



Improving Vector Performance with Vector Advisor

Sergi Siso

Application Performance Engineer

IPCC@Hartree, STFC Daresbury Laboratory

United Kingdom

DL_MESO

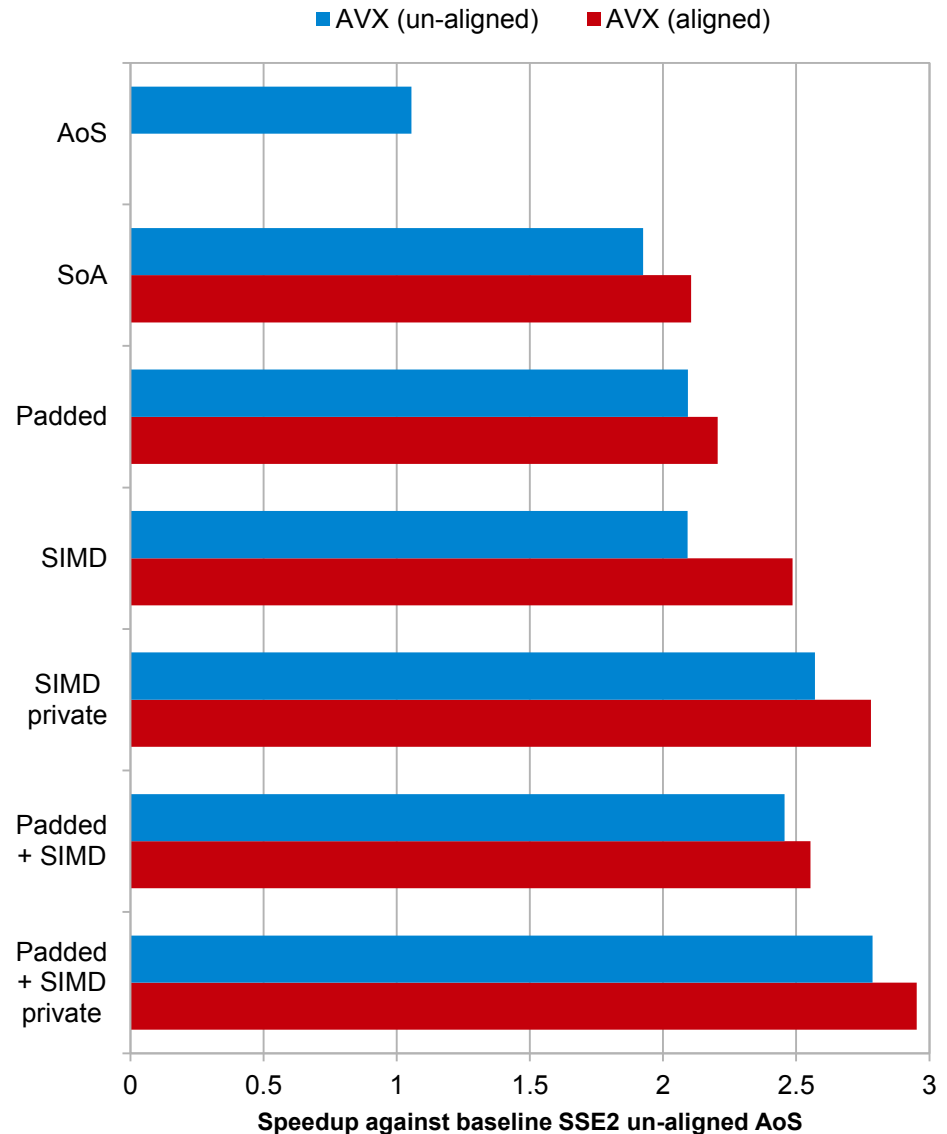
- DL_MESO_LBE: Lattice Boltzmann code
- Models fluids at mesoscale: Hybrid MPI/OpenMP
- Vector Advisor used to analyze loop performance
 - fGetEquilibriumF identified as hotspot
- Baseline characteristics
 - Double precision
 - Array of Structures (AoS) data storage.
 - No data alignment (prevented by data structure).
 - Low trip count loop, not a multiple of vector lengths.
 - Both Peel and remainder loops present.



Performance - fGetEquilibriumF

- V-Advisor recommendations
 - AVX not enabled by default
 - MAP analysis points to AoS -> SoA.
 - Remove Scalar remainders.
 - Align data accesses.
- SoA allowed aligned access and removing peel loops.
- Array padding and #pragma loop count removes remainder loops.
- Additional optimization
 - #pragma SIMD
 - Private SIMD clause allowed additional compiler optimizations.
- Xeon speed up x2.95
- Phi speed up x4.05

Performance of Loop Optimisations



Insights

- Data structure is key to effective vectorization.
- Manual array padding required for Xeon Processors to avoid remainder loops.
- `#pragma SIMD` out performs auto vectorization.
- Vector Advisor enhances existing optimization reports by
 - Easy to filter/browse presentation of loop timing data.
 - Optimization report data embedded in report.
 - Provides direct links to the relevant source code.
 - Useful recommendations.
 - Check possible dependencies and memory access patterns at runtime.
- Need Advisor tool for MIC architecture!



Future work

- Investigate effects of vectorization at different “levels” of the LBE calculation.
 - Currently at inner most loops (lattices)
 - Could vectorize the calculation over lattice points however
- Memory optimizations still possible
- Improve the hybrid OpenMP/MPI parallelization

