Inside NEMO v4.0

ANDREW C. COWARD
NEMO SYSTEM TEAM

National Oceanography Centre
NATURAL ENVIRONMENT RESEARCH COUNCIL

noc.ac.uk

NERC SCIENCE OF THE ENVIRONMENT

NEMO
(200,000 lines not in C)

- The development strategy
- Why do we need a strategy?
- Recent changes and some examples
- Key changes from the user's perspective

https://www.nemo-ocean.eu

NEMO development strategy (summary)

- Develop a new, efficient and scalable NEMO reference code with improved performances adapted to exploit future HPC technologies;

Why do we need a strategy?

Mobilis (NOC)

Nominal end of life is 31 March 2019

72 nodes / 1152 cores - Xeon E5-2650v2 2.6 GHz, 8 core (Ivy Bridge EP) (similar to ARCHER)

Iridis5 University of Southampton HPC

Ranked 354 in Top500 HPC systems worldwide

464 nodes, each with 40 cores (2 x 20 core Xeon Gold 6138 CPUs (Skylake)), total 18,560 cores

Core-for-core performance should be 46% faster than Mobilis

How well does it run NEMO?

NEMO benchmarks

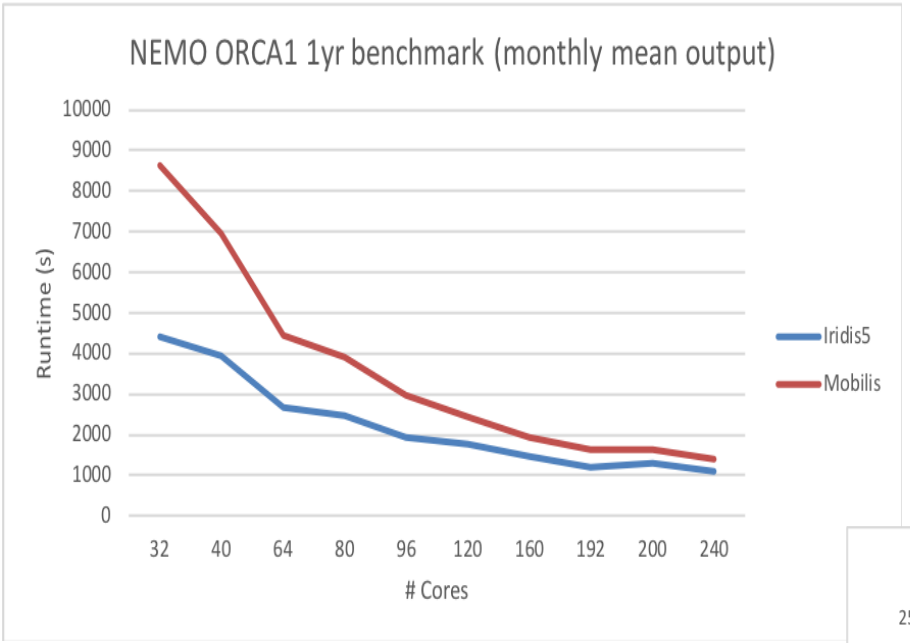

ORCA1:
Performs well at very low core count
Around 30% performance gain at typical job size

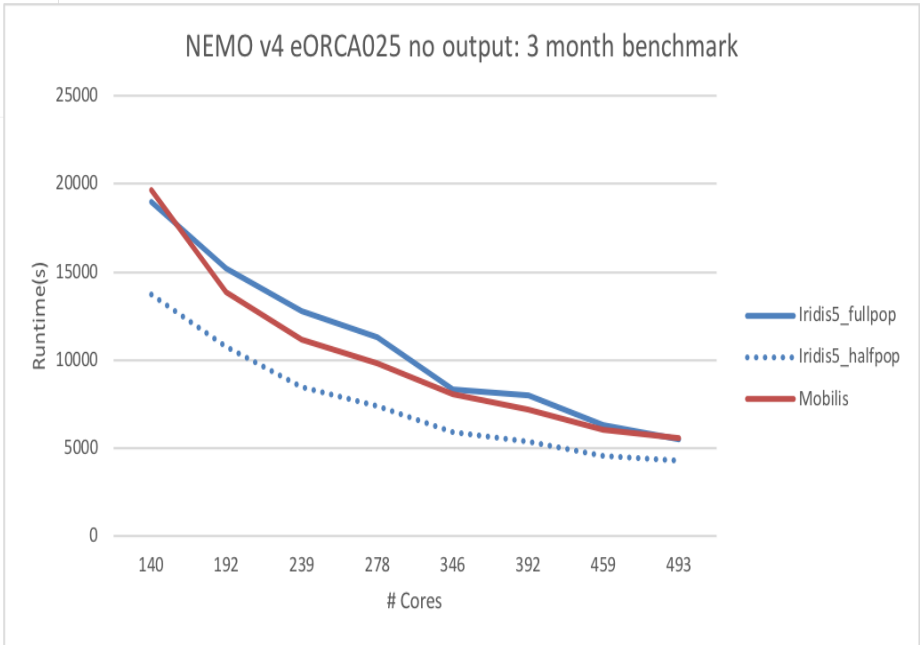

eORCA025:
8% slower than Mobilis
30% performance gain if nodes are half-populated

What have we done during the last year?: 1. The user's perspective:

What have we done during the last year?: 2. The system team perspective:

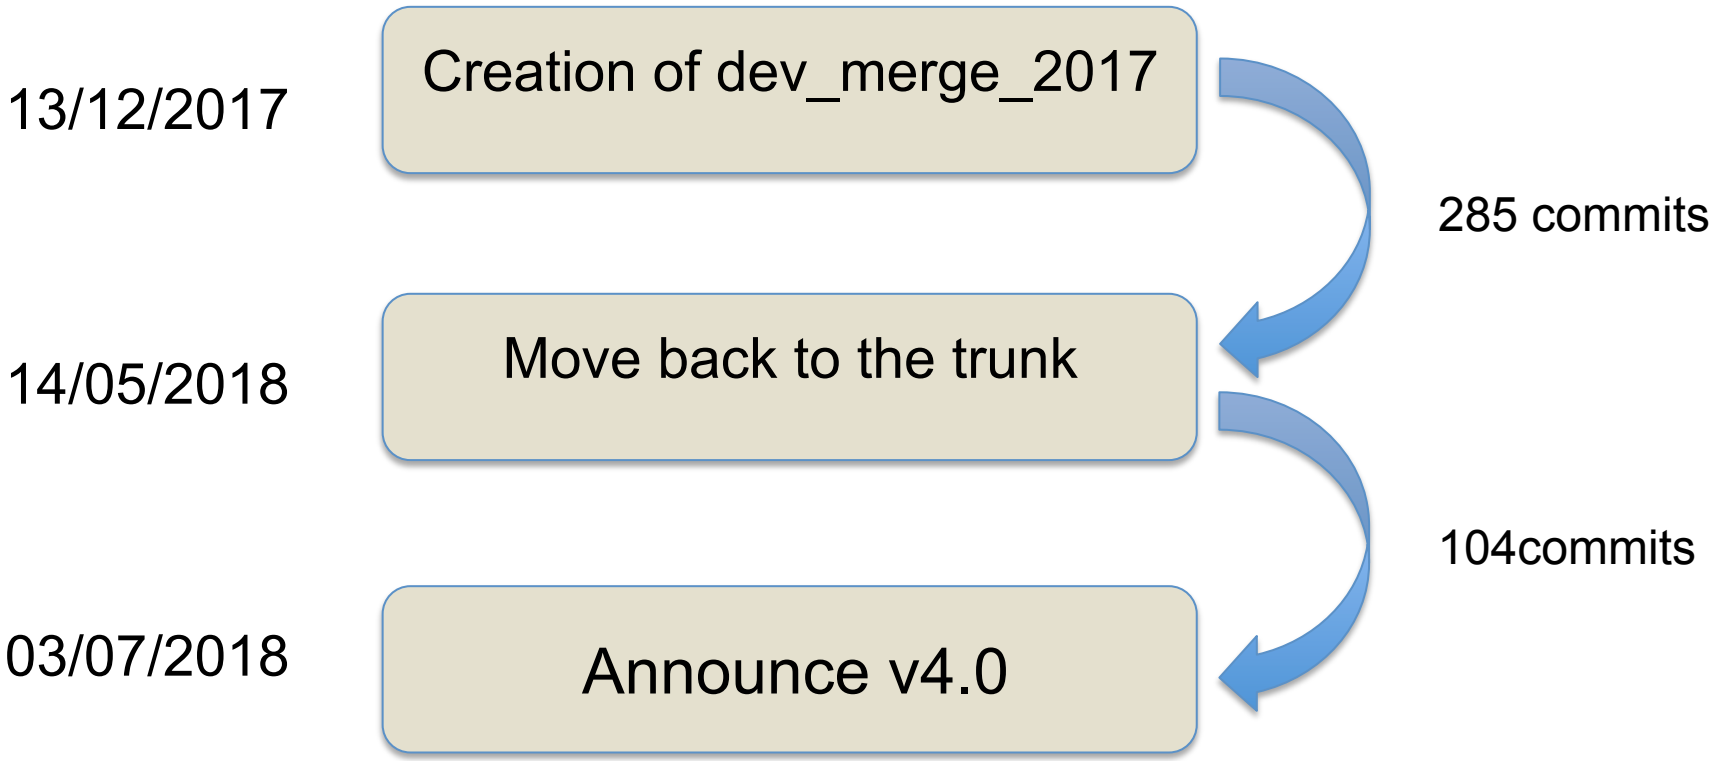

So what has changed?

AGRIF for embedded zooms:

- Now compatible with new sea ice component
- Now compatible with z* coordinate
- Extended ghost cells area to properly handle scheme with spatial order >2
- Added vertical refinement (beta)
- Nesting tools for setup now up to date and working

Enhancements

- Fix for tracer conservation with split explicit free surface
- Bulk formulae : move to aerobulk package (Brodeau et al. 2016), i.e. NCAR, COARE and ECMWF bulk (remove Clio and MFS bulk)
- Wetting and drying
- Added tidal self attraction and loading either read from a file or from usual "scalar" approximation
- Add a 4th order centered (CEN) and Flux Corrected Transport (FCT) tracer advection (using a 4th compact in the vertical)
- iso-neutral mixing (iso and triad operators): add the Method of Stabilizing Correction (MSC) (more accurate calculation) + add a bilaplacian case
- Lateral physics (LDF): scale aware setting of eddy viscosity and diffusivity
- Vorticity: 2 new energy conserving scheme: ENT with Coriolis defined at T-point (better for Flux form) and EET a variant of EEN where e3t is used instead of e3f (solved the issue with e3f specification but is not enstrophy conserving)

Wave coupling

- Coupled interface to external wave model
- Large scale wave interaction process added in momentum and tracer equations

New sea-ice component SI3 (in place of LIMx)

- Improvements in physics (Landfast ice, Lateral melting, Melt ponds etc.)
- All thermodynamics in 1D
- Reduced mpp communications
- Fully compatible with AGRIF

Passive tracer TOP and biogeochemical PISCES components

- The passive tracers transport component was redesigned (greater modularity)

High Performance Computing: performances improvements

- Reduce number of MPI communications (suppression of redundant communications, gather multiple communications into one)
- Back to standard dynamical allocation (remove of wrk_alloc/dealloc statements)
- XIOS software for IOs version 2 as default, and optionally available for restarts

Passive tracer TOP and biogeochemical PISCES components

TOP **(Tracers in the Ocean Paradigm)** is the NEMO hardwired interface for biogeochemical models
The new interface has a modular structure (handled via logical keys):

TRP	Interface to NEMO physical core for computing tracers transport
CFC	Inert carbon tracers (CFC11,CFC12, SF6)
C14	Radiocarbon passive tracer
AGE	Water age tracking
MY_TRC	Template for user-def tracers and external BGC models coupling
PISCES	Built in BGC model demonstrator (see Aumont et al. ,2015))

In addition :

- Full handle on Boundary conditions for tracers (surface, coastal, lateral via BDY)
- Improve parallel execution of transport module and handling of options

http://forge.ipsl.jussieu.fr/nemo/wiki/WorkingGroups/TOP/TOP-UserQuickGuide

Simplification and robustness

- Revised structure of namelist_ref/_cfg and default reference values.
- Lateral physics (LDF): simplification of user interface and removal of CPP keys
- Vertical physics (ZDF) (modularity, share shear production calculation between TKE and GKS, removal of all ZDF CPP keys, removal of avmu & avmv
- Remove the split-explicit ZDF scheme for both TRA and DYN
- Remove the acceleration of convergence
- Generalised lbc_lnk and lbc_nfd
- Unify mppini
- Use non uniform jpi/jpj with dynamic allocation to avoid ghost rows/columns
- MPI Message passing re coded
- Configuration interface completely rewritten (DOM module mainly suppressed , and in place: domain_cfg.nc file, or usr_def module)

New Reference configurations

- AGRIF_DEMO: 2 interlocked zooms (1:4 & 1:3) in the Nordic Seas + 1 zoom (1:1) at the equator
- SPITZ12: regional configuration around the Svalbard

- SPITZ12: regional configuration around the Svalbard

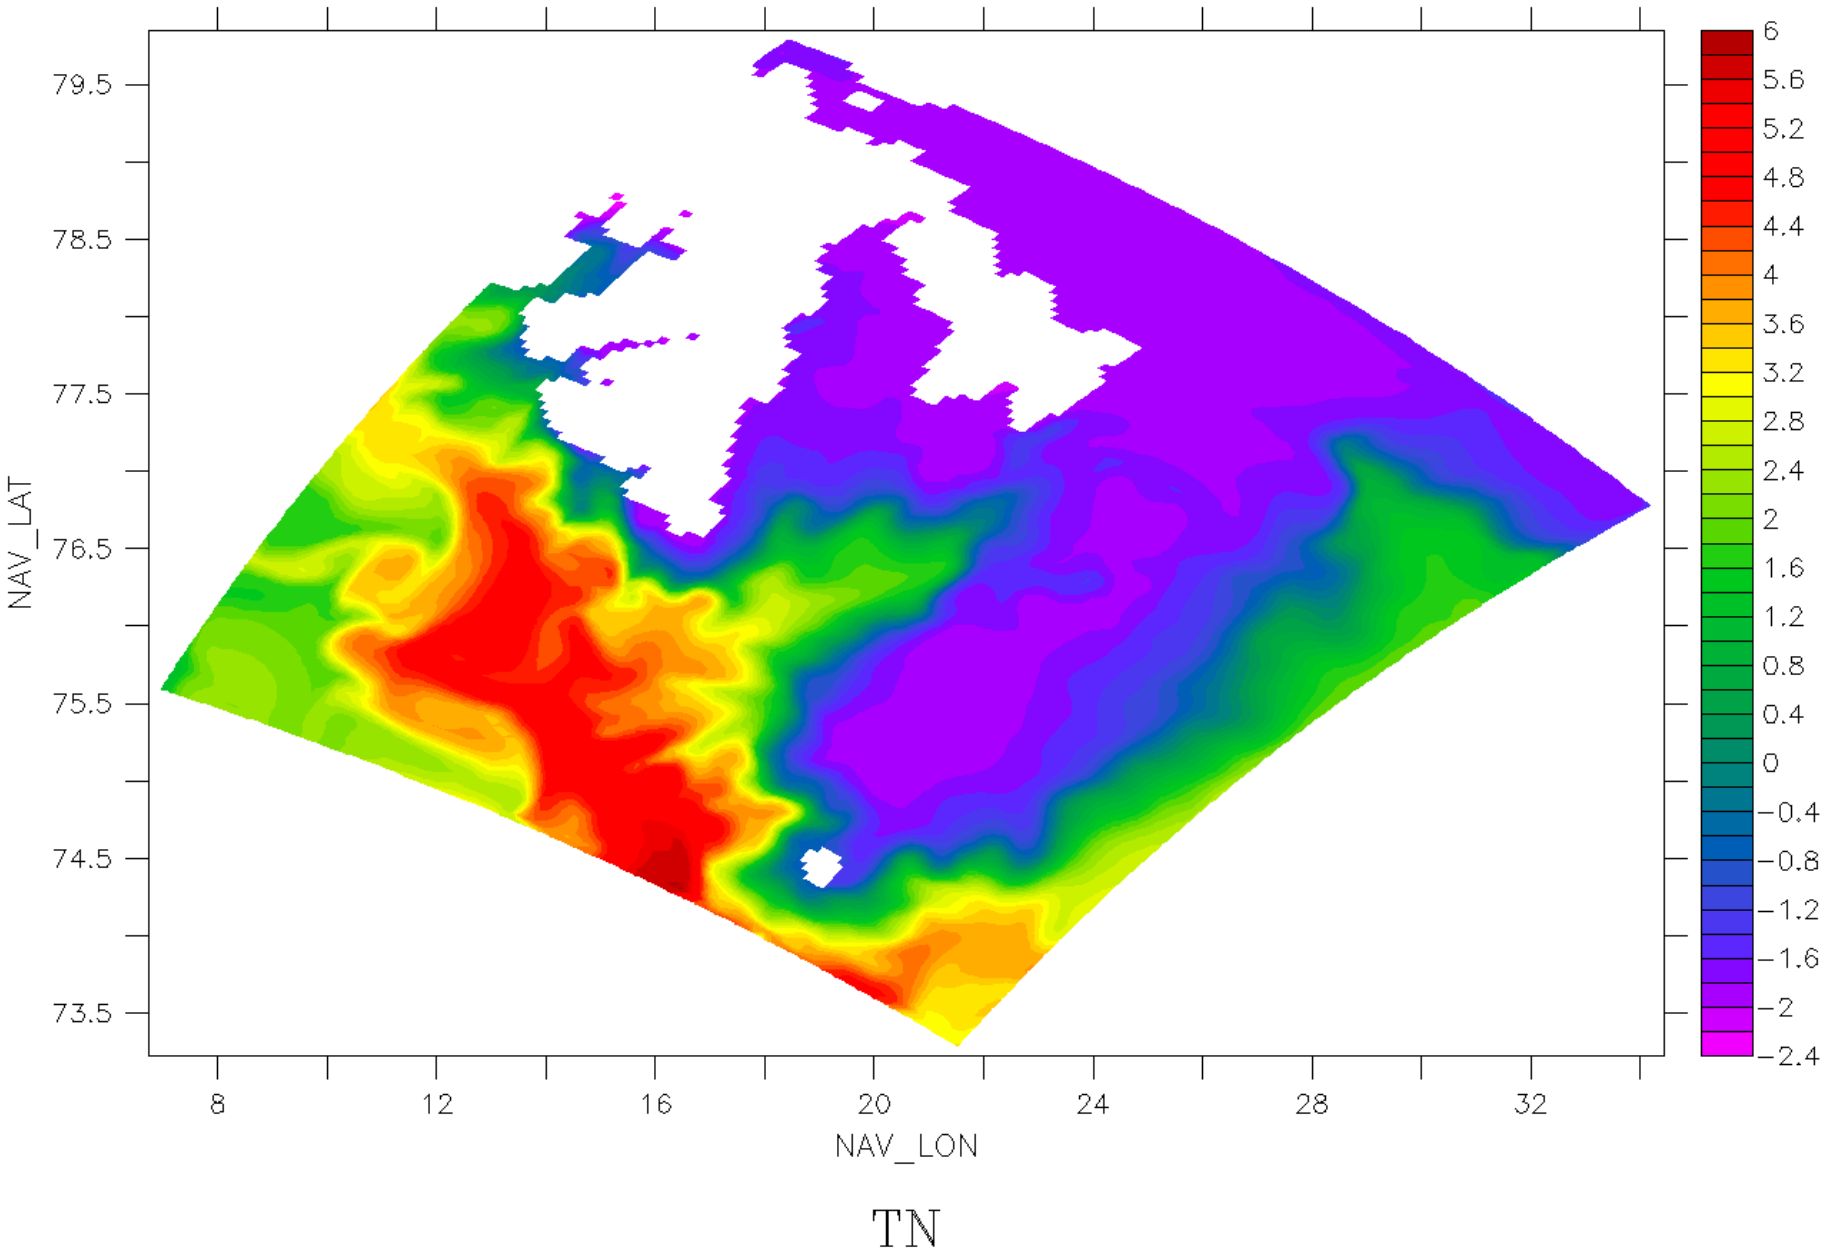

- AGRIF_DEMO: 2 interlocked zooms (1:4 & 1:3) in the Nordic Seas + 1 zoom (1:1) at the equator

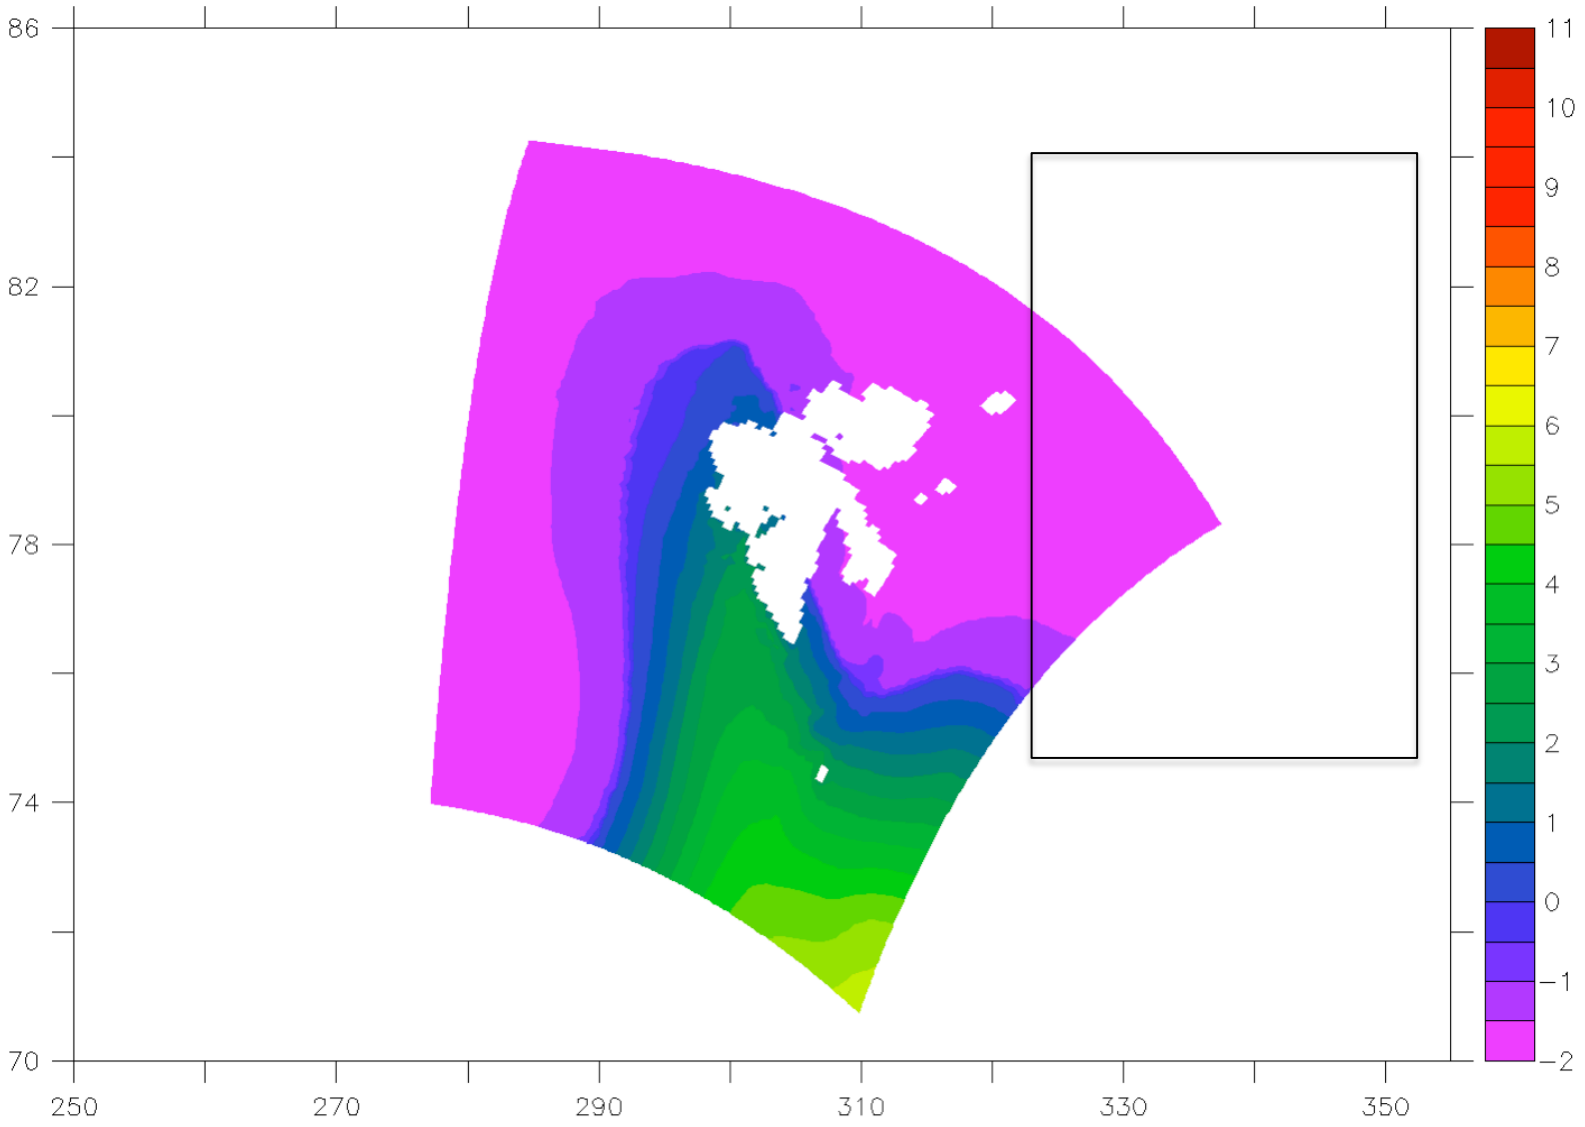

Test Cases (new facility, in addition to reference configurations)

- CANAL: east-west periodic canal of variable size with several initial states and associated geostrophic currents (zonal jets or vortex)
- ICEDYN: east-west + north-south periodic channel. The common configuration includes an AGRIF zoom (1:3) in the middle of the basin to test how an ice patch is advected through it
- ISOMIP: simple box configuration with an ice shelf with simple geometry on top.
- LOCK-EXCHANGE: classical fluid dynamics experiment that has been adapted by Haidvogel and Beckmann (1999) for testing advection schemes in ocean circulation models. It has been used by several authors including Burchard and Bolding (2002) and Ilıcak et al. (2012). Used to illustrate the impact of different choices of numerical schemes and/or subgrid closures on spurious interior mixing
- OVERFLOW Adapted from the non-rotating overflow configuration described in Haidvogel and Beckmann (1999) and further used by Ilıcak et al. (2012)
- VORTEX: illustrates the propagation of an anticyclonic eddy over a Beta plan and flat bottom. It is implemented here with an online refined subdomain (thanks to AGRIF library) out of which the vortex propagates. It serves as a benchmark to diagnose nesting errors as in Debreu et al. (2012), Penven et al. (2006) and Spall and Holland (1991)
- WAD: a set of simple closed basin geometries for testing the Wetting and drying capabilities. Examples range from a closed channel with EW linear bottom slope to a parabolic EW channel with a Gaussian ridge.

- ICEDYN: east-west + north-south periodic channel. The common configuration includes an AGRIF zoom (1:3) in the middle of the basin to test how an ice patch is advected through it

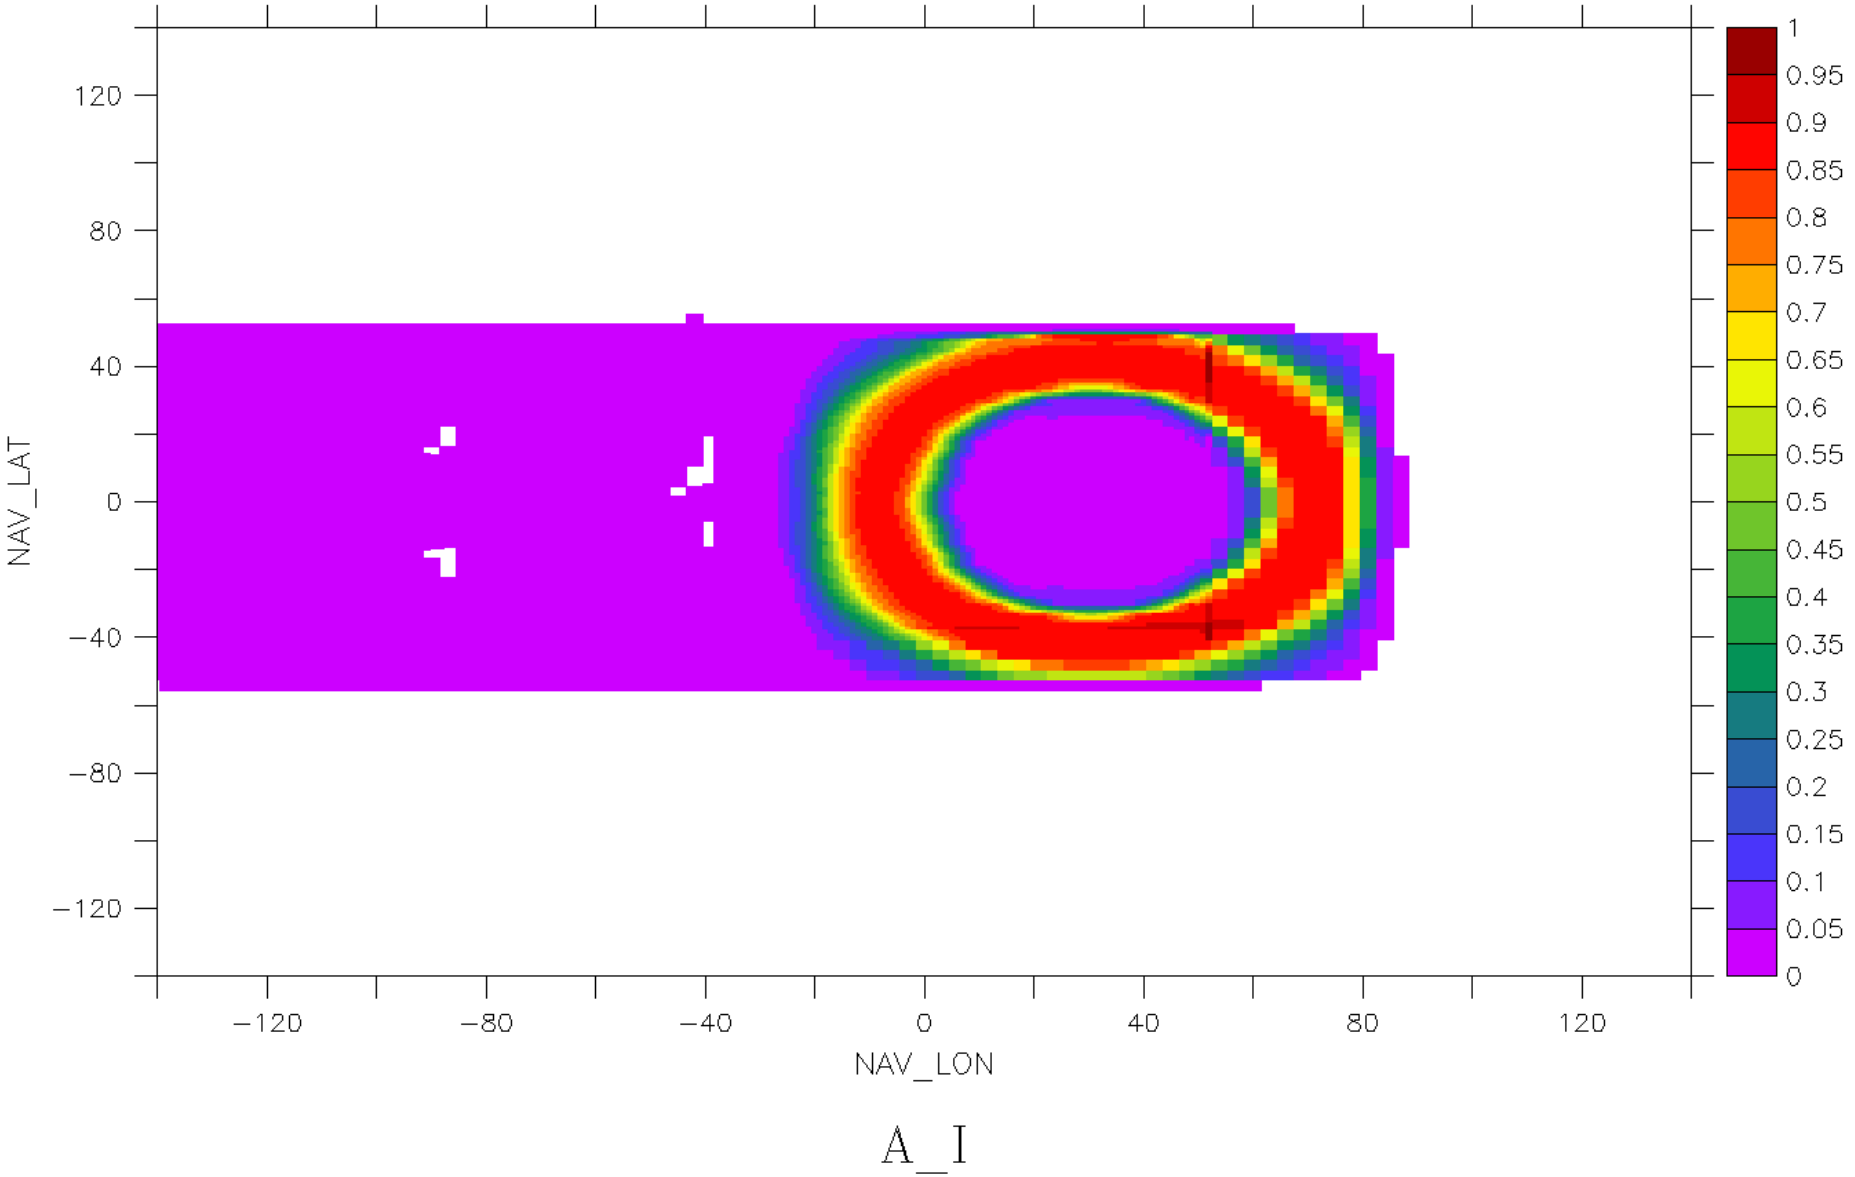

- ISOMIP: simple box configuration with an ice shelf with simple geometry on top.

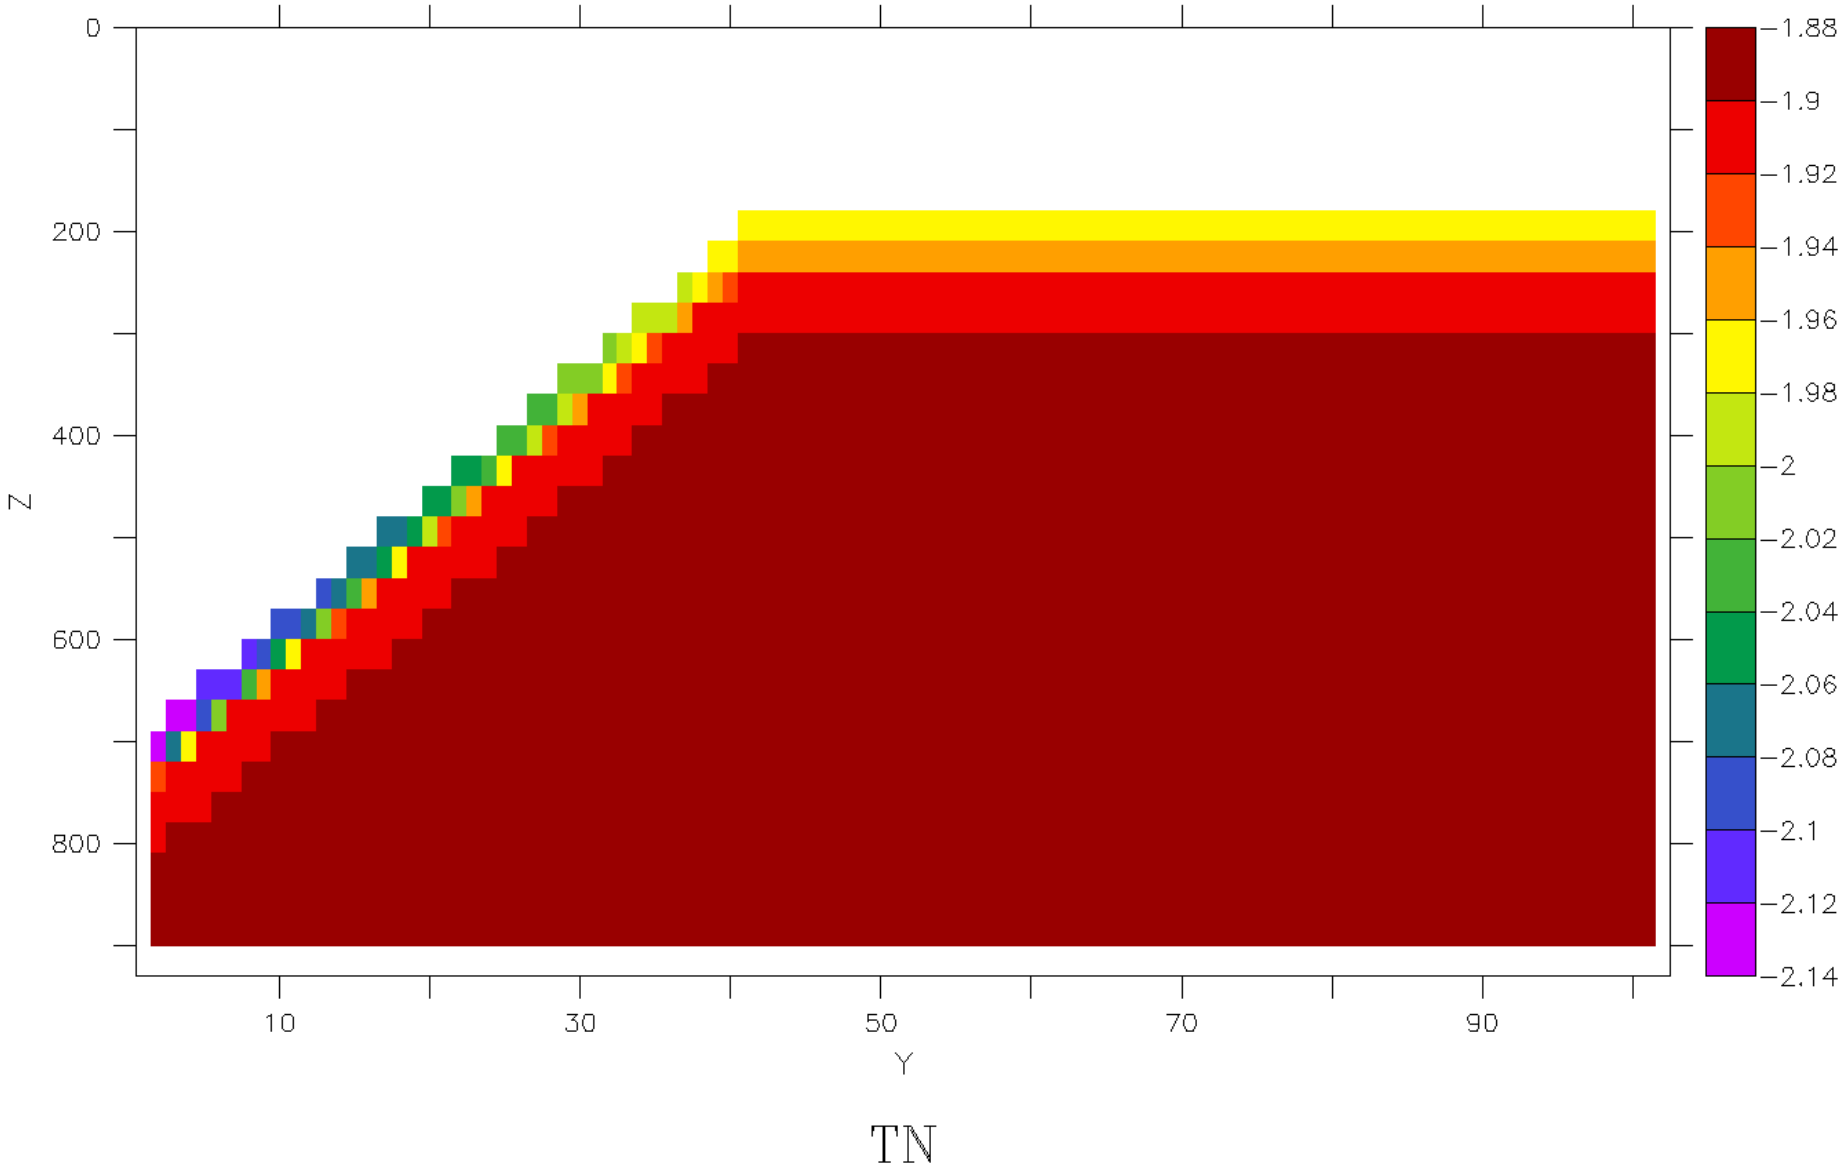

- VORTEX: illustrates the propagation of an anticyclonic eddy over a Beta plan and flat bottom. It is implemented here with an online refined subdomain (thanks to AGRIF library) out of which the vortex propagates.

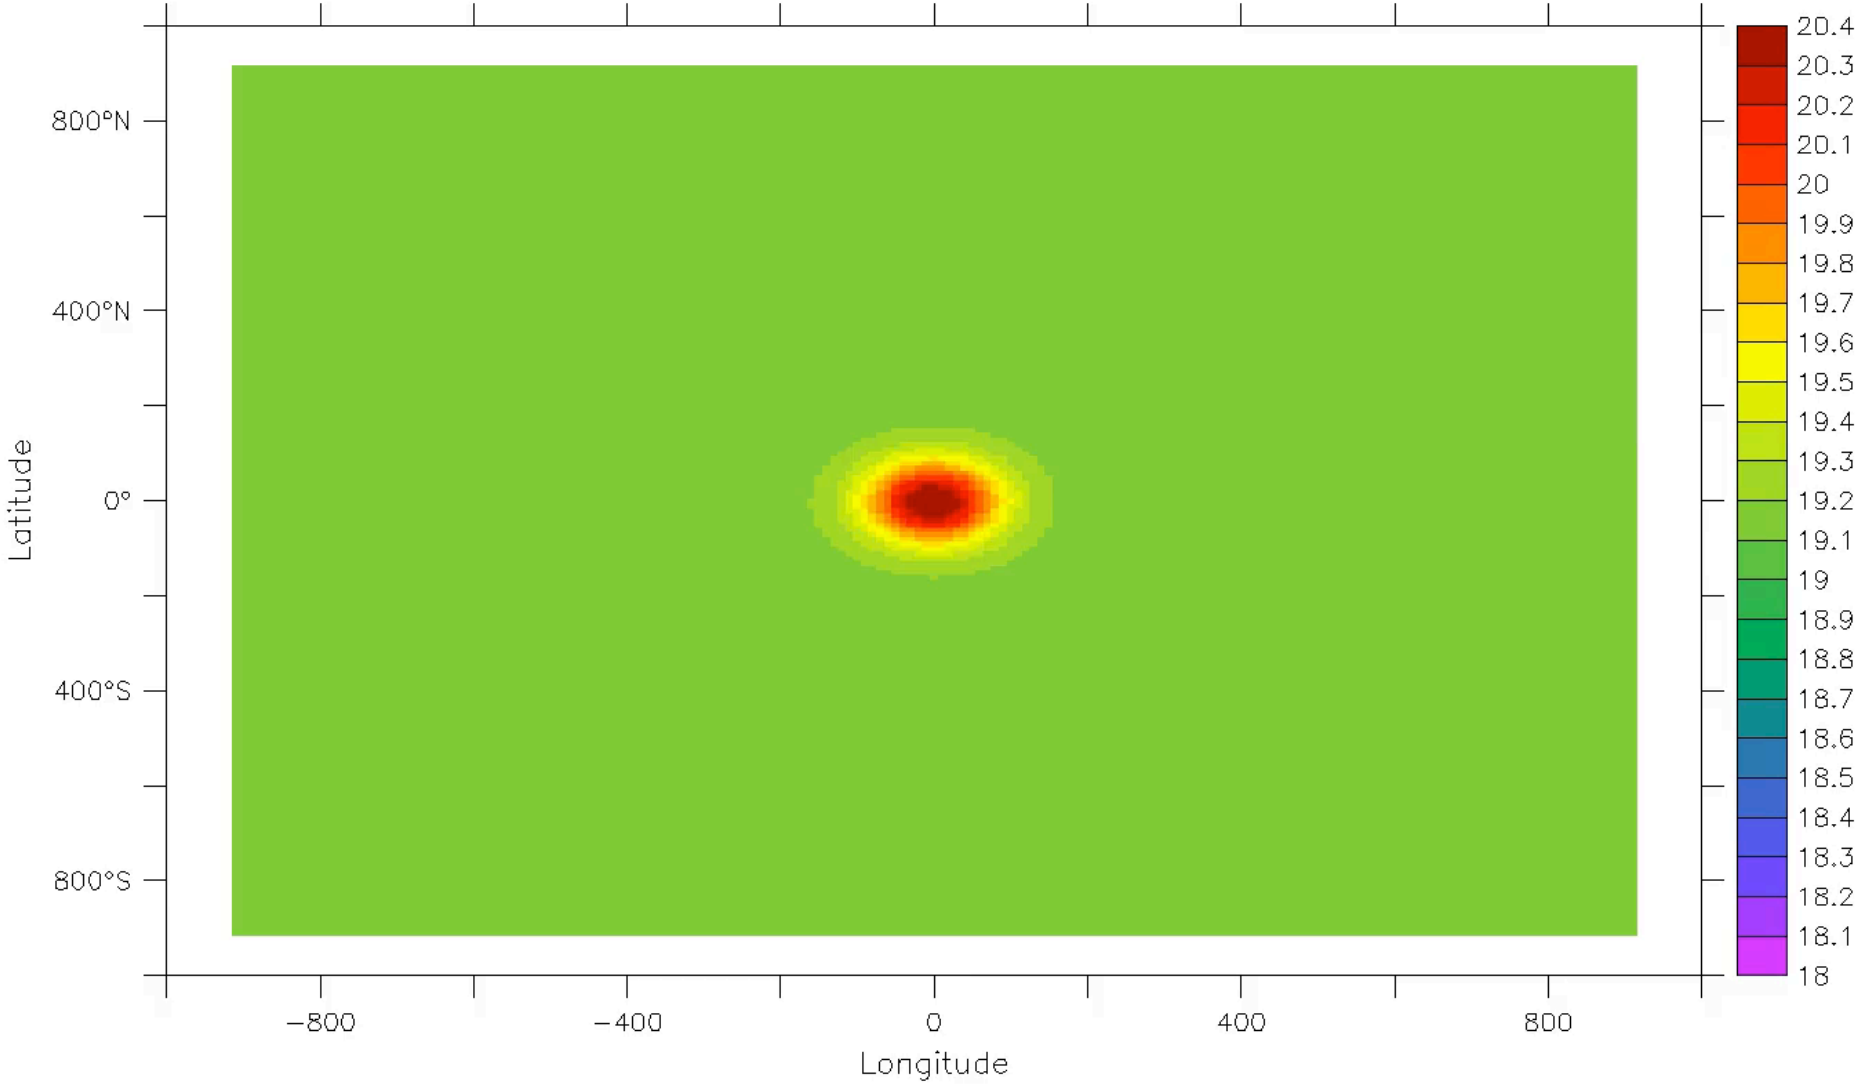

- WAD: a set of simple closed basin geometries for testing the Wetting and drying capabilities. Examples range from a closed channel with EW linear bottom slope to a parabolic EW channel with a Gaussian ridge.

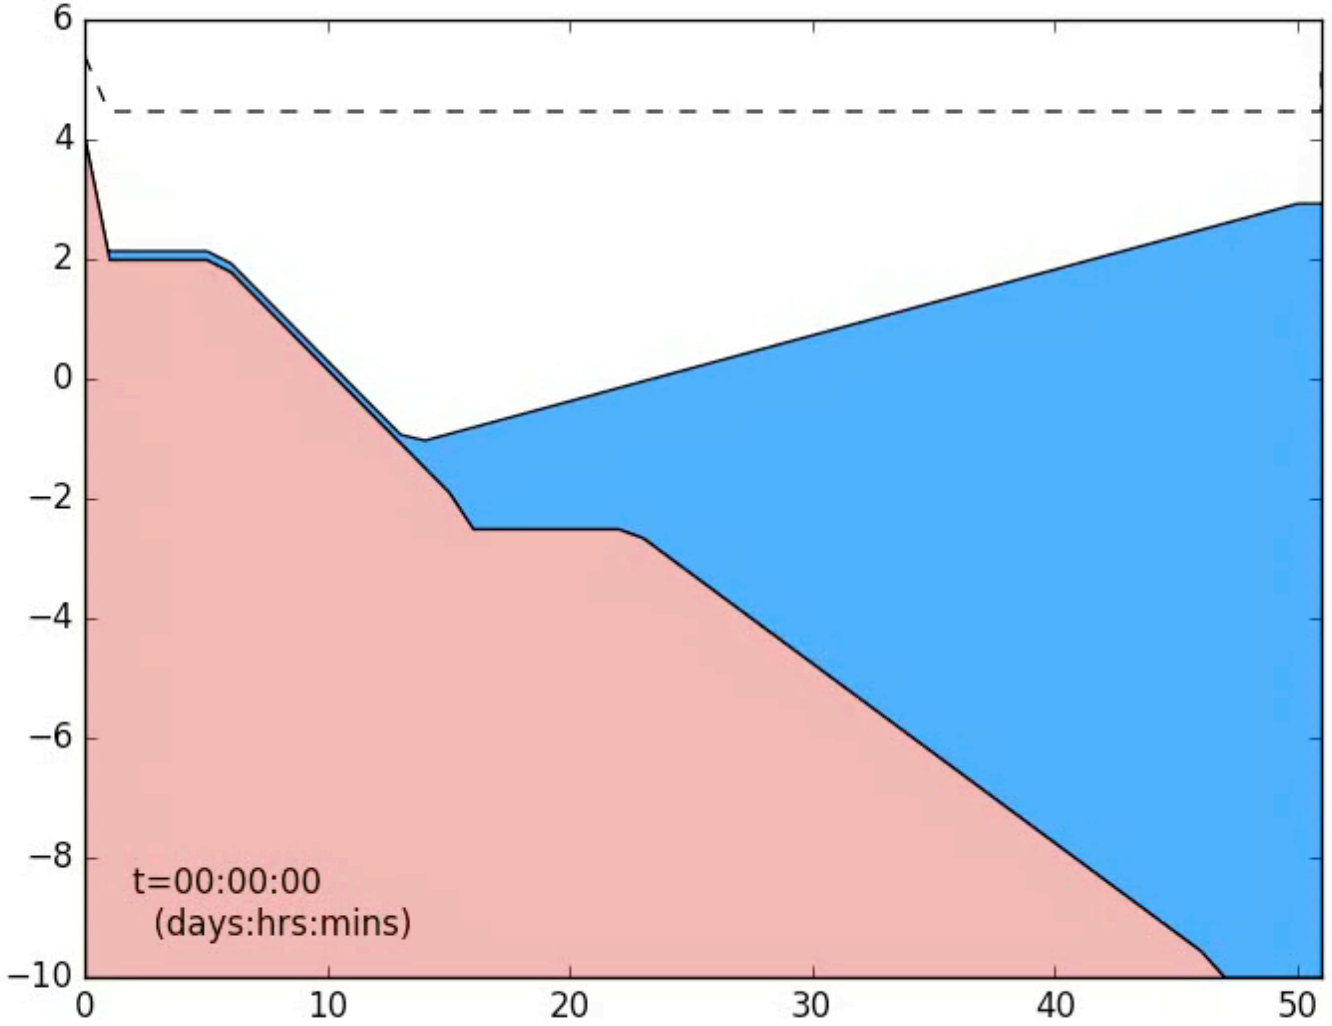

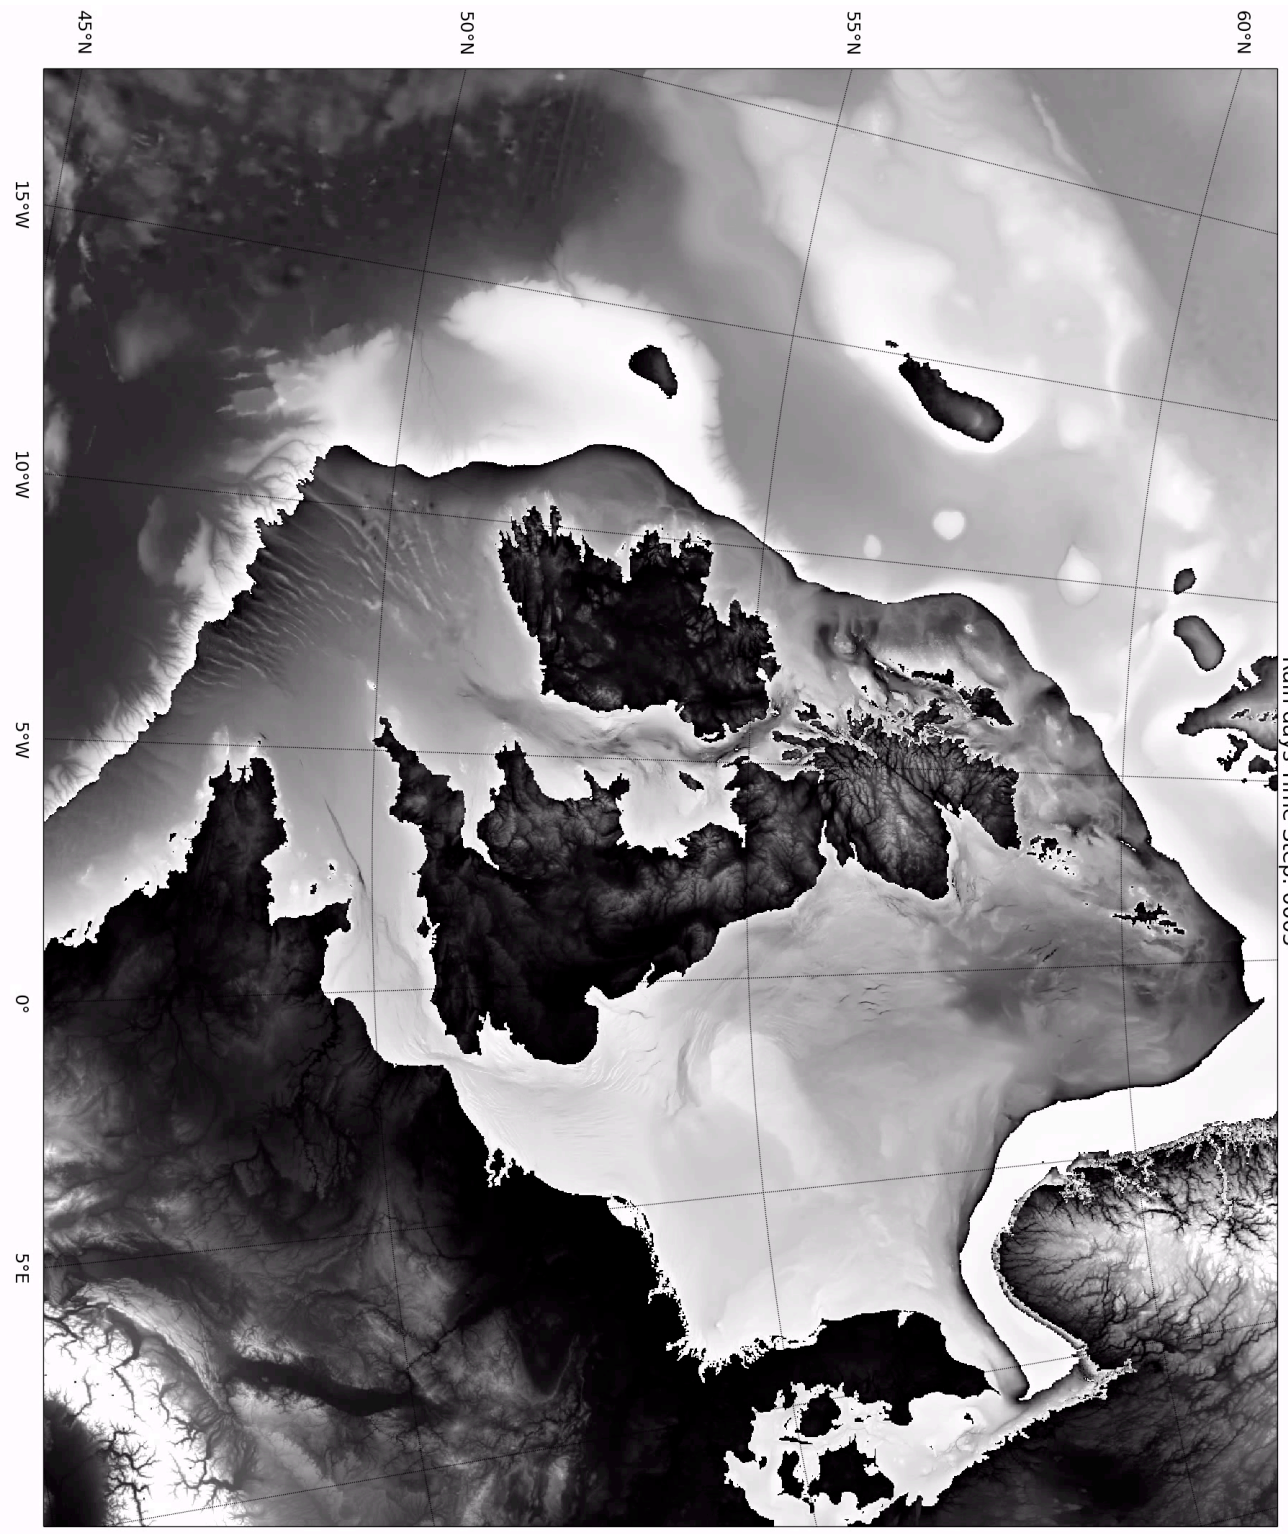

Courtesy of Enda O'Dea, Met, Office

```
 !----restart----!
WGYRE_PISCES            run.stat    restartability  passed :  20180920
WGYRE_PISCES            tracer.stat restartability  passed :  20180920
WORCA2_ICE_PISCES       run.stat    restartability  passed :  20180920
WORCA2_ICE_PISCES       tracer.stat restartability  passed :  20180920
WORCA2_OFF_PISCES       tracer.stat restartability  passed :  20180920
WAMM12                  run.stat    restartability  passed :  20180920
WORCA2_SAS_ICE          run.stat    restartability  passed :  20180920
WAGRIF_DEMO             run.stat    restartability  passed :  20180920
WSPITZ12                run.stat    restartability  passed :  20180920
WISOMIP                 run.stat    restartability  passed :  20180920
WOVERFLOW               run.stat    restartability  passed :  20180920
WLOCK_EXCHANGE          run.stat    restartability  passed :  20180920
WVORTEX                 run.stat    restartability  passed :  20180920
WICEDYN                 run.stat    restartability  passed :  20180920

 !----repro----!
WGYRE_PISCES            run.stat    reproducibility passed :  20180920
WGYRE_PISCES            tracer.stat reproducibility passed :  20180920
WORCA2_ICE_PISCES       run.stat    reproducibility passed :  20180713
WORCA2_ICE_PISCES       tracer.stat reproducibility passed :  20180713
WORCA2_OFF_PISCES       tracer.stat reproducibility passed :  20180920
WAMM12                  run.stat    reproducibility passed :  20180920
WORCA2_SAS_ICE          run.stat    reproducibility passed :  20180920
WORCA2_ICE_OBS          run.stat    reproducibility passed :  20180920
WAGRIF_DEMO             run.stat    reproducibility passed :  20180920
WSPITZ12                run.stat    reproducibility passed :  20180920
WISOMIP                 run.stat    reproducibility passed :  20180920
WVORTEX                 run.stat    reproducibility passed :  20180920
WICEDYN                 run.stat    reproducibility passed :  20180920

 !----agrif check----!
ORCA2 AGRIF vs ORCA2 NOAGRIF run.stat  unchanged  - passed :  20180920
```

Key changes from the user's perspective

- New repository arrangements: lighter and simpler for the basic user; documentation, external packages and sette now handled as svn externals.

Repository changes:

nemo 3.6

```
.
|-- ADM
|   `-- DOC_SCRIPTS
|-- DOC
|   |-- Figures
|   |-- Namelists
|   `-- TexFiles
`-- NEMOGCM
    |-- ARCH
    |-- CONFIG
    |   |-- makenemo
    |-- EXTERNAL
    |-- NEMO
    |   |-- LIM_SRC_2
    |   |-- LIM_SRC_3
    |   |-- NST_SRC
    |   |-- OFF_SRC
    |   |-- OOO_SRC
    |   |-- OPA_SRC
    |   |-- SAS_SRC
    |   `-- TOP_SRC
    |-- SETTE
    |-- TOOLS
    `-- fcm-make
```

nemo 4.0

```
.
|-- arch
|-- cfgs
|-- doc
|-- ext
|-- makenemo
|-- mk
|-- src
|   |-- ICE
|   |-- NST
|   |-- OCE
|   |-- OFF
|   |-- SAO
|   |-- SAS
|   `-- TOP
|-- tests
`-- tools
```

http://forge.ipsl.jussieu.fr/nemo/browser/NEMO/trunk

Key changes from the user's perspective

- New repository arrangements: lighter and simpler for the basic user; documentation, external packages and sette now handled as svn externals.

- Big reduction in CPP keys most options now available through the namelists; pre-processor building of generic routines to reduce code base

key_si3	147	key_diainstant	7
key_agrif	109	key_c1d	7
key_top	72	key_trdtrc	5
key_iomput	63	key_cice4	4
key_mpp_mpi	42	key_diahth	3
key_sed	38	key_diadct	3
key_cice	32	key_cyclone	3
key_vertical	29	key_nosignedzero	3
key_sed_off	18	key_oa3mct_v3	2
key_asminc	18	key_diaharm	2
key_trdmxl_trc	16	key_vectopt_loop	1
key_nemocice_decomp	16		
key_floats	15		
key_oasis3	10		
key_netcdf4	7		

key_si3	147
key_agrif	109
key_top	72
key_iomput	63
key_mpp_mpi	42
key_sed	38
key_vertical	29
key_sed_off	18
key_asminc	18
key_oasis3	10
key_nosignedzero	3
key_vectopt_loop	1

key_cice	32
key_trdmxl_trc	16
key_nemocice_decomp	16
key_floats	15
key_netcdf4	7
key_diainstant	7
key_c1d	7
key_trdtrc	5
key_cice4	4
key_diahth	3
key_diadct	3
key_cyclone	3
key_oa3mct_v3	2
key_diaharm	2

Key changes from the user's perspective

- New repository arrangements: lighter and simpler for the basic user; documentation, external packages and sette now handled as svn externals.

- Big reduction in CPP keys most options now available through the namelists; pre-processor building of generic routines to reduce code base

- Externalisation of domain definition domain configuration is now external to the main code (with the exception of very simple test-cases that can be done through user-defined routines (usrdef)). All metrics (horizontal and vertical) now read in from domcfg netCDF file.

```
netcdf ORCA_R2_zps_domcfg {
     x = 182 ;
     y = 149 ;
     z = 31 ;
     t = UNLIMITED ; // (1
currently)
     float nav_lon(y, x) ;
     float nav_lat(y, x) ;
     float nav_lev(z) ;
     double time_counter(t) ;
     int ORCA ;
     int ORCA_index ;
     int jpiglo ;
     int jpjglo ;
     int jpkglo ;
     int jperio ;
     int ln_zco ;
     int ln_zps ;
     int ln_sco ;
     int ln_isfcav ;
     double glamt(t, y, x) ;
     double glamu(t, y, x) ;
     double glamv(t, y, x) ;
     double glamf(t, y, x) ;
     double gphit(t, y, x) ;
     double gphiu(t, y, x) ;
     double gphiv(t, y, x) ;
     double gphif(t, y, x) ;
     double e1t(t, y, x) ;
     double e1u(t, y, x) ;
     double e1v(t, y, x) ;
     double e1f(t, y, x) ;
     double e2t(t, y, x) ;
     double e2u(t, y, x) ;
     double e2v(t, y, x) ;
     double e2f(t, y, x) ;
     double ff_f(t, y, x) ;
     double ff_t(t, y, x) ;
     double e3t_1d(t, z) ;
     double e3w_1d(t, z) ;
     double e3t_0(t, z, y, x) ;
     double e3u_0(t, z, y, x) ;
     double e3v_0(t, z, y, x) ;
     double e3f_0(t, z, y, x) ;
     double e3w_0(t, z, y, x) ;
     double e3uw_0(t, z, y, x) ;
     double e3vw_0(t, z, y, x) ;
     int bottom_level(t, y, x) ;
     int top_level(t, y, x) ;
}
```

Key changes from the user's perspective

- New repository arrangements: lighter and simpler for the basic user; documentation, external packages and sette now handled as svn externals.

- Big reduction in CPP keys most options now available through the namelists; pre-processor building of generic routines to reduce code base

- Externalisation of domain definition domain configuration is now external to the main code (with the exception of very simple test-cases that can be done through user-defined routines (usrdef)). All metrics (horizontal and vertical) now read in from domcfg netCDF file.

- Namelist changes mostly renaming and reorganising into a more modular form

```
!-----------------------------------------------------------------------
&namzdf        !   vertical physics manager                       (default: NO selection)
!-----------------------------------------------------------------------
   !                       ! type of vertical closure (required)
   ln_zdfcst   = .false.      !  constant mixing
   ln_zdfric   = .false.      !  local Richardson dependent formulation (T =>   fill namzdf_ric)
   ln_zdftke   = .false.      !  Turbulent Kinetic Energy closure       (T =>   fill namzdf_tke)
   ln_zdfgls   = .false.      !  Generic Length Scale closure           (T =>   fill namzdf_gls)
   ln_zdfosm   = .false.      !  OSMOSIS BL closure                     (T =>   fill namzdf_osm)
   !
   !                       ! convection
   ln_zdfevd   = .false.      !  enhanced vertical diffusion
      nn_evdm     =    0          ! apply on tracer (=0) or on tracer and momentum (=1)
      rn_evd      =  100.         ! mixing coefficient [m2/s]
   ln_zdfnpc   = .false.      !  Non-Penetrative Convective algorithm
      nn_npc      =    1          ! frequency of application of npc
      nn_npcp     =  365          ! npc control print frequency
   !
   ln_zdfddm   = .false.   ! double diffusive mixing
      rn_avts  =    1.e-4      !  maximum avs (vertical mixing on salinity)
      rn_hsbfr =    1.6        !  heat/salt buoyancy flux ratio
   !
   !                       ! gravity wave-driven vertical mixing
   ln_zdfiwm   = .false.      ! internal wave-induced mixing            (T =>   fill namzdf_iwm)
   ln_zdfswm   = .false.      ! surface  wave-induced mixing            (T => ln_wave=ln_sdw=T )
   !
   !                       ! coefficients
   rn_avm0     =   1.2e-4     !  vertical eddy viscosity   [m2/s]       (background Kz if ln_zdfcst=F)
   rn_avt0     =   1.2e-5     !  vertical eddy diffusivity [m2/s]       (background Kz if ln_zdfcst=F)
   nn_avb      =    0         !  profile for background avt & avm (=1) or not (=0)
   nn_havtb    =    0         !  horizontal shape for avtb (=1) or not (=0)
/
```


KEEP
CALM
BECAUSE
EVERYTHING
IS OFF

National
Oceanography Centre
NATURAL ENVIRONMENT RESEARCH COUNCIL

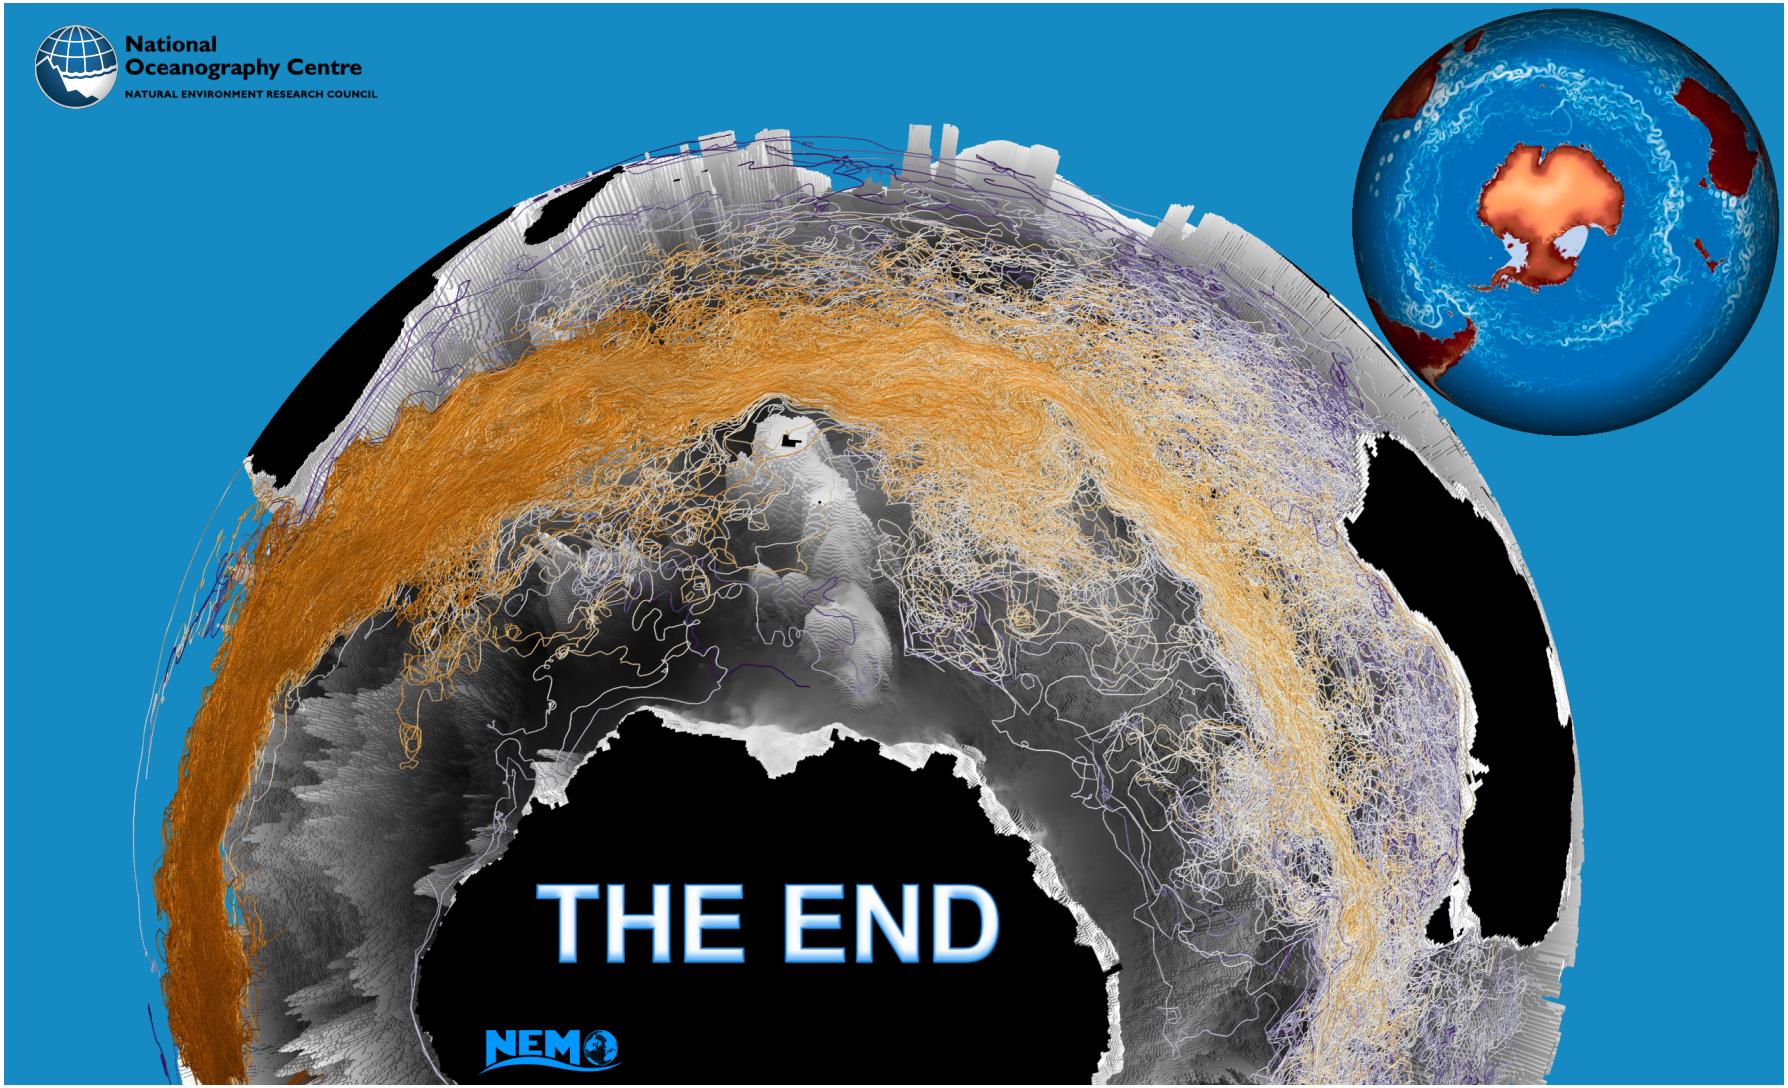

NEMO