

VectorFlow

G. Amadio, A. Gheata, A. G. Rodriguez

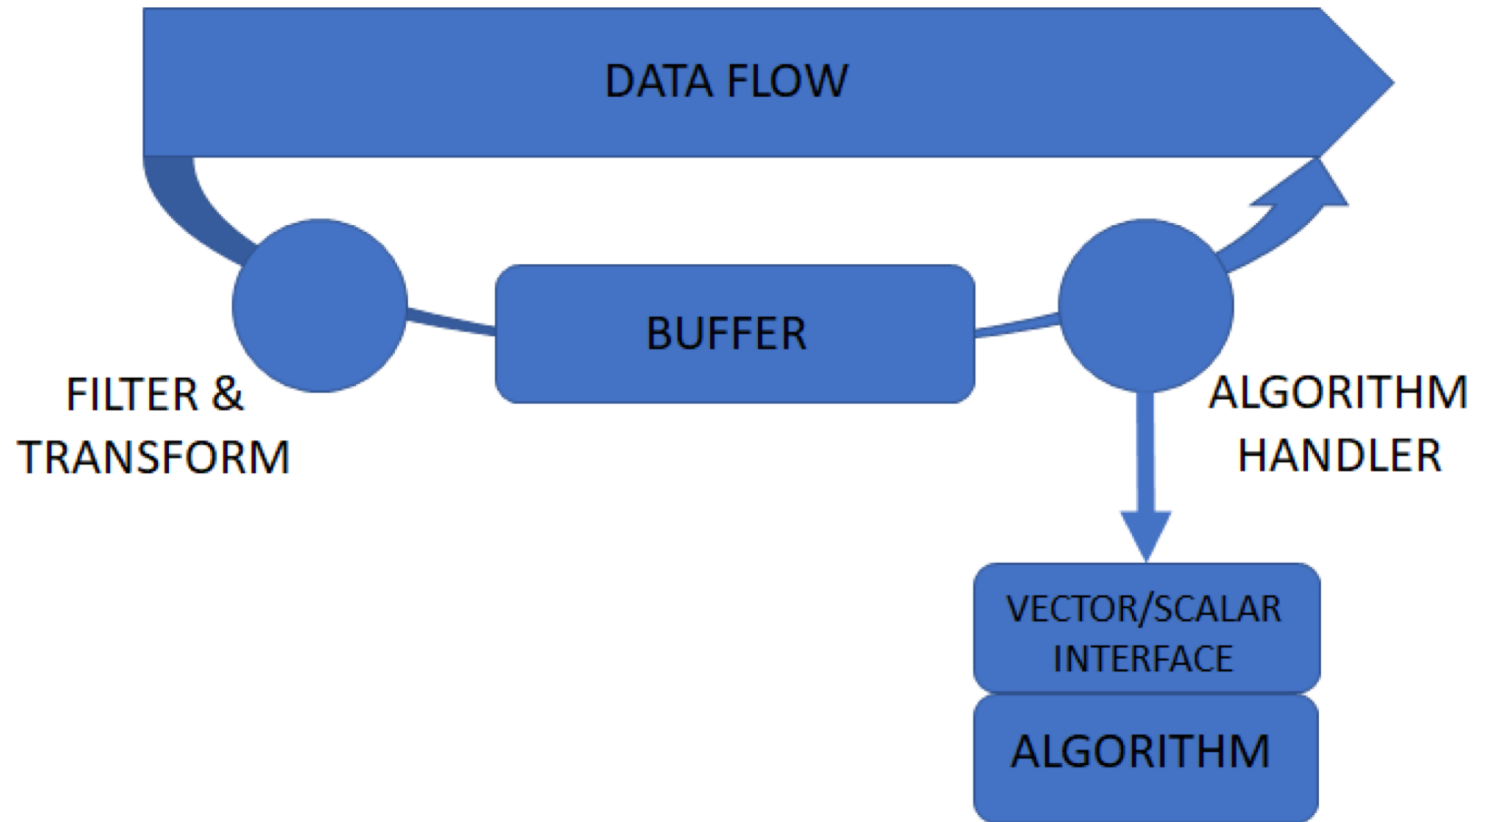
Simulation R&D meeting, Sep 3, 2019

The idea: Integrate vector components in a scalar workflow

Many CPU-intensive high-energy physics algorithms may profit from the vector pipelines of modern processors, they don't because they don't have vectorizable inner loops.

A vector adapter?

- Idea originating from GeantV workflow, but generalized as templated API usable in any workflow
- Do not provide an implementation, but rather a recipe and examples on how to do it
- Using VecCore as underlying vectorization library



Workflow elements


- A ***FILTER*** step to select data of interest.
- An ***ACCUMULATION*** step for the scalar state data used by the computation
- A ***GATHER*** step in aligned memory to prepare the data needed for SIMD processing.
- A ***VECTORIZATION*** step, which integrates with VecCore services.
- A ***SCATTER*** step to re-integrate the output into the framework data flow.

GSoC 2019 project

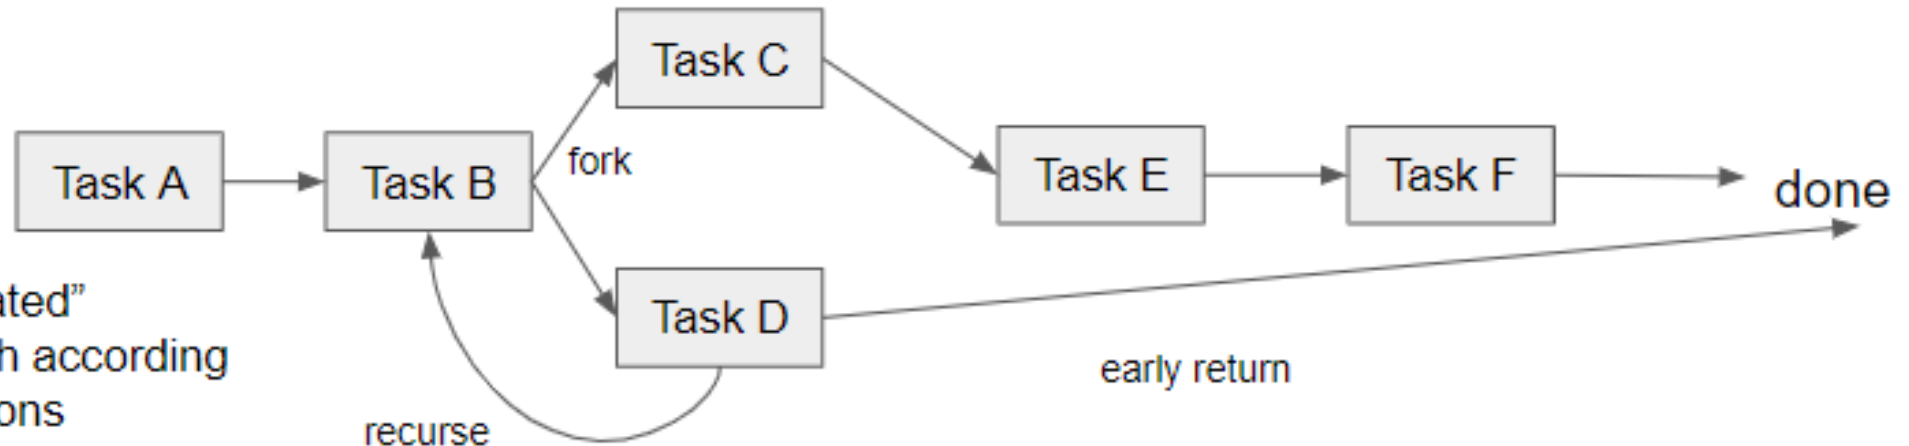
- Implementing VectorFlow interfaces and provide a couple of basic demonstrators
 - Student: Arturo Garza Rodriguez
 - Mentors: G. Amadio and A. Gheata
- New library now available: <https://github.com/agheata/vectorflow>
 - Simple design around the concept of “Work” that can have a scalar and a vectorized implementation
- **In particular the goal was to use this for components in the Geant4 workflow**

Supported workflows

Pipeline flow:



Complex flow:



State is "propagated"
through the graph according
to runtime decisions

- Different flows can be inter-connected
- The data management introduces copy overheads, vectorization has to worth it


Work interface: Base class that provides interfaces to scalar and vector work.

- Interface to arbitrary user algorithm providing abstract interfaces for scalar and vector “Execute” methods
- Templated on user-defined data and container
- Work can have “clients” and can dispatch the processed data to any of those




FLOW interface - Base class that provides support to the different types of flows


- Add each stage to the workflow via *AddWork*.
- Set each stage to be executed either in scalar or vector mode via *SetVectorMode*.
- Two flows to support any variant type of simulation workflows were proposed: the PIPELINE flow & the COMPLEX flow.



Pipeline flow




Complex flow



EXAMPLE

- COMPLEX FLOW: **Generate + Propagate**
- Each layer is translated to a **VectorFlow task** that propagates a **vector** of tracks (previously **gathered** in the correct format for SIMD processing).
- Each particle track can be propagated to an inner or outer layer in the **tube**, i.e. a task **dispatches** data to the other tasks' containers.
- If tracks are propagated outside the geometry, they are **scattered** back to the original data flow.
- The **flow** continues its execution until as long as there is still data in the buffers.



Up to ~2.0x speed-up in AVX2 processor, even though inherent overhead due to data transformations.

Testing VectorFlow with Geant4

- Requires stateless transport engine – ongoing work (WP + AG)
 - Moving the state in the G4Track
 - Done already for managers (WP)
- Pause/resume for a given track w/o overhead
- Writing the vectorflow for the performance-critical components
 - Field propagation, MSC