Bitwise reproducibility and performance of OceanVar

Francesco Carere, Francesca Mele, Italo Epicoco, Mario Adani, Paolo Oddo, Eric Jansen, Andrea Cipollone, Ali Aydogdu

Sharing is encouraged

1. Introduction

- **Bitwise Reproducibility (BR)** is difficult to obtain, especially in distributed computing (Figure 4)
- Although the problem has existed for decades in the W&C community, **no standardised workflows exist** to obtain BR. We investigate the case of OceanVar (MPI BR)
- Proposed alternative to BR: **statistical reproduciblity**

3. MPI BR of OceanVar

MPI BR can be achieved by activating:

- Intel **compiler options**:
  - `-fp-model consistent`
  - `-fimf-use-svml=true`
- **OceanVar code fixes** forcing sequential-like computation

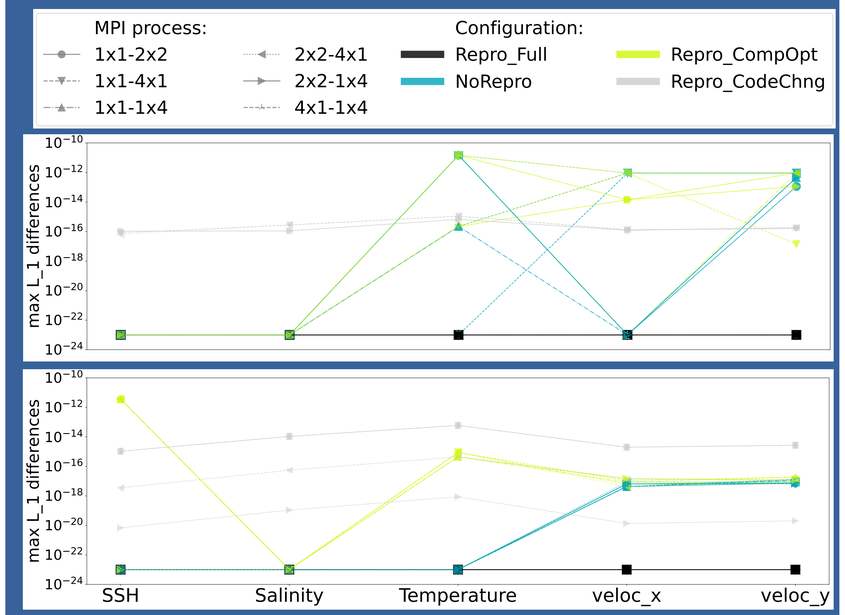

Figure 1: L_1-differences between different MPI processes of different configurations (color) of the *barotropic* (up) and *dynamic height* (down)

2. Definitions

OceanVar is an ocean data assimilation model, tested with the *barotropic model* and *dynamic height* (balance operators)

A full BR model has bitwise equal output across specified configurations (hardware, compilers, MPI processes, ...)

MPI BR is full BR across MPI processes

4. Current workflow

- A testing suite **detects** MPI BR efficiently (Figure 2)
- No tools exist for **root-cause analysis** of MPI BR bugs

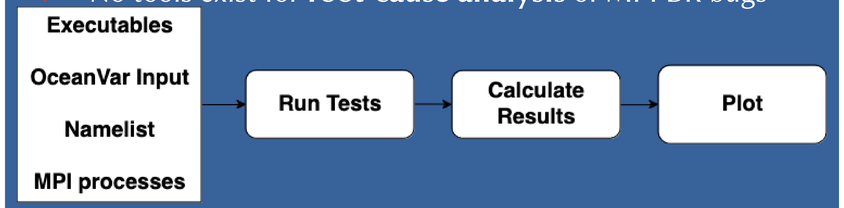

Figure 2: The testing suite automatically runs tests, and calculates and makes plots of differences between configurations (user-defined executables, OceanVar input, options, MPI procs)

6. Performance and MPI BR

- Currently OceanVar scaling is not optimised.
- However, from the Muse ocean model we know that **MPI BR can affect efficiency significantly** (Figure 3)

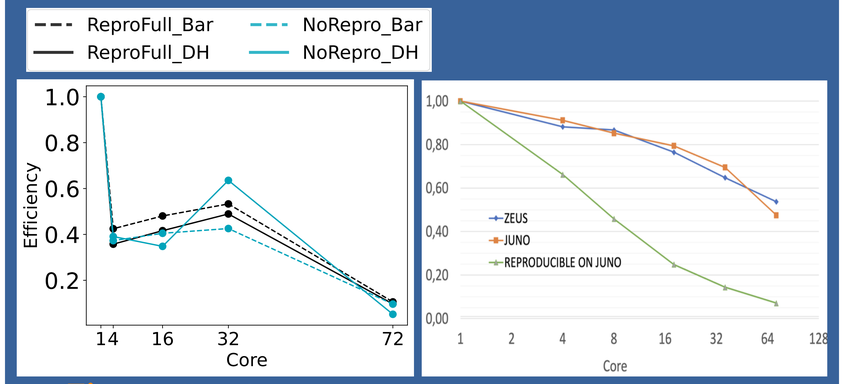

Figure 3: Efficiency of OceanVar (left) and Muse (right)

5. Suggested workflows

- Automated tools exist to debug BR problems due to different compilers (Table 1) but not for other types.
- A major source of non-MPI BR is non-associativity of summation: use stochastic rounding to detect and debug sources with Verrou tool

	Compiler (options)	MPI	SR
Detection	FLiT/pLiner/Ciel	See Fig. 2	CESM-ECT
Root-cause	FLiT/pLiner/Ciel	Verrou(?)	CESM-RUANDA

Table 1: BR debugging tools: detection and root-cause analysis

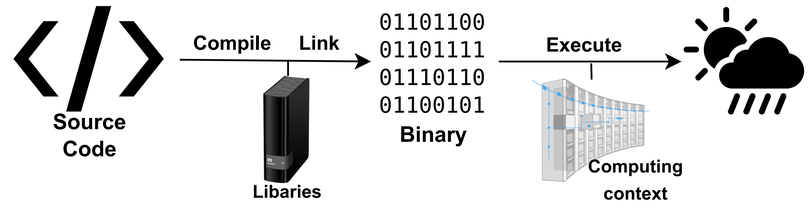

Figure 4: Failure of BR may occur anywhere in the workflow

7. Statistical reproducibility

- In W&C, **statistical reproducibility (SR)** is proposed as an alternative to BR for performance and practical reasons, which has been implemented in the Muse ocean model
- A model is SR if variability across tested configurations is of the order of internal variability
- Currently automated tools are adapted only to MPAS-A and to the atmosphere and ocean parts of CESM (Table 1).

8. Future work

- Implement suggested workflow for MPI BR debugging
- Improve OceanVar scalability, implement SR detection
- Extend test suite to general models. Build automatic MPI

Francesco Carere
francesco.carere@cmcc.it
Centro Euro-Mediterraneo sui Cambiamenti Climatici
Euro-mediterranean center on climate change (CMCC), IESP institute, ADIC Division, Lecce