DIGITAL AUTONOMY IN THE COMPUTING CONTINUUM

From Cloud to Edge to IoT for European Data

11 November 2021

EVENT REPORT

AUTHORS

Giovanni Rimassa *(Martel Innovate)* and Federico M. Facca *(Martel Innovate)*

HORIZON CLOUD

European Commission

NEXT GENERATION IoT

ABSTRACT

Europe's ambition for digital autonomy is to gain leadership in the data economy and to empower our economy and society; this builds on the availability of a secure, interoperable, and sustainable computing infrastructure from cloud to edge to IoT. To gain autonomy in cloud and data supply chains, European industry has to master applications and services across the Computing Continuum. It is particularly necessary to ensure distributed resources and services can be smoothly orchestrated and aggregated on demand to support a variety of data-driven applications. On 11 November 2021, a one-day virtual workshop took place, which aimed at gathering researchers, innovators, industrial stakeholders from both supply and demand sides, SMEs/start-ups, policy makers, standardisation experts, regulators, as well as relevant initiatives and projects, to exchange views on specific R&I topics, challenges, and opportunities and converge on priorities to guide future investments under Horizon Europe, Cluster 4, Destination 3, area "From Cloud to Edge to IoT for European Data".

The morning programme was organised to touch base on the wider research and innovation spaces, emerging challenges and directions linked but not limited to topics under the upcoming Horizon Europe Cloud/Edge/IoT calls for 2022, in the Cloud and IoT areas. Experts and attendees have then been able to discuss topics in three different sessions on programming tools for decentralised intelligence and swarm computing, Cognitive Cloud for an AI-enabled Computing Continuum, and the role of open-source software and hardware in a Computing Continuum supporting Europe's digital autonomy.

The afternoon programme addressed primarily the industrial perspective, represented through various industrial initiatives, alliances, and public-private partnerships. Industrial representatives that are engaged in topics linked to the Continuum Computing and Next-Generation IoT were to present position statements and future perspectives, as a means to gather input for deriving recommendations for research and innovation topics under the Cloud/Edge/IoT Calls within the Horizon Europe Work Programme 2023-2024. These calls aim at providing a medium-term, future-looking perspective to complement Europe's initiatives on common data spaces and federated cloud infrastructures under the Digital Europe Programme. An important baseline for the discussion was established by presenting and moving forward from multiple strategic recommendations by high-level groups, such as the "*European Industrial Technology Roadmap for the Next Generation Cloud-Edge Offering*" of the CEO Roundtable "Shaping the Next Generation Cloud Supply for Europe", and the report from the Strategy Forum on "*Next generation IoT and Edge Computing*".

Specific themes emerged throughout the workshop and express a partial but well-shared common understanding of Cloud/Edge/IoT Continuum, and what are the most advisable directions for Europe. A central theme was to actively support the paradigm shift towards distributed and particularly decentralised systems and systems-of-systems connecting together large sets of highly heterogeneous nodes. Given the complexity of the industrial transition and related value chains, openness was then found to be a common thread of the recommended solution approach: this was at its most clear in the session on Open Source, where the idea of openness encompassed open-source software but the views were extended towards open hardware, data spaces, standards, and especially communities and ecosystems, as stressed by the newly announced European Alliance for Industrial Data, Edge and Cloud or the GAIA-X initiative.

One last common thread in the workshop was the recognition of the interplay between horizontal and vertical-driven Research and Innovation (R&I). While core topics related to computing and cloud technologies, are and remain domain-agnostic, future actions should accelerate the transition of established key industrial sectors, i.e., general cloud, AI and edge computing innovation must be complemented and validated by more domain-specific piloting and validation support in areas such as Energy, Automotive, Farming, Buildings and Living or Manufacturing, in agreement with large-scale European strategic and policy objectives. The Digital Europe or the Connecting Europe Facility 2 programmes are very relevant here, but the Cluster 4 under Horizon Europe may address piloting in verticals at an early stage.

TABLE OF CONTENTS

1. INTRODUCTION

On 11 November 2021, the European Commission (EC) organised a virtual event titled "*Digital Autonomy in the Computing Continuum*". The meeting gathered researchers, innovators, industrial stakeholders, SMEs/start-ups, policy makers, standardisation experts, regulators, from relevant initiatives and related projects. The purpose was to discuss the current situation of the Cloud and IoT domains, the main trends for the near future, connections and influences with the Horizon Europe research and innovation programme and relevance to the Digital Europe Programme.

EC representatives opened the event presenting a new strategy 'From Cloud to Edge to IoT for European Data'. The purpose of this strategy is to allow European companies to regain leadership in the Data Economy. Achieving Digital Autonomy is a cornerstone for this endeavour. The progressive convergence we are witnessing between Cloud Computing (CC) and the Internet of Things (IoT) through Edge Computing is resulting in a Computing Continuum that is the base of future digital solutions. Firstly, the multi-faceted concept of Edge Computing became to represent the middle ground between CC data centres and IoT hyper-local networks of sensors and actuators. Then a much more nuanced paradigm emerged and placed computing infrastructure on a spectrum covering from the Cloud Data Centres to Edge Nodes with many intermediate levels. High Performance Computing (HPC), Artificial Intelligence (AI), 5G/6G networks are also part of this Continuum for which hardware and software need to be jointly considered.

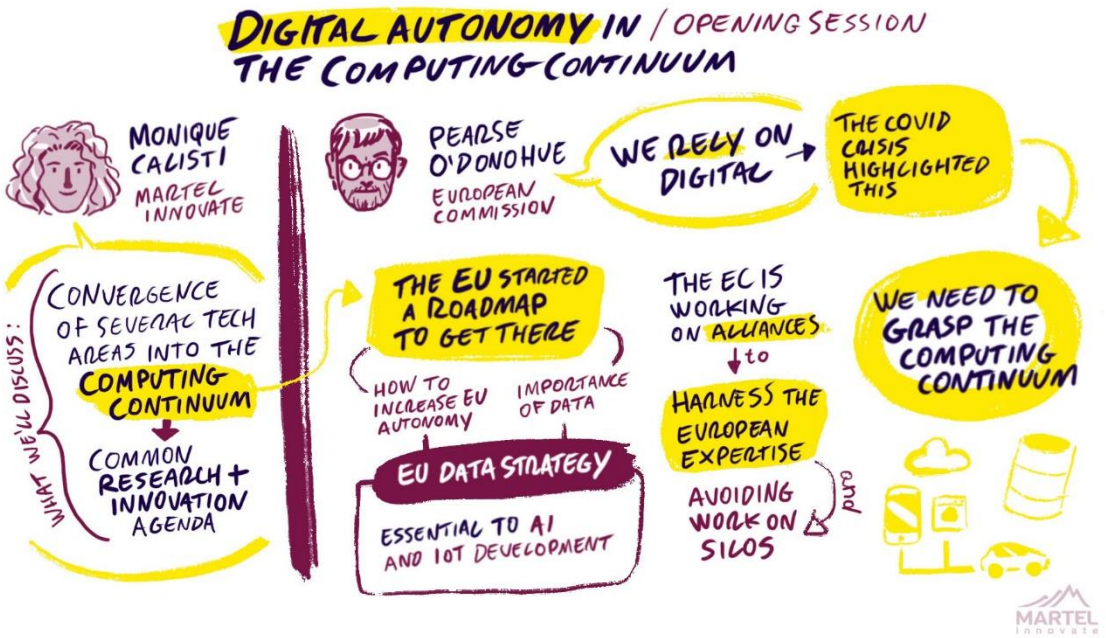

Digital Autonomy has strategic importance for Europe. However, this is not an isolationist trend. The intention is to actively participate in the global digital community by injecting into it a European vision. Digital Autonomy for Europe means the possibility to act independently according to its values. Resilient digital supply chains need to be established while driving digital innovation across all sectors of the economy, being assertive on what are the European values, safeguarding Europe's strategic interests and cooperating with the confidence of knowing that Europe has the means to follow its own path if necessary.

The emergence and continued development of the Data Economy is a key megatrend that affects not only the computing sector but also the economy as a whole. Its combination with the rise of the Computing Continuum is leading to a new landscape where Europe has a role to play. Far from trying to replicate the hyperscalers' approach and business models (i.e., Google, Amazon, or Microsoft), European stakeholders should seek opportunities in the data-driven transformation wave of all economic sectors. The new prospect includes the data volume increasing fivefold with 80 percent of such data produced and processed in highly distributed systems. This game changer shift is redefining the

technological and business context for ICT service provisioning, data management, computing, and network infrastructure for both European and non-European operators.

1.1. European Landscape and Initiatives

The EC has launched a set of coordinated initiatives on multiple levels of funding, stakeholder aggregation, industrial alliances, and partnerships. The core policy document in the area is the *EU Data Strategy*. Major combined efforts currently in focus are the *Important Project of Common European Interest (IPCEI) on Next Generation Cloud* and the *Common European Data Spaces*. To add, a *European Alliance for Industrial Data, Edge and Cloud* has been launched to strengthen the position of EU industry and meet the specific requirements of EU businesses and public administrators processing sensitive categories of data. The Alliance complements the policy and project actions by ensuring stakeholder aggregation, discussion and coordinated work of relevant public and private stakeholders.

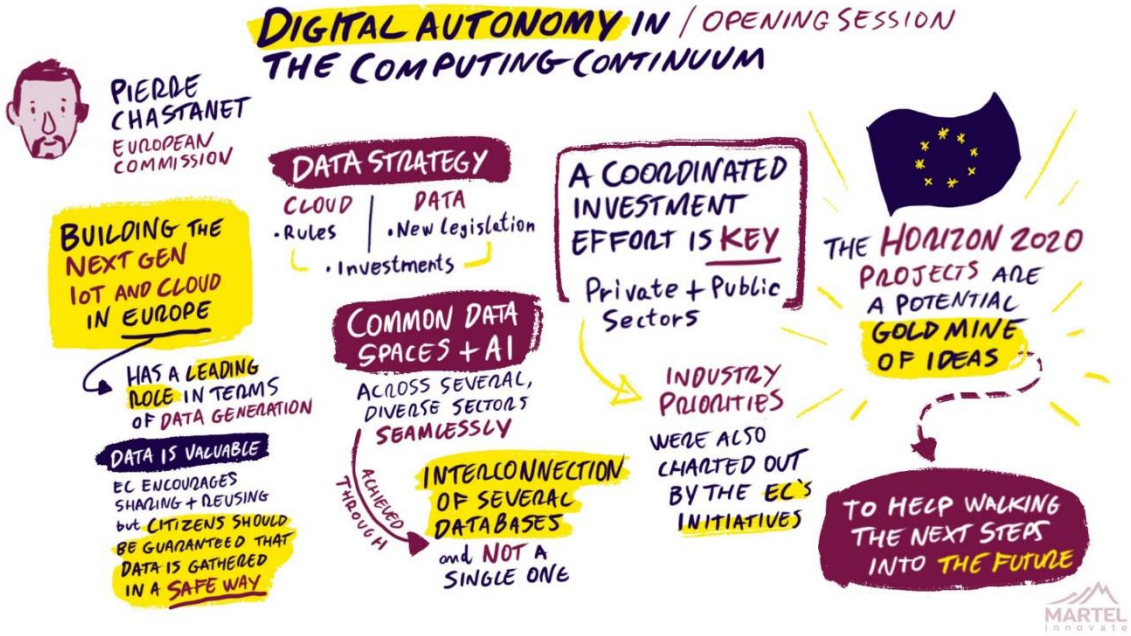

The European investment landscape in next generation Cloud/Edge Computing has *Horizon Europe* as its main pillar. Other funding instruments include the *Digital Europe Programme*, the *Recovery and Resilience Facility* and specific national initiatives. Private investment (through Invest EU or directly from corporations) contributes to a projected total amount of more than 10 billion Euro. A comprehensive set of funding and support instruments has been set up to support the whole technology transfer chain in research, innovation, deployment, market entry, and business growth.

1.2. Workshop Topics and Structure

The morning programme of the event concentrated on three sessions continuing the work on topics depicted by the Horizon Europe Work Programme 2021-2022:

- **Cognitive Cloud**, connected to the *2022-DATA-01-02* topic in Cluster 4, dealing with the adoption and exploitation of Artificial Intelligence techniques to advance automation and dynamic adaptation of resource management in Cloud and Edge systems.
- **Open Source for Digital Autonomy**, related to the *2022-DIGITAL-EMERGING-01-26* topic in Cluster 4 about open-source enablers for Digital Autonomy.
- **Decentralised Intelligence and Swarm Computing**, targeting the *2022-DATA-01-03* topic in Cluster 4 about architectures, programming environments, and tools for distributed multi-device swarms and decentralised intelligence across the Computing Continuum.

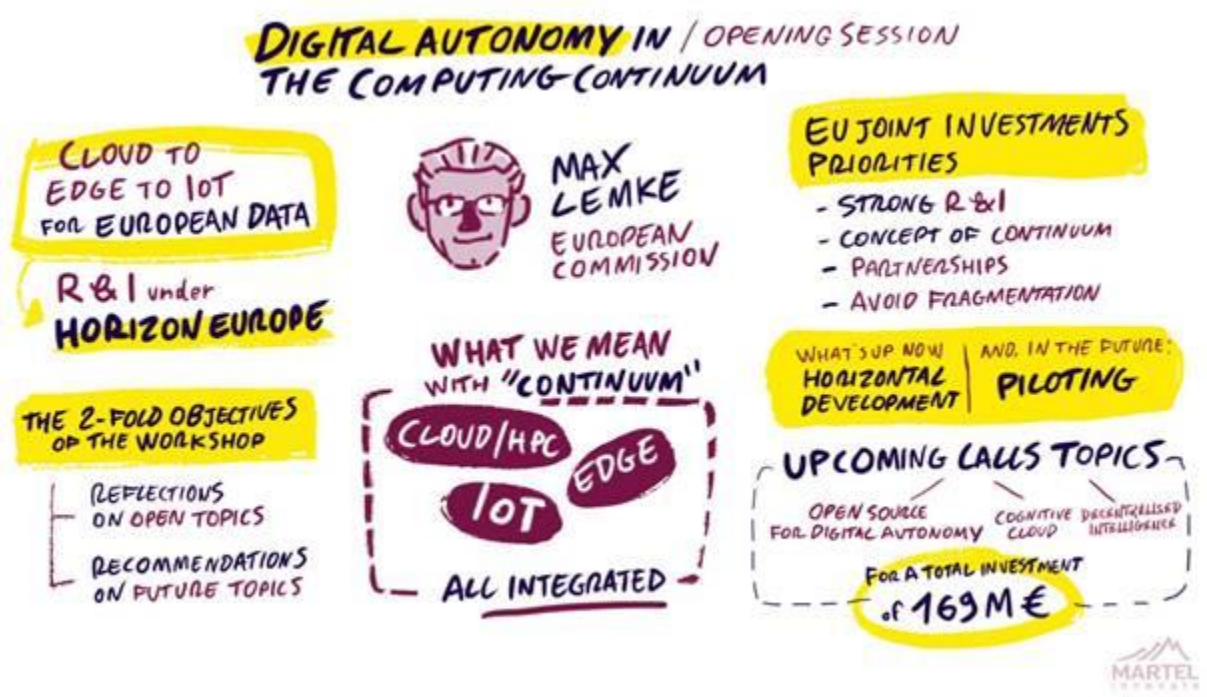

The afternoon programme looked into the future by discussing possible next steps of Horizon Europe, covering the next two years of 2023-2024. This part of the Work Programme is currently under development at the European Commission. The chosen approach for the workshop was for the EC to highlight relevant directions and themes that are currently being discussed and, more importantly, to seek input from representatives of relevant industrial partnerships and alliances through presentations and panels. Finally, input and recommendations were gathered from the expert speakers and attendees. Specific attention was given to those initiatives and organisations that expressed a consolidated view (through a white paper, a strategic roadmap document, or other publication) across their membership on research challenges, technological directions, or investment priorities for the European community and institutions regarding the Computing Continuum as a strategic leadership opportunity for Europe.

Common Themes across the Workshop

Despite the diversity of organisations and stakeholder groups represented by the speakers at the various sessions of the event and the complementary focuses of each session and panel, several themes proved relevant throughout the entire workshop. A well-shared common understanding of the current situation regarding Cloud/Edge/IoT Continuum emerged from the discussions with a set of promising and advisable priorities for Europe.

Europe's contribution to the global economy relies in designing safe and complex physical systems in areas like manufacturing, energy, automotive, farming et al. Workshop participants emphasised distributed and decentralised systems and systems-of-systems as a next step in system design. The advent of the Computing Continuum results in an unavoidable reality of multiple nodes with different computation, storage, network capabilities, belonging to or falling under the jurisdiction and shared governance of multiple principals, each with their own authority, privacy policies, and trust relationships. The traditional, application agnostic, centralised Cloud Computing blueprint is and will remain the dominant topology for at least the next three years, even in IoT applications and systems. On the other hand, defining and providing infrastructure, services and applications in the new complex and heterogeneous environment is a very different task. **Cognitive Cloud is only the beginning.**

Intelligence and swarm computing principles are required throughout the Computing Continuum to handle heterogeneous resources and diverse application requirements.

This heterogeneity raises challenges at multiple levels like interoperability, scalability, and trust. In particular, beyond software and network architecture, the diversity of sensors, devices and edge nodes must be acknowledged. Some speakers highlighted issues at the hardware level and the opportunities for specialised hardware/software codesign. **The need for dynamic adaptability, orchestration and intelligence in a wider sense was underlined in several presentations. Another related theme was the relationship between Cloud/Edge/IoT Continuum, Data and Artificial Intelligence.** This was present well beyond the Cognitive Cloud session and running throughout the event. Research challenges and wider issues of distributed AI, federated learning, decentralised AI model life cycle, data privacy and trust were also presented. Advances in computing and AI algorithms will spur innovation based on embedded and decentralised AI.

Decentralisation and heterogeneity are the context of the emerging Computing Continuum and openness was found to be a common requirement of the recommended solution approach. The complexity in design and the sheer explosive growth of data from devices and objects led to the conclusion that no single European company can master it alone. This recall for openness and working in partnerships with open ecosystems was at its most clear in the session on Open Source.

The idea of openness also included open interfaces and standards relevant for emerging hardware, data spaces, APIs, ABIs and communities. Several associations and alliances of European interest spoke about the vision, roadmap, and recommendations that emerged from engagement and discussions across their membership. On the EC side, emphasis was placed on community and ecosystem openness with, for example, the announcement of the European Alliance for Industrial Data, Edge and Cloud and the mention of its interaction with *GAIA-X*. The overall impression was of a plurality of voices and opinions coming together under a common strategic alignment: a practical echo of European consensus-based mode of operations stemming from the *united in diversity* motto.

One last common trend in the workshop was the recognition of the interplay between horizontal (i.e., industry-agnostic) and vertical (i.e., acknowledging individual industries and partially specialised for them) support. The implementation of core topics such as Computing Continuum or Digital Autonomy are and remain industry- and domain-independent. Horizontal themes must be complemented by specific developments in areas such as Energy or Manufacturing, in agreement with large-scale European strategic and policy objectives. True vertical take-up measures belong to deployment- and business-oriented programmes such as the *Digital Europe Programme* or the *Connecting Europe Facility 2*. Horizon Europe Cloud/Edge/IoT areas have mostly targeted horizontal calls on the 2021-2022 Work Programme. Workshop participants praised for the next Work Programme (2023-2024) to include vertical calls on, e.g., large-scale pilots or industry-specific applications of the general Computing Continuum concepts and technologies, relying on the work done within Horizon Europe and elsewhere in the previous two years.

2. COMPUTING CONTINUUM IN HORIZON EUROPE 2021-2022 WORK PROGRAMME

The three morning sessions were related to topics that are defined in the Work Programme. The goal was not to discuss possible project proposal ideas or approaches, but rather to reflect on the three topics (Cognitive Cloud, Environments and Tools for Decentralised Intelligence at the Edge, Open Source for Cloud-based Services), as key elements of a strategic agenda for Horizon Europe Work Programme 2023-2024.

The *Cognitive Cloud* session elaborated further the existing research challenges covered in the current call and went beyond those challenges to discuss future ones. For example, more intelligent orchestration approaches are needed to effectively decide where data should be processed and in what volumes (especially important when dealing with the resource and energy constrained nodes at the far edge of the Computing Continuum). Moreover, intelligence will play a critical role also in the overall system adaptation, resilience and distributed control capabilities, strategies for data intensive applications and even dynamic resource migration. **All management layers of future continuum-native systems** and applications will have to exhibit **intelligent and cognitive** traits. **Swarm computing** principles should be applied in the whole cloud to edge continuum, to match the expected scale, distribution, and hardware/software heterogeneity. In such a global ecosystem, **international collaboration** will be key in this area as there is a need for joint technological developments and consensus building on future standards for the next generation cloud, edge and IoT platforms.

The *Open Source for Digital Autonomy* session pointed out the relevance of Open Source Software and Hardware for the European strategic objective of achieving digital autonomy. In particular, proper recognition of contributions to Open Source should be fostered at EU-level. Participants pointed out that Open Source communities have managed to develop a complete open stack from the application level to the kernel but the layers below are still closed. The presenters invited the Commission to take adequate measures to expand the open stack to the lower levels. The session reinforced the general consensus that embracing Open Source is a key strategic milestone for Europe to achieve Digital Autonomy. This does not mean that proprietary or closed-source models are to be completely excluded, nor that Open Source on its own guarantees digital autonomy. However, the general and dominant approach for Europe should be to fully engage with open-source projects, processes, and governance to use, to contribute, to lead and to steer global communities when relevant. Hardware architectures such as *RISC-V*, as well as low-level, below-operating-system software support were also discussed, in connection with strong input from industry, where a widespread awareness and adoption of Open Source marks a significant evolution of the past years.

The positive impact of open-source activity on the overall economic development and prosperity was also covered in the session following the results of the study on the "impact of open-source software and hardware on technological independence, competitiveness and innovation in the EU economy". The workshop participants also pointed out the need for adequate EU recognition to contributors to Open Source and the need to use OS assets to regain competitiveness in the computing sector.

The *Decentralised Intelligence and Swarm Computing* session provided evidence that highly distributed swarm systems and applications contribute to a long-term vision of hyper-connected objects and systems. However, swarm concepts provide basic tools to deal with evolving data volumes at the edge, i.e., close to the place where data is created. They allow to dynamically group objects, aggregate data across sub-groups of things and be aligned on underlying security and privacy rules, essential for decision making and autonomous control. First reported applications, like in waste management and logistics, are not very far away, so the EU Data Strategy should consider and include them, as they can be very effective in managing data, and providing dynamic behaviour as well as *self-** autonomic properties. Presentations and discussions explored the main research challenges and directions in this topic, from conceptual characterisation of, e.g., elasticity and future cyber-physical systems, to how to unleash AI at the edge through hardware/software codesign, to many examples and use cases that highlighted the need for virtualisation built on Computing Continuum solid horizontal architectures.

Multiple domains and industries were covered, such as smart agriculture, the automotive industry, the energy sector, logistics, and manufacturing.

2.1. Cognitive Cloud: AI-Enabled Computing Continuum

The session "**Cognitive Cloud: AI-enabled Computing Continuum**" was led by **Maria Tsakali** (Programme Officer, Unit E2 Cloud & Software, European Commission) and saw the participation of **Lutz Schubert** (Head of Research at Ulm University), **Domenico Siracusa** (Head of the Robust and Secure Distributed Computing Research Unit at FBK), **Danilo Ardagna** (Associate Professor at Politecnico di Milano), and **Cristina Aranda** (Cofounder of Big Onion).

The aim of the session was to elaborate on the different perspectives that are developed around the "Cognitive Cloud" concept. Maria Tsakali introduced the session by setting the scene from the EC perspective. Since FP7, the European Commission has constantly supported cloud and edge computing research in Europe; the overall objective has been to support the global competitiveness of Europe and provide the European cloud industry, especially European SMEs, with an alternative model to cloud hyperscalers based on interoperability, portability, and federation of resources.

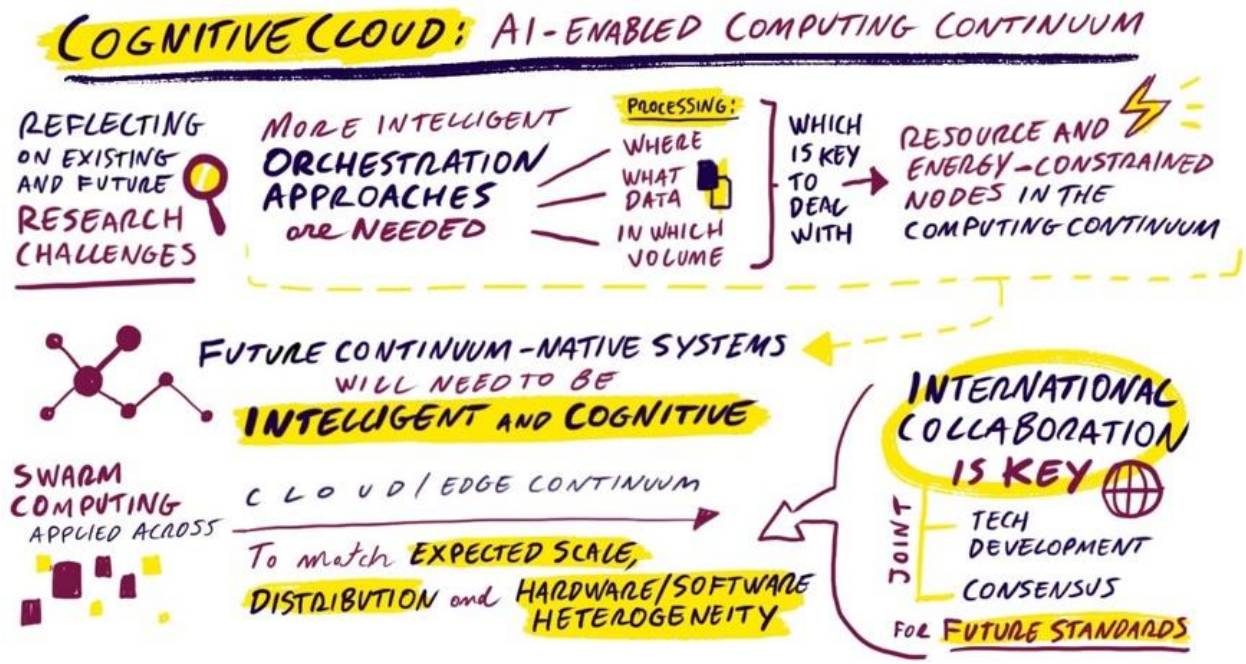

The resulting European ecosystem of actors is a key element for the ongoing digital transformation, including developed solutions for edge computing, multi cloud, and federated clouds that are going to support the Digital Europe Programme. The need to accelerate cloud adoption was addressed by introducing the European Alliance for Industrial Data, Edge and Cloud; the need to enhance cloud portability considering legislation such as GDPR was addressed by the Cloud Rule book; the need to move from single research areas to the harmonised convergence of Cloud, Edge, IoT and networks was already addressed in the Horizon Europe Work Programme for 2021 and 2022, but will remain an important topic to address in the subsequent 2023-2024 Work Programme.

The "Cognitive Cloud" topic in the current call of the Work Programme 2021-2022 was defined thanks to the different inputs from independent experts and the cloud research community (including NESSI, the Future Cloud research cluster and the Horizon Cloud initiative). The topic envisions that AI will transform current clouds into cognitive clouds, thus leading to a higher automation on different aspects of the Computing Continuum by dynamically adapting applications and infrastructure to data and computing variability demands. Following Ms Tsakali introduction, the invited speakers provided a view on current and future challenges to realise a "Cognitive Cloud".

Toward the Smart Edge

Lutz Schubert presented a view on Cognitive Cloud challenges developed in collaboration with Keith Jeffery. Mr Schubert highlighted that Cloud Computing originated from the need to provide resources on demand, anywhere and at any time: some of such goals have yet to be fulfilled, and Edge Computing ultimately derived from this. He also underlined that data is still both the biggest challenge for Cloud and a huge opportunity for our economies. The challenge lies in more than just data size or processing latency: data has different meaning to different people, not just source-specific, but also code-specific, shaped by the processing hardware and bound to a location. According to Mr Schubert, moving the data from the Edge to the Cloud is not enough, data needs to adapt to code and vice versa. **Data and code must be distributed and heterogenous, intelligently managed by computing environments to provide reliability, performance, security, adaptability, and flexibility**.

Mr Schubert concluded his intervention highlighting the key research challenges that need to be tackled to answer these needs: handle locality of data and code; provide adaptability at all levels; enable distributed control; understand what is happening and why.

AI Applications within the Cloud-to-Edge Continuum

Domenico Siracusa complemented Mr Schubert presentation by providing his perspective on future challenges starting from the results of *Decenter* research project, focused on **transitioning from a pure Cloud to a Cloud-to-Edge architecture**, and how to tackle the challenge linked to this transition. The research goal of the project was to **enable management and deployment of AI applications in the continuum following a cloud-native approach**. The team led by Mr Siracusa first investigated how to support AI application developers by adopting a microservice approach, i.e., making AI models smaller, supporting their partitioning, and packaging them in containers ready to be used and combined with other services. The research led to an evolutionary approach of AI lifecycle (AIOps) embracing the DevOps model and applying it to AI.

Based on the lessons learnt in the project, Mr Siracusa identified some key challenges for the future in dynamic programming of environments and adaptation of applications, ability to express application intent, and using this information to program the infrastructure, managing not only cloud and edge servers, but also microcontrollers, thus looking into **Swarm Computing and intelligence**.

Lowering AI Adoption Barriers and Cost

Mr Danilo Ardagna presented the outcomes of the AI-SPRINT project. His presentation analysed how AI is “hungry” for Engineering, i.e., the competences needed to master it, for Cloud, i.e., the elastic capacities of computing, and for Silicon, i.e., the hardware specific capabilities to support AI. Mr Ardagna also highlighted how the hardest part is often not AI per se (AI code account for just 5% of an AI system), but automating the whole system, that is, **enabling intelligent end-to-end system management**.

In line with this understanding, AI-SPRINT focused on supporting AI developers with a programming framework that simplifies their work. The programming framework considers different aspects. On the one side, model engineering in the Computing Continuum cannot be atomic, but needs to be adapted to fit with edge device capacity. AI-SPRINT is exploring solutions to guarantee that AI applications can be deployed fulfilling requirements and constraints in terms of capacity available, energy, costs, and performance of the application. On the other side, neural networks, common building blocks in existing applications, are not designed for the continuum. To tackle this aspect, AI-SPRINT focused on new neural network models that can **support migration across Computing Continuum** layers and partitioning of the AI models as required by resources or application constraints. Other key project outcomes related to AI application performance prediction based on machine learning techniques, resource demand reduction through hyperparameters tuning, and federated learning to increase privacy and security in the training phase, **therefore highlighting the need for Clouds to evolve towards the “Cognitive Cloud” concept**.

A Broader Look at Cognitive Cloud Challenges

The last presenter, Cristina Aranda, covered challenges beyond the technological realm, stating that to build a data-driven society, education is key: technology evolves very fast, and this calls for new strategies to upskill people working in technology arena. Ms Aranda highlighted that the changing role of people was well captured by a recent technology trend analysis from Gartner: centrality of the person, independence from the workplace, flexibility in delivery.

Beyond that, data is a complex and multifaceted asset, requiring a strategic horizon of 5 years, 10 years, or even beyond. In particular, the vast majority of data today is unstructured, and this makes it very difficult to derive information from it. Approaches such as Natural Language Processing (NLP) can help extract knowledge from unstructured data, but need to cope with regulations such as GDPR, prompting for audit mechanisms to ensure that NLP or AI solutions are not biased. According to Ms Aranda, **cognitive platforms have the potential to improve our society**, but it's key that they find a way to manage data complying with regulations and achieving social, not just economic, impact, e.g., being at the service of great challenges such as sustainability and social inclusion. Ms Aranda stated that to achieve such an impact it is important to work as an ecosystem, but that requires governance and an ethical code, which is not yet in place for AI. Three key challenges in this context are: putting citizen experience at the centre of any application; creating knowledge in a controlled and secure infrastructure; defining governance strategy encompassing IT and ethics.

Finally, Ms Tsakali closed the session highlighting how the different speakers provided complementary views on the challenges to deliver Cognitive Clouds and that research on Cognitive Cloud is still in its infancy, thus leaving room for further research in the AI-enabled Cloud to Edge Continuum in the future Horizon Europe Work Programmes.

2.2. Open Source for Digital Autonomy

The session "**Open Source for Digital Autonomy"** was led by **Luis Busquets Pérez** (Programme Officer, Unit E2 Cloud & Software, European Commission) and saw the participation of **Miguel Díez Blanco** (Project Lead – Open Source Programme Office, European Commission), **John Davis** (Director of the Laboratory for Open Computer Architecture, Barcelona Supercomputing Center), **Leire Orue-Echevarria** (Cloud Continuum Project Director, TECNALIA), **Mirko Boehm** (Head of Delivery, MBition/Mercedes-Benz), **Knut Blind** (Coordinator of the Business Unit Innovation & Regulation, Fraunhofer ISI) and **Josef Urban** (NESSI Chairman & Technology Lead in Corporate CTO, Nokia Bell Labs).

The session focused on the role of Open Source to support European Digital Autonomy. Luis Busquets Pérez introduced the session by providing the EC view on the role of Open Source to build a European digital society, and how the EC sees value in Open Source to provide better and more transparent services to citizens. Clearly the focus in E2 Unit is related to open-source solutions linked to Cloud and Software. Open Source is today very relevant in many sectors, witnessing solutions that from the kernel up to the end user application are totally open-source. In the last years, Open Source spread even beyond software to Hardware and Data. While in the early 2000s the EC's view on Open Source was

related to its role to reduce vendor lock-in and promote its adoption via policies, lately this view evolved and now Open Source is perceived as common good on top of which society can prosper. In line with this evolution, the EC sees a major role of Open Source in support of Digital Autonomy to support EU strategic autonomy and European values.

In line with this ambition, the EC supported Open Source through several initiatives. This support will go on also in the next calls, in particular through the call "*HORIZON-CL4-2022-DIGITAL-EMERGING-01-26: Open source for cloud-based services*" focussing on the deployment of developments of the European processor initiative in data centres. Beyond that, the EC launched an Open Source Strategy to support internal IT needs.

The objectives of the session were to validate the importance of Open Source for Digital Autonomy, to identify fields in current computing architectures that require open alternatives, and to define support priorities for open-source development.

Introducing Open Source Governance in the European Commission

Miguel Díez Blanco presented how the EC strategy evolved in relation to Open Source in the last years. While early strategies where mostly focused on specific Directorates, the last strategy released in October 2020, 1) provides a strategy for the overall European Commission; 2) links the strategy with policy goals; 3) changes the work culture of the EC introducing Open Source best practises. The document includes several actions to implement Open Source inside the European Commission. Actions promote the establishment of a Governance of Open Source, ranging from the sharing of code developed internally and increasing and structuring the presence in the Open Source arena, to the reduction of paper work related to the release of code as Open Source, the support for the experimentation of open-source products and their promotion in the European Commission, the engagement of open-source communities, and studies related to Open Source.

RISC-V Open Hardware to innovate EU microelectronics industry

John Davis presented a background on the role of Open Source computing in HPC. The slowdown of Moore's Law is pushing HPC to develop specialised solutions combining hardware and software design. Today x86 architecture leads the Open Source landscape also in HPC. In the future, a leading role is foreseen for RISC-V architecture that provides an open-source hardware alternative to ARM and x86 architectures including Edge and Cloud. Thanks to its characteristics RISC-V may be able to unify and enable the growth of the microelectronics industry in Europe. The delivery of RISC-V System on a Chip requires the growth of an ecosystem and there are a lot of gaps in the picture to be filled by Open Source, to enable existing operating systems to be compiled and run on this new architecture.

Multi disciplinarity and cross sectorial open solutions

Leire Orue-Echevarria started her talk with an overview of the role of Open Source in TECNALIA to deliver value to customers and to European research projects. Open Source usage impacted TECNALIA by increasing the development speed and created an open innovation culture where results of projects are shared as Open Source. Open Source led also to develop new business models to transfer such results to the market. In this shift towards an Open Source strategy, the biggest challenge is the overall governance.

In setting up cloud solutions, TECNALIA faced a limited set of choices from the processor hardware to the virtualisation layer, and this created dependency on these solutions. Achieving Digital Autonomy requires to break such limitations and support the development of interdisciplinary and cross domain solutions: digital solutions are combinations of multiple technologies spanning from Cloud Computing to AI, Security and IoT, and need to allow multiple verticals. Ms Orue-Echevarria identified the following key challenges to deliver multi disciplinarity and cross sectorial open-source solutions: the definitions of open standards, making room for interdisciplinarity in research work programmes, and developments that need to adhere to European values as promoted by the policy framework.

Openness guarantees access, but only contributes to autonomy

Mirko Boehm elaborated on the relationship between autonomy and Open Source. While the licensing of open-source software provides for all needs for being autonomous in managing decisions related to

the software, this is not enough to deliver a service and build a product. There are several conflicts that need to be solved (these conflicts are generally named as “principal-agent problem”) and that require high trust among the parties: e.g., customers need to trust the provider in running the service as they state, including their privacy-preserving management of it; when managing data, there is the need to trust that data are not shared with third parties; a computer that exposes nowadays remote management interfaces, needs to be trusted to execute code as expected and not doing other actions. While open approaches promote openness and transparency, and hence enable trust and collaboration, technology is not enough. Closing the remaining gaps requires agreements, audits, and regulations. Mr Boehm highlighted that Digital Autonomy also requires access to chips and not only to software: while chip design is becoming open, contrary to the software, achieving autonomy to the silicon is limited by several factors, such as access to expertise and production capacity. Finally, Mr Boehm remarked that Open Source is not a tool for autonomy as such, but it’s more a tool to achieve standardisation in markets.

Open Source contributions as a booster for EU economy

As leading author of a study commissioned by EC on the Impact of Open Source in European Economy, Knut Blind highlighted that the impact of Open Source for the European economy ranges from 60 to 100 billion Euro, thus investing in Open Source is an important strategy for EU to reinforce its economy and global positioning. Accordingly, Mr Blind surveyed key recommendations to further support Open Source in Europe: Knowledge Creation, i.e., increasing the level of public R&D funding on strategic open-source initiatives and easing access to Open Source related funding for SMEs; Knowledge Diffusion and Networking, i.e., incentivise the release of project results as Open Source; support the establishment of Open Source networks in the EU; Entrepreneurial Activities, i.e., providing education and fostering a culture of Open Source for EU start-ups, promoting partnerships between Open Source foundations and large industries; Human Capital Development, i.e., developing Open Source skills through Higher Education institutions, supporting certification schemes, growing the diversity in the Open Source community contributors. Mr Blind concluded remarking the importance of coordinated and comprehensive public policies to support Open Source growth in Europe for the benefit of the EU economy.

Open Source and proprietary software are not mutually exclusive

Josef Urban provided the point of view of NESSI as it was presented in a position paper released in 2019. There are many success stories of Open Source, such as Linux, Eclipse, and Kubernetes. The importance of Open Source is globally recognised and today more than fifty million developers are contributing to Open Source and this number will grow up to one hundred million in 2030. To govern Open Source, it is important to understand multiple aspects. E.g., while there are no direct costs (license), the indirect costs are quite high (expertise); Open Source facilitates user engagement; successful communities require good governance and alignment with industry; Open Source is not secure by default, it largely depends on the good management of the project; the proliferation of licenses requires due diligence on them before adoption. Mr Urban concluded his talk highlighting that to adopt Open Source, a proper strategy is required to be able to reap positive effects from Open Source; Open Source is a tool to support Digital Autonomy, but not the only one, also proprietary software can play a key role. Accordingly, while supporting Open Source is important, proprietary software should not be penalised in access to European funding.

Finally, Mr Busquets Pérez closed the session discussing with speakers the importance of incentives, such as Digital Awards, to support open-source development; the dualism between liberal licenses and copy-left licenses and the relevance of this dualism for the industry; business models allowing companies to succeed in leveraging Open Source; how Open Source adds value to the customers by enabling innovation; the difficulty of transferring some Open Source concepts from software to hardware.

2.3. Decentralised Intelligence and Swarm Computing

The session “**Decentralised Intelligence and Swarm Computing**” was led by **Rolf Riemenschneider** (Head of Sector IoT, Unit E4 Internet of Things, European Commission) and saw the participation of **Schahram Dustdar** (Full Professor at Technical University Vienna), **Marc Duranton** (HiPEAC Chair of Vision Group, Research Fellow at CEA), **Fabrizio Del Maffeo** (CEO of Axelera AI), **Purshottam Purswani** (CTO Business Platforms & Solutions at ATOS), **Jerker Delsing** (Professor at Luleå

University of Technology), **Michael ten Hompel** (Professor at Technical University Dortmund), **Alexandra Haas** (Head of IoT Solutions at Swisscom), **Haydn Thompson** (Director at THHINK Wireless Technologies), and **Thomas Berndorfer** (Member of the Executive Board at TTTech Industrial Automation AG).

Rolf Riemenschneider laid out the background, aims and scope of the session in his opening presentation, mentioning the overall trend towards more and more decentralised topologies such as mesh networks, which goes beyond the Cloud/Edge convergence to encompass the full breadth of what is currently understood as IoT, and local data storage and processing. This process is driven by several technological factors, among which security and privacy, decreasing cost of hardware and sensors, and the requirements of real-time decision making, safety, and automation. These trends and vision can be recognised in many sectors and use cases, such as smart living applications, autonomous vehicles, or carbon-aware energy data mesh to support the twin transition.

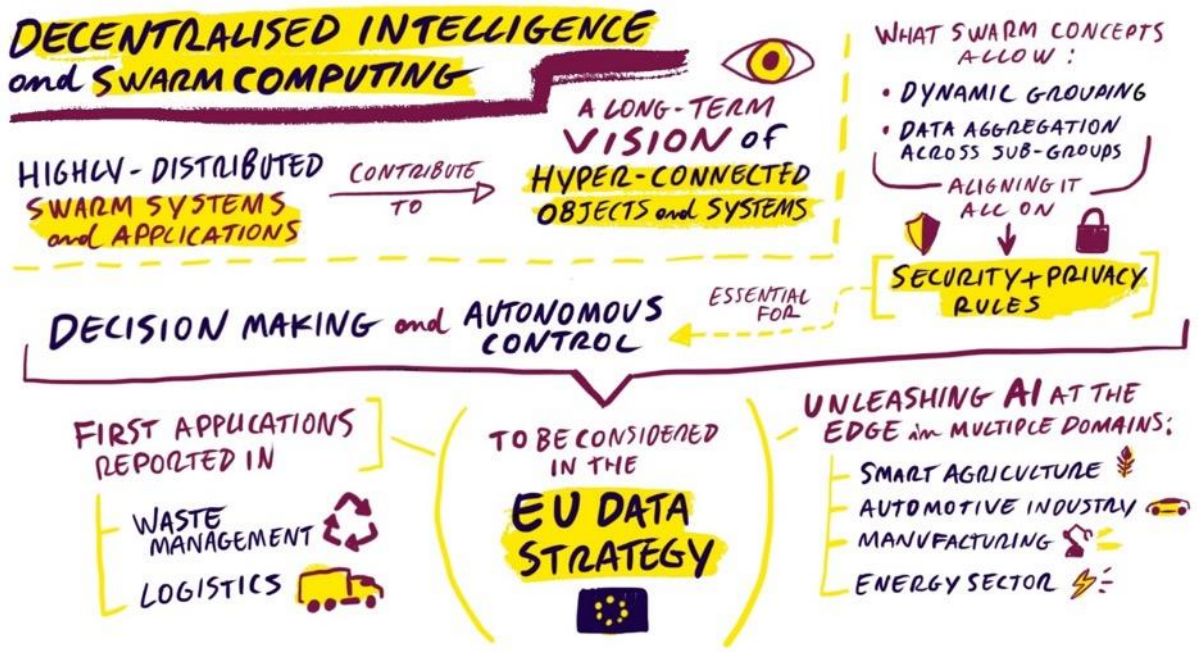

The Cloud/Edge/IoT continuum configures itself as a paradigm shift for distributed computing, raising great challenges and opportunities. This significance is reflected in the Horizon Europe Work Programme for 2021-2022 with more relevant pillars and topics, some already issued and closed, others published in the current WP, and some others that will be defined as part of the subsequent years. The specific topic of the session is the "*Programming tools for decentralised intelligence and swarms*" call, covering among other things agile secure architectures for smart nodes and programming environments for smart edge-connected nodes, with the overarching goal to reinforce Europe's position in next generation smart systems, building on European strengths in embedded devices and wireless networks.

Edge Intelligence as Fabric for the Computing Continuum

Schahram Dustdar pointed to a common trait of many smart systems: they are a composite ensemble of people, services, and things that are arranged in ways that exhibit elasticity. Examples can be found in several domains such as smart homes, smart energy networks, or smart transport; such systems require near-real time local environment data access and resources near the data producer or consumer nodes. This situatedness and context-awareness triggers the understanding of the computing infrastructure as a continuum from a local IoT-centric area throughout edge and fog computing nodes until core Cloud data centres.

An important research challenge and agenda item is the identification of the defining traits of edge computing systems. Mr Dustdar advanced the proposal of elastic diffusion as one such concept. In particular, the idea is to go beyond the identification of elasticity with scalability to consider resource, quality, and cost/benefit as three dimensions of elasticity. Within such a three-dimensional space, the evolution of a continuum-native system can be plotted against time, so that at every instant the resource/quality/cost triad is balanced in a specific, dynamic way. A second research topic is the

continuous and automated adaptation of service deployment across the Computing Continuum, which can be informed by the concept of diffusion, adopting a chemistry-inspired metaphor to signify the intelligent bias that can be introduced in resource management policies. Lastly, Mr Dustdar highlighted some elements of a research roadmap tackling the relationship between Edge Computing and AI, distinguishing in particular between *AI for edge*, where intelligence works as an infrastructure enabler (similarly to Cognitive Cloud) and *AI on edge*, where the goal is to support the full life cycle of application-valuable AI models on an Edge Computing distributed infrastructure.

Digital Twins and Cyber-Physical Systems Come Together

Marc Duranton outlined the importance of digital twins for decentralised intelligence, pointing to some highlights from the 2021 vision document from the HiPEAC Network of Excellence. Applications dealing with decentralised data sources will today orchestrate multiple services and specific interfaces towards the physical environment. Intelligence requirements turn the system into a software agent, sensing data from the physical environment and able in turn to act on it, or into a cyber-physical system (CPS).

On the one hand, some services implement application features; on the other hand, other services will abstract away from elements of the physical environment. These other services are equivalent to partial digital twins, so that the agent/environment interface of the base CPS also occurs between services of the first group (working as software agents in a virtualised environment) and the ones of the second group (providing such a virtualised environment).

Mr Duranton likened such an evolved system model to a concept named by some hyperscalers and computing manufacturers (Microsoft, nVidia) as metaverse or omniverse. Europe should be aware of these trends and be fast in entering these emerging market segments to ensure their compatibility with European values and regulations. Augmenting CPS ideas with service-orientation, strong decentralisation, and pervasive partial digital twins results in *Cognitive* (due to edge AI) *Cyber Predictive* (due to targeted digital twins simulating system evolution) *Physical Systems of Systems* (due to multi-level component aggregation). This dynamic decentralisation poses multiple research and innovation challenges, from interoperability to composability, from scalability to security. The orchestration level itself could enter the next generation with some emerging approaches, such as linguistic programming and very high-level integration with technology-independent APIs and smart service contracts.

Hardware/Software Synergy for AI at the Edge

Fabrizio Del Maffeo concentrated his talk on the specific challenges raised by deploying Artificial Intelligence at the edge, particularly neural network technology, and on how a very promising innovation path blurs the boundaries between hardware and software to reach success through the best blend of both. Effective computing at the edge must deal with and overcome several constraints on, e.g., processing power, storage capabilities, and energy consumption. Focussing on neural networks, a series of techniques have been already applied to adapt a trained network to run more efficiently on resource-constrained nodes (value quantisation and network compression among others). While effective, such techniques are still met with difficult challenges due to the extreme heterogeneity of IoT and far edge hardware: these careful transformations and optimisations must be tuned again and again on each new different hardware node: initiatives like TinyML try to collaboratively tackle such issues.

A very promising trend is to introduce changes in the hardware architecture to suit the specific needs and constraints of AI at the edge, such as memory banks that can also do on-chip data processing, or more radical approaches such as spiking neural networks. Mr Del Maffeo concluded his intervention by underlining how AI at the edge has a specific context that calls for joint application of multiple hardware and software techniques that, together, can provide a true quantum leap to neural network applications.

The Road towards Mainstream Swarm Intelligence

Purshottam Purswani elaborated on the concepts, architecture, and roadmap for Swarm Intelligence. This nature-inspired idea seeks to achieve intelligent emergent behaviour through the interaction of large numbers of relatively simple agents, based on proximity rules and self-organisation. While the research intuition and preliminary work date back from a couple of decades, the new IoT and Computing Continuum context provides new insights and opportunities for such an approach. Self-healing, self-configuration, and self-optimisation are distinguishing properties of intelligent swarms and can be found as well in artificial swarm intelligence systems, such as drone fleets, autonomous vehicles, or networks in advanced industrial applications.

There are both demand requirements (dynamic architectures, security/privacy issues, energy efficiency) and technological enablers (edge computing nodes, decentralised ledgers, advances in wireless networks), but to really ground swarm computing in something practical, Mr Purswani described use cases in the domain of autonomous drones and smart cities. From an architectural point of view, important building blocks include specific algorithms, security frameworks, and decentralised AI; the Edge Computing support bridges the gap from individual edge devices to a fully functional, managed edge cluster relying on low-latency local networking and connectivity. Lastly, a view towards the Computing Continuum was sketched with the help of a series of vision points and research challenges, such as learning model heterogeneity and swarm twins.

Cyber-physical Systems in Industrial Production

Jerker Delsing provided a compelling characterisation of industrial plants as systems of cyber-physical systems, in a similar manner to the model previously described by Marc Duranton. The core point is that in a production plant the material flow and the information flow are distinct but both necessary for the description, monitoring, and operation of the plant as an integrated whole. Starting from an IT integration approach, the recognition of the complexity and number of interdependent dimensions involved in the system of systems, that a modern industrial plant is, leads to considering approaches inspired by biology. In particular, interoperability between IoT systems and traditional Operation Technology (OT) systems is always challenging and has been the subject of EC projects for two decades.

Results from some of these projects have now grown and turned into open-source tools such as Eclipse Arrowhead; after two years, the impact has been very significant, with some metrics, such as cost and lead time reduction, even going beyond the originally planned savings in the domain of industrial plant reconfiguration and management. Mr Delsing also presented a second case study in Smart Home and Smart Energy, where the migration from a proprietary IoT platform to a fully service-oriented and open-source ecosystem resulted in extreme reduction (from 2-3 months to 2-3 days) in the integration time for legacy devices, which can be a major source of effort and delays in these projects.

Decentralised Intelligence in Logistics

Michael ten Hompel opened the second part of the session, which concentrated on applications of decentralised intelligence, with a presentation in the domain of logistics. An upcoming wave of simulation-based AI was pointed at as a first potential breakthrough: autonomous swarm-capable vehicles in the real world were operated by neural networks that were trained in a simulated reality (again a targeted application of digital twin ideas), thus functioning as physical avatars of virtual agents. A second critical issue is temporal coherence, which is decisive for the digital transformation of logistics and other industrial domains and appears in several use cases, such as immersive technologies, Industry 4.0 processes, and coordinated autonomous robots.

Mr ten Hompel also remarked upon the drift from tangible towards intangible assets in logistics and other industries, where competitive advantage and monetisation opportunities now lie much more in know-how and information asymmetry, rather than in freight fleet or warehouse capacity. The upcoming landscape and ecosystem for modern logistics is reaching such a level of integration, jointly relying on the three pillars of IoT, blockchain, and physical fleets, that all major Computing Continuum trends and initiatives, such as Gaia-X or International Data Spaces, become absolute enablers for the industry.

The presentation of a success story in autonomous waste container management, stemming from a cooperation with Rhenus Logistics and Commerzbank, prompted the introduction of the newly created Open Logistics Foundation, an industry-specific example of strong commitment to Open Source and community-wide cooperation, part of the common thread that was woven throughout the workshop.

An Operator and their Clients towards the Computing Continuum

Alexandra Haas continued on the application track by reporting, from the perspective of a telecommunication operator as Swisscom, on what concrete activities are taking place in the industry and how far is the implementation of IoT solutions using Edge Computing, 5G, and Cloud. The main priorities are interoperability and reduction of the complexity of adopting and operating these solutions, particularly when targeting SMEs, which represent the 99% of the Swiss economy, to empower them to compete internationally. Two concrete examples were brought: the first is a company the produces cables and started a transparency and visibility project dealing with their multiple production sites. In this situation, challenges arose such as setting up unified IT operations that include all the production

sites, or provide end-to-end security spanning the connectivity from the production site to the chosen Cloud Core infrastructure. Additionally, to compete internationally some kind of global reach is necessary, as well as the capability to adapt to specific regulatory and network constraints in the different countries of operation.

The second example went forward towards the upcoming Computing Continuum by deploying edge nodes providing CPU and GPU computational resources and a hyperscaler (Microsoft) environment that, in connection with a 5G antenna, was located on a train wagon to be able to follow the field workers from a railway maintenance company. While obstacles such as sensor hardware availability and lack of vendor choice remain, the deployed solution already exhibits traits of future continuum-native and swarm-intelligent systems. Ms Haas concluded by stressing the importance of targeting and empowering SMEs, serving their concrete needs and shorter business time frames to bridge Research and Innovation Action projects with the other, more business-oriented programmes and instruments.

The Importance of Trust for Decentralised Edge Intelligence

Haydn Thompson presented a varied array of applications from his company THHINK Group, an SME that performs research and development and provides consultancy in ICT technologies for the aerospace, automotive, marine, rail transport, energy, and health sectors: from Smart Agriculture to wildland fire detection, to an autonomous shark alerting system. Two general trends that have been identified across their projects are a strongly increased amount of data processing at the source (on-device or in far edge nodes) and the fast growth of edge AI as a common and necessary system feature.

As companies like THHINK add value to their customers by hiding much of the technical complexity intrinsic in demanding Computing Continuum applications, the issue of trust is a critical business enabler: customers need to trust what they see on a data monitoring dashboard without having to be aware of the countless technical details needed to aggregate data from the sources into meaningful views. The shark detection use case, where a system misdetection can result in injury or death, was particularly effective in underlining the multiple relevant aspects of trust: from the specific trustworthiness profile of the employed AI to the GDPR-compliant storage and management of personal data, to the overall perception of the solution as deserving users' trust, all with the associated risks of liability.

A second interesting point made by Mr Thompson was more directly related to networking and connectivity: operating in remote areas, it can be possible to rely on satellite constellations for wireless coverage, but due to their non-stationary orbit the resulting connectivity will be intermittent. This is yet another driver for local data sharing and decentralised intelligence: in such conditions of ubiquitous but intermittent connectivity it is not possible to close a prompt enough remote controlremote-control loop, so that a certain degree of autonomous operation is not just desirable but absolutely necessary. Data cannot simply be uploaded through intermittent links for performance and cost reason, so that a major trade-off in future system design will be task allocation between edge/local and cloud nodes.

Decentralised Autonomy at the Application Level

The session was closed by the presentation from Thomas Berndorfer, providing multiple application examples of how decentralised solutions are being chosen and deployed at TTTech in various industrial domains. The first example was the control system for a wind turbine, and the second one dealt with off-road special vehicles (e.g., snowcats); both cases are industrial safety-critical systems where strict constraints must be satisfied, including reliability and redundancy for control systems that oversee high-power hydraulic actuators. Adopting decentralised autonomy, at least partially and even in the presence of a central monitoring system, becomes therefore an effective and convenient design approach.

A third use case presented by Mr Berndorfer revolves around industrial asset management, that is the conceptually centralised monitoring of multiple, geographically distributed industrial sites. This is a well-known and widely adopted approach of central, cloud-based overseeing of decentralised, autonomous nodes (production plants, distribution centres, or others). The TTTech platform for industrial edge computing, called Nerve, supports customers in a varied set of use cases with enabling features such as networking and multi-tenancy: a shown example targeted hydrogen energy storage, again with a local edge node and a cloud-based management system.

Lastly, Mr Berndorfer highlighted once again the link between decentralised intelligence and interoperable open standards, pointing out how adopted open protocols and architectures represent key enablers and necessary steps toward autonomous and safe industrial systems.

3. SHAPING HORIZON EUROPE 2023-2024 WORK PROGRAMME

The afternoon section looked for inspiration and directions for the 2023-2024 Horizon Europe Work Programme (WP) from a panel of high-level experts in data, Cloud, electronic components and systems, as well as stakeholders from relevant verticals and associations to exchange views on a strategic European vision for the Computing Continuum and potential R&I priorities for Work Programme 2023-2024 and beyond.

The aim of the session was to discuss topics that had already been identified by industry, strategic vision and roadmapping papers, building on previous publications of initiatives like EPOSS, AIOTI, GAIA-X, NESSI, Cloud Alliance, et al., consultations and workshops, and to then add recommendations from these diverse speakers, as well as contributions from interested attendees.

The European Commission representatives set the scene around three pillars, already identified by industry, strategic vision and roadmapping papers, for the Continuum-focussed areas of the Horizon Europe 2023-2024 Work Programme: "*Open Source for Cloud/Edge Digital Autonomy*", "*Piloting emerging Smart IoT Platforms and decentralised intelligence*", and "*Cognitive Computing Continuum: Intelligence and automation for more efficient data processing*". The aim of the session was to discuss and assess these topics, as well as capture input and proposals from a diverse and representative set of European stakeholders, eight distinguished speakers were invited and suitable channels were made available for all participants to intervene, ask questions, and provide comments and opinions.

Two panels of four speakers were organised: the first one (including **Leo Isaac-Dognin**, Director at Capgemini; **Kay Bierzynski**, Technical Project Lead at Infineon; **Hubert Tardieu**, Independent Board Member of GAIA-X; **Josef Urban**, from Nokia) representing industry voices, and the second one (including **Monique Calisti**, CEO of Martel Innovate; **Koen de Bosschere**, Professor at University of Ghent; **Lars Nagel**, CEO of the International Data Spaces Association; **Ovidiu Vermesan**, Chief Scientist at SINTEF) gathering representatives of associations and partnership initiatives.

The panellists were not just speaking for themselves, not for their employer or institution; rather, they were bringing forth the common voice of a consolidated and well identified community, typically relying on a foundational document (e.g., a roadmap, a white paper, the results of a membership-wide consultation) that was prepared and published by the stakeholder group they represented.

Mr Isaac-Dognin represented a group of 27 companies who jointly wrote a "*European industrial technology roadmap for the next generation cloud-edge offering*" report; the topics of interest include cybersecurity, interoperability, and sustainability. Kay Bierzynski, representing the European Technology Platform on Smart Systems Integration (EPoSS), one of the three industrial associations behind the ECSEL and future Key Digital Technologies (KDT) Joint Undertakings, followed up with the findings of their "*AI at the Edge*" white paper, including HW/SW codesign, environmental impact optimisation, and a wide array of use cases. Hubert Tardieu drew from use cases in the automotive sector to emphasise the need for multiple, mature ecosystems that can serve the needs of industry and citizens in various sectors, where the research outcomes from Horizon Europe projects could be effectively put to the test towards an impactful market adoption. Josef Urban, as Chairman of NESSI, concluded the first panel by pointing out the importance of interoperability and multidisciplinary research, identifying technical and non-technical focus areas for future work and recognising Open Source as a critical enabler for digital autonomy, though not the only, exclusive way.

Monique Calisti, coordinating the EU-IoT CSA project that published the "*IoT and Edge Computing – Opportunities for Europe*" report after consultation workshop and activities, highlighted the importance of openness and the sheer diversity of scope within the Computing Continuum infrastructure. Mr De Bosschere, coordinating the HiPEAC European network which publishes a yearly "*HiPEAC Vision*" document, presented some technical, policy, and societal recommendations including supporting open-source hardware activities, interoperability across heterogeneous hardware, and fostering a stronger innovation culture and entrepreneurial mindset across Europe. Lars Nagel laid out IDSA's assessment that data spaces are at an important inflection point where implementation and scale-up actions gain momentum, but this needs a suitable, interoperable, distributed "soft" infrastructure and broad alliances such as the newly announced Data Spaces Business Alliance engaging IDSA, BDVA, FIWARE, and GAIA-X. The second panel was complete with Ovidiu Vermesan's statement, representing AIOTI and highlighting needs and challenges for IoT Edge research, such as large-scale integration and approaches to tackle the scalability gap at the edge.

The first of the three pillars, identified by industry, strategic vision and roadmapping papers and set forth by the EC representatives at the session, focussed on **Open Source**, which was continuously mentioned during the whole event. Open Source Software and Hardware enjoyed cross-cutting nature, going from Cloud services to Edge Computing, heterogeneous nodes, devices, hardware/software synergies and digital autonomy. Most if not all the panellist acknowledged the key role of Open Source and open standards in empowering Europe on the path to the Computing Continuum and digital autonomy: among others, EPoSS, GAIA-X, and the industry group insisted on the importance of open systems and interoperability, EU-IoT pushed for a wide definition of openness that also considers communities, and NESSI linked the value of open-source initiatives to the strength of their ecosystem. Further references to this pillar, in the area of hardware/software codesign, are to be found in the EPoSS and HiPEAC talks as well.

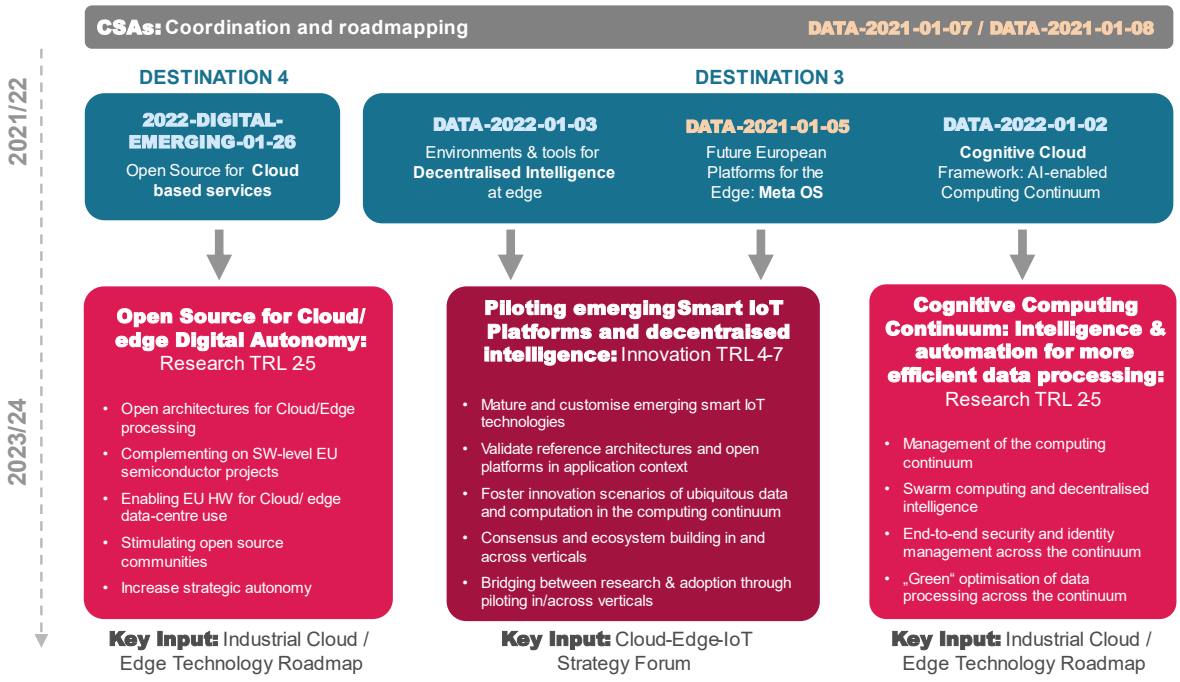

The second pillar considered the challenge of **mainstreaming novel architectures and approaches** across multiple application domains, once again carefully balancing the horizontal research (to drive convergence) with suitable verticalisation when appropriate (to address the current problems and needs of specific industries in a timely way). Large-scale pilots in multiple domains, combining results and technologies from previous R&I actions with state-of-the-art infrastructure and realistic requirements and use cases, were being indicated as necessary step for this mainstreaming process to succeed. This was endorsed by the industry group with their invitation to consider end-users in the overall vision, and

others such as GAIA-X and IDSA, with their stress on verticalised interoperable data spaces, or EPoSS and AIOTI with the importance and challenges of edge intelligence for a great number of industries.

The third pillar highlighted the need for **AI-enabled management principles across the Computing Continuum**. More dynamics and elasticity for groups of local nodes and systems are required. Also, intelligent orchestration mechanisms from the cloud to the edge, computing resources relocation from constrained devices towards central computing centres and hybrid cloud models, anticipating the need to process large data sets and higher levels of resource dynamicity. Experts pointed to the need to **shift from cloud-native architectures and applications, representing more or less the current state of the art, to *continuum-native* ones**, which require **further research** and exploration on multiple aspects, such as end-to-end security, devices heterogeneity, energy consumption optimisation, elasticity to flexibly allocate resources and tasks, and effective languages and semantics to describe it.

The synergy among AI at the Cloud/Edge/IoT, dynamic resource orchestration, and distributed data generation and processing, with the critical underlying goal to support environmental impact optimisation towards climate neutrality, received input and endorsement from multiple angles. The industry group that wrote the "*European industrial technology roadmap for the next generation cloud-edge offering*" report underlined the importance of differentiating the various Computing Continuum segments, each at a different level of maturity, through dynamic and interoperable resource management, to secure European competitiveness and leadership in the flourishing emergent data economy. NESSI explicitly pointed to **Cognitive Computing Continuum as an important research focus**, while EPoSS mentioned **"green" optimisation** as an important driver of the codesign approach for AI at the edge. Moreover, the industry group placed sustainability in their absolute top priorities, possibly with combined hardware/software approaches, and AIOTI underscored the interplay of intelligent processors at the IoT edge, resource- and energy-constrained nodes, and collective behaviour of upcoming Continuum-native systems.

It is evident that for Europe to play a significant role in the emerging data economy, Europe's industry must work in partnership – and seize the opportunities offered by managing the Computing Continuum. A common framework for sharing resources and data, monetising the value of data via transparent cost-contract based data quality schemes and nurturing decentralised ecosystems were reported as the key success factors for Europe to differentiate itself in global markets.

The emerging paradigm of the Computing Continuum requires more collaboration and synchronisation across the different Strategic Research Agendas. To find a suitable balance between top-down and bottom-up, R&I was felt essential to cope with the fact that leading cloud-service providers are quite active in this domain and to exploit the current window of opportunity. **Low-TRL R&I actions would need to be combined with the need for piloting innovative solutions across the Computing Continuum systems in different application scenarios, assessing scalability, interoperability, trust, and sustainability within realistic and concrete situations.** Open standards related to data semantics, data exchange, identity management as well as open architectures were a baseline concept for envisaged platforms and the whole ecosystem of these future piloting initiatives, in line with what was expressed in other sessions. As a key element of the green transition, the energy footprint of computing at central places is on the rise, the trend towards Edge Computing may reverse the tremendous energy consumption of central computing by decentralisation and local computing.

Lastly, human potential aspects were tackled, mentioning critical enablers such as skills, training, talent, and entrepreneurship. The consultation on the 2023-2024 Work Programme will continue with further actions, and the current timeline projects a drafting during the first half of 2022, after the calls of April, to then complete the adoption in the second half of the year.

4. CONCLUDING REMARKS

The workshop closed summarising the main outcomes of the various event's sessions and indicating the expected next steps in view of the Horizon Europe 2022 Calls and the 2023-2024 Work Programme definition. Starting from the opportunities and challenges for Europe that can be directly connected to the upcoming Horizon Europe calls, as well as the pivotal European policies such as the Data Strategy, the closing session painted a comprehensive and integrated picture that covered multiple programmes and investment plans, collaborative alliances, research actions, and the connection between European values and digital autonomy.

The summary of the morning sessions confirmed further the need for an AI-enabled Cloud/Edge/IoT Computing Continuum, the centrality of open computing, including open-source software but also hardware, covering all levels of the complete ICT stack, not only as a viable and effective strategy but also as an irreplaceable cornerstone of any plan towards European digital autonomy. This view has been highlighted from multiple EC, academic, and industrial sources throughout the event and is one of the main takeaways of the whole workshop, with the important caveats that **digital autonomy must not give in to isolationist temptations but safeguard European values and approaches within a collaborative and fair open playing field for all actors**, and that Open Source is fundamentally necessary but not mutually exclusive with other, partially proprietary approaches that can serve European interests in specific situations. The positive correlation between increased activity levels of participation in open-source projects and GDP improvement was also highlighted in the morning sessions, and a significant shift was recognised to have occurred in the overall industry attitude towards open approaches during the last two decades.

The interventions around **swarm and decentralised computing** suggested that these distributed approaches **are actually very close to mainstreaming and have strong market significance**, with concrete examples and success stories presented in a number of different domains and an expected temporal window of less than six years for the establishment of general use applications of these concepts and technologies. This recognition acted as a tie-in to the afternoon session covering the ongoing definition and shaping of the Horizon Europe Work Programme (WP) for the next two years 2023 and 2024.

The consultation nature of the workshop in the afternoon produced several insights and suggestions to be taken on board in the Horizon Europe WP, from new computing architectures and new hardware components to the **shift from cloud-native to continuum-native** and a new generation of flexible, elastic, and secure orchestration approaches and technologies that can enable applications to effectively thrive across the whole Cloud/Edge/IoT R&I spectrum. A different but nevertheless critical angle that emerged in the afternoon session was **the importance of European-wide open and collaborative approaches, often embodied by associations and alliances that gather multiple and diverse actors toward a common roadmap and agenda**. It is now widely recognised that the scope and magnitude of the challenges at hand is such that no single company or tight group of stakeholders can hope to fully tackle and solve them; on the contrary, wide and complementary partnerships are needed to produce multiple integrated approaches and solutions. A correlated topic that was also raised in the discussion was the need of continued progress in **community building, skills and talent development, and the creation of a sustainable and open European ecosystem** that can grow and support the vision of the Computing Continuum and digital autonomy.

The nature of the event, combining EC landscaping and vision with input and recommendations from a varied and representative ensemble of speakers and very pertinent questions and interventions from the participants, set the premise for a successful and fruitful conversation among European stakeholders on R&I in the scope of the Computing Continuum. The *Digital Autonomy in the Computing Continuum* workshop represented an important checkpoint and steppingstone for the European Commission and the whole European research and innovation community on the way to the next Cloud/Edge/IoT calls in 2022 and the upcoming advancements contained in the Horizon Europe 2023-2024 Work Programme.