ETH zürich

EXCLAIM

EXCLAIM

Extreme scale computing and data platform for cloud-resolving weather and climate modeling

https://exclaim.ethz.ch

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Federal Department of Home Affairs FDHA
Federal Office of Meteorology and Climatology MeteoSwiss

C2SM Center for Climate Systems Modeling

CSCS

MeteoSwiss

SDSC

Empa Materials Science and Technology

Extreme scale computing and data platform for cloud-resolving weather and climate modeling

EXCLAIM is an open ETH project that aims to develop an ICON-model based infrastructure that is capable of running kilometer-scale climate simulations at both regional and global scales.

Global climate change is increasingly recognized as one of the main challenges of our time, and there is considerable urgency behind mitigating anthropogenic effects and adapting to the future climate and the associated extremes (IPCC 2021). High-resolution climate models are needed to better constrain uncertainties and to improve the simulation of extreme events such as heatwaves, intense rainfalls, heavy thunderstorms, and droughts.

To develop the weather and climate modeling infrastructure of the future, EXCLAIM brings together the Swiss Federal Institute of Technology in Zürich (ETH Zürich), the Swiss National Supercomputer Centre (CSCS), the Federal Office of Meteorology and Climatology (MeteoSwiss), the Swiss Data Science Center (SDSC) and the Swiss Federal Laboratories for Materials Science and Technology (Empa). The project, which is run under the umbrella of the Center for Climate Systems Modeling (C2SM), is based on the ICON model and aims to build a system that is capable of running kilometer-scale climate simulations at both regional and global scales. ICON is the ICOsahedral Nonhydrostatic Model that was developed primarily by the Deutscher Wetterdienst (DWD) and Max Planck Institute for Meteorology (MPI-M) and is now governed through the ICON consortium that also includes the German Climate Computing Center (DKRZ), the Karlsruhe Institute of Technology (KIT), and the C2SM. Reaching this objective requires not only a two to three-orders of magnitude increase in computational throughput, but also efficient workstreams and data-flows that deal with the resulting data avalanche. In the development of this infrastructure, special attention is given to its usability and transferability.

This project requires a highly collaborative effort that brings together computer scientists with data and domain experts, using innovative coding concepts and exploiting the next generation of supercomputer architectures and the power of novel hardware systems including accelerators such as Graphics Processing Units (GPU). EXCLAIM will also very closely work together with its international partners, especially the DWD and MPI-M in Germany and the European Centre for Medium-Range Weather Forecasts (ECMWF) in the UK/Germany/Italy.

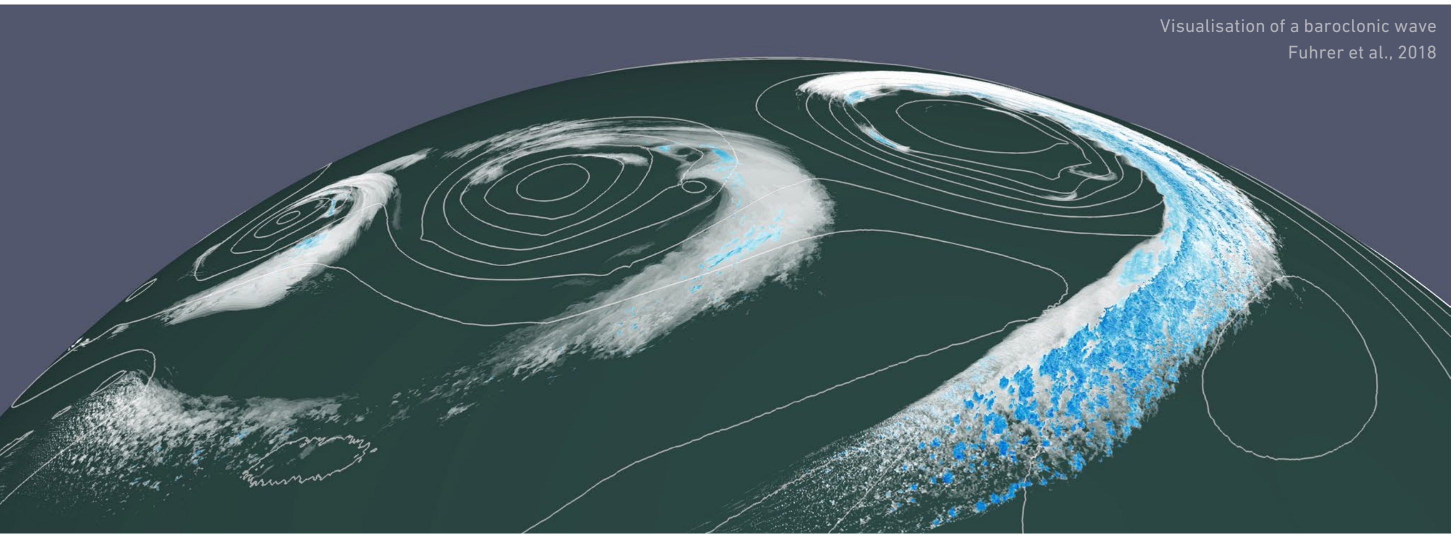

Scientific roadmap

Four core use cases define the scientific roadmap and the associated milestones. These use cases represent steps along a line of increasing complexity.

In the course of the next six years, EXCLAIM develops the corresponding performant model versions and runs them to demonstrate the ability of the platform to reach the specified targets. In the implementation of this roadmap, EXCLAIM minimizes the number of model versions and ensures that throughout the project a small set of ICON versions with increasing computational performance are available. Thorough testing ensures that these versions are functionally equivalent despite transiting from Fortran+ to a Python-based framework. Table 1 provides an overview of the four use cases defining the scientific roadmap of EXCLAIM. Each of those have been identified through an extensive user consultation process within the EXCLAIM community, thus identifying also a set of owners.

"As torrential rain, hailstorms and floods unexpectedly hit the Alpine region and northwest Europe in June and July of 2021, the importance of understanding how extreme weather events are connected to global climate change is rapidly increasing. By representing the underlying fundamental physical processes, weather and climate models are a very powerful tool to understand these interactions. EXCLAIM project aims to dramatically increase the spatial resolution of the models, thereby enhancing their accuracy in simulating weather on a global scale in a future, warmer world.", says Prof. Nicolas Gruber, EXCLAIM lead-PI.

Simulation	Setup	Resolution	Duration
Aqua Planet	Global atmosphere only, no land	10 km 1 km	2 years
Global Uncoupled	Global atmosphere and land, prescribed SSTs	25 km (reference) 3 km	5-10 years
Global Coupled	Global, ocean, sea-ice, land, atmosphere	25 km (reference) 3 km	3 decades to century
Regional Climate Europe (Scenarios CH202X)	Europe (CORDEX domain)	12 km 1 km	century

Table 1: Overview of core scientific use cases

Aqua Planet simulations

Aqua Planet simulations depict the simplest representation of the Earth's climate on a sphere, including its major large-scale atmospheric circulation (Neale and Hoskins, 2000). The use case will consist of a series of simulations with increasing resolution up to 1 km horizontal grid spacing in order to assess the response in radiative forcing, atmospheric circulation features and the water cycle to a uniform warming of the sea-surface. The configuration allows assessing two central hypotheses: convergence of the cloud radiative feedbacks at kilometer scale resolution, and the robustness of the model formulation (Stevens and Bony, 2013; Retsch et al., 2019; Rasp et al., 2018). This configuration also allows us to assess the quality of the representation of the mid-latitude storm tracks, which in current state-of-the-art climate models are significantly biased. Aqua Planet simulations are also an ideal contribution to ensure the efficacy, consistency and accuracy of the ICON development in the transition phase from Fortran to Python.

Global uncoupled simulations for clouds and convection

Kilometer-scale horizontal resolution offers promising prospects for improved simulations of convective clouds. A series of multi-year km-scale simulations with prescribed sea surface temperatures (SSTs) are envisaged. The simulations will use explicit convection and be evaluated in comparison to lower-resolution simulations using parameterized convection. To investigate the effects of changes in climate-relevant forcings and boundary conditions in km-scale simulations, a time-slice approach will be pursued (Wild et al., 2003; Mitchell et al., 2017). Thereby, the observed SSTs and sea-ice distributions are superimposed with corresponding climate-change signals derived from transient climate simulations with coupled atmosphere-ocean global climate models (GCMs) such as available from the Coupled Model Intercomparison Project 6 (CMIP6) or from the HAPPI-MIP project, representative for global warming levels of 2°C or 4°C (Mitchell et al., 2017).

Global coupled climate simulations for extremes

The final step up the complexity chain are global simulations with the fully coupled (atmosphere, ocean, sea-ice, land surface) climate model. A core scientific theme driving this scientific use case are extremes (Seneviratne et al., 2012), especially those that have a connection to the oceanic realm, such as tropical cyclones (Knutson et al., 2010), marine heatwaves (Oliver et al., 2021) and their extension over land regions. With extremes being rare by nature, the baseline simulations need to extend over a few decades. The horizontal resolution will be progressively increased from ~25 km to order of 3 kilometers later in the project.

Regional European climate simulations for scenarios

Climate scenarios at local to regional scales are an indispensable backbone for the development of effective adaptation strategies and for assessing the local and regional climatic effects of different mitigation options. To this end, a regional climate (limited area) implementation will be developed in EXCLAIM on the basis of the uncoupled global setup benefiting from the development of the ICON-22 project at MeteoSwiss. In the first phase, the focus will be on the conventional EURO-CORDEX resolution (Coppola et al., 2021) of around 12 km, increasing in the second phase the resolution to 1 km, thereby providing the full potential of very fine-scale regional climate scenarios for Alpine domains (Schär et al., 2020; Pichelli et al., 2021).

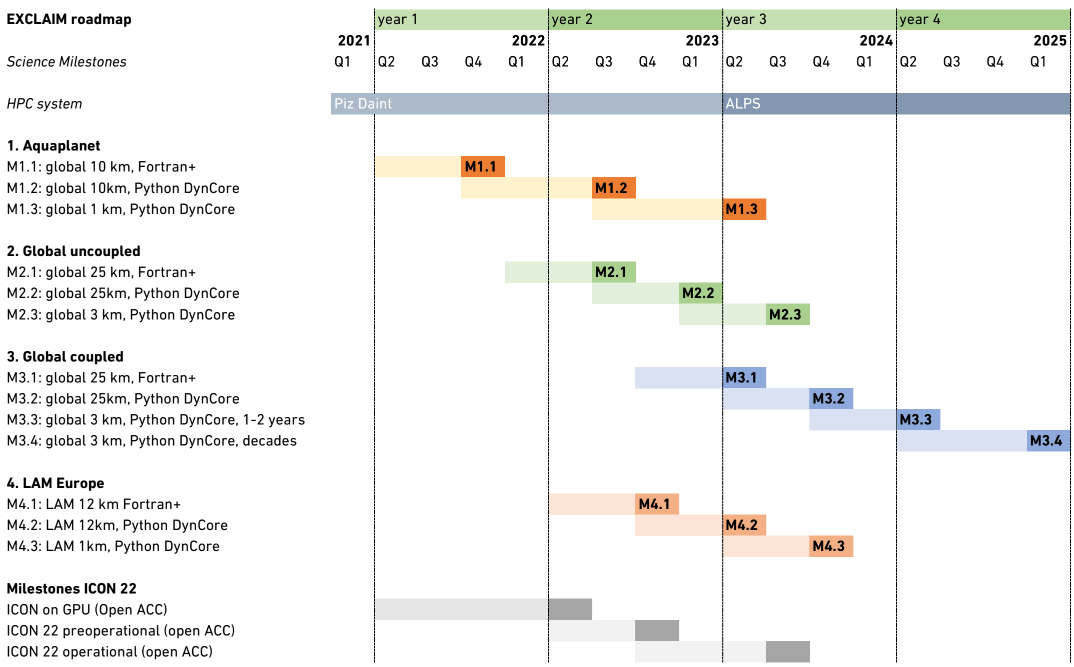

Table 2: Milestones for the implementation of the core use cases.

The baseline application software is the ICON Software for weather and climate simulations as released by the ICON Consortium. The computing and data infrastructure at the Swiss National Supercomputing Centre (CSCS) and the Finnish IT Center for Science (CSC) is being tailored to the requirements of the weather and climate models.

Model system and hardware

The core approach taken by EXCLAIM to reach the performance target is a fundamental redesign and refactoring of the model codes from the current primarily Fortran-based paradigm to one that enables performance portable, architecture agnostic user codes, as previously explored with the COSMO limited-area model (Fuhrer et al., 2014, 2018). This involves the splitting of the currently monolithic codes into a descriptive user code (based on Python) and into a backend that translates to a standard imperative language (e.g., C/C++) for specific architectures (here especially heterogeneous GPU/CPU based architectures). In addition, increases in the performance of a number of numerical schemes, new architectural concepts such as coarse-grained reconfigurable arrays, and the expected increase in memory and hardware performance will help to reach the global kilometer-scale objective within the framework of EXCLAIM.

For the refactoring, EXCLAIM follows a non-disruptive approach that enables scientific productivity during the entire project (Figure 1). This is ensured by a stepwise approach, where individual components of the model system are replaced by the Python-based framework (GT4Py) one by one, and then tested thoroughly before advancing to the next step. Continuous integration/ continuous deployment (CI/CD) methods will be used to ensure the full functioning of the built systems along the way. A well designed verification and testing system ensures scientific validity of the results.

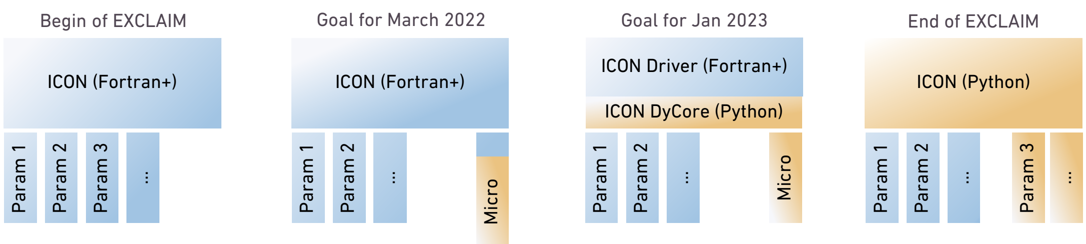

Figure 1: Planned evolution of ICON from Fortran/OpenACC (blue) to Python/GT4Py (orange). The steps are: (1) explore the use of Python/ GT4Py with one or a few parameterizations; (2) replace the dynamical core of ICON with a Python/GT4Py-based implementation; (3) to make available the whole (or most) of ICON in Python.

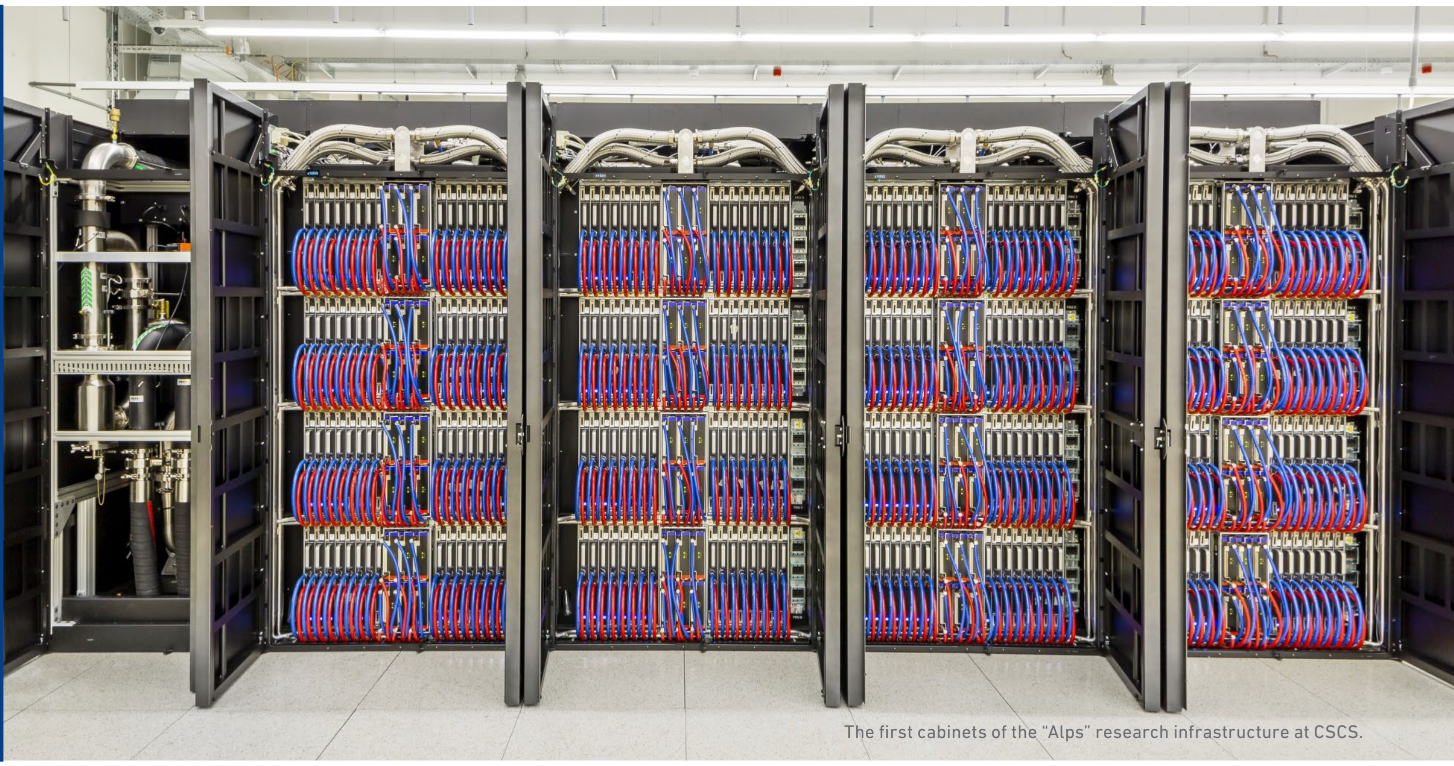

The first cabinets of the "Alps" research infrastructure at CSCS.

Current HPC Performance

The global ~3 km QUBICC version of ICON (Fortran with ACC directives) requires on CSCS supercomputer Piz Daint (using 2000 of the 5700 GPU nodes) about 200 seconds per timestep. This translates into roughly a wall clock time of about one month per simulated year. This performance needs to be accelerated by a factor of about 5 to attain a good performance for the Aqua Planet and global uncoupled simulations, but by a factor of at least 30 for the century-scale global simulations (Schulthess et al., 2019).

The "Alps" research infrastructure is expected to be the world's most powerful AI-capable supercomputer, and is now being built by CSCS, Hewlett Packard Enterprise (HPE) and NVIDIA. "Alps will use the HPE Cray EX supercomputing infrastructure based on a cloud-native software architecture to implement a software-defined research infrastructure, as well as NVIDIA's novel Grace CPU to converge AI technologies and classic supercomputing in one single, powerful data center infrastructure.", says Thomas Schulthess, computational physicist at ETH Zurich and director of CSCS.

"Alps" research infrastructure specifications

Model	HPE Cray EX
TFlops	4719
CPU Type	AMD EPYC 7742
Cores	64
Sockets	2
Interconnect Type	HPE Cray Slingshot
Nodes	1024

Source: CSCS

HPC developments

The roadmap is based on the assumption that the e-infrastructure (networks, computers, storage devices, software, etc.) of the future has a microservice architecture that is accessible through a web-capable interface. Such infrastructures will be available from mid-2023 onward both at CSCS (Alps research infrastructure) and CSC in Finland (infrastructure of the LUMI consortium, of which ETH Zurich/CSCS is a member).

- The LUMI supercomputer will have peak performance of more than 550 Petaflop/s and a high performance storage capacity of more than 100 Petabytes. It will be based on AMD GPU accelerators that can support the OpenACC-based version of the ICON software. The system will be available from mid-2022 onward.

- The performance of the Alps infrastructure will be about an order or magnitude higher than CSCS' current flagship system, Piz Daint, which it will replace in 2023. Alps will support both NVIDIA and AMD GPUs, and is designed to complement the LUMI supercomputer.

Both systems will be based on the new HPE Cray EX supercomputer product line. The initial phase of the Alps infrastructure has been deployed in fall of 2020 and is in operation. This means that the software developments for the EXCLAIM project can already be supported on this new infrastructure.

High-level implementation plan

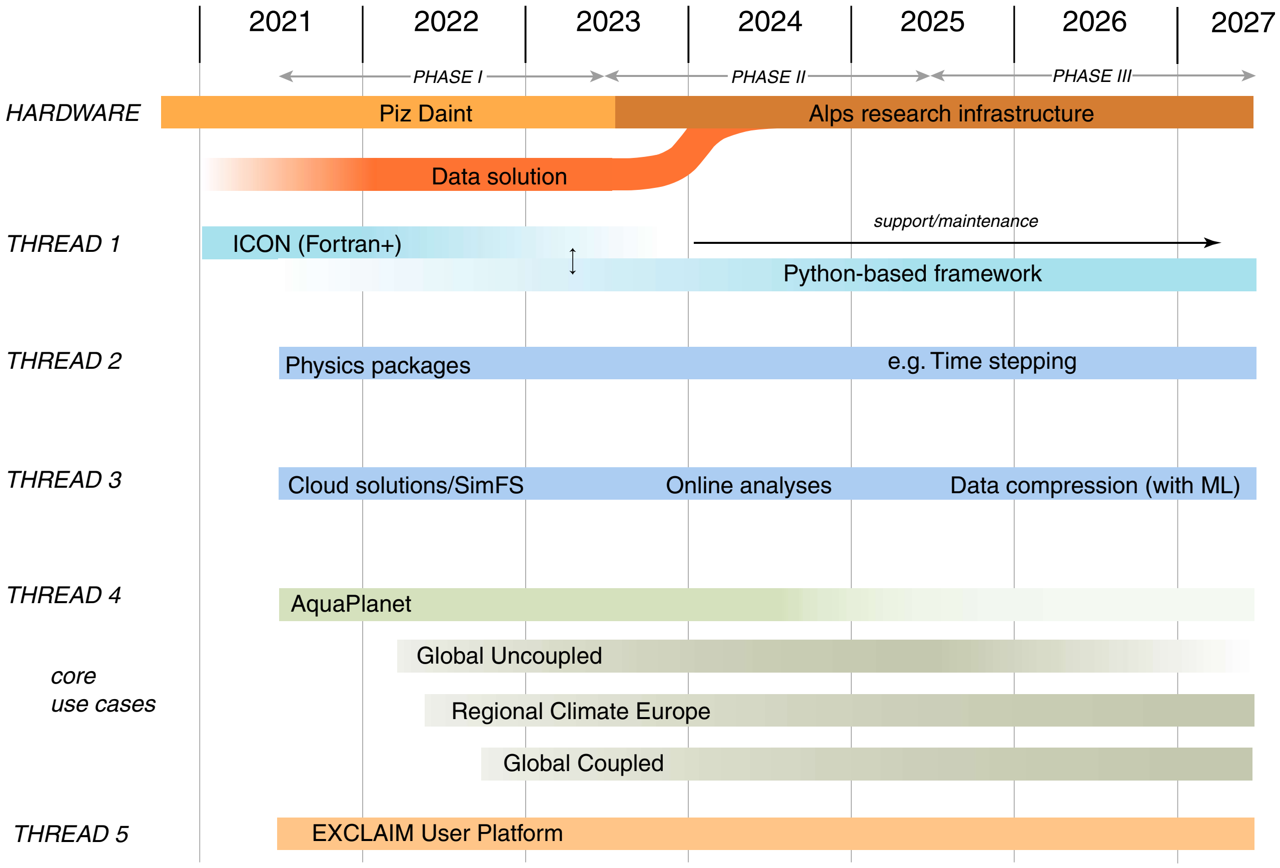

Figure 2: High-level road-map of the EXCLAIM project showing the five key development threads

We implement EXCLAIM along five threads, i.e., (i) refactoring the codes, (ii) exploiting enhancements and efficiency gains at both the computational and algorithmic levels, (iii) developing an analysis and storage framework to deal with the exabytes of data generated, (iv) developing concrete use cases, and (v) building a user support platform in close collaboration with C2SM and CSCS. These five threads will be implemented in parallel in three phases (Figure 2). Phase I will focus on the development of the dynamical core for the atmosphere and the first testing of this framework in the Aqua Planet configuration. Alongside, a first suite of parameterizations will be refactored to Python as well. Phase II is mostly concerned about the scaling of these approaches to the other use cases, with the delivery of the first scalable prototype for global application as the main milestone. Phase III then focuses on the full-scale deployment of the python-based code paradigm for the fully coupled model with the final milestone being the delivery of the global kilometer-scale modeling system that can be employed for century time-scale simulations.

For this roadmap, a number of important concurrent developments are taken into account:

- ❖ the transition of COSMO to ICON for the daily weather forecast at MeteoSwiss, to be in operation from early 2023 (Project ICON 22),
- ❖ the replacement of the Piz Daint HPC infrastructure at CSCS by the ALPS system in early 2023,
- ❖ the development of ICON-Consolidated and ICON SEAMLESS by the ICON consortium, especially in the context of the pre WarmWorld and the WarmWorld activities in Germany, and
- ❖ the European initiative called Destination Earth (DestinE) that will create a digital climate and extremes twin of the Earth.

THREAD 1: Software framework

The software framework thread focuses on the development of the Python framework for our model system ICON. But it also requires attention to the software systems that permit the users to execute scientific simulations and workflows, and to the platform interfaces used by the scientist to deploy those simulations and manage the associated data. Regarding the former, a number of key steps need to be undertaken: (i) Modularization of the ICON Software into system, model and interface components, (ii) Implementation of CI/CD pipelines for the ICON Software to run on different sites and systems, (iii) Development of the Python Framework and versions for the development of the model components, and (iv) Migration of interface and system components to Python to obtain the final EXCLAIM ICON implementation. For the latter, the implementation of CI/CD pipelines to allow for testing and validation is crucial. In addition, special attention need to be given to the data interfaces to store, retrieve, compress, analyze and modify data to and from simulations, and to do data curation (in collaboration with SDSC and SPCL). The focus of Phase I (mid 2021- mid 2023) will be on the development of the above enabling software frameworks and the initial deployment of a Python based implementation of the atmospheric dynamical core.

THREAD 2: Algorithms and parameterizations

The goal of the second thread is to spearhead a strategy towards the future use of algorithms and parameterizations in climate and weather models. Regarding the latter, we will build and extend on ongoing work that uses OpenACC directives to port some of the parameterizations to GPU. EXCLAIM will coordinate and spearhead some of these efforts, e.g., for the a two-moment microphysics scheme (Seifert and Beheng, 2006) and the consideration of large-eddy turbulence (LES) models (Panosetti et al., 2019), which are available in conventional CPU implementations of ICON. As LES models entail horizontal coupling between neighboring grid dells, their implementation into the EXCLAIM framework has to be considered a part of the dynamical core. In addition, we will make an investment into the development of new and improved algorithms, especially with regard to time stepping.

THREAD 3: Data flow and data science

The goal of thread 3 is to ensure that the user can establish an efficient workflow from the preparation of the simulation all the way to its analysis. This requires novel methods in data handling, since the currently deployed data flow models do not scale to the types of problems that EXCLAIM is aiming to tackle. Additional challenges arise e.g., for the use of limited area model configurations at high resolution, since they require the provision of lateral boundary conditions at high frequency as well. During the first phase, thread 3 will focus on the establishment of a new set of data services through converged cloud-HPC-based data services. In addition, the SimFS environment will continue to be developed, bringing it to the roll out phase. Finally, work will be initiated in the areas of machine learning tools for data compression, and online analyses.

THREAD 4: Use cases

The four use cases identified at the beginning of this document form the scientific backbone of the project. These use cases will be implemented following the science roadmap shown in Table 2. EXCLAIM will thereby ensure that these configurations are well tested and scientifically validated, while the scientists will take the lead in executing the simulations and analyzing the results.

THREAD 5: EXCLAIM User Platform

EXCLAIM will support CSCS and C2SM in building jointly an EXCLAIM user platform that will support domain scientists to fully benefit from the developments undertaken by EXCLAIM and other HPC oriented projects. This includes e.g., support in porting and optimizing codes, providing verification and testing suites, managing and analyzing data streams, and the containerization of codes.

List of abbreviations

ACC	Accelerator
AMD	Advanced Micro Devices
C2SM	Center for Climate Systems Modeling
CI/CD	Continuous Integration / Continuous Deployment
CMIP	Coupled Model Intercomparison Project
COSMO	Consortium for Small-scale Modeling
CPU	Central Processing Unit
CSC	IT Center for Science Ltd.
CSCS	Swiss National Supercomputing Centre
DKRZ	German Climate Computing Center
DWD	Deutscher Wetterdienst
ECMWF	European Centre for Medium-Range Weather Forecasts
EMPA	Swiss Federal Laboratories for Materials Science and Technology
ETH	Swiss Federal Institute of Technology in Zurich
EURO-CORDEX	Coordinated Downscaling Experiment - European Domain
EXCLAIM	Extreme scale computing and data platform for cloud-resolving weather and climate modeling
GPU	Graphics Processing Unit
GCM	Global Climate Model
GT4Py	GridTools for Python
HAPPI-MIP	Half a degree Additional warming, Prognosis and Projected Impacts - Model Intercomparison Project
HPC	High-performance computing
HPE	Hewlett Packard Enterprise
ICON	Icosahedral Nonhydrostatic Weather and Climate Model
KIT	Karlsruhe Institute of Technology
LES	Large Eddy Simulation
LUMI	Large Unified Modern Infrastructure
MeteoSwiss	Swiss Federal Office of Meteorology and Climatology
MPI-M	Max-Planck Institute for Meteorology
NWP	Numerical Weather Prediction
QUBICC	The Quasi-biennial Oscillation in a Changing Climate
SDSC	Swiss Data Science Center
SimFS	A Simulation Data Virtualizing File System Interface
SPCL	Scalable Parallel Computing Lab
SST	Sea Surface Temperature

References

Bockelmann, H., Klocke, D., Potthast, R., and Vogel, B. (2021). **Development of a Consolidated ICON.** Technical Report, DKRZ, MPI, DWD, KIT.

Coppola, E., Nogherotto, R., Ciarlo, J. M., Giorgi, F., van Meijgaard, E., Kadygrov, N., Iles, C., Corre, et al. (2021). **Assessment of the European Climate Projections as Simulated by the Large EURO-CORDEX Regional and Global Climate Model Ensemble.** *Journal of Geophysical Research: Atmospheres*, 126/4, Page e2019JD032356, 10.1029/2019JD032356.

Fuhrer, O., Osuna, C., Lapillonne, X., Gysi, T., Cumming, B., Bianco, M., Arteaga, A., and Schulthess, T. (2014). **Towards a performance portable, architecture agnostic implementation strategy for weather and climate models.** *Supercomput. Front. Innov.*, 1, Pages 45–62, 10.14529/jsfi140103.

Fuhrer, O., Chadha, T., Hoefler, T., Kwasniewski, G., Lapillonne, X., Leutwyler, D., Lüthi, D., Osuna, C., Schär, C., Schulthess, T. C., and Vogt, H. (2018). **Near-global climate simulation at 1 km resolution: establishing a performance baseline on 4888 GPUs with COSMO 5.0.** *Geoscientific Model Development*, 11/4, Pages 1665– 1681, 10.5194/gmd-11-1665-2018.

IPCC (2021). **Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change**. Cambridge University Press. https://www.ipcc.ch/report/ar6/wg1/

Knutson, T. R., McBride, J. L., Chan, J., Emanuel, K., Holland, G., Landsea, C., Held, I., Kossin, J. P., et al. (2010). **Tropical cyclones and climate change.** *Nature Geoscience*, 3/3, Pages 157–163, 10.1038/ngeo779.

Mitchell, D., AchutaRao, K., Allen, M., Bethke, I., Beyerle, U., Ciavarella, A., Forster, P. M., Fuglestvedt, J., et al. (2017). **Half a degree additional warming, prognosis and projected impacts (HAPPI): background and experimental design.** *Geoscientific Model Development*, 10/2, Pages 571–583, 10.5194/gmd-10-571-2017.

Neale, R. B. and Hoskins, B. J. (2000). **A standard test for AGCMs including their physical parametrizations: I: the proposal.** *Atmos. Sci. Lett.*, 1, Pages 101–107, 10.1006/asle.2000.0022.

Oliver, E. C. J., Benthuysen, J. A., Darmaraki, S., Donat, M. G., Hobday, A. J., Holbrook, N. J., Schlegel, R. W., and Sen Gupta, A. (2021). **Marine Heatwaves.** *Annual Review of Marine Science*, 13/1, Pages 313–342, 10.1146/annurev-marine-032720-095144.

Panosetti, D., Schlemmer, L., and Schär, C. (2019). **Bulk and structural convergence at convection-resolving scales in real-case simulations of summertime moist convection over land.** *Quarterly Journal of the Royal Meteorological Society*, 145/721, Pages 1427–1443, 10.1002/qj.3502.

Pichelli, E., Coppola, E., Sobolowski, S., Ban, N., Giorgi, F., Stocchi, P., Alias, A., Belušic, D., et al. (2021). **The first multi-model ensemble of regional climate simulations at kilometer-scale resolution part 2: historical and future simulations of precipitation**. *Climate Dynamics*, 56/11, Pages 3581–3602, 10.1007/s00382-021-05657-4.

Potthast, R., Früh, B., Korn, P., and Müller, W. (2021). **Entwicklung eines Erdsystemmodells ICON-Seamless auf Grundlage der Komponenten ICON-NWP und ICON-O – eine Roadmap.**

Rasp, S., Pritchard, M. S., and Gentine, P. (2018). **Deep learning to represent subgrid processes in climate models.** *Proceedings of the National Academy of Sciences*, 115/39, Pages 9684–9689, 10.1073/pnas.1810286115.

Retsch, M. H., Mauritsen, T., and Hohenegger, C. (2019). **Climate change feedbacks in aquaplanet experiments with explicit and parameterized convection for horizontal resolutions of 2,525 up to 5 km.** *Journal of Advances in Modeling Earth Systems*, 11, Page 2070– 2088, 10.1029/2019MS001677.

Schär, C., Fuhrer, O., Arteaga, A., Ban, N., Charpilloz, C., Girolamo, S. D., Hentgen, L., Hoefler, T., Lapillonne, X., et al. (2020). **Kilometer-scale climate models: Prospects and challenges**. *Bulletin of the American Meteorological Society*, 101, 10.1175/BAMS-D-18-0167.1.

Schulthess, T. C., Bauer, P., Wedi, N., Fuhrer, O., Hoefler, T., and Schär, C. (2019). **Reflecting on the Goal and Baseline for Exascale Computing: A Roadmap Based on Weather and Climate Simulations**. *Computing in Science Engineering*, 21/1, Pages 30–41, 10.1109/MCSE.2018.2888788.

Seifert, A. and Beheng, K. D. (2006). **A two-moment cloud microphysics parameterization for mixed-phase clouds. part 1: Model description.** *Meteorology and Atmospheric Physics*, 92/1, Pages 45–66, 10.1007/s00703-005-0112-4.

Seneviratne, S. I., Nicholls, N., Easterling, D., Goodess, C. M., Kanae, S., Kossin, J., Luo, Y., Marengo, J., et al. (2012). **Changes in Climate Extremes and their Impacts on the Natural Physical Environment.** In C. B. Field, Q. Dahe, T. F. Stocker, and V. Barros (Ed.), *Managing the Risks of Extreme Events and Disasters to Advance Climate Change Adaptation: Special Report of the Intergovernmental Panel on Climate Change*, Pages 109–230. Cambridge University Press, Cambridge, https://www.ipcc.ch/site/assets/uploads/2018/03/SREX-Chap3_FINAL-1.pdf.

Stevens, B. and Bony, S. (2013). **What are climate models missing?.** *Science*, 340/6136, Pages 1053–1054, 10.1126/science.1237554.

Wild, M., Calanca, P., Scherrer, S. C., and Ohmura, A. (2003). **Effects of polar ice sheets on global sea level in high-resolution greenhouse scenarios.** *Journal of Geophysical Research: Atmospheres*, 108/D5, 10.1029/2002JD002451.

ETH Zürich
Center for Climate Systems Modeling (C2SM)
Universitätstr. 16
8092 Zürich

https://c2sm.ethz.ch/

Publisher: EXCLAIM Project, https://exclaim.ethz.ch/
Editor: Nicolas Gruber
Contributors: Christof Appenzeller, Marco Arpagaus, Mauro Bianco, Dominik Brunner, Anurag Dipankar, Daniela Domeisen, Erich Fischer, Nicolas Gruber, Sven Kotlarski, Xavier Lapillonne, David Leutwyler, Mark Liniger , Ulrike Lohmann, Fernando Perez-Cruz, Carlos Osuna, Will Sawyer, Christoph Schär, Sebastian Schemm, Christina Schnadt Poberaj, Thomas Schulthess, Sonia I. Seneviratne, Michael Sprenger, Stefano Ubbiali, André Walser, Heini Wernli, Martin Wild, Christian Zeman
Graphic design: Tamara Bandikova
Photos: NOAA/NASA, Tamara Bandikova, CSCS, 123RF