



Navid Constantinou



# Oceananigans.jl

*breakthrough resolution, memory, and energy efficiency  
in global ocean simulations*

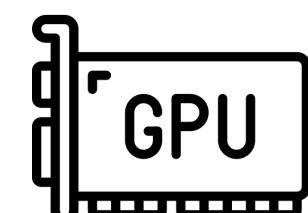
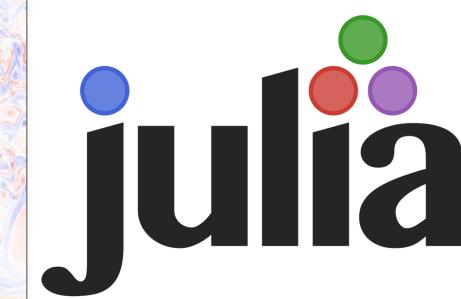
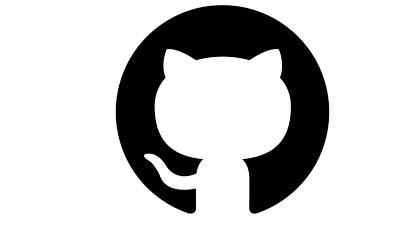
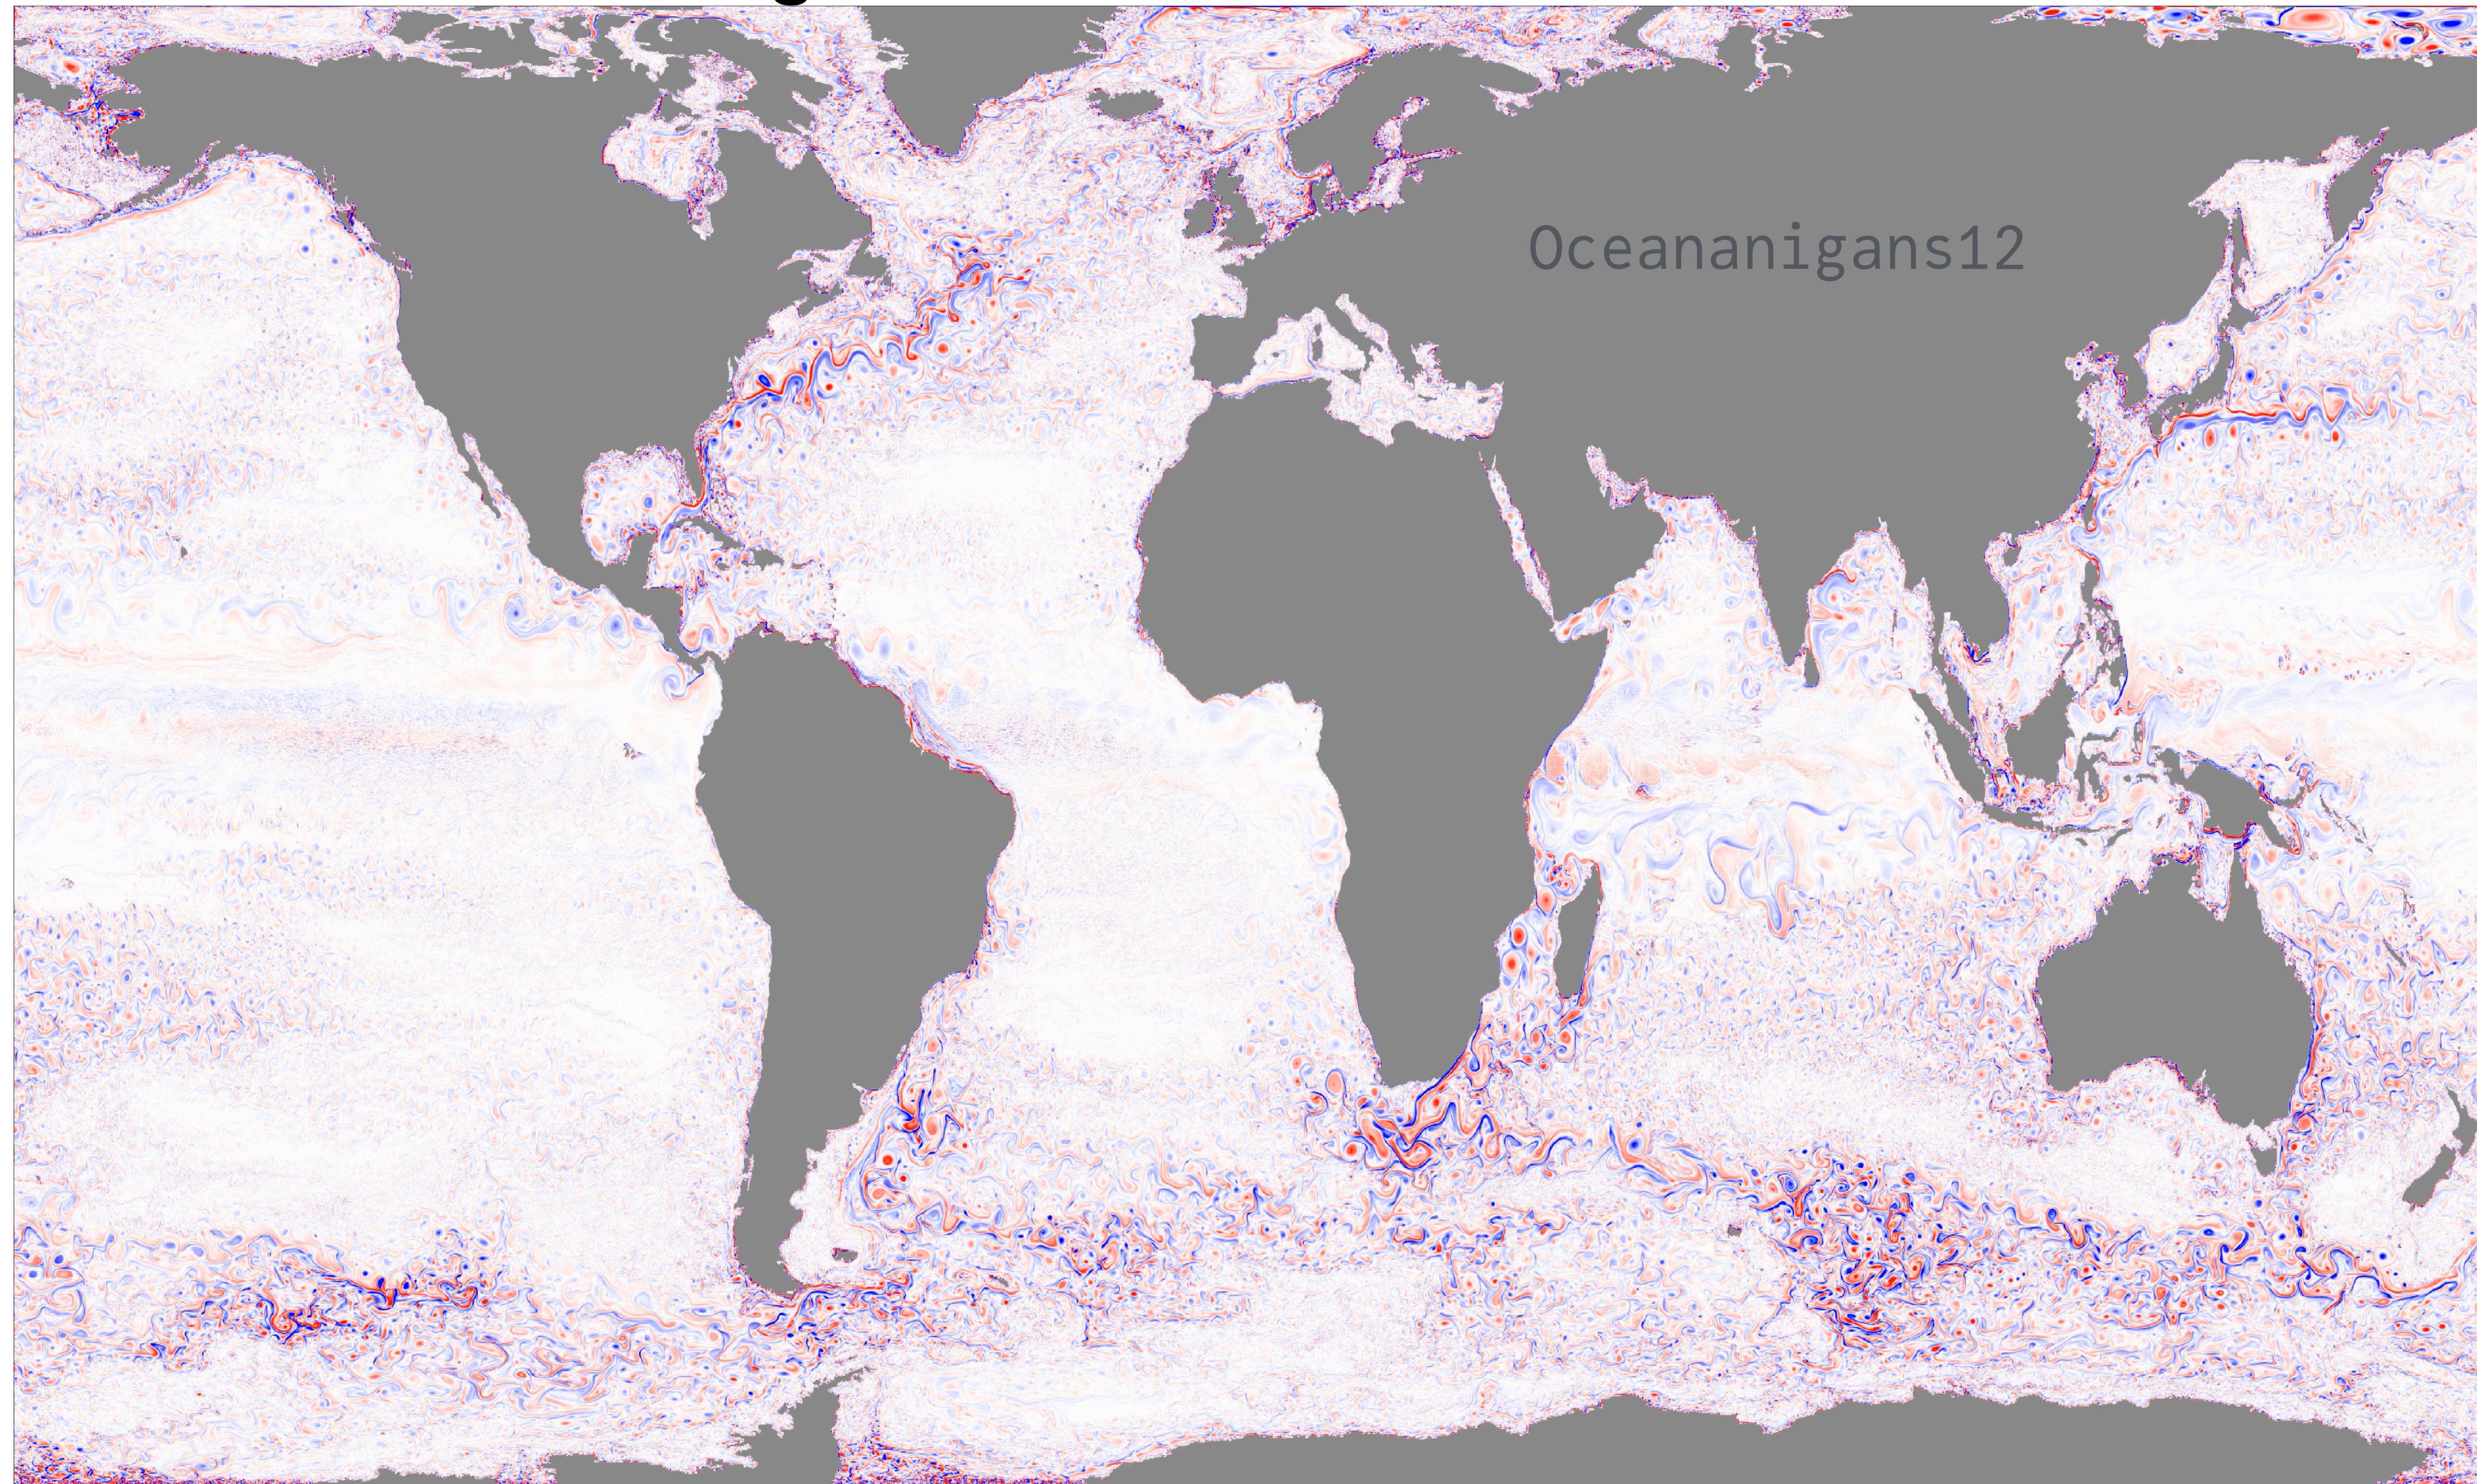


Simone Silvestri



Gregory Wagner

and also all  
Ocean Core  
CliMA Dev Team



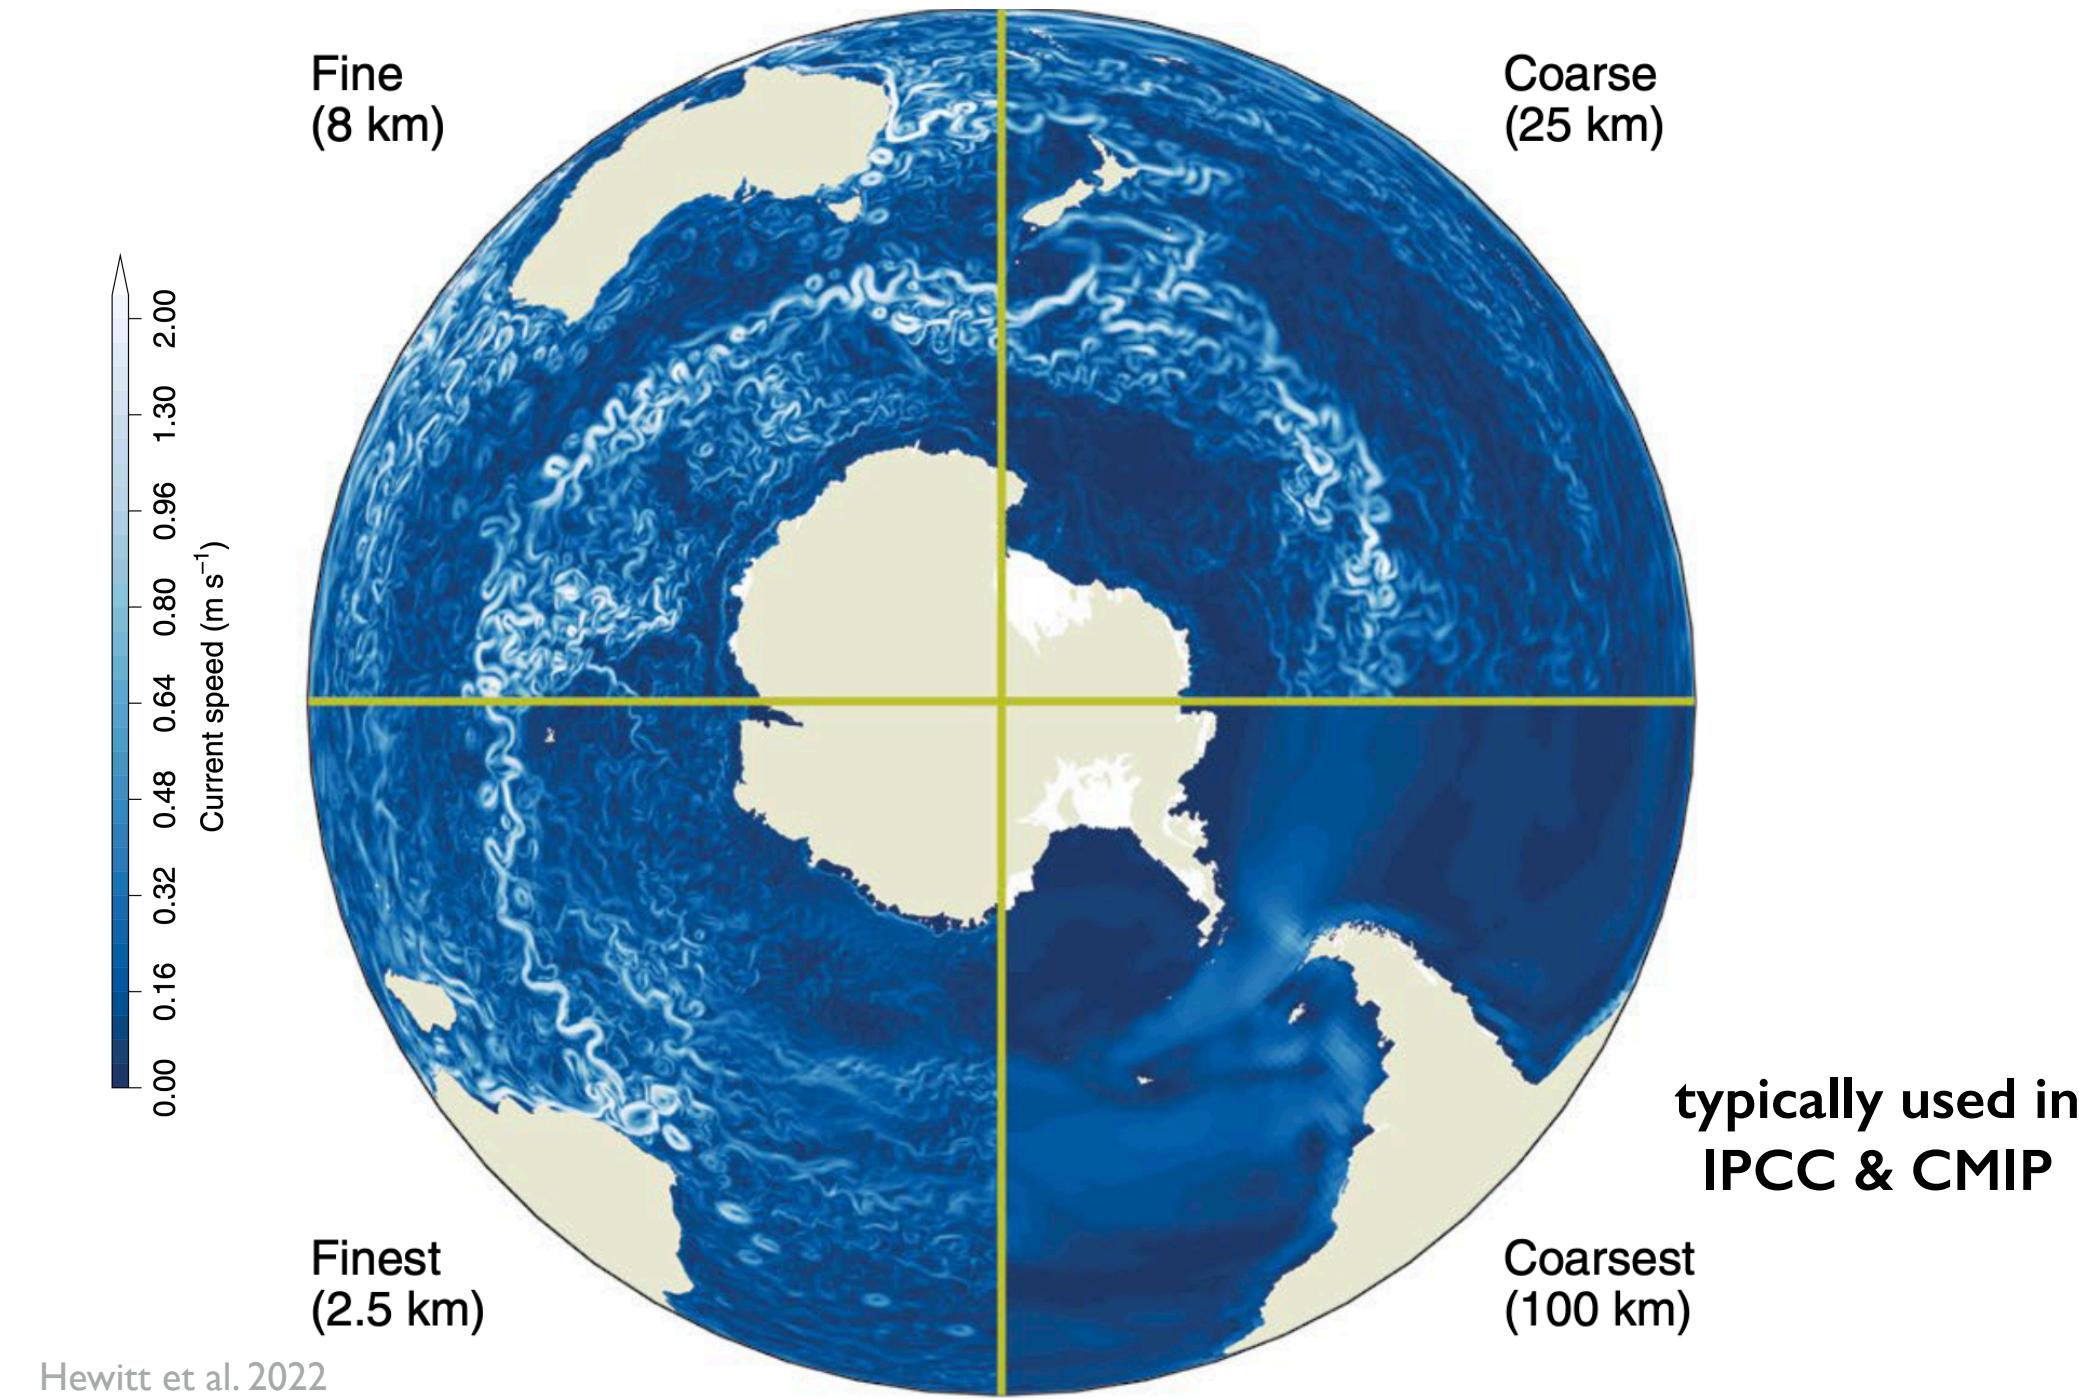
near-global ( $75^{\circ}\text{S}$ – $75^{\circ}\text{N}$ ) ocean simulation at  $1/12^{\circ}$  horizontal resolution, 48 vertical levels @ 68 Nvidia A100 achieving 10 simulated years per day

# breakthrough resolution — why do we keep pushing for higher resolution?

turbulent flows develop fine-scale features... the more we resolve, the more we see



remember Adele's talk on Wed



flows at metre-scales or even 10-100km (eddies ⚡) shape the global ocean circulation at century timescales!

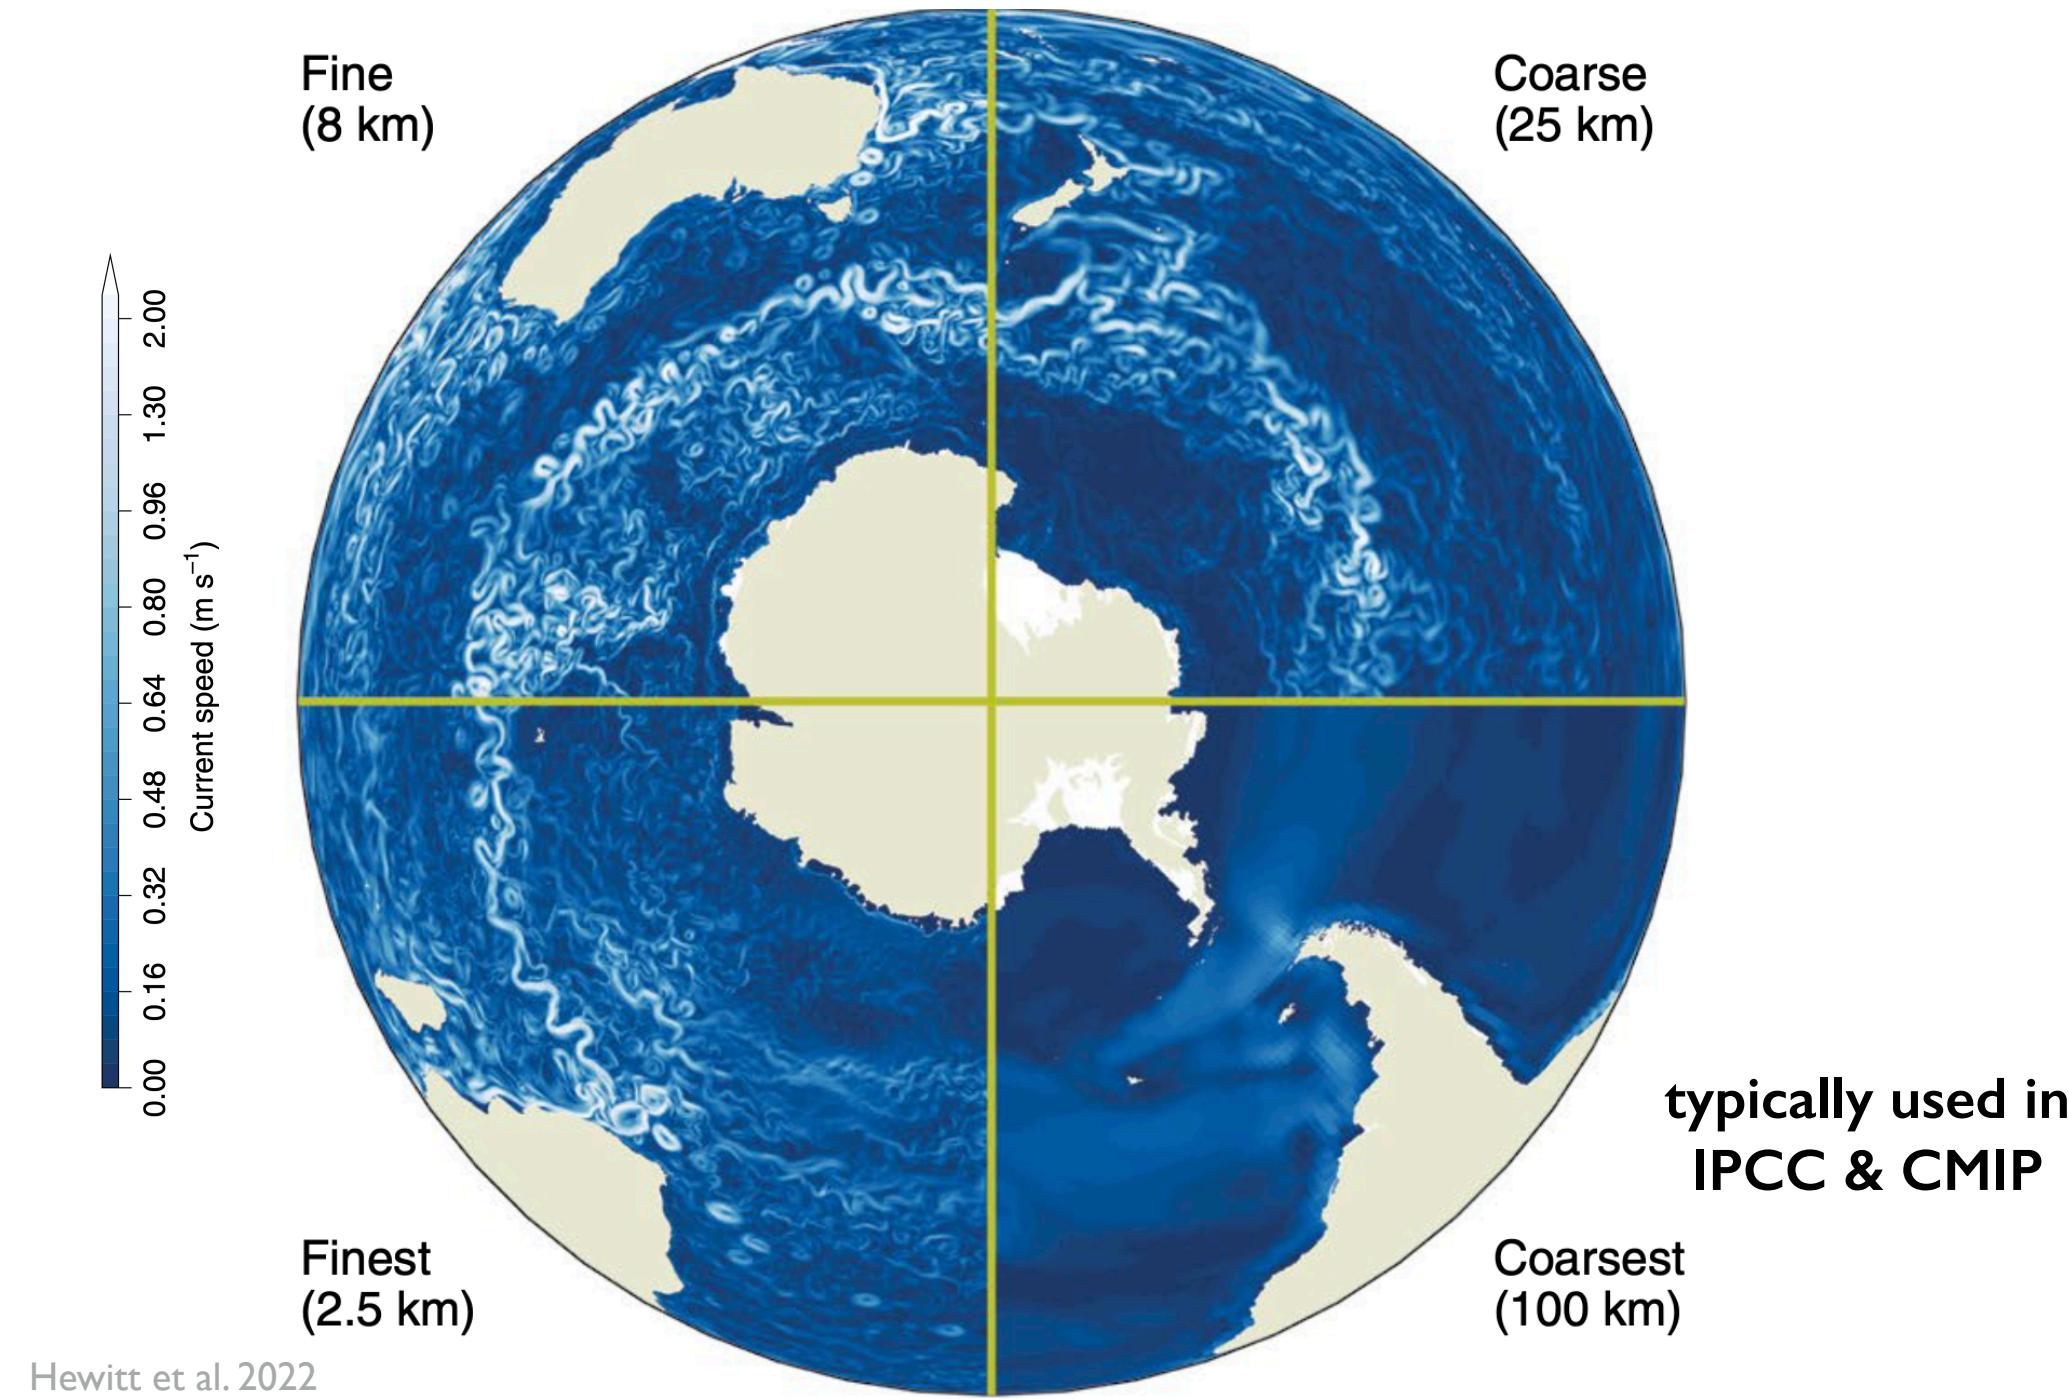
small and short-lived flow features affect global circulation over decades-centuries

# breakthrough resolution — why do we keep pushing for higher resolution?

turbulent flows develop fine-scale features... the more we resolve, the more we see



remember Adele's talk on Wed



*“From little things big things grow”*  
Paul Kelly

*“From little eddies large-scale ocean circulation grows”*  
adage

flows at metre-scales or even 10-100km (eddies ⚡) shape the global ocean circulation at century timescales!

small and short-lived flow features affect global circulation over decades-centuries

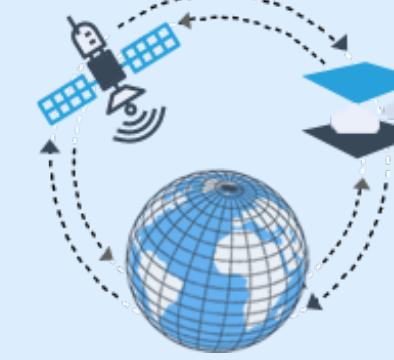
# Climate Modeling Alliance



## ClimateModeling v2.0

*...building a new Earth system model*

*that leverages recent advances in the computational and data sciences  
to learn directly from a wealth of Earth observations from space and the ground.*



Andre Souza



Chris Hill



Simone Silvestri



Greg Wagner



Raffaele Ferrari



Jean Michel  
Campin



Navid  
Constantinou



John Marshall

## Ocean Model Dev Team



Grace O'Neil



Xin-Kai Lee



M. G.



Ali Ramadhan



Ulyana Piterbarg



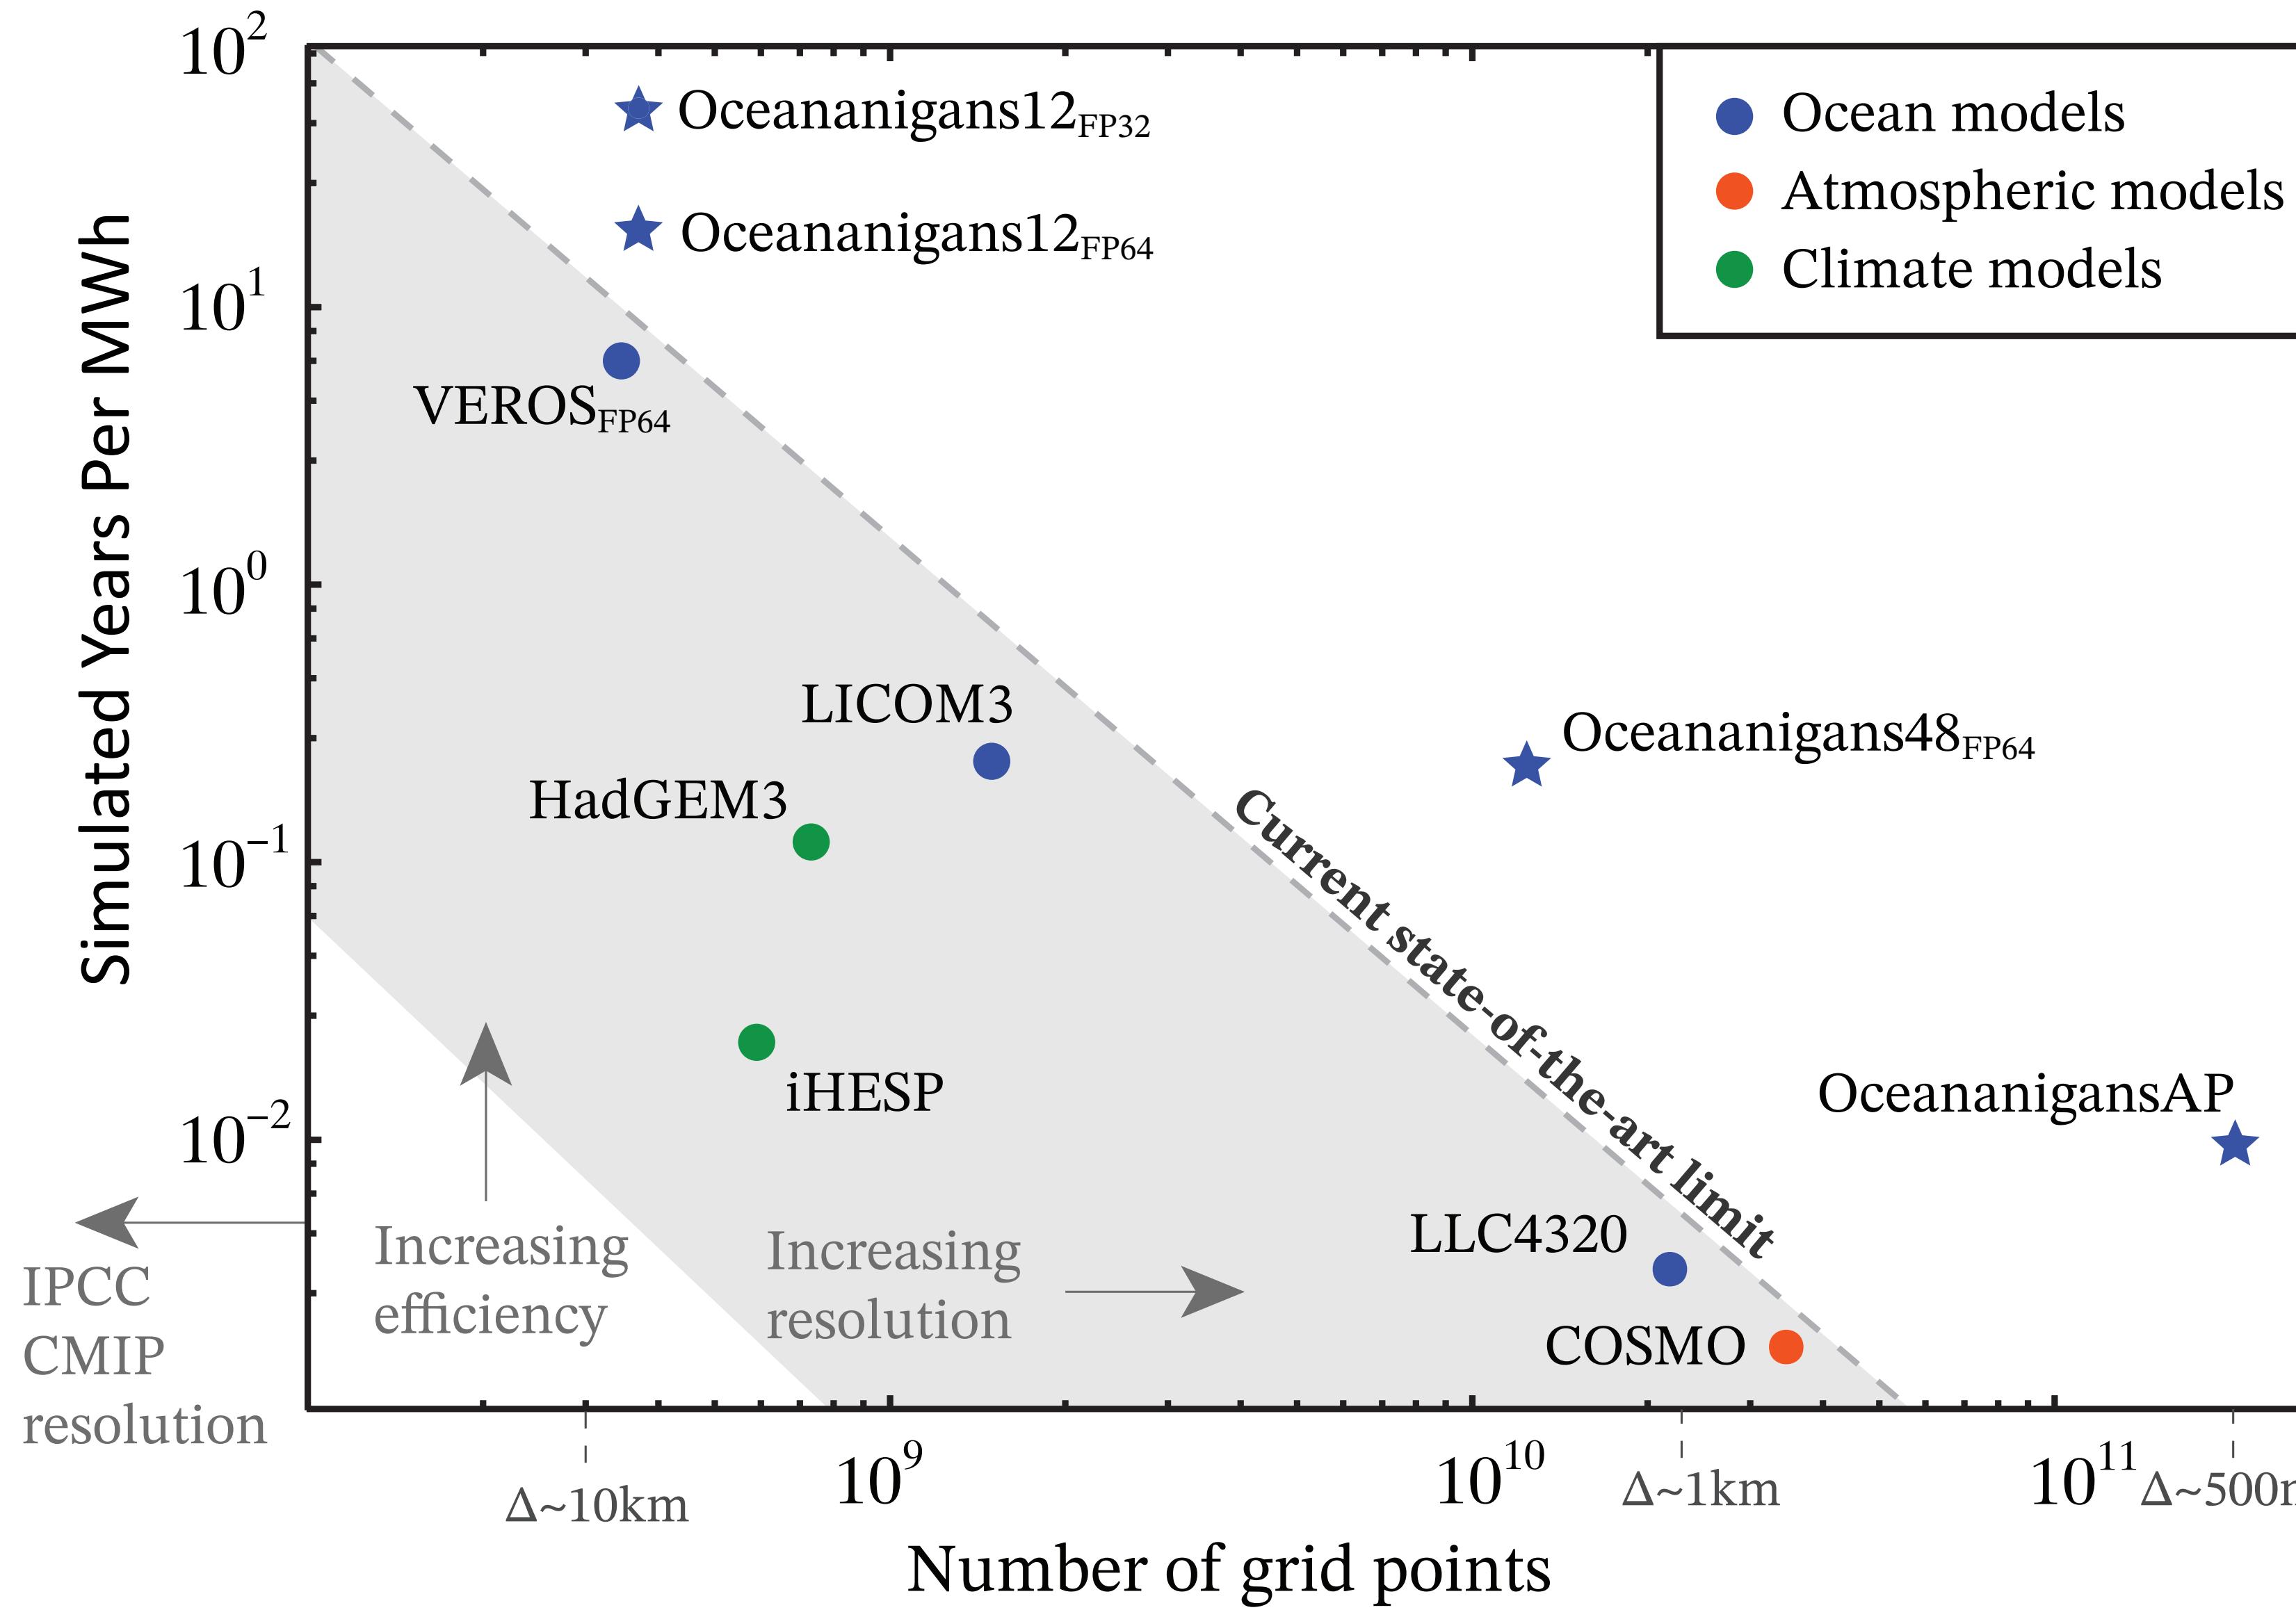
Adeline Hillier



Valentin Churavy

Oceananigans is the ocean core of Clima's Earth System model

# pushing boundaries: Oceananigans on the big scene



near-global ( $75^\circ\text{S}$ – $75^\circ\text{N}$ )

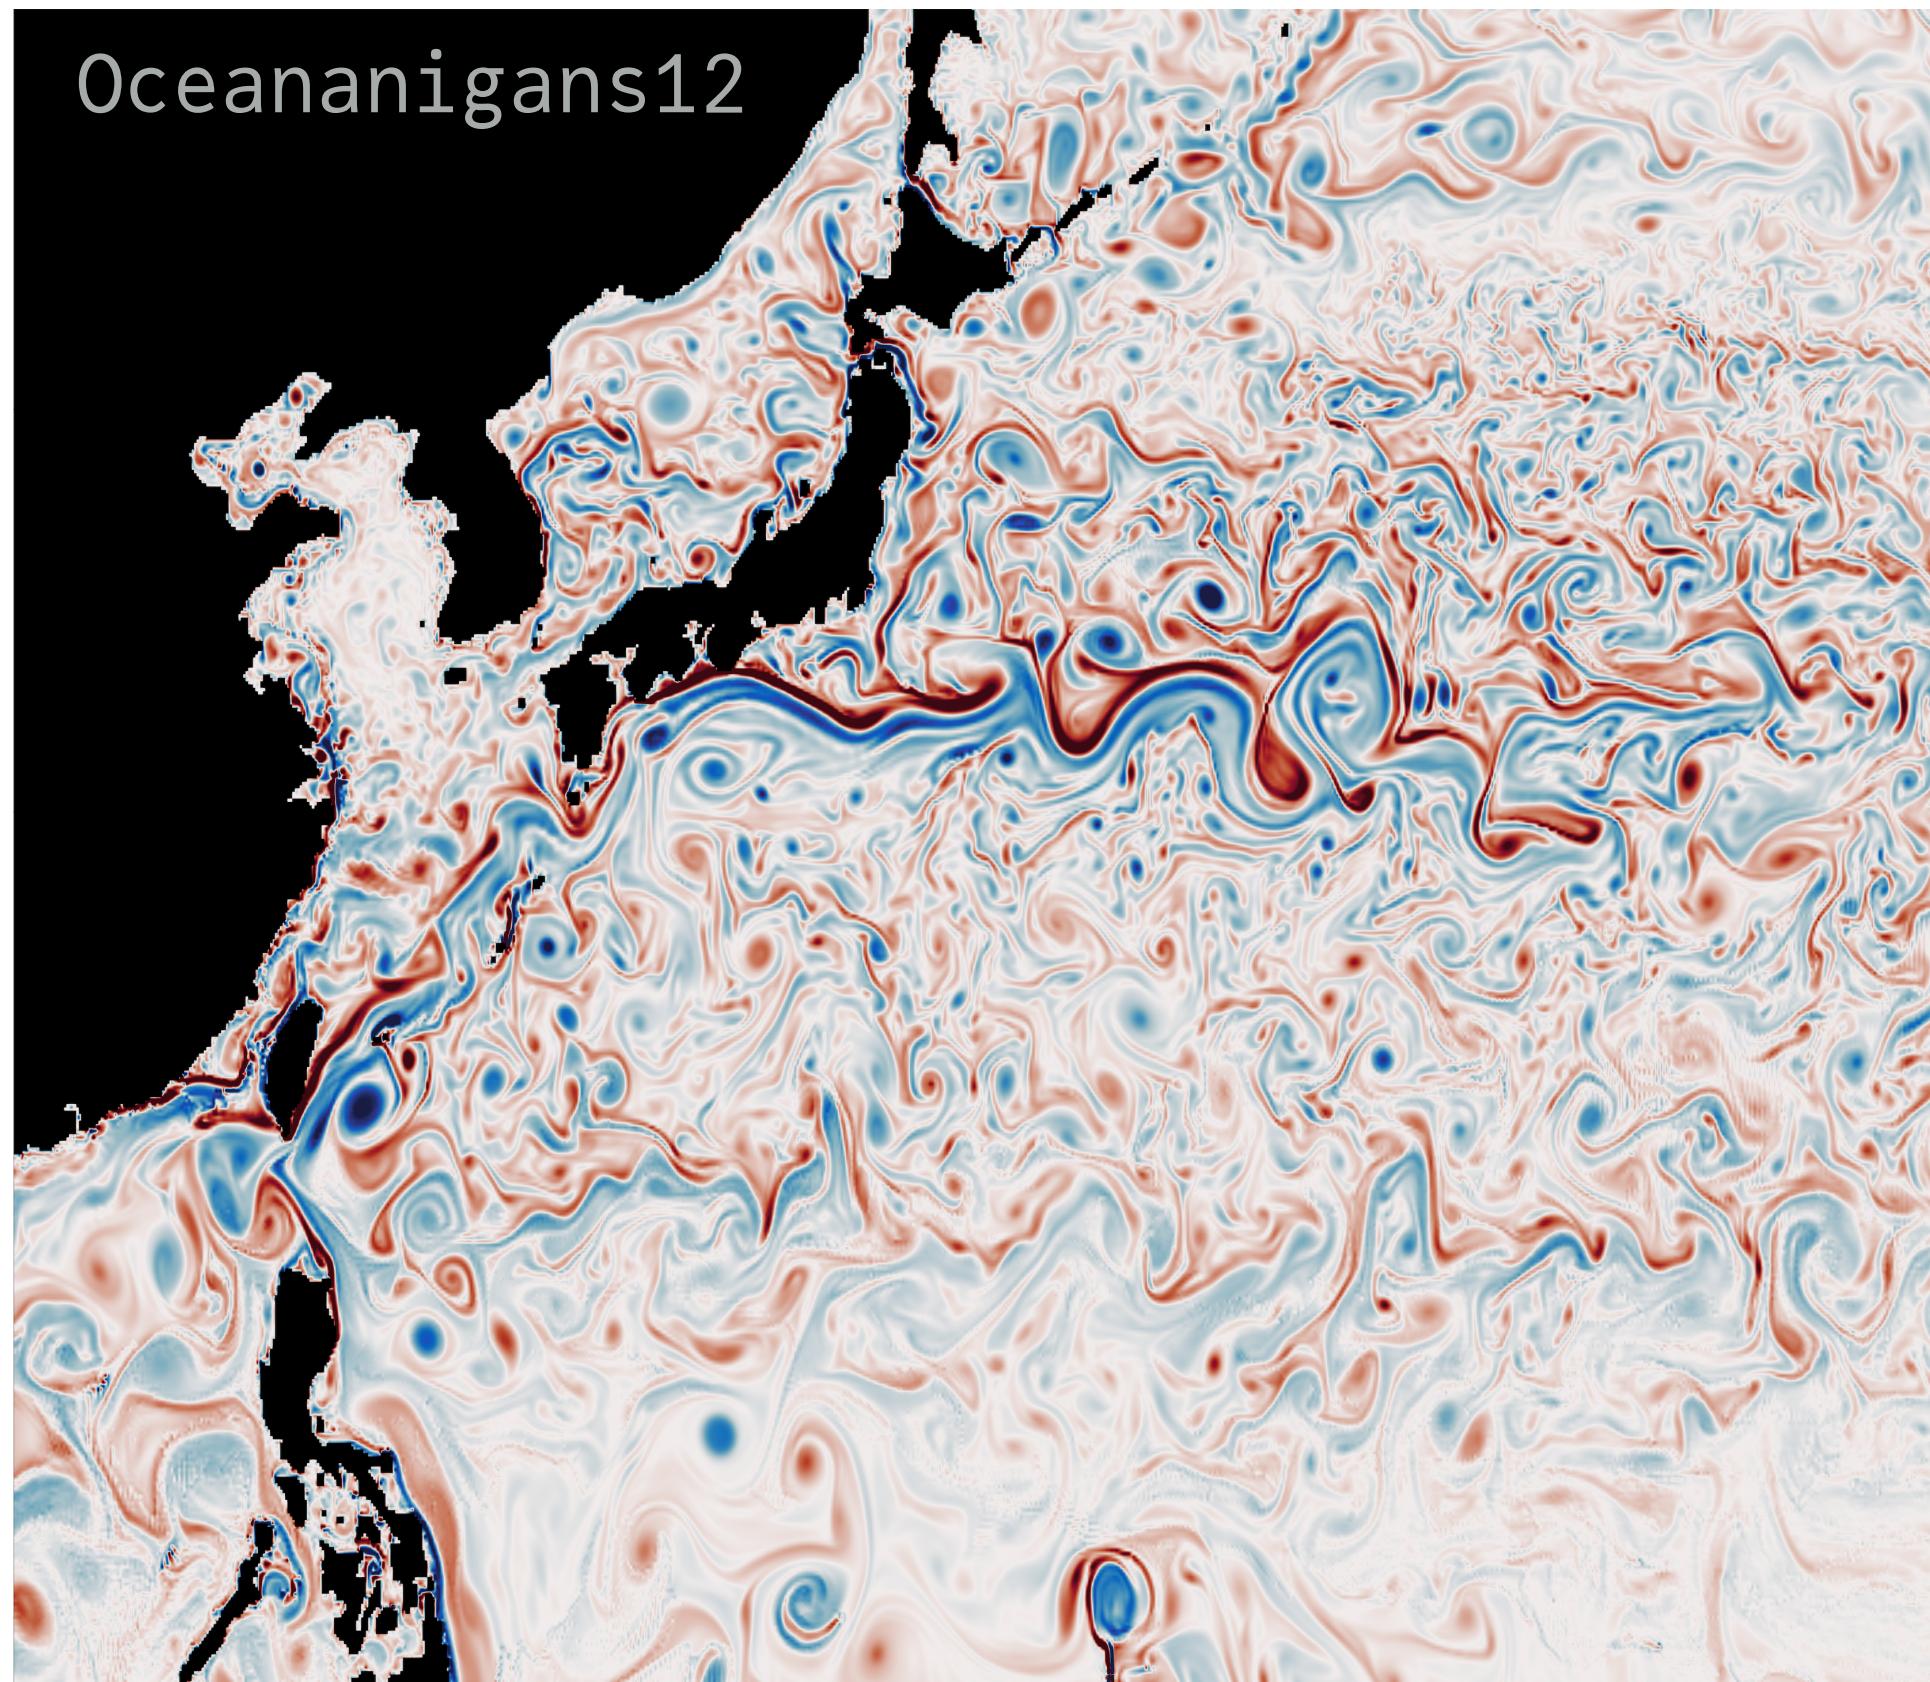
★ Oceananigans12  
 $1/12^\circ$  horizontal resolution  
48 vertical levels  
(video at title slide)

★ Oceananigans48  
 $1/48^\circ$  horizontal resolution  
100 vertical levels

★ OceananigansAP  
idealised Aqua-Planet + 2-continent config  
 $1/192^\circ$  horizontal resolution  
100 vertical levels

# breakthrough energy efficiency

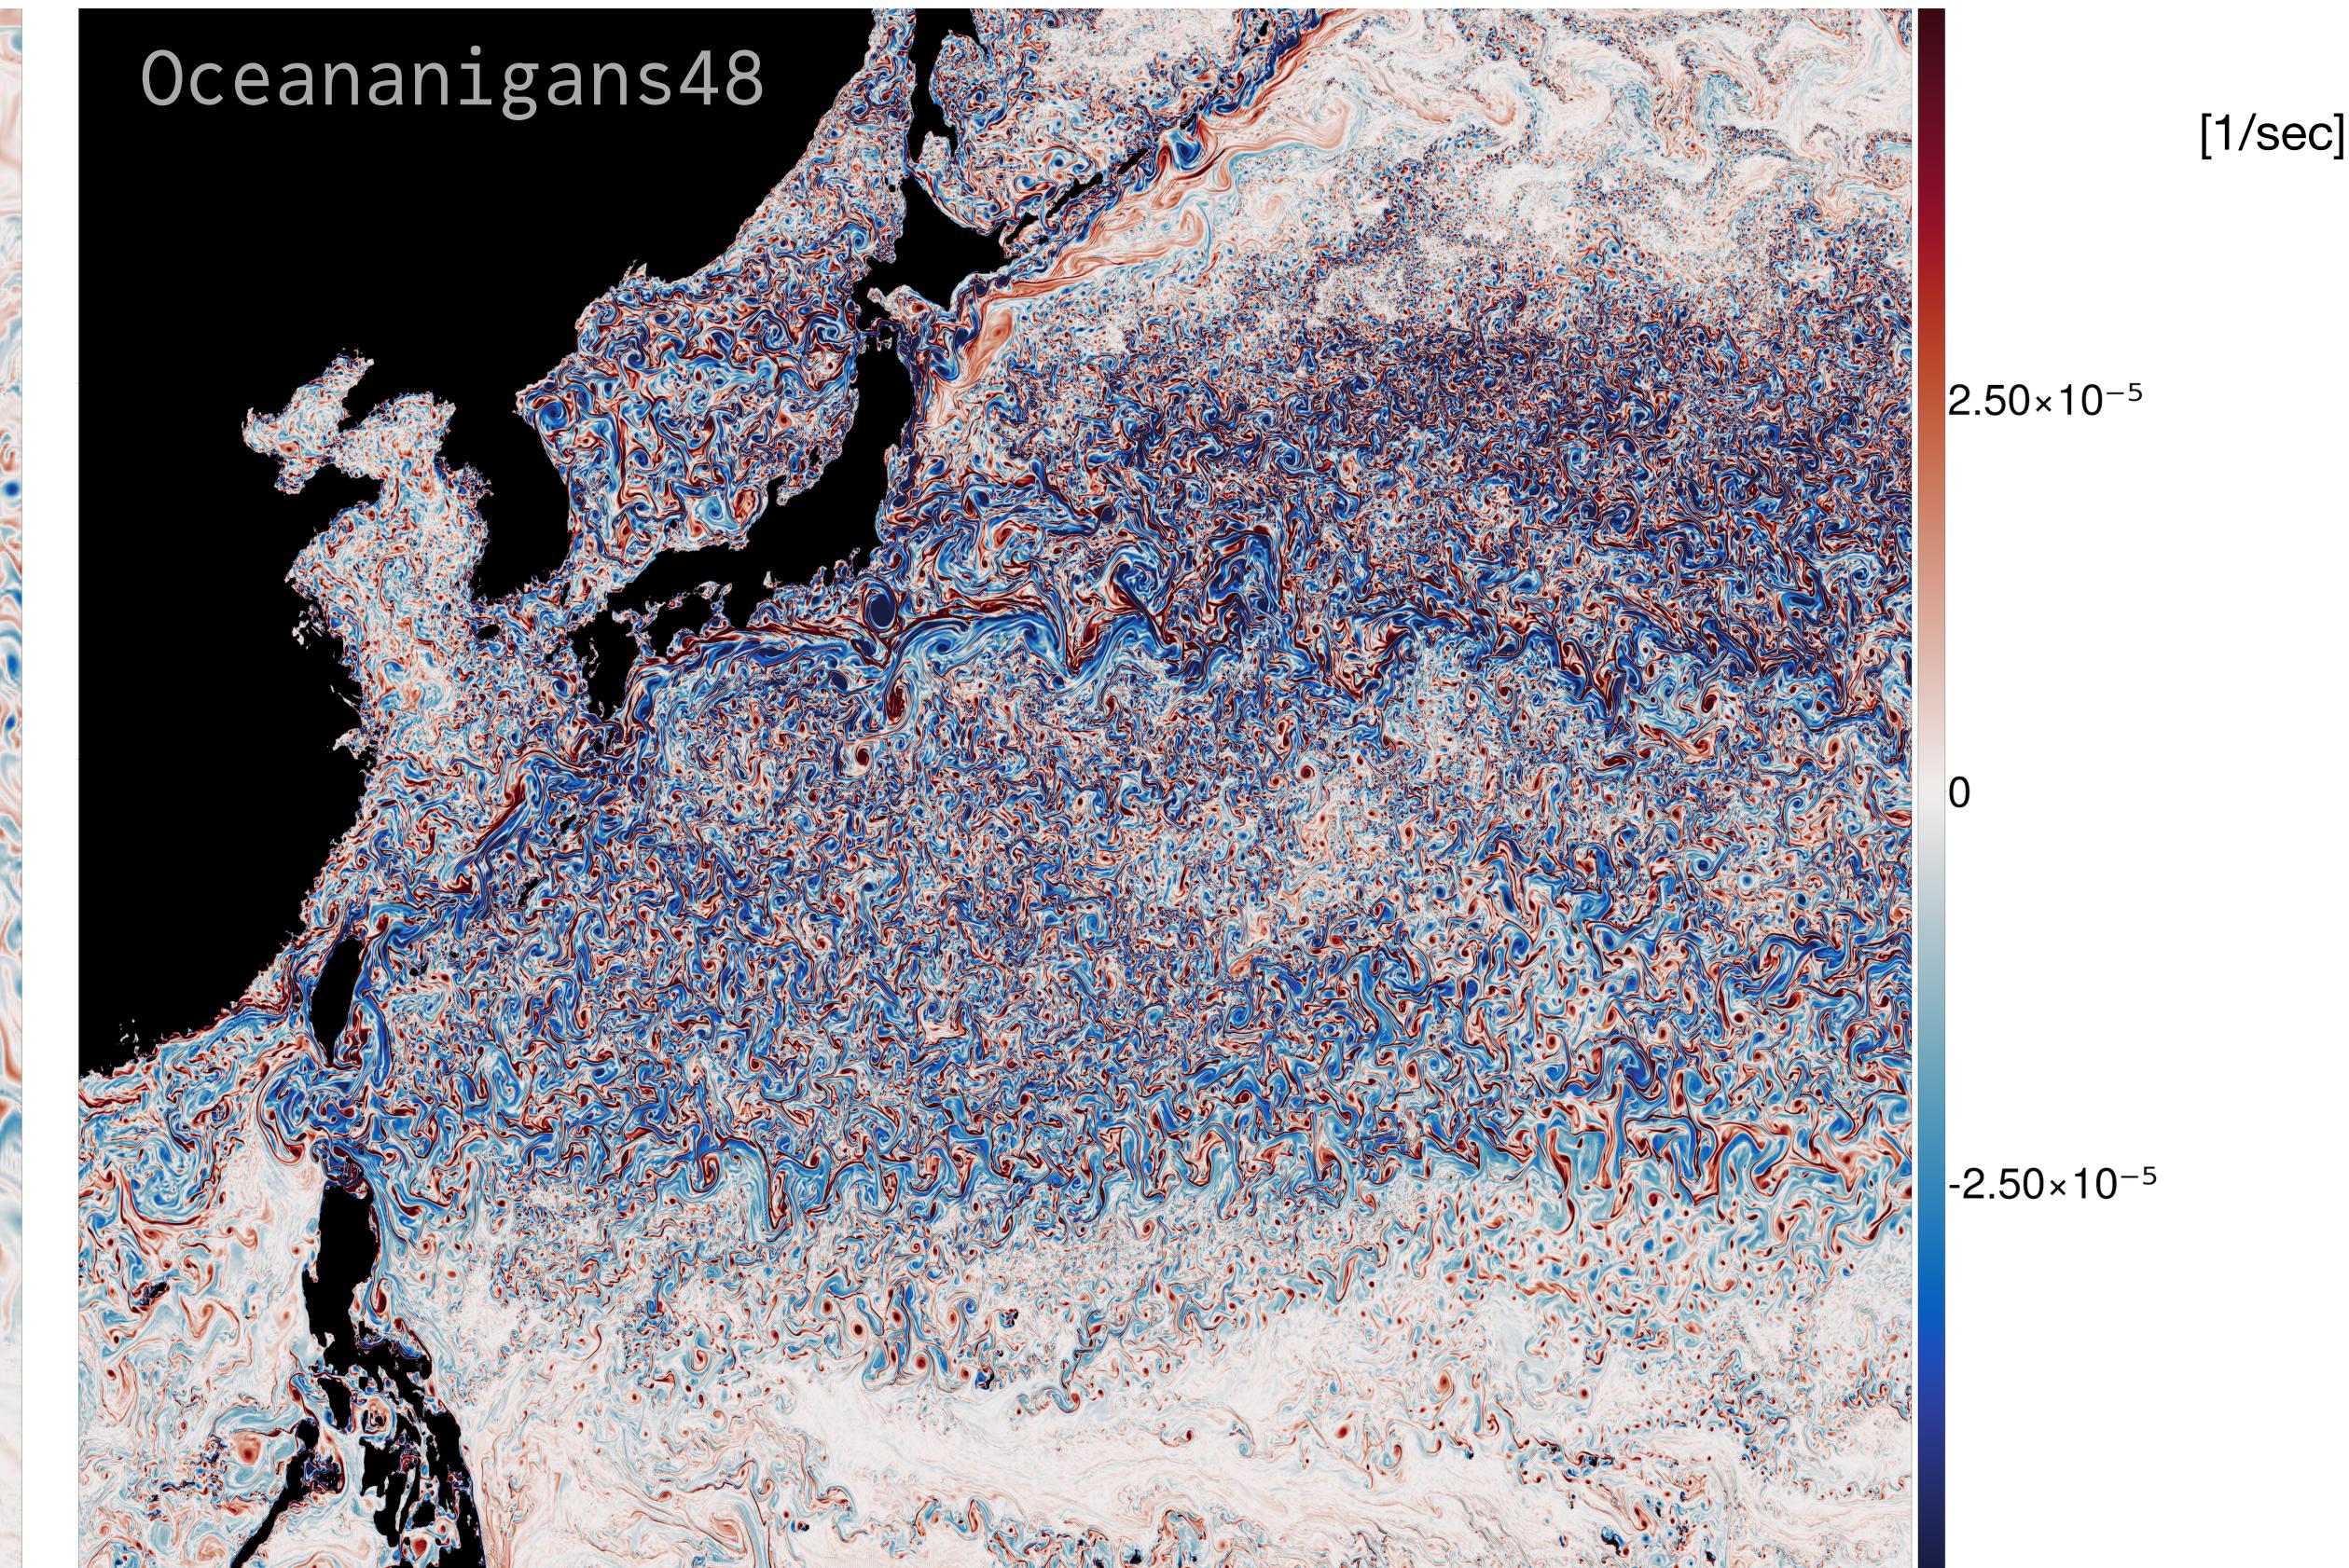
## surface relative vorticity



$1/12^\circ$  horizontal resolution + 48 vertical levels  
 $\sim 10^9$  grid points



10 Simulated Years Per Day  
@ 68 Nvidia A100 GPUs



$1/48^\circ$  horizontal resolution + 100 vertical levels  
 $\sim 5 \times 10^{10}$  points



1 Simulated Year Per Day  
@ 576 Nvidia A100 GPUs

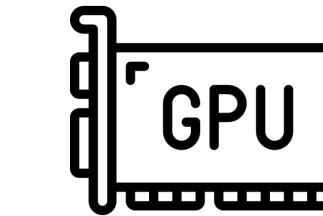


# but why is Oceananigans so *fast* and so *efficient*?



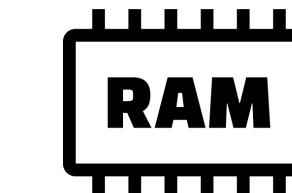
## fast compute

Oceananigans is written from scratch for GPUs  
(we didn't port a CPU-based model onto GPU)



## memory leanness

minimise temporary array creations  
loop over the domain as few times as possible



## scalability

compute is fast, communication is expensive  
novel algorithm for fast 2D motions (barotropic solver) allows  
effective overlap of communication & computation

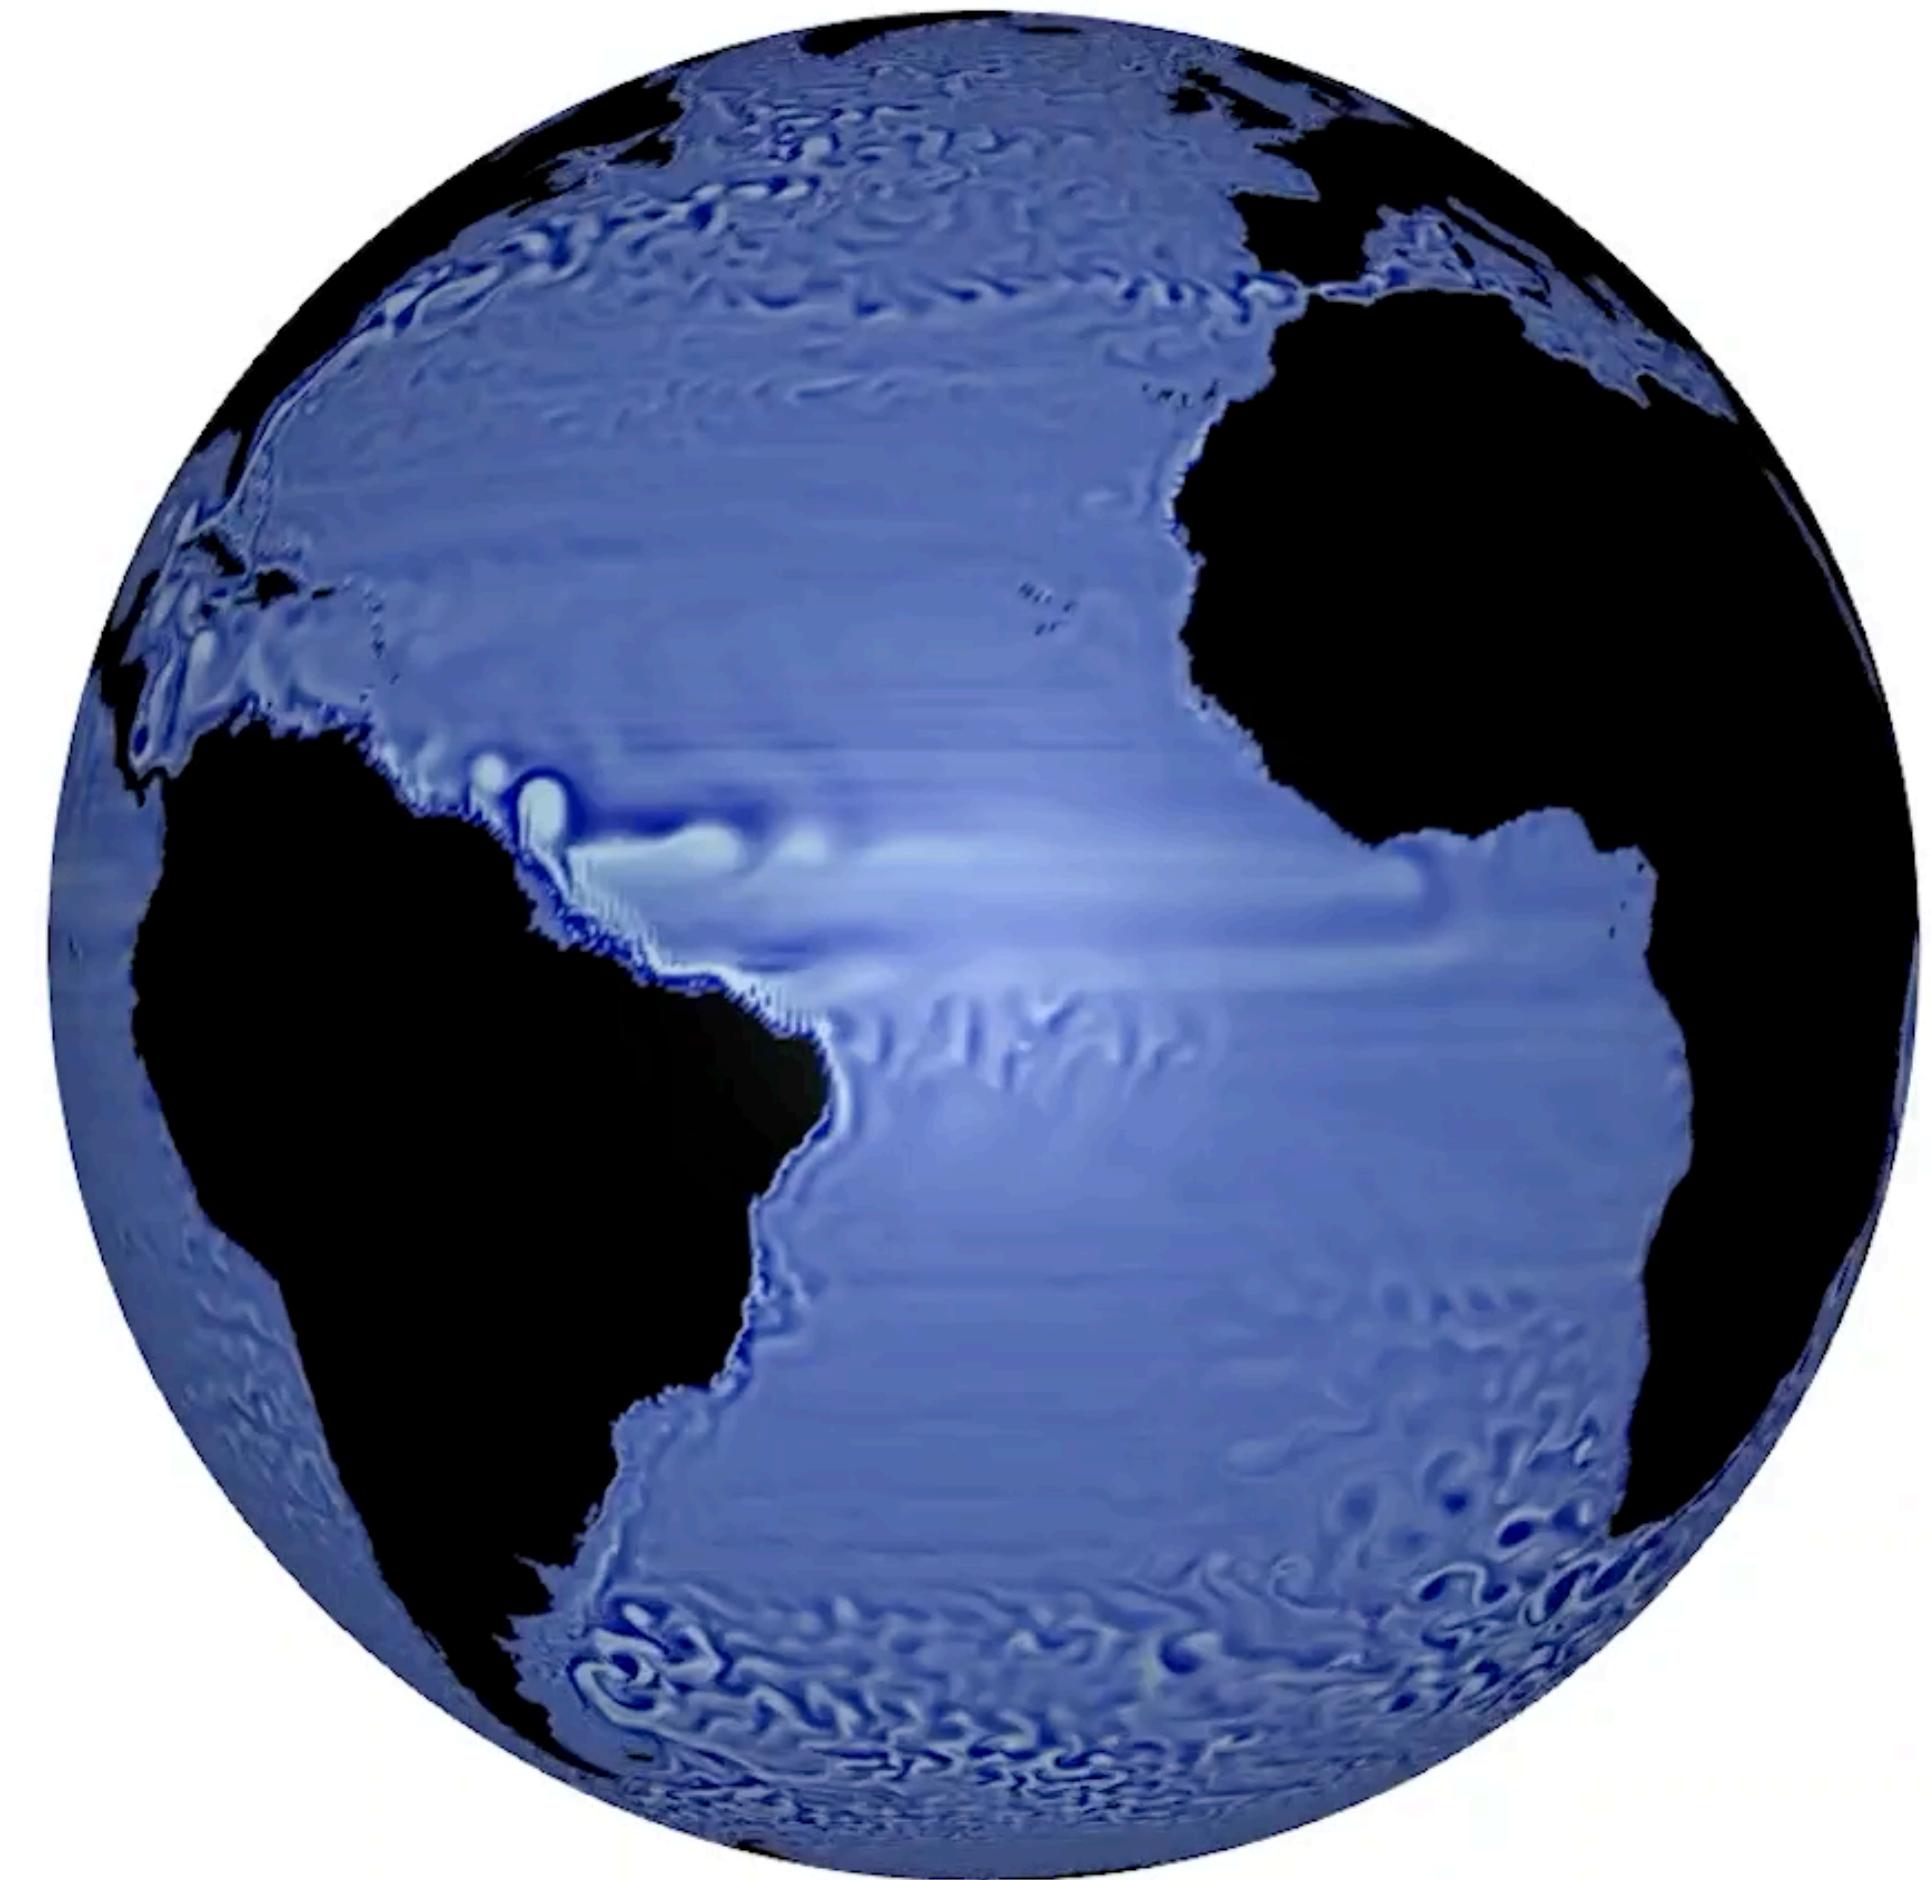


# breakthrough memory efficiency

eddy permitting  
 $1/4^\circ$  horizontal resolution  
48 vertical levels

10 GB memory footprint

fits **easily** on 1 Nvidia V100 GPU



surface relative vorticity



# memory leanness

advection-diffusion  
of temperature

“classic” Fortran-style  
temp arrays on CPU are cheap

$$\partial_t T = \underbrace{-\mathbf{u} \cdot \nabla T}_{\text{advection}} + \underbrace{\partial_z (\kappa \partial_z T)}_{\text{diffusion}}$$

rhs

temporary array → `u_nabla_T = calculate_advection(u, v, w, T)`  
temporary array → `diff_T = calculate_diffusion(kappa, T)`  
temporary array → `rhs_T = - u_nabla_T + diff_T`

# memory leanness

advection-diffusion  
of temperature

“classic” Fortran-style  
temp arrays on CPU are cheap

Oceananigans  
GPU-friendly kernel fusion  
no memory allocation

$$\partial_t T = \underbrace{-\mathbf{u} \cdot \nabla T}_{\text{advection}} + \underbrace{\partial_z (\kappa \partial_z T)}_{\text{diffusion}}$$

rhs

temporary array → `u_nabla_T = calculate_advection(u, v, w, T)`  
temporary array → `diff_T = calculate_diffusion(kappa, T)`  
temporary array → `rhs_T = - u_nabla_T + diff_T`

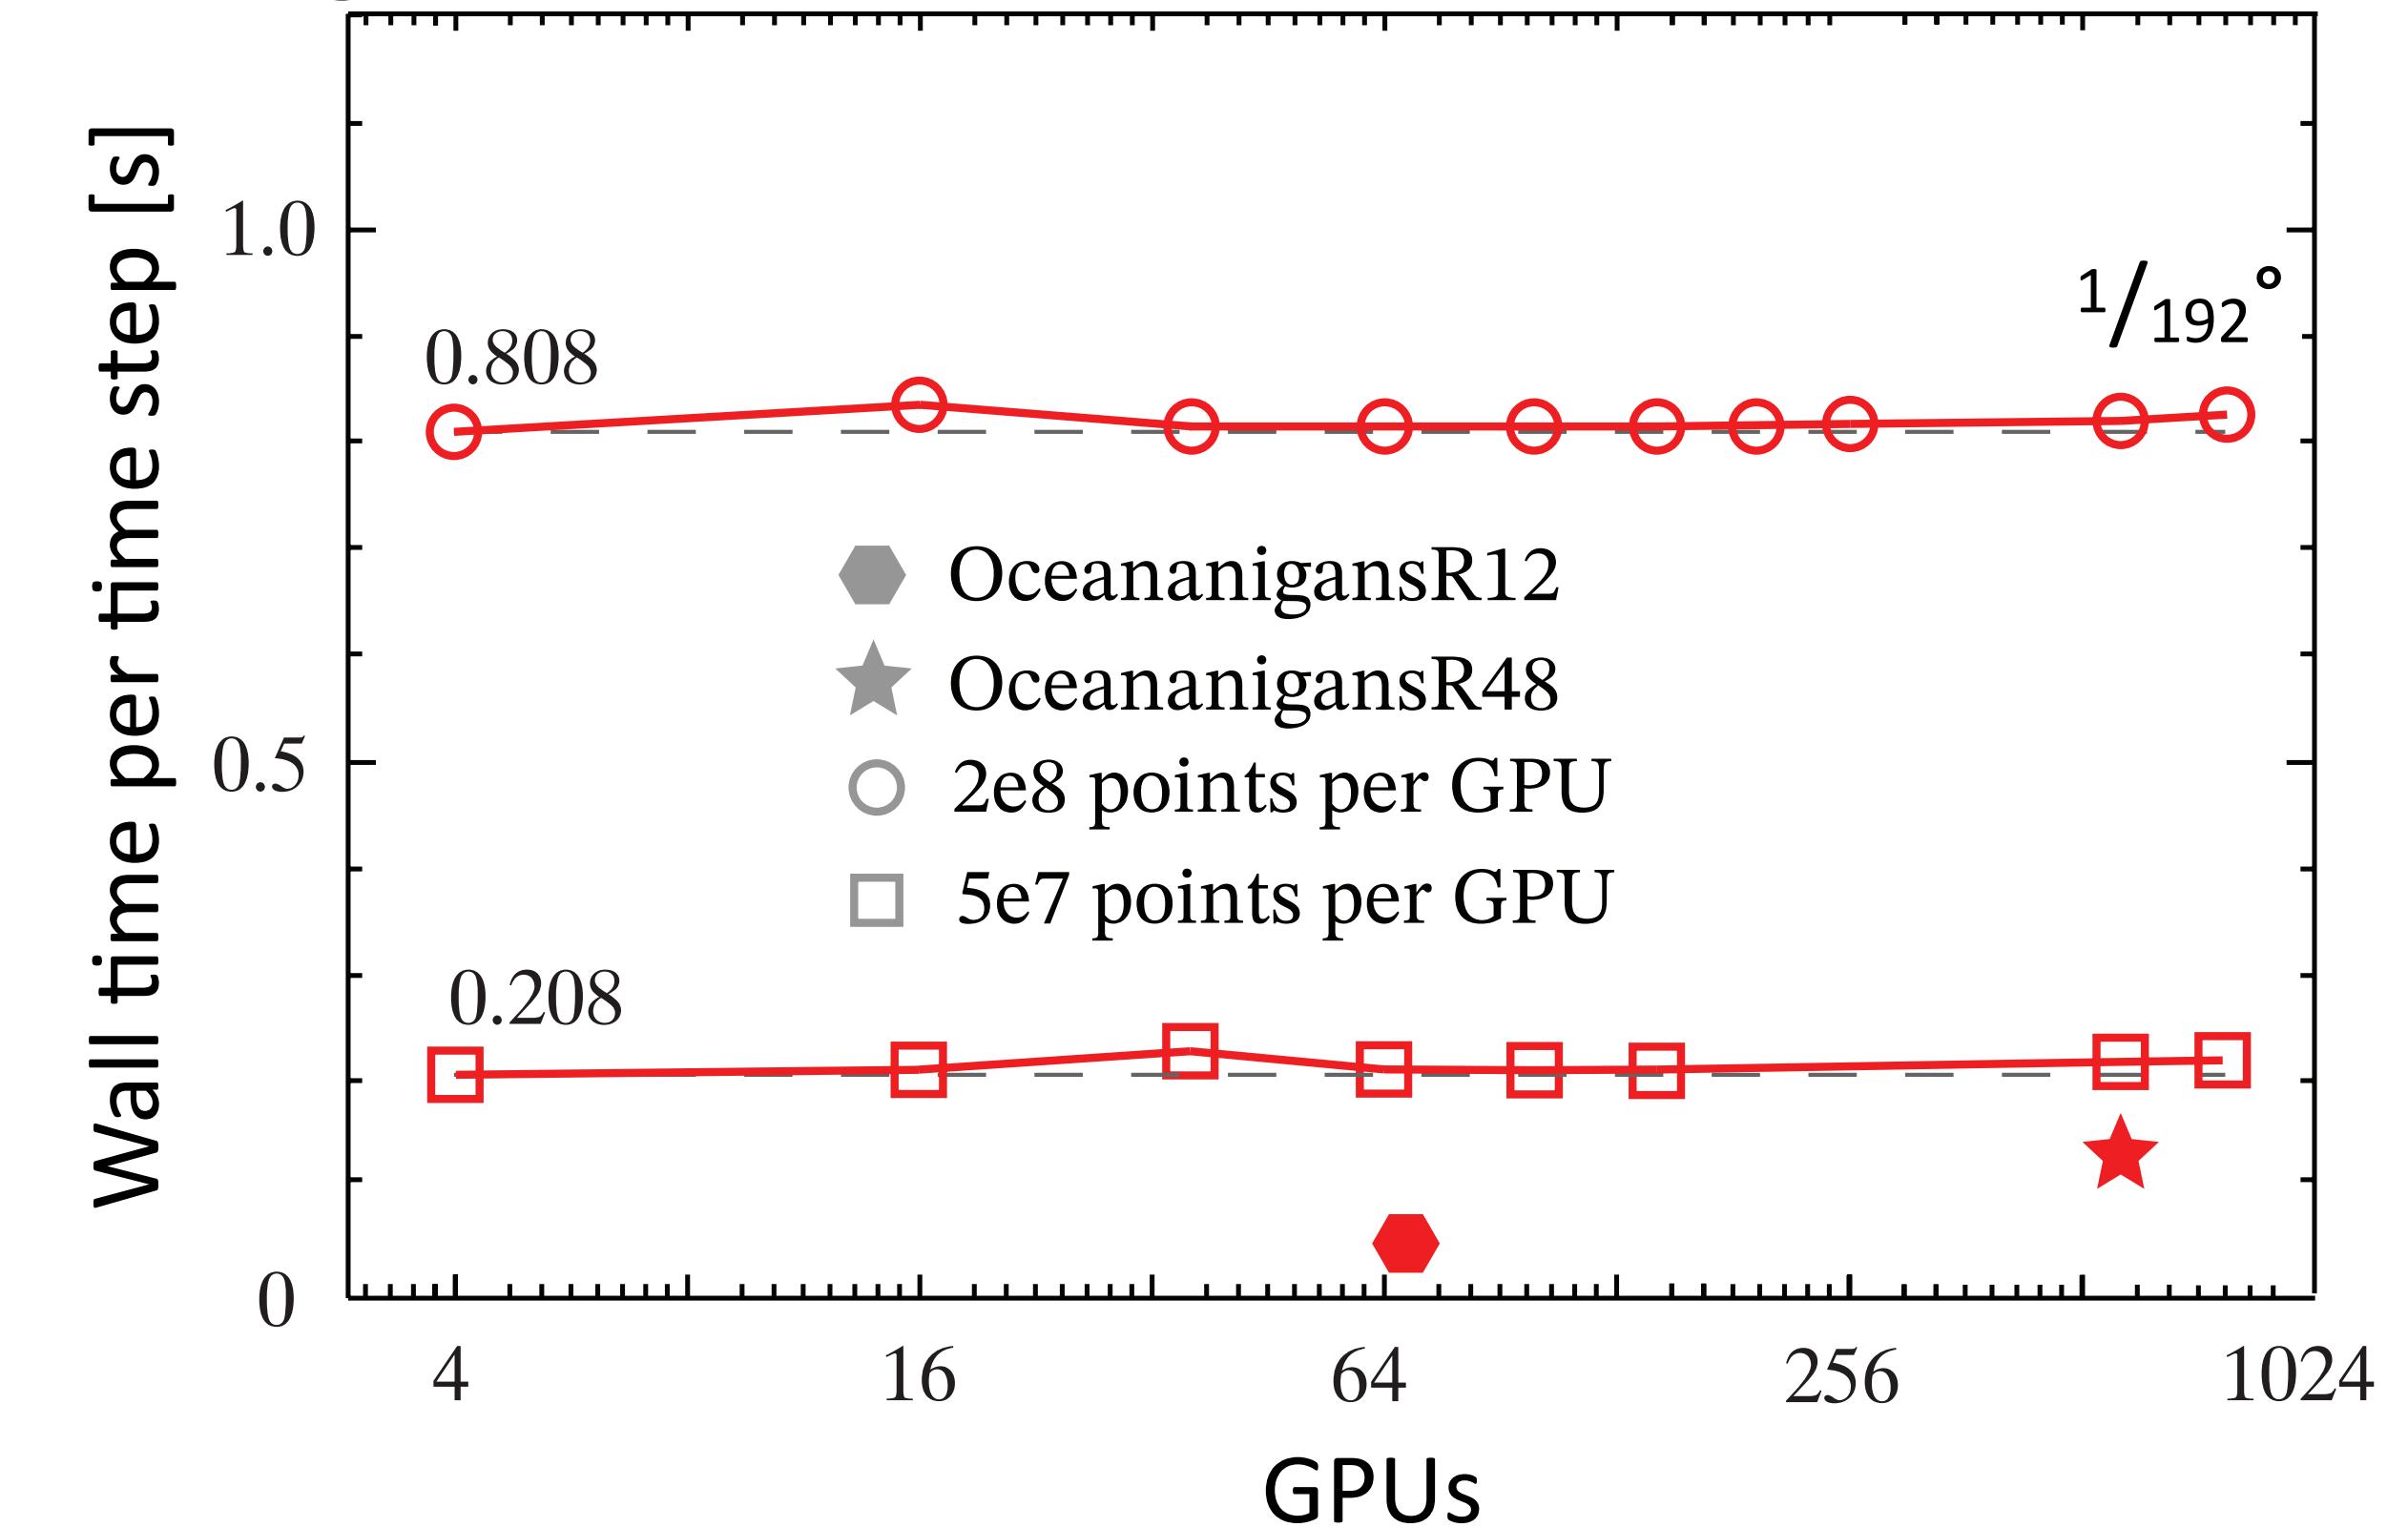
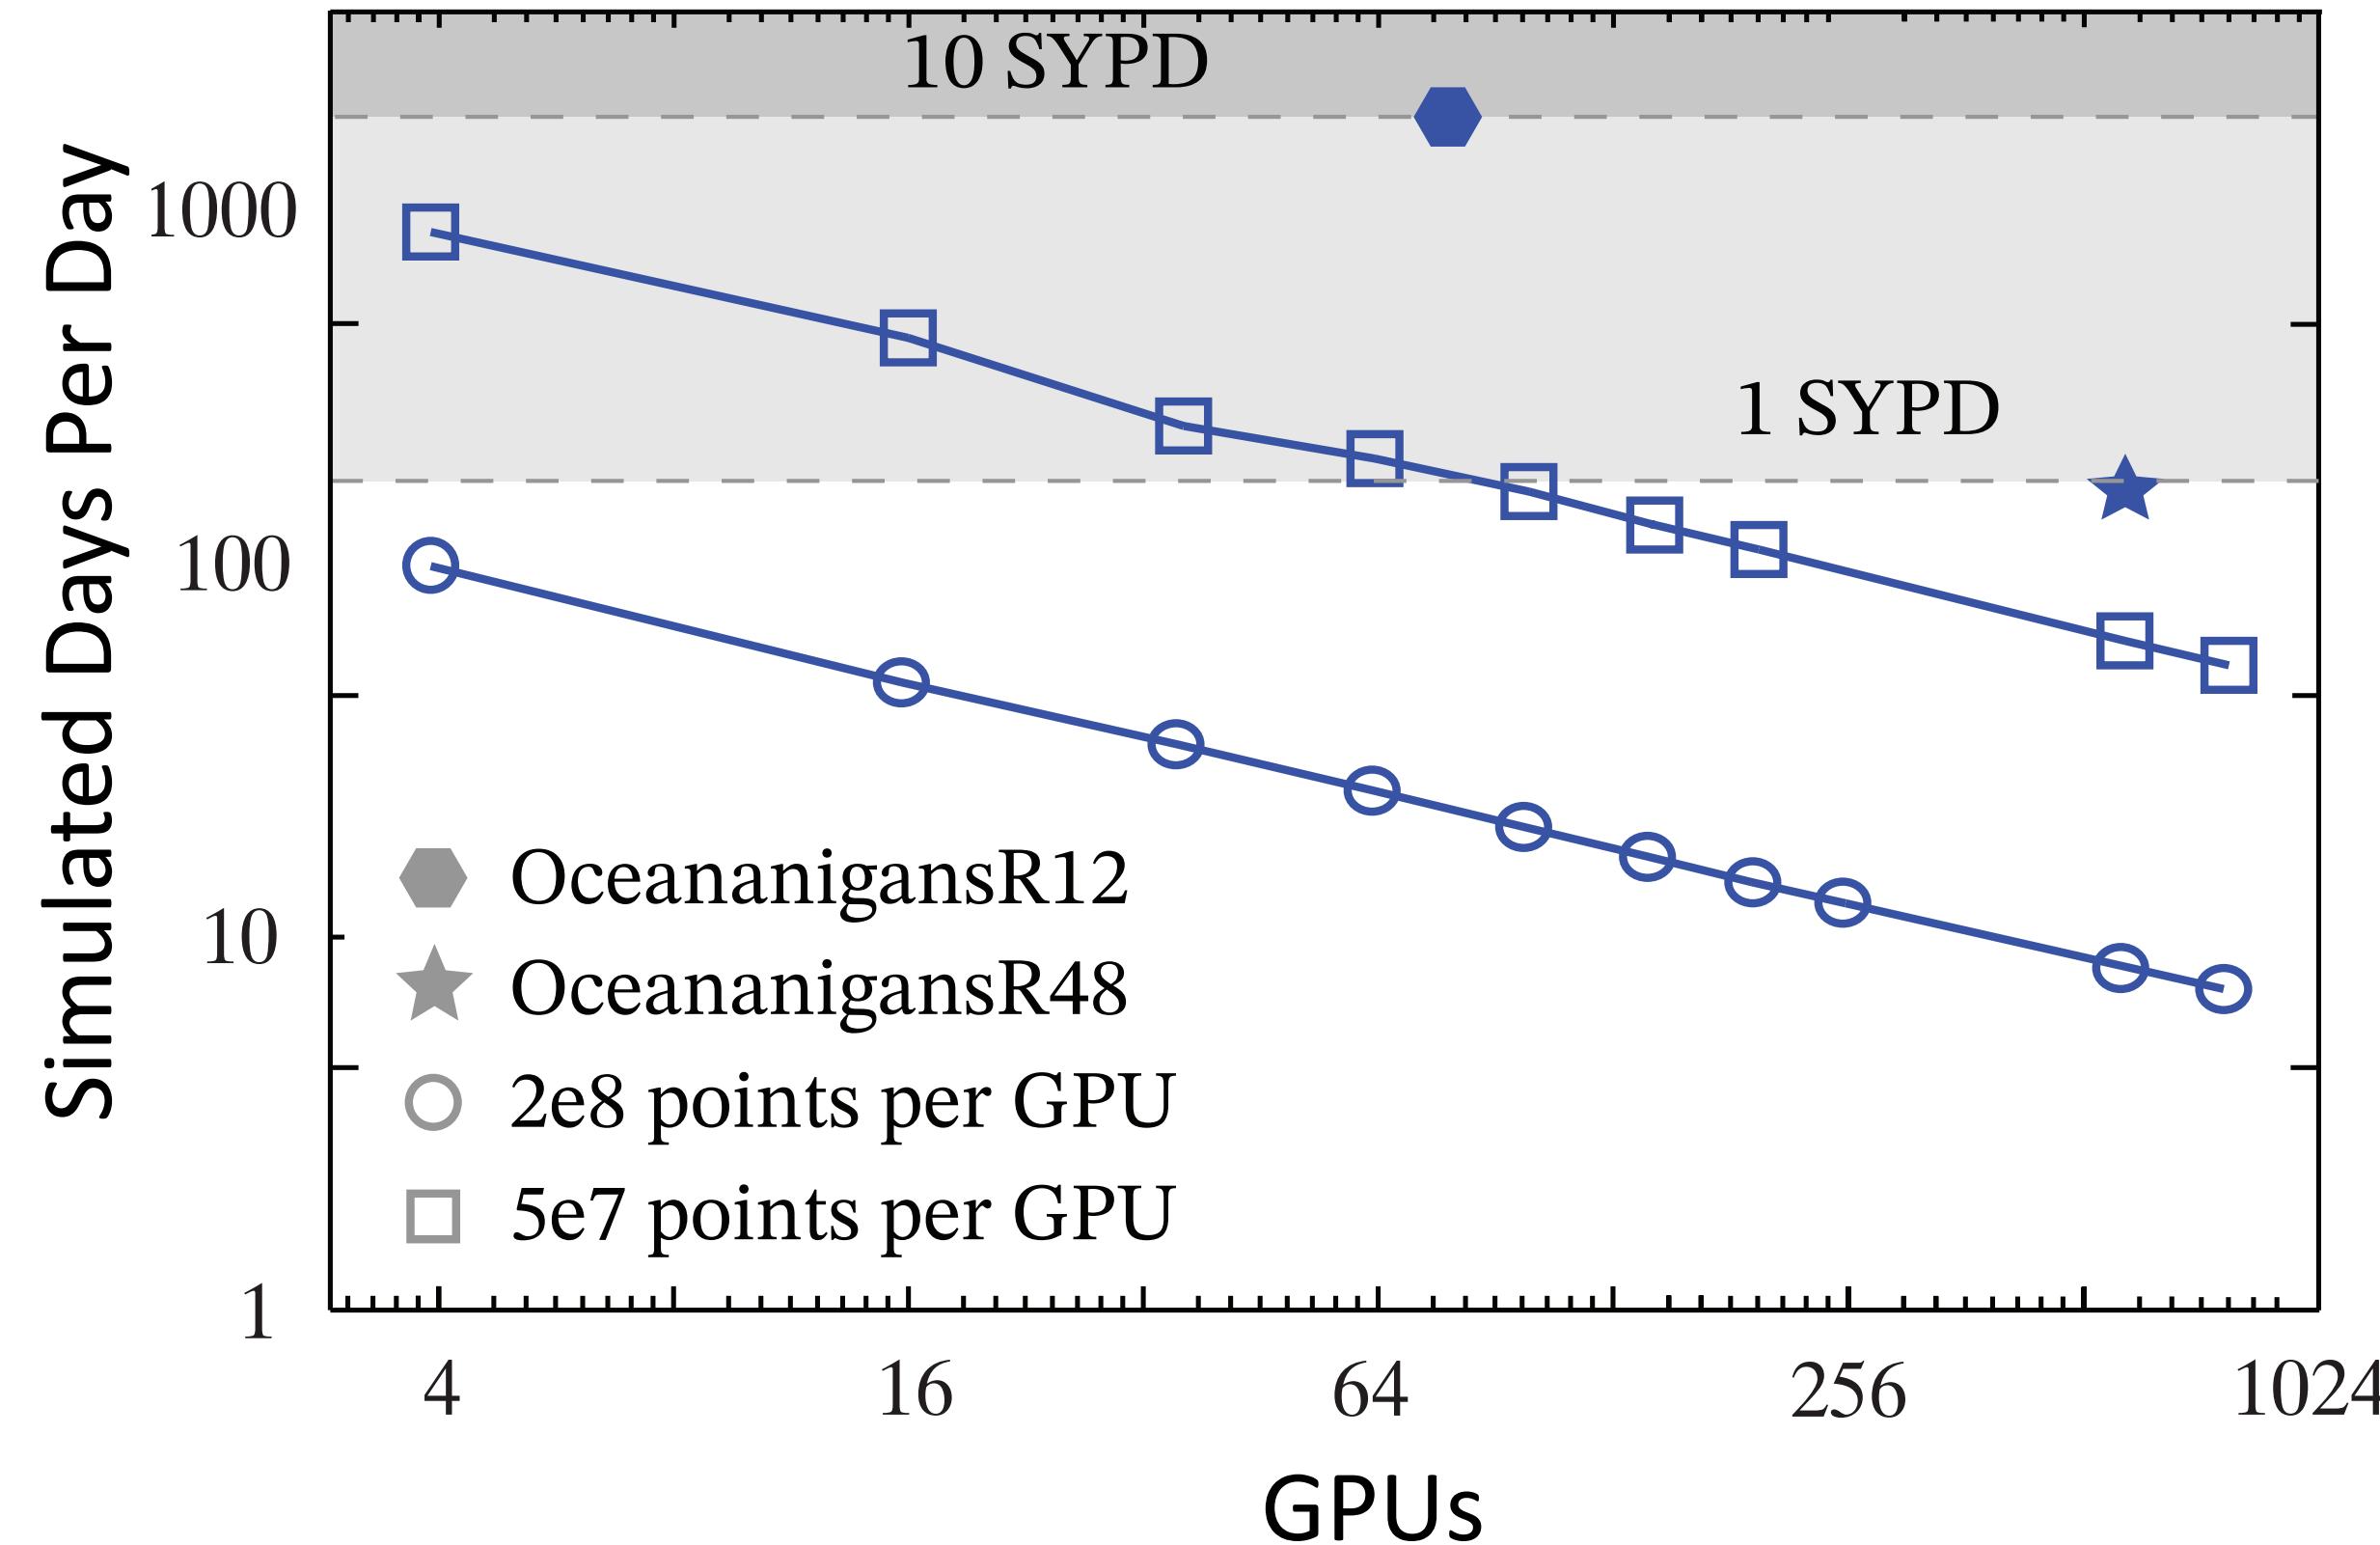
```
tendency_T(i, j, k) = - u_nabla_T(i, j, k) + diffusion(i, j, k)

@kernel function calculate_tendency_T!(rhs_T)
    i, j, k = @index(Global, NTuple)
    rhs_T[i, j, k] = tendency_T(i, j, k)
end
```

we load as few kernels as possible; only one for the tendency of each flow field  
we loop over the grid once for every flow field

# how do we scale?

## weak scalings



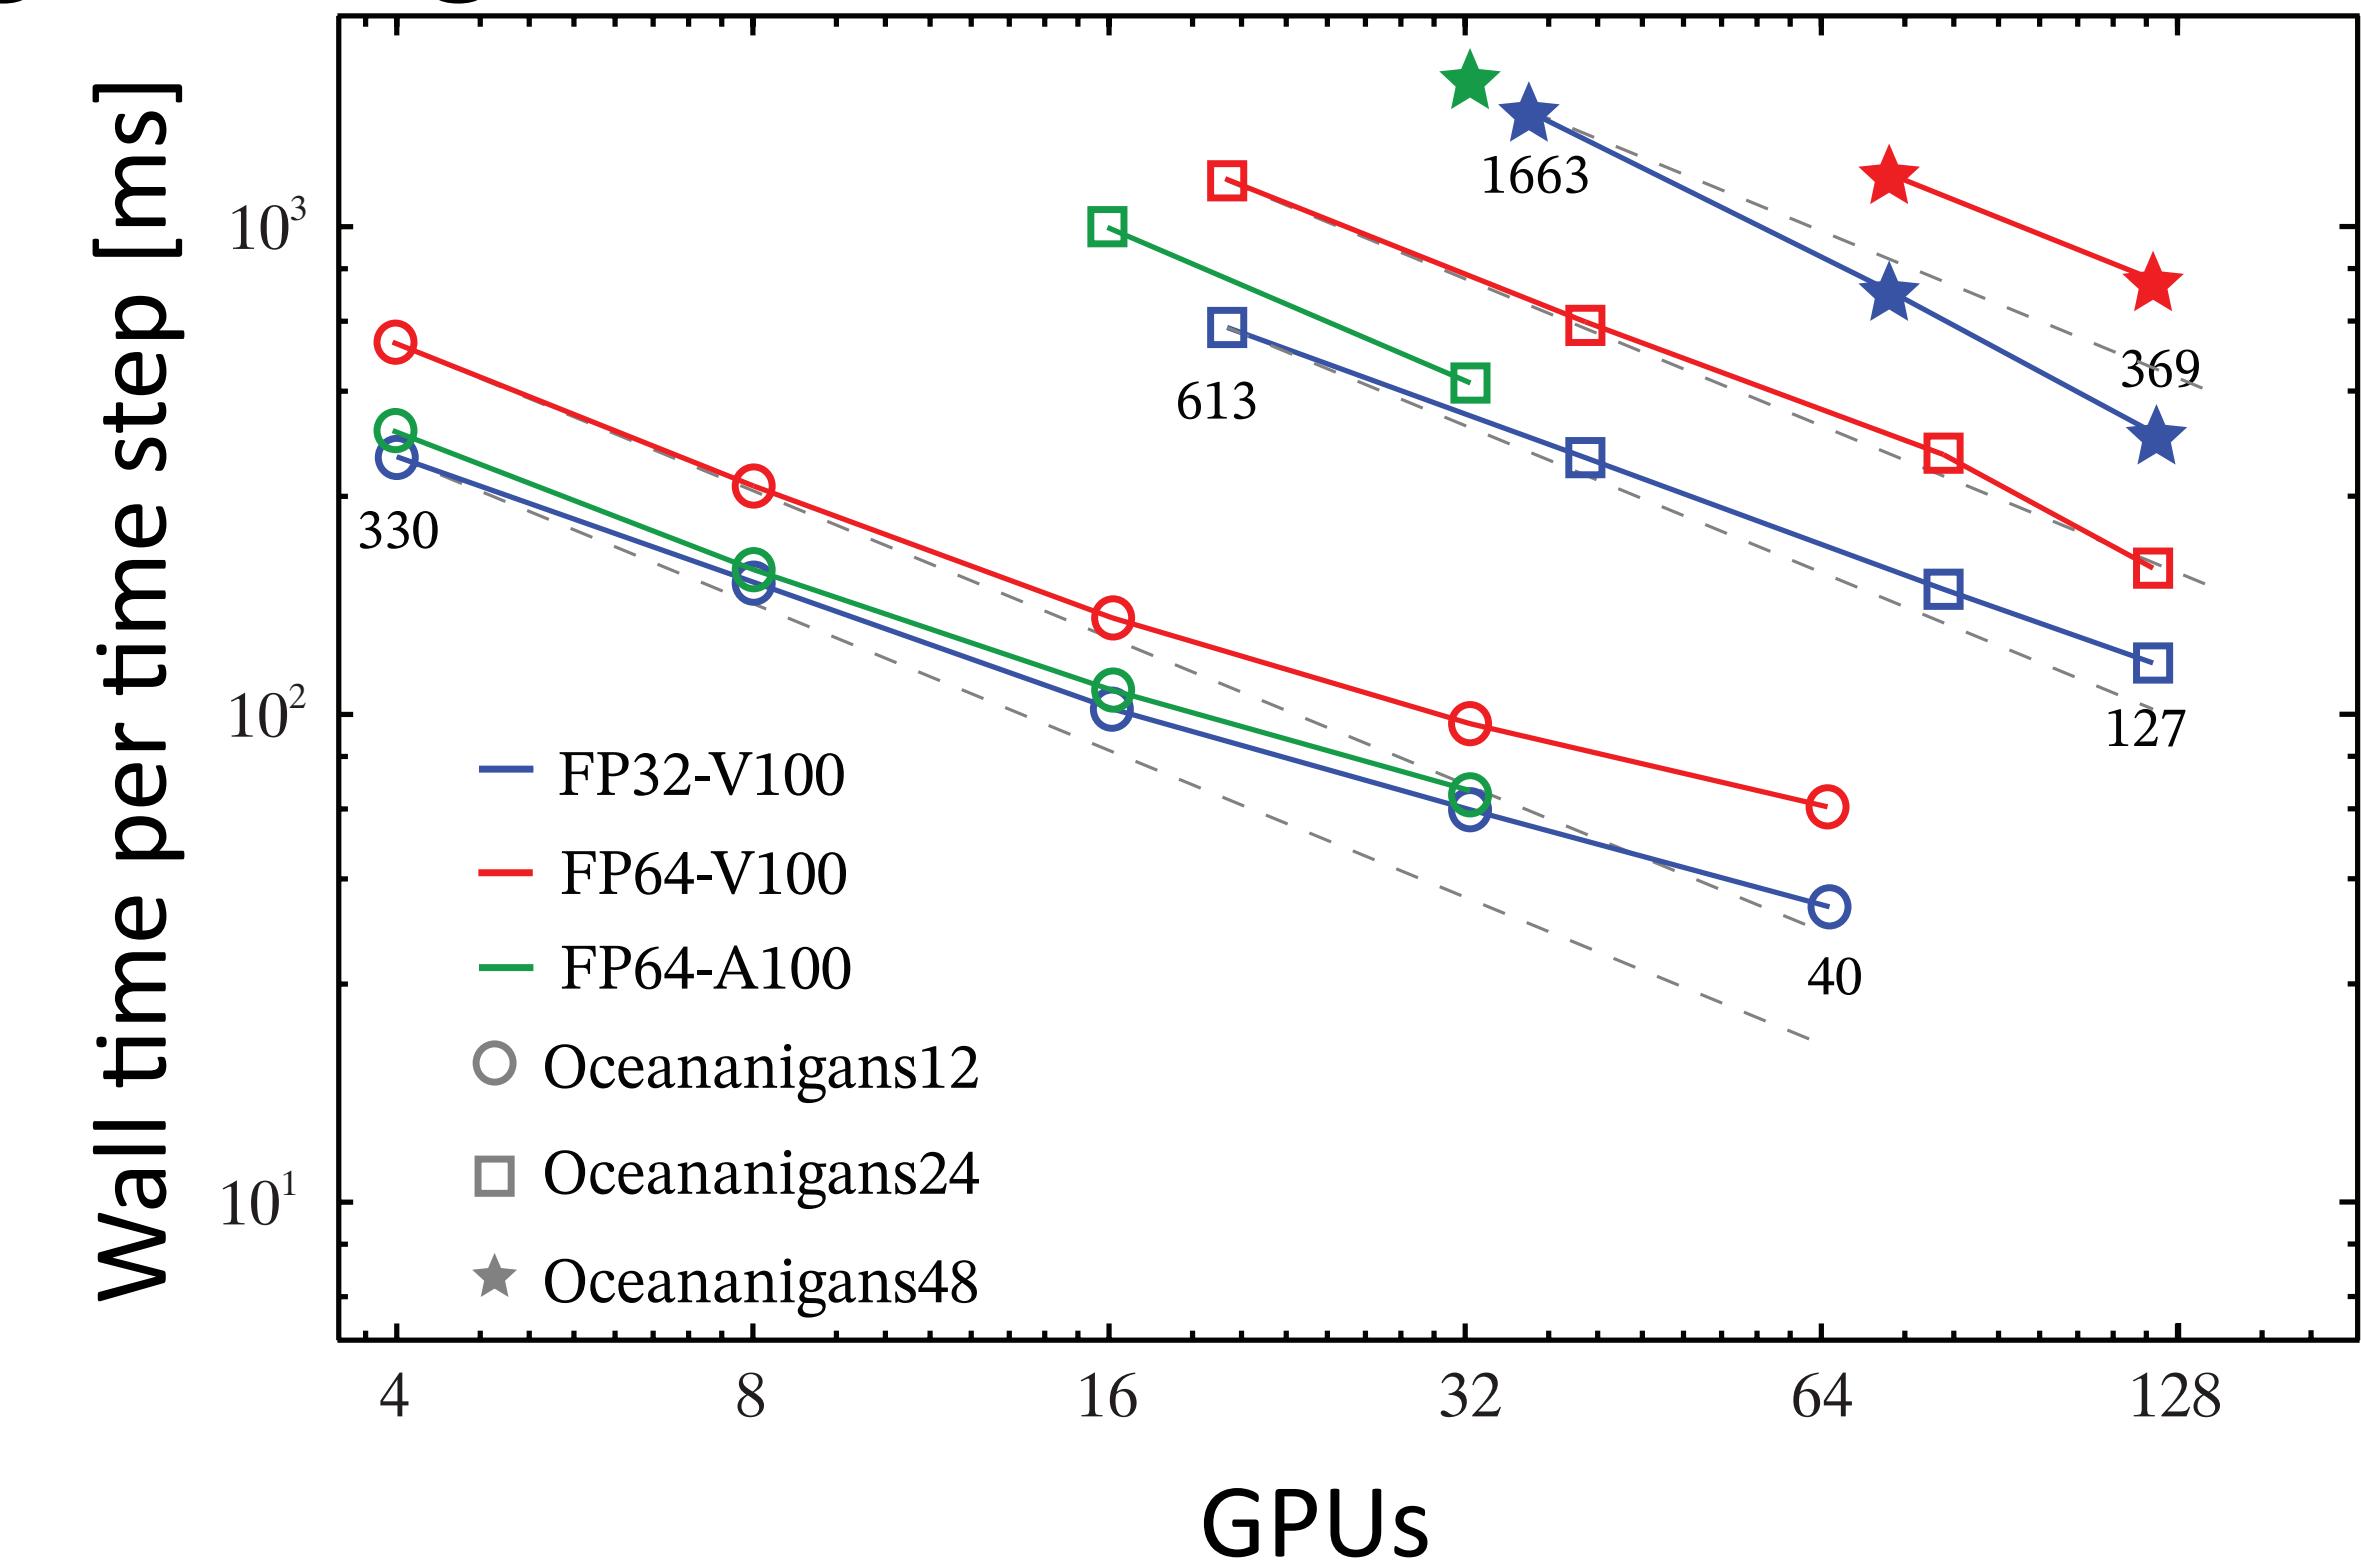
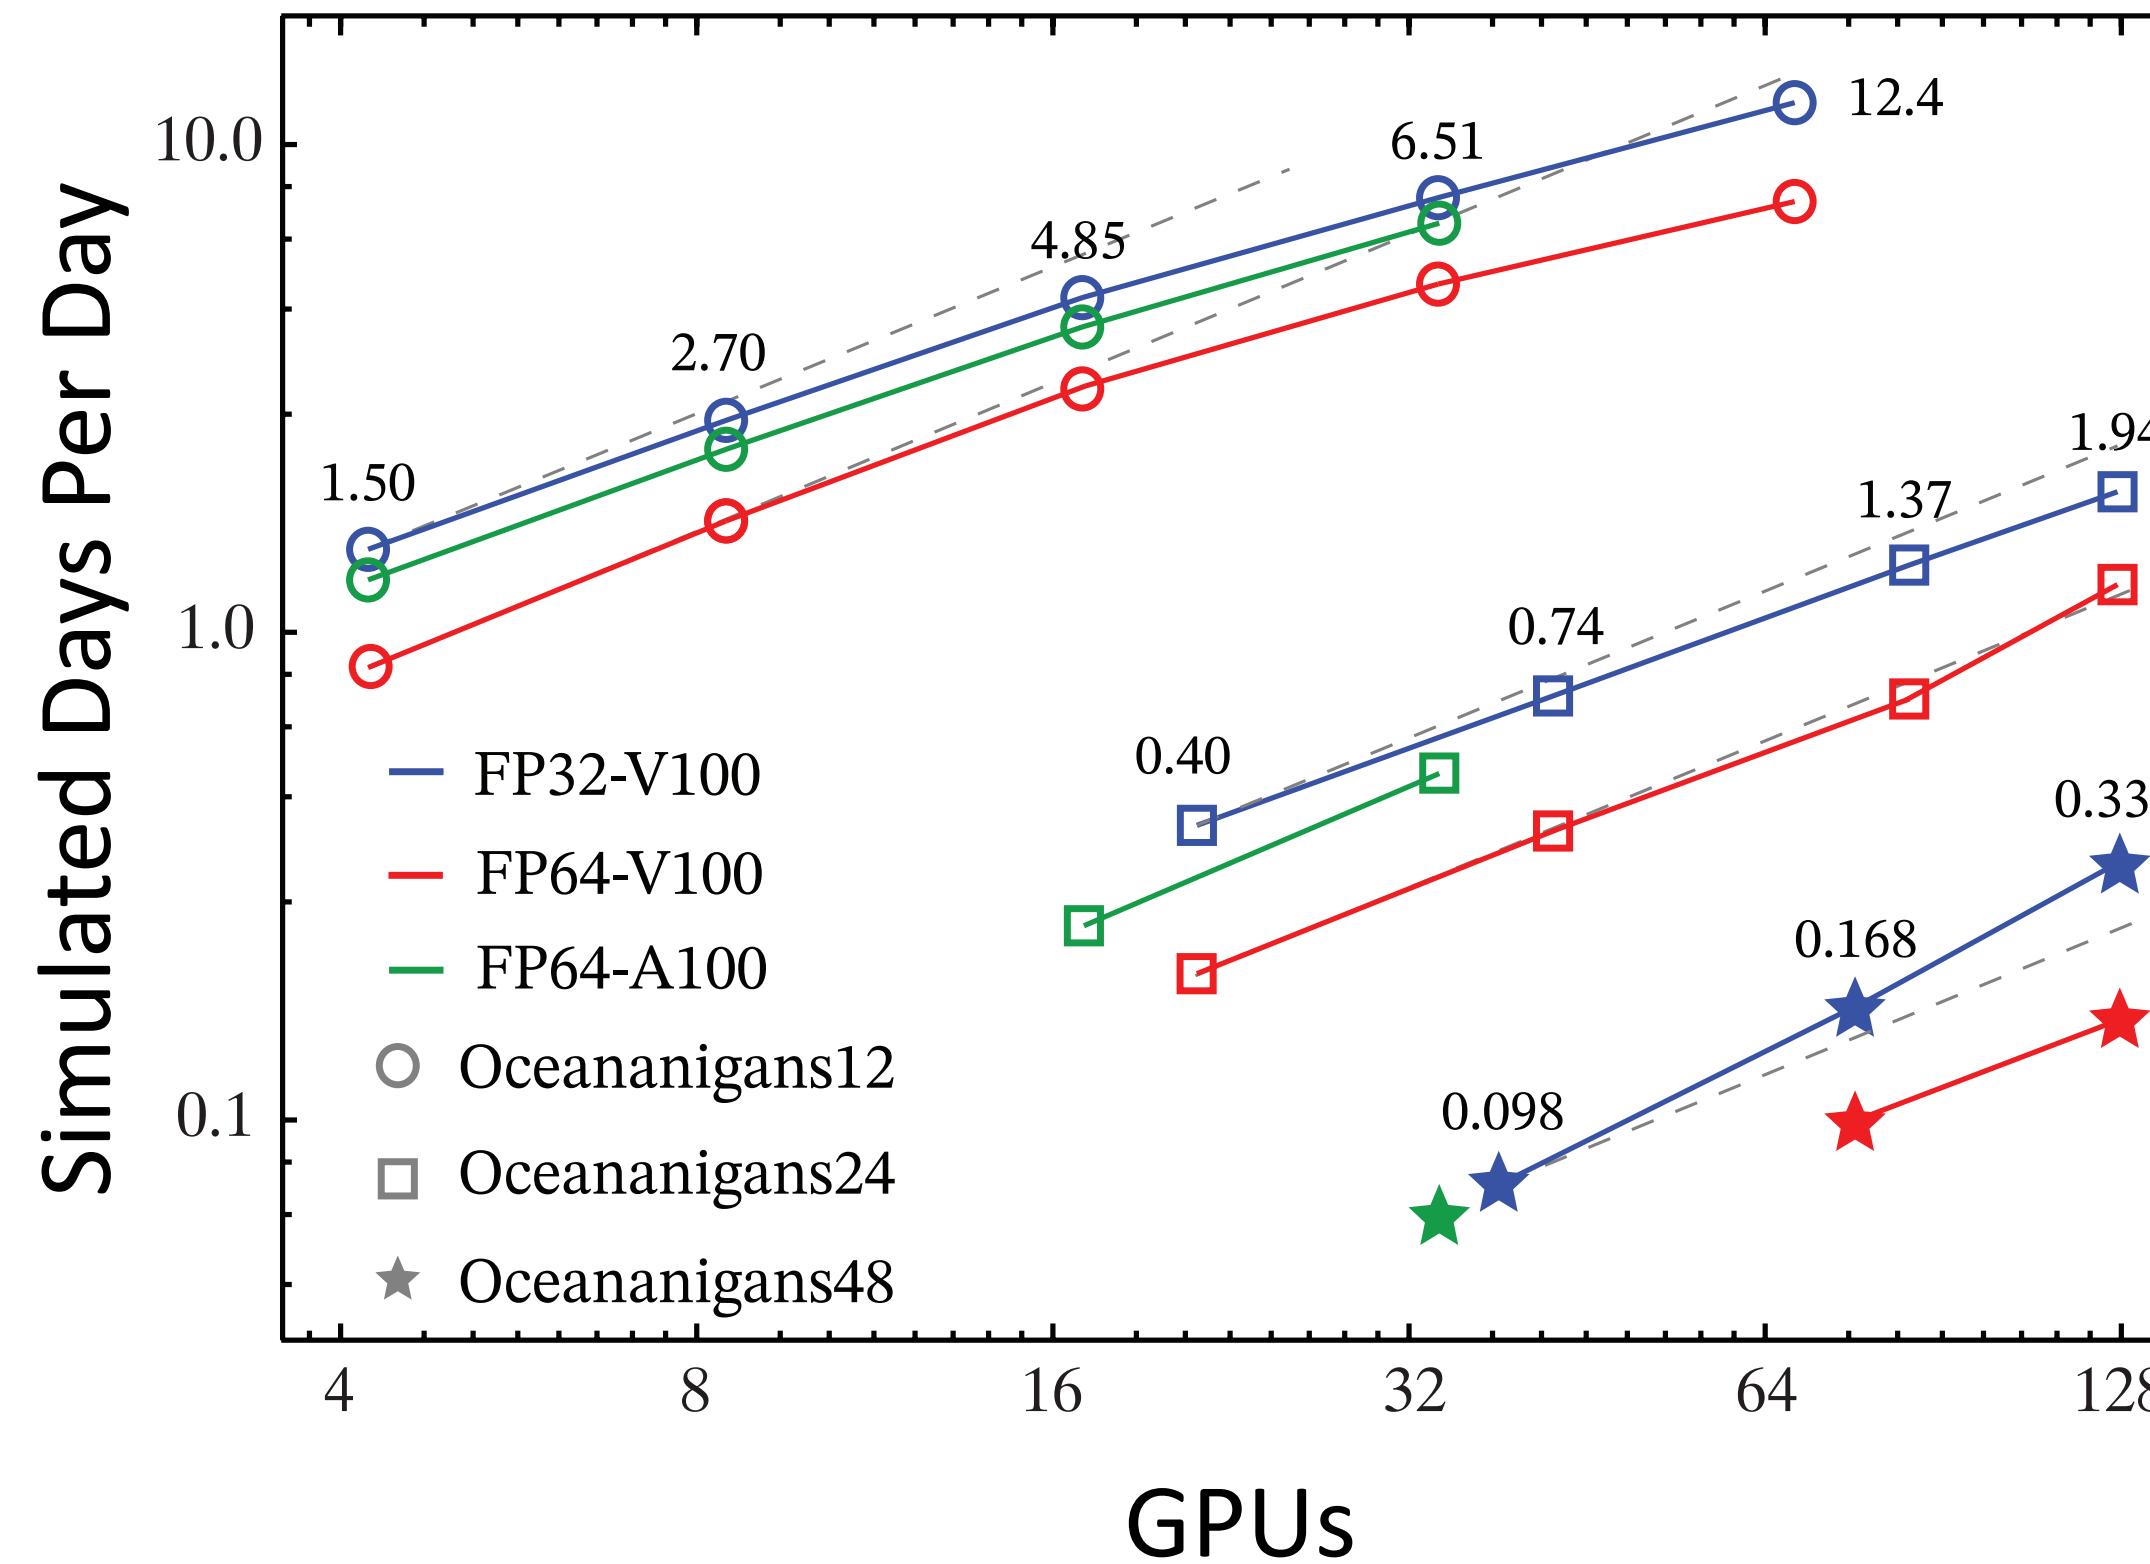
100 Simulated Years Per Day at  $1/2^\circ$  resolution (paleoclimate) @ 1 Nvidia V100 GPU

10 Simulated Years Per Day at  $1/12^\circ$  resolution @ 68 Nvidia A100 GPUs

1 Simulated Year Per Day at  $1/48^\circ$  resolution @ 576 Nvidia A100 GPUs

# how do we scale?

## strong scalings

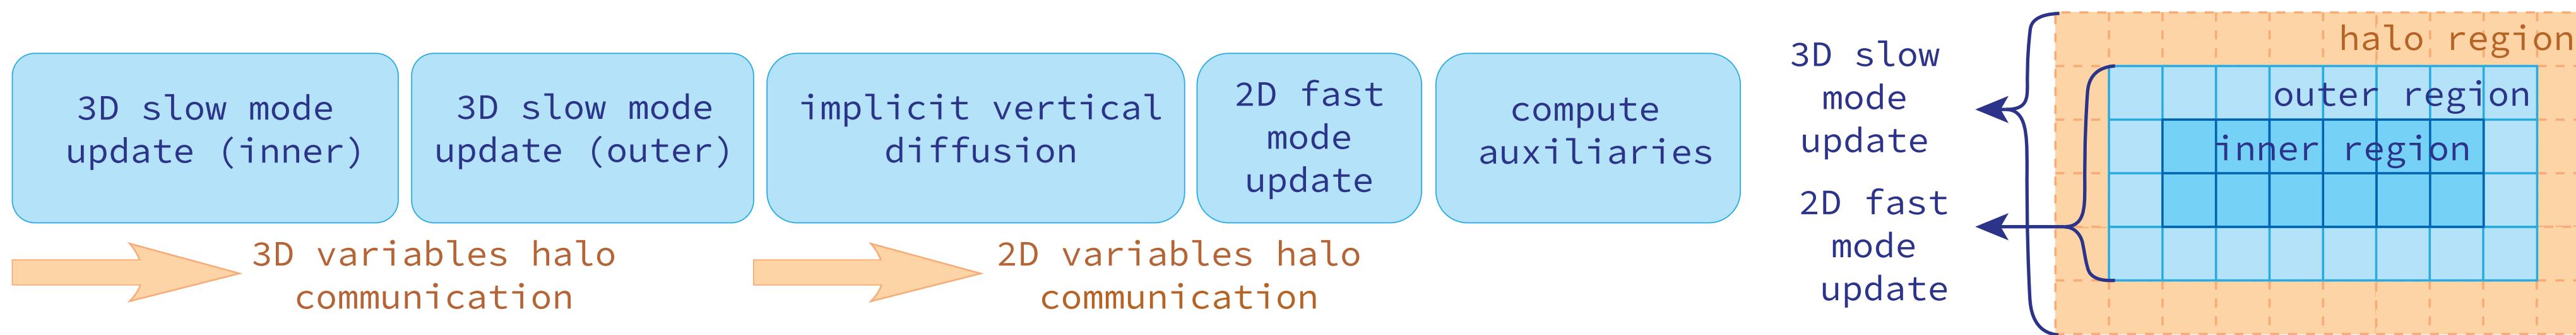
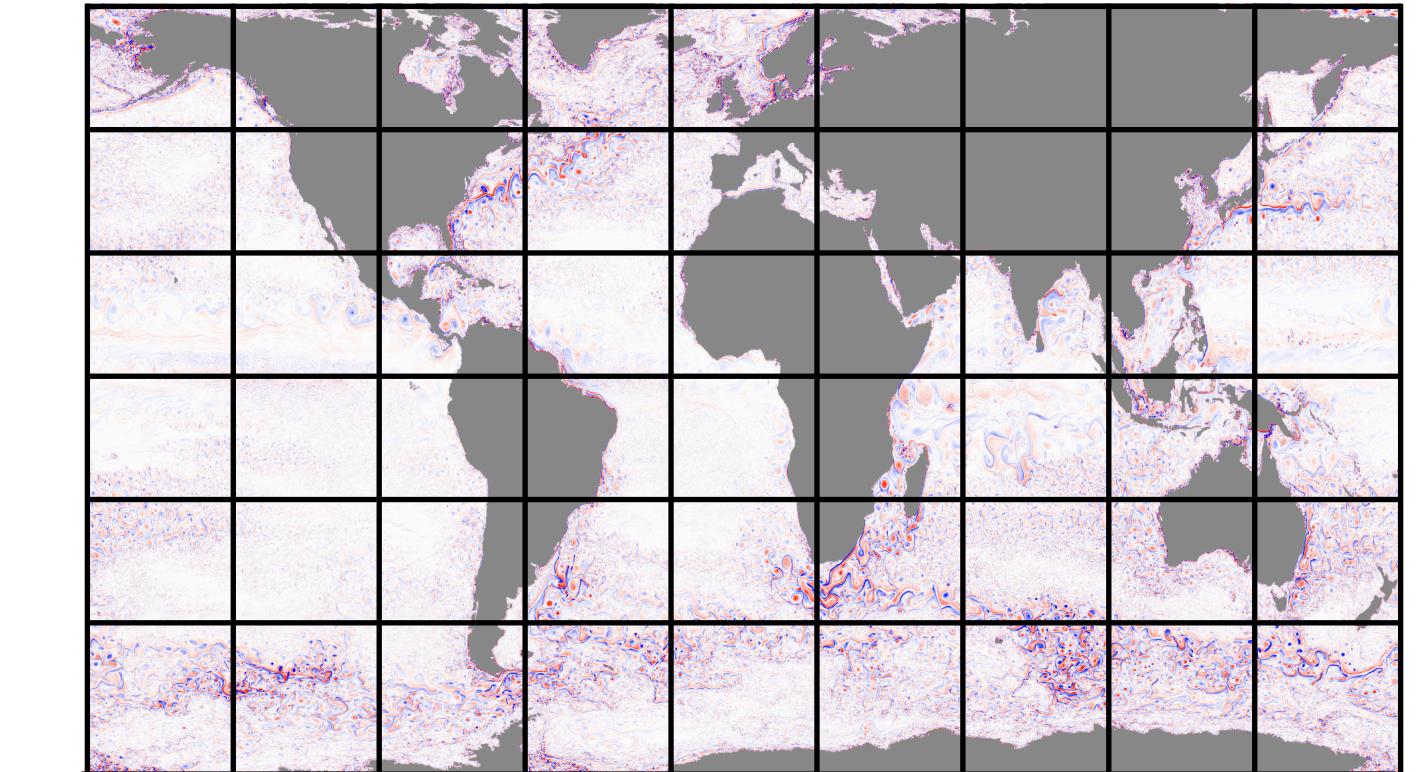
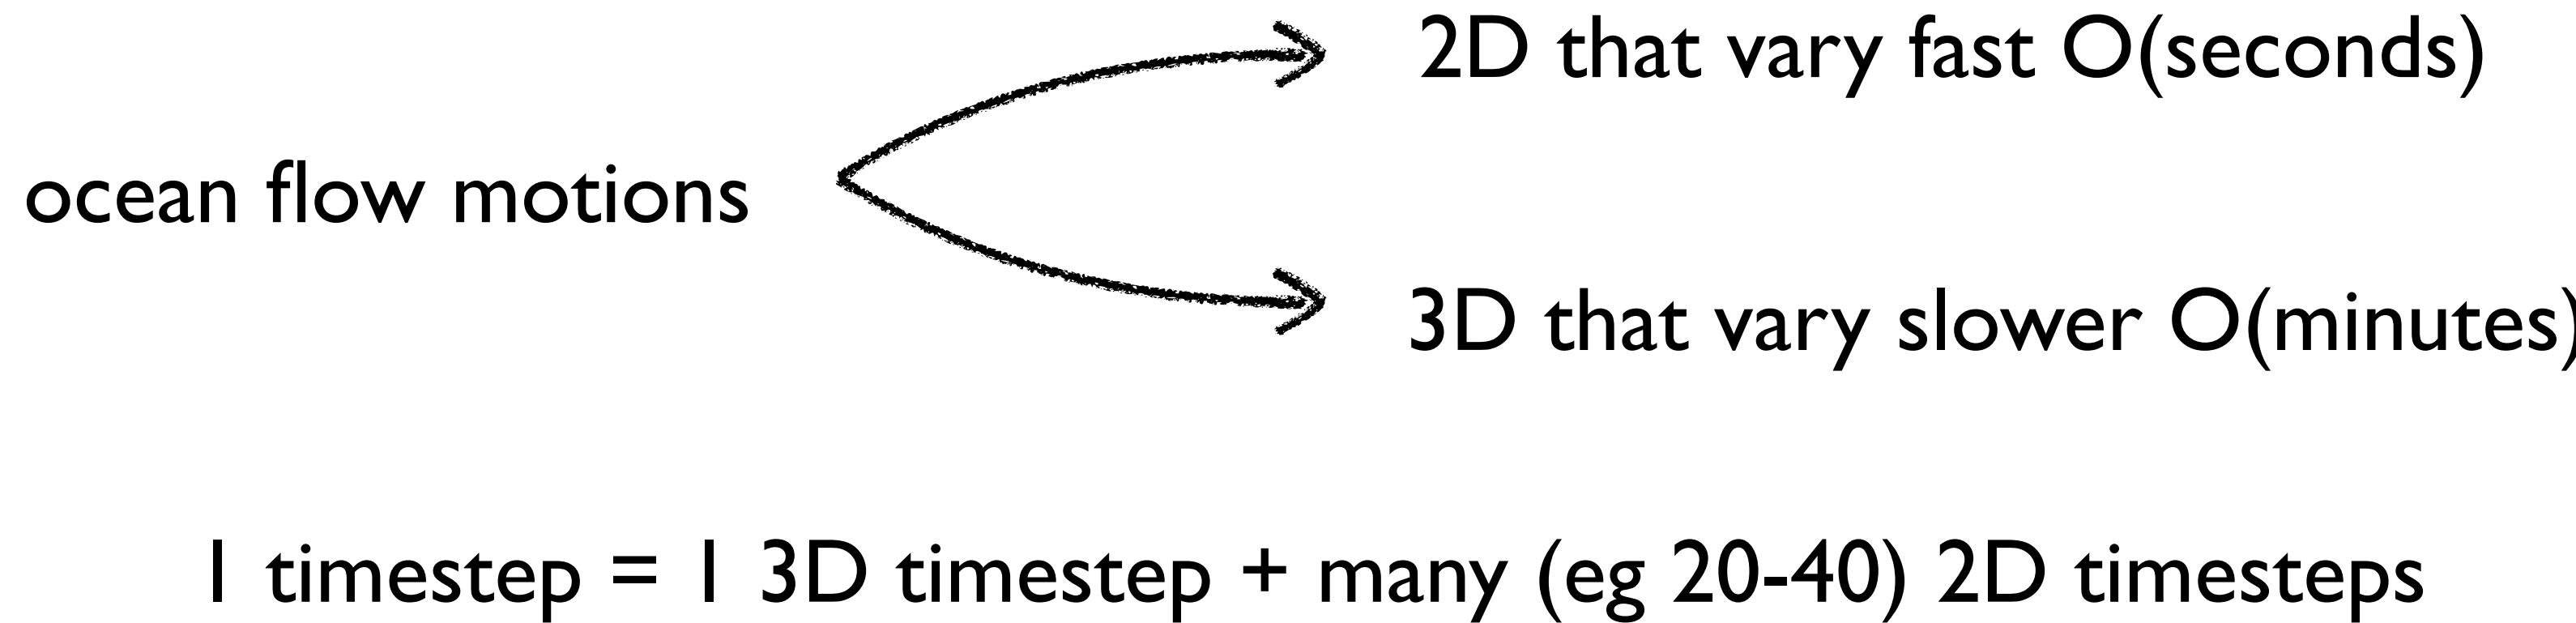


100 Simulated Years Per Day at  $1/2^\circ$  resolution (paleoclimate) @ 1 Nvidia V100 GPU

10 Simulated Years Per Day at  $1/12^\circ$  resolution @ 68 Nvidia A100 GPUs

1 Simulated Year Per Day at  $1/48^\circ$  resolution @ 576 Nvidia A100 GPUs

# scalability



novel algorithm minimises communication in the expense of more computation (cheap on GPU)  
👉 allows overlap of 2D solver with other tasks

# conclusions



Oceananigans enables global ocean modelling  
with unprecedented performance on GPUs



- 👉 start from scratch with GPU in mind
- 👉 software design enables memory leanness
- 👉 novel algorithm allows overlap of computation & communication

thanks