

IN-MEMORY LAYOUT OF STRUCTURED DATA

KIWAKU, CONTAINERS WELL MADE

January 29, 2026

Adrien HENROT

Joël FALCOU, Hadrien GRASLAND, David CHAMONT

Context

- Unified access to data structures from HPC/HTC
- Simplification of their definition and usage
- Facilitate changes of in-memory layout
- Take advantage of memory hierarchies of modern processors

This work aims to be a collaboration between HEP and computer science research.

Goal: Combining Ergonomics and Performance

- Data access should ideally be both **efficient** and **simple**
- Logical definition should not constrain optimization
- In-memory layout should be modifiable without algorithm rewrite
- Solution path: towards a memory layout eDSL

Active research subject (LLAMA, SOACollection, Awkward Arrays,...)

Tools: C++

- Often used in HPC/HTC, powerfull and generalist language
- **Metaprogramming**: zero-cost abstraction (at runtime)
- **Concepts**: Constraints of generic code (templates)
- More compile time tools: constexpr, ...

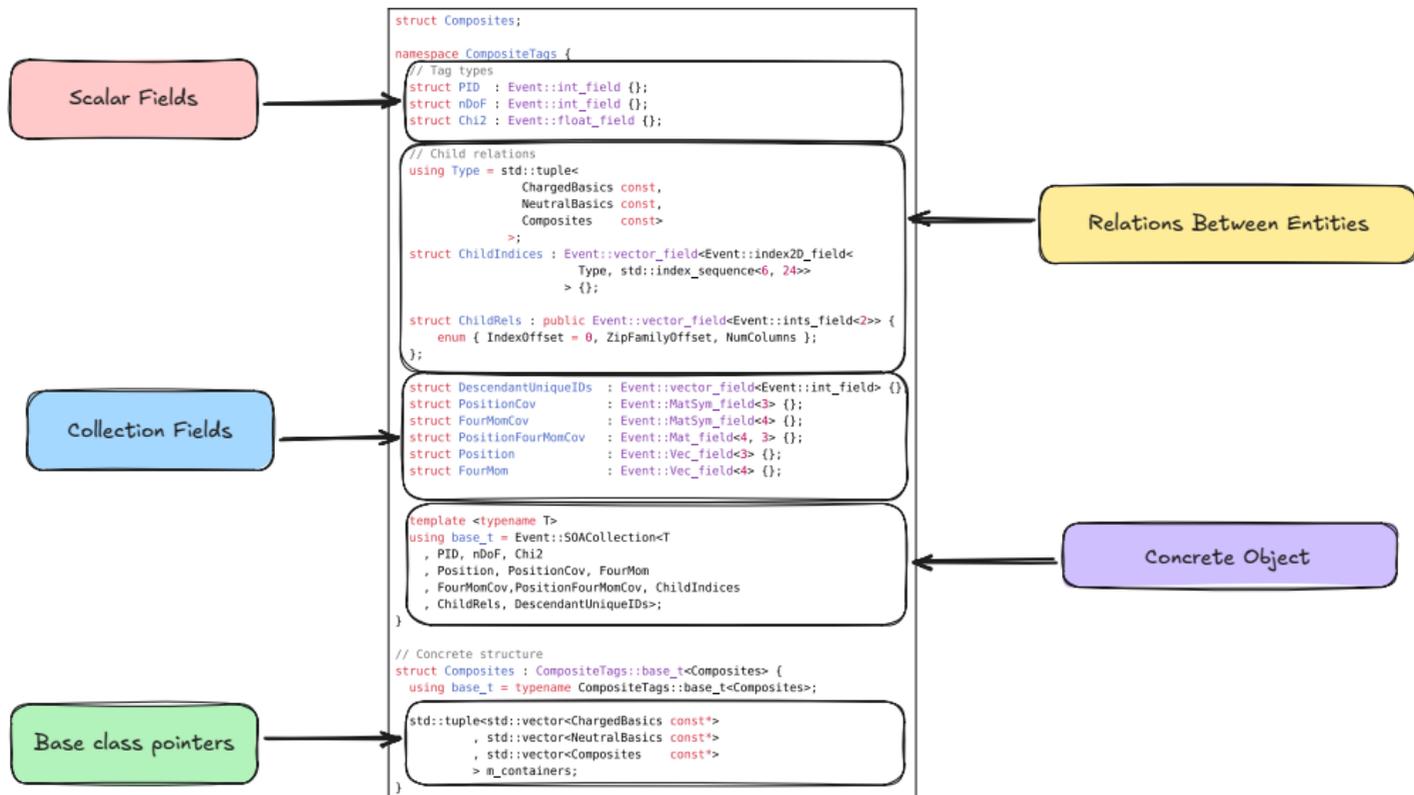
Structured Data

The Problem to be Solved

HPC	Physics
<pre>float matrix[300][40]; float vector[1200];</pre>	<pre>struct Composites; namespace CompositeTags { // Tag types struct PID : Event::int_field {}; struct nDoF : Event::int_field {}; struct ChI2 : Event::float_field {}; // Child relations using Type = std::tuple< ChargedBasics const, NeutralBasics const, Composites const>; }; struct ChildIndices : Event::vector_field<Event::index2D_field< Type, std::index_sequence<0, 24>> > {}; struct ChildRels : public Event::vector_field<Event::ints_field<2>> { enum { IndexOffset = 0, ZipFamilyOffset, NumColumns }; }; struct DescendantUniqueIDs : Event::vector_field<Event::int_field<>> {}; struct PositionCov : Event::MatSym_field<3> {}; struct FourMomCov : Event::MatSym_field<4> {}; struct PositionFourMomCov : Event::Mat_field<4, 3> {}; struct Position : Event::Vec_field<3> {}; struct FourMom : Event::Vec_field<4> {}; template <typename T> using base_t = Event::IGACollection< , PID, nDoF, ChI2 , Position, PositionCov, FourMom , FourMomCov, PositionFourMomCov, ChildIndices , ChildRels, DescendantUniqueIDs>; // Concrete structure struct Composites : CompositeTags::base_t<Composites> { using base_t = typename CompositeTags::base_t<Composites>; std::tuple<std::vector<ChargedBasics const*> , std::vector<NeutralBasics const*> , std::vector<Composites const*> > _containers; };</pre>
f32	?

- Efficiently storing data structure coming from physics?
- Generalization beyond that?

Example from LHCb Code



Classical Data Structures

- **Array of Structure (AoS)**
 - Spatial locality per structure
 - Access by index then by field
- **Structure of Arrays (SoA)**
 - Spatial locality per field
 - Access by field then by index
- **Tiled Structure of Arrays**
 - Mixed locality
 - Access by tile then inside the tile



Example

Structure definition	Data access
<pre>struct pixel { int r,g,b; }; std::array<pixel,10> aos{};</pre>	<pre>for(int i = 0; i < 10; ++i) print("{} ", aos[i].r);</pre>
<pre>template<std::size_t N> struct pixels { std::array<int,N> r,g,b; }; pixels<10> soa{};</pre>	<pre>for(int i = 0; i < 10; ++i) print("{} ", soa.r[i]);</pre>
<pre>constexpr int simd_width = 2; template<std::size_t N> struct block { std::array<int,simd_width> r,g,b; }; std::array<blocks,5> aosoa{};</pre>	<pre>for(int i = 0; i < 5; ++i) for(int j = 0; j < simd_width; ++j) print("{} ", aosoa[i].r[j]);</pre>

Observations

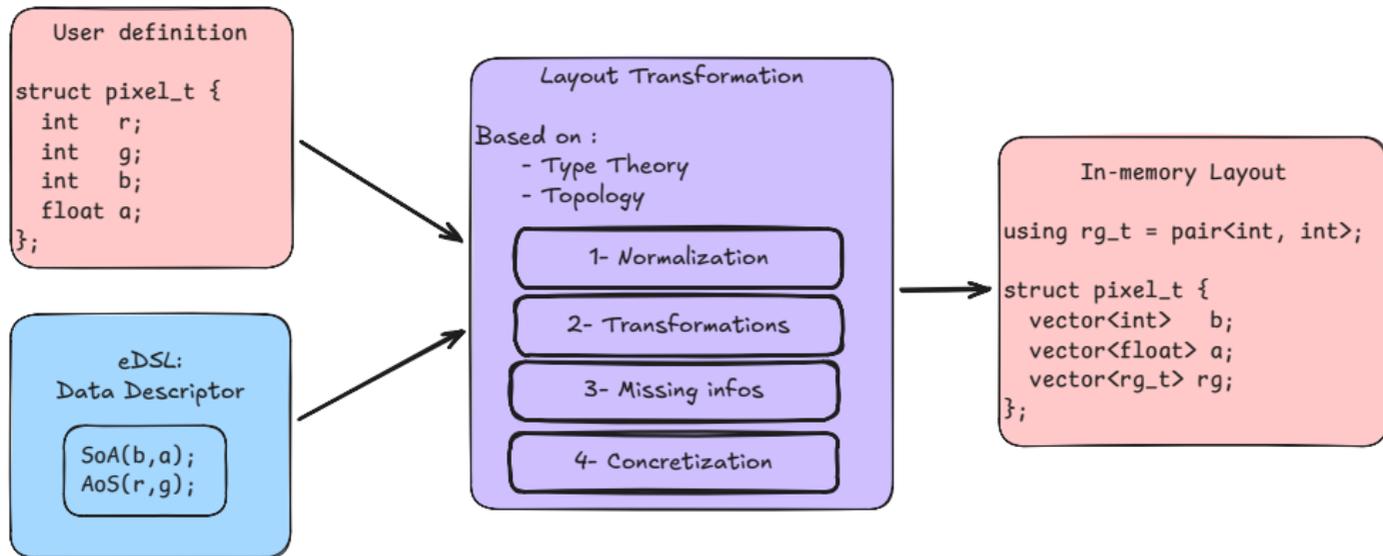
- Fields are fixed: tuple(R,G,B)
- Similar spaces
 - **Logical**: User-defined structure
 - **Physical**: In-memory realization
- The total number of elements is fixed:

$$nb_elements = \sum fields * \sum pixels = 3 * nb_pixels$$

- Different in-memory layout → different ways to access fields

Proposed Solution

Pipeline



- How to handle user-defined structures?
- How to use the final structure?

1- Structure Normalization

- Fields are fused
- Structures are flattened
- Collections are nested

API	Implementation
<pre>using info = pair<char, int>; struct hit { vector<info> data; }; struct tracks { vector<hit> hits; vector<float> momentum; } my_tracks{};</pre>	<pre>struct tracks { vector<float> momentum; vector<vector<char>> hits_data_0, vector<vector<int>> hits_data_1; } new_tracks{};</pre>

2- Applying Transformations

- Fields are reindexed
- Containers are re-attributed
- New names cannot be introduced

API	Implementation
<pre>struct tracks { vector<hit> hits; vector<float> momentum; } my_tracks{}; constexpr auto layout = aos("momentum", "hits.data.1");</pre>	<pre>struct combined { float momentum; vector<int> hits_data_1; } struct part_1 { vector<combined> members; };</pre>

3- Filling Missing Infos

- Unspecified fields are considered (not dropped)
- SoA layout by default
- In reality: Containers are linearized

User Side	Internally
<pre>using info = pair<char, int>; struct hit { vector<info> data; }; struct tracks { vector<hit> hits; vector<float> momentum; } my_tracks{};</pre>	<pre>struct combined {...}; struct final { vector<combined> part; vector<vector<char>> hits_data_0; } final_tracks{};</pre>

4- Concretization

- The structure is built: Index + logical field → Physical element

```
// How it would work at runtime
template<typename T>
constexpr auto final_tracks::resolve(string path, T idxs) {
    if ( path == "hits.data.0" )
        return this.hits_data_0[get<0>(idxs)][get<1>(idxs)];
    else if ( path == "hits.data.1" )
        return this.part[get<0>(idxs)].hits_data_1[get<1>(idxs)];
    else
        return this.part[get<0>(idxs)].momentum;
};

auto item = my_tracks.resolve("hits.data.1", pair{0, 0});
```

Actually handled at compile time: a structure layout can't depend on runtime parameters.

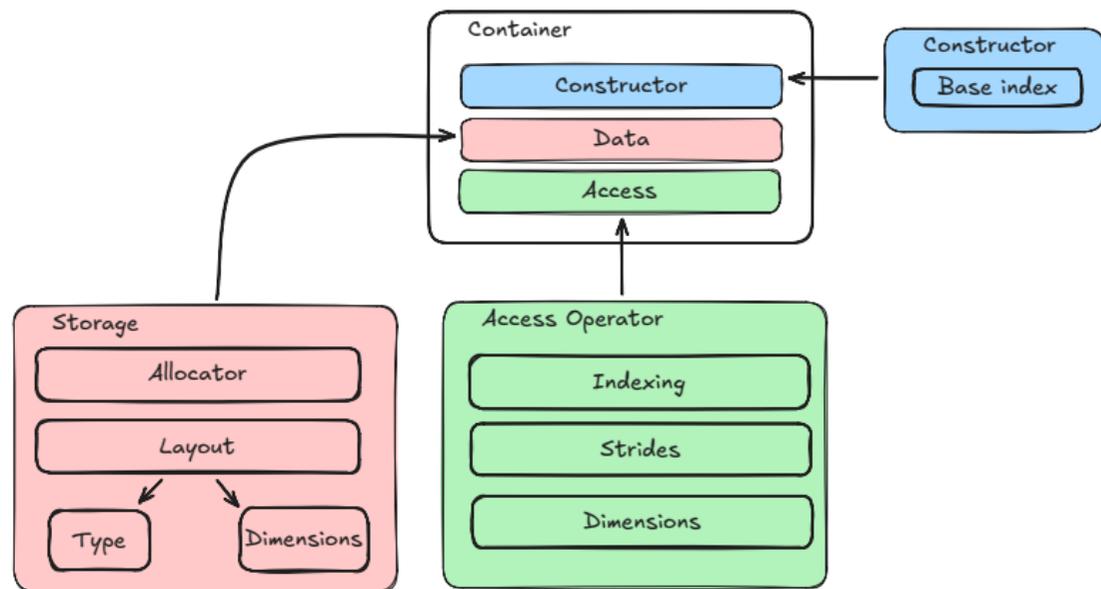
Kiwaku

Multidimensional Arrays Library

- C++20 and onward library
- Two main components
 - `kwk::view` and `kwk::table` (*ownership* semantic)
- Options to define properties

Options	Values
<code>kwk::source</code>	pointers, static arrays, contiguous ranges
<code>kwk::size</code>	<code>kwk::of_size(dn...)</code> , integers: 1D
<code>kwk::label</code>	String like value
<code>kwk::allocator</code>	A type modeling <code>kwk::concepts::allocator</code>

Internal Organization



- Type: element to store: 🏗️
- Allocator: Handles memory: 🚶
- Storage Order:
Order of the dimensions: 🏗️
- Dimensions:
Sizes per dimensions: 🏗️/🚶
- Strides:
Padding/Dimension jump: 🏗️
- Base index: default index value: 🚶
- Indexing: boundary conditions: 🏗️
- Layout: Where are the data: 🏗️

🏗️: Compile Time Info
🚶: Runtime Info

User API

```
Hadamard.cpp clang 20 (Editor #1)
A- Save/Load + Add new... Vim CppInsights Quick-bench clang 20 -std=c++20 -march=skylake-avx512 -O3 -DDEBUG
C++
9 void hadamard(auto &C0, auto const& C1, auto const& C2)
10 {
11     using namespace kwk;
12
13     auto shape = of_size(30,10);
14     auto va = view{ source = C0, shape };
15     auto vb = view{ source = C1, shape };
16     auto vc = view{ source = C2, shape };
17
18     kwk::for_each([](auto &a, auto const& b, auto const &c)
19     {
20         a = b * c;
21     }, va, vb, vc);
22 }
23
24
25
26
27
28
29
30

.LD00_3
30 vmovups ymm0, ymmword ptr [r14 + rax - 112]
31 vmulps ymm0, ymm0, ymmword ptr [r15 + rax - 112]
32 vmovups ymmword ptr [rbx + rax - 112], ymm0
33 vmovsd xmm0, qword ptr [r14 + rax - 80]
34 vmovsd xmm1, qword ptr [r15 + rax - 80]
35 vmulps xmm0, xmm0, xmm1
36 vmovlps qword ptr [rbx + rax - 80], xmm0
37 vmovups ymm0, ymmword ptr [r14 + rax - 72]
38 vmulps ymm0, ymm0, ymmword ptr [r15 + rax - 72]
39 vmovups ymmword ptr [rbx + rax - 72], ymm0
40 vmovsd xmm0, qword ptr [r14 + rax - 40]
41 vmovsd xmm1, qword ptr [r15 + rax - 40]
42 vmulps xmm0, xmm0, xmm1
43 vmovlps qword ptr [rbx + rax - 40], xmm0
44 vmovups ymm0, ymmword ptr [r14 + rax - 32]
45 vmulps ymm0, ymm0, ymmword ptr [r15 + rax - 32]
46 vmovups ymmword ptr [rbx + rax - 32], ymm0
47 vmovsd xmm0, qword ptr [r14 + rax]
48 vmovsd xmm1, qword ptr [r15 + rax]
49 vmulps xmm0, xmm0, xmm1
50 vmovlps qword ptr [rbx + rax], xmm0
```

- **kwk::tiles:**
 - Each tile can have its own iteration pattern
- **Interleaving:**
 - Across flat fields: ordering
 - Across inner structures: hierarchy
- **Heterogeneous data:**
 - Field types (product types)
 - Inner array sizes (jagged arrays)

Proof Of Concept

- Standard structure definition
- Simple way to define a reorganization

```
1  using namespace kumi; using namespace kwk;
2
3  struct pixel { using is_record_type = void; int r,g,b; float a; };
4
5  // Adatation of struct to record protocol
6  std::tuple_size<pixel> : std::integral_constant<std::size_t, 4>;
7  template<> struct std::tuple_element<0, pixel>{ using type = field_capture<"r", int>; };
8  template<> struct std::tuple_element<1, pixel>{ using type = field_capture<"g", int>; };
9  template<> struct std::tuple_element<2, pixel>{ using type = field_capture<"b", int>; };
10 template<> struct std::tuple_element<3, pixel>{ using type = field_capture<"a", float>; };
11
12 using pixel_map      = std::vector<pixel>;
13 constexpr auto layout = soa("g"_f, "a"_f);
14
15 auto pixels = structure(pixel_map{}, layout, 120);
```

With C++26/ROOT Reflection

- Standard structure definition
- Simple way to define a reorganization

```
1 using namespace kumi; using namespace kwk;
2
3 struct pixel { int r,g,b; float a; };
4
5 using pixel_map      = std::vector<pixel>;
6 constexpr auto layout = soa("g"_f, "a"_f);
7
8 auto pixels = structure(pixel_map{}, layout, 120);
```

Data Access

- Uniform access between different layouts
- Algorithms are independent of physical layout
- Efficient data access is guaranteed

```
1 auto to_grayscale = [](auto elt) {
2     auto &[r,g,b,a] = elt;
3     int gray = static_cast<int>(0.2126f * r + 0.7152f * g + 0.0722f * b);
4     r = g = b = gray;
5     return tuple{r,g,b,a};
6 };
7
8 int main() {
9     auto rng = pixels.as_range();
10    std::ranges::transform(rng, rng.begin(), to_grayscale);
11 }
```

Beyond std: Kiwaku Contexts

- Kiwaku supports several execution contexts
 - **simd**: EVE for vectorization
 - **kiwaku**: Native Kiwaku algorithms
 - **standard**: STL compatibility
- Different contexts → specific traversal algorithms

```
auto context = kwk::simd;  
auto context = kwk::kiwaku;  
auto context = kwk::standard;
```

In the end Kiwaku aims at supporting other contexts, for example: *multi-thread*

Conclusion

- High level abstraction with no impact on performance
- Simplified changes of memory layout
- Separation of concerns: Optimization/algorithms
- Simplification to write generic code
- Portable performance and equivalent to specialized code

Next Steps

- Clean up jagged array support
- Sum types support (unions, `std::variant`, ...)
- Sparse data representations (*COO*, *CSC*, *CSR*, ...)
- Context composition (SIMD, Multi-thread, GPU, ...)
- Writing/adaptating traversal algorithms
- Automate layout selection? (C++26 reflection, ROOT fallback)

Thanks for your attention

- Kiwaku: <https://github.com/jfalcou/kiwaku/>
- Kumi: <https://github.com/jfalcou/kumi/>

This work is funded by IN2P3 technical direction.

Annexes

Benchmark: Toy example

