ACCORD

A Consortium for COnvection-scale modelling Research and Development

Status of code refactoring and GPU adaptation

Daan Degrauwe, Piet Termonia

ACCORD ASW 2023

Motivation

- Top 500 HPC systems:

  16 out of 20 top systems have accelerators

- Green 500:

  40 first systems have accelerators

- EuroHPC infrastructure is targeted in the DE_330 project

- Trend towards using external HPC facilities for research and even operations

Rank	System	Cores
1	**Frontier** - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE DOE/SC/Oak Ridge National Laboratory United States	8,730,112
2	**Supercomputer Fugaku** - Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D, Fujitsu RIKEN Center for Computational Science Japan	7,630,848
3	**LUMI** - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE EuroHPC/CSC Finland	2,220,288
4	**Leonardo** - BullSequana XH2000, Xeon Platinum 8358 32C 2.6GHz, NVIDIA A100 SXM4 64 GB, Quad-rail NVIDIA HDR100 Infiniband, Atos EuroHPC/CINECA Italy	1,463,616
5	**Summit** - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592

Adaptation strategy

- Make adaptation as transparent as possible to science developers: principle of “separation of concerns”
- Make sure performance on CPUs is not impacted
- 3 pillars of code adaptation:
  - Smart (hardware-aware) data structures
  - Source-to-source transformation tools
  - Hardware-specific libraries

Activities

- CY48T3
- ARPEGE progress
- ALARO refactoring
- HARMONIE-AROME refactoring and PHYEX repository
- Granularity of parallel loops
- Familiarization with Loki
- Spectral transforms on GPUs

First steps are targeted at increasing the *flexibility* of the code

CY48T3

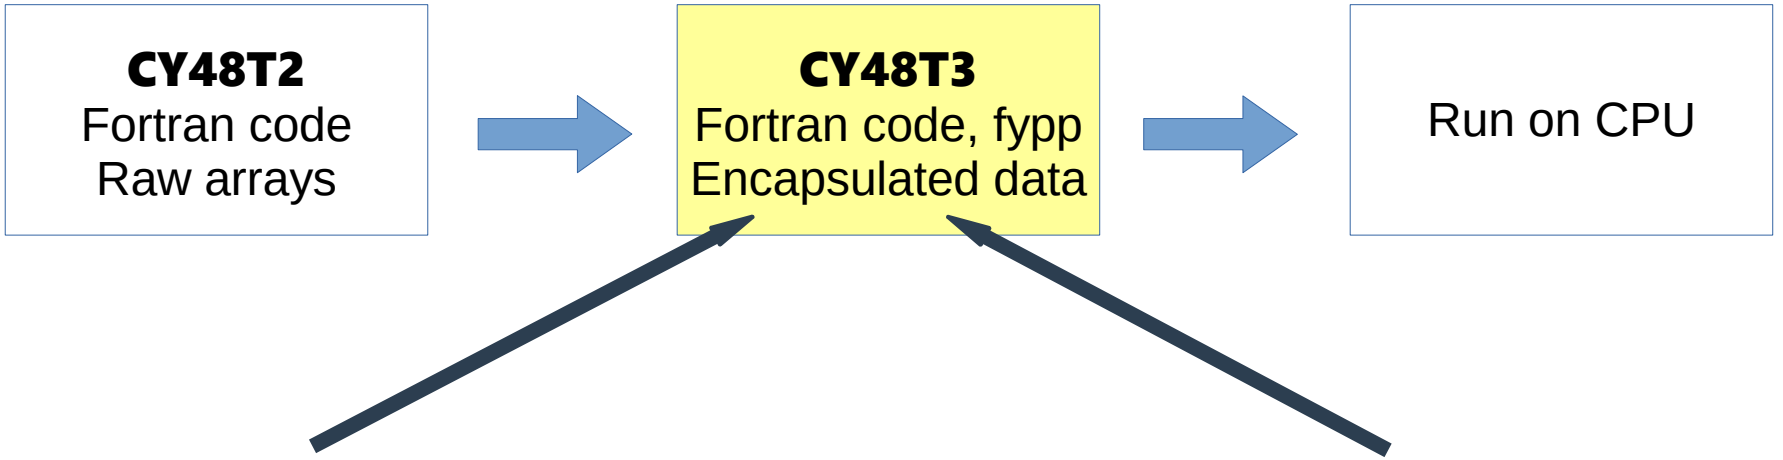

Smart data structures (aka FIELD_API):

- *flexible memory layout*
  separate underlying representation from usage in physics through pointers
- provide interface for accelerator specificities, e.g.
  - OpenMP or OpenACC?
  - Data exchanges between CPU and accelerator

Streamlining of MF_PHYS
Extraction of APL_ARPEGE from APLPAR

ALARO refactoring

CY48T2
Fortran code
Raw arrays

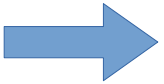

CY48T3
Fortran code, fypp
Encapsulated data

APL_ALARO

Refactoring:

- *Extract APL_ALARO from APLPAR*
  remove non-ALARO and obsolete options

- *Group calculations and subroutine calls belonging together,* e.g.
  - mixing
  - radiation
  - turbulence and surface
  - deep convection
  - microphysics
  - diagnostics

This grouping makes step-by-step porting to accelerators possible

Activities: HARMONIE-AROME work

CY48T2
Fortran code
Raw arrays

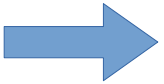

CY48T3
Fortran code, fypp
Encapsulated data

- AROME-MF and HARMONIE-AROME parameterizations are shared with the Meso-NH model
- To avoid divergence, these were put in a common external repository: PHYEX
- Requires changes in HARMONIE build system
- APL_AROME will also be cleaned/streamlined*

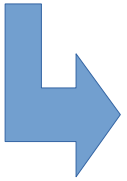

APL_AROME

PHYEX repository

Parallel granularity

The current code is coarsely granular

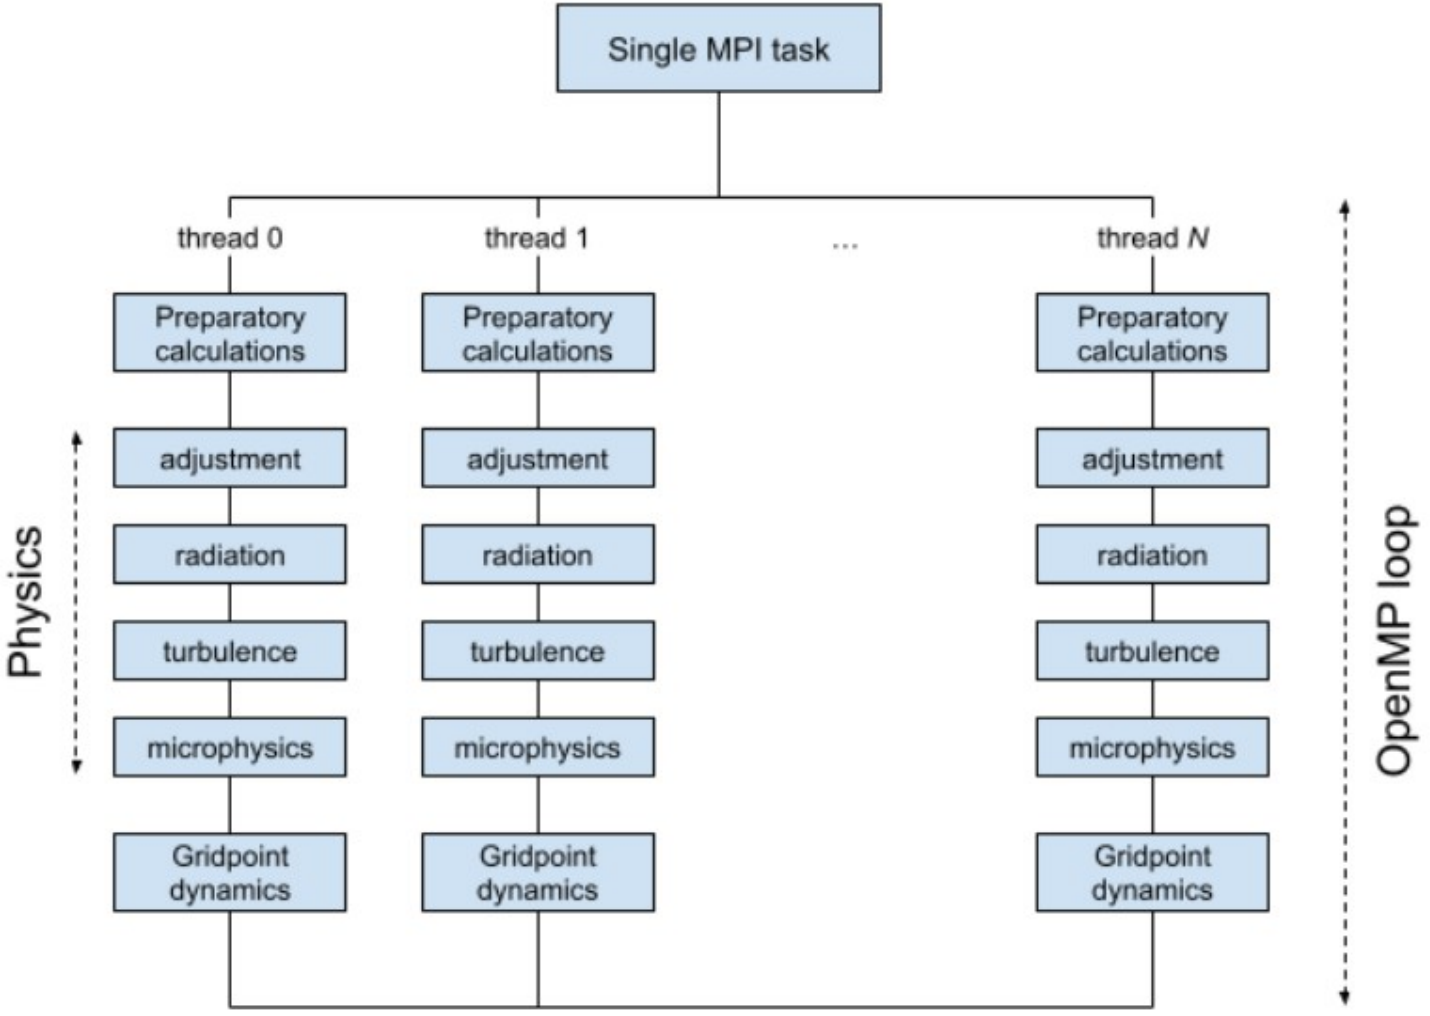

ACCORD
A Consortium for COnvection-scale modelling Research and Development

Parallel granularity

A more flexible code also allows finely granular parallelism

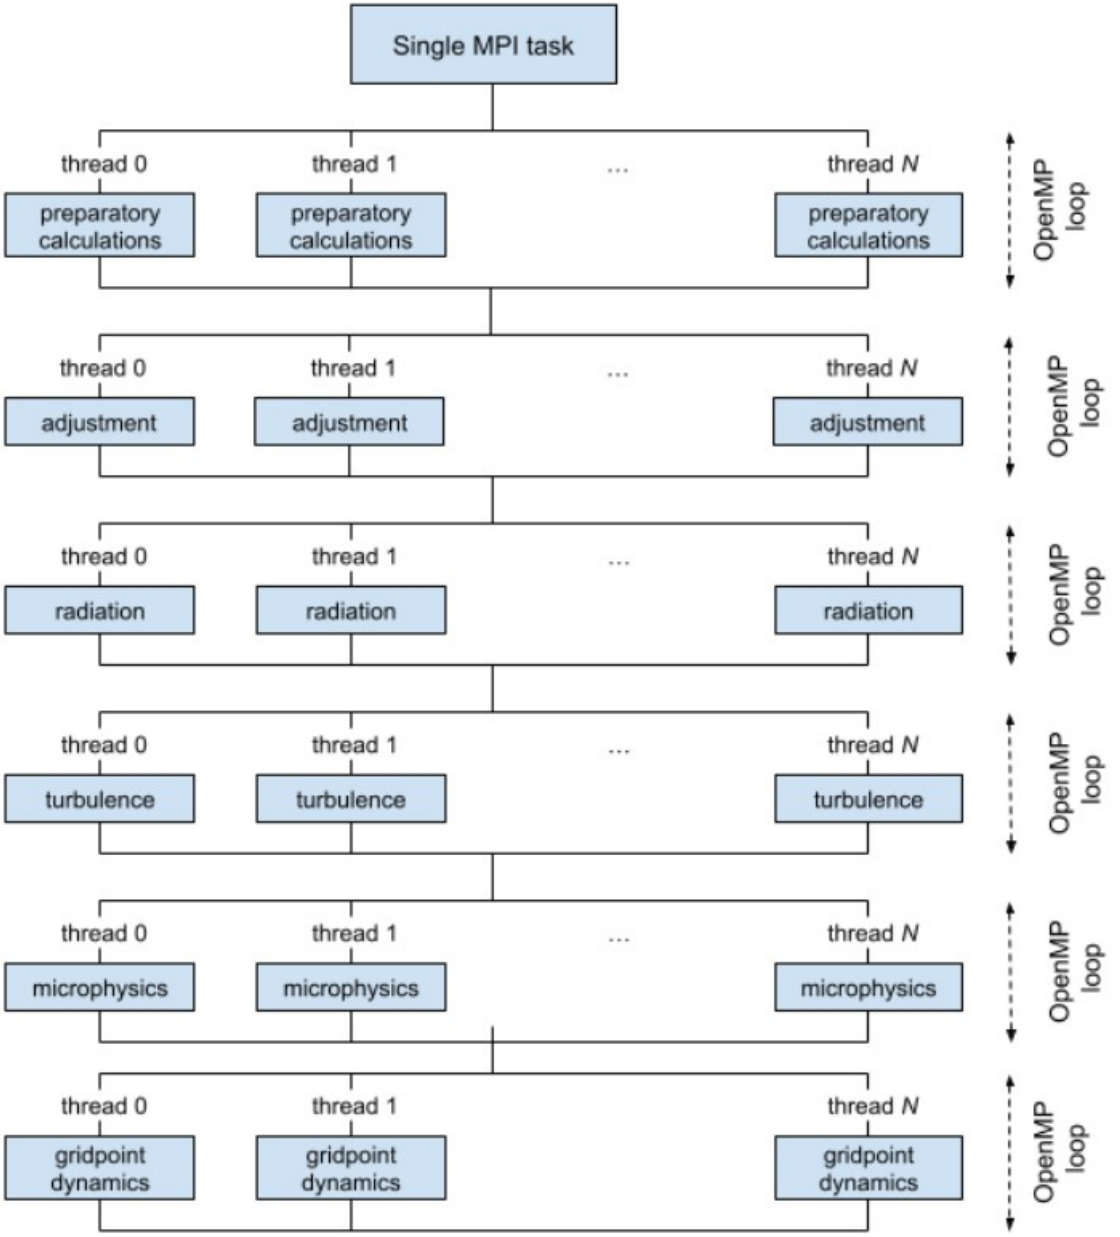

Parallel granularity

CY48T3
Fortran code, fypp
Encapsulated data

APL_ARPEGE/
APL_ALARO/
APL_AROME

Run on CPU

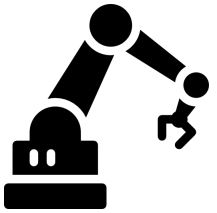

Fxtran-based scripts to

- Split the big OpenMP loop around the gridpoint calculations
- Encapsulate temporary arrays in smart structures

Finely granular parallel code

Verify on CPU

Source-to-source transformation

CY48T3
Fortran code, fypp
Encapsulated data

APL_ARPEGE/
APL_ALARO/
APL_AROME

Run on CPU

Source-to-source transformation tools (Loki, fxtran-based) to generate GPU-targeted code from CPU-targeted code

Finely granular parallel code

Run partially on GPU

Individual parameterization, e.g. ACRANEB2

ACRANEB2
GPU code

Spectral transforms on GPU

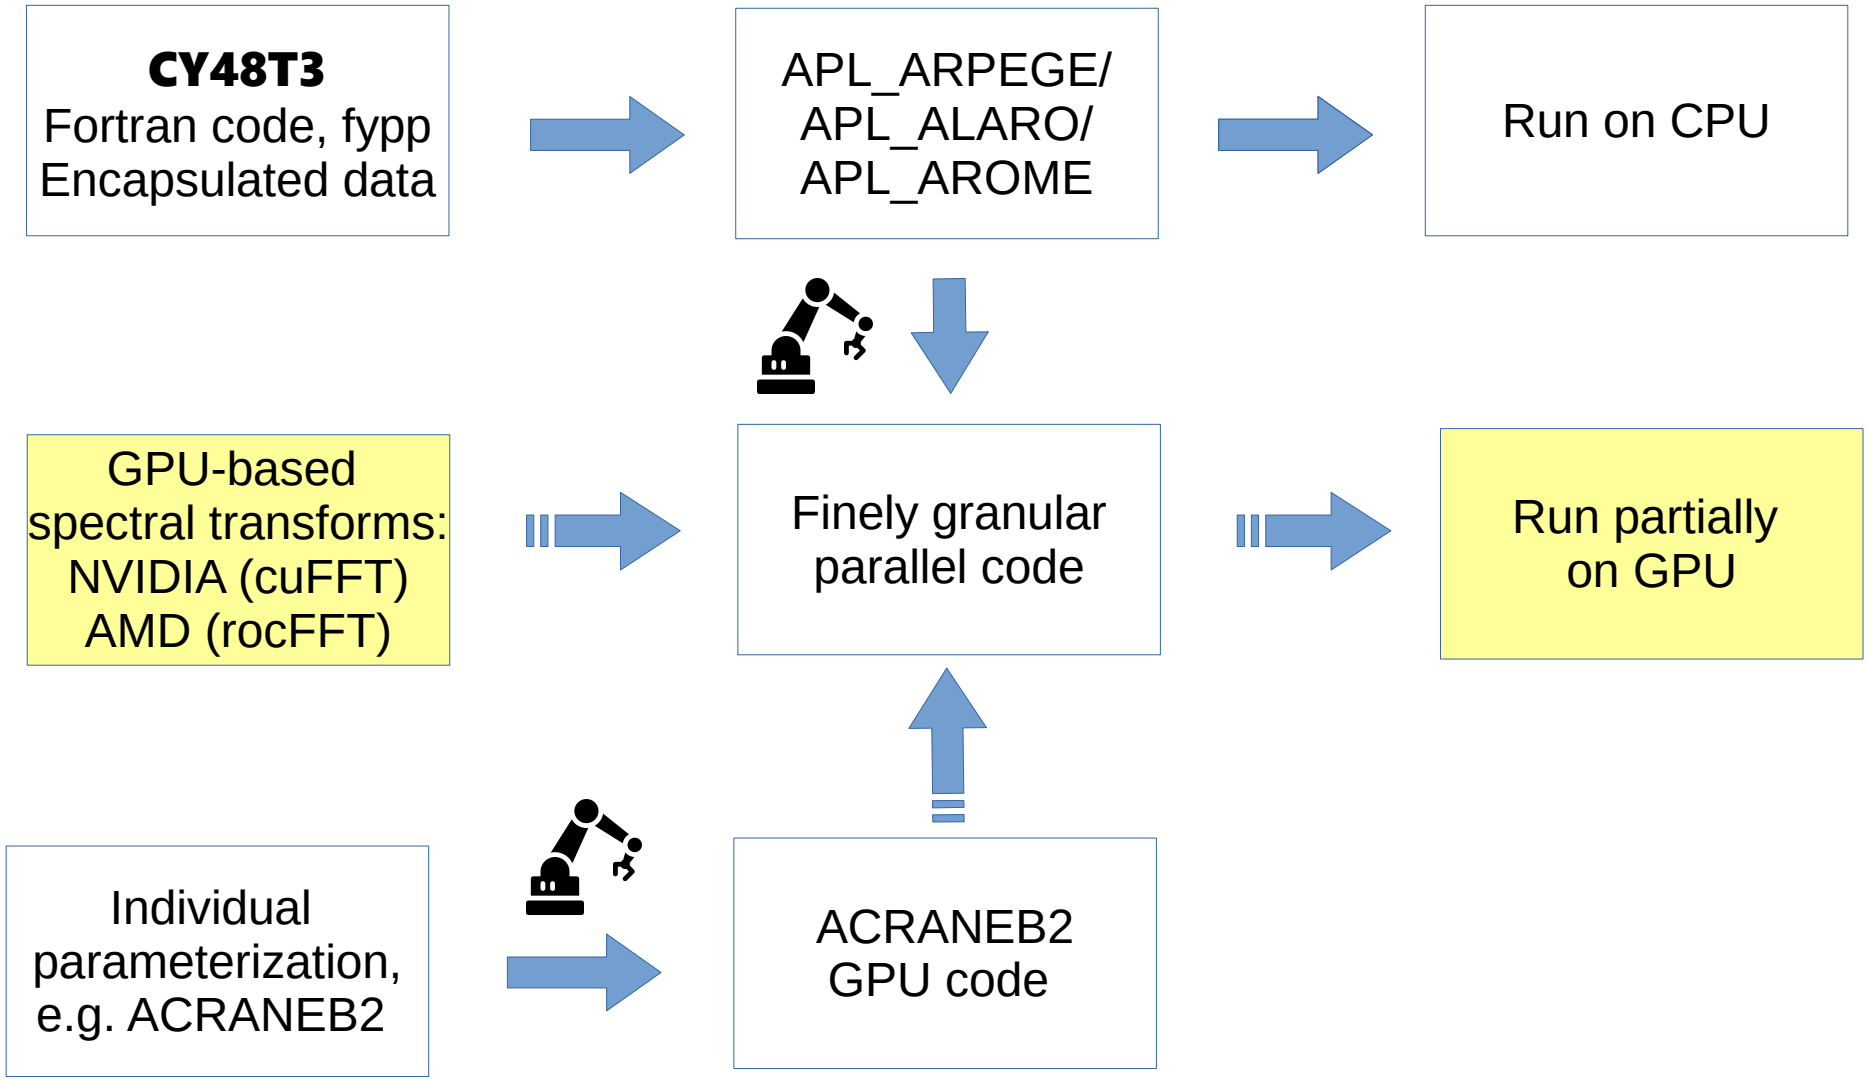

ACCORD
A Consortium for COnvection-scale modelling Research and Development

To conclude

- Adapting code to the future is a complicated task...

  … especially when trying not to disturb the scientists

- Many activities are ongoing, but they fit in a bigger scheme!

To conclude

- Adapting code to the future is a complicated task...

  … especially when trying not to disturb the scientists

- Many activities are ongoing, but they fit in a bigger scheme!

- The biggest highlight of last year is the contribution to code adaptation by so many of you!

Thank you!