Centro Operativo per la Meteorologia C.O.MET.

Italian Air Force Meteorological Centre

WG 6 / PP POMPA, 12 September 2017

Aeronautica Militare

Running the GPU version of COSMO model at C.O.MET.

Cap. Alemanno Marco
NWP Modelling Section

Aeronautica Militare

Outline

- NWP facts and figures

- HPC system (HAL supercomputer)

- Assessment of HAL scalability

- Future plans

Operational Numerical Weather Prediction System

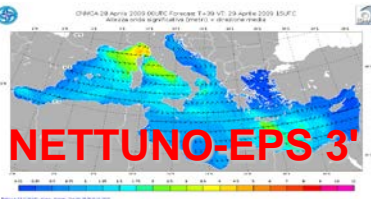**Ensemble Data Assimilation:**
Operational since June 2011

LETKF
(40+1 members)

10 km
45 v.l.

Ensemble Analysis

Deterministic Analysis

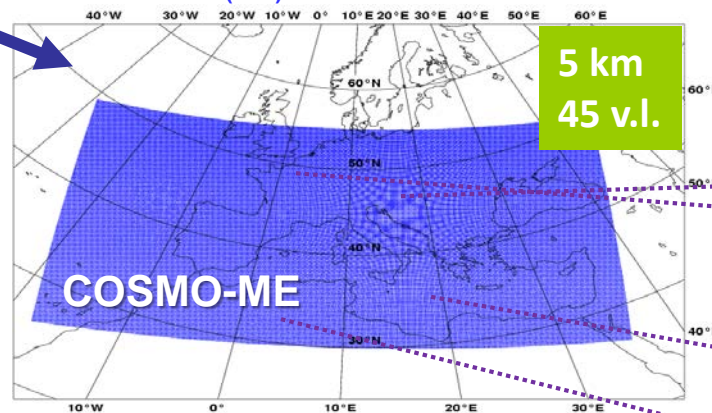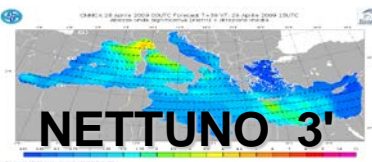ECMWF Boundary Conditions

Local Area Modeling:
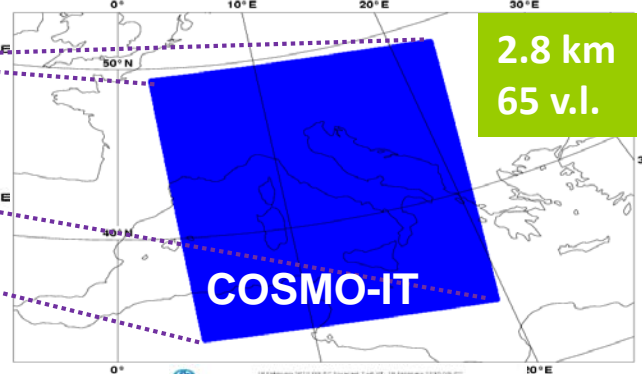
COSMO-ME (5km) ITALIAN MET SERVICE
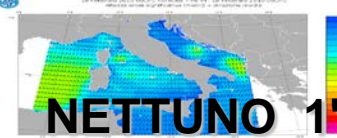
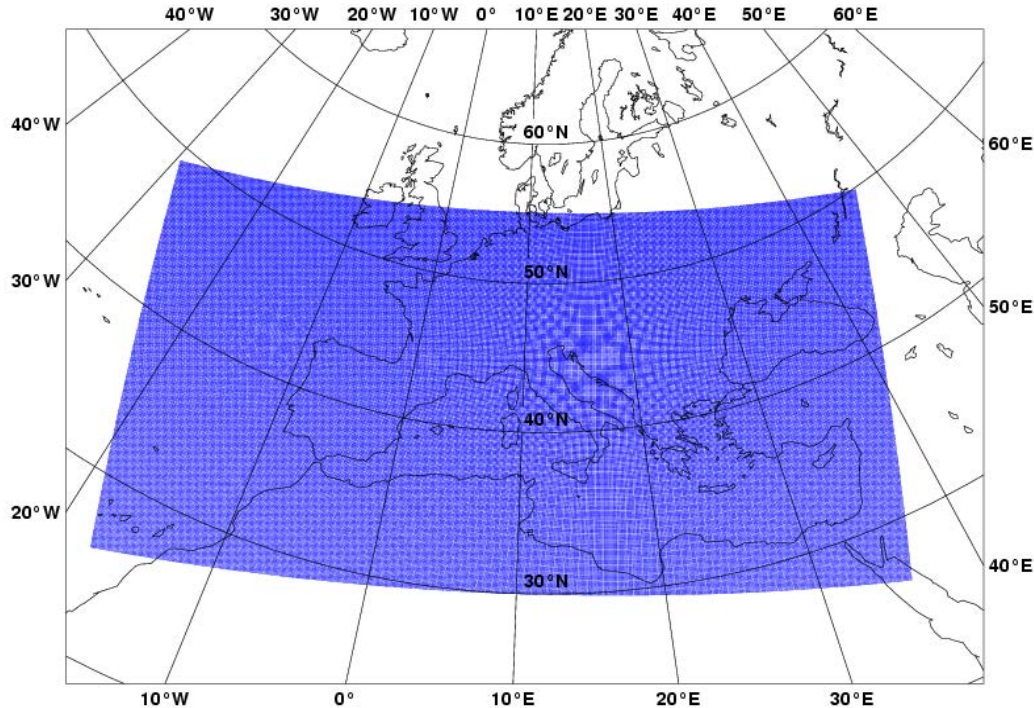5 km
45 v.l.

COSMO-ME

COSMO-IT (2.8Km) ITALIAN MET SERVICE

2.8 km
65 v.l.

COSMO-IT

Ensemble Prediction System:

10 km
45 v.l.

COSMO-ME EPS

NETTUNO-EPS 3'

NETTUNO 3'

NETTUNO 1'

Aeronautica Militare

COSMO-ME
(5 km)

Domain size	1083 x 559
Grid spacing	0.045° (5 km)
Number of layers / top	45 / ~22 km
Time step	60 s
Forecast range	72 hrs
Initial time of model run	00/06/12/18 UTC
Lateral bound. condit.	IFS
L.B.C. update freq.	3 hrs
Initial state	Interpol. LETKF
Initialization	None
External analysis	T,u,v, qv,ps, snow mask
Special features	Filtered topography
Status	Operational

Aeronautica Militare

Our HPC system (now)

**Hybrid CPU/GPU
(CPU Intel, GPU NVIDIA)**
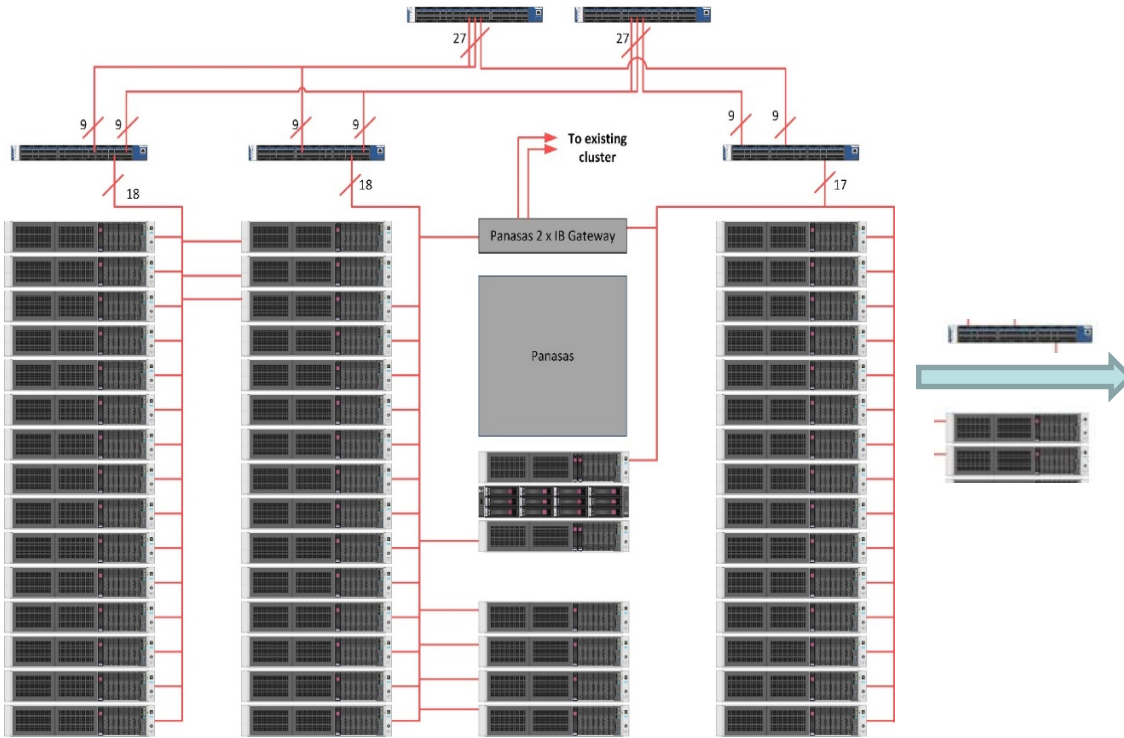
- 56x DL380 G9 Computing Nodes
- 2x DL380 G9 Management Nodes (2x12 Haswell cores - 64 GB)
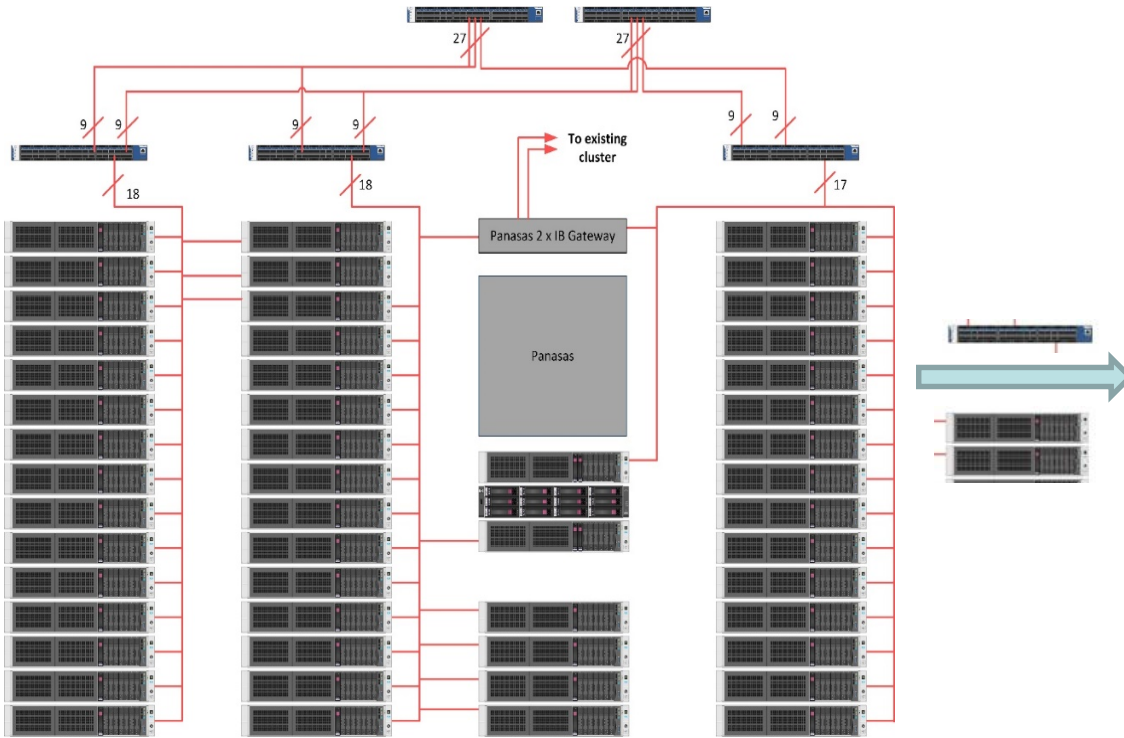- 1x MSA2040 DAS
- 6x Infiniband 36p FDR switches
- 112x Kepler K80 GPUs
- 3.6 TB RAM
- 204 TFLOPS peak (related to xHPL Benchmark)
- 27th place on 111 supercomputers in HPGC benchmark world ranking (1st place in Italy)

Our HPC system (likely in few months...)

**Hybrid CPU/GPU
(CPU Intel, GPU NVIDIA)**

- **66x DL380 G9 Computing Nodes**
- **2x DL380 G9 Management Nodes**
- **1x MSA2040 DAS**
- **6x Infiniband 36p FDR switches**
- **132x Kepler K80 GPUs**
- **8.4 TB RAM**
- **240 TFLOPS peak**

To existing cluster

Panasas 2 x IB Gateway

Panasas

Aeronautica Militare

Our HPC system: how we use it…

- COSMO-ME (vers. 5.04d3) has been running operationally on only-CPU mode.

- COSMO-IT (vers. 5.04d3) has been running experimentally on only-CPU mode with an higher horizontal resolution (2.2 km).

- COSMO-ME EPS (GPU version-40 members) has been running experimentally on GPU-CPU mode.

Our HPC system: how we use it…

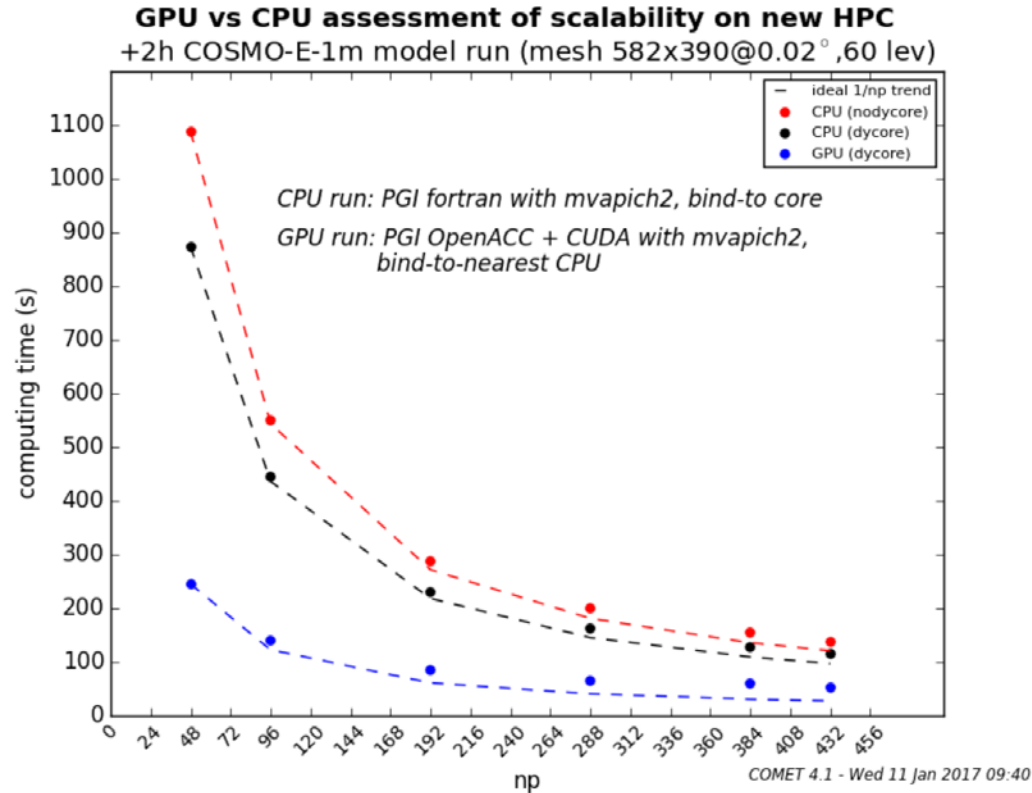- COSMO-ME operational run is completed in only-CPU mode in about 1 hour (it employes 24 nodes in only-CPU-mode)

- COSMO-IT experimental run is completed in only-CPU mode in about 1 hour and 15 minutes (it employes 24 nodes in only-CPU-mode)

- COSMO-ME EPS (GPU version-40 members) experimental run is completed in GPU-CPU mode in about 1 hour and 20 minutes (1 node per each member is employed)

Assessment of HAL scalability

GPU vs CPU assessment of scalability on new HPC
+2h COSMO-E-1m model run (mesh 582x390@0.02°,60 lev)

Legend:
- ideal 1/np trend
- CPU (nodycore)
- CPU (dycore)
- GPU (dycore)

CPU run: PGI fortran with mvapich2, bind-to core

GPU run: PGI OpenACC + CUDA with mvapich2, bind-to-nearest CPU
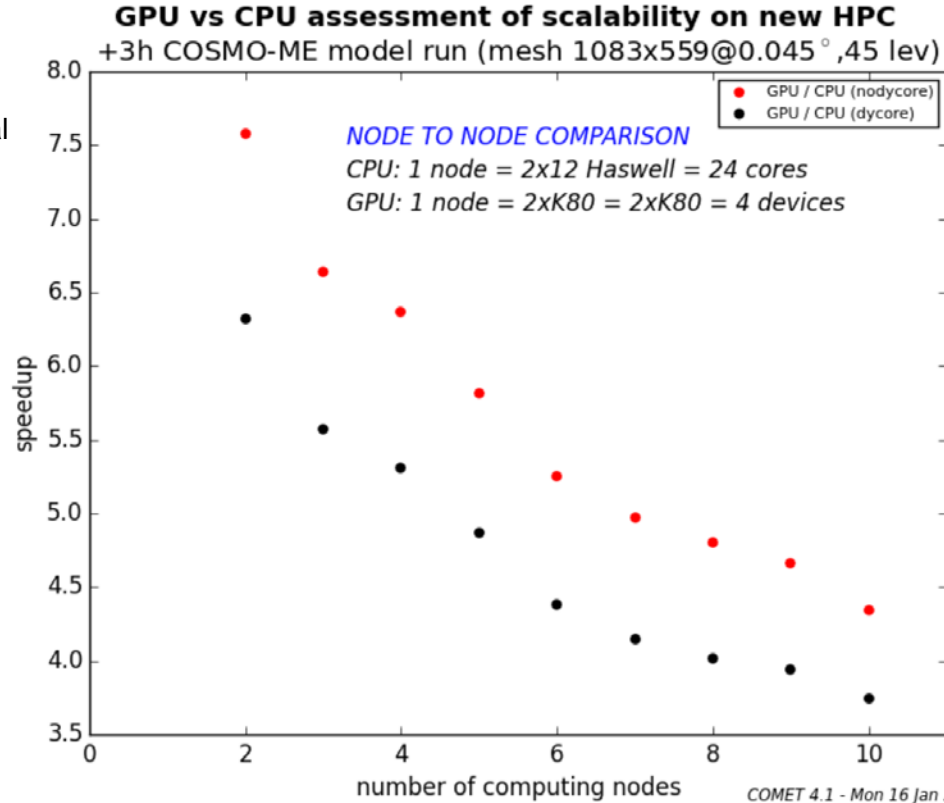
computing time (s)

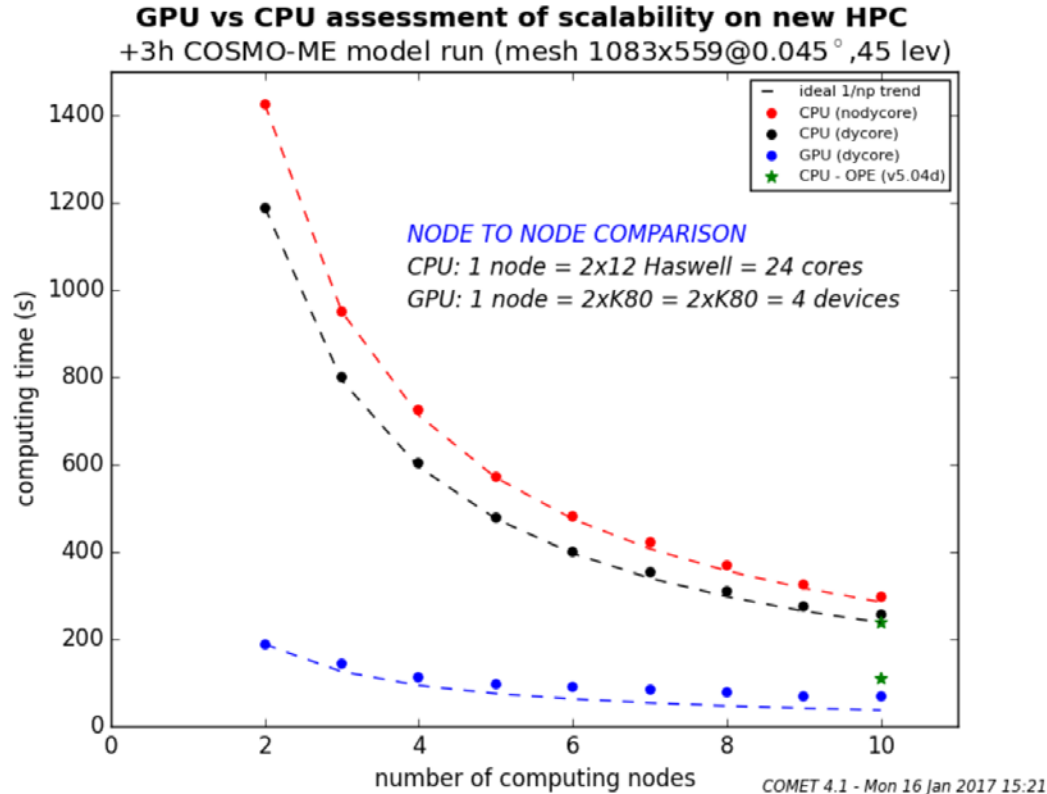np

COMET 4.1 - Wed 11 Jan 2017 09:40

Assessment of HAL scalability (COSMO-ME)

Speedup = only CPU total running time / GPU-CPU total running time

GPU vs CPU assessment of scalability on new HPC
+3h COSMO-ME model run (mesh 1083x559@0.045°, 45 lev)

- GPU / CPU (nodycore)
- GPU / CPU (dycore)

NODE TO NODE COMPARISON
CPU: 1 node = 2x12 Haswell = 24 cores
GPU: 1 node = 2xK80 = 2xK80 = 4 devices

speedup vs number of computing nodes

COMET 4.1 - Mon 16 Jan 2017 15:25

Aeronautica Militare

Assessment of HAL scalability (COSMO-ME)

GPU vs CPU assessment of scalability on new HPC
+3h COSMO-ME model run (mesh 1083x559@0.045°,45 lev)

COSMO-IT (experimental) (2.2 km)
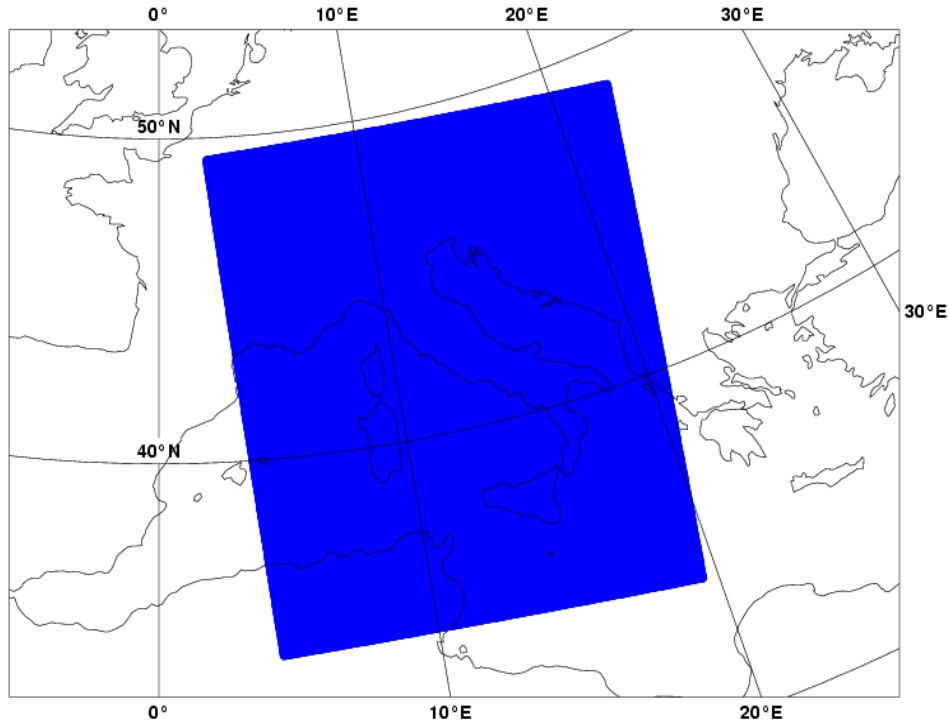
Domain size	576 x 701
Grid spacing	0.02 (2.2 km)
Number of layers / top	65 / ~22 km
Time step	20 s
Forecast range	48/30 hrs
Initial time of model run	00/06/12/18 UTC
Lateral bound. condit.	COSMO-ME
L.B.C. update frequency	1 hr
Initial state	Nudging
Initialization	None
External analysis	None
Special features	Filtered topography
Status	Experimental

Assessment of HAL scalability (COSMO-IT)

GPU vs CPU assessment of scalability of HAL
+3h COSMO-IT model run (mesh 576x701@0.02°, 65 lev)
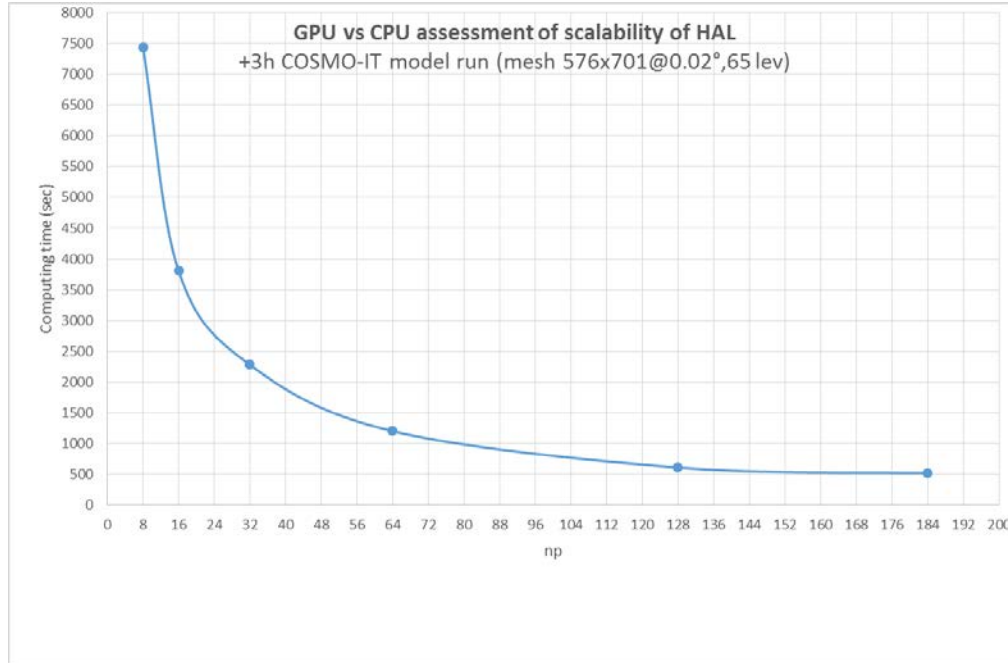
- At least **2** nodes are needed to complete a 3 hour run of COSMO-IT in GPU-CPU mode

- At least **6** nodes are needed to complete a 3 hour run of COSMO-IT in only CPU mode

- Speedup CPU/GPU-CPU= 1.50 for 144 processes and 1.28 for 168 processes

- Speedup CPU/GPU-CPU = 0.30 for 6 nodes and 0.29 for 7 nodes. For the same number of computing nodes, only-CPU run is quicker than GPU-CPU run (number of processes is higher for only-CPU run with the same number of nodes)

Future plans

- Data assimilation processes operational on HAL (both for COSMO_ME and COSMO_IT)

- COSMO-ME EPS (40 members) and COSMO-IT EPS (number of members to be established) operational on HAL in GPU-CPU mode

Any question?

Aeronautica Militare