

CASUS Hands on Software Seminar



LLAMA – A Low Level Abstraction of Memory Access

Bernhard Manfred Gruber, CERN & CASUS

Date: Thursday, 08. October 2020

Time: 10:00 – 12:00 CET

Location: CASUS Lecture Room, Görlitz

Abstract:

Current and future hardware architectures are heterogeneous. Performance portability with a single C++ code base is an increasingly relevant challenge.

Performance portable parallelism to exhaust multi-, manycore and GPU systems is largely addressed, e.g. Alpaka or Kokkos.

Efficient use of a system's memory and cache hierarchies is equally crucial. First attempts like AoS/SoA containers and `std::mdspan` exist, but general solutions are lacking.

LLAMA attempts to fill this gap. It provides an API to define and access a data structure. Underneath, a memory layout infrastructure allows to choose between layout mappings at compile time or even define mappings yourself. This way, LLAMA enables abstraction of memory access patterns without the need to change your code, thus providing a foundation for performance portable memory access.