European HPC application ecosystem

Guntram Berti, scapos

Overview

- CASTIEL2 & the CoEs
- CoE codes
  - The 14 HPC Centres of Excellence
  - CoE Codes starting from outer space ...
  - Some observations on the codes & their contexts
- Inno4scale
  - The 22 studies by domain
  - Approaches used for scalability
  - Some observations & thoughts ...

CoEs & CASTIEL2

- CoEs = Centres of Excellence: pan-European projects funded by the EuroHPC JU working on scaling leading codes for grand challenges
    - current set running until end of 2026
- CASTIEL2 – CSA for helping CoEs and National Competence Centers (NCCs) achieve their goals
    - running until end of 2025 (pot. extension)

A bit of history – CoEs since 2015

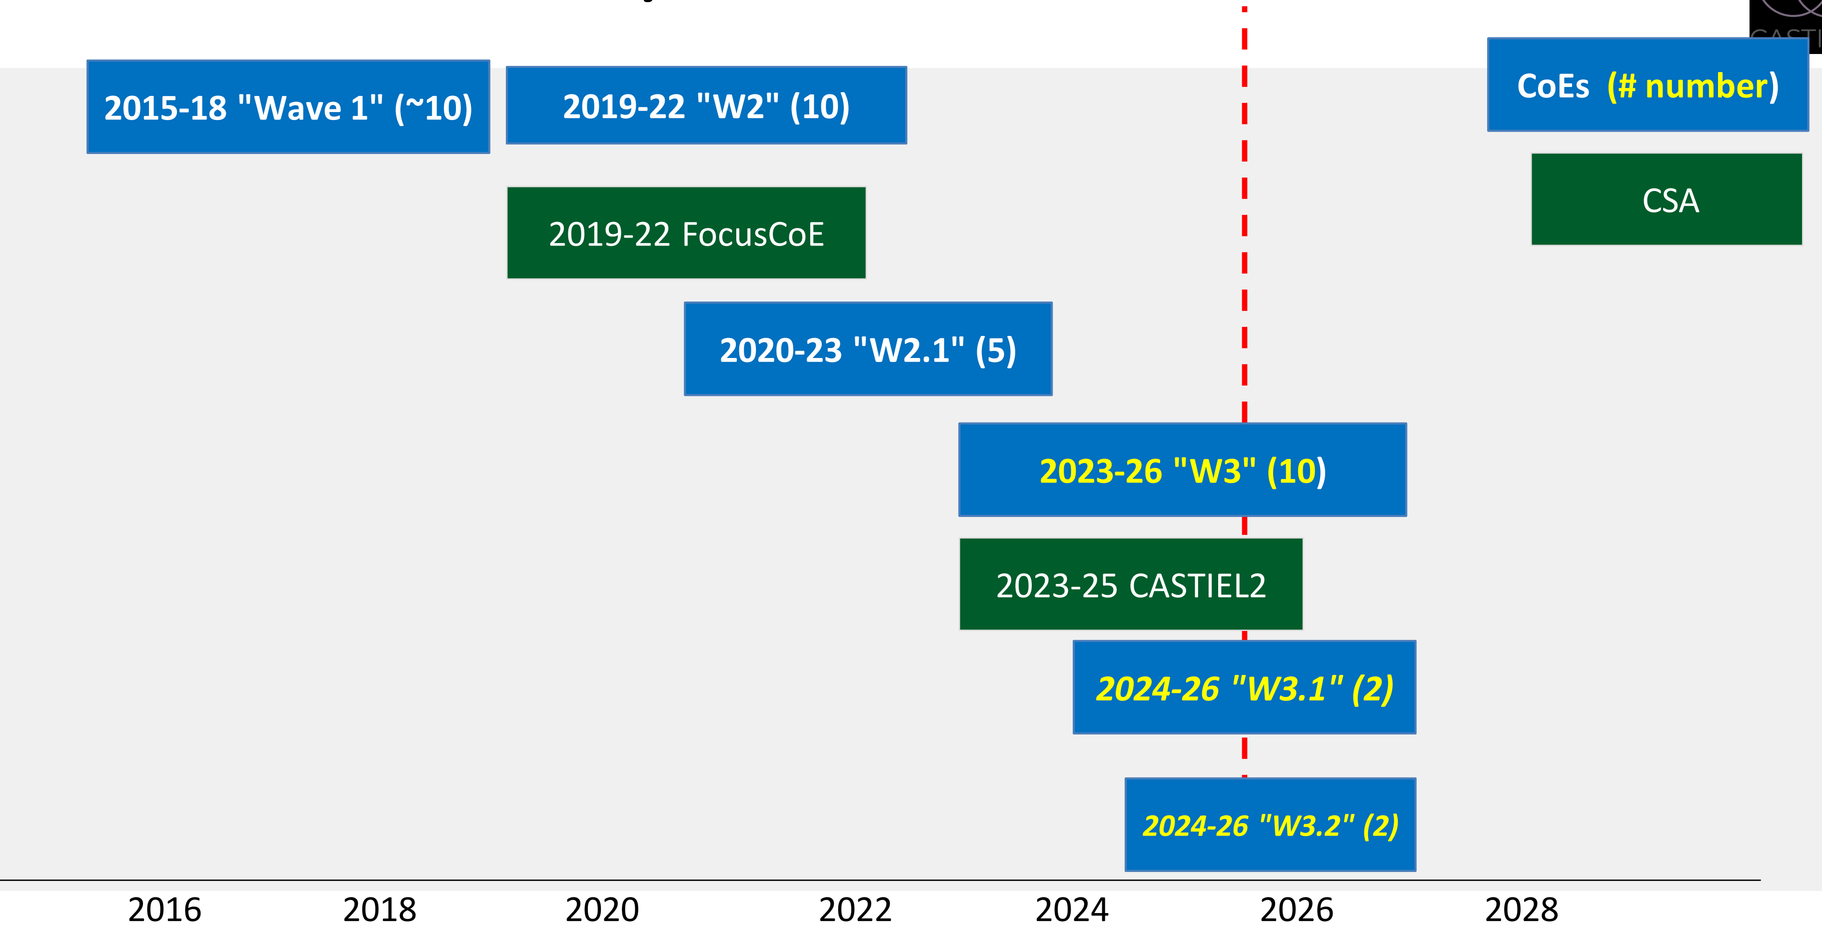

Current CoEs and their areas

CASTIEL 2

Earth & Space Systems

- SPACE
- esiwace – Centre of Excellence in Simulation of Weather and Climate in Europe
- ChEESE – Center of Excellence for Exascale in Solid Earth
- HiDALGO2 – Centre of Excellence

Fundamental Physics

- CEEC
- PLASMA PEPSC

Biology

- dealii X
- bioexcel – Centre of Excellence for Computational Biomolecular Research

Engineering & Tools

- MultiXscale
- POP
- EXCELLERAT

Materials

- MAX

Acronyms

CASTIEL 2

AMR	Adaptive Mesh Refinement (Cartesian)
CFD	Computational Fluid Dynamics
DFT	Density Functional Theory
DG	Discontinuous Galerkin
FD	Finite Differences
FEM	Finite Element Method
FV	Finite Volumes
HD	Hydrodynamics
LES	Large Eddy Simulation

LES	Large Eddy Simulation
MHD	Magnetohydrodynamics
N-S	Navier-Stokes
NWP	Nationa Weather Prediction
PIC	Particle-in-Cell
SEM	Spectral Element Method
SPH	Smoothed Particle Hydrodynamics
UG	Unstructured grid

Astrophysics & cosmology

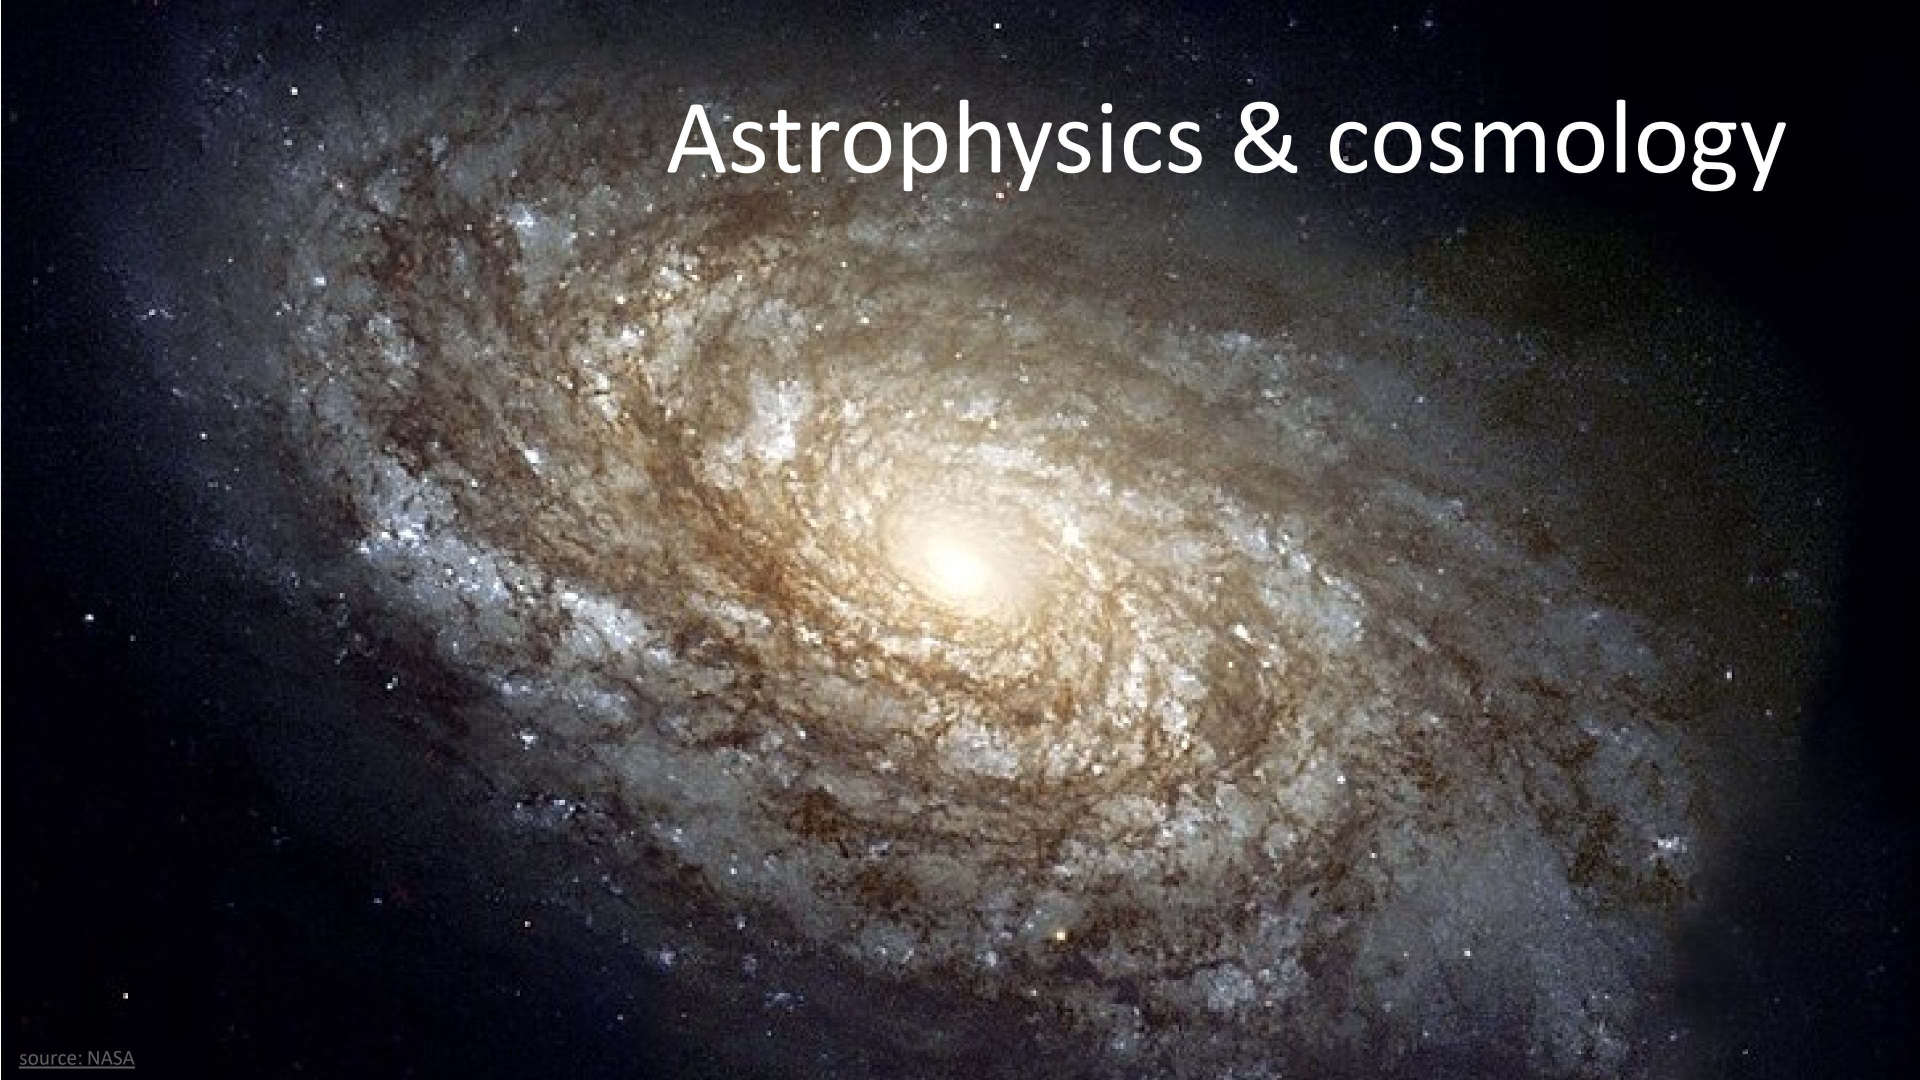

source: NASA

Astrophysics & cosmology

CASTIEL 2

Codes for modelling phenomena on cosmological scales, often including relativistic effects

Code	CoE	Description	Applications
OpenGadget	SPACE	N body simulations of cosmological volumes, galaxy clusters and isolated galaxies	Cosmological simulations of structure formation
CHANGA / GASOLINE	SPACE	SPH with tree-gravity solver for cosmological volumes and zoom-in galaxy simulations. Hydrodynamical (HD) simulations of astrophysical systems at all scales	Cosmology, galaxy formation, planet / planetary system formation
RAMSES	SPACE	Eulerian, cosmology, fully compressible, AMR (cell-by-cell), MHD, radiation hydrodynamics, cosmic rays, dust dynamics, particles (dark matter, stars, dust)	Astrophysical systems, featuring self-gravitating, magnetised, compressible, radiative fluid flows
FIL	SPACE	Magnetised fluids in dynamical curved space-times, magnetic field turbulence & properties of the ejected matter	Binary Neutron Star and Neutron Star - Black Hole mergers w. modelling of gravitational waves
BHAC	SPACE	Accretion flows in fixed space-time metrics, modelling electro-magnetic emission & imaging at radio wavelengths.	Black Hole accretion

Plasma physics

source: Varjisakka

Plasma physics #2

CASTIEL 2

Codes for modelling the 4th aggregate state of matter, characterised by ionised particles and long-range interactions

Code	CoE	Description	Applications
Vlasiator	Plasma PEPSC	Currently the only code able to model Earth's entire Magnetosphere using a hybrid-Vlasov approach, representing the kinetics of ions through their 6-dimensional phase space distribution function.	Space plasmas
iPic3D	SPACE	PIC method. Electrons and ions are described with computational particles and the electric, magnetic and gravitational field with Eulerian or Lagrangian grids. Energy conserving semi-implicit scheme (ECsim).	Plasma Physics, Astrophysics, Space weather, nuclear fusion
PLUTO	SPACE	Eulerian AMR-HD / MHD / RMHD code with Lagrangian module to treat sub-grid electron physics, thermal emission from highly energetic particles and dust grains.	Plasma physics, classical and relativistic magneto-hydrodynamics

Plasma physics

Code	CoE	Description	Applications
BIT	Plasma PEPSC	BIT1 / BIT3 are 1D3V/3D3V electrostatic Particle-in Cell Monte Carlo (PIC MC) codes. Parallel transport of plasma, neutral and impurity particles in the scrape-off layer (SOL) between core confined plasma and solid surface.	Plasma–Material Interfaces (e.g. tokamak fusion reactors)
GENE	Plasma PEPSC	Gyrokinetic turbulence code that can quantitatively describe the underlying physical processes, proven to match experimental data from existing mid-scale fusion devices	Magnetically Confined Plasmas
PIConGPU	Plasma PEPSC	Fully relativistic, manycore, 3D3V code using the particle-in-cell (PIC) method which describes plasma dynamics by computing the motion of electrons and ions based on Maxwell's equations.	Plasma Accelerators, High Energy Density Physics
GYSELA-X	EoCoE	flux-driven full-f nonlinear gyrokinetic code	plasma turbulence & transport from core to edge of Tokamak devices.

Climate & weather

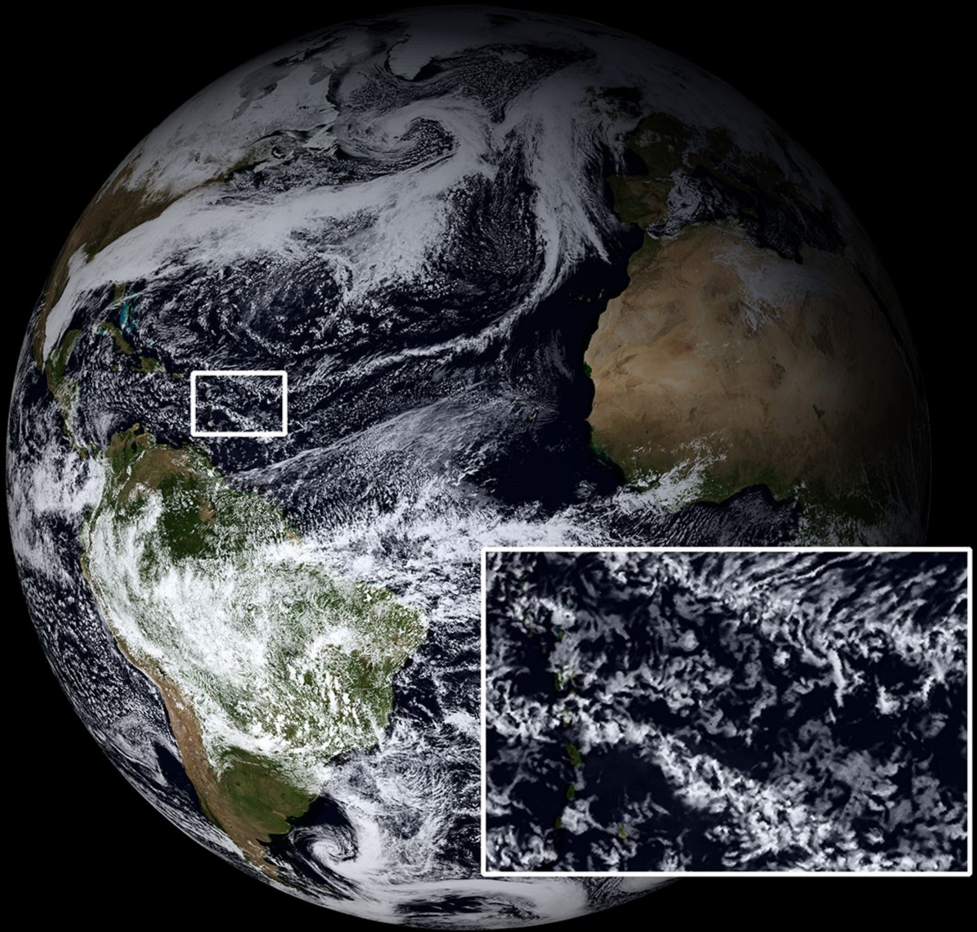

Source: ESiWACE

Climate & weather

CASTIEL 2

Codes / workflows for climate system components

Code	CoE	Description	Applications
ICON	ESiWACE	Climate and Weather prediction model used in production for NWP (DWD) and scientific applications	**Weather forecast** at DWD, part of DE national earth simulation modelling strategy
IFS	ESiWACE	Global weather prediction model used in operational numerical weather predictions at ECMWF	Operational weather predictions at ECMWF. Use for Extremes and Climate Digital Twins of Destination Earth. Open component (OpenIFS) used for in EC-Earth
NEMO	ESiWACE	Nucleus for European Modelling of the Ocean used for oceanographic, forecasting and climate studies	Operational Seasonal Predictions system, Global Ocean, Mediterranean and Black Sea Forecasting & Reanalyses at CMCC
EC-Earth4	ESiWACE	**Earth System** Climate Model, Multi-physics / -components coupled model	Climate projections, climate predictions, process studies

Regional scale:
CFD / coupled workflows

Source: NASA

Cumbre Vieja, La Palma, Dec 2021

CFD-coupled workflows #1

CASTIEL 2

Domain-specific workflows centred around a CFD-type component

WF	CoE	Description	Applications
UAB: Urban Air Project (OpenFOAM)	HiDALGO	OpenFOAM (uap-foam) for urban scale simulation of pollution dispersion	**Pollution spread** analysis: traffic emission, chemical havaries, air quality assessment
Wildfires	HiDALGO	Simulation of wildfire-atmosphere interactions and smoke dispersion at various scales, based on WRF, FireFOAM	Assess **wildfire** risk & impacts near wildland–urban interface (WUI) zones. to improve local resilience and resistance
OpenPDAC (OpenFOAM)	ChEESE	Eulerian-Eulerian compressible multiphase mixture of gasses and fine particles, Lagrangian transport for coarser particles.	Explosive **volcano eruption** phenomena incl. blasts, ash dispersal, pyroclastic flows, ballistic ejection
FALL3D	ChEESE	Advection Diffusion Sedimentation (ADS) equations, structured terrain-following grid using a finite volume (FV) explicit scheme.	Forecast & re-analysis of atmospheric dispersal phenomena & particle ground deposition (**volcano eruptions**)

CFD-coupled workflows #2

CASTIEL 2

WF	CoE	Description	Applications
HySEA suite	ChEESE	Codes for tsunami simulations, incl. seismic sources and landslide generated tsunamis or meteo-tsunamis, optional dispersion, hydrodynamics either one-layer shallow water model or multilayer for vertical resolution	**Tsunamis**: generated by earthquake, aerial, submarine landslide, and meteo-tsunamis
Renewable Energy Sources (RES)	HiDALGO	WRF (community mesoscale weather prediction model) coupled with improved EULAG (all scale geophysical flow solver).	Prediction of **energy produced by wind farms and PV** plants. Prediction of damages to the Distribution System Operator's infrastructure.
Elmer/ICE	ChEESE	Multi-physics package Elmer solving coupled visco-elastic Reynolds flows representing Glacier Outburst Floods (GLOFs).	**Ice-flow**, permafrost and groundwater

Workflows: Thermal, energy, ...

Domain-specific workflows centred around thermal or other component

WF	CoE	Description	Applications
MTW: Material transport in Water	HiDALGO	Coupled Fluid-Particle-Scalar simulation of material transport in water. Based on waLBerla (massively parallel LBM framework for multiphysics applications) coupled with a parallel rigid particle dynamics (RPD) simulation	**Pollutant transport** in aquatic environments: microplastics, chemical- or temperature contaminants on newest HPC technologies.
Parflow	EoCoE	Integrated **hydrology model** that simulates spatially distributed surface and subsurface flow, as well as land surface processes including evapotranspiration and snow	Groundwater modelling framework, combined with the hydropower infrastructure management system HYPERstreamHS high-resolution hydrological forecasting at the continental scale


Computational
Seismology
& Geodynamics
Grindavik
Source: NASA

Computational Seismology #1

CASTIEL 2

Codes / workflows to simulate seismic waves, ruputure / earthquakes

WF	CoE	Description	Applications
SPECFEM3D	ChEESE	Linear seismic wave propagation (elastic, visco-elastic, poro-elastic, fluid-solid) & dynamic rupture, heterogeneous models. High-order spectral-elements, u/s hex meshes	Simulation of earthquake events, seismo-acoustic simulation and seismic imaging
SeisSol	ChEESE	Acoustic, anisotropic-/visco-/poro-elastic rheology, dynamic rupture laws. High-order DG Galerkin & Arbitrary DERivative (ADER) local time stepping, u/s adaptive tetrahedral meshes.	Simulation of complex earthquake and earthquake-tsunami events

Computational Seismology #2

CASTIEL 2

Codes / workflows to simulate seismic waves, ruputure / earthquakes

WF	CoE	Description	Applications
ExaHyPE (ExaSeis)	ChEESE	high-order DG, with ADER time stepping on tree-structured dynamically adaptive Cartesian meshes.	earthquake simulation model, Landslide/tsunami model in development
TANDEM	ChEESE	DG code on unstructured curvilinear grids for linear elasticity problems and sequences of earthquakes and aseismic slip.	Simulation of sequences of earthquakes and aseismic slip

Geodynamics

Codes / workflows simulation of "slower" geodynamics

WF	CoE	Description	Applications
LaMEM	ChEESE	Thermo-mechanical geodynamical processes for rocks w. visco-elasto-plastic or (compressible) poroelastic rheologies. Marker-in-cell with staggered FD	Geodynamics, such as mantle-lithosphere interaction, or the dynamics of magmatic systems
PTatin3D	ChEESE	Package for long time-scale geological processes, using matrix-free stokes operator	Geodynamics incl. continental rifting, subduction and the dynamics of lava domes
xSHELLS	ChEESE	Rotating incompressible flows & magnetic fields in spherical geometries, incl. geodynamo simulations. Combines spherical harmonic transforms & FD in radius to obtain a fast time-stepper.	Earth's core and magnetic field; also liquid interior of other planets

Engineering: CFD & Multi-physics

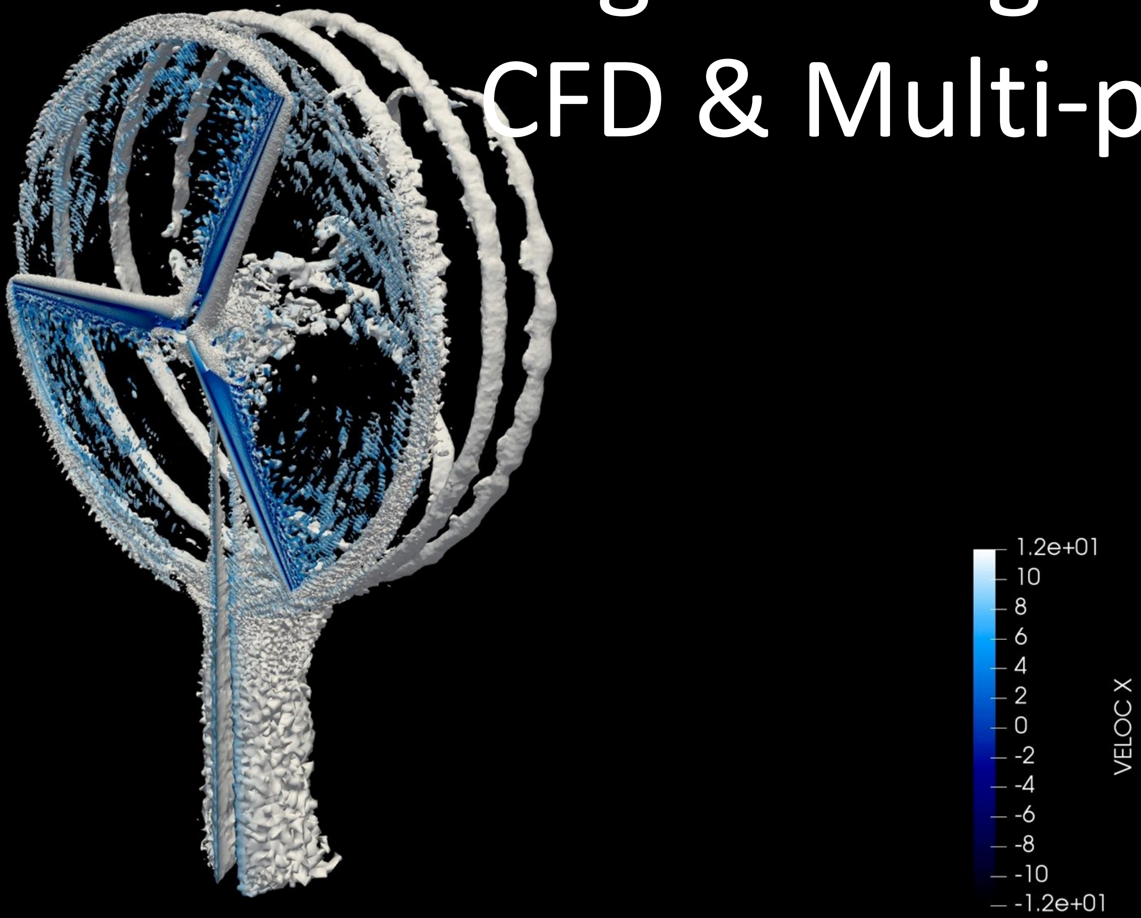

Source: EoCoE II

CFD codes #1

Codes that are entirely or primarily focused on classical CFD (no complex rheology)

Code	CoE	Description	Applications
AVBP	EXELLERAT	Compressible 3D reactive N-S solver using the LES approach on unstructured grids.	Aerospace / Safety / Energy / Transport
CODA	EXELLERAT	Reynolds-Averaged N-S equations on UGs based on 2nd order FV and higher-order DG	Aerospace
FLEW	EXELLERAT	High-order FD solver for compressible flow on structured body-fitted meshes generalized curvilinear coordinate system.	Aerospace
waLBerla	CEEC MultiXScale	Multi-physics simulation software framework using Lattice-Boltzmann method, focus on scalability and large-scale simulations	CFD, particles

CFD codes #2

Codes that are entirely or primarily focused on classical CFD (no complex rheology)

Code	CoE	Description	Applications
Neko	CEEC EXCELLERAT	Portable framework for high-order spectral-element-based simulations, focusing on the incompressible regime	High-fidelity CFD, turbulence, engineering
Nek5000	CEEC	Incompressible and low Mach flow solvers employing the Spectral Element Method, a high-order weighted residual technique for spatial discretization that can accurately represent complex geometries	CFD on complex geometries
NekRS	CEEC	Refactored and GPU oriented version of Nek5000	

Multi-physics codes

CASTIEL 2

Codes that are can couple multiple physical models (often CFD-oriented)

Code	CoE	Description	Applications
Alya	CEEC, EXELLERAT.. EoCoE	Multi-physics coupled: flow, solid mechanics, particle transport, heat, chemistry	Engineering, multi-physics problems
FLEXI	CEEC	High-order solver for general PDEs of hyperbolic / parabolic-type (CFD focus) based on discontinuous spectral element method (DG-SEM). US hexahedral grids.	Multi-physics / multi-scale applications, turbulent flows and shock-turbulence interaction
m-AIA	EXCELLERAT	Multi-physics PDE solver framework, focus on CFD, computational aero-acoustics and structural mechanics.	Engineering applications
deal.II	dealii-X	Wide range of finite element algorithms, multigrid methods	Engineering, physiologie

Workflows: Thermal, energy, ...

Domain-specific workflows centred around thermal or other component

WF	CoE	Description	Applications
OpenFOAM / Raysect	EXCELLERAT	OpenFOAM - FV solver for heat diffusion, Raysect Ray tracer for heat radiation based on Monte Carlo integration	**Nuclear fusion**
Urban Buildings	HiDALGO	Fell++ based toolset for geometry generation and multi-physics models (e.g. thermal, air quality, energy)	**Building energy simulation** from building to city scales and possibly beyond

Living tissues and organs

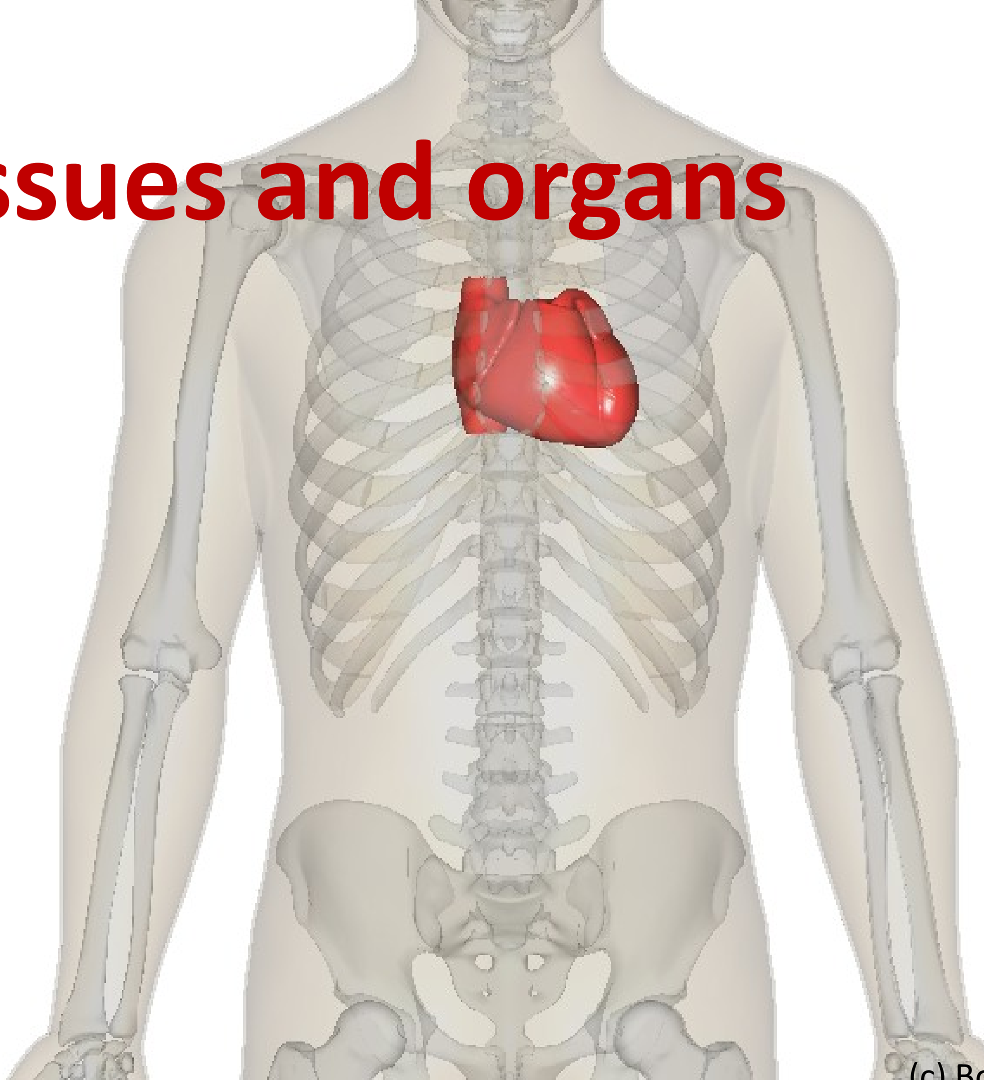

Living tissues and organs

CASTIEL 2

Codes for modelling at cellular, tissue and organ level

Code	CoE	Description	Applications
OpenCARP	MICROCARD2	Cardiac electrophysiology simulator offering single cell as well as multiscale simulations ranging from ion channel to organ level	Cardiac electrophysiology
lifeX	deallii-X	A high performance library for the numerical solution of complex finite element problems	Cardiac function
ExaDG	dealii-X	Numerical solution of partial differential equations in the field of computational fluid dynamics	Physiological applications

Molecular mechanics

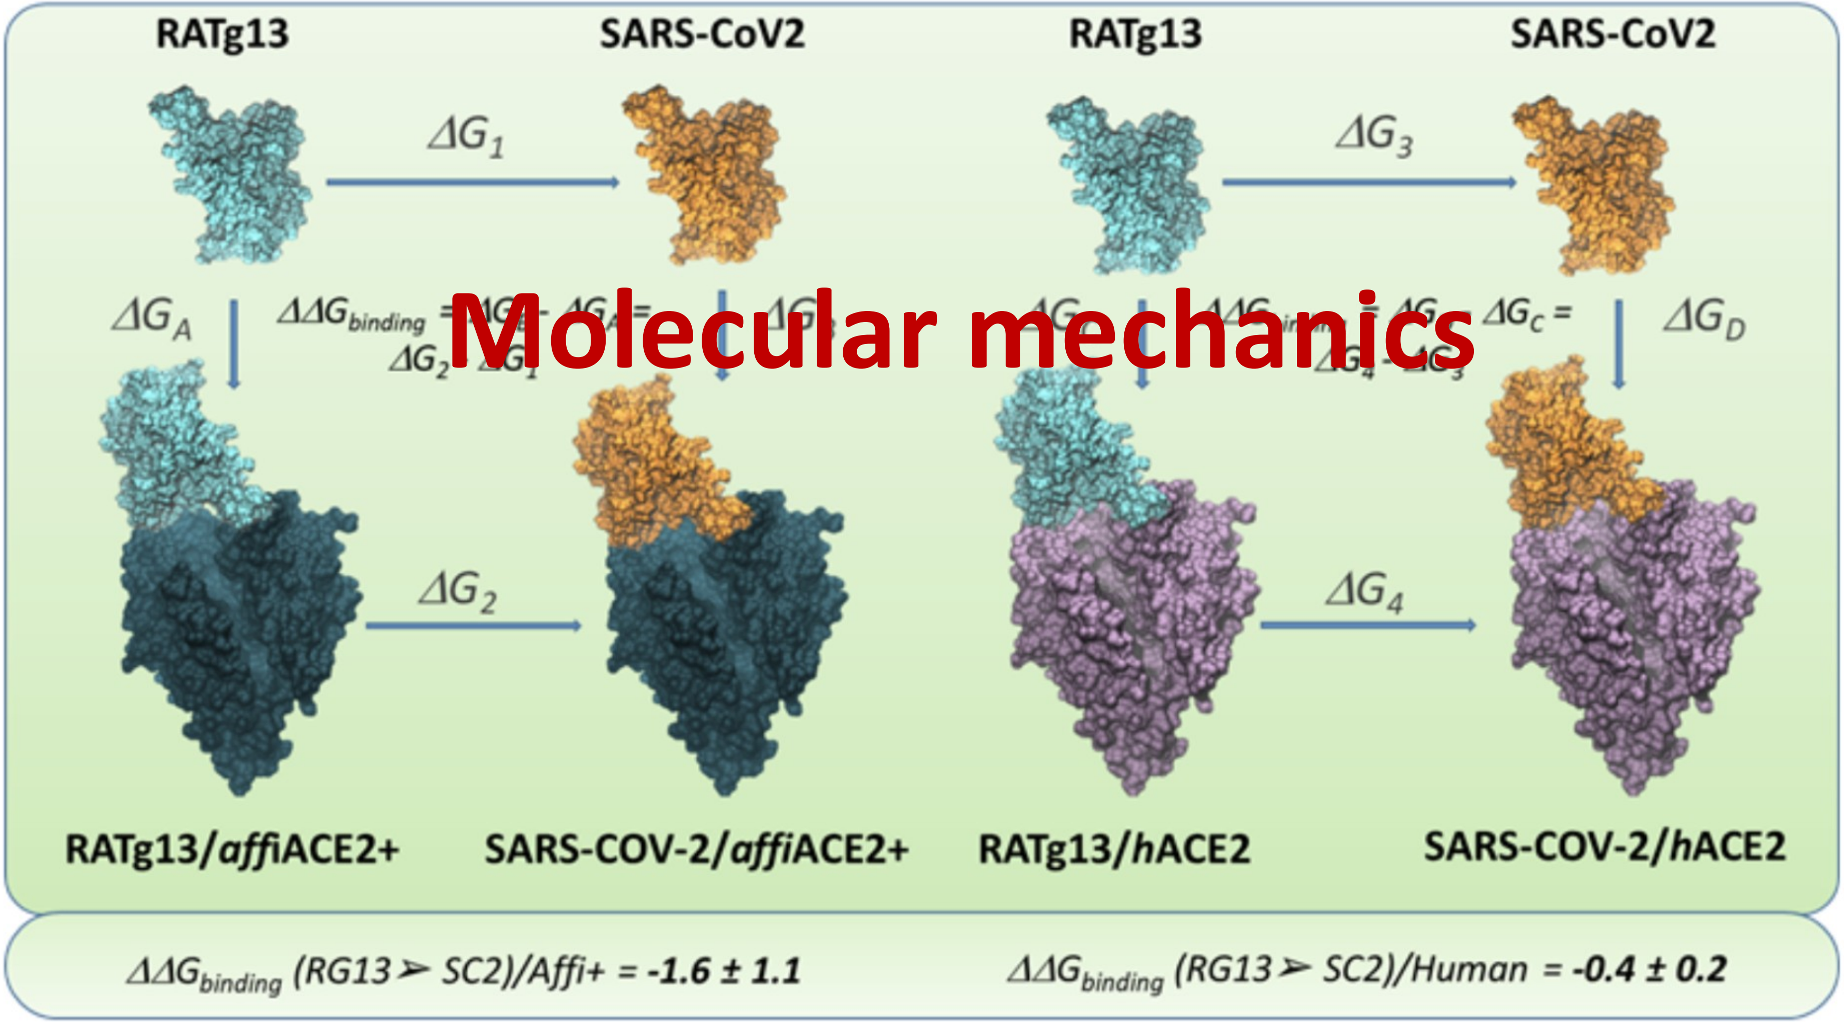

Source: BioExcel

Molecular mechanics

CASTIEL 2

Codes for modelling at molecular level using classical Newtonian mechanics

Code	CoE	Description	Applications
HADDOCK	BioExcel	HADDOCK (High Ambiguity Driven protein-protein DOCKing) is an information-driven flexible docking approach	Modeling of biomolecular complexes
GROMACS	BioExcel	Molecular dynamics simulation suite	Biomedical (biomolecules), polymeres, nanoflow
ESPResSO	MultiXScale	Particle-based simulation package for molecular dynamics, fluid dynamics and Monte Carlo schemes	Electrostatics, magnetostatics, hydrodynamics, electrokinetics and diffusion-advection-reaction
LAMMPS	MultiXScale	Classical molecular dynamics code	Simulating materials on atomic, meso- and macroscopic spatio-temporal scales, solid state materials, biomolecules, soft matter and coarse-grained systems

Quantum mechanics

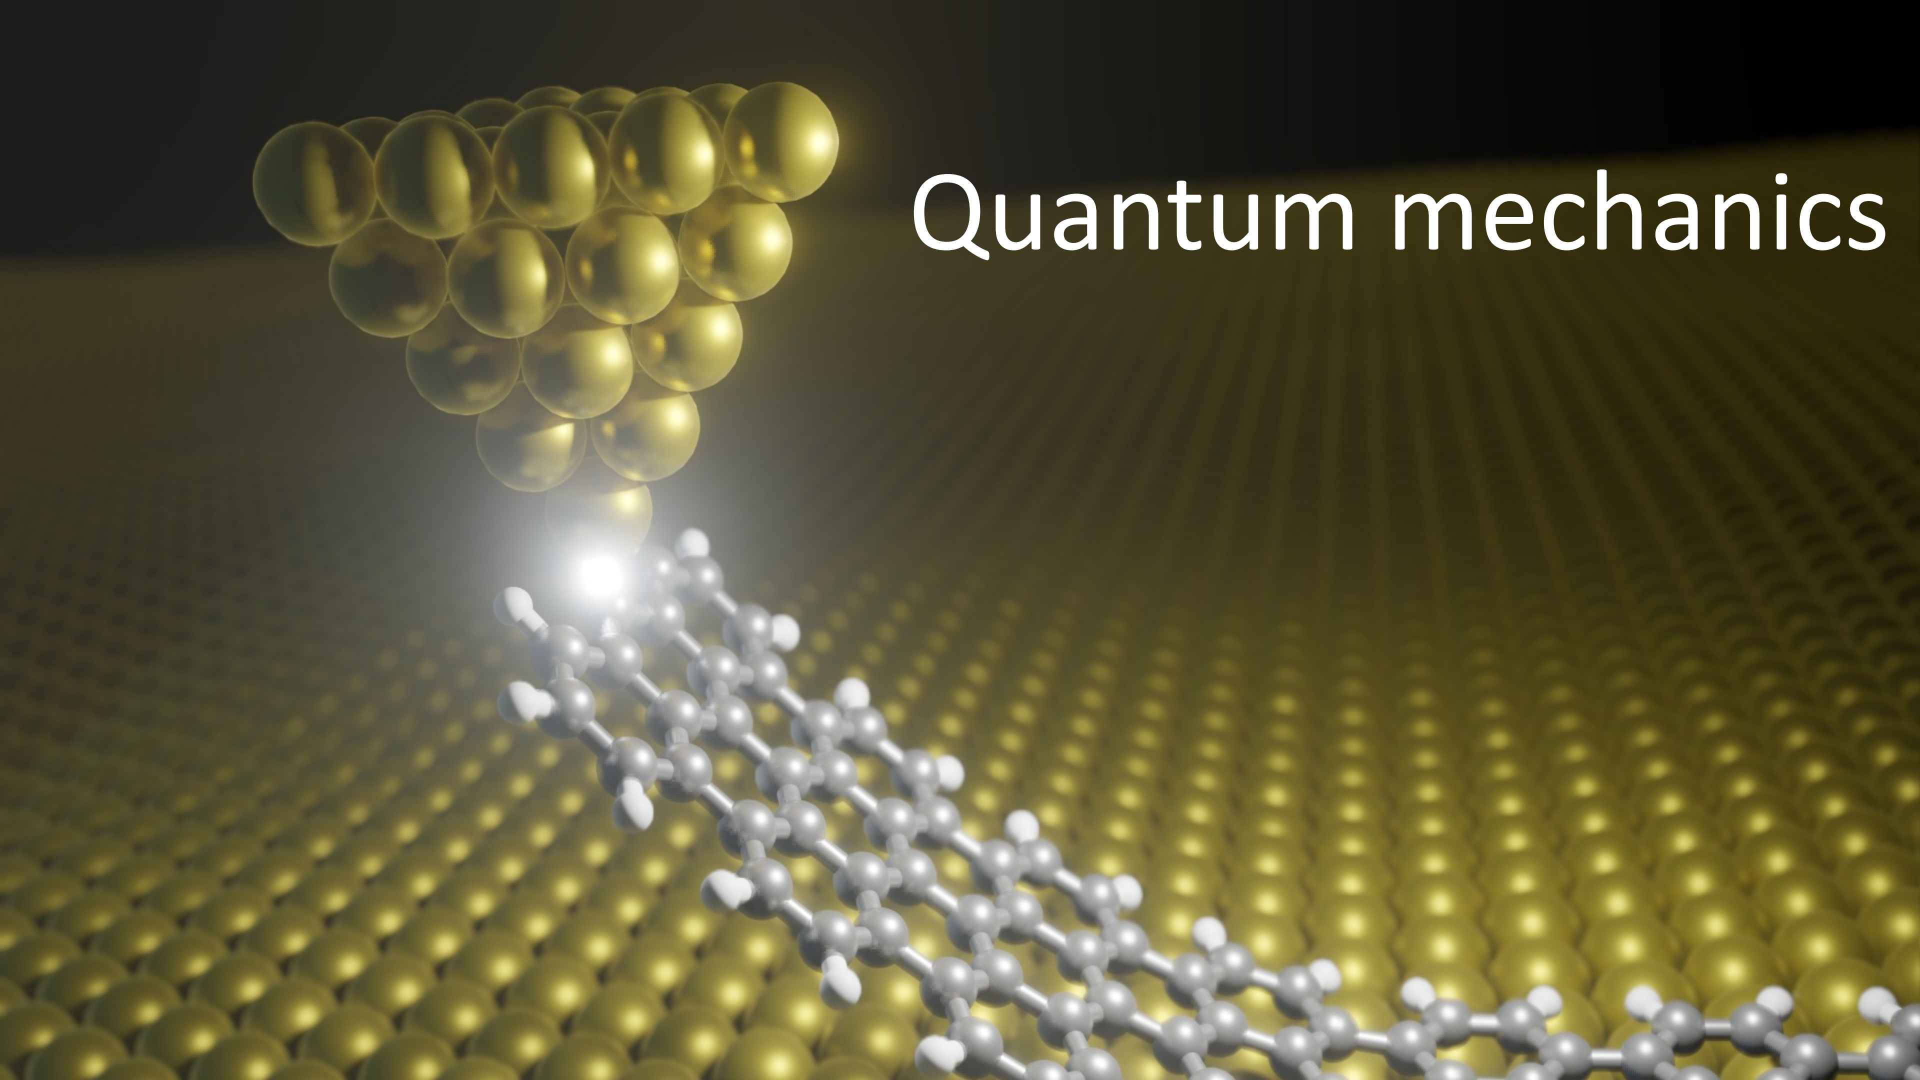

Quantum mechanics #1

Codes for modelling at electron level using quantum mechanics ("ab initio")

Code	CoE	Description	Applications
BigDFT	MaX	Wavelet-based density functional theory (DFT) code	Electronic structure simulations: From Groundstate to advanced properties
FLEUR	MaX	All-electron DFT code based on the full-potential linearized augmented plane wave (FLAPW) method	Electronically and magnetically complex materials, e.g. involving transition metals, heavy or rare-earth elements. Derive magnetic or spin-dependent properties in metals or complex oxide materials
Quantum ESPRESSO	MaX	Integrated suite of codes for electronic-structure calculations and nano-scale materials modelling. Based on density-functional theory, plane waves, & pseudopotentials.	Groundstate calculations. Molecular dynamics Search for transition pathways (NEB). Linear response to ext. perturbations. Dielectric responses, vibrational spectra. Optical, magnons, and EELS spectra, electronphonon coupling coefficients. Self-consistent Hubbard correction parameters (HP)

Quantum mechanics #2

Code	CoE	Description	Applications
SIESTA	MaX	Pseudopotential-based DFT software. Strength: use of atomic-like strictly-localised basis sets: enabling the use of reduced-scaling methods	Calculation of energies and forces, molecular-dynamics simulations, band structures, densities of states, spin-orbit couplings, van der Waals functionals, hybrid functionals.
YAMBO	MaX	The code implements Many-Body Perturbation Theory (MBPT), DFT and Non-Equilibrium Green's Function Theory (NEGF)	Ground-state & excited-state properties Reliable band gaps, band alignments, defect quasi-particle energies, optical and non-equilibrium properties
libNEGF	EoCoE	Non-equilibrium Green's function library, combined with density functional theory (DFT) models of material properties to enable new advances in photovoltaics	models of material properties, e.g. to enable new advances in photovoltaics

~~Tools~~ & mathematical libraries

Code	CoE	Description	Applications
Ginkgo	MICROCARD2	numerical linear algebra library for many-core systems, with a focus on solution of sparse linear systems	
MUMPS	dealii-X	A numerical software package for solving sparse systems of linear equations	
PSCToolkit	dealii-X EoCoE	Framework for solving large and sparse linear systems	

CoE codes - Some observations

- Parallelisation frameworks uses
  - Inter-node communication: MPI use @ (almost) 100%
  - Intra-node communication: diverse picture
- Tools & libraries
  - diversity of tools for parallel development / sw engineering / workflow management / pre- and postprocessing
  - diversity of libraries used for mathematical sub-problems (linear algebra et al)
- Some difficulties / challenges
  - GPU scalability for irregular types of algorithms
  - Load balancing for irregular / dynamic applications
  - Scaling the entire workflow / toolchain
  - Some apps have heavy memory footprints (including needs for huge input data)
  - SW engineering: Maintaining high-performance code for broad range of diverse hardware

CASTIEL 2

Some insights from Inno4scale

Inno4scale - outline

- Goal: Instill development of proof-of-concepts for new / promising algorithmic approaches for the Exascale era
- Approach: Cascading funding for "innovation studies" by external research groups
  - 22 studies funded, running for 12 months, finished Q1/2025

Inno4scale – study overview

CASTIEL 2

- **Sparse LA** – AceAMG (low-comm AMG), aCG (low-comm & directGPU-GPU comm), AMCG (asynchr. CG)
- **Irregularity in applications**: CMB4Scale (compressed adj. matrices), CVolBal (balancing comm), NEOSC (CPU co-processing)
- **LLMs / ANNs** ESPLAG (pruning)
- **CFD** Ex3S (asynch / MixP), ExaSIMPLE (ML for SIMPLE), FLOWGEN (online ML), Scale-TRACK (Euler-Lagrange), STRAUSS (task-parallel, GPU / GPU comm)
- **Parallel-in-time (PinT)** MLMC-PinT4Data (multi-level UQ), NeuralPinT (ML for coarse level), TiPOFlow (sim & opt)
- **Materials (atomic & particle scale)** Exa4GW, ISOLV-BSE (improved Eigensolve), LimitX (ML for eigensolve) , XCALE (ML), MG4ML (particle physics), ASTERIX (compression / Plasma)
- **Uncertainty Quantification (UQ)** ScalaMIDA (Geophysics – multilevel), *MLMC-PinT4Data*

Inno4scale – approaches to scalability

CASTIEL 2

- **Pruning / reducing numerically low-yield computation**
  - Pruning NN models
  - replacing repetitive computations by ML
  - cutting excess precision (mixed precision) or accuracy (weakly contributing to soln)
  - coarsening approximation (multi-level)
  - lossy compression (good enough for restart / analytics)
- **Exploiting special properties** of mathematical problems
- **Cutting through layers** & overheads reduction
  - direct GPU / GPU communication
  - reducing unneeded communication
- **Work load balancing**
  - CPU <-> GPU load balancing e.g Euler vs Lagrange approaches
- **Increasing parallelism**
  - balancing communication, reducing synchronisation
  - parallelisation in time
  - Moving up the ladder, e.g. UQ (parallelisation across multiple runs)

Inno4scale – observations

- Many good results
  - ... but not everything worked (yet?)
- Reduction of total work makes "good" scalability harder
  - serial baseline shifts
  - more complex control flows and data structures vs. SIMD-type accelerators
- Applications are complex & multi-faceted
  - large overall reductions are rare
- "Break-through" type speeds-up often require combination of methods


Conclusion:
Scalability
– no low
hanging
fruits?

Thanks

CASTIEL 2 has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101102047. The JU receives support from the Digital Europe Programme and Germany, Italy, Spain, France, Belgium, Austria.

Inno4scale has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101118139. The JU receives support from the European Union's Horizon Europe Programme.