REQUEST FOR A SPECIAL PROJECT 2026–2028

MEMBER STATE:	Greece
Principal Investigator[1]:	Vassilios D. Vervatis
Affiliation:	National and Kapodistrian University of Athens
Address:	Ocean Physics And Modelling group, Section of Environmental Physics – Meteorology, Department of Physics
Other researchers:	Sarantis Sofianos, John Karagiorgos, Rushit Dallenga
Project Title:	Mediterranean Sea Climate Response (MEDCLIM-R)

To make changes to an existing project please submit an amended version of the original form.)

If this is a continuation of an existing project, please state the computer project account assigned previously.	**SPGRVER2**	
Starting year: (A project can have a duration of up to 3 years, agreed at the beginning of the project.)	2026	
Would you accept support for 1 year only, if necessary?	YES ☒	NO ☐

Computer resources required for project year:		**2026**	**2027**	**2028**
High Performance Computing Facility	[MSBU]	23	22	20
Accumulated data storage (total archive volume)[2]	[TB]	40	65	80

EWC resources required for project year:		**2026**	**2027**	**2028**
Number of vCPUs	[#]			
Total memory	[GB]			
Storage	[GB]			
Number of vGPUs[3]	[#]			

Continue overleaf.

[1] The Principal Investigator will act as contact person for this Special Project and, in particular, will be asked to register the project, provide annual progress reports of the project's activities, etc.
[2] These figures refer to data archived in ECFS and MARS. If e.g. you archive x GB in year one and y GB in year two and don't delete anything you need to request x + y GB for the second project year etc.
[3]The number of vGPU is referred to the equivalent number of virtualized vGPUs with 8GB memory.

Principal Investigator: Vassilios D. Vervatis

Project Title: Mediterranean Sea Climate Response (MEDCLIM-R)

Extended abstract

All Special Project requests should provide an abstract/project description including a scientific plan, a justification of the computer resources requested and the technical characteristics of the code to be used. The completed form should be submitted/uploaded at https://www.ecmwf.int/en/research/special-projects/special-project-application/special-project-request-submission.

Following submission by the relevant Member State the Special Project requests will be published on the ECMWF website and evaluated by ECMWF and its Scientific Advisory Committee. The requests are evaluated based on their scientific and technical quality, and the justification of the resources requested. Previous Special Project reports and the use of ECMWF software and data infrastructure will also be considered in the evaluation process.

Requests exceeding 5,000,000 SBU should be more detailed (3-5 pages).

Introduction and Objectives

A key challenge in generating hindcast simulations and climate projections at regional scale, particularly for coastal applications, lies in the limited resolution and inadequate representation of both basin-wide and coastal ocean dynamics, as well as gaps in understanding of oceanic biogeochemical processes. Most ocean climate models are often aligned with meteorological and climatological frameworks such as those used in CORDEX. At the same time, operational centers run detailed ocean models including physical and biogeochemical components for daily forecasting and reanalysis purposes, which offer a more comprehensive depiction of the ocean state. There is a pressing need for methodologies that bridge the gap between large-scale, long-term climate projections and more detailed, short-term Copernicus Marine models currently employed for ocean monitoring.

The resources of this Special Project (SP) are requested to support the R&D activities of the University of Athens, for a research project recently awarded within the European call HORIZON-CL6-2024-CLIMATE-01-6: *"Ocean models for seasonal to decadal regional climate impacts and feedbacks"*. The call is designed for climate action supporting the development of the Digital Twin Ocean and the improvement of marine EU services. In this context, the awarded project named EU-INTERCHANGE aims at enhancing information of high-resolution regional ocean models and datasets, incorporating methods from both current climate projections and Copernicus Marine forecasting models.

The project focuses on increasing the EU's Copernicus hindcast and climate projection capacity, establishing also an uncertainty quantification protocol for optimal model performance. The main objective of the SP is twofold:

- Develop regional optimized ocean physics, wave and biogeochemical models for the Mediterranean Sea, integrated in a high-resolution Digital Ocean Twin.
- Perform historical simulations and climate projections for the Mediterranean Sea, generating high-fidelity and accuracy model datasets for ocean physics, wave and biogeochemical processes.

Methodology

In this project, we will use the latest version of the ocean circulation model NEMO (https://www.nemo-ocean.eu/; Madec, 2022) coupled with the biogeochemical model PISCESv2 (Aumont et al., 2015) that simulates marine biological productivity, the biogeochemical carbon and oxygen cycles, and the main nutrients. In addition, the third-generation spectral wave model WaveWatch III (WW3; WAVEWATCH III® Development Group, 2019) will be integrated, sending back to the ocean model the Stokes drift surface velocities, the significant wave height and the wind stress due to ocean waves. All models are free available open-source software and will be configured in high-resolution (~2.5 km) grids, optimised for the Mediterranean Sea (Mavropoulou et al., 2020; 2022).

The numerical simulations will be carried out for the following periods:

- Past and present state conditions for the period 1995-2025 performing hindcast:
    - evaluation runs, with atmospheric forcing from ERA5.
    - historical runs, with atmospheric forcing from CMIP6.
- Future state conditions performing climate projections for the period 2025-2050 and/or 2080-2100, with atmospheric forcing from CMIP6.

The set of simulations will be carried out in the same way thus, no data assimilation in hindcast mode is foreseen. Uncertainty quantification and model validation with available observations will be performed (i.e., using in-situ and satellite data).

CMIP6 represents the most recent generation of climate models, with emissions scenarios commonly known as *"Shared Socioeconomic Pathways"* (SSPs), based on various socioeconomic assumptions. For the current research we will use CMIP6 based on SSP1-2.6 and/or SSP2-4.5, deemed as the most probable by IPCC AR6. Available open source CMIP6 models will be evaluated against satellite data in our domain of interest for the historical period.

The implementation plan consists of three phases:

- The development phase of each modelling system, including the configuration setup in high-resolution grid for the Mediterranean Sea, the code compilation and optimization performing sensitivity simulations (1st year of the SP).
- The phase of performing hindcast evaluation and historical runs for the period 1995-2025, validating model results against available observations (1st and 2nd year of the SP).
- The phase of performing climate projections for the period 2025-2050 and/or 2080-2100 (2nd and 3rd year of the SP).

Computational Resources

The ocean model configuration has been previously installed at our national GRNET HPC system[2] [3] [4] (named ARIS) and for the needs of this project will be transferred at ECMWF HPC facilities. We briefly describe the previous installation setup to assess the computational resources and data storage required for the ECMWF HPC system.

The ocean model domain covers the Mediterranean Sea between 9.5-37.1°W and 30.0-44°N, using a 1/36° (~2.5 km) curvilinear Arakawa C-grid and a vertical grid of 75 geopotential levels (**Figure 1**). The model is parallelized by domain decomposition with MPI. The optimal domain decomposition of the configuration is approximately ~800 CPUs, considering also the use of the XIOS library for the model I/O. The computational resources required to run the ocean model component are approximately ~9.000 CPU hours per year simulation. When the biogeochemical model is coupled online the configuration becomes four times more expensive i.e., ~36.000 CPU hours per year simulation for both the ocean physics and biogeochemistry.

The wave model has been previously installed at our national GRNET HPC system[5] at ~1/12° resolution and the resources required to run are approximately ~700 CPU hours per year simulation. In this project, the wave model will be configured in the same NEMO grid at 1/36° resolution, and the resources are estimated to increase approximately by about 8-9 times i.e., at ~6.000 CPU hours per year simulation.

Overall, we are considering approximately ~36.000+6.000=42.000 CPU hours per year simulation for the three models to run. To convert the CPU hours in SBUs based on the ECMWF HPC accounting, we multiply with a fixed proportional factor 17,06 adjusted for the ATOS system i.e., 42.000 × 17,06 ~ 0,72 million SBUs per year simulation for all three models to run. Finally, we consider an approximate ~5% increase of the above estimate, for misused computational resources due to possible erroneous job submissions or other failures of the system. In **Table 1**, we present the computational resources required for each simulation experiment and in total for this SP.

[2] https://www.hpc.grnet.gr/en/
[3] https://www.hpc.grnet.gr/en/awarded-2/10th-production-call/ (cf. INVenTORs-* project)
[4] https://www.hpc.grnet.gr/en/awarder-projects-11th-production-call/ (cf. INVenTORs-* project)
[5] https://www.hpc.grnet.gr/en/εγκεκριμένα-έργα-13th-production-call/ (cf. BOFCOAM project)

Data Storage

The ocean state vector consists of five variables {SSH,T,S,U,V} (also, a few surface 2D atmospheric fluxes will be archived) and approximately a similar number of 2D variables will be kept for the wave model. For the biogeochemical model we are considering archiving at about 10 variables (both in 2D and 3D). The model outputs for the hindcast simulations will be archived: (i) at hourly frequency for the wave model, (ii) at daily frequency for the ocean model, and (iii) at weekly frequency for the biogeochemical model. For the historical simulations and the climate projections with the CMIP6 forcing, we are considering weekly/monthly output archives for the 3D ocean/biogeochemical variables (except the wave model outputs and some other ocean/atm. 2D variables considering a higher frequency).

The data size in netcdf format for the ocean model outputs and for the whole state vector {SSH,T,S,U,V} is approximately ~2,5 GB per file. For the hindcast simulations and daily ocean model outputs this yields a data size at about ~0,9 TB per year, whilst for the climate simulations and weekly model outputs is at about ~0,2 TB per year. The wave model outputs in hourly frequency are estimated at about ~0,1 TB per year (same for the hindcast and climate simulations). The data storage of the biogeochemical state is considerably higher than the ocean state and therefore, we reduce the frequency and the number of variables we archive, estimating at about ~0,4 TB for the hindcast simulations (weekly outputs) and ~0,2 TB for the climate simulations (monthly outputs). Finally, we consider archiving yearly restarts for all models (approx. ~10 GB per file) and for some of the experiments, with data size in total at about ~1,5 TB (for both hindcast and climate sim.). The data storage of each model component and in total for all simulations is presented in **Table 2**.

Research Team

The research team participating in the project consists of the Ocean Physics And Modelling (OPAM) group in the Department of Physics, at the National and Kapodistrian University of Athens. The SP participants are Asst. Prof. Vassilios Vervatis (PI), Prof. Sarantis Sofianos, Dr. John Karagiorgos (Postdoc), Rushit Dallenga (PhD Student) and perhaps another PhD Student (to be decided soon). The OPAM scientists will focus on developing ocean modelling tools and performing hindcast/climate simulations for the Mediterranean Sea. The team has significant contribution and expertise on various aspects of the physical oceanography and marine ecosystem for the Mediterranean Sea, and extensive experience using HPC resources participating in previous ECMWF Special Projects[6] [7] [8] [9].

Table 1. *Computational resources: ocean-wave-bgc hindcast/climate simulations at 1/36º res. for the Mediterranean Sea.*

Simulation	Experiment		Period (sim. years)	Models	Resources
Hindcast	Evaluation	ERA5	1995-2025 (30)	NEMO (ocean) PISCESv2 (bgc) WW3 (wave)	0,72 MSBUs per year sim. ×5% increase misc. job failures ×85 years for all sim. ~**65 MSBUs** in total
	Historical CMIP6	SSP1-2.6 and/or SSP2-4.5	2015-2025 (10)[10]		
Climate	Projection CMIP6	SSP1-2.6 and/or SSP2-4.5	2025-2050 (25) and/or 2080-2100 (20)[11]		

[6] https://www.ecmwf.int/en/research/special-projects/spgrverv-2016
[7] https://www.ecmwf.int/en/research/special-projects/spgrver2-2018
[8] https://www.ecmwf.int/en/research/special-projects/spgrver2-2021
[9] https://www.ecmwf.int/en/research/special-projects/spgrver2-2022
[10] This simulation(s) may serve also as spin-up for the climate projection period 2025-2050.
[11] A 10-year spin-up simulation may be necessary before running the climate projection period 2080-2100.

Table 2. Data storage archive: ocean-wave-bgc variables at 1/36º res. for the Mediterranean Sea.

Models	Hindcast sim. outputs (TB/year)	Climate sim. outputs (TB/year)	Archives
NEMO (ocean)	0,9	0,2	• 1,4 TB/year hindcast ×40 years hindcast sim. ~56 TB
PISCESv2 (bgc)	0,4	0,2	• 0,5 TB/year climate ×45 years climate sim. ~22,5 TB • Model restarts ~1,5 TB
WW3 (wave)	0,1	0,1	Data storage in total ~**80TB**

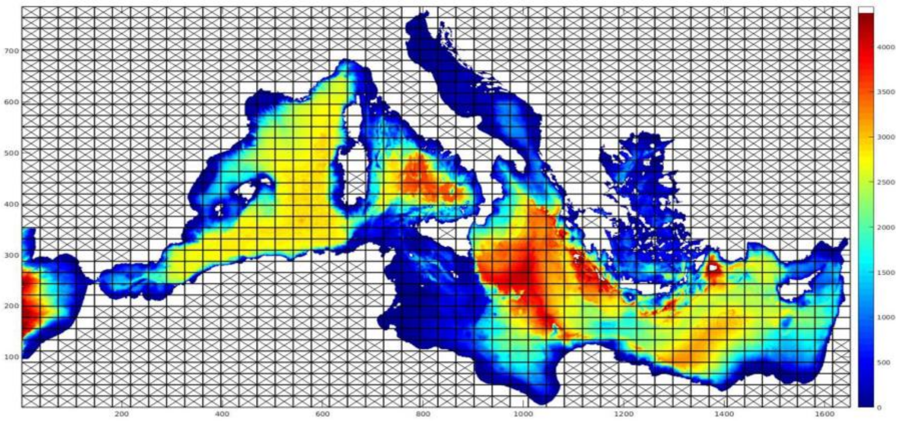

Figure 1. *Mediterranean Sea NEMO domain configuration at 1/36º resolution, suppressing land processors for optimal domain decomposition (~800 CPUs). Colormap: bathymetry (m).*

References

Aumont, O., Ethé, C., Tagliabue, A., Bopp, L., and Gehlen, M. (2015). PISCES-v2: an ocean biogeochemical model for carbon and ecosystem studies, Geosci. Model Dev., 8, 2465–2513, https://doi.org/10.5194/gmd-8-2465-2015.

Madec, G., Bourdallé-Badie, R., Chanut, J., Clementi, E., Coward, A., Ethé, C., et al. (2022). NEMO ocean engine. https://doi.org/10.5281/zenodo.6334656.

Mavropoulou, A.-M., V. Vervatis and S. Sofianos (2020). Dissolved oxygen variability in the Mediterranean Sea, Journal of Marine Systems, Vol. 208, 103348, https://doi.org/10.1016/j.jmarsys.2020.103348.

Mavropoulou, A.-M., V. Vervatis and S. Sofianos (2022). The Mediterranean Sea overturning circulation: A hindcast simulation (1958–2015) with an eddy-resolving (1/36°) model, Deep-Sea Res. Part I, Vol. 187, https://doi.org/10.1016/j.dsr.2022.103846.

The WAVEWATCH III® Development Group (2019). User manual and system documentation of WAVEWATCH III® version 6.07. Tech. Note 333, NOAA/NWS/NCEP/MMAB, College Park, MD, USA, 326 pp. + Appendices.