Funded by the European Union

Co-ordinated by ECMWF

ESCAPE

Energy efficient
SCalable
Algorithms for
weather **P**rediction
at **E**xascale

www.hpc-escape.eu

Peter Bauer

MeteoSwiss · DMI · RMI · NVIDIA · METEO FRANCE · ICHEC · Deutscher Wetterdienst · Optalysys · Bull · PSNC · Loughborough University

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 671627

Funded by the European Union

Future and
Emerging
Technologies-
High
Performance
Computing
(FET-HPC)

ECMWF MeteoSwiss DMI RMI NVIDIA METEO FRANCE ICHEC Deutscher Wetterdienst Optalysys Bull PSNC Loughborough University

ESCAPE partners & expertise

Funded by the European Union

Global operational NWP
Petaflop data centre
Numerical methods, parallel computing

Regional operational NWP
Numerical methods

Key technology development
Parallel computing

Regional operational NWP
Teraflop data centre
Physical processes

Numerical methods
Teraflop data centre
Parallel computing

Teraflop data centre
Parallel computing
Performance modelling

Teraflop data centre
Parallel computing

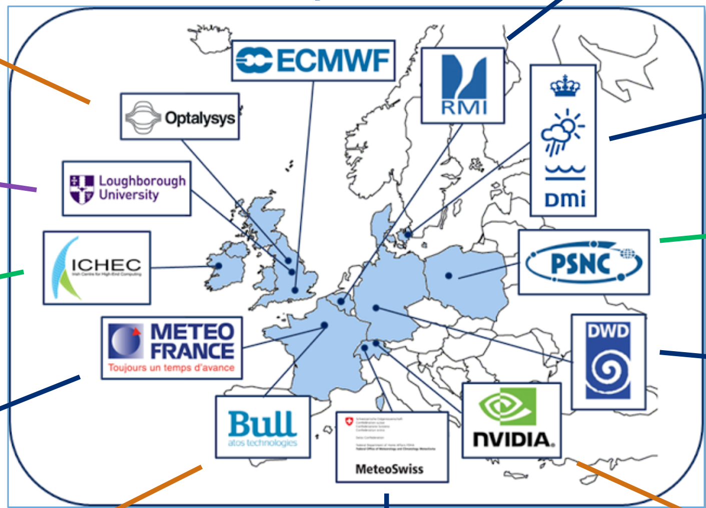

Global/regional operational NWP
Teraflop data centre
Numerical methods

Global/regional operational NWP
Petaflop data centre
Numerical methods

Key technology development
Parallel computing

Regional operational NWP
Petaflop data centre
Parallel computing

Key technology development
Parallel computing

ECMWF · MeteoSwiss · DMI · RMI · NVIDIA · METEO FRANCE · ICHEC · Deutscher Wetterdienst · Optalysys · Bull · PSNC · Loughborough University

ESCAPE European impact map

Funded by the European Union

34 countries
- Member States
- Co-operating States
- Under negotiation

7 countries — COSMO

11 countries — HIRLAM

16 countries — ALADIN/RC-LACE/SELAM

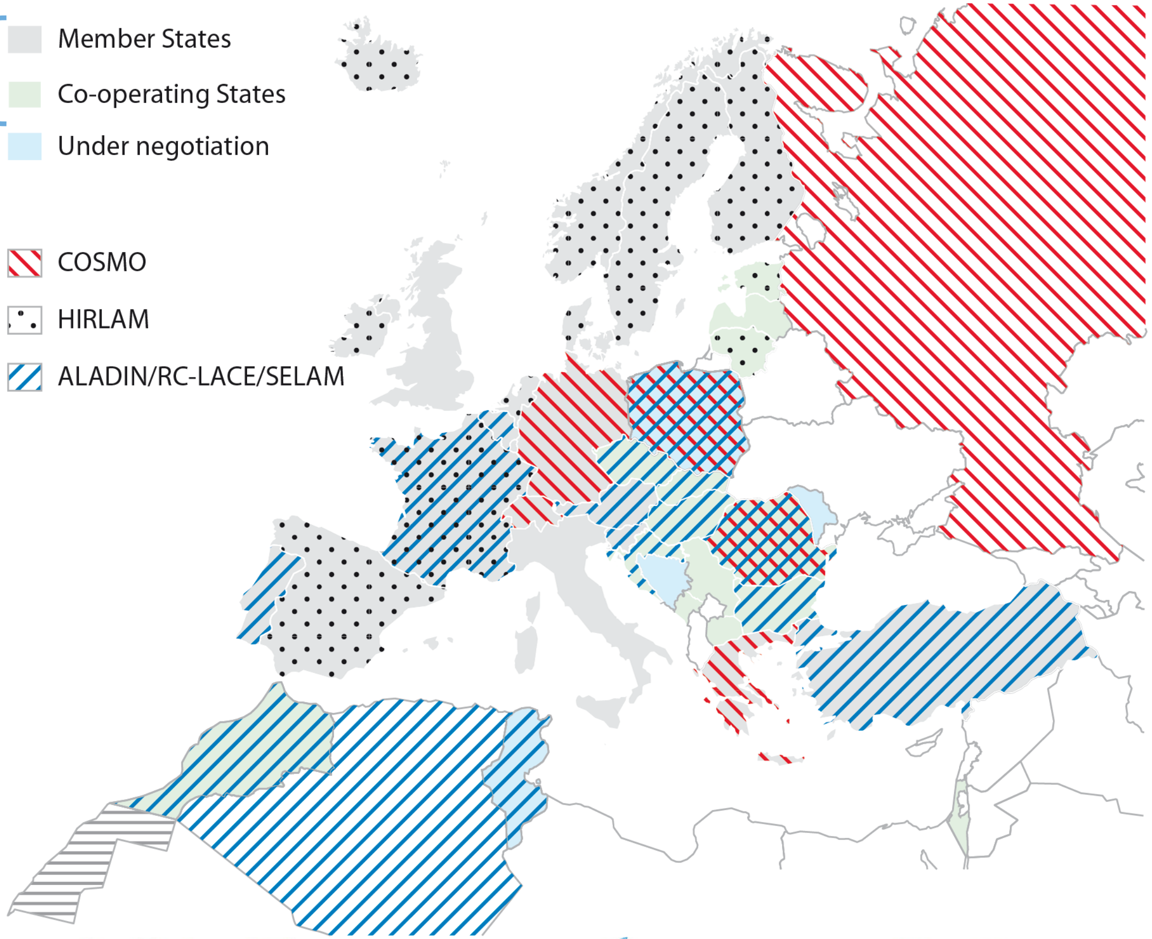

ECMWF · MeteoSwiss · DMI · RMI · NVIDIA · METEO FRANCE Toujours un temps d'avance · ICHEC · Deutscher Wetterdienst Wetter und Klima aus einer Hand · Optalysys · Bull atos technologies · PSNC · Loughborough University

Required: Billion-way parallelism

Funded by the European Union

Hardware

- **hybrid computing** model seems to be here to stay;
- **memory systems** will become ever more complicated;
- **hardware faults** require fast adaptation strategies.

[Dongarra et al. 2015]

→ Interconnection technology is seriously lagging behind computing power: **2-3 orders of magnitude gap!**

communities.″

ESCAPE key objectives

Funded by the European Union

1. Define fundamental **algorithm building blocks** (“*Weather & Climate Dwarfs*”) to co-design, advance, benchmark and efficiently run the next generation of NWP and climate models on energy-efficient, heterogeneous HPC architectures.

2. Diagnose and classify Weather & Climate Dwarfs on **different HPC architectures**

3. Combine frontier research on **algorithm development** and extreme-scale, high-performance computing applications with **novel hardware technology**, to create a flexible and sustainable weather and climate prediction system.

4. Foster the **future design of Earth-system models** and commercialisation of weather-dependent innovative products and services in Europe through enabling open-source technology.

5. Pairing **world-leading NWP** with **innovative HPC solutions**.

ESCAPE workflow

Funded by the European Union

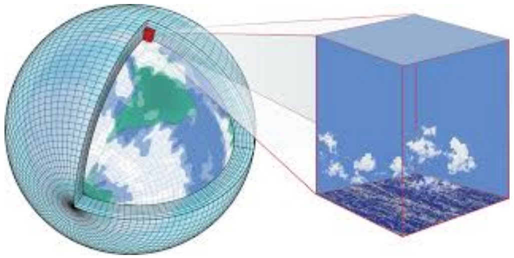

Extract model dwarfs…

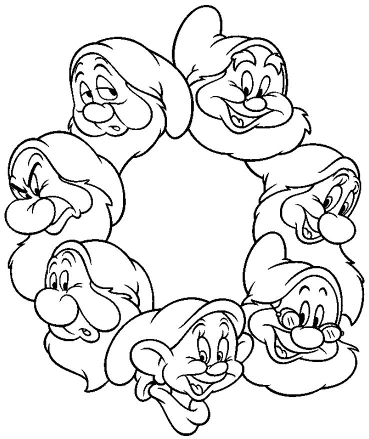

… hardware adaptation …

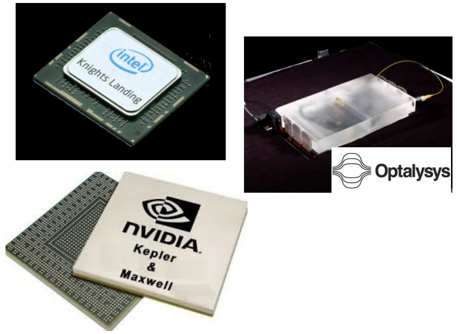

… explore alternative numerical algorithms …

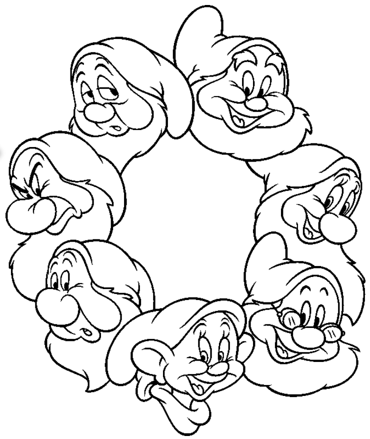

… reassemble model and benchmark

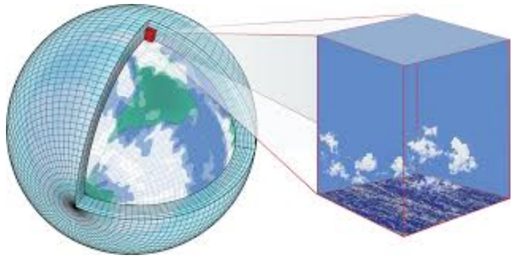

ECMWF MeteoSwiss DMI RMI NVIDIA METEO FRANCE Toujours un temps d'avance ICHEC Deutscher Wetterdienst Wetter und Klima aus einer Hand DWD Optalysys Bull PSNC Loughborough University

What is a dwarf?

Funded by the European Union

Weather & Climate Dwarfs encapsulate basic algorithmic motifs by breaking down numerical weather prediction legacy codes into key functional components - in analogy to the Berkeley Dwarfs.

Weather & Climate Dwarfs are distinctly motivated by the requirement to maximise computing performance, energy consumption, as well as time-and-cost-to-solution.

Why focus on Dwarfs?

Funded by the European Union

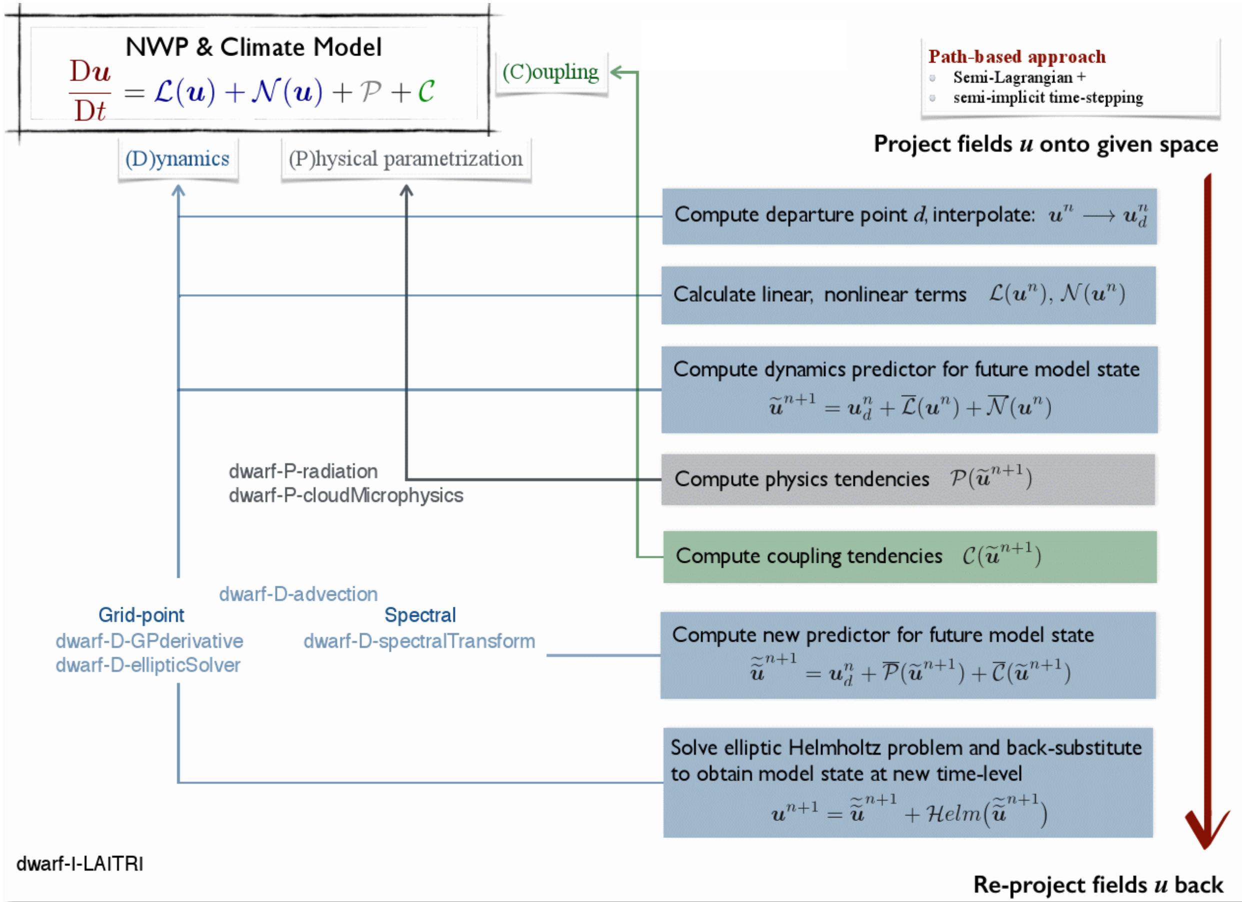

No 1: sequential with time steps (10d x 24h x 3600s / 450s = 1920 @ 9km

Transforms: communication bandwidth, memory

Advection: halo-communication, scalability

Physics: expensive calculations, scalable

Ocean waves: expensive calculations, load balancing

3D-solver: iterations, memory, communication bandwidth

Transforms: communication bandwidth, memory

Relevance of Dwarfs

% of the entire runtime

Funded by the European Union

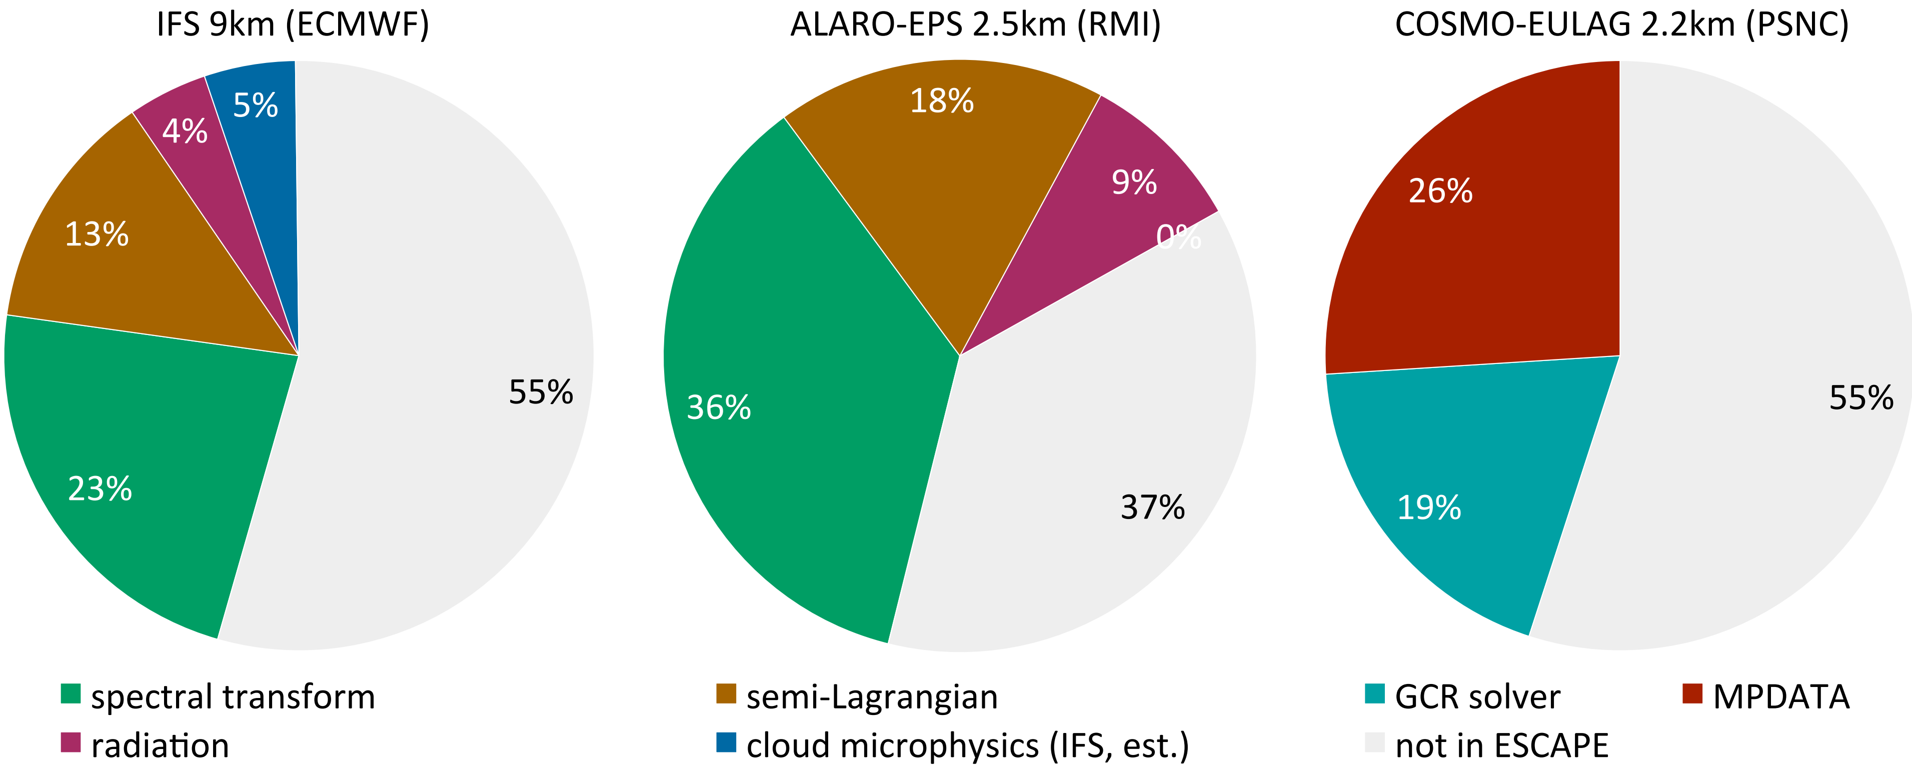

Funded by the European Union

Dwarfs and Programming

Dwarf	prototype implemented	docu-mented	based on Atlas	MPI	Open MP	Open ACC	DSL	Optalysys
D - spectral transform - SH	✅	✅	✅	✅	✅	✅		✅
D - spectral transform - biFFT	✅	✅		✅	✅	✅		✅
D - advection - MPDATA	✅	✅	✅	✅	✅	✅	✅	
D - advection - semi-Lagrangian	✅	✅	✅	✅	✅			
D - elliptic solver - GCR	✅	✅	✅	✅	✅		●	
D - elliptic solver - multigridPrecon	✅	✅	✅	✅	✅			
P - cloud microphysics - CloudSC	✅	✅		✅	✅	✅		
P - radiation scheme - ACRANEB2	✅	✅		✅	✅	✅		
I - LAITRI (3d interpol. algorithm)	✅	✅			✅	✅		

✅: first version running

●: planned

empty cells: not part of ESCAPE

Hybrid Computing – single GPU

Funded by the European Union

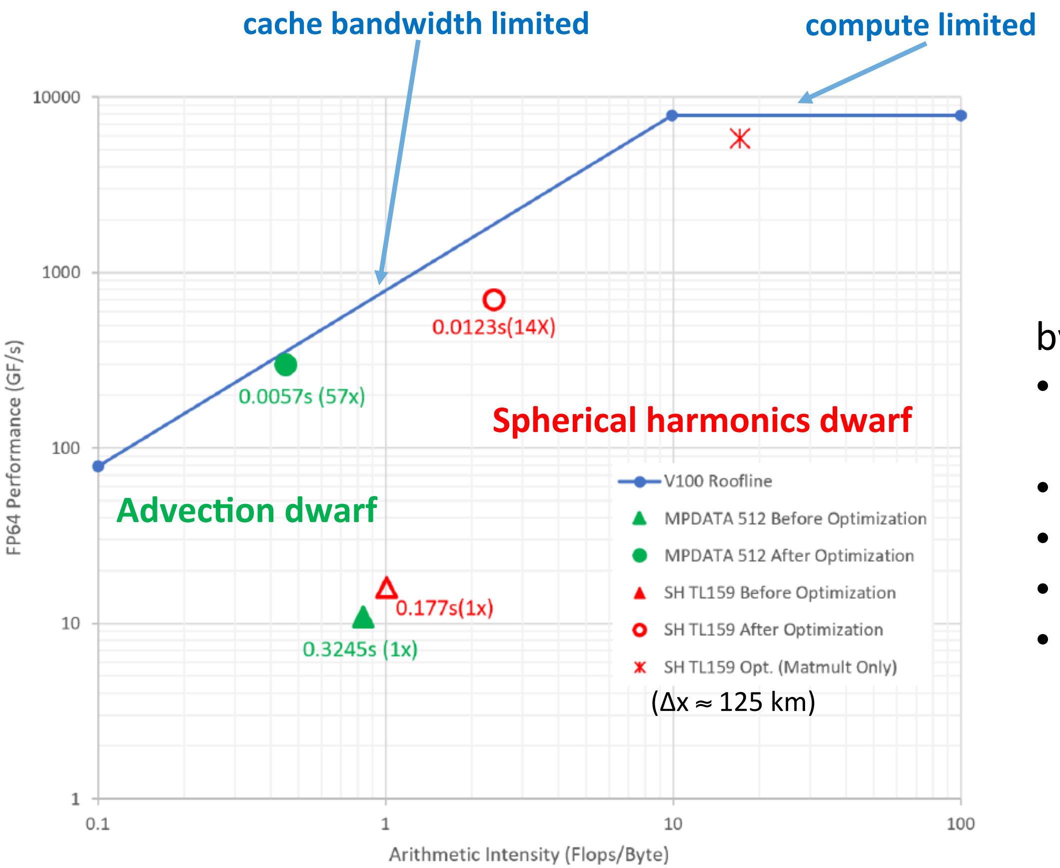

by:

- exposing parallelism in loops for OpenACC mapping
- Kernel optimization by memory mapping
- exploiting CUDA BLAS features
- minimizing data allocation and movement
- (calling C /CUDA from PGI Fortran)

ECMWF MeteoSwiss DMI RMI NVIDIA METEO FRANCE Toujours un temps d'avance ICHEC Deutscher Wetterdienst Wetter und Klima aus einer Hand Optalysys Bull PSNC Loughborough University

Hybrid Computing – multiple GPU

Funded by the European Union

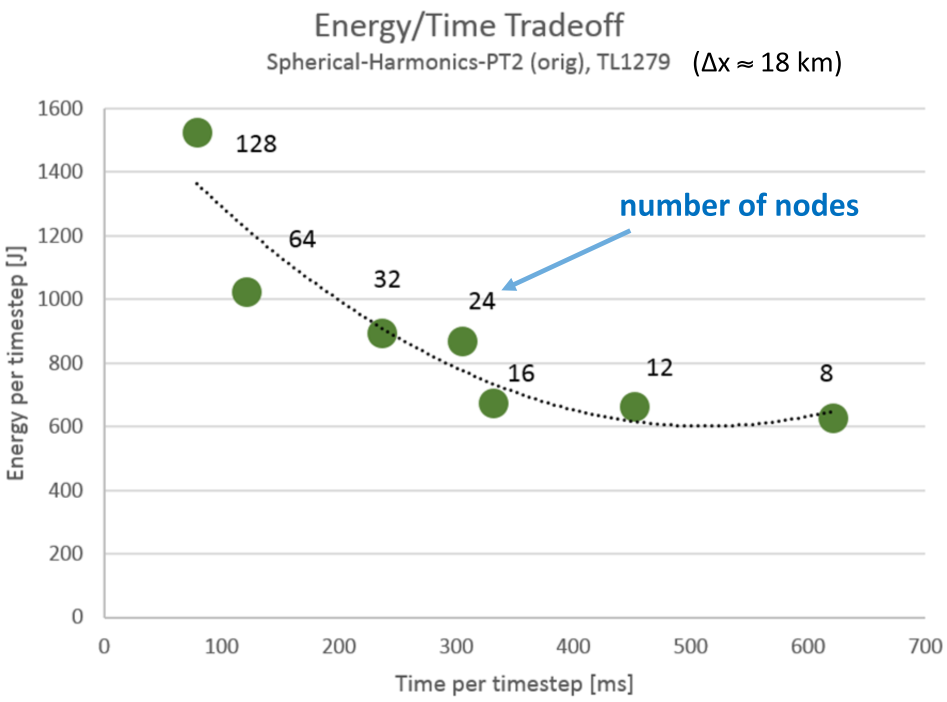

NVIDIA NGX-2 with NVSwitch

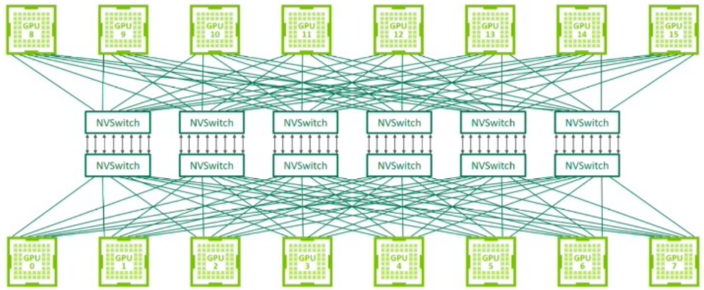

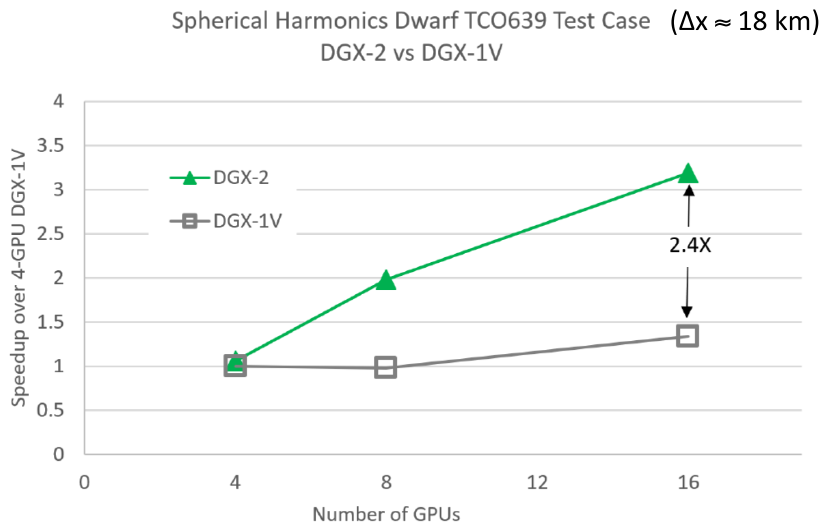

ECMWF MeteoSwiss DMI RMI NVIDIA METEO FRANCE ICHEC Deutscher Wetterdienst Optalysys Bull PSNC Loughborough University

Dwarfs vs Full model – multiple CPU

Funded by the European Union

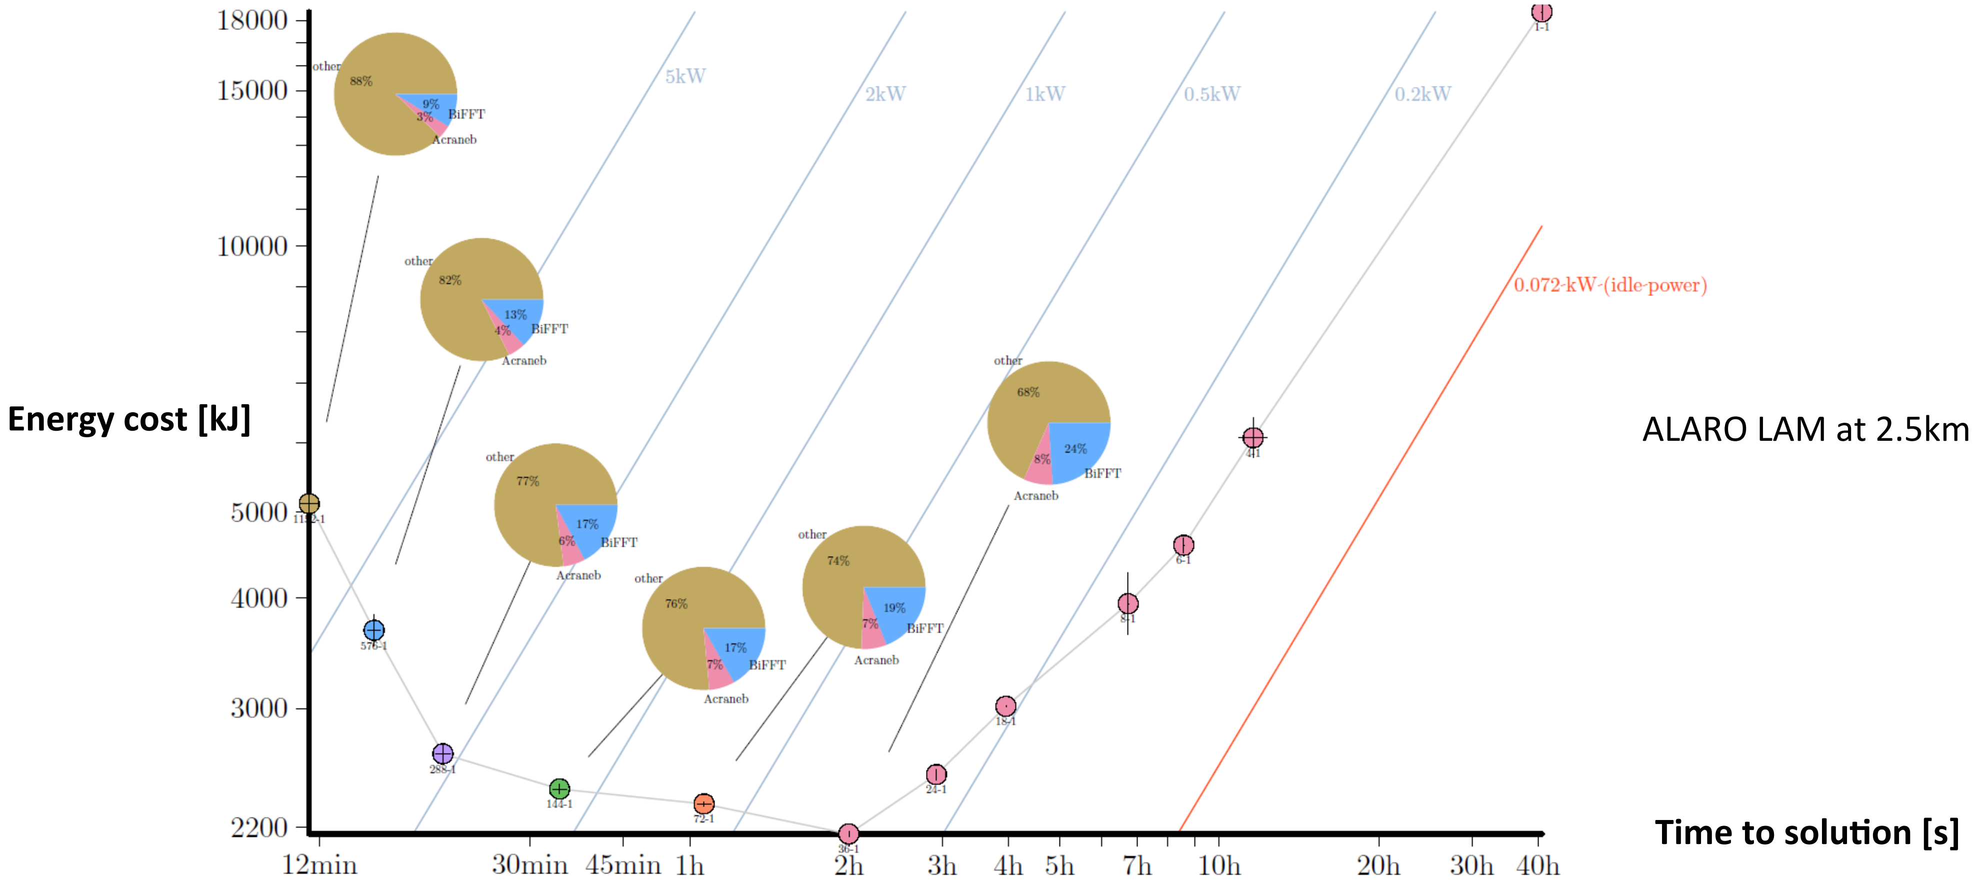

ESCAPE: Separation of concerns

Funded by the European Union

GridTools

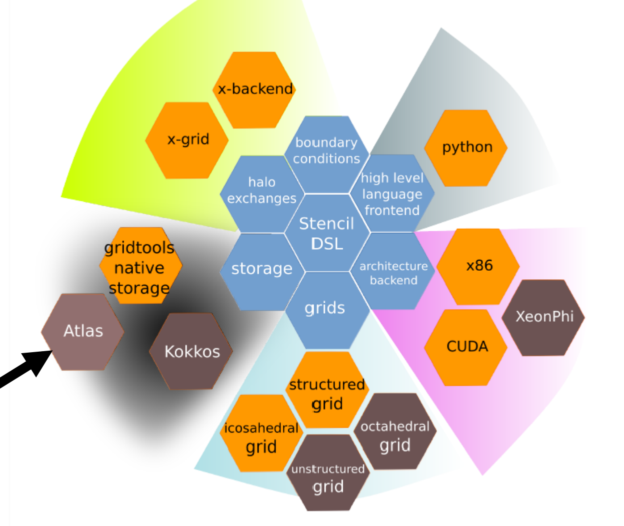

Atlas

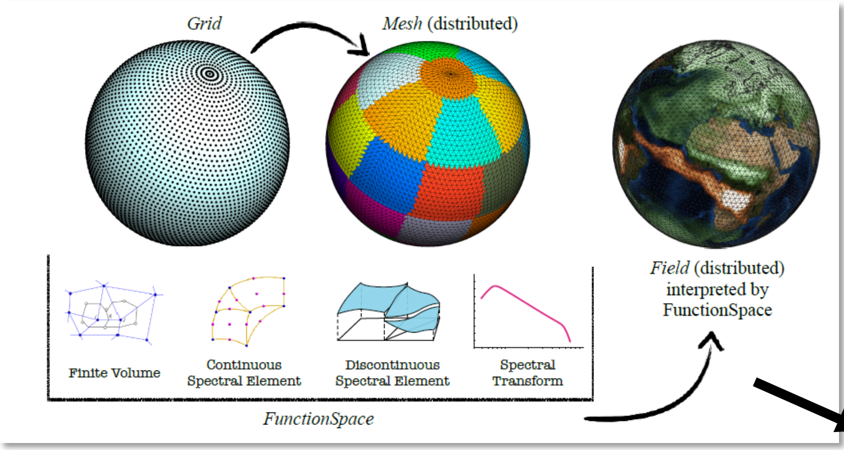

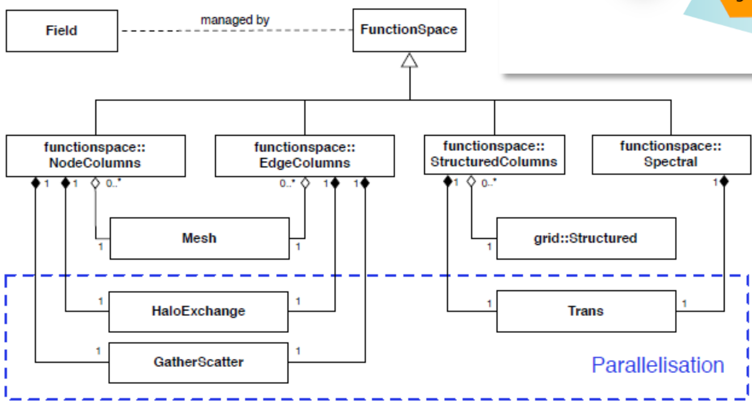

→ **Weather and climate DSL development may be the only solution for enabling efficient science development *and* producing performant and portable code!**

ECMWF MeteoSwiss DMI RMI NVIDIA METEO FRANCE ICHEC Deutscher Wetterdienst Optalysys Bull PSNC Loughborough University

ESCAPE Software collaboration platform

Funded by the European Union

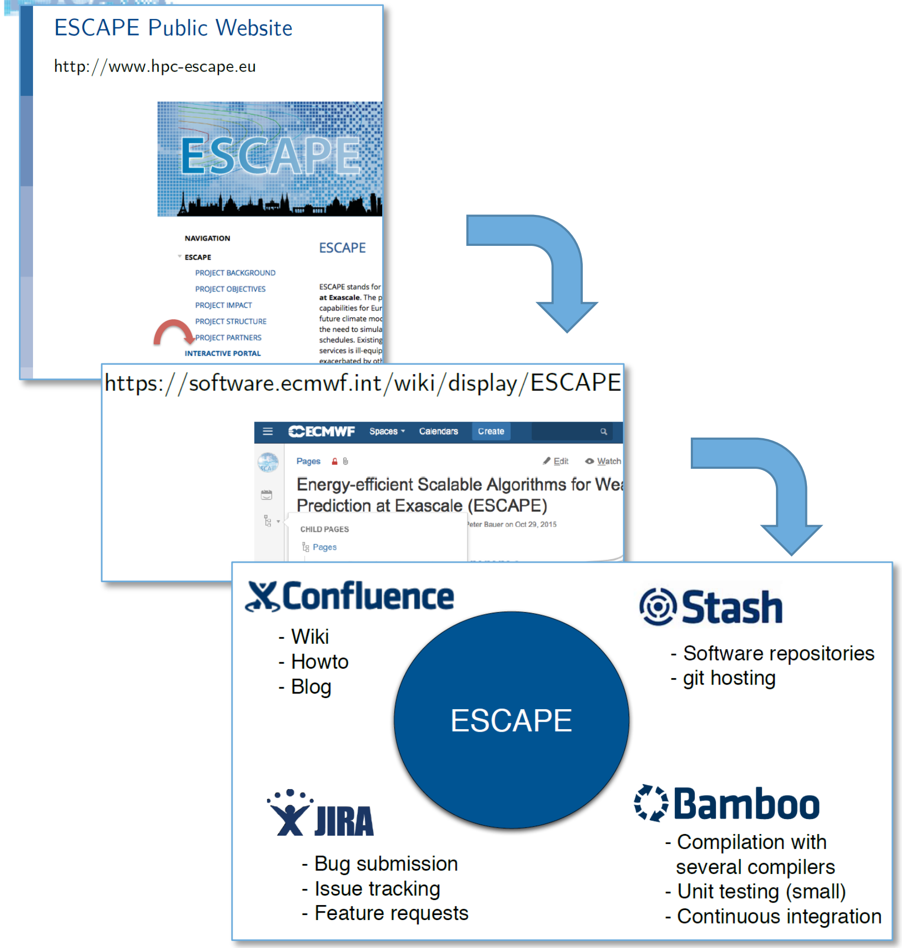

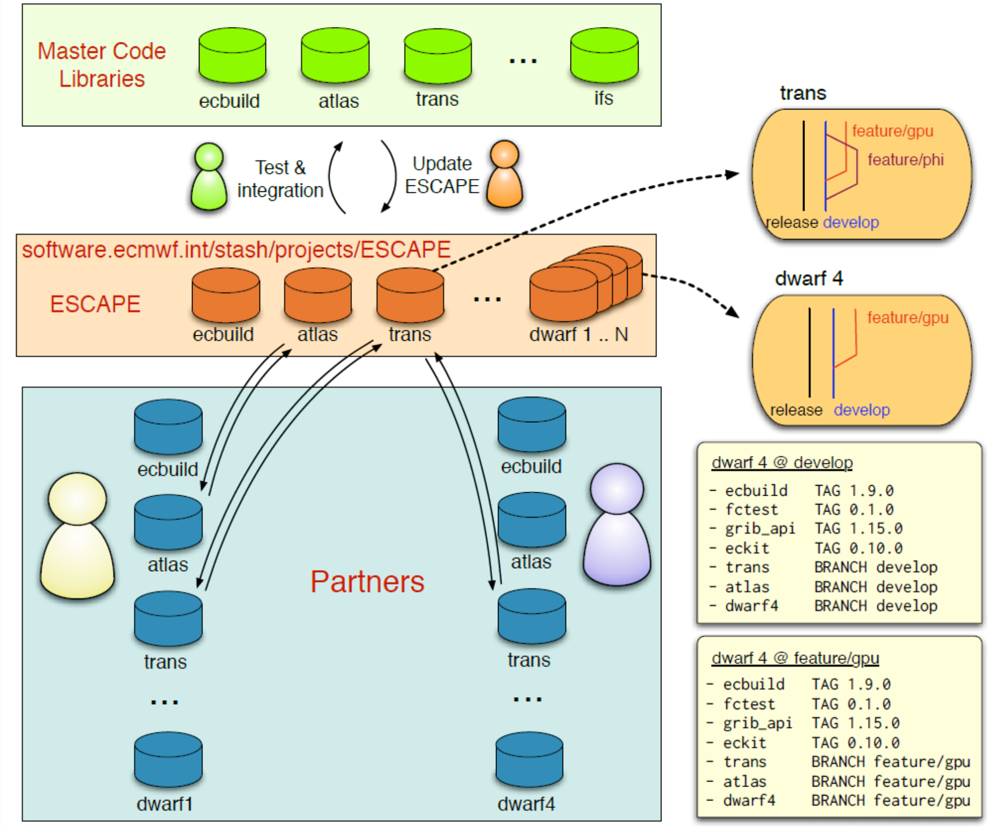

ESCAPE Software stack

Funded by the European Union

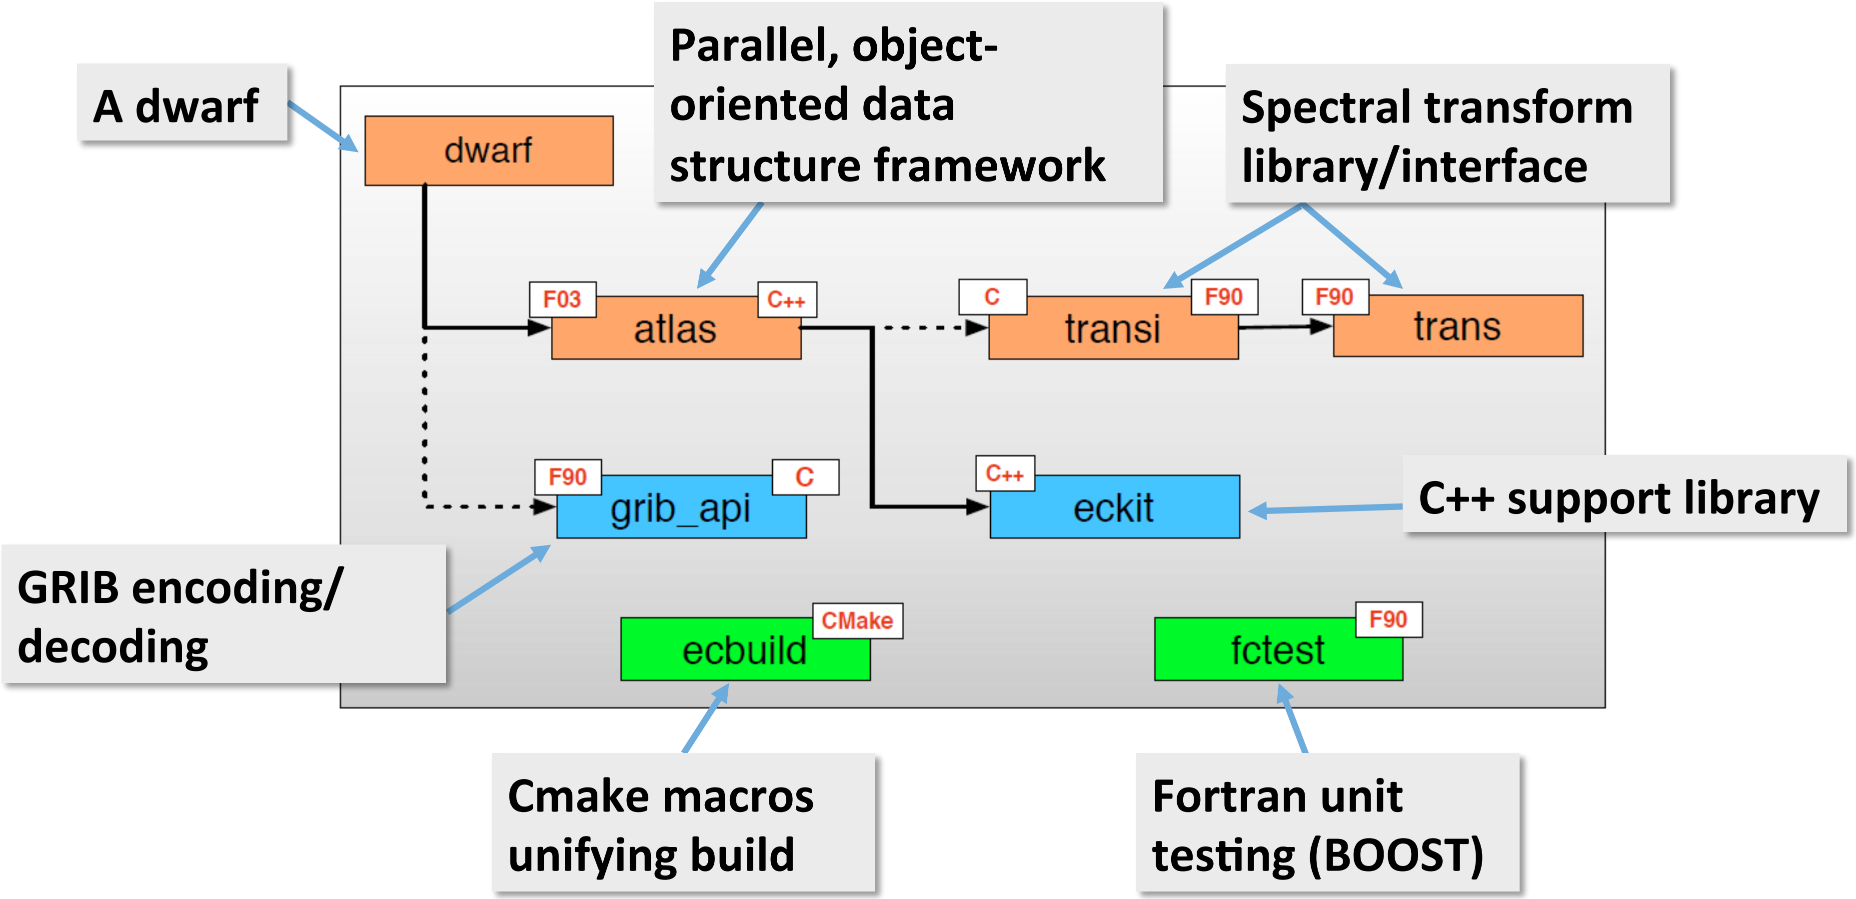

ECMWF MeteoSwiss DMI RMI NVIDIA METEO FRANCE ICHEC Deutscher Wetterdienst Optalysys Bull PSNC Loughborough University

From ESCAPE to ESCAPE-2

Funded by the European Union

Neural networks

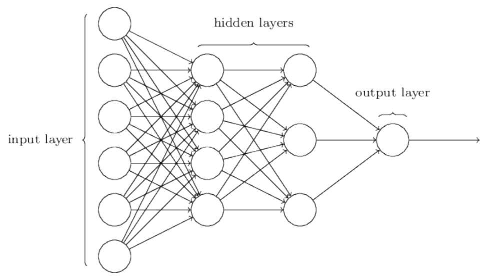

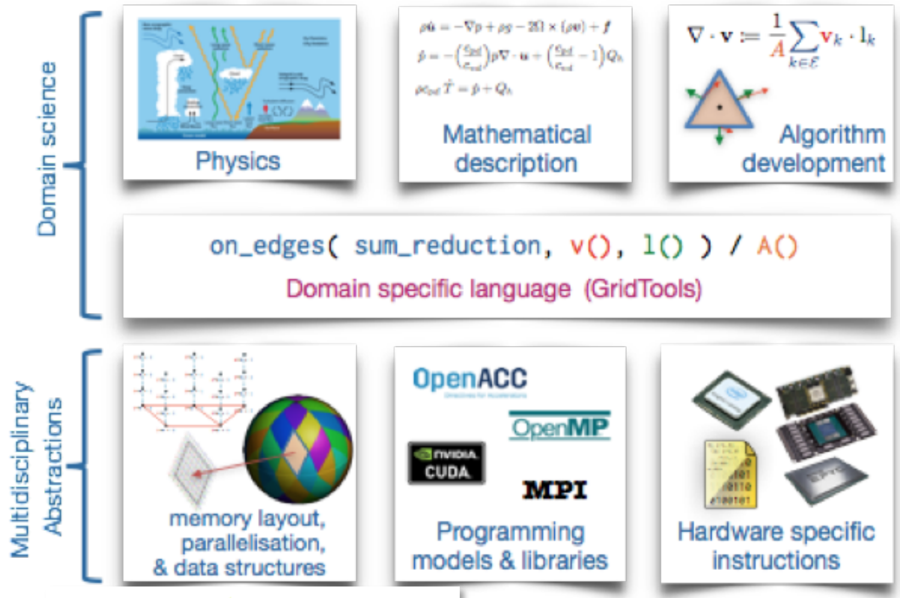

Mathematics&algorithms

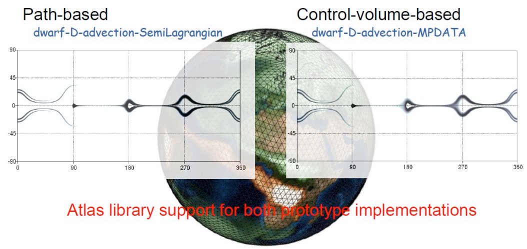

Atlas

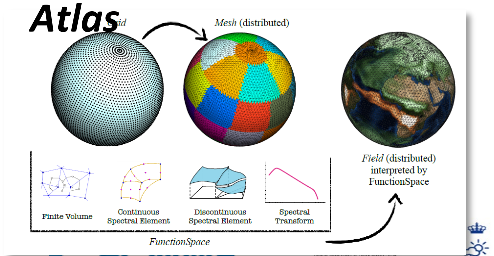

MeteoSwiss

GridTools

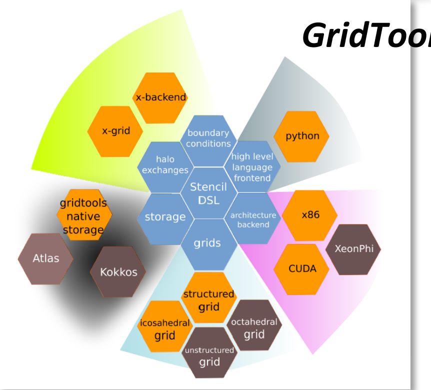

Processors

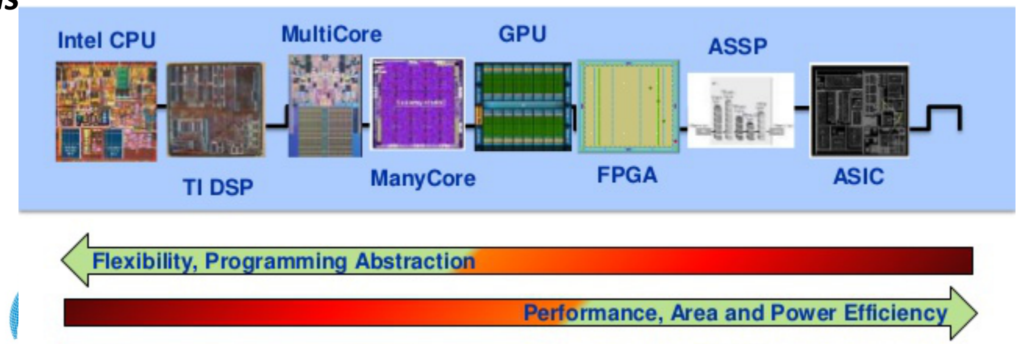

From ESCAPE to ESCAPE-2

Funded by the European Union

Mathematical methods and algorithms:

- Semi-implicit, semi-Lagrangian CG/DG
- Hierarchical multigrid tools
- Fault resilient solver
- Artificial neural networks

and:

- DSL toolchain
- Ensemble based URANIE
- State-of-the-art

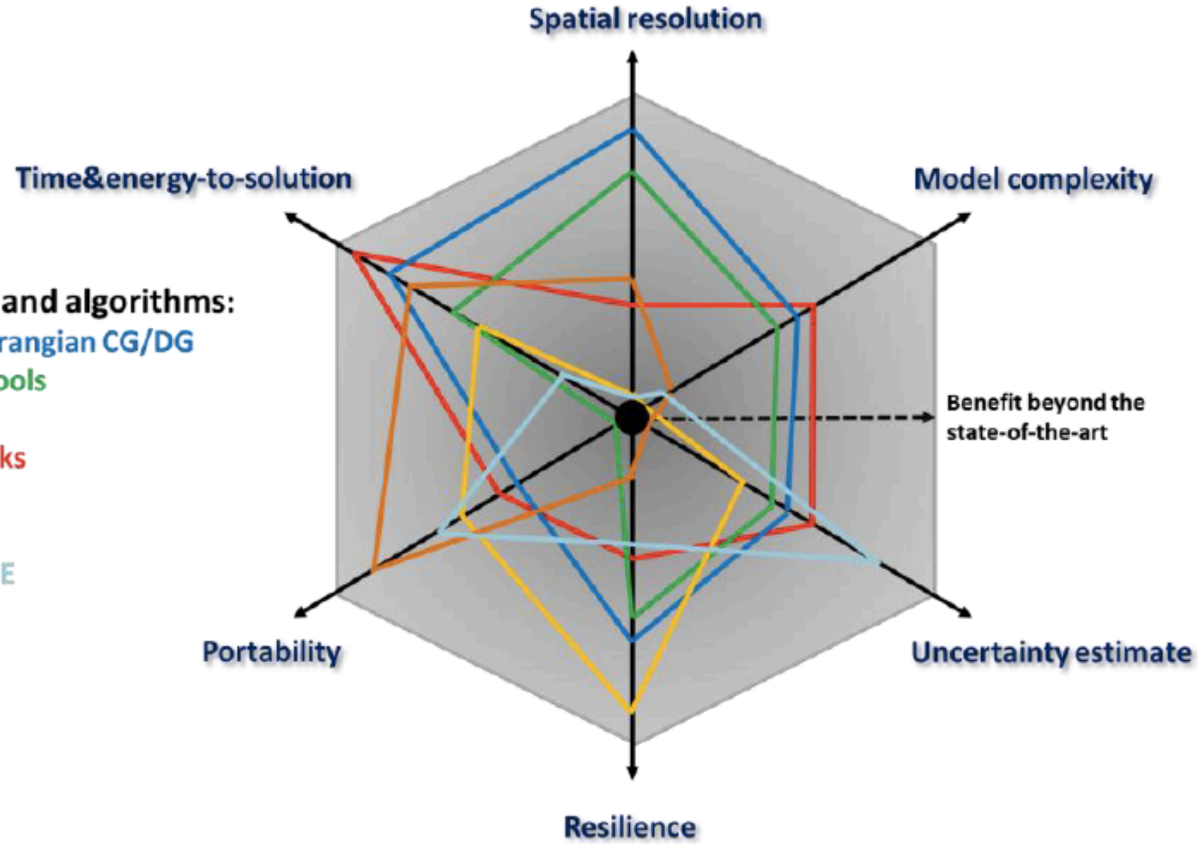

ECMWF · MeteoSwiss · DMI · RMI · NVIDIA · METEO FRANCE Toujours un temps d'avance · ICHEC · Deutscher Wetterdienst Wetter und Klima aus einer Hand · Optalysys · Bull atos technologies · PSNC · Loughborough University

ESCAPE-2: HPCW

Funded by the European Union

Linpack Benchmark → HPCG → **H**igh-**P**erformance **C**limate and **W**eather benchmark

HPCW benchmark tier	Specification	Options for novel developments to be included
Models	ICON ocean FV	Mathematics (finite-difference, time stepping), DSL
	NEMO ocean FD	Mathematics (time stepping), DSL
	IFS atmosphere FV	Mathematics (discretization, time stepping, fault tolerance), DSL
	IFS atmosphere DG	Mathematics (discretization, time stepping, fault tolerance), DSL
	IFS atmosphere ST	N/A (only as reference)
	ICON atmosphere FV	Mathematics (neural networks), DSL
Systems	Kronos workload simulator	Simulating the above

Roadmap for weather & climate computing

Funded by the European Union