Enhancing Autonomy with Blockchain and Multi-Acess Edge Computing in Distributed Robotic Systems

Jorge Peña Queralta *, Li Qingqing *. Zhuo Zou† and Tomi Westerlund *

* Turku Intelligent Embedded and Robotic Systems Lab, Faculty of Science and Engineering, University of Turku, Finland
† School of Information Science and Technology, Fudan University, China
Emails: {jopequ, qingqli, tovewe}@utu.fi, zhuo@fudan.edu.cn

***Abstract*—This conceptual paper discusses how different aspects involving the autonomous operation of robots and vehicles will change when they have access to next-generation mobile networks. 5G and beyond connectivity is bringing together a myriad of technologies and industries under its umbrella. High-bandwidth, low-latency edge computing services through network slicing have the potential to support novel application scenarios in different domains including robotics, autonomous vehicles, and the Internet of Things. In particular, multi-tenant applications at the edge of the network will boost the development of autonomous robots and vehicles offering computational resources and intelligence through reliable offloading services. The integration of more distributed network architectures with distributed robotic systems can increase the degree of intelligence and level of autonomy of connected units. We argue that the last piece to put together a services framework with third-party integration will be next-generation low-latency blockchain networks. Blockchains will enable a transparent and secure way of providing services and managing resources at the Multi-Access Edge Computing (MEC) layer. We overview the state-of-the-art in MEC slicing, distributed robotic systems and blockchain technology to define a framework for services the MEC layer that will enhance the autonomous operations of connected robots and vehicles.**

***Index Terms*—Edge Computing; Robotics; 5G; Computational Offloading; Multi-Access Edge Computing; Autonomous Robots; Network Slicing; Mapping and localization; Edge AI;**

I. Introduction

5G and beyond connectivity has the potential for bringing together the telecommunications, robotics [1], artificial intelligence (AI) [2], Internet of Things (IoT) [3] and blockchain domains [4], all of which share a recent trend in which computation is shifting towards more distributed architectures [5], [6], [7]. This comes together with the concept of network slicing and edge computing, key pillars behind the low-latency and network load optimization in 5G and beyond networks [8]. Through multi-tenant slicing, new business opportunities are being created at the edge of the network [9]. In this study, we provide a vision for the future of 5G-connected robots and vehicles, which will potentially benefit from this connectivity to increase their degree of autonomy and level of intelligence through services offered by third parties at the MEC layer.

We explore the potential for combining the backbone of today's autonomous robotic navigation and localization, the Robot Operating System [10], with the latest development in 5G and slicing strategies at the MEC layer. The MEC layer is an inherently distributed computing platform that enables high-performance computing (HPC) services with minimal latency [11]. The most direct application is to extend existing offloading schemes [12], and integrate them within the 5G stack [13]. This has clear potential in vehicular and robotic navigation, especially when combined with predictive schemes [14]. In addition, we envision that FPGA-based and CGRA-based hardware accelerators at the base stations will provide new levels of reconfigurability, energy efficiency, and processing power within the offloading orchestrators. Moreover, we take into account integration between distributed robotic systems [15], and distributed computation platforms defined within a blockchain [16]. We argue that permissioned blockchains backed by a large public and trusted infrastructure will be a key element of MEC-based services. These blockchains will be able to provide a transparent and secure channel for connected vehicles to interact with third parties.

Slicing at the MEC layer in 5G and beyond can reduce the computational load in connected robots and vehicles. This will allow units with more constrained resources, such as delivery drones, to enhance their situational awareness and increase their autonomy. In terms of safety and reliability in long-term autonomous operation in both self-driving vehicles and autonomous robots, challenges arise from the point of view of (1) localization accuracy [17], (2) situational awareness and level of understanding of the environment [18], and (3) limitations of computational capabilities in smaller robots or drones, with algorithms that might take longer to run depending on the complexity of the environment [19]. Slicing at the MEC layer has potential for providing services to support the operation of connected autonomous vehicles and robots by providing in respect to the above challenges (1) streaming services of high definition (HD) maps for accurate localization with online updates whenever the environment changes; (2) semantic information of the environment, as well as metadata from other connected vehicles; and (3) an adaptive algorithm that autonomously provides in real-time map models and environment data according to the operational and computational capabilities of the vehicle requesting data.

A. Blockchain at the MEC Layer

The integration of Blockchain with slicing at the MEC layer has recently been proposed by different researchers [20], [21], [22]. Nevertheless, we focus on the use of Blockchain to enable services at the MEC layer for autonomous robots and vehicles. Rather than focusing on data integrity and security,

we see that distributed consensus mechanisms in a blockchain are ideal for managing the MEC hardware, services and client-provider interaction. At the same time, the utilization of blockchains in the robotics field has recently shown its potential [6]. We extend it as a framework for integrating external services into distributed robotic and vehicular systems. One of the key challenges in the utilization of a blockchain for applications requiring real-time communication and data processing is the scalability, as indicated by [23], [24]. However, a more wider adoption of blockchain technology across multiple fields, specially the robotics and automation field, is expected with the arrival of low-latency and high-throughput next-generation blockchain networks [6].

B. Supporting Autonomy in Smart Cities

The concept of Smart City has been mostly tied to the IoT since its inception [25]. Nonetheless, the IoT and the robotics domain have since been integrated as connected robots become the standard. The new edge and fog computing paradigms have only increased this synergy between the two domains [26]. Nam *et al.* surveyed the early works on the topic and defined three fundamental dimensions of a Smart City: technological, human and institutional [27]. Self-driving cars or autonomous delivery robots extend the original concept of smart city from passive technology (mainly sensors) to active participators in an increasingly more complex cyber-physical dimension. Nonetheless, little attention has been put on the role of the institutional dimension towards a more widespread penetration of autonomous robots and vehicles in Smart Cities. In a recent work, we argue that institutions and public infrastructure can play a key role in enabling collaboration between autonomous robots with a blockchain [6]. We extend this idea with the introduction of MEC slicing. The key areas of the system are illustrated in Figure 1.

Since the introduction of MEC, network slicing has been seen as a key enabler of future autonomous vehicles [28]. However, the definition of architectures and the specifications of slices have been made mostly from the point of view of the telecommunications domain, taking into account network requirements [29]. In a recent work, Mei *et al.* have proposed an intelligent network slicing framework with differentiated slices for (i) traffic safety, (ii) autonomous driving, (iii) infotainment and vehicular internet, and (iv) service slice that manages the previous ones and measures quality of service [30]. We believe that efficient slicing requires different granularity within the autonomous driving concept, and therefore, we propose a slicing architecture that takes into account both the network requirements and the type of computing resources utilized for different aspects of autonomous driving, with a clear differentiation between offloading services and streaming services. In addition, we replace the top managing slice in [30] with a blockchain-based orchestrating slice that provides a framework for interaction between clients and service providers. Finally, by separating the service orchestration from the data channels and the actual services, which are hosted in their corresponding slice, we argue that the robustness of the system

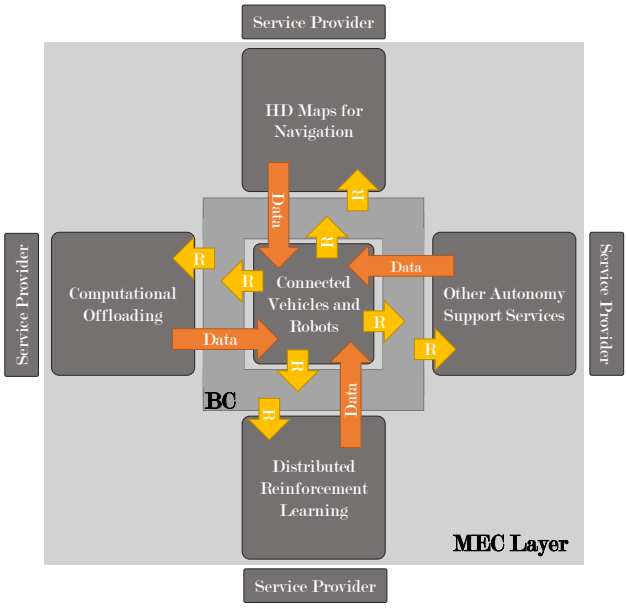

Fig. 1. Autonomy Support Services at the Multi-Access Edge Computing (MEC) layer. Requests (small yellow arrows) go through the blockchain slice (BC), which is in charge of the local service orchestration. Data is streamed directly to end-users to reduce latency and increase throughput.

can be increased, where a failure in any of these slices does not have a significant impact on the performance of the others in the short term. To the best of the authors' knowledge, this is the first work to focus on the point of view of algorithmic requirements for slicing, in comparison to a more traditional focus on network requirements. Moreover, we discuss on the key role that public infrastructure can play in ensuring the fairness and transparency of MEC-based services.

The rest of this paper is organized as follows. In Section II, we introduce the basics of MEC, network slicing, permissioned blockchains and algorithms for robotic navigation. Section III overviews the opportunities for utilizing the Blockchain technology, as a framework to manage services and resources with MEC slicing, together with distributed robotic systems. In Section IV, we explore the potential applications of the proposed approach in Smart Cities, from the streaming of local high-definition maps enabling accurate localization to offloading services to enhance the capabilities of robots with more limited resources. Section V then discusses the challenges and opportunities of the proposed architecture, with an emphasis on the viability and scalability of integrating blockchain technology. Finally, Section VI concludes the work and outlines future research directions.

II. Background and Significance

In this section, we briefly describe the main technologies that are considered in this study: distributed computing platforms at the MEC layer and blockchain technology.

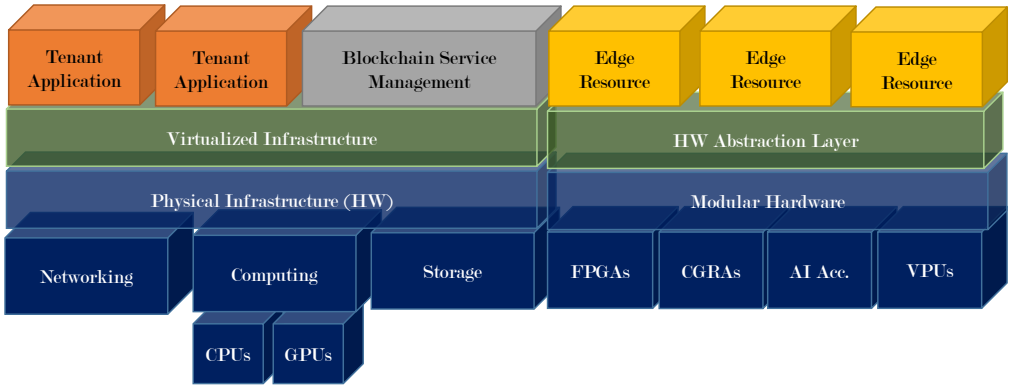

Fig. 2. 5G-MEC Computational Building Blocks

A. Multi-Access Edge Computing and Network Slicing

The standardization of Multi-Access Edge Computing (MEC) has been promoted by the European Telecommunications Standards Institute (ETSI) [31], with the MEC Industry Specification Group (ISG) launched at the end of 2014. The ETSI MEC ISG aims at defining a multi-tenant distributed cloud platform to be located at the edge of the radio access network (RAN) [32]. Moving computation and data intensive tasks towards the edge of the network enables the low-latency and high-bandwidth requirements of 5G and beyond connectivity. Other fundamental technologies towards this end include containerization and virtualization, software defined networking (SDN), and network function virtualization [28].

One of the key pillars enabling multiple verticals within MEC, and opening the RAN edge to a wide variety of industries and users, is network slicing [33]. Network slicing consists of the co-existence of multiple logical software-defined networks (slices) on a common hardware infrastructure, i.e., a multi-tenant cloud infrastructure at the edge of the network, with each of the slices being optimized to meet the requirements of a particular application [34]. We are particularly interested in slicing for the automotive sector, where 5G will be the key in vehicle-to-everything communication [35].

Since its early developments, MEC monetization has been a central topic of discussion [28]. Blockchain can provide a framework to democratize the monetization and utilization of the MEC layer as a platform to offer automation services to connected vehicles or robots.

From the point of view of security, a recent report from the European Union Agency for Cybersecurity (ENISA) on the thread landscape for 5G networks has identified numerous threats [36]. A blockchain can directly provide a higher level of resilience against multiple of these, such as authentication traffic spikes, manipulation of network traffic, malicious diversion of traffic, among others. Our focus is on utilizing a blockchain as a transparent and distributed framework to achieve consensus in terms of MEC resource provisioning. This has a direct impact on preventing threats such as abuse of third party hosted network functions, manipulation of the network resources orchestrator, or opportunistic and fraudulent usages of shared resources, among others.

In summary, MEC offers the benefits of cloud computing at the edge of the network, with the potential to offer new customer experiences. MEC allows for more scalable applications and network infrastructure by ensuring that raw data is processed at the edge, and only the resulting metadata is transmitted over to central cloud servers or other clients, optimizing the network load and reducing unnecessary traffic with information that does not need to be stored.

B. Blockchain and Distributed Ledgers

Blockchain platforms can be classified into two main types, permissionless and permissioned [37]. These can also be denominated public (permissionless), and consortium or private blockchains (permissioned). In public blockchains, there is no authority and all nodes are equivalent. In consortium or private blockchains, there are trusted authorities or nodes in charge of validating transactions [38].

One of the most famous and successful blockchains to date is Hyperledger, a project initiated in 2016 within the Linux Foundation [39]. Hyperledger is a permissioned blockchain which has been successfully utilized in multiple industrial domains [40]. The objective of Hyperledger is the deployment of an open-source and cross-industry framework that can be utilized as a standard platform to run smart contracts within a decentralized ledger.

Consortium blockchains, and Hyperledger in particular, have key advantages that have an impact over business networks such as a MEC-based blockchain [41]: (i) all participants have known identities, and therefore data protection laws can be applied accordingly, with permission required in order to join the network; (ii) data partitioning through channels, ensuring that data is available on a need-to-know basis and is only transmitted to the parties that need; (iii) scalability and adaptable levels of trust, with endorsement policies defining the number and nature of validators required to verify a given transaction, and subsequent network load optimization; (iv) a modular architecture, where identities and other components can be easily extended to meet the various requirements of

third parties or public authorities; and (v) flexible and complex queries over the ledger, simplifying the auditing process.

The data partitioning scheme seamlessly combines with the multi-tenancy at the MEC layer, with different service providers not being necessarily aware of other's transactions and data, even if this information is encrypted, as it would happen in a public blockchain. The endorsement policies can be optimized to manage scalability in large networks, defining validators depending on the client's location or the current network load. Then, alternating nodes across the network could be validated in different locations at different times.

The smart contracts within Hyperledger can be leveraged to manage the interaction between clients and service providers. First, by validating transactions and services before data is actually exchanged. Second, by provisioning and reconfiguring the computing resources within the MEC layer that are dedicated to that given service, optimizing the servers' load to improve users' experience. The data itself does not necessarily go through the blockchain once the service has been approved.

We see the main opportunities as part of smart cities, where the blockchain can be supported by either RAN or 5G-connected public infrastructure, as well as industrial environments where there exists trust. Figure 2 illustrates a generic 5G-MEC architecture with a blockchain to manage services (tenant applications) and edge resources (reconfigurable and on-demand hardware). In a smart city, a public blockchain for data sharing could boost the deployment of autonomous robots from both private and public entities. Besides, the role of the infrastructure should be considered not only as a platform to manage the blockchain lifecycle but also as a static data source and validating platform, where traffic cameras and other sensors that already exist can be integrated.

III. Blockchain-based services at the MEC layer

The integration of blockchain technology at the MEC layer has been proposed by multiple authors [20], [21], [22]. Nonetheless, these have been mostly focusing on the blockchain as a secure way of sharing or data or an immutable ledger to store transactions. However, one of the key applications of blockchains is their utilization as a robust decentralized computer that ensures the validity of execution of pieces of code called smart contracts [16]. We exploit these and the consensus protocols of blockchains to provide a framework for managing edge resources and services.

A. Previous works

Xiong *et al.* proposed the utilization of edge services to offer resource-constrained devices opportunities to join a blockchain by mining at the edge. Then, the end-devices share the data with third-party applications through a pricing scheme, modeling the interactions within the IoT as market activities. While their focus is on utilizing blockchains as a cryptocurrency and auditable platform, we focus on the distributed computation that smart contracts enable instead.

Liu *et al.* presented a similar approach, where the MEC layer is used to offload mining operations [42]. Nonetheless, this was part of a wider offloading framework where the focus was on deciding which offloaded operation would be cached. A similar scheme can be integrated within the offloading slice proposed in this study.

Zhu *et al.*'s EdgeChain is the closest work to this work [43]. EdgeChain is a blockchain-based architecture that is utilized to place third-party applications across the MEC. We extend this idea for dynamic reconfiguration with smart contracts based on client-provider interactions, rather than considering the service providers only.

B. Managing MEC with Consortium Blockchains

A consortium blockchain such as Hyperledger deployed across the MEC layer and connected public infrastructure, which are the nodes acting as validators, can be utilized to manage the interaction between connected clients and service providers, and at the same time orchestrating the hardware resources at the MEC layer. The proposed system architecture is illustrated in Fig. 3

We envision the existence of at least three separate network slices in order to support the aforementioned services. On one side, efficient offloading can be achieved with the on-demand reconfiguration of hardware accelerators, as well as AI accelerators. This slice focuses on low-latency in terms of fast data processing and optimization of computing resources to support as large number of connected devices as possible.

Some key benefits of this architecture are (i) identities are managed directly by the consortium blockchain and all transactions are signed and immutably recorded; (ii) the distribution of hardware resources or computing power is done through smart contracts and agreed across the MEC layer, with blockchain-enabled mobility; and (iii) a connected client, such as an autonomous robot or a self-driving car, requests a service from a third-party through the blockchain slice; if the smart contract approves the service, then the corresponding resources are configured and provisioned at the corresponding slice.

C. Distributed Robotic Systems for blockchain-based services

So far, the proposed architecture considers a distributed network architecture with distributed consensus algorithms. The last part of the piece is a distributed framework for deploying distributed robotic systems and algorithms to provide services to connected vehicles and robots. The Robot Operating System (ROS) has been the de-facto standard in production-ready robotic development for the past years [10]. However, wider adoption requires several challenges to be solved, including automated node discovery, real-time systems, non-ideal networks and distributed multi-robot deployments. These and other use cases are being developed within ROS2 [15]. We believe that ROS2 will be an essential piece in connected vehicles and robots by providing a common framework and standardization to the MEC-based services. ROS2 can solve key challenges in flexible service definitions at the MEC layer. It will provide standardization of data formats, channels and deployment of algorithms, with a common underlying logic for all service providers as will be discussed in the next section.

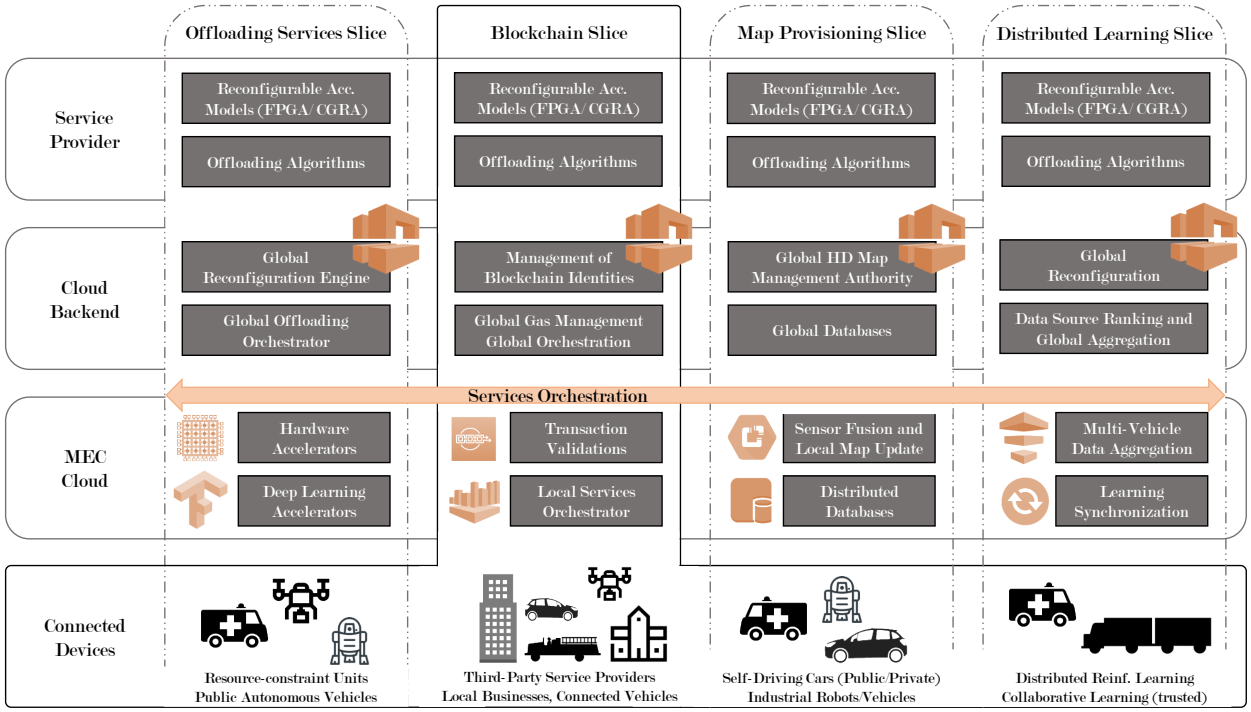

Fig. 3. Architecture for Blockchain-Based MEC Autonomy Services

IV. MEC for Autonomous Robotic Operation in Smart Cities

The proposed architecture can be utilized as a general framework to provide services through MEC slicing. Nonetheless, we focus on how this architecture can be integrated to support the autonomous operation of autonomous robots and vehicles, opening the door to new applications and business opportunities in cities or areas which support such seamless integration of third-party services within the telecommunications network stack. We also outline the role of ROS2 and the blockchain in the different application scenarios.

A. Provision of HD maps in real time

Autonomous navigation in dense urban areas requires self-driving cars and more diverse autonomous robots to have the ability to localize themselves with very high-accuracy. With the current state-of-the-art, this is only possible utilizing high-definition (HD) maps of the environment [44]. However, HD maps are expensive to generate, update and maintain [45]. In terms of generation and real-time updates, the proposed architecture can smoothly integrate within the streaming slice a data fusion module that gathers information from different sources to obtain these maps, as described below. Regarding the maintenance, the main disadvantage of HD maps is that they require large storage on-board the vehicles, and it is impractical to keep maps of large areas within vehicles themselves. An evident solution is to provide streaming services at the MEC layer. However, this requires tight mobility and latency control. We believe that this can be achieved with the combination of ROS2 as a standardization framework, 5G and beyond networks for low-latency and predictive mechanisms to provide in advance data about the areas that robots or vehicles will travel through. ROS messages serve as a standard for data formats, which can be then processed by multiple third parties without requiring extra communication to instruct on the data structure. Moreover, the nature of ROS topics enables consistent integration within streaming services to provide HD maps, with subtopics defining various parameters, e.g. location or point cloud density.

B. Online Update of Local HD Maps

In Smart Cities, administrators can provide a framework to support the online update of HD maps and utilize existing infrastructure as a source of data. Traffic cameras and other sensors utilized for monitoring can be repurposed and their data forwarded to a data fusion scheme within a dedicated MEC slice. Moreover, connected infrastructure with enough processing power can serve to increase the range or capacity of the streaming network. The role of ROS open source algorithms is essential to deploy the state-of-the-art in multi-source and multi-sensor data fusion in public infrastructure. Moreover, the smart contracts within the blockchain can be utilized to rank the available sources of data.

By providing an open framework, city administrations open the door to new local applications such as drone delivery or various types of autonomous robots surveying, monitoring and

performing other tasks across the city. This has the potential to boost both the city's economy and technology innovation.

C. Distributed Reinforcement Learning

As the robotic field has evolved over the past two decades, deep learning has become an essential aspect in complex robotics systems [46]. In particular, reinforcement learning has allowed for traditional dynamics models to be replaced for neural networks that have been able to outperform any previous approaches.

With the first semi-autonomous cars roaming the roads of large cities around the world, humongous amounts of data are being collected to improve the performance of deep learning algorithms. This is a process that requires offline training of neural networks. However, various distributed reinforcement learning algorithms enable robots and autonomous vehicles to have online improvements of their models not only from the real-time data and experiences but also from those of cooperating vehicles.

Offering a distributed reinforcement learning service at the MEC layer would enable connected vehicles to take advantage of the data and experiences of other vehicles to learn faster and better, with more and different experiences being analyzed in shorter periods of time. Nonetheless, such as a service would require a tight control on identities and a mechanism to ensure that model updates are valid and do provide an improvement. A permissioned blockchain provides a transparent and secure identity management framework, while the robustness and vulnerabilities in distributed multi-agent reinforcement learning is still an open problem [47]. If raw data is provided to the learning service, then the model updates can be validated. If data is protected due to privacy concerns and only the model updates are shared, then it becomes considerably more challenging to determine the validity of a given update. A blockchain can provide part of the solution to this problem through its consensus mechanisms. They have been shown to outperform traditional consensus mechanisms in the presence of erroneous or malicious data in other scenarios within the robotics domain [24].

D. Offloading Services

Reliable connectivity and existence of MEC services in a large area opens the doors to robots and vehicles relying on network-enhanced intelligence for their operation. Instead of developing and building complex robotic systems able of long-term self-supported autonomy, local organizations and businesses can build simpler products with similar capabilities, relying on computational offloading to achieve certain functionalities. Not only does this reduce the development and production cost, it also potentially decreases time-to-market, further boosting innovation.

We propose a separate slice for the offloading orchestrator and services because the focus is on optimizing computing power and reducing execution time when possible, compared to the storage and mobility requirements of the streaming slice. Even if the underlying hardware is the same, a different degree of reconfiguration is expected in order to optimize the offloading scheme.

GPU-based and FPGA-based accelerators have been widely used in sensor development and deep learning acceleration over the past few decades, being a perfect match for the requirements of edge computing [48]. More recently, autonomous navigation, localization and mapping algorithms have started to use FPGA-based implementations for real-time matching of HD maps or odometry [49]. Some of these operations are inherently parallelizable, and therefore FPGA-based accelerators have the potential for decreasing the latency by several orders of magnitude. In the proposed architecture, we envision that dynamic reconfiguration of FPGA-based hardware accelerators will play an important role in optimizing edge resources for computational offloading, increasing the number of nodes that can be supported and reducing the execution time of different processes.

ROS services can be directly utilized in offloading schemes, where third party services simply offer these to the network. In addition, there has been a recent interest in developing ROS-compliant accelerators to match the rising computational needs [50]. Having reconfigurable hardware available on-demand at the edge can help third-party service providers integrate these solutions. The reconfiguration and provisioning of resources can then be made through smart contracts executed in the MEC-hosted blockchain. Hardware accelerator models can be naturally abstracted in terms of the number of processing units or resources required, and multiple models can co-exist within a single chip. Finally, data partitioning schemes at the blockchain level and its modular architecture with the aforementioned concepts put together an efficient, open and flexible framework for offering offloading services.

E. Security Concerns

The blockchain is a key piece in the proposed architecture as a source of trust. We consider that the main security concerns in a MEC service framework is not the exposure of data but instead its validity and reliability. This is exemplified by threats identified by ENISA such as the manipulation of the network resources orchestrator (unreliable orchestration or invalid data regarding the resource orchestration) [36]. As we are discussing services that support the operation of autonomous vehicles, we need to take into account that this is a safety-critical application scenario where sending wrong data to a connected vehicle or robot might put in danger pedestrians and drivers. While the blockchain is not able to provide a robust way to validate data by itself, a ranking of the different identities offering services can be created and updater in real-time. Moreover, by implementing the resource orchestration and management of edge resources with smart contracts, the reliability of services providing mission-critical data can be kept under tighter control. Finally, the immutability of the transaction record can be utilized as a posteriori to assign liability and ensure accountability.

TABLE I
CHARACTERIZATION OF MEC-BASED SERVICES THAT SUPPORT THE AUTONOMOUS OPERATION OF CONNECTED VEHICLES AND ROBOTS.

	Critical Parameter		
	Latency	Throughput	Identity/Trust
Offloading Services	✔	✔	
Map Streaming		✔	
Distributed Learning			✔

V. DISCUSSION

The architecture and slicing strategy defined in the previous section are based on the characterization of different services that support autonomous operation of connected robots and vehicles. This characterization is illustrated in Table I from the point of view of the data and network parameters. Offloading services require low-latency and high-throughput data transmission in order to ensure full operational safety and high levels of performance. Latency is not so critical in the streaming of high-definition maps, however, if this is done such that a map of a large enough area around the vehicle is sent at a time. Finally, in distributed learning the amount of data is not critical (as raw data is not shared). Moreover, low-latency is not required because new data provided to vehicles is not critical to their operation (only improves their performance) and the learning process is long. Nonetheless, identity management and trustability are key parameters as erroneous or malicious updates to a model could have a significant impact on performance and operational safety.

The rest of this section describes the main challenges that emerge from the integration of blockchain for real-time services thar require large amount of data exchanges: scalability and storage. Nonetheless, one key point in the proposed architecture is that not all data goes through the blockchain. The purpose of the blockchain is to offer a transparent and reliable resource orchestrator, and as such service requests from connected vehicles are managed by the blockchain. However, once edge resources have been provisioned for a service and a request validated and accepted, the service is provided outside of the blockchain. Therefore, the data that supports the autonomous operation is not stored in the blockchain. In summary, it would not be very different from keeping a distributed database with all service requests.

A. Scalability

Recent advances in blockchain technology show promising results and potential for scalable and low-latency blockchain networks. Luu *et al.* presented Elastico, where sharding in a permissionless blockchains was explored [51]. Sharding is a technique that allows for distributed consensus in a network where nodes are divided in subnetworks or committees. Rather than processing and confirming all transactions globally across the network (for example through a majority consensus), each committee is in charge of processing a disjoint set of transactions, also denominated shard. In Elastico, researchers demonstrated the first sharding protocol that is secure in the presence of byzantine adversaries. Kokoris *et al.* introduced OmniLedger [52], a decentralized and secure ledger that scales linearly with the size of the network and supports transaction confirmation times of under two seconds, potentially being able to match credit card standards in terms of transaction confirmation response time with a large enough network, compared to an average transaction confirmation time (block validation) of around ten minutes in the case of Bitcoin. While Elastico scales almost linearly with the available computation power,OmniLedger does so with the number of validators.

Regarding the scalability of Hyperledger and its channel model, the initial versions did not have a truly scalable performance. However, this has improved considerably sinve Hyperledger Fabric v1.1.0 [53]. In terms of scaling the number of channels, this has shown little performance impact with low to no degradation so far [53].

B. Storage

One of the key concerns when utilizing a blockchain as a service management framework is the exponential storage that will be required along time. While this is a significant factor to take into account with early networks such as Bitcoin or Ethereum, next-generation blockchains have tackled this issue. Omniledger introduced state-blocks decreasing storage costs [52]. More recently, RapidChain achieved a storage saving factor of over 5x when compared to Omniledger, and 16x when compared with Elastico and Bitcoin-like blockchains [54]. In any case, the storage needs of a blockchain and its impact on performance can be controlled by truncating it or defining a fixed lifecycle. While this is a challenge in permissionless open blockchains, a strategy can be defined by public authorities for a permissioned blockchain such as Hyperledger if the infrastructure being used to validate transactions depends on the same public authorities.

VI. CONCLUSION

We have presented a system architecture for offering autonomy support services at the MEC layer in 5G and beyond networks. Our architecture combines network slicing with blockchain technology and distributed robotic systems. In particular, we envision a future of connected vehicles and robots where higher degrees of autonomy will be achieved through third-party services at the MEC layer. These services and the provisioning of MEC resources can be managed with permissioned blockchain networks such as Hyperledger, where offloading orchestrators and streaming orchestrators are implemented through smart contracts.

In future work, we will utilize ROS2 and Hyperledger to provide a proof-of-concept for the proposed architecture, and the bottlenecks and limitations will be analyzed.

ACKNOWLEDGEMENTS

The work leading to these results has been carried out as part of the AutoSOS project (Academy of Finland grant number 328755), the NSFC grant No. 61876039, and the Shanghai Platform for Neuromorphic and AI Chip (NeuHeilium).

REFERENCES

[1] M. Dohler *et al.* Internet of skills, where robotics meets ai, 5g and the tactile internet. In *2017 European Conference on Networks and Communications (EuCNC)*, pages 1–5. IEEE, 2017.

[2] R. Li *et al.* Intelligent 5g: When cellular networks meet artificial intelligence. *IEEE Wireless communications*, 24(5):175–183, 2017.

[3] MR. Palattella *et al.* Internet of things in the 5g era: Enablers, architecture, and business models. *IEEE Journal on Selected Areas in Communications*, 34(3):510–527, 2016.

[4] J Backman *et al.* Blockchain network slice broker in 5g: Slice leasing in factory of the future use case. In *2017 Internet of Things Business Models, Users, and Networks*, pages 1–8. IEEE, 2017.

[5] F. Voigtländer *et al.* 5g for robotics: Ultra-low latency control of distributed robotic systems. In *2017 International Symposium on Computer Science and Intelligent Controls (ISCSIC)*, pages 69–72. IEEE, 2017.

[6] J. Peña Queralta *et al.* Blockchain-powered collaboration in heterogeneous swarms of robots. In *Frontiers in Robotics and AI*, 2020.

[7] X Wang *et al.* In-edge ai: Intelligentizing mobile edge computing, caching and communication by federated learning. *IEEE Network*, 33(5):156–165, 2019.

[8] RA. Addad *et al.* Towards modeling cross-domain network slices for 5g. In *2018 IEEE Global Communications Conference (GLOBECOM)*, pages 1–7. IEEE, 2018.

[9] S. Husain *et al.* Mobile edge computing with network resource slicing for internet-of-things. In *2018 IEEE 4th World Forum on Internet of Things (WF-IoT)*, pages 1–6. IEEE, 2018.

[10] M. Quigley *et al.* Ros: an open-source robot operating system. In *ICRA workshop on open source software*. Kobe, Japan, 2009.

[11] D. Sabella *et al.* Mobile-edge computing architecture: The role of mec in the internet of things. *IEEE Consumer Electronics Magazine*, 5(4):84–91, 2016.

[12] L. Qingqing *et al.* Offloading Monocular Visual Odometry with Edge Computing: Optimizing Image Compression Ratios in Multi-Robot Systems. In *The 5th ICSCC*, 2019.

[13] K. Zhang *et al.* Energy-efficient offloading for mobile edge computing in 5g heterogeneous networks. *IEEE access*, 4:5896–5907, 2016.

[14] K. Zhang *et al.* Mobile-edge computing for vehicular networks: A promising network paradigm with predictive off-loading. *IEEE Vehicular Technology Magazine*, 12(2):36–44, 2017.

[15] A. Bareiś *et al.* Robots that sync and swarm: A proof of concept in ros 2. In *2019 International Symposium on Multi-Robot and Multi-Agent Systems (MRS)*, pages 98–104. IEEE, 2019.

[16] G Wood *et al.* Ethereum: A secure decentralised generalised transaction ledger. *Ethereum project yellow paper*, 151(2014):1–32, 2014.

[17] HG. Seif *et al.* Autonomous driving in the icity—hd maps as a key challenge of the automotive industry. *Engineering*, 2(2):159–162, 2016.

[18] J. Hong *et al.* Rules of the road: Predicting driving behavior with a convolutional model of semantic interactions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8454–8462, 2019.

[19] D. Floreano *et al.* Science, technology and the future of small autonomous drones. *Nature*, 521(7553):460–466, 2015.

[20] Z. Xiong *et al.* When mobile blockchain meets edge computing. *IEEE Communications Magazine*, 56(8):33–39, 2018.

[21] MDA Rahman *et al.* Blockchain-based mobile edge computing framework for secure therapy applications. *IEEE Access*, 6, 2018.

[22] Y Dai *et al.* Blockchain and deep reinforcement learning empowered intelligent 5g beyond. *IEEE Network*, 33(3):10–17, 2019.

[23] EC Ferrer. The blockchain: a new framework for robotic swarm systems. In *Proceedings of the Future Technologies Conference*, pages 1037–1058. Springer, 2018.

[24] V. Strobel *et al.* Managing byzantine robots via blockchain technology in a swarm robotics collective decision making scenario. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 541–549, 2018.

[25] K. Su *et al.* Smart city and the applications. In *2011 International Conference on Electronics, Communications and Control*, pages 1028–1031, Sep. 2011.

[26] J. Peña Queralta *et al.* Collaborative mapping with ioe-based heterogeneous vehicles for enhanced situational awareness. In *IEEE Sensors Applications Symposium (SAS)*. IEEE, 2019.

[27] T. Nam *et al.* Conceptualizing smart city with dimensions of technology, people, and institutions. In *Digital Government Innovation in Challenging Times*, dg.o '11, page 282–291, 2011.

[28] T. Taleb *et al.* On multi-access edge computing: A survey of the emerging 5g network edge cloud architecture and orchestration. *IEEE Communications Surveys & Tutorials*, 19(3):1657–1681, 2017.

[29] C. Campolo *et al.* 5g network slicing for vehicle-to-everything services. *IEEE Wireless Communications*, 24(6):38–45, Dec 2017.

[30] J. Mei *et al.* Intelligent network slicing for v2x services toward 5g. *IEEE Network*, pages 1–9, 2019.

[31] S Kekki *et al.* Mec in 5g networks. *ETSI white paper*, 28:1–28, 2018.

[32] YC Hu *et al.* Mobile edge computing—a key technology towards 5g. *ETSI white paper*, 11(11):1–16, 2015.

[33] 3GPP. Study on architecture for next-generation system rel. 14. *Techical Report*, 23.799, 2016.

[34] N. Alliance. Description of network slicing concept. *NGMN 5G P*, 1:1, 2016.

[35] F. Giust *et al.* Multi-access edge computing: The driver behind the wheel of 5g-connected cars. *IEEE Communications Standards Magazine*, 2(3):66–73, 2018.

[36] The European Union Agency for Cybersecurity. Threat assessment for the fifth generation of mobile telecommunications networks (5g). *ENISA THREAT LANDSCAPE FOR 5G NETWORKS*, 2019.

[37] M. Vukolić. Rethinking permissioned blockchains. In *Proceedings of the ACM Workshop on Blockchain, Cryptocurrencies and Contracts*, pages 3–7. ACM, 2017.

[38] Z Zheng *et al.* An overview of blockchain technology: Architecture, consensus, and future trends. In *2017 IEEE International Congress on Big Data (BigData Congress)*, pages 557–564. IEEE, 2017.

[39] C. Cachin. Architecture of the hyperledger blockchain fabric. In *Workshop on distributed cryptocurrencies and consensus ledgers*, volume 310, page 4, 2016.

[40] E. Androulaki*et al.* Hyperledger fabric: a distributed operating system for permissioned blockchains. In *Proceedings of the Thirteenth EuroSys Conference*, page 30. ACM, 2018.

[41] S. Cocco *et al.* Top 6 technical advantages of hyperledger fabric for blockchain networks. *IBM Developer*, 2018. https://developer.ibm.com/articles/top-technical-advantages-of-hyperledger-fabric-for-blockchain-networks/.

[42] M. Liu *et al.* Joint computation offloading and content caching for wireless blockchain networks. In *IEEE INFOCOM Workshops*, pages 517–522, April 2018.

[43] H. Zhu *et al.* Edgechain: Blockchain-based multi-vendor mobile edge application placement. In *2018 4th IEEE NetSoft*, pages 222–226, June 2018.

[44] L. Qingqing *et al.* Multi Sensor Fusion for Navigation and Mapping in Autonomous Vehicles: Accurate Localization in Urban Environments. In *The 9th IEEE CIS-RAM*, 2019.

[45] Synced. The golden age of hd mapping for autonomous driving. *Medium*, 2018.

[46] N. Sünderhauf *et al.* The limits and potentials of deep learning for robotics. *The International Journal of Robotics Research*, 37(4-5):405–420, 2018.

[47] El Mahdi El Mhamdi, Rachid Guerraoui, and Sébastien Rouault. The hidden vulnerability of distributed learning in byzantium. *arXiv preprint arXiv:1802.07927*, 2018.

[48] S. Biookaghazadeh *et al.* Are fpgas suitable for edge computing? In {*USENIX*} *Workshop on Hot Topics in Edge Computing (HotEdge 18)*, 2018.

[49] L. Qingqing *et al.* Edge Computing for Mobile Robots: Multi-Robot Feature-Based Lidar Odometry with FPGAs. In *The 12th ICMU*. IEEE, 2019.

[50] T. Ohkawa *et al.* Ros-compliant fpga component technology. installation of fpga into ros. In *ROSCon*, 2017.

[51] L. Luu *et al.* A secure sharding protocol for open blockchains. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 17–30. ACM, 2016.

[52] E. Kokoris-Kogias *et al.* Omniledger: A secure, scale-out, decentralized ledger via sharding. In *IEEE SP*, pages 583–598. IEEE, 2018.

[53] C. Ferris. "does hyperledger fabric perform at scale? *Blockchain Pulse: IBM Blockchain Blog*, 2, 2019.

[54] M. Zamani *et al.* Rapidchain: Scaling blockchain via full sharding. In *ACM SIGSAC CCS*, pages 931–948, 2018.