A partnership in weather and climate research

Joint Weather and Climate Research Programme

Weather and Climate Modelling: ready for exascale?

Pier Luigi Vidale

Willis Professor of Climate System Science and Climate Hazards

Marie-Estelle Demory, Reinhard Schiemann,
Jane Strachan, Bryan Lawrence, Hilary Weller

University of Reading

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

Malcolm Roberts

Matthew Mizielinski, Lizzie Kendon, + Met Office

(with thanks to the many MO groups

Involved in model development and elsewhere)

Met Office

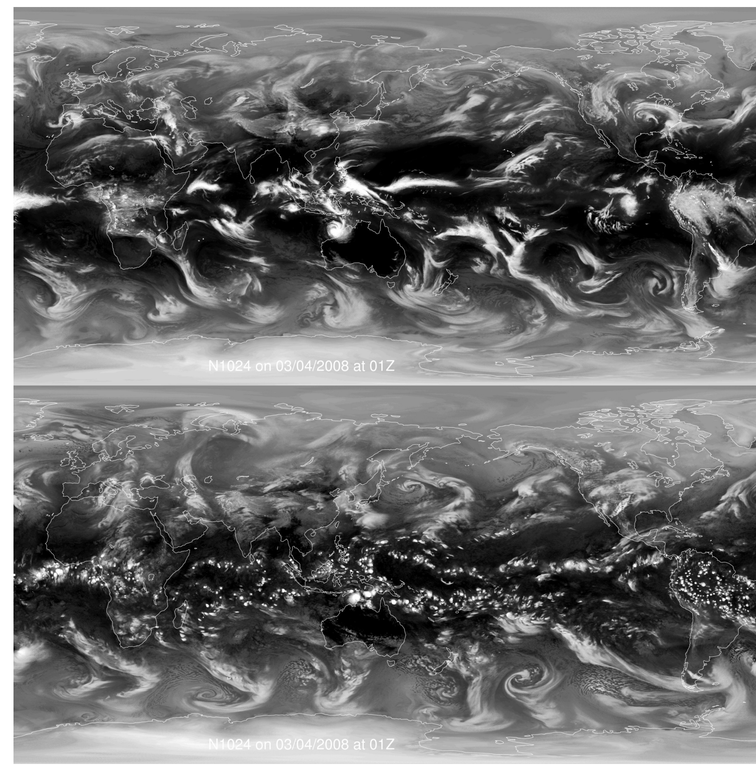

Climate consists of a continuum of time and space scales -

from days to months, years, decades and millennia

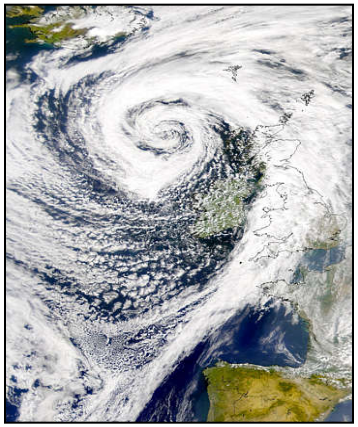

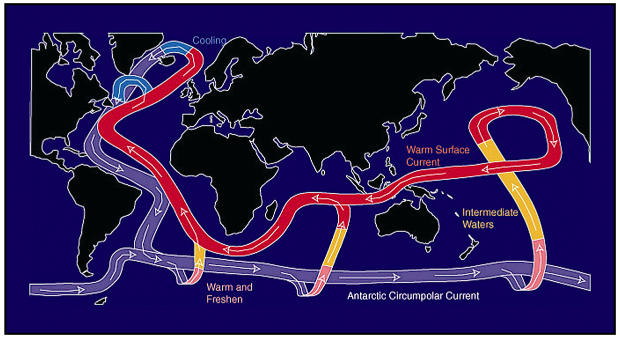

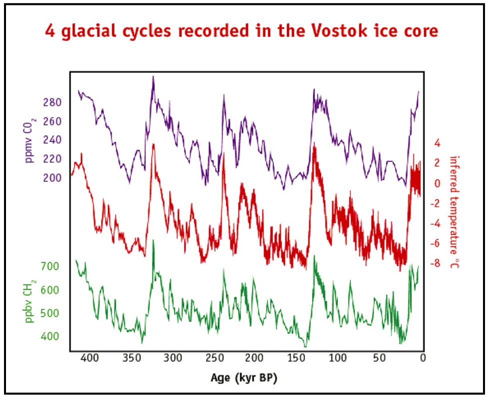

from local to regional, continental and global

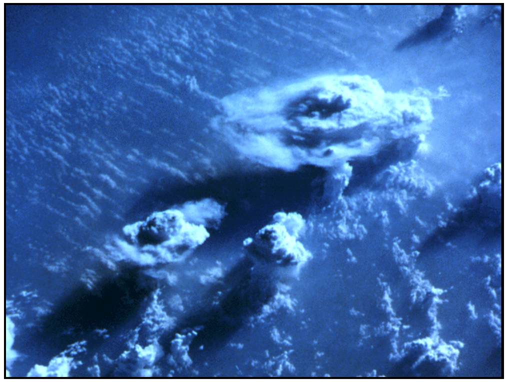

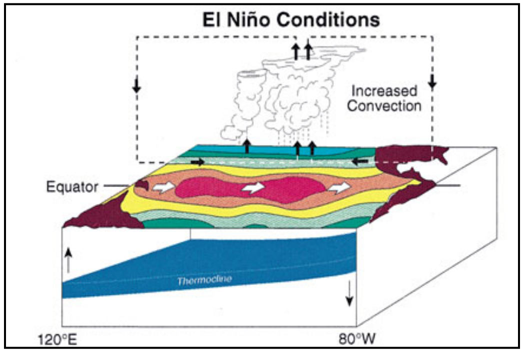

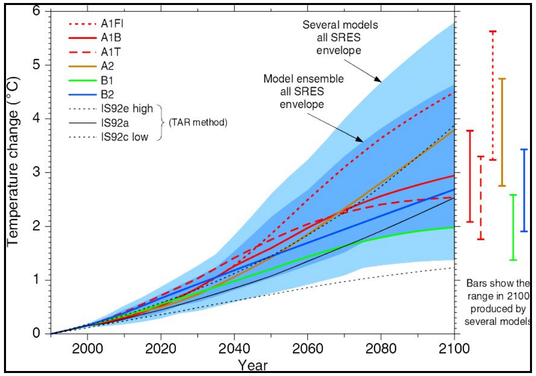

Natural Catastrophes of 2011

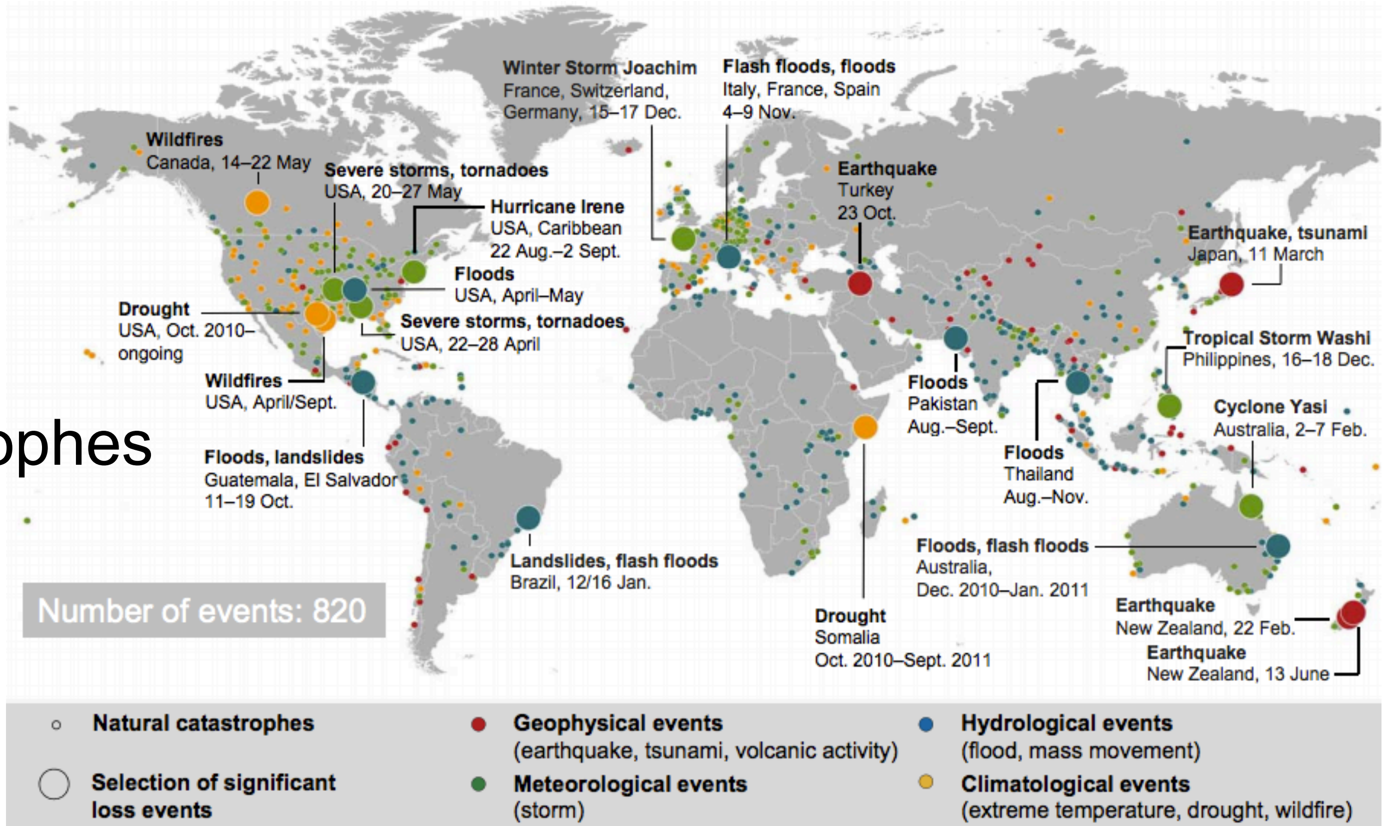

Comparing 2011 with previous years

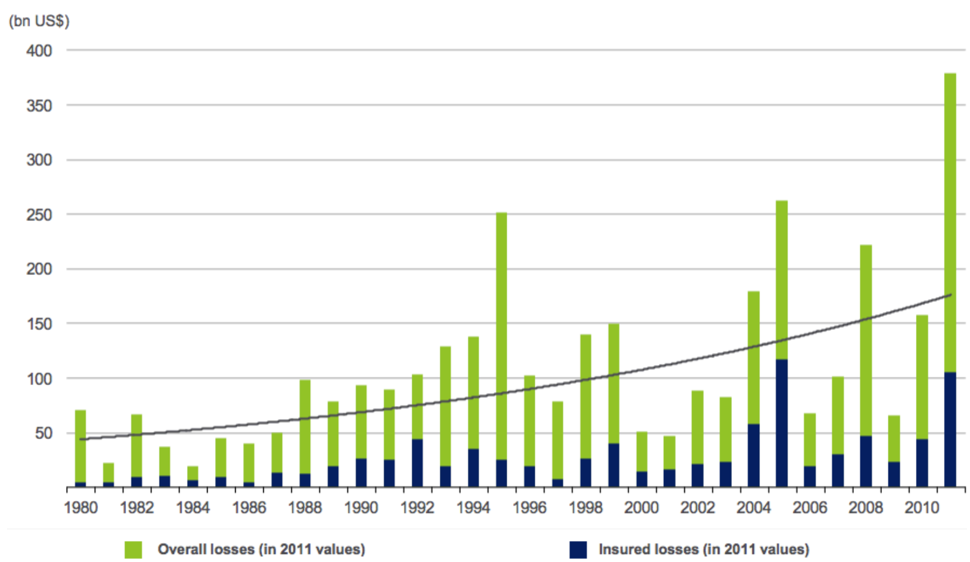

Munich Re 2012

Evolution of N. Atlantic hurricane frequency in past 100+ years: connections with C-Atlantic SSTs.

2005 was a true record year:

15 hurricanes (incl. Katrina), 27 named storms … and some of the most intense storms in US history.

Katrina damage = 1600 dead; 75-200 bn U$.

In the same region, in the Aug-Sep 2007: Dean (cat 5), Felix (cat 5)

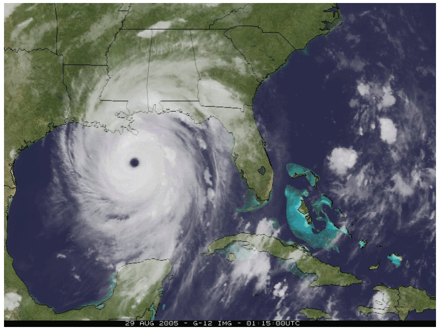

But what about before 1885, or 1950 (for the Pacific) ?

Short/inhomogeneous records of **extremely rare** events:

we need models to complement observations

And yet, coarse GCMs, especially those used for long (e.g. IPCC) integrations, cannot fully represent tropical cyclones

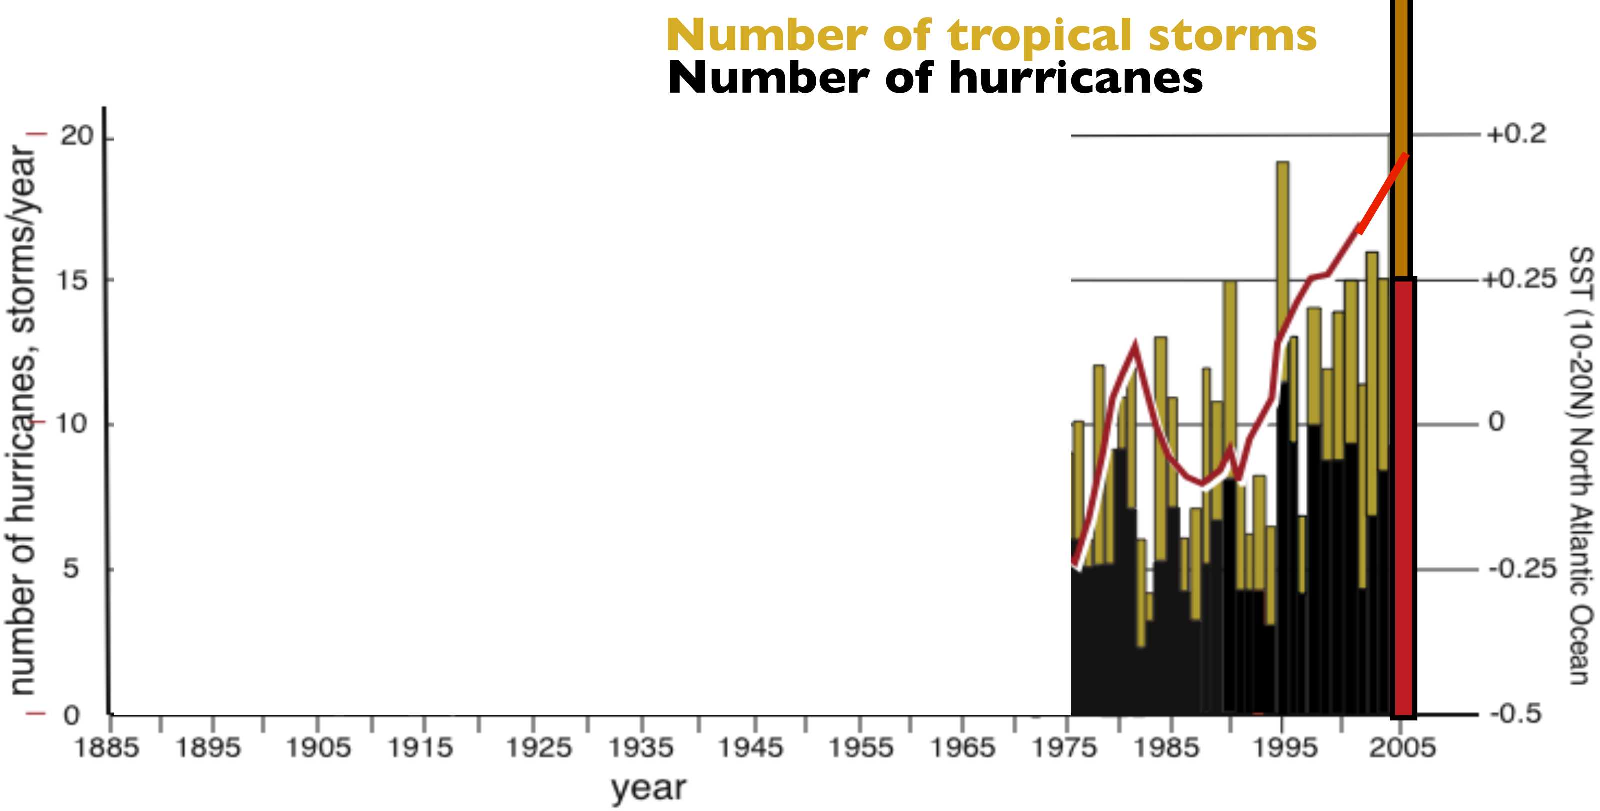

Courtesy of K. Trenberth

Using climate models to complement observational evidence and aid our understanding: decadal variability in hurricane frequency

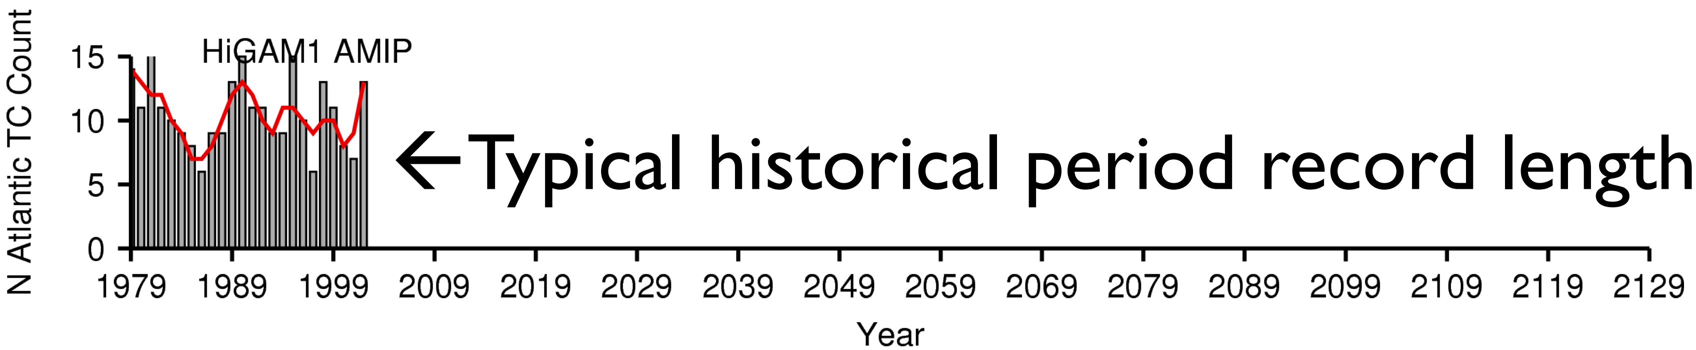

Challenge of a changing climate

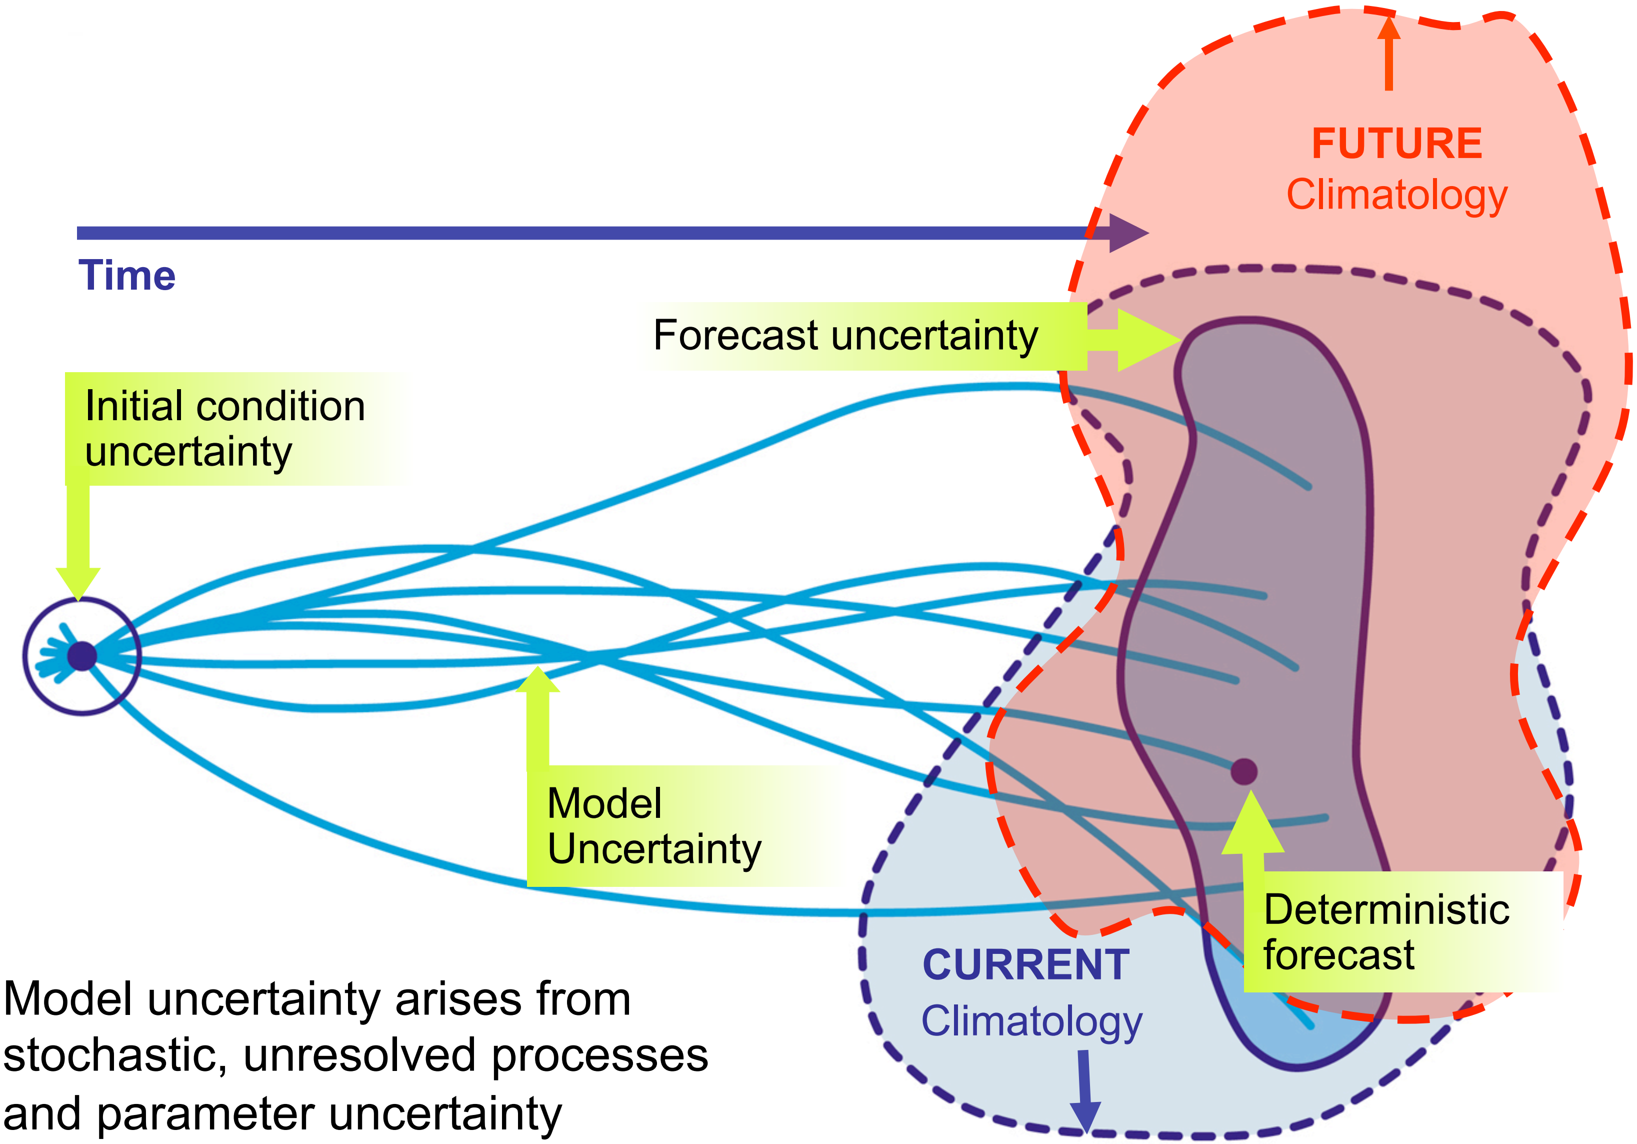

Model uncertainty arises from stochastic, unresolved processes and parameter uncertainty

There is still large uncertainty regarding the regional details of climate change, which is what society really needs.

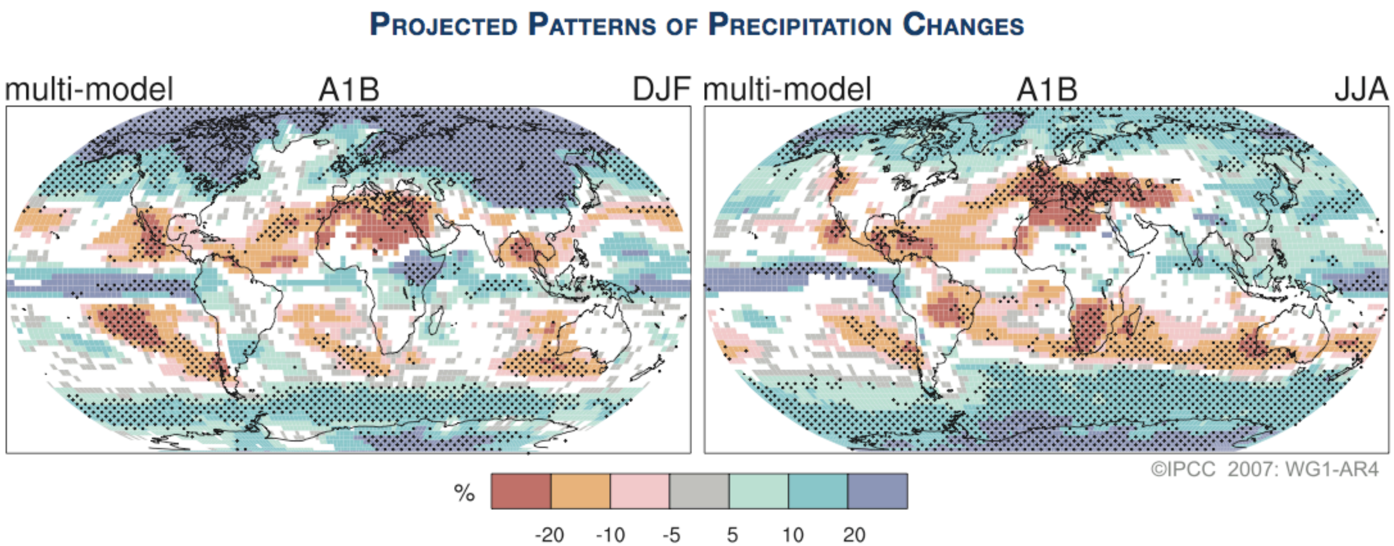

Figure SPM.7. *Relative changes in precipitation (in percent) for the period 2090–2099, relative to 1980–1999. Values are multi-model averages based on the SRES A1B scenario for December to February (left) and June to August (right). White areas are where less than 66% of the models agree in the sign of the change and stippled areas are where more than 90% of the models agree in the sign of the change. {Figure 10.9}*

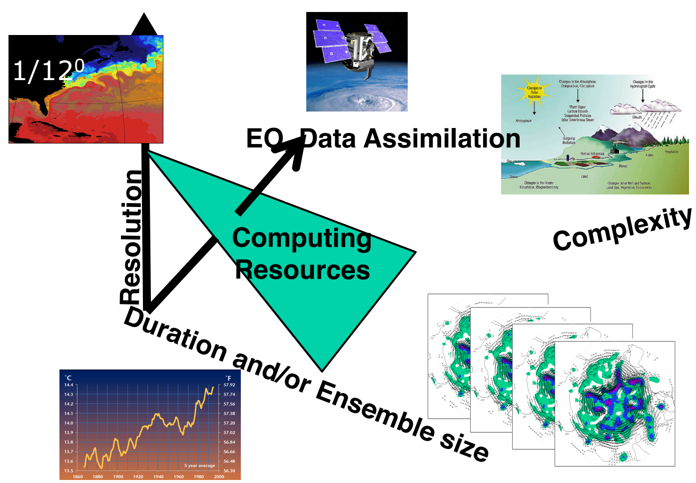

How does a GCM work ?

We slice the planet in boxes… we apply the laws of physics.

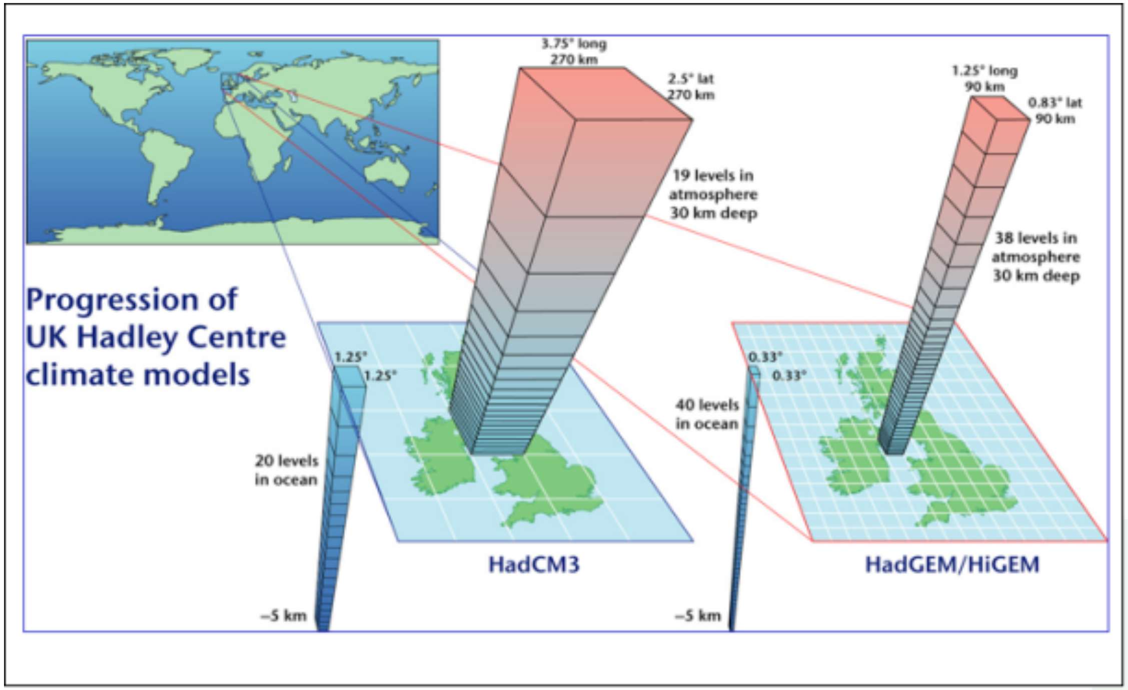

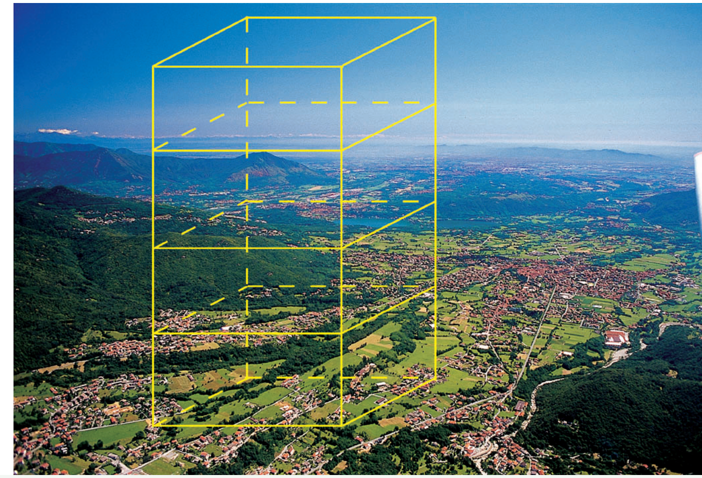

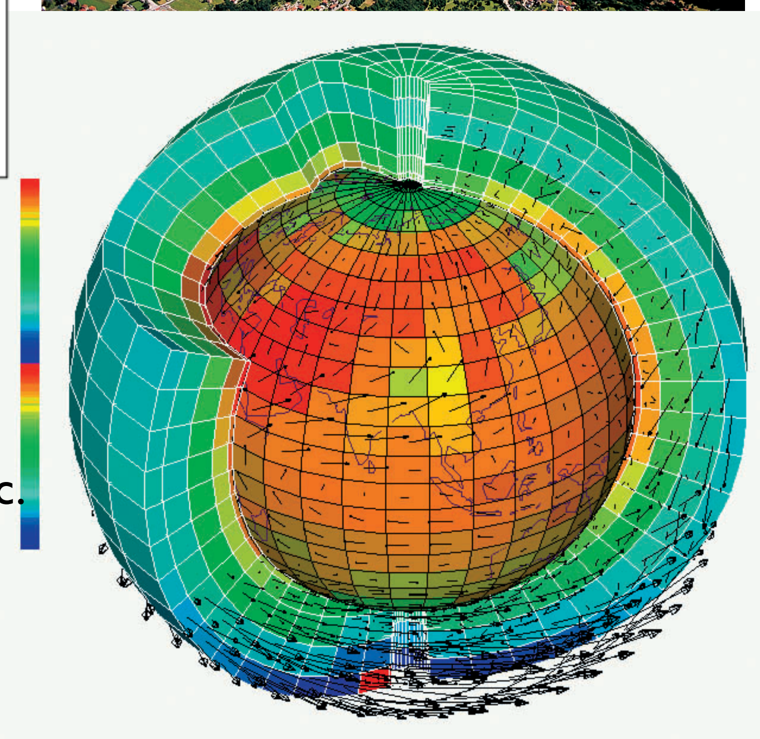

Analytical solutions are too hard: we use discretisation and numerical methods.

At every box location we compute radiation, winds, pressure, precipitation, temperature, using the laws of physics (gravitation, electromagnetism, thermodynamics, fluid dynamics, turbulence), chemistry, biology, ecology, etc.

We do all these computations every 5-30 minutes, for every single box. **We need very large computers !**

Model Grid Size (km) & Computing Capability

	Earth Simulator 2002-2009		PRACE-HERMIT 2012-		
Peak Rate:	10 TFLOPS	100 TFLOPS	1 PFLOPS	10 PFLOPS	100 PFLOPS
Cores	1,400 (2005)	12,000 (2007)	80-100,000 (2009)	300-800,000 (2011)	6,000,000? (20xx?)
Global NWP[0]: 5-1(	18 - 29	8.5 - 14	4.0 - 6.3	1.8 - 2.9	0.85 - 1.4
Se 50-100					1.3
D 5-1(					.2
Change[2]: 20-50 yrs/day	120 - 200	57 - 91	27 - 42	12 - 20	5.7 - 9.1

* Core counts above $O(10^4)$ are unprecedented for weather or climate codes, so the last 3 columns require getting 3 orders of magnitude in scalable parallelization

teraFLOPS = 10^{12} (trillion) floating point operations per second
petaFLOPS = 10^{15} (quadrillion) floating point operations per second
exaFLOPS = 10^{18} (quintillion) floating point operations per second

Range: Assumed efficiency of 10-40%
0 - Atmospheric General Circulation Model (AGCM; 100 vertical levels)
1 - Coupled Ocean-Atmosphere-Land Model (CGCM; ~ 2X AGCM)
2 - Earth System Model (with biogeochemical cycles) (ESM; ~ 2X CGCM)

Thanks to Jim Abeles (IBM)

MetUM global atmosphere/coupled model climate configurations in use

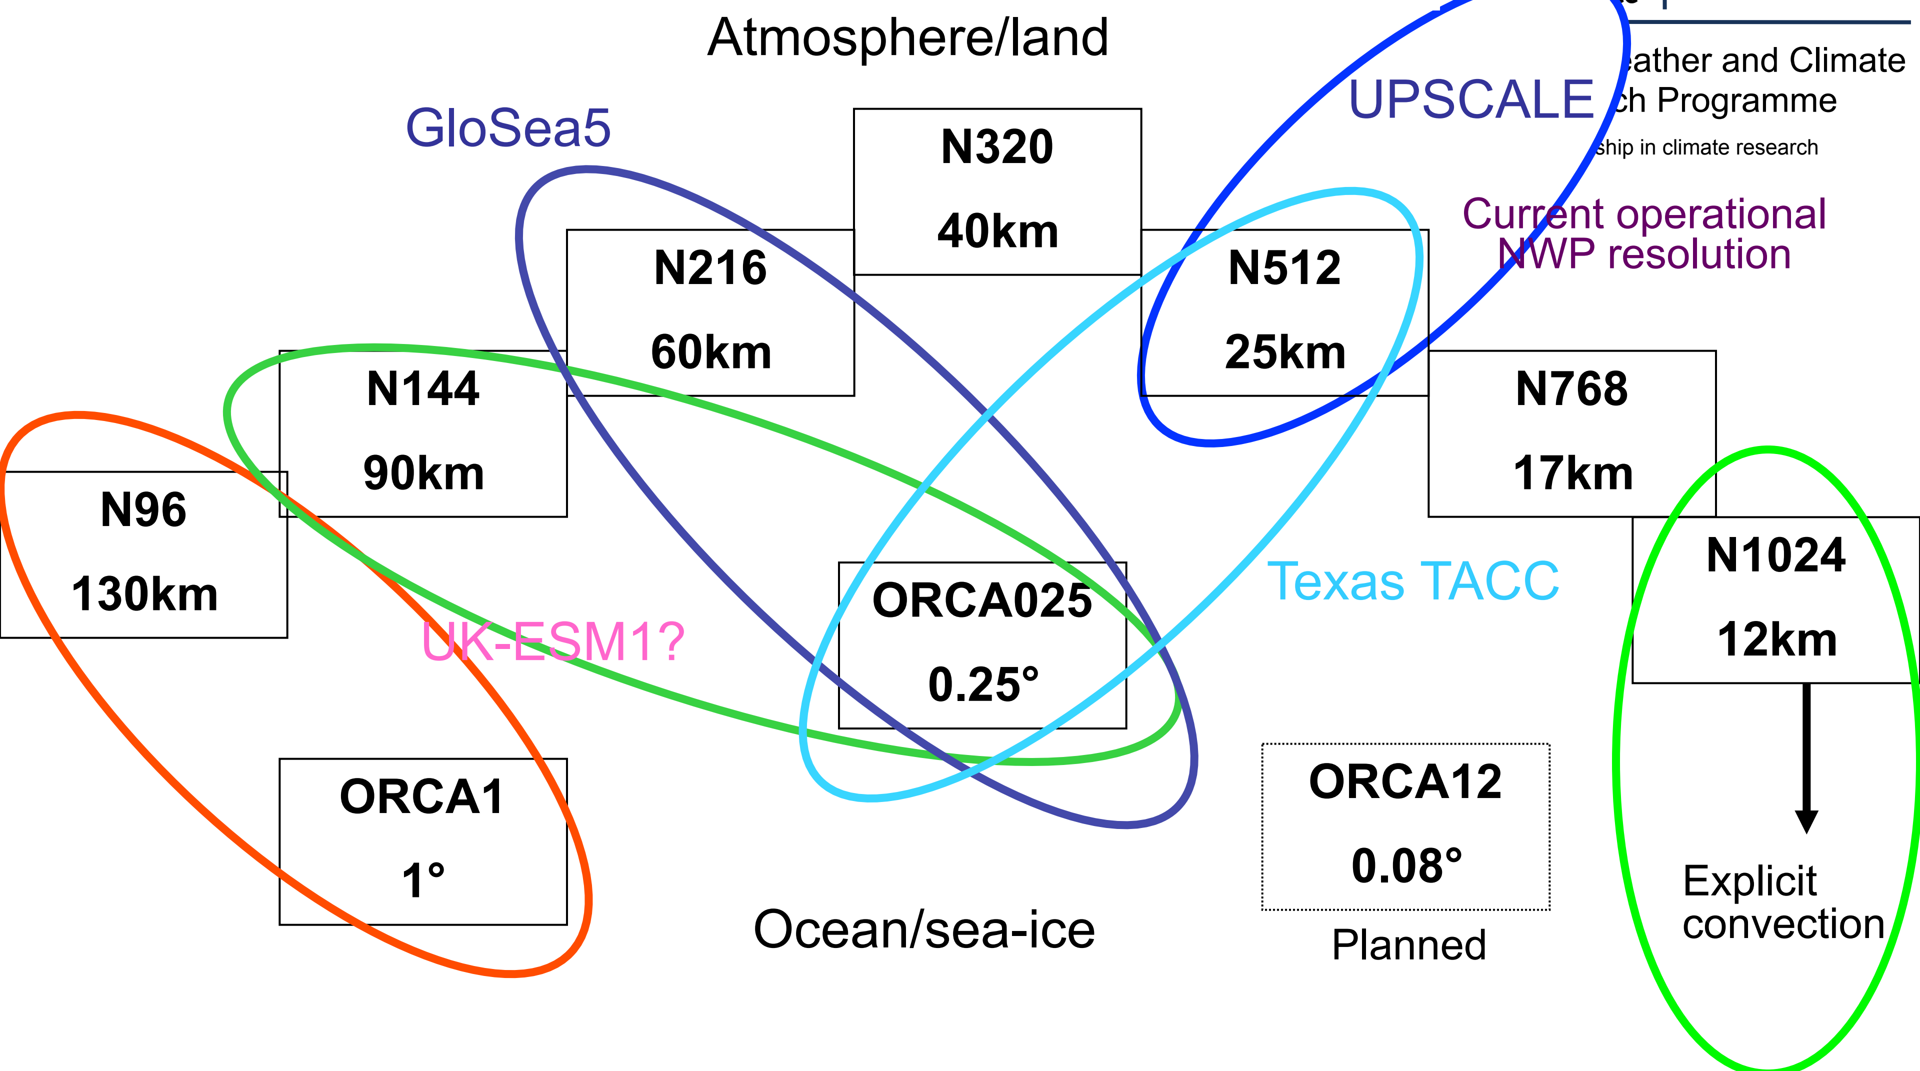

Essentially the same physics/dynamics parameters used throughout model hierarchy

Project to assess impact of global explicit convection

Setup and "best practice" core counts on different architectures

Resolution		Number of cores				Notes
		NEC SX6	IBM P6	IBM P7	CRAY XE6	
N96 = 135km HG1-L38 HG3-L85		1*8 1*8	96	128		For HadGEM1: ES processors about 4x more powerful than P7
	Turnaround	1 sypd 10smo/day	3 sypd	3.5 sypd (EndGame)		
N216 = 60km HG1-L38 HG3-L85		11*8	192	3*32	1024	For HadGEM1: ES processors about 4x more powerful than P7
	Turnaround	1 sypd	8 smo/day	5 smo/day	13smo/day	
N512= 25km				64*32 40*32 (EG) 200*32 (L70 EG)	9408	
	Turnaround			7 smo/day 5.7 smo/day 2 sypd	6 smo/day	Ensemble of 5 runs, concurrent, up to 60K cores
N1024 =				74*32		

The PRACE-UPSCALE Project

Met Office
NATURAL ENVIRONMENT RESEARCH COUNCIL

•Joint Weather & Climate Research Programme
A partnership in climate research

UK on PRACE - weather resolving Simulations of Climate for globAL Environmental risk
Current "numerical mission" of the JWCRP High-resolution climate modelling team
PI: P.L. Vidale, NCAS-Climate, Reading

In 2011 we demonstrated our capability in effectively exploiting 4'800, and up to 12'000 CRAY XE6 cores. As an ensemble of GCMs, we could **concurrently use up to 60'000 cores.**

- Cf. with Earth Simulator: we never managed to effectively use more than 88 cores (out of 5'400 cores in total)

Produced 2-4TB data/day, transferred in real time to the UK, ended up with ~400TB of data

AWARD: 144 million core hours, for 1 year.
Equivalent to : 18x HadGEM2 submission to IPCC (= 8M core hours)
half of the UK HECToR facility

Completed:

1. HadGEM3-A multi-decadal simulations at N96 (130 km) to N512 (25 km)
2. Development of a 12km (N1024) Global Climate Model

•present climate simulations
- forced with OSTIA SSTs
- 1985-2011 (27 years)
- 5 ensemble members, 27 years each

•future climate simulations
- 3 ensemble member, 27 years each
- following RCP8.5
- SST: daily OSTIA + HadGEM2-AO RCP8.5 2100 ΔSST

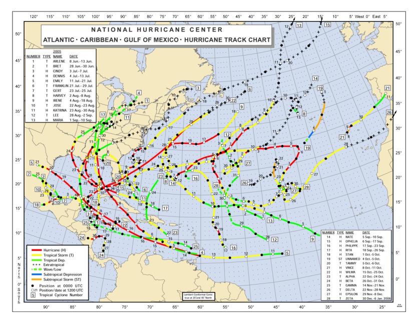

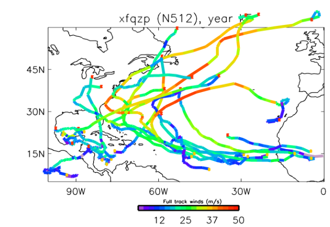

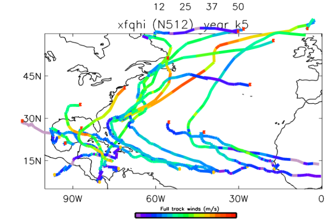

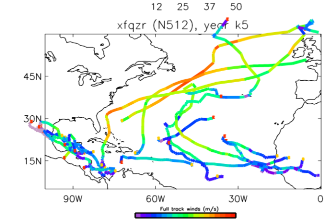

UPSCALE output available on JASMIN@CEDA

Extracting simulated tropical cyclones from model output

TRACK: a "feature tracking" methodology (Hodges,1995)
Independent of model resolution and basin

1) Output 6-hourly global fields (e.g. pressure, winds, humidity, precipitation)
2) Locate and track all centres of high relative vorticity ⇒ **35000 / year**
3) Apply 2-day filter ⇒ **8000 storms / year**
4) Analyse vertical structure of storm for evidence of warm-core (tropical storm structure) ⇒ **120 storms / year**

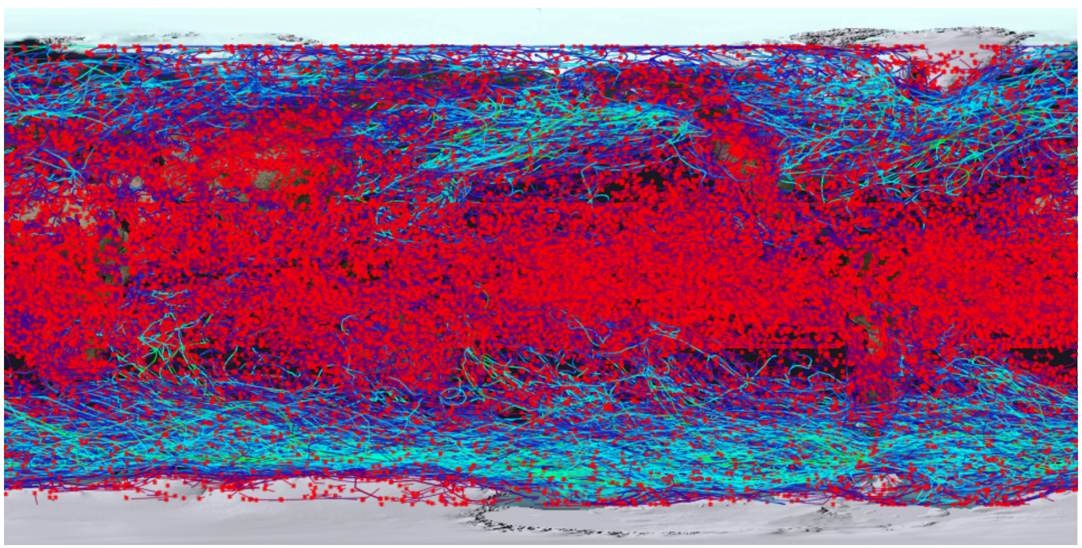

1 year of GCM simulated tropical storms

Consider the **entire life cycle** from genesis, through extratropical transition, to lysis.

Met Office | NATURAL ENVIRONMENT RESEARCH COUNCIL

Joint Weather and Climate Research Programme

A partnership in climate research

PRACE enables Science: Tropical Cyclone Variability

The UPSCALE higher-resolution models simulate these storms far better than IPCC-class models

IPCC class = N96 5 member ensemble

UPSCALE = N512 5 member ensemble

Correlations: N96=0.61 → N512=0.76

Number of hurricanes (ensemble mean and range) over the Atlantic for IPCC-class and UPSCALE models, with HURDAT observations

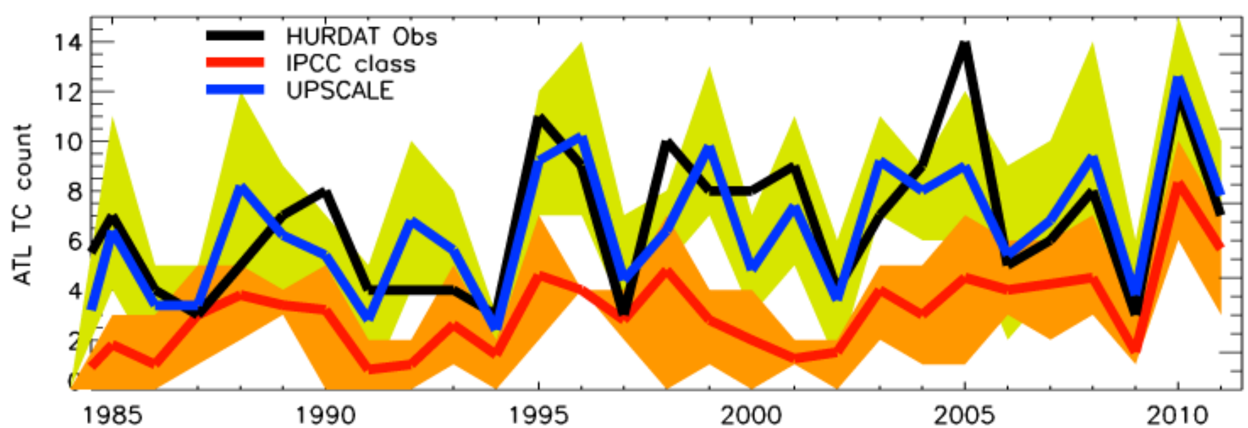

Strachan e et al., J. Clim., 2012 and Vidale et al. 2013, in preparation

Knowledge Transfer Partnerships

Willis Research Network

Where would we like to be?

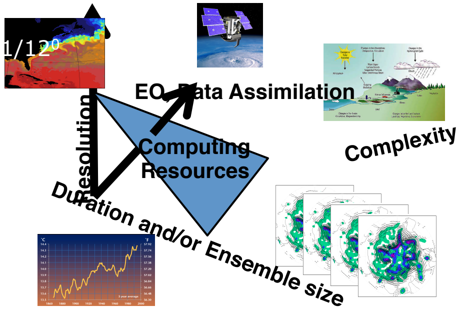

- We are currently running our GCM at 12km
- **Goal**: multiple GCMs at resolution down to 1km, running for decades to centuries
  - Need massive improvements in scalability, from $O(10^4)$ to $O(10^6)$ cores per ensemble member
  - Computational cost 10x larger every time we double the resolution, this would help scalability
- However, data grows even faster: in fact, we could have run a much larger ensemble on HERMIT, but we would never have been able to store/process the data
  - Future models could easily output O(10-100PB) per wall-clock day.

Investigations in the UK's Next Generation Dynamical Core Project
What mesh structure is best (also for local refinement on the sphere) ?

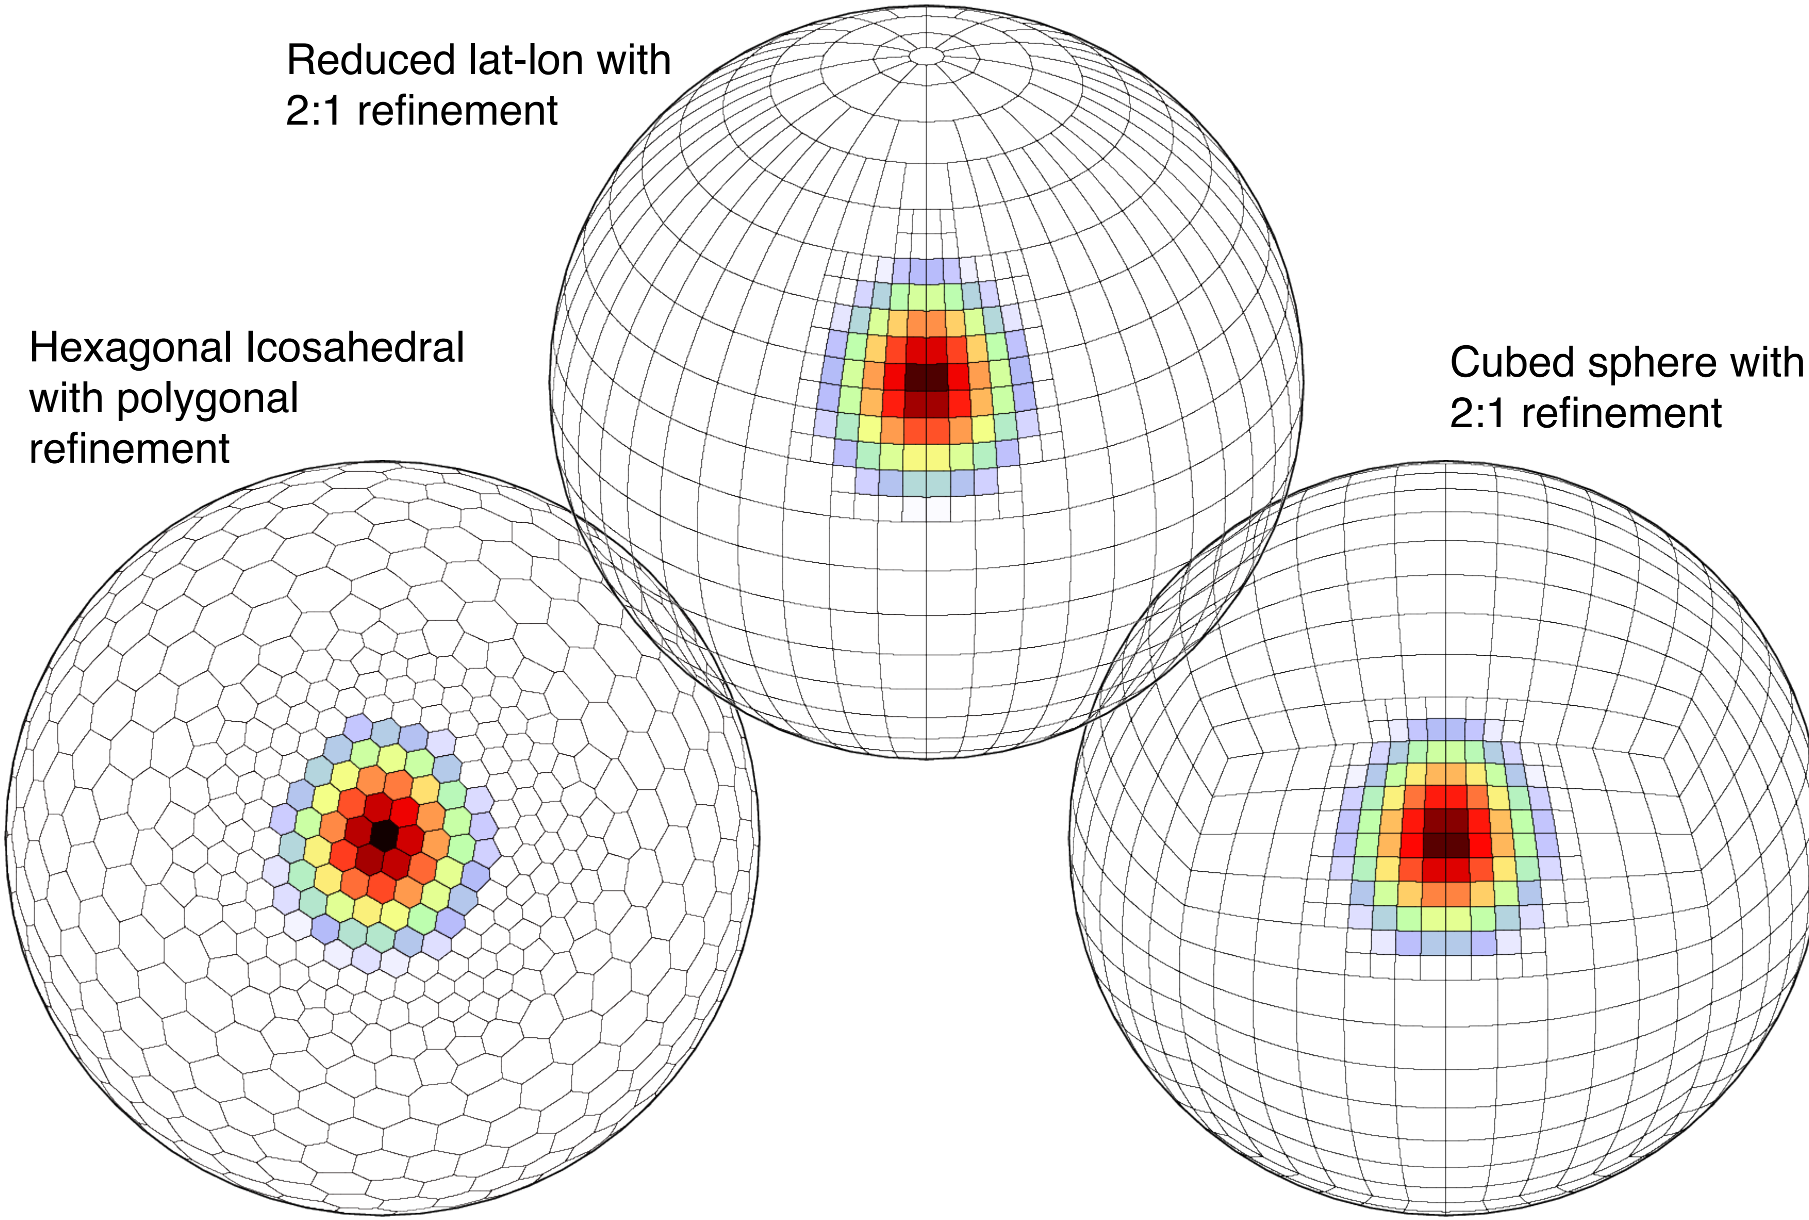

Architecture

- What architectural changes are needed for extreme computing storage systems to make them better suited for BD?
  - Faster communications between cores and between nodes
  - Larger memory caches
  - Much faster access to storage
- What operational changes are needed to support new storage architectures?
  - Special queues that allow for hybrid modelling systems, including couplers and IO servers
  - Support ability to perform analysis locally, where the data are
- Looking at future technologies, what future architectures are possible?

Workflows:

- For extreme computing and big data, describe a forwarding-looking workflow, from simulation to analysis.
  - Significantly reduced porting time
  - Less global communications for the HPC aspect of modelling
  - Keep more data in cache for larger proportion of model integration
  - Perform analysis while model is running: embed diagnostic tools in model framework
  - Perform intra-ensemble analysis in real time, reducing data output/transmission
  - Publish data in a way that allows best scientists and methods to exploit them
- What software is missing to support your workflow?
  - Parallel data processing and analysis
  - Real-time ensemble analysis
  - Standardised analyses, supporting inter-model comparisons
- A plan for achieving interoperability among various systems that one might want to use.
  - Common languages, data models, standard libraries

Taxonomy:

- There are several forms of data-centric computing linked to extreme computing. One outcome of this workshop is to help describe these modes. Please outline how you use your data and how you answer questions about your science using your data.
  - We dump data at every time step (e.g. 10mins), but we also perform time-averaging and we write climatologies while the model runs

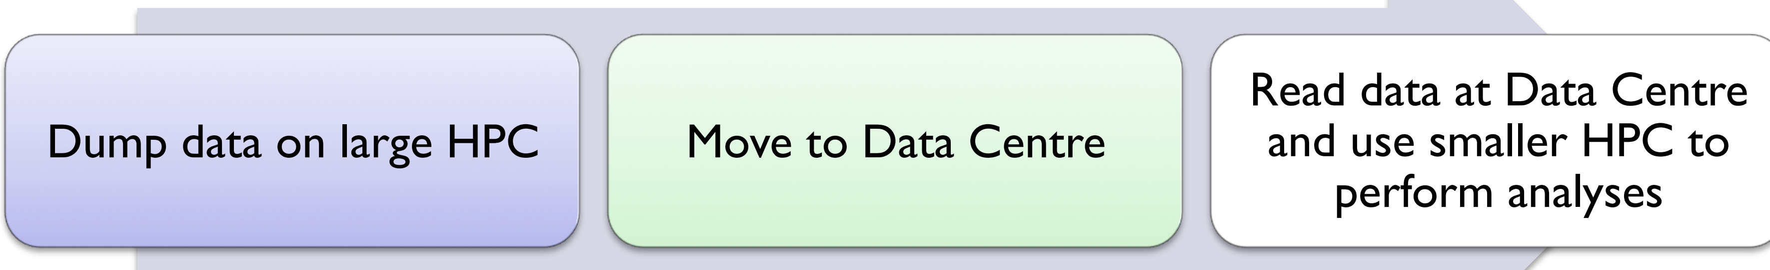

  - Next, we run "feature extraction" algorithms in post-processing, to find the evidence for the science that motivated the numerical experiments
- Do you have a data-driven mini-application that demonstrates a new usage model?
- What are cross-cutting concerns for BD (for example: data integrity)

Software:

- What software are you currently using to manage and explore your data?
  - NetCDF (NCO utilities etc.), IDL, Matlab, R, NCL, SciPy
  - A wealth of custom-made solutions, mostly Fortran and C.
- What algorithms and software libraries/tools need development and improvement to address your big data needs?
  - HDF5, NetCDF4, GRIB need to be better available in exploratory +analysis+graphics packages and they all need to work in parallel mode
- As you look to the future, what are the holes/gaps that have no planned solution?
  - Each group depends too much on custom-made solutions
  - Comparability of model results across groups is quite limited

Interoperability challenges:

- How to handle Data provenance (location, observed/simulated, type of system concerned) from a data representation and IT architectural point of view? How to annotate existing data sets and develop records for data citation and tracking?
  - Extremely important to acknowledge data producers and to make sure that they fully participate in the interpretation/production of science
  - Credit to model developers, for careers
- What Information systems are used for providing semantic capacity to provide effective translation between data and conceptual models used by different communities?
- What IT systems are used for providing information about the actual use of both observational data and simulated data?
  - Web pages: data documentation not automated, so difficult to maintain

W&C Modelling Summary

Ready for exascale?

- On HPC side, not quite, but there is promising on-going work
- On data and data analysis: no
- On working together as a community of W&C scientists, generating hypotheses, asking the same science questions, and being able to answer them? No.

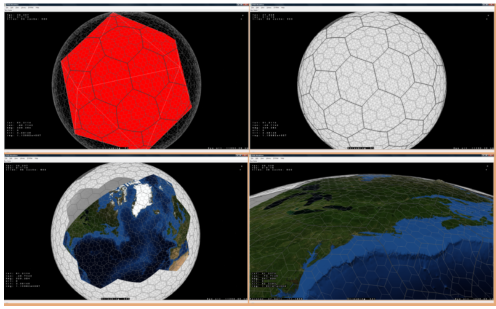

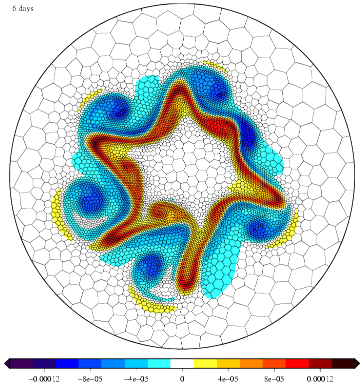

H. Weller, NCAS-Climate Reading