ECMWF Web re-engineering project

Baudouin Raoult

Peter Bispham, Andy Brady, Ricardo Correa,
Sylvie Lamy-Thepaut, Tim Orford,
David Richardson, Cihan Sahin, Stephan Siemen,
Carlos Valiente, Daniel Varela

ECMWF

The web re-engineering project

- **Motivation:**
  - Many of our users rely on our graphical web products for their daily work in their forecast offices, and have requested that our web services be continuously available
  - At the Annual Users' meetings, we have received requests to create tailored products (e.g. control the event threshold on probability maps)
- **Goals:**
  - Redesign the web infrastructure so that the web service is highly available and supported at the same level as the field dissemination
  - Provide more interactivity (e.g. zoom, pan, overlay parameters)
  - Allow product customisation (e.g. control the event threshold on probability maps)
  - Use open (OGC) standards so that ECMWF products can be embedded in users' own software

Requirements

- Highly available – Operationally supported
  - H/A Hardware
  - H/A Software
  - Operator monitoring
- Performance
  - Target: deliver a plot under 1 second
- Interactivity
  - Pan, zoom, overlay (à la GoogleMap)
  - Customisation, plots on demands (e.g. changing event probability threshold)
- Scalability
  - Support any future user load
  - Extensible: easy addition of new products

ECMWF

Gathering of user requirements

- The project has been presented on several occasions:
  - ECMWF Forecast Products Users' Meeting, Computer Representatives Meeting
  - Very positive feedback from forecasters
  - Most forecaster requests focused on the desire to be able to create customised products
  - Requests for new products
- Consultation process will continue throughout the project

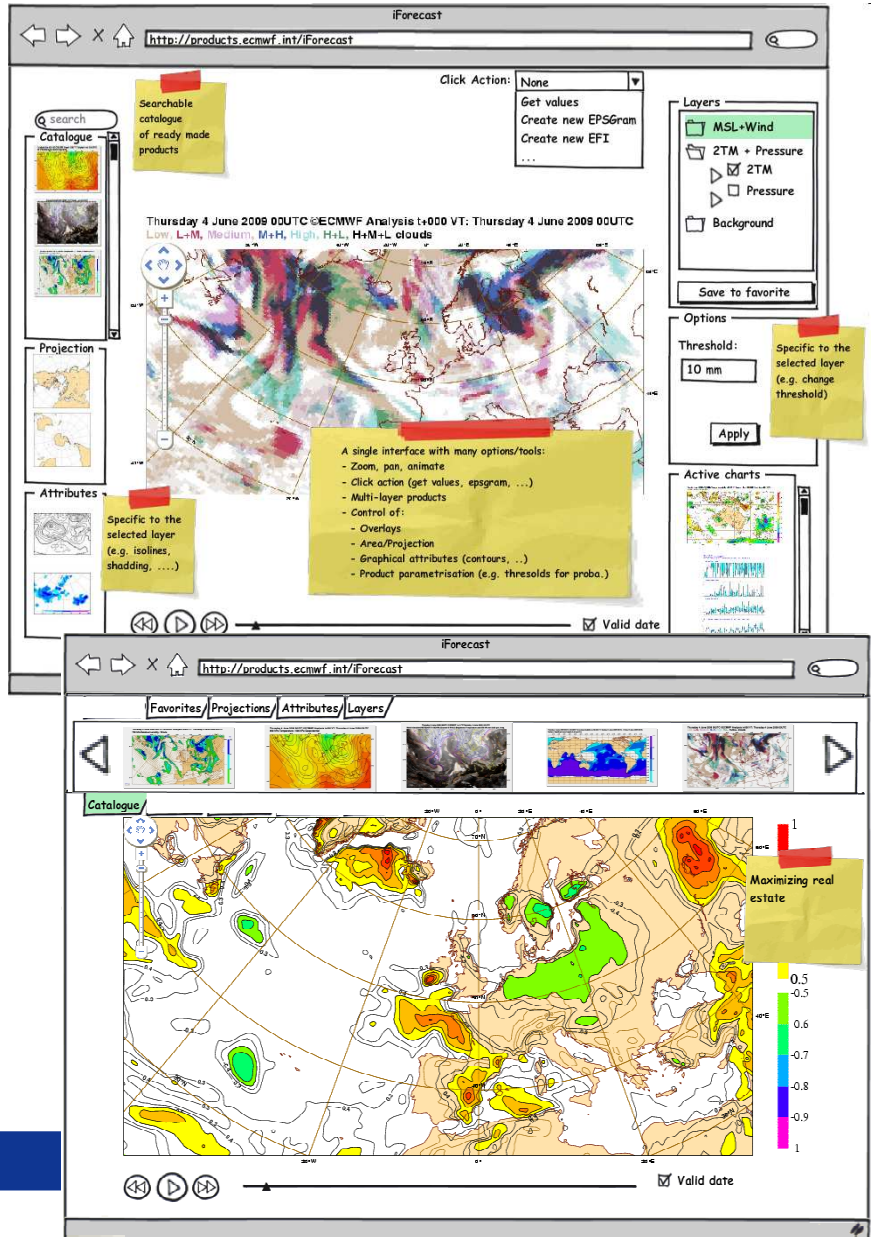

Service Oriented Architecture

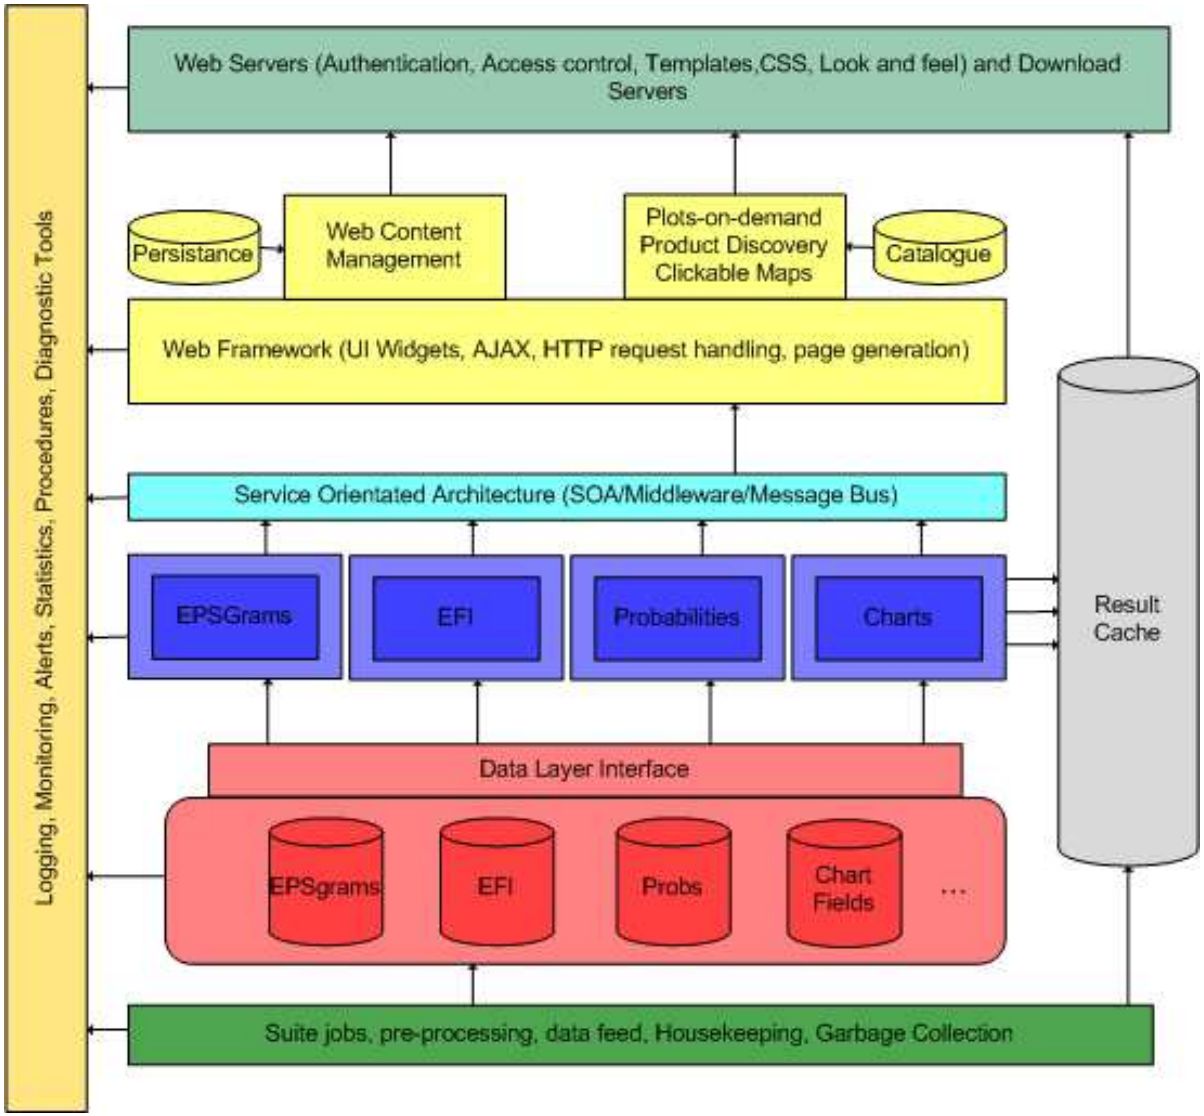

Hardware

- 2 Foundry Load Balancer ServerIronGT
- 3 servers hosting web servers
- 3 servers hosting web application
- 3 servers hosting several virtual machines
- 6 servers hosting storage, compute and plot services
- HP DL360 G5 Dual 2.5Hz Quad Core Xeon
- OpenSuSE Linux 11.1

- Systems located in different parts of the building, attached to different routers and different power sources

ECMWF

Software

- We investigated technologies used by the “big players” (e.g Google, Yahoo, Amazon, Facebook, Wikipedia …):
  - Memcached (Very fast distributed memory)
  - Tokyo Tyrant (Scalable, distributed persistent space)
  - Hadoop (High availability and redundant distributed data)
  - Xen (Virtualisation)
  - DRDB (Network RAID
  - Ganeti (H/A Cluster management)
  - Nagios (Alerts system)
  - Scribe (Distributed logging)

Software (cont.)

- …
  - Ganglia (Distributed monitoring)
  - Django (Python based Web framework, server side)
  - jQuery (JavaScript based web framework, client side)
  - OpenLayers (JavaScript based OGC WMS-client)
  - Apache 2.2 (Web server)
  - MySQL (Database)
- And of course:
  - Magics++
  - grib_api
  - Mars
  - Metview …

ECMWF

Service Oriented Architecture

- Multi-tier architecture, deployed on a series of Linux clusters:
  - Web frontend (Web server)
  - Web backend (Dynamic page generation)
  - Services (Plotting, probability computations, EPSgrams, …)
  - Data layer (Raw fields)
- Cluster approach provides built-in scalability, redundancy and load balancing
- Critical components run on virtual machines that can be redeployed dynamically

Deployment

- Virtual machines for critical components and single points of failure
  - Hadoop name node
  - SOA Broker
  - Spot database
  - Catalogue (MySQL)
  - All virtual machines sized in such a way that they can fit in a smaller number of nodes if necessary
- Physical machines for components with built-in redundancy
  - Hadoop data nodes
  - Memcached servers
  - Services (plot, retrieve, probabilities, epsgrams, …)

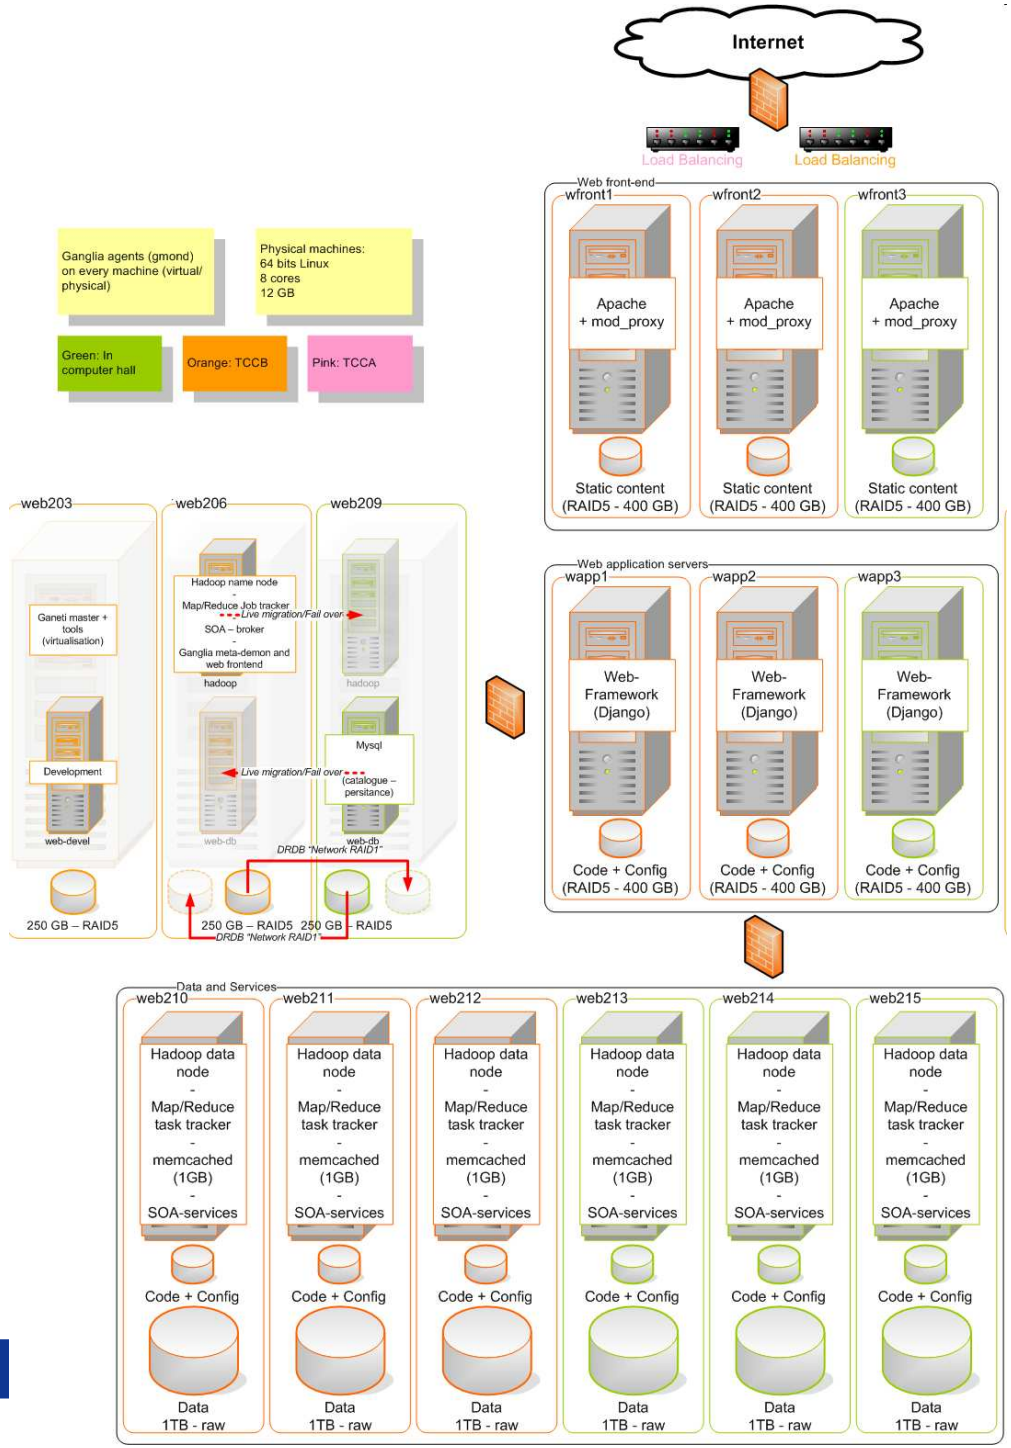

Developing in an SOA environment (is hard)

- Distributed design
- Troubleshooting

- Diagnostics tools
- Instrumentation
- Regression tests

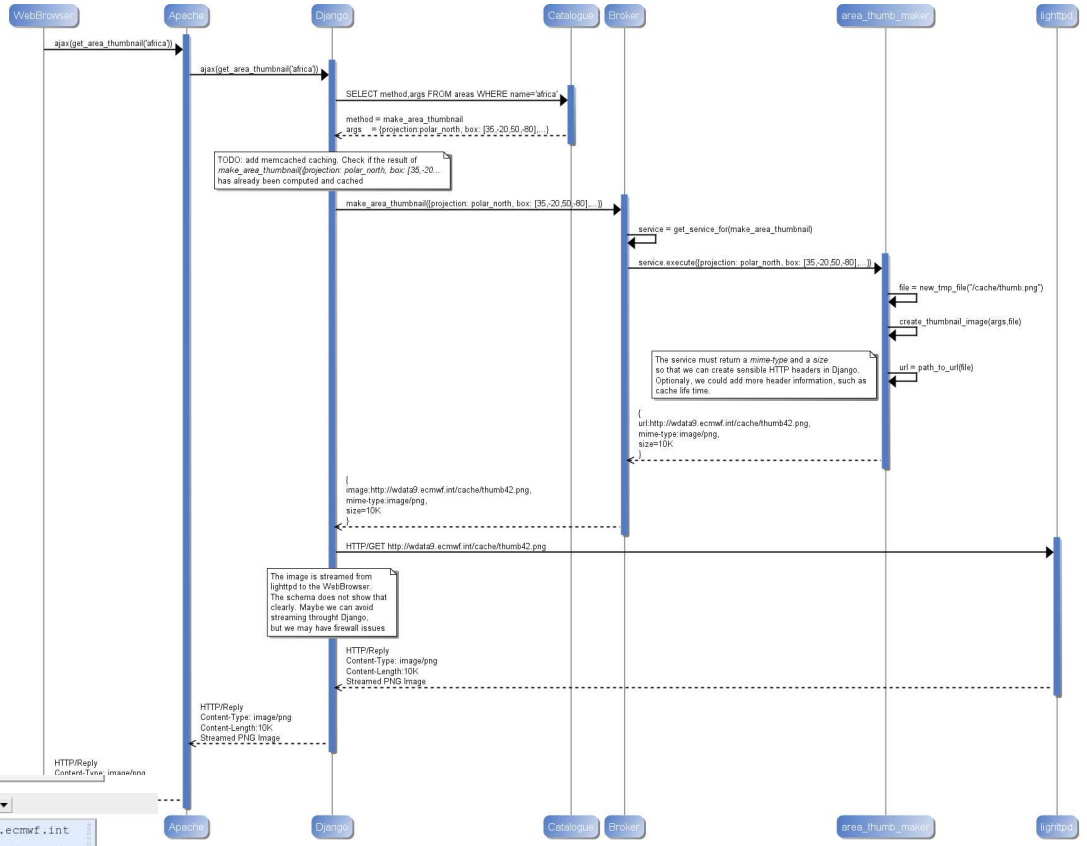

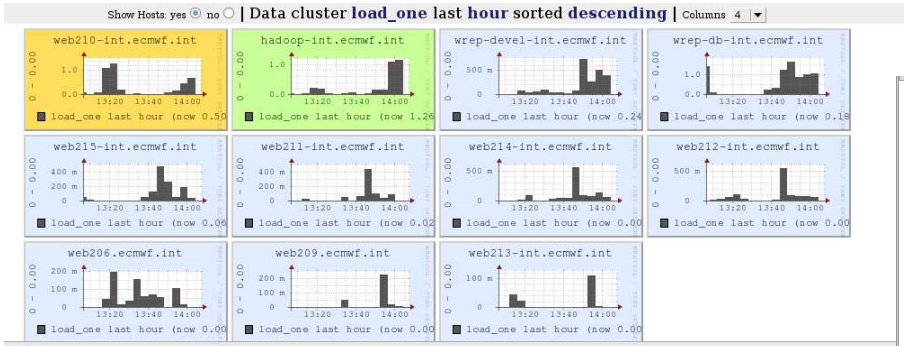

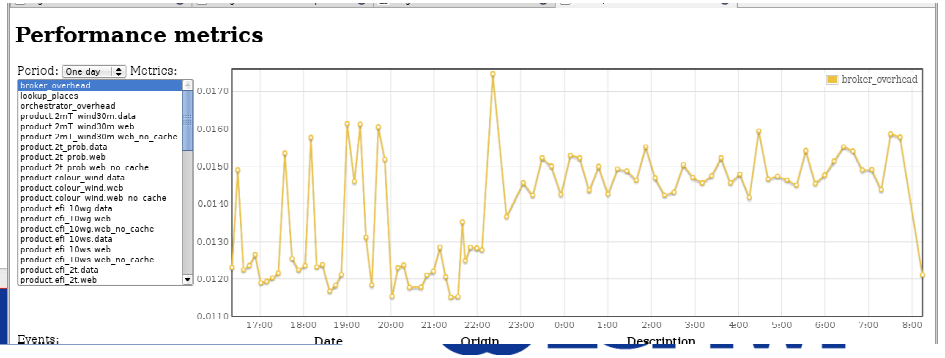

Prototype: Forecasting tool

- Interactivity: zooming, panning, …
- Customisation:
  - Probabilities threshold, …
  - Show/hide, add/remove layers
- Related products: Epsgrams

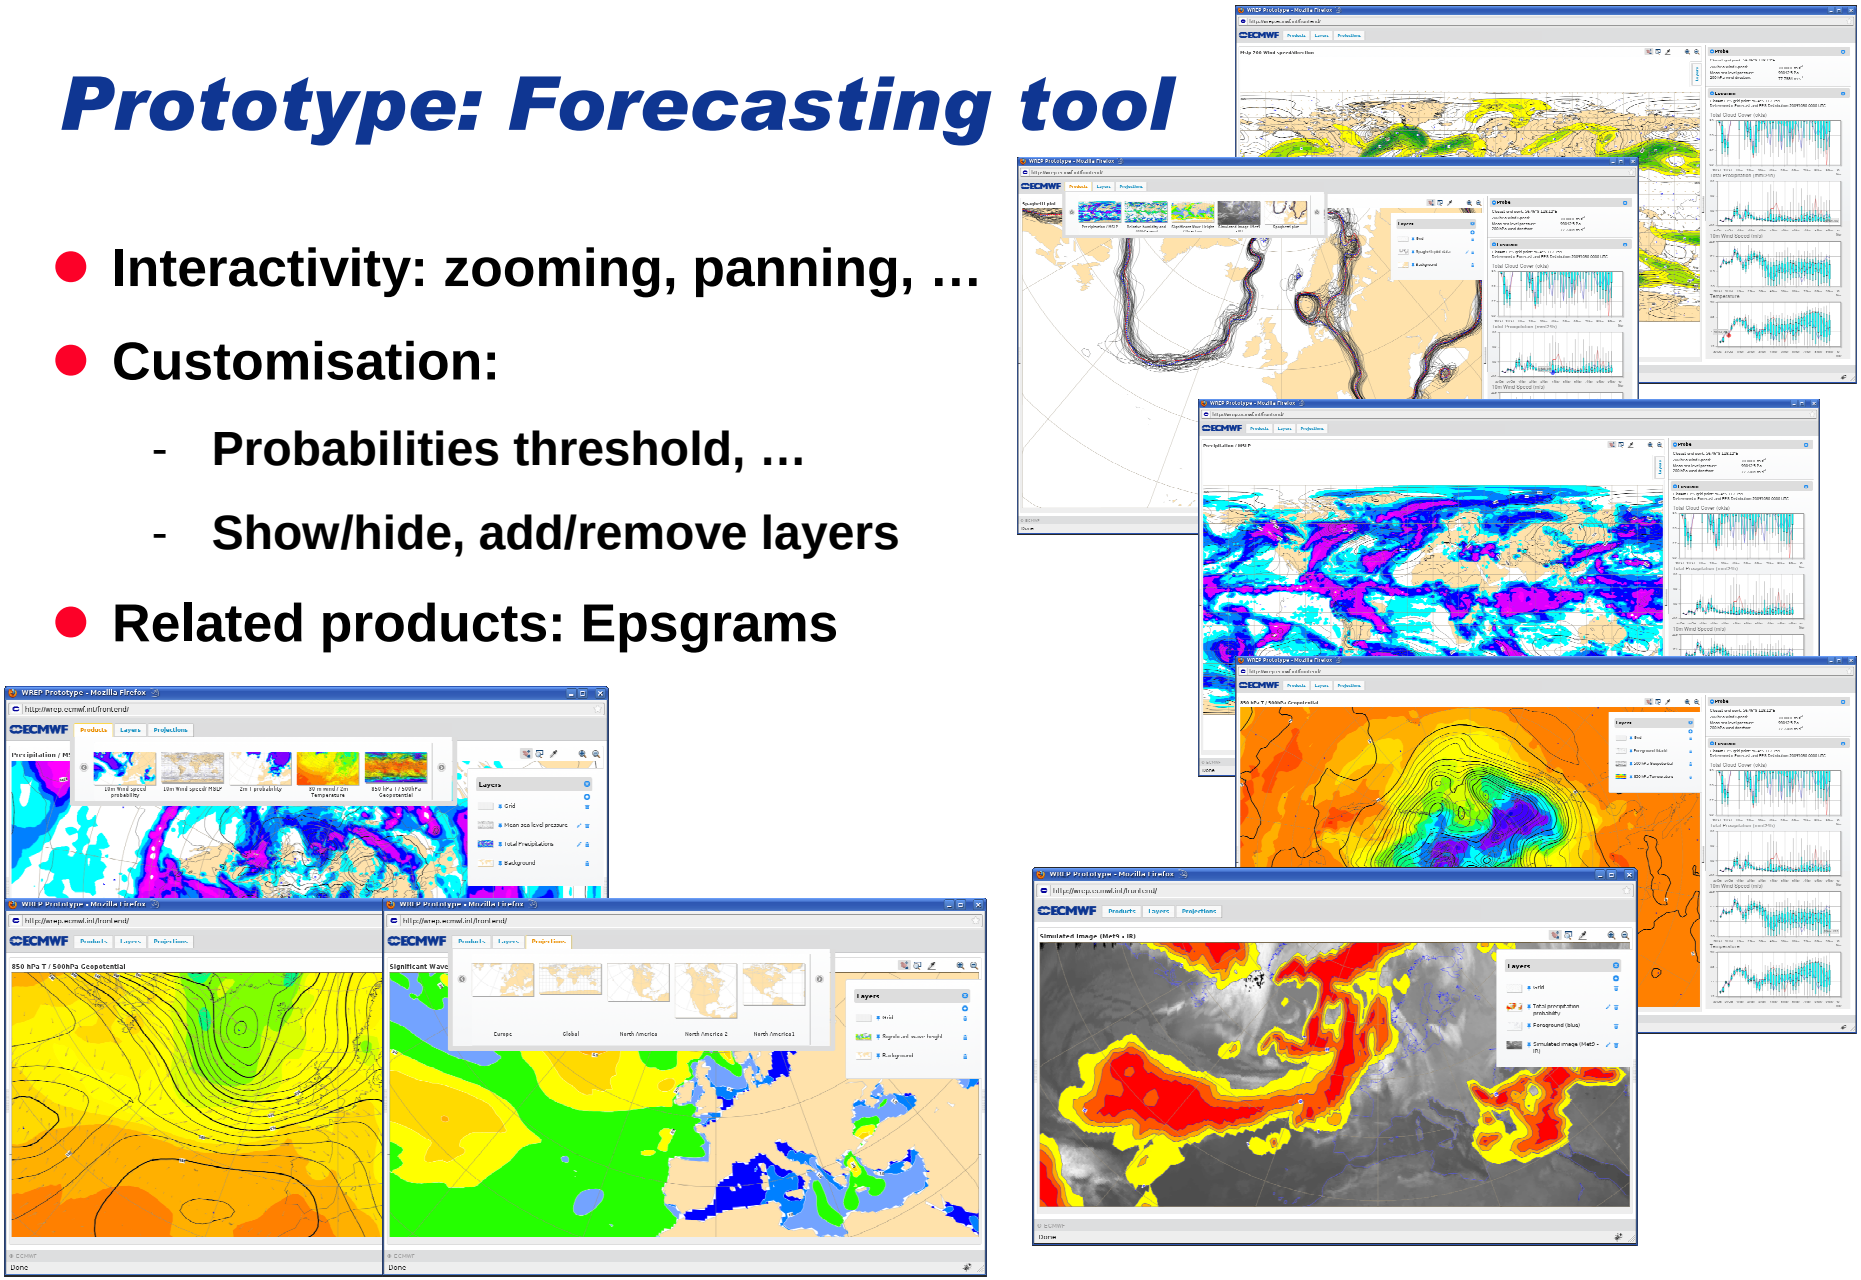

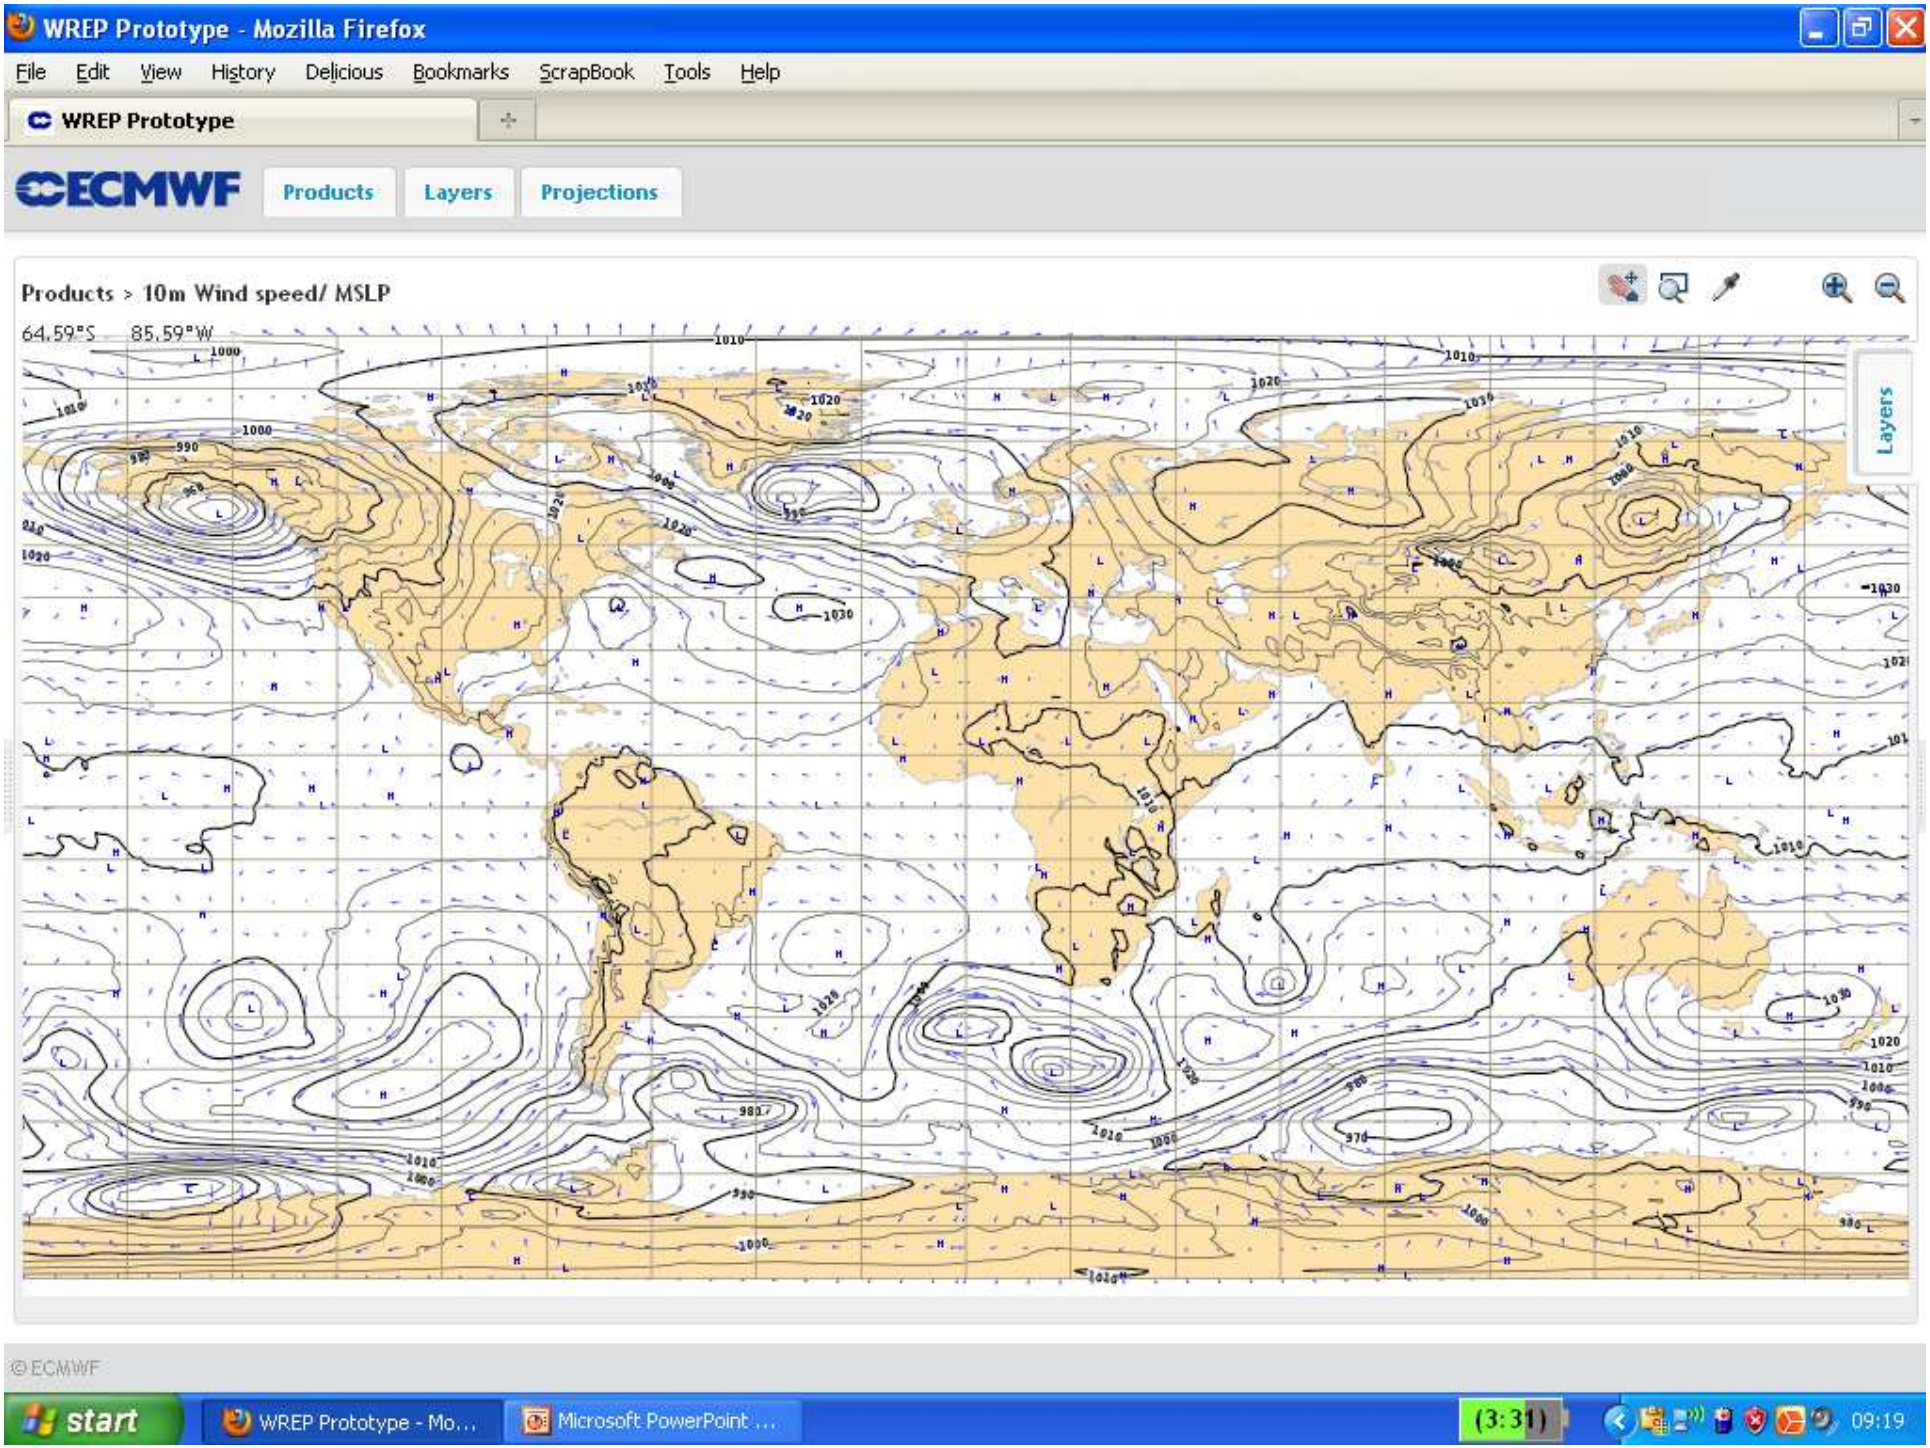
WREP Prototype - Mozilla Firefox
File Edit View History Delicious Bookmarks ScrapBook Tools Help
WREP Prototype
ECMWF
Products
Layers
Projections
Products > 10m Wind speed/ MSLP
64.59°S 85.59°W
Layers
© ECMWF
start
WREP Prototype - Mo...
Microsoft PowerPoint ...
(3:31)
09:19

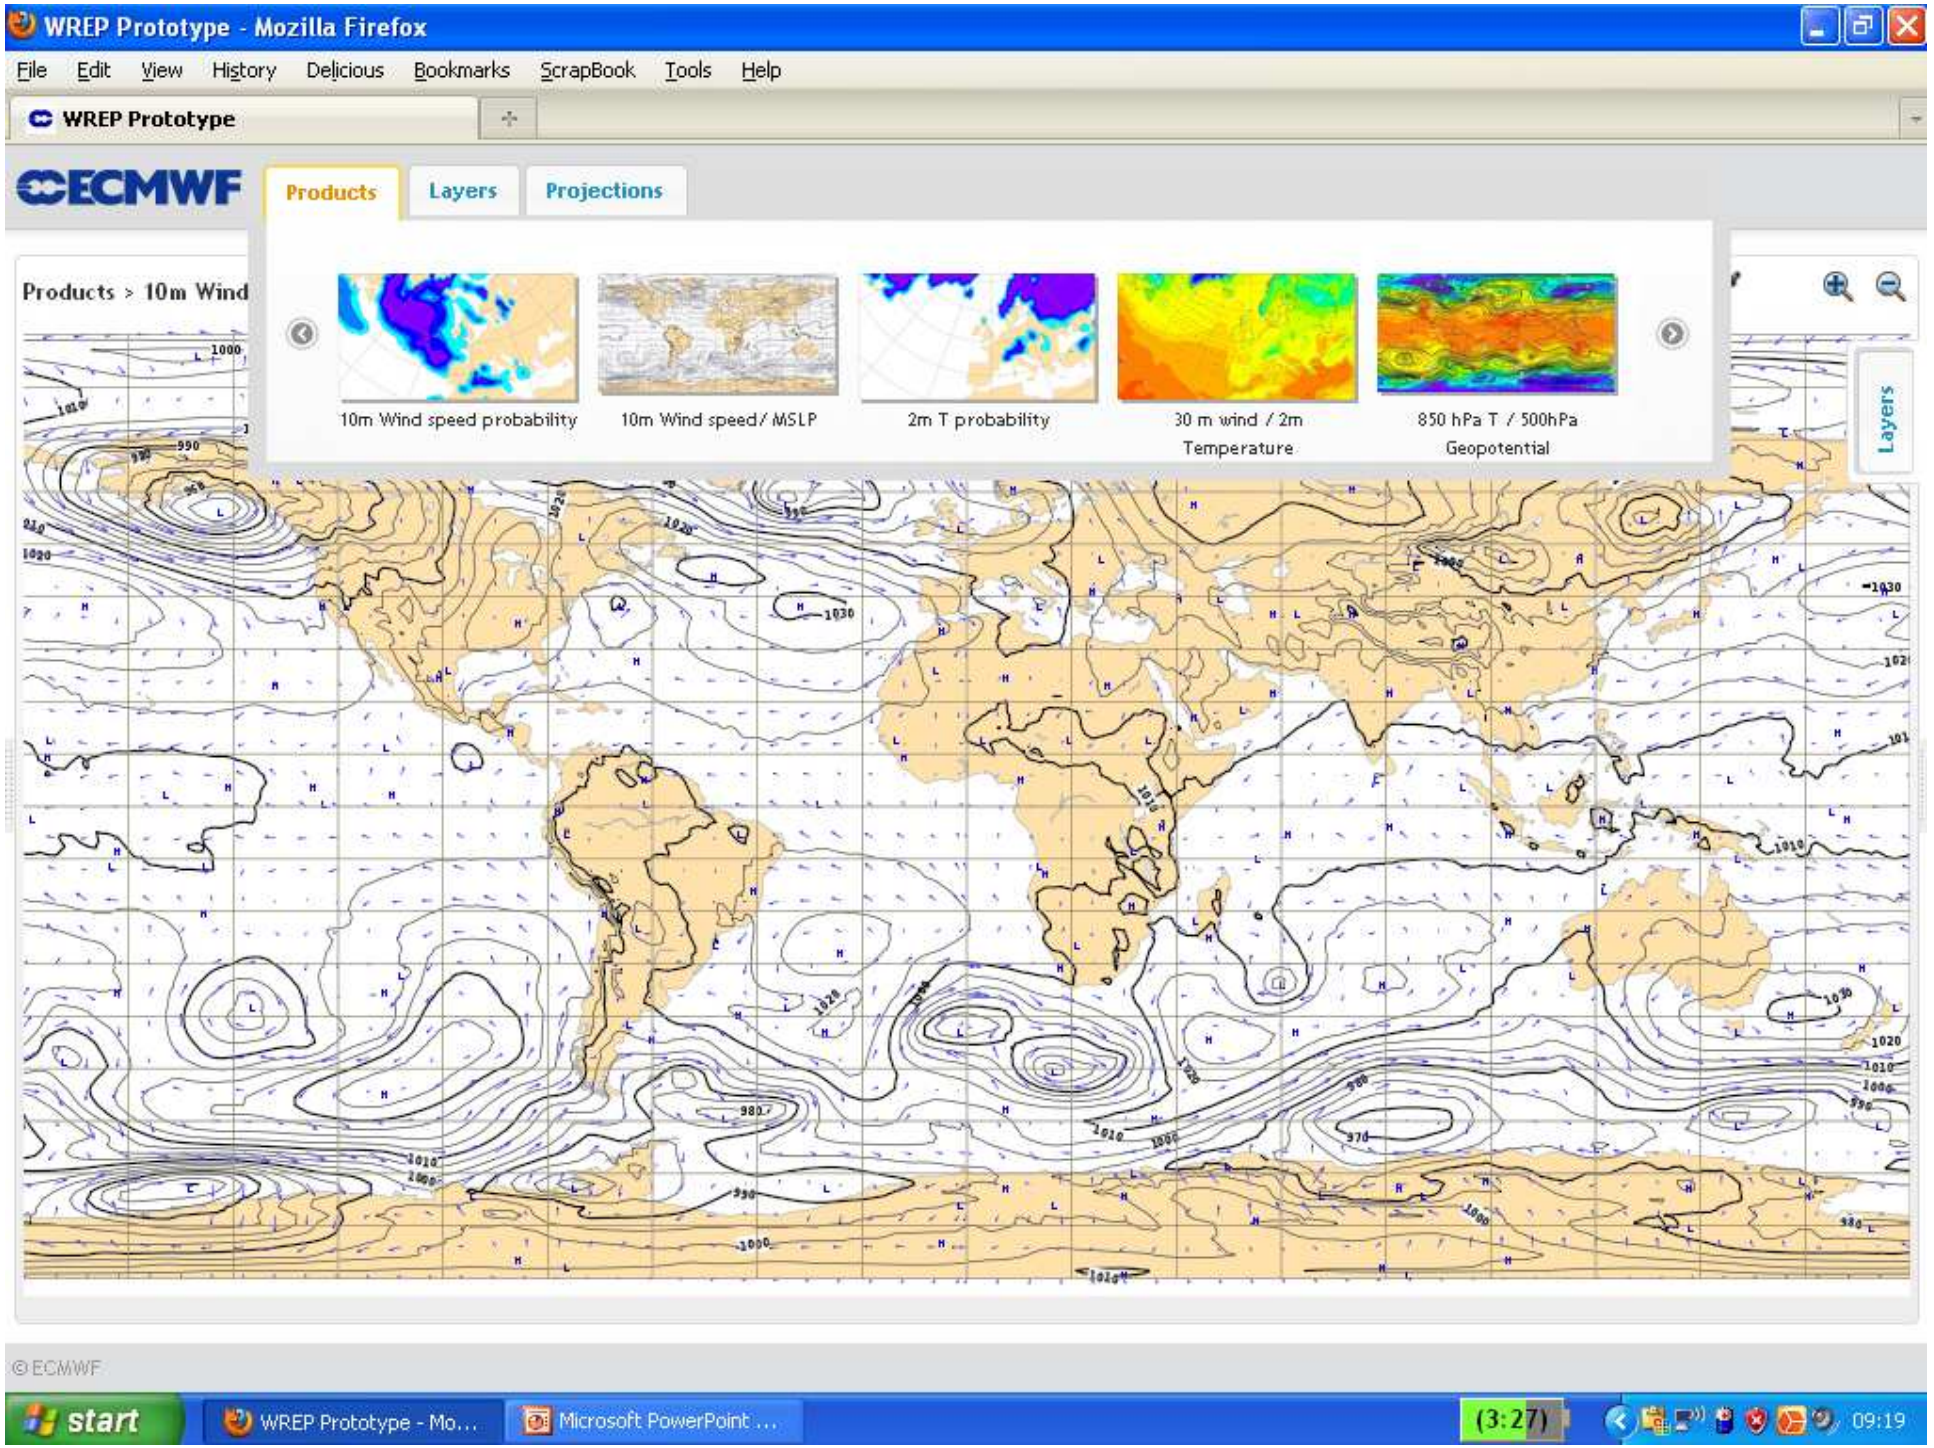
WREP Prototype - Mozilla Firefox
File Edit View History Delicious Bookmarks ScrapBook Tools Help
WREP Prototype
ECMWF
Products
Layers
Projections
Products > 10m Wind
10m Wind speed probability
10m Wind speed/ MSLP
2m T probability
30 m wind / 2m Temperature
850 hPa T / 500hPa Geopotential
Layers
© ECMWF
start
WREP Prototype - Mo...
Microsoft PowerPoint ...
(3:27)
09:19

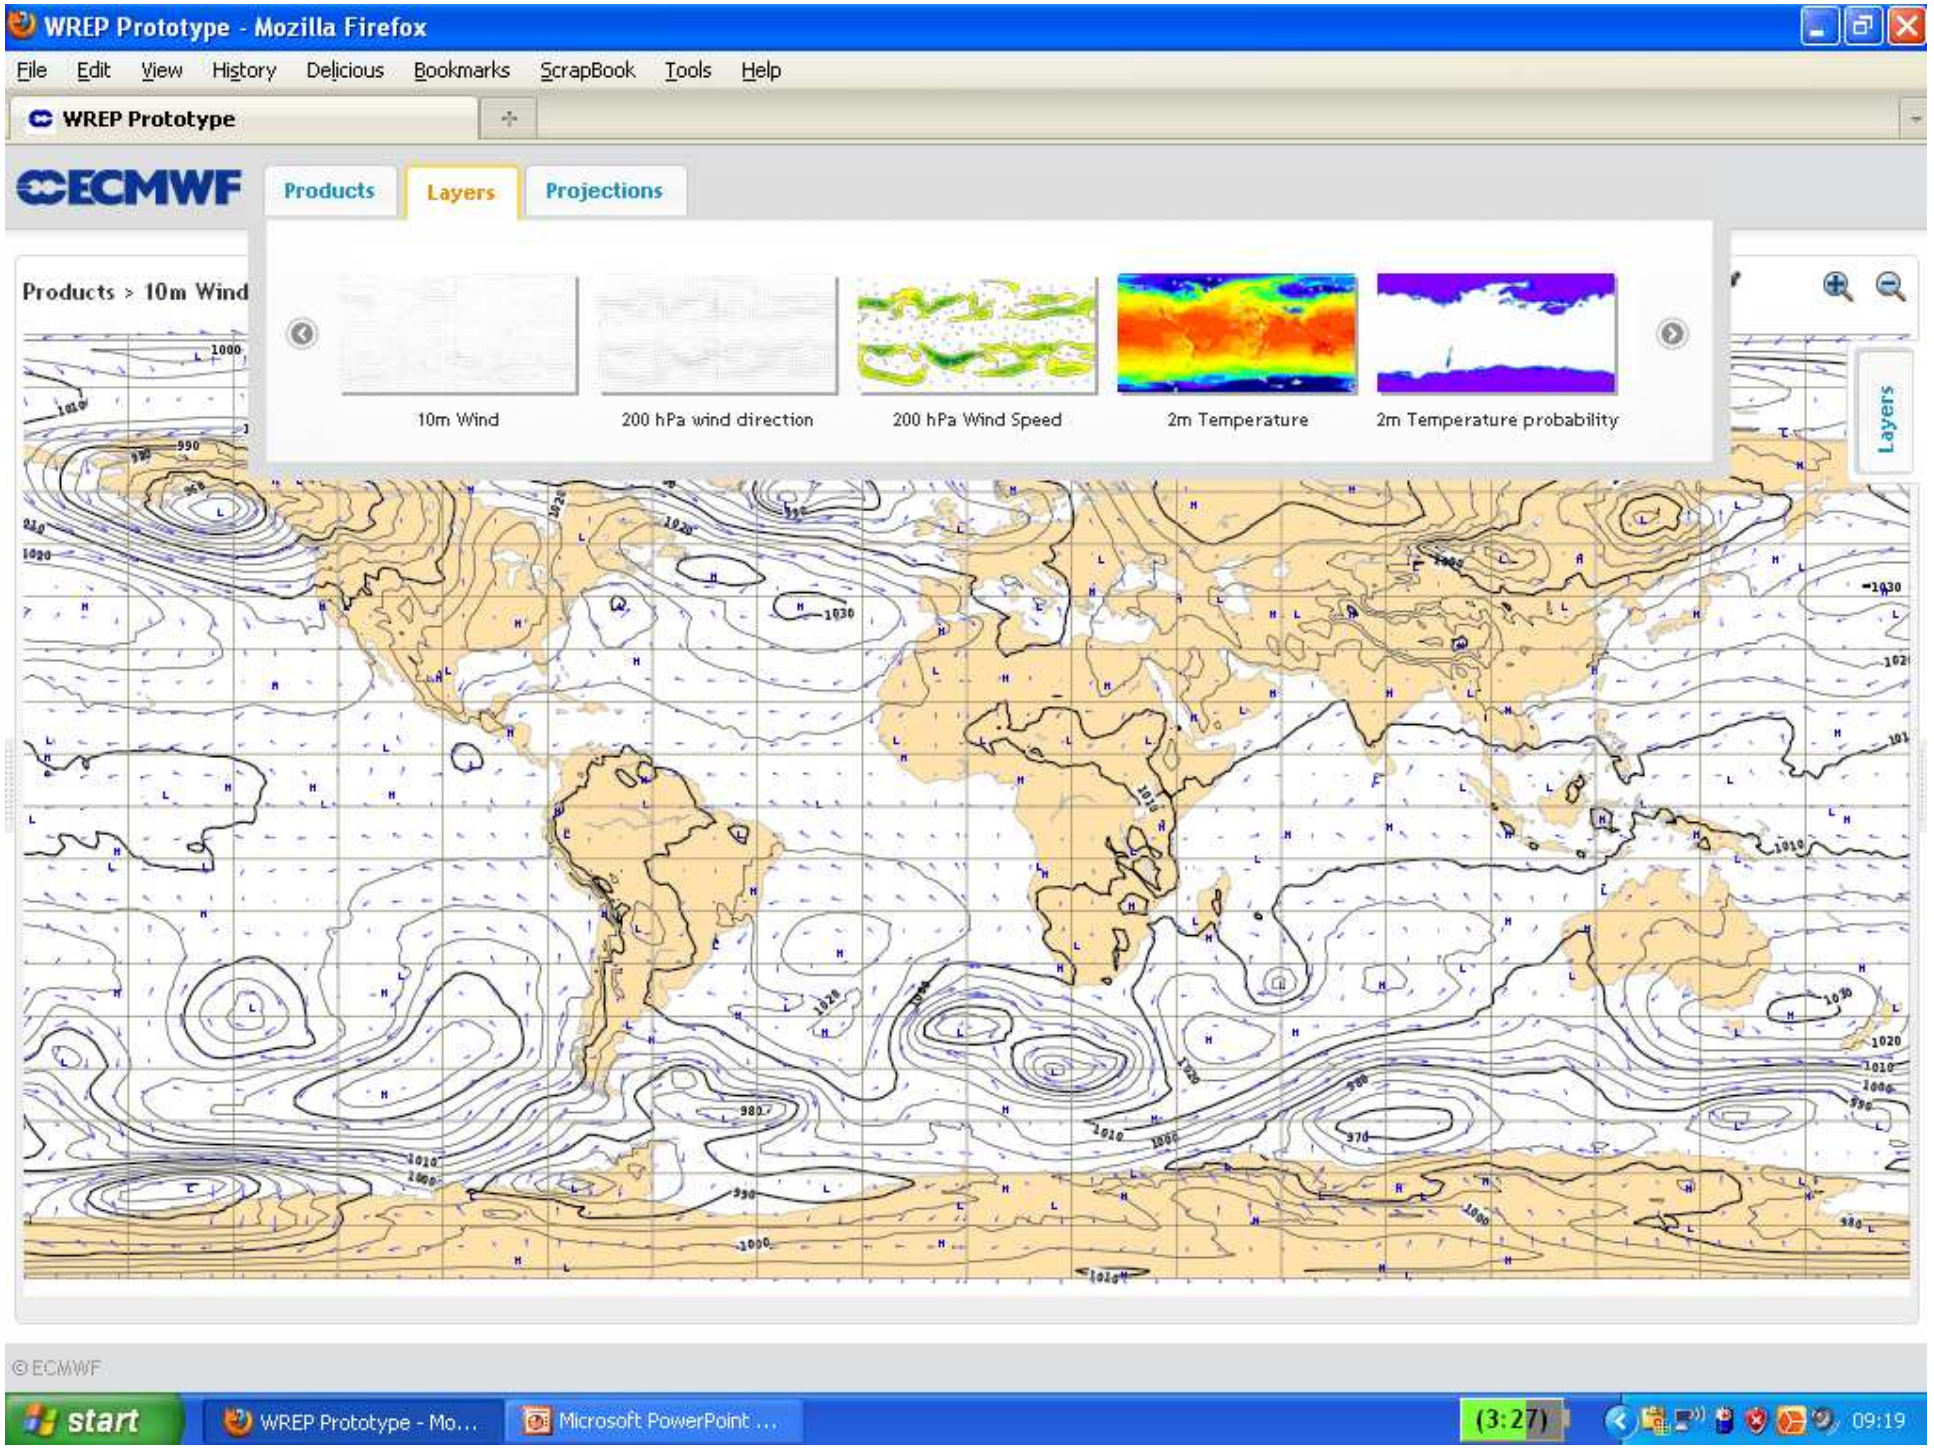
WREP Prototype - Mozilla Firefox
File Edit View History Delicious Bookmarks ScrapBook Tools Help
WREP Prototype
ECMWF
Products
Layers
Projections
Products > 10m Wind
10m Wind
200 hPa wind direction
200 hPa Wind Speed
2m Temperature
2m Temperature probability
Layers
© ECMWF
start
WREP Prototype - Mo...
Microsoft PowerPoint ...
(3:27)
09:19

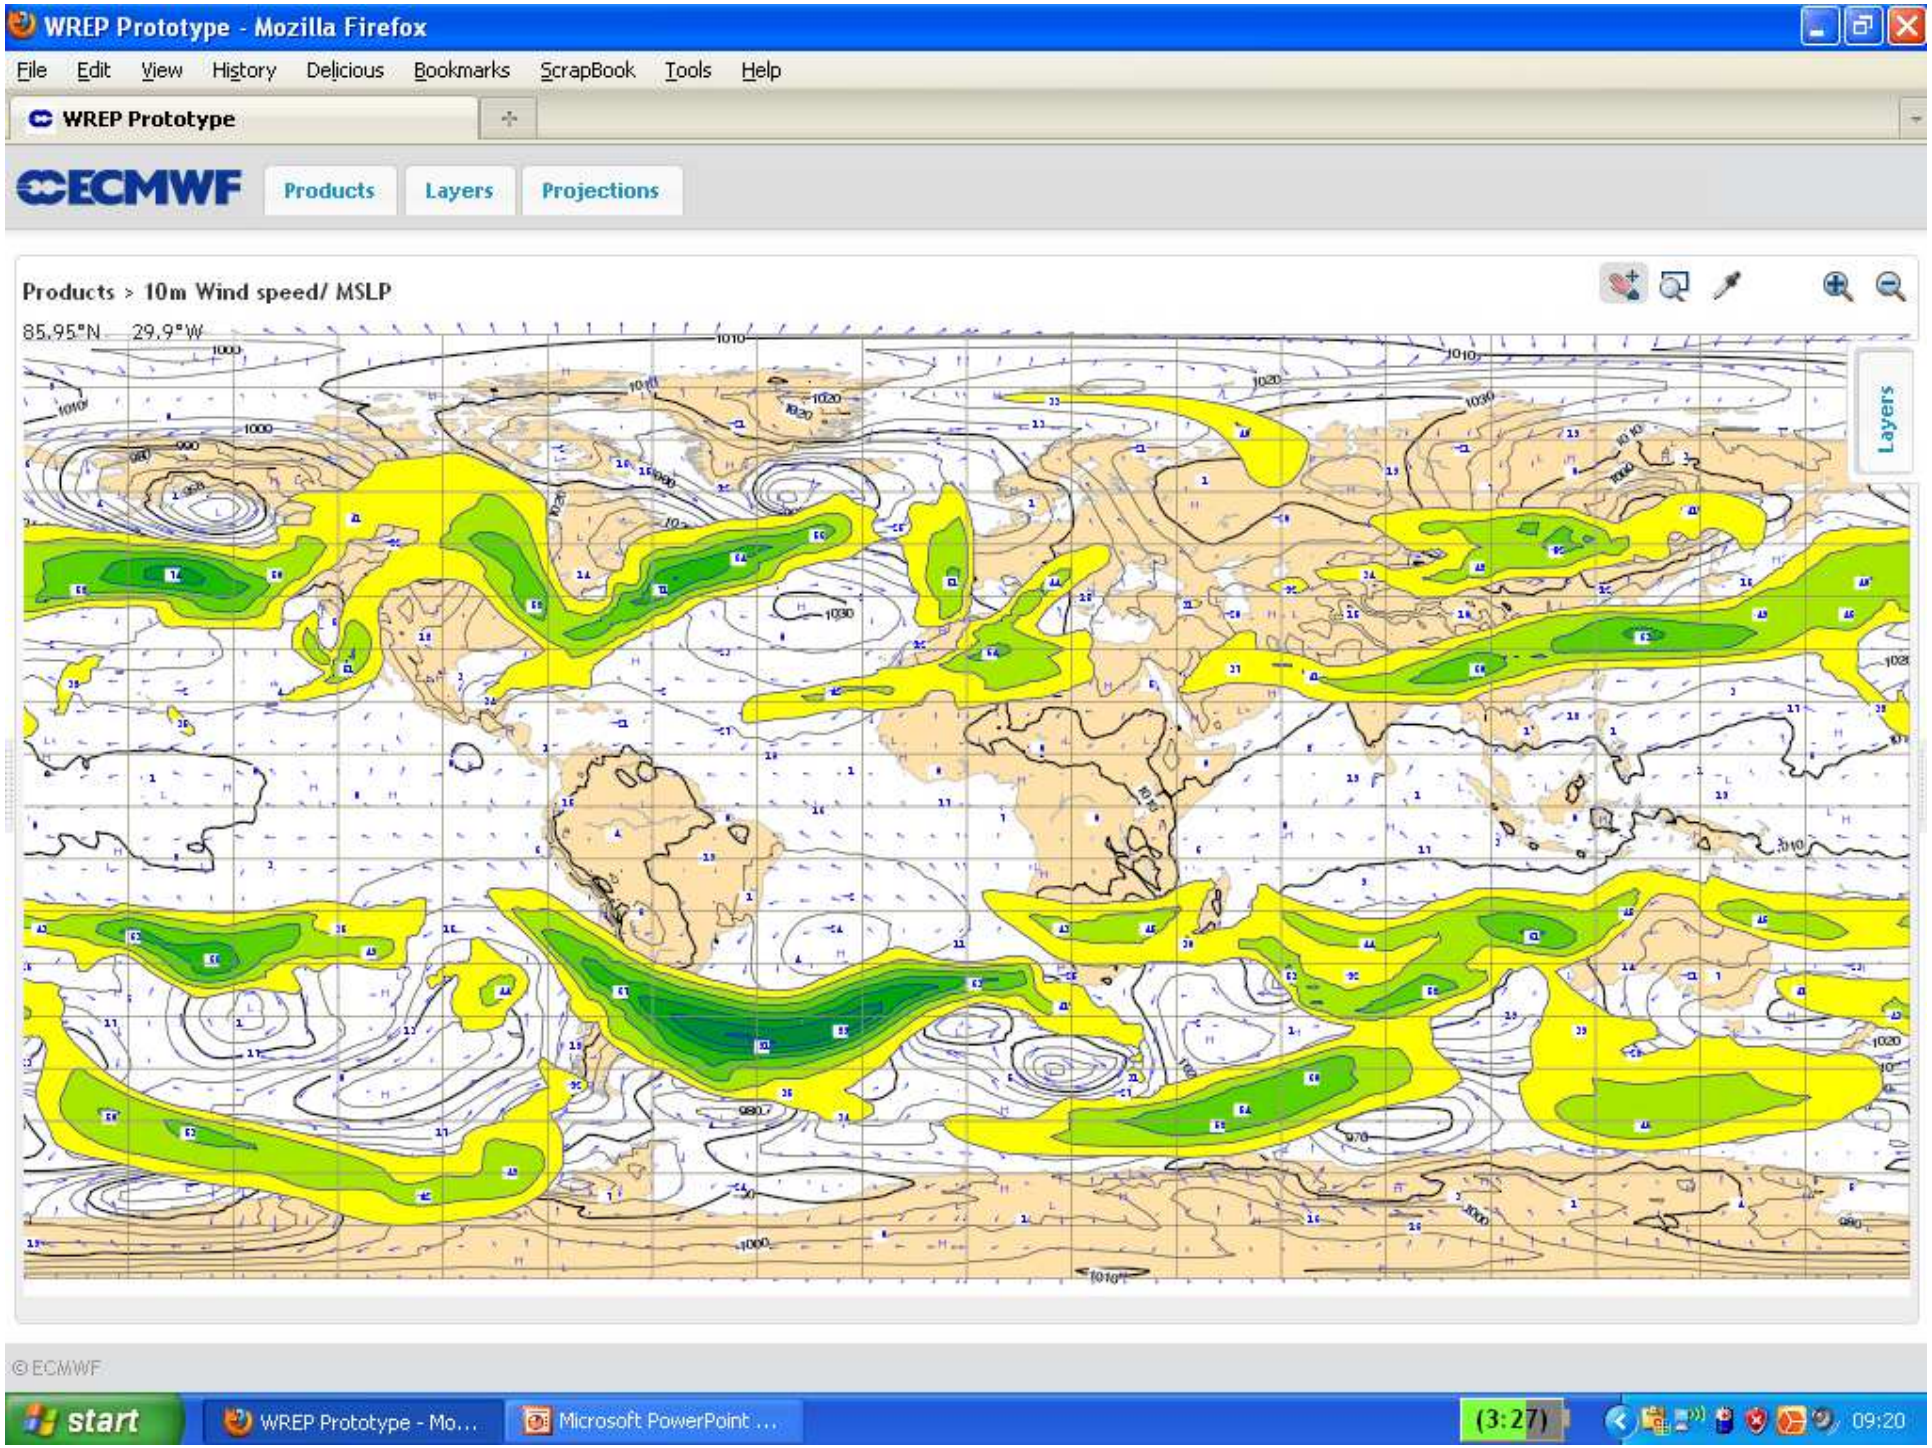
WREP Prototype - Mozilla Firefox
File Edit View History Delicious Bookmarks ScrapBook Tools Help
WREP Prototype
ECMWF
Products
Layers
Projections
Products > 10m Wind speed/ MSLP
85.95°N 29.9°W
Layers
© ECMWF
start
WREP Prototype - Mo...
Microsoft PowerPoint ...
(3:27)
09:20

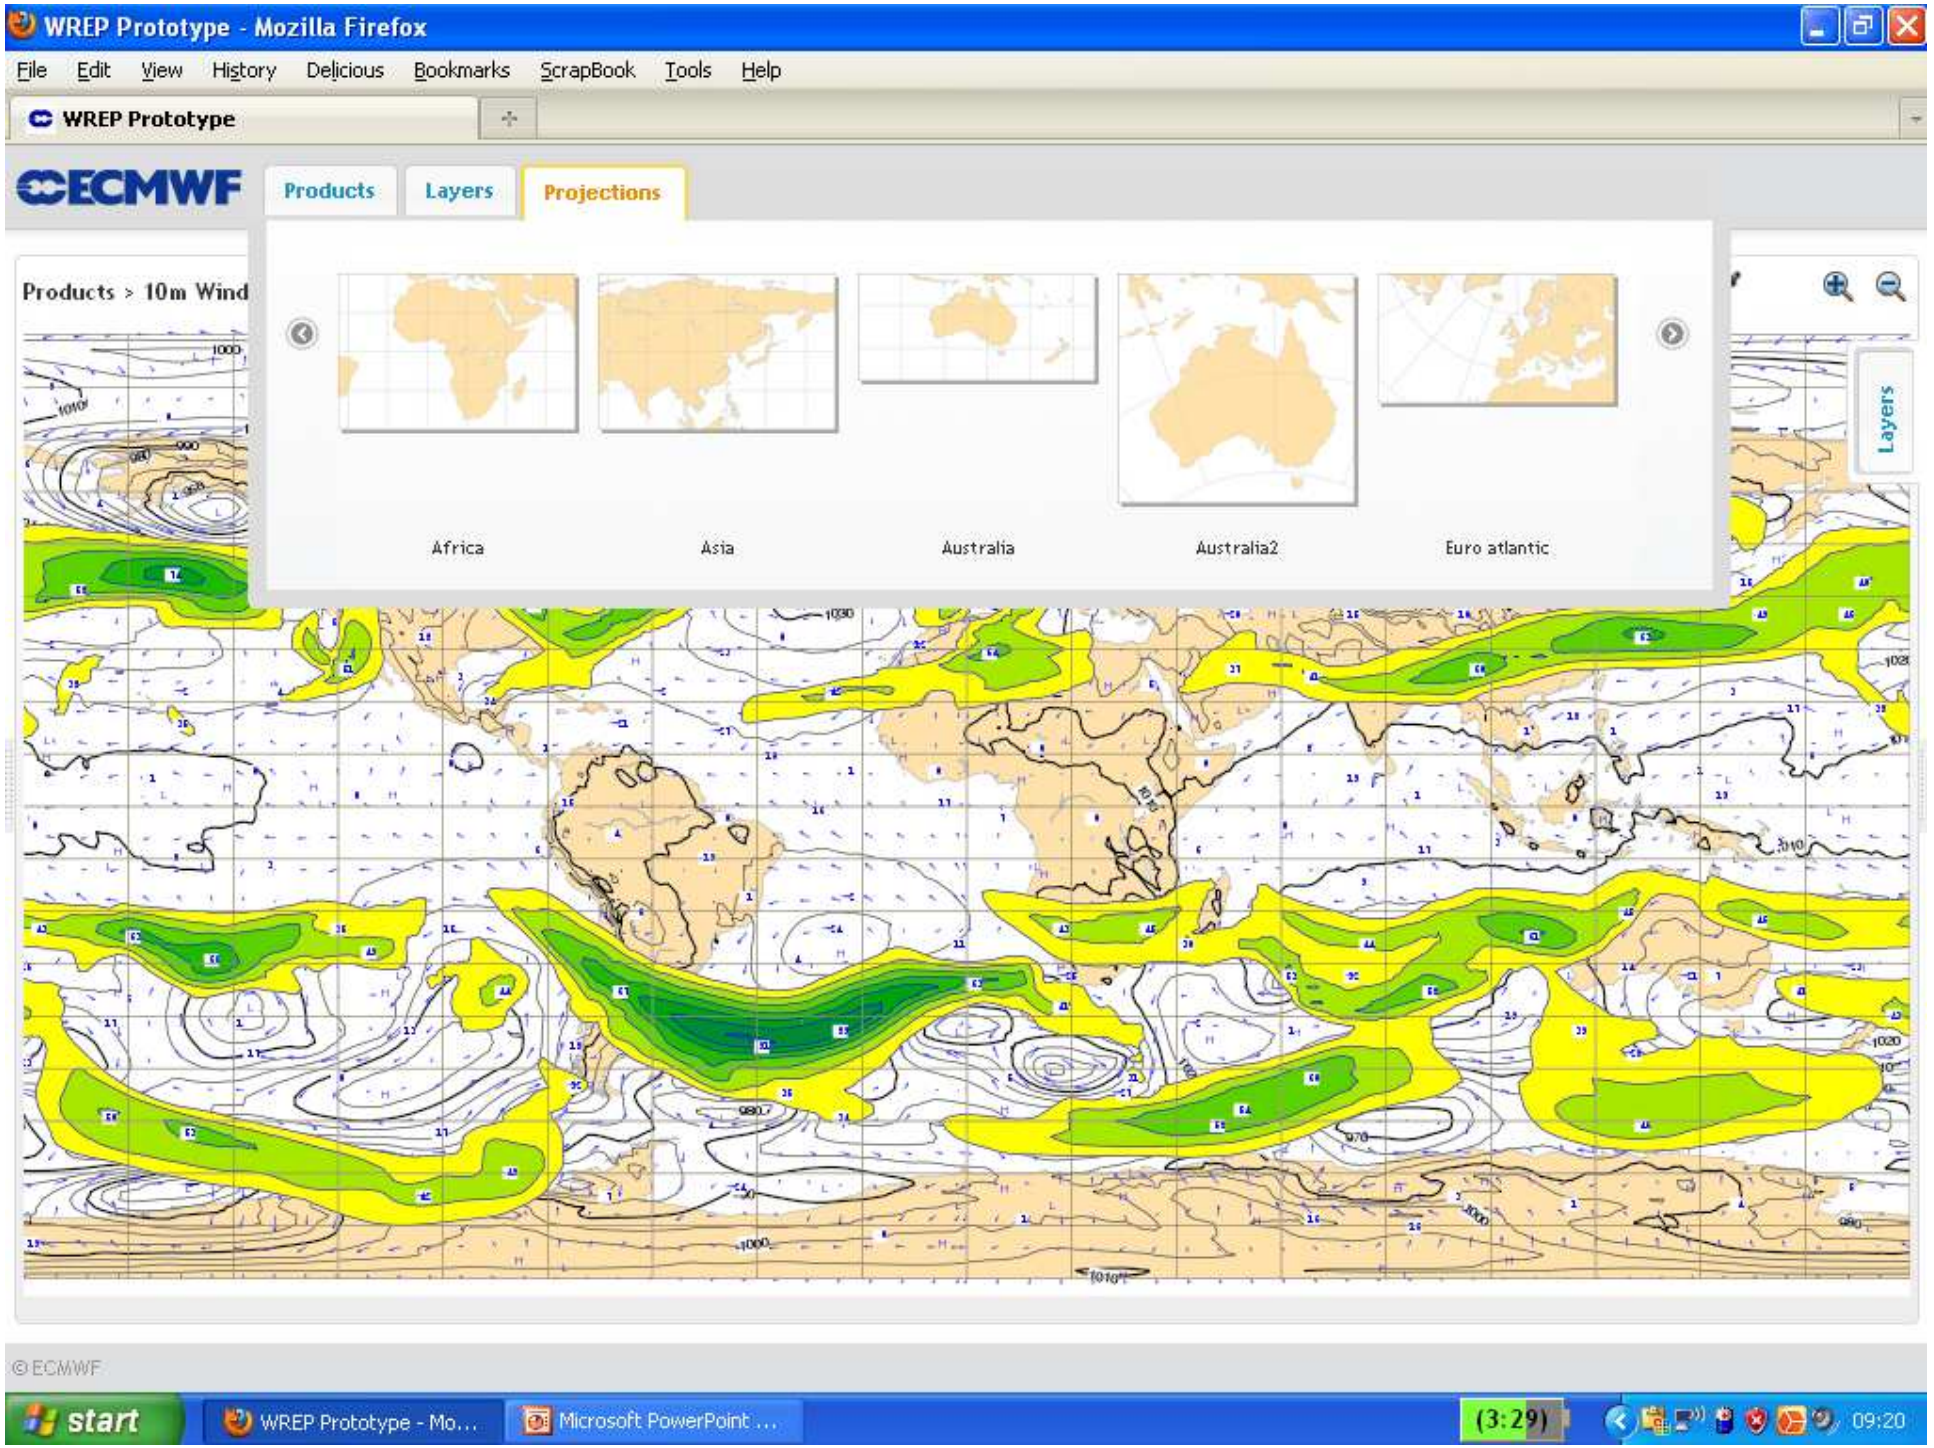
WREP Prototype - Mozilla Firefox
File Edit View History Delicious Bookmarks ScrapBook Tools Help
WREP Prototype
ECMWF
Products
Layers
Projections
Products > 10m Wind
Africa
Asia
Australia
Australia2
Euro atlantic
Layers
© ECMWF
start
WREP Prototype - Mo...
Microsoft PowerPoint ...
(3:29)
09:20

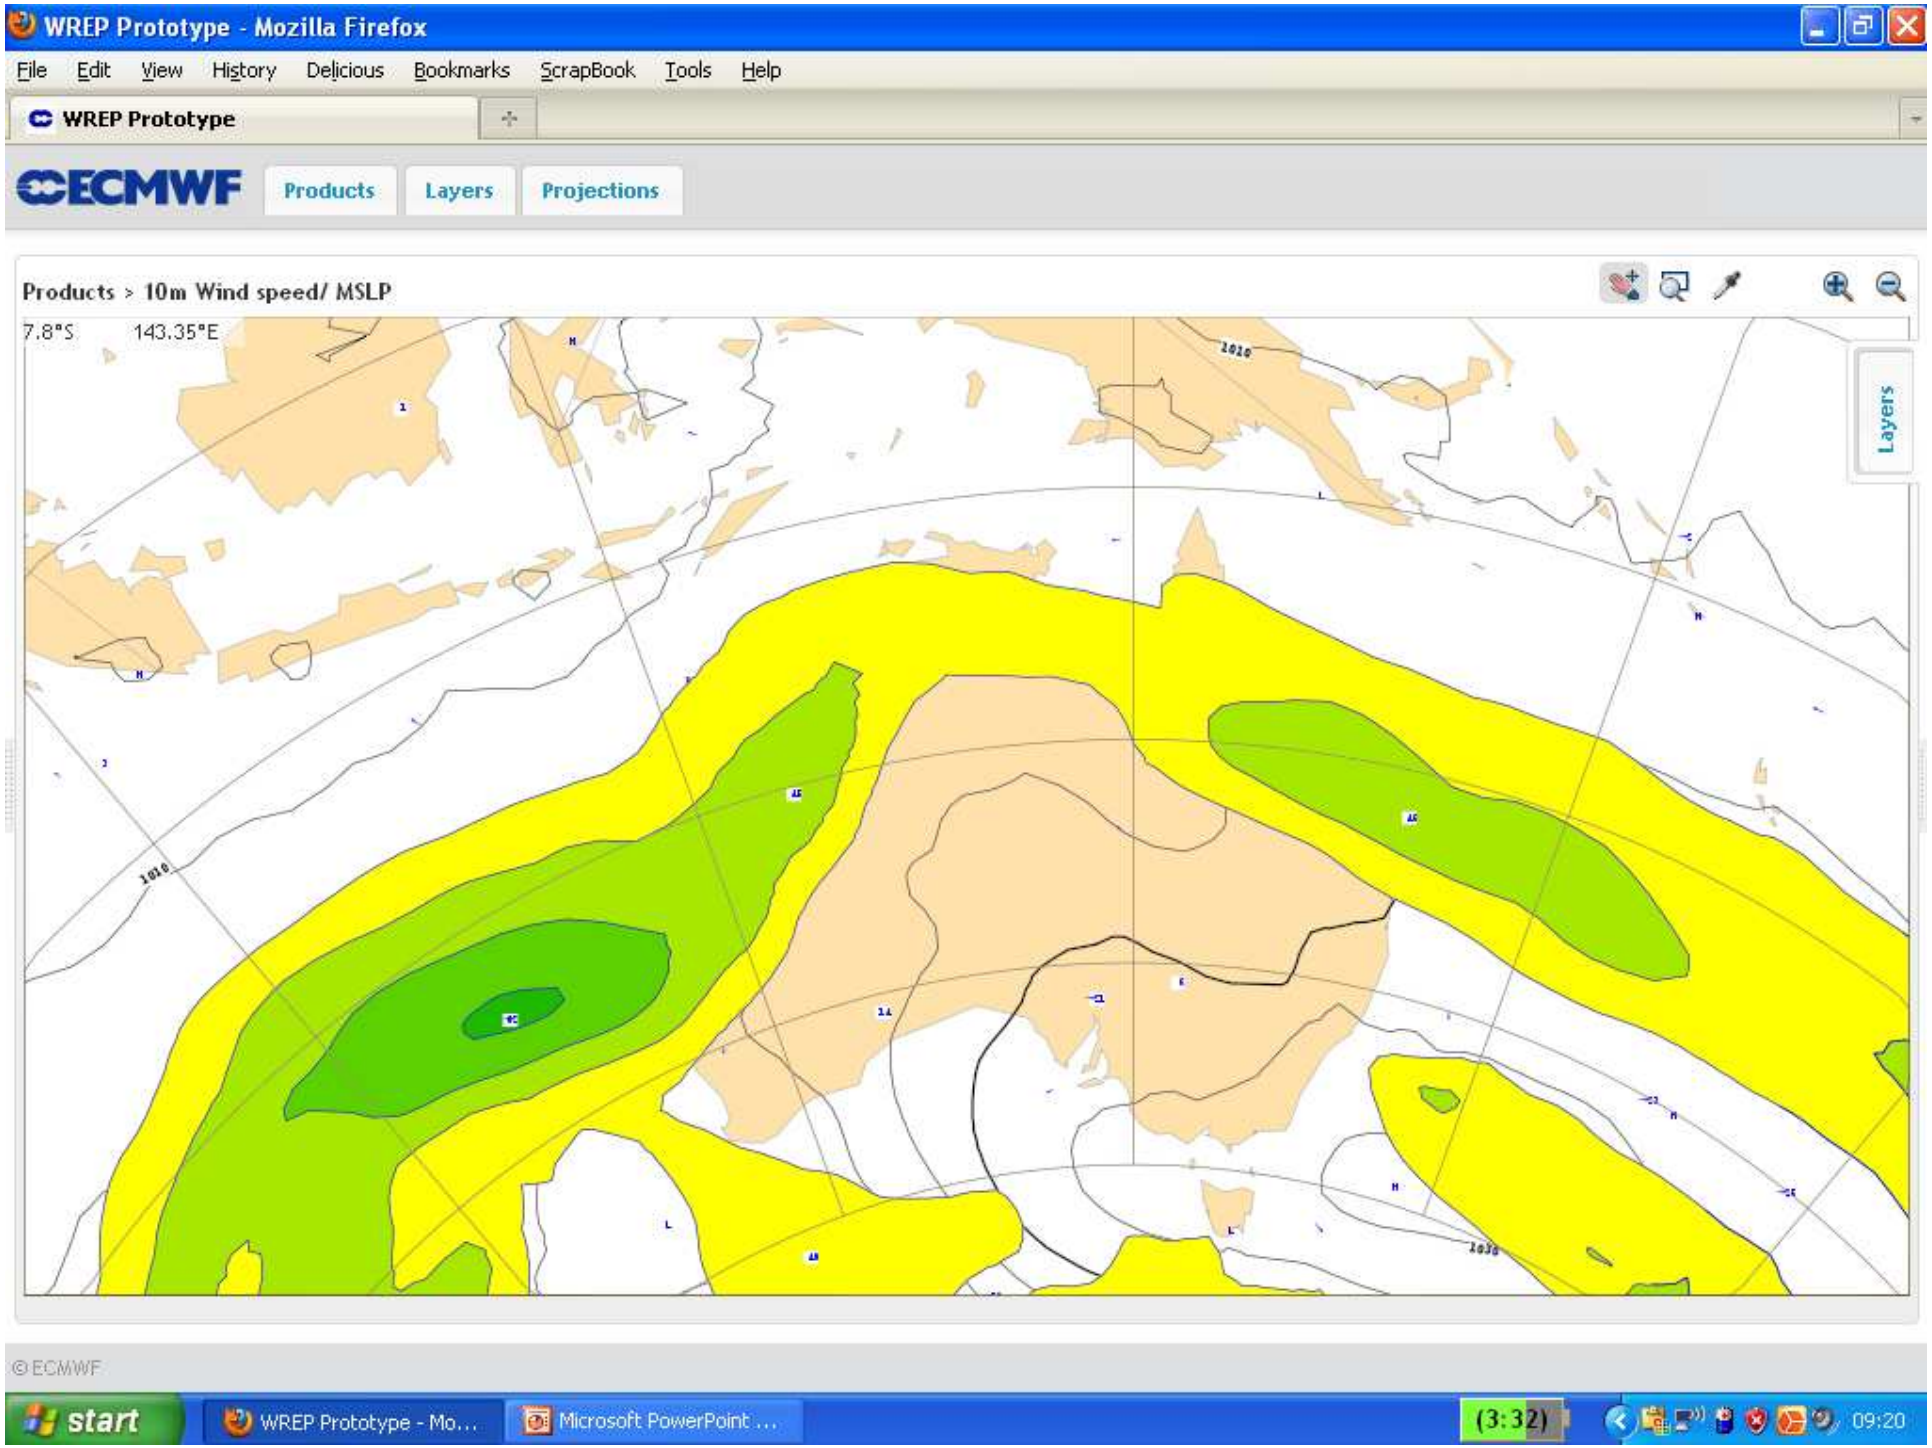
WREP Prototype - Mozilla Firefox
File Edit View History Delicious Bookmarks ScrapBook Tools Help
WREP Prototype
ECMWF
Products
Layers
Projections
Products > 10m Wind speed/ MSLP
7.8°S 143.35°E
1010
1030
Layers
© ECMWF
start
WREP Prototype - Mo...
Microsoft PowerPoint ...
(3:32)
09:20

WREP Prototype - Mozilla Firefox

File Edit View History Delicious Bookmarks ScrapBook Tools Help

WREP Prototype

ECMWF Products Layers Projections

Products > 10m Wind speed/ MSLP

Layers

Probe

Closest grid point:	23.41°S 154.93°E
Mean sea level pressure:	101911 Pa
10m Wind:	-6.29285 $m\,s^{-1}$
200 hPa Wind Speed:	39.7924 $m\,s^{-1}$

Epsgram

Closest EPS grid point: 23.6°S 155.16°E
Deterministic Forecast and EPS Distribution 20091104 0000 UTC

Total Cloud Cover (okta)

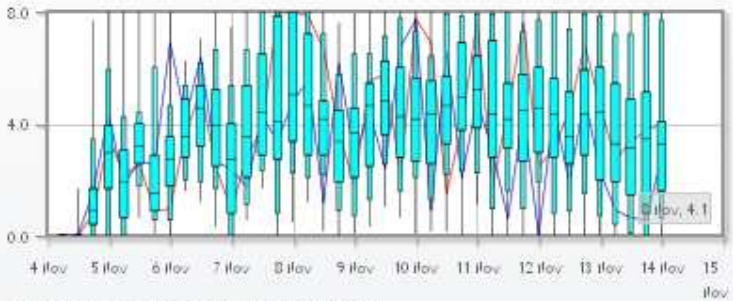

Total Precipitation (mm/24h)

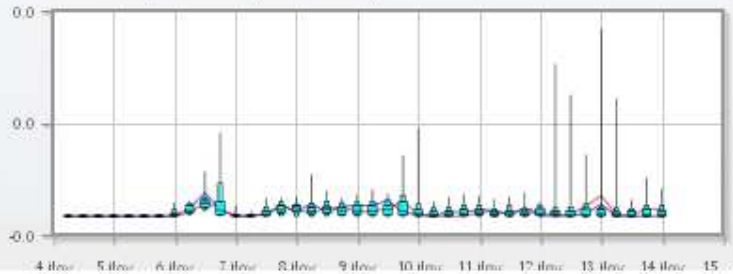

start WREP Prototype - Mo... Microsoft PowerPoint ... (3:30) 09:20

WREP Prototype - Mozilla Firefox

File Edit View History Delicious Bookmarks ScrapBook Tools Help

WREP Prototype

ECMWF

Products Layers Projections

Products > 10m Wind


Simulated image (Met9 - IR)


Spaghetti plot

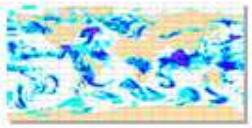
Total precipitation over an interval

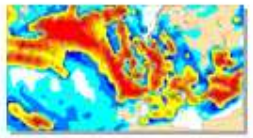
TP / TP probability

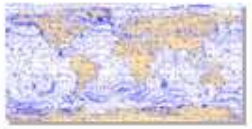
Wind Speed

Epsgram

Closest EPS grid point: 25.39°S 145.44°E

Deterministic Forecast and EPS Distribution 20091104 0000 UTC

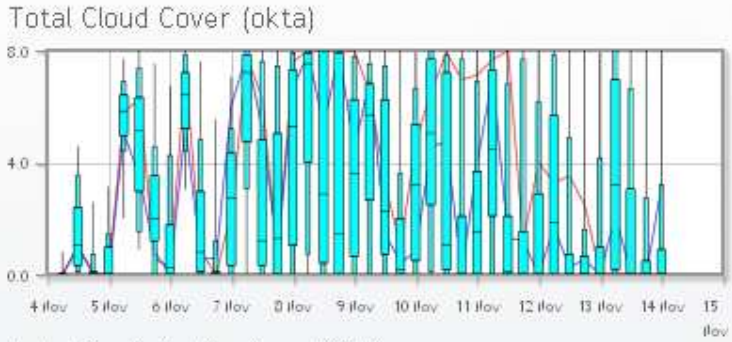

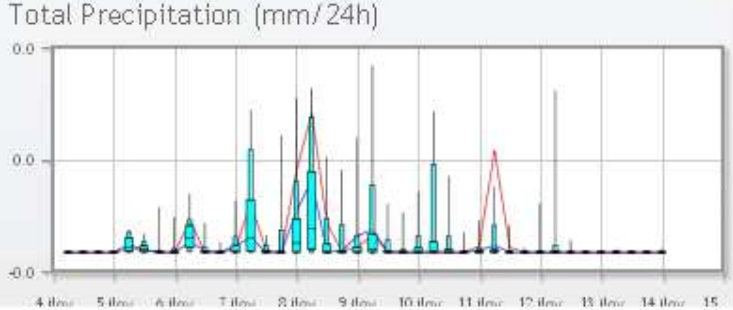

start WREP Prototype - Mo... Microsoft PowerPoint ... (3:01) 09:29

WREP Prototype - Mozilla Firefox

File Edit View History Delicious Bookmarks ScrapBook Tools Help

WREP Prototype

ECMWF Products Layers Projections

Products > TP / TP probability

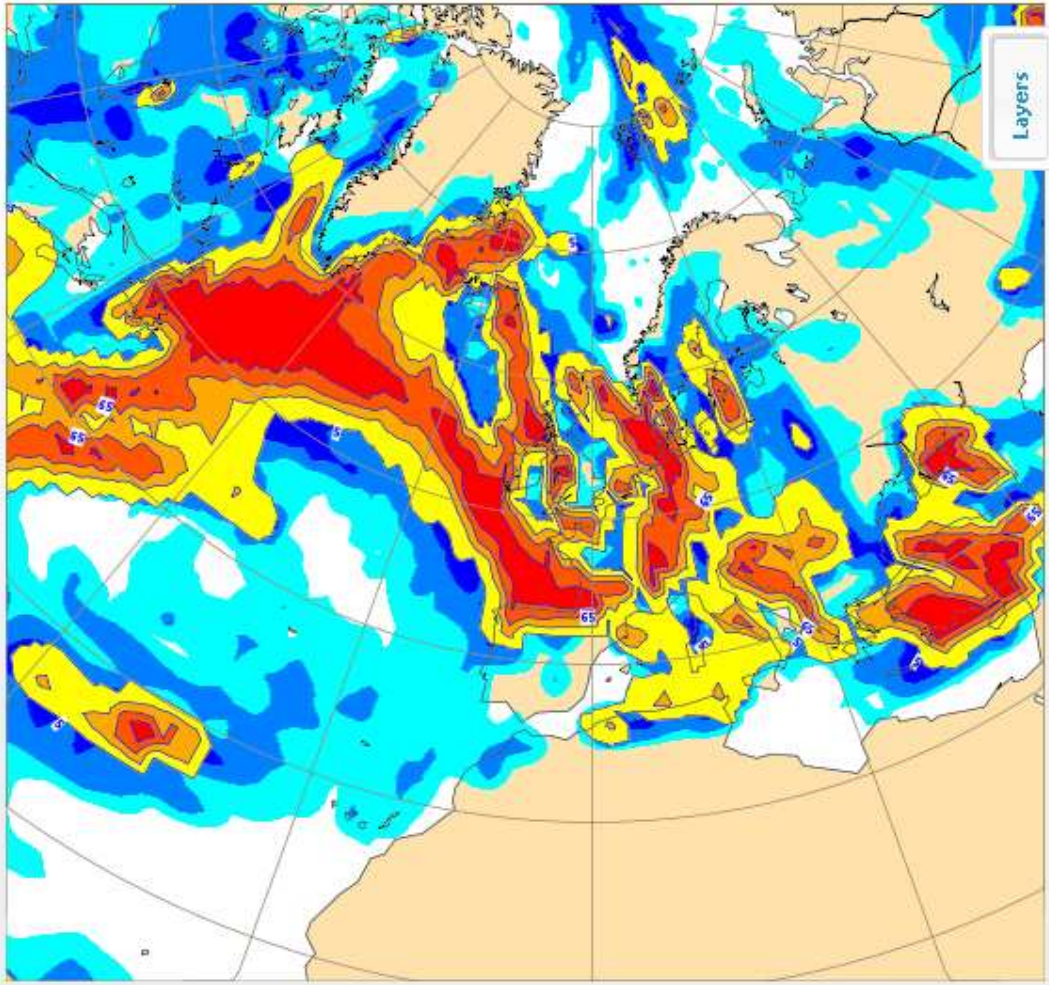

Probe

Closest grid point:	25.23°S 145.37°E
Mean sea level pressure:	101666 Pa
10m Wind:	-3.26648 m s^{-1}
200 hPa Wind Speed:	19.64 m s^{-1}

Epsgram

Closest EPS grid point: 25.39°S 145.44°E
Deterministic Forecast and EPS Distribution 20091104 0000 UTC

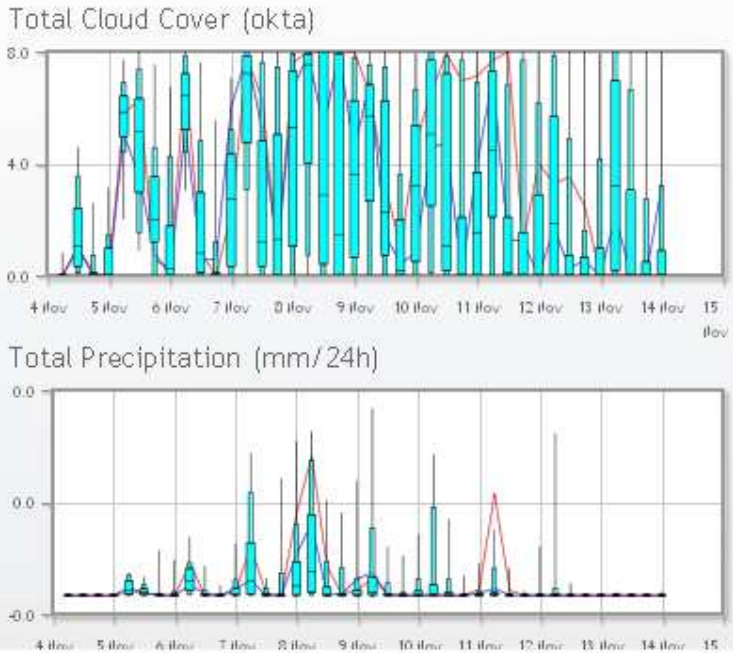

start WREP Prototype - Mo... Microsoft PowerPoint ... (3:08) 09:29

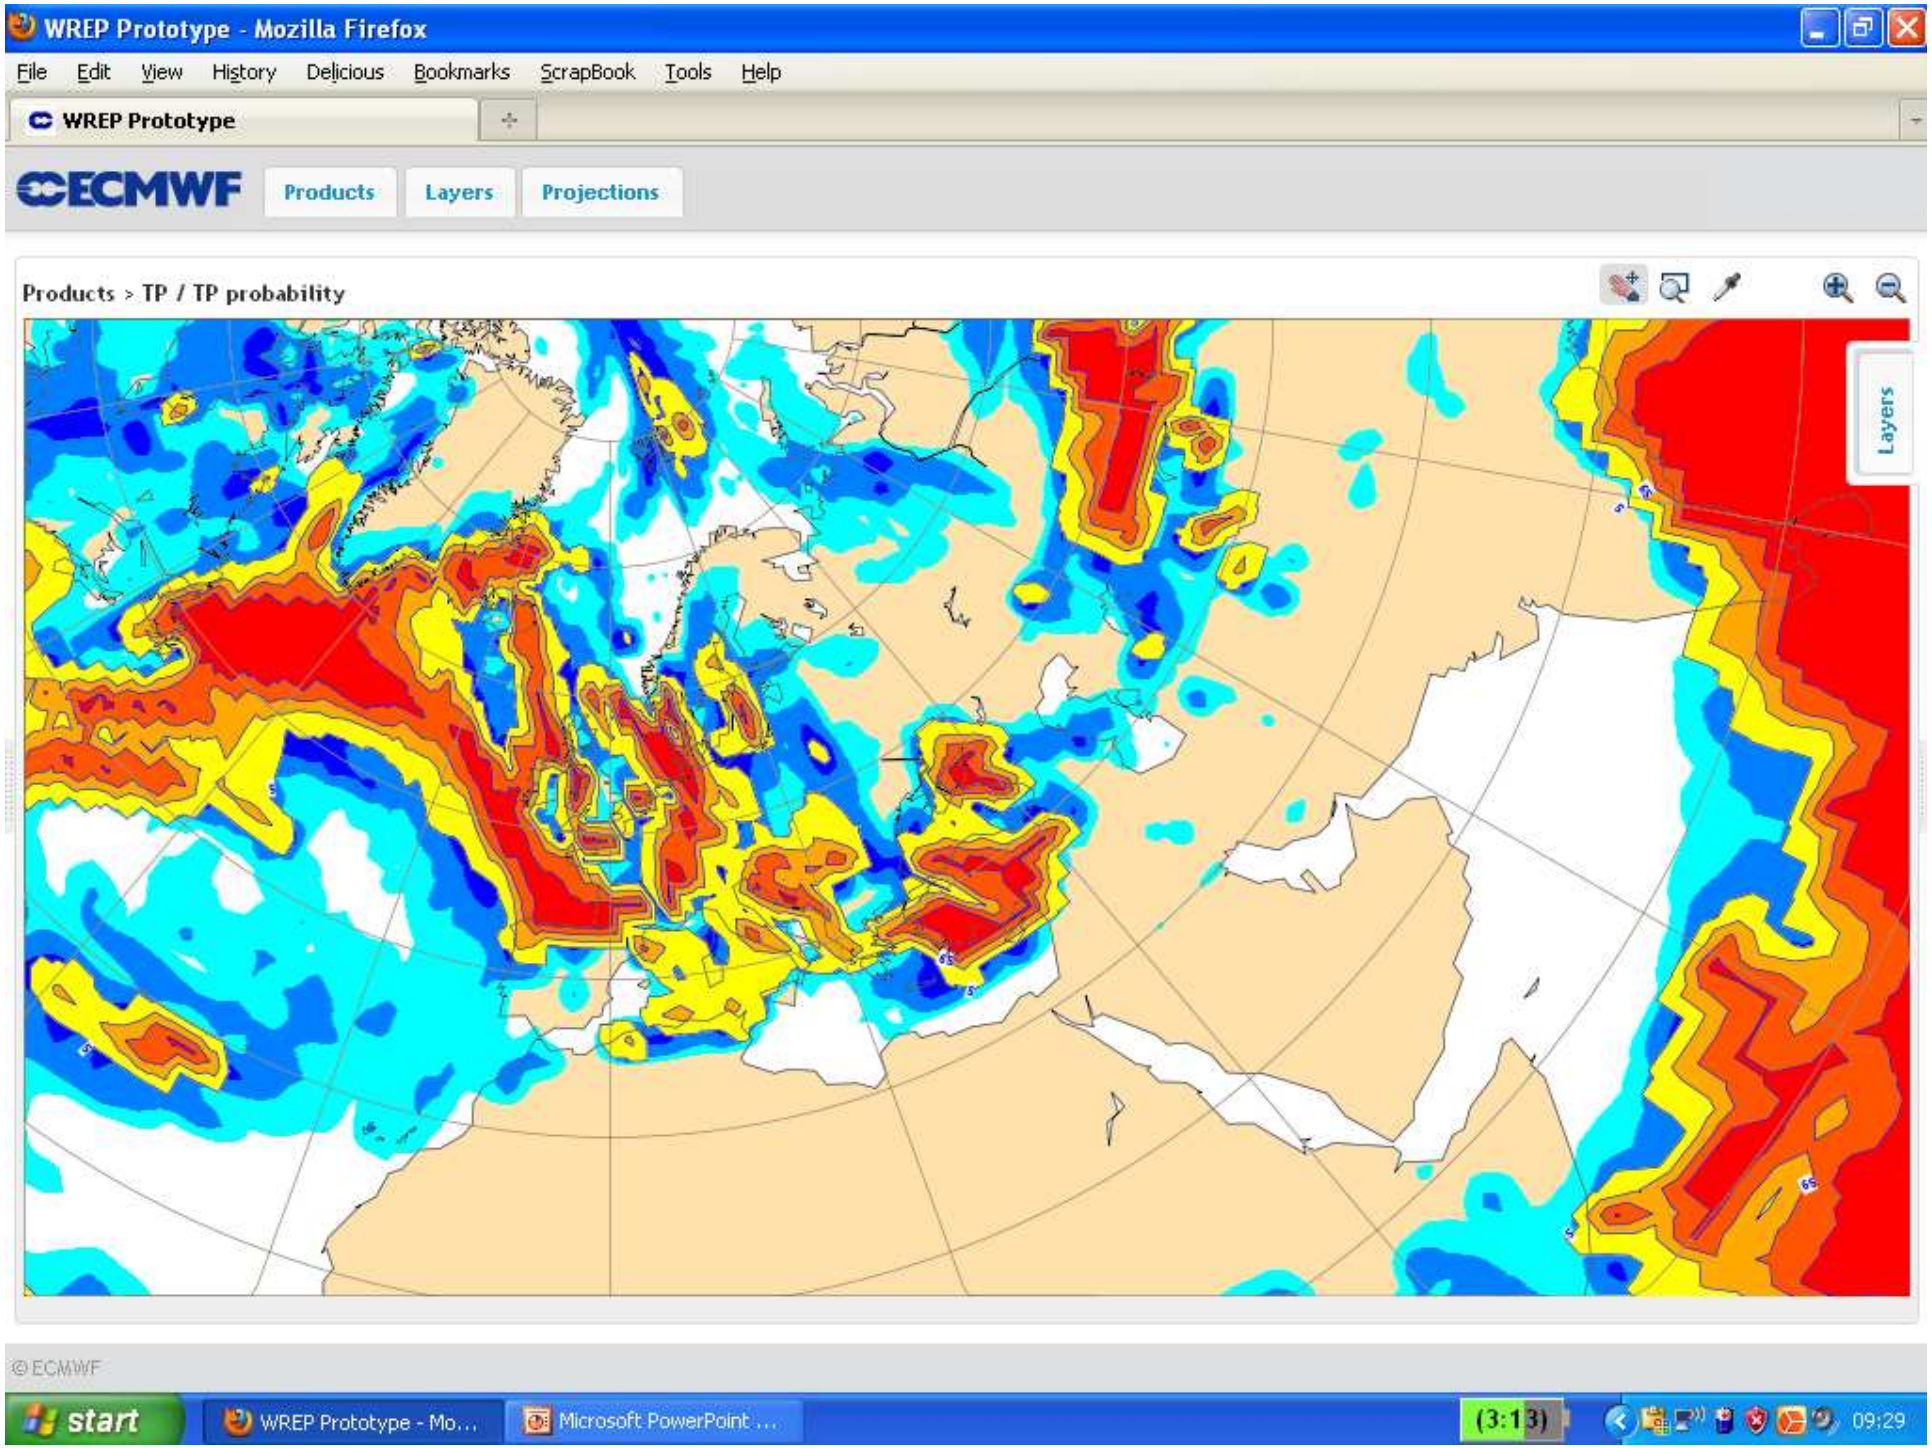
WREP Prototype - Mozilla Firefox
File Edit View History Delicious Bookmarks ScrapBook Tools Help
WREP Prototype
ECMWF
Products
Layers
Projections
Products > TP / TP probability
Layers
© ECMWF
start
WREP Prototype - Mo...
Microsoft PowerPoint ...
(3:13)
09:29

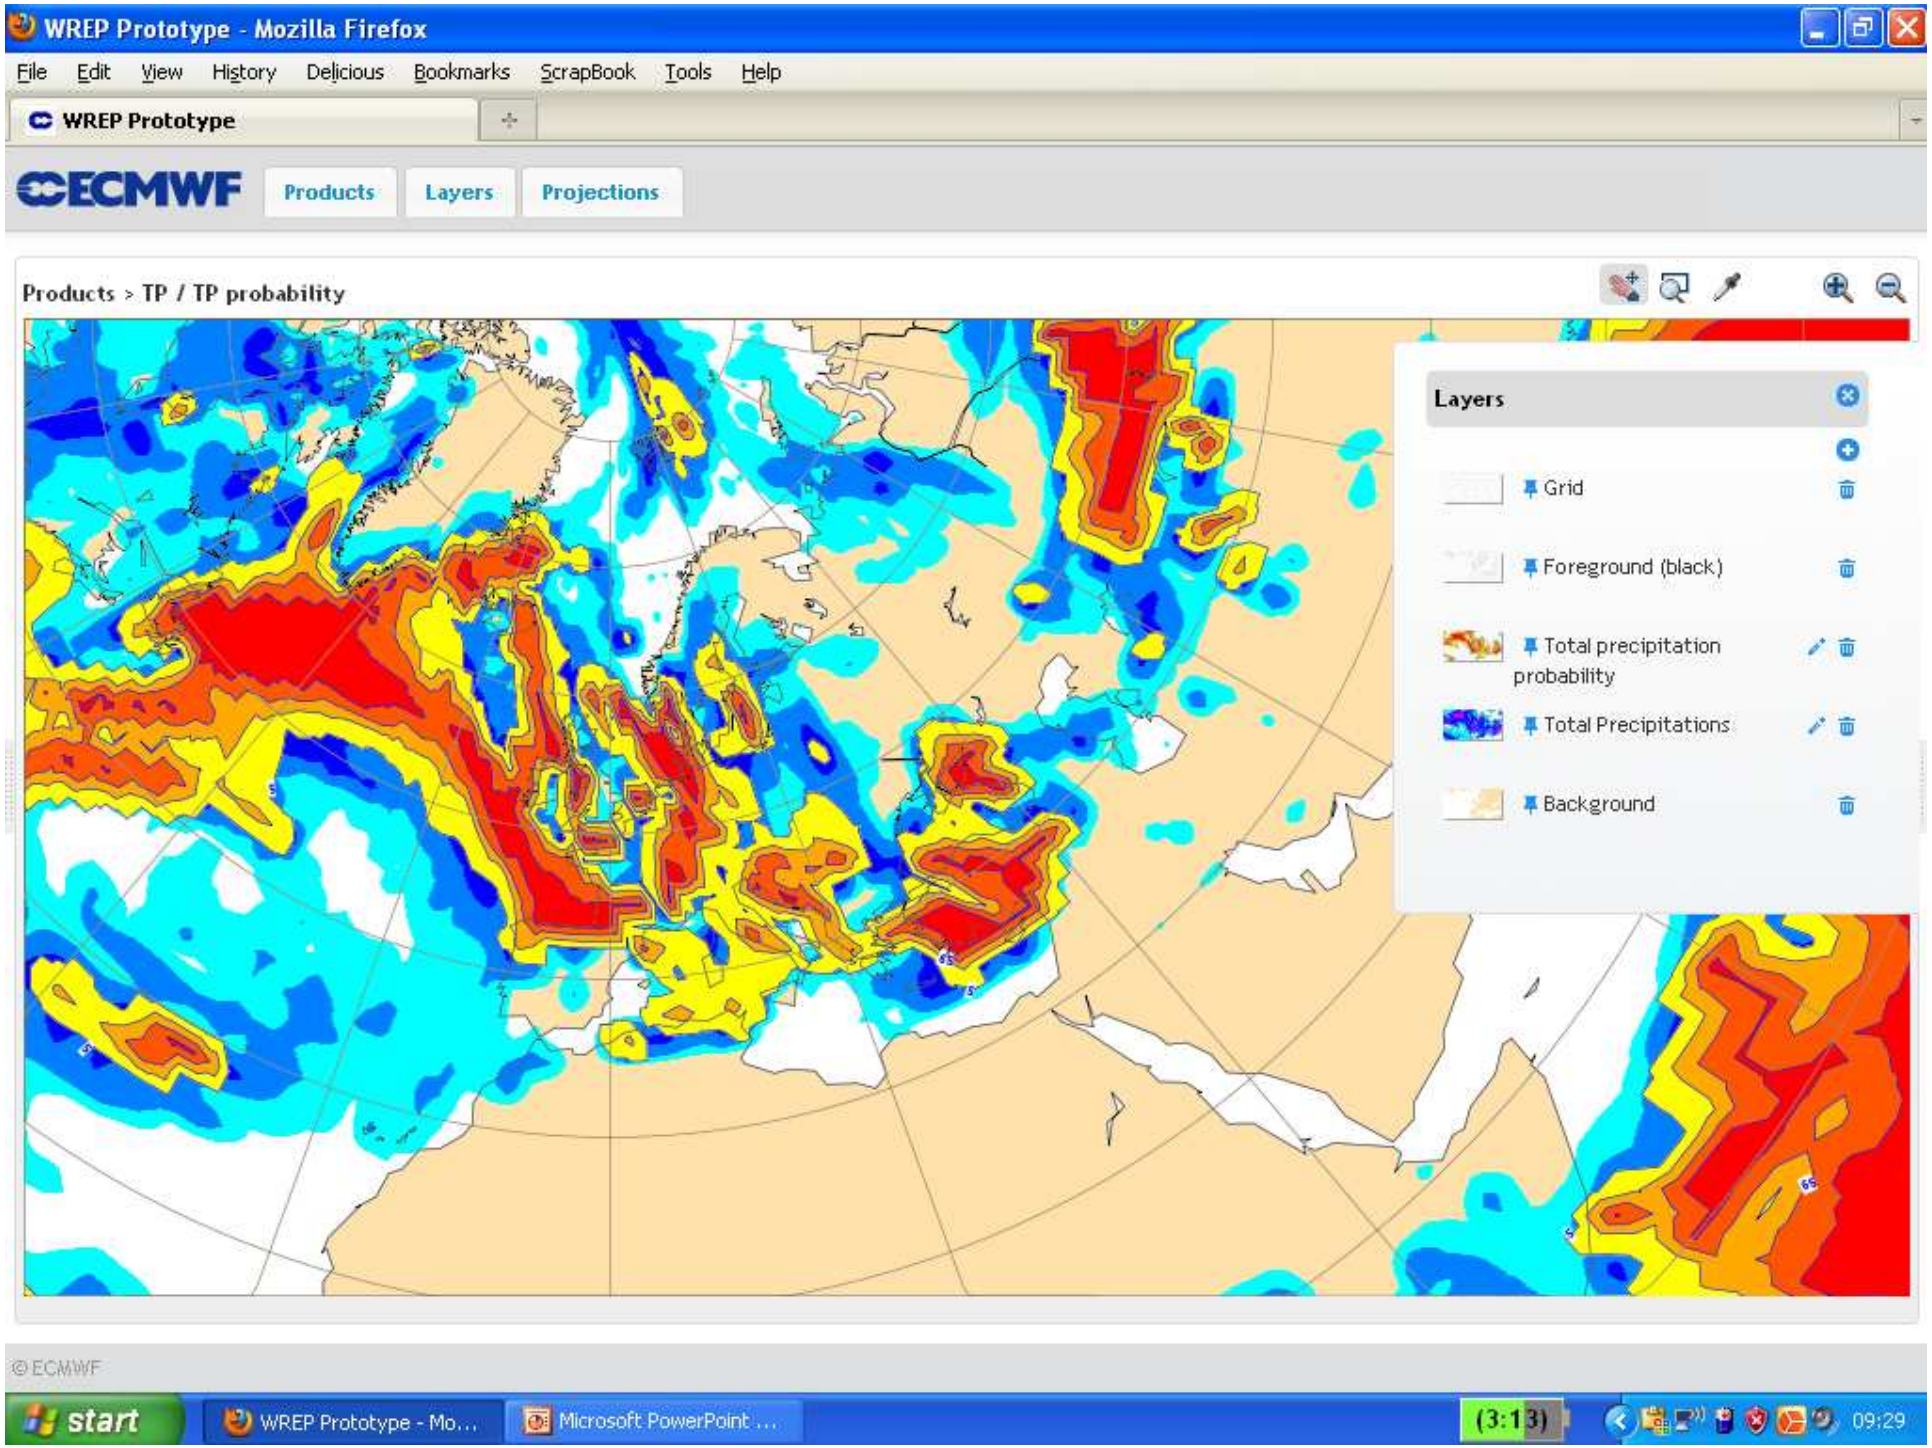
WREP Prototype - Mozilla Firefox
File Edit View History Delicious Bookmarks ScrapBook Tools Help
WREP Prototype
ECMWF
Products
Layers
Projections
Products > TP / TP probability
Layers
Grid
Foreground (black)
Total precipitation probability
Total Precipitations
Background
© ECMWF
start
WREP Prototype - Mo...
Microsoft PowerPoint ...
(3:13)
09:29

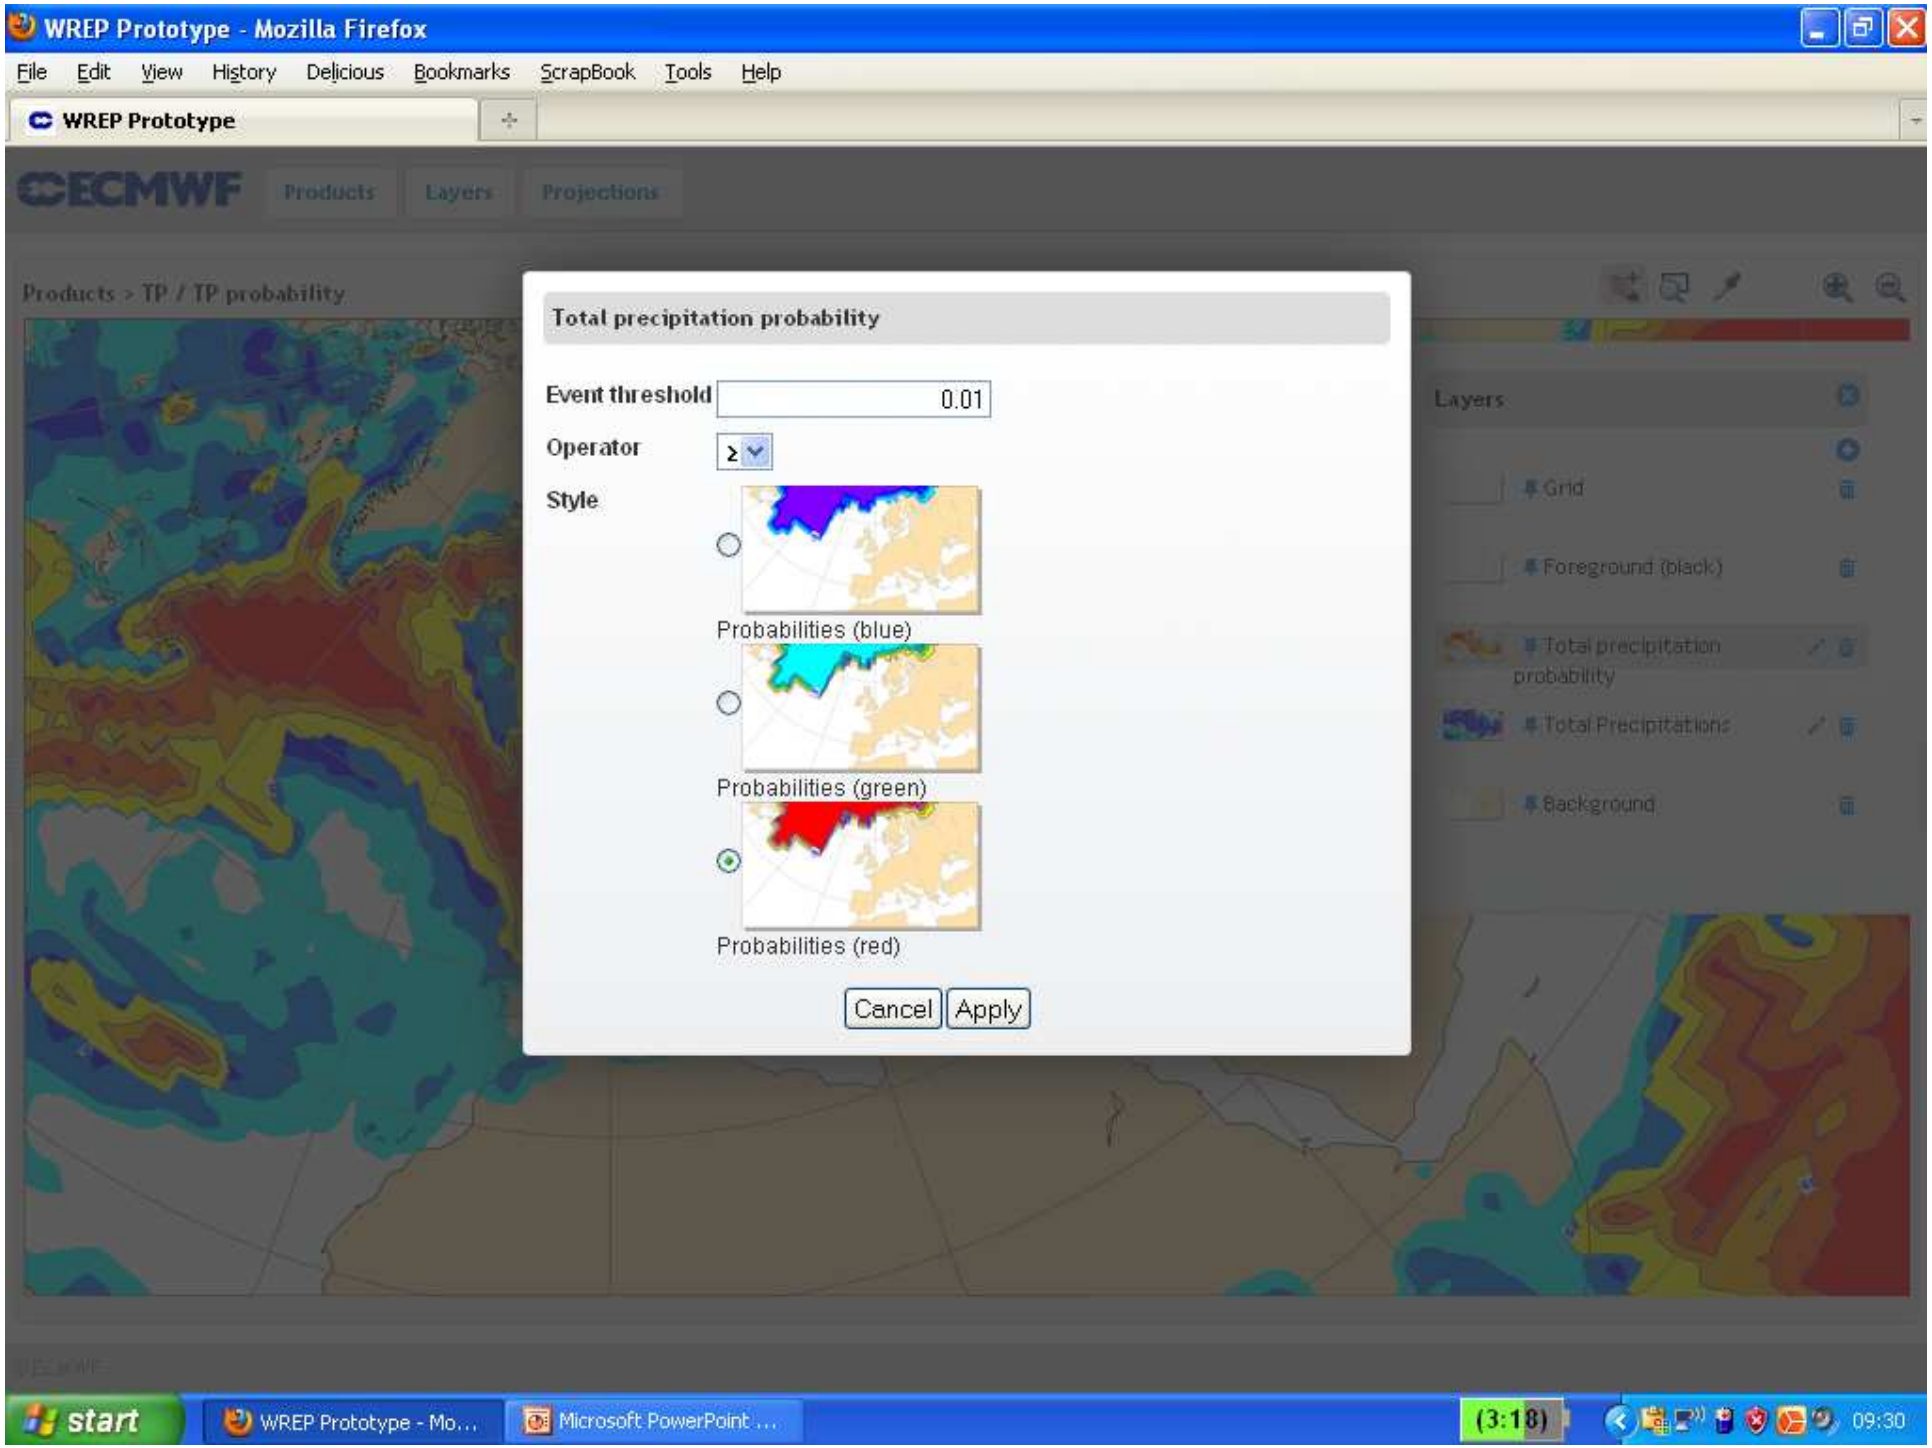
WREP Prototype - Mozilla Firefox
File Edit View History Delicious Bookmarks ScrapBook Tools Help
WREP Prototype
ECMWF
Products
Layers
Projections
Products > TP / TP probability
Total precipitation probability
Event threshold
0.01
Operator
≥
Style
Probabilities (blue)
Probabilities (green)
Probabilities (red)
Cancel
Apply
Layers
Grid
Foreground (black)
Total precipitation probability
Total Precipitations
Background
start
WREP Prototype - Mo...
Microsoft PowerPoint ...
(3:18)
09:30

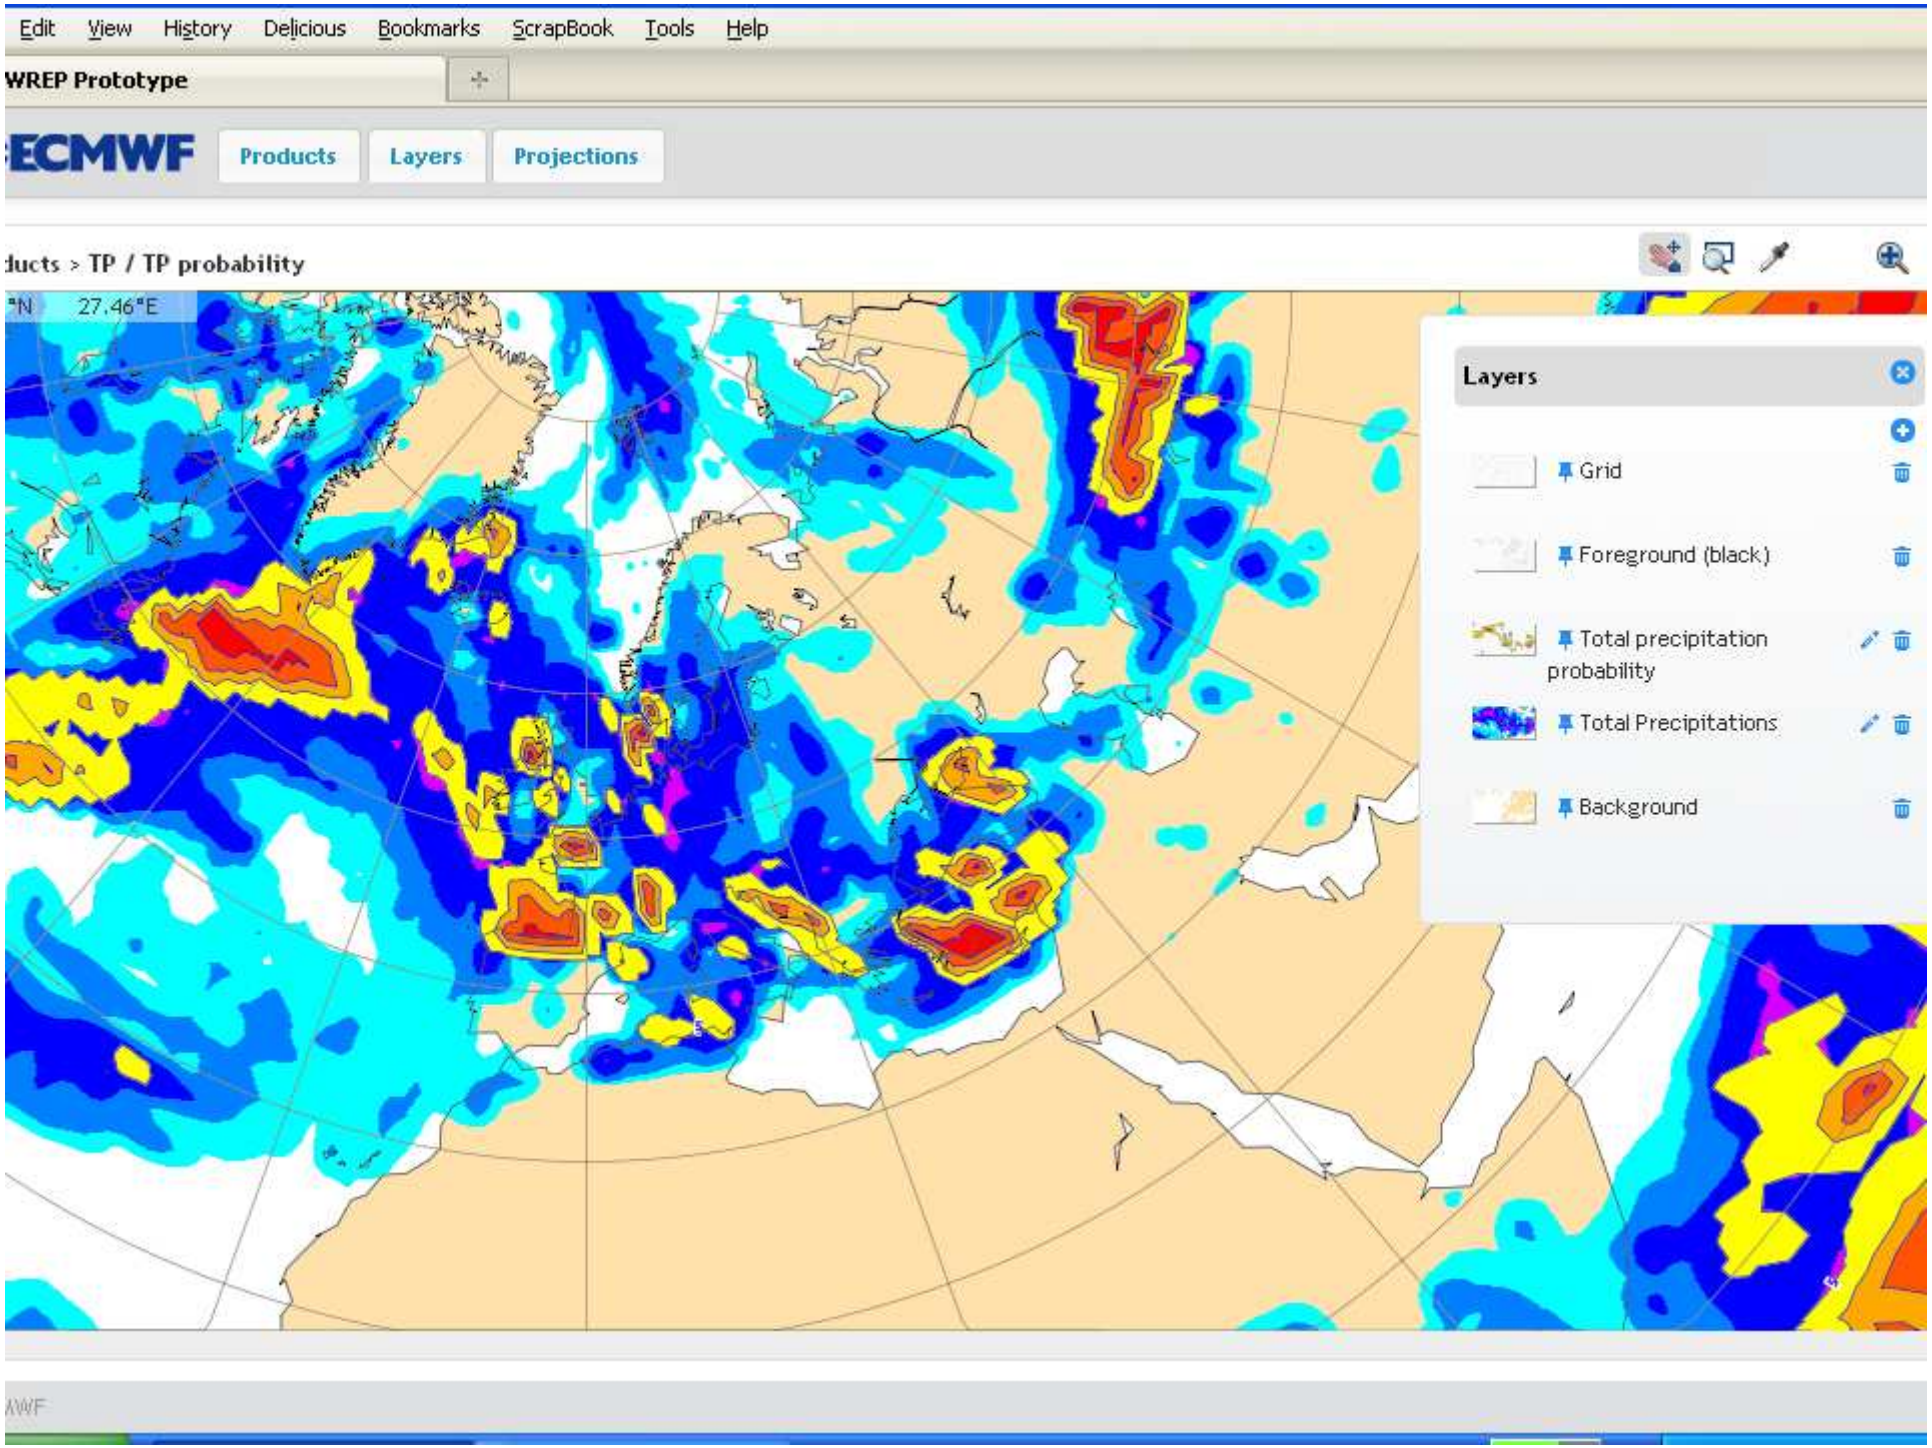
Edit View History Delicious Bookmarks ScrapBook Tools Help
WREP Prototype
ECMWF
Products
Layers
Projections
ducts > TP / TP probability
°N 27.46°E
Layers
Grid
Foreground (black)
Total precipitation probability
Total Precipitations
Background

Prototype: Catalogue browsing

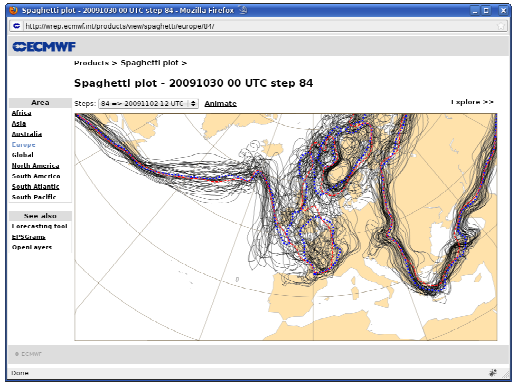

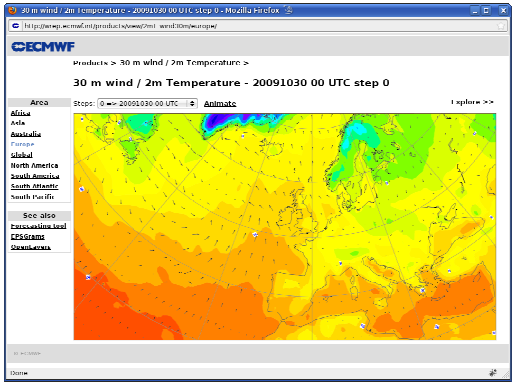

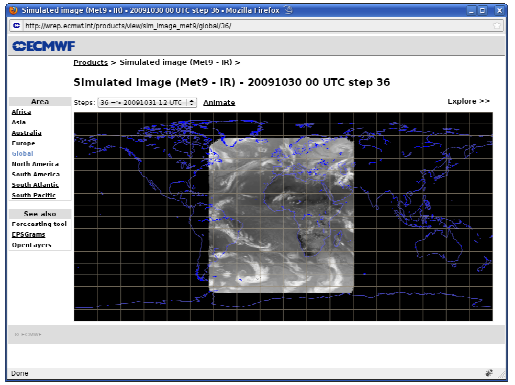

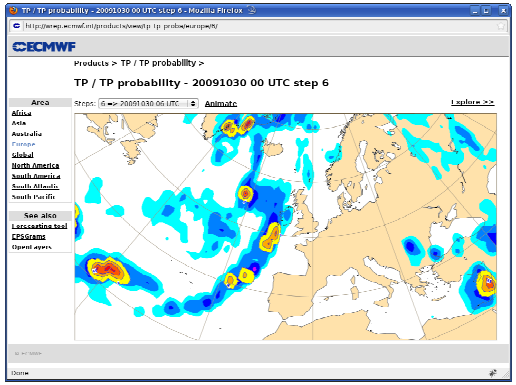

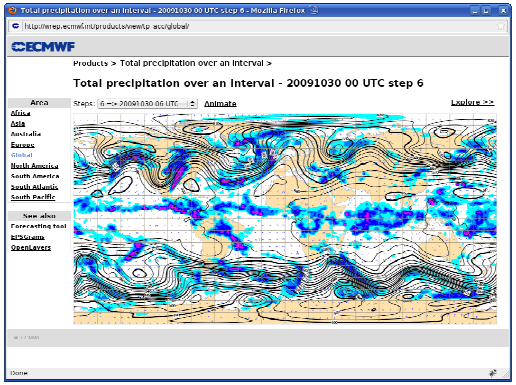

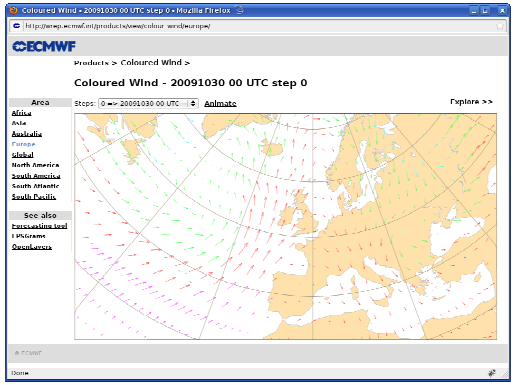

- Browsable catalogue – Link to Forecaster tool
- Limited interactivity – Preset number of projections, animation
- Similar to current web catalogue, but use the WREP infrastructure

Prototype: OGC Web Map Services

- Aim: to make it possible to embed ECMWF products directly in the forecasters' workstations
- On top of WREP infrastructure:
  - "GetCapabilities document" build dynamically from product catalogue content
  - Layers are created on-demand

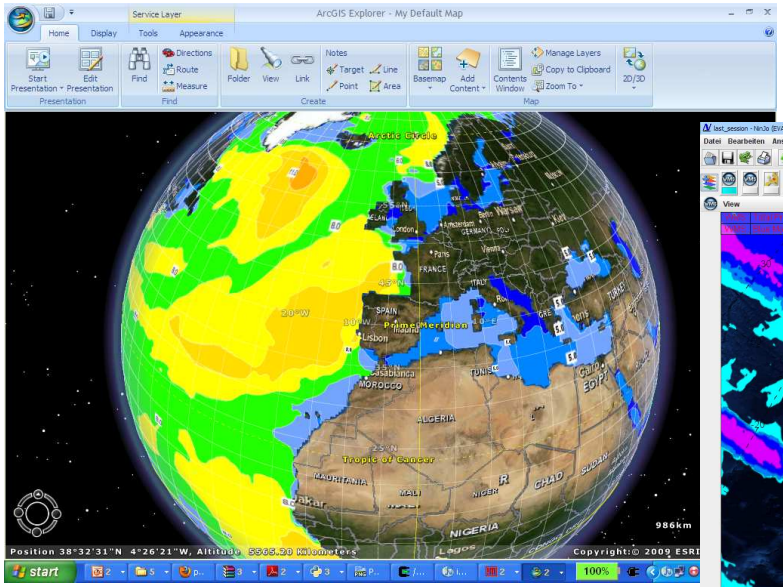

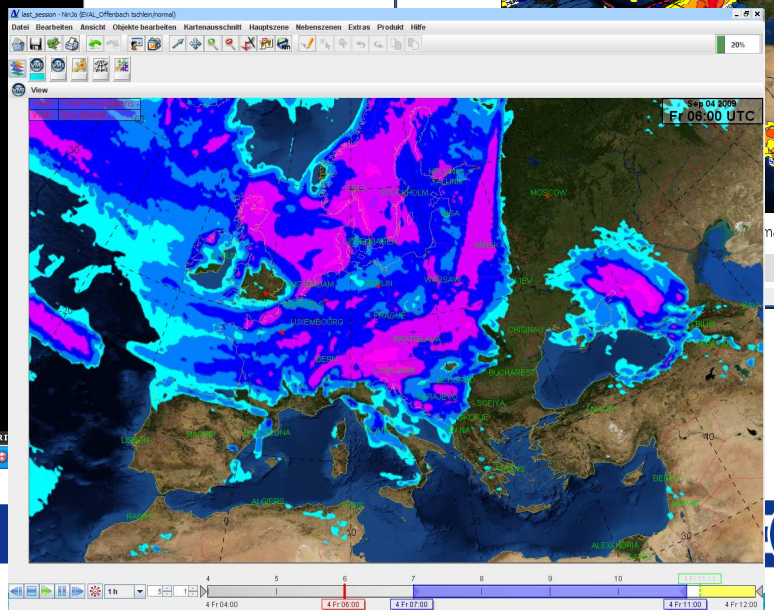

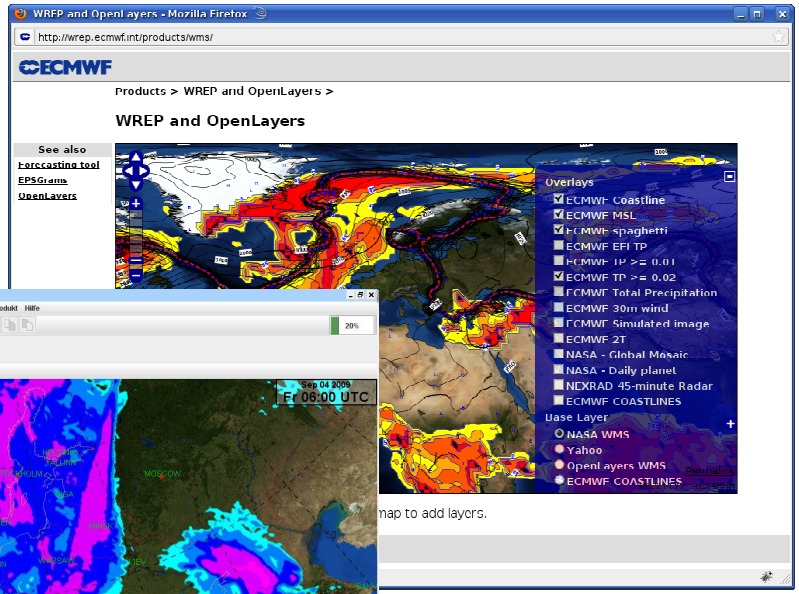

Overlaying layers

Rain in New Orleans

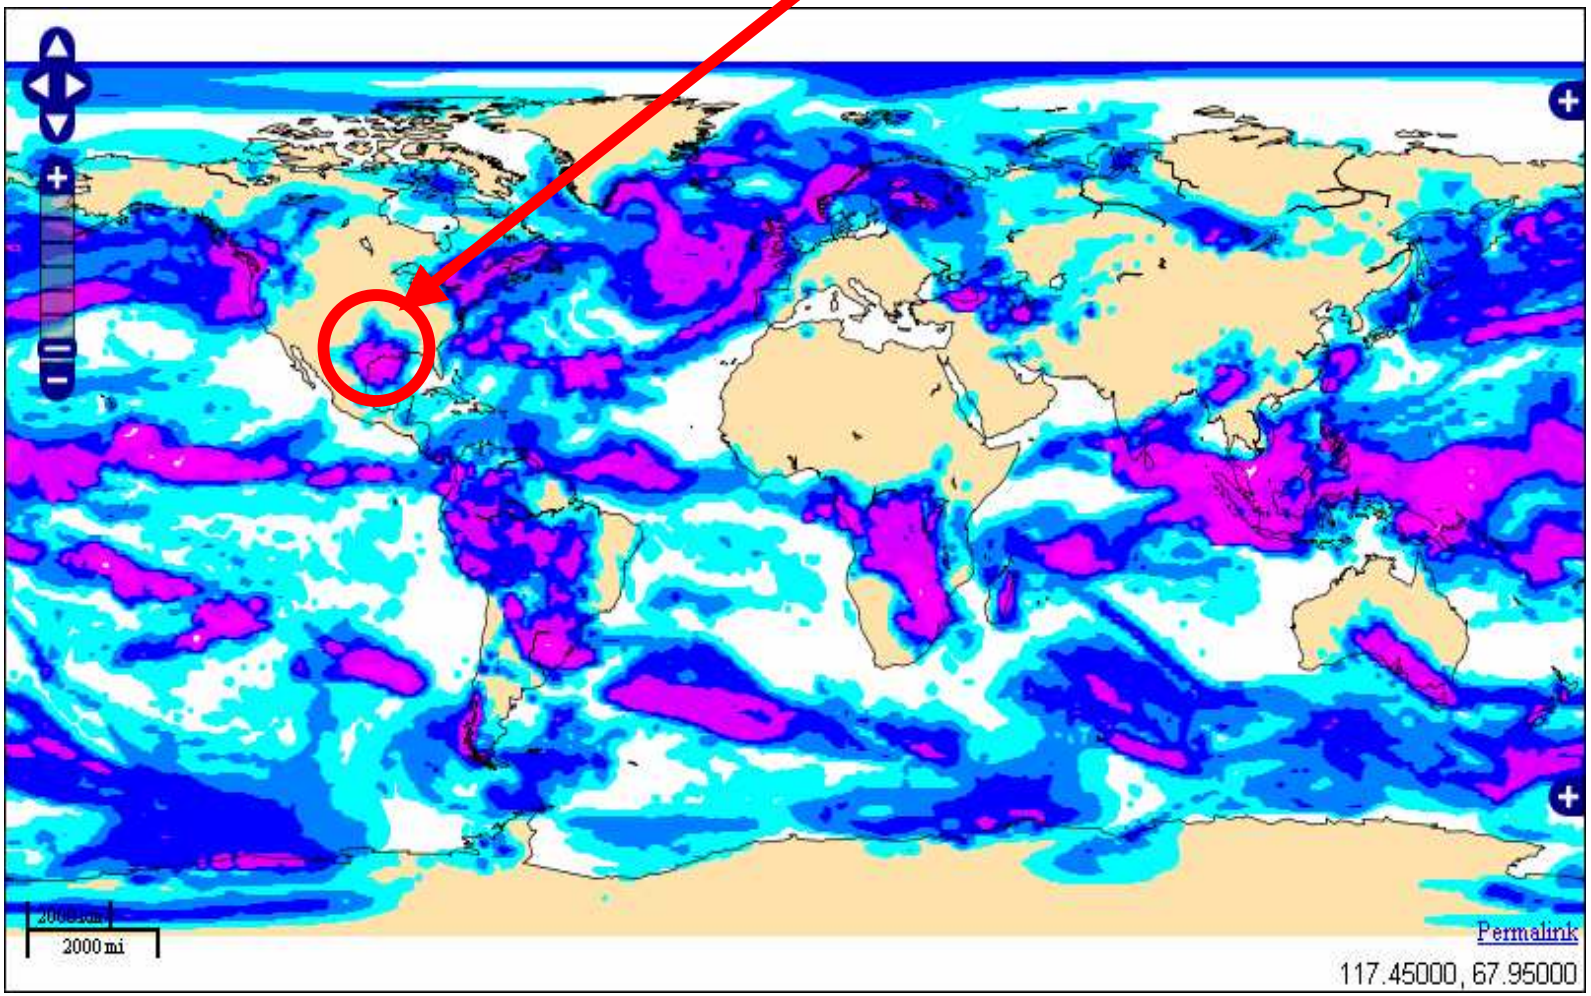

Overlaying layers

Oops, no rain in New Orleans

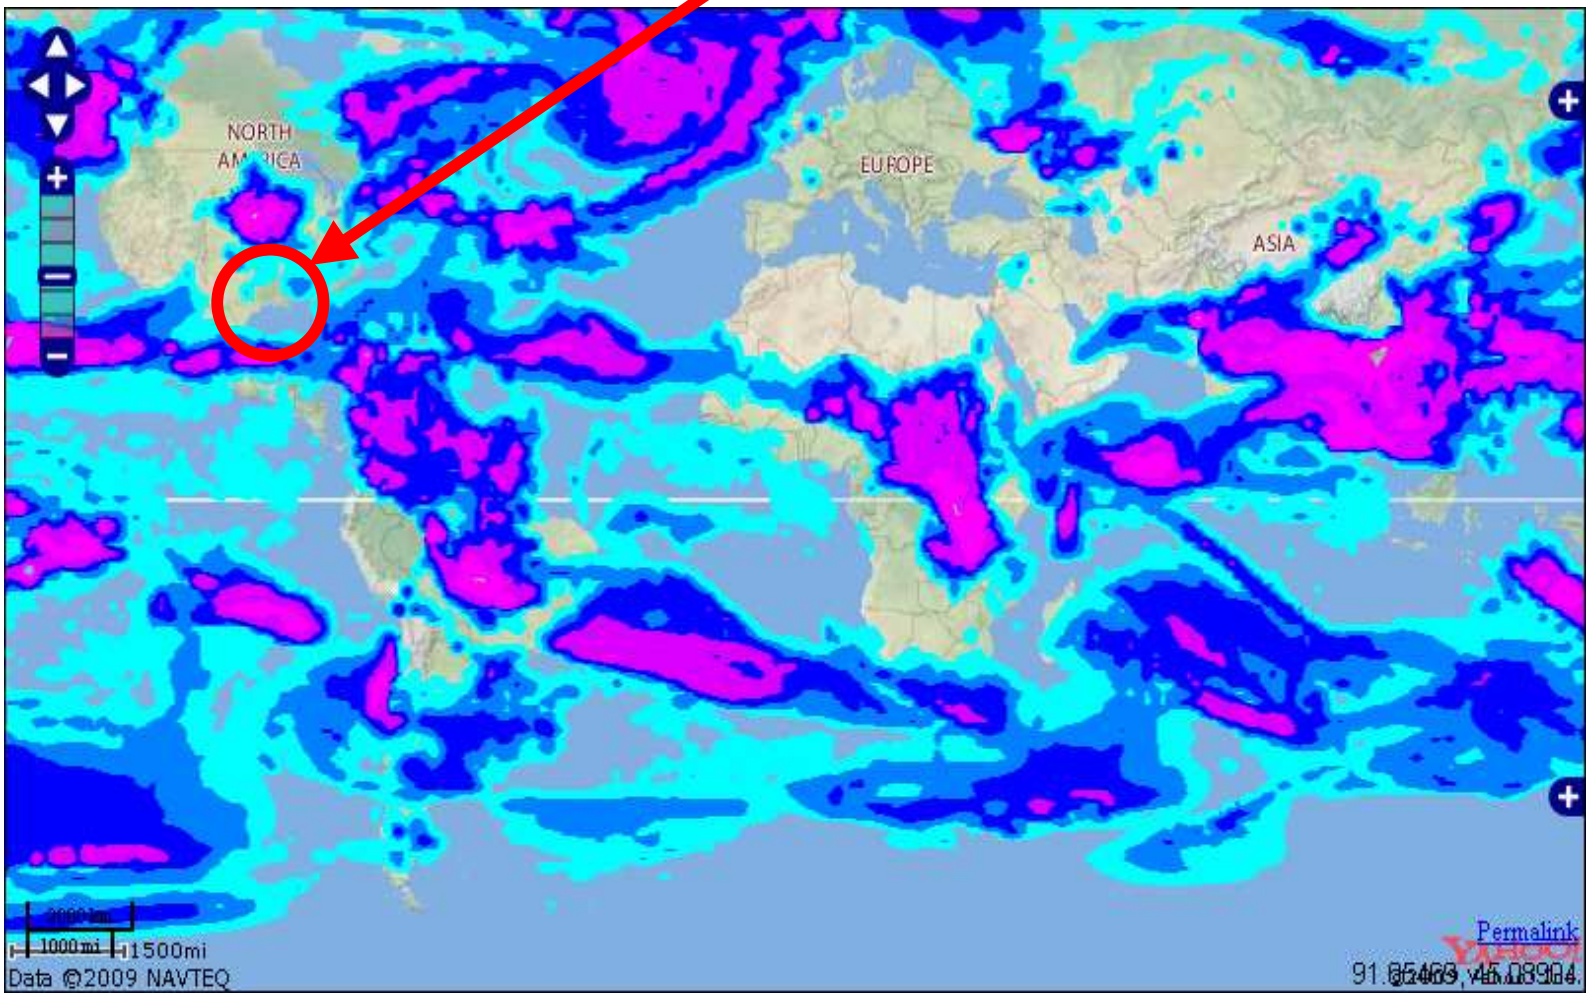

Danger

- Not all WMS servers respect the requested CRS
- WMS client cannot check returned information
  - PNGs are not geo-referenced
- Incorrect information leads to wrong decisions

- How can we ensure that layers are overlaid correctly:
  - In space …
  - … but also in time (sharing the same time dimension)

Tiling and time

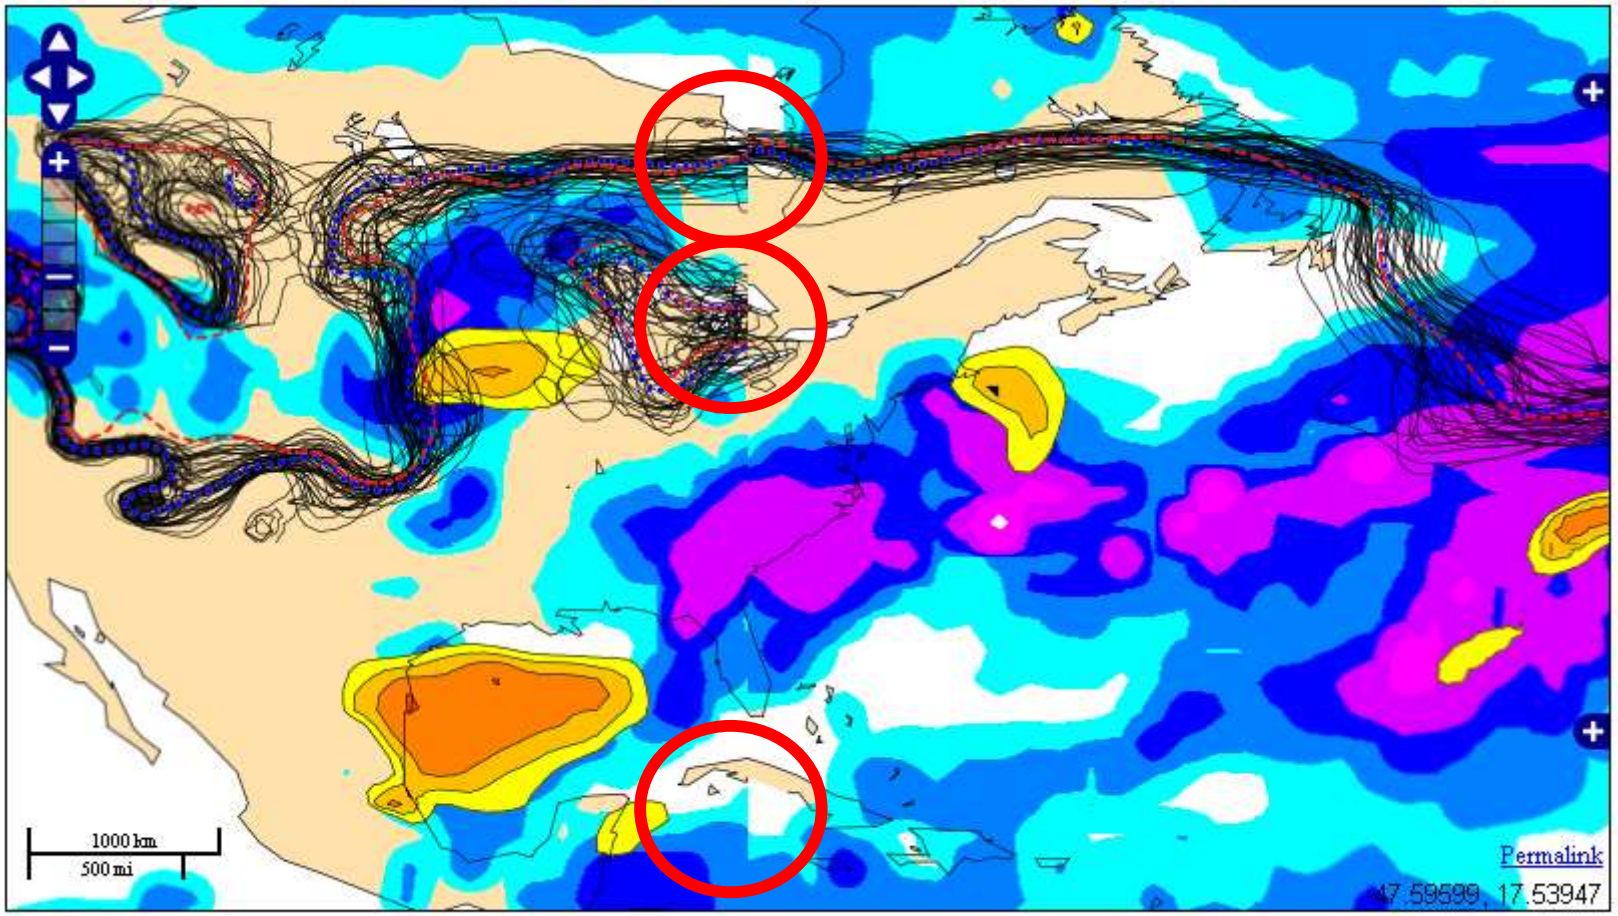

Meteorology is multidimensional

- 2 time dimensions, 3 space dimensions, parameter, model, ensemble number, …
- Layers need to be parameterised (using <Dimension>)
- Example: event probability threshold (e.g. probability of temperature less than a given threshold)

```
<Dimension name="threshold" units="K"/>
<Extent name="threshold" default="273.15">250/400/0</Extent>
```

  - Note that the operator (greater than, less than) must be specified with another <Dimension>
- Other examples: isoline selection in spaghetti plot, percentiles
- GetFeatureInfo: not only field value, but also:
  - Vertical profile, tephigram
  - Time series, metgram, epsgram
  - How does we control which feature info?

Conclusion: fully functional proof of concept

- All products created “on-demand” (2D maps, EPSgrams)
- Zoom, pan, overlay
- Customisation: setting of probability thresholds, contouring
- Browsable catalogue
- Initial user interface
- OGC Web Map Service (WMS)

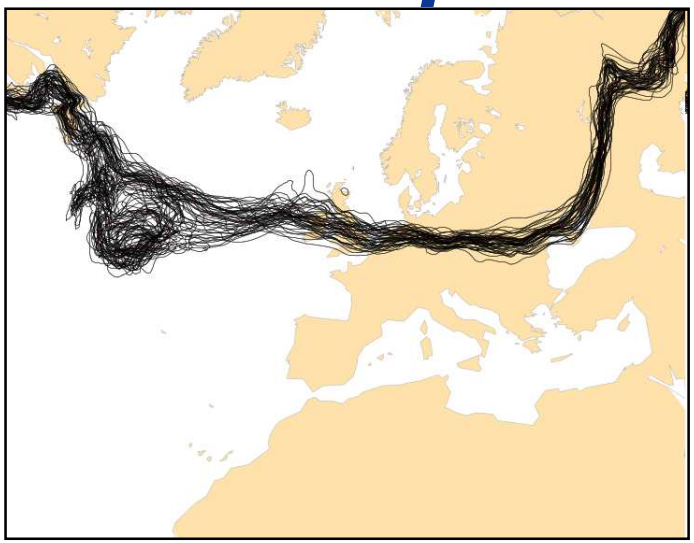

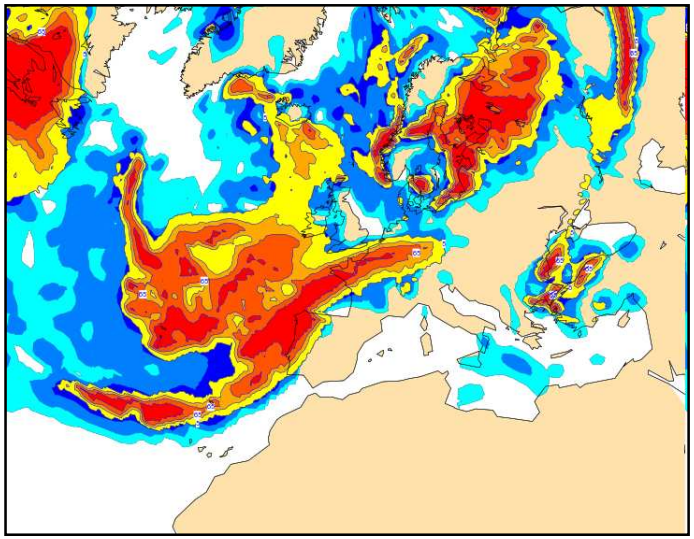

ECMWF

Conclusion: using OGC, challenges

- Meteorological data has many more dimensions than other geospatial data
  - This needs to be captured by the standards
- Moving from testbeds to using OGC standards in operations is challenging
  - Find ways to ensure that the correct information is provided
  - Find ways to ensure that the information is used correctly
  - Implement data policies and access control
- We need to agree on best practices and guidelines

Thank you