Generic lessons from data management in support of climate simulation workflows

Bryan Lawrence & a cast of thousands

NCAS &
University of Reading: Departments of Meteorology and Computer Science

UK Turbulence Consortium meeting — March 2023

Outline

- Motivation: Climate Simulation at Scale
- Standards and Conventions (CF and CMIP)
- Pushing Data Infrastructure
  - Aggregation
  - In-flight Diagnostics
  - Active Storage
- FAIR
- Summary

Movies

Simulation Movies, courtesy of Prof Pier Luigi Vidale
NCAS
University of Reading

Motivation

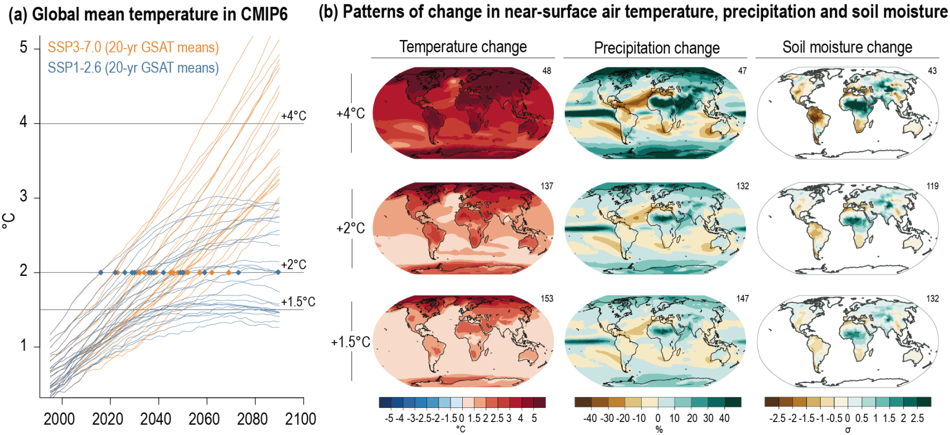

AR6: Warming levels and projected change

What lies beneath?

- Each panel describes a thing called “Model Agreement” and is addressing “changes” for “scenarios”.
- So what is this thing called a model?
- Scenarios?
- And what is agreement?
- There are several panels (many “variables” aka “fields”).
- And what systems exist to get at these data and make these calculations?

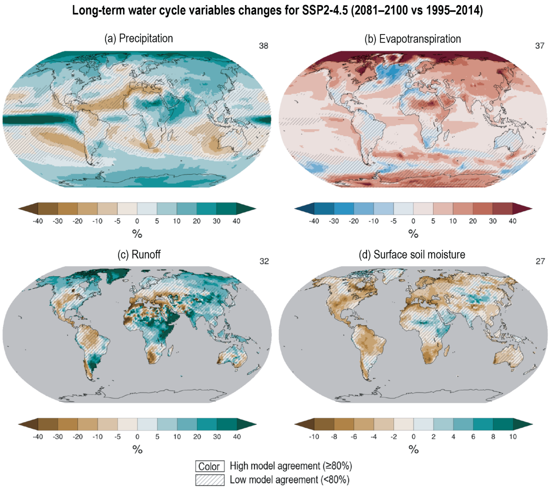

What lies beneath?

- Each panel describes a thing called “Model Agreement” and is addressing “changes” for “scenarios”.
- So what is this thing called a model?
- Scenarios?
- And what is agreement?
- There are several panels (many “variables” aka “fields”).
- And what systems exist to get at these data and make these calculations?

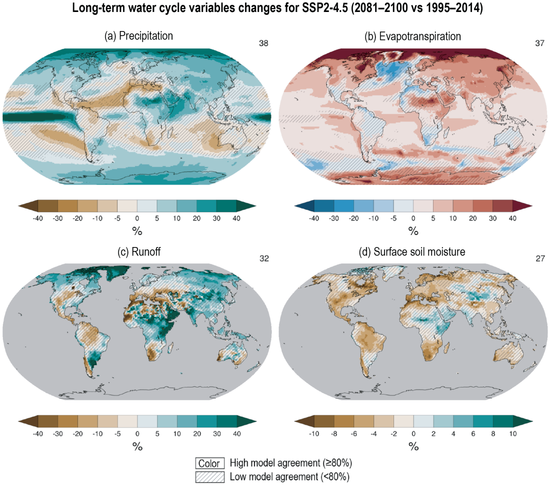

We want to simulate our world

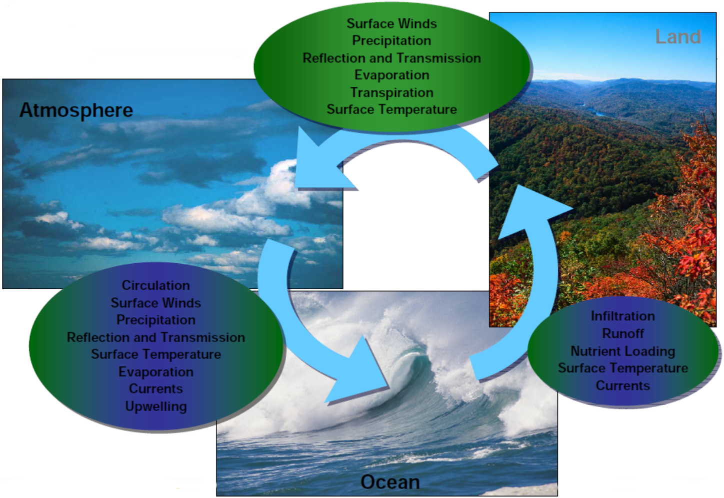

Image: from J. Lafeuille, 2006

Everything is solved on a grid

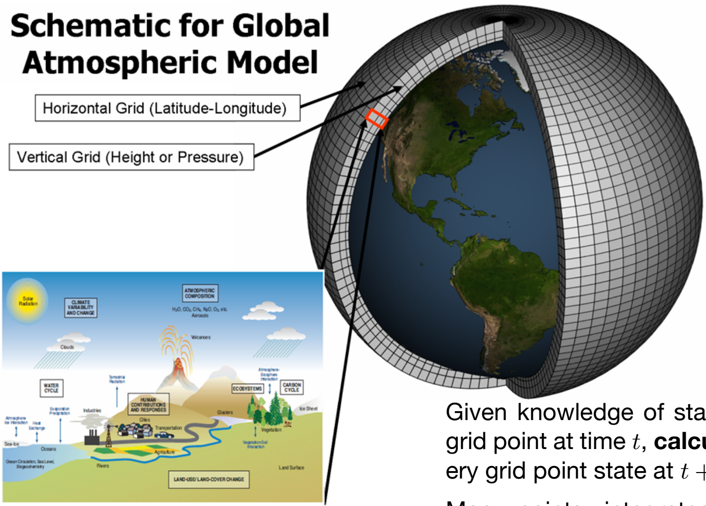

Given knowledge of state at every grid point at time t, **calculate** at every grid point state at $t + \Delta t$.

Many points, integrated for years with timestep of *o(minutes)*!

beyond the fluid atmosphere - Adding more processes

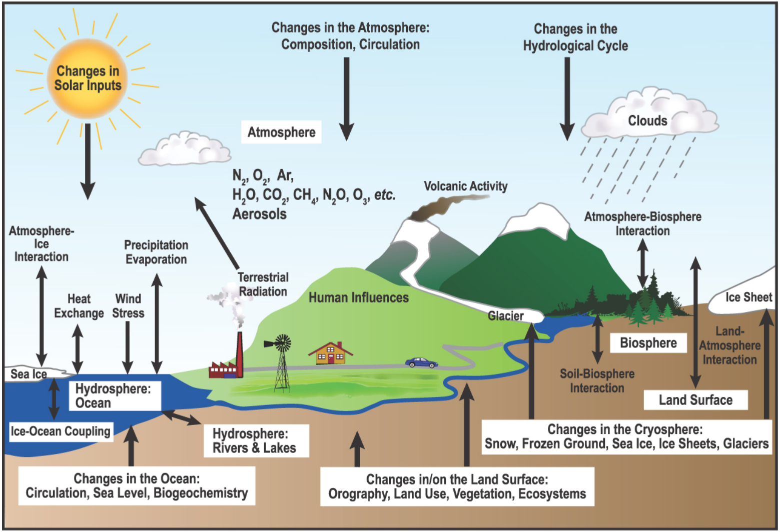

What lies beneath?

- Each panel describes a thing called “Model Agreement” and is addressing “changes” for “scenarios”.
- So what is this thing called a model?
- Scenarios?
- And what is agreement?
- There are several panels (many “variables” aka “fields”).
- And what systems exist to get at these data and make these calculations?

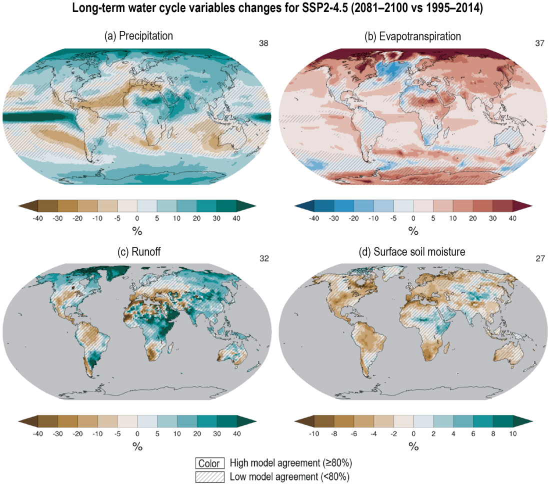

CMIP6 Scenarios

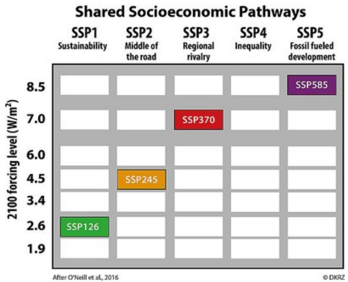

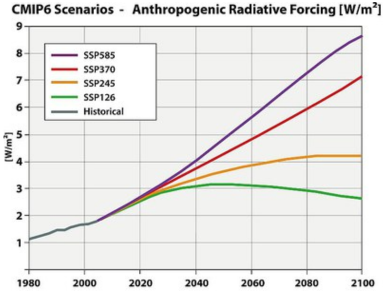

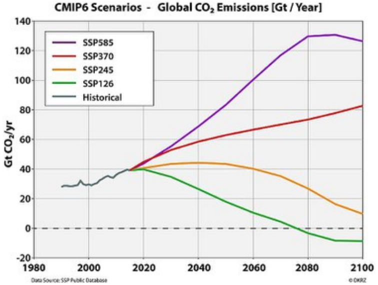

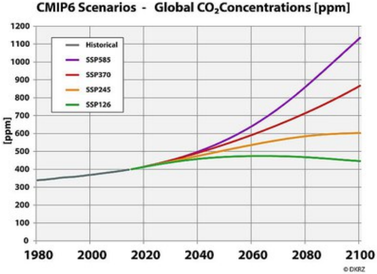

https://www.dkrz.de/en/communication/climate-simulations/cmip6-en/the-ssp-scenarios

Model Intercomparison is about much more than scenarios

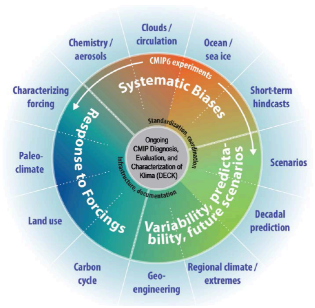

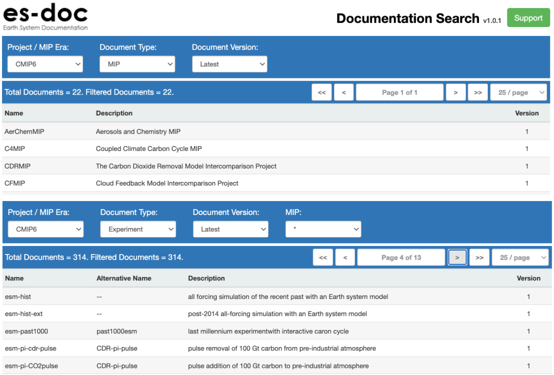
es-doc
Earth System Documentation

Documentation Search v1.0.1 Support

Project / MIP Era: CMIP6 Document Type: MIP Document Version: Latest

Total Documents = 22. Filtered Documents = 22. << < Page 1 of 1 > >> 25 / page

Name	Description	Version
AerChemMIP	Aerosols and Chemistry MIP	1
C4MIP	Coupled Climate Carbon Cycle MIP	1
CDRMIP	The Carbon Dioxide Removal Model Intercomparison Project	1
CFMIP	Cloud Feedback Model Intercomparison Project	1

Project / MIP Era: CMIP6 Document Type: Experiment Document Version: Latest MIP: *

Total Documents = 314. Filtered Documents = 314. << < Page 4 of 13 > >> 25 / page

Name	Alternative Name	Description	Version
esm-hist	--	all forcing simulation of the recent past with an Earth system model	1
esm-hist-ext	--	post-2014 all-forcing simulation with an Earth system model	1
esm-past1000	past1000esm	last millennium experimentwith interactive caron cycle	1
esm-pi-cdr-pulse	CDR-pi-pulse	pulse removal of 100 Gt carbon from pre-industrial atmosphere	1
esm-pi-CO2pulse	CDR-pi-pulse	pulse addition of 100 Gt carbon to pre-industrial atmosphere	1

CMIP6
22 Model Intercomparison Projects
314 Experiments
Each involving multiple models …

European Effort for CMIP6

CMIP6: The sixth phase of the coupled model intercomparison project

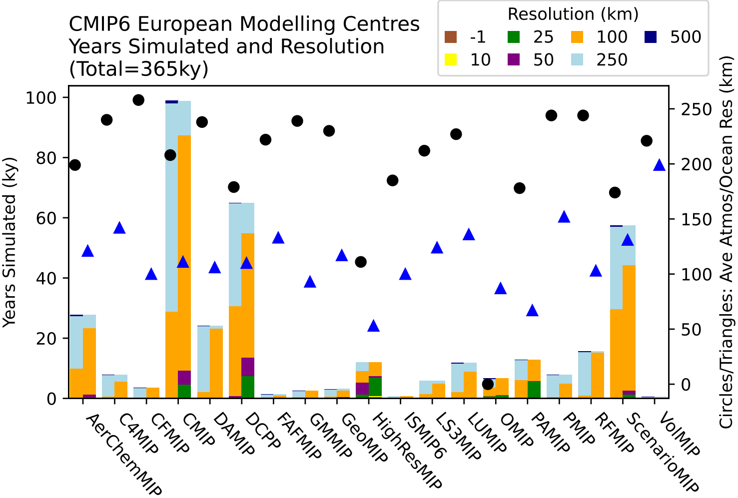

What lies beneath?

- Each panel describes a thing called “Model Agreement” and is addressing “changes” for “scenarios”.
- So what is this thing called a model?
- Scenarios?
- And what is agreement?
- There are several panels (many “variables” aka “fields”).
- And what systems exist to get at these data and make these calculations?

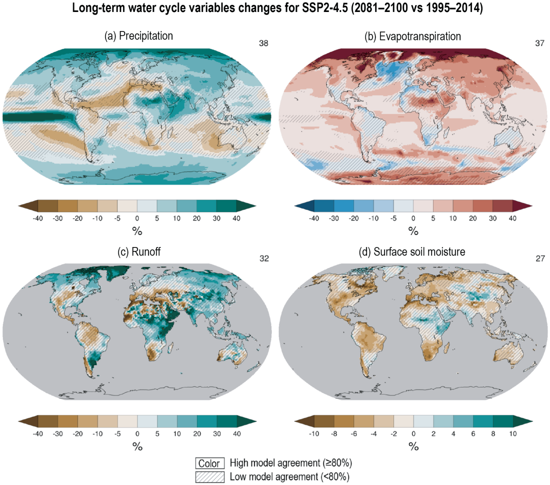

What is agreement? Uncertainty

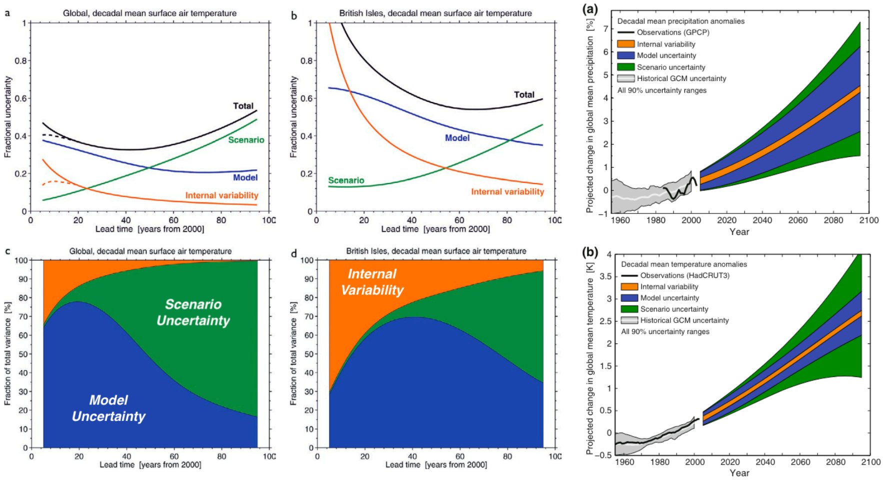

Hawkins, E., & Sutton, R. (2009). The Potential to Narrow Uncertainty in Regional Climate Predictions. Bulletin of the American Meteorological Society, 90(8), 1095–1107. https://doi.org/10.1175/2009BAMS2607.1

Hawkins, E., & Sutton, R. (2011). The potential to narrow uncertainty in projections of regional precipitation change. Climate Dynamics, 37(1), 407–418. https://doi.org/10.1007/s00382-010-0810-6

European CMIP6 Codes

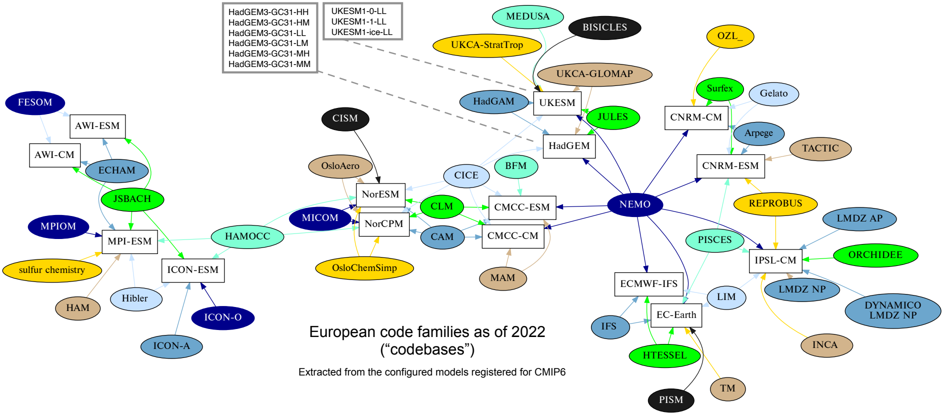

European code families as of 2022 (“codebases”)

Extracted from the configured models registered for CMIP6

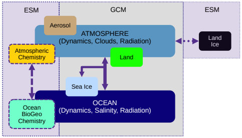

As of March 2023, 50 configured models have been drawn from these modelling families and submitted data to CMIP6.

What lies beneath?

- Each panel describes a thing called “Model Agreement” and is addressing “changes” for “scenarios”.
- So what is this thing called a model?
- Scenarios?
- And what is agreement?
- There are several panels (many “variables” aka “fields”).
- And what systems exist to get at these data and make these calculations?

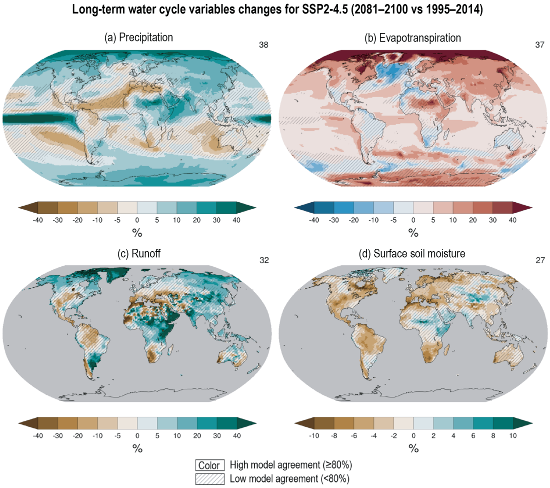

Variables: too many to count?

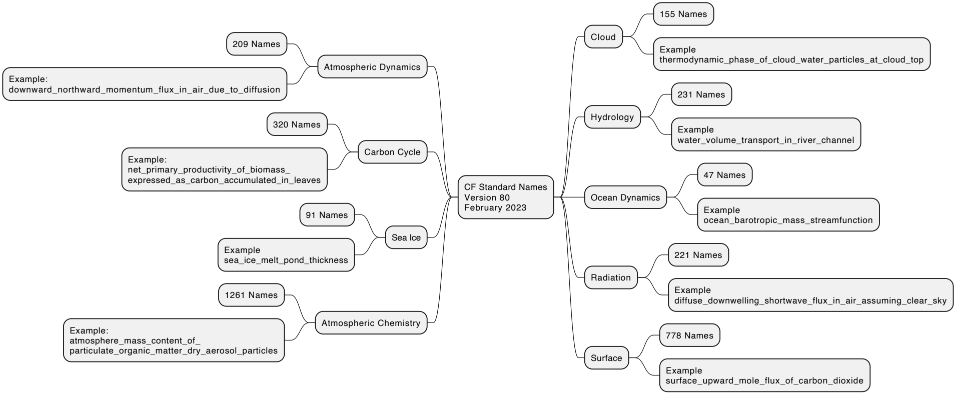

- (Some names are found in multiple categories, but there are still over 2800 standard names.)
- Most earth system simulations will write out hundreds of variables,
- at many different frequencies depending on the variable, from monthly means, to six-hourly point values, etc.

What lies beneath?

- Each panel describes a thing called “Model Agreement” and is addressing “changes” for “scenarios”.
- So what is this thing called a model?
- Scenarios?
- And what is agreement?
- There are several panels (many “variables” aka “fields”).
- And what systems exist to get at these data and make these calculations?

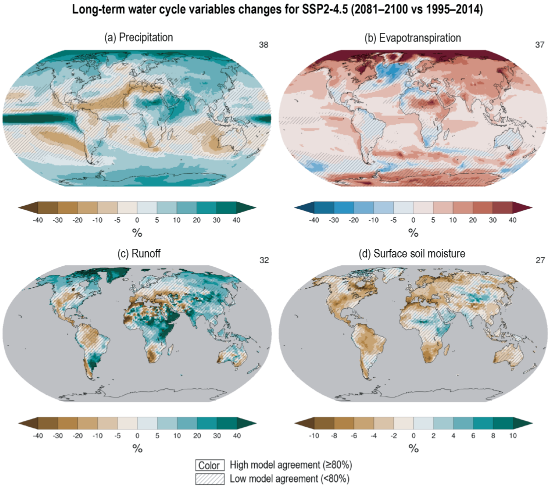

ESGF - Why?

- CMIP3 - 2006 - 35 TB (16 institutes, 24 configured models) was collected and distributed from the U.S.
- (PCMDI at LLNL for those who recognise the acronyms).

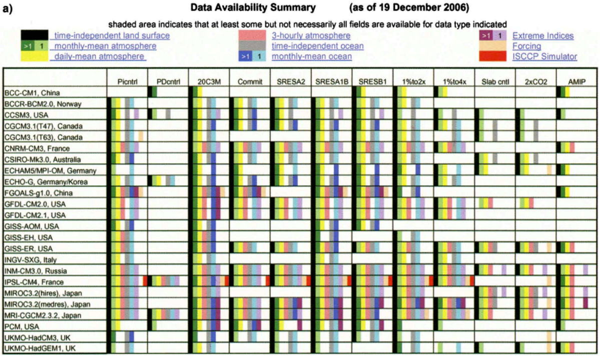

From Meehl et al 2007

CMIP5, first projections for CMIP5 (6+ years later): petabytes!

Central facility not feasible!

Solution: the Earth System Grid Federation (ESGF)

We are now planning CMIP7

ESGF - statistics

CMIP6 13,828,486 total datasets 25,426.81 TB	**CMIP6** 6,871,326 distinct datasets 14,417.12 TB	**CMIP6** 6,957,160 replica datasets 11,009.7 TB
CORDEX 185,906 total datasets 1,459.61 TB	**CORDEX** 185,634 distinct datasets 1,459.05 TB	**CORDEX** 272 replica datasets 0.56 TB
CMIP5 201,130 total datasets 5,293.61 TB	**CMIP5** 52,163 distinct datasets 1,525.07 TB	**CMIP5** 148,967 replica datasets 3,768.55 TB
INPUT4MIPS 11,504 total datasets 19.91 TB	**INPUT4MIPS** 5,660 distinct datasets 9.97 TB	**INPUT4MIPS** 5,844 replica datasets 9.95 TB
OBS4MIPS 210 total datasets 0.2 TB	**OBS4MIPS** 210 distinct datasets 0.2 TB	**OBS4MIPS** 0 replica datasets 0 TB

ESGF - the front end

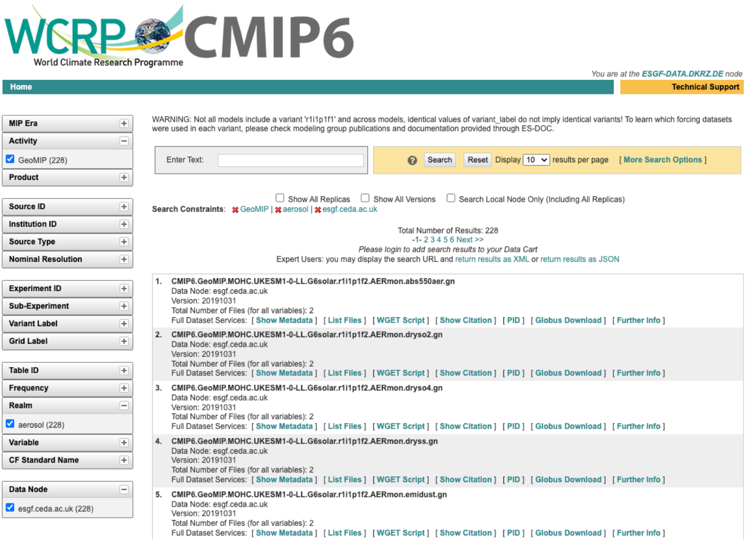

ESGF - globally distributed

Nodes distributed across the planet, but noticeable gaps in the global south.

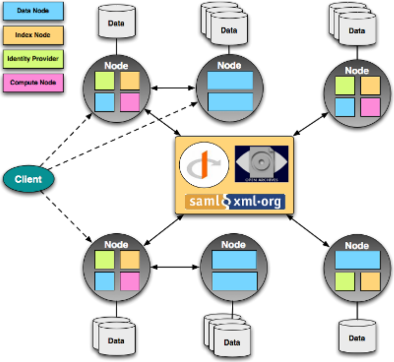

All modelling centres host a data node, some places host index nodes, and replica nodes.

The consequences of data at scale — download doesn't work!

Earth
System
Grid
Experience

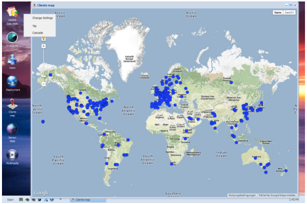

Slide content courtesy of
Stephan Kindermann, DKRZ
and IS-ENES2

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

Started with **Individual End Users**

- Limited resources (bandwidth, storage)

Moved to **Organised User Groups**

- Organize a local cache of files
- Most of the group don't access ESGF, but access cache.

Then **Data Centre Services**

- Provide access to a replica cache
- May also provide compute by data
- CEDA, DKRZ, etc

Trend from download at home, to exploit a cache, to exploit a managed cache with compute!

Simulations and Standards

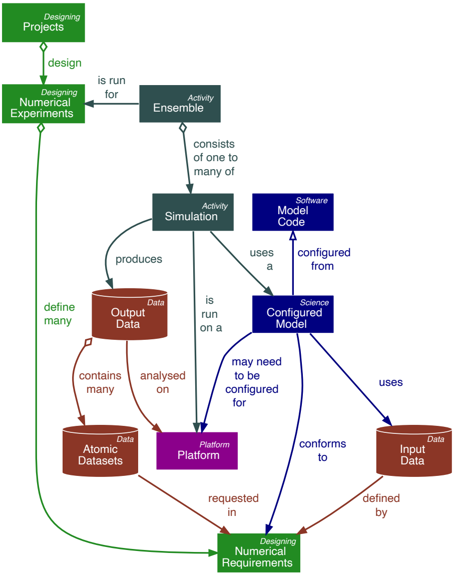

CMIP6 Data

Published to, and available from, the Earth System Grid Federation.

Governed by the WGCM/WIP

Governed by the CF Conventions Committee

Includes many controlled vocabularies for attributes beyond those required by CF

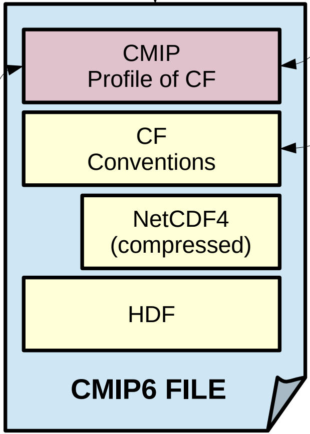

Long-Term Stable

Supported by Institutions with ***mandate***

The Climate Forecast Conventions

CF MetaData Conventions Discussion Governance Training Meetings Vocabularies

CF Metadata Conventions

The CF metadata conventions are designed to promote the processing and sharing of files created with the NetCDF API. The conventions define metadata that provide a definitive description of what the data in each variable represents, and the spatial and temporal properties of the data. This enables users of data from different sources to decide which quantities are comparable, and facilitates building applications with powerful extraction, regridding, and display capabilities. The CF convention includes a standard nam

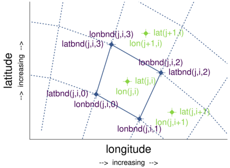

Figure 7.2. Order of `lonbnd(j,i,0)` *to* `lonbnd(j,i,3)` *and of* `latbnd(j,i,0)` *and* `latbnd(j,i,3)` *in the case of two-dimensional horizontal coordinate axes. Tuples* `(lon(j,i),lat(j,i))` *represent grid cell centers and tuples* `(lonbnd(j,i,n),latbnd(j,i,n))` *represent the grid cell vertices.*

The Climate Forecast Conventions

CF MetaData Conventions Discussion Governance Training Meetings Vocabularies

CF Metadata Conventions

The CF metadata conventions are designed to promote the processing and sharing of files created with the NetCDF API. The conventions define metadata that provide a definitive description of what the data in each variable represents, and the spatial and temporal properties of the data. This enables users of data from different sources to decide which quantities are comparable, and facilitates building applications with powerful extraction, regridding, and display capabilities. The CF convention includes a standard nam

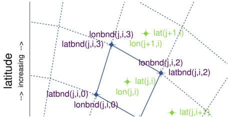

Example 5.11. Latitude and longitude on the WGS 1984 datum + CRS WKT

```
...
 float data(latitude, longitude) ;
   data:grid_mapping = "crs: latitude, longitude" ;
   ...
 int crs ;
   crs:grid_mapping_name = "latitude_longitude";
   crs:longitude_of_prime_meridian = 0.0 ;
   crs:semi_major_axis = 6378137.0 ;
   crs:inverse_flattening = 298.257223563 ;
   crs:crs_wkt =
    GEODCRS["WGS 84",
    DATUM["World Geodetic System 1984",
      ELLIPSOID["WGS 84",6378137,298.257223563,
        LENGTHUNIT["metre",1.0]]],
    PRIMEM["Greenwich",0],
    CS[ellipsoidal,3],
      AXIS["(lat)",north,ANGLEUNIT["degree",0.0174532925199433]],
      AXIS["(lon)",east,ANGLEUNIT["degree",0.0174532925199433]],
      AXIS["ellipsoidal height (h)",up,LENGTHUNIT["metre",1.0]]]
 ...
```

Climate Forecast Conventions and Data Model

Formats and Semantics

- A file format describes how bits and bytes are organised in some sequence on disk.
- Storage Middleware (e.g. NetCDF) has an implicit or explicit data model for what things are stored in that file.
- The Climate-Forecast conventions describe how coordinates and variable properties are stored in NetCDF.
- We have developed an explicit data model so that these can be used for any storage format.

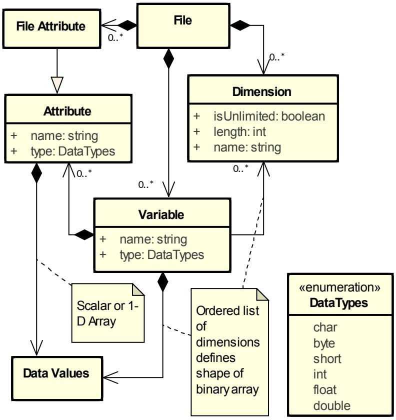

Climate Forecast Conventions and Data Model

Formats and Semantics

- A file format describes how bits and bytes are organised in some sequence on disk.
- Storage Middleware (e.g. NetCDF) has an implicit or explicit data model for what things are stored in that file.
- The Climate-Forecast conventions describe how coordinates and variable properties are stored in NetCDF.
- We have developed an explicit data model so that these can be used for any storage format.

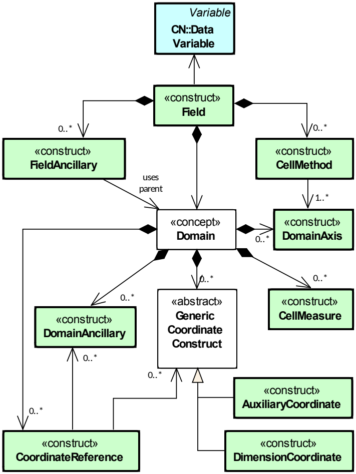

Hassell, D., Gregory, J., Blower, J., Lawrence, B. N., and Taylor, K. E.: A data model of the Climate and Forecast metadata conventions (CF-1.6) with a software implementation (cf-python v2.1), Geosci. Model Dev., 10, 4619-4646, https://doi.org/10.5194/gmd-10-4619-2017, 2017.

cfplot gallery

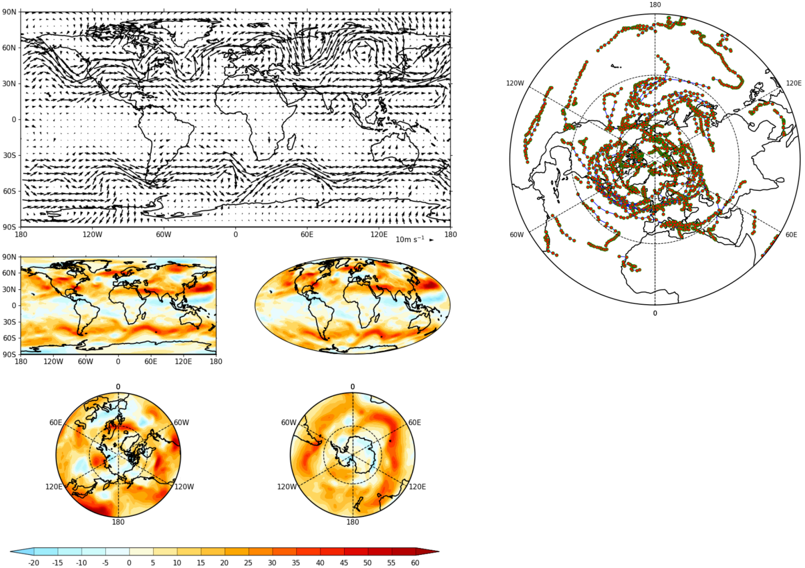

cfplot gallery

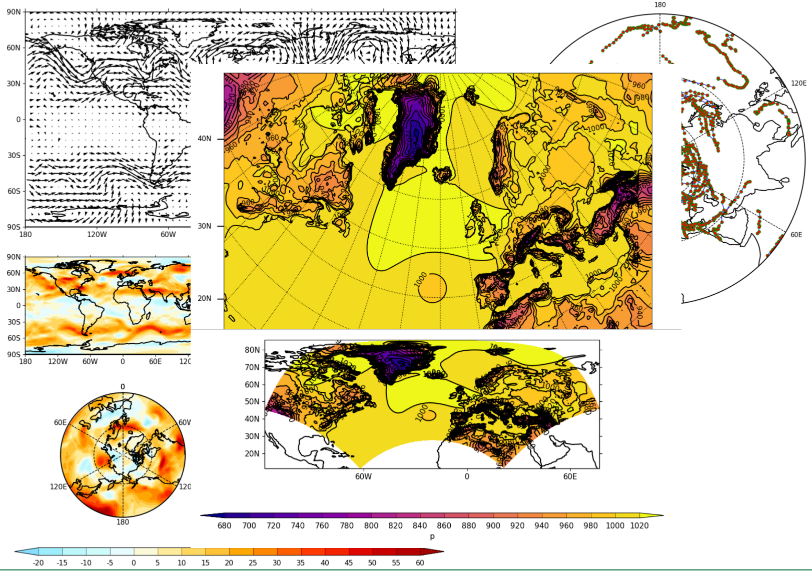

cfplot example

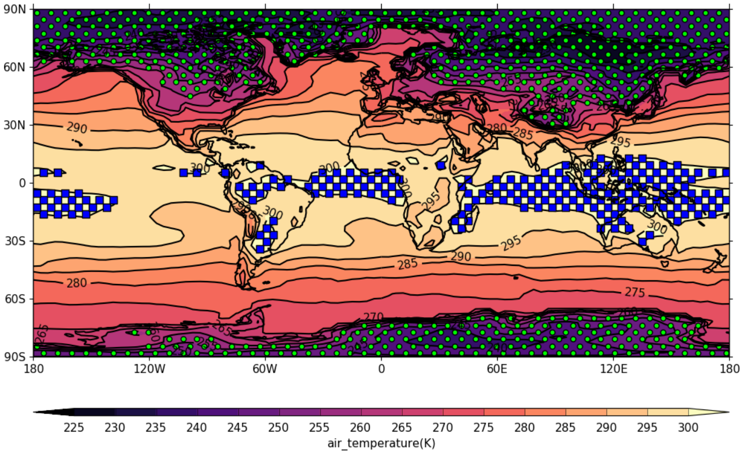

cfplot example

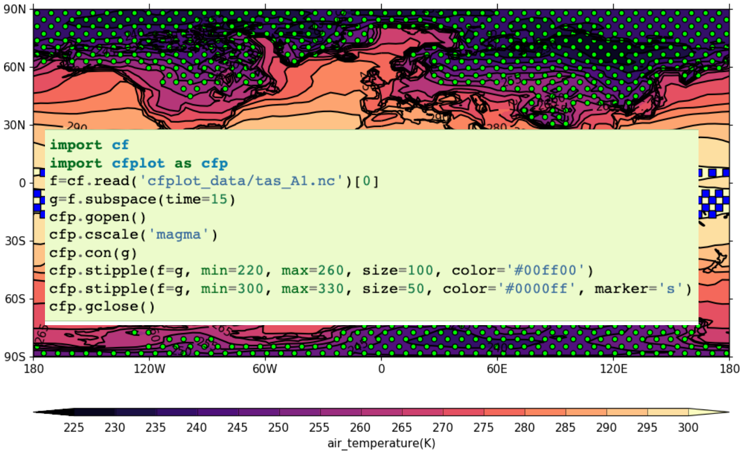


```
import cf
import cfplot as cfp
f=cf.read('cfplot_data/tas_A1.nc')[0]
g=f.subspace(time=15)
cfp.gopen()
cfp.cscale('magma')
cfp.con(g)
cfp.stipple(f=g, min=220, max=260, size=100, color='#00ff00')
cfp.stipple(f=g, min=300, max=330, size=50, color='#0000ff', marker='s')
cfp.gclose()
```

CF Example 1 - Global Mean Anomalies

Consider the following example fairly basic timeseries plot:

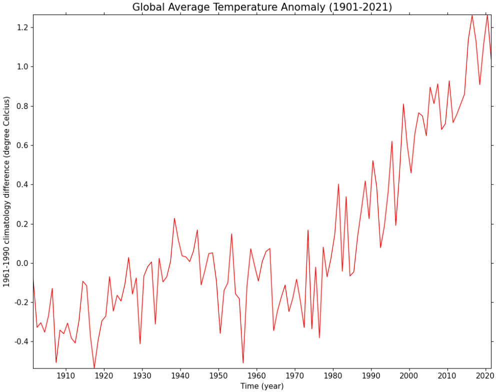

What was involved in doing this? Reading data, geolocation, calculating weights for spatial and temporal averages, averaging, differencing and plotting.

Code and example courtesy of the CF-python team: David Hassell, Sadie Bartholemew, Andy Heaps and Ankit Bhandekar

CF Example 1 - Global Mean Anomalies

Consider the following example fairly basic timeseries plot:

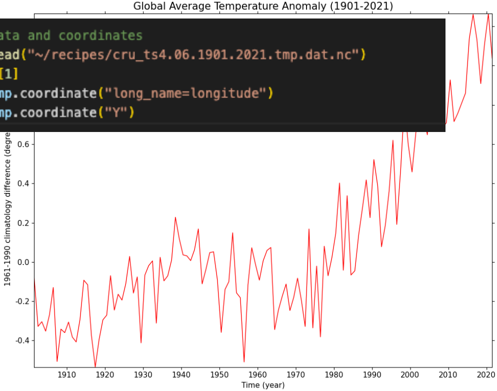


```
# grab data and coordinates
f = cf.read("~/recipes/cru_ts4.06.1901.2021.tmp.dat.nc")
temp = f[1]
lon = temp.coordinate("long_name=longitude")
lat = temp.coordinate("Y")
```

What was involved in doing this? Reading data, geolocation, calculating weights for spatial and temporal averages, averaging, differencing and plotting.

Code and example courtesy of the CF-python team: David Hassell, Sadie Bartholemew, Andy Heaps and Ankit Bhandekar

CF Example 1 - Global Mean Anomalies

Consider the following example fairly basic timeseries plot:

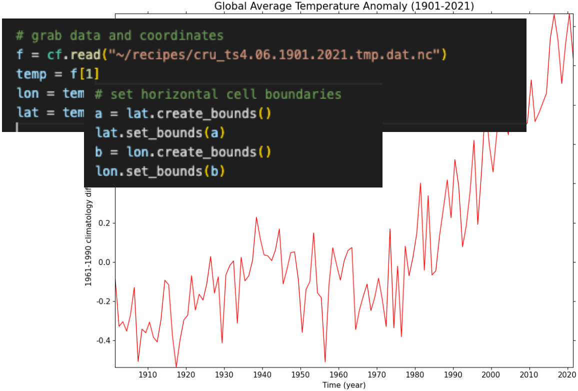

What was involved in doing this? Reading data, geolocation, calculating weights for spatial and temporal averages, averaging, differencing and plotting.

Code and example courtesy of the CF-python team: David Hassell, Sadie Bartholemew, Andy Heaps and Ankit Bhandekar

CF Example 1 - Global Mean Anomalies

Consider the following example fairly basic timeseries plot:

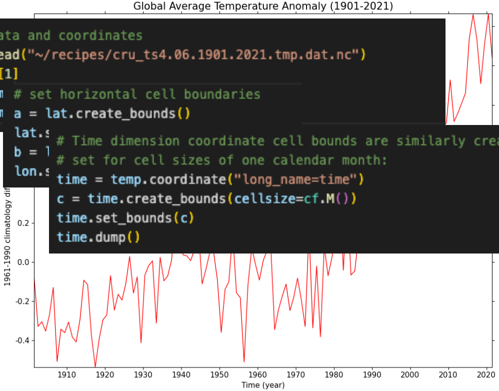


```
# grab data and coordinates
f = cf.read("~/recipes/cru_ts4.06.1901.2021.tmp.dat.nc")
temp = f[1]
lon = tem
lat = tem
```

```
# set horizontal cell boundaries
a = lat.create_bounds()
lat.s
b = l
lon.s
```

```
# Time dimension coordinate cell bounds are similarly created and
# set for cell sizes of one calendar month:
time = temp.coordinate("long_name=time")
c = time.create_bounds(cellsize=cf.M())
time.set_bounds(c)
time.dump()
```

What was involved in doing this? Reading data, geolocation, calculating weights for spatial and temporal averages, averaging, differencing and plotting.

Code and example courtesy of the CF-python team: David Hassell, Sadie Bartholemew, Andy Heaps and Ankit Bhandekar

CF Example 1 - Global Mean Anomalies

Consider the following example fairly basic timeseries plot:

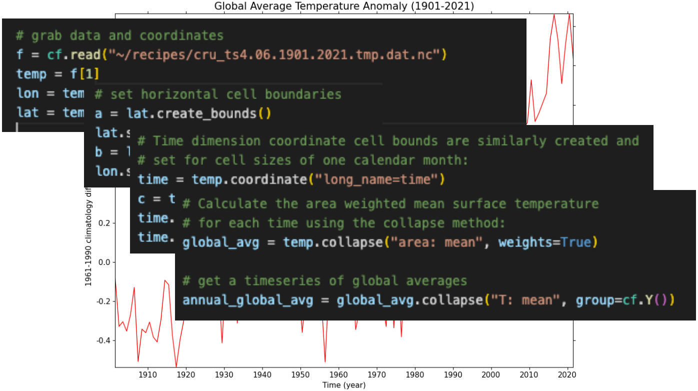

What was involved in doing this? Reading data, geolocation, calculating weights for spatial and temporal averages, averaging, differencing and plotting.

Code and example courtesy of the CF-python team: David Hassell, Sadie Bartholemew, Andy Heaps and Ankit Bhandekar

CF Example 1 - Global Mean Anomalies

Consider the following example fairly basic timeseries plot:

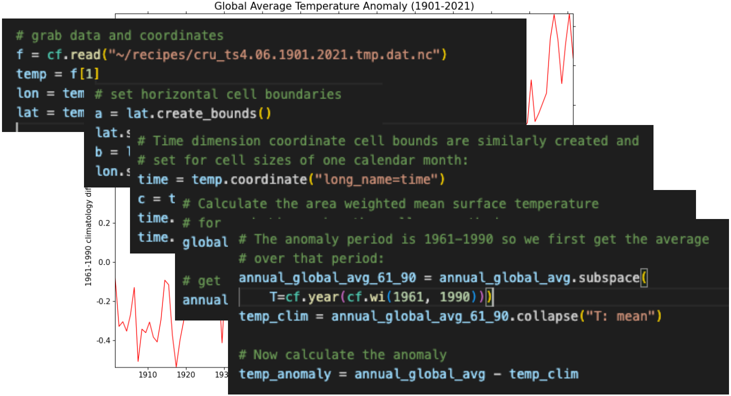

What was involved in doing this? Reading data, geolocation, calculating weights for spatial and temporal averages, averaging, differencing and plotting.

Code and example courtesy of the CF-python team: David Hassell, Sadie Bartholemew, Andy Heaps and Ankit Bhandekar

Application Stack

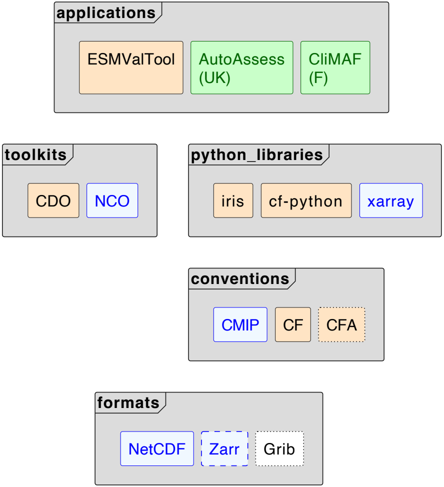

Shared European components in beige, exemplar national applications in green, external dependencies in blue/white.

The Evolution of Resolution: A better global microscope!

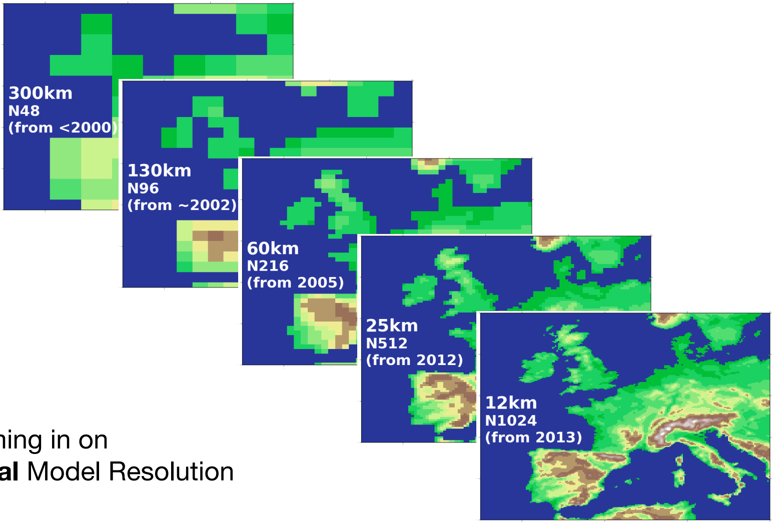

Zooming in on
Global Model Resolution

Volume — the reality of global 1km grids

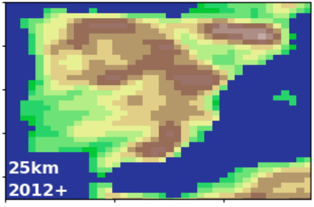

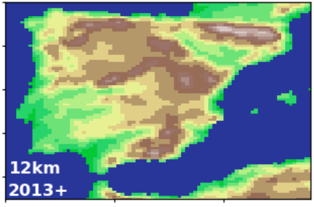

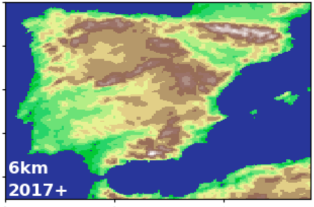

What about 1km? That's the current European Network for Earth System Modelling (ENES) goal!

Consider N13256 (1.01km, 26512x19884)):

- 1 field, 1 year, 6 hourly, 180 levels
- 1 x 1440 x 180 x 26512 x 19884 $=$ 1.09 PB
- 760 seconds to read one 760 GB (xy) grid at 1 GB/s
- but it's worse that that: 10 variables hourly, $>$ 220 TB/day!

Can no longer consider serial diagnostics, and even parallelised is a challenge for the I/O system!

Standards based aggregation

Motivation

- We have lots of files, often lots of files for the same variable.
- Those files can be anywhere on tiered storage, but we tend to talk about “one variable, one-frequency, for one-simulation” as an atomic dataset.
- Can think of an atomic dataset being spread across storage. Individual files can be grouped into quarks each of which is on one tier of storage.

We can have an aggregate view of those files, stored in an aggregation file.

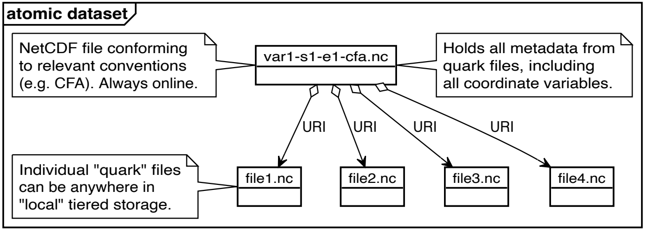

Standards based aggregation

Motivation

- We have lots of files, often lots of files for the same variable.
- Those files can be anywhere on tiered storage, but we tend to talk about "one variable, one-frequency, for one-simulation" as an atomic dataset.
- Can think of an atomic dataset being spread across storage. Individual files can be grouped into quarks each of which is on one tier of storage.

We can have an aggregate view of those files, stored in an aggregation file.

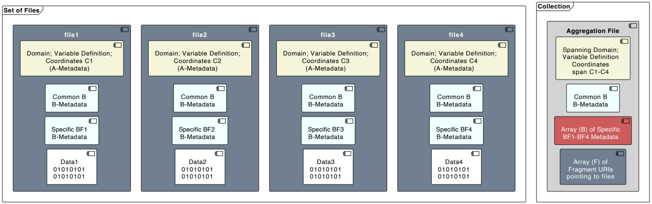

Data movement is expensive

- Moving data over the network, from compute nodes to storage is costly.
- The cost can be thought of as
  - Time or CPU cycles taken to move the data — particularly for read operations, which are always blocking, we can buffer or offload writes;
  - Network bandwidth used (and cost of ensuring the necessary fabric exists);
  - Energy consumed (to move the data).

JASMIN Quobyte Storage Systems (2021)

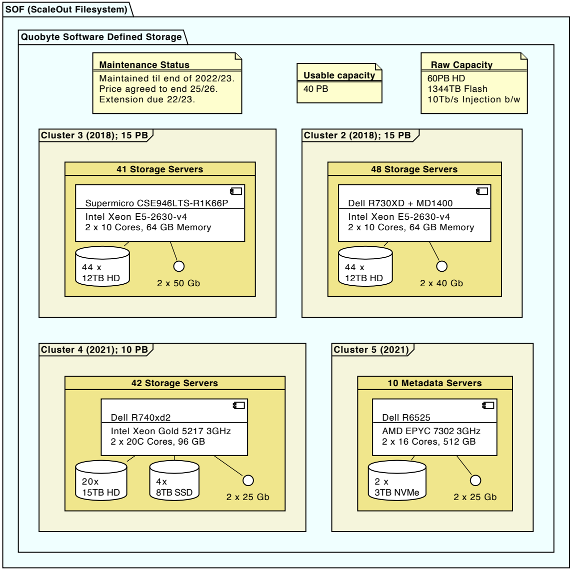

- Modern software defined storage, erasure encoded large files, duplicated small files.
- Latest generation includes SSD for small files and metadata server.
- 40 PB of usable storage includes 131 “compute nodes” in the storage subsystem — that is 3460 compute cores, which is a lot of compute…

Computational aka Active Storage

- Nearly all storage systems now have lots of compute, it's necessary for error correction, and managing which bytes go where, but it's basically underused.
- For a long time now, people have been talking about utilising some of that underused compute, by doing "some compute tasks in storage", something termed Active Storage or Computational Storage.
- Ideas have included shipping functions, shipping VMs or containers, and all sorts of things which have interesting security issues and significant application complexity.
  - If I ship a VM/container, how do I know it is not going to do bad things to the data and/or system?
  - How do I include complex systems in a workflow? How does this VM/container/function fit into my greater workflow?
- We decided there might be a simpler way to do some useful things by leveraging the dask graph and the idea of reductions.

A simple reduction workflow: mean

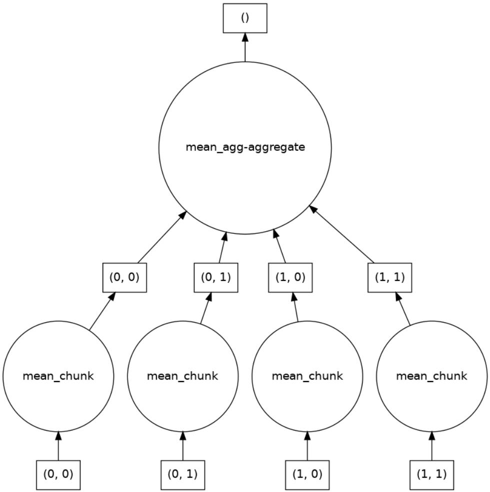

Taking a mean using four computational chunks:

- four lots of data re read from storage
- four means are taken
- means are aggregated result is calculated from the partial means
- requires reading all the data from the storage system, moving it into the compute node(s).

(see also Blelloch algorithm.)

One gory detail: the interface

Each computational chunk is operating on many storage chunks, but we need to know quite a lot about each storage chunk to do the calculations - knowledge held by the application, but not the storage, unless we tell it.

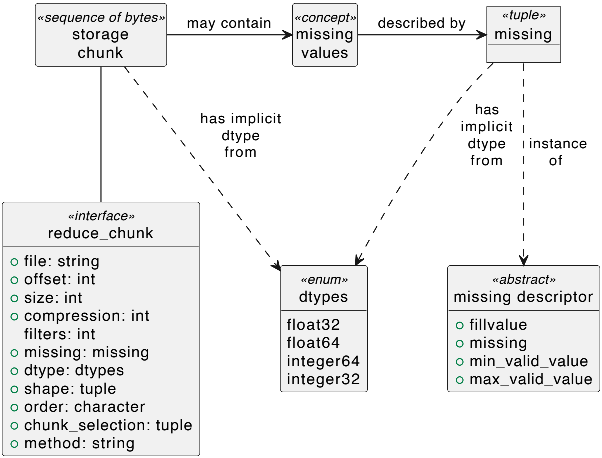

PyActiveStorage

Status Quo:

- We have built a python library that implements standard reductions using a “storage simulator”.
- We have two commercial companies building active storage which conforms to our interface (one for S3, and one for POSIX).
- We know how to “inject” active storage reads into a dask graph where-ever it would otherwise read an array and then immediately reduce it.
- Next step will be to integrate the real storage systems into our library, and benchmark it.
- Has real prospects to reduce the cost of some calculations

Reductions Supported

- Min
- Max
- Count
- Sum
- Mean
- More envisaged.

In-Flight Diagnostics

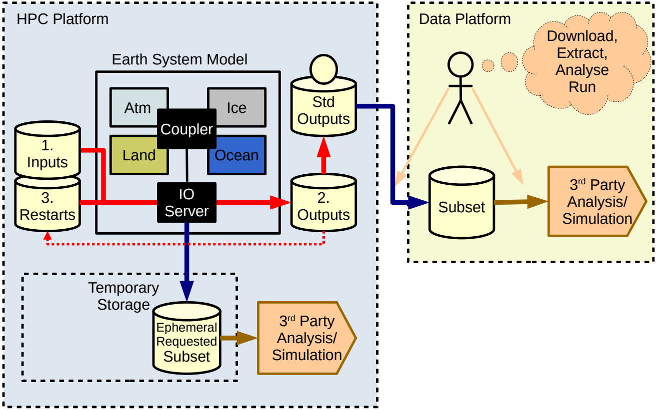

We're all well aware that we need to do things "in-flight" to avoid saving "everything".

Visualisation

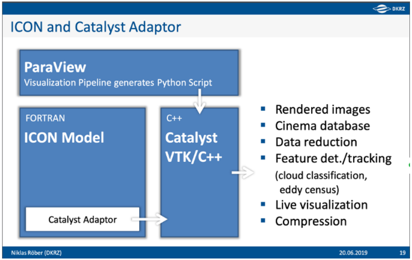

Röber, N., Böttinger, M., & Stevens, B. (2021). Visualization of Climate Science Simulation Data. IEEE Computer Graphics and Applications, 41(1), 42–48. https://doi.org/10.1109/MCG.2020.3043987

- Many analysis tasks require high frequency high resolution data which it will be impossible to store beyond some ephemeral state during the model run.
- Visualisation, Feature Tracking, Forced Models (Hydrology, ML etc)

Can we standardise the interface for third parties doing downstream models or visualisation?

Interesting ideas coming from the EU Destine project!

Ensemble Diagnostics

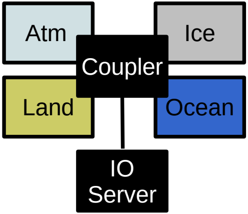

Use the IO server to control multiple ensemble members and carry out "in-flight" ensemble diagnostics.

Can do any kind of reductions "trivially", more complicated options coming with future versions of XIOS (our I/O server)

(It is not trivial to make all the various MPI communicators play nicely.)

Ensemble Diagnostics

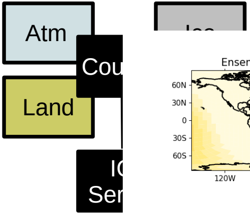

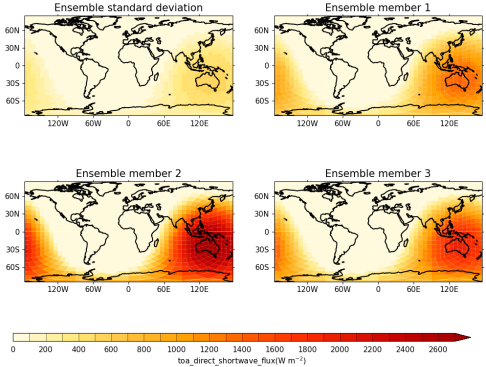

Use the IO server to control multiple ensemble members and carry out "in-flight" ensemble diagnostics.

Can do any kind of reductions "trivially", more complicated options coming with future versions of XIOS (our I/O server)

(It is not trivial to make all the various MPI communicators play nicely.)

Scaling Results on ARCHER2

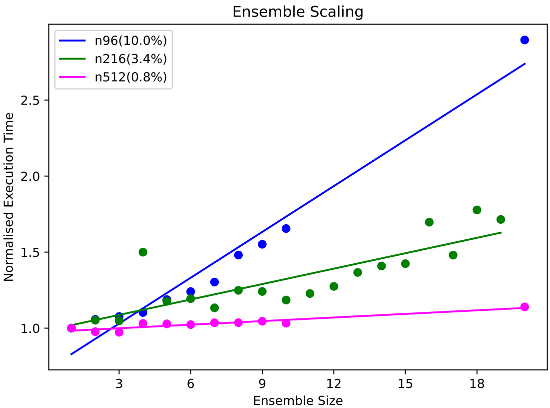

- Ensemble in-flight analysis becomes very cheap at high resolution and very valuable for large ensembles!
- Can consider in-flight ensemble output pruning (ie. avoid writing all ensemble members)

Scaling Results on ARCHER2

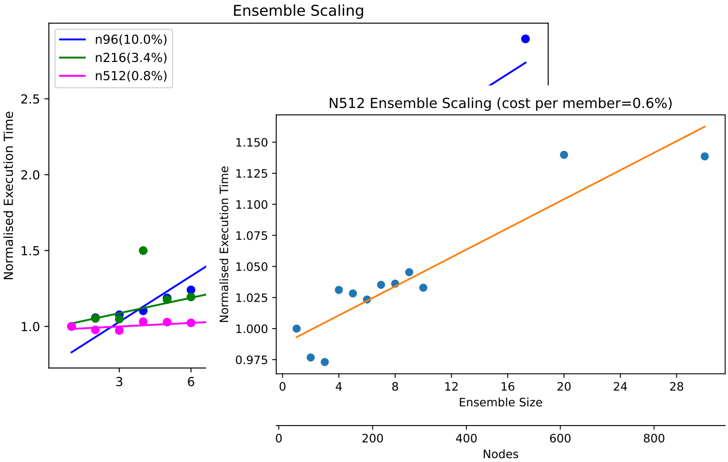

- Ensemble in-flight analysis becomes very cheap at high resolution and very valuable for large ensembles!
- Can consider in-flight ensemble output pruning (ie. avoid writing all ensemble members)

FAIR principles

FINDABLE	F1. **(meta)data are assigned a globally unique and persistent identifier** F2. data are described with **rich metadata** (defined by R1 below) F3. **metadata** clearly and explicitly include the **identifier of the data** it describes F4. (meta)data are **registered or indexed in a searchable resource**
ACCESSIBLE	A1. (meta)data are **retrievable by their identifier** using a **standardized** communications **protocol**, which is **open, free, and universally implementable**, and **allows for an authentication and authorization procedure**, where necessary A2. metadata are accessible, even when the data are no longer available
INTEROPERABLE	I1. (meta)data use a **formal, accessible, shared**, and broadly applicable language for knowledge representation. I2. (meta)data **use vocabularies** that follow FAIR principles I3. (meta)data include qualified references to other (meta)data
RE-USABLE	R1. meta(data) are richly described with a **plurality of accurate and relevant attributes** R1.1. (meta)data are released with a clear and accessible data usage license R1.2. (meta)data are associated with **detailed provenance** R1.3. (meta)data meet **domain-relevant community standards**

FAIR principles

FINDABLE	F1. **(meta)data are assigned a globally unique and persistent identifier** F2. data are described with **rich metadata** (defined by R1 below) F3. **metadata** clearly and explicitly include the **identifier of the data** it describes F4. (meta)data are **registered or indexed in a searchable resource**
ACCESSIBLE	A1. (meta)data are **retrievable by their identifier** using a **standardized** communications **protocol**, which is **open, free, and universally implementable**, and **allows for an authentication and authorization procedure**, where necessary A2. metadata are accessible, even when the data are no longer available
INTEROPERABLE	I1. (meta)data use a **formal, accessible, shared**, and broadly applicable language for knowledge representation. I2. (meta)data **use vocabularies** that follow FAIR principles I3. (meta)data include qualified references to other (meta)data
	R1. meta(data) are richly described with a **plurality of accurate and relevant attributes** R1.1. (meta)data are released with a clear and accessible data usage license R1.2. (meta)data are associated with **detailed provenance** R1.3. (meta)data meet **domain-relevant community standards**

R $\neq$ Reproducible!

Yes,but?

1. What is a dataset? (What is the granularity of a dataset?)
2. What should “searchable” mean?
3. What protocol was that? (The open, free, universally implementable one, with AAA.)
4. When and why should *“data be no longer available”?*
5. Which formal accessible (and useful) metadata schema was that?
   - 5.1 Which relevant metadata attributes?”
   - 5.2 What does provenance mean (for a simulation)?”
   - 5.3 Which domain-relevant community standards)?

What needs to be FAIR?

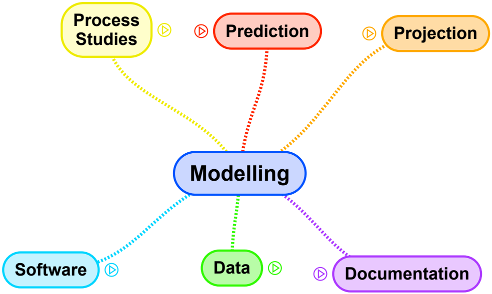

What needs to be FAIR?

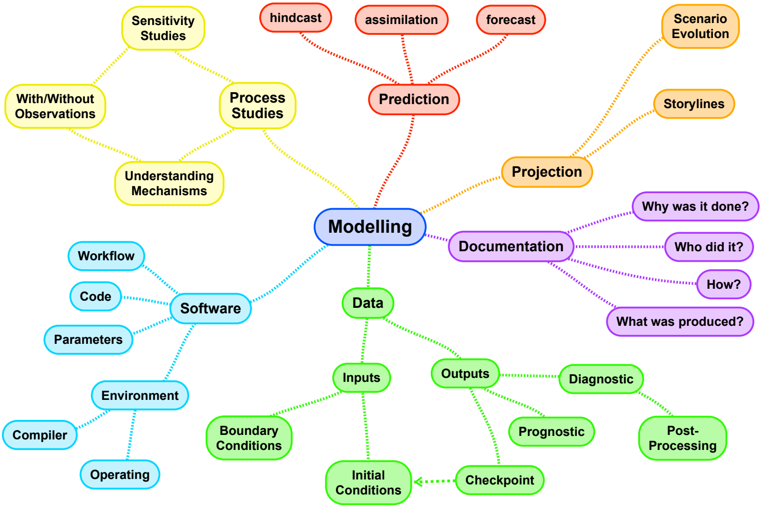

What needs to be FAIR?

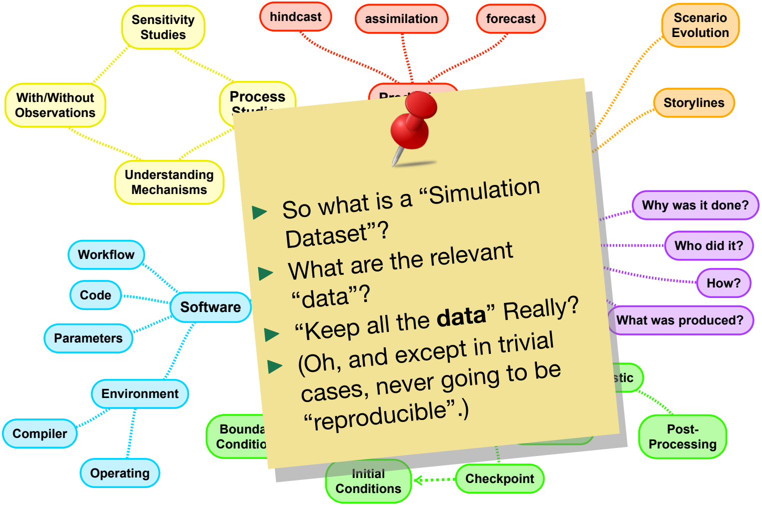

- So what is a “Simulation Dataset”?
- What are the relevant “data”?
- “Keep all the **data**” Really?
- (Oh, and except in trivial cases, never going to be “reproducible”.)

Simulation Context

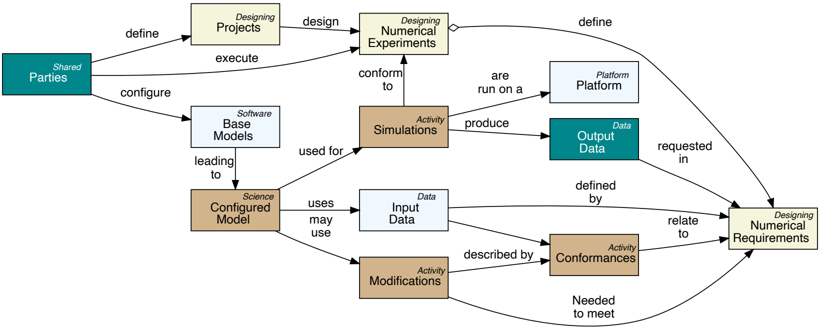

- Traditional Metadata covers only a tiny part of essential documentation!

Simulation Context

Why? What? How?

- Experiments may be devised by one group, and run by another, and analysed by a third.
- Participants are separated in space and time. Metadata crucial to information flow.

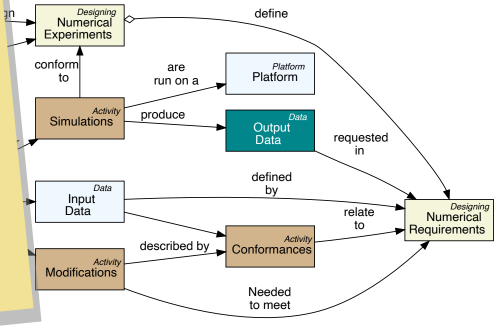

- Traditional Metadata covers only a tiny part of essential documentation!

Simulation Context

Why? What? How?

- Experiments may be devised by one group, and run by another, and analysed by a third.
- Participants are separated in space and time. Metadata crucial to information flow.

Reusable or Reproducible?

- Simulations are generally not reproducible.
- Experiments, if described properly are! So, **Reproducibility** requires Experiment metadata!
- **Reusable** requires all the metadata!

- Traditional Metadata covers only a tiny part of essential documentation!

Summary

- Climate simulation is big computing that generates lots of data.
- Doing collaborative things with data means thinking about metadata and standards.
- Doing things with metadata and standards stops our people wasting time reinventing analytical wheels.
- Sharing data by download isn't enough any more: move computation to the data!
- We are thinking about aggregation across tiered storage, about active storage, and about in-flight diagnostics.
- Oh, and you need to care about being FAIR.

This work arose from NERC national capability funding to the National Centre for Atmospheric Research over many years, as well as the following significant projects:

- The projects ESiWACE and ESiWACE2 have received funding from the European Union's Horizon 2020 research and innovation programme under grant agreements numbers **675191** and **823988**.
- The project IS-ENES3 recieved funding from the European Union's Horizon 2020 research and innovation programme under grant agreement number **824084**.
- The projects ExCALIWork and ExCALIStore have received funding from the UKRI and Met Office Strategic Programme Fund **Excalibur**: Exascale Computing: ALgorithms and Infrastructures Benefiting UK Research.