MYRTUS: Multi-layer 360° dYnamic orchestration and interopeRable design environmenT for compute-continUum Systems

Francesca Palumbo[2,*], Maria Katiuscia Zedda[1,§], Tiziana Fanni[1], Alessandra Bagnato[4], Luca Castello[12], Jeronimo Castrillon[5], Roberto Del Ponte[11], Yansha Deng[10], Bart Driessen[6], Mauro Fadda[3], Tristan Halna du Fretay[13], Julio de Oliveira Filho[6], Veena Rao[9], Francesco Regazzoni[14,15], Alfonso Rodríguez[8], Melanie Schranz[7], Giulia Sedda[2]

[1]Abinsula Srl, [2]University of Cagliari, [3]University of Sassari, [4]Softeam, [5]TU Dresden, [6]TNO, [7]Lakeside Labs, [8]Universidad Politécnica de Madrid, [9]Hiro Microdatacenters, [10]King's College London, [11]Forge Reply Srl., [12]ArubaKube Srl, [13]Canon Research Centre France, [14]University of Amsterdam, [15]Università della Svizzera italiana

ABSTRACT

The MYRTUS Horizon Europe Project embraces the principles of the EUCloudEdgeIOT Initiative and integrates edge, fog and cloud computing platforms, leveraging a cognitive engine based on swarm intelligence and federated learning to orchestrate collaborative distributed and decentralised components. Components are augmented with interface contracts covering both functional and non-functional properties.

CCS CONCEPTS

• **General and reference → Design**; **Performance**; • **Hardware → Methodologies for EDA**; • **Computer systems organization → Embedded and cyber-physical systems**.

KEYWORDS

Computer hardware and architecture, Design environment, Dynamic orchestration, Computing continuum, Interoperability, AI.

ACM Reference Format:

Francesca Palumbo[2,*], Maria Katiuscia Zedda[1,§], Tiziana Fanni[1], Alessandra Bagnato[4], Luca Castello[12], Jeronimo Castrillon[5], Roberto Del Ponte[11], Yansha Deng[10], Bart Driessen[6], Mauro Fadda[3], Tristan Halna du Fretay[13], Julio de Oliveira Filho[6], Veena Rao[9], Francesco Regazzoni[14,15], Alfonso Rodríguez[8], Melanie Schranz[7], Giulia Sedda[2]. 2018. MYRTUS: Multi-layer 360° dYnamic orchestration and interopeRable design environmenT for compute-continUum Systems. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation emai (Conference acronym 'XX)*. ACM, New York, NY, USA, 6 pages. https://doi.org/XXXXXXX.XXXXXXX

MYRTUS is funded by the European Union, by grant No. 101135183. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them. [§]*Project Coordinator (katiuscia.zedda@abinsula.com)*, [*]*Scientific Coordinator (francesca.palumbo@unica.it)*.

Conference acronym 'XX, June 03–05, 2018, Woodstock, NY

ACM ISBN 978-1-4503-XXXX-X/18/06
https://doi.org/XXXXXXX.XXXXXXX

1 INTRODUCTION

The last few years have been fraught with world changing incidents, including extreme weather events due to climate change and COVID-19 pandemic, resulting in major investments in digitalization and virtualization. Embedded and ubiquitous computing and smart anything everywhere were already at our disposal and represented an answer to support remote work and distant monitoring of industry verticals and to allow people's lives to go forward. This has entailed an extraordinary data deluge to be mastered, with applications requiring to monitor/analyse data flows in different ways: close to source when strict real-time requirements are there, or in large data centres/cloud facilities when batch-data analyses are involved. Novel approaches to design, deploy, and orchestrate such a complex computing ecosystem, where a diversity of fog- and edge-level devices converge with the cloud to form a computing continuum, are yet to be defined to *"provide the trustworthy data processing infrastructure and services that public administrations, businesses and citizens require in Europe"*[1]. Such strategies should also cope with the semiconductor shortage[2], which is still affecting some market segments[3]. The urgency to adopt a sustainable computing approach calls then for a strategy to avoid, or at least delay, technology obsolescence, as well as, to balance the use of resources, and to minimise energy consumption.

Extreme digitalization and virtualization brought also unprecedented levels of computing pervasiveness and human-computer interaction. As examples, the smart factory 5.0, with cooperation of humans and cognitive machines, is just around the corner, and the metaverse is expected to break the barriers between virtual and real worlds [4]. Cyber-Physical entanglement is at the base of Cyber-Physical Systems (CPSs) that are *"built from, and depend upon, the seamless integration of computational algorithms and physical components"*[4], and are about to evolve towards a living dimension [10], expected to become capable of both cooperating with their human mates and of inducing positive human behaviours. Such a disruptive change requires addressing, not only the usual CPS challenges, related to modeling, design, architectural efficiency, and security,

[1]https://digital-strategy.ec.europa.eu/en/news/towards-next-generation-cloud-europe
[2]https://inews.co.uk/news/technology/global-chip-shortage-silicon-2023-intel-laptop-supply-chain-1263242
[3]https://www.jpmorgan.com/insights/research/supply-chain-chip-shortage
[4]https://www.nsf.gov/pubs/2017/nsf17529/nsf17529.htm

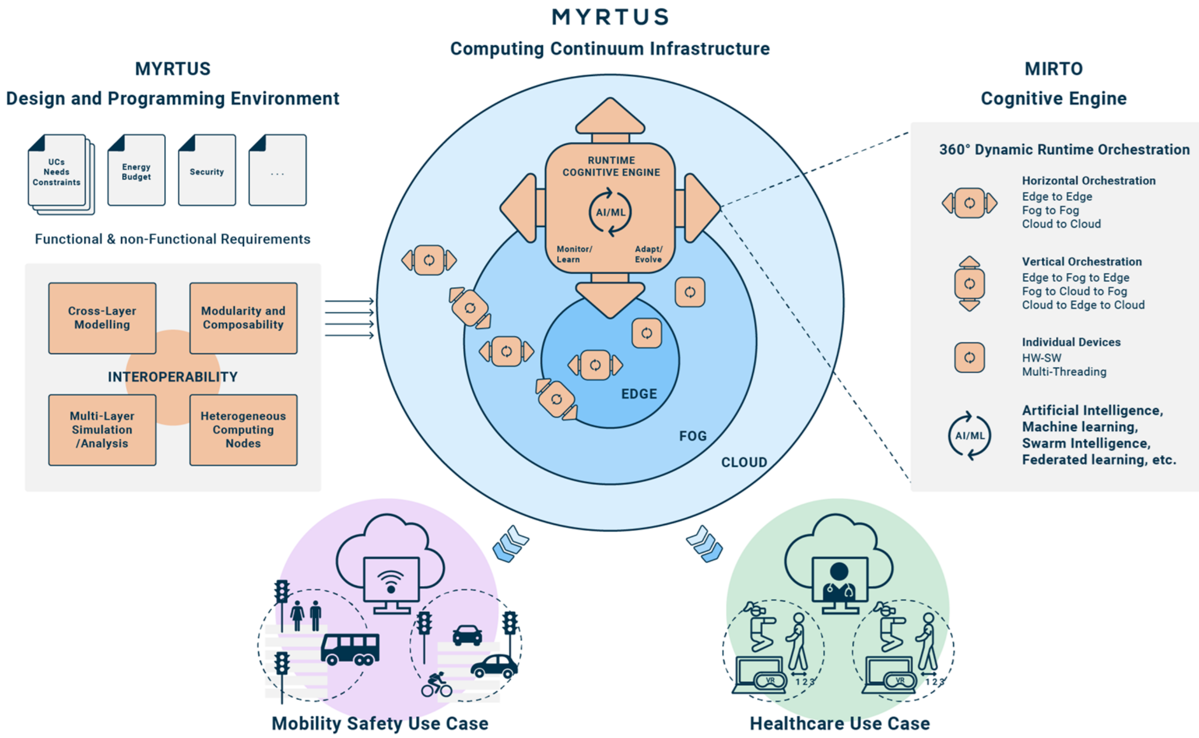

Figure 1: MYRTUS general overview

but also facing the emergence of novel issues, due to the distributed self-evolving nature of a living heterogeneous computing infrastructure, where boundaries among humans and devices fade away, pushing the concept of autonomic systems to the extreme.

2 MYRTUS PROJECT AMBITION

The ***"Multi-layer 360°dYnamic orchestration and interopeRable design environmenT for compute-continUum Systems (MYRTUS)"*** project[5] aims to provide the technology to make CPS evolving towards a living dimension, embracing the principles of the EUCloudEdgeIOT Initiative[6] by (1) contributing to pull edge, fog and cloud computing platforms closer together into a seamless execution environment, and (2) providing languages and tools to orchestrate collaborative distributed and decentralised components. MYRTUS technologies pursue also sustainable and responsible computing, openness, security, and trustworthiness. Technology is assessed within two challenging scenarios: healthcare and mobility.

Figure 1 illustrates a schematic overview of MYRTUS and its technological pillars that are discussed in the rest of the section:
Pillar 1 *- MYRTUS computing continuum infrastructure.*
Pillar 2 *- Multi-layer 360°dynamIc RunTime Orchestration (MIRTO) Cognitive Engine.*
Pillar 3 *- MYRTUS Design and Programming Environment (DPE).*

[5]https://myrtus-project.eu/
[6]https://eucloudedgeiot.eu/

Pillar 1: MYRTUS computing continuum infrastructure

Technical Challenge: Centralised computing solutions struggle to meet scalable and low-latency application scenarios, and privacy-related matters. Data exposure can improve learning and decision-making processes but may violate data protection regulations. Edge computing helps to address security and privacy with proximity, but hardly satisfies intense computing and storage requirements.
Technical Objective: MYRTUS defines a reference infrastructure where a diversity of fog-level and edge-level devices converge with the cloud to form a computing continuum capable of addressing the needs of complex and dynamic systems, including CPS with a living dimension.
Key Enabling Technology: MYRTUS onion-like computing continuum reference infrastructure will represent the enabler for seamless optimal execution of complex computational workflows across a heterogeneous continuum, in an energy-efficient and trustworthy manner.

Edge processing entails widely acknowledged advantages, including minimal data-to-result processing time, savings in bandwidth, energy and costs for data transfer (as well as in data storage), improved resiliency to network failures, increased data security and privacy. Data can be immediately transformed into actionable items, upon which to make decisions. On the downside, edge computing offers limited computing resources, storage capacity and energy budget. Cloud processing provides scalable and quasi-unlimited data storage, also offering centralization capabilities with easier data accessibility to any organisation entity. To bring those extremes closer, the fog intermediate layer was defined by CISCO [2]. In MYRTUS, the resources are considered as a continuum and organised as a composable, layered cloud-fog-edge onion-like infrastructure, integrating heterogeneous, autonomous, federated, and collaborative computing nodes, as appears in the centre of Figure 1, enabling the seamless execution of complex computational workflows. The goal is to achieve close-to-zero latency processing when needed while still having a guaranteed high-availability connection to the cloud.

MYRTUS infrastructure builds atop strong competences and expertise at the edge on a plethora of diverse devices including off-the-shelf multicores, Heterogeneous Multiprocessor Systems-on-Chip FPGA-based accelerators [15, 13], and adaptive RISC-V processors with custom computing units [18]. These devices allow advanced complex computation, including Artificial Intelligence (AI) algorithms at the edge and the implementation of self-adaptive behaviours [9]. To provide additional analytics services on medium to long-term data, edge resources are flanked by smart gateways [5] and Fog Micro Data Centres[7]. The fog layer, besides providing medium to long-term analytics and storage, will also monitor and supervise the local subsystems and provide network virtualization. The MYRTUS infrastructure is completed by the cloud layer providing storage and analytics with long-term capabilities, as well as deep data mining, all-around subsystems monitoring and coordination and embedded cloud federation functions. Indeed, data and workload distribution over a large set of heterogeneous resources from different providers requires an automation viewpoint on trust and compliance that, in MYRTUS, will be accomplished by leveraging the integration with the Gaia-X trust model.

Pillar 2: MIRTO 360°Cognitive Engine

Technical Challenge: Cloud, edge and end-devices are currently handled as different isolated silos, preventing an application to be seamlessly deployed and dynamically updated for continuous optimization across the continuum.
Technical Objective: through the AI-powered MIRTO cognitive engine, MYRTUS features 360°dynamic runtime orchestration, guaranteeing high performance and energy efficiency, while preserving security and trust.
Key Enabling Technology: MIRTO will enable continuous workload/system optimization, strategic to foster a sustainable computing approach, across the continuum.

MIRTO (Figure 2) is a distributed cognitive engine orchestrating a scalable, dynamic, and heterogeneous infrastructure. It exploits self-awareness through feedback control loops [14] to implement self- reconfiguration [20], featuring self-healing/-optimization capabilities. MIRTO aims at: i) ensuring dynamic functional composition, efficient resource allocation and task execution and avoiding network congestion, managing the service deployment; ii) guaranteeing energy efficiency, and iii) improving trustworthiness, securing computation and data protection; through a multi-layer cognitive self-adaptation loop implementing the following steps [9].

[7]https://hiro-microdatacenters.nl/

SENSING - internal and external triggers are sensed/monitored. For a device, internal stimuli refer to the status of the monitored resource (e.g., energy consumption, throughput) and external stimuli are environmental or human triggers. For a layer, internal stimuli are information related to that hierarchy level (e.g., load of other involved resources) and external stimuli are coming from a different layer (e.g., failure of a subsystem). Sensory information is fused at an ontological knowledge base within MIRTO.
EVALUATION - Sensory information enables performance evaluation and comparison to expected indicators. Evaluation identifies performance degradation or opportunities for improvement.
DECISION - Informed decisions are made based on the computed indicators and the knowledge base and will affect both functional (e.g., executed mission or algorithm) and non-functional (e.g., QoS, security, energy, etc.) requirements execution aspects.
RECONFIGURATION - Runtime reconfiguration strategies are put in place per layer and in a target-dependent manner.

MIRTO will build upon containerization, interface abstraction, and modularization concepts, being able to deploy in heterogeneous environments and monitor the deployment process in a homogeneous way. It will result in a set of layer- and target-dependent cooperative collection of orchestration agents realising the above-mentioned self-adaptation loop with specific adaptation goals, leveraging on an interoperable hierarchy of distributed monitors and managers allowing for cross-layer information flow. MIRTO will allow dynamic 360° orchestration, by means of its distributed agents, enabling both vertical (inter-layer) and horizontal (intra-layer) orchestration, as well as individual resource management. All layers share one ontological knowledge base (logical view), which can be distributed in different layers (implementation view). This way, data remain close to their consumers, but can be shared between layers for analysis and decision-making. To pursue interoperability and easier uptake of the proposed solution, TOSCA[8] standard language will be adopted. TOSCA is an open-source domain specific language (DSL) providing model-based descriptions of services, platforms, infrastructure, and data components, along with their relationships, requirements, capabilities, configurations, and operational policies.

Pillar 3: MYRTUS Design and Programming Environment

Technical Challenge: The more a system is heterogeneous, complex, and adaptive, the more designers rely on partially integrated toolchains/methodologies tuned for specific aspects. There is no effective interoperability.
Technical Objective: MYRTUS DPE for continuum computing systems, features interoperable support for cross-layer modelling, threat analysis, design space exploration, modelling, components synthesis, and code generation.
Key Enabling Technology: The DPE judiciously integrate tools and standard formats, from high-level application specification to low-level deployment tools for heterogeneous nodes with a strong focus on interfaces and interoperability, facilitating the MYRTUS solutions uptake at large scale.

[8]https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=tosca

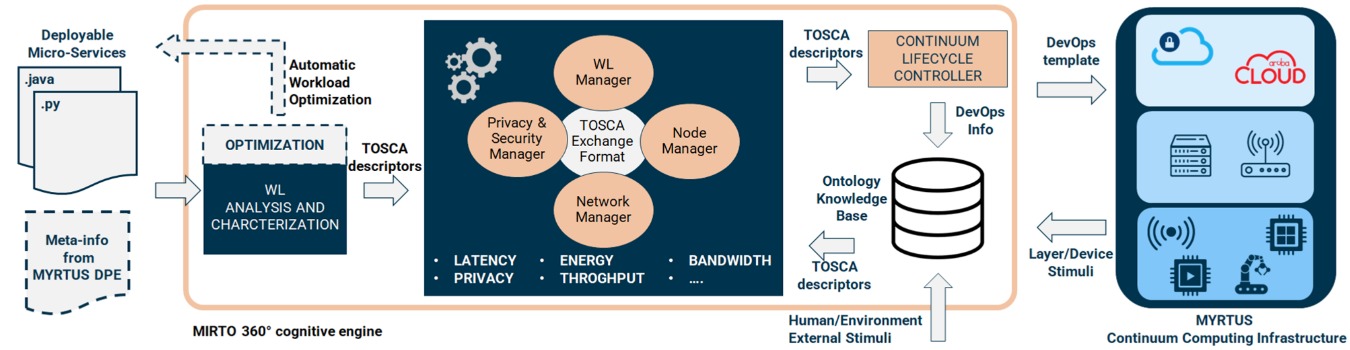

Figure 2: MIRTO 360°Cognitive Engine

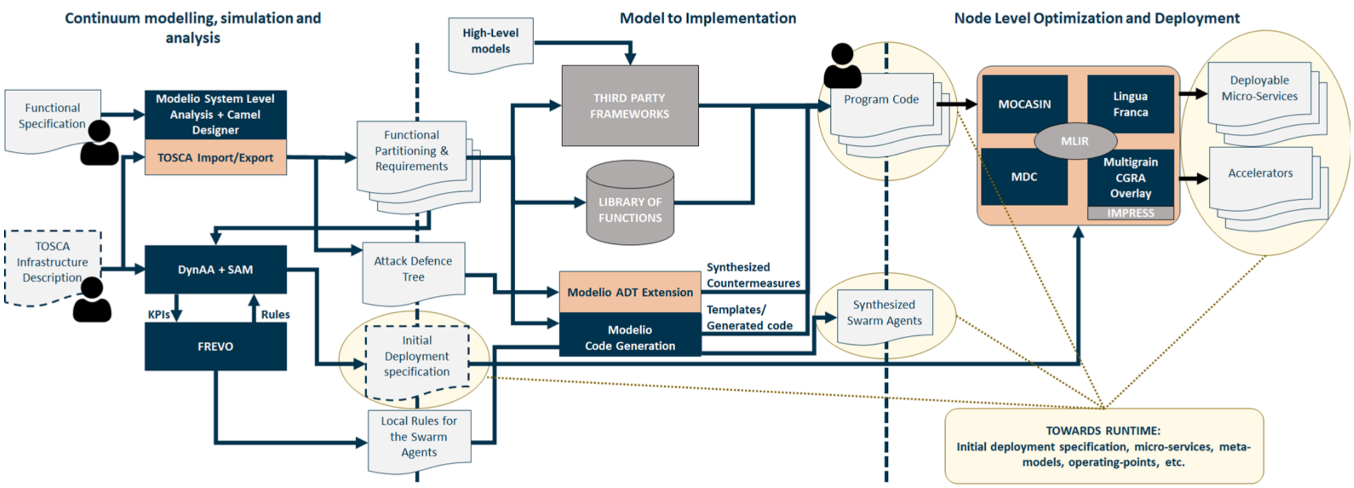

Figure 3: MYRTUS Design and Programming Environment

Figure 4 depicts MYRTUS DPE that integrates both consortium-provided and third-party tools. As inputs, the users define the high-level functional specifications of the scenario/service to be deployed over the continuum and, if available, an initial description of the infrastructure using standard languages, like a TOSCA-compliant description. Outputs depend on the DPE flow steps and are described below. Along with them, MYRTUS DPE exports meta-information to seamlessly integrate with the MIRTO cognitive engine:1) a design-time characterization of the operating points to jump-start the workload analysis and characterization within MIRTO, and 2) the model-level information to enable model-based re-optimization of the deployable microservices without performing costly analysis of the binaries at runtime, leveraging on promising preliminary works for adaptive execution on data centres [6]. The chosen microservice approach fosters better runtime composability, allowing for an easier deployment of Hardware (HW) and Software (SW) tasks where most appropriate to the given constraints and infrastructure status at deployment time. Modularity is meant to improve testability, maintainability and portability of the microservices. MYRTUS DPE flow is composed of the following three main steps:

Continuum modelling, simulation and analysis: This step will produce i) the functional partitioning of the scenario starting from UML description, ii) the functional level requirements, exploiting the model-based performance indicators estimation provided by Modelio[9], iii) an initial deployment specification model in TOSCA, iv) the local rules for the Swarm Intelligence (SI) agents to be used within MIRTO agents, and v) the Attack Defence Tree representing the analysis of the possible threats to which the system is exposed.
Model to Implementation: Going from the modeling level to deployment implies defining the Program Code. Modelio will be used to synthesize code, the swarm agents, and the code to mitigate the security threats. Moreover, interoperability mechanisms will be also guaranteed so that implementations can be derived from external DSLs or coded using external libraries.
Model to Implementation: This step results in the deployable microservices, along with the specification of the specialised HW accelerators. A common interoperability framework based on Multi-Level Intermediate Representation (MLIR) will be used to allow i) importing third-party codes (from DSLs like [19] or TensorFlow

[9]https://www.modelio.org/index.htm

models), ii) having access to third-party tools (like polyhedral compilers for optimization purposes), and ii) compiling code for different targets (as reconfigurable accelerators, CPUs, or customizable RISC-V cores). This will build atop the MLIR infrastructure of the EVEREST project [12, 11].

3 MYRTUS ASSESSMENT

MYRTUS techniques will be tested in smart mobility and virtual telerehabilitation scenarios, ensuring that the project delivers real-world benefits to users.

Swarm-inspired cooperative telerehabilitation: In home rehabilitation, people achieve higher performance interacting with a partner [17]. Collaborative telerehabilitation has been studied for neuromotor recovery, but MYRTUS will provide the technologies to investigate a novel swarm-inspired paradigm for telerehabilitation promoting attention, visuomotor coordination, visuospatial exploration, movement recovery, pace in a group task, empathy, and engagement, while doing neurorehabilitation exergames (videogames with physical activity) at home through a self-managed rehabilitation framework enabling remote supervision and interaction between therapists, patients and virtual nodes. Human-only interaction presents two limitations: 1) the number of participants simultaneously connected might not be sufficient, and 2) triggering specific reactions and promoting adaptation, learning, and interaction among patients might be unfeasible. Here is where the living dimension of CPS becomes handy. Defining the exergame enabling an heterogenous composition of human beings and virtual agents, the latter fitting specific patient/therapist models, will overcome these issues by guaranteeing training no matter the number of connected people, while supporting the weaker patients and progressively challenging the strongest ones. The rehabilitation goal is a collaborative solution of a task in a virtual environment, according to the rules imposed by the exergame, whose specs (e.g., purpose of the rehabilitation, type of exercise, etc.) will be defined by bioengineers and rehabilitation specialists.

Intersection Safety: Managing traffic safety at road intersections will become increasingly more complex and adaptive with the deployment of Connected, Cooperative Automated Mobility (CCAM). In addition to traffic rules and traffic light controllers, roadside systems will autonomously communicate and cooperate intensively with Automated Vehicles (AV), Vulnerable Road Users (VRU) with personal devices, and traffic safety and information services in the cloud. Roadside systems will need to be more reliable and accurate to detect/predict imminent conflicts (e.g. red-light violations) and collisions with VRU. Connecting AV to each other and to the infrastructure allows for collaboration on the perception of other road users (Collective Perception Messages services) to improve perception performance and reduce occlusion effects, and coordination of intentions and manoeuvres (Maneuver Coordination Message services) in order to prevent/mitigate safety hazards. This requires continuous adaptation and learning in a changing environment with endless safety critical scenarios that could occur. Depending on the traffic situation, computational tasks can be relatively simple, or extremely complex. Dynamic offloading of tasks from edge to fog or cloud can ensure these tasks can always be performed whilst alleviating the computational burden for the different components. It also allows sharing the "lessons learned" across all other smart infrastructures. Indeed, this kind of CPS is meant to become a living one with continuous exchange, both at the infrastructure level and with users in a seamless manner and learning processes. MYRTUS technologies will enable a next evolution in CCAM with the ability to advise and coordinate manoeuvres to avoid collisions.

4 MYRTUS EXPECTED IMPACT

MYRTUS contributes to create new knowledge in the computing continuum domain, with methodologies and tools for node execution and processing portability over edge-fog-cloud, including dynamic and seamless orchestration. The goal is to become a reference in the computing continuum, offering solutions that overcome the problem related to vendor/platform lock-in, promoting and facilitating the adoption of MYRTUS technologies among startups and SMEs, and reducing their development time and cost. MYRTUS embraces the sustainable and responsible computing paradigm, promoting obsolescence avoidance (supported by MYRTUS principle of openness, interoperability, and portability) and resource saving and energy efficiency (supported by HW specialisation and optimization techniques). Collaboration is a key driver of innovation and knowledge exchange that can lead to more efficient research outcomes and a better understanding of the broader research landscape. MYRTUS has a strategy to establish synergies with other projects and initiatives, including important associations (HIPEAC, INSIDE, Gaia-X, etc.), technology communities, the IPCEI initiatives and the other research projects.

5 BEYOND STATE OF THE ART

MYRTUS solutions will address challenging research issues related to design, programming, and operation of highly distributed and scalable systems along the continuum, and will provide important breakthroughs in specific domains (such as AI at the edge, security in federated environment, etc.).

Centralised cloud solutions struggle to meet scalable and low latency scenarios and privacy, while edge computing hardly satisfies intense computing and storage requirements. Additionally, cloud, edge, and end-devices are currently handled as isolated silos. MYRTUS intends to offer an ecosystem of different computing resources spanning layers. Heterogeneity and composability of customizable and commercial solutions will be key to fit different performance vs energy profiles. Additionally, MIRTO AI-powered orchestrator will enable the dynamic creation, management, and run-time optimization of the continuum, offering i) a unique virtual space spanning across multiple administrative/technological domains, ii) a hosting environment for distributed applications that can be seamlessly deployed across continuum.

Heterogeneity is mastered by leveraging on SI at the fog/edge, which supports seamless integration of different devices via agent-based swarm modelling enabling high reactivity for timely allocation of resources and energy efficiency, security, and privacy by keeping processing of data as local as possible. Moreover, the MIRTO cognitive engine will support evolution leveraging on the runtime learning and federated approaches implemented within the AI-powered managers (visible in Figure 2). As an example, the node manager is supposed to evolve at runtime, continuously improving

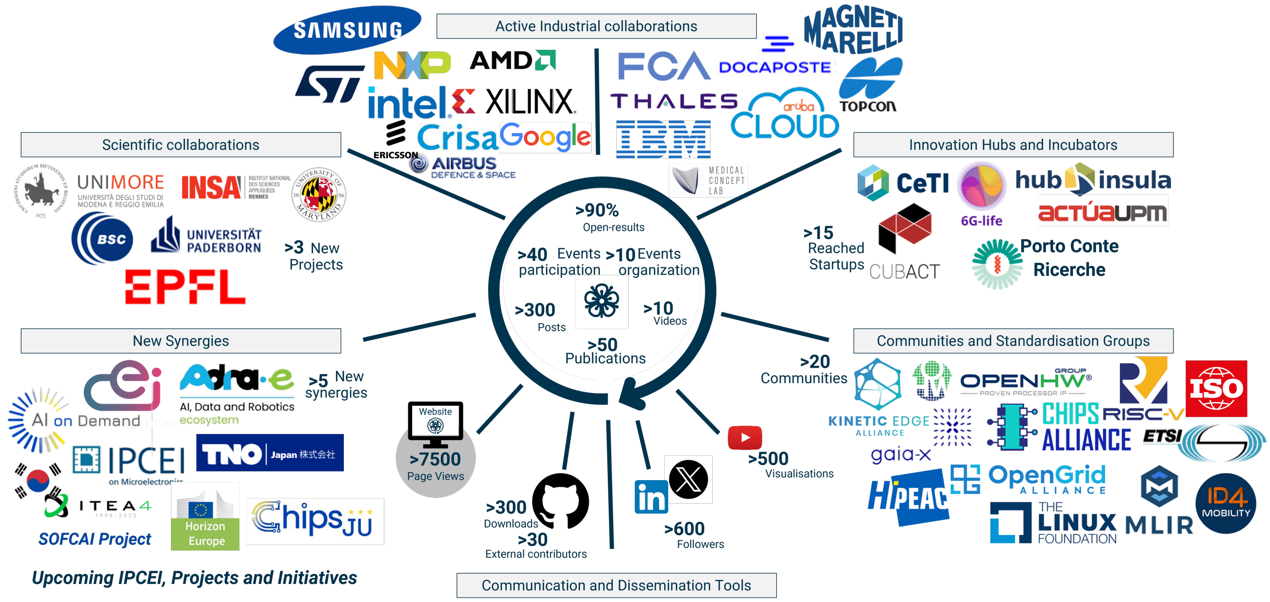

Figure 4: MYRTUS Impact Strategy

the models while using them (intra-node) and exploiting federated learning to create consensus among distributed nodes and share the evolved models (inter-node).

Finally, MYRTUS aims at tackling the interoperability chimera. The more heterogeneous and complex a system is, the more designers rely on diverse non-integrated toolchains/methodologies. There are good examples of frameworks for CPS [8] and swarm of CPS [16] design and operation, and for porting AI at the edge [3]. The issue is how to get the most out of them, making them talk. Effective interoperability, despite promises of meta-model [1] and semantic [7] approaches, is still to be reached. MYRTUS DPE will concretely tackle this issue leveraging open-source approaches. The model-side of the DPE is based on Modelio, which is open-source and counts on a vast community. The implementation-side aggregates different consortium and third party's tools around MLIR that addresses SW fragmentation, improving compilation for heterogeneous HW by reducing the cost of building domain-specific compilers, and aiding in connecting existing compilers together.

REFERENCES

[1] Balarin et al. 2003. Metropolis: an integrated electronic system design environment. *Computer*, 36, 4, 45–52. DOI: 10.1109/MC.2003.1193228.

[2] Bonomi. et al. 2012. Fog computing and its role in the internet of things. In *ACM Workshop on Mobile Cloud Computing*, 13–16. DOI: 10.1145/2342509.2342513.

[3] Busia et al. 2021. ALOHA: A unified platform-aware evaluation method for cnns execution on heterogeneous systems at the edge. *IEEE Access*, 9, 133289–133308. DOI: 10.1109/ACCESS.2021.3115243.

[4] Dwivedi et al. 2022. Metaverse beyond the hype: multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. *Int. J. Inf. Manag.*, 66, 102542. DOI: 10.1016/J.IJINFOMGT.2022.102542.

[5] Fanni et al. 2022. The multi-sensor gateway, a unified communication scheme and orchestration actor for heterogeneous systems. In *CPS Workshop 2022*. Vol. 3252. CEUR-WS.org. https://ceur-ws.org/Vol-3252/paper3.pdf.

[6] Fettweis et al. 2019. Architecture and advanced electronics pathways toward highly adaptive energy- efficient computing. *Proceedings of the IEEE*, 107, 1, 204–231. DOI: 10.1109/JPROC.2018.2874895.

[7] Larsen et al. 2018. Features of integrated model-based co-modelling and co-simulation technology. In *Software Engineering and Formal Methods*. Springer International Publishing, 377–390. ISBN: 978-3-319-74781-1.

[8] Palumbo et al. 2019. CERBERO: Cross-layer model-based fRamework for multi-oBjective dEsign of reconfigurable systems in unceRtain hybrid envirOnments: invited paper: CERBERO teams from uniss, unica, ibm research, tase, insa-rennes, upm, usi, abinsula, ambiesense, tno, st, crf. In *Conf. on Computing Frontiers*, 320–325. DOI: 10.1145/3310273.3323436.

[9] Palumbo et al. 2019. Hardware/Software Self-adaptation in CPS: the CERBERO Project approach. In *Embedded Computer Systems: Architectures, Modeling, and Simulation*. Springer International Publishing. ISBN: 978-3-030-27562-4.

[10] Palumbo et al. 2023. Towards a living dimension: The future of cyber- physical systems. *HiPEAC Vision 2023*, (Jan. 2023). DOI: 10.5281/zenodo.7461786.

[11] Pilato et al. 2024. A system development kit for big data applications on FPGA-based clusters: the EVEREST approach. In *Design, Automation and Test in Europe Conf.* 6pp.

[12] Pilato et al. 2021. EVEREST: a design environment for extreme-scale big data analytics on heterogeneous platforms. In *Design, Automation and Test in Europe Conf.* DOI: 10.23919/DATE51398.2021.9473940.

[13] Ratto et al. 2022. Multithread Accelerators on FPGAs: A Dataflow-Based Approach. In *PARMA-DITAM Workshop 2022*. Vol. 100. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 6:1–6:14. DOI: 10.4230/OASIcs.PARMA-DITAM.2022.6.

[14] Rutten et al. 2017. Feedback control as mape-k loop in autonomic computing. In *Software Engineering for Self-Adaptive Systems III. Assurances*. Springer International Publishing, 349–373. ISBN: 978-3-319-74183-3.

[15] Sau et al. 2017. Challenging the best hevc fractional pixel fpga interpolators with reconfigurable and multifrequency approximate computing. *IEEE Embedded Systems Letters*, 9, 3, 65–68. DOI: 10.1109/LES.2017.2703585.

[16] Schranz et al. 2018. Modelling a cps swarm system: a simple case study. In *Conf. on Model-Driven Engineering and Software Development*. SCITEPRESS - Science and Technology Publications, Lda, 615–624. DOI: 10.5220/0006731106150624.

[17] Takagi et al. 2019. Individuals physically interacting in a group rapidly coordinate their movement by estimating the collective goal. *eLife*, 8, (Feb. 2019), e41328. DOI: 10.7554/eLife.41328.

[18] Vázquez et al. 2022. Extending RISC-V Processor Datapaths with Multi-Grain Reconfigurable Overlays. In *Conf. on Design of Circuits and Integrated Circuits*, 01–06. DOI: 10.1109/DCIS55711.2022.9970069.

[19] Sven Karol, Tobias Nett, Jeronimo Castrillon, and Ivo F. Sbalzarini. 2018. A domain-specific language and editor for parallel particle methods. *ACM Transactions on Mathematical Software (TOMS)*, 44, 3, Article 34, (Mar. 2018), 32.

[20] Kephart and Chess. 2003. The vision of autonomic computing. *Computer*, 36, 1, 41–50. DOI: 10.1109/MC.2003.1160055.

Received 6 March 2024; revised xxx; accepted xxx