Mixed-precision sensitivity in Earth-System modelling

Sam Hatfield*, Kristian Mogensen, Peter Düben, Nils Wedi

ECMWF

Oriol Tinto, Mario Acosta, Miguel Castrillo

BSC

*samuel.hatfield@ecmwf.int

ESiWACE2 has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823988

Impact of ocean on medium-range weather forecasts

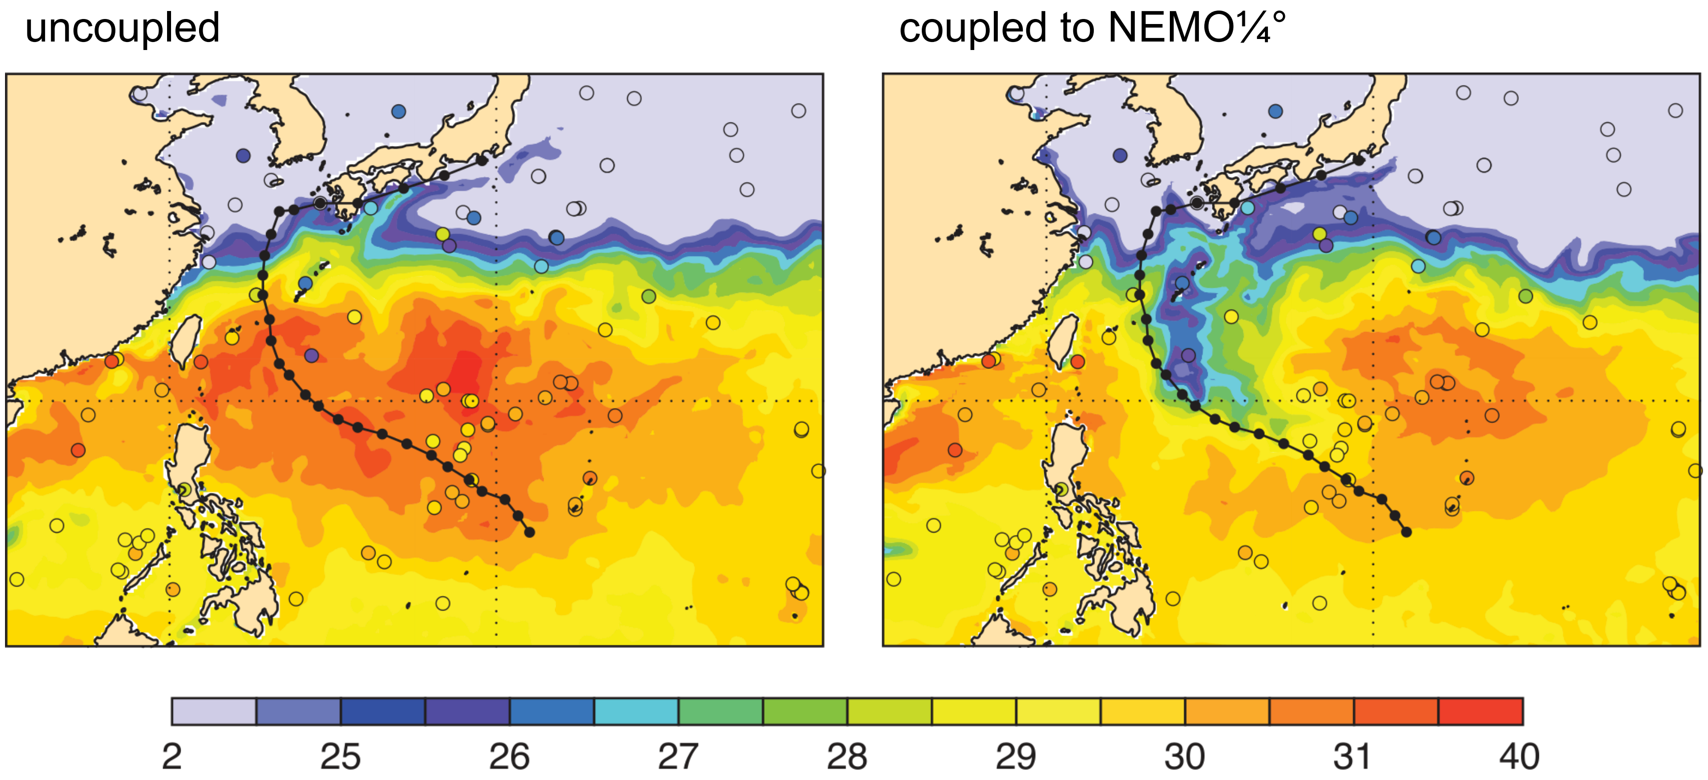

SST anomaly (C°), Typhoon Neoguri

● Best track

Mogensen et al. (2017)
J. Geophys. Res. Oceans

Cost of ocean modelling

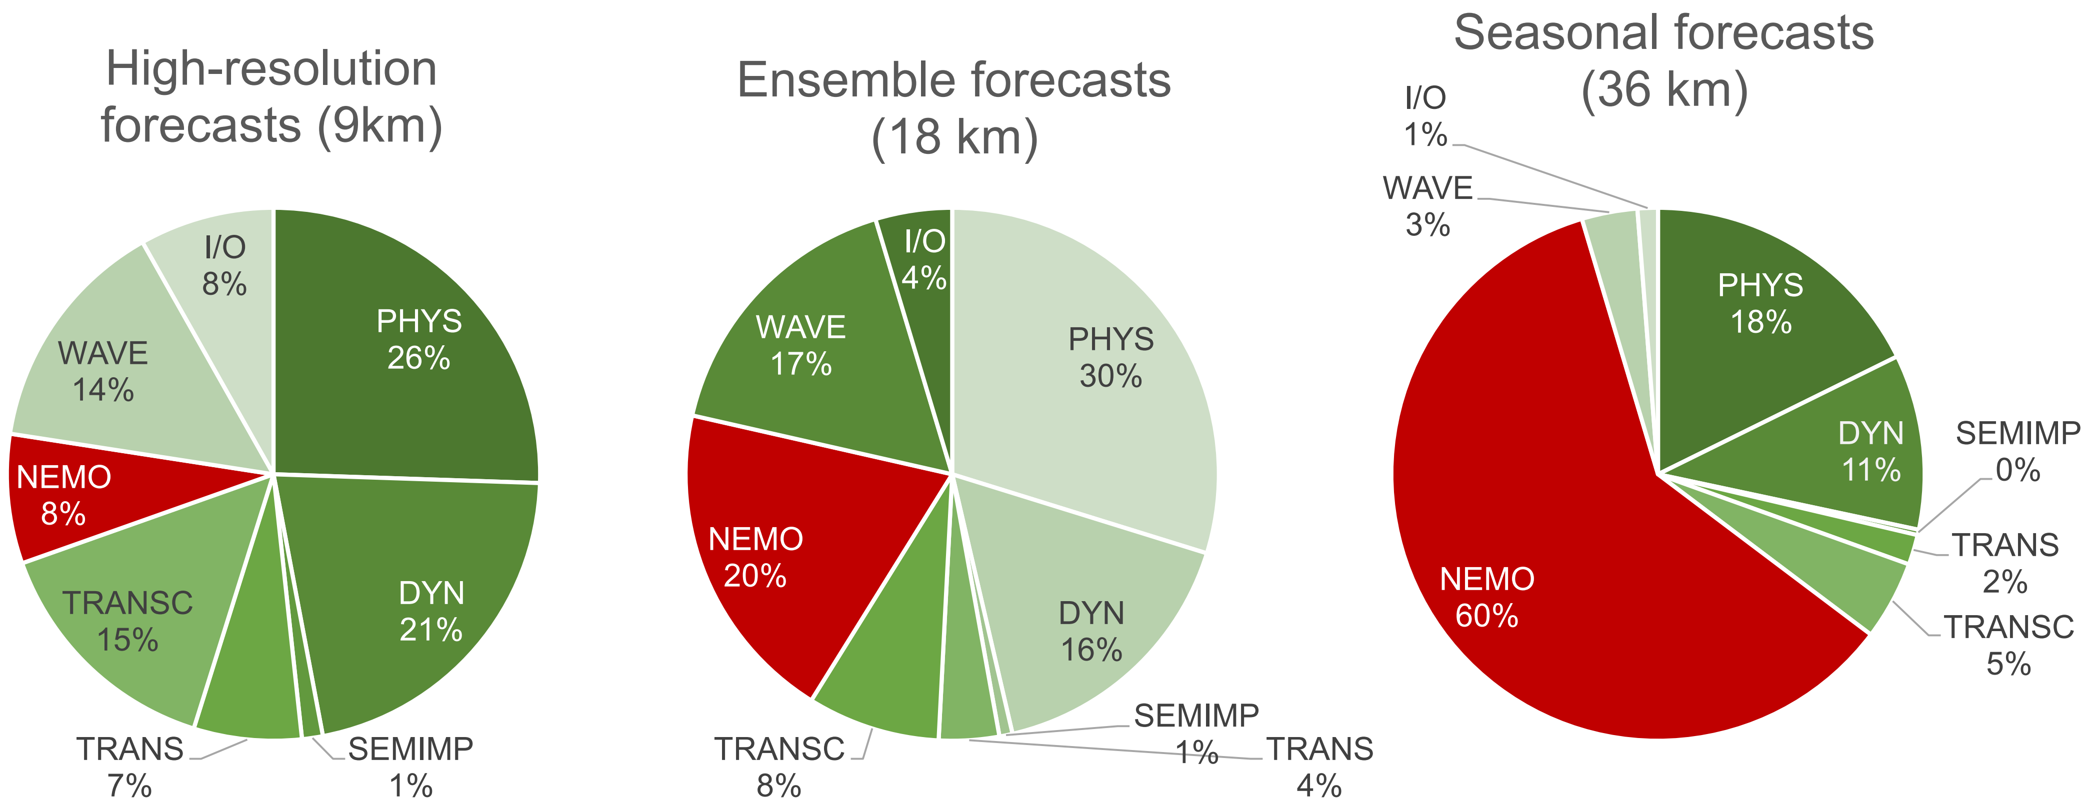

Single-precision in the atmosphere

Z500

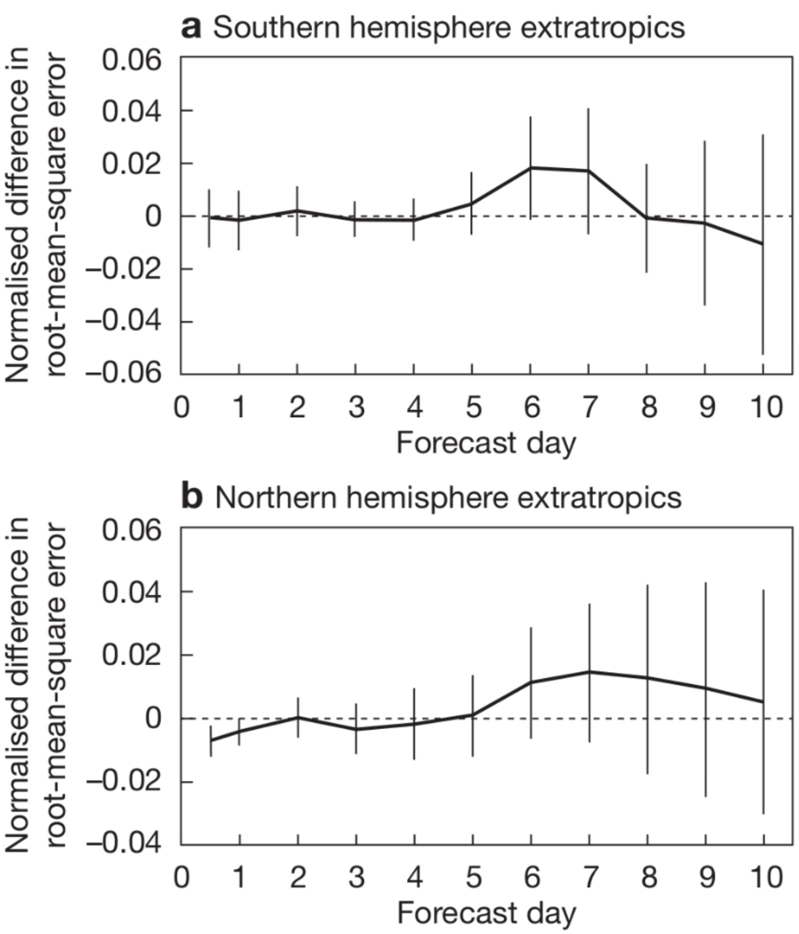

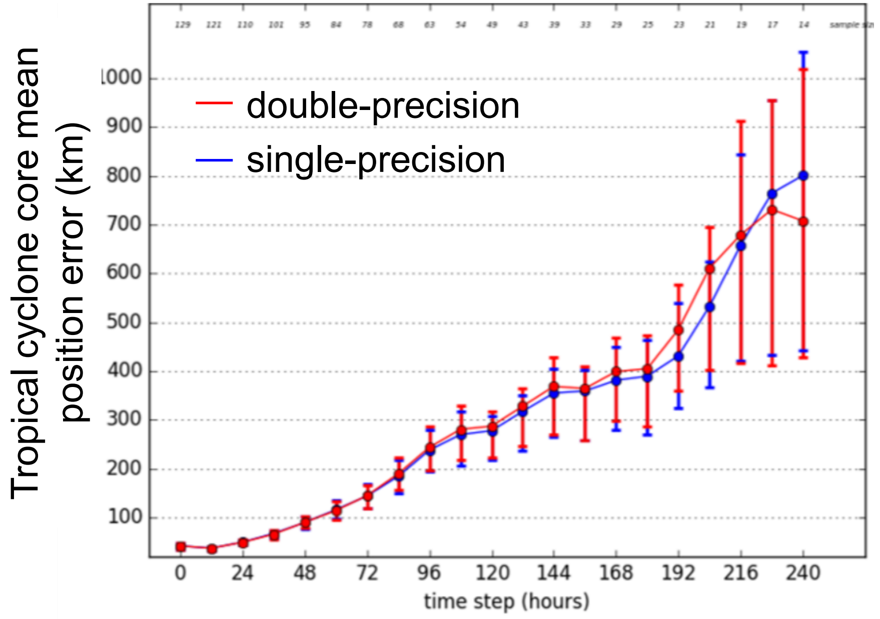

1.7x speed-up (40% reduction in wall-clock time)

Default for 1.5 km IFS experiments

Data assimilation not considered yet

How low can you go? Half-precision in the atmosphere

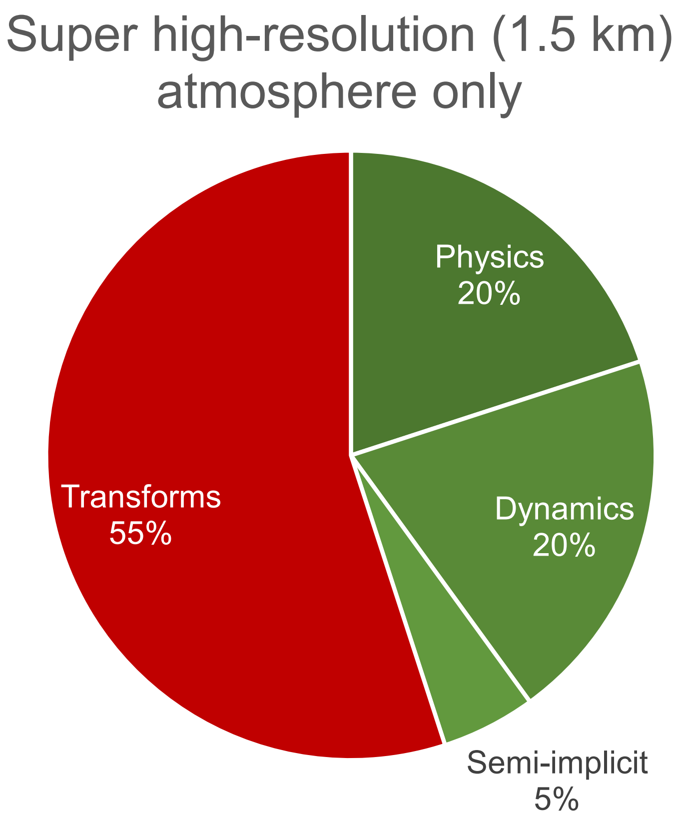

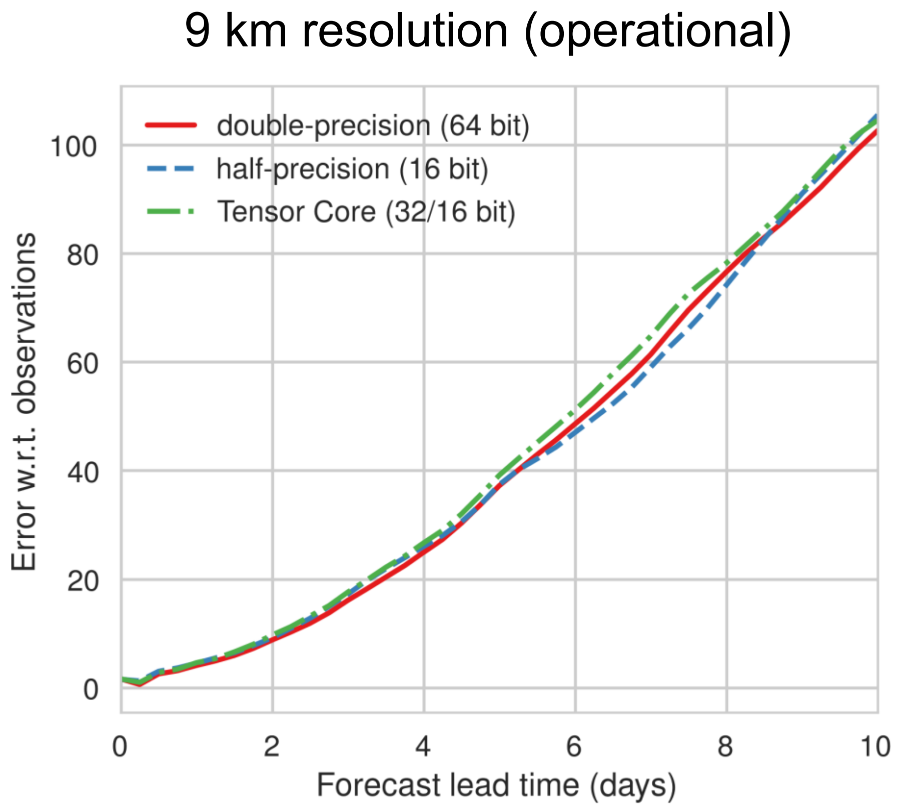

Single-precision ocean modelling at ECMWF and BSC

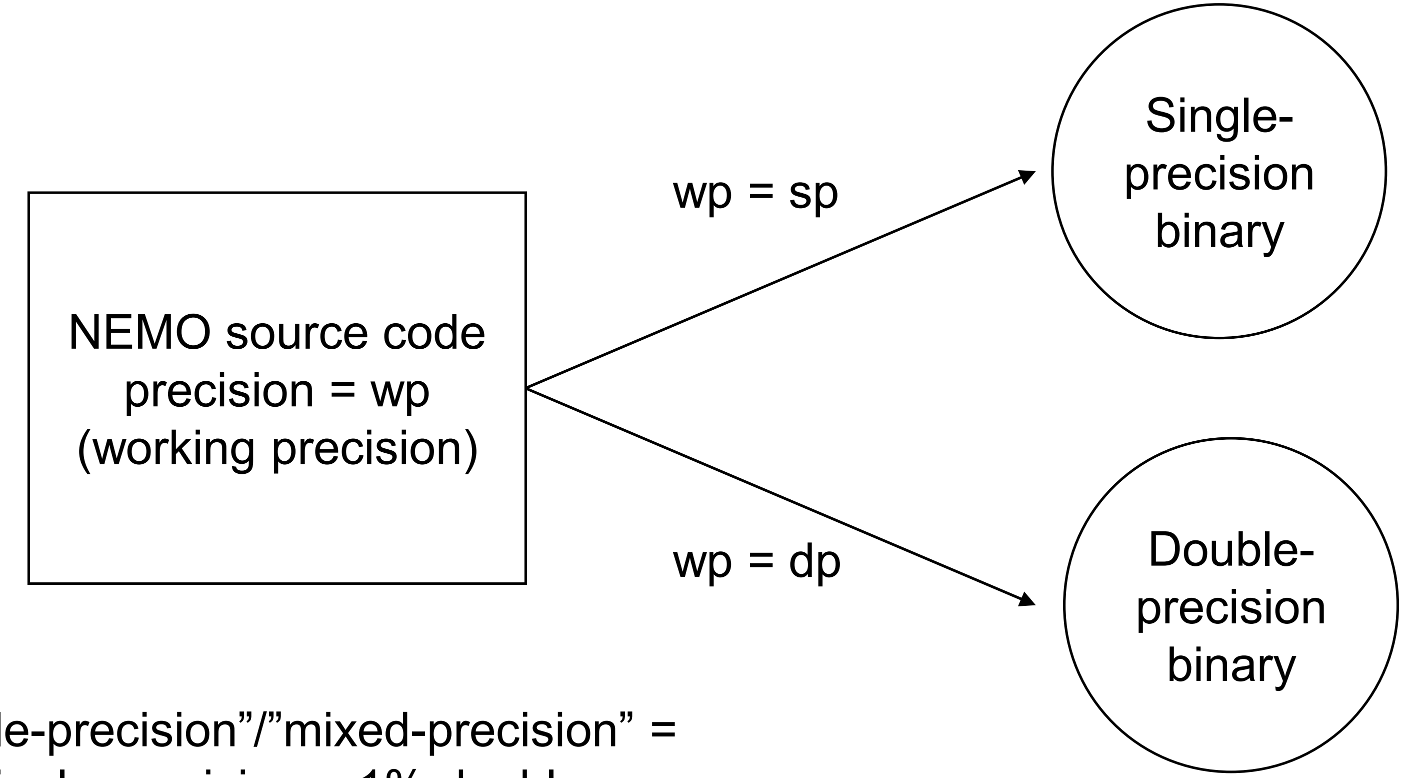

Note: “single-precision”/”mixed-precision” = ~99% single-precision, ~1% double-precision

Two types of error when reducing precision

"Catastrophic" errors

- Divide-by-zero from small sea-ice concentrations
- Overflows from comparisons with large numbers
- Cause model crashes

Assumption: these errors are edge cases that have no physical significance and can be eliminated with careful recoding

"Graceful" errors

- Slow unavoidable build-up of rounding errors
- Loss of conservation
- Don't cause model crashes

Assumption: these errors are small compared with model/observation uncertainty

Examples of catastrophic error (NEMO 4.0.1)

Example #1 (Lagrangian floats trajectories)

```
! Original code
ztxfl(jfl) = 1.E99 ! <- overflow!

! New (single-precision compatible) code
ztxfl(jfl) = HUGE(0.0_wp)
```

Example #2 (sea-ice thickness distribution)

~mitochondrion

```
! Original code
WHERE (sea_ice_conc >= 10**-20)
    t_surf = zaTsfn / sea_ice_conc
ELSEWHERE
    t_surf = 273.15
END WHERE
```

~tennis court

```
! New code
WHERE (sea_ice_conc >= 10**-6)
    t_surf = zaTsfn / sea_ice_conc
ELSEWHERE
    t_surf = 273.15
END WHERE
```

GYRER27 comparison with double-precision

Which is double-precision?

Relative vorticity after 2 years spin-up

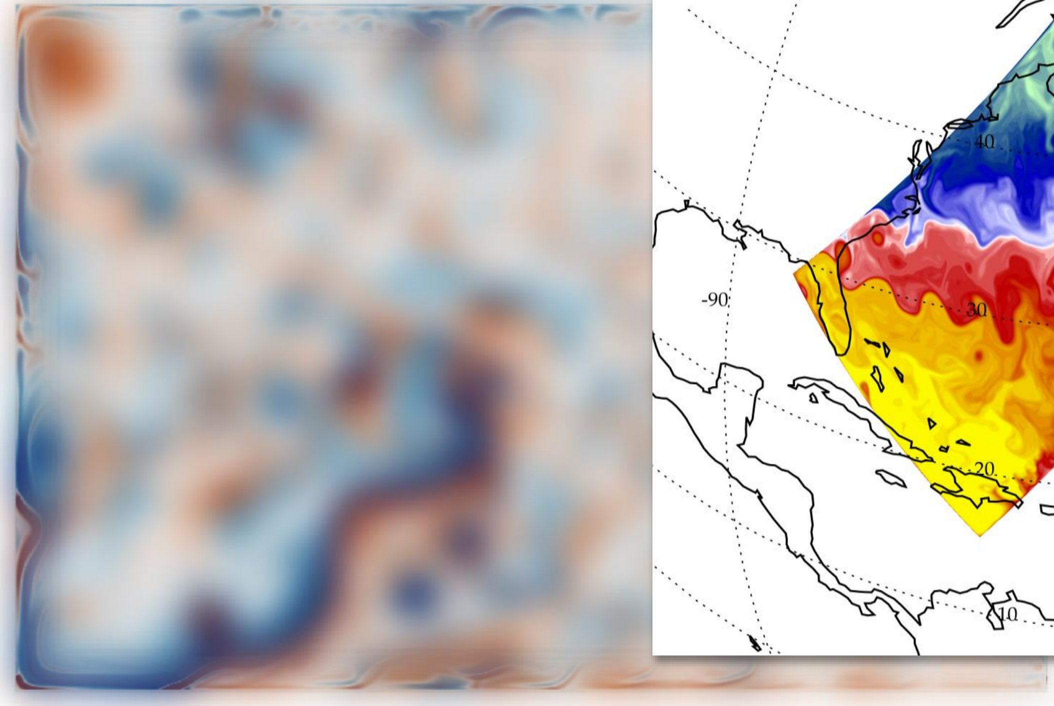

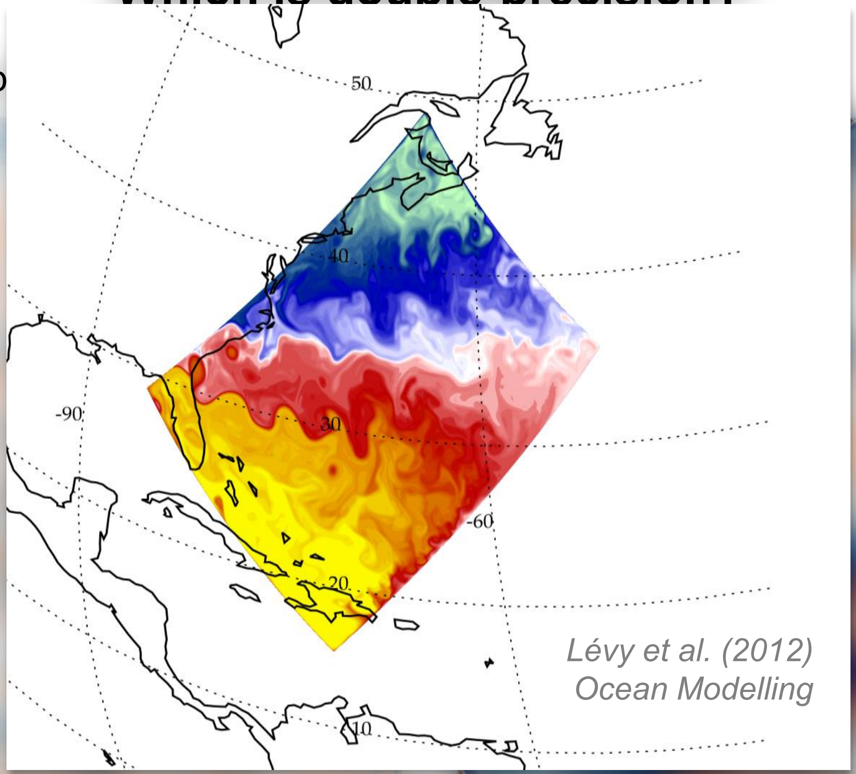

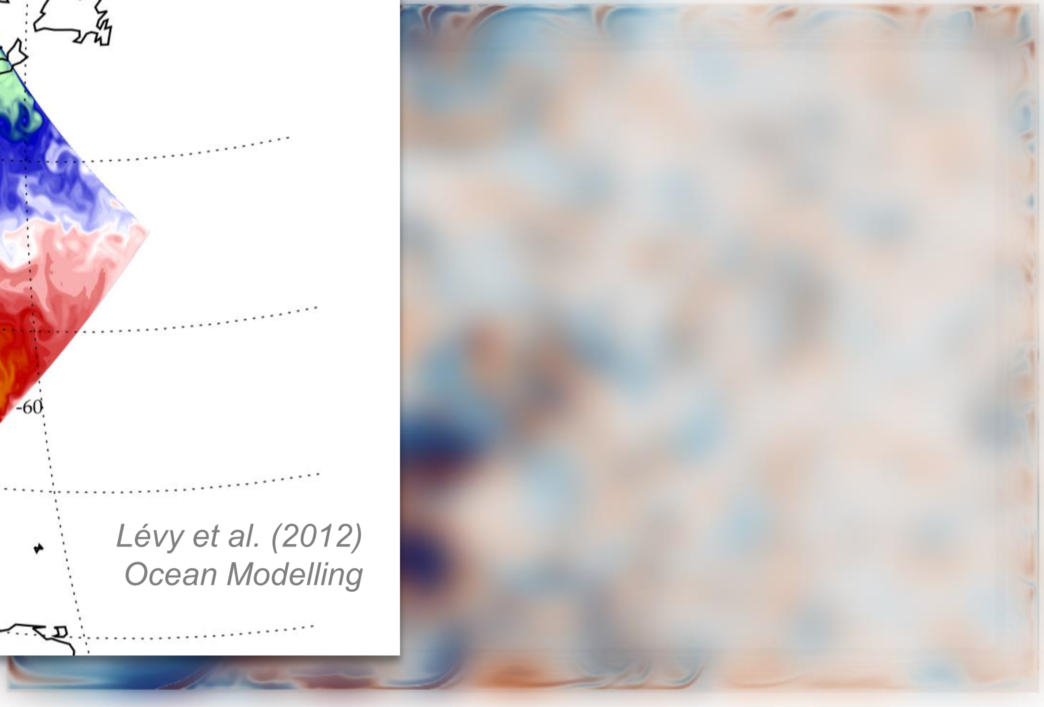

Mixed-precision

Double-precision

GYRER27 comparison with double-precision

Which is double-precision?

Relative vorticity after 2 years spin-up from rest

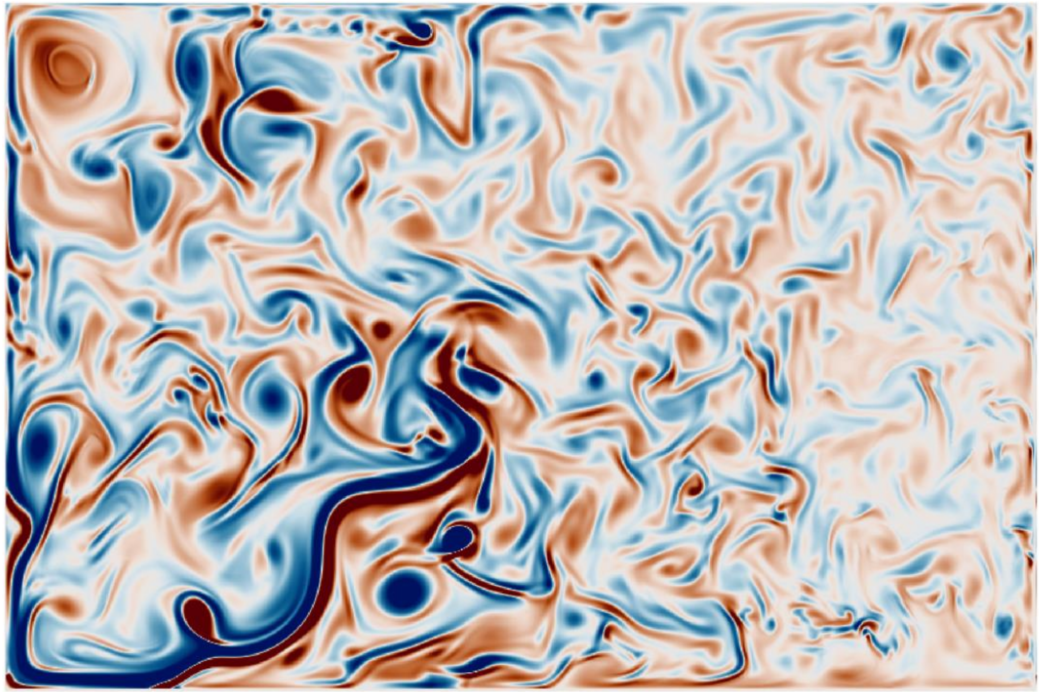

Mixed-precision

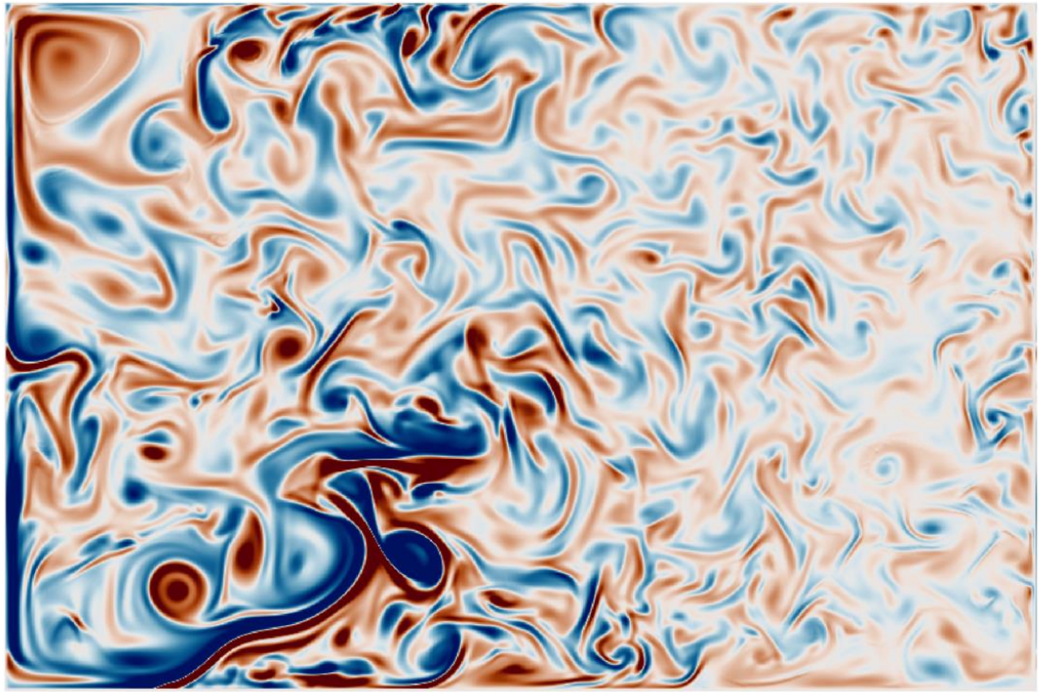

Double-precision

ORCA¼° sea-surface salinity

double-precision, day 0

mixed-precision, day 0

Strong-scaling

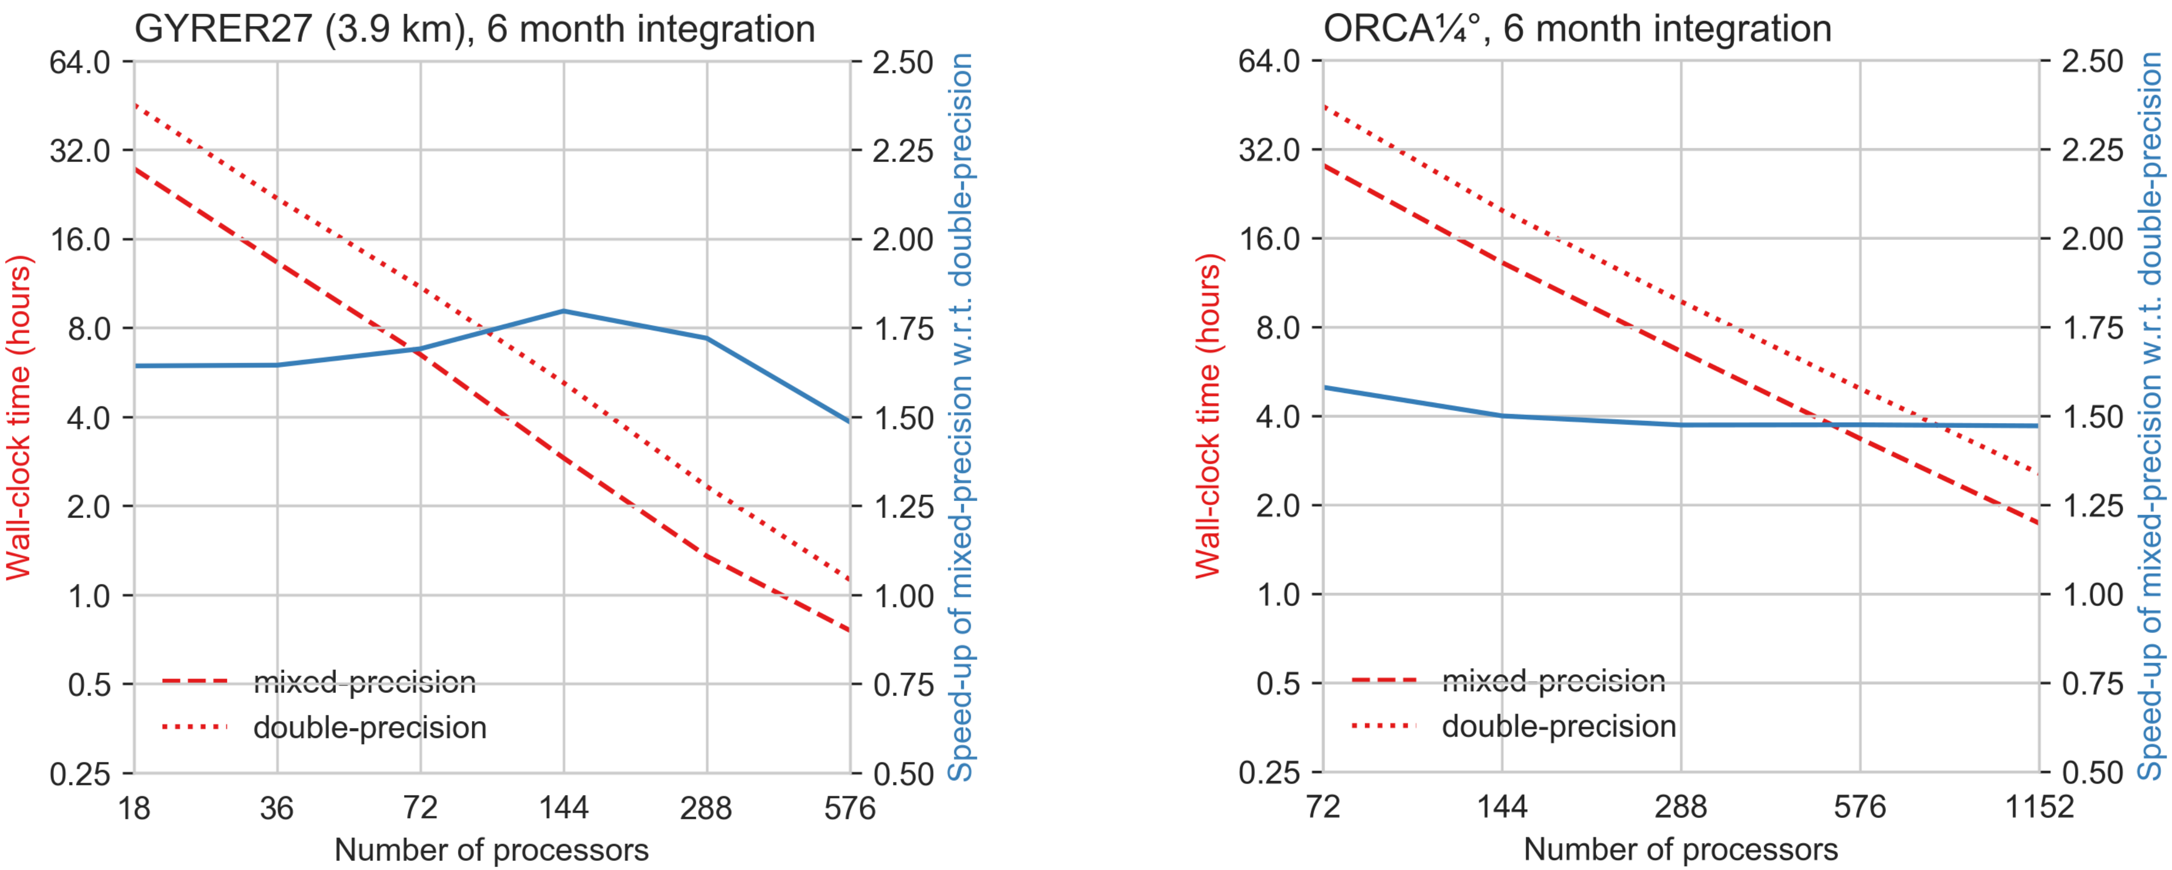

ORCA ¼° profile

Subroutine	Purpose	% of DP cost	Speed-up MP:DP
`icedyn_rhg`	Sea-ice rheology	11%	**1.17**
`tra_adv`	Tracer advection	9%	**1.48**
`zdf_phy`	Vertical ocean physics	9%	**2.24**

576 cores, 6 month integration

Overall speed-up from mixed-precision: **1.5x**

i.e. ~35% reduction in cost

10 free extra ensemble members in seasonal forecast

But “minor question”: **what about the quality of the simulation?**

Testing strategy for the atmosphere

1. Run "perfect model" tests
   - Does it even run?
   - Is it worth it? (Speed-up…)

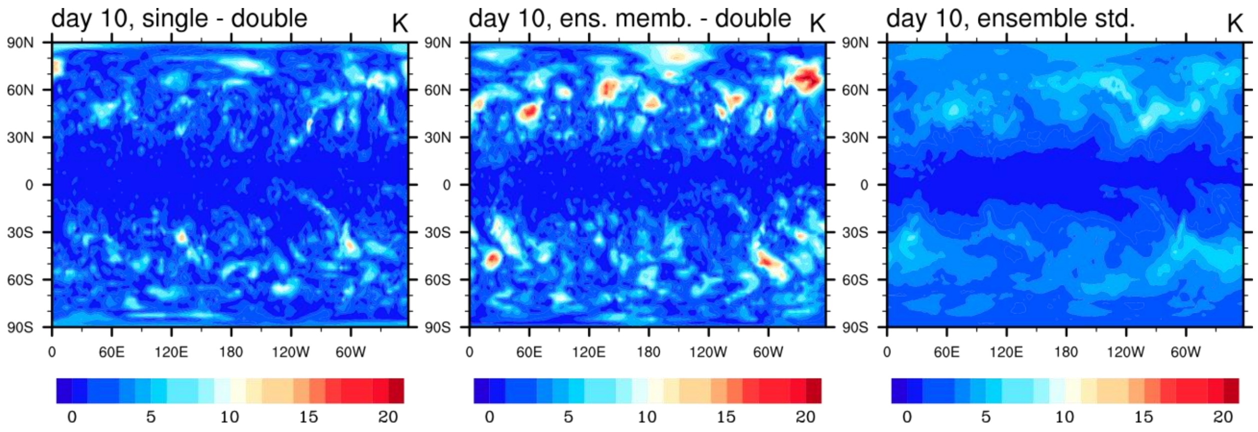

850 hPa temperature

Dueben and Palmer (2014) Mon. Wea. Rev.

Testing strategy for the atmosphere

1. Run "perfect model" tests
   - Does it even run?
   - Is it worth it? (Speed-up…)
2. Run hindcast tests
   - How does diff. between single/double compare with model/observation uncertainty?

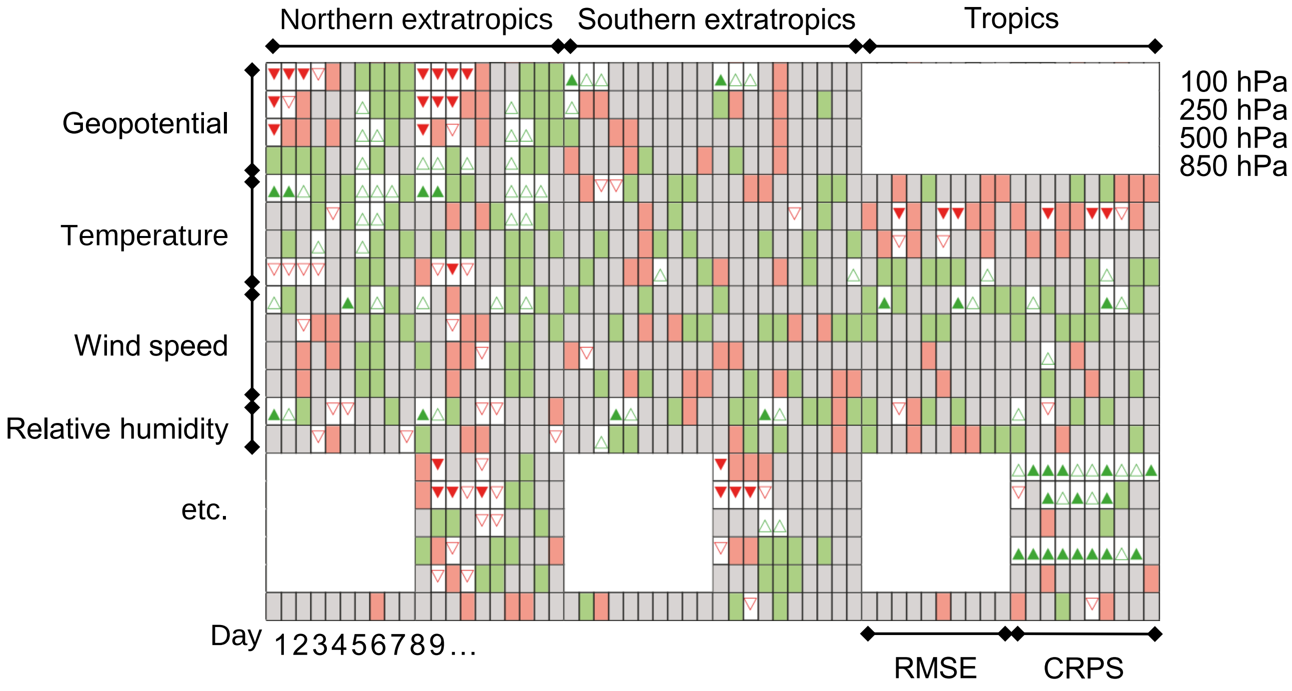

Testing strategy for the atmosphere

1. Run “perfect model” tests
   - Does it even run?
   - Is it worth it? (Speed-up…)
2. Run hindcast tests
   - How does diff. between single/double compare with model/observation uncertainty?
3. Iron out wrinkles for operational use
   - Check biases
   - Check mass etc. conservation

What about the ocean?

3 day MSLP forecasts
single-precision - double-precision

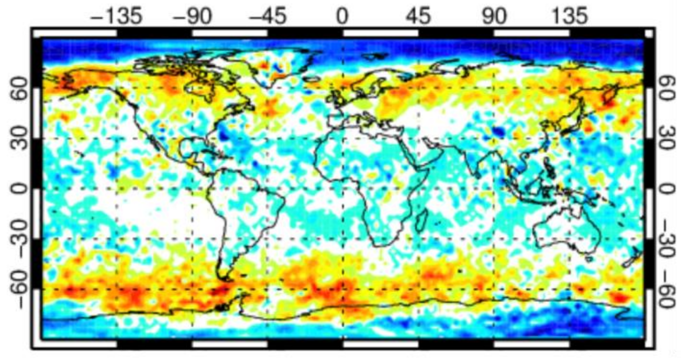

Without m=0 Legendre transform fix

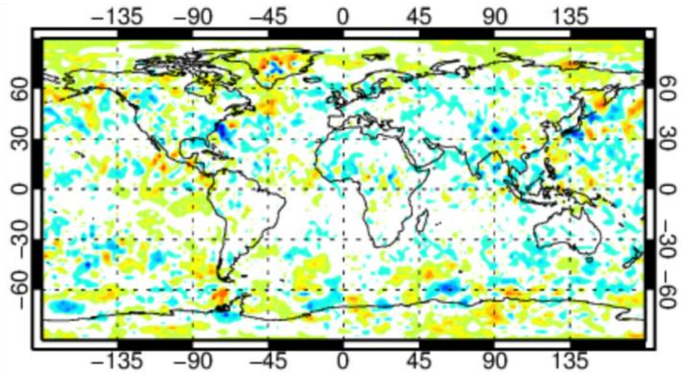

With m=0 Legendre transform fix

Testing strategy for the ocean

1. Run “perfect model” tests
   - Does it even run? ✔ (but…)
   - Is it worth it? (Speed-up…) ✔ (1.5×)
2. Run hindcast tests, forced by reanalysis
   - Long (40 year) runs
   - Medium-range forecasts

↑ Initialised!

Verifications:

- Transport (RAPID)
- SSTs (e.g. OSTIA, CCI)
- Sea-ice (e.g. OSTIA)
- Double-precision (reducing precision is **not** a “model upgrade”!)

A DCMIP for the ocean

Nakano et al. (2018) Mon. Wea. Rev.

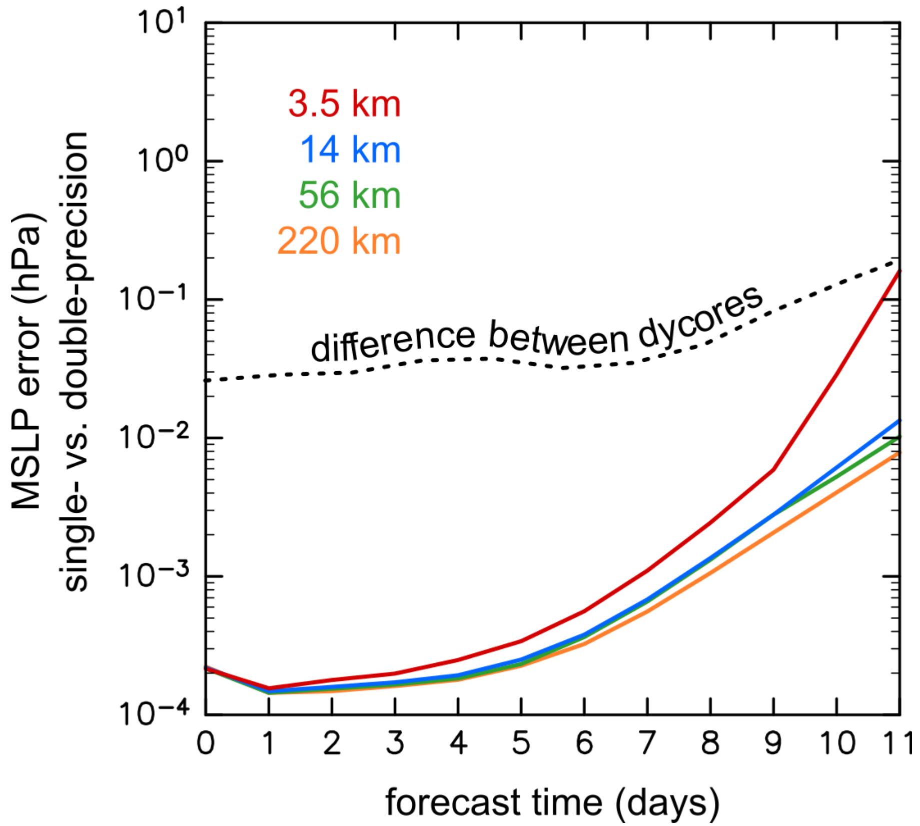

baroclinic wave test case

increasing complexity

forced global with tropical cyclone?

double-gyre

Lévy et al. (2012) Ocean Modelling

lock exchange

Ilıcak et al. (2012) Ocean Modelling

Conclusion

- Mixed-precision has been used successfully in the atmosphere at ECMWF, with ~1.7× speed-up
- Mixed-precision in the ocean provides ~1.5× speed-up, which could significantly accelerate seasonal runs
- But quantifying impact on model is not easy
  - Another motivation for testcase model intercomparisons