



Norwegian
Meteorological
Institute

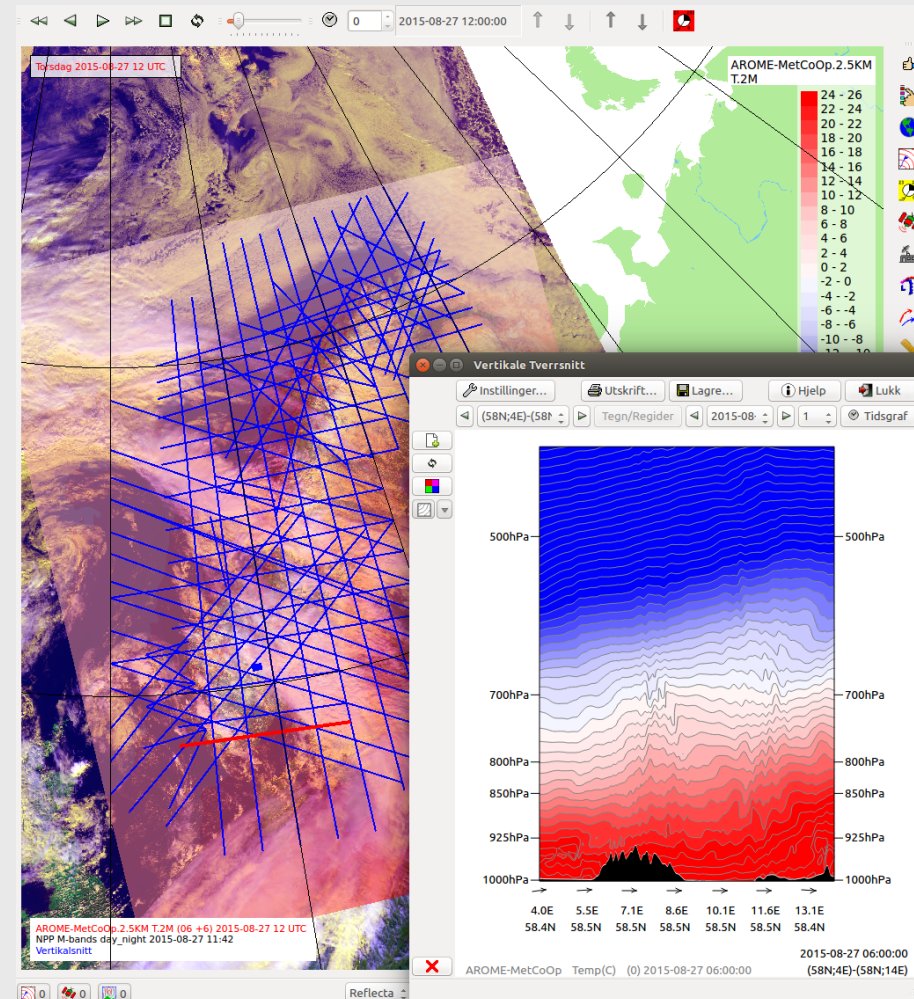
Experiments on Splitting User Interface and Data Handling in Diana

MET Norway / Alexander Bürger

15.10.18

Digital Analysis

- meteorological workstation software
 - MET Norway, SMHI, ...
 - forecasters + researchers
 - GUI and batch versions
 - perl interface
- implemented using C++ + Qt
- open source:
github.com/metno/diana
- connected to other forecaster applications via network protocol
 - ted, modfly, mimage, ...

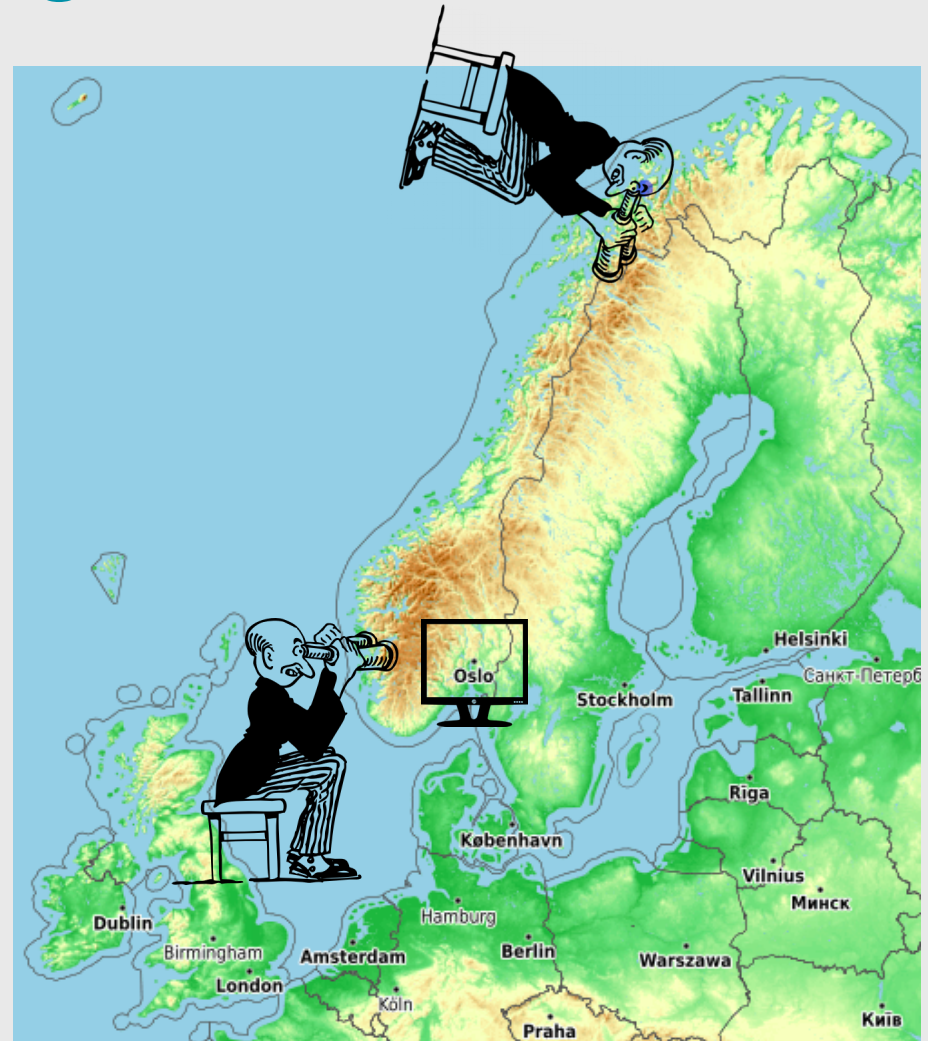


Handling huge model outputs

- model size is continuously increasing
- forecasters cannot wait
 - pilots or journalists want answers now
 - efficiency has to increase
 - waiting is boring and annoying

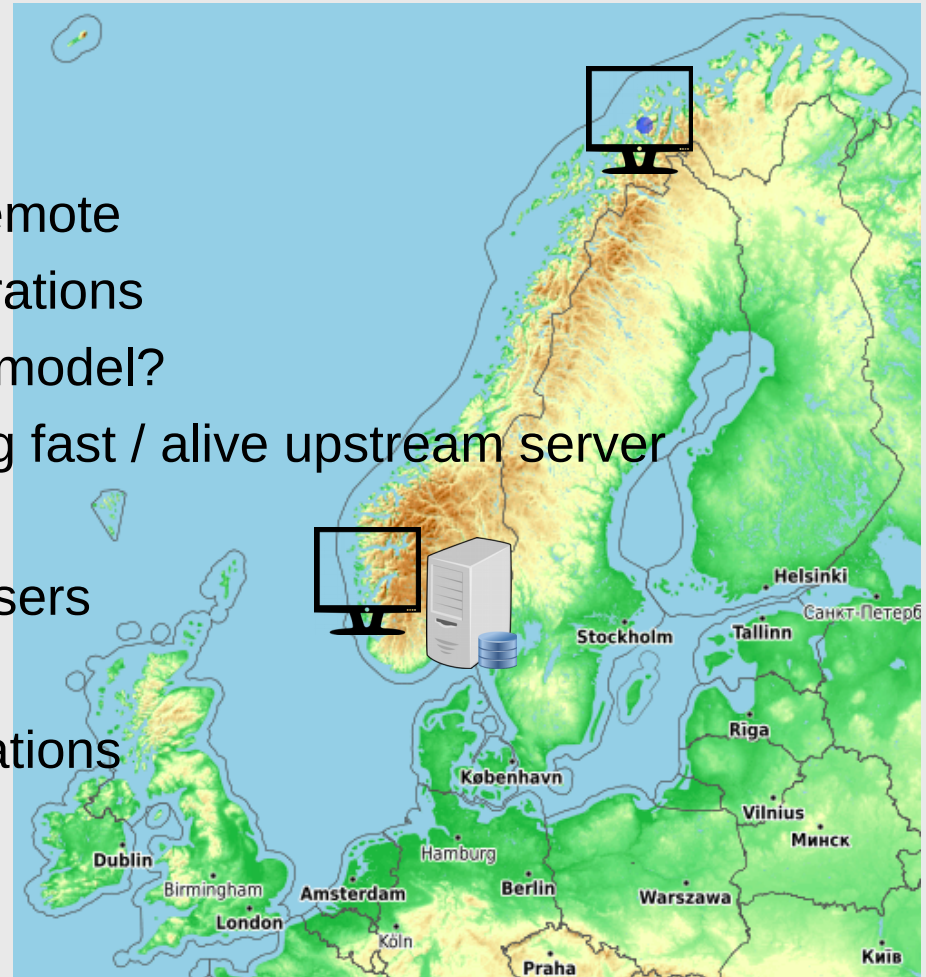
Today, everything is remote

- complete diana application runs remotely on dedicated servers
 - VirtualGL + ssh
 - slow startup, big latency for GUI elements
 - difficult integration (sound, printers, “disk”)
 - everything breaks when the ssh tunnel breaks
 - it can hang because of network, io, bug



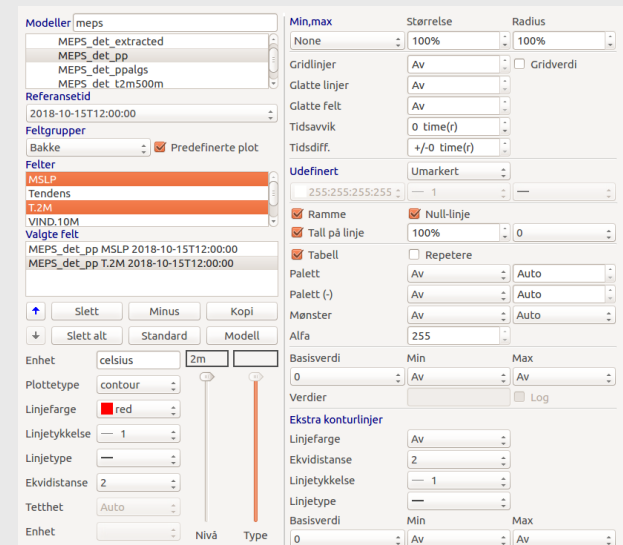
Partly local, partly remote?

- idea: split
 - GUI local
 - data servers (possibly) remote
- many possible server configurations
 - several servers, one per model?
 - caching proxies, selecting fast / alive upstream server
 - ...
- one server handles multiple users
 - better caching
 - fewer reprojection calculations
 - ...



Why not WMS / WPS / ...?

- typical use
 - either **known** model + reference time + field + style
 - or **step by step** model, reference time, field, style
- WMS: huge capabilities document
 - 50 models
 - each has 20 reference times
 - each has 50 fields
 - (each field has 50 style options)
 - 50000 layers
 - how often to reload?
- WPS: maybe, later
- possibility for WMS/WPS frontends
- evolution not revolution, for now
 - stay close to current implementation for the experiment

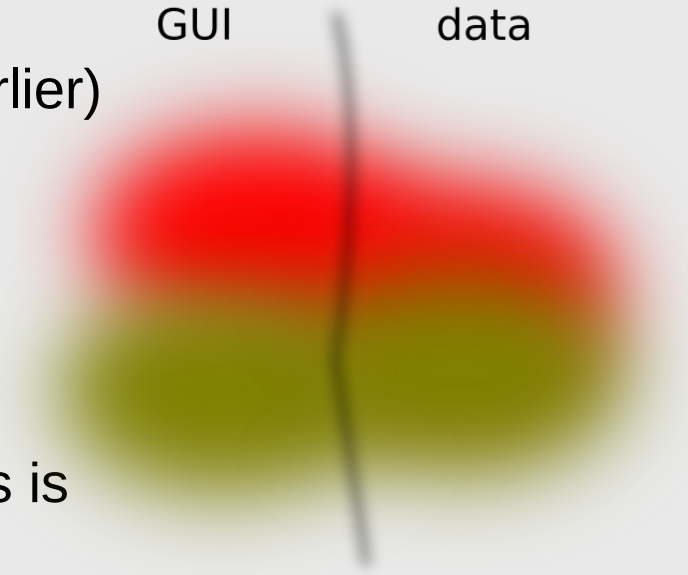


All beginnings are difficult

- Diana is not “new” (born 1999 or earlier)
- lots of good stuff
- lots of mixed responsibilities
- lots of preparations before starting splitting-ui-and-data experiment
 - even if the experiment fails, this is good


GUI

data



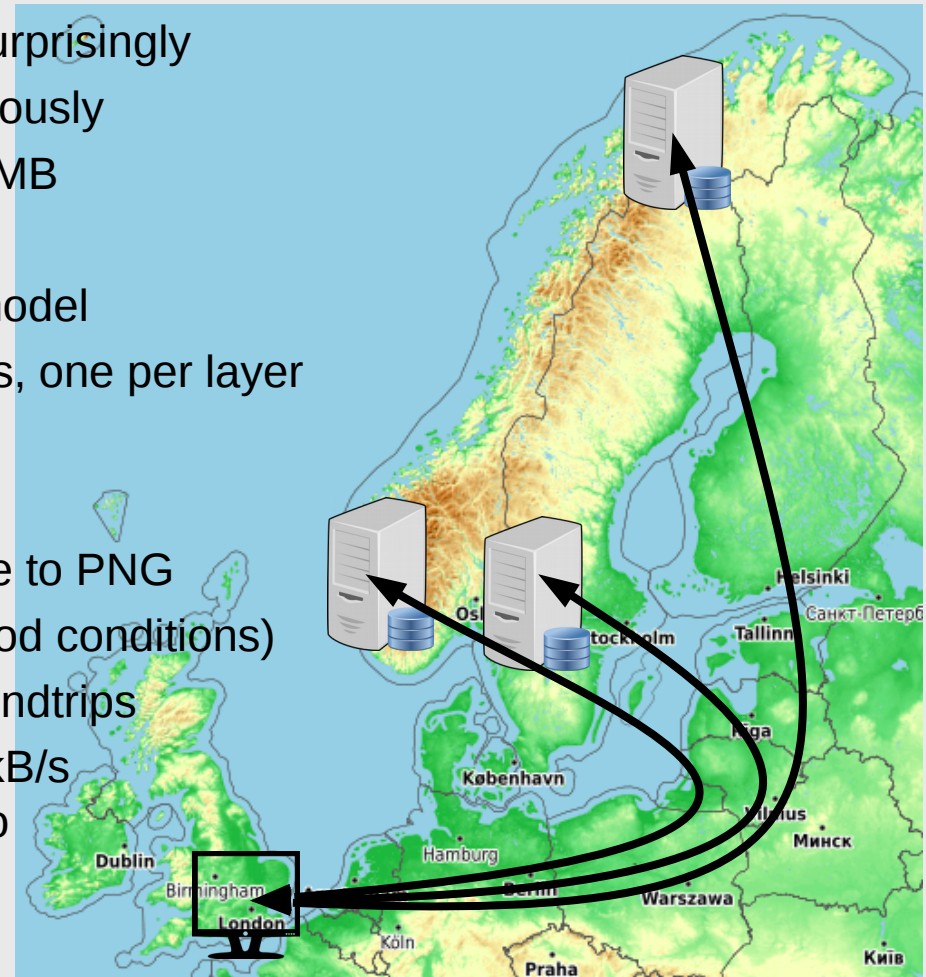
Experiments so far

- restricted to model data (FIELD)
- one client, multiple servers
 - serving data or images
- fetching data / images in multiple threads
- some challenges
 - requirements: “as now or similar in every detail”
 - lack of tests
 - multi-threading problems
- based on GRPC – simple, but not good for our purpose



Preliminary results

- it works, somehow and sometimes, surprisingly
- display size is also increasing continuously
 - 2kx2k pixel ARGB32 image = 16MB
 - 0.16s on a 100MB/s network
 - might be smaller or larger than model
 - 1 field might need several images, one per layer
 - 1 field needs data only once
- some things are slow
 - 30ms or more to compress image to PNG
 - 24ms latency Oslo – Tromsø (good conditions)
 - must fetch in parallel, without roundtrips
- very slow but still functional over 200kB/s
VPN – WLAN – Oslo – Tromsø – Oslo



Side effects & future

- possibility for different clients
 - WMS server might use diana data servers as backend
 - lightweight client might offer a simplified user interface
- possibility for different servers
 - servers in multiple data centers
 - servers coming and going
 - serve from data center until model results are copied, then switch to “local” server
 - servers in other languages than C++
 - servers doing some processing
- diana application does not die if the network is slow or broken
- need to solve some issues before testing
- need to evaluate possible optimisations