

# Compiler support for semi-manual AoS-to-SoA conversions with data views

Pawel K. Radtke<sup>1</sup>[0009–0009–8613–3632] and Tobias  
Weinzierl<sup>1</sup>[0000–0002–6208–1841]

Department of Computer Science, Durham University, Durham, United Kingdom  
{pawel.k.radtke,tobias.weinzierl}@durham.ac.uk  
<https://scicomp.webspace.durham.ac.uk>

**Abstract.** The C programming language and its cousins such as C++ stipulate the static storage of sets of structured data: Developers have to commit to one, invariant data model—typically a structure-of-arrays (SoA) or an array-of-structs (AoS)—unless they manually rearrange, i.e. convert it throughout the computation. Whether AoS or SoA is favourable depends on the execution context and algorithm step. We propose a language extension based upon C++ attributes through which developers can guide the compiler what memory arrangements are to be used. The compiler can then automatically convert (parts of) the data into the format of choice prior to a calculation and convert results back afterwards. As all conversions are merely annotations, it is straightforward for the developer to experiment with different storage formats and to pick subsets of data that are subject to memory rearrangements. Our work implements the annotations within Clang and demonstrates their potential impact through a smoothed particle hydrodynamics (SPH) code.

**Keywords:** Array-of-structs, struct-of-arrays, memory layout transformations, data views, compiler, vectorisation

## 1 Introduction

Loops over sequences of data are the workhorses of scientific codes. Modern high-level languages provide us with the concept of structures to model our data elements. They are convenient to represent particles, mesh cells, and so forth. Due to its `structs`, the C++ language leans towards an array-of-structs (AoS) storage for sequences. The class is the primary modelling entity for the programmer, and the language favours sequences over class instances, i.e. objects [6, 7, 10, 11].

In many cases, implementations over structure-of-arrays (SoA) outperform their SoA counterpart. They facilitate the efficient usage of vector instructions [4, 9, 14] and are less sensitive to cache effects [7, 8, 14]: With AoS, compute kernels over sequences of structs have to gather and scatter vector register content (shuffling), vector parallelism is not blatant to the compiler, and structs for multiple loop iterations might not fit into the cache. With SoA, vector registers

can be filled with coalesced memory accesses, vector computations are exposed explicitly, and data from subsequent loop iterations is likely to reside in cache. While CPUs improve their gather-scatter efficiently with every new generation, the observations remain valid and apply GPUs, too [15].

SoA is not always superior to AoS. Tasks such as boundary data exchange, particle movements over meshes, or, in general, any algorithm that has to alter or permute struct sequences or access it in a non-continuous way [15] benefit from AoS. The size of characteristic sequences [7] and the memory footprint per struct further affect which storage format performs better. Finally, any runtime difference depends upon how successful the compiler vectorises an algorithm [11, 19] and what the target architecture looks like. The choice of an optimal data structure is context-dependent. There is no “one format rules them all”.

Refactoring code to accommodate memory rearrangements is error-prone and laborious. Wrappers allow developers to write their algorithms in a memory layout-agnostic way. C++ template meta programming combined with specialised containers are popular to achieve this [7, 10, 11, 14, 15]. This is a *static*, *global* approach. The data layout remains fixed. Any static strategy fails to react to the algorithmic context. A *dynamic*, *local* approach reorders data prior to the loop, such that only a local code block is aggressively optimised while the global data layout stays invariant.

To determine the permutation of data relative to plain AoS, we distinguish *manual* (user-driven) from *automatic* workflows. A *guided* one [11, 19] is a hybrid, where the actual transformation is carried out by the compiler, but only at the explicit request of the user. We propose a *dynamic*, local, guided approach: The code remains written in plain C++. Through additional annotations, programmers specify alternative data layouts for particular loops. A compiler takes the annotations and manually reorders data in a separate, temporary memory block prior to the loop. This is an out-of-place reordering, i.e. does not alter the original data layout [4, 5, 18]. The compiler also alters all corresponding data accesses such that they fit to the reordered copy of the data. A counterpart annotation ensures that data modifications are copied back at the end of the block, i.e. data are kept consistent.

Our approach is lightweight and non-invasive: The code remains correct if the annotations are unknown to the compiler. Different to manual optimisations or the implementation through future C++ generations supporting reflections [2], no code rewrites are required. This facilitates experimenting with different data layouts and separating algorithm development from performance tuning [3]. Our approach is not lightweight behind the scenes, as it introduces data movements. To reduce data movements, we rely on *views*: Developers can apply the AoS-to-SoA and SoA-to-AoS permutations to a subset of the structs’ attributes.

We demonstrate the potential of the idea by means of selected SPH kernels [13]. Individual SPH interactions can either be strongly memory-bound, or rely on compute-intense kernels. Since we allow the code base to stick to AoS overall, we do not negatively impact algorithmic phases such as the particle boundary exchange or any sorting, but still show that some kernels—depending

on the context and the algorithmic character—perform better due to temporal data reordering. We cannot yet provide a heuristic when conversions pay off. Our data suggests that common knowledge that it pays off only for Stream-like [12] kernels [7] or large arrays [15] is not unconditionally true.

Dynamic data structure transformations within a code have been used by several groups successfully [4, 18]. Our contribution is that we clearly separate storage format considerations from programming, introduce the notion of views, and move all data conversions into the compiler. As this approach requires no rewrite of existing code, additional libraries or new features such as reflections [2], it has the potential to streamline code development and code tuning. We notably do not increase the syntactic complexity of the code. As we challenge the common knowledge that temporal data reordering prior to loops or computational kernels is problematic [8, 9], we lay the foundations of making these data layout optimisations a standard optimisation step within a compiler pipeline.

The remainder is organised as follows: We sketch our use case first to motivate our work (Section 2), though all ideas are more widely applicable. We next introduce our code annotations, and then discuss their semantics (Section 3). This allows us to realise the annotations in Section 4, before we finally study their impact. A brief discussion and outlook in Section 6 close the discussion.

## 2 Demonstrator use case

We motivate and illustrate our ideas by means of the compute kernels of a simple smoothed particle hydrodynamics (SPH) code. SPH is a Lagrangian technique. The particles interact, move and carry properties. This way, they represent the underlying fluid flow. We zoom into the elementary operations of any SPH code, which is the hydrodynamic evolution with leapfrog as SPH’s predominant time integrator. More complex physical models typically start from there [13].

*Algorithmic kernels.* We focus on  $N$  particles and assume that they can interact with any other particle. This resembles an SPH code which clusters the computational domain into control volumes each hosting a small, finite number of particles. The control volumes are chosen such that particles only interact with other particles within the volume and its neighbours, as SPH typically works with a finite interaction radius. We assume to know its upper bound. The core compute steps of the algorithm then read as follows:

1. We determine the smoothing length of each particle, i.e. the interaction area, and its *density*. This initial step evaluates two-body interactions, i.e. is in  $\mathcal{O}(N^2)$ . It studies all nearby particles within the volume and decides if to shrink or increase the interaction radius. The process then repeats. The density computation overall is iterative, but we study one iteration step only.
2. We compute the *force* acting on each particle. The force is the sum over all forces between a particle and its neighbours within the interaction radius. Overall, the force calculation is asymptotically of quadratic complexity.

3. We *kick* the particle by half a time step, i.e. we accelerate it. This is a mere loop over all particles within a cell and hence of linear complexity.
4. We *drift* a particle, i.e. update its position. This is another loop in  $\mathcal{O}(N)$ .
5. We *kick* a second time, i.e. add another acceleration.

*Memory layout and data flow.* Our SPH particle is modelled as a struct. We work with AoS. The hosted structs aka particles can hold from a few quantities up to hundreds of doubles. Our benchmark code induces a memory footprint of 256 bytes per particle. Some of these bytes encode administration information, others store physical properties. For many steps, only few attributes enter the equations. The others are, within this context, overhead [1, 7].

The density calculation starts from the density and smoothing length of the previous time step and updates those two quantities, and others such as the neighbour count, the rotational velocity vector, and various derivative properties. Smoothing length, density and further properties feed into the force accumulation which eventually yields an acceleration per particle. Kicks are relatively simple, i.e. add scaled accelerations onto the velocity. The second kick in our implementation also resets a set of predicted values which feed into the subsequent density calculation. Drifts finally implement a simple Euler time integration step, i.e. a tiny daxpy.

SPH is a Lagrangian method and hence works with particles which are scattered over the computational domain. While we tend to hold the particles continuously in memory per cell to facilitate vectorisation, the nature of the moving particles implies that we have to resort frequently. In a parallel domain decomposition environment, we furthermore have to exchange few particles per time step as part of the halo exchange.

### 3 Source code annotations

Converting AoS into SoA is an evergreen in high-performance computing, once we have committed to AoS as development data structure (cmp. for example [3, 4, 7–11, 14–16, 18, 19]). Vector computing in a SIMD or SIMT sense including coalesced memory accesses, cache and TLB effects drives such rewrites. A sophisticated conversion takes into account whether we have to convert only a subset of the struct, i.e. if we can peel a struct or open a view [8].

We propose that developer focus exclusively on either SoA or AoS. For the demonstrator from Section 2, AoS is, in line with many codes [6, 7, 10, 11], a natural choice. The (temporal) data permutation into SoA then is delegated to the compiler. For this we introduce C++ attributes (Algorithm 1):

- Whenever the compiler encounters a `[[clang::soa_conversion_target]]` attribute, we instruct the translator to convert the array named as argument into SoA prior to entering the following loop. All attribute accesses within the loop will be redirected to the converted data. We instruct the compiler to perform a temporary out-of-place transformation [16].

---

**Algorithm 1** The drift of all particles within a control volume is annotated with instructions to the compiler to convert parts of the underlying AoS data into SoA prior to the actual loop invocation.

---

```

1: function VOID DRIFT(Particle *particles, int size)
2:   [[clang::soa_conversion_target(particles)]]
3:   [[clang::soa_conversion_target_size(size)]]
4:   [[clang::soa_conversion_inputs(pos, vel, updated)]]
5:   [[clang::soa_conversion_outputs(pos, updated)]]
6:   for (int i = 0; i < size; i++) do
7:     auto &p = particles[i];
8:     p.pos[0] += p.vel[0] * dt;
9:     p.pos[1] += p.vel[1] * dt;
10:    ...
11:    p.updated = true;
12:   end for
13: end function

```

---

- To trigger the conversion, users have to specify the size of the target array through `[[clang::soa_conversion_target_size]]`, as plain C arrays do not come along with such meta information.
- We can convert the whole struct, i.e. all attributes, hosted within an AoS, or we can restrict the conversion to particular attributes of these structs through `[[clang::soa_conversion_inputs]]`. The input annotation opens a *view* on the struct. It peels the struct [8].
- Alterations to the implicitly created SoA view become invalid once we leave the code block. If changes should be copied back into the original data, i.e. *synchronised*, users have to add `[[clang::soa_conversion_outputs]]`.

### 3.1 Transformation semantics

Let  $S$  be a struct and  $\mathbb{S} = [S_0, S_1, S_2, \dots]$  be a sequence, i.e. array of these structs (AoS). We assume that  $\mathbb{S}$  is identified through a pointer to the first element of the sequence. As we work with raw data in a C sense, we require explicit information on the size  $|\mathbb{S}|$  of  $\mathbb{S}$  from the user, though passing `s.size()` is admissible if we work with C++'s `std::array` or `std::vector`, e.g.

`[[clang::soa_conversion_target]]` over  $\mathbb{S}$  identified by its pointer informs us that  $\mathbb{S}$  is held as AoS, and it introduces a mapping  $aos\_to\_soa : \mathbb{S} \mapsto \hat{\mathbb{S}}$ .  $\hat{\mathbb{S}}$  is a tuple of data, where each entry is an array of a primitive type with the size  $|\mathbb{S}|$ . With `[[clang::soa_conversion_inputs]]`, the set  $\mathbb{A}$  of attributes that are mapped can be restricted. Not every attribute of  $S$  has to be copied into the reordered data area.

Our compiler does not automatically copy content from the temporary data structure back. It does not automatically synchronise  $\hat{\mathbb{S}}$  and  $\mathbb{S}$ . Instead, we require users to specify which attribute set  $\hat{\mathbb{A}}$  to copy back through `[[clang::soa_conversion_outputs]]`. `[[clang::aos_conversion_target]]` in-

roducing *soa\_to\_aos* and its counterpart `[[clang::aos_conversion_outputs]]` are well-defined the other way round.

*Out-of-place temporary data.* Let  $\hat{\mathbb{S}}$  hold attributes from  $\mathbb{A} \cup \hat{\mathbb{A}}$ . That is, the temporary data hold all attributes which are either defined through input or output statements. We have no subset relation and allow  $\mathbb{A} \cap \hat{\mathbb{A}} \neq \emptyset$ . With this,  $\text{mem}(\hat{\mathbb{S}}) \leq \text{mem}(\mathbb{S})$  if *mem* yields the memory footprint. We furthermore note that  $\text{soa\_to\_aos} = \text{aos\_to\_soa}^{-1}$  whenever  $\mathbb{A} = \hat{\mathbb{A}}$ .

*Prologue.* The target statement adds a preamble to the following loop. The preamble opens with the creation of a temporary data structure  $\hat{\mathbb{S}}$ , and then immediately fills the image with

$$\hat{S}.a_i \leftarrow S_i.a \quad \forall 0 \leq i < |\mathbb{S}|, \forall a \in \mathbb{A}. \quad (1)$$

Only attributes  $a \in \mathbb{A}$  of  $\mathbb{S}$  are copied into the SoA helper data structure due to (1). For  $\mathbb{A} \neq \hat{\mathbb{A}}$ , we do not fill all attributes in  $\hat{\mathbb{S}}$ . Some fields hold garbage.

*Redirection.* Any subsequent access to  $S_i.a$ ,  $a \in \mathbb{A} \cup \hat{\mathbb{A}}$  within the loop is redirected to  $\hat{S}.a_i$ . Our compiler extension redirects accesses to the structs within  $\mathbb{S}$  to the structure of arrays. Any follow-up optimisation pass within the compiler pipeline will hence work with SoA. Through an explicit specification of  $\mathbb{A}$  and  $\hat{\mathbb{A}}$ , we can shrink the memory footprint of  $\hat{\mathbb{S}}$  and reduce the copy overhead. Attribute accesses  $a \notin \mathbb{A} \cup \hat{\mathbb{A}}$  continue to hit  $\mathbb{S}$  directly.

*Epilogue.* The epilogue adds

$$S_i.a \leftarrow \hat{S}.a_i \quad \forall 0 \leq i < |\mathbb{S}|, \forall a \in \hat{\mathbb{A}}. \quad (2)$$

to the outcome code, followed by a free of  $\hat{\mathbb{S}}$ . In principle, the epilogue uses *soa\_to\_aos*. The difference between (2) and (1) is the temporal order: An output statement inserts the mapping into an epilogue, while the target statement tackles the prologue.

### 3.2 Extensions

Our language extension is focused on one-dimensional data sets which are logically continuous, i.e. allow for coalescent data access: A container traversed via `for (auto& p: ...)` can technically be realised with scattered data  $\mathbb{S}$ . Our conversion yields one continuous excerpt  $\hat{\mathbb{S}}$ .

We implicitly exploit that any loop imposes an ordering. Therefore, we can extend our annotations by `[[clang::soa_conversion_start_idx()]]` which restrict the conversion to a subset of the range. There is no need for a `[[clang::soa_conversion_end_idx()]]`, as the size is manually specified.

---

**Algorithm 2** Pseudo code which illustrates how the compiler rewrites the example in Algorithm 1. Embedded arrays where the size has to be known at compile time and types of all converted data can be analysed by the compiler from the source code.

---

```

1: function VOID DRIFT(Particle *particles, int size)
2:   pos0_tmp = new double[size];           ▷ temporary out-of-place array
3:   pos1_tmp = new double[size];
4:   ...
5:   for (int i = 0; i < size; i++) do
6:     pos0_tmp[i] = particles[i].pos[0];    ▷ out-of-place AoS-SoA conversion
7:     ...
8:   end for
9:   for (int i = 0; i < size; i++) do
10:    pos0_tmp[i] += vel0_tmp * dt;          ▷ p.pos[0] += p.vel[0] * dt;
11:    ...
12:    updated_tmp[i] = true;                 ▷ p.updated = true;
13:   end for
14:   for (int i = 0; i < size; i++) do
15:    particles[i].pos[0] = pos0_tmp[i];     ▷ SoA-AoS conversion due
16:    ...                                   ▷ to outputs statement
17:   end for
18:   delete[] pos0_tmp;                     ▷ free temporary arrays
19:   ...
20: end function

```

---

## 4 Compiler realisation

We implement our proposed techniques prototypically such that we can highlight their potential impact for the SPH demonstrator. For this, we plug into the Clang LLVM front-end. Clang takes C++ code and emits unoptimised LLVM IR. As we work within Clang, our modifications affect the resultant unoptimised intermediate representation (IR) output, and do not alter any downstream steps within the translation pipeline. Notably, the generated IR code still benefits from all LLVM-internal optimisation passes. Clang’s lexer, parser and semantic analyser yield an abstract syntax tree (AST). The AST then is subsequently consumed to produce LLVM IR. These steps constitute the compiler’s **FrontendAction**.

We realise our functionality with a new **FrontendAction**. It traverses the AST top down and searches for our annotations. When it encounters a convert, it inserts allocation statements for the out-of-place memory allocations once all information on the data view are available, and it adds the actual data copying from (1). The names of the temporary variables are mangled to avoid variable re-declarations. Following the prologue step, we query the actual loop body and redirect the memory accesses to  $\mathbb{A} \cup \hat{\mathbb{A}}$ . Finally, we add an epilogue that synchronises the data according to (2) and frees the temporary buffers. As AST modifications are not possible in Clang, i.e. as the AST is immutable, our realisation pretty prints the AST into C++ subject to our modifications and then in-

vokes Clang’s original `FrontendAction`. This process happens in-memory, and is transparent to the end user. The behaviour can be disabled on the command line through `-fno-soa-conversion-attributes-language-extension`. Our compiler optimisation pass materialises in a source-to-source compiler (cmp. rewrite of Algorithm 1 into 2) which feeds into subsequent translation sweeps [18].

As we precede Clang’s original `FrontendAction`, it is possible to dump our output into source files explicitly instead of implicitly passing it on into the subsequent `FrontendAction`. While originally designed for debugging, this feature is particularly appealing in environments where the modified compiler is only available locally, while the compilers on the target production platform cannot be modified.

As we navigate the source code using the AST, the effect of our annotations is scoped at the translation-unit level: The data transformations cannot propagate into user translation units (object files) and notably fail to propagate into libraries. Instead, the annotations will break the code if library functions are used internally. While functions with signatures similar to `foo(particles[i])` (cmp. Algorithm 1) cannot be used in code blocks subject to the code transformations, functions capturing attributes individually are supported. A function `foo(particles[i].pos[0], particles[i].pos[1], ...)` will continue to work. A sequence of `foo` calls will even vectorise if translated accordingly (`declare simd` in OpenMP for example).

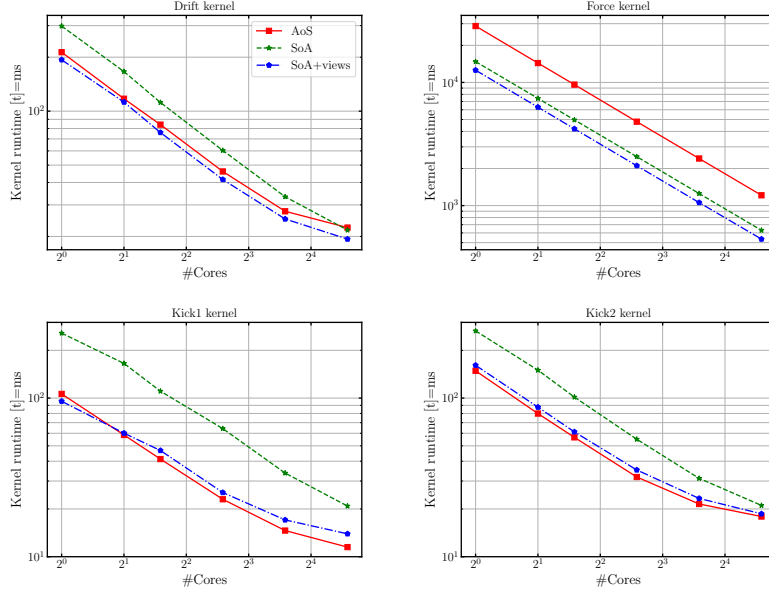
## 5 Benchmark results

We assessed the impact of our compiler prototype on two architectures. Our first system is an AMD EPYC 9654 (Genoa) testbed. It features  $2 \times 96$  cores over  $2 \times 4$  NUMA domains spread over two sockets, hosts an L2 cache of 1,024 KByte per core and a shared L3 cache with 384 MByte per socket. Our second system is an Intel Xeon Gold 6430 (Sapphire Rapid). It features  $2 \times 32$  cores over two sockets. They form two NUMA domains with an L2 cache of 2,048 KByte per core and a shared L3 cache with 62 MByte per socket. To avoid NUMA effects, we confine all our Intel and AMD tests to a single NUMA domain through `numactl`.

*Upscaling.* We start our studies with a classic strong scaling configuration for  $4.19 \cdot 10^6$  particles. The particles are held as AoS and we iterate over them with a parallel loop invoking one of our kernels only. Our measurements distinguish three different variants: A plain AoS serves as vanilla version. This is the code base actually written down in C++. We compare this to SoA where the whole particle sequence is converted prior to loop and synchronised back afterwards, i.e.  $A = \hat{A}$  with all attributes of the struct involved. Finally, we assess a version where views narrow down the attribute sets to the number of variables which are actually read or written, respectively.

The force calculation scales almost perfectly as it is compute-bound, while kicks and drifts tail off (Figure 1). This implies that the latter ones become bandwidth-bound once we use a sufficiently high number of cores. Converting





**Fig. 1.** Scalability plots for  $2048^2$  particles on a single node for four different kernels (Genoa). We benchmark the baseline code (AoS) against a version which converts all of the particle data (SoA) against a version which works with views.

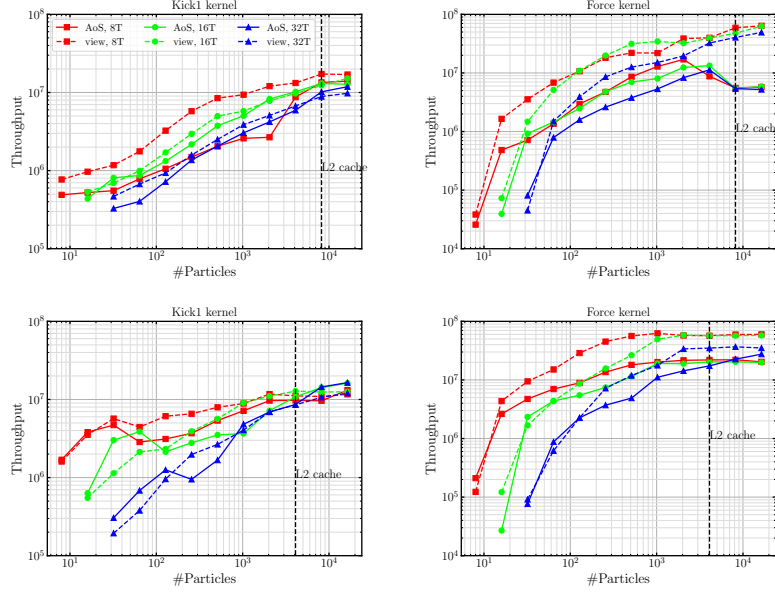
into a SoA leads to a higher compute time, but narrowing down the conversion to views eliminates this penalty. For the drift, the converted version becomes even slightly faster than the AoS vanilla version. The situation is fundamentally different for the computationally demanding force kernel. Here, the temporary conversion into SoA pays off, and the views then help to reduce the runtime even further. All measurements for the Sapphire Rapid yield qualitatively similar data.

Adding additional data movements due to AoS–SoA conversions to computationally cheap kernels is poisonous: Even an improved AVX512 vectorisation cannot compensate for the logical latency that our out-of-place conversion introduces. We delay the actual loop launch as we first have to convert data. If we employ many cores and operate overall in the saturation regime, the difference starts to disappear. As a core has to wait for the memory controller anyway, it can as well use the cycles to convert the data structure in a nearby cache.

If a kernel exhibits sufficient compute load, a temporary conversion pays off. This contradicts the intuition that a conversion is particularly beneficial for Stream-like operations [7, 9]. It aligns with other papers where temporary transformations are used for expensive code parts [4].

*Impact of particle count.* We continue with a test where we alter the number of particles per kernel, draw 16 samples per run to avoid noise, and study the throughput, i.e. the number of particle updates or the number of particle–particle

interactions per second, respectively. As the scalability of the benchmark is confirmed, we stick to three characteristic thread counts. As we have demonstrated the impact of the views, we focus on a comparison of the views to the vanilla AoS storage.



**Fig. 2.** Dependency on particle counts for fixed thread numbers (8, 16 and 32). Sapphire Rapid (top) vs. Genoa data (bottom). The L2 cache size denotes the size for a single thread, i.e. how many particles would fit into one single L2 cache.

For the cheap kernels, the conversion pays off as long as the particle count is contained within a machine-specific optimal range. The efficiency gap closes on Intel systems as we use more threads, and becomes rather erratic on AMD (Figure 2). “Optimal range” notably implies that we are not in a scaling regime yet, where using all threads pays off. For the compute-heavy force kernel, the gap between sole AoS and the views widens on the Intel system as the particle count increases. We notably observe a “sudden” throughput deterioration for the plain AoS which does not occur for the converted algorithm. On AMD, the curves for both schemes are relatively smooth, yet they close again if we use too many threads: The SoA version with views stagnates, while the AoS version catches up.

The large caches on both testbeds imply that we never fall out of L3 cache once the first warm-up run out of 16 is completed. We hence rule out last level cache misses to explain the throughput behaviour. Our L3 is filled with the AoS data once per test, and these data are then converted 16 times into SoA with views and back. We manually verified that the machine code (AVX vectorisation)

is comparable on both target systems and is efficient. Runtime behaviour and qualitative differences in the data thus have to be explained through the memory access characteristics.

We need a decent number of particles per thread to give the vectorisation the opportunity to unfold its potential. For very small particle counts per thread, a conversion never pays off. For small particle counts per thread, we however benefit within the loop body from improved memory access characteristics. Once the particle count increases, the advantage is eaten up by the single-threaded conversion in the epilogue and preamble of the loop if we encounter an algorithm with linear complexity and low computational load. If we have “too many” threads, they have to synchronise through the shared L3 cache, which adds a further penalty and renders the conversion useless.

For high computational load and quadratic complexity, the improvement in better memory access characteristics outweighs the conversion penalty. The performance gap widens. On the Intel system, the original AoS version takes a hit once we fall out of the L2 cache on a core. The converted SoA memory comprises only the hot data [18], i.e. the data we actually work on. It continues to fit into the L2. On the AMD system, we however quickly start to suffer from NUMA effects once we employ too many threads. We have to give up on our performance wins.

*Context.* Our demonstrator uses structs that are embedded into other structs [14], as we work with coordinate vectors within the particle structure. It supports data hierarchy through struct composition, but seems to restrict itself exclusively to AoS and SoA. It is important to realise however that we actually support more sophisticated data structures, too. This observation results directly from the fact that our conversion is local and temporary.

Our underlying code base [13] and many other codes advocate for the use of AoSoA [3] or ASTA [16] to compromise between speed and flexibility. As we may convert subarrays of a given AoS structure, our approach can logically yield an AoSoA/AoS hybrid. The overall data are stored as AoS, but subsections flip to SoA. Along the same lines, our techniques allow for the manipulation of data structures where chunks of AoS data are scattered over the heap. It does not require one global data space. Both properties are important for our SPH code base, where cells for example host a pointer. The data per cell are held as AoS, but the cells’ data are scattered over the heap.

As long as SPH codes run through the individual algorithmic phases step by step and synchronise after each step in a fork-join manner, only the algorithmic phases with quadratic complexity per control volume seem to benefit. However, once we work with codes that work with functional decomposition (task-based parallelism) or employ multiple ranks per node, simulations tend to run different algorithmic steps on different threads concurrently. As a result, fewer threads than the whole node configuration are effectively available to a compute step. It might be reasonable to parallelise the conversion into SoA and back. In practice, it is not clear if this adds significant value as a step rarely “owns” many or all threads.

## 6 Conclusion and outlook

We introduce a local, guided approach to convert from AoS into SoA and vice versa within simulation codes. When users add our novel C++ annotations to their code, the compiler rewrites its underlying data structure out-of-place into a temporary scratchpad for the affected code block, redirects all source code operations to this altered memory location, and eventually synchronises (parts of) the temporary data structure with its original again. Our approach allows for the incremental, non-disruptive optimisation of code, as all code remains valid if a compiler does not support our annotations, and as users can continue to design their implementation with their favourite data layout in mind [3]. In object-oriented codes, this will likely be AoS. Our experimental studies challenge the wide-spread assumption that such temporary, local reordering is not affordable. It is affordable once we introduce the notion of a view, i.e. the opportunity to restrict transformations only to some attributes of a struct.

Our data are obtained through a prototypical Clang compiler extension. A more mature, yet more heavy-weight long-term realisation might implement the transformations within LLVM’s MLIR layer and hence make it independent of the used front-end. The promise here is that the conversion would automatically benefit from other memory optimisations such as the automatic introduction of padding and alignment. Within the C++ domain, it is an interesting challenge for future work to discuss how the data transformations could be integrated with C++ views or memory abstractions as pioneered within Kokkos [17], as the conversion gain is clearly tied to the underlying problem sizes and data layouts. Both directions of travel for the compiler construction are timely and reasonable, as our work has demonstrated that temporal out-of-place conversions have the potential to speed up code. In particular, it is reasonable to assume that our idea unfolds its full potential once we apply it to GPU-based codes [16]. It also feeds into a high-level discussion whether data transformations should be offered to developers [2] or realised by and within the compiler subject to developer guidance.

There is an elephant in the room: We leave the responsibility to insert appropriate data reordering instructions with the user and do not try to move the decisions themselves into the compiler [8]. Designing an automatic data transformation requires appropriate heuristics. Our studies suggest that such a heuristic is inherently difficult to construct, as it would require a “what if” analysis: what would the gain of the transformation be, i.e. how much performance improvement would unfold due to better vectorisation, reduced cache misses and reduced TLB, and could those improvements compensate for any transformation overhead. Each of these questions has to be answered over a huge space of potential reorderings. We hence hypothesise that only feedback optimisation (static) or on-the-fly adjourning and just-in-time compilation can deliver beneficial automatic reordering [8, 18].

A second train of thought arises from the insight that conversions benefit greatly from the concept of views, i.e. partial permutations and copying of the data structures. Future work has to study how we can weaken the notion of

temporal, localised transformations [6] further: It is a natural idea to preserve converted data over multiple code blocks, i.e. not to free and re-allocate them but to synchronise selectively. This would lead to a setup where data is held both in AoS and SoA—and both subject of different views—and the compiler automatically synchronises these representations upon demand. We expect this to reduce overhead massively, trading memory for speed.

## Acknowledgments

Tobias’ research has been supported by EPSRC’s Excalibur programme through its cross-cutting project EX20-9 *Exposing Parallelism: Task Parallelism* (Grant ESA 10 CDEL) and the DDWG projects *PAX-HPC* (Gant EP/W026775/1) as well as *An ExCALIBUR Multigrid Solver Toolbox for ExaHyPE* (EP/X019497/1). His group appreciates the support by Intel’s Academic Centre of Excellence at Durham University. The work has received funding through the eCSE project ARCHER2-eCSE11-2 “ExaHyPE-DSL”.

The code development relied on experimental test nodes installed within the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility ([www.dirac.ac.uk](http://www.dirac.ac.uk)). The equipment was funded by BEIS capital funding via STFC capital grants ST/K00042X/1, ST/P002293/1, ST/R002371/1 and ST/S002502/1, Durham University and STFC operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure.

## References

1. Bungartz, H.J., Eckhardt, W., Weinzierl, T., Zenger, C.: A precompiler to reduce the memory footprint of multiscale pde solvers in c++. *Future Generation Computer Systems* **26**(1), 175–182 (Jan 2010). <https://doi.org/10.1016/j.future.2009.05.011>
2. Childers, W., Dimov, P., Katz, D., Revzin, B., Sutton, A., Vali, F., Vandevoorde, D.: Reflection for c++26 (published 2024-08-14, retrieved 2024-08-29), programming Language C++, P2996R5, <https://www.open-std.org/jtc1/sc22/wg21/docs/papers/2024/p2996r5.html>
3. Gallard, J.M., Krenz, L., Rannabauer, L., Reinarz, A., Bader, M.: Role-Oriented Code Generation in an Engine for Solving Hyperbolic PDE Systems, vol. 1190, p. 111–128. Springer International Publishing, Cham (2020). [https://doi.org/10.1007/978-3-030-44728-1\\_7](https://doi.org/10.1007/978-3-030-44728-1_7)
4. Gallard, J.M., Rannabauer, L., Reinarz, A., Bader, M.: Vectorization and minimization of memory footprint for linear high-order discontinuous galerkin schemes. In: 2020 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW). p. 711–720. IEEE, New Orleans, LA, USA (May 2020). <https://doi.org/10.1109/IPDPSW50202.2020.00126>
5. Gustavson, F., Karlsson, L., Kågström, B.: Parallel and cache-efficient in-place matrix storage format conversion. *ACM Transactions on Mathematical Software* **38**(3), 1–32 (Apr 2012). <https://doi.org/10.1145/2168773.2168775>

6. Hirzel, M.: Data layouts for object-oriented programs. In: Proceedings of the 2007 ACM SIGMETRICS international conference on Measurement and modeling of computer systems. p. 265–276. ACM, San Diego California USA (Jun 2007). <https://doi.org/10.1145/1254882.1254915>
7. Homann, H., Laenen, F.: Soax: A generic c++ structure of arrays for handling particles in hpc codes. *Computer Physics Communications* **224**, 325–332 (Mar 2018). <https://doi.org/10.1016/j.cpc.2017.11.015>
8. Hundt, R., Mannarswamy, S., Chakrabarti, D.: Practical structure layout optimization and advice. In: International Symposium on Code Generation and Optimization (CGO’06). p. 233–244. IEEE, New York, NY, USA (2006). <https://doi.org/10.1109/CGO.2006.29>
9. Intel: Memory layout transformations, <https://www.intel.com/content/www/us/en/developer/articles/technical/memory-layout-transformations.html>
10. Jeffers J, Reinders J, S.A.: Intel Xeon Phi processor high-performance programming. Knights landing edition. Morgan Kaufman (2016)
11. Jubertie, S., Masliah, I., Falcou, J.: Data layout and simd abstraction layers: Decoupling interfaces from implementations. In: 2018 International Conference on High Performance Computing & Simulation (HPCS). p. 531–538. IEEE, Orleans (Jul 2018). <https://doi.org/10.1109/HPCS.2018.00089>
12. McCalpin, J.: Memory bandwidth and machine balance in high performance computers. IEEE Technical Committee on Computer Architecture Newsletter pp. 19–25 (12 1995)
13. Schaller, M., Borrow, J., Draper, P.W., Ivkovic, M., McAlpine, S., Vandenbroucke, B., Bahé, Y., Chaikin, E., Chalk, A.B.G., Chan, T.K., Correa, C., van Daalen, M., Elbers, W., Gonnet, P., Hausammann, L., Helly, J., Huško, F., Kegerreis, J.A., Nobels, F.S.J., Ploekinger, S., Revaz, Y., Roper, W.J., Ruiz-Bonilla, S., Sandnes, T.D., Uyttenhove, Y., Willis, J.S., Xiang, Z.: Swift: a modern highly parallel gravity and smoothed particle hydrodynamics solver for astrophysical and cosmological applications. *Monthly Notices of the Royal Astronomical Society* **530**(2), 2378–2419 (2024). <https://doi.org/10.1093/mnras/stae922>
14. Springer, M., Sun, Y., Masuhara, H.: Inner array inlining for structure of arrays layout. In: Proceedings of the 5th ACM SIGPLAN International Workshop on Libraries, Languages, and Compilers for Array Programming. p. 50–58. ACM, Philadelphia PA USA (Jun 2018). <https://doi.org/10.1145/3219753.3219760>
15. Strzodka, R.: Abstraction for AoS and SoA Layout in C++, p. 429–441. Elsevier (2012). <https://doi.org/10.1016/B978-0-12-385963-1.00031-9>
16. Sung, I.J., Liu, G.D., Hwu, W.M.W.: Dl: A data layout transformation system for heterogeneous computing. In: 2012 Innovative Parallel Computing (InPar). pp. 1–11 (2012). <https://doi.org/10.1109/InPar.2012.6339606>
17. Trott, C.R., Lebrun-Grandie, D., Arndt, D., Ciesko, J., Dang, V., Ellingwood, N., Gayatri, R., Harvey, E., Hollman, D.S., Ibanez, D., Liber, N., Madsen, J., Miles, J., Poliakoff, D., Powell, A., Rajamanickam, S., Simberg, M., Sunderland, D., Turcksin, B., Wilke, J.: Kokkos 3: Programming model extensions for the exascale era. *IEEE Transactions on Parallel and Distributed Systems* **33**(4), 805–817 (Apr 2022). <https://doi.org/10.1109/TPDS.2021.3097283>
18. T.V., V., N., V.: Implementing data layout optimizations in the llvm framework, <https://llvm.org/devmtg/2014-10/Slides/Prashanth-DLO.pdf>
19. Xu, S., Gregg, D.: Semi-automatic Composition of Data Layout Transformations for Loop Vectorization, vol. 8707, p. 485–496. Springer Berlin Heidelberg, Berlin, Heidelberg (2014). [https://doi.org/10.1007/978-3-662-44917-2\\_40](https://doi.org/10.1007/978-3-662-44917-2_40)

## A Compiler download

The forked Clang/LLVM project is publicly available at <https://github.com/pradt2/llvm-project>, our extensions are available in the `hpc-ext` branch. To avoid conflicts, it is strongly recommended to remove any existing Clang/LLVM installations before proceeding.

Once cloned, we create a build folder in the cloned repository, and afterwards execute the following command:

```
cmake
-DMAKE_BUILD_TYPE="RelWithDebInfo"
-DMAKE_C_COMPILER="gcc"
-DMAKE_CXX_COMPILER="g++"
-DLLVM_ENABLE_PROJECTS="clang;openmp"
-DLLVM_TARGETS_TO_BUILD="host"
-DBUILD_SHARED_LIBS="ON"
-G "Unix Makefiles"
../llvm
```

`make -j$(nproc) && make install` compile and install the compiler. From hereon, `clang++` directs to our modified compiler version.

Its command-line interface (CLI) is backwards compatible with the mainline Clang/LLVM version. The support for the new attributes discussed in this paper is enabled by default, no additional compiler flags are needed. Yet, the features can be switched off (Section 4).

If the use of any of the new attributes leads to a compilation error, a common troubleshooting starting point is to inspect the rewritten source code. To see the rewritten code, we can add `-fpostprocessing-output-dump` to the compilation flags. This flag causes the post-processed source code be written to the standard output.

## B Download, build and run testbench

All of our code is hosted in a public git repository at <https://gitlab.lrz.de/hpcsoftware/Peano>. Our benchmark scripts are merged into the repository's main, i.e. all results can be reproduced with the main branch. Yet, to use the exact same code version as used for this paper, please switch to the `particles-aos-vs-soa` branch.

The test benchmarks in the present paper are shipped as part of Peano's benchmark suite, which implies that Peano has to be configured and built first. The code base provides CMake and autotools (Alg. 3) bindings. Depending on the target platform, different compiler options have to be used. Once configured, the build (`make`) yields a set of static libraries providing the back-end of our benchmarks.

---

**Algorithm 3** Cloning the repository and setting up the autotools environment.

---

```

1: git clone https://gitlab.lrz.de/hpcsoftware/Peano
2: cd Peano
3: libtoolize; aclocal; autoconf; autoheader; cp src/config.h.in .
4: automake --add-missing
5: ./configure CXX=... CC=... CXXFLAGS=... LDFLAGS=...
6: --enable-loadbalancing --enable-particles --enable-swift
7: --with-multithreading=omp
8: make

```

---

The actual benchmark can be found in Peano's subdirectory **benchmarks/swift2/hydro/aos-vs-soa-kernel-benchmarks**. Within this directory, a sequence of Python commands (Alg. 4) produces the actual kernel benchmark executable **kernel-benchmarks**. The Python script provides further options available through the argument **--help**. Internally, it parses the configuration used to build the static library, creates all glue code required by the benchmark, and then yields a plain Makefile. By default, this Makefile is immediately invoked.

---

**Algorithm 4** Building the benchmark itself.

---

```

1: cd benchmarks/swift2/hydro/aos-vs-soa-kernel-benchmarks
2: export PYTHONPATH=../../python
3: python3 kernel-benchmark.py -d 2

```

---

With the executable at hand, we can run the benchmark and obtain an output similar to

```

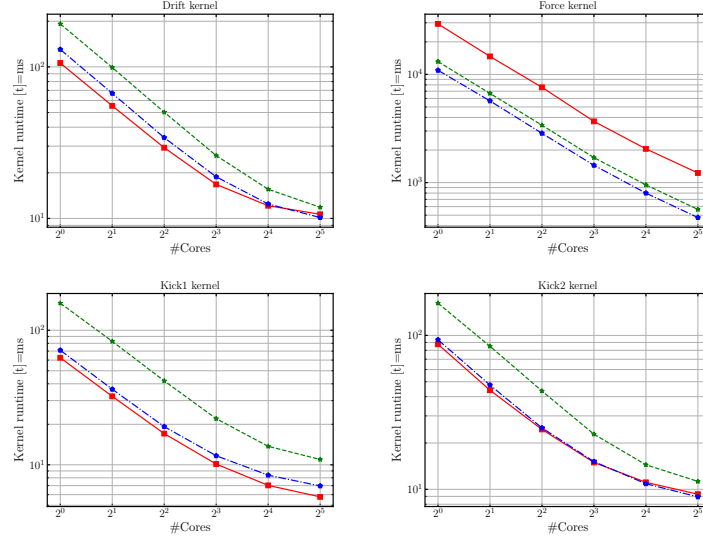
=====
4096 particles (16 samples)
=====
density kernel: 0.802495 (avg=0.802495,#measurements=16,max=1.52808(value #12),min=0.705117(value #4),+90.4159%,-12.1344%,...
force kernel: 0.640393 (avg=0.640393,#measurements=16,max=1.17305(value #13),min=0.576625(value #5),+83.1761%,-9.95752%,...
kick1 kernel: 4.29028e-05 (avg=4.29028e-05,#measurements=16,max=0.000144243(value #12),min=3.4756e-05(value #4),+236.209%,...
kick2 kernel: 7.64441e-05 (avg=7.64441e-05,#measurements=16,max=0.000265726(value #12),min=6.2503e-05(value #4),+247.608%,...
drift kernel: 1.86415e-05 (avg=1.86415e-05,#measurements=16,max=4.9618e-05(value #12),min=1.632e-05(value #4),+166.17%,...

```

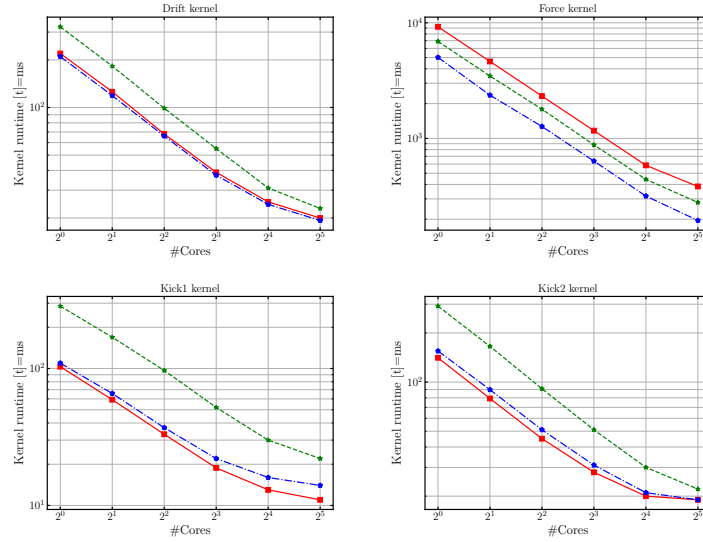
which we can postprocess further.



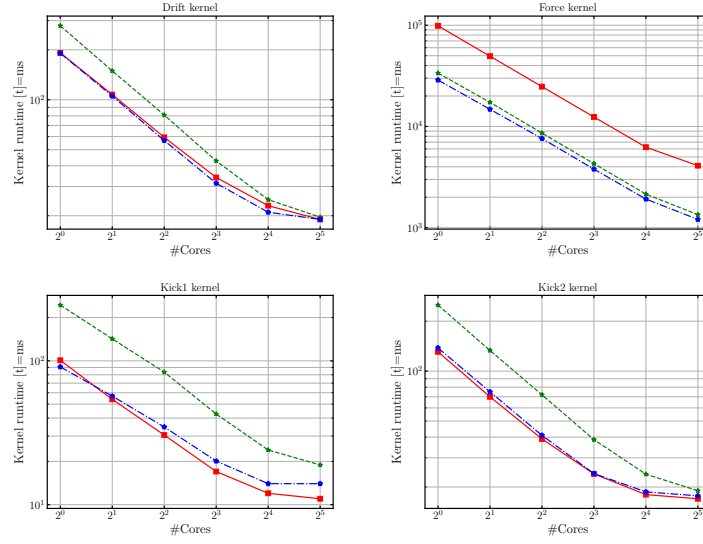
## C Additional scalability tests



Scalability plots for 2048<sup>2</sup> particles, i.e. the same experiment as in Figure 1 on the Sapphire Rapid testbed.

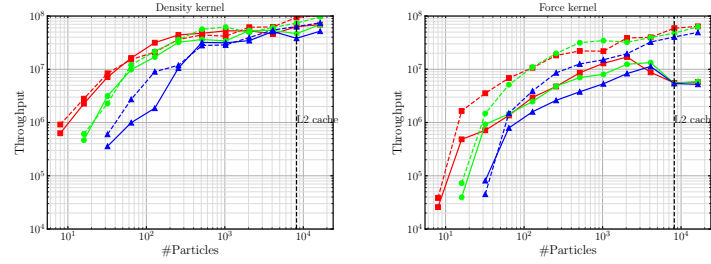


Experiment from Figure 1 for 4 · 2048<sup>2</sup> particles.




Experiment from Figure 1 for  $16 \cdot 2048^2$  particles.

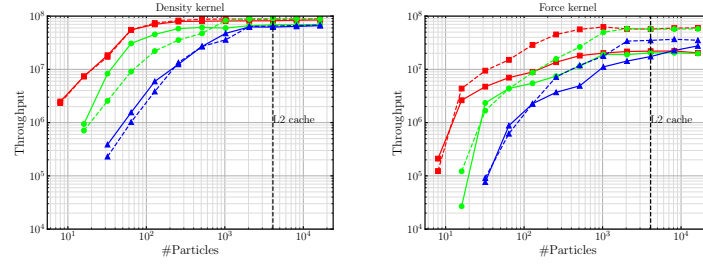
## D Additional throughput data



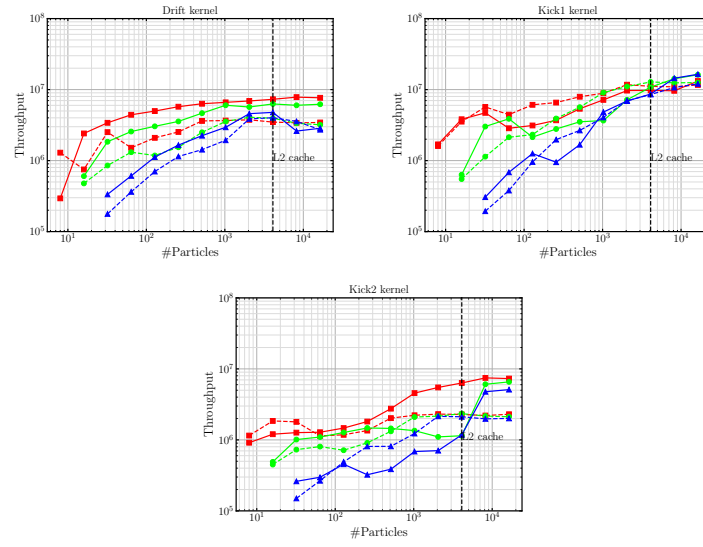
All Sapphire Rapid data for the algorithmic steps with quadratic complexity.



Sapphire Rapid data for the algorithmic steps with linear complexity which are not show in main manuscript.



Genoa data for algorithmic steps with quadratic complexity.



Genoa data for algorithmic steps with linear complexity.



**Citation on deposit:**

Radtke, P., & Weinzierl, T. (2024, September).  
Compiler support for semi-manual AoS-to-SoA  
conversions with data views. Presented at PPAM  
2024 - 15th International Conference on Parallel

Processing & Applied Mathematics, Ostrava, Czech Republic

**For final citation and metadata, visit Durham Research Online URL:**

<https://durham-repository.worktribe.com/output/2582837>

**Copyright Statement:** This accepted manuscript is licensed under the Creative  
Commons Attribution 4.0 licence.

<https://creativecommons.org/licenses/by/4.0/>