Resiliency for Exascale Climate Modelling

An exploratory work on NEMO ocean model

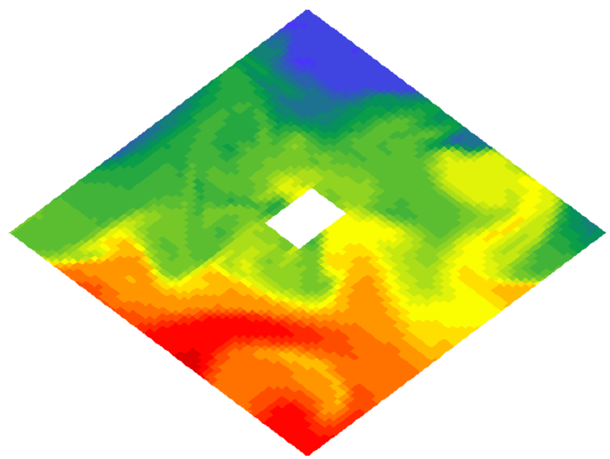

E. Maisonnave (CERFACS), V. Slavnic (IPB)

Rationale

Vector to massively parallel transition for climate modelling bothered by code complexity: adaptation rather than rewriting

Importance to include now in new developments compliant with fault tolerance

Necessary to make assumptions on exascale machine characteristics, kind of failure, Middleware (MPI) adaptation

Code modification: first test on NEMO ocean (high parallelism, near to exascale code structure)

Assumptions

Exascale computer features	
Million to Trillion of cores	
Expensive communication cost	
Expensive disk access	
M.T.B.Failures reduced	

Assumptions

Exascale computer features	Climate models on Exascale
Million to Trillion of cores	Higher parallelism > higher resolution > selected studies (seasonal, decadal ?)
Expensive communication cost	No collective exchanges
Expensive disk access	No frequent checkpointing
M.T.B.Failures reduced	Fault tolerant MPI + Fault tolerant model algorithm

NEMO present parallelism

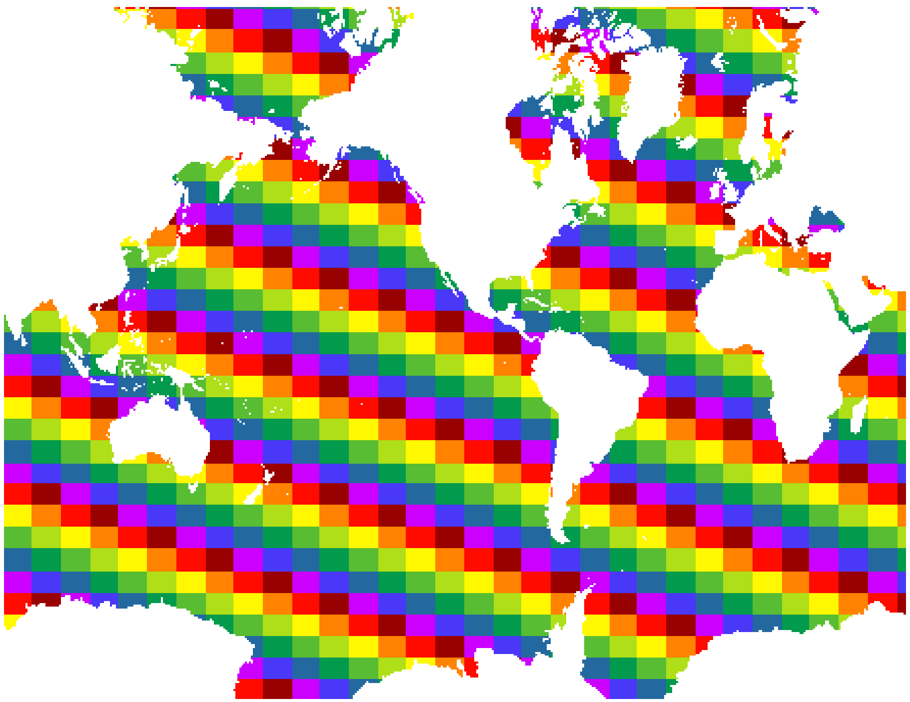

1000 sub−domains decomposition

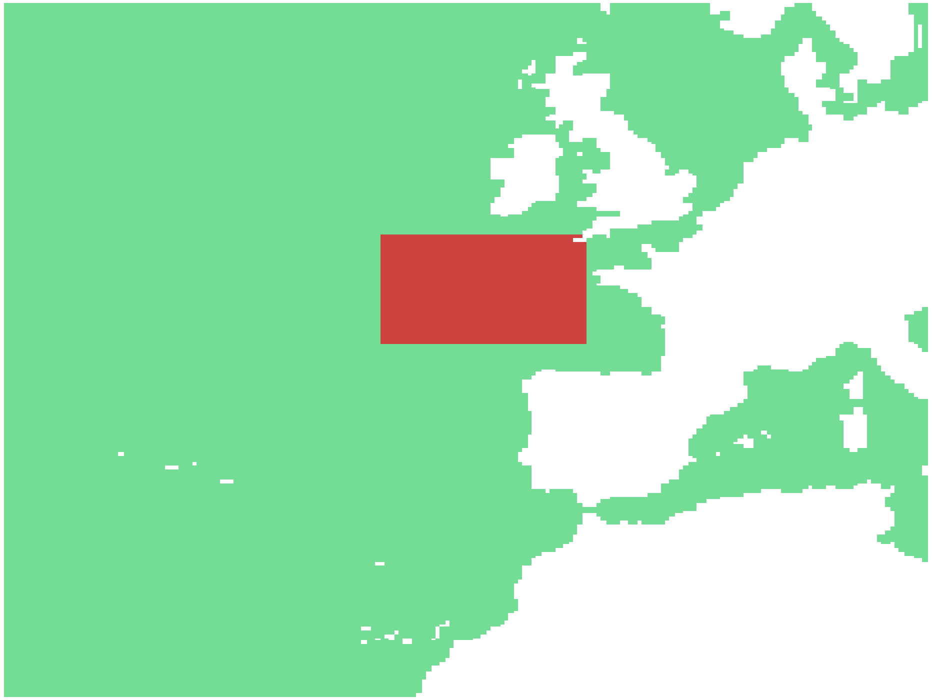

ORCA025 – 1 of 1000 sub_domains

NEMO exascale parallelism ?

Sub-domain area << global domain:
We suppose that to approximated calculations on one subdomain leads to

an error << a model perturbation

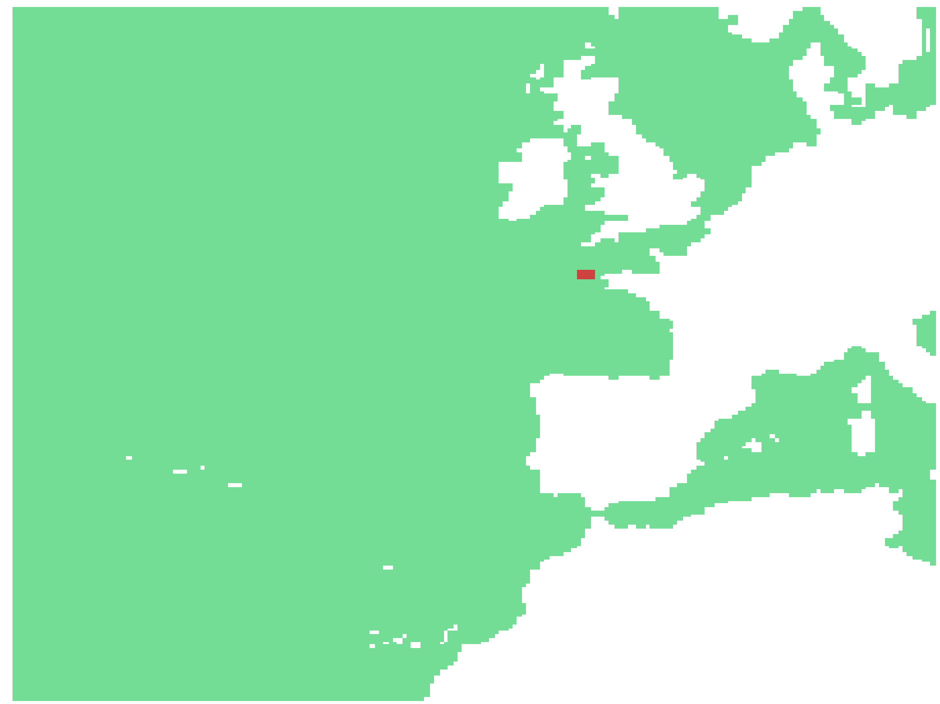

1 of 100.000 sub_domains

NEMO parallelism

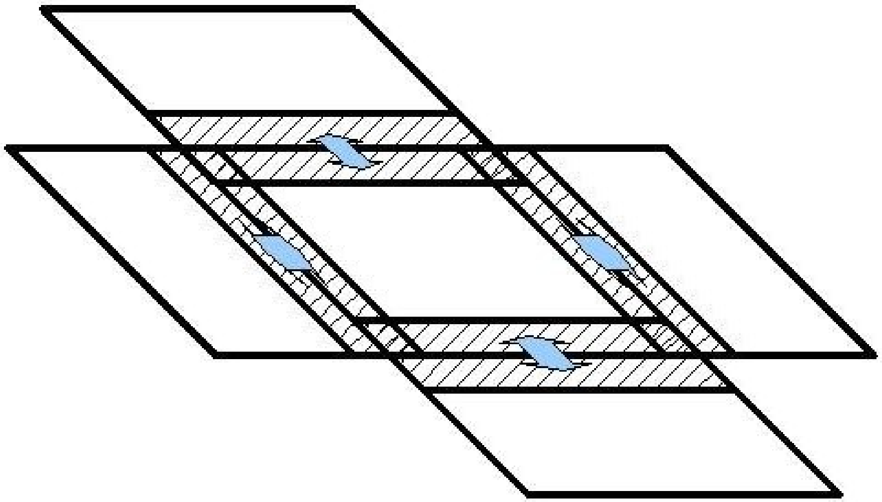

Variables exchanged from neighbour subdomain to halos through MPI communications

NEMO parallelism with fault tolerance

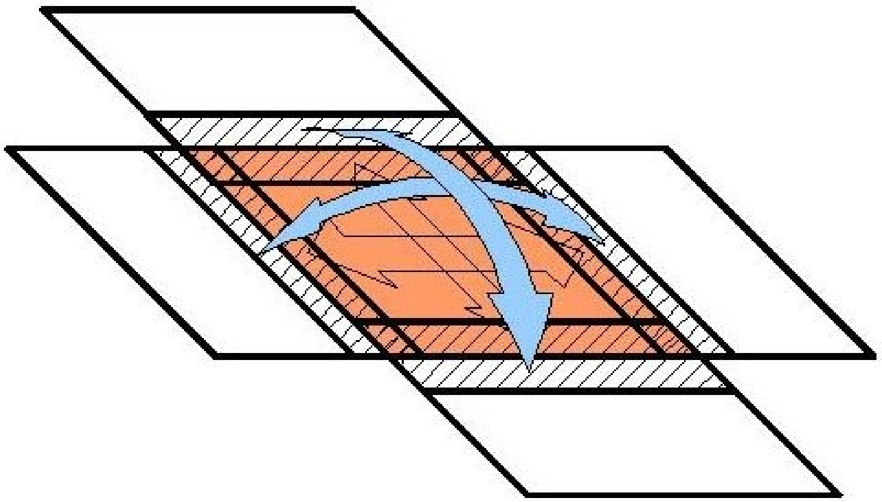

Repair during simulation:

(1) Model error handling: redefine MPI communications
to exchange boundary conditions through missing sub-domain
(2) Downgraded calculations at redirected boundaries
until next checkpoint (< 1 hour)

NEMO parallelism with fault tolerance

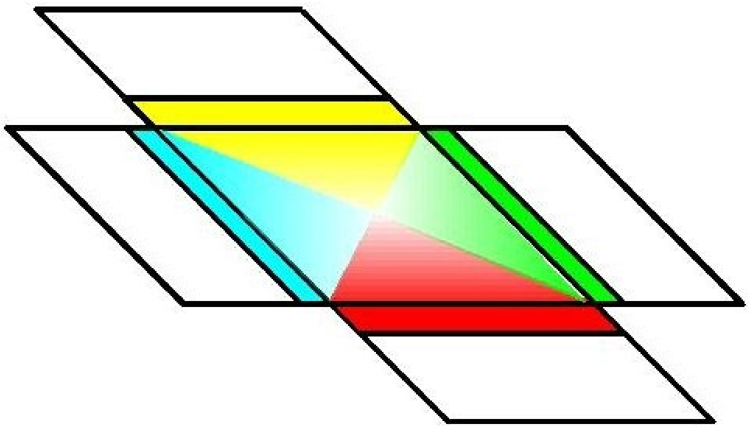

Repair after regular checkpointing:

(3) missing grid point values are extrapolated from values of grid point neighbors.
Simulation can be resumed on all sub-domains

Test case

NEMO GYRE 1/12th degree
1,000 sub-domains on 1,000 cores
Simulates 1,000,000 sub domains on 1,000 nodes of 1,000 cores
Regional grid: simulates a 100 times greater domain (global)
100,000,000 sub domains on 100,000 nodes of 1,000 cores (Exascale)

Starting point: initial conditions from a previous standard spin up (10 years)

Three 10-year long simulations:
reference (without failure) + 2 simulations with 2 different failure rates
1 sub domain is switched off (simulated 1 node failure):
a) every month (100 simultaneous failures on global grid)
b) every 6 months (15 simultaneous failures on global grid)

Failure occurs at very beginning of the 1 month long simulation:
<u>failure rate is voluntarily over estimated in both case</u>

Results

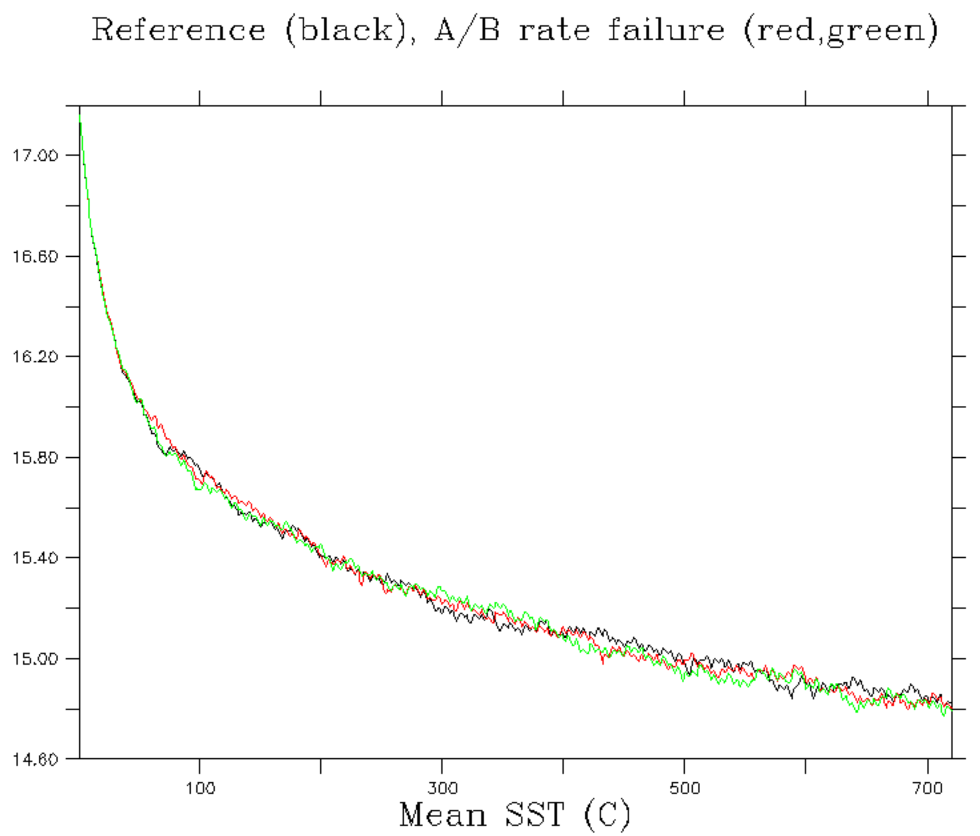

Sea surface temperature: same drift

Results

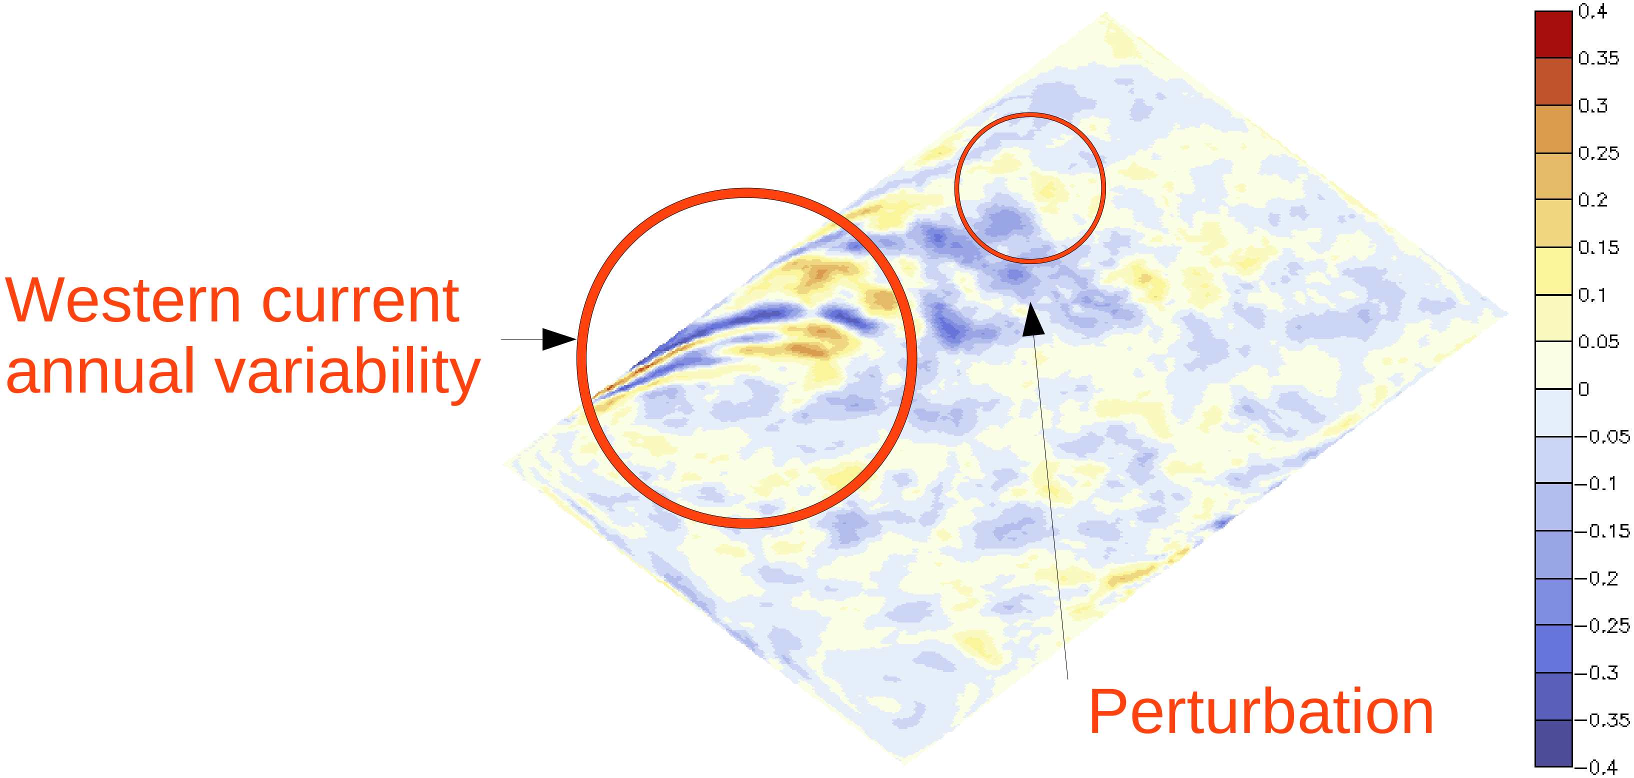

SST Std deviation (B failure rate – Reference)

Sea surface temperature: annual or weekly perturbation lower than simulated variability

Results

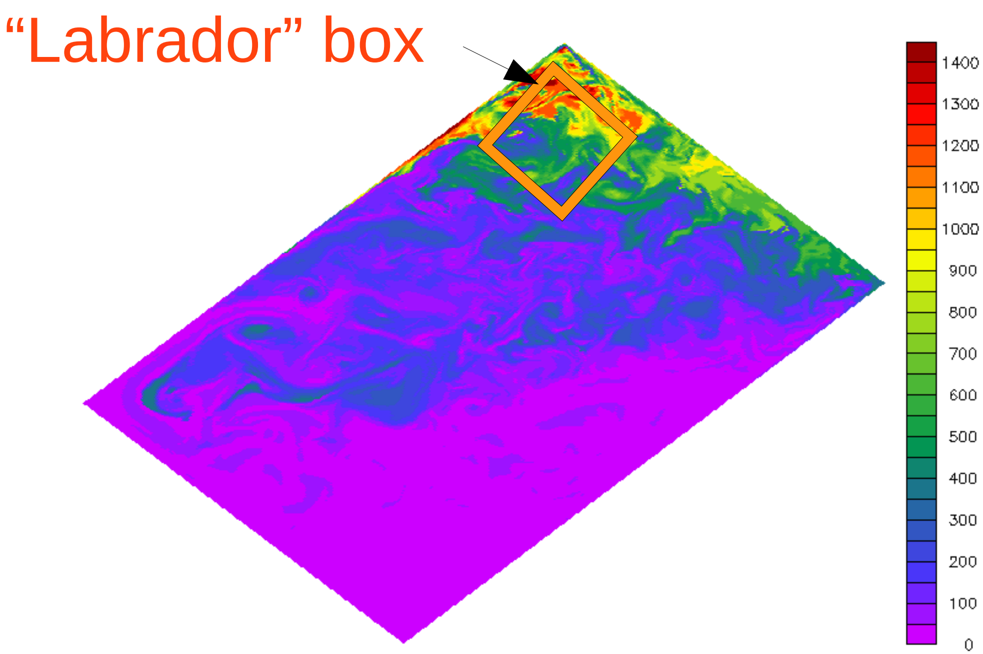

Mixed layer depth (m), exemple of 5d mean, B-rate failure

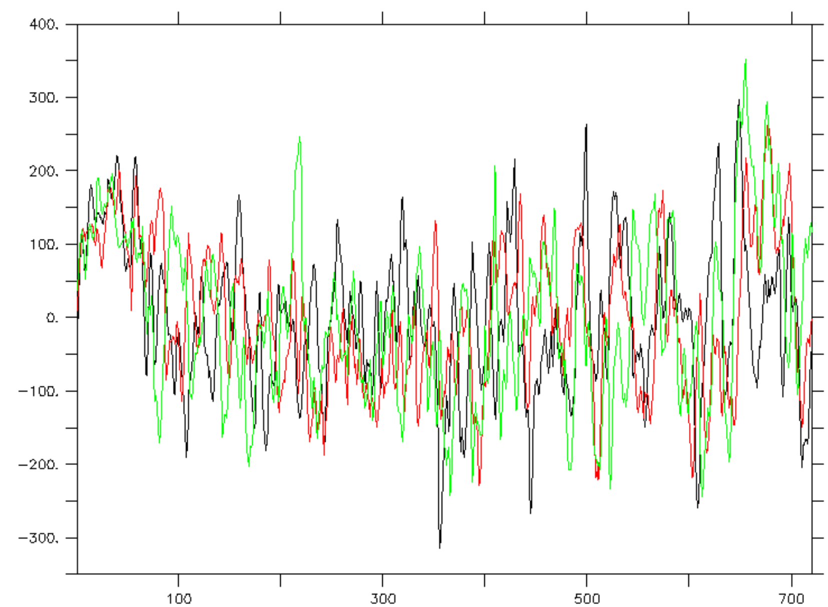

Mixed layer depth: series on anomalies for reference and A/B failure rate simulations

Results

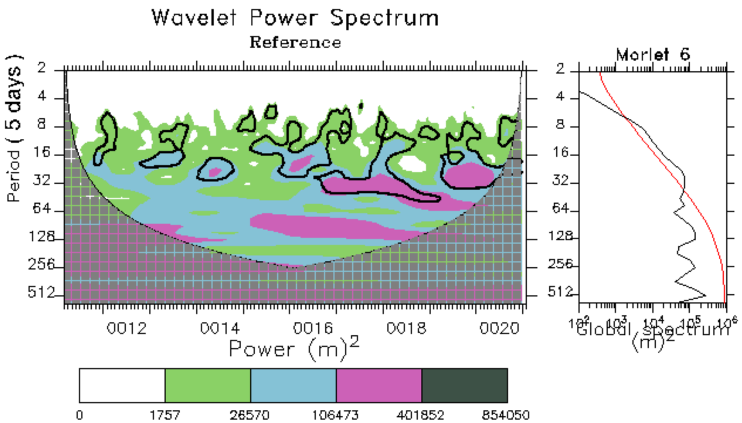

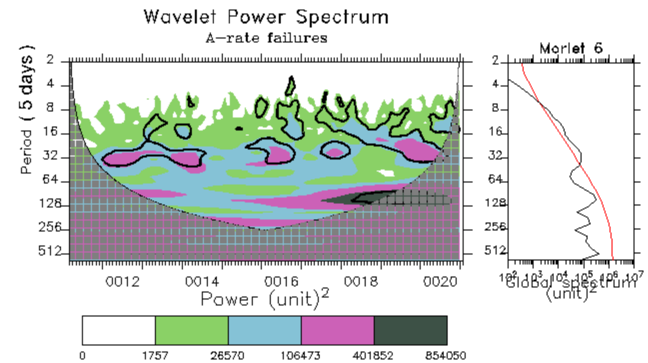

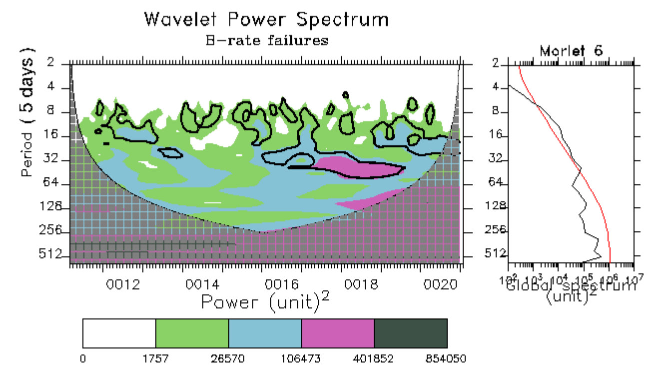

Wavelet analysis (5 days unit), Reference and A/B rate failures
No major differences (1.5 year ?)

Results

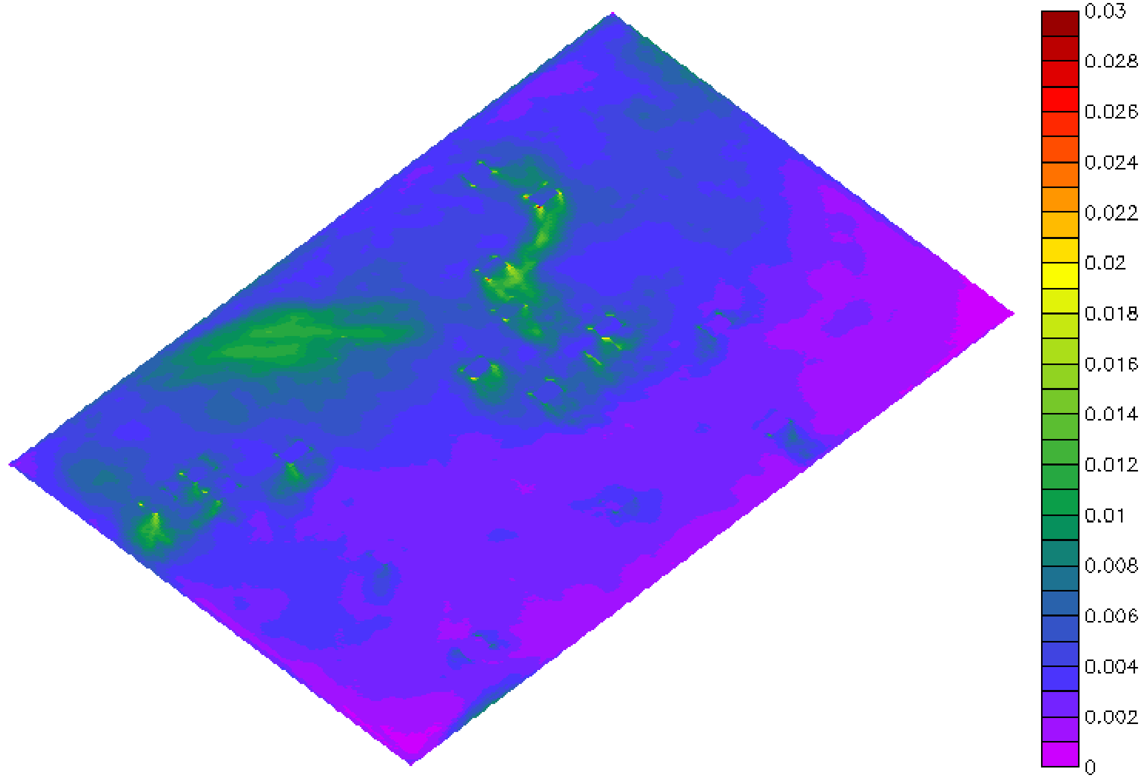

Std dev of Turbulent Kinetic Energy, log(cm2/s2): Reference/ A-rate failure simulations

Perturbation higher than simulated variability but briefly absorbed by turbulence (see movie)
No statistic significance (strong event but local, temporary and not recurrent)
Simulated Western boundary current unchanged

Conclusion

Implementation done to realize how difficult is error handling of fault tolerance

- Simple solution leads to 10 year long fault tolerant simulations
- No significant changes in variability or mean state (SST, mixed layer depth, TKE)
- Dedicated to specific configuration (no reproducibility, no conservativeness)

Now:

- To implement on Exascale compliant models (new dynamical cores ?)
- Better on-line repair algorithm (conservativeness ?)
- Better checkpointing repair (to reduce gradients)
- Face bathymetry issues (ORCA12 grid instead of GYRE)
- Interface to MPI3 (fault tolerant version)
- Fault tolerance on atmosphere model and coupler to build an FT-climate model