SPECIAL PROJECT FINAL REPORT

Project Title:	Scaling properties of the simulation speed: a testbed case for three limited-area models
Computer Project Account:	spitcape
Start Year - End Year :	2022 - 2023
Principal Investigator(s)	Valerio Capecchi
Affiliation/Address:	LaMMA - Laboratorio di Meteorologia e Modellistica Ambientale per lo sviluppo sostenibile Via Madonna del Piano, 10 - Sesto Fiorentino (Firenze) - Italy
Other Researchers (Name/Affiliation):	Alberto Ortolani / LaMMA

The following should cover the entire project duration.

Summary of project objectives

(10 lines max)

The SPITCAPE Special Project (year 2022) had the main target to study how the simulation speed of the WRF-ARW, MOLOCH and Meso-NH limited-area numerical models scales as the computational power increases. A secondary goal was to study how the number of horizontal grid points assigned to each core affects the speed-up of the simulation speed. Intermediate steps of the project were:
1) To define a domain of integration common to the three limited-area models;
2) To set a common framework to build the executables of the three limited-area models;
3) To determine the simulation speed of the three limited-area models as the number of computing elements spans from 128 to 6400;
4) To determine the tipping point from the strong to the weak scaling regime.

Summary of problems encountered

(If you encountered any problems of a more technical nature, please describe them here.)

No relevant technical issues were encountered.

Experience with the Special Project framework

(Please let us know about your experience with administrative aspects like the application procedure, progress reporting etc.)

We do not report any issues regarding the administrative aspects. We had, as in previous experiences, fully support from the Special Projects Staff.

Summary of results

(This section should comprise up to 10 pages, reflecting the complexity and duration of the project, and can be replaced by a short summary plus an existing scientific report on the project.)

The targets of the SPITCAPE project were achieved as detailed below.

1) A common domain of integration to the three limited-area models was define, and its workload was named CT2500m. This corresponds to a single domain of integration covering the Italian area with a grid spacing of 2500 m and 60 vertical levels. Number of grid points are 450 in the East-West direction and 540 in the North-South direction. Starting date is 12 UTC 24 October 2011 and the forecast length is set to 30 hours (that is ending date is 18 UTC 25 October 2011). A snapshot of the domain for the three models is shown in Figure 1(a-c).

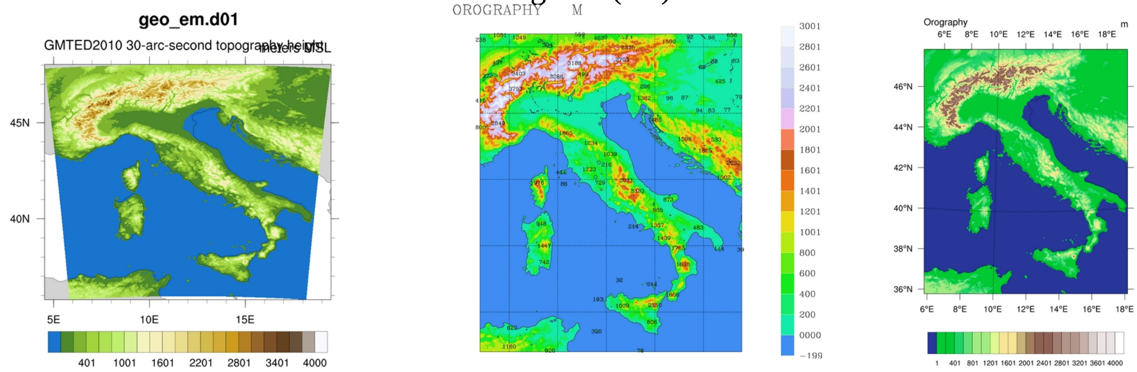

Figure 1: the CT2500m workload: domain of the integration for the WRF (left), MOLOCH (center) and Meso-NH (right) model, respectively.

We underline that, as regards the MOLOCH model only, some constraints exist regarding the relationship between the extent of the domain and the number of computing elements used (Cioni, G., 2014: *Thermal structure and dynamical modeling of a Mediterranean tropical-like cyclone*. M.S. thesis, Fisica del Sistema Terra, Universitá di Bologna, 130 pp.). As a consequence, the number of grid points (and thus the exent of the domain) is slightly different across the simulations when increasing the number of computing elements. However, we found that the difference between the biggest and the smallest domain of the MOLOCH model is approximately 7%.

2) One challenge of the SPITCAPE project was to compile the three limited area models on the new ATOS BullSequana XH2000 supercomputer, using a common framework. This was done only for the Meso-NH and MOLOCH models (using the default Intel compiler, Parallel Studio XE Cluster Edition), whereas the precompiled version fo the WRF model was used to speed-up the completation of the project. However, we underline that the precompiled version of the WRF model is built within the Intel environment.

3) To determine the simulation speed of the three limited-area models as the number of computing elements increases, the following experiments were foreseen in the SPITCAPE request:

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.

Due to the saturaion in speed performances (see detail in the following paragraph), only the first 8 numerical experiments were carried out.

4) To determine the tipping point from the strong to the weak scaling regime, we show in Figure 2 the wall-clock time needed to perform the CT25000m workload for each model.

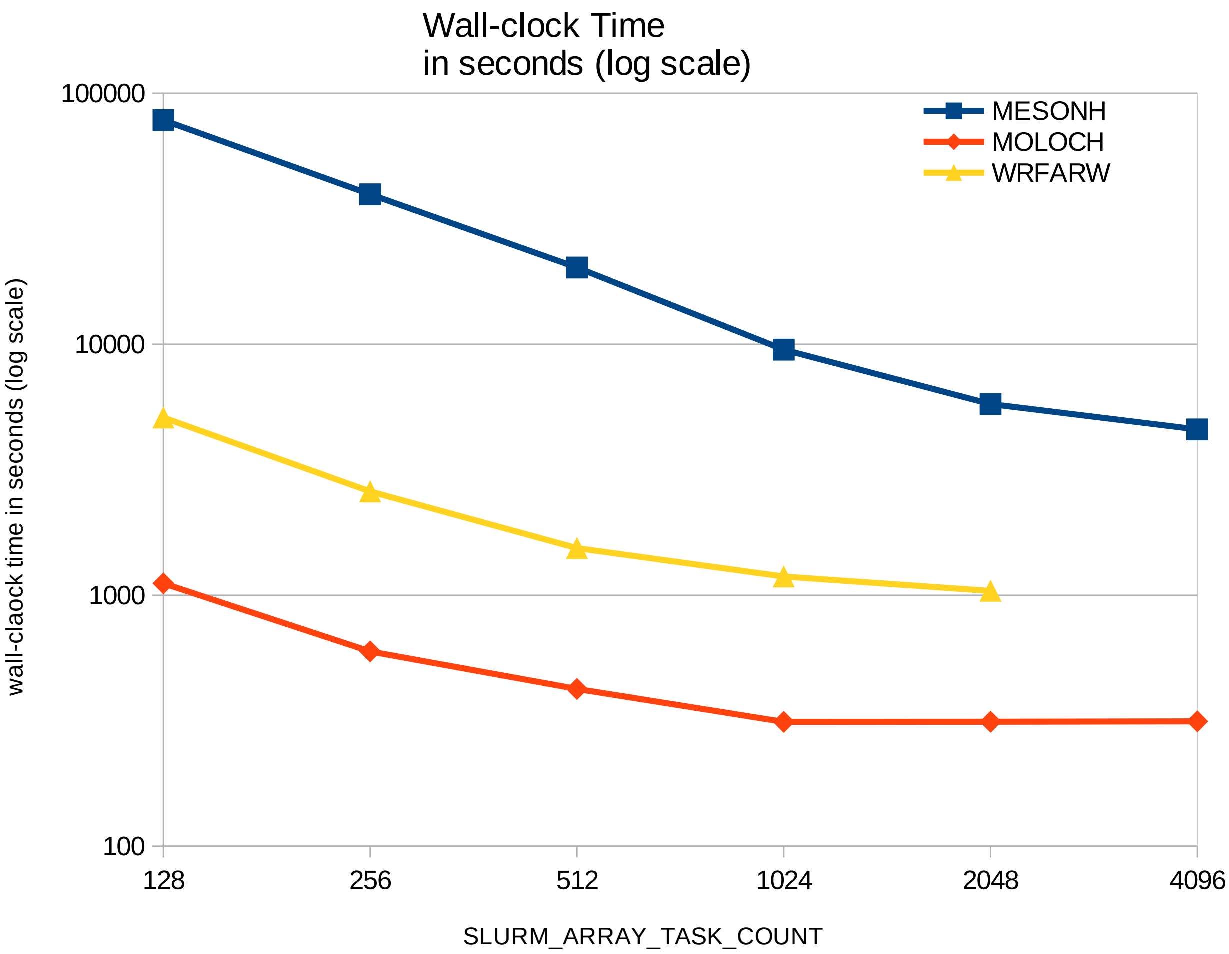

Figure 2: wall-clock time (in seconds) as a function of the computing elements to accomplish the CT2500m workload.

We stress the fact that for computing elements greater than 2048, the WRF-ARW model crashes complaining about the excessive number of cores used. The error message is:

For domain   1 , the domain size is too small for this many processors, or the decomposition aspect ratio is poor.
Minimum decomposed computational patch size, either x-dir or y-dir, is 10 grid cells.

Figure 3 shows the speed-up (in %) of the wall-clock time when doubling the number of computing elements.

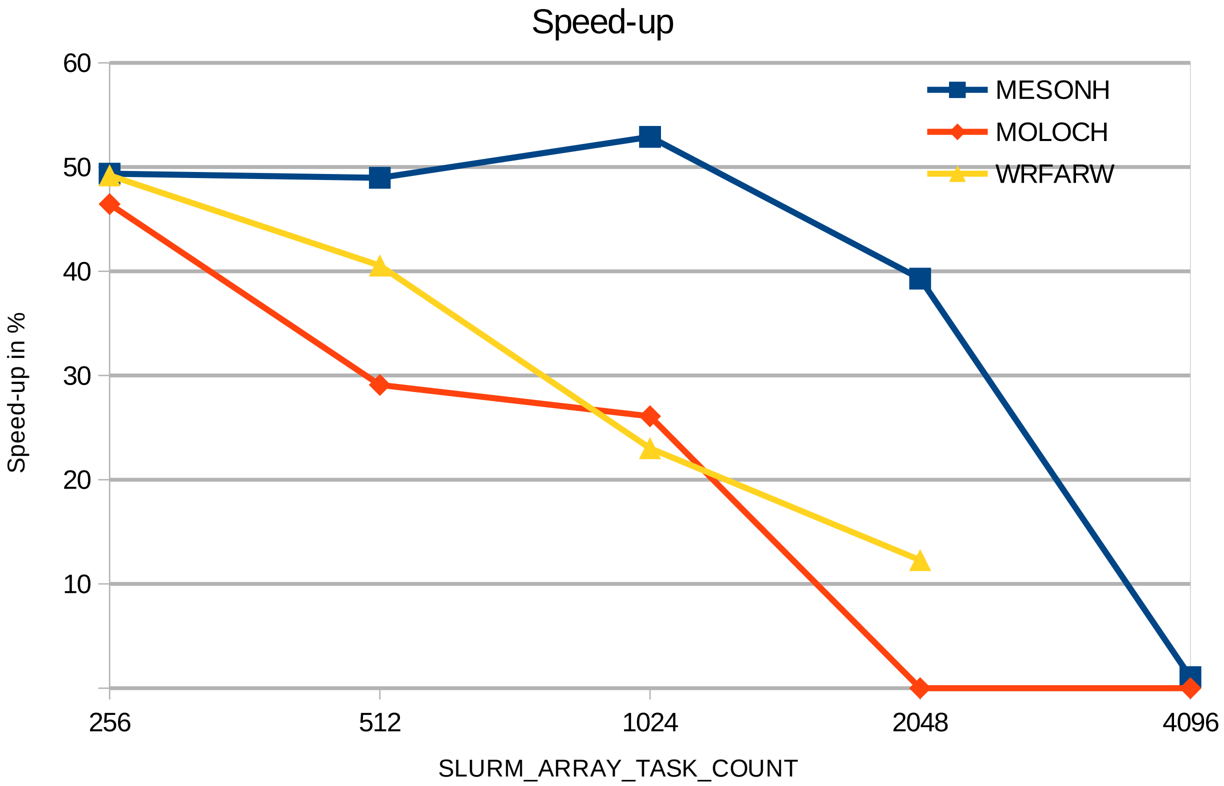

Figure 3: speed-up (in %) as a function of the computing elements

We underline the fact that in a strong scaling regime the colored lines of Figure 3 should approximate the 50% horizontal line. Figure 3 demonstrates that the Meso-NH model scales well (speed-up approximately 40%) up to 2048 computing elements. The performance of the MOLOCH model drops below 30% for computing elements greater than 512, is in the range 25%-30% up to 1024 tasks and is null for computing elements equal or greater than 2048. The WRF-ARW model is in the strong scaling regime up to 512 computing elements (speed-up greater than 40%), then the yellow line drops below 25% and 15% for computuing elements equal to 1024 and 2048, respectively. As stressed above, it crashes when setting more than 2048 cores, indicating that the performance in speed-up are poor.

The secondary goal of the project was detecting how the number of horizontal grid points, named N(x,y), assigned to each core affects the speed-up of the simulation speed. Looking at Figure 3, we state that a model is in the strong scaling regime when the speed-up is greater than the, arbitrary, threshold of 40%. This happens for computing elements up to 2048, 256 and 512 for the Meso-NH, MOLOCH and WRF model, respectively. Considering the geometry of the integration domain set in task 1), that is number of grid points 450 in the East-West direction and 540 in the North-South direction, N(x,y) is 118, 960 and 474 for the Meso-NH, MOLOCH and WRF model, respectively.

List of publications/reports from the project with complete references

none

Future plans

(Please let us know of any imminent plans regarding a continuation of this research activity, in particular if they are linked to another/new Special Project.)

In the request for the SPITCAPE 2022 Special Project, we stated that “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”. In other words, the aim was to assess on which extent the effective spatial resolution differs from the grid one, and how this difference varies across the three models. This can be obtained by computing the kinetic energy (KE) spectra on horizontal wind as a function of scale (Skamarock, "Evaluating mesoscale NWP models using kinetic energy spectra." Monthly weather review 132.12, 2004). So far, we didn’t perform this task; however to achieve the target, we don’t foresee to apply for a new Special Project in the near future. In fact, to compute the KE spectra on horizontal wind, we only need the outputs of the three limited-area models. Thus in the next future we will need to access the files stored in the ECFS, PERM and HPCPERM filesystems.