

ExaHyPE: Code updates

- **Updated results**

- **Improvements and GPU porting**

- Optimisation of grid traversal bottlenecks as identified in second code audit (CPU stalls; GPU underutilized)
 - Addition of a Kokkos backend supporting the major vendors
 - Improved GPU reduction with SYCL; also supported by AdaptiveCpp **and** Intel oneAPI

- **Mixed precision (WP3)**

- **Use (integration with code/workflow) of TALP in performance monitoring activities** No



ExaHyPE: Codes updates to come

- **Mandatory section**

- **please, specify if you still have additional effort in WP2 for your code**
 - Development project on Leonardo is expired
 - Planned to request new 1-year project

- **Planned code updates**

- **Improvements and GPU porting**
 - Avoiding GPU-CPU memory transfers by introducing memory pools for ExaHyPE grid entities: AoS -> SoA (first prototype is implemented and validated)
 - Improving parallel efficiency of grid traversals by eliminating memory copies between subdomains in shared memory
 - Mixed precision on GPUs
 - Overlapping of grid traversal and compute kernels
- **Use (integration with code/workflow) of TALP in performance monitoring activities** No activities yet planned

