A Roadmap to the Implementation of 1km Earth System Model Ensembles

Deliverable D1.2

This project has received funding from the European Union's
Horizon 2020 Research and Innovation Programme
under Grant Agreement No 675191

About this document

Work package in charge: WP1 Governance, Engagement & Long-term Sustainability
Actual delivery date for this deliverable: 1 August 2019
Dissemination level: The general public (PU)

Lead authors
Deutsches Klimarechenzentrum GmbH (DKRZ): Joachim Biercamp
European Centre for Medium-Range Weather Forecasts (ECMWF): Peter Bauer, Peter Düben
The University of Reading (UREAD): Bryan Lawrence

Other contributing authors:
Centre National de Recherche Scientifique - Institut Pierre-Simon Laplace (CNRS-IPSL): Sylvie Joussaume
Max-Planck-Institut für Meteorologie (MPI-M): Jochem Marotzke
Deutsches Klimarechenzentrum GmbH (DKRZ): Philipp Neumann
BULL/ATOS (Bull): Erwan Raffin

Contacts: esiwace@dkrz.de
Visit us on: www.esiwace.eu
Follow us on Twitter: @esiwace

Table of contents

1. Abstract /publishable summary 4
2. Conclusion & Results 4
3. Project objectives 4
4. Detailed report on the deliverable 5
5. References *(Bibliography)* 5
6. Dissemination and uptake 5
    6.1 Uptake by the targeted audience 5
    6.2 This is how we are going to ensure the uptake of the deliverables by the targeted audience5
7. The delivery is delayed: ☒ Yes ☐ No 5
8. Changes made and/or difficulties encountered, if any 6

1. Abstract /publishable summary

The community of Earth system modelling has a strong interest to reach a horizontal resolution of about 1km for ensemble simulations with global coupled weather and climate models. Within the ESiWACE project, we have demonstrated that 1 km simulations with global atmosphere models are technically feasible (see ESiWACE Deliverable D2.1). However, we are still far from being able to run a 1 km coupled ensemble in an operational mode even if some of the fastest supercomputers that are currently available are used.

This deliverable presents a concise roadmap that lays out the technical work, and the financial and organisational framework that is needed to run ensembles of Earth system models (ESM) at 1km resolution with enough throughput to perform operational weather forecasts and climate projections.

2. Conclusion & Results

The implementation of the necessary changes in the weather and climate models themselves and in the associated handling of large data volumes requires a significant effort of several hundred person years per model and therefore a significant investment. It is therefore recommended to share and coordinate as much of the required work at European level as possible and to limit the number of individual exascale models (by maintaining the scientifically required diversity of high-level model formulations).
Open key questions are: how to fund the critical mass that is necessary for radically new developments, how to achieve and sustain economy-of-scales developments serving the community, and how to exploit funding opportunities at European and national level without losing sight of the big goals? These questions can only be answered at international science policy level but require immediate attention.

3. Project objectives

This deliverable contributes directly and indirectly to the achievement of all the macro-objectives and specific goals indicated in section 1.1 of the Description of the Action:

Macro-objectives	**Contribution of this deliverable?**
Improve the efficiency and productivity of numerical weather and climate simulation on high-performance computing platforms	Yes
Support the end-to-end workflow of global Earth system modelling for weather and climate simulation in high performance computing environments	Yes
The European weather and climate science community will drive the governance structure that defines the services to be provided by ESiWACE	Yes
Foster the interaction between industry and the weather and climate community on the exploitation of high-end computing systems, application codes and services.	Yes
Increase competitiveness and growth of the European HPC industry	Yes

Specific goals in the workplan	**Contribution of this deliverable?**
Provide **services** to the user community that will impact beyond the lifetime of the project.	Yes
Improve **scalability** and shorten the time-to-solution for climate and	Yes

operational weather forecasts at increased resolution and complexity to be run on future extreme-scale HPC systems.	
Foster **usability** of the available tools, software, computing and data handling infrastructures.	No
Pursue **exploitability** of climate and weather model results.	Yes
Establish governance of common software management to avoid unnecessary and redundant development and to deliver the best available solutions to the user community.	Yes
Provide **open access** to research results and **open source** software at international level.	No
Exploit **synergies** with other relevant activities and projects and also with the global weather and climate community	Yes

4. Detailed report on the deliverable

The full detailed report can be found in Annex 1. The deliverable is a document in form of a white paper laying out lays out the technical work, and the financial and organizational framework that is needed to run ensembles of Earth system models (ESM) at 1km resolution with enough throughput to perform operational weather forecasts and climate projections.

5. References *(Bibliography)*

See the references in the document.

6. Dissemination and uptake

6.1 Uptake by the targeted audience

As indicated in the Description of the Action, the audience for this deliverable is*:*

x	The general public (PU)
	The project partners, including the Commission services (PP)
	A group specified by the consortium, including the Commission services (RE)
	This reports is confidential, only for members of the consortium, including the Commission services (CO)

6.2 This is how we are going to ensure the uptake of the deliverables by the targeted audience

The content of this deliverable will be disseminated through workshops, newsletters and personal communications.

7. The delivery is delayed: ☒ Yes ☐ No

The delivery was originally planned for June 2019. Due to limited availability of the authors it was delayed by one month.

8. Changes made and/or difficulties encountered, if any

n.a.

A Roadmap to the Implementation of 1km Earth System Model Ensembles

Main authors: Joachim Biercamp, Peter Bauer, Peter Düben, Bryan Lawrence

Contributions by Sylvie Joussaume, Jochem Marotzke, Philipp Neumann, Erwan Raffin

1 Introduction

The ability to make reliable predictions of weather and climate has a significant social value. Essential tools for these predictions are numerical models that represent the most important components of the global Earth system including atmosphere, ocean, sea ice, and land surface. These models have seen an astonishing increase in complexity and accuracy over the past decades (Bauer et al. 2015), but the quality of predictions is still limited by errors in model formulation and discretisation method, inaccurate initial and boundary conditions, and, finally, the aspect we are focussing on in this paper: finite spatial resolution.

Deep convection is an important feature of the atmosphere that is dominating the vertical transport of energy in the tropics and therefore driving the global circulation. Simulations with a horizontal grid spacing used in operational models (10km or more) require the parametrisation of deep convection. Simulations with a couple of kilometres resolution, the so-called "grey" zone, make the bulk formulation of the parametrisation schemes, that assume that there are always several convective cells within each grid-cell, invalid. This is a challenge for weather and climate modellers. However, there is evidence that around 1km resolution will allow the explicit resolution of deep convection making convective parametrisation schemes obsolete with the potential for a significant improvement in model fidelity. Further, 1km resolution would bring additional benefits regarding the representation of topographic gravity waves and surface drag, and the ability to assimilate satellite data at the native resolution (e.g. Stevens et al., 2019a, 2019b). In the ocean, the size of the eddies and the width of the boundary currents are governed by the Rossby radius which in the Arctic is about 5km, and in northern shelf regions down to the order of 1km (Nurser and Bacon, 2014) . Simulations with O(1km) resolution in the ocean would allow for the explicit representation of meso-scale eddies. These are essential for realistic simulations of sea-surface variability globally, of eddy transports governing the ocean uptake of heat and carbon in the southern oceans, and of deep-water formation that is crucially influenced by eddy-induced restratification. We therefore need O(1km) resolution also in the ocean at least in polar regions.

The community of Earth system modelling has therefore a strong interest to reach a horizontal resolution of O(1km) for ensemble simulations with global coupled models as soon as possible.

Within the ESiWACE project, we have demonstrated that 1 km simulations with global atmosphere models are technically feasible (see ESiWACE Deliverable D2.1). However, we are still far from being able to run a 1 km coupled ensemble in an operational mode even if some of the fastest supercomputers that are currently available are used (Neumann et al., 2019, Fuhrer et al., 2018, Schulthess et al. 2019, Dueben et al. 2019).

In this paper, we present a concise roadmap that lays out the technical work, and the financial and organisational framework that is needed to run ensembles of Earth system models (ESM) at 1km resolution with enough throughput to perform operational weather forecasts and climate projections as soon as possible. In this context, we will also briefly discuss the recent developments for two new projects that are driven by the European Earth System modelling community: ESiWACE2[1] and ExtremeEarth[2]. The ESiWACE2 Centre of Excellence aims to support developments of current weather and climate models to enable their execution on exascale supercomputers in the future. ESiWACE2 was selected for funding and we have already started to go the first "miles" on our roadmap. ExtremeEarth, which is looking at a longer time-scale as it addresses the science-technology-impact challenge more holistically, has not received external funding yet, but the community is considering several pathways to sustain the efforts that should ultimately allow to achieve our O(1km) ensemble prediction goal, and to develop the entire technology infrastructure for translating such predictions to socio-economic value in all related impact sectors.

2 Background

Weather and climate modelling has traditionally required access to the largest possible computing resources. What we can simulate, and thus the scientific questions we can address, has always been limited by computing resources. Although available computing power increased steadily during the last decades it became increasingly difficult to make efficient use of the available supercomputers since complexity of both the computers (and thus the way of programming) and the models is increasing rapidly. This has already been addressed in the infrastructure strategy of ENES (Mitchel et al., 2012, Joussaume et al., 2017).

Unfortunately, the acceleration of nominal peak compute performance (per unit money and energy needed to run the system) between generations of computer architecture has slowed down and might even stop in the near future. Future generations of supercomputers will rely on more heterogeneous hardware architectures (including CPUs, GPGPUs, TPUs, and FPGAs) and programming methods. This will make it necessary to adapt existing weather and climate models to the changing computing landscape. A step that may require fundamental changes or even entire model revisions. Lawrence et al. (2018) present a comprehensive analysis of the gaps and challenges on this path. The presence of heterogeneous hardware may also allow to develop purpose-built HPC architectures for ESM (Carman et al., 2017) and to co-design software and hardware to achieve increased performance.

Recent publications suggest that the performance increase necessary for the next generation of high-resolution weather and climate models would need to lie in the range of 100 to 1000 to allow for global atmosphere simulations at O(1km) resolution with sufficient throughput for operational weather and climate predictions. Fuhrer et al. (2018) conclude that the global atmosphere model COSMO, which is used for weather forecast in several European countries, would need to run 100 times faster to achieve a throughput of 3-5 simulated years per day. Schulthess et al. (2019),

---

[1] https://www.esiwace.eu/new_overview/new_esiwace2_project_setting

[2] https://extremeearth.eu/

conclude that both COSMO and IFS (the atmosphere model of ECMWF) execute about 100-250 times too slow for operational throughput rates at a horizontal resolution of 1 km, even when executed on the fastest supercomputer in Europe (Piz Daint). Carman et al. (2017), state that ESM will need over 1000 times the computational power of today.).These results are backed up by scalability studies with the ESiWACE demonstrators. (Deliverable 2.1). For example, if we could scale the ICON demonstrator to run on about 10 Million compute-cores (which means completely filling the largest super-computer available today) we could only achieve a throughput of about 0.05 simulated years per day. Even if we had still bigger core counts available, the model would not be able to run faster because its so-called strong scaling limit is reached.

All these estimates are still underestimating the real requirements for ensembles of ESM at O(1km) resolution since they do not consider ensemble simulations and model testing (increase in required computational power depending on ensemble size x50-100), coupling of multiple components, like atmosphere, ocean, wave, land surface models (ca. x2, depending on number of components and processes), and model IO (ca. x2).

It needs to be emphasized, that we do not give precise numbers for computational short falls here. The exact number is very much dependent on the use case. For example for weather forecast, one might need shorter simulation times then for climate studies. The complexity of models, the number of components of a coupled system, the output frequency and the number of stored variables, the ensemble size, the required capability for downstream processing etc. will vary between use cases. Moreover, performance increase can be achieved by i) increasing capacity, i.e. adding more computational cores, which can be used to run additional ensemble members or ii) capability, i.e. improving the speed of computation, leading to shorter turn-around times for individual simulations.

In any case, we argue that we need several orders of magnitude in overall performance increase to enable useful science or services using cloud and eddy resolving Earth system models.

3 The roadmap

When preparing for the future of ESM in the view of the changing HPC landscape, we need to discriminate between two time horizons. The next generation of supercomputers (“pre-exascale systems”) will be available around 2021 and will presumably run until about 2025. These systems will still be based on conventional technology (including GPGPU) but – since the performance yield we can achieve with ESM is mainly limited by memory bandwidth - will most probably deliver an even smaller percentage of their nominal peak performance than today’s systems (around 5%) when used with ESM. In the long run, i.e. exascale and beyond, we will likely see new architectures and new programming paradigms that may be very different from the state-of-the-art. Our community will need to develop a new code basis for our models to be able to make efficient use of these machines. Furthermore, the amount of data produced by simulations will increase dramatically and mandates new methods and strategies for data storage, management and exploitation.

Today’s generation of models is very hard to port and only a very small number of models is capable of running on non-standard CPU hardware, and it can be argued that the pre-exascale period will be our only chance to develop a new generation of models that is sufficiently flexible and performance portable to adapt quickly to the yet unforeseeable developments in the exascale supercomputing era. Unfortunately, the pre-exascale period is very short when compared to typical model development cycles so that code adaptation may still rely on existing programming approaches.

3.1 Technical steps to prepare for the exascale

The new generation of models will need to be able to:

1. make efficient use of millions of processing units in parallel;
2. run efficiently on very different hardware (CPUs, GPGPUs, FPGAs, machine learning accelerators such as TPUs, embedded system/RISC hardware…) while keeping a single realisation of the source code which is easy to read by domain scientists;
3. allow for energy aware executions that are maximising simulated years per Joule or throughput, depending on the requirements of runs;
4. be fault resilient in a computing environment that can suffer from hardware faults or soft-errors;
5. minimize model IO to disc/tape via trade-offs between re-computation and storage of model fields, online evaluation of model diagnostics, and the capability to perform output at different levels of resolution;

The following sub-sections will outline how these requirements can be met.

3.1.1 A new generation of models

This section will list developments that can help to optimise weather and climate models for exascale supercomputers and to address the model requirements that are listed above. We also provide estimates for performance gains that can be expected from each development and an approximate level of person years to achieve this. However, these estimates can only be very rough and should not be taken literally.

Refactoring, using domain specific languages and new programming models
Addressing requirements 2 and 4; Effort: O(10) person years per model; Performance gain O(x2)

Domain specific languages (DSLs) such as GridTools (formerly known as Stella; Gysi et al. 2015) and PSyclone (https://github.com/stfc/PSyclone) or programming models such as Kokkos (https://github.com/kokkos/kokkos) can enable performance portability. The main idea is to write the model based on library functions, which render hand-written specific optimization specific to a hardware (by e.g. optimising stencils or loop-orders) unnecessary, thus making the code that is visible to domain scientists performance portable and independent of specific hardware platforms. While the model is developed in the “front-end” that is kept as simple as possible, technical details, such as the use of directives to run models on GPGPUs, will be “hidden” in the code level of the DSL that has a “back-end” which is used to optimise the DSL for use on different hardware. If the model is ported to new hardware, only the back-end needs to be adjusted. This configuration allows a separation of concerns between domain scientist and computer scientists. The refactoring of the code that is required when introducing the DSLs typically allows for improvements of model efficiency.

We would like to emphasize that the estimate of O(10) person years per model is for technical code refactoring only, not code optimization or adaptation of back-ends to emerging hardware. This effort is much larger and included in the following section.

Performance modelling, data-flow and co-design
Addressing requirement 2; Effort: O(100) person years per model; Performance gain O(x6)

To develop software and hardware simultaneously to achieve the optimal solution is known as co-design. Today’s weather and climate centres already consider the requirements of their models (for example regarding the memory bandwidth) when buying supercomputers, and hardware needs are considered in model developments (for example regarding memory use or scalability). However, the development of customised hardware for the specific needs of ESM is still a long way in the future due to the complexity of models and the use of many different algorithms within the models. To change this would require, firstly, developing a detailed performance model for weather and climate models that would allow reliable performance estimates for simulations with complex models on different hardware with no need to actually run the model. Secondly, data flow would need to be minimized within model simulations to reduce data movement. This could be realised via an automated minimisation of Direct Acyclic Graphs (DAGs)[3]. Finally, it would require the ability to project knowledge about performance and data-flow onto existing hardware components in an optimal way. If the co-design approach is taken to an extreme, it may be possible to configure field programmable gate arrays (FPGAs) or to even develop application-specific integrated circuits (ASICs) for each model component and to calculate a model time step by streaming the data through these chips. However, each of those steps would be very difficult and will require significant staff commitments, and the development of automated tools. To find the optimal solution for data flow in complex programmes is one of the seven millennium price problems stated by the Clay Mathematics Institute (a so-called NP-hard problem).

Algorithmic developments

Addressing requirements 1,4 and 5; Effort: O(20) person years per model; Performance gain O(x2)

The scalability of state-of-the-art weather and climate models has improved significantly over the past decade and some of the models are already able to scale well to the size of today's supercomputers if resolution is sufficiently high (Fuhrer et al, 2018, Neumann et al, 2019, Dueben et al. 2019). However, there are still bottlenecks that remain such as the sea-ice model or the computations of barotropic modes for ocean models (Koldunov et al., 2019) that need to be addressed. There is potential that semi-implicit or implicit time-stepping schemes can allow for the use of larger time steps and therefore deliver performance improvements with respect to explicit schemes but these schemes require the use of solvers that need global communication and further improvements via pre-conditioning and multi-grid methods. To increase arithmetic intensity and thus improve time-to-solution data re-use might be improved. The concurrent integration of several model components to improve scalability as well as the development of fault resilient models (Dueben and Dawson 2017), that identify and mitigate soft faults and allow dynamic node allocation, will require further changes to the model formulation. Finally, parallel-in-time integration schemes may allow improving efficiency further (Ruprecht, 2018).

Workflow management, dynamic allocation of compute nodes and test suites

Addressing requirements 3 and 4; Effort: O(10) person years per model; Performance gain O(x1)

Since single-model simulations will become more expensive (in terms of computing time, energy demand and number of processing units) it will be a requirement to allow for the re-allocation of idle compute nodes to large model simulations that suffer from hardware faults to avoid model crashes. It will also be necessary to synchronise simulations in terms of model IO to avoid overflows of buffers. Furthermore, the workflow management needs to be aware of requirements for scalability, memory use, model output and time-to-solution for all simulations to optimise energy consumption and over-all throughput. It needs to be possible to test model changes in simplified model settings (such as a single-column or small-planet versions of the full model) to reduce the demand for model testing at O(1km) global resolution as much as possible.

---

[3] https://en.wikipedia.org/wiki/Directed_acyclic_graph

Machine learning

Addressing requirements 1 and 2; Effort: O(20) person years per model; Performance gain O(x2)

It is yet to be seen how the recent boom of machine learning will impact on ESM. Potential applications range from improvements of parametrisation schemes (Voosen, 2018) to to learning the equations of motion (Dueben and Bauer, 2018). However, in particular deep learning may help to improve model efficiency via the emulation of existing model components, in particular emulating parametrisation schemes such as the radiation scheme (Chevallier et al., 2000). Once the neural network emulators are of sufficient quality to replace conventional parametrisation schemes, they can use efficient deep learning libraries such as Tensorflow and be easily ported to different hardware.

Mixed precision

Addressing requirements 1,2 and 3; Effort: O(20) person years per model; Performance gain O(x3)

It was shown recently, that it is sufficient to use single precision instead of double precision for many variables in atmosphere models (Dueben and Palmer, 2014), at least for weather applications. This generates a reduction of model runtime by approximately 40% and also reduces memory requirements by a factor of two to allow for the execution on a smaller number of compute nodes. Additional savings can be expected if half precision is used to speed-up expensive model kernels such as the Legendre transformation of spectral atmosphere models (Hatfield et al., 2019). This would allow to make use of deep learning accelerators that are typically optimised for high performance at low numerical precision.

If all of the developments that are listed above are realised with maximal success, this would yield an estimated performance increase of 2x2x1x6x2x3=x144 and make simulations with single ESM with the required throughput realistic. If this performance increase could be combined with an additional efficiency increase of computing hardware over time, in particular through higher core-counts by keeping the (electrical) power envelope on the same level, this would allow for operational simulations of ensembles of weather and climate simulations on exascale supercomputers. The implementation of the changes listed above requires a significant effort of several hundred person years ***per model*** *(the persons in question being top experts in computer science and computational science) and therefore a significant investment. It is therefore recommended to share and coordinate as much of the required work at European level as possible and to limit the number of individual exascale models (by maintaining the scientifically required diversity of high-level model formulations). Additionally, several of the proposed measures, like the DSL approach, will need significant and sustained maintenance, which calls for the active involvement of technology providers.*

However, this investment would be definitely justified given the cost for hardware and electricity of exascale supercomputers and the economic value of reliable weather and climate predictions.

3.1.2 New ways of dealing with data

Increasing the resolution of the models and the size of ensembles, naturally leads to also to an increase of the data volumes, which in ESM already today, belong to the largest in the science. We will most probably reach exa-bytes of data even before we will be able to make efficient use of exa-flops.

The following sub-sections will outline how the data challenge can be met.

Reduced resolution output

O(1-5) person year per model, plus education. Data savings O(10x) with a possible multiplicative (2-5x).

Numerical solutions of dynamics equations are not accurate near the grid scale, these scales are needed to get the larger scales (o(5x) larger) right, but are known to be inaccurate at small scales - yet we customarily output all data at full grid resolution. If this data were not written to disk, $O(5x)^2$ of data to disk could be avoided. For some variables, resolution near the grid scale may be strongly influenced by underlying physical differences such as the presence of land or sea or orography, and still provide useful information – estimated at 60% of variables. In all cases, data does not need to be written at the same precision as used in calculations, with additional savings of 2 to 4x in data volume if reduced precision output is introduced (some models already do this to some extent, but most do not).

Tiered storage

O(2) person years per year per supercomputer. O(10x) saving in long-term spinning disk archives.

Current workflows consist primarily of three streams of data: apart from data used for starting and restarting execution, the output is stored on disk for analysis and distribution, with most (if not all) data subsequently archived on tape for longer term re-use. High volumes of data reside on disk to support both immediate use in post-processing, including input to offline models, and longer-term usage by large communities of downstream users. Often the longer-term users exploit a fraction of the data. Post-processing could be accelerated by using fast storage close to the model (e.g. burst buffers) and then sending only some data to disk with the rest going immediately to tape. The initial problem here is to get such tiered storage implemented at all key supercomputing sites, then workflow solutions (e.g. as discussed below) would need to be implemented for each model.

On-the-fly diagnostics

O(5) person years for library support, plus O(1) person year per diagnostic "family" per model, plus behaviour change to a "campaign approach" for model evaluators to exploit data when it is live.

With initial output available in a burst buffer, it will be possible to do a significant amount of post-processing before data is flushed to disk. This would allow calculations where full precision and/or high spatial resolution are necessary (e.g. budget studies) to be carried out more quickly and without the full data ever getting to spinning disk. Similarly, visualisation and much post-processing via offline models could be carried out directly from the burst buffer, avoiding the need for disk storage contention within (for example) the temporal post-processing windows associated with numerical weather prediction. This mode of working would be relatively familiar to the operational weather prediction community, but for core climate scientists, would need a new approach around organising the exploitation of climate simulations into specific periods of activity (as opposed to just organising around "requested output").

Ensemble analysis

O(5-10) person years for library support, plus O(5-10) person years to develop new ensemble analysis techniques. Potential savings are O(ensemble-size) in output reductions.

In most cases existing practice is that ensemble members are executed independently, whether or not they are executed at the same time. If ensembles were executed as one job, it would be possible to do a considerable amount of the analysis either directly using "an ensemble manager/diagnostic sub-model" or utilising the burst buffer and in-flight diagnostics. Such diagnostics could reduce the need to write all the individual ensemble member output to disk, with ensemble statistics calculated directly and only “important” ensemble member behaviour identified and recorded. Considerable work would be needed to efficiently manage premature termination of ensemble members and to identify what “important” characteristics should be recorded. Work on cross-ensemble member compression also offers options to minimise output.

Active storage
It is not yet possible to quantify the effort, but the potential savings in disk storage and processing time are vast.

An ideal storage system would maximise the data on tape, and only bring data onto disk as it was needed. It would also have internal “data reduction” operations, allowing standard operations to be complete while data was in-transit between storage tiers. To be really efficient, such an “active” storage system, would be aware of both data formats, and data semantics, reformatting data for maximal speed of analysis on the fly. Some of this requires vendor support, some will require computing centre support, and some will need to be handled in workflow.

Without some or all of the enhancements listed above, the output requirements of a 1km model could be so large as to make the model prohibitively slow (waiting for input/output) and the output storage prohibitively expensive. If all of the developments that are listed above can be realised with maximal success, this would yield an estimated saving in data volume that needs to be stored by two orders of magnitude, with clear benefits in terms of necessary I/O rates and storage costs. In particular, the I/O rates become physically accessible, which they would not be if all the data had to be written to spinning disk in real time.

3.2 Funding and support instruments

The research effort at the interface between weather-climate and computational science requires a substantial investment in software development and science coordination. The ultimate goal is to create prediction workflows that are usable with different ESM and in different configurations of say spatial resolution, model complexity and ensembles that are portable between different, most likely heterogeneous hardware architectures and that can be applied for both research experimentation and operational production. ESiWACE therefore assumes a central role in funnelling distributed work efforts towards this goal and to create a sustainable software environment given the—at present—rapidly evolving European hardware infrastructure, in which the EuroHPC Joint Undertaking will play a key role.

The ExtremeEarth project concept addresses these issues as a wholistic research and innovation platform that is based on a longer-term, joint European effort providing the necessary key technologies, namely high-performance computing, big data handling and interactive work spaces that provide a break-through in user-driven workflow design. These workflows will deliver the best available information for European society to deal responsibly with environmental extremes. ExtremeEarth has to be seen as the product of many years of investment in research that have led to the recognition that only a large-scale, concerted action can provide this break-through. However, the given funding programmes do not cater for such a focused effort, and the weather and climate community either needs to reduce its ambition, or funding agencies have to create a new opportunity to achieve what ExtremeEarth has laid out in a single solution to a long-standing and critical societal challenge.

Given the dimension of the task, funding and support instruments need to be based on both national and European funding instruments, but also include unfunded collaboration with international efforts in the US and Asia.

At European level, the weather and climate community has greatly benefited from the European Commission's (EC) Horizon 2020 programme where significant investments are being made to address grand societal challenges and to support the development of future and emerging technologies.

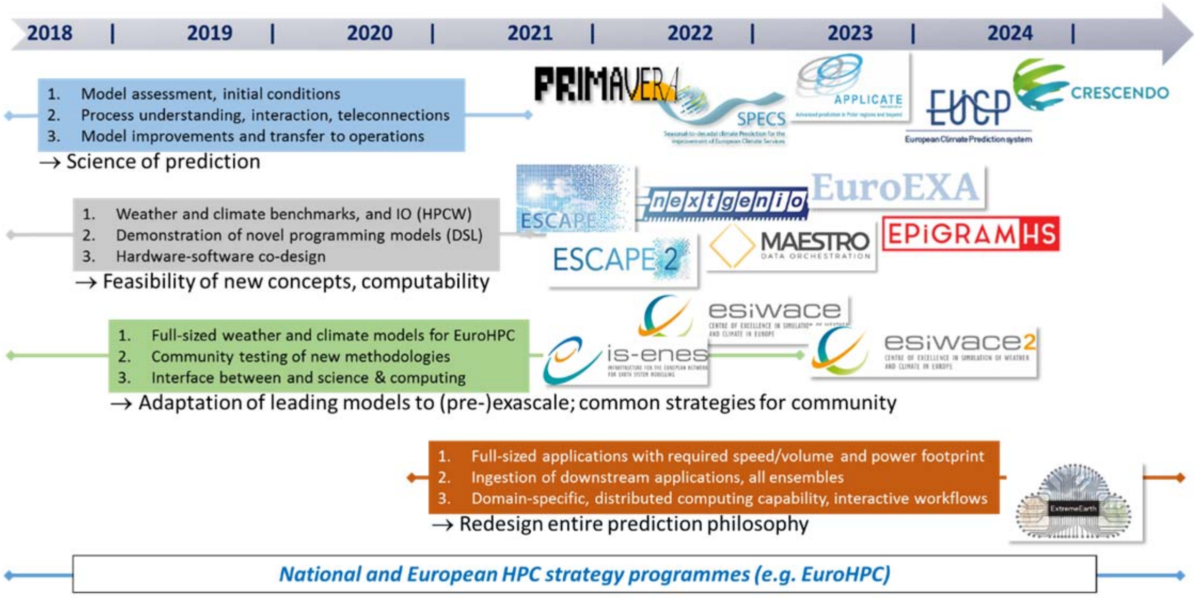

Figure 1: Current streams of (selected) weather/climate science and computational science projects funded by and proposed to the European Commission (for brief project descriptions see Appendix).

As shown in Fig. 1, these funding programmes permitted collaborative research along several thematic streams, namely advancing the science of ESM to produce more realistic simulations, testing the feasibility of new mathematical, programming and workflow concepts to advance computing and data handling performance, ingesting these concepts in the leading European models and preparing operational workloads for exascale HPC infrastructures, and defining the strategic approach for community-wide progress. By including hardware vendors in the project consortia they allow for knowledge transfer and co-design exercises. This project framework is complemented to a large extent by in-kind commitments throughout the community, and by national initiatives.

At national level, there are several activities that complement the existing European projects, which shall in particular foster developments at the applied-computational science interface similar to the first three streams shown in Fig. 1. For example in Germany, a National Modelling Strategy is being discussed, which should leverage common developments of the Max-Planck-Institute for Meteorology and the German Weather Service, the Earth System Modelling programme and Pilot Lab for Exascale Earth System Modelling of the Helmholtz Association, as well as contributions from other parties, like universities. CLIMERI-France is the French national research infrastructure for climate modelling, joining efforts from the Institut Pierre Simon Laplace, Météo France and Cerfacs for common simulations and joined software developments. It is complemented by national projects to support new model developments. In the UK, several initiatives governed by the Joint Weather and Climate Research Programme (JWCRP, jointly involving the Met Office and the Natural Environment Research Council) are addressing numerical weather prediction and Earth System models for exascale. The Swiss Platform for Advanced Scientific Computing (PASC) is a structuring project supported by the Council of Federal Institutes of Technology and is coordinated by CSCS in collaboration with other Swiss universities and the EPFL. Its overarching goal is to position Swiss computational sciences in the emerging exascale era, and it was essential for example, to transition the regional operational weather forecast model COSMO to hybrid CPU-GPU HPC systems.

The present Horizon-2020 funding programme will be replaced by Horizon Europe and Digital Europe in the near future. Key research will be performed under the former while the latter will provide new infrastructures with relevance for weather and climate prediction, for example in HPC and Artificial Intelligence. The new Mission concept under Horizon Europe is only being defined now but will require time to evolve grand Mission themes into effective programmes given the wide range of thematic areas and the fact that the supporting instrument options (like research projects or partnerships) need to be implemented and spun up. There is therefore a risk that the present momentum cannot be sustained and the full value of the European investment into infrastructures like EuroHPC be extracted.

While the weather and climate community has been successful at creating a co-funded strategy and research/innovation framework involving expertise from academia, HPC centres, operational services and industry (as shown in Fig. 1), the strong reliance on external funding at both national and European level adds significant risk to the sustainability of the framework. Also, the uncertain future of large-scale efforts like ExtremeEarth does not allow investing in a more substantial, future-proof and effective infrastructure that would fulfil the needs of European society for adapting to climate change and preparing crucial socio-economic impact areas like energy, food and water for weather and climate extremes.

Therefore, some key questions remain open such as: how to fund the critical mass that is necessary for radically new developments, how to achieve and sustain economy-of-scales developments serving the community, and how to exploit funding opportunities at European and national level without losing sight of the big goals encapsulated in ExtremeEarth? These questions can only be answered at international science policy level but require immediate attention.

4 Where we stand

ESiWACE2 is set to deliver prototypes of global coupled atmosphere-ocean models with a resolution of about 5km by 2021. This will include first technical steps towards exascale in both, the source code of the models and the data middleware. As a follow-up on to the DYMOND project (Stevens et al., 2019) the scientific quality of the models will be assessed and compared with similar international models. ESiWACE2 will also foster the integration of community developments like the high-performance climate and weather benchmark (HPCW) developed by the ESCAPE-2 project and the DSL work performed by Swiss and UK partners of the project. The ESiWACE2 service activities will support the adaption of additional codes from the community to pre-exascale and exascale technology.

To proceed from there towards operational ensembles of true 1km ESM, we need to address the sustainability issues highlighted above. This can only be successful in cooperation with and supported by a wide community of stakeholders, as described in more detail in the ESiWACE deliverable 1.2 “Business Plan”.

References

Bauer, P., Thorpe, A., Brunet, G., 2015: The quiet revolution of numerical weather prediction, Nature, 525, https://doi.org/10.1038/nature14956.

Carman, J., Clune, T., Giraldo, F., Govett, M., Gross, B., Kamrath, A., Lee, T., McCarren, D., Michalakes, J., Sandgathe, S., Whitcomb, T., 2017: Position paper on high performance computing needs in Earth system prediction. National Earth System Prediction Capability, https://doi.org/10.7289/V5862DH3

Chevallier, F., Morcrette, J.-J., Chéruy, F., and Scott, N. A., 2000: Use of a neural-network-based long-wave radiative-transfer scheme in the ECMWF atmospheric model, Quarterly Journal of the Royal Meteorological Society, 126, 761–776

Dueben, P. D., and Dawson, A., 2017: An approach to secure weather and climate models against hardware faults, J. Adv. Model. Earth Syst., 9, 501– 513, doi:10.1002/2016MS000816.

Dueben, P. D., Saarinen, S. and Wedi, N., 2019: Global simulations of the atmosphere at 1.45 km resolution with the Integrated Forecast System, submitted to the Journal of the Meteorological Society of Japan

Dueben, P. D. and Bauer, P., 2018: Challenges and design choices for global weather and climate models based on machine learning, accepted in Geoscientific Model Development. Available under: https://www.geosci-model-dev-discuss.net/gmd-2018-148/

Dueben, P. D., and Palmer, T. N., 2014: Benchmark tests for numerical forecasts on inexact hardware, Monthly Weather Review, 142, 3809–3829

Fuhrer, O., and Co-authors, 2018: Near-global climate simulation at 1 km resolution: establishing a performance baseline on 4888 GPUs with COSMO 5.0, Geoscientific Model Development, 11, 1665-1681

Gysi, T., Osuna, C., Fuhrer, O., Bianco, M. and Schulthess, T.C., 2015: STELLA: a domain-specific tool for structured grid methods in weather and climate models, SC '15: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, Austin, TX, 1-12

Hatfield, S., Chantry, M., Dueben, P. and Palmer, T.N., 2019: Accelerating high-resolution weather and climate models with deep-learning hardware, PASC2019

Joussaume S., Lawrence B. and Guglielmo F., 2017: Update of the ENES infrastructure strategy 2012-2022, ENES Report Series 2, 20 pp.

Koldunov, N. V., Aizinger, V., Rakowsky, N., Scholz, P., Sidorenko, D., Danilov, S., and Jung, T., 2019: Scalability and some optimization of the Finite-volumE Sea ice-Ocean Model, Version 2.0 (FESOM2), Geoscientific Model Development Discussions, https://doi.org/10.5194/gmd-2018-334, in review

Lawrence, B., and Coauthors, 2018: Crossing the chasm: how to develop weather and climate models for next generation computers?, Geoscientific Model Development, 11, 1799-1821

Mitchell J., Budich R., Joussaume S., Lawrence B. and Marotzke J., 2012: Infrastructure strategy for the European Earth System Modelling community 2012-2022, ENES Report Series 1, 33 pp., https://doi.org/10.5285/ca90b281d6ff4cffb9a9bbdeb5fa63f3

Neumann, P., and Coauthors, 2019: Assessing the scales in numerical weather and climate predictions: will exascale be the rescue, Philosophical Transactions of the Royal Society, A377, 20180148

Nurser, A. J. G. and Bacon, S., 2014: The Rossby radius in the Arctic Ocean, Ocean Science, 10, 967-975, https://doi.org/10.5194/os-10-967-2014

Ruprecht, D., 2018: Wave propagation characteristics of parareal, Computing and Visualization in Science, 19(1-2), 1-17

Schulthess T., and Coauthors, 2019: Reflecting on the Goal and Baseline for Exascale Computing: A Roadmap Based on Weather and Climate Simulations, Computing in Science & Engineering, 21(1), 30-41

Stevens, B., and Coauthors, 2019a: DYAMOND: The DYnamics of the Atmospheric general circulation Modeled On Non-hydrostatic Domains. Progress in Earth and Planetary Science, 1-24

Stevens, B., and Coauthors, 2019b: Large-eddy and Storm Resolving Models for Climate Prediction - The Added Value for Clouds and Precipitation. Journal of the Meteorological Society of Japan. Ser. II, 1-115

Voosen, P., 2018: The Earth Machine, Science, 361, 6400, 344-347

Appendix - Externally funded project directory

APPLICATE (https://applicate.eu/): Enhance weather and climate prediction capabilities not only in the Arctic, but also in Europe, Asia, and North America, given the arctic amplification of climate change and unprecedented impacts on high latitude societies, transport and lower latitude extreme weather events.

CRESCENDO (https://www.crescendoproject.eu/): Improve our knowledge of the Earth’s climate processes and provide the best possible future projections to governments and decision-makers, in particular better informs a number of key Model Intercomparison Projects (MIPs) where biogeochemical and aerosol components are of critical importance to delivering realistic future projections.

EPiGRAM-HS (https://epigram-hs.eu/): Validate programming environments for selected applications, extend the programmability and maximise the productivity of application development for large-scale heterogeneous computing systems.

ESCAPE/ESCAPE-2 (http://www.hpc-escape.eu/, http://www.hpc-escape2.eu/): Define the next-generation IFS, ICON and NEMO numerical building blocks and compute intensive algorithms, apply compute/energy efficiency diagnostics, identify new approaches and implementation on novel architectures, perform testing in operational configurations, and prepare benchmarks for the EuroHPC HPC infrastructure.

ESiWACE/ESiWACE2 (https://www.esiwace.eu/): Provide support, training, services, for fostering community models, tools and software, work towards enhanced code performance at pre-exascale and exascale demonstrated by high-resolution simulations.

EUCP (https://www.eucp-project.eu/): Develop an innovative European regional ensemble climate prediction system based on a new generation of improved and typically higher-resolution climate models, covering timescales from seasons to decades.

EuroEXA (https://euroexa.eu/): Co-design a balanced architecture for both compute- and data-intensive applications using a cost-efficient, modular-integration approach enabled by novel inter-die links and the tape-out of a resulting EuroEXA processing unit with integration of FPGA for data-flow acceleration.

ExtremeEarth (http://www.extremeearth.eu/): Revolutionise Europe's capability to predict and monitor environmental extremes and their impacts on society enabled by the imaginative integration of edge and exascale computing and beyond, and the real-time exploitation of pervasive environmental data.

IS-ENES (https://is.enes.org/): Pursue the integration of the Earth’s climate system modelling community and prepare the sustainability of its infrastructure. Foster the common development of models and tools, and the efficient use of HPC. Support the exploitation of model data by the Earth system science community, the climate change impact community and the climate service community.

MAESTRO (https://www.maestro-data.eu/project/): Build a data-aware and memory-aware middleware framework that addresses ubiquitous problems of data movement in complex memory hierarchies and at many levels of the HPC software stack.

NextGenIO (http://www.nextgenio.eu/): Define exascale IO requirements across demanding applications, hardware and data architectures, develop IO workload simulators, support tools for NVRAM, the necessary system-ware and a prototype hardware.

PRIMAVERA (https://www.primavera-h2020.eu/): Develop a new generation of advanced and well-evaluated high-resolution global climate models, capable of simulating and predicting regional climate with unprecedented fidelity, for the benefit of governments, business and society in general.

SPECS (http://www.specs-fp7.eu/): Deliver a new generation of European climate forecast systems, with improved forecast quality and efficient regionalisation tools to produce reliable, local climate information over land at seasonal-to-decadal timescales, and provide an enhanced communication protocol and services to satisfy the climate information needs of a wide range of public and private stakeholders.