



HAL
open science

Towards Integrated Hardware/Software Ecosystems for the Edge-Cloud-HPC Continuum

Gabriel Antoniu, Patrick Valduriez, Hans-Christian Hoppe, Jens Krüger

► **To cite this version:**

Gabriel Antoniu, Patrick Valduriez, Hans-Christian Hoppe, Jens Krüger. Towards Integrated Hardware/Software Ecosystems for the Edge-Cloud-HPC Continuum. 2021, 10.5281/zenodo.5534464 . hal-03358930

HAL Id: hal-03358930

<https://hal.science/hal-03358930>

Submitted on 1 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Towards Integrated Hardware/Software Ecosystems for the Edge-Cloud-HPC Continuum

Supporting integrated applications across the Edge-Cloud-Supercomputer layers to address critical scientific, engineering and societal problems

White Paper

Gabriel Antoniu (Inria), Patrick Valduriez (Inria),
Hans-Christian Hoppe (SCAPOS), Jens Krüger
(Fraunhofer ITWM)

28/09/2021

etp4hpc.eu

[@etp4hpc](https://twitter.com/etp4hpc)



EUROPEAN
TECHNOLOGY
PLATFORM
FOR HIGH
PERFORMANCE
COMPUTING



Introduction

Modern use cases such as autonomous vehicles, digital twins, smart buildings and precision agriculture, greatly increase the complexity of application workflows. They typically combine physics-based simulations, analysis of large data volumes and machine learning and require a hybrid execution infrastructure: edge devices create streams of input data, which are processed by data analytics and machine learning applications in the Cloud, and simulations on large, specialised HPC systems provide insights into and prediction of future system state. From these results, additional steps create and communicate output data across the infrastructure levels, and for some use cases, control devices or cyber-physical systems in the real world are controlled (as in the case of smart factories). All of these steps pose different requirements for the best suited execution platforms, and they need to be connected in an efficient and secure way. This assembly is called the *Computing Continuum* (CC) [1]. It raises challenges at multiple levels: at the application level, innovative algorithms are needed to bridge simulations, machine learning and data-driven analytics; at the middleware level, adequate tools must enable efficient deployment, scheduling and orchestration of the workflow components across the whole distributed infrastructure; and, finally, a capable resource management system must allocate a suitable set of components of the infrastructure to run the application workflow, preferably in a dynamic and adaptive way, taking into account the specific capabilities of each component of the underlying heterogeneous infrastructure.

To address the challenges, we foresee an increasing need for integrated software ecosystems which combine current “island” solutions and bridge the gaps between them. These ecosystems must facilitate the full lifecycle of CC use cases, including initial modelling, programming, deployment, execution, optimisation, as well as monitoring and control. It will be important to ensure adequate reproducibility of workflow results and to find ways for creating and managing trust when sharing systems, software and data. All of these will in turn require novel or improved hardware capabilities. This white paper provides an initial discussion of the gaps. Our objective is to accelerate progress in both hardware and software infrastructures to build CC use cases, with the ultimate goals of accelerating scientific discovery, improving timeliness, quality and sustainability of engineering artefacts, and supporting decisions in complex and potentially urgent situations.

Key insights

- There is a clear trend to combine numerical computations, large-scale data analytics and AI techniques to improve the results and efficiency of traditional HPC use cases, and to advance new use cases in fields such as autonomous vehicles, digital twins, smart buildings/towns etc. Such use cases are typically implemented as complex workflows and will require the coordinated use of supercomputers, cloud data centres and edge-processing devices.
- Today, separate, efficient software ecosystems [2] exist for the management of computation, communication and data on supercomputer facilities, cloud infrastructures or edge-based systems. Yet, these address the specific requirements of their infrastructure layer and typically fail to smoothly interoperate and cooperate. The same is true for the different machine learning (in particular Deep Learning) ecosystems.
- Complex workflow orchestration across the whole continuum leads to challenges at multiple levels: application/algorithmic level (programming paradigms), middleware level (deployment, execution, scheduling, monitoring, data storage and transfer, processing and analysis) and resource management level.
- CC workflows will likely have specific hardware requirements e.g. for on-the-fly encryption, efficient and low latency communication, transparent compression, automatic brokerage, in-network data processing and generally for energy efficient computation and communication as large scale systems will have significant energy footprints [3] [4].
- Co-design and cooperation among experts in the different areas involved (e.g. HPC, data analytics, AI/Deep Learning, cybersecurity, mobile communication) will be a key prerequisite for building Computing Continuum hardware and software and support complex application workflows in an effective, efficient and secure way. The TransContinuum Initiative (TCI)¹ is providing such a cross-domain cooperation framework.
- HPC datacentres today feel the pressure of Cloud computing vendors like AWS and MS Azure [5]. It will be attractive for small to medium scale HPC workloads to leverage Cloud interfaces and instantly available Cloud resources which offer almost unlimited scalability and compute throughput.
- HPC centres will have to adopt such Cloud interfaces, and they are poised to gain additional customers from CC use cases and from the deep learning and data analytics communities.

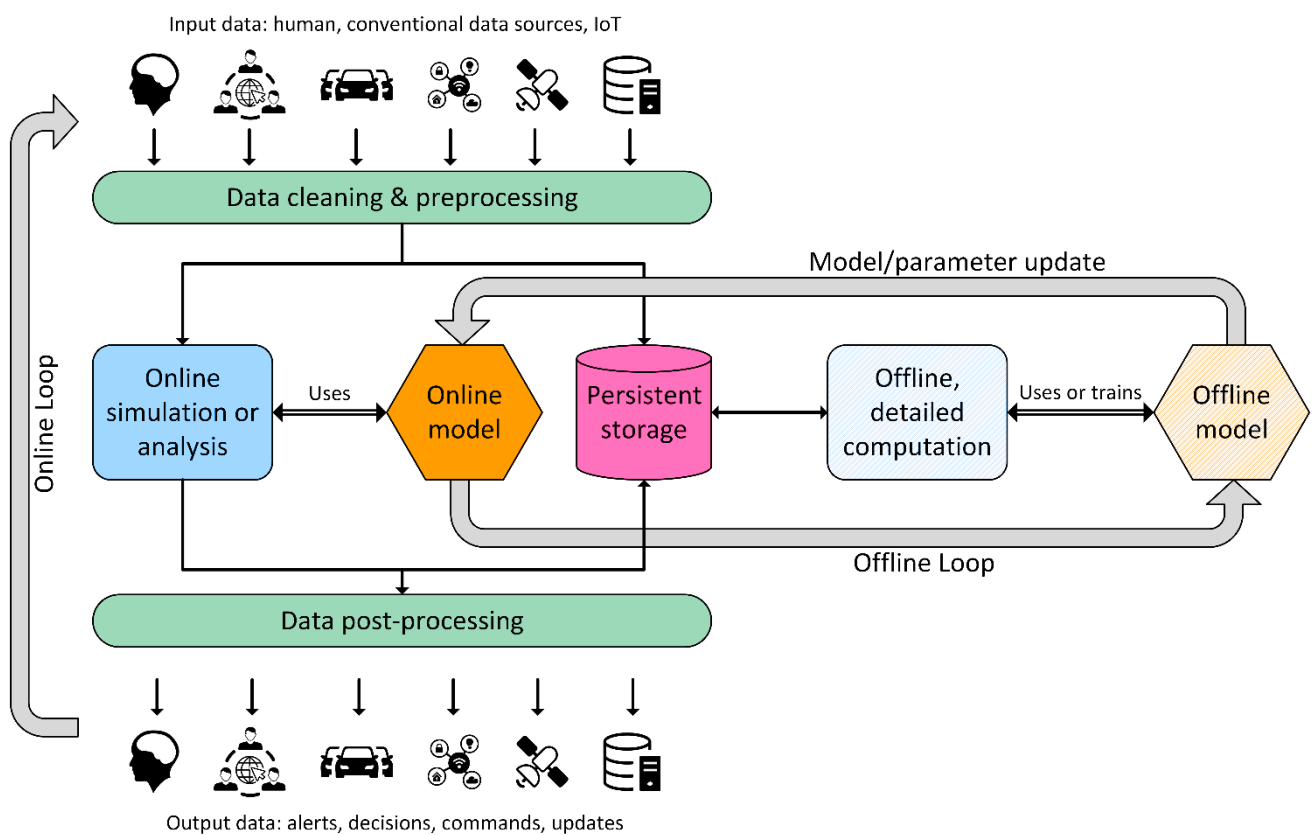
¹ <https://www.etp4hpc.eu/transcontinuum-initiative.html>

Key recommendations

- Applications/algorithms: develop programming models and systems which integrate HPC, AI/ML and data analytics processing and facilitate hybrid applications such as AI-enabled simulations. This is not a CC requirement by itself, yet it will help in programming the steps in a CC workflow.
- Data storage, transfer and processing: unified abstractions must enable interoperable data storage and processing across the continuum, and facilitate data analytics at all levels (HPC center, cloud, edge). In addition, automated data placement and transformations in transfer should be supported.
- Workflow programming: programming models and systems are required which can support large, dynamically evolving workflows, with workflow steps combining HPC, data analytics and AI/ML processing.
- Workflow deployment, orchestration and monitoring: deployment and orchestration of workflows must be seamless and dynamically adapt to the load and the condition of the distributed infrastructure. Fine-grained monitoring of running workflows and reproducibility of results (to a degree adequate for the actual application) are important. This could, for instance, be achieved by extending recent advances in the Cloud Computing field.
- Software interoperability and composability: this concerns the established HPC, data analytics and AI/ML stacks for which interoperable and composable implementations will be required .
- Authentication, authorisation and accounting: these have to be interoperable across all layers of a CC infrastructure, and be combined with a secure, pay per use billing mechanism.
- Data security: interoperable mechanisms for data encryption and transfer, data access control and monitoring will be needed.
- Management of heterogeneous systems: a CC infrastructure will include a wide variety of heterogeneous compute, storage & communication systems, including accelerators like GPGPUs, FPGAs and others. It will be critical to find an integrated way to manage these resources in an efficient and secure manner.
- HPC data centres must rethink their offerings, adapt to new usage models and customer requirements and evolve their business models accordingly. De-facto cloud standard interfaces APIs for storage and computation must be considered.
- HPC data centres must address the increasing heterogeneity of computing hardware with new approaches for resource management and user management.
- Energy efficiency: the energy consumption of CC use cases must be reduced as far as possible, for instance by the development and use of highly efficient hardware, the minimisation of data movement and communication between systems and layers, and the adoption of novel algorithmic approaches, such as mixed precision computing.

What is the Computing Continuum and why is it important

There is an extremely rapid proliferation of digital devices generating valuable data in many fields of application (science, meteorology, autonomous vehicles, industry 3.0/4.0, social media, etc.). The combination of physics-based simulation (classic HPC) and data-driven modelling (using machine learning techniques) has been shown to create impressive results, besides improving time and energy required [6] [7]. Large scale distributed supercomputer/Cloud/Edge infrastructures demonstrate how compute and data throughput can be scaled [8]. Given the above, it now seems possible to create new data-driven use cases which address important scientific, commercial and societal challenges in novel ways, by combining simulation, analytics and learning. Such use cases will make full use of CC infrastructures, and to a large degree depend on their existence.



A Use Case in Precision Agriculture

Modern greenhouses have sophisticated ways to control the vegetable environment to increase their production, based on models of plant growth. However, they require more and more parameters to be set by the grower (e.g., over 200 in a standard soil-less glasshouse for tomatoes)². To be relevant, these models need to be adapted to the location of the farm, the needs of the species and of the cultivar, their potential in yield and quality (dry matter and sugar content). A promising approach consists in combining models of plant and climate development with actual monitored data processed through machine learning algorithms, to accurately model the greenhouse costs and its production potential. To introduce near-real time data assimilation of climate and plant observation to correct the simulations scenarios of plant needs, data inputs will be computed by a model (Digital Twin) executed in the cloud; then fusion with recent data dynamics and projections will be done by the models deployed on the edge.

A major challenge in such a scenario is that the complexity of the system represented (indoor climate, greenhouse management systems, plants interacting and reacting to uncertain weather) can make the simulation results drift significantly from the monitored data dynamics (energy failure, climatic extremes, changes in grower production strategy...). When this occurs, the models need to be corrected according to the observed data dynamics (calibration of the models, learning of machine learning algorithms). When such drifts are detected for the Digital Twin, the process should automatically launch corrections of part of the models, which in turn calls for on-demand, dynamic resource allocations in the cloud or at the edge. For example, if a breakdown in greenhouse climate management occurs and is fixed during the day, the quasi-real-time forecast and decisions on the Edge should be reviewed while the global strategy computed mainly in the cloud stays valid. In contrast, on-demand Cloud simulations need to be deployed when significant changes in recent data dynamics is detected.

Overall, this use case calls for the deployment and the execution of a distributed, cross platform workflow from the edge to the cloud, with the potential of the unpredictable data dynamics triggering on-demand, computationally intensive simulations in the cloud and related processing at the edge.

The use case scales by including more greenhouses, growers and regions.



A sensorized greenhouse

² This use case is proposed by CYBELETECH <https://www.cybeletech.com/en/home/> in collaboration with Ctifl <https://www.ctifl.fr>.

A Use Case in Smart Buildings & Cities

Buildings and infrastructure are pivotal in the socio-economic transition towards a sustainable and climate-neutral economy. People spend most of their life in buildings, which are responsible for approximately 50% of the global energy consumption across their whole life-cycle. Optimizing their energy consumption is therefore highly important to ensure sustainable development.

Additionally, more and more sensors are installed over built-up areas to form so-called smart city environments. This way real time air pollution monitoring, traffic and citizen dynamics, road conditions and other events can be monitored and analyzed. Based on that data and on detailed meteorological forecasts, prescriptive measurements can be taken e.g. to control building lighting and heating/cooling/ventilation or the traffic flow to reduce energy use and air pollution, trigger warnings or take action with predictive maintenance.

To address this goal, based on historical operation data, machine-learning techniques are largely deployed. This kind of data-driven analytics can be carried out on a distributed platform. Real-time sensor data will be used as an input, integrating Edge processing for data preprocessing with Cloud-based analytics. Further computation-bound simulations will be used for predicting effects. To enable high-resolution computations, access to an HPC facility may be required. Thus data movement and data-driven calculations will continuously be performed on the hybrid platform.

A key aspect is that the input data can vary in frequency, relevance, amount of available data, etc. This can lead to variations in data provisioning, that will induce variations of the computational/processing workloads; this requires dynamic reallocation of processing tasks and resources across the CC; at the same time, this will result in a more/less intensive computation for learning behavioural models, and, lastly, to a modification in the forecast and hindcast calculation due to better match actually observed data. This variability does require support for seamless data processing across the CC and also smart, dynamic scheduling, as well as allocation of processing tasks, to react to data variations.



AFRISOL³: a sensorized smart building in Madrid (this is a pilot building for the project investigating the use case described here)

³ <https://www.ciemat.es/portal.do?IDM=61&NM=2&identificador=92>

Earth System Digital Twin

The “Destination Earth” initiative⁴ is creating detailed, digital replicas of the Earth, and will model a wide range of natural systems (including atmosphere, oceans and rivers, the cryosphere and subterranean water) at a very high and hitherto unattained level of detail and precision, using proven simulation methods in combination with AI/ML techniques. Besides its obvious value for all Earth-related scientific activity, this digital twin will enable detailed investigations of the full range of climate change effects, guide the definition of effective mitigation measures and support data-driven decisions by political and administrative bodies. In addition, the platform will be used to forecast environmental extremes and disasters (like hurricanes, wildfires or floods) and enable effective responses to limit their natural, economic and societal impact.

The system relies on a large, distributed network of sensors providing real-time input data, which includes space assets, meteorological measurement stations, sensors in planes and ships, buildings and in perspective also sensors in mobile devices. The data coming in is of a very diverse nature - it includes time-series of scalar and vector data (like temperature, wind direction and speeds) all across the globe (and up to the tropopause), digital multi-spectral images and videos, plus potentially acoustic and seismic data. In addition, much of the incoming data has to be calibrated and assimilated with the more controlled traditional meteorological data. This task clearly requires significant data processing on the edge and in cloud layers, combined with a massive and reliable storage system, since no byte should ever be left behind.

The simulation of the natural phenomena considered based on the input data is a classical HPC task, and proven technology does exist in this area. However, the scale of the simulations is unheard of due to a significant increase in resolution/detail, and capturing the interactions between processes such as weather, ocean dynamics, and hydrology is very challenging. All simulation results will be annotated and stored forever, to serve as inputs for science, basis for policy decisions, or as a prerequisite for improving and calibrating models. The amount of data generated here will likely dwarf the input data. To reduce equipment costs, as well as time and energy to solution, advanced AI/ML techniques will be used. The natural location for the simulation activities are supercomputer centres, working together with large Cloud-based storage centres. The AI/ML models used will be trained on suitable systems in the CC, and regular validation and if necessary re-training will be performed.

To explore effects of human activities, be it the long-term impact of different economic and societal measures on climate change or the short-term effects of disaster relief, reduced-order models will be created and kept up-to-date. Such models can run for inference in smaller, Cloud-based data centers, or at the edge, close to, for instance, a flooded area. The initial creation and update of such models will likely happen in other parts of the CC. The number of end users of such models, or of the simulation data, will be very large, so effective data distribution and replication techniques will need to be implemented, and secure authenticated access has to be provided using scalable and intuitive Cloud mechanisms.

Taken together, Destination Earth will exercise all CC infrastructure and software layers, integrating data acquisition and assimilation, simulation of multiple scenarios, creation of forecasts, creation and calibration of both forecast and reduced order “decision” models, reliable storage of all input and simulation data, and handling of unforeseen user requests with highest priority (such as for disaster relief). The Destination Earth infrastructure will be mission critical, requiring special attention to ensure high availability, and it will need to handle tens of thousands of users from science, government, and industry.

⁴ <https://digital-strategy.ec.europa.eu/en/library/destination-earth>

Autonomous Vehicle Management


Autonomous, self-driving vehicles need to rely on powerful edge computing support for their operation. On-board sensors deliver a continuous, high-bandwidth stream of data about the vehicle's environment (through videos or radar/lidar), neighboring vehicles might share their data, and internal sensors reflect the condition of the vehicle itself. All this information is processed under hard real-time requirements by AI/ML based software components, resulting in specific steering decisions which will protect the safety of the vehicle and its environment while proceeding on a course to the eventual target. Due to the real-time schedule, such AI/ML inference has to happen on an Edge system, be it local systems in the vehicle or in a close-by location (near to a rail track/road).

The AI/ML models used have to work for a truly overwhelming number of complex scenarios, and it is clear that regular updates of such models will be necessary. Anticipating all conditions a priori (or even after a long test driving campaign) seems unfeasible. In particular, accidents or "near misses" will no doubt happen, and "safe fail" mechanisms will be triggered if the main AI/ML engine cannot reach an optimal decision. In these cases, the underlying models must be updated, using the actual incident data, to prevent accident scenarios from being repeated, and to cover corner cases by the main AI/ML engine.

For a large fleet of autonomous vehicles in day-to-day operation, a CC use case can be derived: a large number of edge devices send incident reports and large-volume data to large HPC installations, which will update or re-train the AI/ML models used by the edge systems, and regularly download updates to all managed edge systems.

The system will process private/protected data, and for a multi-tenant scenario, confidentiality about AI/ML models and vehicle-specific data has to be guaranteed. Should severe accidents happen, their data has to be fed into a high-priority training and update operation, whereas incidents relating to possible optimisations (such as avoiding triggering safe-fail) might conceivably be batched together.

The recently disclosed supercomputing installation of Tesla [9] shows the scale of HPC platform needed for the initial creation of models for self-driving automobiles by a single vendor. Our CC use case scales by the number of client edge systems/vehicles, the number of different versions (we would anticipate that such a CC infrastructure would be used across vendors), and the aggregated usage time (or mileage) of the vehicles. The resolution and complexity of environment data will play a role in scaling, too. This clearly shows that leading-edge HPC capacity will be needed for re-training, in addition to a highly capable data transmission and aggregation system handling many millions of endpoints, ways to prioritize urgent updates, storage systems able to keep all transmitted vehicle and environment data, and a scheduling/orchestration system which will safeguard the creation and distribution of urgent updates.



Where are we now?

Today's landscape: an archipelago of disconnected solutions. Today's software approaches available to address the needs of CC use cases consist of separate software stacks optimised for different goals, specific to the target infrastructure (supercomputer, cloud datacenter, edge devices, respectively). There is no simple solution making it possible to deploy and orchestrate a combination of consistent interoperable components across the full continuum. Moreover, the existing stacks are often specialised to be applied in their prevalent use cases; for example, combining an ab-initio QCD simulation with an AI engine for assessing the resulting, macroscopic characteristics of materials characteristics is extremely hard, even if the entire code was to run on a single system. Below, we list the major hurdles to overcome:



- Multitudes of software development stacks are tailored to specific use cases, with no guarantee of interoperability and composability between them. This greatly impedes application software development for integrated CC use cases.
- Across HPC, data analytics and AI, the existing software stacks have very different requirements for their execution infrastructure, and cannot be run efficiently on a single, homogeneous system.
- The different compute, storage and communication systems of a complex CC installation will belong to different owners and be operated according to different rules:
 - Common abstractions of AAA (Authentication, Authorisation and Accounting) plus resource usage policies and pricing are needed; lack of this has limited the uptake of grid computing in the past.
 - Similarly, compatibility and interoperability across all parts of a CC infrastructure must be assured; this includes data formats, communication, security/encryption, data processing paradigms, etc. Suitable standards and their wide-spread adoption will play a very important role here.
 - Large-scale heterogeneity has to be managed in an effective and efficient way. This again cuts across compute, storage and communication systems, and the scheduling/orchestration middleware has to optimize the mapping of workflows with regards to performance and energy use.
- Security of platforms, communication links and data has to be assured, and specific rules & restrictions for the access to private or confidential data have to be enforced.
- On the hardware and systems side, the basic building blocks are common between the HPC, data analytics and AI/ML field, yet the system architectures and configurations are tailored to their respective usage area and do differ significantly. Flexible and efficient operation of CC infrastructures will require either reconfigurability at the system level or very agile mapping of processing steps to different systems.

Building Integrated Software Ecosystems for the Continuum: Challenges

Moving from an archipelago of disconnected software and system solutions towards integrated ecosystems requires addressing the aforementioned hurdles. This leads to challenges at multiple levels.



Application-level challenges

At the application level, traditional physics-based simulations (traditionally executed on HPC systems) need to smoothly cooperate with data-driven, learning-based analytics and prediction engines (typically run on clouds). Programming the workflow at the highest level requires the ability to consistently combine all these components in a unified framework. This requires innovative algorithmics and flexible programming models and supporting environments, which also safeguard performance and energy-efficiency.

Composability (the ability to combine multiple programming models or software stacks for a single application with defined rules) will be needed. Workflows could then combine the use of different robust programming models to enhance usability and achieve efficiency using standardized interfaces within and between workflow steps.

Storing and processing data across the continuum

Mastering the data flows from edge devices to clouds and supercomputers requires flexible, **shared data abstractions and unified mechanisms for data storage, to support distributed processing and analytics** across the whole CC infrastructure. This further increases the relevance of the classic Big Data challenges (e.g., Volume, Velocity, Variety), which can be projected in specific ways on the continuum:

Coping with Extreme Volume. Hybrid workflows running across the CC might need to process both synthetic data generated by simulations and real, sensor-originated data. Therefore, data volumes are virtually infinite.

The storage infrastructure needs to support the access and processing of “cold”, historical data and “hot”, real-time data (possibly streaming data), which further accentuates the volume pressure. Data lifecycle management becomes a key aspect which must consider technical, legal and application requirements and restrictions.

Coping with Extreme Velocity. The data processing ecosystem typically needs to combine processing of historical data with real-time data processing (e.g., in-situ/in-transit processing on HPC systems and cloud-based or edge-level stream-based processing) in a unified way. Unifying data processing approaches in a common software ecosystem becomes a challenge. Focusing on energy efficiency will require data movements to be kept to a minimum.

Coping with Extreme Variety. Unified data storage abstractions and systems enabling efficient data sharing to enable distributed processing and analytics across the Computing Continuum will have to overcome the extreme variety challenge: data have to be exchanged from Edge devices to HPC-class machines, therefore the data should be presented in a coherent and easy to use form for all machines in the “continuum”. This requires:

- Interoperability of the data exchange formats.
- "Semantic interoperability" through shared ontologies understandable at all levels.
- Storage interfaces should match the needs of the application, e.g. by supporting object or key-value stores.

Managing computation across the continuum

At the middleware level, a major challenge is to facilitate the development of tools for seamless deployment, orchestration, scheduling, and execution of complex workflows across hybrid, heterogeneous CC infrastructures. This in turn calls for:

- Dynamic scheduling and orchestration of workflows which evolve at runtime, optimizing time and energy to solution on a dynamically evolving system.
- Support of a wide variety of processors, accelerators, storage devices and systems, and communication systems. A CC infrastructure will be deeply heterogeneous, and very likely so even within each of its layers.
- Support seamless deployment and migration of workflows or workflow steps, maybe aided by containerisation methods.
- Monitoring systems which provide the data required by adaptive, dynamic schedules to efficiently run workflows and workflow steps; recent advances in the support of micro-services on Cloud data centres may be relevant here.
- Definition and automatic derivation of performance models, which would enable a priori scheduling decision to be taken with a high degree of confidence.
- HPC centres should offer APIs to be integrated into workflows and for monitoring.

Managing dynamic workflows with ad-hoc load variation

In complex workflows, combining simulation and real-time data analytics, there are situations when one must react to certain events, depending on data contents or depending on interactive requests. This typically happens when on-demand simulations are triggered with real-time constraints for rapid decision making. The use cases described above exhibit a need for efficient management of dynamic load variations across the continuum.

In complex workflows executed over a CC infrastructure, changes in data characteristics and dynamic changes in the infrastructure create the necessity to dynamically adapt the mapping of the workflow onto the infrastructure and swiftly reconsider on which resources the different workflow steps are to be executed. In some applications, such as disaster warning and response [10], this need may appear when parts of the infrastructure suddenly become unavailable. This requires efficient coupling between Cloud-oriented dynamic orchestrators and traditional batch-based resource management systems, as a step towards more integrated software approaches to dynamic resource management across the continuum. It might even be necessary to question the batch-oriented job scheduling in HPC systems.

AI-related challenges

AI-powered workflows running on HPC-enabled infrastructure are gaining momentum. Their potential execution on hybrid infrastructures create new challenges:

- Optimising resource usage for AI workflows requires strategies to improve resource utilisation by automatically rescheduling jobs to best suited hardware option which in fact maximises the energy efficiency and reduces unnecessary allocation of hardware.
- The new heterogeneity of use cases and hardware must be taken into account;
- The deep learning software stacks must be supported (python dependencies handling, containerisation);
- Ad-hoc training and inference runs with tight timing constraints must be supported (urgent and interactive computing);

Cybersecurity challenges

To effectively leverage a CC infrastructure, significant challenges relate to cybersecurity [11], when it comes to federated authentication, authorisation and accounting, monitoring, resource allocations, encryption, user insulation, container certification, etc. These important considerations are currently impeding the deployment of large scale workflows. This is especially true for handling GDPR-related data. HPC centres that consider being involved in workflows handling such data must make sure to provide all tools necessary to address regulatory requirements.

Cooperation challenges

Addressing the aforementioned challenges requires the collaboration of several expert communities (HPC, Big Data, AI, cybersecurity, IoT, 5G, etc.). Establishing commonly agreed, shared goals and priorities, as well as a common vocabulary and common roadmaps appears as a key strategic objective. Achieving this goal will enable strong synergies and well-founded, sound formulation of relevant and useful research directions. This is precisely the core motivation underlying the TransContinuum Initiative (TCI)⁵, whose efforts are building the framework for such a cooperation. The open directions discussed in this paper on how to build an integrated ecosystem to leverage the CC are contributing to this effort.

Conclusions

In an increasing number of areas we are witnessing the emergence of complex workflows combining simulations, data analytics and learning, running on hybrid infrastructures where supercomputers are connected to cloud data centres and edge devices. To address the requirements of such workflows, we foresee a growing need for integrated software ecosystems which should build on state-of-the-art “island” solutions to bridge the gaps between them. This paper discusses the associated challenges and formulates recommendations to explore specific research topics which will contribute to this direction.

⁵ <https://www.etp4hpc.eu/transcontinuum-initiative.html>

References

- [1] D. Balouek-Thomert, E. Gibert Renart, A. R. Zamani, A. Simonet and M. Parashar, "Towards a computing continuum: Enabling edge-to-cloud integration for data-driven workflows," *The International Journal of High Performance Computing Applications*, vol. 33, no. 6, 2019.
- [2] G. Nguyen, S. Dlugolinsky, M. Bobák, V. Tran, A. López García, I. Heredia, P. Malík and L. Hluchý, "Machine Learning and Deep Learning frameworks and libraries for large-scale data mining: a survey," *Artificial Intelligence Review*, vol. 52, p. 77–124, 2019.
- [3] O. Peckham, "15 Years Later, the Green500 Continues Its Push for Energy Efficiency as a First-Order Concern in HPC," *HPCwire*, no. <https://www.hpcwire.com/2021/07/15/15-years-later-the-green500-continues-its-push-for-energy-efficiency-as-a-first-order-concern-in-hpc/>, 17 07 2021.
- [4] E. Masamet, A. Shehabi, N. Lei, S. Smith and J. Koomey, "Recalibrating global data center energy use estimates," *Science*, vol. 367, no. 6481, pp. 984-986, 2020.
- [5] Hyperion Research, [Online]. Available: <https://hyperionresearch.com/wp-content/uploads/2021/07/Hyperion-Research-HPC-Briefing-Slides-During-ISC-June-2021.b.pdf>.
- [6] S. Dewitte, J. P. Cornelis, R. Müller and A. Munteanu, "Artificial Intelligence Revolutionises Weather Forecast, Climate Monitoring and Decadal Prediction," *Remote Sensing*, vol. 13, no. 16, p. 3209, 2021.
- [7] M. Feldman, "HPC in 2020: AI is no longer an experiment," *The Next Platform*, no. <https://www.nextplatform.com/2020/01/09/hpc-in-2020-ai-is-no-longer-an-experiment/>, 09 01 2020.
- [8] Google Cloud, "Helping researchers at CERN to analyze powerful data and uncover the secrets of our universe," [Online]. Available: <https://cloud.google.com/customers/cern>.
- [9] O. Peckham, "Ahead of 'Dojo,' Tesla Reveals Its Massive Precursor Supercomputer," *HPCwire*, no. <https://www.hpcwire.com/2021/06/22/ahead-of-dojo-tesla-reveals-its-massive-precursor-supercomputer/>, 22 06 2021.
- [10] K. Fauvel, D. Balouek-Thomert, D. Melgar, P. Silva, A. Simonet, G. Antoniu, A. Costan, V. Masson, M. Parashar, I. Rodero and A. Termier, "A Distributed Multi-Sensor Machine Learning Approach to Earthquake Early Warning," in *34th AAAI Conference on Artificial Intelligence, Outstanding Paper Award - Special Track for Social Impact*, New York, <https://ojs.aaai.org//index.php/AAAI/article/view/5376>, 2020.
- [11] M. Asch, F. Bodin, M. Beck, T. Moore, M. Taufer, M. Swany and J.-P. Vilotte, "Cybercosm: New Foundations for a Converged Science Data Ecosystem," <https://arxiv.org/abs/2105.10680v3>, 2021.



Authors:

Gabriel Antoniu is a Senior Research Scientist at Inria, Rennes. He is currently the Head of the KerData research team, which focuses on scalable data storage, I/O, processing and in situ visualization on extreme-scale systems (HPC, cloud, edge infrastructures), supporting the convergence of HPC, Big Data analytics and AI.

Patrick Valduriez is a Senior Research Scientist at Inria, France. He is currently the head of the Zenith team that focuses on data science, in particular data management in large-scale distributed and parallel systems and scientific data management.

Hans-Christian Hoppe has worked as a Principal Engineer with Intel, creating the Intel Cluster tools, managing the Intel Visual Computing institute and leading the Intel Exascale Lab at Research Center Jülich. In 2021, he joined Scapos AG to lead the Fortissimo4EuroHPC and FocusCoE projects on advancing the industrial use of HPC.

Jens Krüger is a Scientist at Fraunhofer ITWM. He works on heterogeneous HPC frameworks and extreme-scale data analytics solutions. He focuses on bringing software and technologies to cloud environments (e.g. Amazon AWS) and developing corresponding business models. He also leads the Fraunhofer activities of the STX processor development with the European Processor Initiative (EPI) project family.

The authors would like to thank Rafael Mayo-García from [CIEMAT](#) and Marion Carrier from [CybeleTech](#) for their help in describing relevant use cases for the computing continuum. This work was partially funded by the ACROSS project of the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 955648.

Cite as: G. Antoniu et al., « Towards Integrated Hardware/Software Ecosystems for the Edge-Cloud-HPC Continuum », ETP4HPC White Paper, 2021, doi 10.5281/zenodo.5534464.

DOI: 10.5281/zenodo.5534464

© ETP4HPC 2021