 x86_64_v1  
2 imul    0x8(%rsi),%rdx  
3 add     %rcx,%rdx  
4 imul    0x10(%rsi),%rdx  
5 add     %rdx,%r8  
6 mov     0x18(%rsi),%rdx  
7 lea     (%r8,%r8,2),%rax  
8 lea     (%rdx,%rax,8),%rax  
9 movdqu (%rax),%xmm0  
10 mov    0x10(%rax),%rax  
11 movups %xmm0,(%rdi)  
12 mov    %rax,0x10(%rdi)  
13 mov    %rdi,%rax  
14 ret
```

# Compile-Time Composition

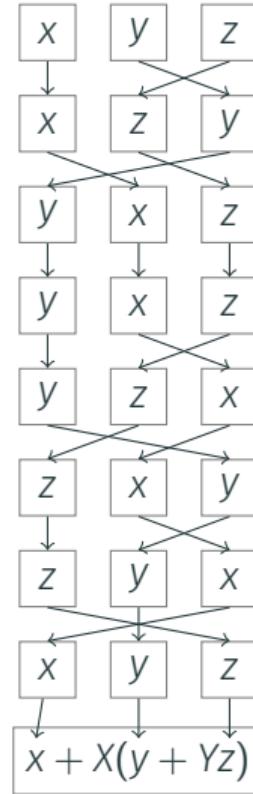
```
1 using storage = shuffle<lexicographic  
<ulong3, ulong3>, perm<1, 2, 0>>;
```



```
1 ; g++ 12.2.0 x86_64_v1  
2 imul    0x8(%rsi),%rcx  
3 add     %r8,%rcx  
4 imul    0x10(%rsi),%rcx  
5 lea     (%rcx,%rdx,1),%rax  
6 mov     0x18(%rsi),%rdx  
7 lea     (%rax,%rax,2),%rax  
8 lea     (%rdx,%rax,8),%rax  
9 movdqu (%rax),%xmm0  
10 mov    0x10(%rax),%rax  
11 movups %xmm0,(%rdi)  
12 mov    %rax,0x10(%rdi)  
13 mov    %rdi,%rax  
14 ret
```

# Compile-Time Composition

```
1 using storage = shuffle<shuffle<
  shuffle<shuffle<shuffle<shuffle<
    shuffle<lexicographic<ulong3,
    ulong3>, perm<2, 1, 0>>, perm<0,
    2, 1>>, perm<1, 2, 0>>, perm<0,
    2, 1>>, perm<0, 1, 2>>, perm<2,
    0, 1>>, perm<0, 2, 1>>;
```



```
1 ; g++ 12.2.0 x86_64_v1
2 imul 0x8(%rsi),%rdx
3 add  %rcx,%rdx
4 imul 0x10(%rsi),%rdx
5 add  %rdx,%r8
6 mov   0x18(%rsi),%rdx
7 lea   (%r8,%r8,2),%rax
8 lea   (%rdx,%rax,8),%rax
9 movdqu (%rax),%xmm0
10 mov   0x10(%rax),%rax
11 movups %xmm0,(%rdi)
12 mov   %rax,0x10(%rdi)
13 mov   %rdi,%rax
14 ret
```

## Empirical Evaluation (HEP Example): Setup

- We imagine a high-energy physicist
- Goal: find the best implementation of a vector field in her software
- Select one of the access patterns that models the application
  - Not exactly the same as the real application, but sufficiently close
- Make a selection of candidate implementations
  - Construct these from the “building blocks” discussed before
- This generates a set of benchmarks, which can be run on one or more systems
- In this case: six setups in the DAS-6<sup>1</sup> compute cluster!

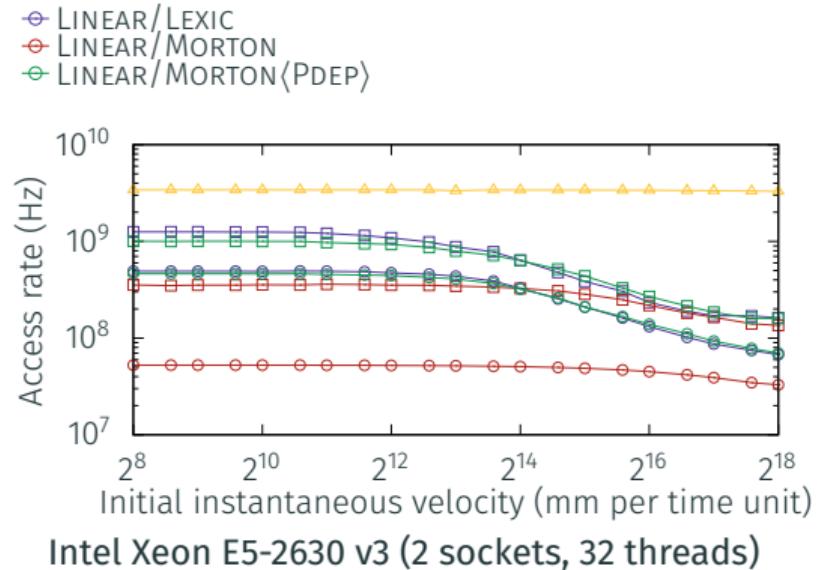
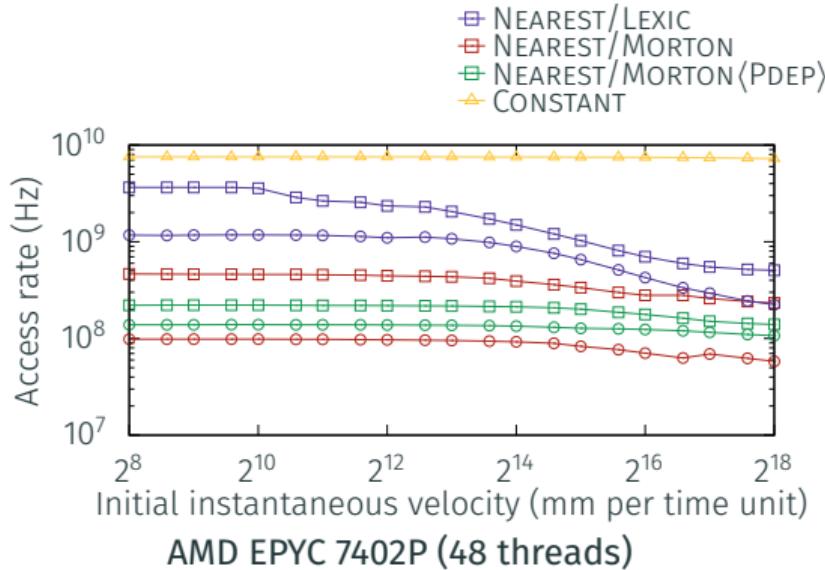
---

<sup>1</sup><https://www.cs.vu.nl/das6/>

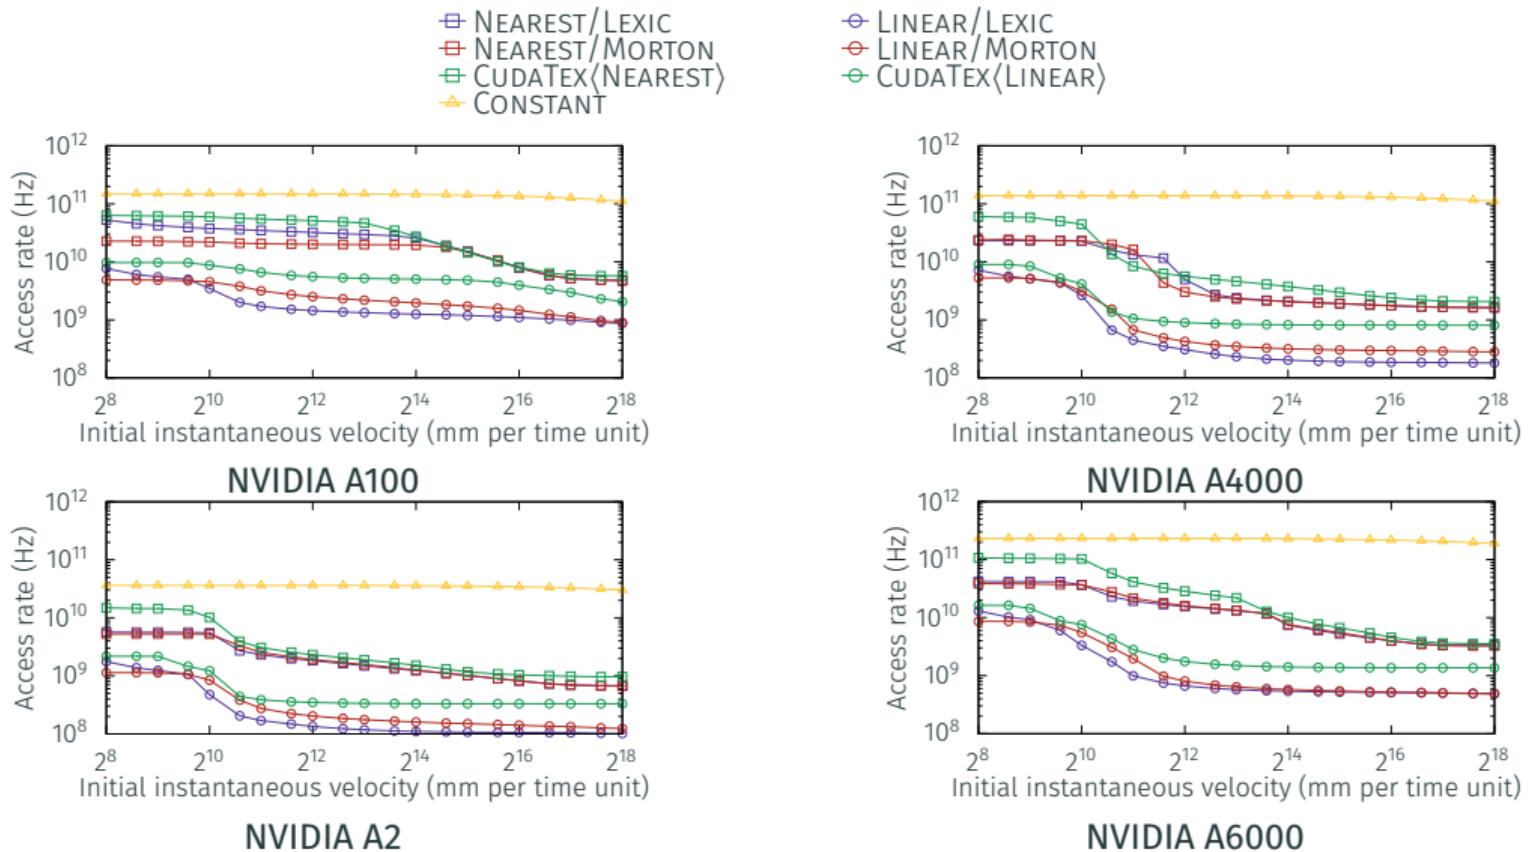
# Empirical Evaluation (HEP Example): Setup

```
1 benchmark::register_product_bm<
2     // Seven access patterns
3     boost::mp11::mp_list<
4         Lorentz<Euler>,
5         Lorentz<RungeKutta4>,
6         RungeKutta4Pattern,
7         EulerPattern,
8         Random,
9         Scan>,
10    // Seven composite implementations
11    boost::mp11::mp_list<
12        FieldConstant,
13        FieldTex<TexInterpolateLin>,
14        FieldTex<TexInterpolateNN>,
15        Field<InterpolateNN, LayoutStride>,
16        Field<InterpolateNN, LayoutMorton>,
17        Field<InterpolateLin, LayoutStride>,
18        Field<InterpolateLin, LayoutMorton>
19    >
20    // Generates a large number of benchmarks!
21 >();
```

# Empirical Evaluation (HEP Example): Results



# Empirical Evaluation (HEP Example): Results



## Usage in Real-World Applications

---

- Our benchmark suite provides a range of components and infrastructure
- This translates (with some effort on our part) to a *user-facing library*
- Makes the results of benchmarks immediately *applicable*
- Eliminates ambiguity between benchmark results and manual re-implementation
- Library available under MPL-2.0 license
- Currently in use in experimental HEP software!

# Artifact



- Artifact is permanently available at  
<https://zenodo.org/record/7540593>
- Provided with an *optional* Docker file
- Specific care taken to ensure code works without Docker
  - But using the Docker file will guarantee sane and reproducible environment

January 15, 2023

Software Open Access

Benchmarking Software for "Systematically Exploring High-Performance Representations of Vector Fields Through Compile-Time Composition". It allows for the reproduction of our results, as well as the generation of new ones. To do so, we provide a detailed description of the benchmarks and their execution environment. The total process should take no more than ten hours (most of which is non-interactive compute time) for a quick reproduction, and around two days for a full reproduction.

This artifact contains the benchmarking code for our paper "Systematically Exploring High-Performance Representations of Vector Fields Through Compile-Time Composition". It allows for the reproduction of our results, as well as the generation of new ones. To do so, we provide a detailed description of the benchmarks and their execution environment. The total process should take no more than ten hours (most of which is non-interactive compute time) for a quick reproduction, and around two days for a full reproduction.

None: This ZIP file does not contain a top-level directory. To unpack it into a single directory, use `unzip -d /path/to/unzip /path/to/archive_name.zip`.

DOI: <https://doi.org/10.5281/zenodo.7540593>

License (for files): [MIT License](#)

Publication date: January 15, 2023

Versions

Version	Date
Version 3	Jan 15, 2023
Version 2	Jan 15, 2023
Version 1	Aug 24, 2022

File (3.0 MB)

Name	Size
10.5281/zenodo.7571923/v3.zip	4.3 MB

Path to artifact: [https://zenodo.org/record/7540593/files/10.5281/zenodo.7571923/v3.zip](#)

Citations (0)

Show only:  Literature (0)  Dataset (0)  Software (0)  Unknown (0)  
 Citations to this version

No citations.

Start typing a citation style...

Export

BibTeX Cite DataCite Dublin Core DCAT JSON JSON-LD GeoJSON MARCXML OGMendley

288 views 14 downloads See more details...

OpenAIRE

SW-1 Benchmarking suite:

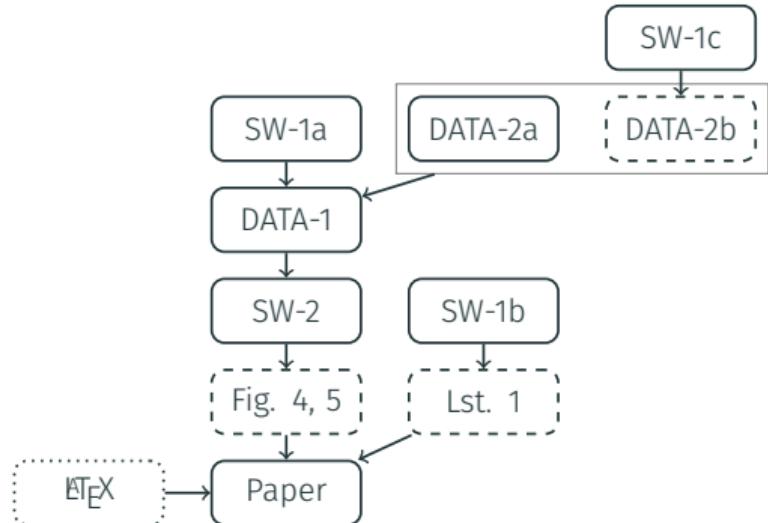
- a Primary benchmarking code
- b Code for assembly inspection
- c Dummy field generation code

SW-2 Plotting code.

DATA-1 Results from our experiments

DATA-2 Vector field data

- a Real-world data from HEP
- b Dummy field data



# Conclusion

---

- We present a *benchmark suite* for vector fields on structured grids
- *Composition* at compile-time achieves native performance while being *extensible* and allowing *design space exploration*
- Gives rise to a *novel library*, available publicly now
- *Artifact* available to fully reproduce our results, and to generate new data