Guimaraes, C., Groshev, M., Cominardi, L., Zabala, A., Contreras, L. M., Talat, S. T., Zhang, C., Hazra, S., Mourad, A. & de la Oliva, A. (2021). DEEP: A Vertical-Oriented Intelligent and Automated Platform for the Edge and Fog. *IEEE Communications Magazine, 59*(6), 66–72.

DOI: 10.1109/mcom.001.2000986

DEEP: A Vertical-Oriented Intelligent and Automated Platform for the Edge and Fog

Carlos Guimarães*, Milan Groshev*, Luca Cominardi†, Aitor Zabala‡, Luis M. Contreras§,
Samer T. Talat¶, Chao Zhang**, Saptarshi Hazra††, Alain Mourad‡‡, Antonio de la Oliva*
*Universidad Carlos III de Madrid, Spain, †ADLINK Technology, France, ‡Telcaria Ideas S.L., Spain,
§Telefónica I+D, Spain, ¶Industrial Technology Research Institute, Taiwan, **Lund University, Sweden,
††RISE, Sweden, ‡‡InterDigital Europe, UK

***Abstract*—The 5th generation (5G) of mobile communications introduces improvements on many fronts when compared to its previous generations. Besides the performance enhancements and new advances in radio technologies, it also integrates other technological domains, such as cloud-to-things continuum and artificial intelligence. In this work, the 5G-DIVE Elastic Edge Platform (DEEP) is proposed as the linking piece for the integration of these technological domains, making available an *Intelligent Edge and Fog 5G End-to-End (E2E)* solution. Such solution brings numerous benefits to the vertical industries by enabling a streamlined, abstracted, and automated management of their vertical services, thus contributing to the introduction of novel services, cost savings, and improved time-to-market. Preliminary validation of the proposed platform is performed through a proof-of-concept, along with a qualitative analysis of its benefits to Industry 4.0 and Autonomous Drone Scouting vertical industries.**

I. INTRODUCTION

The previous generations of mobile communications firstly aimed at supporting data services, along with traditional voice services, and later at improving spectral efficiency and enhanced services at lower costs. However, the network infrastructure followed an *one-size-fits-all* paradigm. Prospecting significant changes on its vertical and horizontal structure, the 5th generation (5G) of mobile communications arises as a more disruptive evolution [1] to shift towards a *custom-fit* paradigm. Novel advances in radio technologies are achieved as well as improvements in the infrastructure by incorporating multiple technology domains that include physical and virtual computational, storage and networking resources spanning across several administrative domains. Moreover, network softwarization, network function virtualization, and network slicing are embraced to enable virtual and logically isolated networks (i.e., network slices) [2], tailored to the requirements of different vertical services, over the same and shared network infrastructure.

This new paradigm is arousing the interest of vertical industries (e.g., manufacturing, public safety, public infrastructures, transportation, and energy sectors) as the opportunity to truly implement a digital transformation in their industries, thus contributing to the introduction of novel services, cost savings, and improved time-to-market. However, it comes with the cost of extremely heterogeneous infrastructures belonging to several stakeholders, and vertical industries still lack overall knowledge about Information and Communications Technologies (ICTs) to manage them on their own. As such, to ease the adoption of 5G technologies by vertical industries, this complexity must be hidden. On one hand, abstracted and simplified interfaces that conceal unnecessary details and complexity are required, so that vertical industries mostly focus on their business domain. On the other hand, supporting tools that ease and automate the management of vertical services must be devised, as an alternative to the complex and very lengthy manual configuration and validation of traditional management tools.

This work addresses the aforementioned challenges by proposing a supporting platform, named 5G-DIVE Elastic Edge Platform (DEEP), as an add-on on top of distributed 5G-powered Edge and Fog infrastructures. It arises as the missing piece to bridge different technology domains in a unified *Intelligent Edge and Fog 5G End-to-End (E2E)* solution that offers abstracted services to the vertical industries. First, it hides the underlying complexity through vertical-oriented interfaces tailored to the needs of vertical industries. Second, it takes over of the lifecycle management of vertical services on behalf of the vertical industries, enforcing automation of their business processes. Ultimately, the proposed DEEP platform contributes to a faster adoption of 5G technologies by different stakeholders [3], noticeably the vertical industries.

The remainder of this paper is structured as follows. Section II analyzes the main challenges for the adoption of 5G technologies by the vertical industries. The DEEP platform and its innovations are proposed in Section III, of which an experimental validation is presented in Section IV. Its benefits to selected Industry 4.0 and Autonomous Drone Scouting use cases are identified in Section V. Finally, Section VI highlights the main conclusions, pointing out to future work.

II. CHALLENGES FOR THE VERTICAL INDUSTRIES

5G networks, Edge and Fog computing, and Artificial Intelligence (AI) are key enablers to support the digital transformation of vertical industries. This section presents their main shortcomings from the vertical industries perspective.

A. Private 5G Networks

The concept of private 5G network, referred as Non-Public Network (NPN), is introduced to accelerate the digital transformation of vertical industries. Several deployment options for

NPNs are envisioned [4]: *(i)* standalone NPNs, implementing fully private networks with dedicated and secure capabilities; and *(ii)* Public Network Integrated-NPNs (PNI-NPNs), leveraging on capabilities available at the Public Network (PN) to complement the NPN. Different degrees of integration are foreseen, which might depend completely or partially on the infrastructure and the functions of the PN. The implementation of PNI-NPN has been achieved by: *(i)* setting up dedicated network slices; *(ii)* integrating NPN as a non-3GPP access network; and *(iii)* setting up dedicated Access Point Names or Packet Data Network.

Challenge #1: How to ease the vertical industries to dynamically set up the desired PNI-NPN, especially when different functions and resources are controlled by a single domain or shared among them, and when different trust and business agreements are in place?

B. Edge and Fog Computing

ETSI Multi-Access Edge Computing (MEC) [5], Industrial Internet Consortium (IIC) (now incorporating OpenFog) and 5G-CORAL [6] are three reference architectural frameworks targeting Edge and Fog computing environments. Although each opted for different architectural approaches, they share a common characteristic: they all aim to provide a cohesive infrastructure-level abstraction in a distributed environment. While this is undoubtedly a cornerstone for building the Edge and Fog, this approach alone prevents the effective support to the vertical industries: *(i)* the low-level abstraction exposed by these systems; *(ii)* the static definition of application and service requirements; *(iii)* the black-box approach adopted in which the infrastructure is agnostic to the application logic; and *(iv)* the lack of integration with the underlying communication network.

Challenge #2: How to abstract the vertical industries from the complexity of a highly distributed infrastructure to ease their adoption in the vertical industries' business process?

C. Towards Intelligent and Automation

Artificial Intelligence (AI) and Machine Learning (ML) are promising alternatives to transit towards an intelligent and automated operation [7][8].

The ITU-T Focus Group on Machine Learning for Future Networks including 5G (FG ML5G) [9] and the ETSI Experiential Networked Intelligence (ENI) group [10] are both proposing alternative architectures for the integration of AI/ML techniques in the network management task with an emphasis on 5G and vertical services. Both proposals follow a service-oriented approach to enable network automaticity. Similarly, ETSI Zero-touch network and Service Management (ZSM) [11] explores different approaches to achieve automation. Another initiative is being developed by the O-RAN alliance [12] for the applicability of AI/ML techniques in the control of the Radio Access Network (RAN) radio functional split components, for real-time and non-real-time processing control of the radio components hosted at the O-RAN Distributed and Centralized Units. O-RAN promotes a key industrial trend towards a larger disaggregation of RAN system components, especially for lowering costs and covering the specific needs and demands of vertical industries for their deployments.

Challenge #3: How to conjugate the variety of data spanning across multiple technologies domains, algorithms and components into a single, operational platform permitting an automated network and service management on behalf of the vertical industries?

III. 5G-DIVE ELASTIC EDGE PLATFORM (DEEP): CONCEPT OVERVIEW

The previous challenges are perceived as a barrier to the adoption of 5G technologies by the vertical industries. As such, the 5G-DIVE Elastic Edge Platform (DEEP) platform (as presented in Figure 1) aims at supporting the vertical industries in day-by-day operations, management, and automation of business processes. In other words, it contributes to the development, execution, and management of vertical services, including the incorporation of intelligence capabilities, abstracted from the complexity of building, and maintaining the infrastructure associated with the delivery of the application.

When integrated into an E2E solution comprising different technology and administrative domains, the DEEP platform facilitates the exploitation of an *Intelligent Edge and Fog 5G E2E* solution by the vertical industries.

A. Supporting Strata

To accomplish its vision, the DEEP platform defines three supporting strata, which are detailed as follows.

1) Business Automation Support Stratum (BASS): simplifies and automates the creation and management of vertical services, being the logical entry point for the vertical industries and their OSS/BSS systems. The vertical industry delegates (part of) its business automation to this stratum, which takes over the deployment of E2E vertical services across multiple technologies and administrative domains, and ensures that SLAs and policies are met at each moment. The BASS comprises the following elements:

- **Vertical Service Coordinator:** handles vertical-oriented requests, translates them into network services, and coordinates their E2E deployment.
- **Vertical Service Blueprint Catalogue:** consists of a repository containing vertical services blueprints, providing a vertical service template with service-specific parameters.
- **SLA & Policy Management:** monitors the SLAs and policies for different vertical services and, in case of a violation, triggers corrective actions. It can be enhanced with intelligence and forecasting capabilities, so that violations are predicted and dealt preemptively.
- **Active Monitoring:** manages the deployment of monitoring probes required for performance monitoring.
- **External Federation:** handles the deployment of E2E vertical services across different administrative domains.

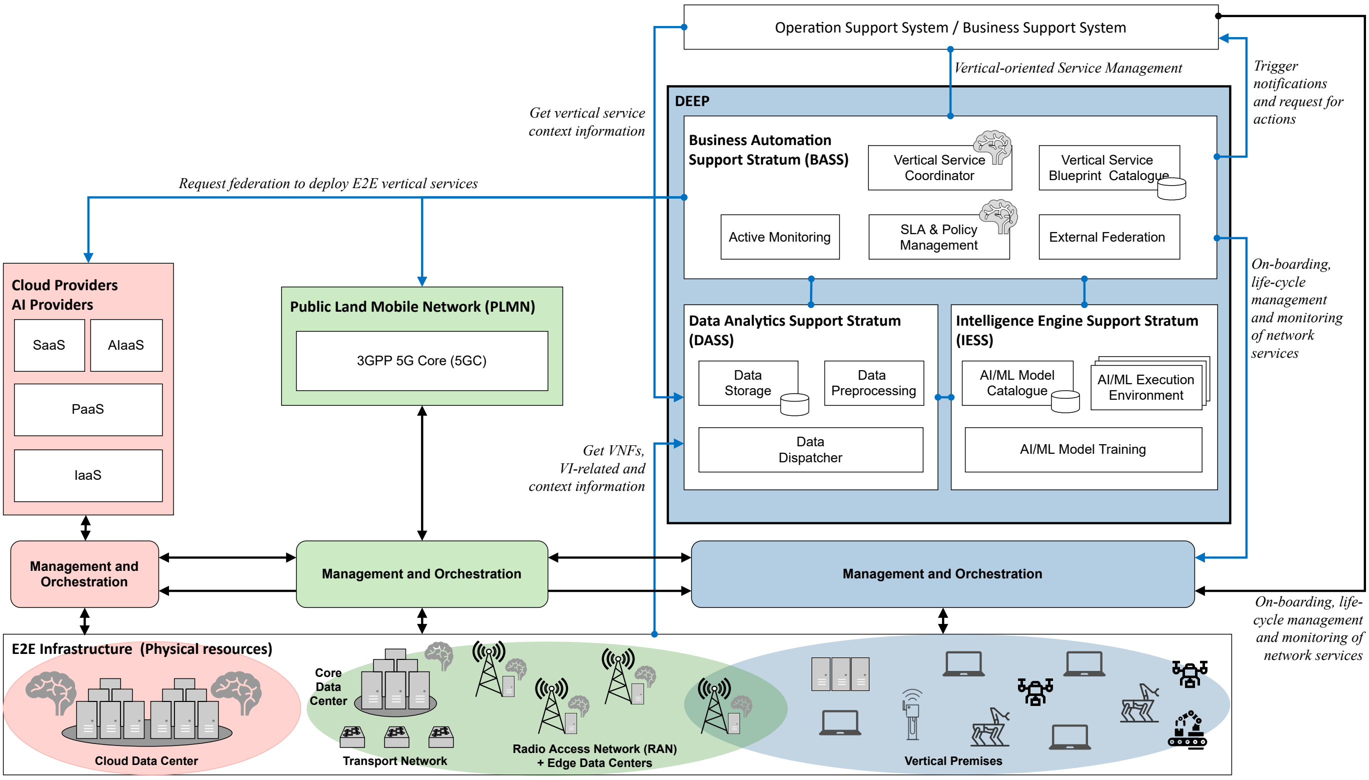

Fig. 1. DEEP Platform Overview in an E2E Solution

2) Data Analytics Support Stratum (DASS): eases gathering, pre-processing, storage, and sharing of data. Data is gathered from the vertical services operation and potentially enriched with a variety of context information, from both the virtualization infrastructure and the surrounding environment. The DASS comprises three main elements:

- **Data Dispatcher:** gathers data from the different sources and delivers it to all the subscribed entities, enforcing authorization mechanisms if required.
- **Data Pre-processing:** transforms raw data into an understandable and common format (e.g., cleaning, filtering, grouping, normalization, anonymization, and compression).
- **Data Storage:** stores gathered data. The decision whether to store the data is up to its source which, optionally, also provides its lifespan.

3) Intelligence Engine Support Stratum (IESS): facilitates the training and provisioning of intelligent models following an intent-driven approach and AutoAI solutions. Envisioned as an AI/ML engine, it leverages on data-driven AI/ML models to support complex decision-making systems, predict events, forecast demands, find patterns and anomalies, perform classification, among others. The IESS comprises three main elements:

- **AI/ML Model Catalogue:** consists of a repository of (un)trained AI/ML models made available by the DEEP platform. These are annotated with metadata describing their features, requirements, and suitable applications.
- **AI/ML Model Training:** handles the procedures for training an AI/ML model, from the selection of one or more AI/ML algorithms, their hyperparameters and training dataset, up to the cross-validation their accuracy.
- **AI/ML Execution Environment:** comprises the runtime environment to train the selected AI/ML model(s).

B. DEEP Innovations

The DEEP platform implements a set of innovations designed to abstract complexity from the vertical industries and to ease automation tasks.

1) Vertical Service Abstraction: The DEEP platform provides an abstraction layer, together with vertical-oriented interfaces, that enables the vertical industries to define their services solely based on their knowledge domain. Departing from available blueprints, they fill service-specific parameters related to their domain expertise, including SLAs, and privacy, computing, intelligence, and automation requirements, following an intent-driven approach.

2) AI/ML-based Intelligence Support: The DEEP platform comprises tools that facilitate the definition, online/offline training, and cross-validation of AI/ML models. Trained AI/ML models are then used for: *(i)* automated vertical service management (e.g., auto-scaling, self-healing and preemptive migration); and *(ii)* intelligent vertical applications (e.g., object recognition or movement prediction). Departing from the available catalog, the selection of suitable AI/ML models for a given intent is also automated using AutoAI-driven solutions.

3) Data Distribution and Unification: The DEEP platform defines a data distribution service to handle the scale at which data is produced and consumed, especially when an increasing number and heterogeneity of devices compose the computing, storage, and networking infrastructure. It unifies data in motion, data in-use, data at rest and computations through traditional publish/subscribe primitives and geographically distributed storage, queries, and computations. Other than benefits in terms of performance and efficiency, the vertical industries retain control over the privacy of their data by storing it closest to its source and making it accessible by any authorized entity.

4) Enhanced Monitoring: The DEEP platform extends the physical and virtual resource monitoring by integrating vertical-application level monitoring, provided directly by the applications themselves or collected through monitoring probes, to accurately assess the performance of the vertical services. The purpose of such monitoring data is three-fold: *(i)* train AI/ML models in the IESS; *(ii)* use as input for the automation tasks in the BASS; and *(iii)* provide additional insights to the vertical industries.

5) External Federation: The DEEP platform facilitates the integration of federated services and resources as part of the vertical industry domain. Additionally, secure, auditable, and on-demand federation mechanisms are defined using distributed ledger technologies (DLT). This has an increasing importance to enable the creation and management of E2E vertical services, which may involve multiple administrative domains (e.g., implementation of the selected PNI-NPN option, or the usage of third-party Cloud and Edge facilities to reduce costs).

6) Locality and Privacy: The DEEP platform exploits the locality offered by an Edge and Fog infrastructure to process and analyze sensitive data where it is generated, thus enabling strict privacy and low latency response. The DASS implements a low overhead and high throughput data distribution service designed to support low power networks and constrained devices, like those expected in the vertical premises. The IESS explores solutions such as federated learning and transfer learning that boost collaborative AI/ML without centralized training data, thus retaining privacy and locality of the private data. Moreover, local services such as those offered by ETSI MEC enhance contextual awareness via radio network, location, or any other information relevant to the changing environment.

IV. AUTOMATED VERTICAL SERVICE MANAGEMENT: PRELIMINARY VALIDATION

Leveraging on the initial code release of the DEEP platform, a proof-of-concept (Figure 2) centered on the *SLA & Policy Management* element of the BASS is developed, hereinafter referred to as *SLA Enforcer*. This component is an integral part of implementation of the DEEP platform to automate the lifecycle management of vertical services.

A. Proof-of-Concept Design

The *SLA Enforcer* implements the *SLA & Policy Management* capabilities as its core, integrating interfaces towards the *Vertical Service Coordinator* and *Active Monitoring* for, respectively, SLAs management and monitoring capabilities as well as interfaces towards the DASS, IESS and orchestrator.

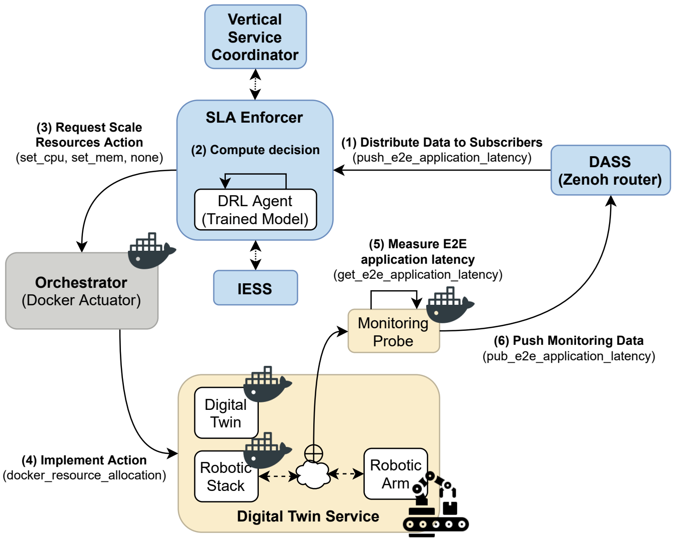

Fig. 2. *SLA Enforcer Closed-Loop Workflow*

A Digital Twin service (as defined in [13]) is the target of this proof-of-concept, to which the *SLA Enforcer* minimizes the resource allocation without incurring on the violation of its SLAs.

1) Integration with DEEP Platform: The *SLA Enforcer* is agnostic to the vertical service, thus possible to be applied to use cases in distinct vertical industries, including those described in Section V. However, their specificity impacts *(i)* SLAs definition; *(ii)* data to monitor; *(iii)* actions to take; and *(iv)* AI/ML model. During service creation, the DEEP platform configures the *SLA Enforcer* based on the vertical intent:

1) The *Vertical Service Coordinator* extracts any intent-based SLAs from the vertical-oriented descriptor, forwarding them to the *SLA & Policy Management.*
2) The *SLA & Policy Management* validates the intents, translating them into a set of policies that identify their scope, thresholds and validation data. If the intent validation fails, manual input from the administrator is required.
3) Based on the extracted policies, the *SLA & Policy Management* requests *(i)* an AI/ML model towards the IESS; and *(ii)* the deployment of monitoring probes towards the *Active Monitoring*.
4) Using AutoAI, the *IESS* selects the AI/ML model from its catalogue, performs its training and cross-validation, delivering the trained model to the *SLA & Policy Management*.
5) The *SLA & Policy Management* triggers a new instance of *SLA Enforcer* tied to the vertical service's SLAs.

Note that, the *SLA Enforcer* is not restricted to the set of actions described in this proof-of-concept. It not only supports vertical and horizontal scaling strategies for computing, storage, and network resources, but also other lifecycle management actions such as (sub-)service migration and hibernation, or even trigger Radio Access Technologies (RATs) handover.

2) Scenario and Testbed Setup: The *SLA Enforcer* implements a Deep Reinforcement Learning (DRL) agent based on the Deep Q-Network (DQN) algorithm. The reward function is computed based on the: *(i)* CPU and MEM consumption; *(ii)* CPU and MEM limits; *(iii)* E2E application latency; and *(iv)* SLA-inferred latency thresholds. Docker is used to orchestrate and manage the Digital Twin service (i.e., *Digital Twin* and *Robotic Stack*) through Docker containers. The Robotic Stack container is deployed in a Banana Pi Single Board Computer, which represents a constrained device available in the Fog, while the *Digital Twin* is deployed in a laptop representing the user equipment used by the worker. The *Robotic Arm* emulates the Niryo One Robot Manipulator, deployed through a virtual machine with 1 vCPUs and 2 GB of RAM in an Edge Server (Dell PowerEdge R430). The *Monitoring Probe* is deployed in a man-in-the-middle fashion between the *Robotic Stack* and the *Robotic Arm*. The DASS implements Zenoh[1] as its core. Finally, all nodes are connected via Gigabit Ethernet.

3) Workflow: The *SLA Enforcer* follows a closed-loop workflow to continuously minimize the resource allocation (as shown in Figure 2). For every decision, the *SLA Enforcer* fetches the E2E application latency computed by the *Monitoring Probe* (step 1). Hence, it verifies if *(i)* SLAs are fulfilled; and *(ii)* resources need to be scaled (step 2). Scaling decisions are forwarded to the Orchestrator (step 3), which enforces their implementation in the virtualization infrastructure (step 4). In this proof-of-concept, it configures the CPU and MEM limits that a given container can use. The workflow closes with the *Monitoring Probe* computing and publishing via DASS the E2E application latency (steps 5 and 6).

4) Experiment: The validation experiment is composed of a Digital Twin service, where the *Digital Twin* continuously sends commands to the *Robotic Arm* and the *Robotic Arm* sends its current pose to the *Digital Twin*. The *Robotic Stack* acts as the middleware bridging both the virtual and the physical instances, performing e.g. command validation and path computation. As such, the resources allocated to the *Robotic Stack* impacts the E2E application latency, which consists of the time between a command is issued by the *Digital Twin* and it is fully executed by the *Robotic Arm*. The *Monitoring Probe* inspects both flows to compute the E2E application latency. Data is published via the DASS and, sequentially, fed into the *SLA Enforcer*. The *SLA Enforcer* minimizes resource allocation by executing the necessary scaling actions. These actions include increasing or decreasing the CPU and MEM limits of the *Robotic Stack* container, carried out by using the Docker API. Two different tasks, performed through the *Digital Twin*, are considered: *(i)* precision tasks, with a strict E2E application latency requirement of 500ms; and *(ii)* screening tasks, where the E2E application latency can be relaxed up to 1000ms. In this proof-of-concept, the SLA defines the maximum latency that the E2E service must guarantee 95% of the time.

B. Dynamic Resource Allocation

Results of the learning process (Figure 3) show that, for different E2E application latency requirements, the *SLA Enforcer* learns the best values for the CPU and MEM that satisfy the vertical service's SLA.

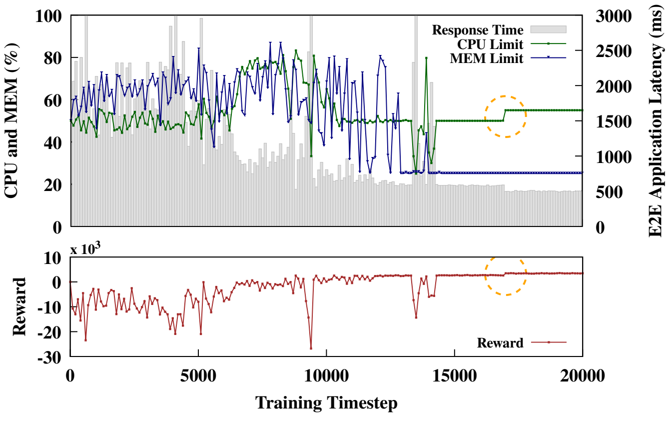

(a) Strict latency requirement (500ms)

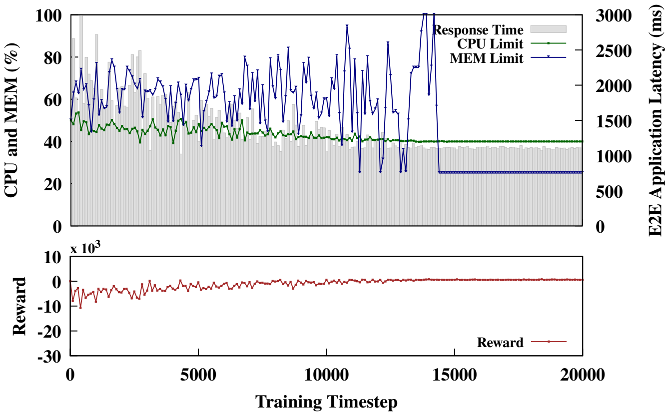

(b) Relaxed latency requirement (1000ms)

Fig. 3. Learning Process of the *SLA Enforcer*

The *SLA Enforcer* learns how to stabilize the E2E application latency below the specified requirement after the first 14500 training timesteps. Upon the learning process, the *SLA Enforcer* exploits the optimal CPU and MEM configurations since the overall reward is maximized. Nevertheless, the *SLA Enforcer* still explores other configurations aiming to increase its overall reward (e.g., timestep 17000 in Figure 3a), but such cases becomes more sporadic as the agent converges.

The optimal CPU and MEM limits are around 55% and 25% for Figure 3a, and 40% and 25% for Figure 3b. By defining a continuous action space, instead the current discrete action space with intervals of 5%, *SLA Enforcer* can adjust the limits closer to the real CPU and MEM usage.

C. Monitoring Data Distribution

Figure 4 shows a benchmark of the DASS in a peer-to-peer configuration deployed over a 10 GbE scenario, namely its performance in terms of messages per second and goodput for different payload sizes. Each data point represents the median of 10 runs.

The DASS can route around 1.1M messages per second for payload sizes lower than 512 bytes. For higher payload sizes, the bottleneck is mainly due to the link capacity which is achieved around 2048 bytes, as reflected in the goodput

---

[1] http://zenoh.io/

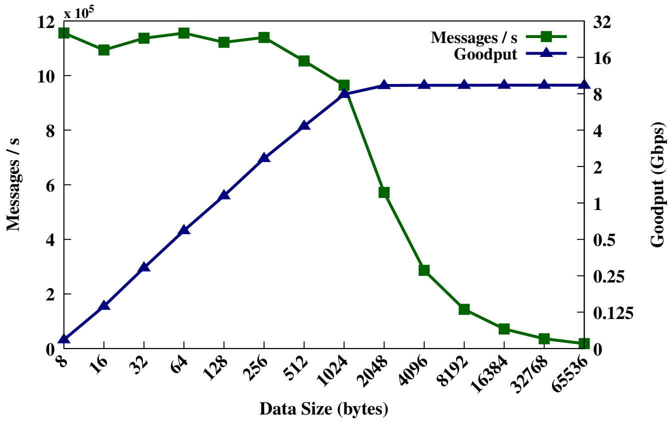

Fig. 4. DASS Benchmark

results. Finally, hardware accelerators or network offloading mechanisms can potentially increase the performance of the DASS.

V. Benefits of an Intelligent Edge and Fog 5G E2E Solution

In 5G-DIVE project[2], vertical pilots on manufacturing (i.e., Industry 4.0) and public safety sectors (i.e., Autonomous Drone Scouting) will be used to assess the technical merits and business value proposition of the *Intelligent Edge and Fog 5G E2E* solution enabled by the DEEP platform (which are presented in Table I). They comprise a wide range of requirements, including different types of 5G service profiles, indoor and outdoor scenarios, computing requirements and AI capabilities.

Industry 4.0 is shifting the business processes towards an increasingly connected infrastructure and smart manufacturing. 5G connectivity, cloud-to-things continuum, data analytics and automation are some of the key enablers of this transformation [14]. It includes the following use cases: *(i)* Digital Twin enables a virtual representation to act as a mirror of its real-world counterpart; *(ii)* Augmented Zero Defect Manufacturing improves the quality control processes and increases the flexibility to cope with mass customization; and *(iii)* massive Machine-Type Communications provides intelligent and scalable connectivity for dense heterogeneous with multiple RATs deployments.

Autonomous Drone Scouting focuses on the needs of the public safety sector to deploy emergency solutions in the aftermath of disaster situations [15]. Drones scout a disaster area to rapidly assess the real situations. It includes the following use cases: *(i)* Drones Fleet Navigation improves the success rate of rescue missions, where fleets perform horizontal flight while avoiding obstacles; and *(ii)* Intelligent Image Processing for Drones automates the detection of critical areas or emergencies.

VI. Conclusions and Summary

The environment providing 5G services for vertical industries is composed of a multitude of technology domains under different administrative domains, creating a unified system extremely complex to handle. In this work, the DEEP platform is proposed as the bridge to link several of these technology domains together, in the pursue of an *Intelligent Edge and Fog 5G E2E* solution. Leveraging on network and computing abstractions, AI/ML-based automation, and scalable and low overhead distribution data services, the DEEP platform aims at finding sustainable ways to abstract and handle the complexity related to the E2E deployment by the vertical industries, becoming an important piece in an increased number of scenarios to be targeted by 5G.

Within the 5G-DIVE project, the DEEP platform will set out as the common platform to address TRL 5/6 pilots on Industry 4.0 and Autonomous Drone Scouting. The outcome of these pilots might raise the need for other AI/ML usages within the DEEP platform, such as the generation of novel insights about the operation of vertical services, forecasting (e.g., traffic demands and mobility of resources), or prediction of anomalies.

Acknowledgment

This work has been (partially) funded by the H2020 EU/TW joint action 5G-DIVE (Grant #859881). We want to thank Javier Sacido, Matteo Pergolesi, Olivier Hécart, and Julien Enoch for their support in the proof-of-concept integration.

References

[1] M. Cave, "How disruptive is 5G?" *Telecommunications Policy*, vol. 42, no. 8, pp. 653–658, 2018.

[2] L. U. Khan, I. Yaqoob, N. H. Tran, Z. Han, and C. S. Hong, "Network Slicing: Recent Advances, Taxonomy, Requirements, and Open Research Challenges," *IEEE Access*, vol. 8, pp. 36 009–36 028, 2020.

[3] 5GPPP, "View on 5G Architecture - Version 3.0," Feb. 2020.

[4] J. Ordonez-Lucena, J. F. Chavarria, L. M. Contreras, and A. Pastor, "The use of 5G Non-Public Networks to support Industry 4.0 scenarios," in *5th IEEE Conference on Standards for Communications and Networking (CSCN 2019)*, Granada, Spain, Oct 2019.

[5] T. Taleb, K. Samdanis, B. Mada, H. Flinck, S. Dutta, and D. Sabella, "On Multi-Access Edge Computing: A Survey of the Emerging 5G Network Edge Cloud Architecture and Orchestration," *IEEE Communications Surveys Tutorials*, vol. 19, no. 3, pp. 1657–1681, 2017.

[6] D. Rapone, R. Quasso, S. B. Chundrigar, S. T. Talat, L. Cominardi, A. De la Oliva, P. Kuo, A. Mourad, A. Colazzo, G. Parmeggiani, A. Z. Orive, C. Lu, and C. Li, "An Integrated, Virtualized Joint Edge and Fog Computing System with Multi-RAT Convergence," in *2018 IEEE International Symposium on Broadband Multimedia Systems and Broadcasting (BMSB)*, Jun 2018.

[7] Y. Wu, "Cloud-Edge Orchestration for the Internet-of-Things: Architecture and AI-Powered Data Processing," *IEEE Internet of Things Journal*, 2020.

[8] D. M. Gutierrez-Estevez, M. Gramaglia, A. D. Domenico, G. Dandachi, S. Khatibi, D. Tsolkas, I. Balan, A. Garcia-Saavedra, U. Elzur, and Y. Wang, "Artificial Intelligence for Elastic Management and Orchestration of 5G Networks," *IEEE Wireless Communications*, vol. 26, no. 5, pp. 134–141, 2019.

[9] ITU-T, "ITU-T Y.3172 - Architectural framework for machine learning in future networks including IMT-2020," Jun. 2019.

[10] ETSI, "ETSI GS ENI 005 V1.1.1 - Experiential Networked Intelligence (ENI); System Architecture," Sep. 2019.

[11] ——, "ETSI GR ZSM 005 V1.1.1 - Zero-touch network and Service Management (ZSM); Means of Automation," May 2020.

[12] ORAN, "ORAN-WG2.AIML.v01.00 - AI/ML workflow description and requirements," Oct. 2019.

[13] L. Girletti, M. Groshev, C. Guimarães, C. J. Bernardos, and A. de la Oliva, "An Intelligent Edge-based Digital Twin for Robotics," in *2020 IEEE Globecom Workshop on Advanced Technology for 5G Plus (GC 2020 Workshop - AT5Gp)*, Taipei, Taiwan, Dec. 2020.

---

[2] 5G-DIVE Project - https://5g-dive.eu/

TABLE I
BENEFITS OF THE *Intelligent Edge and Fog 5G E2E* SOLUTION ENABLED BY THE DEEP PLATFORM

Use Cases	DEEP Platform	Edge and Fog Computing	5G Connectivity	Automation and Intelligence
Digital Twin	Makes available an Intelligent Edge and Fog 5G E2E solution, bridging virtual and physical network functions, RAN and mobile core elements, on-premise, edge, and cloud computing, and off-premise, metro and core transport resources. Enables a dynamic and on-demand lifecycle management of vertical services, according to their business needs, through automation and intelligence. Empowers monitoring capabilities to the vertical industries. Eases the creation and configuration of AI-based applications as part of vertical services. Optimizes data distribution services, including the Fog comprising low power and constrained networks and devices.	Virtualizes and offloads software components, resulting in lightweight, low cost, smarter and shared devices. Improves response times, saves bandwidth, and preserves locality and privacy.	Fulfills the strict connectivity requirements in terms of latency, reliability, availability required by time-sensitive tasks (e.g., real-time physical-digital synchronization).	Predicts the next states of the physical entity for a lag-free experience. Enables autonomous AI-based operation of the physical entity.
Augmented Zero Defect Manufacturing		Offloads object recognition tasks from the video cameras to close and more powerful nodes (e.g., AI-accelerators). Preserves locality and privacy.	Fulfills the strict connectivity requirements in terms of latency, reliability, availability and guaranteed throughput required for a extremely fast identification of defective pieces.	Detects defective pieces independently of their position. Learns and adapts to changes in the objects. Detects anomalies in the production line.
Massive Machine-Type Communications		Offloads baseband functions from the radio heads to virtualized software functions in the Edge, providing scalable and flexible connectivity to trillions of devices. Preserves locality and privacy.	Fulfills the strict connectivity requirements in terms of application layer end-to-end latency, link reliability and throughput required by specific RATs.	Detects traffic load, allowing dynamic scalability and availability with orchestration features.
Drones Fleet Navigation		Offloads the navigation and collision avoidance functions for efficient, safe, and accurate autonomous drone flight.	Fulfills the strict connectivity requirements in terms of latency, reliability, availability, and guaranteed throughput required for video and sensor data streaming.	Enables autonomously navigation using geolocation and image analytics capabilities.
Intelligent Image Processing for Drones		Offloads the virtualized object detection algorithm with GPU-enabled computing capability. Performs 2D aerial image stitching in real-time at the edge.	Enables the traffic breakout at the edge, reducing the latency between a drone and its serving applications.	Enables latency-aware disaster detection with a guaranteed real-time response.

[14] J. Cheng, W. Chen, F. Tao, and C.-L. Lin, “Industrial IoT in 5G environment towards smart manufacturing,” *Journal of Industrial Information Integration*, vol. 10, pp. 10 – 19, 2018.

[15] H. Shakhatreh, A. H. Sawalmeh, A. Al-Fuqaha, Z. Dou, E. Almaita, I. Khalil, N. S. Othman, A. Khreishah, and M. Guizani, “Unmanned Aerial Vehicles (UAVs): A Survey on Civil Applications and Key Research Challenges,” *IEEE Access*, vol. 7, pp. 48 572–48 634, 2019.

BIOGRAPHIES

Carlos Guimarães: Postdoctoral researcher at Universidad Carlos III de Madrid (Spain).

Milan Groshev: Ph.D. student at University Carlos III Madrid (Spain).

Luca Cominardi: Senior Technologist at ADLINK (France).

Aitor Zabala: Chief Technical Officer at Telcaria Ideas S.L. (Spain).

Luis M. Contreras: Technology expert at Telefónica I+D / Global CTIO (Spain).

Samer T. Talat: Researcher at Industrial Technology Research Institute (Taiwan).

Chao Zhang: Ph.D. student at Lund University (Sweden).

Saptarshi Hazra: Project-Leader / Researcher at Research Institutes of Sweden (Sweden).

Alain Mourad: Director Engineering R&D at InterDigital.

Antonio de la Oliva: Associate professor at University Carlos III Madrid (Spain).