NorESM – High resolution scalability experiments

Alok K. Gupta[1], Øyvind Seland[2], Mats Bentsen[1]

[1]Uni Research Klima, Bergen
[2]Meteorologisk institutt

Overview

- Introduction to NorESM
- Resources requirements for CMIP6
- Changes in atmosphere components and scalability
- Scalability of MICOM + CICE
- Coupled scalability results

NorESM – an introduction

- NorESM is based on Community Earth System Model (CESM) from the University Corporation for Atmospheric Research and National Center for Atmospheric Research, Boulder, USA.
- Presently, used in five different countries – Norway, Sweden, Denmark, Finland and India
- Various model components are partially developed and modified such as – MICOM, HAMOCC, CAM extended with own aerosol-chemistry-cloud package, modified energy correction of CAM, new optional air-sea turbulent flux scheme.
- We plan to use quarter degree CAM for HighResMIP, 1 degree CAM for most CMIP6 experiments and mostly, we are going to use quarter degree ocean with most experiments.

NorESM framework and model components

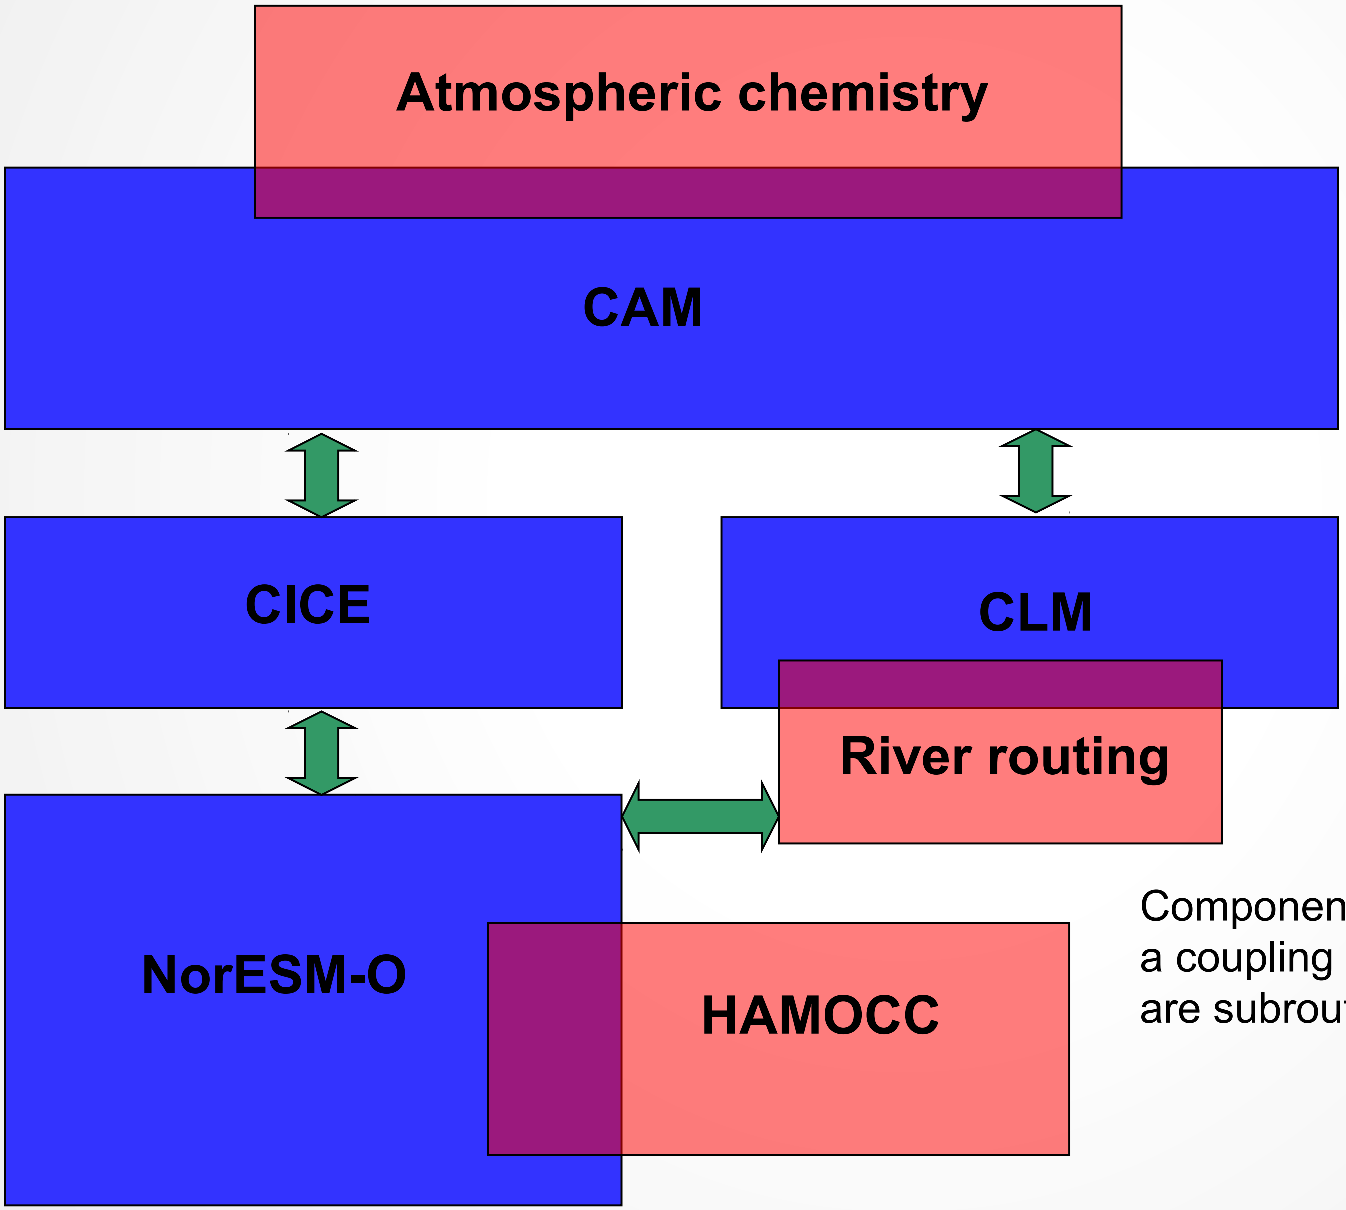

Components in blue communicate trough a coupling component. Components in red are subroutines of blue components.

HighResMIP

- Quarter degree atmosphere + quarter degree ocean

  48 Levels in atmosphere + 53 level in ocean
- 100 years AMIP - tier 1 + 200 years coupled simulation -tier 2
- Require 70 M CPU hours on Hexagon + approx. 1 PB of data
- A common European application through HiPrACE is underway for using PRACE resources for HighResMIP

Total MIPS CPU time + storage

Computationally, two important resolutions:
MH - 1º atmosphere, 1/4º ocean
HH - 1/4º atmosphere, 1/4º ocean
With atmosphere level 30
But, we are planning to go in between 32 to 48

C= Coupled
A= Atmosphere only
O=Ocean only
L=Land only

HighResMIP require approx. 40 % of Total CPU time and 57% of storage

Short name of MIP	Simulation years	CPU hours in M	Storage in TB
Spin-Up+DECK	4673 C	41	179
AerChemMIP	1101 C + 475 A	4.08	11.13
C4MIP	225 C	0.3	69.14
CFMIP	174 A	.04	33.11
DAMIP	1881 C	2.26	42.68
DCPP	6100 C	7.4	1.07
GeoMIP	1029 C	1.5	34.57
HighResMIP	100A +200 C	44	598.98
LS3MIP	165 L+242 C	0.4	1.95
LUMIP	165 L+520 C	0.8	2.54
OMIP	620 O	3.8	36.13
PDRMIP	100 C	0.1	12
PMIP	1500 C	1.8	23.24
RFMIP	180 A+513 C	1.7	1.46
ScenarioMIP	688 C	6.6	15.63
VolMIP	255 C	0.3	9.18
Total		**116.08**	**1071.81**

Changes in atmosphere

- Aerosol parameterisation include a mixture / process addition to log-normal assumptions. Aerosol mixtures used for calculating optical properties and CDNC are a combination of log-normal modes with added process.
- Somewhat higher number of tracers than default CAM5.3 with MAM3 aerosols, 34 vs 25. The majority of these is due to explicit secondary organic aerosol treatment.
- Parameterisation of ice nucleation has been updated to a more recent version of CAM5.X
- Improved treatment of energy conservation and averaging of albedo

PRACE and Norwegian machines

PRACE - Partnership for Advanced Computing in Europe

Preparatory Access for Tier-0 computers (Oct 2014 – Mar-2015) – investigated scalability got access to following three machines

- **SuperMUC**

  IBM System x iDataPlex - Intel® Xeon® Processor E5-2680 2.7 GHZ; cores per node 16; memory per node 32 GB

  Interconnect - Infiniband FDR10

- **Curie TN**

  5040B510 bullx nodes - Intel® Xeon® Processor E5-2680 2.7 GHZ; cores per node 16; memory per node 64 GB

  Interconnect - InfiniBand QDR

- **HORNET**

  Cray XC40 - Intel® Xeon® Processor E5-2680 v3 2.5 GHZ; cores per node 24; memory per node 128 GB

  Interconnect – Aries

- **Hexagon**

  Cray XE6-200 - AMD Opteron 6276 Interlagos CPU 2.3 Ghz, cores per node 32; memory per node 32 GB

  Interconnect – Cray Gemini Interconnect

- **Vilje**

  SGI Altix 8600 - Intel Xeon E5-2670 CPU 2.6 Ghz, cores per node 16; memory per node 32 GB

  Interconnect: Mellanox FDR infiniband, Enhanced Hypercube Topology

Atmosphere scalability
Finite volume Vs Spectral element

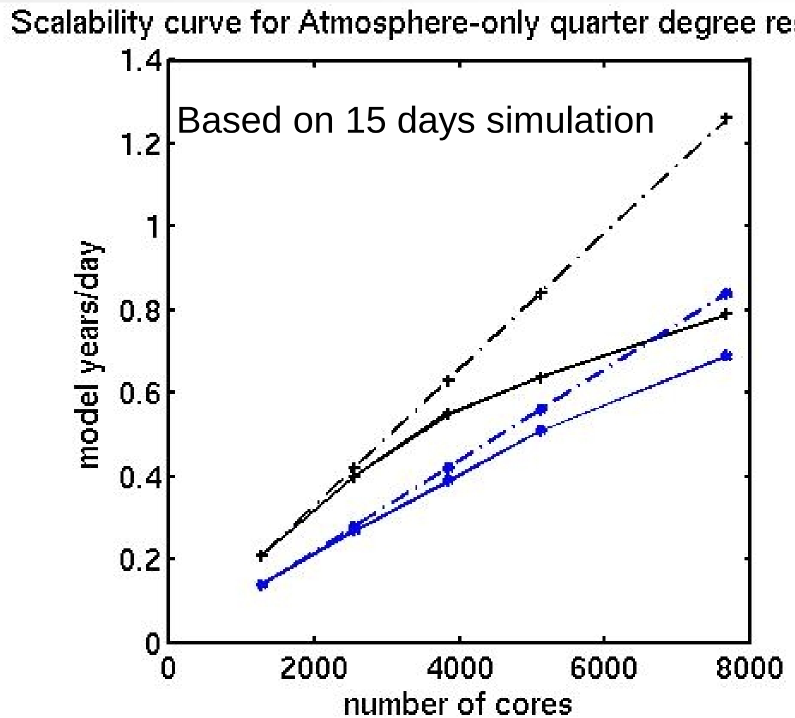

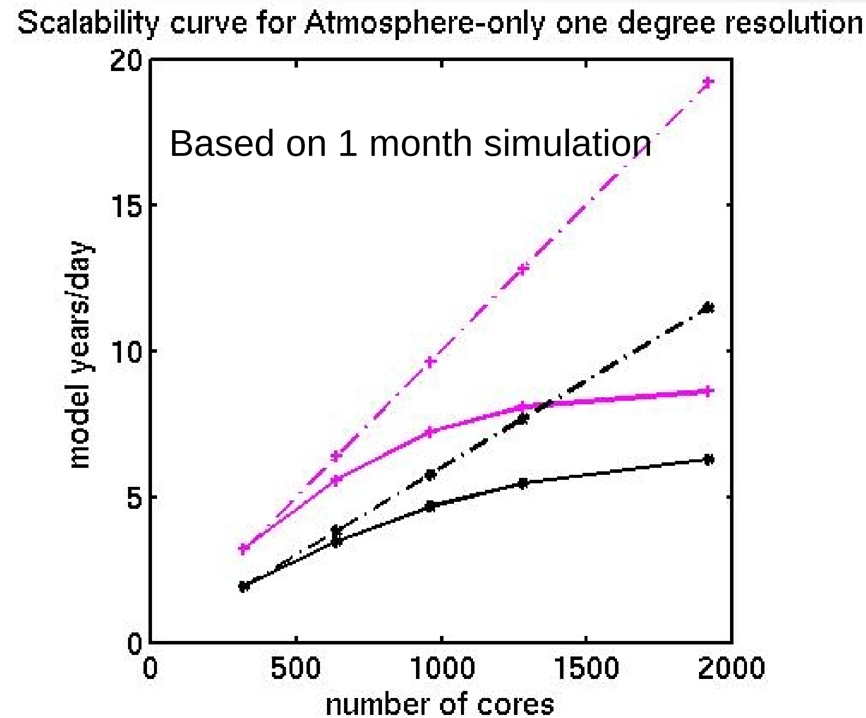

black line - CAM5.3+aerosol chemistry FV grid + Energy correction
blue line - CAM5.3+aerosol chemistry SE grid
magenta line - CAM5.3 FV grid
dashed line - ideal speed up

Aerosol chemistry is approx. 60% computationally costly and follow same scalability curve
Problem scales well for half of the maximum MPI tasks can be used
SE scales well but FV is having a better speed-up

Scalability on PRACE System

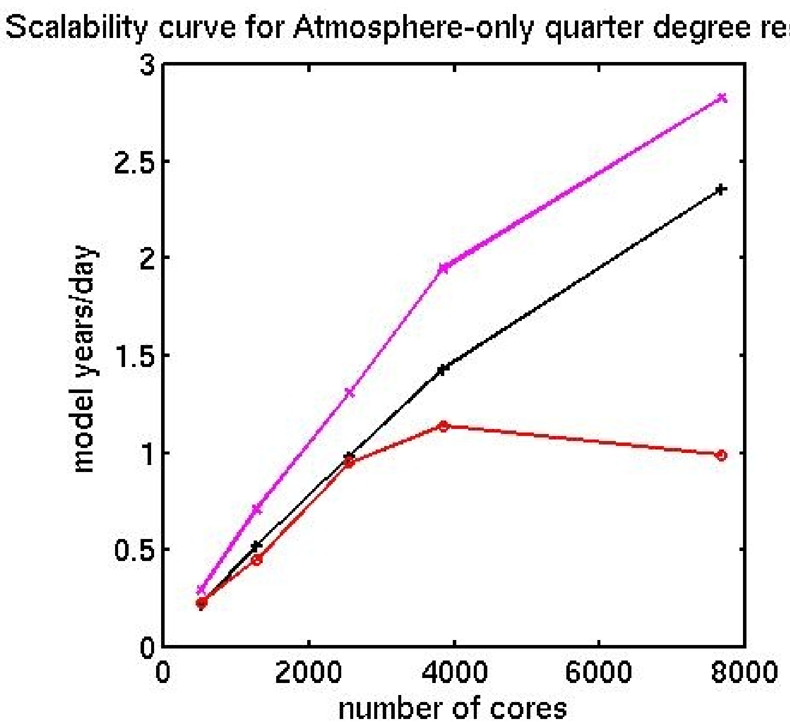

CAM5.3+aerosol chemistry FV grid

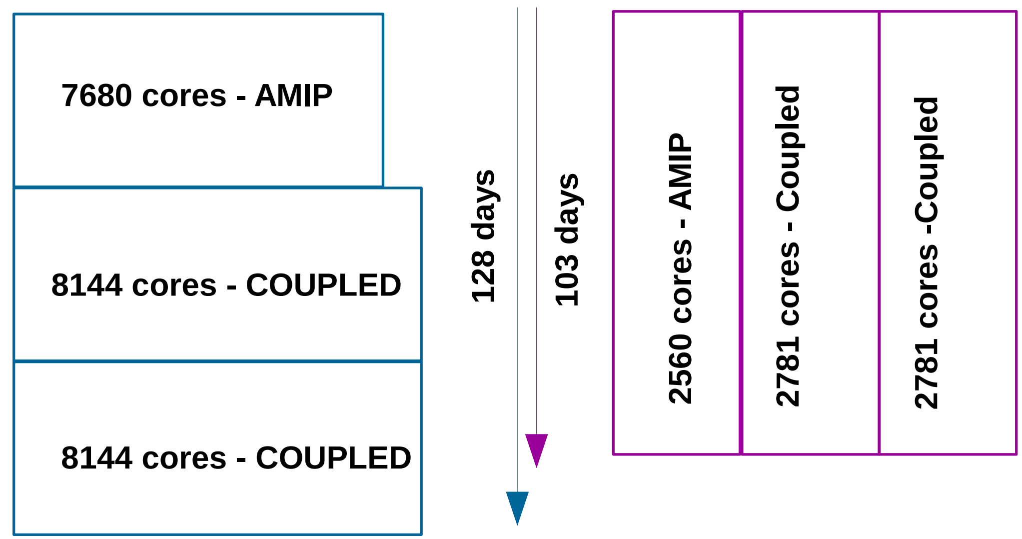

Example- Resource utilisation and management using SuperMUC

No. of Cores	SuperMUC	Curie TN	HORNET
512	.22	.23	.29
1280	.52	.45	.71
2560	.98	.95	1.31
3840	1.43	1.14	1.95
7680	2.36	.99	2.83

MICOM

- MICOM uses the parallel frame work of HYCOM; but, most of the code is changed
- Independent on processors count – explicit solver- same bit-wise results
- Recently, upgraded to parallel I/O – two algorithms are tested
- MICOM+HAMOCC is two times more expensive; HAMOCC is using now 17 tracers
- Able to solve unto 18.5 simulated years per day for quarter degree resolution
- Recently, substantially modified for eddy and turbulent mixing
- We are using tripolar grid

Horizontal grids

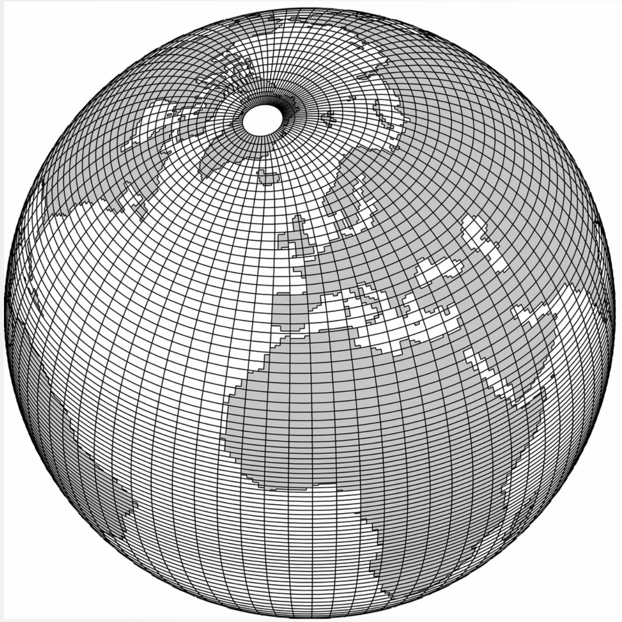

1.125° bipolar grid (every 4th grid line shown):

- 320 × 384 grid cells.
- Used for the NorESM CMIP5 experiments.
- Enhanced meridional resolution near the equator (f_e = 1/4).

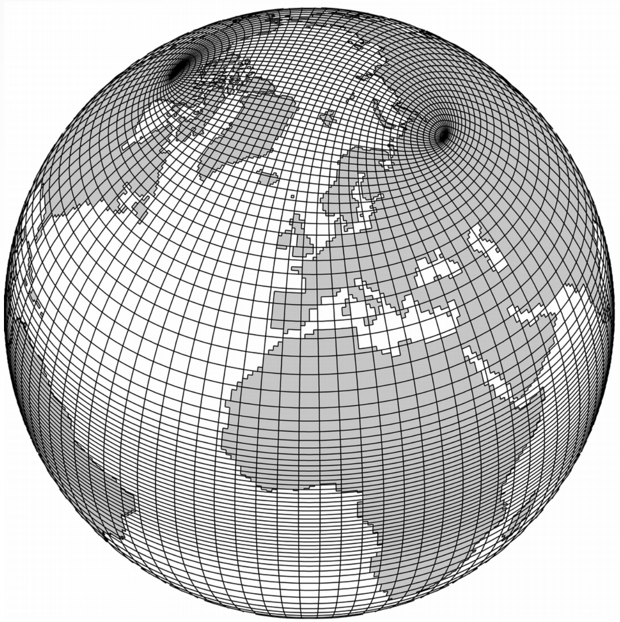

1° tripolar grid (every 4th grid line shown):

- 360 × 384 grid cells.
- Used for the Nor Esm-O CORE2 experiments.
- Enhanced meridional resolution near the equator (f_e = 1/4).

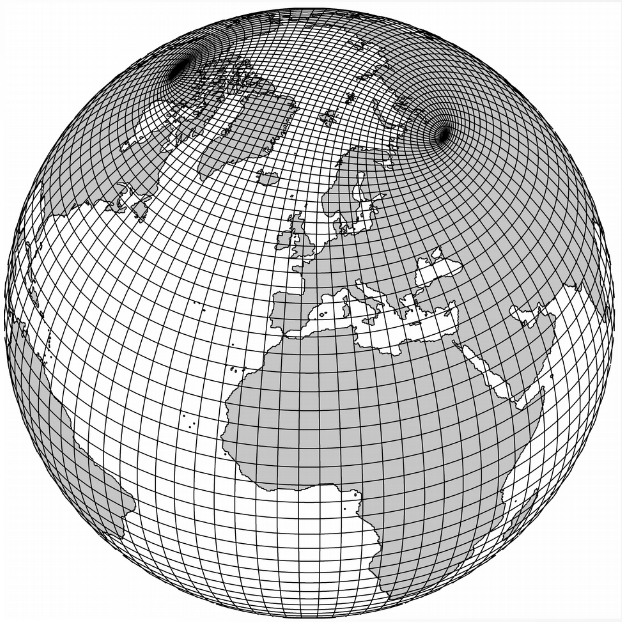

0.25° tripolar grid (every 16th grid line shown):

- 1440 × 1152 grid cells.
- Isotropic grid near equator.
- Target resolution for NorESM CMIP6 experiments.

MICOM decomposition

MPI TASK =221 NX=1440, NY=1152 MPI TASK =2895

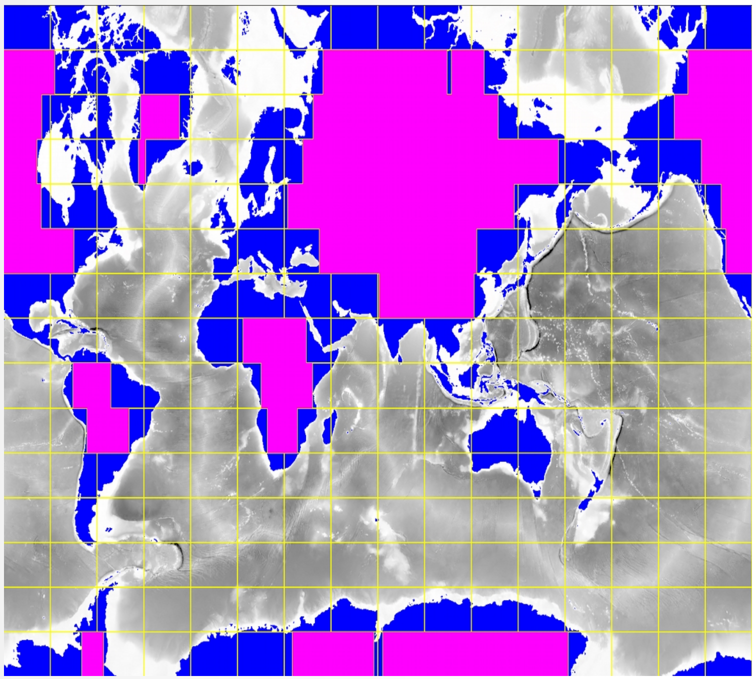

NXP=16 ; NYP=16

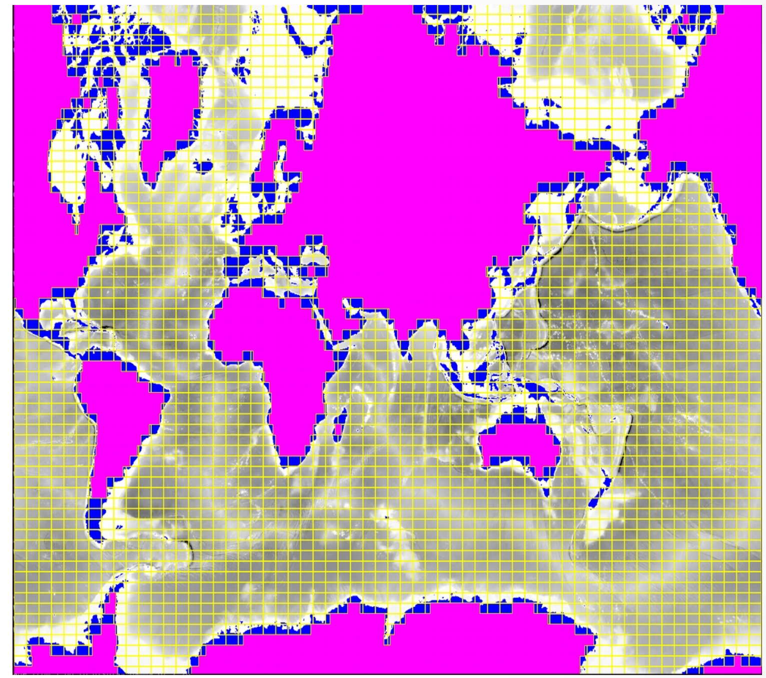

NXP=60 ; NYP=68

Two conditions for decomposition:
NXP should be factor of NX
Row NXP should be symmetric based on ON or OFF
NXP=max number of processors in X-direction
NYP=max number of processors in Y-direction

Parallel I/O – two algorithms

- Serial I/O – horizontal grid is collected on first processor and then written to file in loop along z- direction
- Algorithm 1: Parallel I/O- When every processors is writing its own part in 2-D or 3-D block
- Algorithm 2: Parallel I/O- Whole horizontal row of decomposed domain is gathered on one first processor of that row and then, written; if it is 3-D matrices then whole block

MICOM scalability

8 stripes with 1 MB size, time-step 30 minutes, 3 months simulation
End of simulation day 21 MB, End of month 7.8 GB, restart file – 38 GB

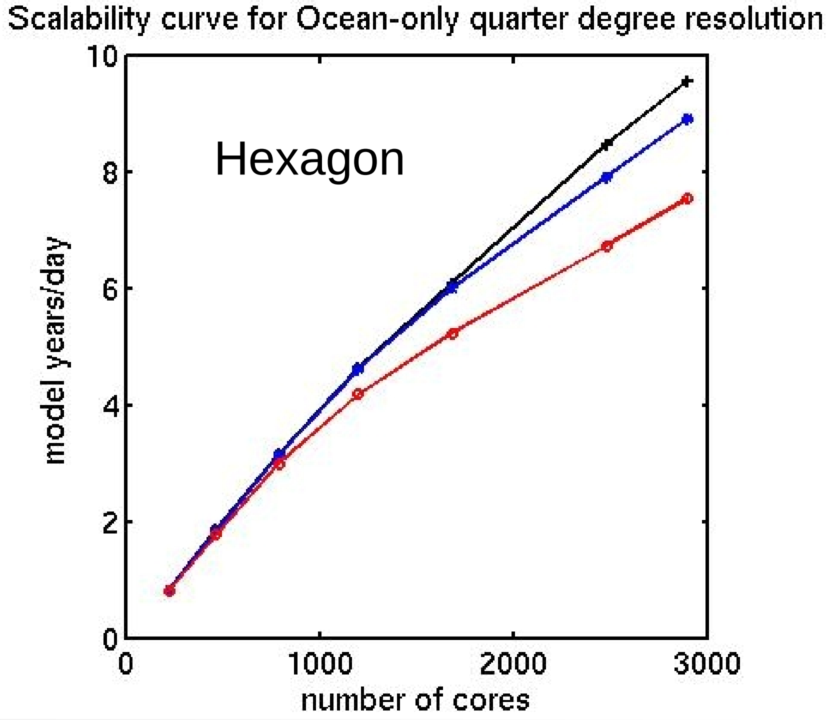

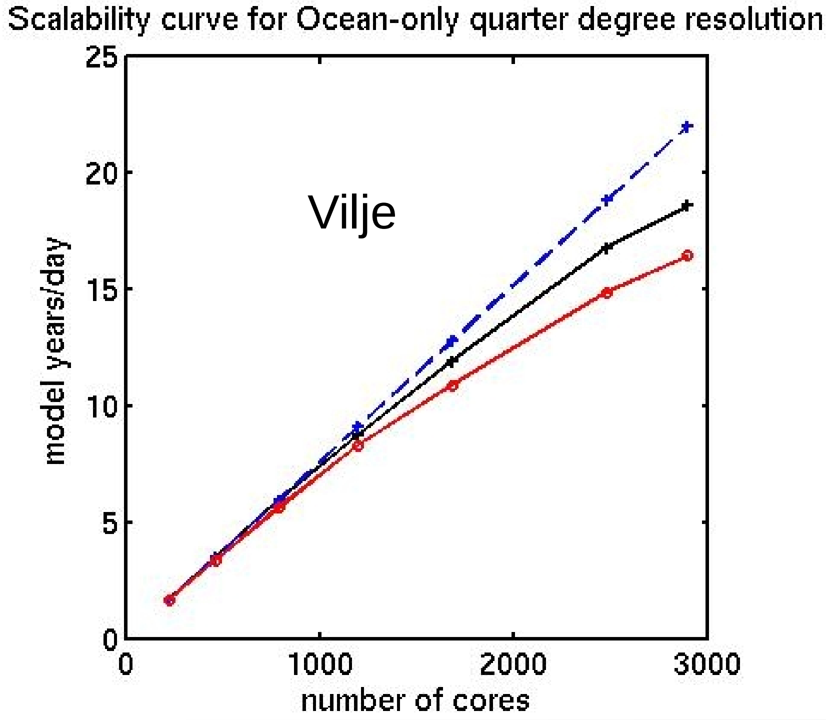

Serial I/O - red line
Algorithms 1 - blue line
Algorithms 2 - black line
Ideal curve - blue dashed line

Stripes	Sim-years/day - Vilje	Sim-years/day - Hexagon
8	18.60	9.57
16	18.64	9.70
32	18.44	
64	18.22	

MICOM on PRACE machines

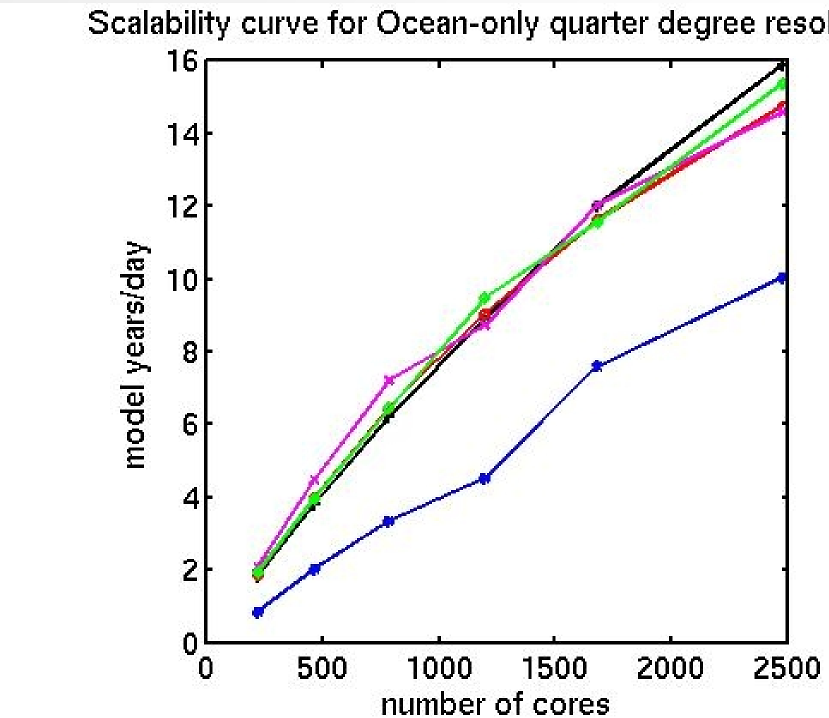

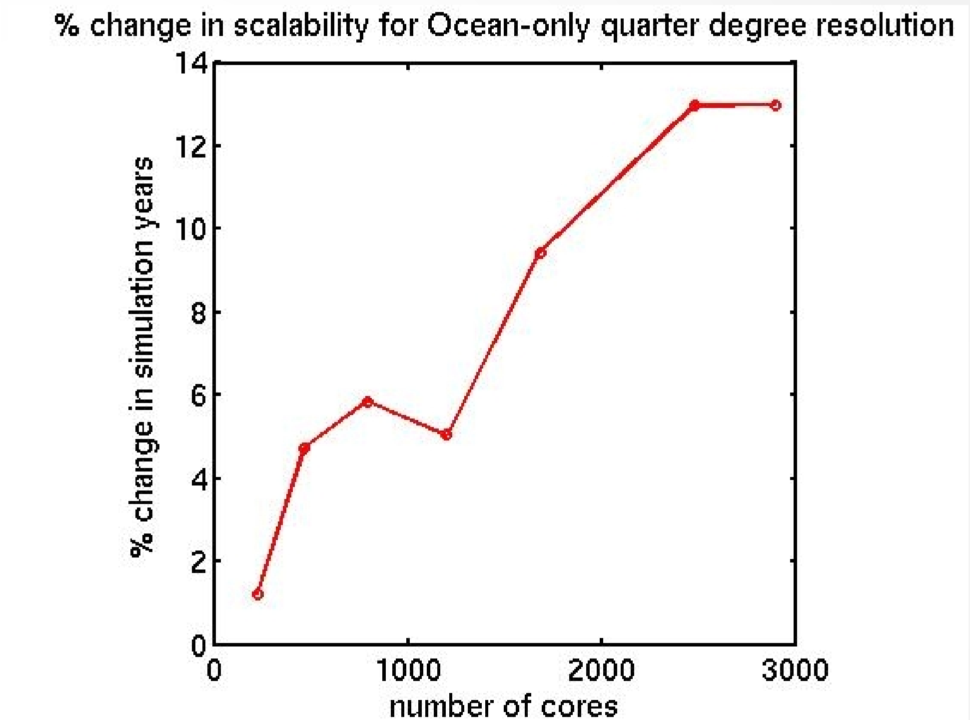

No. of Cores	Hexagon	Vilje	SuperMUC	Curie TN	HORNET
221	.83	1.92	1.82	1.86	2.06
464	2.02	3.96	3.79	3.97	4.46
785	3.33	6.42	6.23	6.45	7.20
1197	4.51	9.47	8.9	9.05	8.71
1681	7.61	11.56	12	11.62	12.02
2475	10.05	15.36	15.89	14.76	14.57

Additional 384 processors are assigned to ice and land components

CICE scalability

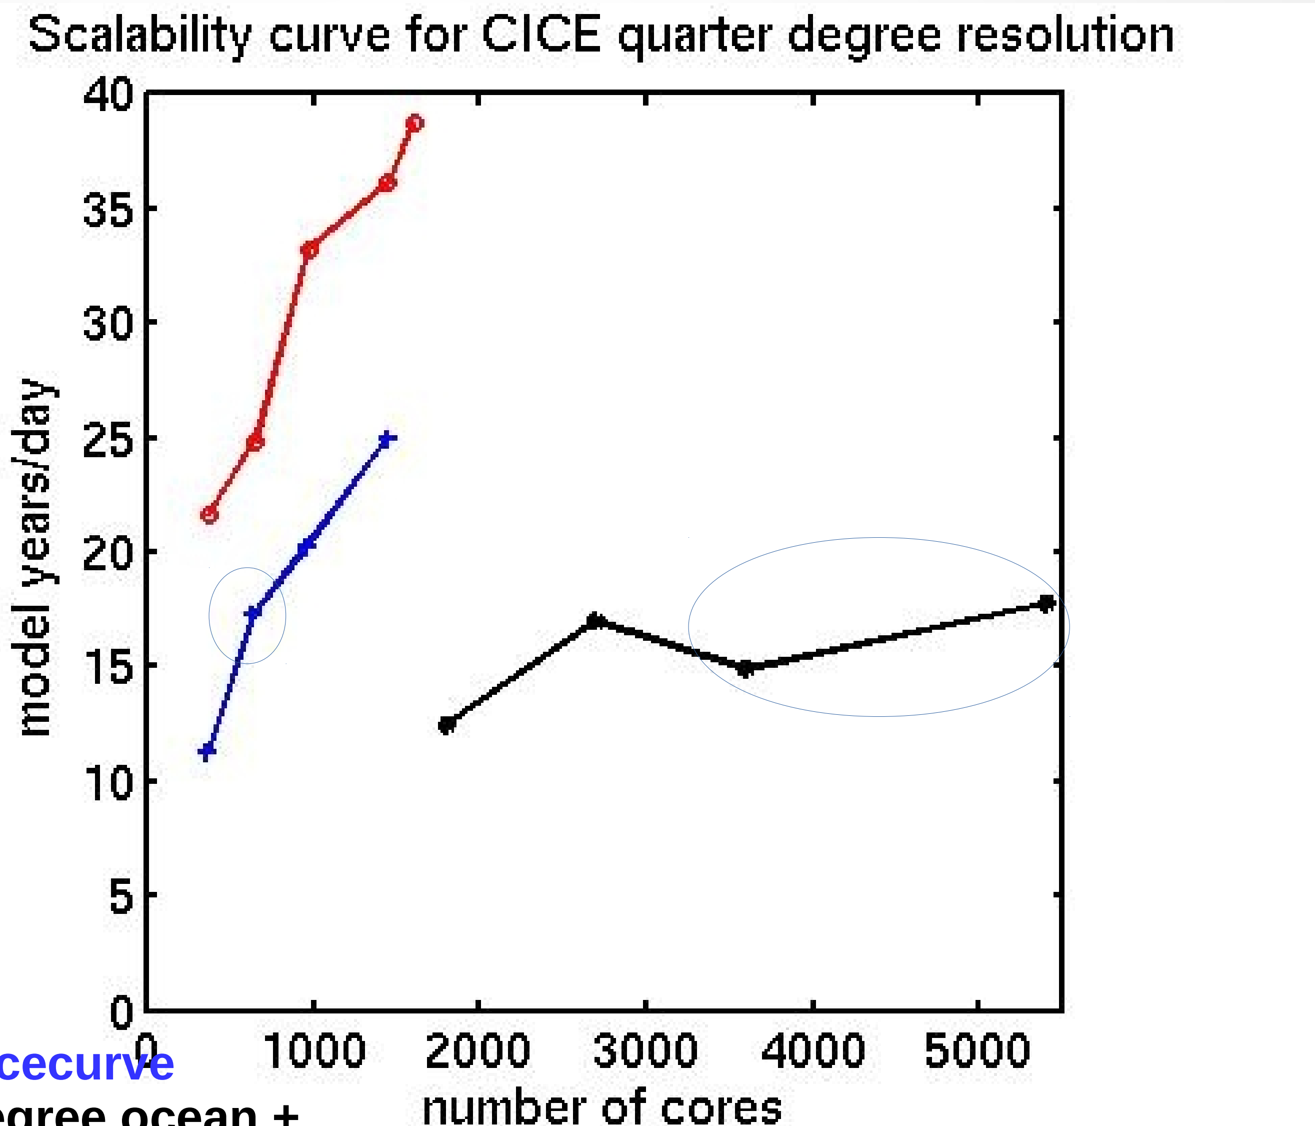

Red line – blkrobin
Blue line – slenderX2 + spacecurve
Black line – from quarter degree ocean +
quarter degree atmosphere coupled simulation
Blkrobin + spacecurve
Points within oval shape - spacecurve

Coupled experiments

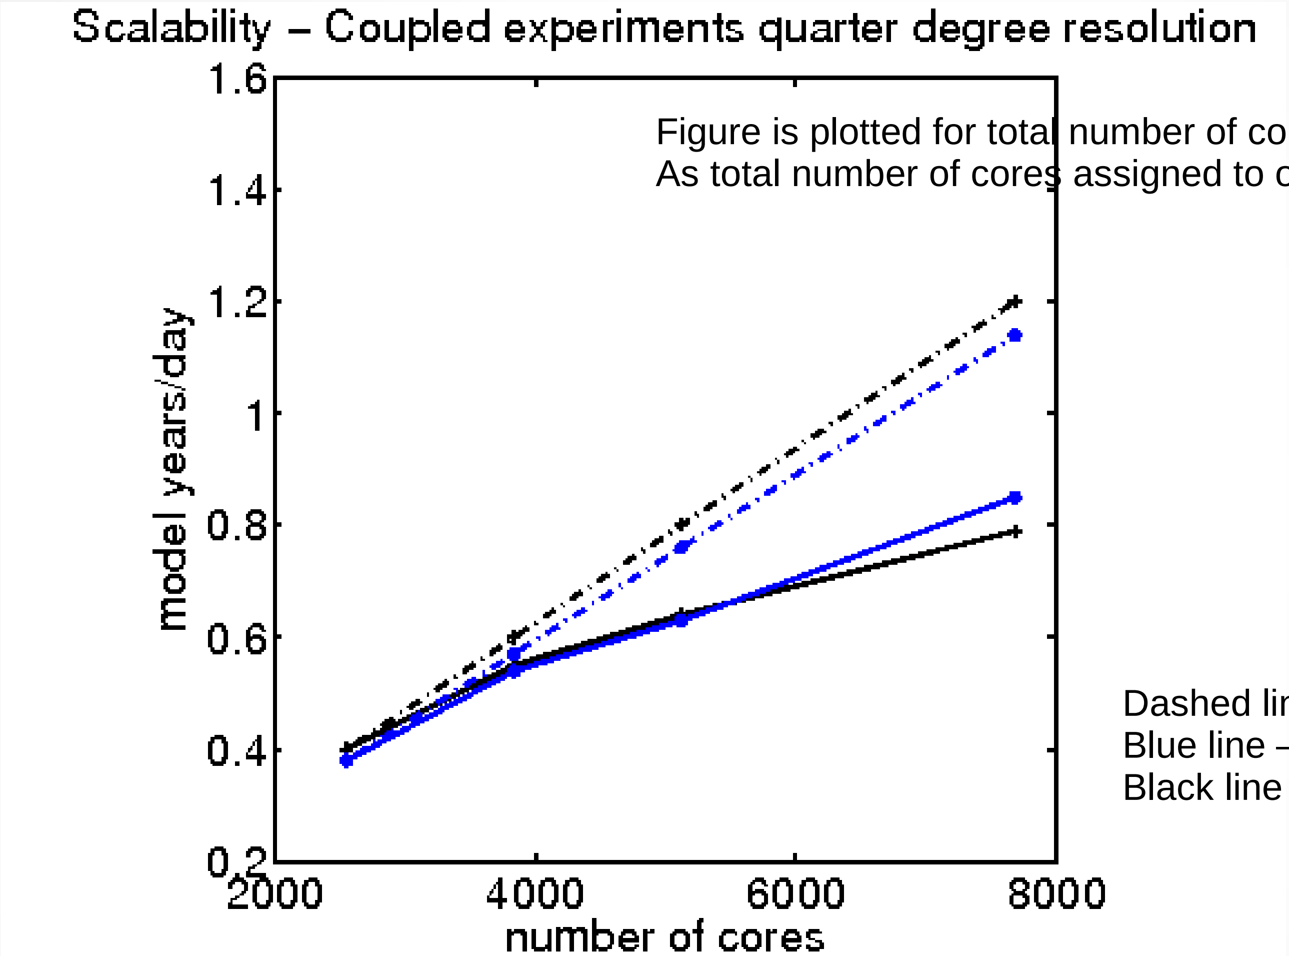

Figure is plotted for total number of cores assigned to atmosphere
As total number of cores assigned to ocean is negligible

Dashed line – approx. Ideal speed up curve
Blue line – coupled simulation
Black line – Atmosphere only simulation

Coupled one degree atmosphere and quarter degree ocean with 3729 processors

Machine	Model years/day- blkrobin	Model years/day - slenderX2
Hexagon	5.92	5.08
Vilje	10.74	9.39

Summary

- Presently, atmosphere component is bottelneck for further scalability of couple system; with aerosol chemistry it becomes approx. 60% more expensive for us.
- Better resource utilisation and managment could save time and resources both
- Climate simulation for high resolution would require huge amount of resources
- Only CPU resources will not be sufficient enough; better algorithms and solver will be more important