ECMWF Workshop on Scalability
April 14-15th, 2014, Reading UK

HPC for climate models: Lessons from IS-ENES projects

Sylvie JOUSSAUME,
CNRS, Institut Pierre Simon Laplace,
Coordinator of IS-ENES2

enes
EUROPEAN NETWORK FOR EARTH SYSTEM MODELLING

http://enes.org/

ENES
European Network for Earth System modelling

is-enes
INFRASTRUCTURE FOR THE EUROPEAN NETWORK FOR EARTH SYSTEM MODELLING

http://is.enes.org/

A network of European groups in climate/Earth system modelling
Launched in 2001 (MOU)

Ca 50 groups from academic, public and industrial world

Main focus :
discuss strategy to accelerate progress in climate/Earth system modelling and understanding

Several EU projects
ENSEMBLES, COMBINE, EUCLIPSE, EMBRACE, SPECS
PRISM, METAFOR, IS-ENES (1& 2)
Collaboration with PRACE

IS-ENES
Infrastructure for ENES

FP7 European projects
IS-ENES 2009-2013
IS-ENES2 2013-2017

Infrastructure
Models & their environment
Model data (ESGF)
Interface with HPC ecosystem

Users :
Climate modelling community
(Global & regional)
Impact studies

Earth System modelling in Europe

NCC	NorESM1-M NorESM1-ME
MPI-M	MPI-ESM-LR MPI-ESM-MR MPI-ESM-P
MOHC	HadCM3 HadGEM2-A HadGEM2-CC HadGEM2-ES
EC-EARTH	EC-EARTH
IPSL	IPSL-CM5A-LR IPSL-CM5A-MR IPSL-CM5B-LR
CNRM-CERFACS	CNRM-CM5
CMCC	CMCC-CESM CMCC-CM CMCC-CMS

CMIP5 in Europe
7 European modelling groups
17 models

CMIP5
Evaluate/Understand/Projections
3400 simul. yrs up to > 12000 yrs
50 expts up to > 160 expts
1000 – 3000 Tbytes (CMIP3: 36)

29 modelling groups
61 models

IPCC AR5

Earth System modelling in Europe

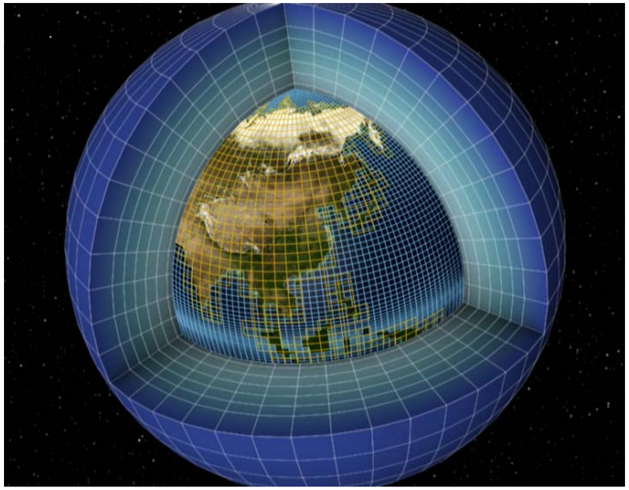

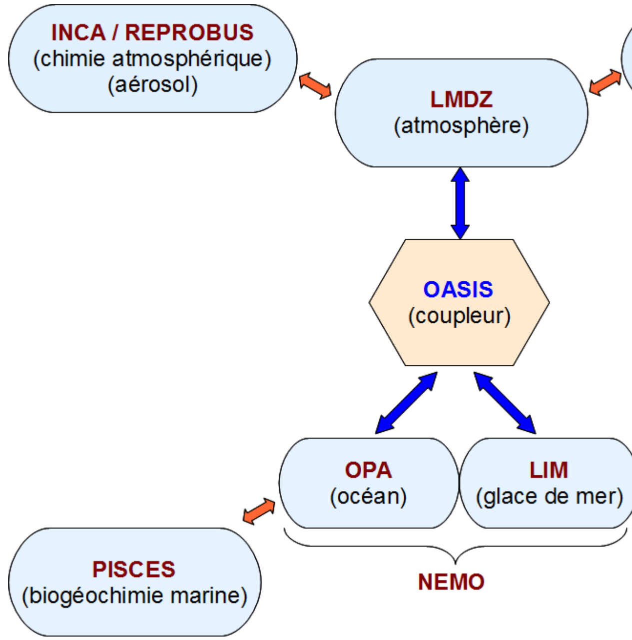

http://enes.org

Country	name of model (CMIP5)	Atmosphere	Ocean	Sea Ice	Coupler	Land Surface *Vegetation	Atmospheric Chemistry	Ocean Bio-geochemistry
Consortium	EC-EARTH	IFS	NEMO	LIM	OASIS	HTESSEL	TM5	
France	IPSLCM5	LMDz	NEMO	LIM	OASIS	ORCHIDEE	INCA	PISCES
France	CNRM-Cerfacs	ARPEGE	NEMO	GELATO	OASIS	SURFEX		
Germany	MPI-ESM	ECHAM5	MPIOM	MPIOM	OASIS	JSBACH*	HAM	HAMOCC
Italy	C-ESM	ECHAM5	NEMO	LIM	OASIS	SILVA		PELAGOS
UK	HadGEM2	UM	UM	CICE	OASIS	TRIFFID*	UKCA	diat-HADOCC
Norway	NorESM	NCAR	MICOM	CICE	CPL7	CLM	Chemistry	HAMOCC

EC-Earth Con Netherlands, Sweden, Ireland, Denmark, Spain, Portugal, Italy, Belgium

IS-ENES : Infrastructure for ENES

FP7 project « Integrating Activities »

http://is.enes.org/

1rst phase: March 2009- Feb 2013 (7.6 M€), 18 partners
2nd phase: Apr 2013- March 2017 (8 M€), 23 partners

Better understand and predict climate variability & changes

Foster:

- The integration of the European ESM community
- The development of ESMs and their environment
- High-end simulations
- The application of ESM simulations for climate change impacts

Support to international coordinated experiments for **IPCC**

CMIP5
7 european models

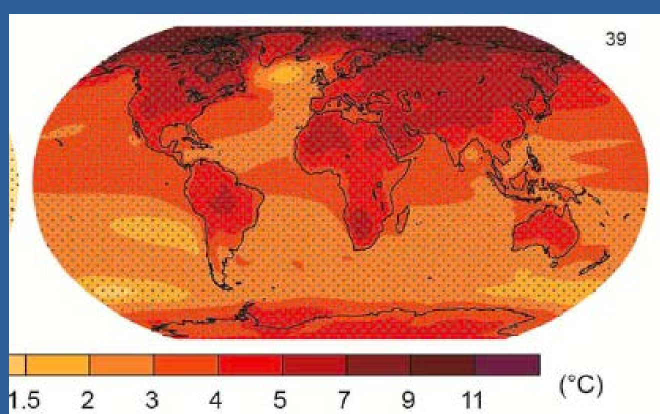

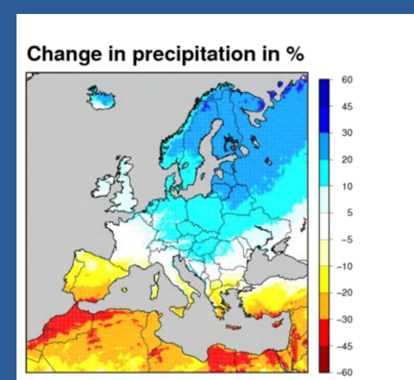

CORDEX

Euro-cordex
Med-cordex
Africa Cordex

Infrastructure Strategy for the European Earth System Modelling Community 2012-2022

Drivers : Science & Society

From understanding to development of "Climate Services"

Grand challenge: towards global 1 km climate models
(resolve convective clouds)

Recommandations:

1) Access to world-class HPC for climate - «tailored » up to « dedicated »
2) Develop the next generation of climate models
3) Set up data infrastructure (global and regional models) for large range of users from impact community
4) Improve physical network (e.g. link national archives)
5) Strengthen European expertise and networking

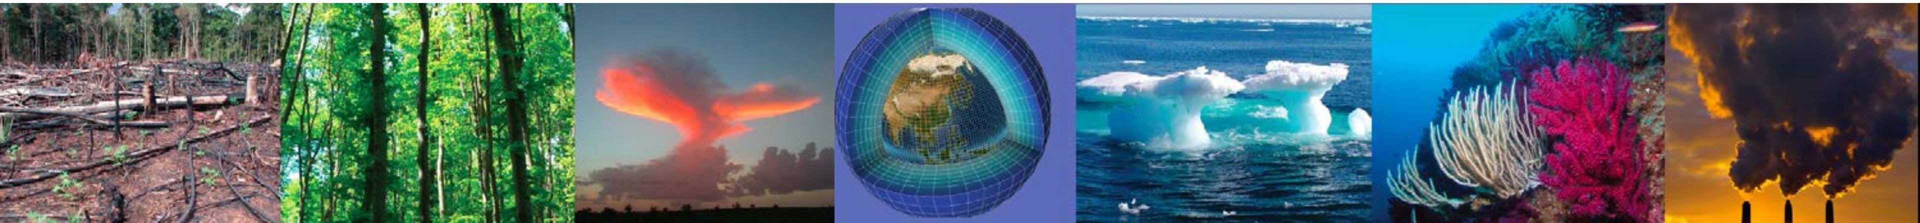

HPC for climate models

HPC facilities

Resolution, Complexity, Ensembles, Duration

Tier 1: Mainly national facilities, dedicated or general-purpose

Tier 0: projects on PRACE

Towards an agreement with PRACE for high-end experiments for CMIP6

ENES HPC Task Force

Models and their environment

Improve model performance

parallel I/O, coupler, tests

Future model generation :

dynamical cores, numerics, algorithms

Physics (eg radiation)

Share best practices: Workflows

Data

storage, distribution & analyses

Parallel I/O

Efficiency of post-processing: CDO

Distributed database (ESGF) :

data & metadata standards, developts

http://is.enes.org/

Technology Tracking

Share experience on accelerators

Strengthen interactions with vendors

IS-ENES HPC Workshops

Lecce 12/2011

Toulouse 01/2013

Hamburg, 03/ 2014

ENES Workshops on HPC for climate models

http://is.enes.org/

2nd Workshop Toulouse, 30/01-01/02/2013

Model performance, dynamical cores,
use of PRACE
EU exascale projects

Improve model performance:

- to be driven by science
- Performance intercomparison needed : identify & share best practices
- Need interdisciplinary teams: climate and computational

Technology tracking: not convinced of GPUs

Prepare future models

- Need for new dynamical cores to enhance scalability
- Separate science from technical software ?

Data challenge : exabyte even more challenging than exaflops

BAMS, workshop report, André et al., May 2014

Scalability issue

Joint Weather and Climate Research Programme

A partnership in climate research

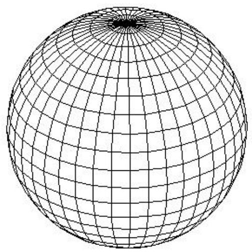

Limitation of Scalability
e.g. at resolution 25-30 km for the atmosphere

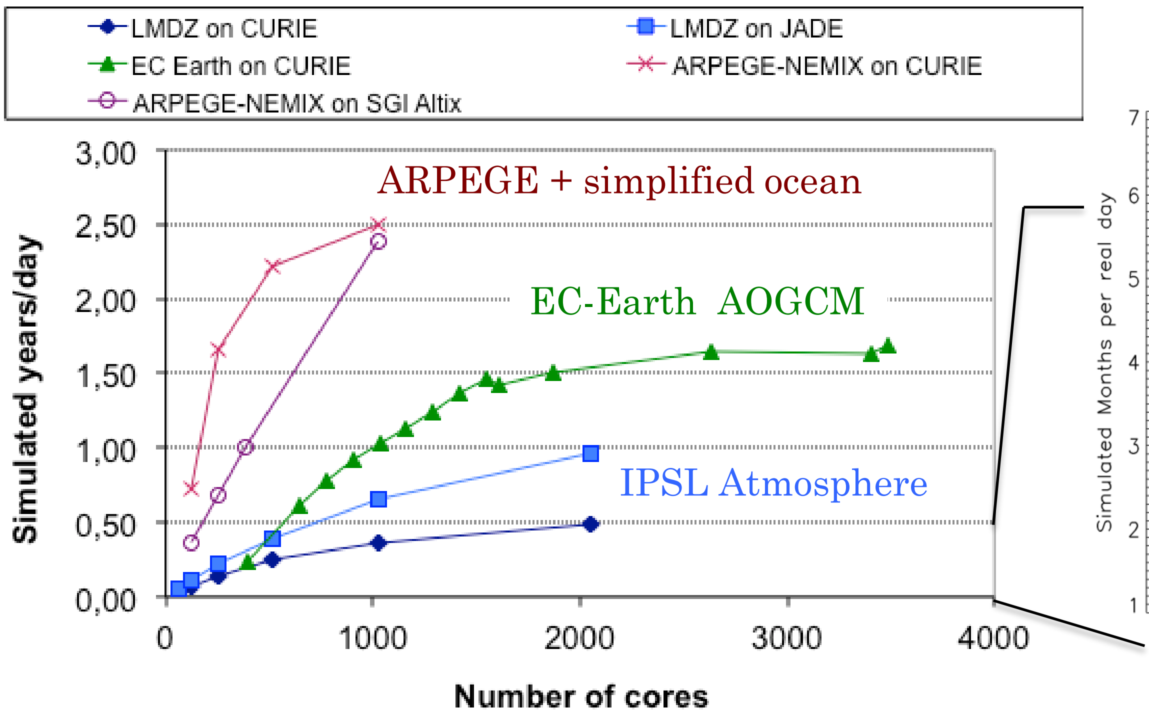

G. Riley et al., IS-ENES (2012)

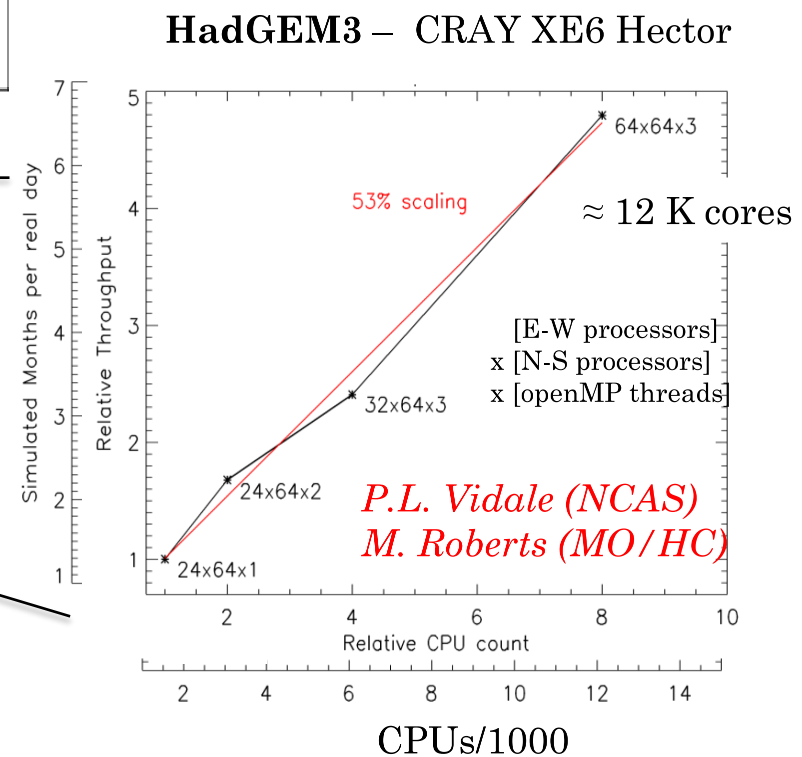

Model performance:
Need to revisit dynamical cores

On-going international projects:
G8 exascale project ICOMEX
Dynamical Core MIP

Cubed-sphere
(CAM-SE)

Collaboration NCAR-Sandia,
Dennis et al. (IJ HPC appl, 2012)

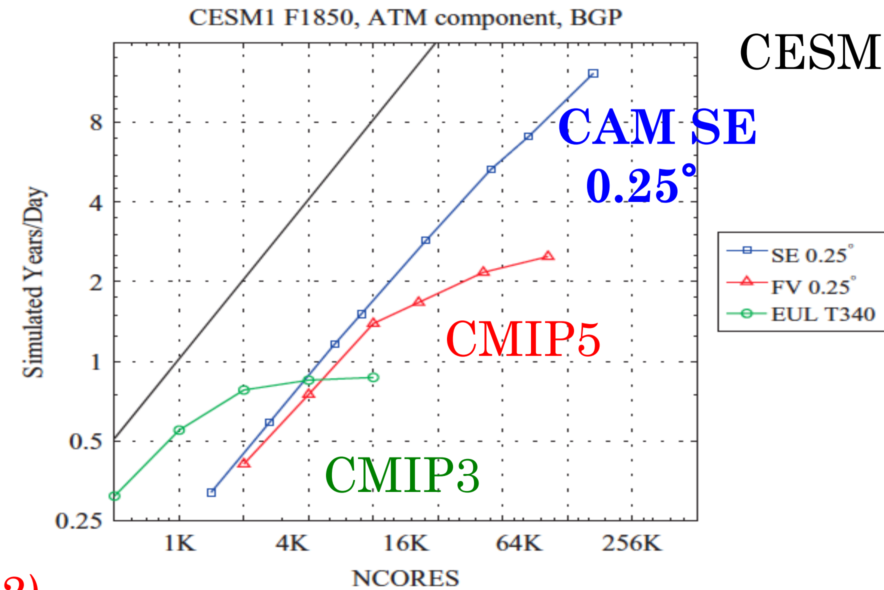

CESM1, 0.25° , BGP

Europe: 3 projects
ICON (DE)
DYNAMICO (FR)
LFRIC (UK)

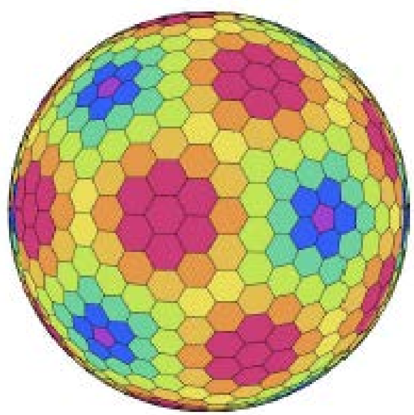

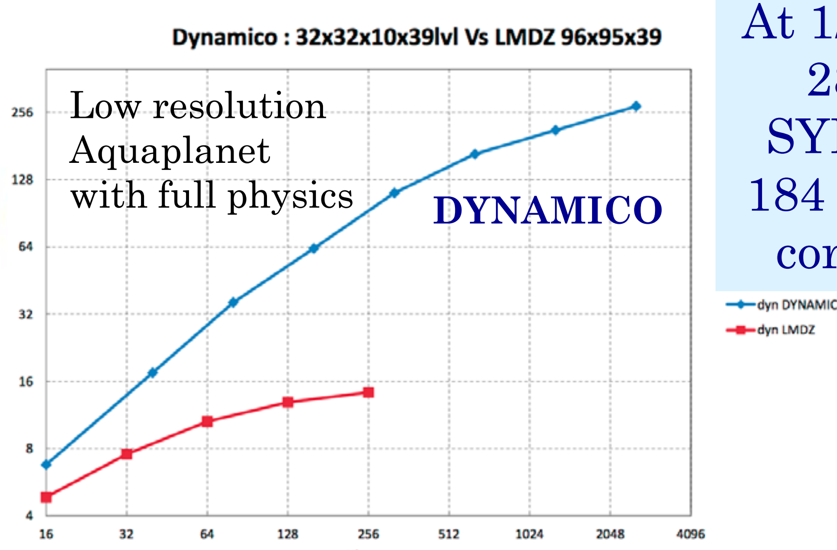

At 1/3°
28
SYPD
184 000
cores

Court. T. Dubos et Y. Meurdesoif (IPSL)

Computing performance : comparison

Real model performance: some considerations

- Productions runs may be configured for capability (minimizing time to solution or SYPD) or capacity (minimizing allocation or CHSY).
- Computing resources can be applied to resolution or complexity: what is a good measure of model complexity?
- ESM architecture governs component concurrency: need to measure load balance and coupler cost.
- Codes are memory-bound: locate bloat (memory copies by user or compiler).
- Models configured for scientific analysis bear a significant I/O load (can interfere with optimization of computational kernels). Data intensity (GB/CH) is a useful measure for designing system architecture.
- Actual SYPD tells you if you need to devote resources to system and workflow issues rather than optimizing code.

Balaji Talk

Agreed: Common diagnostics

Get in metadata

Computing performance : comparison

Balaji Talk

Model	Resolution	Cmplx.	SYPD	CHSY	Coupler	Load Imb.	I/O	MBloat	ASYPD
CM2.6 S	A0.5L32 O0.1L50	18	2.2	212,465	5.71%	20%		12%	1.6
CM2.6 T	A0.5L32 O0.1L50	18	1.1	177,793	1.29%	60%	24%	12%	0.4
CM2.5 T	A0.5L32 O0.25L50	18	10.9	14,327	17%	0%			6.1
FLOR T	A0.5L32 O1L50	18	17.9	5,844	0%	57%	5.1%	31%	12.8
CM3 T	A2L48 O1L50	124	7.7	2.974	0.5%	41%	14.76%	3%	4.9
ESM2G S	A2L24 O1L50	63	36.5	279	8.91%	1%		34%	25.2
ESM2G T	A2L24 O1L50	63	26.4	235	2.63%	22%		34%	11.4

These are very preliminary results (many thanks: Rusty Benson, Seth Underwood, Niki Zadeh!) but seem to provide a basis for analyzing results across models and platforms.

Technology tracking : accelerators

Rich Loft (NCAR): missing x 20 at 2020 for 1 km
Scalability, Memory,
Node performance (accelerators ?)

Talks on Xeon Phi

Still difficult to optimise

Review by W. Sawyer

Summary, final thoughts

- Disappointment depends on your expections...
- Numerous efforts in porting to GPUs, many are demonstrative and consists of bits and pieces
- Dynamics: GPU performance mirrors increase in memory bandwidth, e.g., 2x for K20x vs. dual-socket SNB
- Physics: increased computational intensity gives larger performance benefit on GPU
- At least one model is near production status on GPU
- Acceptance of GPU paradigms by model developers is a problem

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

CSCS
Swiss National Supercomputing Centre

Order of x 2
Better with regards to energy

New ways of structuring codes ?

Separating Science/ technology layers

Data issues

CMIP5 ≈ 2 to 3 PB
CMIP6 x 30 ?

Earth System Grid Federation

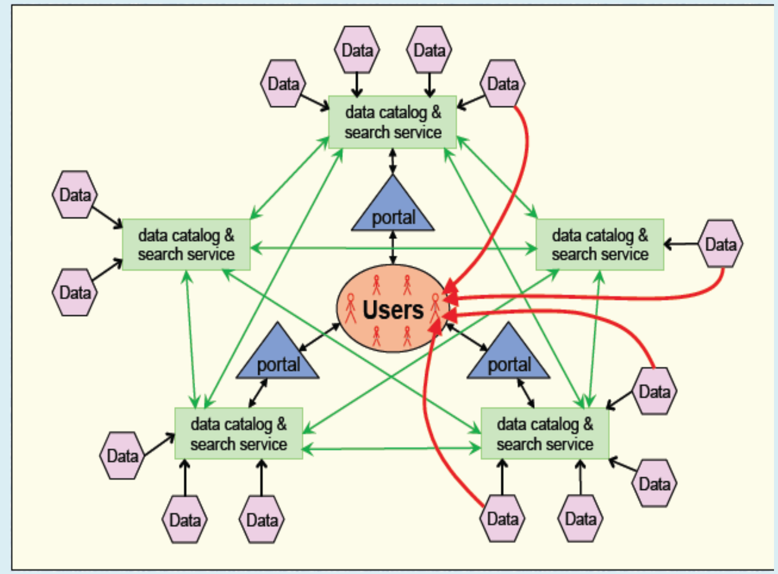

Data challenge :

exabyte even more challenging than exaflops

Different levels

I/O servers (XIOS, CDI-IO);
Compression; Storage;
Post-processing;
Distributed archive (ESGF)

XIOS

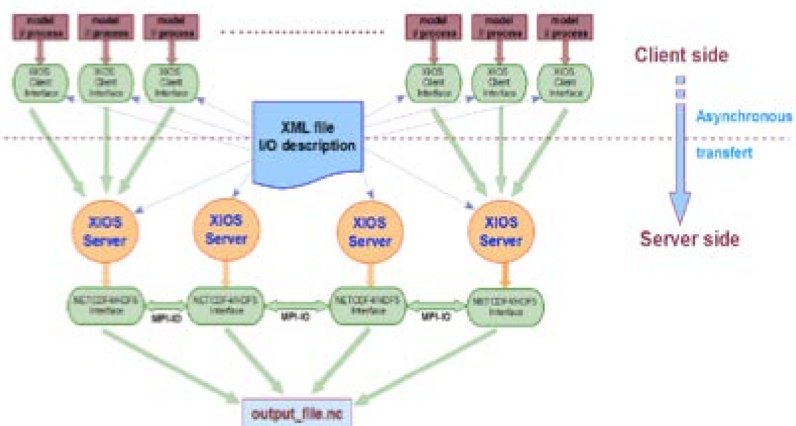

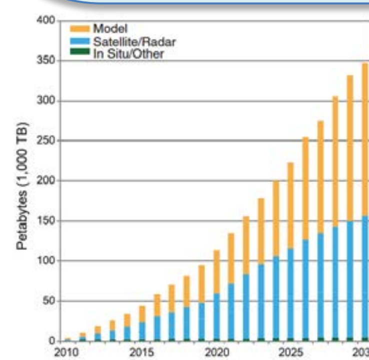

Overpeck et al. (Science 2011)

G8 ExArch project

Climate analytics on distributed exascale data archives

Europe

Strategic nature of HPC for Europe

Technology (ETP4HPC)/ HPC ecosystem (PRACE) / Applications (CoE)

Start of H2020 European Framework 2014-2020

Centers of excellence & Emerging technologies (exascale)

Climate

Launch of next international coordinated experiments CMIP6

National / Europe / International

Prepare next generation of climate models

Dynamical cores …

H2020 opportunities for climate on HPC issues :
Center of Excellence & Future and Emerging Technologies

Feb 27th 2014 meeting with DG Connect on climate and weather

https://ec.europa.eu/digital-agenda/news-redirect/14942

Towards a sustained European infrastructure

SeaWIFS Project (NASA/GSFC et Orbimage)