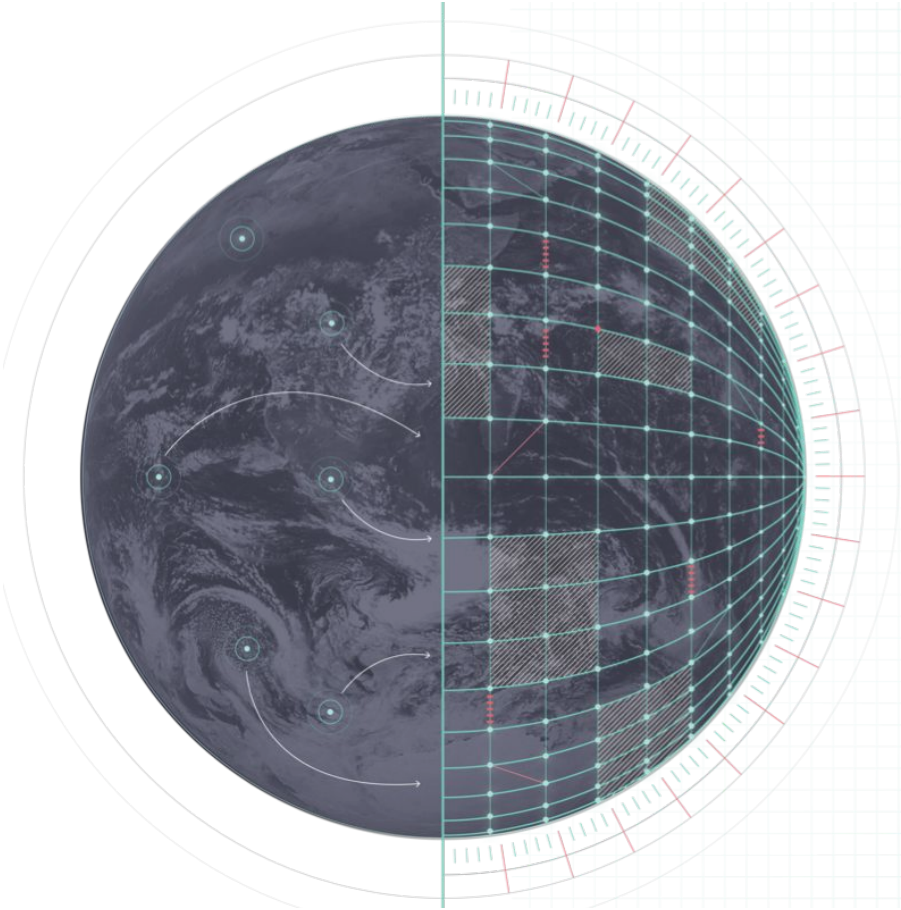

DESTINATION EARTH

Using LUMI to improve the prediction of extreme weather events

Haumont Denis, RMI

LUMI-BE Users days, 06 November 2023, Brussel

Funded by the European Union
Destination Earth
implemented by
ECMWF esa EUMETSAT

INTRODUCTION

Numerical Weather Prediction
(NWP)

Numerical Weather Prediction - Atmosphere modeling

A complex (and chaotic!) system, with many interacting physical phenomena

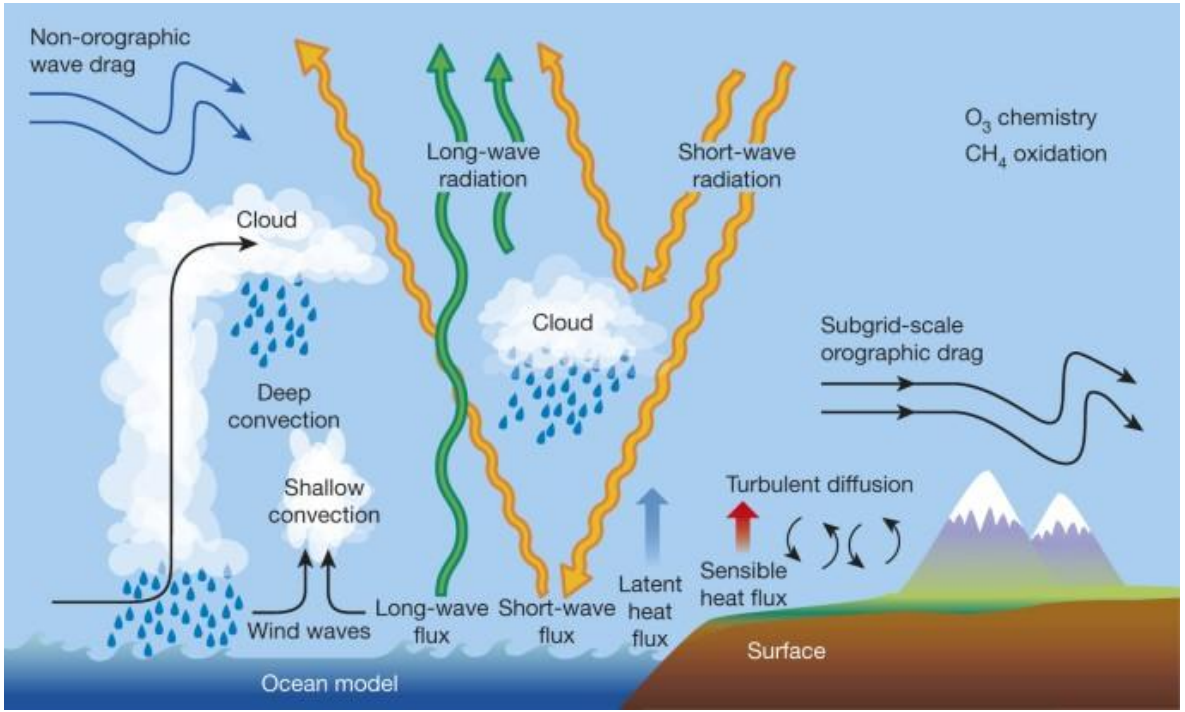

Atmosphere modeling is complex…leading to complex equations

Conservation of momentum

$$\frac{D\vec{v}_h}{Dt} + f\vec{k} \times \vec{v}_h + \frac{1}{\rho}\nabla_z p = \vec{F}_{rv}$$

$$n\frac{Dw}{Dt} + g + \frac{1}{\rho}\frac{\partial p}{\partial z} = nF_{rw}$$

Conservation of energy

$$\frac{DT}{Dt} + \frac{RT}{c_v}\nabla \cdot (\vec{v}_h, w) = \frac{Q}{c_v}$$

Conservation of mass

$$\frac{D\rho}{Dt} + \rho\nabla \cdot (\vec{v}_h, w) = 0$$

Ideal gas law

$$p = \rho RT$$

How do we solve those complex equations?

- Atmosphere is discretized in 3D (columns)
- Most equations are solved in each column individually (Good for parallelization!)
- Exchanges between columns are computed

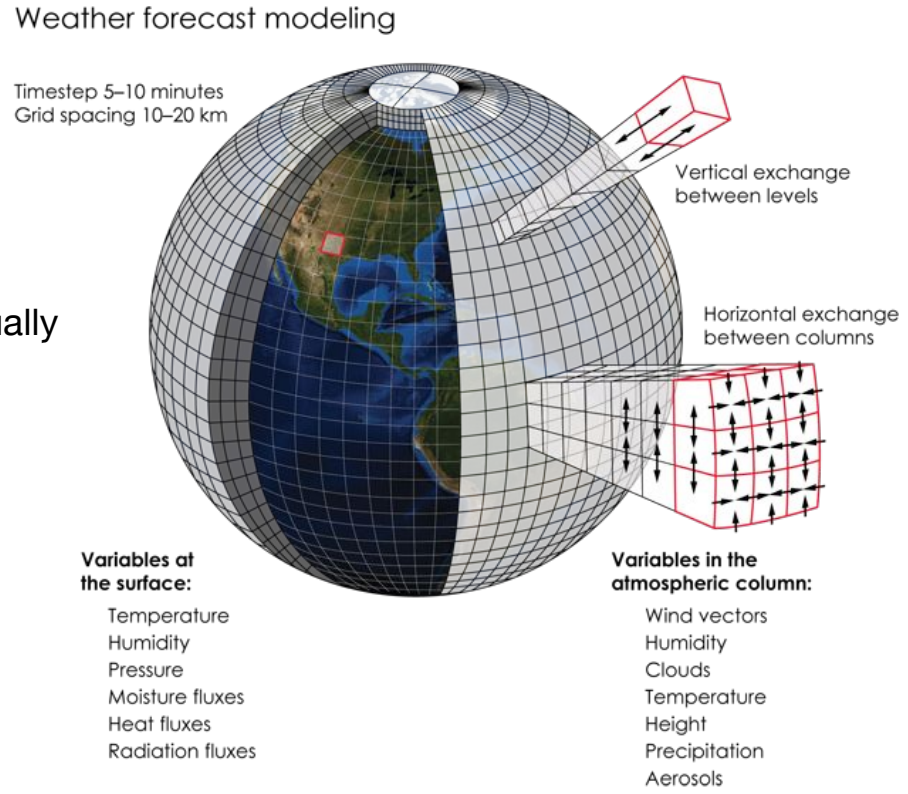

NWP for operational predictions: combining 2 levels

Global Model (IFS)

All around the globe
9 km resolution
~1 billions grid points (!)

+

Local Area Model (LAM)

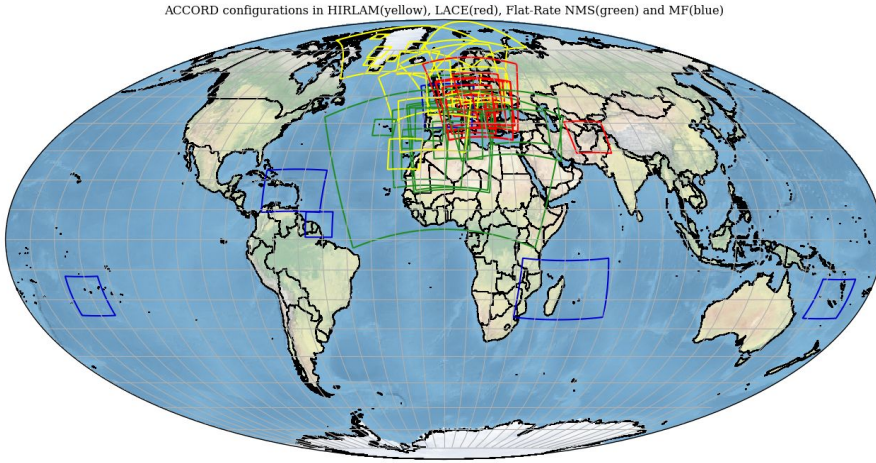

~ Size of a country
Up to 1.3 km resolution
Pre-definite areas

Local Area Model in Belgium: The ACCORD model

- Spectral semi-Lagrangian semi-implicit non-hydrostatic dynamical core
- Code shared with other countries of the ACCORD consortium

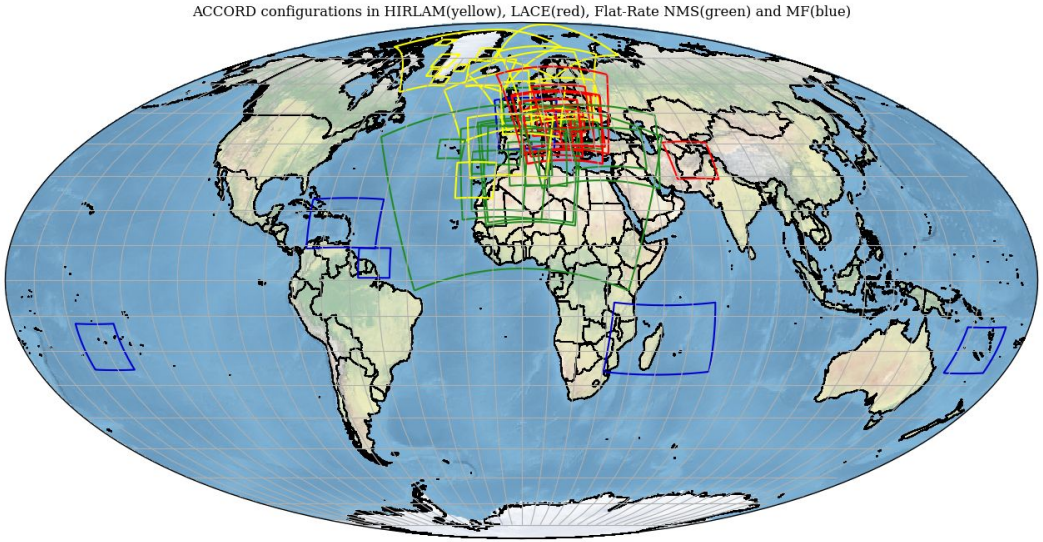

D. Degrauwe & P. Termonia; AMA 2022

EXTREME WEATHER EVENTS

High resolution requirements

NWP - Extreme Weather Events

In Belgium

- Floodings Pepinster, 2021
- Storm Pukkelpop, 2011
- Hittegolf 2003

Gaining more and more importance (climate change)

Extreme Weather Events: high resolution requirements (1/3)

The hurricane-scale wind on 9 March 2021 which escaped operational predictions

Faroe Islands

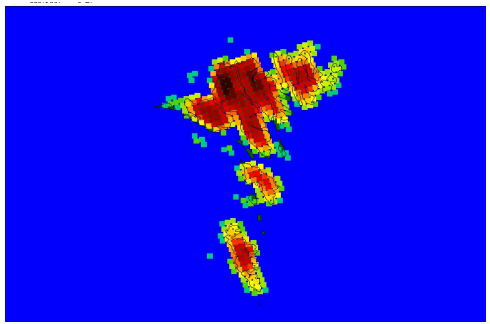

500 m grid

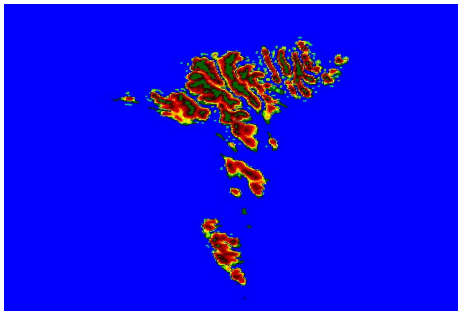

150 m grid

Extreme Weather Events: high resolution requirements (2/3)

The hurricane-scale wind on 9 March 2021 which escaped operational predictions

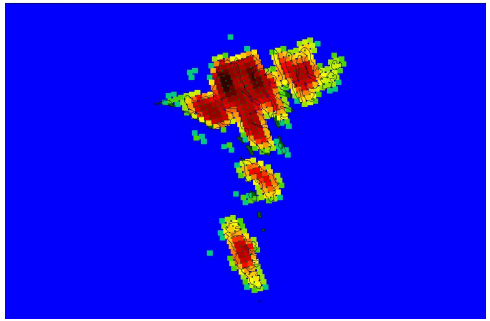

500 m grid

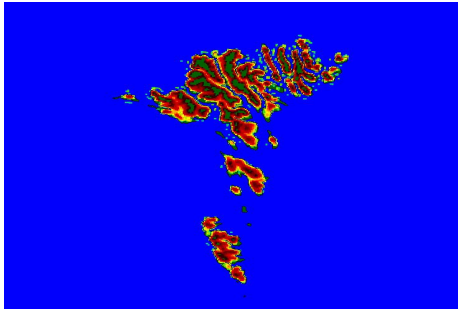

150 m grid

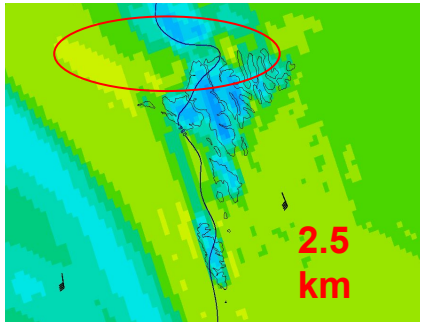

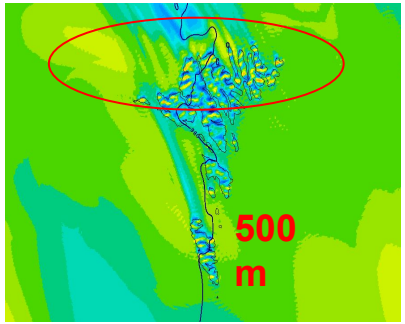

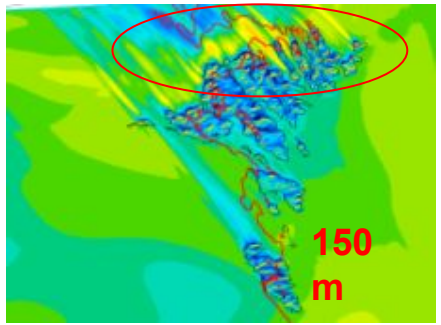

DMI-Harmonie

Storm only visible at 150m resolution

Extreme Weather Events: high resolution requirements (3/3)

- Higher resolution = higher computation cost
  - From 2.5 km to 500m: x100 computational cost

- Consequences:
  - Not possible to compute everywhere: on demand
  - Use EuroHPC infrastructure: heterogeneous hardware (CPU, GPU, vector processors, …)

Extreme Weather Events: high resolution requirements (3/3)

- Higher resolution = higher computation cost
  - From 2.5 km to 500m: up to x100 computational cost

- Consequences:
  - Not possible to compute everywhere: on demand
  - Use EuroHPC infrastructure: heterogeneous hardware (CPU, GPU, vector processors, …)

DESTINATION EARTH

On Demand Extreme Digital Twin

Destination Earth: On Demand Extremes Digital Twin

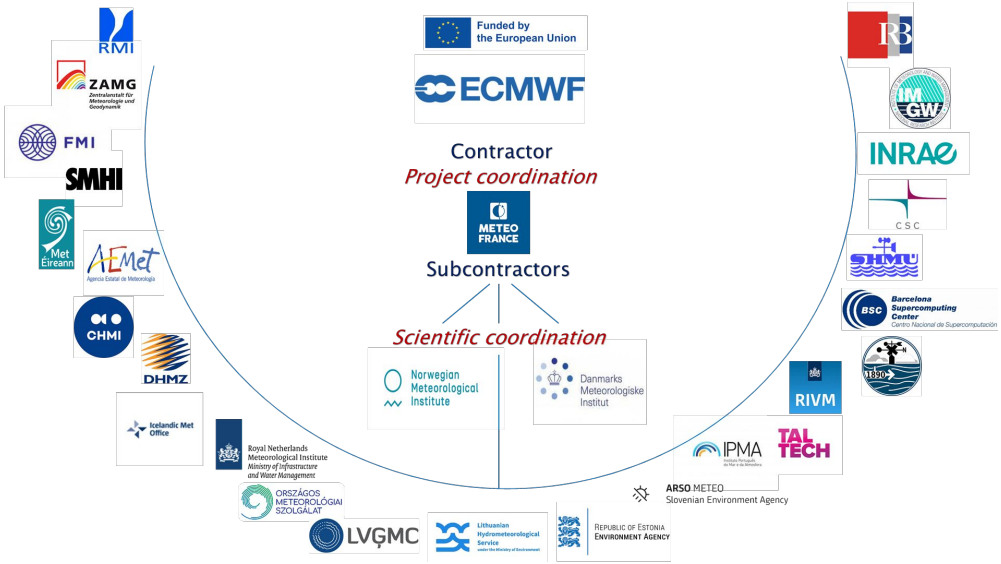

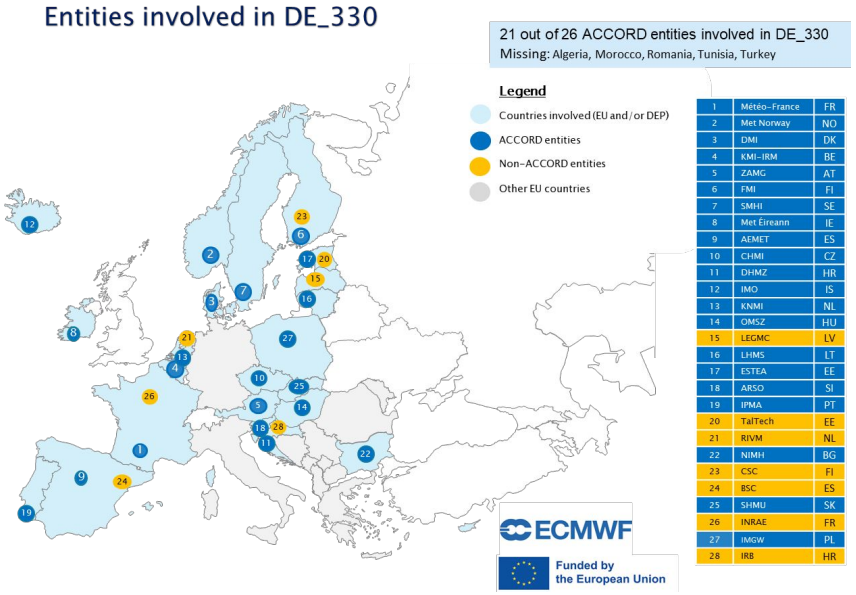

A collaboration between many countries (members of ACCORD consortium)

Digital Twin of the Earth

Using Science....

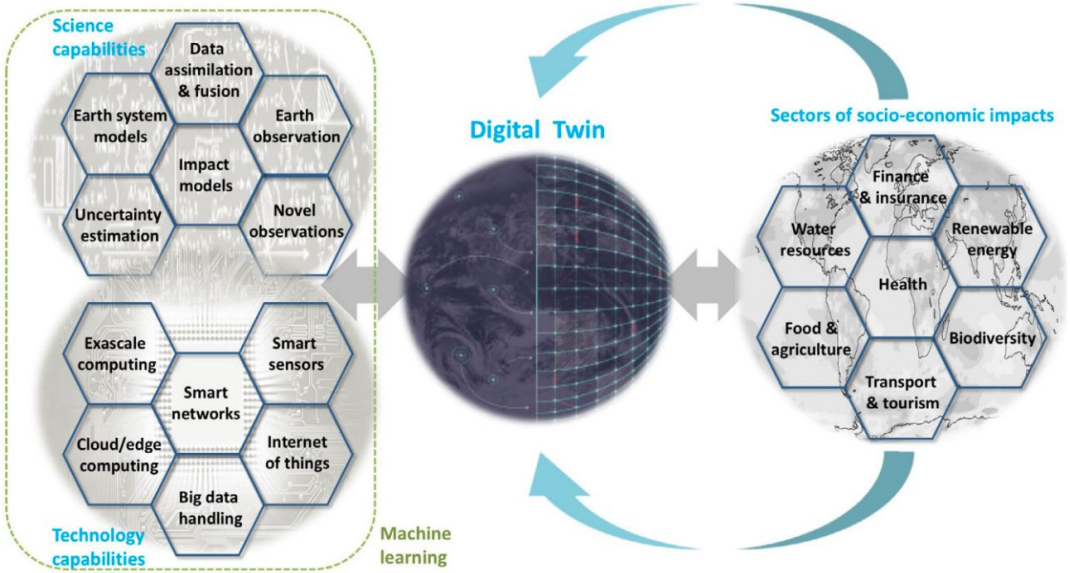

... for policy decisions support

Two Components

- Global continuous Digital Twin (Global Model)
- On-demand Extremes Digital Twin (Local Area Model)

Digital Twin of the Earth

Using Science....

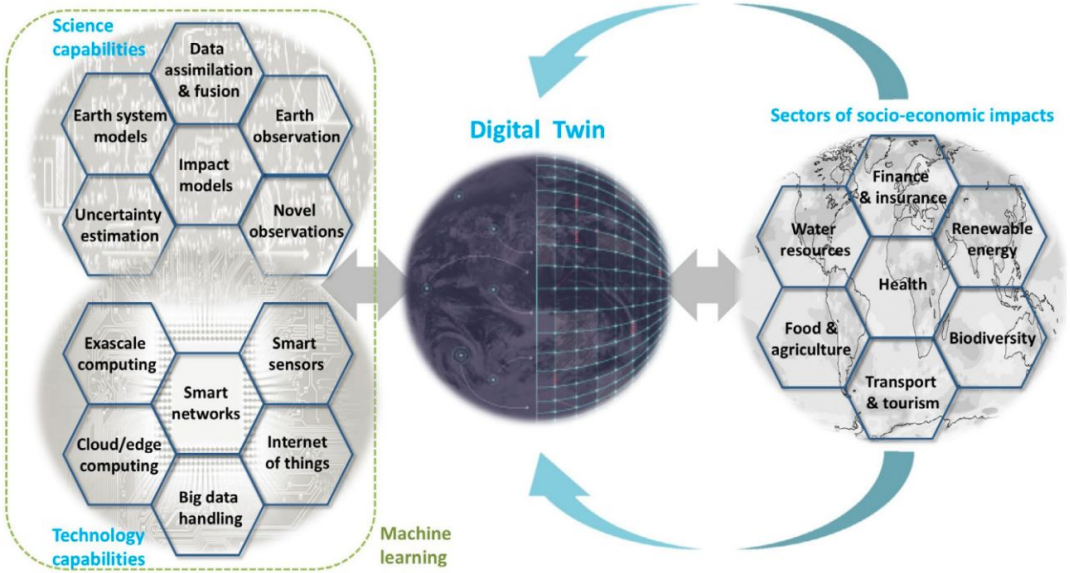

... for policy decisions support

Two Components

- Global continuous Digital Twin (Global Model)
- On-demand Extremes Digital Twin (Local Area Model)

On Demand Extremes Digital Twin

STEP 1

Detect a possible extreme event from the global continuous Digital Twin

On Demand Extremes Digital Twin

STEP 1

STEP 2

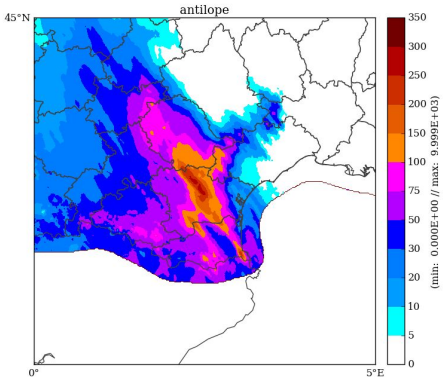

The Aude flooding 2018

Create a on-demand LAM computation over the zone of interest
(< 1km grid resolution)

On Demand Extremes Digital Twin

STEP 1

STEP 2

STEP 3

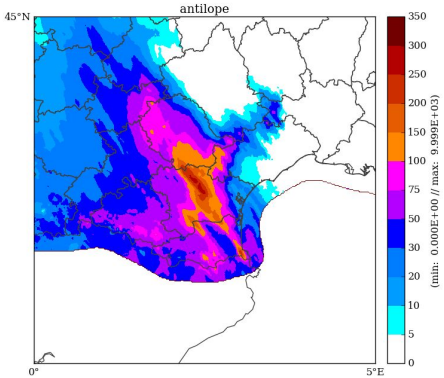

The Aude flooding 2018

Run the relevant impact model for decision-making support

On demand application: storm tracking

MSG SEVIRI CH. 9 (IR)

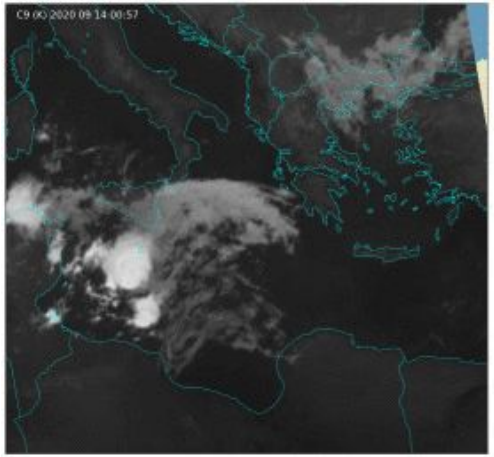

ACCORD NWP **2.5 km**

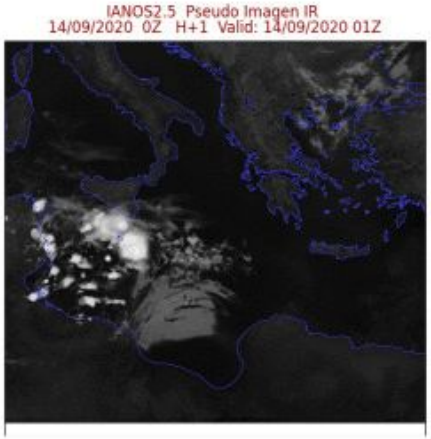

ACCORD NWP **500 m**

Ianos Medicane Sept 2020

(Courtesy Javier Calvo, AEMET)

PORTING CODE TO GPU

LUMI-G for High Resolution Forecasts

GPU porting: Objectives

- Increase the resolution of forecast to allow accurate assessment of extreme events

- Exploit GPU computation power

  AMD GPU MI250x peaks at 42.2 TFLOPS !

Gpu porting: Constraints

- Current performance on CPU (used for operations) must not degrade
- GPU porting as transparent as possible for Scientists
- Allow to extend the code in the future
  - keep a single code base
  - code remain readable to scientists
- Codebase is large: allow progressive porting, by phases

Our Strategy for GPU code adaptation

GPU adaptation of existing code is based on 3 pillars:

- Using hardware-optimized libraries where possible (i.e. for parts of the code that don't change too often): hipBLAS, hipFFT, …
- Increase flexibility by improving code layout (i.e refactoring)
- Source-to-source translation of code that regularly undergoes scientific changes (automated)

  We use Loki scripts to implement an OpenACC porting of our Fortran code

  Loki is an open source Python package for source-to-source translation

  https://github.com/ecmwf-ifs/loki/

Example of what we would like to achieve

- Flexibility: Hybrid execution on GPU and CPU
- Source-to-source translation for scientific code porting
- Optimized vendor libraries (FFT)

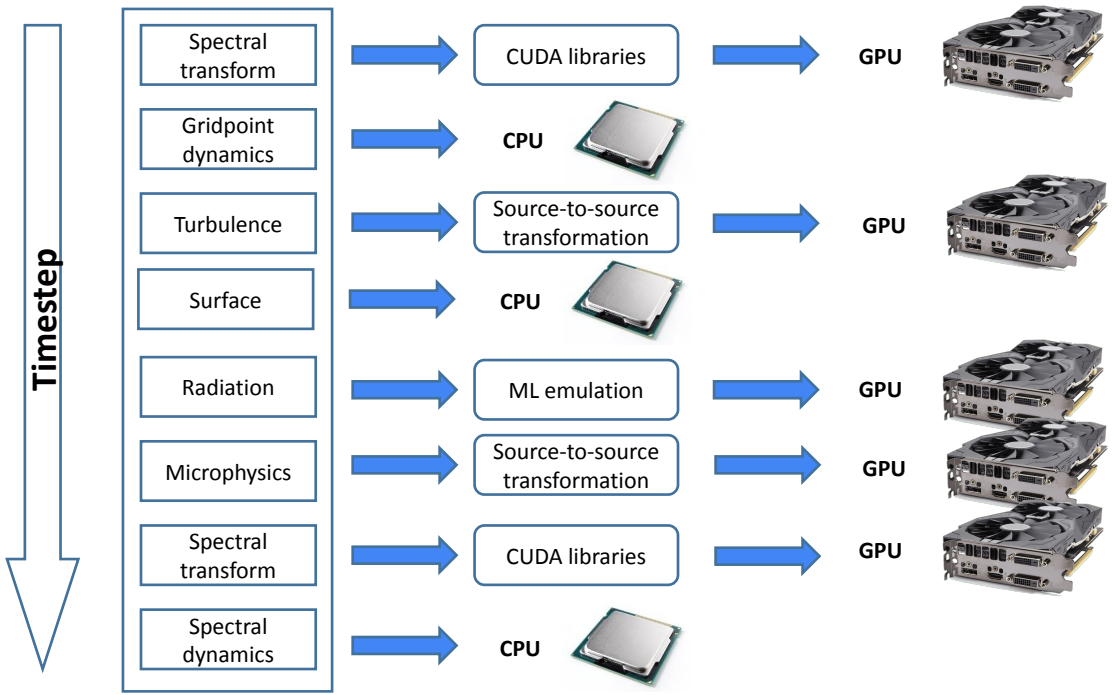

D. Degrauwe & P. Termonia; AMA 2022

Note: Multi-GPU implementation

- NWP is very communication-intensive

  e.g. MPI_ALLTOALL in the spectral transforms

- Performance wille rely on efficient GPU to GPU direct MPI transferts

Preliminary results - Radiation scheme on LUMI (1/2)

Comparison of different parallelization strategies for ACRANEB2 (NVIDIA vs AMD)

VERSION			LUMI		ATOS	
		Thread	Time (sec.)	Throughput (Mpts/sec)	Time (sec.)	Throughput (Mpts/sec)
MANUAL	cpu_ref	16	1,76	11,6	2,21	9,24
	small_kernels (***)	16	50,55	0,4051	3,08	6,64
	vector_routines	1	TODO	TODO	1,25	16,34
	stack_driver (**)	1	N/A	N/A	TODO	TODO
FXTRAN	openacc-kernels (default)	1	0,9	22,4	1,14	17,94
	openacc-vector-stack (**)	1	N/A	N/A	0,83	24,65
	single-directive-redim-sp (*)	1	0,68	30,06	0,3	68,04
Loki	loki-kernels (***)	16	32,44	0,6313	2,88	6,64
	loki-scc	1	1,81	11,34	1,93	10,62
	loki-scc-hoist	1	TODO	TODO	TODO	TODO

Preliminary Results - Profiling Spectral Transform (2/2)

- Spectral transform ported to LUMI GPU based on hipFFT
- Profiling: memory transfert and transposition of matrices are very costly

(Spectral semi-Lagrangian semi-implicit non-hydrostatic dynamical core)

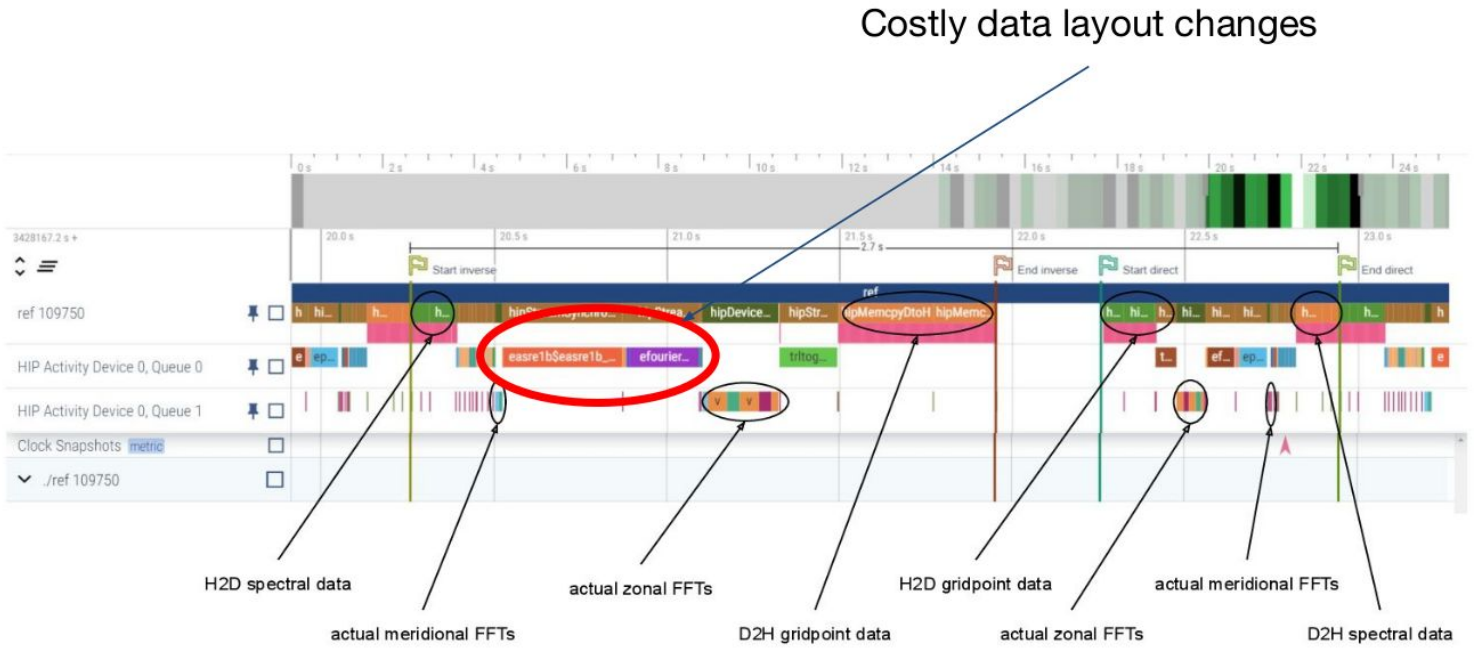

CONCLUSION

Some lessons learned during GPU porting

Conclusion : Some lessons learned during GPU porting

- Performances still far from expected
  - Memory layout crucial…but hard to change
  - source to source translation are limited in the kind of optimizations they can perform…manual optimization required
- Our scientist community is quite conservative (for a reason)
  - Hard to enforce coding rules, implement refactoring or adopt new technologies
- GPU: Compilers and technology stack are still very experimental
- LUMI not suitable for operational work (short notice maintenance, system hiccups, etc)

Conclusion : Destination Earth perspectives

- Destination Earth Project will continue:
  - Phase 1 is finishing in April 2024
  - Phase 2
    - Porting other parts of the NWP to the GPU
    - Improve performance
    - Investigate ML approaches
  - Phase 3
    - Make it operational


LUMI
Thanks !