REQUEST FOR A SPECIAL PROJECT 2022–2024

MEMBER STATE: Italy

Principal Investigator[1]**:** Valerio Capecchi

Affiliation: LaMMA - Laboratorio di Meteorologia e Modellistica Ambientale per lo sviluppo sostenibile

Address: Via Madonna del Piano, 10 - Sesto Fiorentino (Firenze) - Italy

Other researchers: Alberto Ortolani - LaMMA

Project Title: Scaling properties of the simulation speed: a testbed case for three limited-area models

If this is a continuation of an existing project, please state the computer project account assigned previously.	**SP ITCAPE**	
Starting year: (A project can have a duration of up to 3 years, agreed at the beginning of the project.)	2022	
Would you accept support for 1 year only, if necessary?	YES ☒	NO ☐

Computer resources required for 2022-2024: (To make changes to an existing project please submit an amended version of the original form.)		**2022**	**2023**	**2024**
High Performance Computing Facility	(SBU)	3 800 000		
Accumulated data storage (total archive volume)[2]	(GB)	1000		

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.

Principal Investigator: Valerio Capecchi

Project Title: Scaling properties of the simulation speed: a testbed case for three limited-area models

Extended abstract

The completed form should be submitted/uploaded at https://www.ecmwf.int/en/research/special-projects/special-project-application/special-project-request-submission.

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.

Following submission by the relevant Member State the Special Project requests will be published on the ECMWF website and evaluated by ECMWF as well as the Scientific Advisory Committee. The evaluation of the requests is based on the following criteria: Relevance to ECMWF's objectives, scientific and technical quality, disciplinary relevance, and 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 asking for 1,000,000 SBUs or more should be more detailed (3-5 pages). Large requests asking for 10,000,000 SBUs or more might receive a detailed review by members of the Scientific Advisory Committee.

1. Project's description

The ability to perform detailed weather simulations for real-time applications depends on the ability of the numerical model to effectively use large computational resources, which are becoming increasingly more common even at relatively small supercomputing centres. However, to what extent the code of limited-area models scales as the computational power increases, results somehow vague, and, to our knowledge, no systematic investigation exists in this regard. On the other hand, the scalability performances are a crucial information when evaluating the computer power needed to meet time-critical deadlines in operational weather agencies. To fill this research gap, we propose to investigate the scaling properties of the simulation speed for three state-of-the-art limited-area weather models.

2. Scientific and technical plans

The simulation speed is defined as the simulation length over the wall-clock elapsed time, and it is usually used as a measure of the effectiveness of the code of a numerical model. The simulation speed (hereinafter *SSP*) is a function of the time step and depends also on the number of variables to be processed at each time step and on the number of calculations to be made at each time step, which is an indirect measure of the computational cost required by the numerical method used to solve the equations.
Following the notations of Coiffier (2011), we have:

$$SSP=(\Delta t \cdot P) / (N_V \cdot N_C)$$

where Δt is the time step, which depends on the grid spacing to satisfy the Courant–Friedrichs–Lewy stability condition, the term N_V is the number of variables to be processed at each time step Δt and depends on the number of grid points (i.e., the number of rows, columns, and vertical levels) of the integration domain. The term N_V represents the number of calculations to be made at each time step and is a function of the computational cost required by the numerical method used to solve the equations. The numerator P, is a measure of the computational speed (e.g., the number of processing elements or the floating operations per second). In our experimentations, the computational speed P will be used as a parameter, and will be smoothly increased to estimate the *SSP*, as shown in Figure 1. We will look at the ratio of increase in simulation speed to the increase in the number of computing elements, and define such rate the scaling property of the model.

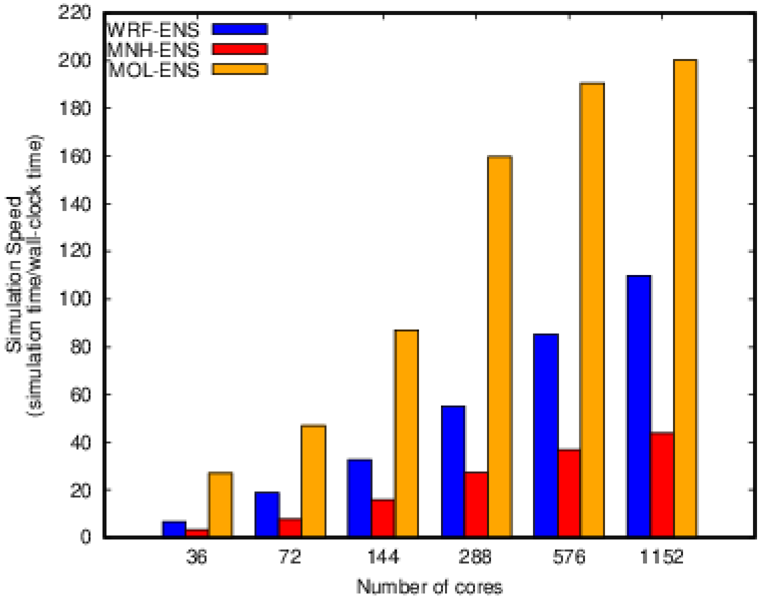

Figure 1: the simulation speed (y-axis) is plotted against the number of cores deployed to perform a 36-hour-long simulation for three limited-area models: the WRF (blue bar), Meso-NH (red bar) and MOLOCH (orange bar) model.

We stress the fact that to estimate the simulation speed, the wall-clock time taken into account, considers only the period spent to compute the evolution of the state variables and not that spent for reading the initial conditions and postprocessing the model outputs.

The preliminary results shown in Figure 1 were obtained in the framework of the SPITCAPE 2018-2020 Special Project and are reported in Capecchi (2021). The three limited-area models chosen for the preliminary numerical experiment were the WRF-ARW, Meso-NH and MOLOCH models; see Section 4 for a brief description of such models. Previous investigations also demonstrated that the reduction (in percent) of the elapsed wall-clock time scales as shown in Figure 2 when doubling the number of computing cores,:

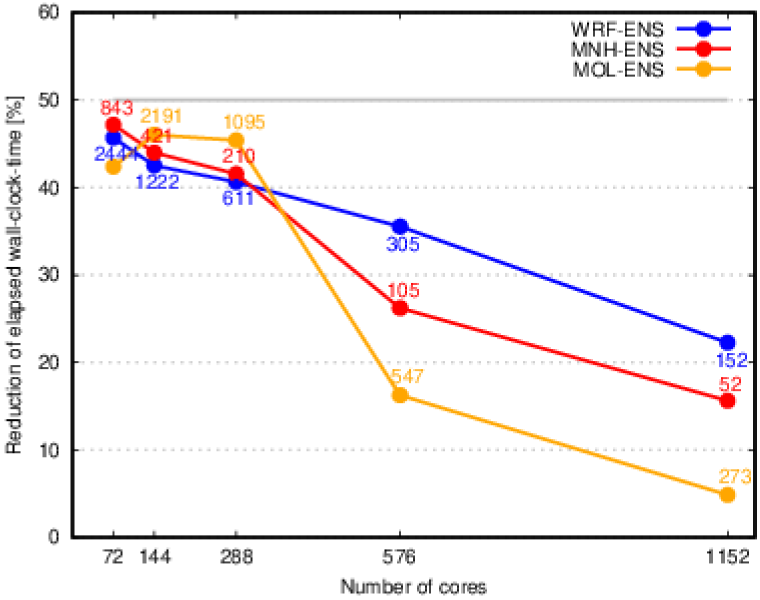

Figure 2: On the y-axis it is shown the reduction (in percent) of the elapsed wall-clock time when doubling the number of computing cores to used to perform a 36-hour-long simulation. The coloured labels indicate the number of grid points assigned to each core.

In a strong scaling regime, the coloured lines shown in Figure 2 should approximate the 50% horizontal grey line. However, during the SPITCAPE 2018-2020 Special Project, we noted that the number of horizontal grid points assigned to each core, say $N_{x,y}$, affects the speed-up of the simulation speed ($N_{x,y}$ is indicated in Figure 2 with the coloured labels). In fact, limited-area models decompose the domain into horizontal patches, and each computing element is responsible for a single patch; when $N_{x,y}$ is relatively low, the time required to perform the calculations on the perimeter of each patch overwhelms the computational time.

The present Special Project is aimed at overcoming the shortcomings of the SPITCAPE 2018-2020 Special Project (see Capecchi, 2021), namely:

1. To define a domain of integration common to the three limited-area models. This domain has to cover all of Italy to assess the computational effort needed to deploy an ensemble or a deterministic higher-resolution system at the national level.
2. To set a common framework to build the executables of the three limited-area models. In other words, we will use the same compiler vendor and compilation flags (whenever feasible). This step is needed because it is known that compilation options impact the simulation speed.
3. To determine the simulation speed of the three limited-area models as the number of computing elements spans from 128 (that is SLURM_JOB_NUM_NODES=1) to 6400 (that is SLURM_JOB_NUM_NODES=50), that is by doubling the variable SLURM_JOB_NUM_NODES up to 16, then by approximately incrementing the number of computing nodes by 25% with respect to the previous experiment. This would yield the ten experiments shown in Table 1 below.
4. To determine, for each limited-area model, the tipping point from the strong to the weak scaling regime. This will be achieved by:
    - Investigating the profile of the speed-up of the simulation speed as a function of the number computing elements and determining the point when such profile diverts significantly (i.e., absolute difference greater than a predefined threshold) from the theoretical expected behaviour.
    - Determining the corresponding number of grid points assigned to each computing core. This number will mark the threshold below which the time spent for inter-patch communication overwhelms the time spent for the model dynamics.

Experiment number	**SLURM_ARRAY_TASK_COUNT**	**SLURM_JOB_NUM_NODES**	**Increment in the number of nodes (%)**
1	128	1	100
2	256	2	100
3	512	4	100
4	1024	8	100
5	2048	16	100
6	2560	20	25
7	3200	25	25
8	4096	32	22
9	5120	40	25
10	6400	50	25

Table 1: SLURM variables settings (i.e., total number of computing elements and corresponding computing nodes) of the ten numerical experiments aimed at assessing the simulation speed and scaling properties of the WRF-ARW, Meso-NH and MOLOCH models. Note: in this table we assume that the SLURM variable SLURM_NTASKS_PER_NODE is equal to 128 for all the experiments.

A secondary goal of the project is to define a benchmark dataset, which is intended to provide a means for comparing the performance of the three different limited-area models in terms of simulation speed and efficient use of the available computational resources. The proposed benchmark event is the flooding of the Cinque Terre (hereinafter CT) UNESCO site, occurred in the north-western Italy on the 25th of October 2011. Two different workloads will be proposed: CT2.5 and CT1.25: they are single domain simulations and differ because of the grid spacing and number of vertical levels. The first workload has a grid spacing equal to 2.5 km and 50 vertical levels; it is thought to be the basis for ensemble member simulations (performed via dynamical downscaling of global forecasts). The size of the mesh of the second workload, CT1.25, is 1.25 km and the number of vertical levels is 60; this domain is thought to provide a basis for high-resolution deterministic forecasts. The estimated size of the 2D grille is 500X550 and 1000X1100 for the CT2.5 and CT1.25 workload, respectively. Starting date for both workloads is 12 UTC 24 October 2011 and the common forecast length is set to 30 hours (that is ending date is 18 UTC 25 October 2011).
In the next future, international centres are supposed to produce global ensembles with horizontal resolution of approximately 3-4 km by the end of the decade (“ECMWF Strategy to 2030”). Thus, our scalability study of limited-area models with a mesh size of approximately 2.5 and 1.5 km seems appealing at most for the next few years. That’s why we ask resources for 2022 only. However, we claim that determining $N_{x,y}$ is a resolution-off problem and it would provide meaningful insights for future higher-resolution domains.
We also stress the fact that investigating the simulation speed of the CT2.5-like domain would offer also an estimate of the computational efforts needed to downscale climate projections at the convection-permitting scale, which is arising as an appealing topic in regional climate-change studies (Ban et al, 2021; Pichelli et al, 2021).

A final goal of the project is to compare the effective resolution of the three models, that is their ability “to resolve features at the limits of the grid resolution” (Skamarock, 2004). The computation burden related to each model and the scaling properties are in fact compared on the assumption that models reproduce the same spatial resolution. However, it is crucial to assess on which extent the effective spatial resolution differs from the grid one, and how this difference varies across the three models. This task will be achieved by computing the kinetic energy (KE) spectra on horizontal wind as a function of scale. In fact, by assessing the departure of KE spectra from the theoretical slope, namely a wavenumber dependence of $k^{-5/3}$ for wavelength shorter than 10^3 *km*, we can have an idea of the effective resolution of the models (see the seminal works by Skamarock, 2004 and Ricard et al, 2012). A comparative analysis of KE spectra, although not new, would add meaningful insights regarding the effective resolutions of the three models, because of results obtained over a common domain and with the same mesh size. Such results will be discussed against the findings of previous works: i.e., an affective resolution of approximately $7\cdot\Delta x$ for the WRF-ARW model (Skamarock, 2004) and approximately $4\cdot\Delta x$ (Ricard et al, 2012) for the Meso-NH model (here Δx indicates the mesh size of the model in km).

3. Justification of the computer resources requested

Computational resources requested leverage previous experiences on the Cray XC40 supercomputer during the SPITCAPE 2018-2020 Special Project. A quick test performed on TEMS, showed that we have to consider an inflating factor of at least 1.2 when converting SBUs consumption from the Cray XC40 supercomputer to the forthcoming ATOS BullSequana XH2000 supercomputer. Numerical experiments performed will be ten for each limited-area model and for each workload for a total number of 60 model runs.
In Table 2 below, we report the SBUs requests for CT2.5:

Experiment number	**SLURM_JOB_NUM_NODES**	**WRF-ARW Estimated SBUs consumption**	**Meso-NH Estimated SBUs consumption**	**MOLOCH Estimated SBUs consumption**
1	1	3000	4200	600
2	2	3700	5900	1100
3	4	4900	8200	2000
4	8	6800	13500	3700
5	16	9500	22300	6900
6	20	11500	29500	9700
7	25	13700	39100	13900
8	32	16500	51800	19800
9	40	19800	68700	28300
10	50	23700	91000	40200
TOTAL		**106 400**	**334 200**	**126 200**

Table 2: estimated cost of the CT2.5 workload for the three limited-area models and for the ten experimental runs.

The cost of CT2.5 is estimated in 567 000 SBUs.
We estimate the CT1.25 cost to be five times the CT2.5 cost, which approximately corresponds to 2 835 000 SBUs. By taking into consideration an overhead of approximately 10% the total SBUs request is **3 800 000**.
Since we're mainly interested in the SLURM log files indicating the simulation elapsed runtime, and we're not interested in the model outputs (except for the wind variables needed for calculating the KE spectra), we estimate that the SCRATCH/ SCRATCHDIR/HPCPERM quotas should be enough. However, to save initial and boundary conditions data from the global model and to store the results of the CT1.25 workload for potential future use and comparison with previous paper (i.e., Buzzi et al 2014), we request 1000 GB of data storage.

4. Technical characteristics of the code

The Mesoscale and Microscale Meteorology (MMM) Laboratory at the National Center for Atmospheric Research (NCAR) has led the development of the WRF Model since its inception in the late 1990s. WRF is a fully compressible, Eulerian, nonhydrostatic mesoscale model; we will implement the Advanced Research version of WRF (ARW), which is described in Skamarock et al. (2019). The microphysics option adopted is detailed in Thompson and Eidhammer (2014); it is a mixed-moment scheme with the activation of aerosols as condensation nuclei and explicitly predicts the number concentration of rain and cloud ice in addition to the mixing ratios of five hydrometeor species: cloud water, rain, cloud ice, snow, and a hybrid graupel–hail category. The version 4.0 or above will be used in the Project.

Meso-NH is a French research community model, jointly developed by the Centre National des Recherches Météorologiques and Laboratoire d'Aérologie at the Université Paul Sabatier. It is designed to simulate the time evolution of several atmospheric variables ranging from the large meso-alfa scale (about 2000 km) down to the micro-gamma scale (about 20 m), typical of the large eddy simulations. For a general overview of the Meso-NH model and its applications, see Lac et al. (2018). Regarding microphysics, we will set the one-moment ICE3 scheme (Caniaux et al.

1994), taking five water species into account: cloud droplets, raindrops, pristine ice crystals, snow or aggregates, and graupel. We stress the fact that, in contrast to what has been done during the 2018-2020 SPITCAPE Special Project, the numerical scheme adopted will be the the WENO5, which is suggested for mesoscale applications (Lac et al, 2018). This will considerably reduce the computational cost of the Meso-NH model. The source code of the Meso-NH model is written in standard Fortran 90. The required libraries to run Meso-NH are NetCDF because the output files are in nc4 format, MPI and grib_API/ecCodes tool. The version 5.4.3 or above will be used in the Project.

MOLOCH is a nonhydrostatic, fully compressible model that uses a hybrid terrain-following coordinate, relaxing smoothly to horizontal surfaces. It is developed at the Institute of Atmospheric Sciences and Climate (ISAC) of the Italian National Research Council (CNR). Details about the model can be found in Malguzzi et al. (2006). It was initially developed for research purposes, but today it is being used operationally by various regional meteorological services both in Italy and abroad. The microphysical scheme is based on the parameterisation proposed in Drofa and Malguzzi (2004), which describes the interactions of cloud water, cloud ice, rain, snow, and graupel. The MOLOCH model is based on a single Fortran 90 code; it is fully parallelised, applying the domain splitting technique, and is compatible with MPICH2 and OpenMP parallel computing environments. The version released in 2019 (or later) will be used in the Project.

The WRF-ARW, Meso-NH and MOLOCH models were successfully compiled and implemented on the Cray XC40 supercomputer during the SPITCAPE 2016-2018 and 2019-2021 Special Projects. As regards the forthcoming ECMWF HPC infrastructure, WRF is available on TEMS as a module for Intel programming environment and Open MPI because this combination has given the best performance; therefore, no difficulties are foreseen for its compilation on the ATOS system. The compilation of the Meso-NH model on the new ATOS system is somehow a challenge; however, model developers report that the Meso-NH model compiles and runs correctly at Météo-France on an ATOS supercomputer. MOLOCH relies on few additional libraries (namely, the NetCDF and ecCodes libraries) and building its executables was proven to be “smooth” enough on several architecture using the Intel Fortran compiler.

5. References

- Ban, Nikolina, et al. "The first multi-model ensemble of regional climate simulations at kilometer-scale resolution, part I: evaluation of precipitation." *Climate Dynamics* (2021): 1-28.
- Buzzi, A., et al. "Heavy rainfall episodes over Liguria in autumn 2011: numerical forecasting experiments." *Natural Hazards and Earth System Sciences* 14.5 (2014): 1325-1340.
- Caniaux, G., J. Redelsperger, and J. P. Lafore, 1994: A numerical study of the stratiform region of a fast-moving squall line. Part I: General description and water and heat budgets. *J. Atmos. Sci*., 51, 2046–2074, https://doi.org/10.1175/1520-0469(1994)
- Capecchi, V. (2021). Reforecasting Two Heavy-Precipitation Events with Three Convection-Permitting Ensembles, *Weather and Forecasting*, *36*(3), 769-790
- Chan, Steven C., et al. "Europe-wide precipitation projections at convection permitting scale with the Unified Model." *Climate Dynamics* 55.3 (2020): 409-428.
- Coiffier, J. (2011): Fundamentals of Numerical Weather Prediction. Cambridge: Cambridge University Press.
- Drofa, O., and P. Malguzzi, 2004: Parameterization of microphysical processes in a nonhydrostatic prediction model. Proc. 14th Int. Conf. on Clouds and Precipitation (ICCP), Bologna, Italy, ICCP, 19–23.
- Lac, C., and Coauthors, 2018: Overview of the Meso-NH model version 5.4 and its applications. *Geosci. Model Dev*., 11, 1929–1969, https://doi.org/10.5194/gmd-11-1929-2018.
- Malguzzi, P., G. Grossi, A. Buzzi, R. Ranzi, and R. Buizza, 2006: The 1966 ‘‘century’’ flood in Italy: A meteorological and hydrological revisitation. *J. Geophys. Res*., 111, D24106, https://doi.org/10.1029/2006JD007111.
- Pichelli, Emanuela, et al. "The first multi-model ensemble of regional climate simulations at kilometer-scale resolution part 2: historical and future simulations of precipitation." *Climate Dynamics* 56.11 (2021): 3581-3602.
- Ricard, Didier, et al. "Kinetic energy spectra characteristics of two convection-permitting limited-area models AROME and Meso-NH." *Quarterly Journal of the Royal Meteorological Society* 139.674 (2013): 1327-1341.
- Skamarock, William C. "Evaluating mesoscale NWP models using kinetic energy spectra." *Monthly weather review* 132.12 (2004): 3019-3032.
- Skamarock, William C., et al. A description of the advanced research WRF model version 4. *National Center for Atmospheric Research: Boulder, CO, USA*, 2019, 145.
- Thompson, G., and T. Eidhammer, 2014: A study of aerosol impacts on clouds and precipitation development in a large winter cyclone. *J. Atmos. Sci*., 71, 3636–3658, https://doi.org/10.1175/JAS-D-13-0305.1