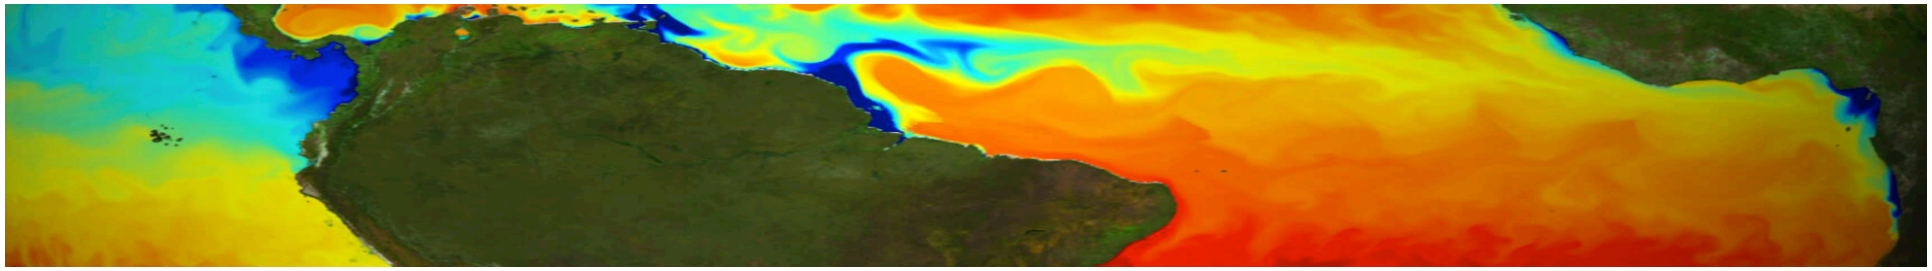

High Performance Computing
for

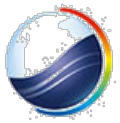

Mercator Ocean
Ocean Forecasters

cmcc
Centro Euro-Mediterraneo sui Cambiamenti Climatici

Science & Technology Facilities Council

UNIVERSITY OF EXETER

National Oceanography Centre
NATURAL ENVIRONMENT RESEARCH COUNCIL

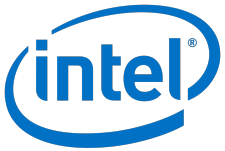

Atos

2017 Work done

Intra-node performance

- Real vs peak performance
  - Investigation of a benchmark configuration on different HPC systems (Silvia, Tim & Martin, Maff)
  - Vectorization improvement
    - Use of SIMD directives → execution time gain **~3-4%** (Cyril)
    - Improvement of memory allocation strategy → up to **~20%** faster on high-resolution configurations on CRAY system (Tim)
- Hybrid parallelization
  - OpenMP shared memory approach (Silvia)
    - Fine grained (loop-level) approach → execution time gain **~5-6%**
    - Coarse grained approach (tested on vertical physics) → intra-node parallel efficiency gain **~9%** (Silvia & Gurvan)

I/O performance improvement

- Reading/writing restart files with XIOS is working but it is not fast → further development by XIOS team needed (Meric)

2018 Workplan

Intra-node performance

- Real vs peak performance (Silvia)
  - Memory access overhead
    - Cache blocking → need to automatically set the best cache block size

- Hybrid parallelization (Silvia)
  - OpenMP shared memory approach:
    - Comparison between the coarse-grained approach and the cache blocking technique
    - Investigation of the coarse-grained approach on kernels more affected by the communication overhead

Scalability improvement

- Reducing the communication overhead
  - Decreasing the communication frequency (Andrew)
  - Reducing collective communications (started by Tim, Silvia)
  - Overlapping inter-processes exchange (Silvia)

Other HPC activities

Performance portability on heterogeneous architectures

- Lightweight DSL approach (Andy, Silvia)
  - Parsing the NEMO code to make it PSyclone compliant
  - Automatic optimizations (also including OpenMP and OpenACC directives) to run the code on heterogeneous architectures

Reduced precision (Miguel, Oriol)

- Investigation on which parts of the code could be safely integrated in single precision without having an impact on the model accuracy