Met Office

Progress with integrating METplus into the operational and model development process at the Met Office

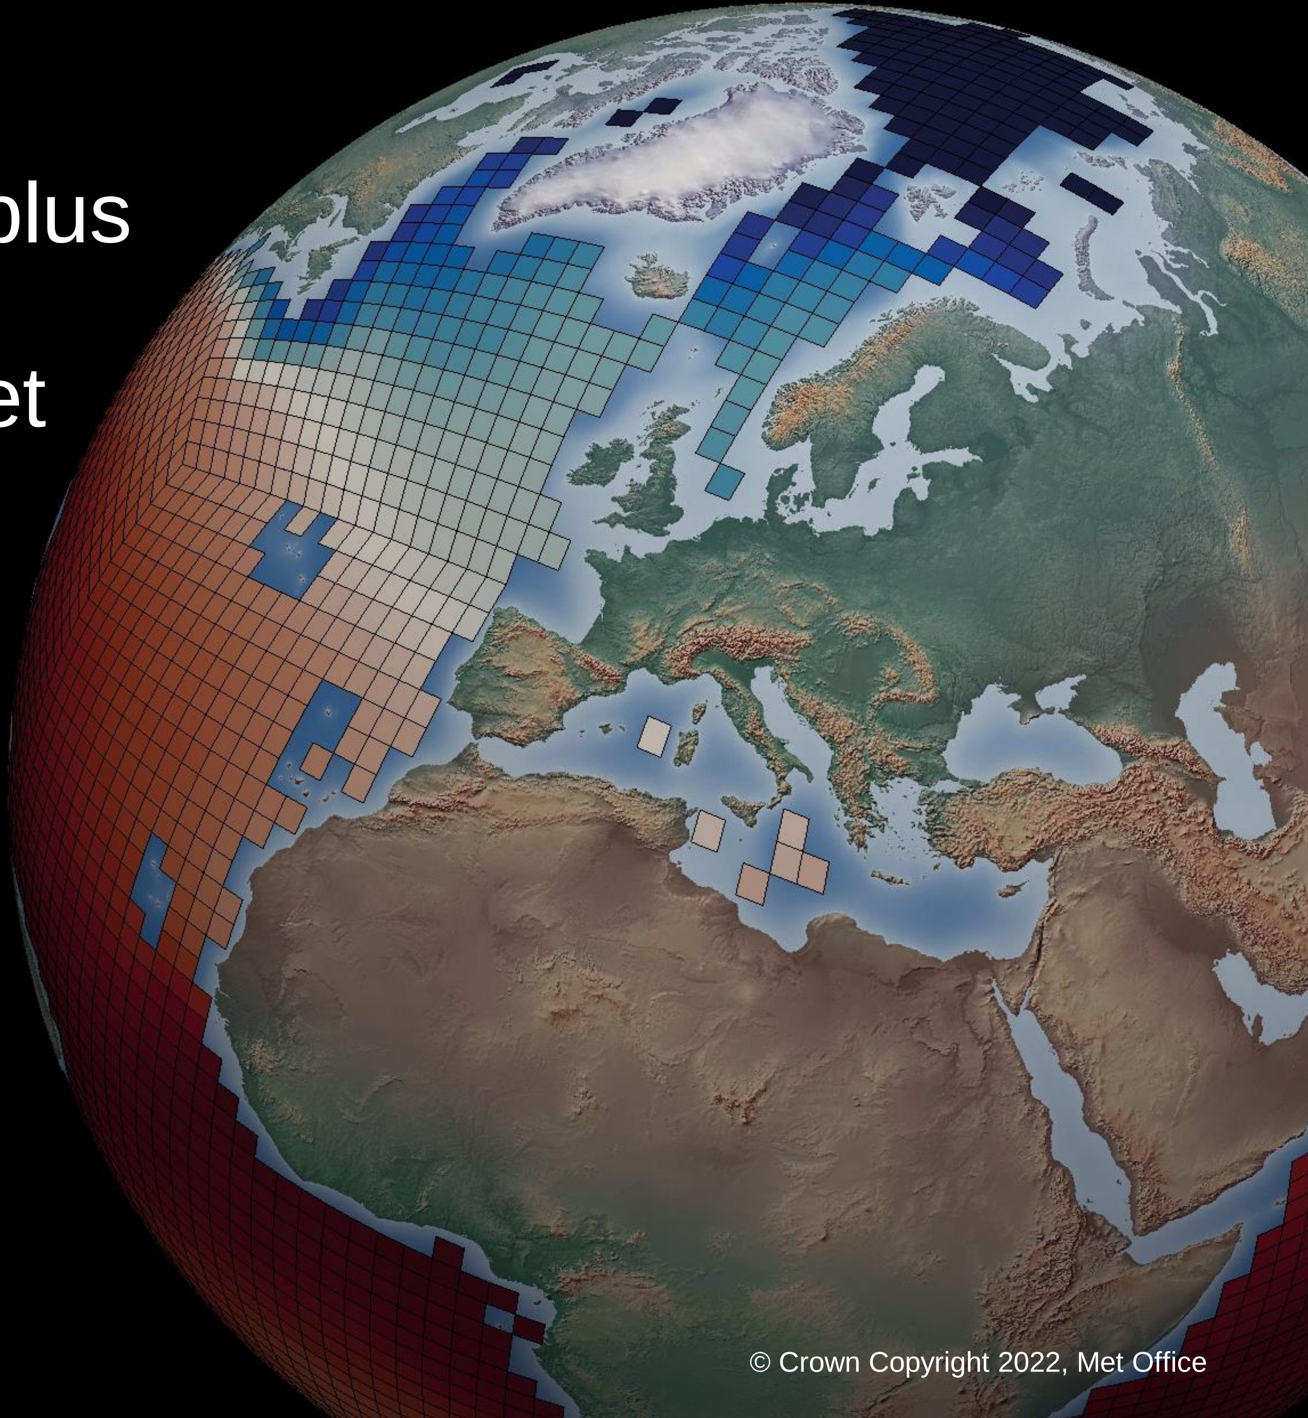

Marion Mittermaier

on behalf of the METplus implementation team

EWGLAM September 2022

Met Office

The boundaries are getting blurred….
(…. and increasingly unfamiliar??)

- I remain the **Science lead for the NG-Ver project**, which sits under the Next Generation Modelling System (NGMS) umbrella.
- The timeline for NG-Ver requires us to go live on the HPC+1 with **METplus** as **our operational system in autumn 2024 (with the UM) → that's the latest!**

- With the move to **km-scale *global* NWP** (we are just about to accept a 5 km prototype into the main development and evaluation stream) we need to **transition the tools and methods** we used for limited area models and convection-permitting NWP **to the global stage**. It is part of my "new" role to ensure that this transition happens.

Implementation into operations & research

Time scales not well aligned!

Rose suite

METplus use case configuration

Integration with results visualisation

How to keep the visual changes to our users to a minimum?!

NG-Ver: METplus

- Decision to adopt in 2019.
- Well on the way to implementation but the road to full operationalisation and use across Science will take many more years.
- Have a **contract with NCAR DTC** to assist with pertinent development and developed good bilateral working arrangements over last 18 months.
- **Observation ingest/provision**: ODB and JEDI. Multiple routes.
- **Optimal integration of METplus wrappers with Rose/cylc** (also on github) for scheduling on HPC
- We are also contributing code, e.g. the introduction of OpenMP to `grid_stat` by a Met Office software engineer has led to a **50% speed up**, which is of benefit to all METplus users.
- **Governance structure for METplus development** with representation from NOAA (ESRL, EMC, GFDL), NRL, USAF and Met Office to help guide development plans, develop protocols etc.

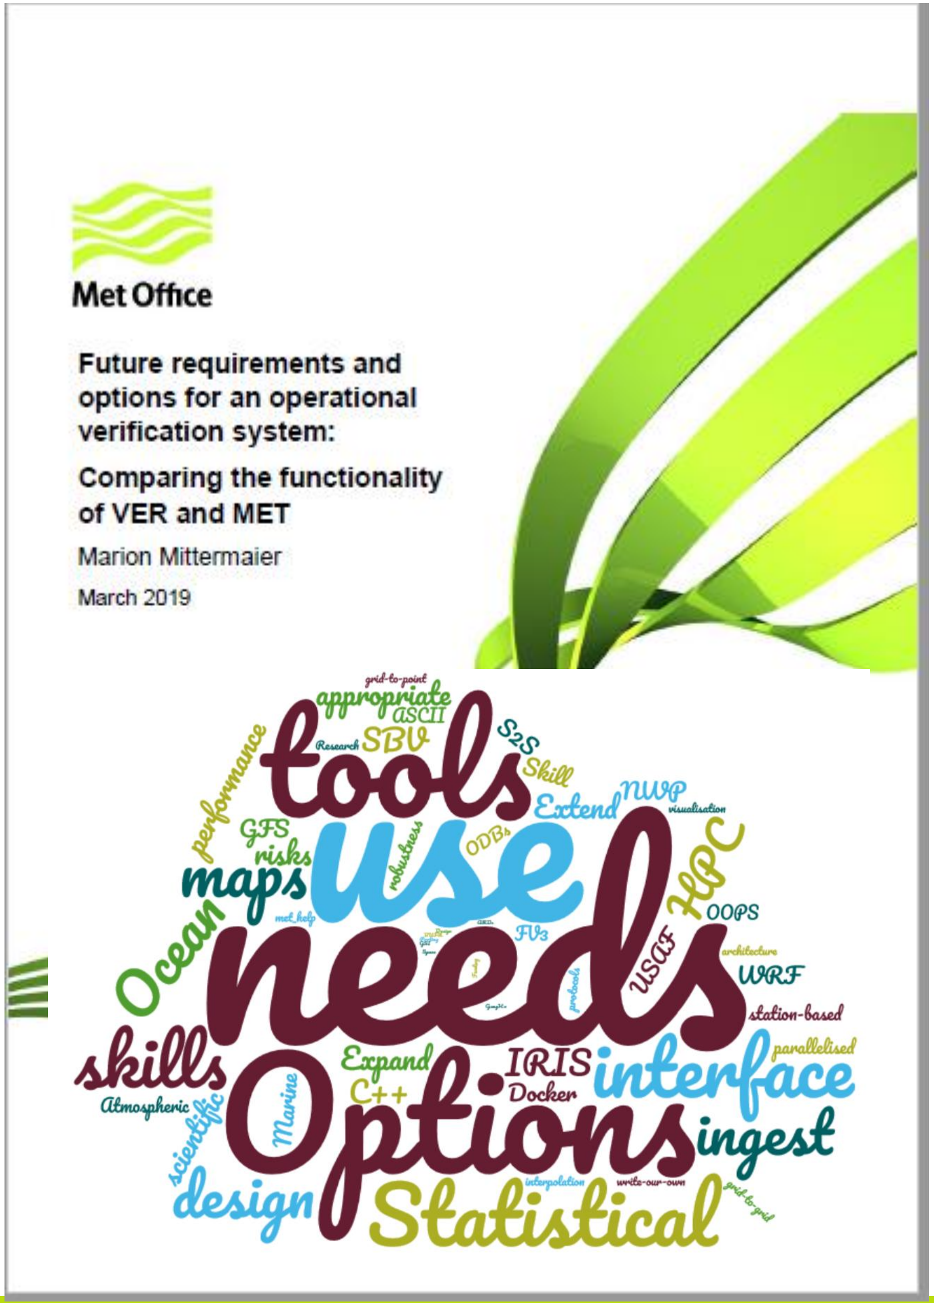

Word cloud created from the report

Met Office

Implementing METplus

1) METplus wrappers only

2) Ignore METplus wrapper

3) Incorporate with cylc workflows

4) Use for trials and development

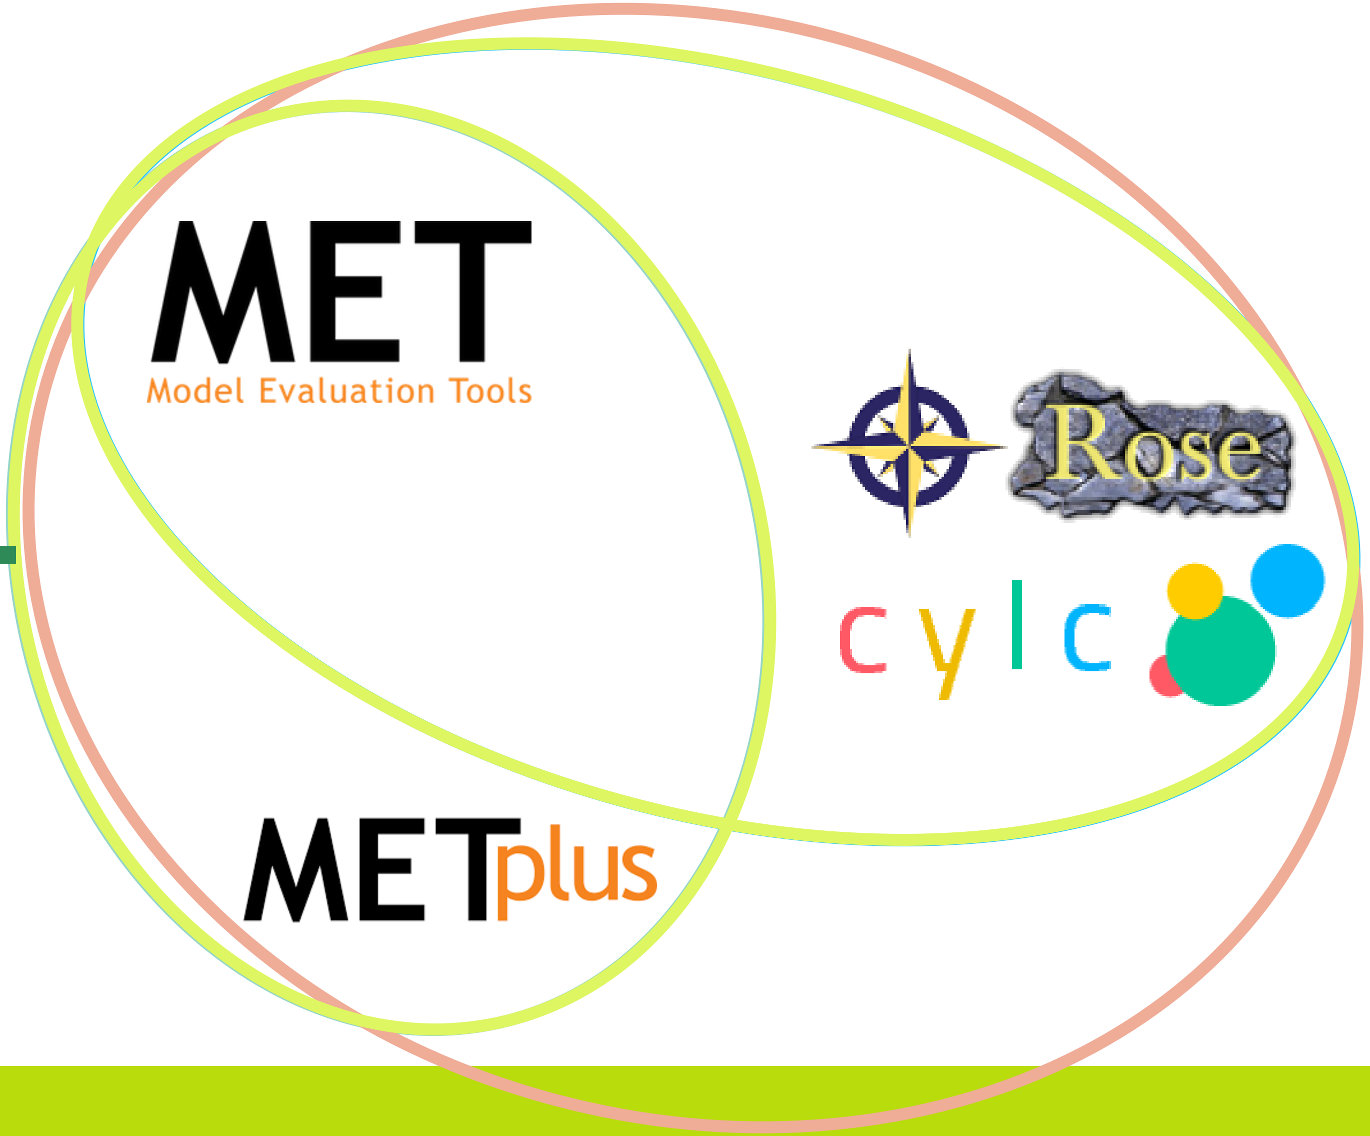

Met Office

rose-suite.conf

```
[jinja2:suite.rc]
####
# General settings
####
RUNAHEAD_LIMIT="P1D"
SCITOOLS_VERSION="scitools"
#WORKFLOWS options are 'POINTSTAT',
#of these
WORKFLOWS='GRIDSTAT'
#START_CYCLE and FINAL_CYCLE are gi
# WORKFLOW specific section. Choose
START_CYCLE="20210304T00Z"
FINAL_CYCLE="20210306T18Z"
```

app directory

Gridstat

run_gridstat.sh

user_system.conf

gridstat_variables.conf

```
# GridStat METplus Configuration

# section heading for [config] variables - all items below this line and
# before the next section heading correspond to the [config] section
[config]

# List of applications to run - only GridStat for this case
PROCESS_LIST = GridStat

# time looping - options are INIT, VALID, RETRO, and REALTIME
# If set to INIT or RETRO:
#   INIT_TIME_FMT, INIT_BEG, INIT_END, and INIT_INCREMENT must also be set
# If set to VALID or REALTIME:
#   VALID_TIME_FMT, VALID_BEG, VALID_END, and VALID_INCREMENT must also be set
LOOP_BY = INIT

# Format of INIT_BEG and INT_END using % items
# %Y = 4 digit year, %m = 2 digit month, %d = 2 digit day, etc.
# see www.strftime.org for more information
# %Y%m%d%H expands to YYYYMMDDHH
INIT_TIME_FMT = %Y%m%d%H

# Start time for METplus run - must match INIT_TIME_FMT
INIT_BEG=2005080700
```

Implementing system “MET” in VerPy

- MET has been added as a system (alongside VER and SBV) in VerPy to enable MET output to be ingested and plotted using our standard tools.
- MPR files can be ingested into Oracle station-based database (historical pair-wise database of global station and forecast data back to ~1997) but questions remain about aspects of database usage.

```
gm_opts = [
    # Analysis
    {'jobid' : 'anl',
     'system': 'MET',

     'source': MET_source + 'cnt.db',
     'type'  : 'cnt',
     'fcrs'  : [1200, 2400, 3600, 4800],
     'plottheme' : theme.

     'areas' : [2, 101, 102, 103],
     'truth' : 0,
     'params': scalar_param,


     'stats' : [6051, 6752]},
```

With this change the “front end” that scientists see doesn’t change but the “back end” has changed (from VER to METplus) https://www-nwp/~gltrials/GES/PS45_FinalPkg_DJF1920M/imt/page.html?DDV

NWP Science | Helen Titley (nee Watk... | Scientific Software Sta... | GrADS Documentation | VerPy 4.0 documentati... | METplus 4.0.0 docume... | Operational Suite Strat... | ParallelSuites – OPCH... | Other Bookmarks

ntroduction | Deterministic Summary | Deterministic Verification | Assimilation Stats Summary | Deterministic Assimilation Stats | Tropical Cylcone Verification | Surface Diagnostics | Ensemble Summary | Ensemble Verification | Ensemble Assimilation Stats | Documentation

Global Evaluation

% Difference (PS45 Final Pkg DJF19M vs. PS44 control) - overall 0.98%,
RMSE against observations for Equalized,
20191201 12:00 to 20200228 12:00

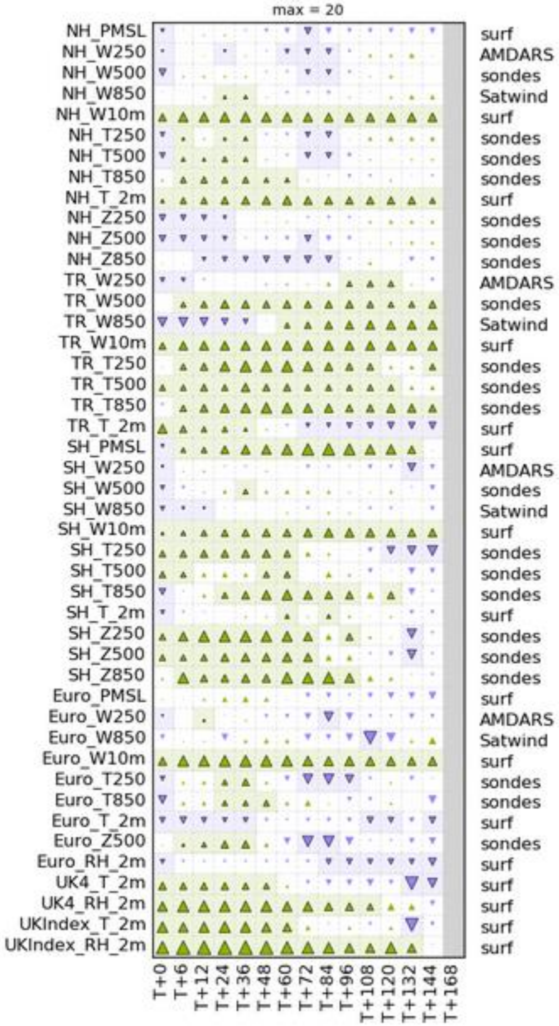

% Difference (PS45 Final Pkg DJF19M vs. PS44 control) - overall 0.09%,
RMSE against ownanal for Equalized,
20191201 12:00 to 20200228 12:00

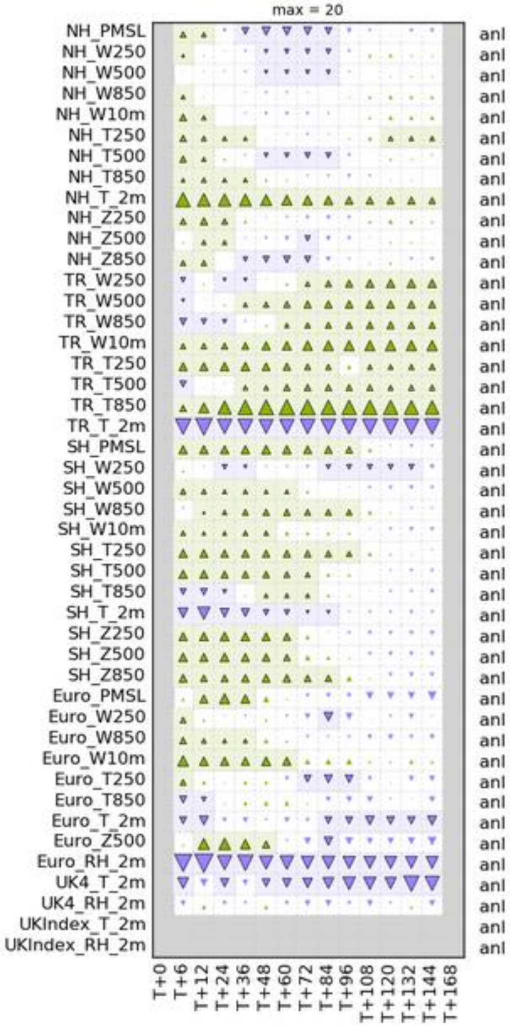

% Difference (PS45 Final Pkg DJF19M vs. PS44 control) - overall 0.47%,
RMSE against ecanal for Equalized,
20191202 00:00 to 20200228 12:00

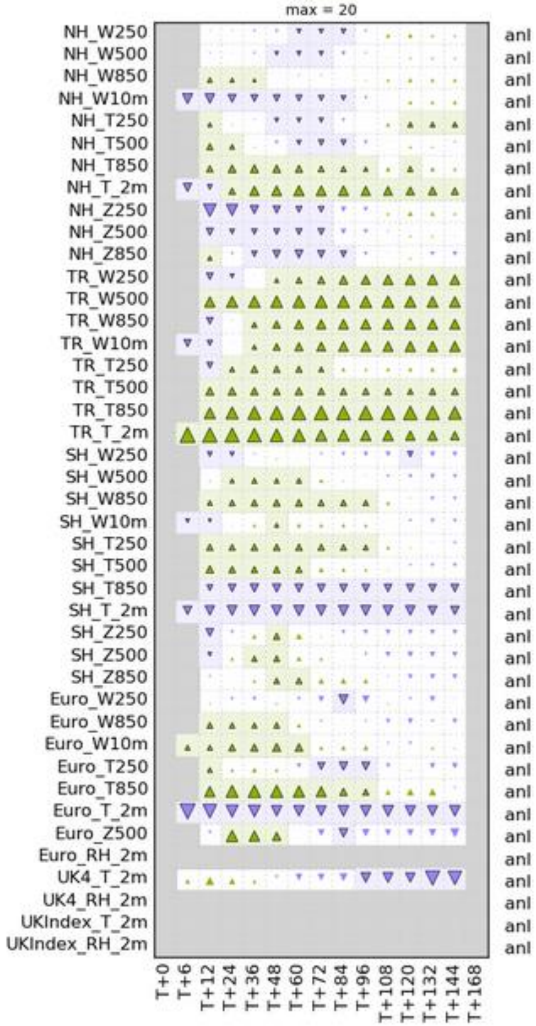

Met Office

Reconciling the old and the new

Harder task than one might think!

We need the results to agree with the old system, and be able to explain any instances where the two systems do not agree.

Fundamental differences have arisen between grib and netcdf; C++ and Fortran conventions etc → some of these are irreconcilable!

Aside: this reconciliation of results is only important for maintaining historical time series.
When comparing model A with model B with METplus these code/software differences have no impact.

IMPORTANT: these results highlight how difficult it is to inter-compare models when the results are NOT computed using the same software… and highlights the potential importance of community codes.

Using grid_stat & gen_vx_mask

gen_vx_mask boundaries

- Using grid_stat & gen_vx_mask to replicate the capabilities of the Met Office's current operational verification system (VER) in area based verification for WMO CBS areas.
- Masks generated in gen_vx_mask contained fewer grid points than VER counterparts using the same coordinates – contributing to a difference in RMSE of ~4%.
- To include the same number of grid points, an extension of 0.6˚ was applied. This was the minimum extension possible consistent on both 1.5 & 2.5 grids.
- **RMSE relative to VER still as much as 2%...**

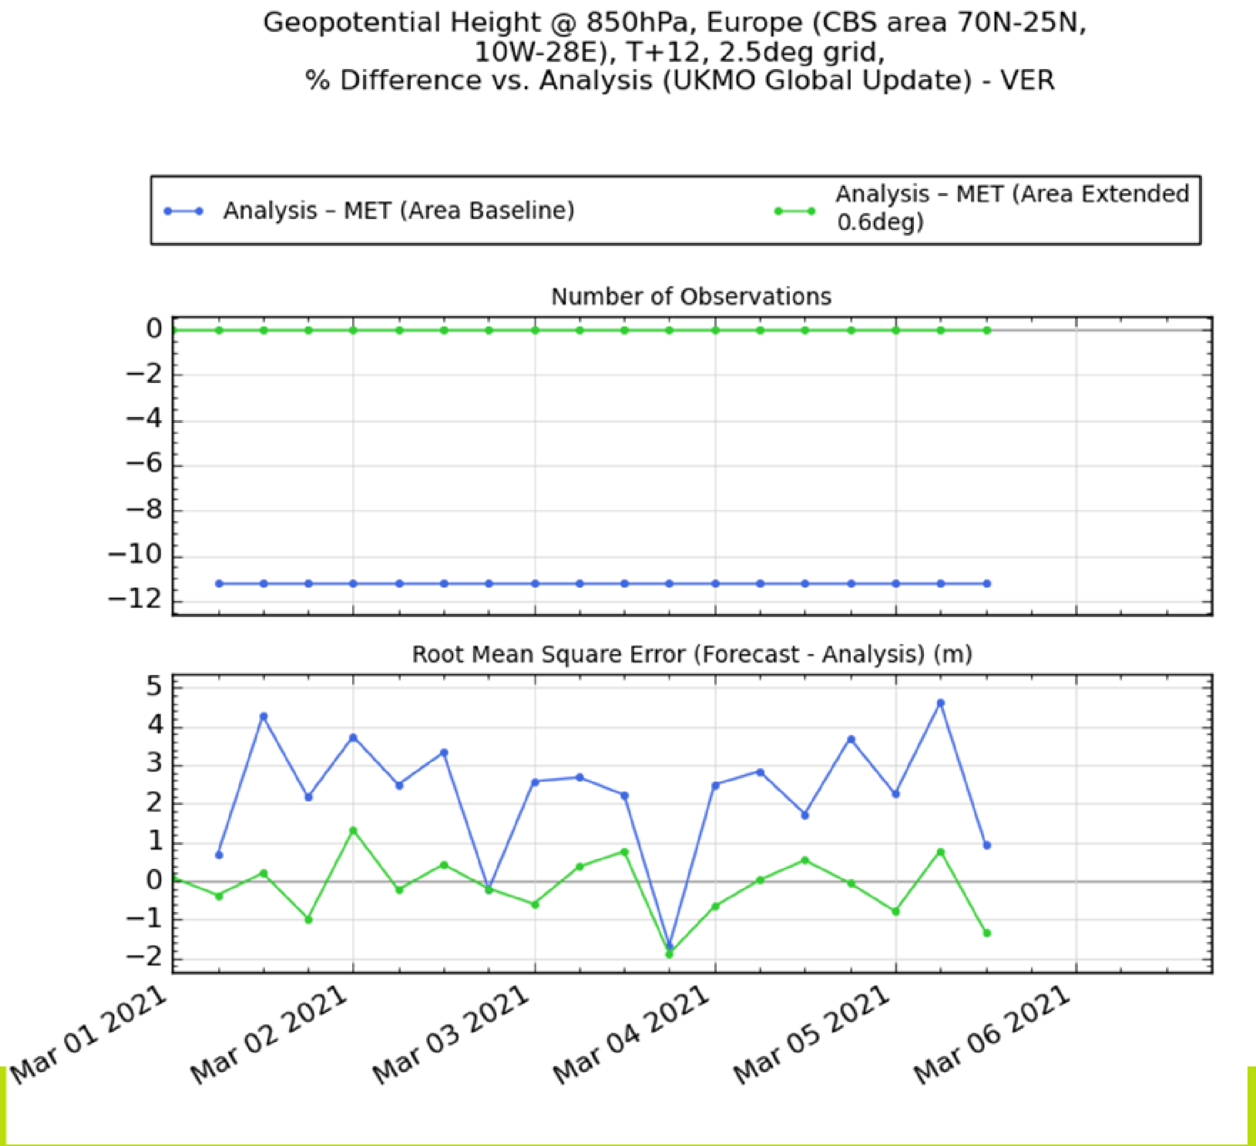

See Seb Cole's posters!

Using grid_stat & gen_vx_mask

Opposing orders of processing

- Looking at the how differences between MET and VER were distributed in space, errors were clustered around certain latitudes…
- 60˚ for example, interpolating using nearest neighbour method lead to the interpolation of different points in VER and MET to verification grid

Latitude				
59.859:	1579.0	1579.0	1580.0	1580.0
59.953:	1577.0	1578.0	1578.0	1579.0
60.047:	1575.0	1576.0	1576.0	1577.0
60.141:	1573.0	1574.0	1574.0	1575.0

Geopotential heights at 850hPa on Unified Model grid

- Manipulating VER to run S to N as MET does accounts for the 2% RMSE.

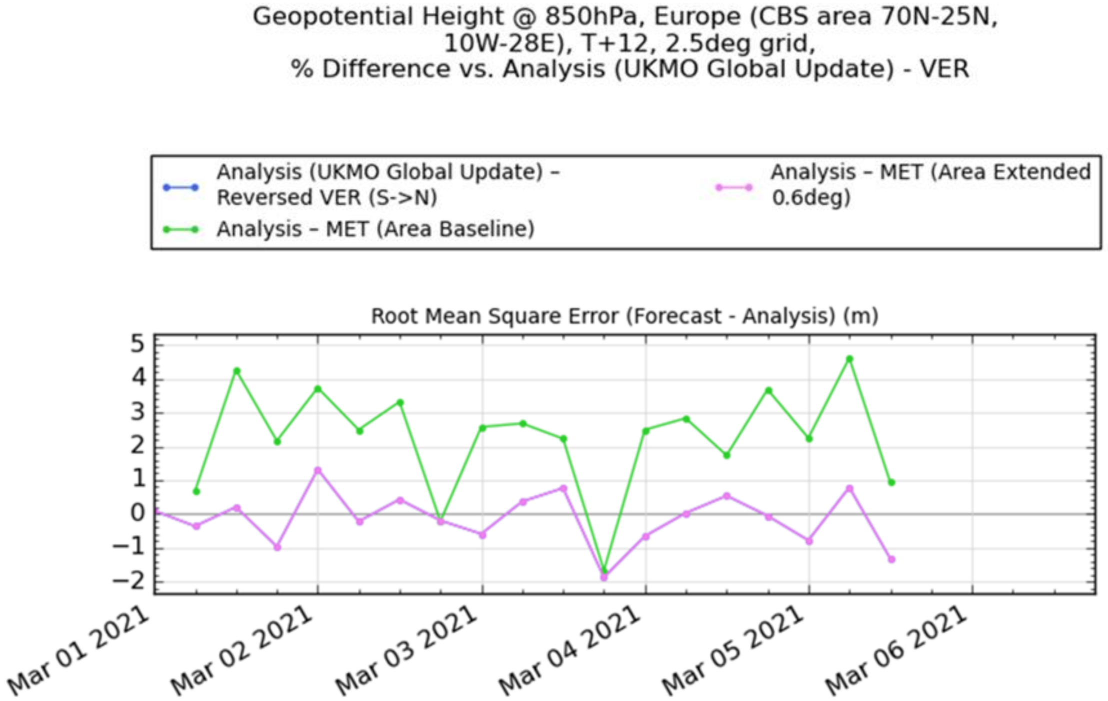

(Blue Reversed VER line beneath pink MET Area Extended)

On the operational implementation of MET:
- Verification stats will not be directly comparable with VER.
- Proceed with extended areas to reduce the differences seen?

See Seb Cole's posters!

A METplus use case example

Illustrate how different tools are run together (as the code is a lot more modular):

- Regrid-Data-Plane used to ensure all models for comparison and observations on the same grid
- PCPCombine used to create observation dataset from half-hourly GPM IMERG snapshots.
- GridStat to compute the FSS with FBIAS==1
- Plot in VerPy

Data pre-processing (forecast data)

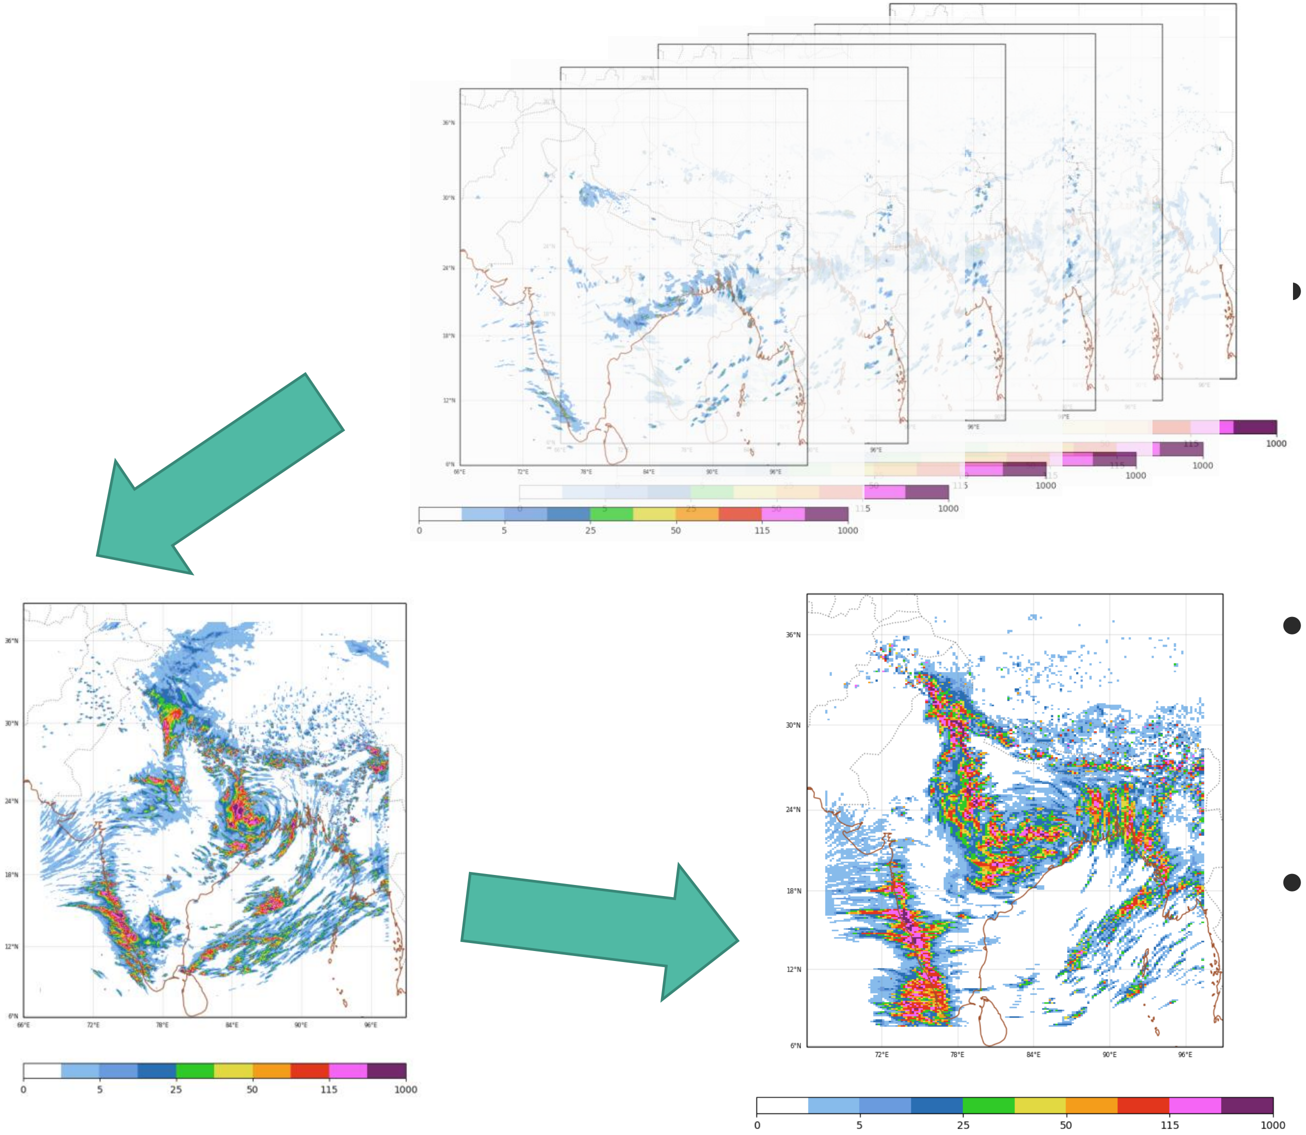

- PCPCombine used to create 24-hour accumulations from hourly rainfall amounts for one of the forecast models
- Data then regridded (Regrid-Data-Plane) as required for pushing through Grid-Stat
- Would like to attempt this with METplus!

Met Office

Data pre-processing (observations)

- Regrid-Data-Plane
  used to get observations and model onto the same grid

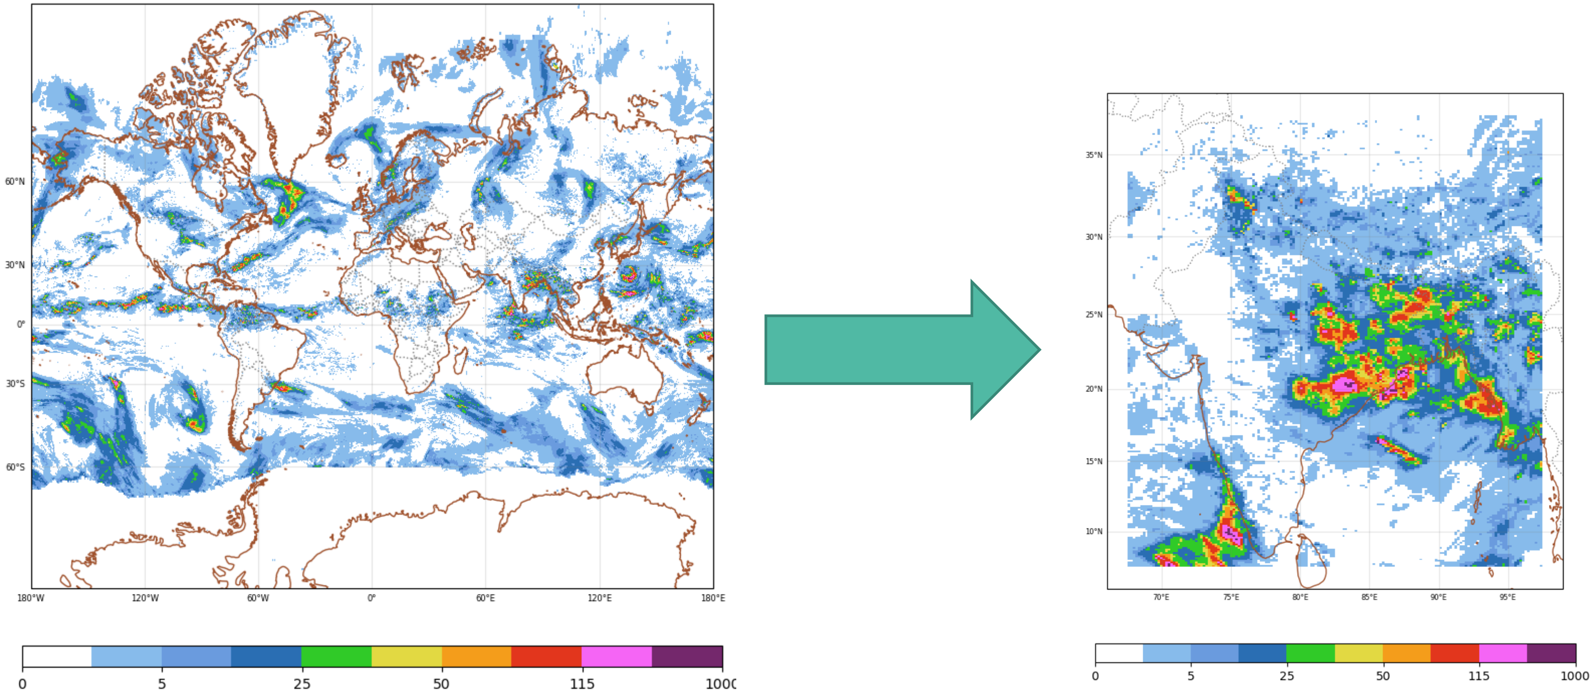

Grid-Stat to produce required verification statistics

- Four months, 1 initialisation time per day, 3 lead times, 4 thresholds plus FBIAS==1
- Timing information initially from cylc workflow
  - Used to test parallelisation improvements to Grid-Stat
- METplus wrapper set-up subsequently

Model 1

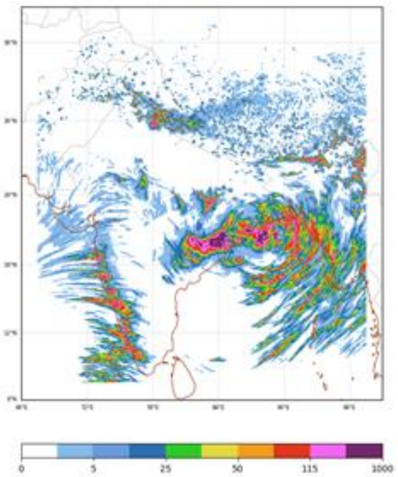

Model 2

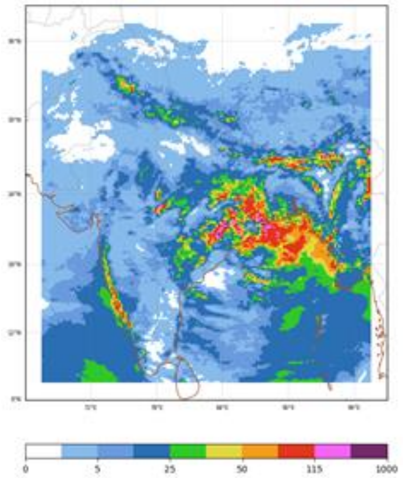

Forecast exceeds threshold

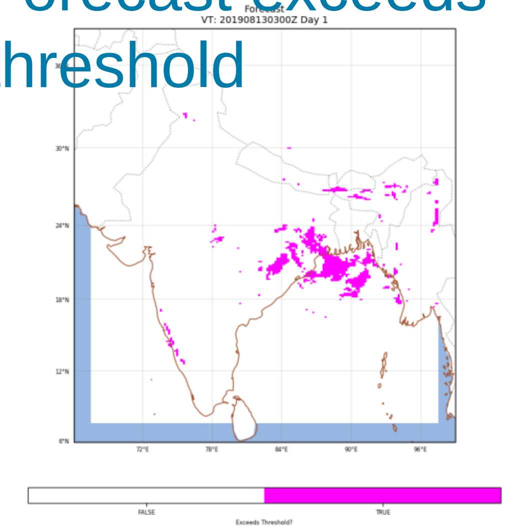

Observations

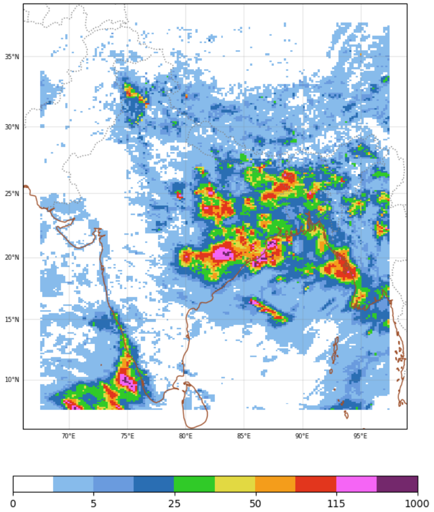

Observations exceed threshold

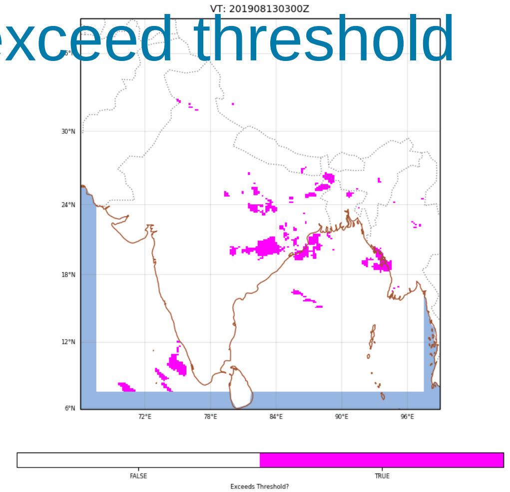

Rachel North

Assessment of daily precipitation with Fractions Skill Score

- Evaluate Day 1,2 and 3 forecasts from control members of two (different resolution) ensemble systems
- Day 1 example shown here
- Using **FBIAS==1** threshold setting in GridStat
- Several different observed thresholds
- Different biases in different scale ensembles

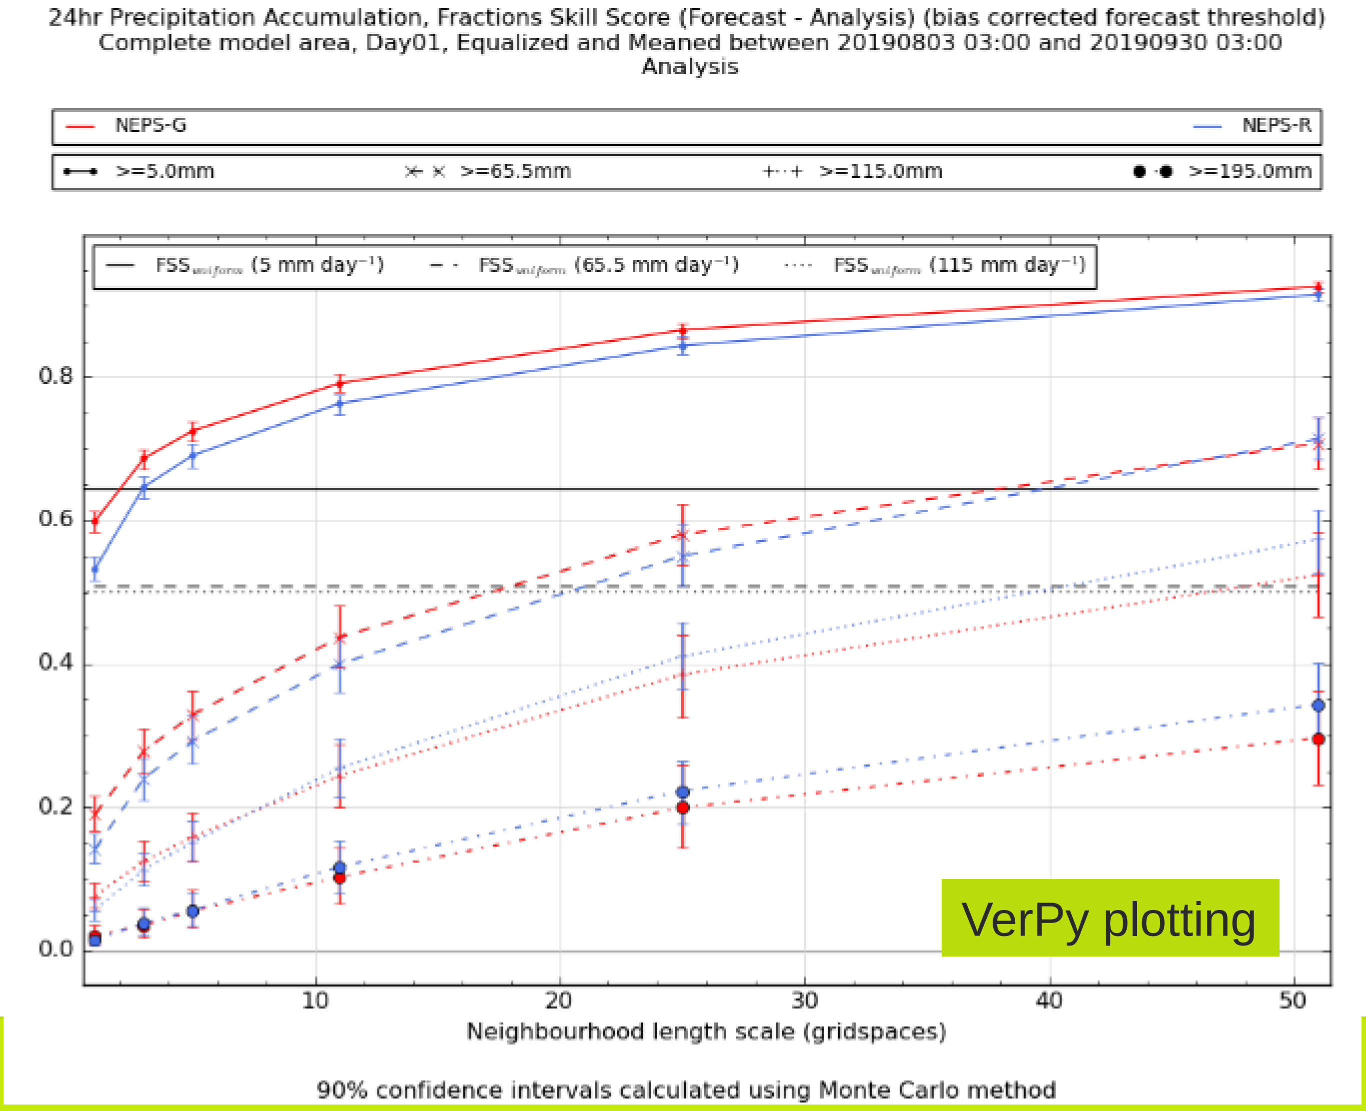

Rachel North

Met Office

To date we have….

.. extended the prototype Rose suite w
into a **cycling proto-verification suite**
model verification against analysis th
(though it is not yet running on HPC).

This includes the **pp to netcdf convers**
has not been implemented yet. Four r
incrementally added.

The global model verification against
forecasts **4 times a day, processing 76**
cycle over 25 different areas as well a

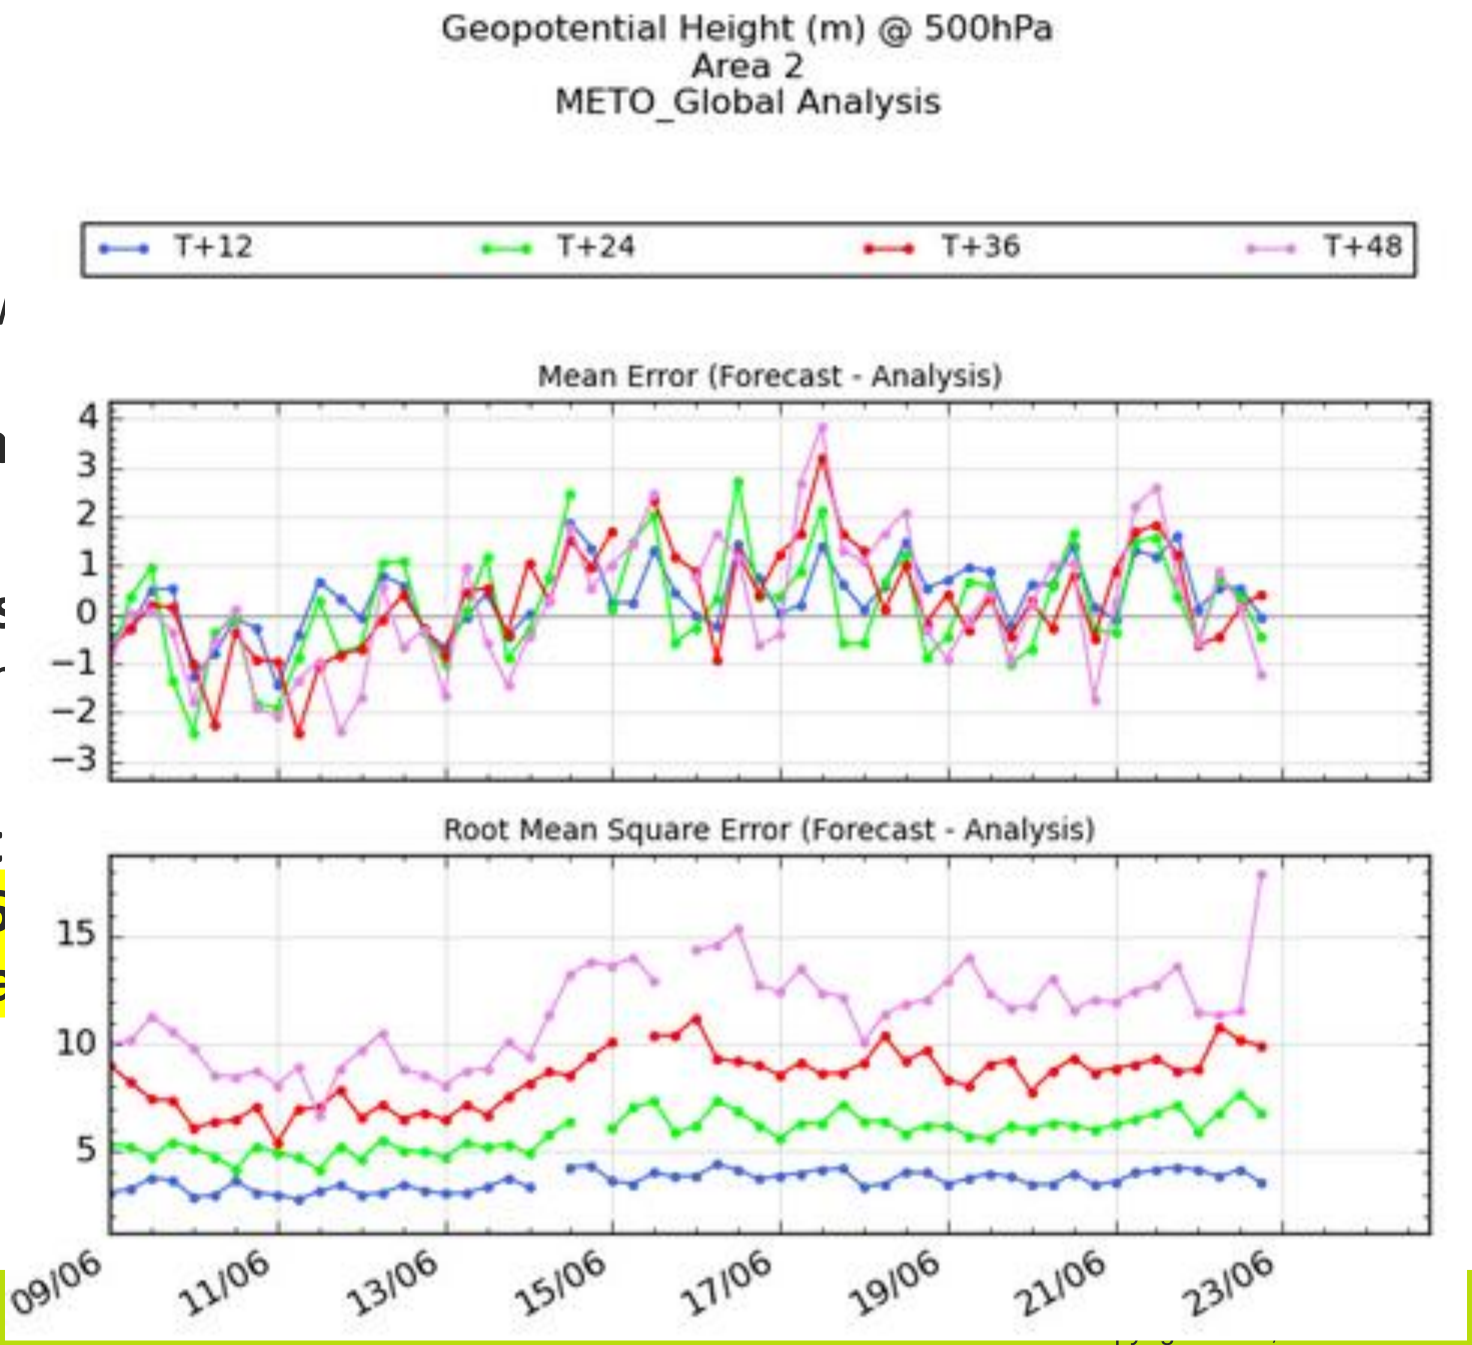

Rob Darvell

We have also…

- Support for rotated grids
- Tested HiRA for deterministic
- Tested ensemble_stat vs analyses

But supporting UK models and observations is still at the early stage.

What the future holds…

KSCALE → the need for speed

LFRic → a “rocky” road….

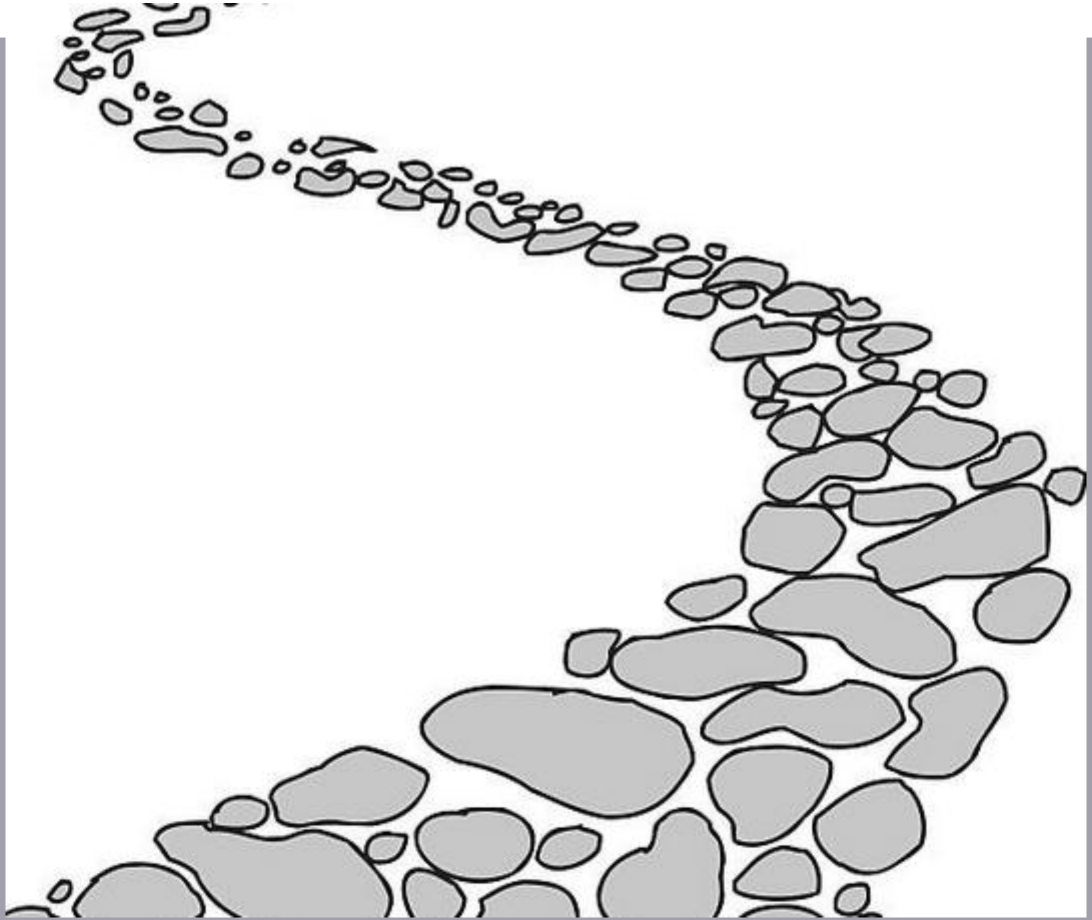

K-Scale – A brief overview

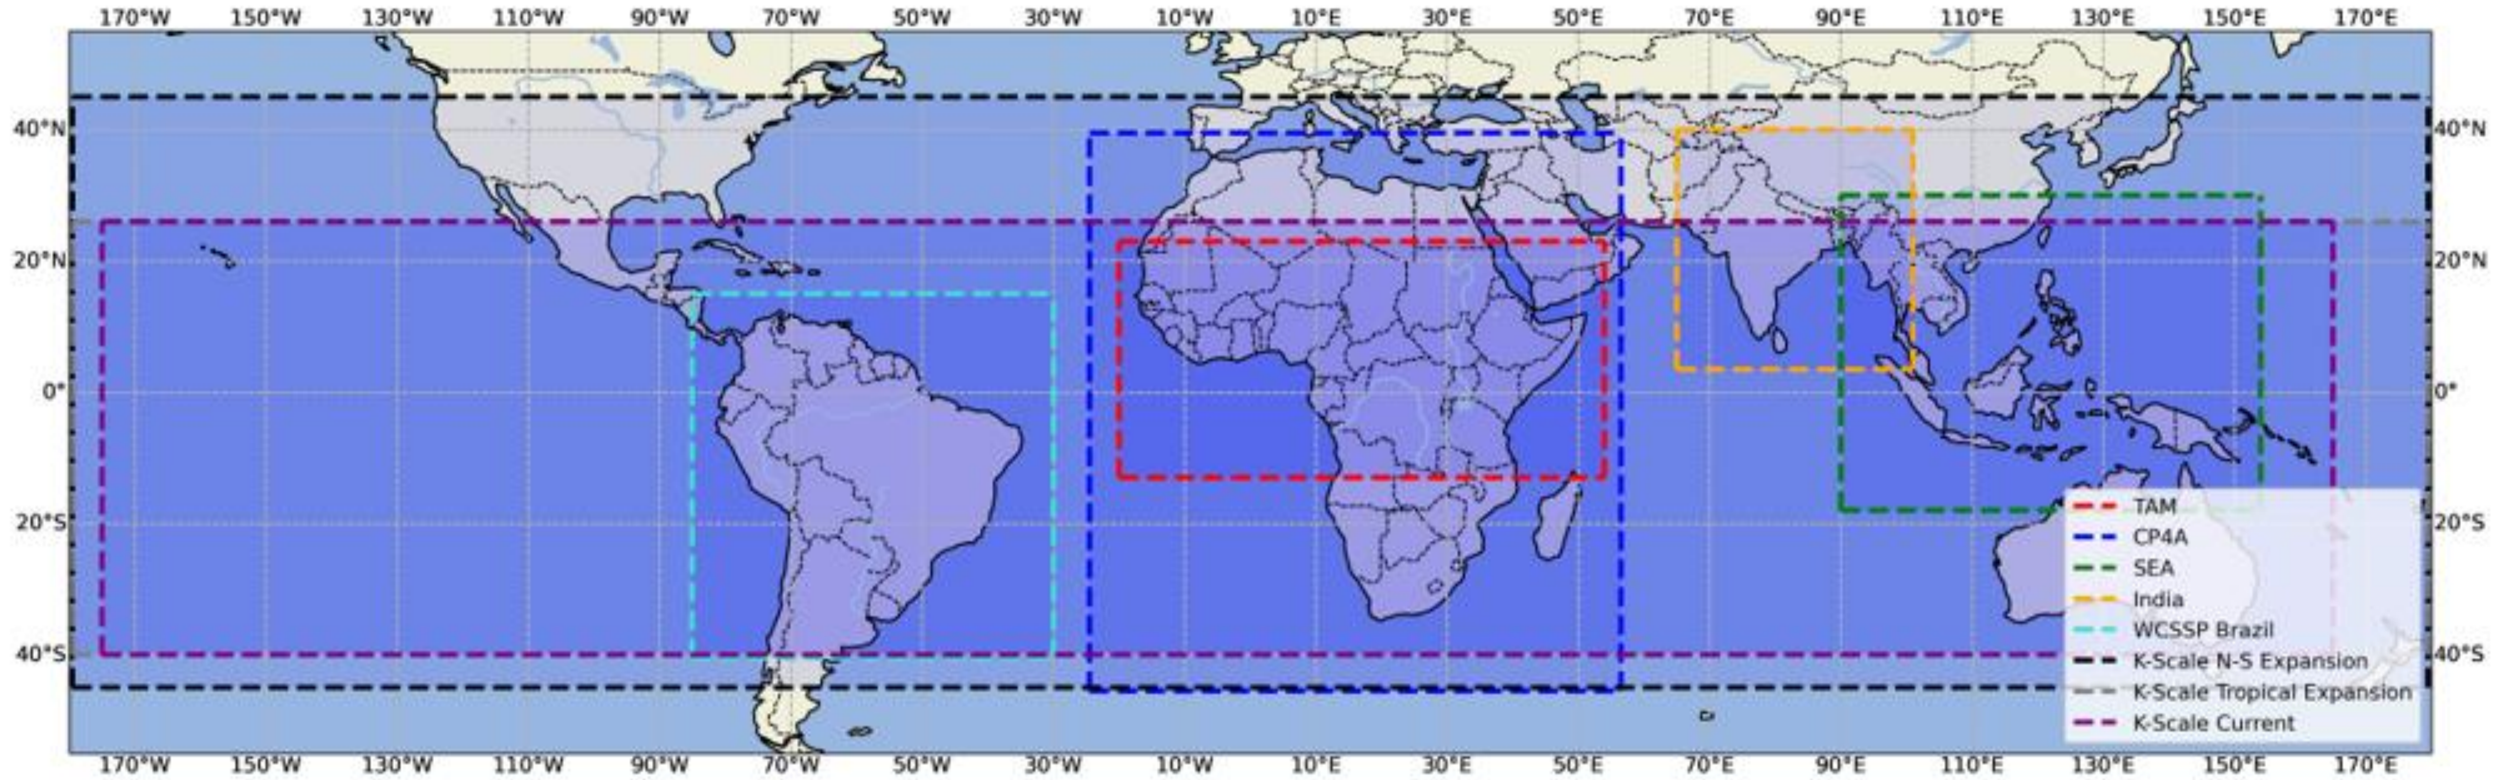

- K-Scale is part of the path to high resolution theme - one of the key aims is to develop global convection permitting atmospheric models that are coupled to eddy-resolving ocean models.
- The current 2.2 km K-Scale domain (purple outline) is 17000 by 3300 grid points.
- Four different initialisation times all currently using the same set of fixed OSTIA SSTs

Acknowledgement: James Warner

A LOT more to do…

Met Office

Contribution: OpenMP in grid_stat for fractional statistics

Code: Maff Glover   Testing: Rachel North

TIME	ORIGINAL	OPTIMISED	
		OMP_NUM_THREADS=4	OMP_NUM_THREADS=1
Start	Tue 2 Nov 17:57:24 GMT 2021	Mon 8 Nov 17:23:16 GMT 2021	Mon 8 Nov 17:47:27 GMT 2021
End cycle 1	Tue 2 Nov 18:14:36 GMT 2021	Mon 8 Nov 17:32:00 GMT 2021	Mon 8 Nov 17:56:24 GMT 2021
FIRST VALID TIME	17m 12s	8m 44s	8m 57s
End	Tue 2 Nov 18:30:41 GMT 2021	Mon 8 Nov 17:41:07 GMT 2021	Mon 8 Nov 18:05:17 GMT 2021
LAST VALID TIME	16m 5s	9m 7s	8m 53s
FULL RUN	**33m 17s**	**17m 51s**	17m 50s

~50% speedup

The Met Office cube-sphere revolution

1. LFRic and global km-scale changes everything.
2. To do model analysis **we need tools that can do what we need and can cope**: e.g. Iris and METplus
3. **Tools don't develop or maintain themselves**.
4. Community code development presents many opportunities (and a few challenges)

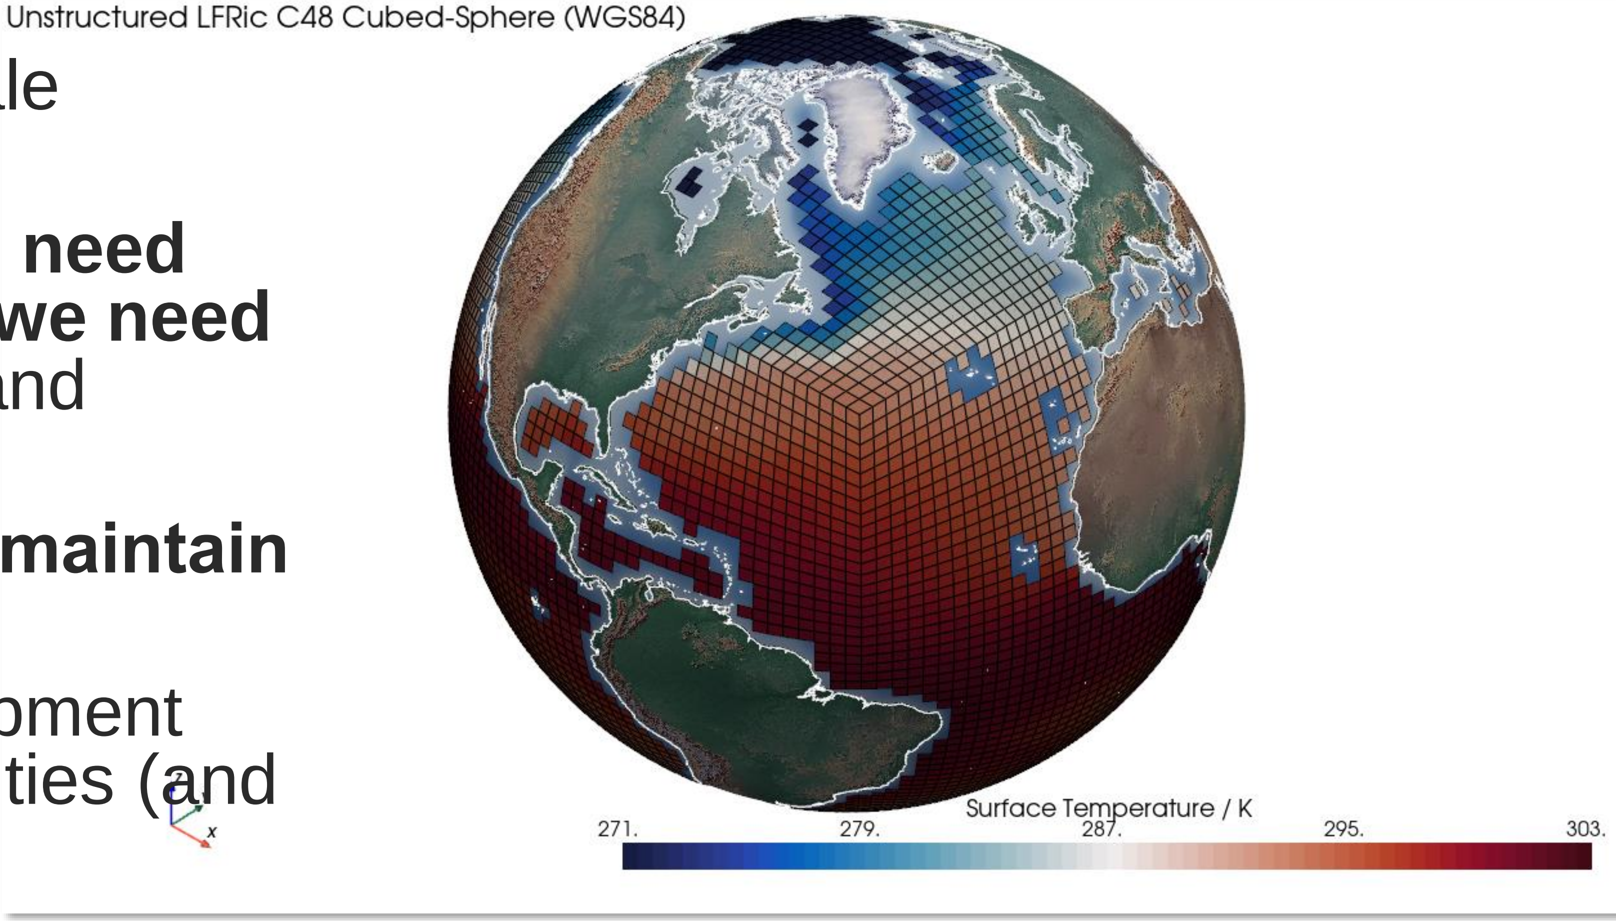

LFRic C48 unstructured cubed-sphere loaded with Iris and rendered with GeoVista – land mask cells have been removed to reveal the texture mapped base layer

With 1:10m Natural Earth coastlines and a 1:50m Natural Earth cross-blended hypsometric tints base layer with shaded relief and water

Projected onto a WGS84 (Geographic) S2 sphere

Met Office

Unstructured Meshes: How?

- **Nodes** from individual coords
- **Edges/faces** from node **Connectivity** (using indices)
- Data on each node / edge / face

- Nodes / edges / faces need not align
- Can mix 3- / 4- / 5- / n-sided faces

All this “richness” comes with a cost… typically metadata is 4x greater than actual data → vastly increased file sizes with impact on I/O.

	X	Y
node0	19.2	8.1
node1	19.9	11.0
	…	…

faces:	data:
2 0 1 3	8.598
4 2 3 -	20.43

Data on faces:

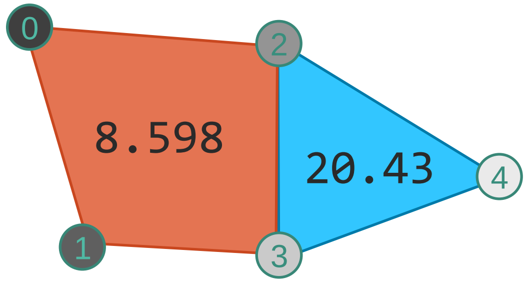

Data on nodes:

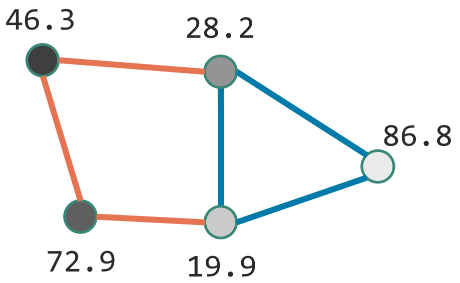

data:
46.3
72.9
28.2
19.9
86.8

Met Office

Unstructured Meshes: Bigger Files

How many numbers in a **~5 million** cell dataset?

UM 2560 × 1920 grid	
Data	4,900,000
Geospatial Info	
Point	4,480
Bound	8,960
Geospatial Total	**13,400**

LFRic C905 cube-sphere mesh	
Data	4,900,000
Geospatial Info	
Node	9,800,000
Face-Node Connectivity	19,700,000
Face Centre	9,800,000
Geospatial Total	**39,300,000**

6 seconds
Reading a leaflet

40 minutes
Reading a newspaper

5 hours
Reading a book

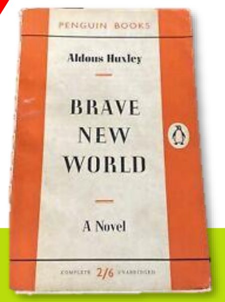

Big numbers! Here’s a scale analogy…

Met Office

Unstructured Meshes: Terminology

Grids

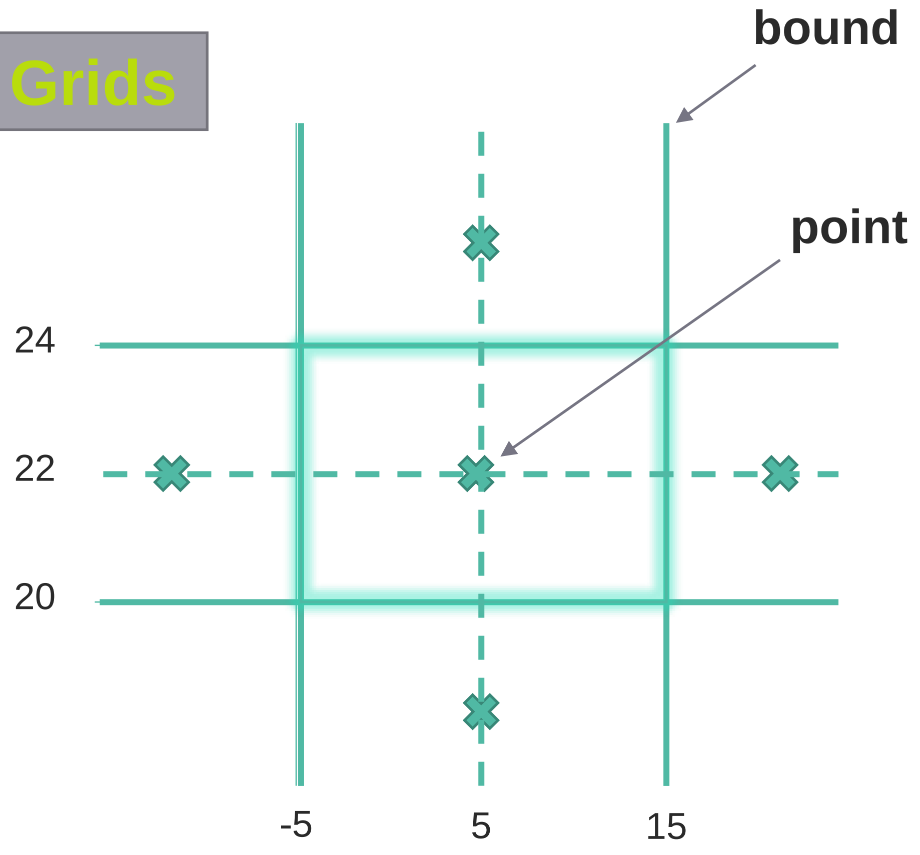

With only 4 pieces of information you can derive the entire grid

Meshes

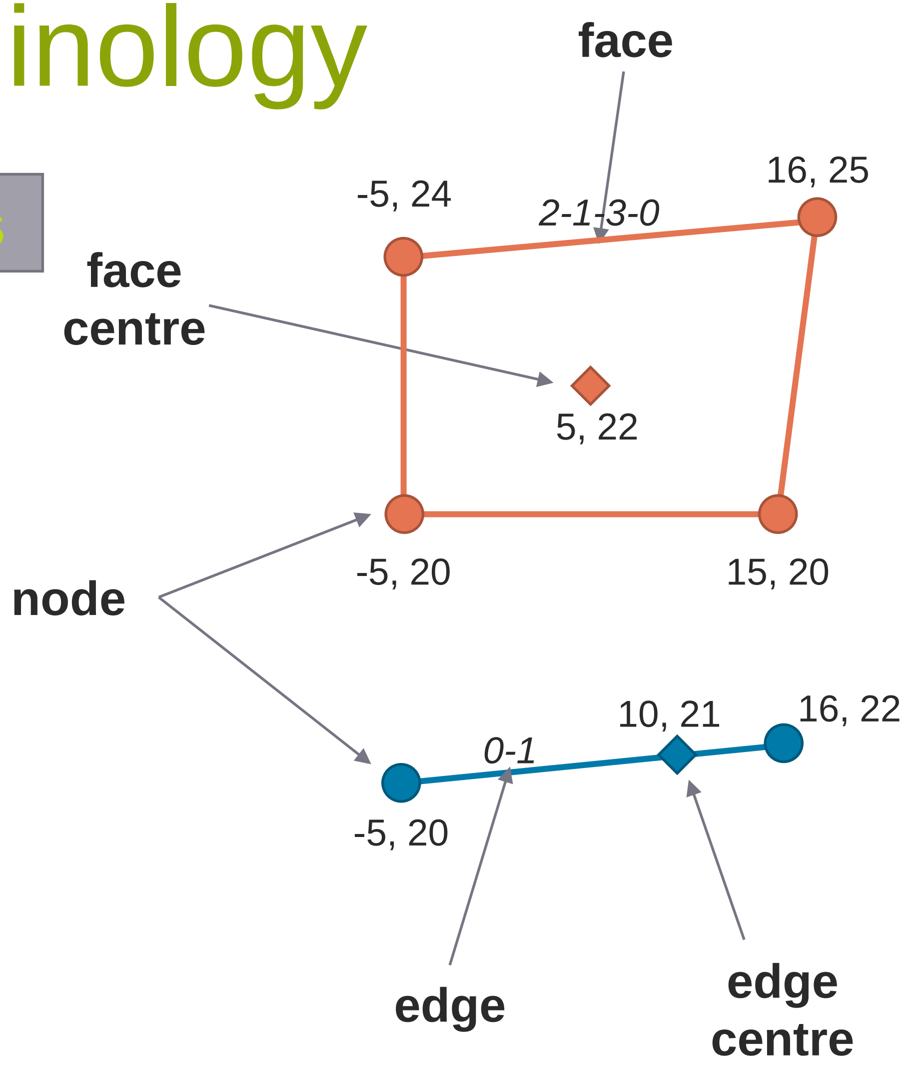

Bill Little, AVD team

faces are made up of **nodes** (not edges)

NG-VAT: Iris UGRID Support

Iris 3.2

- Has the ability to read and write UGRID netcdf under the conda environment
- Includes the capability to convert from UGRID → regular → UGRID
- This can and will form the basis of all diagnostics developed on the native UGRID
- The code also works for ocean grids (e.g. tripolar) grids and provides a comprehensive set of basic tools for manipulating model output on non-rectilinear native grids (that look and behave similarly to rectangular grids.)

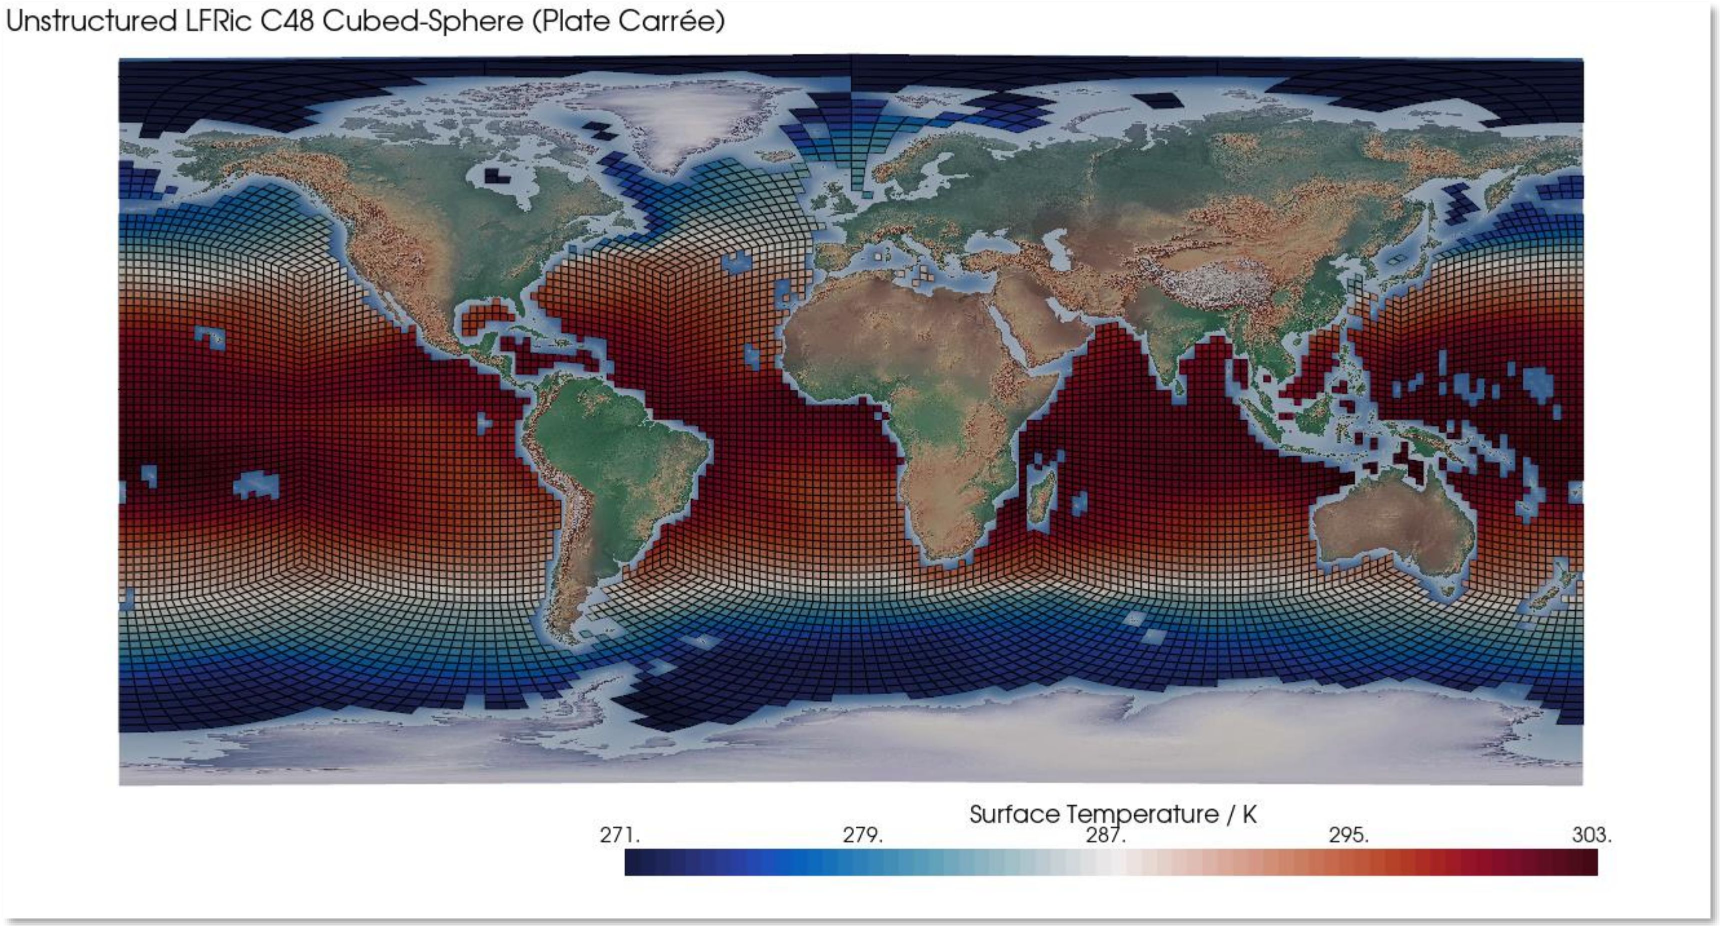

“Geovista” developed to **enable the visualisation of UGRID** since the standard python libraries such as matplotlib and cartopy (in particular) don’t work well with UGRIDs (don’t scale)

We envisage Iris and METplus to be used together for diagnosing model errors and evaluation.

Met Office

Iris

Unstructured Tools: Iris – Load from UGRID Files

```
>>> # Get a sample datafile with a C48 cubesphere mesh.
>>> file_path = "lfric_surface_mean.nc"
>>> with PARSE_UGRID_ON_LOAD.context():
...     unstruct_cube = iris.load_cube(file_path, “rainfall_flux” )
...
>>> print(unstruct_cube)
rainfall_flux / (kg m-2 s-1)     (time : 12; -- : 13824)
    Dimension coordinates:
        time                            x       -
    Mesh coordinates:
        latitude                        -       x
        longitude                       -       x
    Attributes:
        Conventions                 'UGRID'
        description                 'Created by xios’
        ...
```

Why we do need native-grid verification!

Sticking points

Unstructured grids

Grids grids and more grids!

Observations ingest

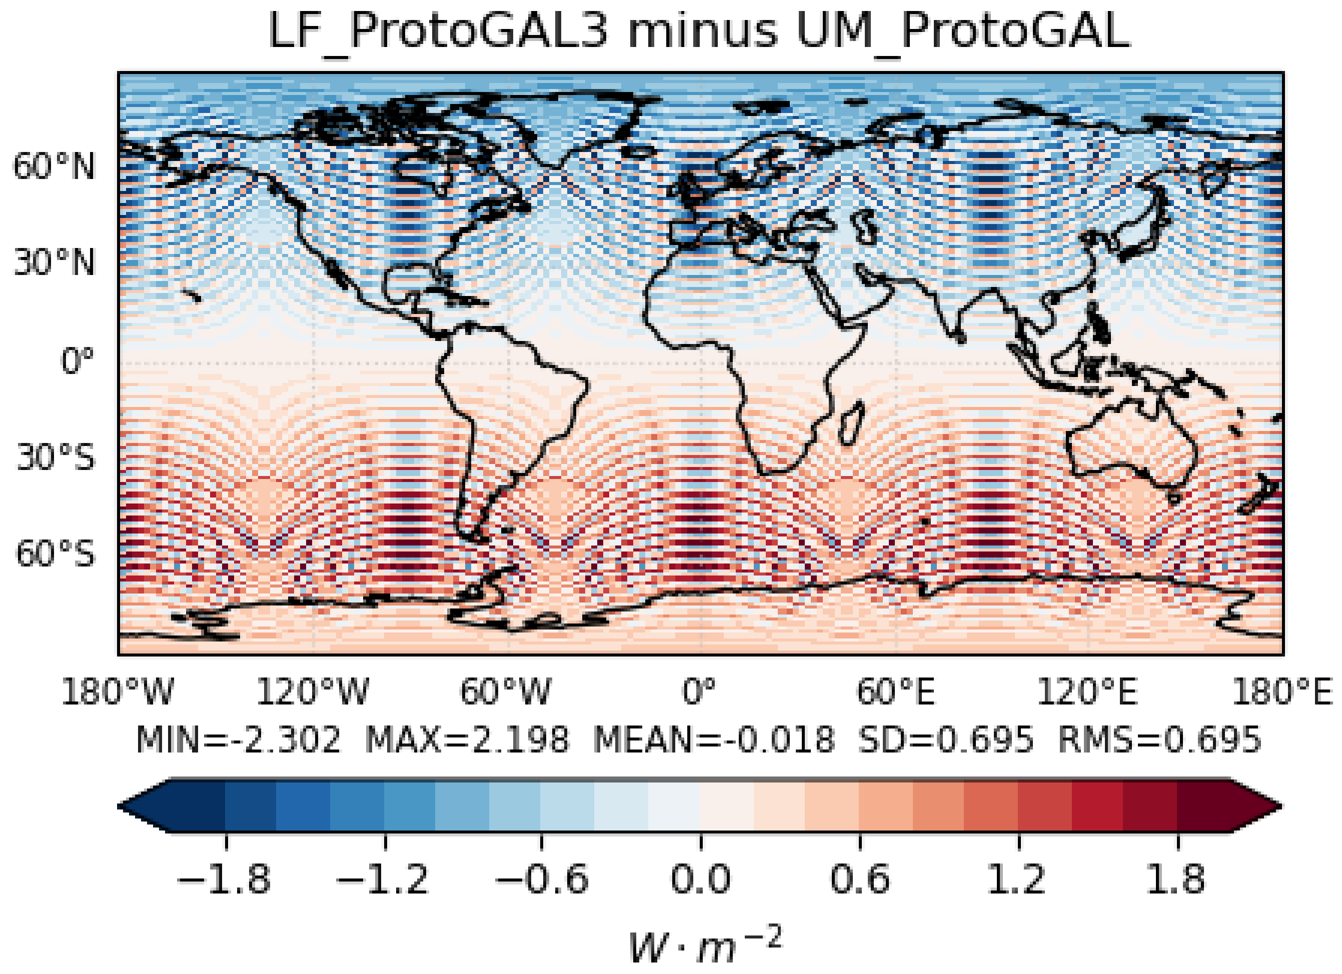

Met Office

Some progress

Regridded LFRic using Iris UGRID capability

Verified here against ECMWF using grid_stat

For model development we MUST:

- Limit the amount of interpolation for verification
- Be able to verify on the native grid, especially against point observations

But,

- None of the (15) interpolation methods currently handle UGRID;
- **At least some of them must be reimagined.**

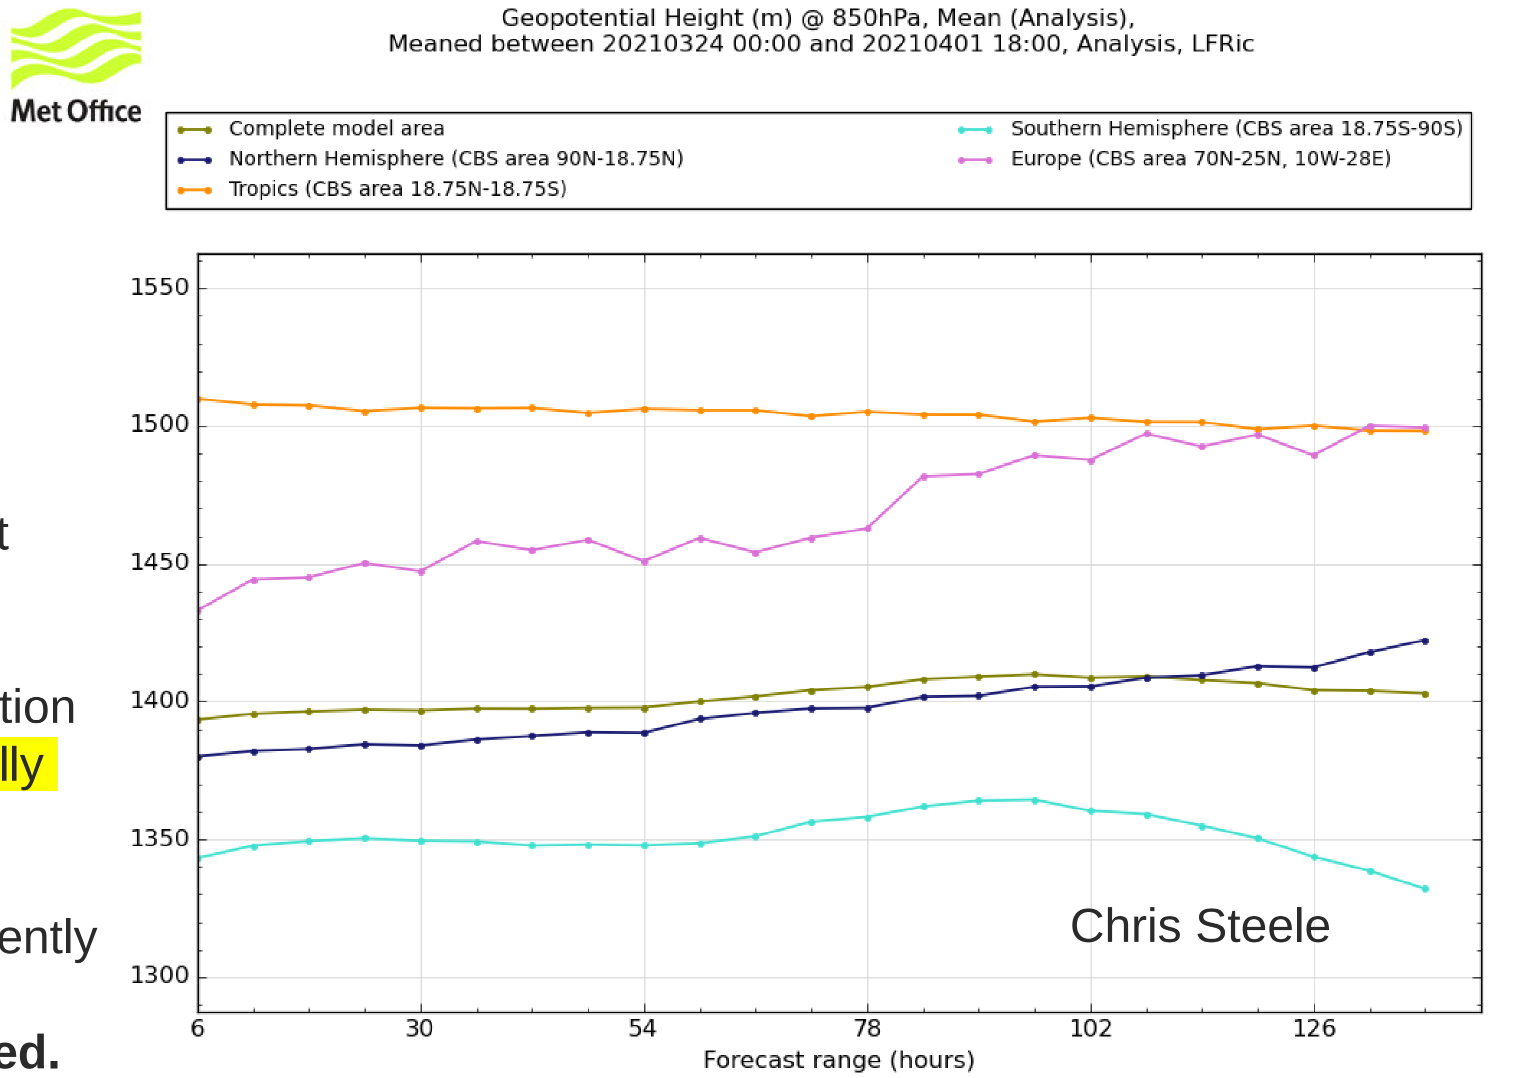

We have demonstrated that externally regridded LFRic data can be verified using MET.
It could even be done using python embedding but will not get us over the line when it comes to verifying against the native analysis or against observations.
Internalising this in the C++ code base still seems some way off, but we are actively seeking solutions that do not increase dependencies.

A few clarifying points about UGRID

Most users of forecast data will consume output after interpolation to a regular grid.
For model-intercomparison (e.g. CBS) this is also the case.
For score cards this is also the case.

However, for model development assessing the model upstream from the interpolation step is advisable!

Met Office

Using ODB with python embedding point_stat

The observations also remain a difficult nut to crack.

Here are some initial tests using python embedding the compiled ODB libraries to query the ODBs and extract the observations required for a point_stat job.

Initial results suggest that the observation counts between VER and MET are not identical.

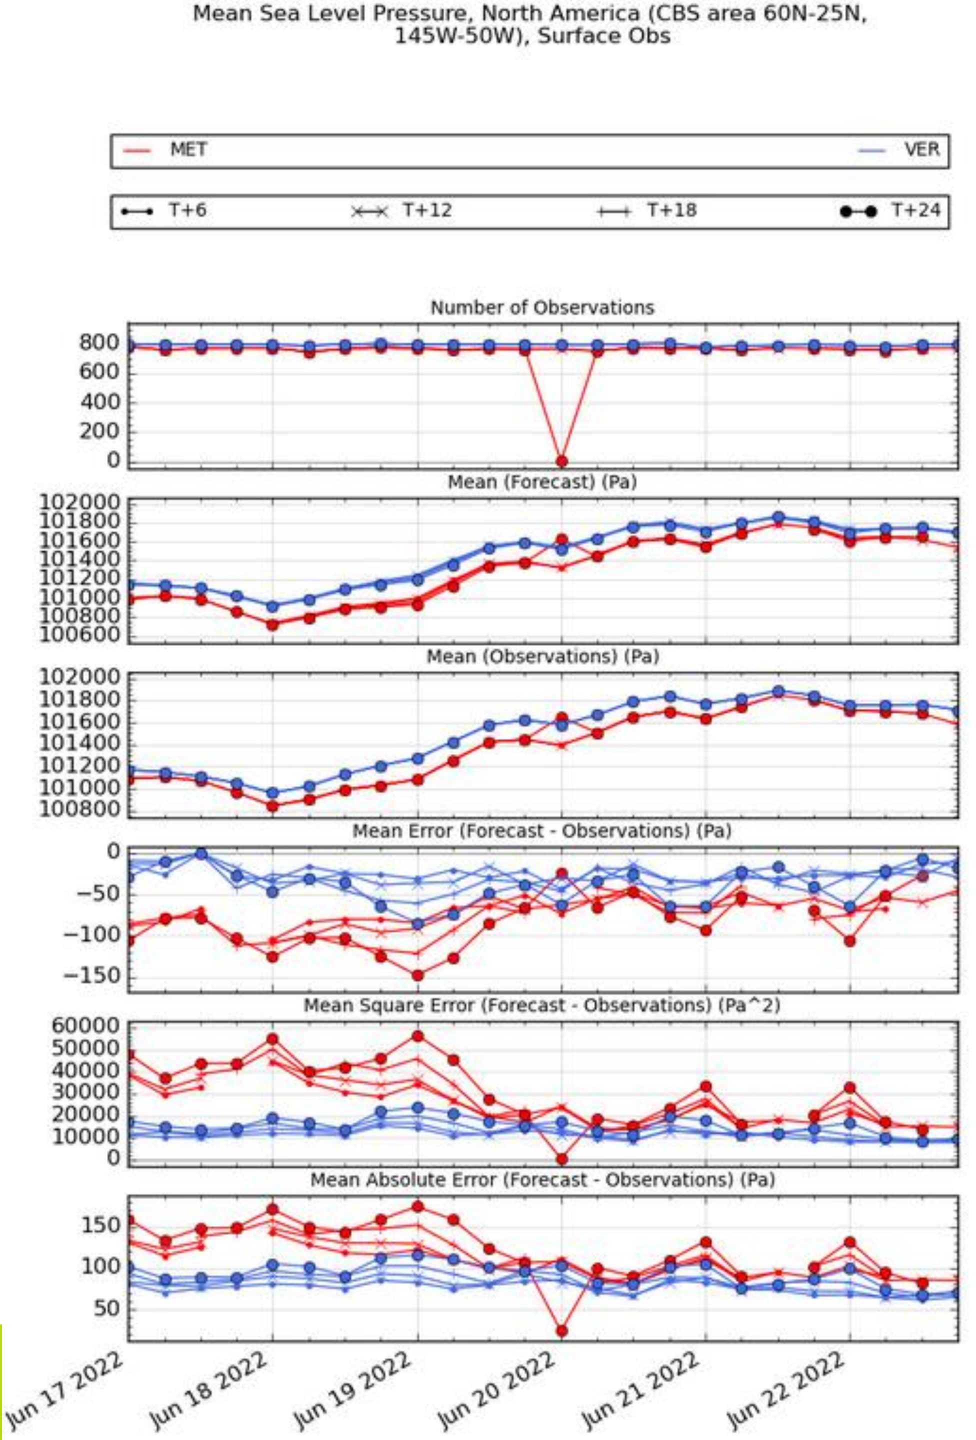

Rob Darvell

Met Office

Questions?