

SYSTEMS

Cray X1E Configuration

- 16 physical nodes X1E
- 8 MSP each
 - 1,2 GHz, 19,2 Gflops -64 bits- by MSP
 - 32 logical nodes
 - 31 application nodes + 1 support node
- 128 MSP / 512 SSP
- 512 GB memory
- 2,304 Tflops theoretical peak performance for applications.
- Cross-compiler based on a linux cluster



FMI's Cray XT5 Configuration

(thanks to Sami Saarinen and Niko Sokka)

- 2 identical clusters, 17.3 TFlop/s peak for each, ca. 35TF total
- Hex-core AMD Opteron 2.2GHz Istanbul chip
 - 12 (= 2 x 6) cores in a shared memory node
 - 8.8 GFlop/s peak per core, 105.6 GFlop/s peak per node
 - 164 nodes x 12 cores = 1968 cores per each cluster
 - 16 GB shared memory per node (~1.3GB per core)



ECMWF HPCF Configuration

- Two identical clusters (cluster c1a and cluster c1b) each with 272 p6-575 standard compute nodes
- 264 standard nodes with 32 CPUs, 64 GB memory
- 8 high memory compute nodes with 32 CPUs, 256 GB memory
- ~ 8500 physical processors in total
- 9 p6-575 network-I/O nodes, 32 CPUs, 32 GB memory
- 0.6 Petabytes of disk space

HARMONIE (AROME Physics) Experiment

AROME – Spanish Area

Harmonie version 35h1.2
 NH Dynamics
 AROME Physics
 60 secs time step
 1152 x 864 x 60
 Horizontal resolution: 2.5 km.
 Sfc I/O every 15 minutes
 Upper & Sfc I/O every 1 hour
 Building the model with Sami's makeup
 12 hours forecast

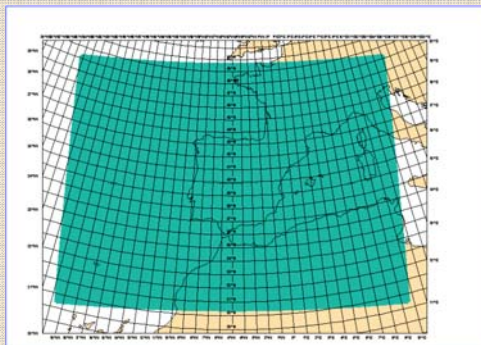
Lots of compilation and linking issues (see Sami's presentation)

Extensive use of memory (only half of the PEs can be used for computations)

Only 1 hour of integration at AEMET Cray X1e in 32 MSP's (128 SSP's)


Huge penalization of I/O time steps

Integration area



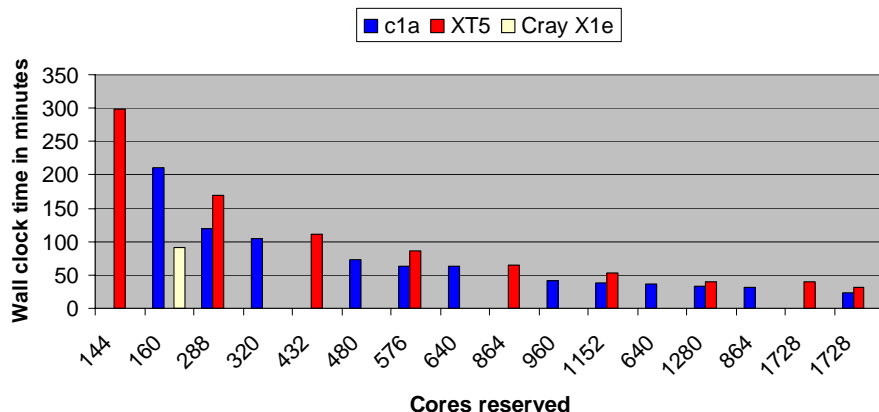
Orography

Arome_2.5km Orografia (grid [m])



Scalability

Scalability of AROME model



For many more details see Sami Saarinen's presentation on Thursday 15th, at the Plenary session 5: "System aspects and verification"

Details

- 1) Cores reserved mean than only half were used for computations
- 2) Wall-clock time for Cray X1e computed from 1 hour integration.
- 3) Different values at the horizontal axis due to different configuration of the nodes.
- 4) For details see table below.

Grid	Cores or PEs	Cores Reserved	c1a	XT5	Cray X1e
6 x 12	72	144		298	
8 x 10	80	160	211		90,98'
12 x 12	144	288	120	170	
8 x 20	160	320	105		
12 x 18	216	432		111	
10 x 24	240	480	73		
12 x 24	288	576	63	86	
16 x 20	320	640	63		
18 x 24	432	864		64	
16 x 30	480	960	41		
24 x 24	576	1152	38	53	
16 x 40	640	640	37		
16 x 40	640	1280	33	39	
24 x 36	864	864	32		
36 x 48	864	1728		39	
36 x 48	1728	1728	24	32	

FUTURE

- New Computer's ITT Autumn 2010
- Benchmark based in Harmonie (AROME Physics)

- Complete set of tests at AEMET Cray X1e
- Tests at SGI HPCF
- Tests at some linux clustes at the University of the Balearic Islands

Acknowledgments

Author would like to acknowledge the useful advice of Sami Saarinen and the support of FMI and, in particular, Niko Sokka for the very fruitful visit paid to FMI and for allowing the tests at their Cray XT5.