The extremes prediction use case

The TransContinuum initiative: exploiting the full range of digital technologies for the prediction of weather and climate extremes

By Peter BAUER, Marc DURANTON and Michael MALMS

Dealing responsibly with extreme events requires not only a drastic change in the ways society addresses its energy and population crises. It also requires a new capability for using present and future information on the Earth system to reliably predict the occurrence and impact of such events. A breakthrough in Europe's predictive capability can be made manifest through science and technology solutions delivering as yet unseen levels of predictive reliability with real value for society.

The "TransContinuum Initiative", initiated by ETP4HPC, ECSO, BDVA, 5GIA, EU MATHS IN, CLAIRE, AIOTI and HiPEAC, offers unprecedented opportunities to overcome the technological limitations currently hampering progress in this area. Beyond providing this use case with better technology solutions, the initiative offers a foundation for an Earth system – computational science collaboration that will eventually lead to science and technology being truly co-developed, and thus to sustainable benefit for one of today's most relevant applications and for European technology providers.

Key insights

- The reliable prediction of weather and climate extremes represents an extreme-scale computing and data handling challenge.
- The European and global Earth system science community has established a solid network of technology solutions for its complex and time critical workflows; however, they will not scale to meet future requirements.
- Climate change and the urgency for societies to adapt to future extremes require new technology solutions that provide breakthroughs for how we observe and simulate the Earth system.
- The cross disciplinary "TransContinuum Initiative (TCI)" promises true co-design opportunities exploiting the entire breadth of digital technologies for the benefit of skilful extremes prediction capabilities in Europe.

Key recommendations

- Future systems will require at least 100 times more computational power for producing reliable predictions of Earth-system extremes with short lead times. It will require implementing an Earth-system digital twin – a cyber-physical entanglement. A solution is a layered federation with fewer elements near the heavy workloads and more elements at the observational data pre-processing front-end and the data analytics post-processing back-end.
- The extensive use of edge computing needs low-cost yet high-performance computing facilities and to overcome the data-transfer bottlenecks between the computing and data intensive parts of the digital twin and downstream applications
- Interoperable machine learning toolkits facilitating the portability of data processing in the cloud and workflow management options in the cloud for orchestrating the rather complex data assimilation and Earth-system simulation workloads should be developed
- Several domains need to be simultaneously promoted:
  - for software: interactive workflows, mathematical methods and algorithms, high-productivity programming environments, performance models and optimization tools.
  - for hardware: heterogeneous processor configurations through accelerators and data-flow engines, high-bandwidth memory, deep memory hierarchies for I/O and storage, super-fast interconnects and configurable computing including the supporting system software stack.

The use case

Natural hazards represent some of the most important socio-economic challenges our society faces in the decades to come. Natural hazards have caused over 1 million fatalities and over 3 trillion Euros of economic loss worldwide in the last twenty years, and this trend is accelerating as a result of the drastic rise in demand for resources and population growth. The combination of likelihood and impact makes extremes and climate action failure the leading threats for our society [1] [2] [3].

The "Earth system extremes" use case relies on very complex numerical Earth system simulation models that ingest hundreds of millions of observations per day to help improve the formulation of the physical process representation as well as produce the initial conditions used for launching predictions of the future. This logic applies equally to weather timescales of days or weeks and climate timescales of decades. These systems are also run as ensembles whereby each ensemble member represents both initial condition and model uncertainty. The result is a prediction of state, but also a prediction of uncertainty, which is crucially important for decision-making on tight schedules [4].

Europe currently leads the world in medium-range weather prediction and is also a major competence centre for global climate prediction [5] [6]. For decades, HPC has been a key enabler for this track-record and has led to weather and climate prediction becoming one of the leading use cases for computing and data handling at large scale. The apparent change in computing architectures has stimulated a wide range of programmes that aim to prepare the operational Earth-system monitoring and prediction infrastructures for future technologies [7]. These programmes increasingly realize that it will take more than progress in HPC to fulfil future extremes prediction targets. This is where the need to exploit the opportunities offered by the entire digital TransContinuum [8] comes into play, and where new ways of co-developing Earth system and computational science need to be found.

Implementing an Earth-system digital twin – called the cyber-physical entanglement in Fig. 1 – through these technologies is the ultimate goal. The loop shown in the figure would be open as humans are continually influencing nature, for example through CO^2 emissions at global scale or irrigation in agriculture at small scales. But natural variability can also affect the system and needs to be replicated, for example extreme events such as volcanic eruptions injecting large amounts of ash into the atmosphere leading to a change in radiative forcing.

Beyond advancing present-day Earth-system simulation and observation capabilities, the digital twin would close the gap between Earth system science and socio-economic sectors in which it might be applied and facilitate a high level of flexible intervention by non-science/technology experts. Hiding the complexity of the

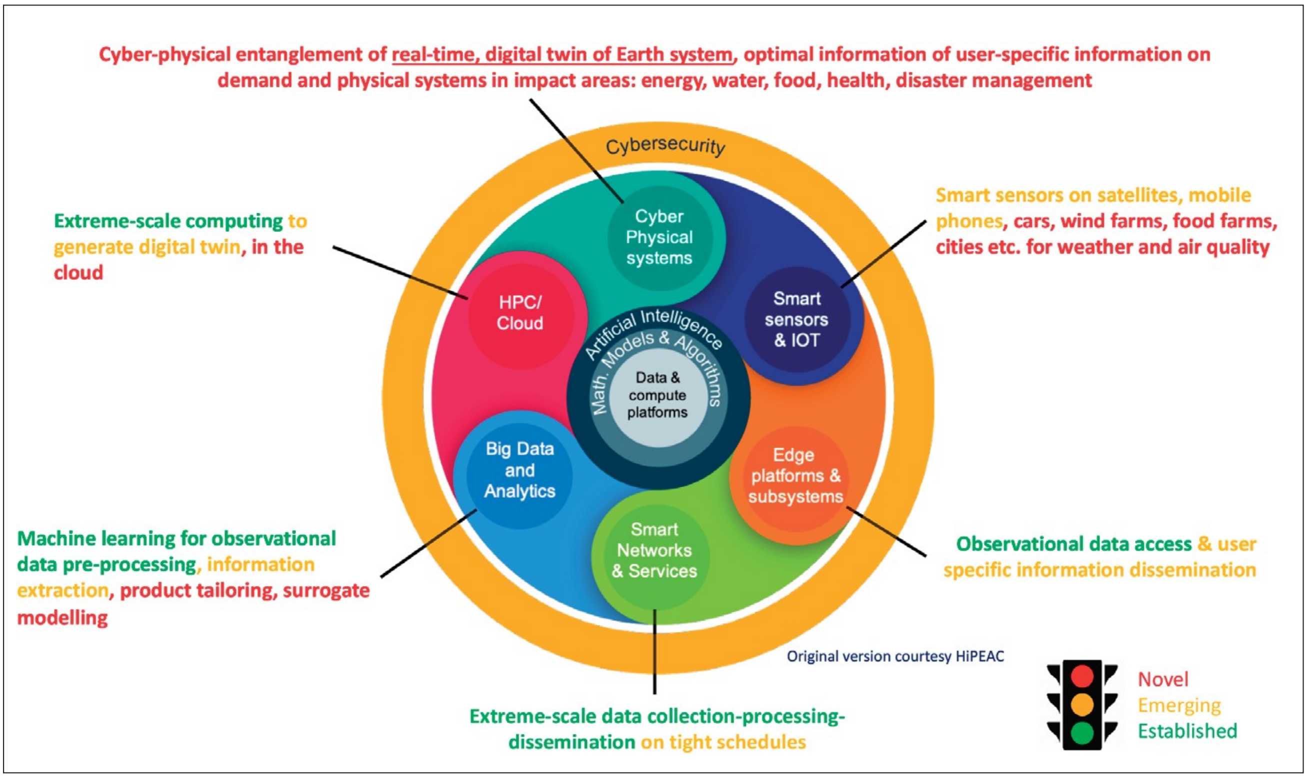

Figure 1: Main elements of digital continuum and relevance for extremes prediction use case including readiness of both application and technologies (green = established in production, but not optimal in performance; yellow = research, not in production mode; red = novel and unexplored)

digital TransContinuum from the user is key to success for user-driven digital twins [9].

As shown in Fig. 1, the TransContinuum embeds smart sensors and IoT within smart networks supplying data to extreme-scale computing and big data handling platforms. These platforms exploit cloud services across workflows giving access to additional distributed computing and data analytics capabilities. Mathematical methods and algorithms as well as artificial intelligence and machine learning provide the glue between the elements of the TransContinuum. Cybersecurity acts as a shield for the entire loop. Such a methodological framework based on extremes prediction also offers gains for other domains.

Challenges along the TransContinuum

Assessing the challenges for the extremes prediction use case with respect to the TransContinuum requires a closer look at the production workflow in today's systems and how they are likely to evolve in the future.

Smart sensors and Internet of Things

At present, Earth-system observations that are used operationally already comprise hundreds of millions of observations collected daily to monitor the atmosphere, oceans, cryosphere, biosphere and the solid Earth, the largest data volumes being provided close to a hundred satellite instruments [10]. This volume is expected to increase by several orders of magnitude in the next decade, bringing with it a need to ingest such observations in digital twin systems within hours. Smart sensor technology is highly relevant for satellite instruments that can perform targeted observations and perform on-the-fly data pre-processing.

However, observations from commodity devices deployed on e.g. phones, car sensors and specialized industrial devices monitoring agriculture, renewable energy sources and infrastructures also offer data that is currently inaccessible to operational services and that can fill vast observational gaps in less developed countries, offers much finer resolution in densely populated areas and in regions of significant socio-economic interest [11]. Beyond technical challenges, a generic approach for fast access to commercial data for public use via public-private partnerships is required, which is already on the agenda of intergovernmental organisations like the World Meteorological Organization. This element of the TransContinuum is only developing now and has huge potential.

Smart networks and services

Collecting and transferring massive amounts of diverse data from devices scattered in multiple areas via distributed pre-processing centres to centralized digital platforms and computing centres requires a suitable network infrastructure. Evolved networks and services should offer secure and trustable solutions that will

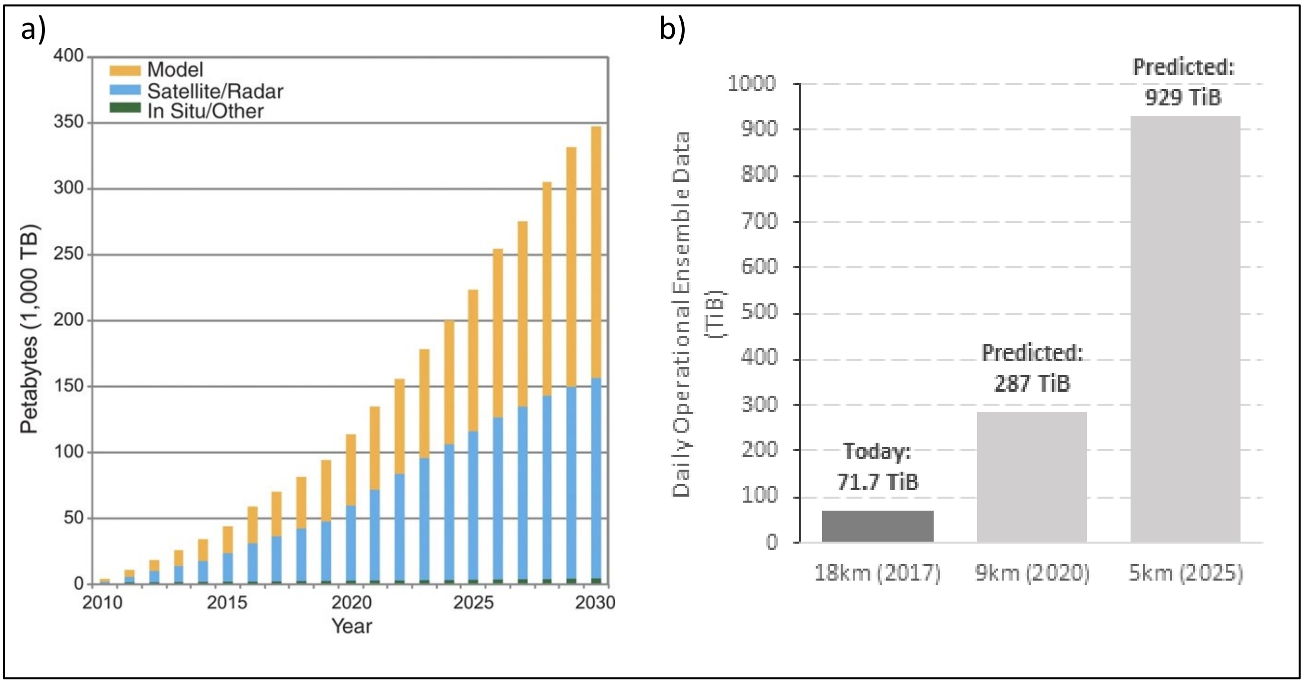

Figure 2: a) The volume of worldwide climate and daily weather forecast data is expanding rapidly, creating challenges for both physical archiving and sharing, as well as for ease of access, particularly for non-experts. The figure shows the projected increase in global climate data holdings for climate models, remotely sensed data, and in situ instrumental/proxy data (from [12]). b) Daily data volumes produced by the 50-member ECMWF weather forecast ensemble at today's spatial resolution of 18 km, the next upgrade to 9 km and the expected configuration in 2025

support the desired quality of service for different data flows. The extremes prediction use case is particularly HPC- and big data-driven and creates a significant footprint for the supporting communication networks. Pre-processing in edge devices can already offload some of this burden as this use case requires near real-time data availability at the HPC facilities. The extensive use of edge computing needs low-cost yet high-performance computing facilities that will interact with end devices as well as with one another.

Given the decade-long evolution of interconnected data collection, transmission, pre-processing, centralized HPC and post-processing, and dissemination in weather prediction workflows there is little room for optimization of present-day systems. However, workloads, data volumes and data diversity grow much faster than in the past and the demand for providing more skilful predictions is more urgent than ever before. Therefore, network and service solutions need to orchestrate and dynamically manage data routing and computing resources with much more flexibility and scalability. Whether both performance and cost-effectiveness need dedicated solutions or whether it can be achieved by a mixture of commercial and institutional systems is an as yet unsolved question.

Big data and data analytics

The volume of Earth-system observation and simulation data in weather prediction already exceeds hundreds of TBytes/day while climate projection programmes produce PByte-sized archives, which take climate scientists years to analyse (see Fig. 2) [12]. The extrapolation of data volumes and production rates to advanced Earth-system digital twins will prohibit the effective and timely information extraction that is critical for timely actions to anticipate and mitigate the effects of extremes. Both simulations and observations need to be generated and combined in the Earth-system digital twin within minutes to hours of time-critical workflows towards near-real-time decision-making.

Overcoming the data-transfer bottlenecks between the computing and data intensive parts of the digital twin and downstream applications is crucial, and future workflow management needs to make such applications an integral part of the observation and prediction infrastructure. Powerful data analytics technology and methodologies offer the only option to make the effective conversion from PBytes/day of raw data to Mbytes/day of usable information. This conversion must be tailored to individual sectors needing to prepare and respond to extremes, namely water, food, energy, health, finance and civil protection. Machine learning-based quality control and error correction, information extraction and data compression as well as processing acceleration will be key.

High-performance and cloud computing

Today, experimental and operational Earth-system simulations use petascale HPC infrastructures, and the expectation is that future systems will require at least 100 times more computational power for producing reliable predictions of Earth-system extremes with lead times that are sufficient for society and industry to respond [13]. This need translates into a new software paradigm to gain full and sustainable access to low-energy processing capabilities, dense memory

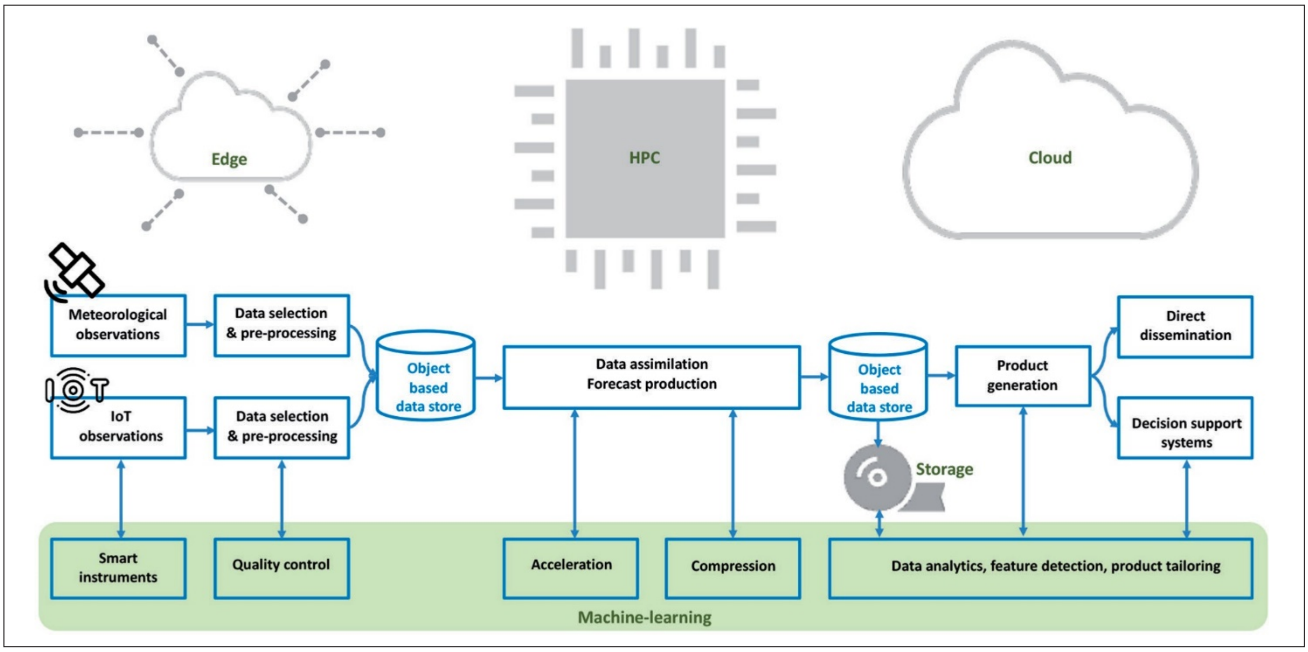

Figure 3: Main components of extremes prediction workflow, their alignment with edge, high-performance and cloud computing elements of the TransContinuum, and key contributions of machine learning to workflow enhancements.

hierarchies as well as post-processing and data dissemination pipelines that are optimally configured across centralized and cloud-based facilities. European leadership in this software domain offers a unique opportunity to turn European investment in HPC digital technology into real value.

A big challenge is the definition of the evolution of today’s strongly centralized workload and data lifecycle management towards a more distributed system – still being data-centric to keep big data near extreme-scale computing – that allows the decomposition of workflows such that the production chain can exploit the full processing and analysis potential of the TransContinuum. A likely solution is a layered federation with fewer elements near the heavy workloads and more elements at the observational data pre-processing front-end and the data analytics post-processing back-end. This distribution also needs to be driven by the urgency of product delivery as, particularly for extremes prediction, processing speed trumps all other considerations.

Adding flexibility without sacrificing performance through the TransContinuum is where most of the strategic recommendations for research and innovation compiled by the European Technology Platform for HPC (ETP4HPC) [14], the Big Data Value Association (BDVA) [15] and High Performance Embedded Architecture and Compilation (HiPEAC) [16] are positioned. For software, these are interactive workflows, mathematical methods and algorithms, high-productivity programming environments, performance models and optimization tools. For hardware, heterogeneous processor configurations through accelerators and dataflow engines, high-bandwidth memory, deep memory hierarchies for I/O and storage, super-fast interconnects and configurable computing including the supporting system software stack are part of the agenda [17].

For the prediction of extremes, the main computing tasks are currently performed on centralized, dedicated systems where the main data storage facilities also lie. The observational input data flow crosses all levels of edge, cloud and centralized computing during which selected pre-processing steps are exercised, for example, for satellite data collection and pre-processing at distributed receiving stations of the ground segment, their further dissemination and processing by space agencies and meteorological centres, and assimilation into models at prediction centres. Similar data flow mechanisms exist for ground-based meteorological observation taken from networks, stations and (few) commercial providers.

Increasingly however, the back-end of model output data processing interfacing with commercial service providers is being placed into the cloud, which also offers better access to machine learning-based data analytics. Here, the biggest gaps are interoperable machine learning toolkits facilitating the portability of data processing in the cloud and workflow management options in the cloud for orchestrating the rather complex data assimilation and Earth-system simulation workloads (see Fig. 3).

Conclusions

Predicting environmental extremes clearly has an extreme-scale computing and data footprint. It has reached a point where only through significant investment in all parts of the TransContinuum can the future needs of our society for reliable information on the present and future of our planet be fulfilled. These investments must go hand in hand with investments in basic Earth system science, to better understand key processes and their interaction, and for service provision ensuring that Earth system data is translated into useful information for society.

While the present digital infrastructures support weather and climate prediction well enough, the chosen domain-specific solutions will not scale to meet future requirements. True near-sensor, edge processing does not yet exist and IoT devices are not yet used; smart and configurable networks as well as flexible HPC and cloud solutions are not available; and even for the extreme-scale computing and data handling tasks, the present algorithmic and programming frameworks do not allow the exploitation of the real potential of emerging technologies.

Building on past European science and technology programmes such as the Future and Emerging Technology for HPC (FET-HPC) [18] under Horizon 2020, the recent implementation of the EuroHPC Joint Undertaking, and the Copernicus Programme, Europe has actually recognized the present gaps and needs for serious investments. This has motivated the ambitious Destination Earth action [19]. In support of the Green Deal and the European strategy for data [20], Destination Earth promises to bring the Earth system science, digital TransContinuum and service development strands together assigning the highest priority to extremes prediction and climate adaptation. The programme has adapted the concept of digital twins of the Earth system for this purpose as promoted by HiPEAC and ETP4HPC and its strong digital infrastructure contribution has huge potential for achieving European technology leadership by addressing one of the principal challenges for today's society.

Peter Bauer is the Deputy Director of the Research Department at the European Centre for Medium-Range Weather Forecasts, Reading, UK.

Marc Duranton is researcher at the Research and Technology Department of CEA (Alternative energies and Atomic Energy Commission), France and the coordinator of the HiPEAC Vision 2021.

Michael Malms is the SRA lead editor within the office team of ETP4HPC and main initiator and coordinator of the TransContinuum Initiative.

This article was initially published as part of the HiPEAC Vision available at hipeac.net/vision.
V.1, January 2021.
Cite as: P. Bauer, M. Duranton, and M. Malms. The extremes prediction use case. In M. Duranton et al., editors, HiPEAC Vision 2021, pages 44-49, Jan 2021.
The HiPEAC project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement number 871174.
© HiPEAC 2021
This document was prepared with the support of the EXDCI project, which receives funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 800957.

---

[1] UNISDR, "Economic losses, poverty & disasters, 1998- 2017", (UNISDR, Geneva, 2018)

[2] AON, "Weather, climate & catastrophe insight", Annual report, (AON, Chicago, 2020)

[3] WEF, "The global risks report 2020", Insight report 15th edition, (World Economic Forum, Geneva, 2020)

[4] P. Bauer, A. Thorpe, and G. Brunet (2015). The quiet revolution of numerical weather prediction. Nature, 525(7567), 47-55

[5] R. A. Kerr (2012). One Sandy forecast a bigger winner than others. Science, 736-737

[6] P. Voosen (2019). New climate models predict a warming surge. Science, 16

[7] P. Neumann, J. Biercamp, (2019, September). ESiWACE: On European Infrastructure Efforts for Weather and Climate Modeling at Exascale. In 2019 15th International Conference on eScience (eScience) (pp. 498-501). IEEE

[8] "TransContinuum Initiative (TCI): our vision", https:// www.etp4hpc.eu/transcontinuum-initiative.html

[9] P. Bauer, B. Stevens, & W. Hazeleger (2020). A digital twin of Earth for the green transition. Nature Climate Change, submitted

[10] W. Balogh, and K. Toshiyuki. "The World Meteorological Organization and Space-Based Observations for Weather, Climate, Water and Related Environmental Services." In Space Capacity Building in the XXI Century, pp. 223-232. Springer, Cham, 2020

[11] J. M. Talavera, L. E. Tobón, J. A. Gómez, M. A. Culman, J. M. Aranda, D. T. Parra, ... & L. E. Garreta, (2017). Review of IoT applications in agro-industrial and environmental fields. Computers and Electronics in Agriculture, 142, 283-297

[12] J. T. Overpeck, G. A. Meehl, S. Bony & D. R. Easterling (2011). Climate data challenges in the 21st century. Science, 331(6018), 700-702

[13] T. C. Schulthess, P. Bauer, N. Wedi, O. Fuhrer, T. Hoefler & C. Schär (2018). 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

[14] ETP4HPC, https://www.etp4hpc.eu/pujades/files/ETP4HPC_SRA4_2020_web(1).pdf

[15] BDVA, https://www.bdva.eu/BDVA_SRIA_v4

[16] HiPEAC, https://www.hipeac.net/vision/#/latest/

[17] P. Bauer, P. D. Dueben, T. Hoefler, T. Quintino, T. Schulthess & N. P. Wedi (2020), The digital revolution of Earth-system science. Nature Computational Science, submitted.

[18] "European HPC technology research projects", https:// www.scientific-computing.com/white-paper/european- hpc-technology-research-projects

[19] "Destination Earth (DestinE)", https://ec.europa.eu/ digital-single-market/en/destination-earth-destine

[20] "A European Strategy for Data", https://ec.europa.eu/ digital-single-market/en/european-strategy-data