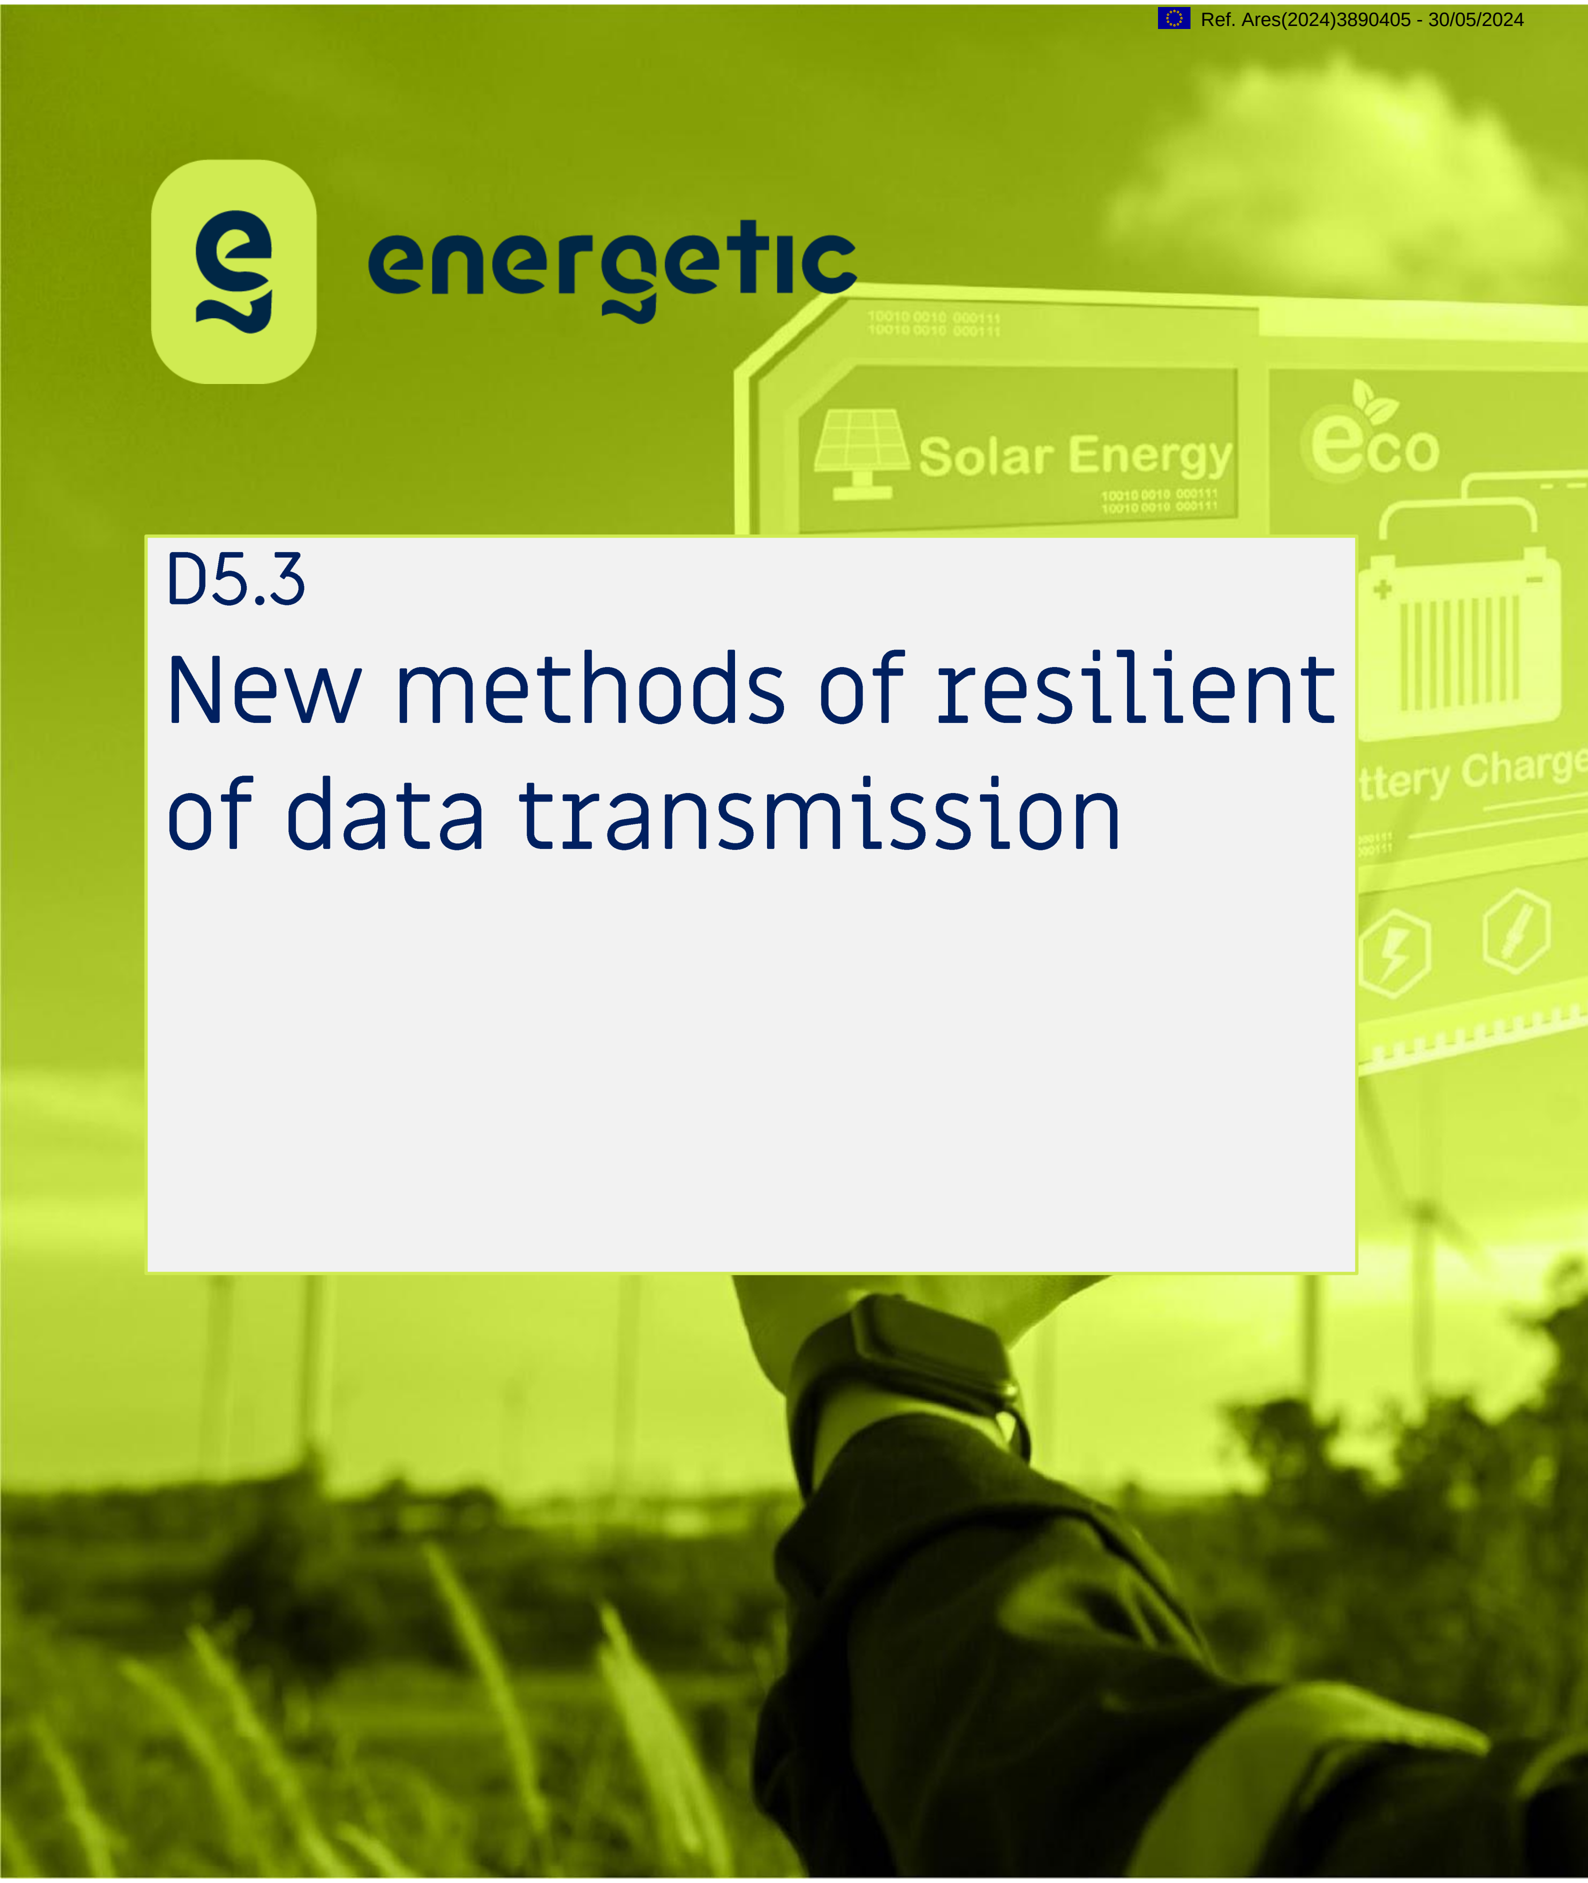

energetic

D5.3

New methods of resilient of data transmission

Funded by the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or CINEA. Neither the European Union nor the granting authority can be held responsible for them. Contract No. 101103667

D5.3	Work Package No.	WP 5	Task/s No.	Task 5.1.2	
Work Package Title	Battery Cloud predictive maintenance				
Linked Task/s Title	Data transmission				
Status	Final	(Draft/Draft Final/Final)			
Dissemination level	PU - Public	Grant Agreement Nº 101103667			
Due date deliverable	31/05/2024	Submission date		30/05/2024	
Deliverable version	New methods of resilient of data transmission				

Document Contributors

Deliverable responsible		Tallinn University of Technology (Taltech)	
Contributors	**Organization**	**Reviewers**	**Organization**
Sadok Ben Yahia	Tallinn University of Technology	Charles Courtney	University of Bath
Bassem Sellami	Tallinn University of Technology		

Document History

Version	Date	Comment
1.1	26-04-2024	Initial version
1.2	06-05-2024	A version update for architecture
1.3	21-05-2024	The draft version shared with internal reviewers
1.4	29-05-2024	Final version with integrated comments and changes from reviewers

Table of Contents

List of Figures

List of Tables

List of abbreviations

Abbreviation	Description
CA	Consortium Agreement
DMP	Data Management Plan
DoA	Description of Action
DT	Digital twin
EB	Executive Board
EC	European Commission
FAIR	Fundable Accessible Interoperable Reusable
GA	Grant Agreement
GDPR	General Data Protection Regulation
GeAs	General Assembly
HTTPS	Hypertext Transfer Protocol Secure
IEC	INSA Ethics Committee
IAB/EAB	International Advisory Board/ External Advisory Board
ICT	Information and Communication Technology
KOM	Kick-off Meeting
KPI	Key Performance Indicator
LEC	Local Ethics Committee(s)
MoM	Minutes of Meetings
MQQT	Message Queuing Telemetry Transport
MT	Management Team
PM	Project Manager
PSC	Project Scientific Coordinator
POC	Proof of Concept
R&D	Research and Development
RP	Reporting Period
SQL	Structured query language
TL	Task Leader
TLS	Transport Layer Security
UN	United Nation
WBS	Work Breakdown Structure
WP	Work Package
WPL	Work Package Leader
VNF	Virtual Network Functions
NFV	Network Functions Virtualization
PBFT	Practical Byzantine Fault Tolerance
PoW	Proof of Work
PoA	Proof of Authority

SPECIFIC ACRONYMS	
A3C: Asynchronous Actor-Critic Agent **AI:** Artificial Intelligence **API:** Application Programming Interface **AWS:** Amazon Web Service **BMS:** Battery Management System **CB:** Cell Balancing **DRL:** Deep Reinforcement Learning **DoS:** Denial of Service **DDoS:** Distributed DoS **DNN:** Deep Neural Network **EV:** Electric Vehicle **GAN:** Generative Adversarial Network **GCP:** Google Cloud Platform **GDPR:** General Data Protection Regulation **HAL:** Hardware abstraction Layer **HIL:** Hardware-in-the-loop **IoT:** Internet of Things	**IoV:** Internet of Vehicles **ML:** Machine Learning **PaaS:** Platform as a Service **SDN:** Software-Defined Networking **SIL:** Software-in-the-loop **SoA:** State-of-the-Art **SOA:** Safe Operation Area **SoH:** State of Health **SoC:** State of Charge **SoX:** State of Health or Charge or other **RUL:** Remaining Useful Life **RSU**: Roadside Unit

1. Executive Summary

The deliverable D5.3 presents an innovative architecture designed to manage battery data within the ENERGETIC project, focusing on resilience, security, and efficiency. This architecture integrates edge, fog, and cloud layers, incorporating Blockchain and Software-Defined Networking (SDN) technologies to enhance transparency and optimize battery management.
The architecture emphasizes continuous connectivity and adaptability to accommodate the dynamic nature of deployment and evolving requirements. Deliverable D5.3 lays the groundwork for effective battery data management in the ENERGETIC project, paving the way for further development and implementation phases.

2. Summary of ENERGETIC Project

ENERGETIC is a project funded by the Horizon Europe Programme of the European Commission whose goal is to develop the next generation Battery Management System (BMS) for optimizing batteries' system utilization in the first life (transport) and the second life (stationary) on a path towards more reliable, powerful, and safer operations.

The ENERGETIC project contributes to the field of translational enhanced sensing technologies, exploiting multiple Artificial Intelligence models supported by Edge and Cloud computing. With a digital twin, ENERGETIC can not only keep an eye on and predict how much longer a Li-ion battery will work, but they can also diagnose problems by looking into the reasons for degradation using AI models that can be explained.

This involves the development of new technologies for sensing, the combination and validation of multiphysics and data-driven models, information fusion through Artificial Intelligence, Real-time testing, and smart Digital Twin development. Based on a solid and interdisciplinary consortium of partners, the ENERGETIC R&D project develops innovative physics and data-based approaches both at the software and hardware levels to ensure optimized and safe utilization of the battery system during all modes of operation.

The project's objectives encompass a comprehensive approach to revolutionizing battery management. This involves integrating cost-effective sensors into the BMS and, supplying enriched physical data like temperature and ultrasonic details for optimized battery use. A scalable Hardware Abstraction Layer (HAL) platform is designed to gather, synchronize, and standardize sensor data for comparison with simulation results, enhancing understanding of battery aging.

Multiphysics modeling tools will look at the assessment of different KPIs, called State-of-X (SoX), where "X" stands for Charge, Energy, or health, as well as remaining useful life (RUL) of different types of batteries. AI models will use neural networks and explainable AI to predict SoX and figure out what is wrong. An innovative Digital Twin (DT)-based BMS, incorporating Edge and Cloud computing, aims to monitor batteries more intelligently. The project will also set standards for predictive maintenance in Cloud-based energy storage, fostering future services. Finally, the innovative smart DT-based BMS will be validated experimentally under various usage scenarios, with efforts to disseminate and exploit project outcomes to elevate the technology's impact and readiness.

3. Introduction

Based on what was covered in deliverable D5.2 about the data management architecture of battery management systems (BMS) in electric vehicles (EVs) and how they connect to cloud technologies, D5.3 goes into more detail about these systems' security holes and how they need to change as cyber threats rise. As EVs become more interconnected, the potential impact of cyberattacks on the BMS become more severe, threatening the safety and functionality of the vehicles.

D5.2 highlighted how the BMS collects critical data on battery health, charge levels, and performance metrics from EVs and transmits this information to a cloud layer for further processing and analysis. This setup, while effective for centralized data analysis and management, introduces significant risks, particularly concerning the integrity and security of the data transmitted and stored. Malicious actors could exploit vulnerabilities to manipulate or intercept data, potentially leading to incorrect charging processes, overheating, or even catastrophic battery failures.

Building on this scenario, D5.3 introduces a blockchain-integrated architecture designed to enhance the security, integrity, and transparency of data transactions between the BMS and the cloud layer. Blockchain technology, with its inherent characteristics of decentralization, immutability, and consensus-based validation, offers a robust solution to the security challenges identified in D5.2.

The proposed blockchain-based BMS2Cloud, represented in Figure 1, ensures that data transmitted from the BMS to the cloud is first hashed and validated through a consensus mechanism before being permanently recorded on a decentralized ledger. This process not only secures the data against tampering but also ensures that only verified and trustworthy data is stored in the cloud for further analysis [1]. This layer of security is crucial for maintaining the operational integrity of the BMS and the overall safety of the EV.

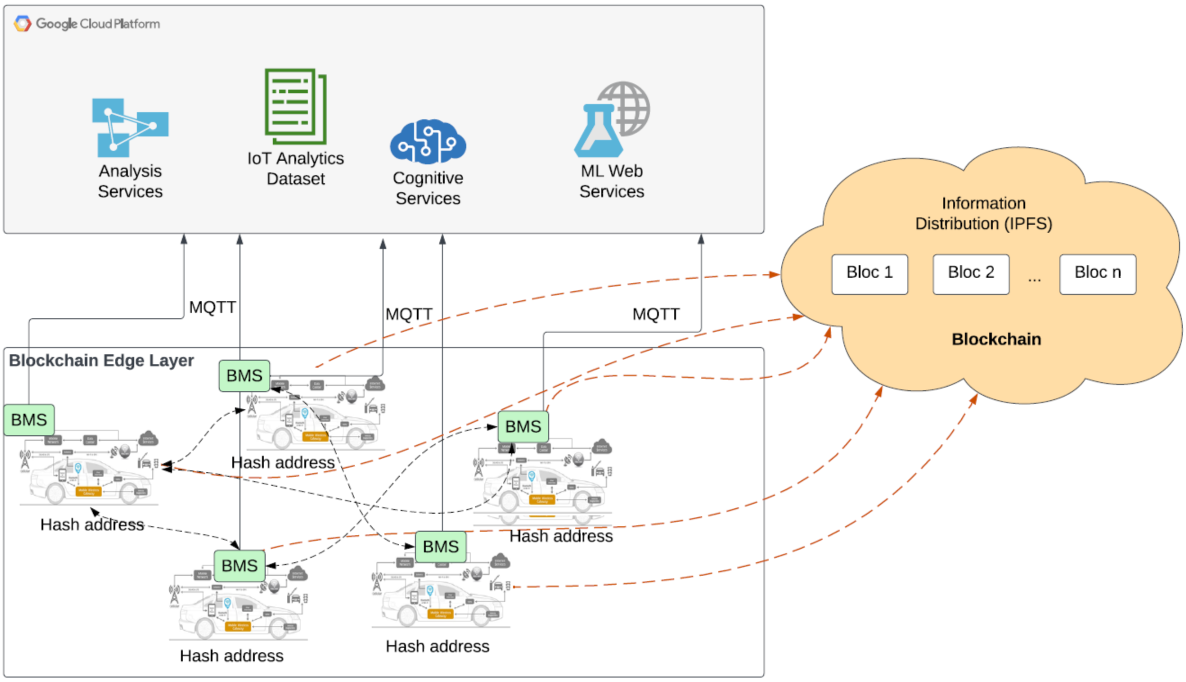

Figure 1: Blockchain-based BMS2Cloud architecture.

Additionally, the integration of blockchain technology facilitates the implementation of smart contracts, which can automate and secure interactions such as dynamic pricing for EV charging, peer-to-peer energy transactions, and automated maintenance notifications. This not only enhances the functionality of the BMS but also opens up new avenues for more efficient and transparent operations within the EV ecosystem.

Through this enhanced architecture, this deliverable aims to address the critical security issues raised in D5.2 by leveraging the strengths of blockchain technology to fortify the BMS against cyber threats. This approach not only underscores the necessity for advanced security measures in the rapidly evolving field of smart transportation but also illustrates the potential of emerging technologies like blockchain to revolutionize the management and security of critical infrastructure systems in connected electric vehicles.

4. Scope and content definition of the deliverable

4.1. Item Definition and overall system architecture

While the blockchain-enabled EV network presented in Figure 1 offers robust data integrity and security, it encounters several limitations that highlight the need for executing certain tasks closer to the EVs, ideally on fog nodes. These limitations include scalability challenges, latency issues, and the computational overhead associated with maintaining the Blockchain infrastructure. To overcome these obstacles, we propose an innovative architecture that integrates software-defined networking (SDN) with blockchain technology [2], incorporating fog nodes to execute tasks nearer to the EVs. This approach capitalizes on SDN's dynamic network management capabilities, optimizing bandwidth usage and reducing latency through real-time data path optimization [3]. By processing data on fog nodes, the architecture offloads computational tasks from the central cloud, enhancing efficiency and responsiveness. Furthermore, SDN's programmability allows for the dynamic implementation of security policies, which, when combined with blockchain's immutable ledger and decentralized consensus, create a robust and flexible security framework. This integration not only addresses scalability and latency issues but also reduces operational costs by optimizing network resources. As a result, the proposed framework's architecture offers a secure, scalable, and effective way to handle the complex data and operational needs of a large EV network, ensuring timely and reliable data transfer while improving the overall performance of the network. This makes it an ideal solution for the evolving needs of modern EV infrastructures.

This section gives an overview of the necessary content of this deliverable. Therefore, the following overall architecture of the framework is shown in Figure 2.

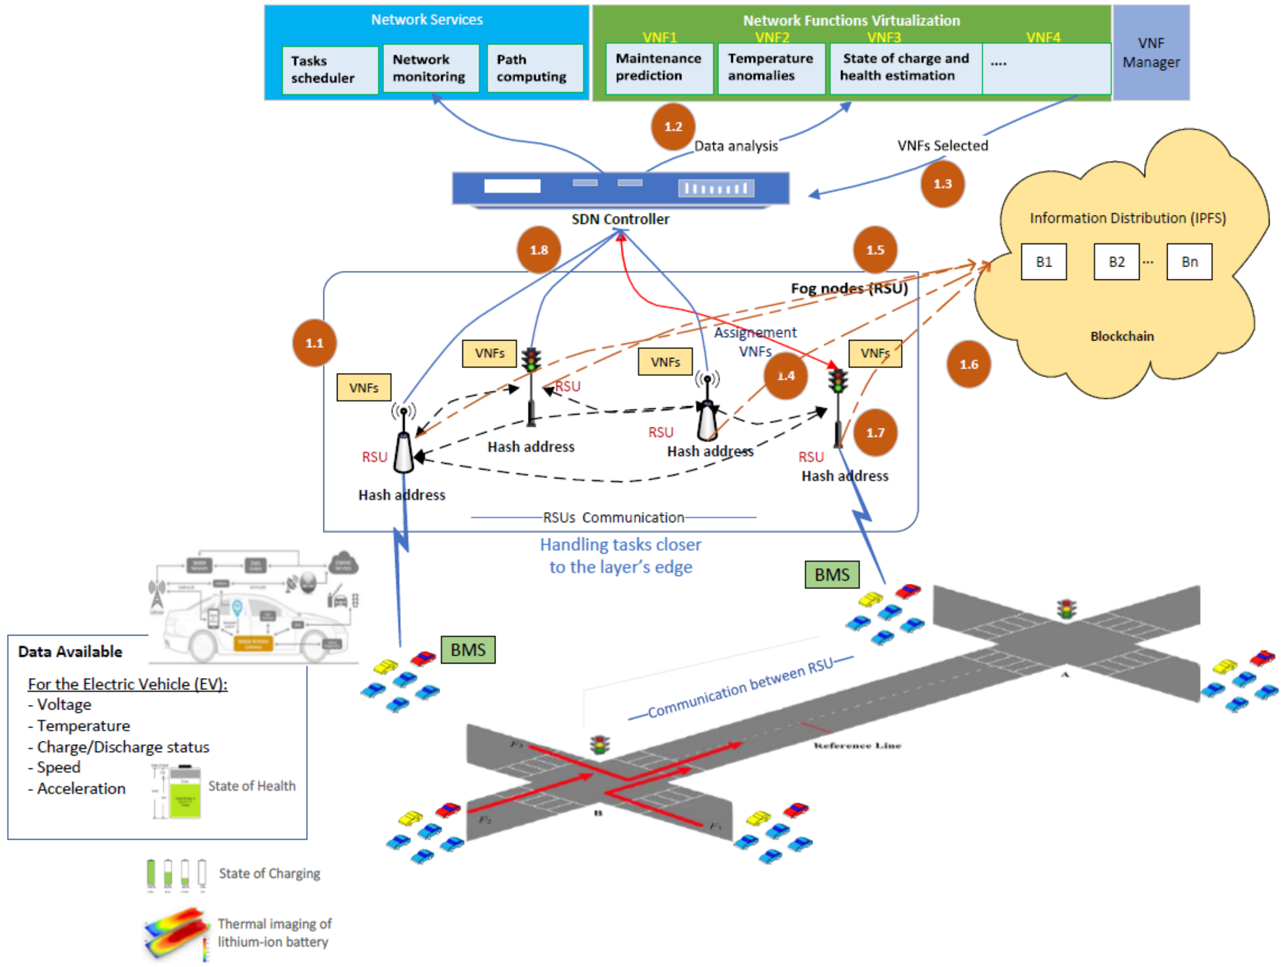

Figure 2: Blockchain-SDN-based EV Network Architecture at a Glance.

The overall system consists of the following steps:

1.1) Data Collection from Vehicles: The EV is equipped with BMS sensors that monitor and transmit real-time data, including charge status, voltage, temperature, and current flow of the battery cells. This data is collected and analyzed to ensure optimal performance, longevity, and safety of the battery system within the vehicle's ecosystem.

1.2) Data Processing and SDN Control: The SDN controller, which acts as the central processing unit for battery health management, receives the data that the Fog node (Roadside Unit; RSU) has collected**.** The SDN controller processes and analyzes the data, assessing the battery's operational conditions based on predefined parameters and algorithms. It identifies potential risks such as overcharging, deep discharge, or temperature anomalies, ensuring the battery operates within safe and efficient parameters.

1.3) VNF Function Selection: The SDN controller evaluates the processed data to determine the appropriate management functions to ensure battery safety and efficiency. These functions can include cell balancing, temperature regulation, state of charge and health estimation, and charging control. The selection of VNF functions is based on the current operational conditions of the battery and the overall requirements of the battery system within the vehicle.

1.4) Fog Node (RSU) Deployment: The SDN controller utilizes its oversight of the battery system to activate the selected VNF functions. It identifies the fog node that has the necessary monitoring capabilities, communication systems, and proximity to the battery cells involved. The VNF functions are activated and managed on the selected fog node, employing advanced algorithms and control strategies. These fog nodes now serve as localized points for battery health management, ready to maintain and optimize battery performance within the vehicle's operational range.

1.5) Blockchain Integration: The activated VNF functions generate critical battery-related events, such as charge level warnings or temperature alerts. These events, along with pertinent metadata (e.g., timestamp, battery cell information), are recorded as transactions on the Blockchain. Each fog node can act as a node within the Blockchain network, enhancing the system's transparency and reliability through a decentralized and distributed ledger. This ensures a secure and tamper-proof record of all battery management activities.

1.6) Consensus and Verification: The Blockchain network achieves consensus among the participating nodes, ensuring the validity and integrity of recorded transactions. A subset of network nodes verifies the recorded transactions relating to safety events, ensuring transparency and preventing tampering.

1.7) Real-Time Management Services: The VNF functions, once activated, continuously process incoming data from the battery cells in real-time. They perform tasks such as charge level monitoring, temperature regulation, and health assessment based on predefined algorithms and rules. The management services provided by the VNF functions aim to maintain battery efficiency, warn of potential issues, or trigger protective measures (e.g., temperature control) if necessary.

1.8) Feedback and System Adaptation: The SDN controller receives feedback from the active VNF functions, including battery status alerts, performance updates, or health metrics. Based on this feedback, the SDN controller can dynamically adjust the VNF functions, modify their settings, or recalibrate them to optimize battery performance and longevity. This adaptive approach ensures the battery management system remains responsive to the changing conditions of the battery cells.

The proposed architecture is designed to ensure resilience, security, and service availability with low latency, addressing the critical needs of an extensive EV network. Edge computing is built into each EV's Battery Management System (BMS). This means that real-time data collection and initial processing happen at the source, cutting down on latency and making the system more responsive to immediate problems with battery health and performance. The deployment of Roadside Units (RSUs) and other local computing nodes in the fog layer enables localized data processing, allowing for quick, responsive actions such as adjusting charging rates or activating cooling systems based on the BMS data. This decentralization of processing tasks not only reduces the load on the central cloud but also ensures that essential services remain available even if part of the network experiences disruptions, thus enhancing resilience.

Integration of Software-Defined Networking (SDN) further strengthens the architecture by providing centralized control and dynamic optimization of network resources. The SDN controller prioritizes BMS data transmission, ensuring timely and reliable communication, while its programmability allows for the dynamic implementation of security policies, mitigating potential threats in real-time. This centralized control combined with dynamic resource management ensures that network performance is optimized continuously, maintaining low latency and high service availability.

Incorporating blockchain technology for secure data recording adds another layer of security and integrity to the system. Significant BMS events and decisions are recorded as immutable transactions on the blockchain, including metadata such as timestamps and vehicle identifiers, creating a transparent and tamper-proof record of all activities. This ensures that all data is secure from unauthorized alterations and can be reliably audited, enhancing trust in the system's operations.

Network Functions Virtualization (NFV) also lets the SDN controller set up necessary network functions on the fly, like traffic shaping and security measures. This makes sure that network resources are managed well and services are always available. Continuous real-time monitoring and response capabilities make sure that the system can quickly fix any problems that are found. This keeps the battery within safe operating parameters and changes NFV deployments as needed to keep the network running at its best.

In general, this architecture is strong because it uses edge and fog computing, advanced cloud analytics, and real-time network management to make sure that services are available and secure, with low latency, thanks to distributed processing, blockchain, and dynamic SDN policies.

5. Discussion and Analysis of Results

5.1. Evaluating distributed communication based on a decentralized Ledger

We evaluate the performance and scalability of our proposed architecture for Battery Management Systems (BMS) in Electric Vehicles (EVs), which integrates blockchain and Software-Defined Networking (SDN). Performance is defined by transaction latency and throughput, where an IoT transaction from the BMS is considered valid only once it is committed to the blockchain. The block interval—the amount of time between publishing subsequent blocks—and block size influence these parameters, which set an upper limit on transaction throughput. Scalability is defined as the blockchain network's ability to handle varying workloads in relation to the number of nodes.

To implement a prototype, we used the Ethereum blockchain with 20 nodes acting as blockchain miners, each running the leader-election consensus algorithm. We compared our solution against well-known baseline consensus algorithms, namely Proof of Work (PoW) and Practical Byzantine Fault Tolerance (PBFT) [4]. We evaluated the message exchange using these different baseline algorithms and compared them with our proposed architecture, which leverages SDN for dynamic network management and optimization.

We first evaluated the PBFT approach, examining its message exchange characteristics and how it handles the BMS data. Next, we discussed the message exchange in PoW, noting its distinct performance traits and impact on BMS data transactions. Finally, we assessed the message exchange in our proposed architecture, comparing its overhead against both PoW and PBFT. This comprehensive evaluation allowed us to determine the efficiency and effectiveness of our approach in terms of transaction latency, throughput, and scalability within a decentralized ledger framework tailored for BMS in EVs.

5.1.1. Message Exchange in PBFT Baseline Algorithm

PBFT was designed to work efficiently in asynchronous systems with low overhead. Specifically, in PBFT, a specific leader proposes the order of the transactions, and then the blockchain nodes communicate with each other in several steps to reach an agreement.

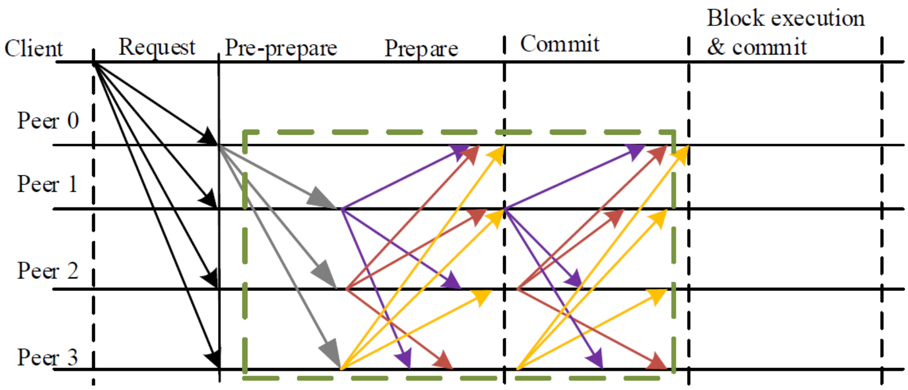

Figure 3: Message Exchange in Voting-Based PBFT.

Figure 3 shows how the PBFT voting-based consensus in the PoW algorithm exchanges messages. The PBFT is introduced to break the performance bottleneck of PoW-based blockchain systems.

5.1.2. Message Exchange in the PoW Baseline Algorithm

The PoW-based consensus does not require a fixed leader to validate the blocks. Instead, a group of leaders competes to validate a block of transactions. First, a node can propose a block of transactions and should solve a PoW (i.e., a mathematical challenge) from the previous transactions and get rewards, if that proposal is accepted. Then, the node generates a pseudorandom number, the so-called nonce, broadcasted to all connected nodes as shown in Figure 4. Afterward, nodes compete to become the next leader, i.e., miners, by selecting transactions and generating a hash. The node generating a more minor hash than the nonce value becomes the next leader.

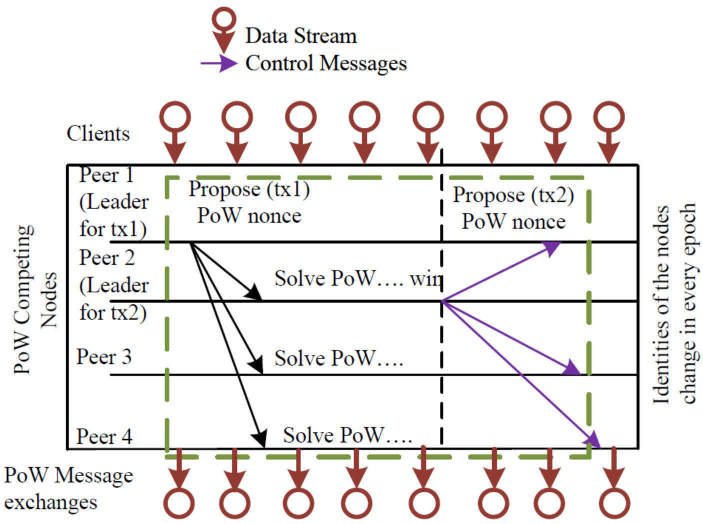

Figure 4: PoW Message Exchanges.

Compared to the PBFT-based approach, the PoW algorithm induces less message exchange overhead to validate transactions, as depicted in Figure 4. Specifically, the PoW consensus algorithm requires four message rounds to commit a block. Before a new IoT transaction block is confirmed, most network nodes should verify and approve it. Additionally, in PoW, all unverified IoT transactions are put together in a poll. Then all miners work to check that those transactions are legitimate by solving a complex mathematical puzzle. Thus, the PoW consensus algorithm is the most reliable and secure among the three algorithms.
The problem occurs when more than one node simultaneously solves the mathematical puzzle. In such a case, all other connected miners detect and trigger a fork operation of the hash chain, which causes an overhead explosion of the message exchange. Such a situation involves the selection of the longest hash chain and takes a long time for each transaction to be validated. It often takes more than six blocks for a transaction to be finalized. The drawback of this approach, the de facto scenario in a distributed

blockchain, is the limited number of transactions processed and concluded per second. Thus, scalability becomes an issue because the block size is too small to sustain thousands of transactions, lowering the overall blockchain network's throughput and increasing the energy consumption required to validate these transactions.

5.1.3. Message Exchange in our Architecture

Figure 5(a) shows a leader node (Authority 0) receiving a request from an IoT device, such as a connected car, to validate its transactions. The leader then broadcasts the block to a group of pre-approved authority nodes, which in this context are RSU1, RSU2, and RSU3, to validate IoT transactions and commit them to the blockchain. The proposed validation process involves the election of authority verifier nodes that have the best Quality of Service (QoS) settings, specifically the lowest transaction latency. This process requires minimal computation as it only needs one round to validate and commit a new block to the blockchain, assuming bounded transaction latency expressed in time steps.

Figure 5(b) illustrates a scenario where a leader authority node (a2) broadcasts a new block (b1) to the blockchain while a non-leader authority node (a3) simultaneously broadcasts another block (b2). The first newly created block (b1) reaches nodes a1 and a5 before block (b2) arrives at these nodes. Conversely, block (b2) reaches nodes a3 and a4 before they receive the first block (b1). When blocks from various miners become misaligned and network desynchronization occurs, as shown on the right side of Figure 5(b), each node in the blockchain performs a fork operation. Authority nodes a3 and a4 decide to continue using block b1 as the first block and reference it as the previously reacted block, with block b2 as the next arriving block.

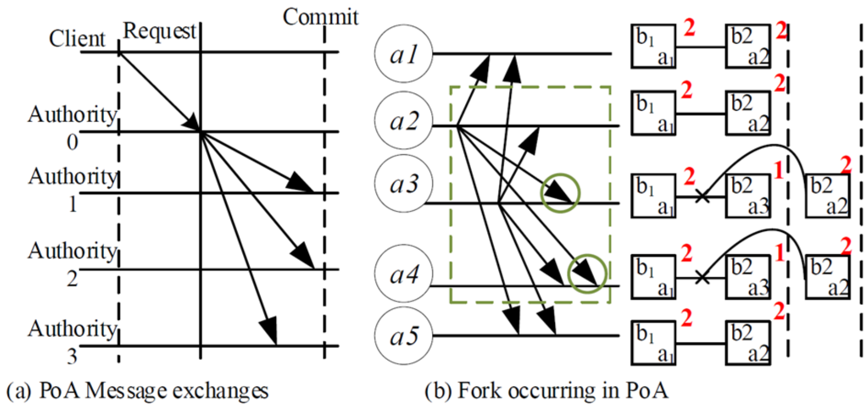

Figure 5: Latency during Message Exchange in PoA

From this discussion, it is evident that our architecture based on Proof-of-Authority (PoA) outperforms both PoW and PBFT, enhancing the performance of the blockchain IoT network. Our approach selects representative mining nodes based on their reputation. Once the reputation system is established, preselected nodes have the authority to propose new transactions and notify others of the results. The IoT service providers have preselected, trusted, and deployed the RSUs, which serve as representative mining nodes, for latency sensitive IoT scenarios like our case of Internet of Vehicles (IoVs). These providers assume that the RSUs have no incentive to build up their reputation and then corrupt the entire system. Malicious nodes can be easily detected and blacklisted based on their IDs, such as physical or local addresses. Nodes with suspicious behavior can be quarantined and cross-checked against their respective IDs.

In the following, we gauge the blockchain overhead performance of the proposed approach in terms of transaction latency and throughput. An IoT transaction is only considered valid once it is committed to the blockchain, establishing an upper bound on transaction throughput. The block interval, which is the amount of time between publishing subsequent blocks, and the block size both have an impact on performance. Scalability is defined as the blockchain network's ability to handle varying workloads based on the number of nodes.

5.1.4. Evaluating the latency

The transaction latency involves the commitment of subsequent blocks and includes a safety check against double-spending to ensure that the transaction is irrevocable within the chain for at least the subsequent six mined blocks.

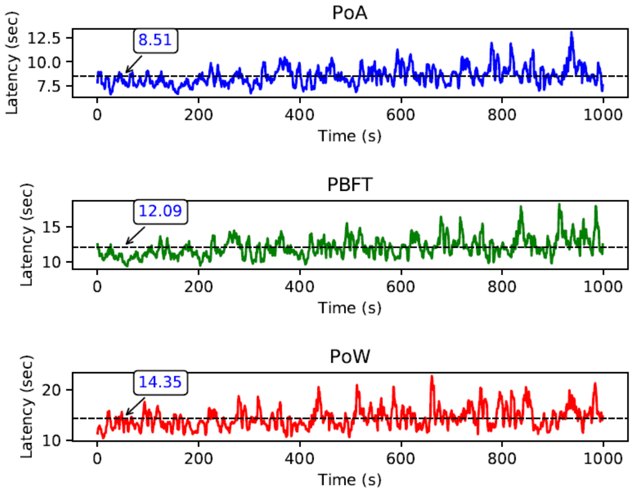

Figure 6: Latency of different consensus algorithms: our approach (PoA) versus PoW and PBFT over time

The evaluation of latency for three types of consensus protocols (PoW, PBFT, and the proposed PoA approach) is depicted in Figure 6. The results of experiments show that the latency of the PoW consensus is higher compared to the other consensus algorithms. PoW consumes substantial computational resources to confirm a block. In the same way, the practical Byzantine Fault Tolerance (PBFT) has higher latency than our proposed PoA-based approach, even though it has lower latency than PoW because it gives equal weight to all participating peer-to-peer network nodes. Specifically, our approach achieves an average latency of 8.51 seconds for confirming an IoT transaction. In contrast, the voting-based PBFT consensus achieves an average latency of 12.09 seconds to validate a transaction, while the PoW consensus validates a new IoT transaction with an average time delay of 14.35 seconds. This means that our method works better than the standard ones (PoW and PBFT) and reduces network latency. This shows that the PoA consensus works well for BMS data in EV networks.

5.1.5. Evaluating the throughput

There are two subcategories of blockchain throughput: read throughput and transaction throughput. Read throughput measures the rates at which data are read, i.e., the number of reading operations completed in a given period, formally expressed as reads per second (RPS). Read throughput is often not considered a critical metric for evaluating the quality of service in a blockchain network because blockchain nodes typically achieve significantly higher read and query efficiency compared to the write operations in the ledger. Transaction throughput, on the other hand, measures how fast the blockchain can process incoming IoT transactions, expressed as transaction rates per second (TPS). Blockchain network throughput is not measured at a single node but reflects the overall performance of the blockchain network across all nodes.

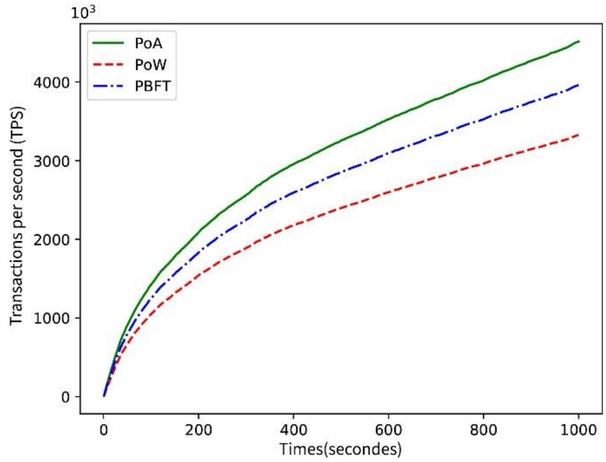

Figure 7: Transaction rate (TPS) of the proposed approach against the PoW and PBFT algorithms.

Figure 7 illustrates the transaction rates for PoW, PBFT, and our proposed PoA consensus algorithm. Transaction throughput increases linearly with the increase in block size. Specifically, our approach achieves 5 million transactions per second, compared to PoW and PBFT. PoW achieves the lowest transaction throughput, with an average of 3 million transactions per second, while PBFT achieves an average throughput of 4 million transactions per second. Therefore, our approach shows significant improvements compared to both PoW and PBFT, demonstrating its superior capability in processing the high volume of BMS data in EV networks efficiently.

5.1.6. Energy Consumption

Blockchain energy consumption has received significant attention due to its potential cost-ineffectiveness and the substantial computing power required. The cost of ensuring transaction trustworthiness can be the inefficiency of mining, which involves a competitive process where all blockchain nodes vie for energy- and resource-intensive cryptographic lottery.

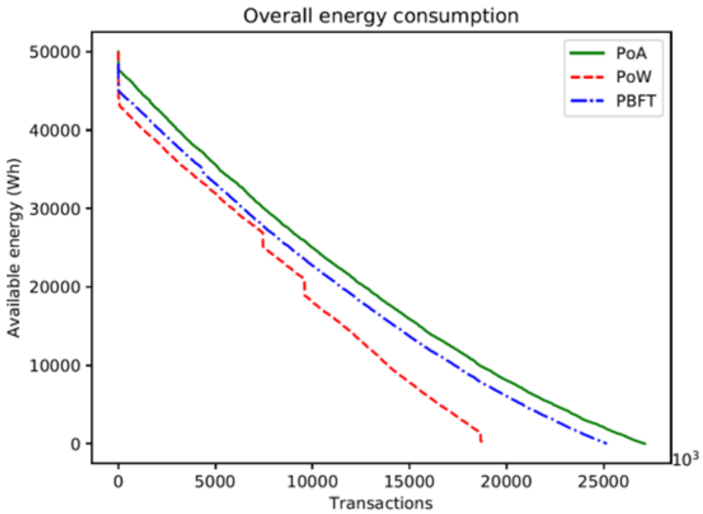

Figure 8: Evaluating the energy efficiency of battery-powered IoT nodes.

Figure 8 shows that the PoW consensus algorithm performs the worst, with battery capacity depleting after processing, confirming, and validating 18 million transactions. Similarly, the PBFT approach performs slightly better than PoW, managing to process and confirm an average of 28 million transactions before running out of energy at low levels. Based on PoA, our approach achieves better energy efficiency than PoW and PBFT. Therefore, our approach significantly reduces resource utilization and energy consumption.

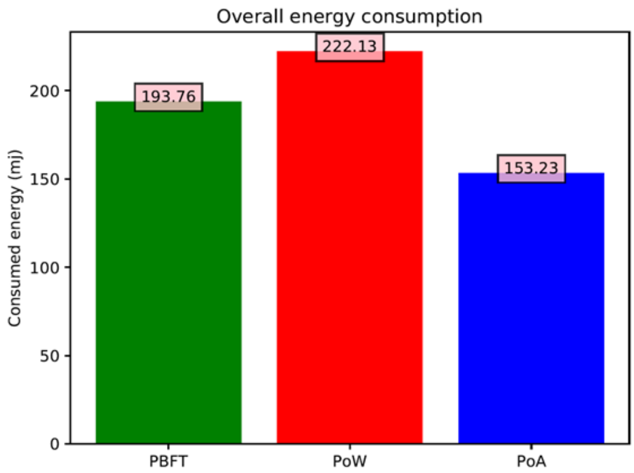

Figure 9: Comparison of the energy consumption of the proposed architecture with the PoW and PBFT baselines.

To further investigate the most energy-efficient blockchain consensus algorithm, Figure 9 illustrates the overall energy consumption of each approach. The PoW approach is highly energy-intensive, consuming 222.13 mJ to deter frivolous or malicious attacks. The PBFT approach consumes 193.76 mJ, avoiding the complex mathematical computations required by PoW. Finally, our PoA-based approach outperforms both, with an energy consumption of 153.23 mJ. This demonstrates that our proposed method not only enhances energy efficiency but also ensures the sustainability of the blockchain network in managing BMS data for EVs.

5.1.7. CPU Resource Utilization

Resource usage is a critical aspect of blockchain technology, determining how smoothly distributed peer-to-peer (P2P) network nodes operate. CPU resources are essential in blockchain communication as they dictate how many transactions can be validated and included in a block added to the blockchain.

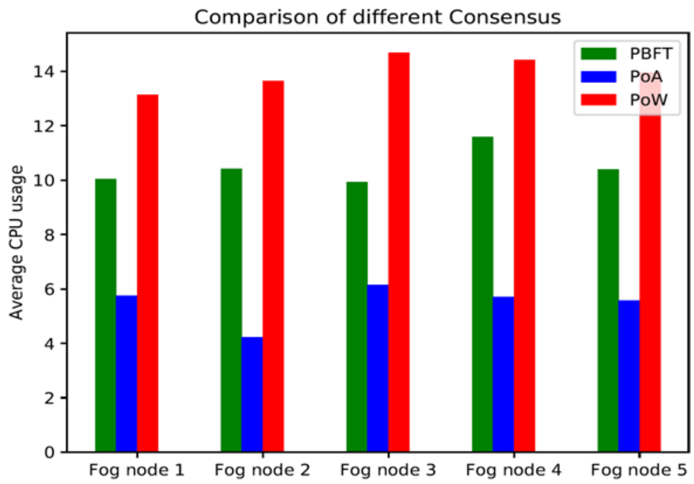

Figure 10: Comparison of the energy consumption of the proposed architecture with the PoW and PBFT baselines.

Figure 10 illustrates the CPU usage across different approaches. While the PoW approach shows a CPU usage of around 14%, which, although under the 50% threshold required to validate a block, is still high enough to make straightforward tasks laboriously slow,. The PBFT approach consumes about 10% of CPU resources on different nodes. In contrast, our proposed approach outperforms both by demonstrating a CPU usage of just 5%. Our method does not rely on computational power to validate IoT transactions; rather, fog nodes use CPU computation to process incoming requests. Our method uses an election-based consensus algorithm to choose blockchain validator nodes ahead of time. This process does not use a lot of CPU power, so it uses fewer resources than other consensus algorithms. This low CPU consumption ensures that the blockchain network operates smoothly and efficiently without burdening the nodes with excessive computational demands. Our method improves the speed of blockchain validation processes by using less CPU power. It also supports the real-time needs of BMS data management in EV networks, making it a very useful and efficient way to combine blockchain technology with SDN in EV applications.

5.2. Services Deployment across Decentralized Ledger-Based Communication Layers

5.2.1. Harnessing Hybrid Digital Twin Technology at the Edge Layer for Enhanced BMS Management

In our proposed architecture, we introduce a groundbreaking approach to BMS management within electric vehicles (EVs) by harnessing cutting-edge hybrid digital twin technology at the edge layer. This innovative solution revolutionizes how EV batteries are monitored and optimized in real time. By integrating real-time sensor data from the BMS with advanced predictive models and machine learning algorithms, we create a virtual representation, known as a digital twin, of the physical EV battery system.

The hybrid digital twin serves as a dynamic and adaptive model that closely mirrors the behavior and characteristics of the actual EV battery in real time. It continuously receives and processes sensor data from the BMS, allowing for accurate and up-to-date insights into battery health and performance. Through sophisticated algorithms, the digital twin can predict future states and behaviors of the battery, enabling proactive management strategies to be implemented.

One of the key advantages of deploying the hybrid digital twin at the edge layer is its ability to empower EVs to autonomously manage their battery health and performance. By leveraging the digital twin's capabilities, EVs can make informed decisions in real time, optimizing energy usage and maximizing battery lifespan. This includes the ability to perform simulations, predictive maintenance scheduling, anomaly detection, and adaptive control strategies.

For example, the digital twin can simulate various scenarios and predict the impact of different charging patterns or environmental conditions on battery performance. It can also schedule maintenance tasks based on predictive analytics, identifying potential issues before they escalate into costly failures. Moreover, the digital twin enables adaptive control strategies, allowing EVs to adjust charging and usage patterns dynamically based on changing conditions.

Overall, the deployment of hybrid digital twin technology at the edge layer can enhance the reliability and longevity of EV batteries. By enabling EVs to optimize their energy usage and ensure optimal battery performance, this transformative technology contributes to the overall efficiency and sustainability of electric transportation systems. It represents a paradigm shift in BMS management, paving the way for more intelligent and proactive approaches to battery optimization in EVs.

5.2.2. Optimizing Electric Vehicle Charging Station Recommendations Using Fog Layer Predictive Analytics

Within our architecture, the fog layer plays a pivotal role in optimizing electric vehicle (EV) charging station recommendations through advanced predictive analytics. By harnessing the computational capabilities of roadside units (RSUs) and local computing nodes, we enable real-time forecasting of charging station availability. Leveraging historical data, traffic patterns, and environmental factors, the fog layer predicts future demand for charging stations at specific locations. This predictive analytics-driven approach allows us to recommend optimal charging stations to EV drivers, minimizing wait times and optimizing charging schedules. By providing accurate and timely recommendations, we enhance the usability and efficiency of EV charging infrastructure, thereby promoting the widespread adoption of electric vehicles and facilitating the transition to sustainable transportation systems.

In evaluating the performance of our fog layer predictive analytics service, we conducted extensive testing across various scenarios and traffic conditions. Our results demonstrate a significant improvement in the efficiency and effectiveness of EV charging station recommendations. By accurately forecasting charging station demand and availability, we observed a notable reduction in wait times for EV drivers, leading to improved user satisfaction and utilization of charging infrastructure. Additionally, our predictive analytics-driven approach enables more effective utilization of charging resources, optimizing the allocation of charging stations based on anticipated demand. As a result, we observed increased efficiency and reduced congestion at charging stations, contributing to a smoother and more seamless charging experience for EV owners. Overall, our fog layer predictive analytics service has proven to be a valuable asset in enhancing the usability, efficiency, and accessibility of EV charging infrastructure, thereby driving the widespread adoption of electric vehicles and advancing sustainable transportation systems.

In the following, we unveil the outcomes of our rigorous evaluation of multi-class classification machine learning and deep learning models, focusing on their performance in predicting charging station availability. Precision, recall, and F1 score metrics were employed to assess and compare the effectiveness of these models.

Table 1 presents the results derived from applying machine learning/deep learning (ML/DL) methodologies to the Paris dataset [5]. Notably, the Random Forest (RF) model emerged as the top performer among the ML baseline models, boasting an accuracy rate exceeding 95%. Following closely, the K-Nearest Neighbors (KNN) model attained a commendable rate of 94% across all metrics. On the other hand, Logistic Regression came in second with an F1 score of 60.73%, closely following the Support Vector Machine (SVM) model at 59.54%.

Despite marginally trailing behind the K-Nearest Neighbors (KNN) and Random Forest (RF) models in performance scores, the artificial neural network (ANN) model displayed impressive precision (89.46%), recall (87.94%), and F1 score (88.69%).

Upon careful analysis of the results depicted in Table 1 and Figure 11, it becomes evident that the Random Forest and K-Nearest Neighbor models consistently outperform others on this dataset, exhibiting high precision, recall, and F1 scores. Conversely, Logistic Regression and Support Vector Machine models demonstrate weaker performance. However, the Artificial Neural Network model showcases commendable performance, positioning itself between the top-performing models and the weaker ones. Notably, while the ANN model facilitates real-time predictions, the instance-based learning of KNN limits its ability to offer instantaneous forecasts. Furthermore, RF provides results solely lacking rates for output classes.

APPROACH	RECALL	PRECISION	F1-SCORE
ARTIFICIAL NEURAL NETWORK	87.94	89.46	88.69
K-NEAREST NEIGHBOURS	94.13	94.10	94.11
LOGISTIC REGRESSION	72.16	66.65	60.44

RANDOM FOREST	95.56	95.54	95.41
SUPPORT VECTOR MACHINE	71.90	66.41	59.54

Table 1: Performance results (%) per model for the Belib Dataset

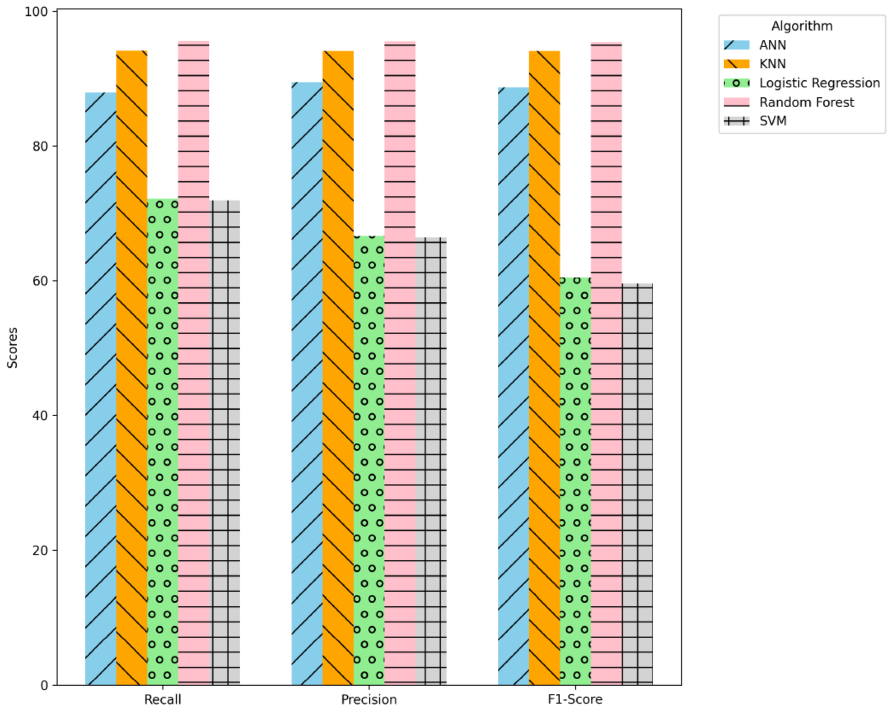

Figure 11: Performance results (%) per model for the Belib Dataset.

Figure 12 provides a visual representation of the initial five prediction outcomes generated by the artificial neural network (ANN) prediction engine, utilizing the Paris dataset “Belib” [5]. The percentages displayed in the figure denote the proportions of status outputs, with each percentage corresponding to a specific interpretation. These ideas are shown in more detail in Figure 12, which shows the predicted results of the model using ANN for the Belib dataset. In Figure 12a, the prediction indicates that the charging spot is 93.89% Available ("Disponible"). Moving to Figure 12b, the prediction suggests that the charging spot is 95.26% in maintenance ("En maintenance") and 0.326% Available ("Disponible"). Similarly, in Figure 12c, the prediction portrays the charging spot as 79.79% Available ("Disponible") and 0.842% in the process of commissioning ("En cours de mise en service"). Figure 12d depicts the charging spot as 64.62% Available ("Disponible") and 16.53% Busy ("En charge"). Lastly, in Figure 12e, the prediction indicates that the charging spot is 99.92% Available ("Disponible"). These visual representations provide valuable insights into the ANN model's predictions, aiding in understanding and interpreting the status of charging spots with high accuracy and reliability.

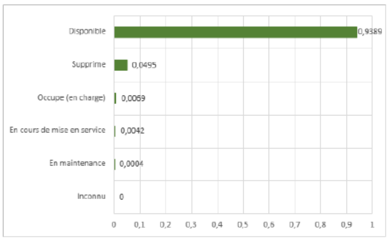

(a) First prediction outcome

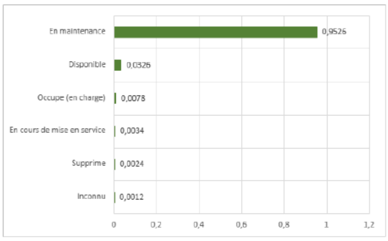

(b) Second prediction outcome

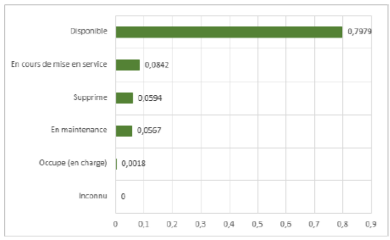

(c) Third prediction outcome

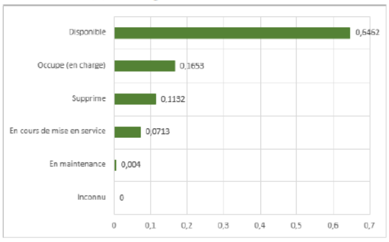

(d) Fourth prediction outcome

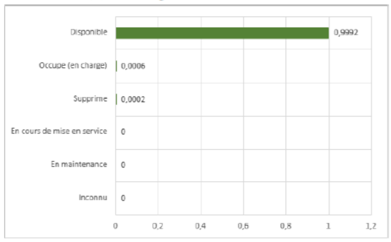

(e) Fifth prediction outcome

Figure 12: Results from the prediction model using ANN for the Belib database

Table 2 unveils the findings obtained from the application of machine learning (ML) and deep learning (DL) approaches to the Estonian data [6]. Notably, the K-Nearest Neighbours (KNN) model emerges as a consistent performer across all metrics. It demonstrates high precision, recall, and F1-score, highlighting its robust ability to correctly classify instances of different classes. Following closely, the Logistic Regression model exhibits competitive performance, with high recall indicating its proficiency in identifying positive cases and maintaining a balanced F1 score. Moreover, the Random Forest model showcases solid performance across all metrics, boasting high recall and precision, leading to a strong F1 score. Similarly, the Support Vector Machine (SVM) model delivers commendable results with elevated precision and recall values, thereby reflecting a well-balanced F1 score. In addition, the Artificial Neural Network (ANN) model gets great results, especially in accuracy, with high recall and F1-score, showing that it works well for classification tasks on this dataset.

Upon thorough examination of the results presented in Table 2 and Figure 13, it becomes apparent that all models exhibit robust performance on the "Enefit Volt Dataset" [6]. Specifically, the ANN, Logistic Regression, and SVM models consistently achieve high precision, recall, and F1 scores. Although the K-Nearest Neighbours and Random Forest models also perform admirably, their F1 scores are slightly lower compared to the others. Overall, these models demonstrate suitability for accurately classifying instances in the "Enefit Volt Dataset," underscoring their efficacy in leveraging ML and DL approaches for insightful analysis and decision-making.

APPROACH	RECALL	PRECISION	F1-SCORE
ARTIFICIAL NEURAL NETWORK	95.74	96.04	95.89
K-NEAREST NEIGHBOURS	94.19	94.11	94.15
LOGISTIC REGRESSION	96.63	94.35	94.81
RANDOM FOREST	96.02	94.78	95.17
SUPPORT VECTOR MACHINE	96.32	95.15	95.75

Table 2: Enefit Volt Dataset Performance Results (%)

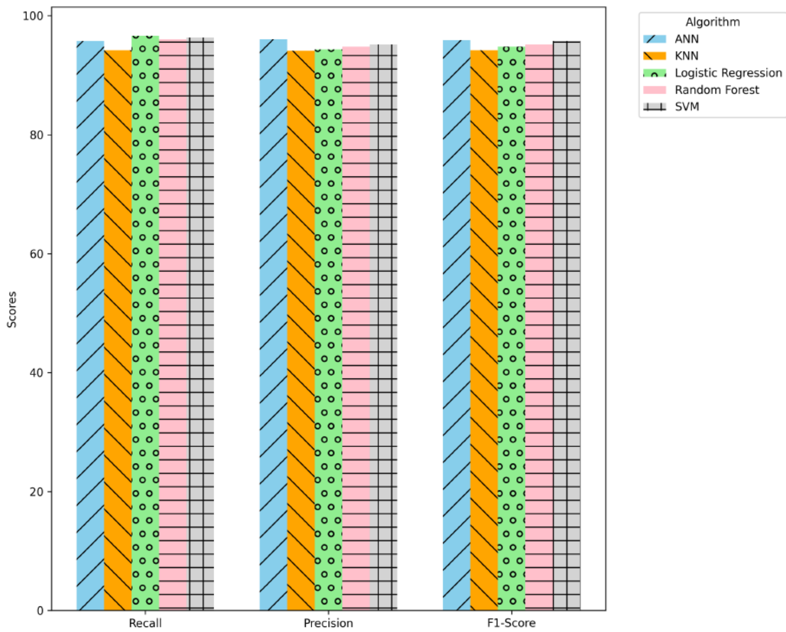

Figure 13: Enefit Volt Dataset Performance Results (%) by Model

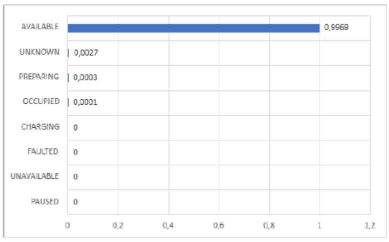

(a) First prediction outcome

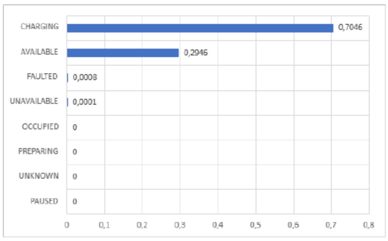

(b) Second prediction outcome

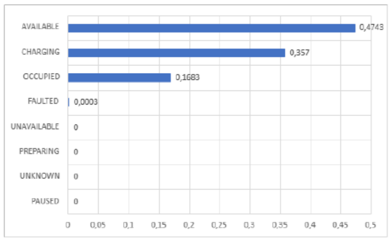

(c) Third prediction outcome

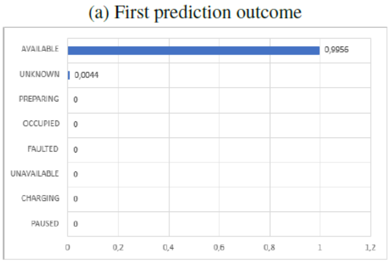

(d) Fourth prediction outcome

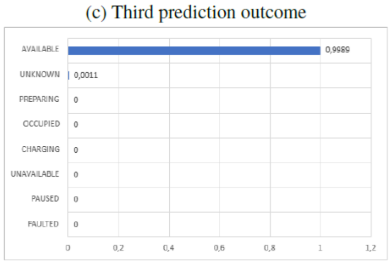

(e) Fifth prediction outcome

Figure 14: Results from the prediction model using ANN for Enefit Volt dataset

Figure 14 visually represents the first five prediction outcomes obtained by applying the artificial neural network (ANN) prediction engine to the Estonian dataset. Each figure within the visualization corresponds to percentages linked to different status outputs, as elucidated in Figure 14. Specifically, this figure illustrates the prediction outcomes of the model utilizing ANN for the Enefit Volt dataset.

In Figure 14a, the prediction indicates that the charging spot is 99.69% available and free to use. Transitioning to Figure 14b, the prediction suggests that the charging spot is 70.46% in a charging session and 29.46% available. As a result, in Figure 14c, the prediction portrays the charging spot as 47.43% available, with 35.7% in a charging session and 16.83% occupied by someone else. Moving to Figure 14d, the prediction indicates that the charging spot is 99.56% available and free to use. Lastly, in Figure 14e, the prediction suggests that the charging spot is 99.89% available and free to use. These visual representations offer valuable insights into the ANN model's predictions, facilitating a comprehensive understanding of the status of charging spots with high accuracy and reliability.

The results showcased in our study signify a significant advancement in the realm of BMS management within electric vehicles (EVs) by leveraging advanced machine learning techniques. Our precision, recall, and F1 scores outshine those reported in existing literature, underscoring the effectiveness of our proposed architecture in enhancing BMS management.

The performance outcomes observed on both the "Belib Dataset" and the "Enefit Volt Dataset" validate the efficacy of our architecture in optimizing electric vehicle charging station recommendations. Notably, the inclusion of additional features in the Estonia dataset, such as price per kWh and outlet types, contributes to its superior performance compared to the Paris dataset, aligning with the goals of our project.

While our evaluation metrics for the Paris dataset indicate slightly lower scores for the Artificial Neural Network (ANN) model compared to Random Forest (RF) and K-Nearest Neighbors (KNN), it's crucial to emphasize the unique strengths of the ANN model. Unlike KNN's lazy-learning approach, ANN offers the advantage of providing real-time predictions, enhancing the adaptability and responsiveness of our system.

Overall, our results demonstrate the capability of our architecture to accurately manage BMS within EVs and optimize electric vehicle charging station recommendations. The selection of the most suitable model depends on the specific requirements of the task at hand, and the ANN model's ability to present outcomes as percentages for each distinct category is a valuable feature for real-world implementation. With the potential for future datasets to incorporate more complex features, our architecture is poised to deliver even greater performance enhancements.

5.2.3. Enhancing Battery Maintenance through Cloud-layer Predictive Analytics

In our design, the digital twin technology at the edge layer and the predictive analytics capabilities of the fog layer work together with the battery maintenance prediction service in the cloud layer to create a full ecosystem for managing EV batteries. While the edge layer focuses on real-time monitoring and initial processing of battery data, the fog layer optimizes charging station recommendations based on predictive analytics, and the cloud layer takes a proactive approach to battery management.

By harnessing advanced machine learning algorithms and analyzing data from EV Battery Management Systems (BMS), the cloud layer's battery maintenance prediction service enhances the overall efficiency and longevity of EV batteries. It complements the edge layer's real-time monitoring by proactively identifying potential maintenance needs and performance degradation trends. This predictive capability enables timely interventions, preventing costly breakdowns and maximizing battery lifespan.

Moreover, by leveraging the scalability and computational power of cloud computing, the service can handle large volumes of data and perform complex analyses to extract valuable insights. The integration of predictive analytics into the cloud layer not only ensures optimal battery health and performance, but also contributes to the reliability and sustainability of electric transportation systems.

Together, these layers create a synergistic architecture that addresses various aspects of EV battery management, from real-time monitoring to predictive maintenance, ultimately advancing the adoption of electric vehicles and promoting sustainable transportation solutions.

6. Conclusion

Deliverable D5.3 marks a significant milestone in the ENERGETIC project, presenting an advanced architecture designed to manage battery data efficiently while prioritizing resilience and security. Our overarching aim is to establish a robust framework that meets diverse stakeholder requirements, ensures platform usability, and fosters general acceptability.

This report provides a comprehensive overview of the system, highlighting critical components such as edge devices, data interfaces, and cloud management systems. We emphasize seamless connectivity and data flow from collection to transmission and consumption, underscoring the importance of continuous connectivity and adaptability.

Security is a paramount concern, and our proposed measures, including Transport Layer Security (TLS) for confidentiality and certificate-based authentication for data integrity and authenticity, ensure robust protection throughout transmission and storage. We carefully analyze the choice of transmission protocol, considering factors such as hardware support, real-time processing requirements, data volume, and cost constraints.

A key innovation introduced in this deliverable is the integration of Blockchain and Software-Defined Networking (SDN) technologies, enhancing security, transparency, and efficiency. Furthermore, we present innovative services deployed across edge, fog, and cloud layers, leveraging cutting-edge technologies to optimize battery management, predict maintenance needs, and enhance charging station recommendations.

These services signify a paradigm shift in battery data utilization, ensuring optimal performance and longevity across various use cases. We emphasize the dynamic nature of architecture deployment, stressing continuous adaptation to evolving requirements and end-user needs.

In conclusion, Deliverable D5.3 lays the groundwork for the effective utilization of battery data in the ENERGETIC project, providing a resilient and secure data management framework tailored to our specific needs. It sets the stage for subsequent development, testing, and implementation phases, highlighting the importance of resilience and security in advancing sustainable energy solutions.

7. References

- ENERGETIC Consortium Agreement

- ENERGETIC Grant Agreement

[1] H. R. Hasan, K. Salah, I. Yaqoob, R. Jayaraman, S. Pesic, M. Omar, Trustworthy iot data streaming using blockchain and ipfs, IEEE Access 10 (2022), 17707–17721. doi:10.1109/ACCESS.2022.3149312.

[2] B. Sellami, A. Hakiri, S. Ben Yahia, Deep reinforcement learning for energy-aware task offloading in join sdn-blockchain 5g massive iot edge network, Future Generation Computer Systems (2022).

[3] A. Rahman, M. J. Islam, S. S. Band, G. Muhammad, K. Hasan, P. Tiwari, Towards a blockchain-sdn-based secure architecture for cloud computing in smart industrial iot, Digital Communications and Networks 9 (2) (2023) 411–421.

[4] W. Zhao, S. Yang, X. Luo, On consensus in public blockchains, in: Proceedings of the 2019 International Conference on Blockchain Technology, ICBCT 2019, Association for Computing Machinery, New York, NY, USA, 2019, p. 1–5.

[5] b. City of Paris, Belib API. Available at: https://opendata.paris.fr/explore/dataset/belib-points-de-recharge-pour-vehicules-electriques-disponibilite-temps{%}0Aapi/?disjunctive.statutpdc{&}disjunctive.arrondissement.{&}disjunctive.arrondissement.

[6] Enefit VOLT. URL: https://enefitvolt.com/en/elektriauto-avalik-laadimine.