



FP7 Support Action - European Exascale Software Initiative

DG Information Society and the unit e-Infrastructures

EESI

EESI WP3 – Application Grand Challenges

Task 3.2: Working Group on WCES

Weather, Climatology and Earth Sciences

Chair: ENES-CMCC (Giovanni Aloisio)

Vice-Chair: INGV (Massimo Cocco)

Speaker: prof. Giovanni Aloisio
CMCC & University of Salento
Lecce - Italy

IESPI Meeting, San Francisco April 6, April 2011

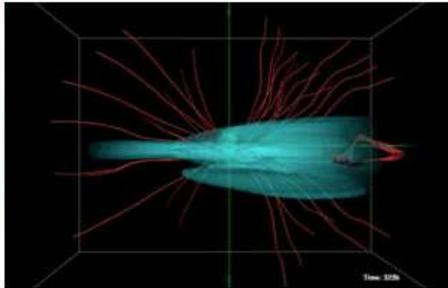




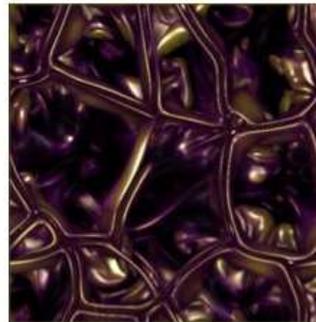
Geoscience really does need Petascale Computing (and Beyond)



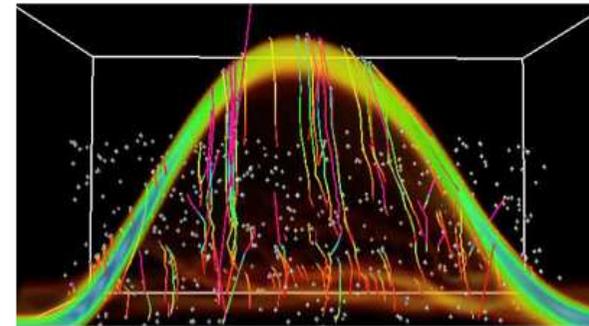
In a wide variety of domains...



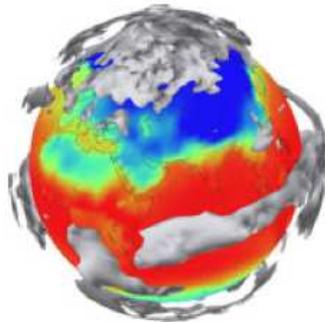
Space Weather



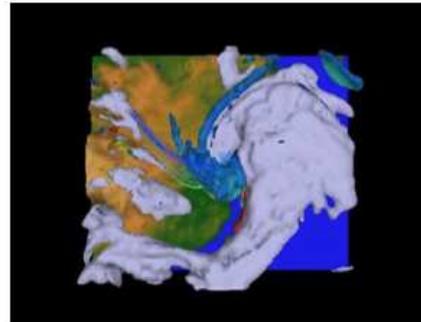
Turbulence



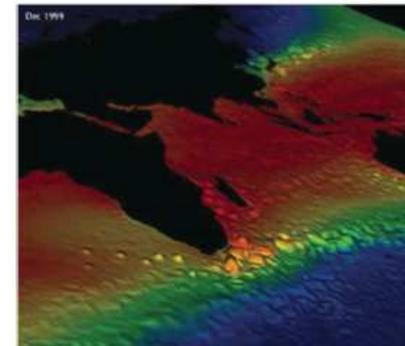
The Sun



Climate



Weather

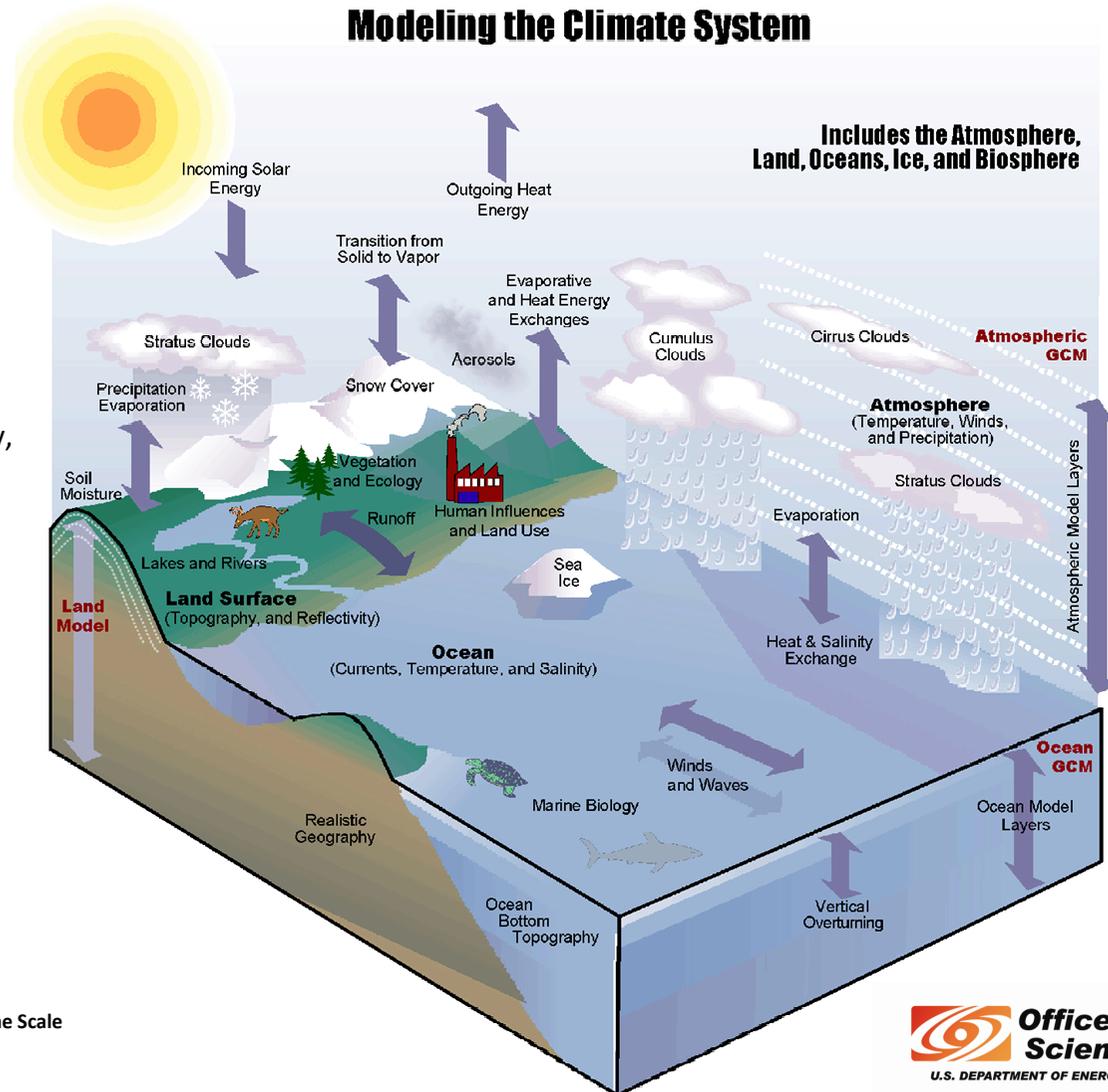


Oceanography

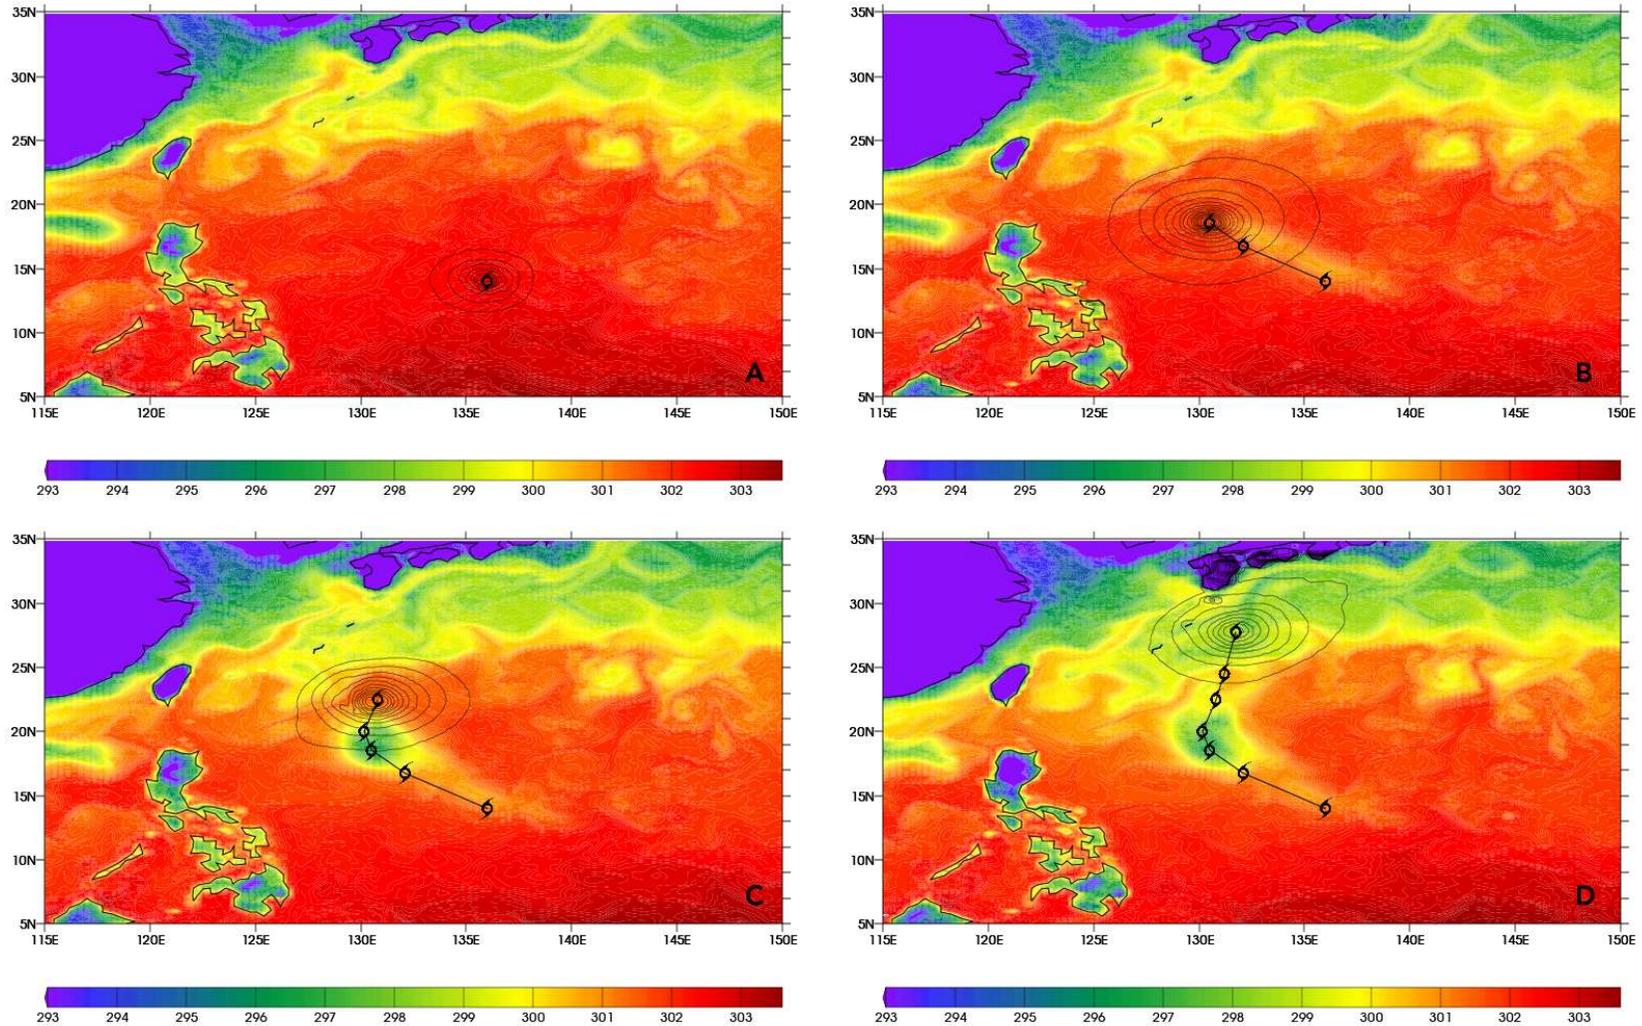
Climate Modeling

The big challenge is to model this complex system

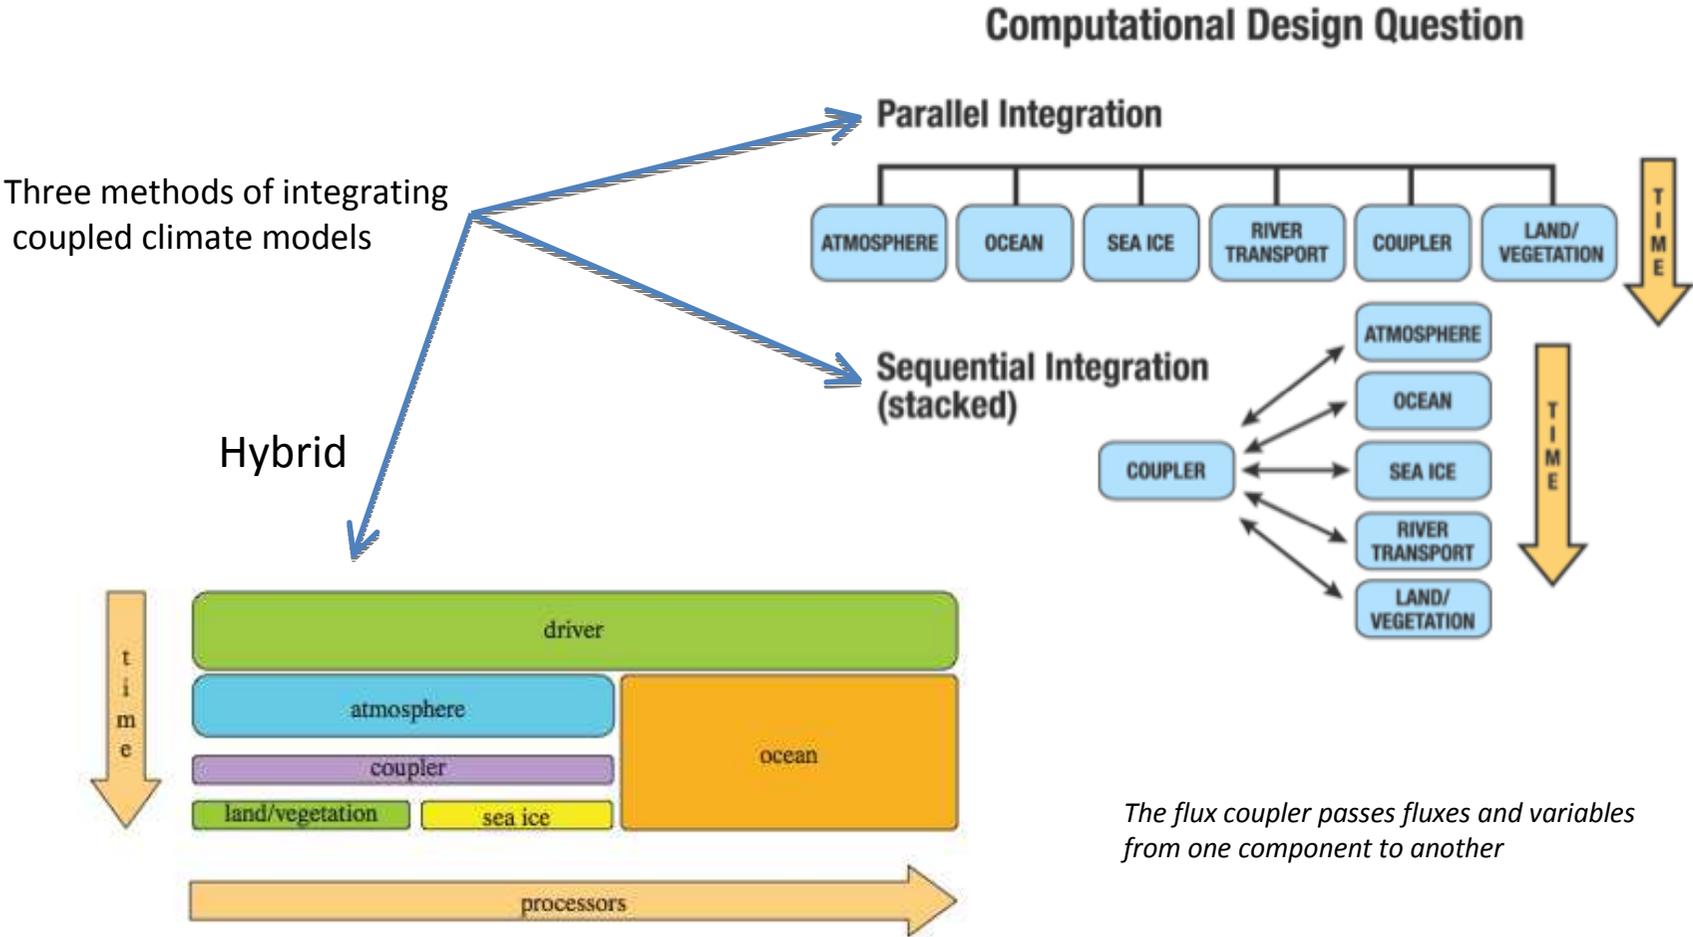
- Several complex process to be simulated
- Several interacting processes
- Great range of time scales to be analyzed
- Great range of spatial scales to be considered
- Need interdisciplinary sciences (physics, chemistry, biology, geology,...)
- Inherently non-linear governing equations
- Need sophisticated numerics
- Need huge computational resources



Modelers work on **Global Models of each climate component** to capture **small-scale features** (such as tropical cyclones) and **simulate how they interact with other components** (such as oceans)



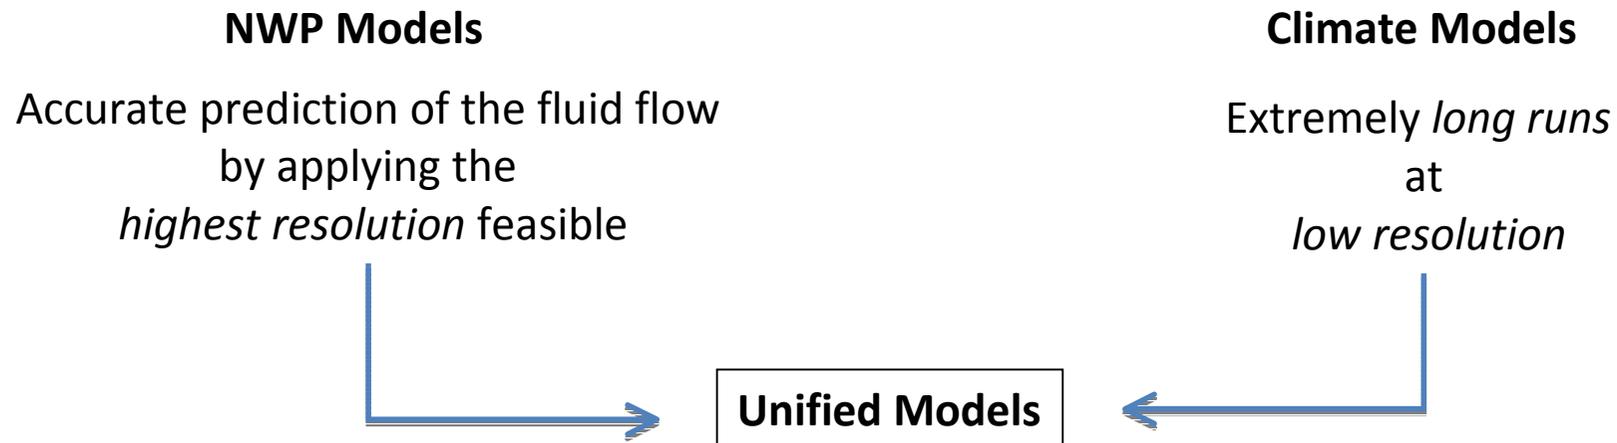
Coupling components in a climate model



Global Atmospheric Models

The first developed (1960's) and evolved for application to the problems of

- deterministic weather forecasting → **Numerical Weather Prediction (NWP)**
- seasonal prediction
- climate simulations



The need to combine global and regional climate models is increasing

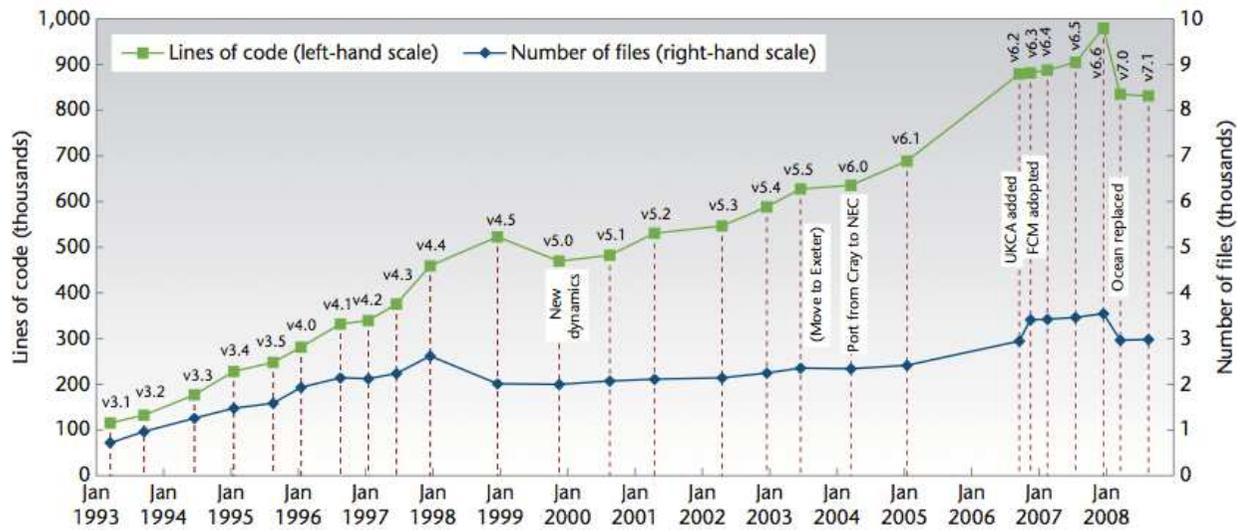


Figure 2. The Unified Model's growth over the past 15 years. Discontinuities in the growth curve indicate major component replacements, such as at versions 5.0 and 7.0, when the UM's dynamical core and ocean model were replaced, respectively. UM adopted the UK atmospheric chemistry model (UKCA) at version 6.2 and flexible configuration management (FCM) at version 6.3.

Growth over the past 15 years

The Met-Office Unified Model

Driver of change

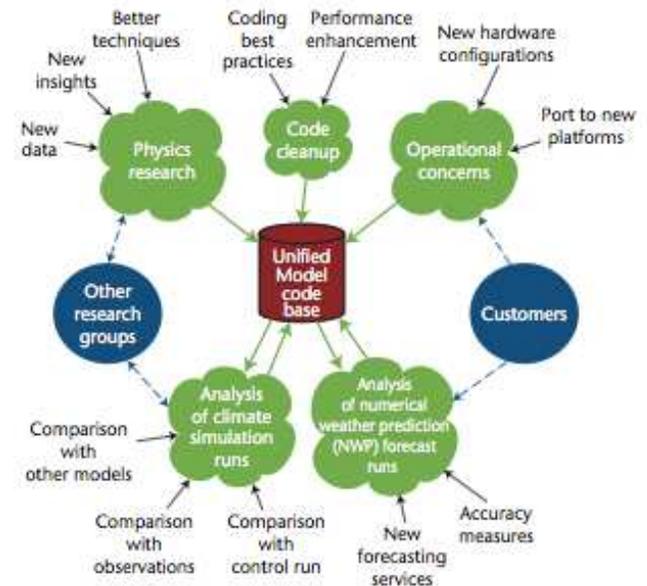


Figure 3. Drivers of change in UM code. The model's steady evolution is influenced by many factors, from advances in how scientists represent physical processes in the model to operational concerns, such as hardware upgrades. (Courtesy of Damian Wilson, UK Meteorological Office.)

1521-9615/09/\$26.00 © 2009 IEEE
 COPIED BY THE IEEE CS AND THE AIP

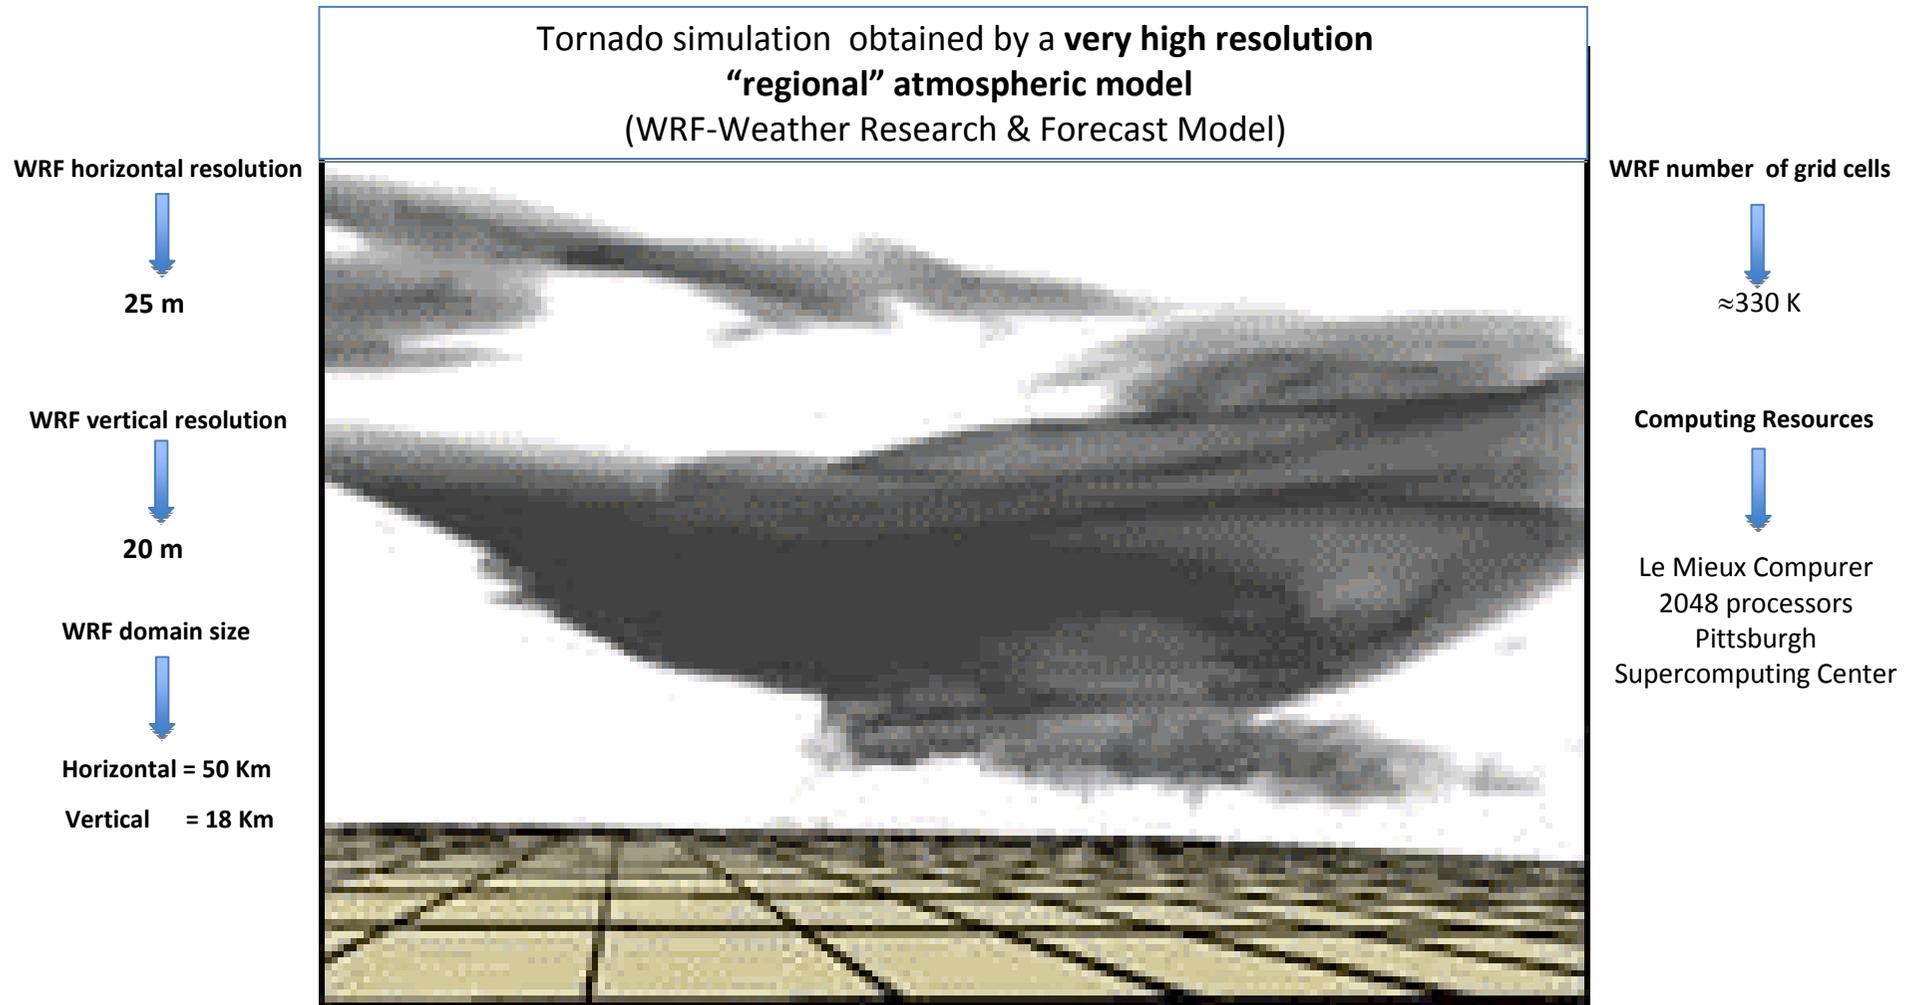
STEVE M. EASTERBROOK

University of Toronto

TIMOTHY C. JOHNS

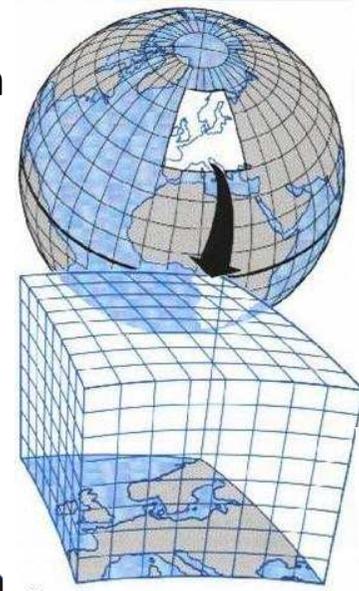
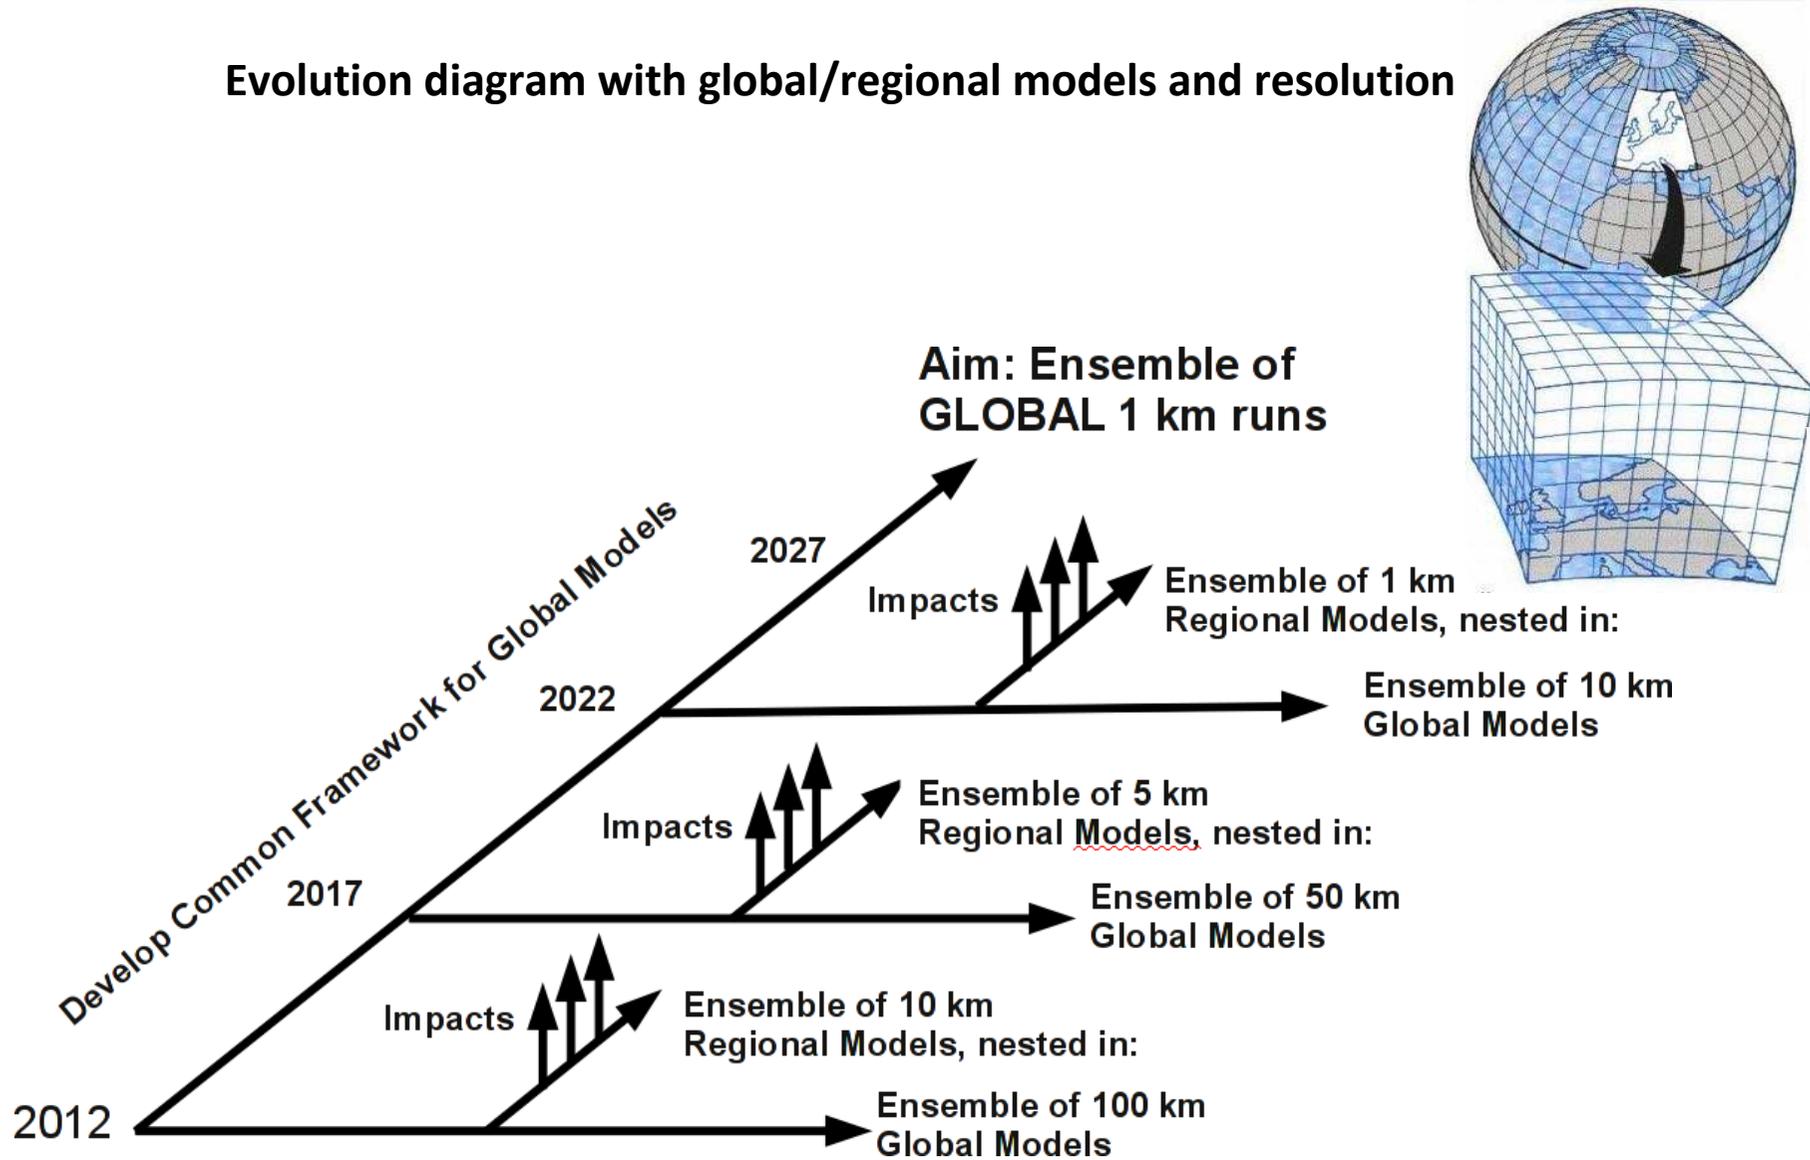
Hadley Centre for Climate Prediction and Research

The dream of future global climate models



The simulations was performed at the University of Pittsburgh (From Droegemeier and Xu of the University of Oklahoma’s Center for Analysis and Prediction of Storms (CAPS))

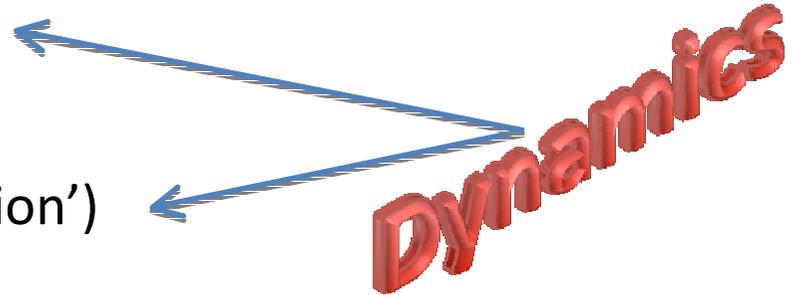
Evolution diagram with global/regional models and resolution



Ensemble runs are needed for quantifying uncertainty

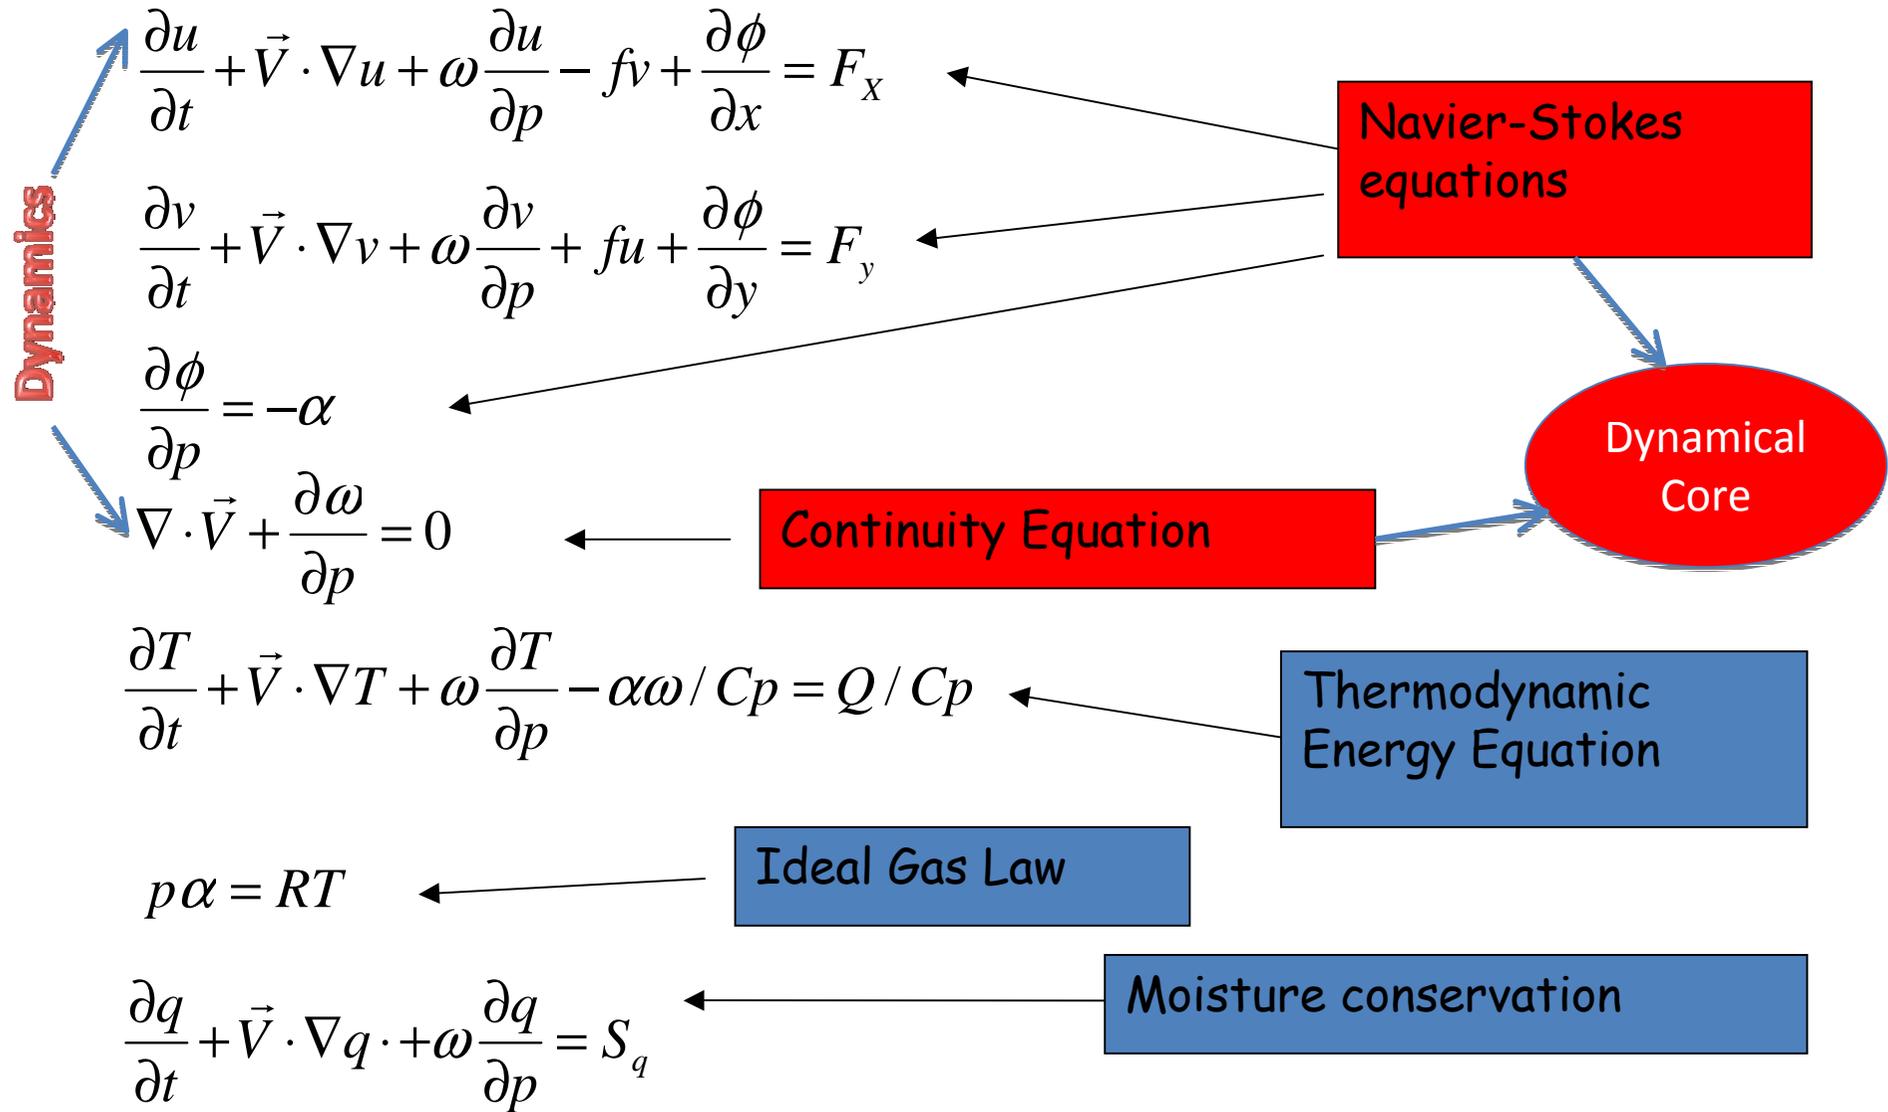
Climate - basic equations

- Momentum conservation ('Newton's Second Law' → Navier-Stokes equations)
- Mass conservation ('Continuity equation')
- The ideal gas law: The thermodynamic state of the atmosphere, at any point, is determined by its pressure, temperature and density
- Energy conservation ('The first law of thermodynamics')



(Conservation of water (vapour, liquid), salinity, ...)

The atmosphere in "Primitive Equations"



$u, v, \omega, T, \alpha, \Phi,$ and q

- In principle: possible to solve (#unknowns = #equations).
- In practice: analytical solutions not possible (e.g. non-linearity)
- (Various filtered forms of the equations of motion...)

A very complex natural system to be simulated

- Vast temporal (secs to B yrs) and spatial (1000s of km to microns) scales
- Highly nonlinear behavior

Newton's Law, applied to the atmosphere

$$\rho \mathbf{g} - \nabla p + \nu \nabla^2 \mathbf{u} = \rho \left(\frac{\partial}{\partial t} + \mathbf{u} \cdot \nabla \right) \mathbf{u}$$

Unpacks into..



10,000km

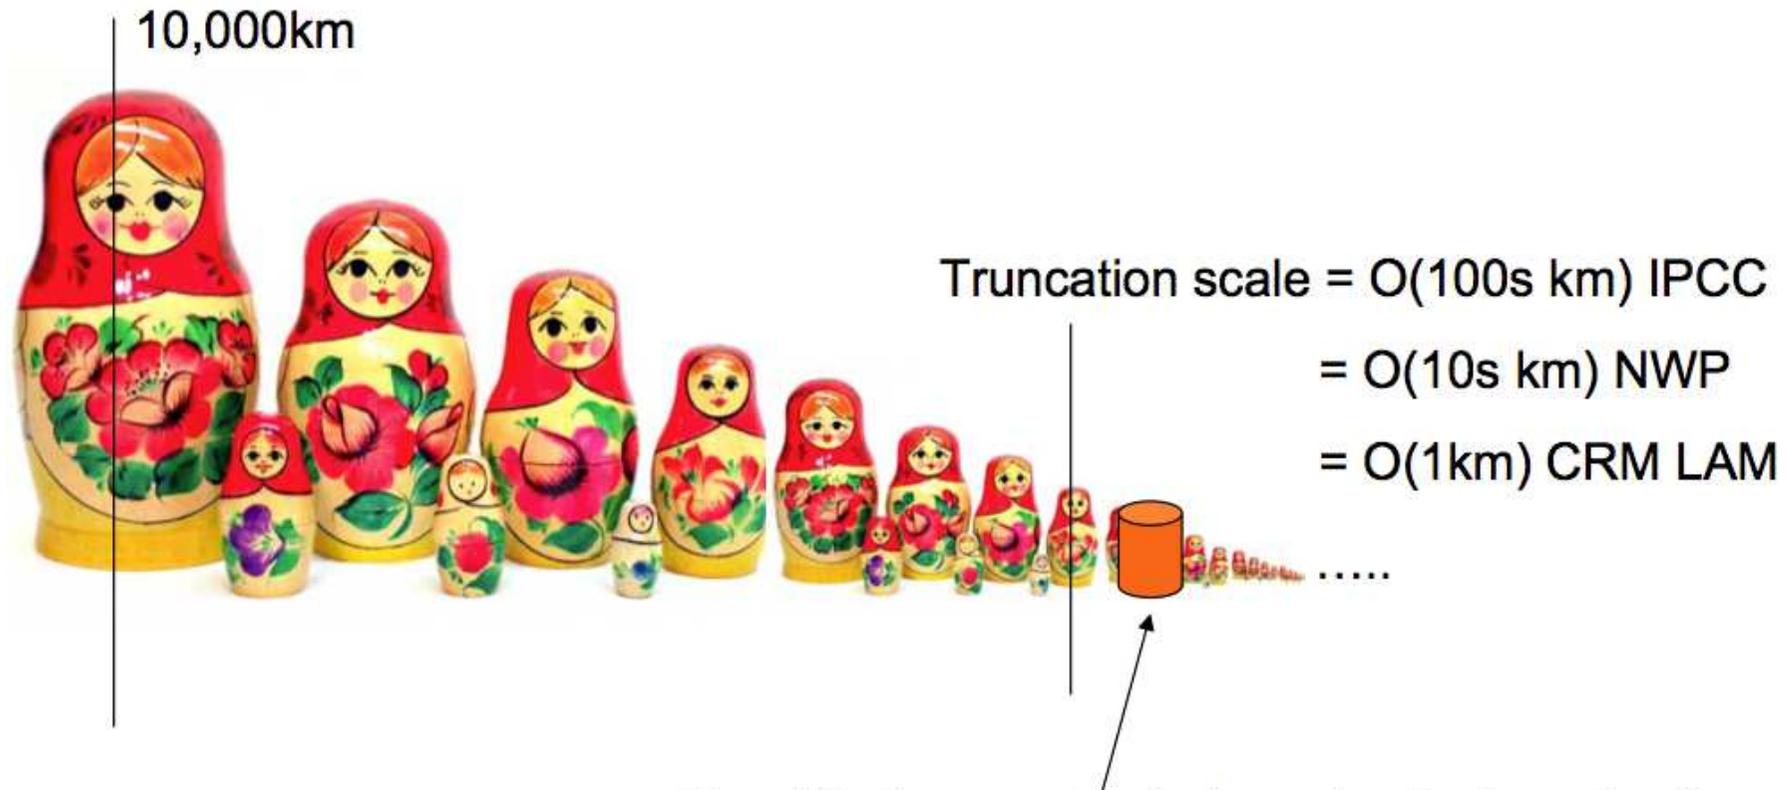
<1 mm

$$\rho \mathbf{g} - \nabla p + \nu \nabla^2 \mathbf{u} = \rho \left(\frac{\partial}{\partial t} + \mathbf{u} \cdot \nabla \right) \mathbf{u}$$

← **nonlinear**

Unpacks into billions of individual equations, describing scales of motion from planetary scales to microscopic scales.

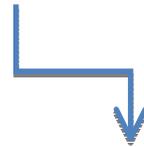
Even the world's biggest computer aren't big enough to represent all scales of motion in the atmosphere....



Simplified approximate formulae to describe the bulk effect of motions (eg clouds) that the computer model can't resolve.

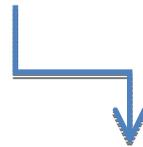
Dynamical cores issues at Exascale

- Dycores are the most challenging problem within geo-physical applications for future Exascale computing



In today's ESM, the dynamical cores are the main scalability bottlenecks

- The fundamental problem is the sequential nature of the time-stepping algorithm used to solve the PDEs involved



As the resolution of the models increases, the time-step used normally needs to shorten in order to control errors caused by the time-discretization and to avoid instabilities.

Dynamical cores issues at Exascale

- In the timeframe of Exascale computing, we will enter the cloud-resolving scales for global atmospheric models (~1km)



This will necessitate moving to *non-hydrostatic dynamical cores*

Non-hydrostatic atmospheric model - model core formulation

prognostic equations

$$\left\{ \begin{aligned} \frac{\partial \vec{v}}{\partial t} &= - \frac{\vec{\omega}_a}{\rho} \times \rho \vec{v} - \nabla(K + \Phi) - c_{pd} \theta_v \nabla \Pi \\ \frac{\partial \rho}{\partial t} &= - \nabla \cdot (\rho \vec{v}) \\ \frac{\partial \Pi}{\partial t} &= - \frac{R_d \Pi}{c_{vd} \rho \theta_v} \nabla \cdot (\theta_v (\rho \vec{v})) \\ \frac{\partial \rho \theta_v}{\partial t} &= - \nabla \cdot (\theta_v (\rho \vec{v})) \end{aligned} \right.$$

|·pV (to obtain energy equ.)

Π = Exner pressure
 θ_v = virtual pot. temperature
 ρ = density
 \vec{v} = 3D velocity vector
 K = spec. kinetic energy
 Φ = geopotential
 $\vec{\omega}_a$ = 3D abs. vorticity vector
 R_d = gas constant for dry air
 c_{vd} = spec. heat capacity at constant volume for dry air
 c_{pd} = spec. heat capacity at constant pressure for dry air

+ Transport equations for specific moisture quantities.

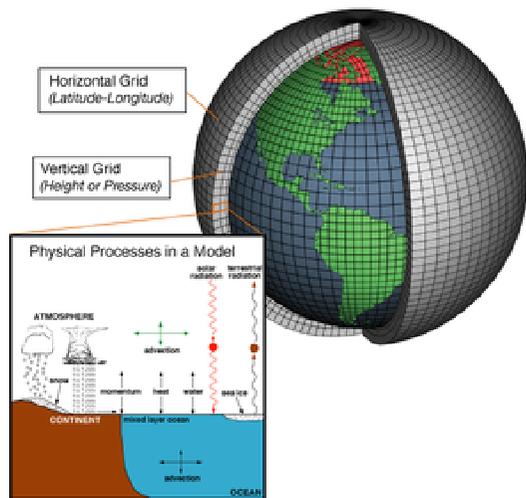
Hamiltonian description

- discretisation of Poisson backets
- symplectic time integration

$$\frac{\partial \mathcal{F}}{\partial t} = \{ \mathcal{F}, \mathcal{H} \}$$

Dynamical cores issues at Exascale

- The representation used for the physical quantities, like the *latitude longitude grid* seem to present special problems when going into the Exscale regime



The latitude longitude grid *suffers from the convergence of the longitudes when approaching the poles. More promising seem the various quasi-uniform grids* currently being developed to solve

The Pole problem !

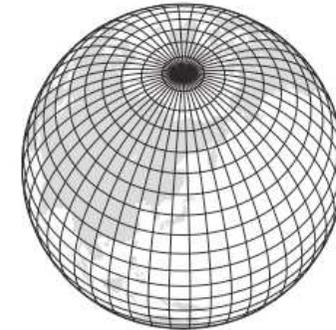
- Some currently used numerical scheme used to solve the Pole problem, like the spectral method (with its associated Gaussian grid)

requires *global communications* (within the transforms between the spectral representation and the Gaussian grid) and may *scale poorly at very large core counts*

The Pole Problem

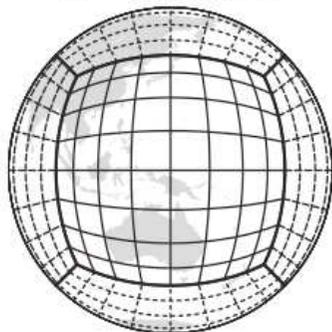
- The latitude longitude grid suffers from the convergence of the longitudes when approaching the poles

LATITUDE-LONGITUDE GRID

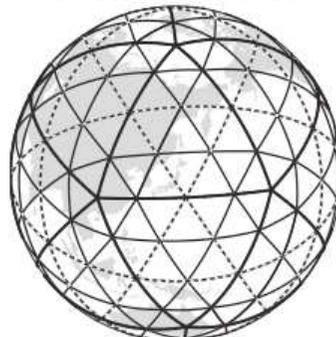


- Promising approaches are based on quasi-uniform grids (cubed sphere, icosahedral, Yin-Yang, Fibonacci ...)

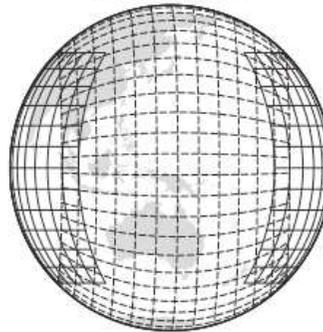
CUBED SPHERE GRID



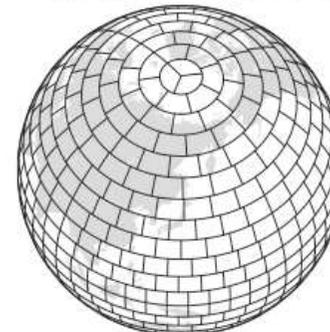
SPHERICAL GEODESIC
OR ICOSAHEDRAL GRID



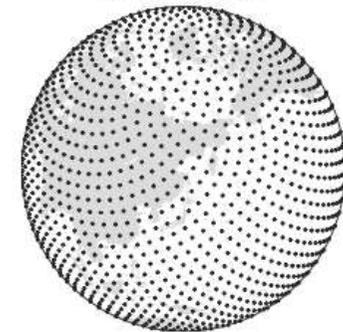
YIN-YANG GRID



KURIHARA OR REDUCED GRID



FIBONACCI GRID

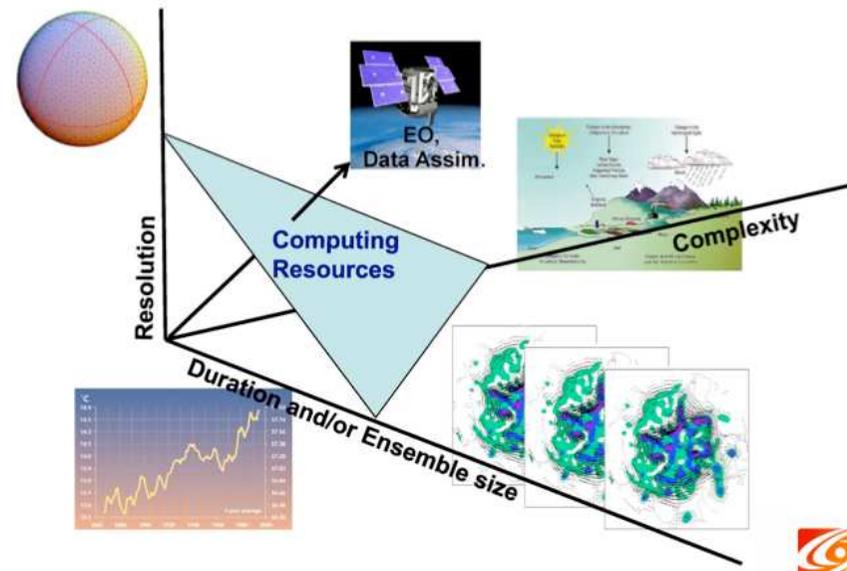


The 5 HPC dimensions of climate prediction

- Resolution
- Earth-System complexity
- Duration and/or Ensemble size
- Paleo-timescale integration
- Data assimilation and initial value forecasts (seamless prediction)

Capacity

Capability



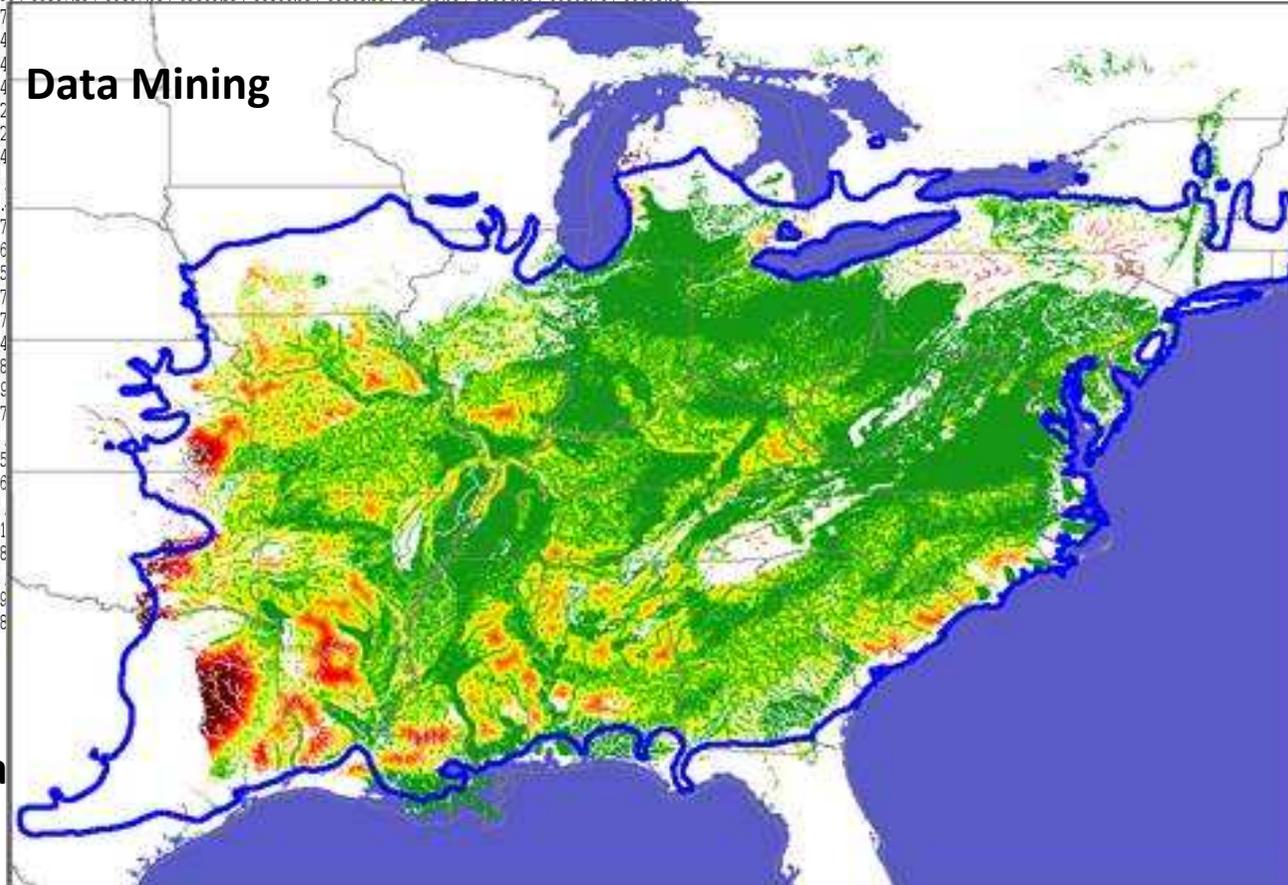
Data access, analysis and mining

```

{
  99859.23, 99859.92, 99860.53, 99861.2, 99861.8, 99862.34, 99862.91, 99863.48, 99864.02, 99864.53, 99864.97, 99865.48, 99865.91,
  99866.34, 99866.71, 99867.09, 99867.45, 99867.78, 99868.1, 99868.35, 99868.59, 99868.87, 99869.11, 99869.29, 99869.46, 99869.62,
  99869.8, 99869.85, 99869.89, 99869.95, 99869.91, 99869.98, 99869.93, 99869.8, 99869.73, 99869.66, 99869.5, 99869.32, 99869.15,
  99868.75, 99868.54, 99868.25, 99867.95, 99867.59, 99867.05, 99866.59, 99866.05, 99865.55, 99865.05, 99864.53, 99863.78, 99863.05,
  99862.25, 99861.48, 99860.63, 99859.7,
  99847.85, 99846.48, 99845.03, 99843.4,
  99824.9, 99822.84, 99820.67, 99818.54,
  99796.14, 99793.9, 99791.57, 99789.34,
  99768.8, 99766.97, 99765.27, 99763.52,
  99749.93, 99748.91, 99748.01, 99747.2,
  99742.5, 99742.36, 99742.24, 99742.24,
  99744.77, 99745.37, 99745.9, 99746.5,
  , 99754.82, 99755.52, 99756.4, 99757.,
  99767.55, 99768.63, 99769.68, 99770.7,
  99782.27, 99783.35, 99784.49, 99785.6,
  99797.68, 99798.9, 99800.11, 99801.25,
  99813.12, 99814.2, 99815.48, 99816.77,
  99828.0, 99829.09, 99830.17, 99831.27,
  99841.66, 99842.63, 99843.55, 99844.4,
  99853.27, 99854.1, 99854.88, 99855.56,
  {
    99824.98, 99826.53, 99828.03, 99
    99841.48, 99842.32, 99843.09, 99843.7,
    99847.5, 99847.83, 99847.82, 99847.9,
    99850.88, 99851.41, 99852.0, 99852.75,
    99863.16, 99864.35, 99865.62, 99866.6,
    99871.2, 99870.59, 99869.62, 99868.4,
    99837.77, 99832.82, 99827.66, 99822.1,
    99756.61, 99749.63, 99742.68, 99735.6,
    99676.4, 99671.84, 99667.6, 99663.62,
    99641.27, 99640.58, 99640.18, 99639.9,
    99647.72, 99649.16, 99650.7, 99652.26
  }
}
    
```

Data access

Data Mining



Data

Average	Variance
7.27002589	118.20635699
7.20344162	118.71782431
7.14841461	118.93681027
7.08705139	119.32944630
7.03001658	119.66761535
6.96573385	120.08689023
6.90545019	120.45400045
6.84642792	120.79659695
6.78913752	121.12353909
6.72494253	121.61722149
6.66145833	122.09671745
6.59991582	122.48782104
6.54492569	122.75315922
6.48511887	123.20487250
6.42746862	123.56921073
6.37406540	123.81541520
6.31255722	124.28107138
-88.928	23.9063
-88.928	25.3125
-88.928	26.7188



FP7 Support Action - European Exascale Software Initiative

DG Information Society and the unit e-Infrastructures

EESI

EESI - WP 3 - Application Grand Challenges Task 3.2: Working Group on WCES Weather, Climatology and Earth Sciences

WG 3.2 Roadmap



Name	Organization	Email	Country	Area of Expertise
Giovanni Aloisio	ENES-CMCC	giovanni.aloisio@unisalento.it	IT	Exascale Computing
Massimo Cocco	INGV	massimo.cocco@ingv.it	IT	Seismology
Nadia Pinardi	CMCC	n.pinardi@sincem.unibo.it	IT	Oceanography
Sylvie Joussaume	IPSL	sylvie.joussaume@isce.ipsl.fr	FR	Earth-system Modeling
Sophie Valcke	CERFACS	valcke@cerfacs.fr	FR	Coupled Climate Models
Jean-Pierre Vilotte	IPG	vilotte@ipgp.jussieu.fr	FR	Solid Earth, Seismology
Marie-Alice Foujols	IPSL	Marie-Alice.Foujols@ipsl.jussieu.fr	FR	Climate Modeling
Bryan Lawrence	BADC	bryan.lawrence@stfc.ac.uk	UK	Climate Data Management
Graham Riley	Manchester Univ.	graham.riley@manchester.ac.uk	UK	High Performance Computing
Mats Hamrud	ECMWF	nar@ecmwf.int	UK	Data Assimilation
Tim Palmer	ECMWF	Tim.Palmer@ecmwf.int	UK	Weather/Climate Modeling
Reinhard Budich	MPI	reinhard.budich@zmaw.de	DE	Earth System Modeling
Joachim Biercamp	DKRZ	biercamp@dkrz.de	DE	High Performance Computing
Heiner Igel	LMU	heiner.igel@geophysik.uni-muenche	DE	Solid Earth, Geophysics
Colin Jones	SMHI	colin.jones@smhi.se	SE	Climate Change Modeling
Johan Silen	FMI	johan.silen@fmi.fi	FI	Geophysics & Computing
Jose Baldasano	BSC	jose.baldasano@bsc.es	ES	Earth Sciences

WG 3.2 Roadmap - Table of contents (I)

Table of contents

1 EXECUTIVE SUMMARY (G. ALOISIO, M. COCCO)	4
2 GOALS OF THE WCES WORKING GROUP (G.ALOISIO, M.COCCO)	5
3 WCES SCIENTIFIC ROADMAP	6
3.1 EARTH SYSTEM MODELS	6
3.1.1 <i>Introduction (Chair, co-chair and Editorial leaders) (half a page)</i>	6
3.1.2 <i>Scientific and technical vision</i>	6
3.1.2.1 <i>Coupled Earth System Models optimization (J.Biercamp) (6 pages max)</i>	6
3.1.2.1.1 <i>Dynamical cores (M.Hamrud) (2 pages max)</i>	7
3.1.2.1.2 <i>Flexible couplers and coupling frameworks (S. Valcke) (2 pages max)</i>	7
3.1.2.3 <i>Multidisciplinary data assimilation (M. Hamrud) (tbd) (2 pages max)</i>	8
3.1.2.4 <i>Uncertainty and predictability (C. Jones) (2 pages max)</i>	8
3.1.2.5 <i>Evaluation of predictive models (J. Silen) (2 pages max)</i>	9
3.1.3 <i>Conclusions (Chair, co-chair and Editorial leaders) (half a page)</i>	9

WG 3.2 Roadmap - Table of contents (II)

3.2	DATA INTENSIVE APPLICATIONS IN EARTH SCIENCE	10
3.2.1	<i>Introduction (Chair, co-chair and Editorial leaders) (half a page)</i>	<i>10</i>
3.2.2	<i>Scientific and technical vision</i>	<i>10</i>
3.2.2.1	Scientific Data Management (B. Lawrence, M. Juckes) (6 pages max).....	10
3.2.2.3.1	Data Discovery (B. Lawrence) (2 pages max)	10
3.2.2.3.2	Input/Output (M. Juckes) (2 pages max).....	11
3.2.2.2	Data analysis and Visualization (H.Higel, J. Biercamp) (2 pages max).....	12
3.2.3	<i>Conclusions (Chair, co-chair and Editorial leaders) (half a page)</i>	<i>12</i>
3.3	CROSS-CUTTING DIMENSIONS	13
3.3.1	<i>Introduction (Chair, co-chair and Editorial leaders) (half a page)</i>	<i>13</i>
3.3.2	<i>Scientific and technical vision</i>	<i>13</i>
3.3.2.1	Performance (M-A. Foujols) (2 pages max)	13
3.3.2.2	Programmability (tbd) (2 pages max)	14
3.3.2.3	Resilience (tbd) (2 pages max)	14
3.3.2.4	Workflows (R. Budich) (2 pages max)	15
3.3.3	<i>Conclusions (Chair, co-chair and Editorial leaders) (half a page)</i>	<i>15</i>

WG 3.2 Roadmap - Table of contents (III)

3.4	WCES IMPACTS AND FINDINGS AT EUROPEAN LEVEL	15
3.4.1	<i>Societal, environmental and economical impact (all) (1 page max)</i>	15
3.4.2	<i>Strengths and weaknesses (all) (1 page max)</i>	15
3.4.3	<i>Sources of competitiveness for Europe (all) (1 page max)</i>	16
3.4.4	<i>Needs of education and training (all) (1 page max)</i>	16
3.4.5	<i>Potential collaborations outside Europe (all) (1 page max)</i>	16
3.4.6	<i>Agenda, need of HR, provisional costs (all) (1 page max)</i>	16
3.4.7	<i>Communication plans, dissemination actions, identification of potential target public (all) (1 page max)</i>	16
3.4.8	<i>European stakeholders and exascale prototype (all) (1 page max)</i>	17
4	CONCLUSIONS (G. ALOISIO, M. COCCO)	18

Critical Exascale Software for WCES

<i>X-Stack Components</i>	<i>Needed Capabilities</i>	<i>Exascale Uniqueness</i>	<i>Exascale Criticality</i>
Frameworks	Resilience API and utilities	3	C
	Multi-institutional/multiproject collaboration plan	2	U
	Tool chain development/selection	2	U
	Programming model evaluation/adoption	2	C
	Data placement	2	C
	Multicomponent simulation utilities	2	U
	Access to third-party libraries	1	C
Numerical Libraries	→ Fault-oblivious, error-tolerant software	3	C
	→ Asynchronous methods	2	C
	→ Overlap data and computation	3	U
	Self-adapting hybrid and hierarchical based algorithms	1	C
	→ Hybrid and hierarchical-based algorithms (e.g., linear algebra split across multicore and GPU)	1	U
	→ Algorithms that minimize communications	3	C
	Architecture-aware algorithms/libraries	3	C
	Autotuning-based software	1	U
	Standardization activities	1	U
	Energy-efficient algorithms	2	U
	Mixed arithmetic	1	U

Critical Exascale Software for WCES

<i>X-Stack Components</i>	<i>Needed Capabilities</i>	<i>Exascale Uniqueness</i>	<i>Exascale Criticality</i>
Algorithms	→ Scalability	2	N
	→ Fault tolerance/resilience	1	N
	→ Conforming to architectural requirements	3	N
	→ New areas/uses of algorithms	1	U
Debugging	→ Concurrency and architecture driven high frequency of errors/failures	3	C
	→ Scalability of debugger methodologies (data volumes and interfaces)	3	C
	→ Focus on multilevel debugging, communicating details of faults between software layers	3	U
	→ Synthesis of fault information into understanding in the context of apps and architecture	3	C
	→ Specialized lightweight operating systems	2	N
	→ Automatic triggers, need compile time bridge to debugging that removes need to rerun	2	N
	→ Scalable clustering of apps process states and contexts, filter/search within debugger	2	N
	→ Vertical integration of debug and per information across software layers	2	N
	→ Excision of buggy code snippets to run at lower concurrencies	1	N
	→ Heterogeneity	1	N

Critical Exascale Software WCES

<i>X-Stack Components</i>	<i>Needed Capabilities</i>	<i>Exascale Uniqueness</i>	<i>Exascale Criticality</i>
I/O	Customization with I/O, purpose-driven I/O	3	C
	 New I/O models, SW, runtime systems and libs	3	C
	 Intelligent/proactive caching mechanisms for I/O	3	N
	 Fault-tolerant mechanisms	3	C
	I/O into programming models and languages	3	N
	Balanced architectures with newer devices	2	N
	File systems or alternative mechanisms	2	N
	 Active storage	2	N
	 Wide-area I/O and integration of external storage systems	2	N
	Special-purpose network protocols for parallelism	2	N
	Balanced architectures with newer devices embedded with the node	1	N
	Scientific data management	 Scalable data analysis and mining SW and tools	3
 Scalable data format and high-level libraries		3	C
Scientific workflows tools		2	C
 Search and query tools		2	N
 Wide-area data access movement and query tools		2	N
 Scientific databases		2	N

Critical Exascale Software WCES

<i>X-Stack Components</i>	<i>Needed Capabilities</i>	<i>Exascale Uniqueness</i>	<i>Exascale Criticality</i>
Programming models	→ Exascale programming model	3	C
	Scalable, fault-tolerant MPI	3	C
	Applications development tools	3	N
	→ Heterogeneous node programming model	2	C
	Domain-specific programming models	2	N
	Language features for massively parallel I/O	2	U
	Language support for adaptive computation	2	N
	Interoperability between models	1	2N
Compilers	Implement exascale languages	3	C
	Support for resilience	3	C
	→ Implement heterogeneous programming models	2	C
	→ Support for massive I/O	2	C
	New optimization frameworks	2	N
	Interactions between compilers and tools, runtime	2	C
	Dynamic compilation, feedback optimization	2	N
	→ Autotuning-based software	2	N
	Enhancements to existing languages/APIs	1	N
	→ Automatic parallelization	1	N

Critical Exascale Software WCES

<i>X-Stack Components</i>	<i>Needed Capabilities</i>	<i>Exascale Uniqueness</i>	<i>Exascale Criticality</i>
Operating Systems	Define the base OS (Standard API)	3	C
	APIs for resilience (access to RAS, etc)	3	C
	Collective OS operations	3	N
	Scalable system simulation environment	2	C
	Improved APIs for scalable performance monitoring and debugging	2	C
	New APIs for energy management	2	U
	Improved APIs for explicit memory management	1	C
	Improved APIs for threading	1	U
Performance	→ Extremely scalable performance methods and tools	3	C
	→ Performance measurement and modeling in presence of noise/faults	3	C
	→ Automated/automatic diagnosis and autotuning	2	N
	Predictive future large-scale system design	2	C
	Vertical integration across SW layers	2	N
	→ Performance-aware design and implementation	2	U
	Performance optimization for other metrics than time	2	U
	→ Support for heterogeneous hardware and hybrid programming models	1	C

<i>X-Stack Components</i>	<i>Needed Capabilities</i>	<i>Exascale Uniqueness</i>	<i>Exascale Criticality</i>
Programmability	New models of computation	3	C
	New runtime/OS interface for environment aware programming	2	C
	Programmability to decouple exascale system issues from applications programming	1	C
Power	→ Power performance monitoring and aggregation that scales to 1 billion core system	3	C
	→ Power control system	3	C
	Scalable control algorithms to bridge gap between global and local power models	2	C
	→ Power-aware and scalable resource control and scheduling	2	C
	→ Optimally tuned system power based on control loop	1	N
	Power instrumentation and control standardization	1	N
Runtime Systems	→ Load balance	3	C
	→ Asynchrony, overlap	2	C
	Hierarchical execution models and scheduling	3	N
	→ Scaling/optimization of communications	3	C
	→ Memory management and locality scheduling	2	C
	Heterogeneity: scheduling	2	U
	Fine-grained mechanisms at node level	1	U

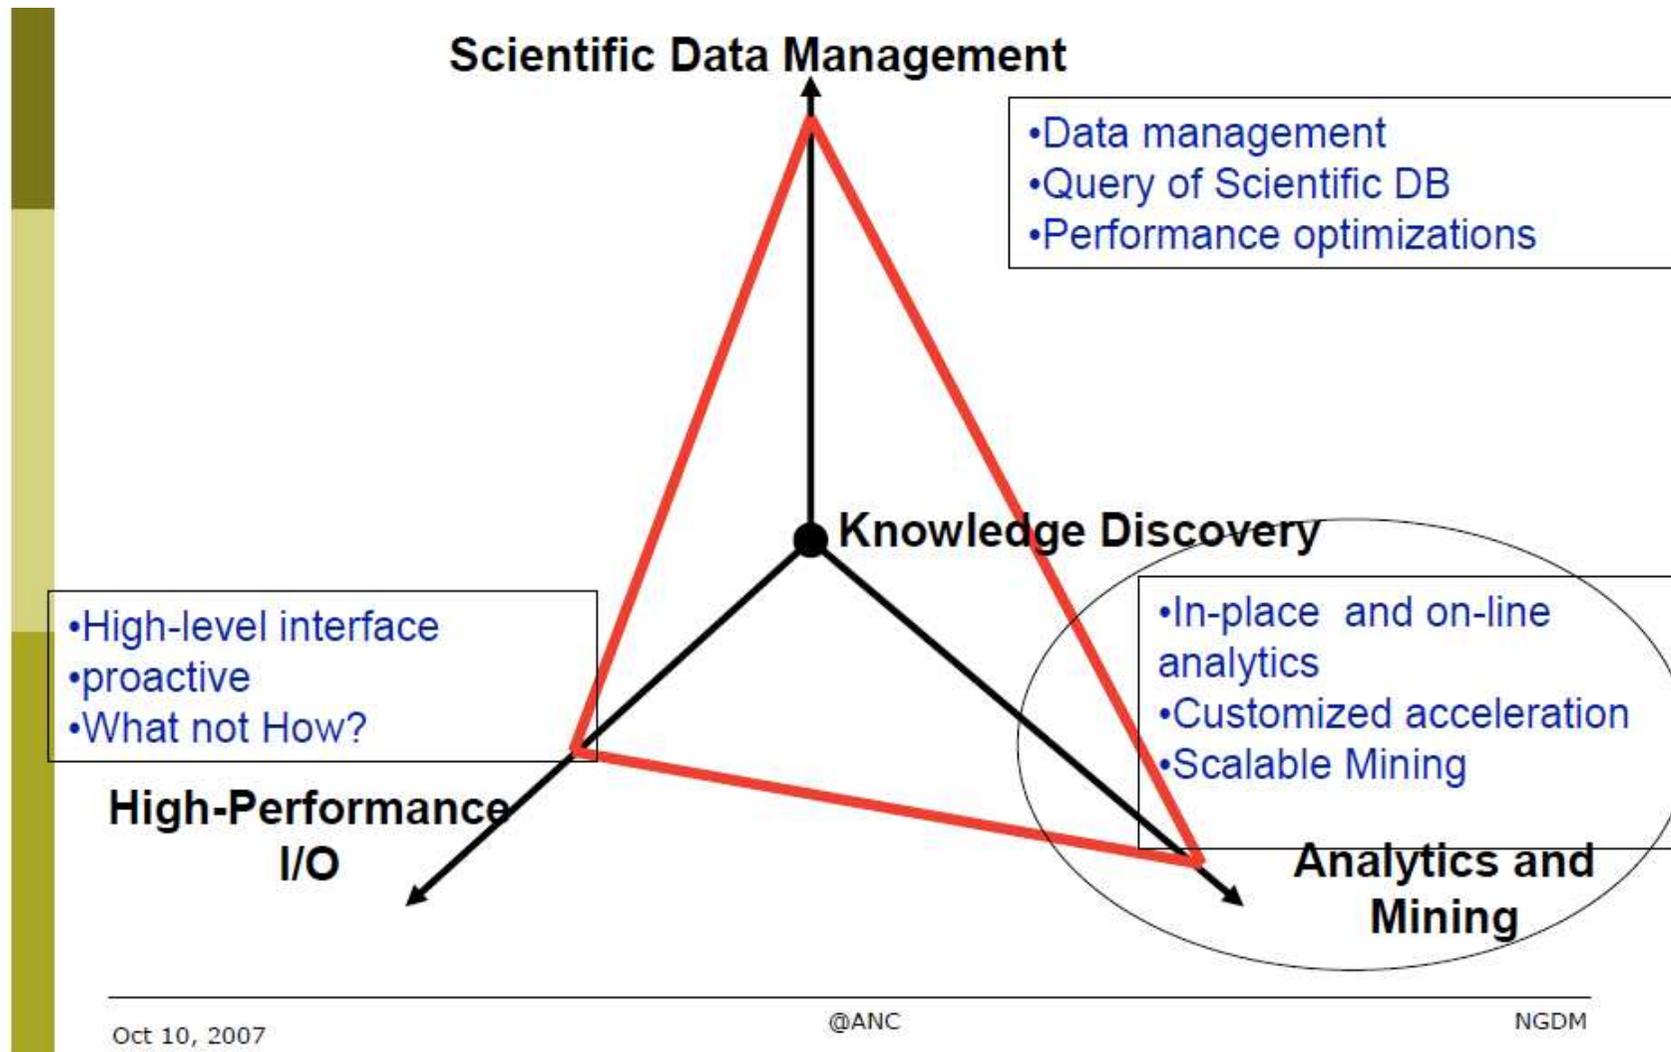
Critical Exascale Software

<i>X-Stack Components</i>	<i>Needed Capabilities</i>	<i>Exascale Uniqueness</i>	<i>Exascale Criticality</i>
Resilience	Performance measurement and modeling in presence faults	3	C
	Better fault tolerance protocols	2	C
	Fault isolation/confinement	2	C
	NV-RAM for local state, cache of file system	2	C
	Replication (TMR, backup core)	2	U
	Proactive actions (migration)	2	U
	Domain-specific API and utilities for frameworks	2	C
	Applications-guided fault management	2	C
	Language/compiler/runtime support for resilience (fault-aware programming, API from OS, RAS)	2	C
	Fault-tolerant MPI	2	C
	→ Fault-oblivious, error-tolerant numerical libraries	2	C
	Resilient applications and algorithms	1	N
	→ Fault-oblivious system software	2	C
	Fault-aware system software and API for resilience	2	C
	→ Prediction for time optimal checkpoint/migration	2	U
	Fault models, event log standardization root cause analysis	2	C
	→ Resilient I/O, storage, and file systems	2	C
Situational awareness	2	C	
Experimental environment	2	C	
Fault isolation/confinement + local management	2	C	

Critical Exascale Software for WCES

X-Stack Components	Needed Capabilities
Numerical libraries	Fault-oblivious, error-tolerant software, Asynchronous methods, Overlap data and computation, Hybrid and hierarchical based algorithms, Algorithms that minimize communications
Algorithms	Scalability, Fault-tolerance/resilience, New areas/uses of algorithms
Debugging	Scalability of debugger methodologies, Focus on multilevel debugging, Synthesis of fault information into understanding in the context of apps, Heterogeneity
I/O	New I/O models, SW, run time systems and libs, Intelligent/proactive caching mechanism for I/O, Fault-tolerant mechanisms, Active storage, Wide-area I/O and integration of external storage systems
Scientific data management	Scalable data analysis and mining SW and tools, Scalable data format and high-level libraries, Search and query tools, Wide-area data access movement and query tools, Scientific databases
Programming models	Exascale programming models, Heterogeneous node programming model
Compilers	Implement heterogeneous programming models, Support for massive I/O, Auto-tuning based software, Automatic parallelization
Performance	Extremely scalable performance methods and tools, Performance measurement and modeling in presence of noise/faults, Automatic diagnosis and auto-tuning, Performance-aware design and implementation, Support for heterogeneous hardware and hybrid programming models
Power	Power performance monitoring and aggregation, Power control system, Power-aware and scalable resource control and scheduling, Optimally tuned system power based on control loop
Runtime system	Load balance, Asynchrony, overlap, Optimization of communications, Memory management and locality scheduling
Resilience	Fault-oblivious system software, Prediction for time optimal checkpoint, Resilient I/O, storage and file systems

Challenges in Scientific Knowledge discovery



Knowledge discovery from climate change data

Main challenge: large scale, scalable and efficient scientific knowledge discovery from data

Climate data: multi dimensional, multi model, multi resolution,...

Mining ensembles (new dimension!) to “face” uncertainty

Multiple challenges at different levels:

- Efficient data indexing
- Multi-layer parallel I/O design and implementation
- MPI-based analytics functions/libraries
- MPI-based mining functions/libraries
- Active storage primitives (move analytics and mining functions to the storage nodes) & parallel file systems

CMCC Research Activity on Scientific Data Management and possible contribution to the X-Stack software

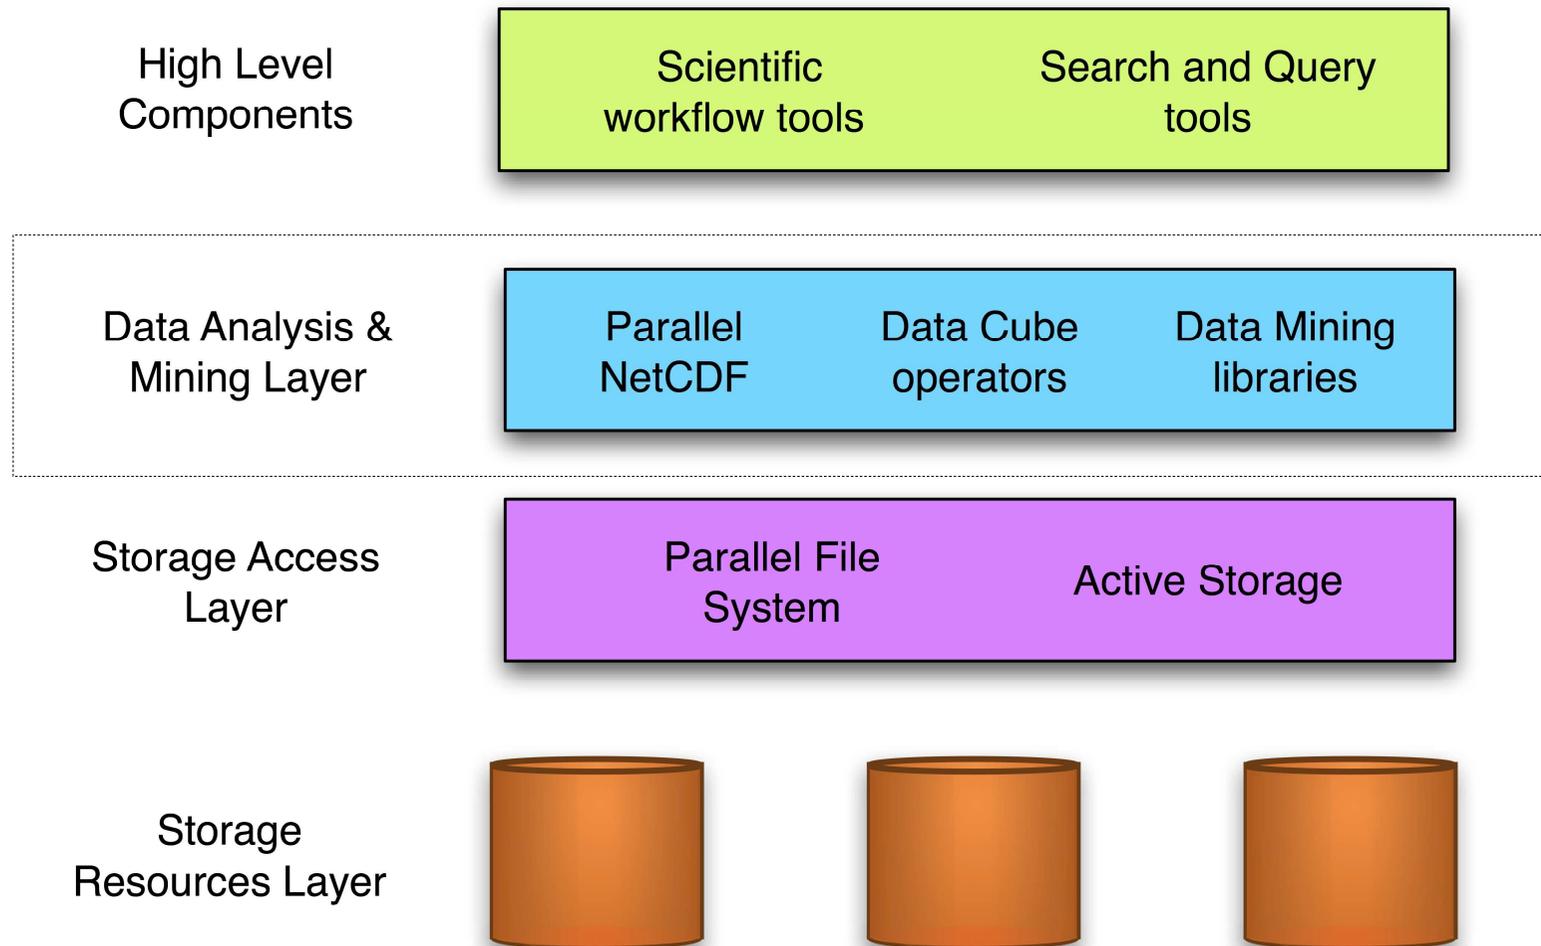
Strong interest in Scientific data management activities:

Scientific data management	Scalable data analysis and mining SW and tools
	Scalable data format and high-level libraries
	Scientific workflows tools
	Search and query tools
	Wide-area data access movement and query tools
	Scientific databases

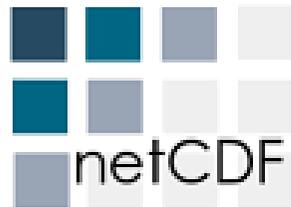
and some I/O related topics:

I/O	Active storage
	Wide-area I/O and integration of external storage systems

CMCC Research Activity on Scientific Data Management and possible contribution to the X-Stack software



From NetCDF data to data cube virtualization



Data Formats

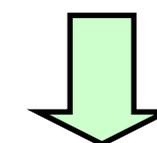
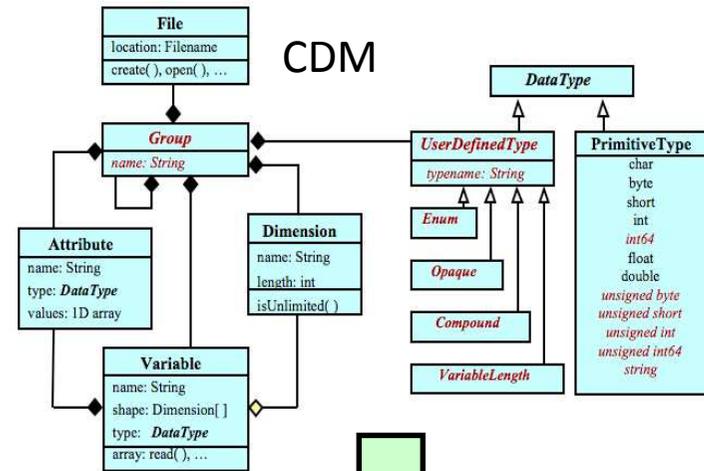
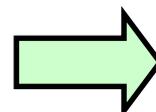
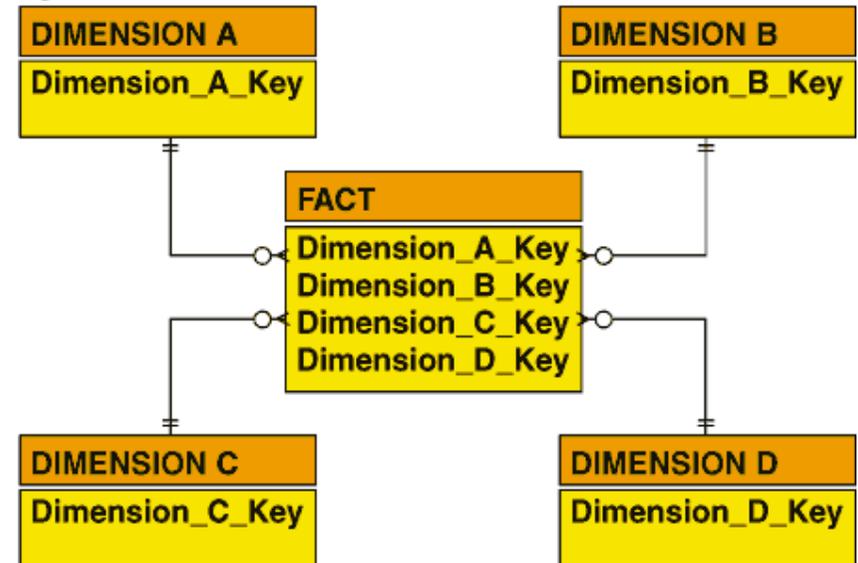
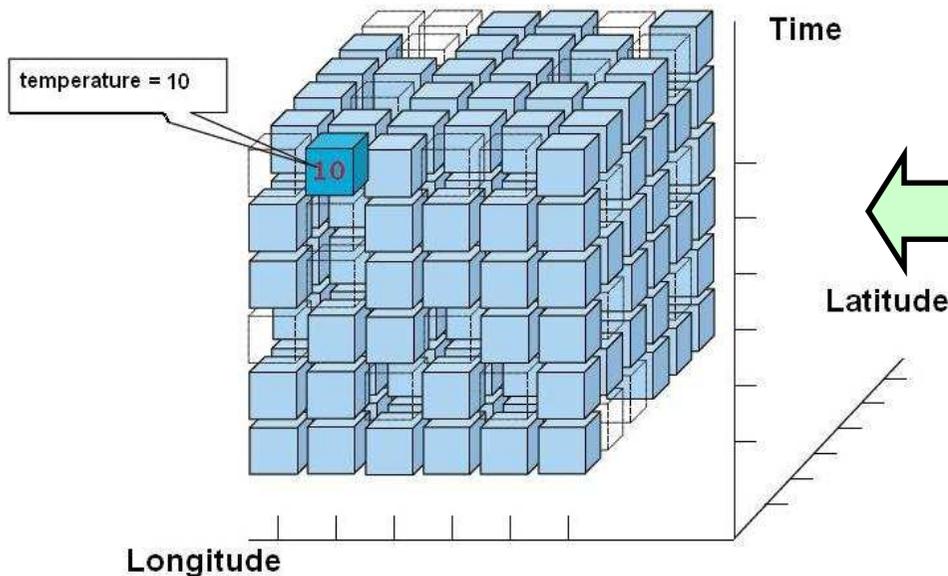


Figure 1

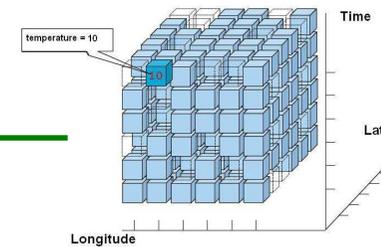
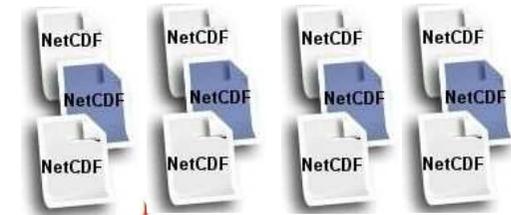
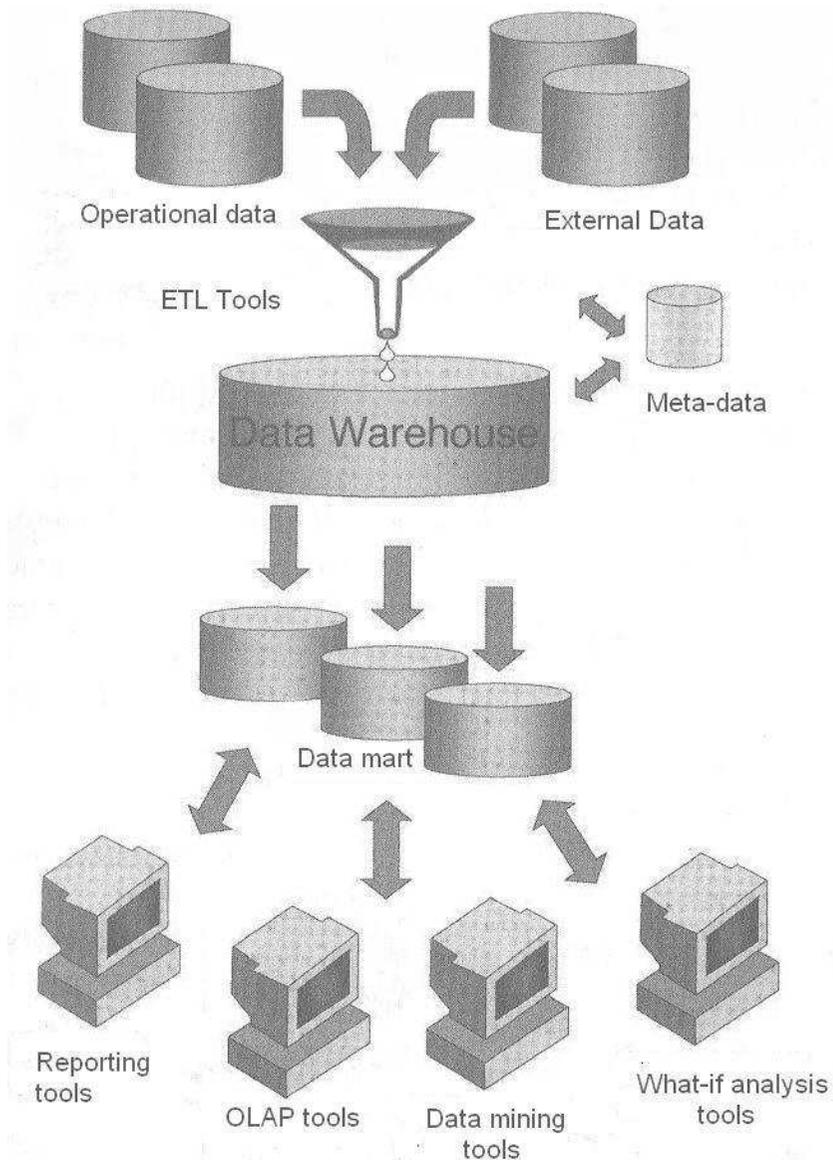


Star Schema



Data cube implementation

Extraction, transformation & Loading



Thanks