

Centres of Excellence in HPC in CASTIEL 2

A summary of their consortium, objectives, area of activity, codes and use cases.

Authored by CASTIEL 2 - Work Package 2

This work was/is partially supported by CASTIEL 2. CASTIEL2 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.

01.01.2023 - 31.12.2026
Life Sciences
https://bioexcel.eu/
3rd Phase - Type 2 (Apps)

Coordinator: KTH Royal Institute of Technology - SE (PI: Rossen Apostolov)	University of Utrecht NL
Barcelona Supercomputing Centre ES	IRB Barcelona ES
EMBL-EBI DE	MPG DE
Nostrum Biodiscovery ES	Normal Richard Albert NO

BioExcel's mission is to provide applications, tools, support, and networking opportunities to Life Science researchers, enabling them to address grand scientific challenges by fully exploiting the power of data and computing e-infrastructure. Powerful and sophisticated software packages are critically important for fundamental and applied research, including industrial development, in particular in the areas of drug design, biotechnology, food and chemical industries.

- Service Portal: https://bioexcel.eu/services/
- Trainings: https://bioexcel.eu/services/training/
- Potential Collaboration with Industry (not partners): Yes

USE CASES

- Antibody Design through Biomolecular Interactions Engineering
- High-Throughput Modelling of Interactomes
- Electronic Interaction Phenomena: Proton Dynamics and Fluorescent Proteins
- Rational Drug Design

CODES & WORKFLOW ENGINES

- **GROMACS:** one of the most used molecular dynamics codes. Highly tuned assembly kernels with hybrid CPU/GPU execution.
- **HADDOCK:** used for integrative modelling. Python wrapper to execute a series of CNS (Crystallography and NMR system) scripts that also run pre and post-processing analysis.
- **PMX:** used to automate the free energy calculations and assess binding between drug and protein. It uses GROMACS as a backend engine.
- **BioBB:** Python library consisting of a collection of wrappers on top of biomolecular tools (e.g. GROMACS, HADDOCK, PMX), adding interoperability between them, and allowing an easy creation of biomolecular simulation workflows. The library combined with BSC PyCOMPSs workflow manager can be used to launch the workflows in HPC systems using thousands of cores.

CEEC

Centre of Excellence in Exascale CFD

01.01.2023 - 31.12.2026
Exascale CFD
https://ceec-coe.eu/
1st Phase - Type1 (Exascale)

Coordinator: KTH Royal Institute of Technology - SE (PI: Niclas Jansson)	AUTh EL
Barcelona Supercomputing Centre ES	FAU - University Erlangen DE
University of Stuttgart (HLRS) DE	Umea University SE
Denmark Technical University DK	BAM DE

CEEC's ambition is to enable the use of exascale computers for key computational fluid dynamics (CFD) applications and demonstrate their capabilities through key light-house cases. Understanding the turbulent behavior of gases and liquids has direct societal impact. In fact, it will be key for enabling the transition to a carbon-free economy through optimizations like more fuel efficient airplanes and cargo ships or wind turbines that can be more easily and sustainably installed on the ocean floor. CEEC aims at: implement exascale-ready workflows to solve grand scientific challenges, develop new or improves algorithms that can efficiently exploit exascale systems, improve energy efficiency of simulations, and demonstrate workflows on lighthouse cases that are relevant for academia and industry.

- Service Portal: -
- Trainings: https://ceec-coe.eu/training/
- Potential Collaboration with Industry (not partners): Yes, but not primary target

USE CASES

- Shock – Boundary layer interaction and buffet on wings at the edge of the flight envelope. - FLEXI
- High fidelity aeroelastic simulation of the SFB 401 wing in flight conditions - ALYA
- Topology optimization of static mixers - Neko
- Localized erosion of an offshore wind-turbine foundation - waLBerla
- Simulation of Atmospheric Boundary Layer (ABL) flows - NekRS/Nek500
- Merchant ship hull - Neko

CODES

- **ALYA**: Multi-physics / multi-scale simulation framework
- **FLEXI**: DG-Spectral element based CFD code
- **Nek5000 & NekRS**: Spectral element based CFD Codes
- **Neko**: Spectral element based CFD code
- **waLBerla**: Lattice Boltzmann based multiphysics framework

ChEESE

01.01.2023 - 31.12.2026
Solid Earth
https://cheese-coe.eu/
2nd phase - Type 2 (Apps)

Coordinator: CSIC ES (PI: Arnau Folch)	
INGV IT	Ludwig Maximilian Univerity of Munich DE
Barcelona Supercomputing Centre ES	University of Mainz DE
CINECA IT	University of Malaga ES
Icelandic Metereological Office IS	Norwegian Geotechnical Institute NO
ETH Zurich CH	IPGP F
EVIDEN F	CNRS F
University of Stuttgart (HLRS) DE	Ruder Boscovic Institute HR
Technical University of Munich DE	CSC FI
Sorbonne University F	E4 Analytics IT

ChEESE-2P aims to prepare 11 community flagship codes to address 12 domain-specific Exascale Computational Challenges (ECCs) on, eg., computational seismology, magnetohydrodynamics, physical volcanology. Codes will be optimized in terms of performance on different types of accelerators, scalability, deployment, containerization, and portability. Codes and workflows will combine to form a new generation of 9 Pilot Demonstrators underpinned by concepts like multi-scale, multi-source, and multi-physics, that will materialise in 15 Simulation Cases representing capability and capacity use cases of particular relevance in terms of science, social relevance, or urgency.

- Service Portal: -
- Trainings: https://cheese-coe.eu/events/training
- Potential Collaboration with Industry (not partners): Yes - also public administration targeted

EXASCALE COMPUTATIONAL CHALLENGES

- Full-waveform inversion and high resolution tomography (capability computing)
- Near real-time seismic scenations (capacity computing)
- Physics-based Probabilistic Seismic Hazard Assessment (capacity & urgent computing)
- Near real-time tsunami source inversion and characterization (capability computing)
- Faster Than Real Time tsunami scenarios (capacity & urgent computing)
- Probabilistic Tsunami Hazard Assessment (capacity computing)
- Volcanic Explosions and Interactions (capability computing)
- Probabilistic Volcanic Hazard Assessment (capacity computing)
- The dynamo model and the Earth's magnetic field evolution (capability computing)
- Landslide triggering (capability computing)
- Lithospheric stress states and induces seismicity (capability computing)

CODES

- **SeisSOL:** complex earthquake simulation
- **SPECFEM3D:** linear seismic wave propagation
- **ExaHyPE:** special emphasis on **ExaSeis** collection of seismic models
- **Tandem**: linear elasticity + sequences of earthquakes and seismic slips
- **xSHELLS**: flows, magnetic fields including Geodynamo simulations
- **HySEA**: earthquake –induced Tsunamis
- **FALL3D:** atmospheric transport & ground deposition
- **OpenPDAC:** simulation of volcanic fluids
- **LaMEM:** thermo-mechanical geodynamic modelling
- **pTatin3D**: simulation of long time-scale geodynamics processes
- **ELMER/ICE:** Multi-physics package

esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

01.01.2023 - 31.12.2026
Weather - Climate
https://www.esiwace.eu/
3rd Phase - Type 2 (Apps)

Coordinator: Barcelona Supercomputing Centre **ES** **(PI: Mario Acosta)**	NLeSC NL
ECMWF UK	SMHI SE
ATOS F	CSC FI
DKRZ DE	University of Helsinki FI
CMCC IT	Juelich Forschungzentrum DE
Max Planck Institute for Metereology DE	Latest Thinking DE

ESiWACE3 focuses on the community's support to reach a higher readiness level regarding exascale supercomputing and knowledge transfer between the different Earth System modelling centres and teams across Europe. The project focuses on three main pillars: (i) the transfer and establishment of knowledge and technology for efficient and scalable simulations of weather and climate across the Earth system modelling community, (ii) closing common technology knowledge gaps and providing toolboxes for high-resolution Earth system modelling via joint developments and (iii) serve as a sustainable community hub for training, communication and dissemination for high-performance computing for weather and climate modelling in Europe. ESiWACE3 will bring the various approaches to address these challenges from the different modelling groups together to transfer knowledge across the weather and climate domain, generate synergies between the local efforts, provide targeted support to modelling groups via customised high-performance computing services, and provide training to educate the next generation of researchers.

- Service Portal: https://www.esiwace.eu/services
- Training events: https://www.esiwace.eu/training
- Potential Collaboration with Industry (not partners): Yes, but it is not the primary target.

OBJECTIVES

- Increase the efficiency of weather and climate simulations on state-of-the-art supercomputers.
- Design tools to close technology gaps for high-performance computing.
- Develop tools to tackle the data challenge of high-resolution weather and climate modelling.
- Support the wider weather and climate modelling community using state-of-the-art supercomputers via targeted services, training and capacity building.
- Build a well-connected and inclusive community for high-resolution Earth System modelling across Earth system science and HPC.

CODES & BENCHMARKING TOOL

- **EC-Earth:** coupled Earth System Model for climate simulations
- **NEMO**: ocean model for climate simulations
- **ICON:** coupled Earth System Model for climate simulations
- **IFS:** Atmospheric model for weather simulations
- **HPCW:** Set of weather & climate benchmarks
-

EXCELLERAT P2

01.01.2023 - 31.12.2026
Engineering - Manufacturing
https://www.excellerat.eu/
2nd phase - Type 2 (Apps)

Coordinator: Univ. of Stuttgart (HLRS) - DE (PI: Bastian Koller)	
Barcelona Supercomputing Centre ES	DLR DE
CINECA IT	University of Roma "La Sapienza" IT
SIPEARL F	RWTH AACHEN DE
University of Ljubljana SI	SICOS DE
KTH Royal Institute of Technology SE	SSC DE
ARCTUR SI	Teratec F
CERFACS F	Fraunhofer SCAI DE

EXCELLERAT P2 spans the whole domain of engineering – ensuring that innovations produced are applicable to other engineering challenges, that co-design is handled properly, that success stories are created and distributed. EXCELLERAT P2 strives to show the contribution of HPC, HPDA and AI in engineering to a low-carbon, more environmentally friendly and socially responsible product development and manufacturing as well as the mobility and energy sector. The holistic view of EXCELLERAT P2 to the usage of HPC (and HPDA and AI) is needed, if engineering applications shall be supported to increase the European competitiveness.

- Service Portal: https://services.excellerat.eu/
- Trainings: https://services.excellerat.eu/searchevents/training
- Potential Collaboration with Industry (not partners): Yes, IAC concept (Industry-Academia Collaboration): coordination, funding and association from industry-academia collaboration project

USE CASES

1. External aircraft aerodynamics - Aircraft and Aerospace - DLR - CODA
2. Hydrogen combustion for propulsion - Automotive, Aircraft & Aerospace - CERFACS - AVBP
3. Migration of aeroacoustic noise - Automotive, Aircraft & Aerospace, Energy - RWTH Aachen - m-AIA
4. Fully integrated aircraft simulations with emission models - Aircraft & Aerospace - BSC - Alya
5. High-fidelity simulations of rotating parts -Aircraft & Aerospace, Energy KTH - Neko
6. Active control for drag reduction of transonic airfoils - Aircraft & Aerospace- CINECA - FLEW
7. Engineering design of digital twin of the first wall of tokamak fusion reactor - Energy - University of Ljubljana - OpenFOAM/Raysect

CODES

- **m-AIA**: CFD-related, multi-physics code in C++ with hybrid MPI/OpenMP parallelization and some solvers ported to NVIDIA GPUs
- **AVBP**: 3D parallel Navier-Stockes solver for reactive multi-phase multi-physics problems using Large Eddy Simulation on unstructured grids (runs on any CPU and NVIDIA GPUs).
- **Neko**: SME code running on CPU and GPUs using native accelerator kernels (HIP, Cuda, OpenCL).
- **Alya**: Navier-Stokes solver exploiting a Spectral Element discretization. Runs fully on GPUs (mostly tested on Nvidia V100 and A100 GPU).
- **CODA**: CFD software, Python frontend, C++ backend, supports hybrid MPI/GASPI + OpenMP + NVIDIA GPUs
- **OpenFOAM:** free, open source CFD software
- **Mitsuba2**: Monte Carlo Ray tracing for optical and radiative heat transfer simulations (runs on GPUs)
- **FLEW:** FVM solver for DNS of compressible flow, written in Fortran and OpenACC.

Coordinator: PSNC - PL (PI: Marcin Lawenda)	Univeristy of Strasbourg F
ATOS ES	ICCS EL
University of Stuttgart (HLRS) DE	Meteogrid SL ES
SZE HU	Future Needs CY

01.01.2023 - 31.12.2026
Simulations for global challenges
http://www.hidalgo2.eu/
2nd phase - Type 2 (Apps)

HiDALGO2 explores synergies between modelling, data acquisition, simulation, data analysis and visualisation. It will also efficiently utilise current and future HPC and AI infrastructures to develop highly scalable solutions to global climate and social challenges (e.g. violent weather, floods, pollution). HIDALGO2 aims at bringing together advanced solutions (HPC-AI-HPDA) to provide decision makers and stakeholders tools that would mitigate tragic consequences of climate and civilization phenomena, by delivering necessary knowledge.

- Service Portal:
- Trainings: -
- Potential Collaboration with Industry (not partners): Yes (also targets public services)

USE CASES & CODES

- Urban Air Project: evolution of air in urban areas considering pollution, wind, comfort and planning. The core is the Urban Air Flow (UAP.AF) computational model, based on two main software: OpenFOAM and Fluid-Solver. For both solvers, modelling of more physical properties will be developed including thermal convection, solar heat radiation, transport and reaction between several species.
- Urban Building: advanced building models for better integration with architecture and for pollutants and heat models.
- Renewable Energy Sources: Solutions for production of renewable energy accustomed to urban and rural areas. Multiscale weather prediction based on Fortran and MPI coupled energy production estimation (based on AI/HPDA)
- Wildfires: Simulations of wildfires and atmosphere interaction and smoke dispersion in forest and urban areas. Forecast of weather conditions from WRF and LES at very detailed scales, generating ensembles of weather scenarios over the landscape and coupling of CFD solutions for the detailed modeling of air flow at the settlement scale.

All HiDALGO2 Pilots will be integrated into a newly developed dashboard, a key entry point for internal and external stakeholders of these global challenges.

Additional transversal activities:

- Setup an efficient benchmarking and profiling process
- Optimize pilot towards efficient execution
- Combine simulations and AI models
- Evaluate impact of emerging architectures and technologies

MAX DRIVING THE EXASCALE TRANSITION

01.01.2023 - 31.12.2026
Material Design
http://www.max-centre.eu/
3rd phase - Type 1 (Exascale)

Coordinator: CNR IT (PI: Elisa Molinari)	ICN2 ES
Barcelona Supercomputing Centre ES	SISSA IT
CINECA IT	E4 IT
SIPEARL F	Jozef Stefan Institute SI
CEA F	IT4I CZ
ATOS BULL F	Leonardo IT
Julich Forschungzentrum DE	University of Modena&Reggio IT
CSIC ES	University of Bremen DE

MaX-3 targets lighthouse codes for material simulations aimed at understanding, predicting, and discovering the properties and performance of materials. The aim is: to upscale the MaX codes and their performance to multiple heterogeneous exascale architectures; to endow these codes with innovative capabilities enabled by such architectures; to co-design the hardware and software in collaboration with the relevant European stakeholders; to enable turn-key simulation capabilities that meet the power of exascale resources and deliver the resilience needed; to disseminate the entire ecosystem of codes, workflows, and data; and to train and engage developers and users in fully leveraging such powerful instruments for discovery and innovation.

- Service Portal: http://www.max-centre.eu/services
- Trainings: http://www.max-centre.eu/training-and-education
- Potential Collaboration with Industry (not partners): Yes (but not primary target)

CODES

- **Quantum ESPRESSO** is an integrated suite of open-source computer codes for electronic-structure calculations and materials modeling at the nanoscale. It is based on density-functional theory (DFT), plane waves, and pseudopotentials.
- **SIESTA** is a code that performs efficient electronic structure calculations and ab initio molecular dynamics simulations of molecules and solids. SIESTA's efficiency stems from the use of a basis set of strictly-localized atomic orbitals.
- **YAMBO** is an open-source code implementing first-principles methods based on Green's function theory to describe excited-state properties of realistic materials.
- **FLEUR** is a code family for calculating ground-state as well as excited-state properties of solids within the context of DFT. A key difference lies in the treatment of all electrons on the same footing.
- **BigDFT** is an electronic structure pseudopotential code that employs Daubechies wavelets as a computational basis, designed for usage on massively parallel architectures.

ADDITIONAL RESOURCES

- Max Libraries: http://www.max-centre.eu/software/libraries
- Features and algorithms: http://www.max-centre.eu/software/features-and-algorithms
- Workflows (AiiDA): http://www.max-centre.eu/software/workflows
- Data: http://www.max-centre.eu/data

01.01.2023 - 31.12.2026
Multiscale Modelling
https://www.multixscale.eu/
1st Phase - Type 2

Coordinator: Kemijski Institut SI (PI: Matej Praprotnik)	IIT IT
Barcelona Supercomputing Centre ES	University of Ghent BE
University of Barcelona ES	SURF NL
Max Planck Society DE	HPC Now! ES
University of Stuttgart (HLRS) DE	Forschungzentrum Julich DE
University of Bergen NO	Leonardo IT
Sorbonne University F	Univerity of Groningen NL
CNR IT	Univerity of Tolouse F

MultiXscale aims at increasing performance, productivity and portability in the domain of multiscale simulation. It couples the scientific expertise of the CECAM network, represented by leading experts in multiscale simulations from different European institutions, with the technical expertise of the EESSI collaboration, and targets the computational laboratories of EuroHPC and beyond. It will shoulder much of the technical burden of developing and distributing domain-relevant applications for (pre-)exascale through application co-design for exascale technologies, and the provisioning of exascale-oriented libraries and services that nudge the community to adopt battle-tested, future-oriented, scalable workflows and portable technologies. These will allow application developers to pursue domain-relevant scientific innovation without being over-burdened by technical detail, and empower industrial and academic application users to painlessly adopt bleeding-edge technologies from the domain on whatever computational resource they may have access to.

- Service Portal: -
- Trainings: -
- Potential Collaboration with Industry (not partners): Yes (no PoC budget)

USE CASES

- Helicopter design and certification for civil transport
- Battery applications to support the sustainable energy transition
- Ultrasound for non-invasive diagnostics and biomedical applications

BUILD & DEVELOPMENT PLATFORM

- **EESSI** (European Environment for Scientific Software Installation): complete build and deployment environment for all MultiXscale applications and beyond. Provides a streaming service for scientific software, automatically distributing hardware-optimised applications and updates to all hardware types: supercomputers, servers, laptops.

Applications used and contributed to that enter into the multiscale workflows:

- **ESPResSo**: highly versatile software package for performing and analyzing scientific Molecular Dynamics many-particle simulations of "coarse-grained" bead-spring models as they are used in soft-matter research in physics, chemistry and molecular biology.
- **LAMMPS**: classical molecular dynamics code with a focus on materials modeling.
- **waLBerla**: Lattice Boltzmann based multiphysics framework

01.01.2023 - 31.12.2026
Plasma Simulations
https://plasma-pepsc.eu/
1st Phase - Type 1 (Exascale)

Coordinator: KTH Royal Institute of Technology - SE (PI: Stefano Markidis)	FORTH EL
Barcelona Supercomputing Centre ES	HZDR DE
Max Planck Society DE	IPP CZ
SIPEARL F	University of Helsinki FI
Technical University of Munich DE	University of Ljubljana SI

Plasma-PEPSC aims at reaching scientific breakthroughs in plasma science Grand Challenges through exascale computing and extreme-scale data analytics, by enabling unprecedented simulations on current pre-exascale and future exascale platforms to control plasma-material interfaces, optimize magnetically confined fusion plasmas, design next-generation plasma accelerators and predict space plasma dynamics in the Earth's magnetosphere. Plasma-PEPSC aims at achieving these goals by maximizing the parallel performance and efficiency of four European flagship plasma codes with a large user base: BIT, GENE, PIConGPU, and Vlasiator, by building on algorithmic advances as well as on programming model and library developments (MPI, accelerator and data movement APIs and runtimes, in-situ data analysis). Plasma-PEPSC ensures an integrated HPC software engineering approach for deploying, verifying, and validating extreme-scale kinetic plasma simulations that can serve as a community standard.

- Service Portal: -
- Trainings: -
- Potential Collaboration with Industry (not partners): Yes, but more likely at later stage

GRAND CHALLENGES

- Controlling Plasma-Material Interfaces - BIT1
- Optimizing Magnetically Confined Plasmas - GENE
- Enabling Next-Generation of Plasma Accelerators for Real World Applications - PIConGPU
- Predicting Near-Earth Space Dynamics - Vlasiator

CODES

- **GENE**: Plasma microturbulence based on Eulerian calculation and MPI domain decomposition.
- **BIT1:** full-size kinetic modelling of plasma edge
- **Vlasiator**: Semi-Lagrangian simulation of the Vlasov Equation and MPI domain decomposition. OpenMP/MPI hybrid parallelizations, sparse grids.
- **PIConGPU** plasma accelerators at full resolution and scale

01.01.2023 - 31.12.2026
Astrophysics / Cosmology
https://www.space-coe.eu
1st phase - Type 1 (Exascale)

Coordinator: University of Turin (PI: Andrea Mignone)	
INAF IT	University of Oslo NO
CINECA IT	E4 IT
KU Leuven BE	FORTH EL
IT4I CZ	EnginSoft IT
LMU DE	Eviden (Bull) F
CNRS F	HITS DE
Johann Wolfgang Goethe University DE	BSC ES

SPACE CoE aims to extensively re-engineer astronomy and cosmology (A&C) codes for Exascale, adopting new computational solutions and innovative programming paradigms, software solutions, and libraries, to be able to extract information from the extremely complex and voluminous data that will be produced by the forthcoming missions/telescopes (e.g. SKA, Euclid). SPACE will address the high-performance data analysis of the data torrent produced by exascale A&C simulation applications, also with machine-learning and visualization tools.

- Service Portal: -
- Trainings: -
- Potential Collaboration with Industry (not partners): NO

CODES

- **PLUTO**: an Eulerian, finite-volume, shock-capturing code based on high-order Godunov methods providing several integration algorithms. It supports both classical and relativistic flows and it is equipped also with a particle module for hybrid MHD-PIC simulations.
- **OpenGADGET**: one of the most used codes It is a N-body code that solves the gravitational and hydrodynamical equations in their Lagrangian form for a large ensemble of particles.
- **CHANGA/GASOLINE:** Tree and Smoothed Particle Magnetohydrodynamics codes that are widely used for cosmological galaxy formation simulations. The two codes have the same implementation, but CHANGA features parallelisation design with Charm++ which enables scaling up to 128K cores for highly-clustered datasets.
- **iPic3D**: fully electromagnetic massively parallel particle in cell based on the semi-implicit approach to address multiple scale problems in plasma physics.
- **RAMSES**: code to to model astrophysical systems, featuring self-gravitating, magnetised, compressible, radiative fluid flow. Based on the Adaptive Mesh Refinement (AMR) technique.
- **WhiskyTHC**: code aimed at solving the general-relativistic hydrodynamics equations on AMR grids by means of high-resolution shock capturing methods. It leverages a high-level C++ programming paradigm to achieve a high degree of efficiency.
- **FIL**: it employs a fourth-order accurate finite differencing scheme to solve the equations of general-relativistic magnetohydrodynamics. It includes a framework for realistic microphysics through the use of temperature and composition dependent Equation of State tables.
- **BHAC**: multidimensional general relativistic magnetohydrodynamics code based on the MPI-AMRVAC framework. BHAC solves the equations of ideal general relativistic magnetohydrodynamics in one, two or three dimensions on arbitrary stationary spacetimes.