British Atmospheric Data Centre

NATIONAL CENTRE FOR ATMOSPHERIC SCIENCE
NATURAL ENVIRONMENT RESEARCH COUNCIL

ExArch: Climate analytics on distributed exascale data archives

Martin Juckes, V. Balaji, B.N. Lawrence,
M. Lautenschlager, S. Denvil, G. Aloisio, P. Kushner, D. Waliser,
S. Pascoe, A. Stephens, P. Kershaw, F. Laliberte, J. Kim, S. Fiore

UCLA

Institut Pierre Simon Laplace

DKRZ

UNIVERSITY OF TORONTO

Princeton University

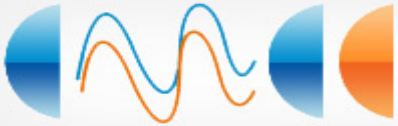

Centro Euro-Mediterraneo per i Cambiamenti Climatici

The G8 exa-scale research initiative

- A joint initiative by research councils of Canada, France, Germany, Japan, Russia, UK, and USA;
- Research into exploitation of "exa-scale" computational resources;
- Focus on 10-year time horizon;
- A unique opportunity for funded international collaboration;
- But with restrictions – only funding to 7 participating countries; restricted eligibility within participating countries (US labs not eligible);

British Atmospheric Data Centre
NATIONAL CENTRE FOR ATMOSPHERIC SCIENCE
NATURAL ENVIRONMENT RESEARCH COUNCIL

Funded projects

- *Enabling Climate Simulation at Extreme Scale*,
  Marc Snir, Uni. Illinois
- *Icosahedral-grid Models for Exascale Earth System Simulations*,
  Günther Zängl, DWD
- *ExArch: Climate analytics on distributed exascale data archives,*
  Martin Juckes, BADC
- *Using next generation computers and algorithms for modelling the dynamics of large biomolecular systems,*
  Makoto TAIJI, AICS
- *Nuclear Fusion Simulations at Exascale*,
  Graeme Ackland, Edinburgh Uni.
- *Modeling Earthquakes and Earth's Interior based upon Exascale Simulations of Seismic Wave Propagation,*
  Jeroen Tromp, Princeton

The ExArch project

What

- ExArch is a research project, studying the provision of exa-scale climate data analysis services;
- The research will be supported by some infra-structure and strategy development;
- Infra-structure development will be within the GO-ESSP framework, linked to the CMIP5 and CORDEX archives.

How

- Scalable infrastructure
- Using a Web Processing Service to take the calculation to the data

ExArch: overview

Start: March 1st (ish), 2011
Duration: 39 months
Budget: 1.44 million Euros
Effort: 246 staff months

British Atmospheric Data Centre
NATIONAL CENTRE FOR ATMOSPHERIC SCIENCE
NATURAL ENVIRONMENT RESEARCH COUNCIL

Princeton University

Institut Pierre Simon Laplace

DKRZ

UNIVERSITY OF TORONTO

UCLA

Centro Euro-Mediterraneo per i Cambiamenti Climatici

- Web processing services
- Query syntax
- Common information model
- Processing operators and quality control
- Scientific diagnostics
- EO data for model evaluation
- Grid computing

Advisory Board

- Dean Williams (PCMDI)
- Colin Jones (co-chair of the WCRP Task Force on Regional Downscaling)
- Ghassem Asrar (WCRP director)
- Pierre-Philippe Mathieu (ESA Climate Change Initiative)

Work Packages

Strategy

- 2 Workshops
- Interactions with GCOS, ESA and NASA
- Governance structures
- Accessibility

Informatics

- Software management
- Robust metadata
- Query management
- Near archive processing

Climate Science

- Quality assurance
- Climate science diagnostics

Science drivers

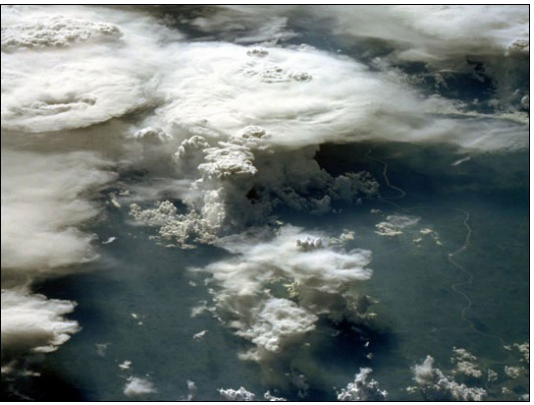

Earth Science Image Analysis Lab.

Improving precision
Dealing with complexity
Sampling uncertainty:
- *Initial condition*
- *Internal variability*
- *Model uncertainty*

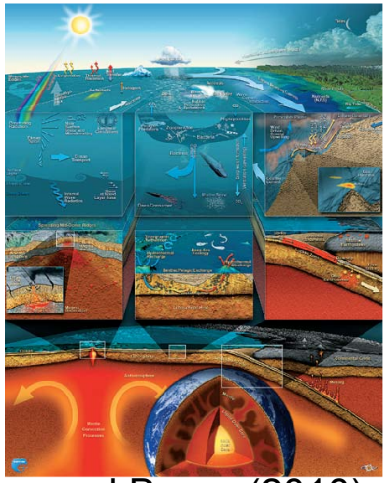

Delaney and Barga (2010)

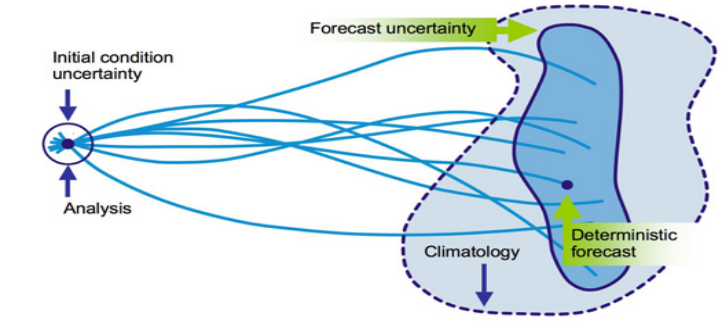

Steve Easterbrook

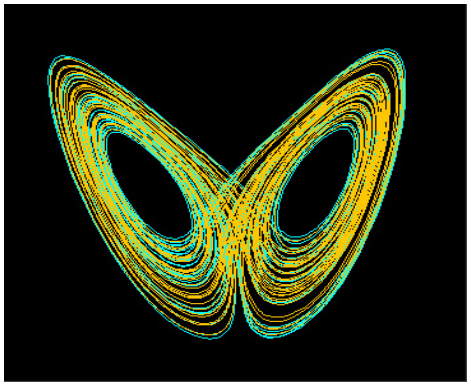

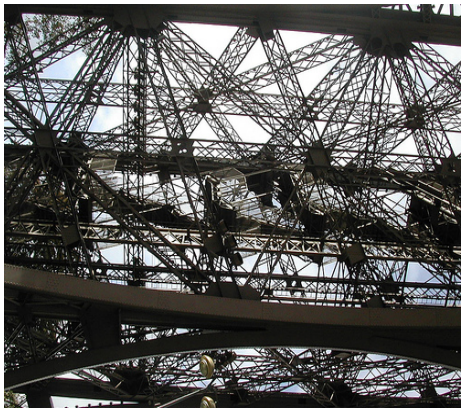

Technology drivers and constraints

- Exa-flop computers
  - Accelerators (GPUs) will provide rapidly increased computational power in the coming decade;
- DM (data movement) bound computing
  - The costs (power, time, infrastructure) of moving data (across a chip, chip to chip, core to cache, cache to work space, work space to archive, archive to user) are overtaking those of the initial computation;

	CMIP5	CMIP6	CMIP7
Year	2012	2017	2022
Computation power (normalised)	1	30	1000
Npp (points pole-to-pole)	200	357	647
Resolution [km]	100	56	31
Number of mesh points [millions]	3.2	18.1	108.4
Ensemble size	100	178	324
Number of variables	800	1068	1439
Interval of 3-dimensional output (hours)	6	4	3
Years simulated	90000	120170	161898
Storage density	0.00004	0.00004	0.00004
Archive size (Pb) (atmosphere)	**5.31**	**143.42**	**3766.99**

N_{pp} = Number of mesh points pole to pole

N_g = Total number of spatial mesh points = $O(N_{pp}^3)$

N_v = Number of variables ~ $\sqrt{N_{pp}}$

N_e = Ensemble size ~ N_{pp}

N_t = Time steps per simulated year ~ N_{pp}

N_y = Years simulated per intercomparison ~ $\sqrt{N_{pp}}$

$$\text{Cost} \sim N_{pp}^6$$

Some hardware trends – ballpark figures

Analysis by Kryder and Soo Kim (2009) suggests hard drives will not be replaced by solid state or other new media before 2020.

	Change per year	Change per decade
Data centre storage	+60%	110-fold increase
Data centre energy use	+25%	9-fold increase
Energy use/unit capacity	-22%	10-fold decrease

	2010	2020
Purchase cost/Tb	200 USD	3 USD
Operating power	10 W/Tb	1W/Tb
Electricity cost (UK)	90 GBP/MWh	120 GBP/MWh
Cost of 1Tb* 3 years	200 + 46	3 + 6
Size at constant funding	1Pb	27Pb

Data trends

Overpeck et al. (2011), Science forecast a 100-fold increase in data volume over 10 years.

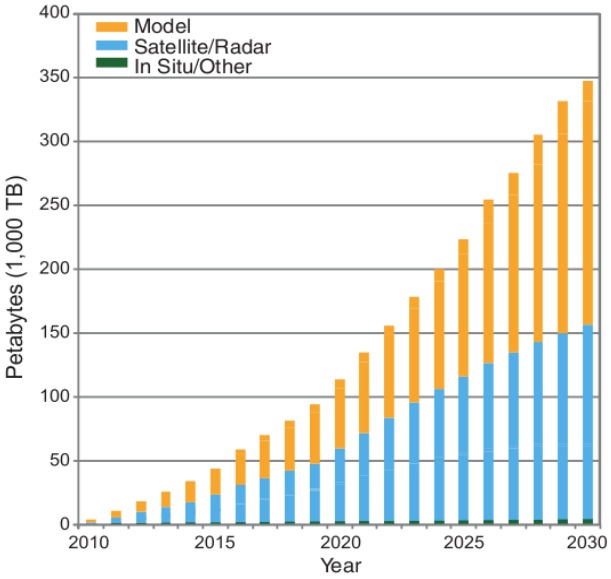

The projections on previous slides can be made consistent with trends projected by Overpeck et al. if:
(1) we archive more selectively by a factor 5;
(2) compress the data;
(3) use parked disks and an efficient caching system;

Governance: why?

- Need a federated archive structure, with multiple funding streams;
- Multiplicity of masters creates conflicts;
- Multiple interacting communities implies multiple governance structures (e.g. NetCDF, NetCDF CF, ESGF, OGC, METAFOR CIM, INSPIRE, OPeNDAP, CSML, DRS(s));
- Getting everyone round the table doesn't scale well.

Objectives

- Efficient cooperation among ESGF members – exploiting synergies, avoiding duplication;
- Facilitate co-design partnerships;

What sort of analysis?

Prototypical query:

The projected frequency of intense tropical cyclones in some region of the globe for input into an impacts model?

Query execution:

- evaluation of provenance and quality control meta-data to determine which datasets to include;
- despatch of queries to processing nodes, negotiating authentication and access control layers;
- collection of results from the processing nodes, evaluation of return codes for fault detection;
- further calculations to combine collected results;
- archive results for re-use;
- delivery of processed results to the end-user, perhaps in deferred fashion if the associated computation needs to be scheduled on a "cloud".

Workshops

1. A workshop on network capacity and data transport software will be held in Europe in of 2012 Autumn, coinciding with a GO-ESSP meeting. The workshop will define high, medium and low scenarios for expected services in the coming decade.

2. A workshop on service requirements and delivery options will be held in San Francisco, 2013 (coinciding with the fall AGU), coordinated by PCMDI.

Inputs: some benchmarks on current systems

Themes: taking the calculation to the data

- Exploit OGC framework for scalable and interoperable services;
- Use existing libraries for repeatable processing (e.g. CDO);
- Exploit existing standards – NetCDF CF, CMOR, METAFOR CIM;
- Support data ingestion into ESGF archives (CORDEX, NASA and ESA climate datasets);
- Delegated (i.e. chained) security (MashMyData → ESGF);

Core objective

Install CDO library behind an OGC Web Processing Service interface.

Web processing Services

The CEDA OGC Web Services package has been used in the UK Climate Impacts Programme, providing access to climate projections with uncertainty information.

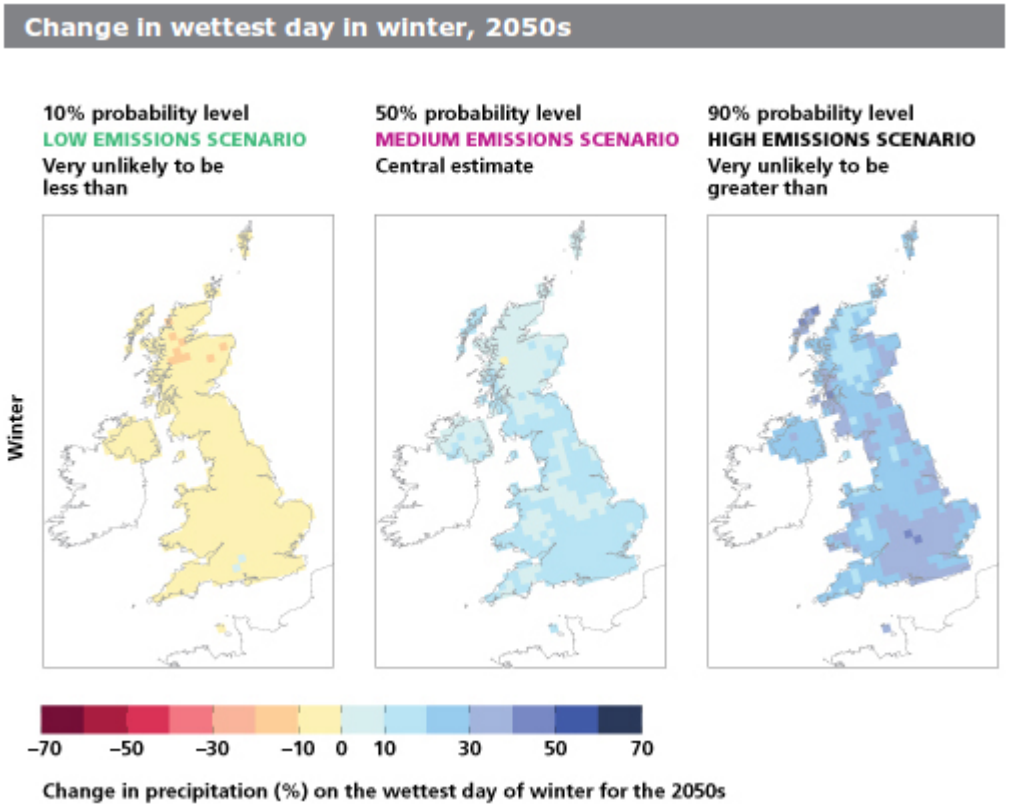

Some other topics addressed by ExArch

- Software management;
- Automated meta-data collection;
- METAFOR CIM profile for CORDEX;
- EO data;
- Query scope and syntax;
- Data categorisation (aka data reference syntax);
- Caching;
- Access from grid resources;
- Security: delegation;
- Quality assurance;
- CDO: scalability and portability;
- Comparing model and EO data;
- Case studies:
  - cyclones;
  - projections in seasonal snow cover;
  - moist thermodynamics.

The end

How will Earth System Science for Climate evolve?