

CQRS and Blockchain with Zero-Knowledge Proofs for Secure Multi-Agent Decision-Making

Ayman NAIT CHERIF^{1*}, Mohamed YOUSSEFI², Zakariae EN-NAIMANI³,
Ahmed TADLAOUI⁴, Maha SOULAMI⁵, Omar BOUATTANE⁶
2IACS Laboratory, ENSET, Hassan II University, Casablanca, Morocco^{1, 2, 4, 5}
EESS Laboratory, ENSET, Hassan II University, Casablanca, Morocco^{3, 6}

Abstract—Autonomous decision-making in decentralized multi-agent systems (MAS) poses significant challenges related to security, scalability, and privacy. This paper introduces an innovative architecture that integrates Decentralized Identifiers (DIDs), Zero-Knowledge Proofs (ZKPs), Hyperledger Fabric blockchain, OAuth 2.0 authorization, and the Command Query Responsibility Segregation (CQRS) pattern to establish a secure, scalable, and privacy-focused framework for MAS. The use of DIDs and ZKPs ensures secure, self-sovereign identities and enables privacy-preserving interactions among autonomous agents. Hyperledger Fabric provides an immutable ledger, ensuring data integrity and facilitating transparent transaction processing through smart contracts. The CQRS pattern, combined with event sourcing, optimizes the system's ability to handle high volumes of read and write operations, enhancing performance and scalability. Practical applications are showcased in Smart Grids, Healthcare Data Management, Secure Internet of Things (IoT) Networks, and Supply Chain Management, highlighting the architecture's ability to address industry-specific challenges. This integration offers a robust solution for ensuring trust, verifiability, and scalability in distributed systems while preserving the confidentiality of agents.

Keywords—Decentralized multi-agent systems; decentralized identifiers; zero-knowledge proofs; hyperledger fabric; OAuth 2.0; CQRS; smart grids; healthcare data management; IoT; supply chain management

I. INTRODUCTION

In recent years, the rise of data-driven and compute-intensive applications has significantly increased the demand for scalable and decentralized systems capable of processing vast amounts of data in real-time. Multi-Agent Systems (MAS) have emerged as a robust solution for managing such complex environments. Comprising autonomous agents that collaborate to achieve common goals, MAS are integral to domains such as smart grids, supply chain management, and distributed computing, where they enable efficient, distributed decision-making [1]. However, these systems also introduce challenges related to security, scalability, and privacy [2]. Despite their potential, MAS face significant challenges in ensuring secure and private interactions among agents. Traditional security mechanisms, such as Public Key Infrastructure (PKI), are often inadequate for decentralized environments, leaving MAS vulnerable to attacks such as man-in-the-middle (MitM) and data breaches [3][4]. Additionally, the increasing need for privacy in distributed systems complicates interactions, particularly when sensitive data might be exposed during agent communication and

transaction processing. These challenges lead to critical research questions: how can secure communication be ensured in MAS to prevent MitM attacks and maintain data integrity? What mechanisms can provide privacy-preserving capabilities in decentralized systems while maintaining scalability? How can blockchain and Zero-Knowledge Proofs (ZKP) be effectively integrated to address these challenges in MAS? To address these questions, the objectives of this research are to develop a secure and privacy-preserving framework for MAS to safeguard agent interactions, leverage blockchain technology for data integrity and decentralized processing, employ ZKP to enhance privacy while maintaining system scalability and security, and optimize real-time decision-making through the adoption of Command Query Responsibility Segregation (CQRS) principles in MAS.

Blockchain technology offers a promising solution to these challenges by providing decentralized, tamper-proof ledgers that enhance data integrity [5]. This approach is valuable in sectors like finance and logistics, where accountability and transparency are critical. The immutability of blockchain transactions ensures that once data is recorded, it cannot be altered, providing a verifiable record of actions. However, blockchain's transparency can also reveal sensitive information, creating further privacy challenges [6]. To address both security and privacy, this paper proposes a novel framework integrating CQRS, blockchain, and ZKP. CQRS, introduced by Greg Young, separates commands (write operations) from queries (read operations), optimizing system scalability, especially in high-volume environments like MAS that require real-time processing. While blockchain addresses data integrity, it does not inherently protect agent identities. Zero-Knowledge Proofs offer a solution by enabling agents to verify transaction authenticity without revealing the underlying data. Integrating ZKP with blockchain ensures that only authorized agents can execute transactions while maintaining anonymity [7]. Recent advancements in ZKP, including zk-SNARKs, have bolstered blockchain privacy, enabling privacy-preserving solutions for sectors such as healthcare, IoT, and decentralized finance [8][9].

This paper presents a comprehensive framework that integrates CQRS, blockchain, and ZKP to address critical challenges in MAS, including security, scalability, and privacy. It offers a theoretical analysis of the proposed architecture's features and demonstrates its real-world applicability in areas such as Smart Grids, Healthcare Data Management, Secure IoT Networks, and Supply Chain Management. Additionally, the framework's effectiveness is validated through comparisons with existing state-of-the-art solutions, highlighting its

contributions to the field. The remainder of the paper is organized as follows. Section II and Section III provide a detailed overview of the background and related work, highlighting the limitations of existing solutions. This is followed by a presentation of the proposed framework, describing its architecture and components in Section IV. Section V and Section VI analyze the framework's security, privacy, and scalability features, demonstrate its application in real-world scenarios, and discuss its performance. The paper concludes in Section VII with an outline of potential future work.

II. RELATED WORK

The decentralized nature of Multi-Agent Systems (MAS) has led to their extensive use in distributed environments such as IoT networks, smart grids, and autonomous systems. MAS enable collaborative, autonomous decision-making but face critical challenges concerning security, scalability, and privacy. Numerous research efforts have been made to address these concerns, with a focus on authentication, data protection, and performance optimization through read/write operations.

A. Authentication and Authorization in MAS

Authentication is a cornerstone of MAS security, ensuring that agents are who they claim to be and that communications are protected. Traditionally, centralized methods like PKI have been used for authentication, but these introduce single points of failure and scalability issues [10]. Blockchain-based authentication mechanisms offer an alternative by leveraging decentralized smart contracts to eliminate trust intermediaries and prevent unauthorized access [11]. This decentralized approach is especially useful in environments where trust between agents cannot be guaranteed [12].

Despite these advancements, blockchain's high computational overhead for processing smart contracts poses challenges for real-time decision-making environments like IoT or smart grids. Moreover, these systems often lack fine-grained authorization mechanisms, which are crucial for handling multi-level access in dynamic, distributed environments. A solution to this problem, as proposed by He et al., involves a blockchain-based authentication scheme designed for mobile cloud computing, which introduces dynamic access control but struggles with resource efficiency [13].

B. Data Integrity and Privacy in Distributed MAS

In MAS, protecting data from tampering or loss is crucial. Blockchain's immutability and tamper-proof properties have been proposed as solutions to ensure data integrity [14]. Off-chain storage solutions, such as IPFS, have also been introduced to reduce on-chain storage costs, while still maintaining data integrity through cryptographic hashing. However, these systems face limitations, particularly when data stored off-chain is not protected by the same level of integrity as on-chain data.

Moreover, protecting data from loss due to network failures or agent disconnections is not adequately addressed in current off-chain storage models, which can lead to data tampering risks. To address these concerns, multi-copy data integrity auditing and batch auditing techniques in blockchain can help

ensure that data availability and integrity are preserved, even in distributed environments [15].

Oliveira et al. [16] develop a modular MAS architecture for distributed data mining, effectively handling scalability but lacking stronger privacy mechanisms. Qasem et al. [17] also focus on MAS-integrated data mining, highlighting issues with data privacy and suggesting the need for secure communication protocols like ZKPs. Similarly, Nait Cherif et al. [18] propose a blockchain-based solution for data integrity, though it requires significant computational resources, which may be alleviated by the CQRS-based system in this work. Ge et al. [19] survey advancements in distributed sampled-data cooperative control of MAS, emphasizing different sampling mechanisms to improve performance, yet these methods often lack real-time adaptability, posing a scalability bottleneck. In comparison, the CQRS pattern in this work supports high throughput by separating read and write operations, thus enhancing real-time performance under high loads.

C. Scalable Resource Management in Multi-Agent Systems

Efficient read/write operations are critical in distributed MAS, especially as the number of transactions scales. The CQRS pattern has been widely adopted to separate read and write operations, thereby optimizing system performance. This is particularly beneficial in high-volume transaction environments, where large amounts of data must be processed in real-time [20] [21].

However, while CQRS improves scalability, it does not inherently provide protection against man-in-the-middle (MitM) attacks or secure write operations, exposing the system to security threats [22]. Additionally, ensuring data consistency across distributed agents, especially in real-time applications, remains an unresolved issue, as CQRS alone does not guarantee that agents will have synchronized access to the latest data [23].

Dynamic and scalable MAS architectures benefit from integrating CQRS with blockchain technology. For example, Dashti et al. [24] introduce a MAS framework with dynamic agent capabilities but struggle with issues related to agent turnover. Our approach leverages blockchain for immutable tracking and secure data integrity, mitigating this issue. Similarly, Breugnot et al. [25] propose a distributed graph structure for load balancing in MAS but encounter challenges with data synchronization. By implementing CQRS in our framework, we enhance the synchronization process, enabling seamless data processing across distributed nodes.

Control and optimization are also central to MAS design, particularly in adversarial or resilience-focused settings. Wang [26] highlights the importance of distributed control but does not adequately address security against adversarial threats. In contrast, our approach integrates Decentralized Identifiers (DIDs) for secure agent identification and Zero-Knowledge Proofs (ZKPs) for verifiable, tamper-resistant interactions. This architecture not only improves resilience against data tampering but also maintains robust state management through CQRS, addressing challenges like those noted by Rust et al. [27] in state persistence across agents. Furthermore, Fanitabasi [28] discusses optimization in MAS under adversarial conditions, underscoring the need for resilience. By combining CQRS with

DIDs and ZKPs, our framework provides enhanced security and reduces vulnerabilities to malicious agents.

Lastly, transparency and secure communication in resource allocation are vital for MAS efficiency. Fu and Zhou [29] present a MAS solution for multi-project scheduling, although they struggle with transparency in resource allocation. Our framework addresses this by using blockchain to ensure verifiable transaction records, reducing information asymmetry. Additionally, Costa et al. [30] offer secure communication protocols that are limited in scalability under high agent loads. Leveraging CQRS, our architecture improves message throughput, enhancing secure, efficient communication across distributed systems.

D. Thesis for Improvement

Despite the progress made, several weaknesses in the current research need to be addressed:

- **Scalability and Real-Time Challenges:** Current blockchain-based authentication models suffer from scalability issues, as high computational costs limit their applicability in environments where real-time.
- **Data Integrity for Off-Chain Storage:** The reliance on off-chain storage solutions, such as IPFS, creates vulnerabilities in data tampering and data loss, as these systems lack the same level of cryptographic protection as on-chain storage.
- **Privacy Concerns:** While blockchain ensures data transparency, it falls short in protecting the privacy of agent interactions. Even with the use of Zero-Knowledge Proofs (ZKP), the high computational burden and the complexity of implementing ZKP protocols in large-scale systems make it difficult to achieve both privacy and performance.

This paper proposes a framework that combines CQRS, blockchain, zero-knowledge proofs (ZKP), and OAuth2 to address the limitations identified in existing solutions:

- **Optimized Scalability:** By enhancing ZKP protocols and integrating OAuth2 for efficient authentication, we reduce computational overhead and improve scalability, making the system suitable for high-frequency, real-time applications.
- **Enhanced Data Protection:** Utilizing blockchain's inherent immutability and tamper-proof properties, we ensure that all data remains secure and unaltered on-chain, eliminating the need for off-chain storage and its associated risks.
- **Securing Write Operations:** Through advanced cryptographic techniques and the integration of OAuth2 for secure authentication, we protect against man-in-the-

middle attacks and ensure consistent data synchronization across distributed agents.

By addressing these gaps, our framework presents a more scalable, secure, and privacy-preserving solution for high-frequency, real-time multi-agent systems. This integration of advanced cryptographic methods and established design patterns offers a robust approach to the challenges facing MAS in complex, distributed environments.

III. BACKGROUND

In this section, we explore the foundational technologies and concepts that form the basis of the proposed framework: Multi-Agent Systems (MAS), Blockchain, Zero-Knowledge Proofs (ZKP), Command Query Responsibility Segregation (CQRS), Event-Driven Architecture (EDA), and OAuth 2.0. Together, these technologies address critical challenges like scalability, security, data integrity, and privacy in decentralized systems. Each concept builds upon the others to enhance the robustness and efficiency of MAS.

A. Multi-Agent Systems (MAS)

Multi-Agent Systems (MAS) consist of several autonomous agents working collaboratively or competitively to achieve objectives that are often too complex for a single agent or centralized system to handle. These agents interact with one another to solve problems in distributed environments, such as IoT networks, smart grids, and supply chain management. MAS are particularly effective in situations where decentralized decision-making is required, and the agents can operate independently to respond to changing conditions in real-time [31].

However, while MAS offer significant advantages in scalability and flexibility, they also introduce substantial challenges—most notably, ensuring secure communication and anonymity between agents. As agents in MAS frequently exchange sensitive information, protecting this data from unauthorized access or attacks becomes paramount. Traditional security measures often fall short in such decentralized environments, necessitating more advanced cryptographic techniques like Zero-Knowledge Proofs (ZKP). In this context, ZKP provides a solution that enhances both security and privacy within MAS by allowing authentication and access control without revealing underlying data, thereby ensuring secure and anonymous exchanges.

B. Event-Driven Architecture (EDA) and Event Sourcing

To enable more flexible and scalable interactions within MAS, many systems adopt an Event-Driven Architecture (EDA). Event-sourcing and event-driven architecture are two related but distinct concepts that are often used together in software systems. Event-driven architecture (EDA) is a design pattern that involves building a system in which different components communicate with each other by generating and reacting to events [32] (Fig. 1).

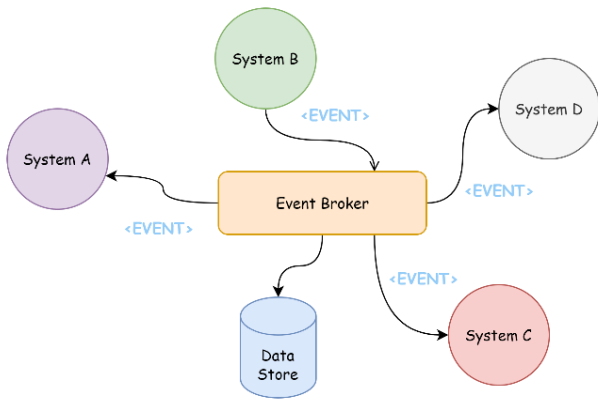


Fig. 1. Event driven architecture.

For example, when an agent in MAS detects a change in the environment or receives new data, it can trigger an event that other agents can subscribe to and react to as needed. This decoupling of event producers and consumers enables the system to scale more effectively and enhances its fault tolerance [33].

Closely related to EDA is Event Sourcing, which provides a mechanism to maintain a historical log of all state changes within the system. Rather than storing the current state of the system, event sourcing records every change as an event, allowing the system's state to be reconstructed by replaying these events in sequence. This feature is particularly useful in MAS for auditing, debugging, and recovering from failures. By integrating event sourcing with EDA, MAS can ensure that state changes are both traceable and resilient to failure [34].

The combination of EDA and Event Sourcing complements the scalability of MAS while providing a reliable mechanism to track and react to changes across distributed agents. However, while these architectures improve the system's flexibility, they do not address the security and data integrity challenges associated with state changes. This is where Blockchain plays a crucial role.

C. Command Query Responsibility Segregation (CQRS)

Building on the flexibility provided by EDA and event sourcing, Command Query Responsibility Segregation (CQRS) further enhances the scalability of MAS by separating read operations from write operations (Fig. 2).

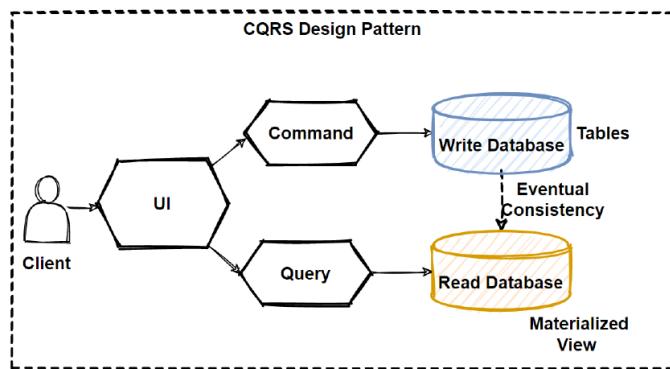


Fig. 2. CQRS design pattern.

In MAS, agents often need to perform a large number of both read (query) and write (command) operations to maintain and update the system state in real-time. By segregating these responsibilities, CQRS allows systems to scale independently for these two functions, optimizing system performance and reducing bottlenecks [35].

While CQRS improves the efficiency of MAS, particularly in high-transaction environments, it does not inherently provide security for write operations. Without adequate safeguards, the system remains vulnerable to tampering during state changes, which could lead to unauthorized alterations of critical data. At this point, Blockchain becomes essential for securing these write operations, as its tamper-proof ledger ensures that all changes to the system's state are immutably recorded and can be verified by all agents in the system [36].

The integration of CQRS with Blockchain strengthens the system's security, but it also introduces privacy concerns, particularly in public blockchains where data is visible to all participants. This necessitates the use of Zero-Knowledge Proofs (ZKP), which provide privacy-preserving mechanisms for MAS, ensuring that agents can interact securely without exposing sensitive information [37].

D. Blockchain Technology

Blockchain is a decentralized, distributed ledger that keeps a growing list of records called blocks. In a chain of blocks (hence the name "blockchain"), each block has multiple attributes including a timestamp, a link to the previous block, and its own data [38] (Fig. 3). Cryptographic techniques secure this chain of blocks, making it nearly impossible to alter its data. [38].




Fig. 3. Block structure.

Without the need for a centralized authority, transactions can be made using blockchain technology in a secure and transparent manner. Due to the distributed nature of blockchain data, it is virtually impossible for one entity to alter it without being detected. Because of this, blockchain technology is well suited for uses like supply chain management and financial transactions that demand a high level of transparency and trust [39]. In a typical blockchain system, transactions are initiated by users and are broadcast to the network. The integrity of these transactions is ensured by network nodes. A block is then added to the blockchain after a transaction has been verified and added to it [40].

However, while blockchain excels at maintaining data integrity and transparency, it introduces a new set of challenges regarding privacy. In a traditional blockchain system, transaction details are visible to all participants, which may not

be acceptable in scenarios where agents need to exchange sensitive data. To mitigate this issue, Zero-Knowledge Proofs (ZKP) are integrated with blockchain, allowing agents to prove that they have valid data or authorization without revealing the underlying information. This combination ensures that blockchain retains its transparency and security, while also protecting the privacy of the agents involved [41] [42].

E. Zero-Knowledge Proofs (ZKP)

Zero-Knowledge Proofs (ZKP) offer a solution to the privacy and anonymity concerns that arise in MAS when agents exchange sensitive information. It was initially brought forth in the 1980s by Shafi Goldwasser, Silvio Micali, and Charles Rackoff [43]. ZKP allows an agent to prove the validity of a claim (such as their identity or authorization) without revealing any additional details. This cryptographic technique is particularly valuable in blockchain-based MAS, where agents need to maintain both transparency and privacy during transactions.

The Zero-Knowledge Proof (ZKP) workflow, depicted in Fig. 4, illustrates the key phases of the ZKP process: setup, commitment, challenge, response, and verification. This workflow ensures that agents can authenticate their claims without revealing sensitive information, preserving both privacy and security.

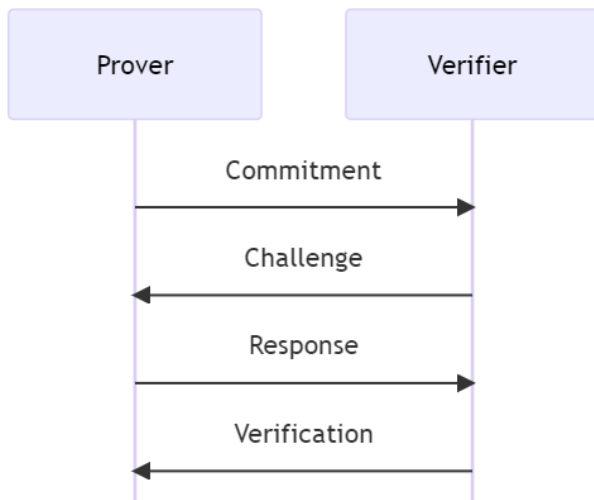


Fig. 4. ZNP workflow.

For instance, in a system where agents need to prove ownership of certain data or assets, ZKP enables them to authenticate their claims without exposing the data itself. This ensures that even as agents participate in a decentralized system like blockchain, they can retain their privacy, preventing sensitive information from being exposed to other participants [44]. However, implementing ZKP at scale presents computational challenges, as the process can be resource-intensive, making optimization a key area of research for large-scale MAS.

The benefits of ZKP in ensuring privacy and security are further enhanced when combined with OAuth 2.0, an authorization framework that allows controlled access to resources without exposing user credentials [18].

F. OAuth 2.0

OAuth 2.0 is an authorization framework widely used to secure access to resources by enabling third-party applications to interact with systems on behalf of users. In the context of MAS, OAuth 2.0 can manage access control, allowing agents to interact securely with external services without sharing their credentials directly. However, OAuth 2.0 alone does not guarantee anonymity, as it can still link access tokens to specific agents, potentially exposing their identities [45].

Fig. 5 demonstrates the OAuth 2.0 protocol flow, highlighting the interaction between agents, the authorization server, and the resource server. This process facilitates secure access token issuance and validation, ensuring controlled and authenticated access to system resources. The integration of OAuth 2.0 with ZKP in the proposed architecture further strengthens privacy by decoupling sensitive user credentials from access authorization.



Fig. 5. OAuth 2 protocol flow.

To address this limitation, integrating ZKP with OAuth 2.0 provides a powerful solution for secure and anonymous resource access. By using ZKP to authenticate agents without revealing their identities, OAuth 2.0 can enable secure communication while preserving agent anonymity. This combination not only enhances the security of MAS but also ensures that agents can interact with external services without compromising their privacy.

The integration of OAuth 2.0 and ZKP represents a novel approach to solving the challenges of access control and privacy in MAS. This hybrid mechanism allows for secure data access while maintaining the anonymity of the agents, offering a robust solution to the security concerns present in distributed multi-agent systems [46].

IV. PROPOSED ARCHITECTURE

The proposed architecture combines Command Query Responsibility Segregation (CQRS), blockchain, and Zero-Knowledge Proofs (ZKP) to address the core challenges of security, scalability, and privacy in Multi-Agent Systems (MAS).

Existing architectures often fall short in meeting these requirements due to their inherent limitations, such as inadequate privacy mechanisms, reliance on centralized infrastructures, and performance bottlenecks in high-volume environments.

Blockchain ensures data integrity through its immutable ledger, but its transparency can compromise sensitive information, highlighting the need for enhanced privacy measures.

To overcome this, the integration of ZKP enables secure agent authentication and transaction verification without exposing sensitive data, ensuring privacy-preserving interactions. CQRS further enhances the framework by segregating read and write operations, optimizing performance and scalability in real-time, high-demand scenarios.

Additionally, the architecture addresses vulnerabilities associated with centralized Public Key Infrastructure (PKI) by incorporating Decentralized Identifiers (DIDs), reducing single points of failure and enhancing system resilience. Unlike existing systems that struggle with consistent data synchronization and computational inefficiencies, the modular and interoperable design of the proposed framework ensures adaptability across diverse domains such as IoT, healthcare, and smart grids. This holistic integration of technologies provides a scalable, secure, and privacy-preserving solution tailored to the evolving demands of decentralized MAS.

A. Need for Enhancements in MAS Architecture

Understanding the limitations of traditional MAS architectures highlights the necessity for a paradigm shift. Addressing these challenges is crucial for developing systems that meet modern requirements.

The proposed architecture aims to:

- **Eliminate Single Points of Failure:** By decentralizing identity and access control mechanisms, the system enhances resilience and reduces dependency on any single component.
- **Improve Scalability:** Through asynchronous communication and efficient processing models like CQRS, the architecture supports the seamless addition of agents without compromising performance.
- **Strengthen Security and Privacy:** Using advanced cryptographic techniques, secure communication protocols, and privacy-preserving authentication methods, the system safeguards agent interactions.
- **Ensure Data Integrity:** Leveraging blockchain technology for immutable and verifiable data storage ensures that all agents have a consistent and trusted view of the system's state.

B. System Components and Architecture

1) *Agent:* Agents serve as autonomous entities within the decentralized system, representing unique participants such as users, services, or applications. Each agent is capable of performing actions, communicating with other agents, and managing its own state. The primary purpose of an agent is to interact seamlessly within the network. Functionally, agents handle identity management by generating and managing cryptographic key pairs and Decentralized Identifiers (DIDs). They utilize Zero-Knowledge Proofs (ZKP) and OAuth 2.0 tokens for authenticating and authorizing interactions.

Communication between agents is secured through encrypted and signed messages transmitted via Hyperledger Fabric channels. Agents manage their state by executing commands that alter their state and recording corresponding events on the blockchain. Additionally, they can reconstruct their internal state by replaying events from the blockchain, using snapshotting techniques for efficiency. In terms of interactions, agents communicate with other agents through secure channels, interact with the Decentralized Discovery Facility (DDF) for agent discovery, and submit and retrieve events from the blockchain for state management.

2) *Decentralized Identifier (DID) and DID document:* DIDs provide a decentralized and self-sovereign identity framework for agents, enabling secure and verifiable interactions without relying on centralized authorities. Agents generate unique DIDs following standardized specifications, such as Hyperledger-compatible DID methods. Each DID is associated with a DID Document, which contains public keys, authentication methods, and service endpoints. These documents are published on the blockchain for public reference, facilitating secure communication and verification among agents. DIDs are shared among agents during interactions, utilized during registration with the Decentralized Discovery Facility (DDF), and play a crucial role in authentication processes.

3) *Zero-Knowledge Proofs (ZKP) processing service:* The ZKP Processing Service enhances privacy and security by enabling agents to prove possession of certain information, such as ownership of a DID, without revealing the information itself. This service assists agents in generating ZKPs to demonstrate control over their private keys and DIDs. It facilitates the various phases of the ZKP protocol, including setup, commitment, challenge, response, and verification. The service collaborates with agents during registration and authentication processes and interfaces with the Authorization Server during the issuance of OAuth 2.0 tokens.

4) *Decentralized Discovery Facility (DDF):* The DDF acts as a decentralized registry and discovery service, allowing agents to locate and interact with other agents and their services within the network. It handles the registration of agent DIDs and associated metadata after verifying ZKPs. The discovery functionality provides searchable access to registered agents and their services, enabling agents to find peers and available functionalities efficiently. The DDF receives registrations from agents along with their ZKPs and DID Documents and responds to discovery queries from agents seeking information about other participants.

5) *OAuth 2.0 authorization server:* The OAuth 2.0 Authorization Server is responsible for managing the issuance and validation of access tokens, thereby facilitating secure authorization for agents to access resources and communication channels within the system. It processes OAuth 2.0 authorization requests from agents, incorporating DIDs and ZKP commitments. Upon successful verification of ZKPs and issuance of authorization codes, the server generates JWT

access tokens. It also validates incoming tokens during resource access requests to ensure their authenticity and permissions. The Authorization Server interacts with agents during the token issuance process and interfaces with Resource Servers to validate access tokens presented by agents.

Fig. 6 presents the high-level interaction among components within the proposed architecture. It illustrates the communication flow between agents, the Decentralized Discovery Facility (DDF), blockchain, and resource servers. This diagram emphasizes how the architecture integrates multiple technologies to ensure seamless operation, secure agent authentication, and efficient resource management.

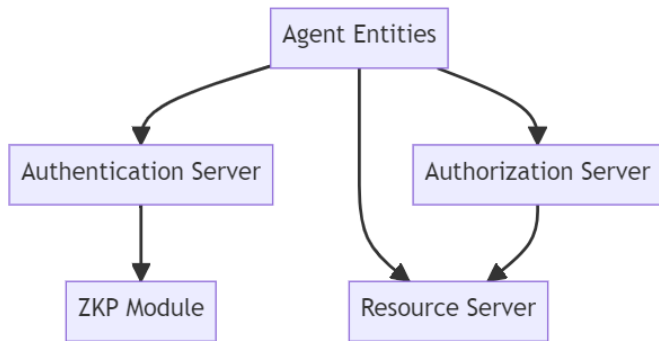


Fig. 6. Component interaction.

6) *Hyperledger fabric network*: The Hyperledger Fabric Network provides a permissioned blockchain infrastructure that ensures secure, immutable, and scalable communication and state management among agents. It utilizes private communication channels designated for specific groups of agents or types of interactions. The network comprises peers, which are nodes that host the ledger and execute smart contracts (chaincode), maintaining the blockchain's state, and orderers, which are nodes responsible for ordering transactions into blocks to ensure consistency and reliability across the network. Subcomponents include smart contracts that define the business logic for processing transactions, handling commands, and managing state changes, as well as the ledger, which serves as an immutable record of all transactions and events, ensuring data integrity and transparency. The Hyperledger Fabric Network facilitates message exchange between agents through its channels, records state-changing events submitted by agents to ensure immutability and verifiability and provides a reliable ledger for agents to reconstruct their state through event sourcing.

7) *Resource servers*: Resource Servers host protected resources such as data stores, computation services, or external APIs, and enforce access controls based on OAuth 2.0 tokens. They handle resource requests from agents seeking access to these protected resources by receiving and processing these requests. The servers validate the authenticity, validity, and permissions of JWT access tokens included in resource requests. Based on the validated tokens and associated

permissions, Resource Servers grant or deny access to the requested resources. They communicate with the Authorization Server to validate access tokens and interact with agents to provide access to requested resources upon successful validation.

8) *Command Query Responsibility Segregation (CQRS) components*: The CQRS Components enhance system performance and scalability by separating read and write operations, allowing independent optimization of data handling. Command Handlers process state-changing commands submitted by agents, invoking smart contracts to record events on the blockchain. Query Handlers manage read operations by accessing optimized, separate data stores that provide quick and efficient data retrieval. An Event Store maintains a log of all events generated by command executions, enabling state reconstruction and auditability. These components receive commands from agents via Fabric channels, update read models based on events recorded in the blockchain, and support agents in querying the current state without interference from write operations.

9) *Hardware Security Modules (HSMs)*: Hardware Security Modules (HSMs) are critical for securely storing cryptographic keys and performing sensitive operations, thereby protecting against unauthorized access and tampering. They safeguard private keys used by agents for signing messages and events, ensuring that these keys are never exposed in plaintext. HSMs execute cryptographic functions within a protected environment, maintaining the integrity and confidentiality of keys. Agents utilize HSMs to securely manage their cryptographic materials, and HSMs integrate with the DID generation and signing processes to ensure the security and trustworthiness of cryptographic operations within the system.

The integration of these components results in an architecture that significantly enhances the security, scalability, and privacy of Multi-Agent Systems. By addressing the limitations of traditional approaches through decentralized identity management, secure communication protocols, optimized processing models, blockchain integration, advanced cryptographic techniques like ZKP combined with OAuth, and decentralized access control mechanisms, agents can interact in a trustworthy and efficient manner. The proposed architecture lays a robust foundation for developing MAS that are resilient, scalable, and capable of meeting the complex demands of modern distributed environments.

C. Detailed Workflow from Agent Authentication to Message Exchange

1) *Agent initialization and DID generation*: When an agent joins the system, it begins by generating a public-private key pair. This key pair is fundamental to the agent's cryptographic identity, enabling secure communication and authentication within the network. To ensure the private key remains secure, the agent stores it in a Hardware Security Module (HSM) or a Trusted Execution Environment (TEE). These storage solutions protect the key from unauthorized access and tampering.

Next, the agent creates a Decentralized Identifier (DID) following standard specifications, such as those compatible with Hyperledger. This DID represent a unique, self-sovereign identity that operates independently of any centralized authority. By generating multiple DIDs, the agent enhances its privacy and minimizes the risk of being tracked across different interactions.

The agent then constructs a DID Document, which includes its public keys and service endpoints. This document is published on the blockchain, making the agent’s identity accessible to other agents within the network. The DID Document serves as a public reference, describing the agent’s DID and associated metadata to facilitate secure and verifiable interactions. This initialization process ensures that the agent maintains full control over its identity, promotes interoperability through standardized identification formats, and establishes a secure foundation for decentralized interactions.

2) *Registration with Decentralized Discovery Facility (DDF) using Zero-Knowledge Proofs (ZKP)*: To register with the Decentralized Discovery Facility (DDF), an agent must prove ownership of its DID without revealing its private key. This is achieved using Zero-Knowledge Proofs (ZKP). The agent collaborates with a ZKP Processing Service to generate a proof that demonstrates control over its private key in a way that preserves confidentiality.

The ZKP process involves several key steps:

- **Setup Phase:** Public parameters are defined, including a large prime number p , a generator g , and a derived parameter h (1), where s is the agent’s secret.
- **Commitment:** The agent selects a random nonce r and computes the commitment C (2), binding itself to the secret without revealing it.
- **Challenge:** The verifier issues a random challenge C to the agent.
- **Response:** The agent calculates z (3), proving knowledge of the secret without disclosing it.
- **Verification:** The verifier checks (4). If the equality holds, the verifier is convinced that the agent knows the secret without learning its value.

$$h = g^s \text{ mod } p \tag{1}$$

$$C = g^r \text{ mod } p \tag{2}$$

$$z = r + c \cdot s \text{ mod } (p - 1) \tag{3}$$

$$g^z \text{ mod } p = C \cdot h^c \text{ mod } p \tag{4}$$

After successfully generating the ZKP, the agent submits its DID Document along with the ZKP to the DDF via smart contracts. The DDF verifies the proof and records the successful registration on the blockchain. This process ensures that only legitimate agents can register, maintaining the integrity and security of the decentralized system.

3) *Agent discovery*: When an agent wants to discover other agents and their services, it queries the DDF for information about other agents’ DIDs and available services. This querying

process facilitates collaboration and the utilization of services within the network. Importantly, agents have the option to selectively disclose specific attributes or services, allowing them to control the information they share and protect sensitive data. By enabling discovery while preserving privacy, the system fosters collaboration among agents without compromising their autonomy or exposing unnecessary information.

4) *Token issuance with OAuth 2.0 including DIDs*: After successfully registering and authenticating, an agent can obtain an access token through the OAuth 2.0 protocol, which includes its DID. The issuance of tokens is a critical component for several reasons. Firstly, tokens provide a secure and standardized method for managing permissions and access controls within the decentralized system.

By issuing tokens that encapsulate DIDs, the system ensures that agents can be reliably identified and granted specific permissions to access various resources without the need for repeated authentication processes. This streamlines access management, reduces the overhead associated with continuous verification, and enhances overall system security by limiting the exposure of sensitive information.

The token issuance process begins with the agent preparing an OAuth 2.0 authorization request, incorporating its DID and ZKP commitment. The Authorization Server then issues a challenge C , and the agent responds with (3), adhering to the ZKP protocol.

Upon verifying the ZKP, the server issues an authorization code, which the agent exchanges for a JWT access token containing its DID and permissions. This integration of OAuth 2.0 with ZKP ensures a secure and standardized method for managing permissions, allowing reliable agent identification without exposing sensitive information.

The integration of Zero-Knowledge Proofs (ZKP) and OAuth 2.0 is depicted in Fig. 7. This diagram shows how ZKP enhances the OAuth 2.0 authorization process by enabling agents to authenticate their interactions without revealing sensitive data. The secure issuance of access tokens, combined with privacy-preserving proof mechanisms, reinforces the system's security and ensures robust resource access control.

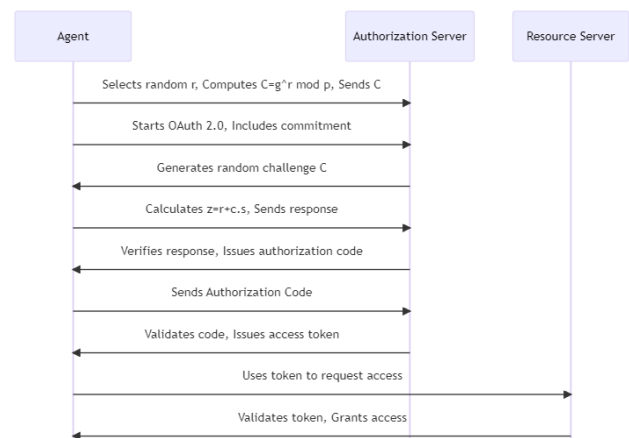


Fig. 7. ZKP and OAuth 2.0 flow.

5) *Resource access and verification using DIDs*: When an agent needs to access protected resources, it presents its JWT access token containing its DID in the resource request, typically within the Authorization header as a Bearer token. The resource server then validates the token's signature, ensures it has not expired, checks the scopes, and verifies the agent's DID. If the token is valid and the agent possesses the necessary permissions, the resource server grants access to the requested resources. This mechanism ensures that resource access is secure and controlled, leveraging DIDs for reliable identification and OAuth 2.0 for robust authorization management.

6) *Agent communication, data sharing, and state management workflow*: If an agent wants to send a message to other agents or read messages, it utilizes Hyperledger Fabric channels alongside the Command Query Responsibility Segregation (CQRS) pattern to ensure efficient and secure communication and state management. The agent begins by authenticating using ZKP during registration and obtaining OAuth 2.0 access tokens, which grant it permissions to interact with specific Fabric channels designated for communication.

Fig. 8 illustrates the detailed workflow for agent communication, data sharing, and state management within the architecture. By leveraging Hyperledger Fabric channels and the CQRS pattern, the diagram demonstrates how agents efficiently interact and manage their states while ensuring data integrity and consistency. This flow is critical for maintaining real-time performance and scalability in high-demand applications.

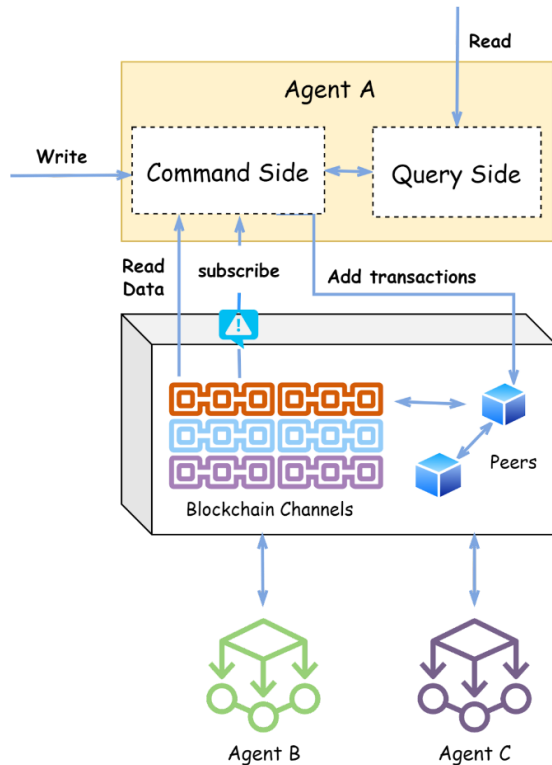


Fig. 8. Agent communication, data sharing, and state management workflow.

a) *Sending a message*: To send a message, the agent connects to the appropriate Fabric channel using its access token. It encrypts the message content with the recipient's public key and signs the message with its private key to ensure confidentiality and authenticity. The message, along with relevant metadata, is then published to the Fabric channel, where the blockchain handles routing based on the message's metadata and the channel's configuration.

b) *Receiving a message*: Receiving agents connect to the relevant Fabric channels using their access tokens, subscribe to specific channels, and asynchronously receive messages. Upon receiving a message, the agent verifies the signature using the sender's public key from the DID Document and decrypts the content using its private key. The message is then processed according to the agent's business logic.

c) *Event recording and state management*: When an agent executes a command that changes its state, it generates an event representing that change. The event includes details such as the event type (e.g., "OrderPlaced", "BalanceUpdated"), a timestamp recording when the event occurred, a data payload containing relevant information, and metadata like causation ID, correlation ID, or version. To ensure authenticity and prevent tampering, the agent signs the event with its private key.

The signed event is then packaged into a transaction proposal and submitted to the blockchain network. This involves creating and signing a transaction proposal that reflects the internal state changes in the form of an event. The proposal is sent to the network for validation and inclusion in the blockchain. Blockchain nodes, or peers, validate the transaction proposal by ensuring proper formatting, verifying the agent's signature, and performing authorization checks to confirm that the agent has the rights to perform the action represented by the event.

Once validated, the ordering service sequences the transactions and packages them into blocks. These blocks are then disseminated across the network and added to each node's copy of the blockchain ledger, making the event part of the immutable history of the system.

d) *State reconstruction*: To maintain data integrity and enable agents to recover their state independently, the system employs event sourcing and blockchain integration. When an agent needs to reconstruct its internal state, it retrieves the sequence of relevant events from the blockchain and replays them in chronological order. This process ensures that the agent's state is consistent with the system's history. For efficiency, agents may use snapshotting, creating snapshots of their state after processing a certain number of events. Future state reconstructions can begin from the latest snapshot and replay only subsequent events, ensuring consistency and integrity with the system's history.

This integrated approach ensures that communication is secure and efficient while maintaining data integrity and providing a reliable audit trail through event sourcing.

7) *Conclusion*: This comprehensive workflow integrates decentralized identity management, secure authentication, scalable communication, and robust data integrity mechanisms.

By leveraging Zero-Knowledge Proofs, Decentralized Identifiers, Hyperledger Fabric, OAuth 2.0, and the CQRS pattern, the system ensures that agents can interact securely, efficiently, and privately within a decentralized multi-agent environment.

Security and privacy are maintained throughout the workflow using advanced cryptographic techniques. Zero-Knowledge Proofs and OAuth 2.0 provide robust authentication and authorization without exposing sensitive information. Blockchain technology ensures data integrity by offering immutable and tamper-proof storage of events and identities, while encryption safeguards data both in transit and at rest. The CQRS pattern and event sourcing facilitate efficient state management and fault tolerance, allowing agents to recover their state independently by replaying blockchain-recorded events.

The system's design supports scalability and performance through asynchronous communication and the segregation of read and write operations. This allows the network to handle high transaction volumes and numerous agents efficiently. Additionally, the use of standardized protocols and identifiers promotes interoperability, enabling seamless integration with external systems. Compliance and auditability are achieved through transparent and immutable records on the blockchain, facilitating regulatory adherence and providing comprehensive audit trails for accountability and governance.

D. Additional Security Measures

To ensure the integrity, confidentiality, and availability of its components, the system incorporates robust security measures. These measures cover key management, smart contract security, and comprehensive monitoring and incident response protocols.

1) *Key management and credential security*: Key management protocols are designed to safeguard cryptographic keys, ensuring that only authorized entities can perform sensitive operations within the system.

a) *Secure storage*:

- **Hardware Security Modules (HSMs) and Trusted Execution Environments (TEEs)**: Private keys are securely stored within HSMs, preventing unauthorized access and extraction. These modules offer physical and logical protections against theft, tampering, and malware attacks.
- **Access Controls**: Strict access control policies are in place, with role-based access controls (RBAC) ensuring that only authorized personnel and processes can access sensitive keys.

b) *Key rotation and revocation*:

- **Regular Key Rotation**: Keys are rotated periodically to minimize exposure risk. Upon rotation, new keys are generated, and corresponding DID Documents are promptly updated on the blockchain.
- **Immediate Revocation**: Protocols are established for rapid key revocation in case of a suspected compromise. All agents are ensured access to the latest DID Documents to verify key validity and prevent the use of revoked keys.

c) *Access control policies*:

- **Least Privilege Principle**: Agents and users are granted only the permissions necessary to perform their roles, reducing the attack surface.
- **Segregation of Duties**: Responsibilities are divided among different agents to prevent conflicts of interest and limit the impact of compromised credentials.

2) *Smart contract security*: Smart contracts are rigorously secured to prevent vulnerabilities that could lead to financial losses or unauthorized state changes.

a) *Development best practices*:

- **Secure Coding Standards**: Smart contracts adhere to established coding standards to prevent common vulnerabilities, with regular updates based on the latest security research.
- **Use of Established Libraries**: Trusted libraries and frameworks are used to minimize the risk of bugs and vulnerabilities.
- **Modular Design**: Smart contracts are structured into smaller, manageable modules for easier analysis, testing, and maintenance.
- **Independent Security Audits**: Third-party security experts conduct comprehensive reviews of smart contract code to identify vulnerabilities and verify implemented security measures.

b) *Continuous testing*: Automated testing pipelines with unit tests, integration tests, and fuzz testing are implemented to identify vulnerabilities early.

c) *Upgradability and governance*:

- **Secure Upgrade Paths**: Smart contracts include upgradeable patterns to allow controlled updates when necessary, with mechanisms in place to prevent unauthorized modifications.
- **Governance Mechanisms**: Clear governance processes are established for smart contract changes, involving relevant stakeholders to ensure transparency and collective decision-making.

3) *Monitoring and incident response*: Continuous monitoring and incident response protocols are established to detect and address security breaches promptly, ensuring system resilience.

a) *Continuous monitoring*:

- **Intrusion Detection Systems (IDS)**: IDS monitors network traffic and system logs for signs of unauthorized access or malicious activities, using both signature-based and anomaly-based detection.
- **Logging and Alerts**: Comprehensive logs of system activities are maintained, with real-time alerting mechanisms to notify administrators of critical events.

b) *Threat Intelligence and Updates*

- **Stay Informed:** The system stays updated on emerging threats through security advisories and threat intelligence feeds.
- **Patch Management:** Security patches and updates are applied promptly based on vulnerability severity and potential system impact.

V. SECURITY, SCALABILITY AND PERFORMANCE ANALYSIS

Ensuring that a decentralized multi-agent system operates securely, scales efficiently, and maintains high performance is paramount for its success and reliability. This analysis delves into how the system's architecture and components collectively achieve these objectives, highlighting strengths, potential challenges, and the interplay between different system aspects.

A. Security Analysis

1) *Comprehensive security framework:* The system employs a multi-layered security approach, integrating advanced cryptographic techniques, secure authentication and authorization protocols, and robust key management practices. By leveraging Zero-Knowledge Proofs (ZKP) and Decentralized Identifiers (DIDs), agents can authenticate and authorize interactions without exposing sensitive information, significantly reducing the risk of credential theft and unauthorized access.

2) *Immutable ledger and data integrity:* Hyperledger Fabric's blockchain infrastructure ensures that all transactions and state changes are immutably recorded, providing a tamper-proof audit trail. Digital signatures and secure storage mechanisms (HSMs/TEEs) further reinforce data integrity and authenticity, ensuring that only authorized agents can perform state-altering actions.

3) *Smart contract security:* Adhering to secure coding standards, utilizing established libraries, and implementing formal verification and independent audits fortify smart contracts against common vulnerabilities. The incorporation of modular design and secure upgrade paths allows the system to adapt and evolve without compromising security, addressing potential threats proactively.

4) *Proactive threat detection and incident response:* Continuous monitoring through Intrusion Detection Systems (IDS) and anomaly detection algorithms enables real-time identification of suspicious activities. A well-defined incident response plan ensures that the system can swiftly contain and mitigate threats, minimizing potential damage and maintaining operational integrity.

5) *Privacy preservation:* By enabling selective disclosure and employing encryption for all data exchanges, the system upholds strong privacy guarantees. Agents maintain control over their personal information, and the use of multiple DIDs prevents tracking and profiling, aligning with privacy-by-design principles and regulatory requirements.

B. Scalability Analysis

1) *Command Query Responsibility Segregation (CQRS) pattern:* The implementation of the CQRS pattern effectively

separates read and write operations, allowing each to be optimized independently. This segregation reduces contention and enhances system throughput, enabling the system to handle a high volume of transactions and queries without performance degradation.

2) *Asynchronous communication via hyperledger fabric channels:* Hyperledger Fabric channels facilitate parallel and isolated communication pathways, allowing multiple groups of agents to interact concurrently without interference. This design supports horizontal scaling, as additional channels can be created to accommodate growing numbers of agents and diverse interaction requirements.

3) *Event sourcing and snapshotting:* Event sourcing records all state changes as discrete events, enabling efficient state reconstruction and facilitating scalability. Snapshotting further enhances performance by allowing agents to rebuild their state from recent snapshots rather than processing an extensive event history, reducing computational overhead and speeding up state initialization.

4) *Distributed ledger technology:* The decentralized nature of Hyperledger Fabric allows the system to scale horizontally by adding more nodes to the network. This distribution enhances fault tolerance and load balancing, ensuring that the system remains resilient and performs consistently as it grows.

5) *Optimized data stores for read operations:* By maintaining separate, optimized data stores for read queries, the system ensures that read-heavy operations do not impede write performance. This optimization is crucial for maintaining low latency and high availability, especially as the number of agents and interactions increases.

C. Performance Analysis

1) *High throughput and low latency:* The system is designed to handle a substantial number of transactions per second (TPS) with minimal latency. Hyperledger Fabric's consensus mechanisms and efficient transaction ordering services contribute to maintaining high throughput, while the use of CQRS and optimized data stores ensures rapid data retrieval and processing.

2) *Efficient state management:* Event sourcing, combined with snapshotting, allows for swift state reconstruction and reduces the time required for agents to initialize or recover their state. This efficiency is vital for maintaining real-time responsiveness and ensuring that agents can operate seamlessly, even under heavy load.

3) *Resource optimization:* The segregation of read and write operations, along with asynchronous communication channels, ensures optimal utilization of system resources. By distributing workloads effectively and minimizing bottlenecks, the system maintains consistent performance levels regardless of scaling demands.

4) *Resilience and fault tolerance:* The decentralized architecture and immutable ledger ensure that the system remains operational even in the face of individual node failures or attacks. Redundant data storage and distributed consensus

protocols contribute to the system's ability to recover quickly and maintain performance standards during adverse conditions.

5) *Continuous improvement and adaptability:* The system's architecture allows for continuous optimization and scaling without necessitating significant overhauls. The modular design of smart contracts, combined with secure upgrade paths, enables the integration of performance enhancements and new features, ensuring that the system can evolve in response to changing demands and technological advancements.

D. Trade-offs and Considerations

1) *Complexity vs. security and scalability:* While the system's advanced security measures and scalable architecture provide significant benefits, they also introduce additional complexity. Managing multiple DIDs, implementing ZKP, and maintaining a distributed ledger require sophisticated infrastructure and expertise, potentially increasing the initial setup and maintenance overhead.

2) *Resource consumption:* The use of encryption, secure storage mechanisms, and continuous monitoring can lead to increased resource consumption. Balancing security and performance with resource efficiency is essential to ensure that the system remains cost-effective and sustainable as it scales.

3) *Latency in event processing:* Although event sourcing and snapshotting enhance state management efficiency, there can be inherent latency in processing and recording events on the blockchain. Optimizing transaction throughput and implementing efficient event handling protocols are necessary to minimize delays and maintain real-time responsiveness.

4) *Governance and upgrade management:* Ensuring secure and controlled upgrades to smart contracts and system components is crucial for maintaining system integrity. Establishing robust governance mechanisms and clear procedures for implementing changes helps mitigate the risks associated with upgrades but requires ongoing coordination and management.

E. Conclusion

The system's architecture robustly addresses critical aspects of security, scalability, and performance through the integration of advanced technologies and best practices. By leveraging Hyperledger Fabric's decentralized ledger, employing the CQRS pattern for efficient state management, and implementing comprehensive security measures, the system ensures secure, scalable, and high-performing operations.

While the system presents inherent complexities and resource demands, these are outweighed by the substantial benefits of enhanced security, efficient scalability, and reliable performance. Continuous monitoring, proactive incident response, and adaptable governance frameworks further strengthen the system's resilience and capacity to evolve in response to emerging challenges and growth demands.

Overall, the system exemplifies a well-balanced approach to building a decentralized multi-agent environment that is both secure and capable of handling significant scalability and

performance requirements, making it well-suited for a wide range of decentralized applications and use cases.

VI. USE CASES AND REAL-WORLD APPLICATIONS

In this chapter, we explore practical applications of the proposed architecture across various industries. By demonstrating how the architecture can be applied to real-world scenarios, we highlight its versatility, effectiveness, and potential impact on modern decentralized systems.

A. Smart Grids

Smart grids represent the modernization of traditional electrical grids by integrating advanced communication and control technologies. They enable efficient energy distribution, real-time monitoring, and dynamic management of energy resources. The decentralized nature of smart grids, with numerous energy producers and consumers, makes them an ideal candidate for the proposed architecture.

1) Application of the Architecture

a) Decentralized energy management:

- **Autonomous Agents:** Each energy producer (e.g., solar panels, wind turbines) and consumer (households, businesses) is represented by an autonomous agent.
- **Decentralized Identity Management:** Agents generate Decentralized Identifiers (DIDs), ensuring secure and self-sovereign identities without reliance on centralized authorities.
- **Secure Communication:** Agents communicate securely using mutual TLS and encrypted channels, exchanging data such as energy production, consumption, and pricing information.

b) Energy transactions and trading:

- **Blockchain Integration:** The blockchain serves as an immutable ledger for recording energy transactions, such as energy generation records, consumption data, and peer-to-peer energy trades.
- **Smart Contracts:** Automate energy trading agreements, settlement of payments, and enforcement of contractual obligations between agents.
- **Event Sourcing and CQRS:** Efficiently handle high volumes of transactions, separating command and query responsibilities for optimal performance.

c) Privacy-preserving data sharing:

- **Zero-Knowledge Proofs (ZKP):** Allow agents to prove certain attributes (e.g., energy surplus availability) without revealing sensitive data like exact energy usage patterns.
- **Data Confidentiality:** Encryption ensures that only authorized agents can access specific data, protecting user privacy and complying with regulations.

2) Benefits:

- **Enhanced Efficiency:** Real-time data exchange and autonomous decision-making optimize energy distribution and reduce waste.

- **Increased Reliability:** Decentralized control reduces single points of failure, improving grid resilience.
- **Cost Savings:** Direct peer-to-peer energy trading lowers transaction costs and enables dynamic pricing models.
- **Regulatory Compliance:** Secure and privacy-preserving mechanisms align with data protection laws and industry standards.
- **Challenges and considerations:**
- **Scalability:** Managing a large number of agents and transactions requires robust scalability, addressed by the architecture's design.
- **Interoperability:** Integration with existing grid infrastructure necessitates adherence to industry protocols and standards.
- **Security Threats:** Protecting against cyber-attacks is critical, necessitating ongoing security assessments and updates.

B. Healthcare Data Management

Managing sensitive healthcare data requires stringent security, privacy, and compliance with regulations such as HIPAA. The proposed architecture offers a secure and interoperable framework for handling electronic health records (EHRs), ensuring data integrity and patient privacy.

1) Application of the Architecture:

a) Patient records:

- **Decentralized Identifiers (DIDs):** Each patient is assigned a DID, and their EHRs are stored as events on the blockchain.
- **Event Sourcing:** Records every access and modification to patient data, providing an immutable audit trail.

b) Data access control:

- **Zero-Knowledge Proofs (ZKP):** Patients use ZKPs to grant healthcare providers access to specific parts of their medical records via OAuth 2.0 tokens.
- **OAuth 2.0 Integration:** Manages permissions, allowing patients to control who can view or update their health data.
- **Immutable Logging:** All access and modifications to EHRs are recorded on the blockchain, ensuring accountability and traceability.
- **Smart Contracts:** Enforce data access policies and automate consent management, reducing administrative overhead.

c) Efficient data retrieval:

- **Command Query Responsibility Segregation (CQRS) Pattern:** Enables healthcare providers to query patient data efficiently without impacting the system's write operations.

- **Optimized Read Models:** Ensure rapid access to necessary data, enhancing the responsiveness of healthcare services.

2) Benefits:

- **Enhanced Privacy:** Patients maintain control over their medical data, deciding who can access specific information.
- **Data Security:** Blockchain immutability and robust authentication mechanisms protect against unauthorized access and data breaches.
- **Regulatory Compliance:** Immutable audit trails facilitate compliance with healthcare regulations and standards.

3) Challenges and considerations:

- **Data Interoperability:** Ensuring compatibility with existing healthcare information systems requires adherence to industry standards.
- **Scalability:** Managing large volumes of healthcare data necessitates efficient storage and retrieval mechanisms.
- **User Adoption:** Encouraging healthcare providers and patients to adopt the new system involves overcoming resistance to change and ensuring ease of use.

C. Secure Internet of Things (IoT) Networks

IoT networks consist of numerous interconnected devices that collect and exchange data. Securing these networks is challenging due to the sheer number of devices and the potential for vulnerabilities. The proposed architecture ensures secure device authentication, data integrity, and efficient management of IoT interactions.

1) Application of the architecture:

a) Device authentication:

- **Decentralized Identifiers (DIDs):** Each IoT device is assigned a unique DID, ensuring secure and authenticated interactions within the network.
- **Zero-Knowledge Proofs (ZKP):** Devices use ZKPs to authenticate themselves without revealing sensitive credentials.

b) Data integrity:

- **Blockchain Integration:** Data generated by IoT devices is recorded on the blockchain, ensuring its integrity and preventing tampering.
- **Immutable Logging:** Critical events and data exchanges are stored immutably, providing a reliable audit trail.

c) Scalable communication:

- **Asynchronous Message Queues:** Manage the high volume of data exchange between devices and the blockchain, ensuring efficient processing and minimal latency.
- **Secure Protocols:** Ensure that all communications are encrypted and authenticated, protecting against unauthorized access and data breaches.

d) Access control:

- OAuth 2.0 Tokens: Regulate which devices can access or modify specific data streams, maintaining strict access control policies.
- Capability-Based Access Control: Use cryptographically secure tokens to manage permissions dynamically and securely.

2) Benefits:

- Secure Authentication: Decentralized authentication mechanisms prevent unauthorized devices from joining the network.
- Data Integrity: Immutable records on the blockchain ensure that IoT data remains accurate and trustworthy.
- Scalability: The system efficiently handles large-scale IoT deployments, accommodating a growing number of devices without compromising performance.

3) Challenges and considerations:

- Device Heterogeneity: Managing diverse IoT devices with varying capabilities requires adaptable and flexible security protocols.
- Resource Constraints: Ensuring that security measures are lightweight enough to operate on resource constrained IoT devices.
- Cybersecurity Threats: Protecting IoT networks from sophisticated cyber-attacks necessitates ongoing security enhancements and monitoring.

D. Supply Chain Management

Supply chains involve the movement of goods and services from suppliers to end consumers, encompassing various processes like manufacturing, logistics, and retail. The complexity and need for transparency in supply chains make them suitable for leveraging blockchain and decentralized technologies.

1) Application of the architecture:

a) Transparent tracking and traceability:

- Immutable Record Keeping: Blockchain records every event in the product lifecycle, from raw material sourcing to final delivery, ensuring data integrity.
- Event Sourcing: Each action (e.g., shipment, quality check) is recorded as an event, allowing real-time tracking and historical analysis.
- Decentralized Agents: Manufacturers, suppliers, logistics providers, and retailers operate as agents within the system, communicating securely.

b) Secure and efficient communication:

- Asynchronous Message Queues: Facilitate communication between agents, handling asynchronous updates and ensuring timely information flow.
- Secure Protocols: Mutual authentication and encrypted channels prevent unauthorized access and data breaches.

c) Confidentiality and competitive advantage:

- Zero-Knowledge Proofs (ZKP): Enable agents to verify compliance with standards or certifications without revealing proprietary information.
- Selective Disclosure: Agents can share necessary data with partners or regulators while keeping sensitive business details confidential.

2) Benefits:

- Enhanced Transparency: Consumers and stakeholders can verify the authenticity and origin of products, building trust.
- Fraud Reduction: Immutable records prevent tampering, reducing counterfeit goods and unethical practices.
- Operational Efficiency: Automation and real-time data exchange streamline processes, reducing delays and costs.
- Compliance and Reporting: Simplifies regulatory compliance by providing verifiable records and audit trails.

3) Challenges and considerations:

- Data Standardization: Ensuring consistent data formats and standards across diverse participants is essential for interoperability.
- Adoption Barriers: Convincing all supply chain participants to adopt the new system may require demonstrating clear value propositions.
- Privacy Concerns: Balancing transparency with the need to protect sensitive business information requires careful design.

E. Conclusion

By addressing real-world problems with a secure, scalable, and privacy-focused approach, the proposed architecture paves the way for innovative decentralized applications that can transform industries and enhance the way individuals and organizations interact in the digital age.

VII. CONCLUSION

In this paper, we have presented a comprehensive architecture designed to enhance security, scalability, data integrity, and privacy in Multi-Agent Systems (MAS). By integrating advanced technologies such as Blockchain, Zero-Knowledge Proofs (ZKP), OAuth 2.0, and Decentralized Identity Management (DID), we have addressed the critical challenges inherent in decentralized environments where autonomous agents interact and collaborate.

The architecture's contributions address key security challenges by ensuring data integrity and immutability, with blockchain technology providing an immutable ledger for all transactions, events, and identity information. Cryptographic linkages between blocks and consensus mechanisms like PBFT and DPoS prevent tampering. Privacy-preserving authentication is achieved with Zero-Knowledge Proofs, allowing agents to prove knowledge without disclosing information. Secure

authorization is reinforced by combining ZKP with OAuth 2.0, enabling fine-grained access control, while decentralized identity management leverages DIDs and DPKI to reduce reliance on centralized providers, empowering agents to manage their identities independently.

To enhance scalability and performance, the architecture uses Command Query Responsibility Segregation (CQRS) and Event Sourcing, enabling independent scaling of read and write operations. Asynchronous communication through message queues supports a high volume of non-blocking interactions, while optimized cryptographic operations using zk-STARKs, Bulletproofs, and hardware acceleration improve performance. Layer-2 solutions such as state channels and sidechains boost throughput and reduce latency by offloading transactions from the main blockchain.

Data integrity and privacy are safeguarded with an immutable event store on the blockchain, providing a tamper-proof history for auditing, compliance, and state recovery. Selective disclosure mechanisms and ZKPs empower agents to share only necessary information while preserving anonymity. The architecture aligns with global data protection regulations, such as GDPR, enhancing its compliance and adaptability.

The architecture's versatility allows it to be applied across various domains, including smart grids, supply chains and secure internet of things (IoT) networks, demonstrating its broad applicability. Its modular and interoperable design, adhering to standards like OAuth 2.0 and W3C's DIDs, promotes seamless integration with existing systems.

Despite these advancements, the architecture faces limitations. Cryptographic operations, particularly ZKPs, impose computational demands that can be challenging for devices with limited processing power. Consensus mechanism scalability may be hindered by network growth and latency, while data management requires balancing efficient storage with integrity. The complexity of implementing these advanced technologies may increase the learning curve, and decentralized system adoption may face resistance due to regulatory uncertainties.

Future research will focus on optimizing cryptographic protocols, exploring post-quantum cryptography, and enhancing scalable and energy-efficient consensus mechanisms. Improvements in identity management, integration with AI for intelligent agent behavior and real-time threat detection, and collaboration with regulatory bodies to establish supportive frameworks will further refine the architecture.

REFERENCES

- [1] Applied Sciences. "Research Progress on the Application of Multi-agent Systems." MDPI (2021). MDPI
- [2] P. Rajasekar, K. Kalaiselvi, R. Shanmugam, S. Tamilselvan and A. P. Pandian, "Advancing Cloud Security Frameworks Implementing Distributed Ledger Technology for Robust Data Protection and Decentralized Security Management in Cloud Computing Environments," 2024 Second International Conference on Advances in Information Technology (ICAIT), Chikkamagaluru, Karnataka, India, 2024, pp. 1-6, doi: 10.1109/ICAIT61638.2024.10690718.
- [3] Aldweesh, A., Alauthman, M., & al-Qerem, A. Revolutionizing Cryptography: Blockchain as a Catalyst for Advanced Security Systems. IGI Global, 2024. IGI Global.
- [4] Baptista, Frederico & Dehez Clementi, Marina & Detchart, Jonathan. (2024). DFLy: A Publicly Auditable and Privacy-Preserving UAS Traffic Management System on Blockchain. *Drones*. 8. 10.3390/drones8080410.
- [5] Kalbantner, Jan & Markantonakis, Konstantinos & Hurley-Smith, Daren & Shepherd, Carlton. (2024). ZKP Enabled Identity and Reputation Verification in P2P Marketplaces. 591-598. 10.1109/Blockchain62396.2024.00087.
- [6] Ballesteros-Rodríguez, Alberto & Sánchez-Alonso, Salvador & Sicilia-Urbán, Miguel-Ángel. (2024). Enhancing Privacy and Integrity in Computing Services Provisioning Using Blockchain and Zk-SNARKs. *IEEE Access*. PP. 1-1. 10.1109/ACCESS.2024.3447785.
- [7] Xiong, S., Chen, P., Ge, S., & Ni, Q. SFOM-DT: A Secure and Fair One-to-Many Data Trading Scheme Based on Blockchain. *IEEE Xplore*, 2024. *IEEE Xplore*.
- [8] Lavin, Ryan & Liu, Xuekai & Mohanty, Hardhik & Norman, Logan & Zaarour, Giovanni & Krishnamachari, Bhaskar. (2024). A Survey on the Applications of Zero-Knowledge Proofs. 10.48550/arXiv.2408.00243.
- [9] Ma, Shengchen & Zhang, Xing. (2024). Integrating blockchain and ZK-ROLLUP for efficient healthcare data privacy protection system via IPFS. *Scientific Reports*. 14. 11746. 10.1038/s41598-024-62292-9.
- [10] Harz, D., & Boman, M. (2018). The Scalability of Trustless Trust. *ArXiv*, abs/1801.09535. https://doi.org/10.1007/978-3-662-58820-8_19.
- [11] Lad, K., Dewan, M., & Lin, F. (2020). Trust Management for Multi-Agent Systems Using Smart Contracts. 2020 IEEE Intl Conf on Dependable, Autonomic and Secure Computing, 414-419. <https://doi.org/10.1109/DASC-PICom-CBDCCom-CyberSciTech49142.2020.00080>.
- [12] Yang, T., Liu, Y., Yang, X., & Kang, Y. (2019). A Blockchain based Smart Agent System Architecture. *Proceedings of the 4th International Conference on Crowd Science and Engineering*. <https://doi.org/10.1145/3371238.3371244>.
- [13] Leng, J., Sha, W., Lin, Z., Jing, J., Liu, Q., & Chen, X. (2022). Blockchain smart contract pyramid-driven multi-agent autonomous process control for resilient individualised manufacturing towards Industry 5.0. *International Journal of Production Research*, 61, 4302 - 4321. <https://doi.org/10.1080/00207543.2022.2089929>.
- [14] Wei, P., Wang, D., Zhao, Y., Tyagi, S., & Kumar, N. (2020). Blockchain data-based cloud data integrity protection mechanism. *Future Gener. Comput. Syst.*, 102, 902-911. <https://doi.org/10.1016/j.future.2019.09.028>.
- [15] Zhang, Q., Zhang, Z., Cui, J., Zhong, H., Li, Y., Gu, C., & He, D. (2023). Efficient Blockchain-Based Data Integrity Auditing for Multi-Copy in Decentralized Storage. *IEEE Transactions on Parallel and Distributed Systems*, 34, 3162-3173. <https://doi.org/10.1109/TPDS.2023.3323155>.
- [16] Oliveira, G., Silva, J., Cortes, O., & Coutinho, L. (2022). A Multi-Agent Architecture for Distributed Data Mining Systems. *Anais do XVI Brazilian e-Science Workshop (BRESCI 2022)*. <https://doi.org/10.5753/bresci.2022.222487>.
- [17] Qasem, M., Obeid, N., Hudaib, A., Almaiah, M., Al-Zahrani, A., & Al-Khasawneh, A. (2021). Multi-Agent System Combined With Distributed Data Mining for Mutual Collaboration Classification. *IEEE Access*, 9, 70531-70547. <https://doi.org/10.1109/ACCESS.2021.3074125>.
- [18] Nait Cherif, A., Achir, Y., Youssfi, M., Elgarej, M., & Bouattane, O. (2023). Ensuring security and data integrity in Multi Micro-Agent System Middleware with Blockchain Technology. 2023 3rd International Conference on Innovative Research in Applied Science, Engineering and Technology (IRASET), 1-6. <https://doi.org/10.1109/IRASET57153.2023.10152950>.
- [19] Ge, X., Han, Q., Ding, D., Zhang, X., & Ning, B. (2018). A survey on recent advances in distributed sampled-data cooperative control of multi-agent systems. *Neurocomputing*, 275, 1684-1701. <https://doi.org/10.1016/j.neucom.2017.10.008>.
- [20] Long, Z. (2017). Improvement and Implementation of a High Performance CQRS Architecture. 2017 International Conference on Robots & Intelligent System (ICRIS), 170-173. <https://doi.org/10.1109/ICRIS.2017.49>.
- [21] Debski, A., Szczepanik, B., Malawski, M., Spahr, S., & Muthig, D. (2018). A Scalable, Reactive Architecture for Cloud Applications. *IEEE Software*, 35, 62-71. <https://doi.org/10.1109/MS.2017.265095722>.

- [22] Mansky, W., Appel, A., & Nogin, A. (2017). A verified messaging system. *Proceedings of the ACM on Programming Languages*, 1, 1 - 28. <https://doi.org/10.1145/3133911>.
- [23] Shao, J., Yin, B., Chen, B., Wang, G., Yang, L., Yan, J., Wang, J., & Liu, W. (2016). Read Consistency in Distributed Database Based on DMVCC. 2016 IEEE 23rd International Conference on High Performance Computing (HiPC), 142-151. <https://doi.org/10.1109/HiPC.2016.025>.
- [24] Dashti, Z., Oliva, G., Seatzu, C., Gasparri, A., & Franceschelli, M. (2022). Distributed Mode Computation in Open Multi-Agent Systems. *IEEE Control Systems Letters*, 6, 3481-3486. <https://doi.org/10.1109/LCSYS.2022.3185419>.
- [25] Breugnot, P., Hermann, B., Lang, C., & Philippe, L. (2021). A Synchronized and Dynamic Distributed Graph structure to allow the native distribution of Multi-Agent System simulations. 2021 29th Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP), 54-61. <https://doi.org/10.1109/PDP52278.2021.00017>.
- [26] Wang, J., & Elia, N. (2010). Control approach to distributed optimization. 2010 48th Annual Allerton Conference on Communication, Control, and Computing (Allerton), 557-561. <https://doi.org/10.1109/ALLERTON.2010.5706956>.
- [27] Rust, P., Picard, G., & Ramparany, F. (2020). Resilient Distributed Constraint Optimization in Physical Multi-Agent Systems. , 195-202. <https://doi.org/10.3233/FAIA200093>.
- [28] Fanitabasi, F. (2018). A Review of Adversarial Behaviour in Distributed Multi-Agent Optimisation. 2018 IEEE/ACM International Conference on Utility and Cloud Computing Companion (UCC Companion), 53-58. <https://doi.org/10.1109/UCC-Companion.2018.00034>.
- [29] Fu, F., & Zhou, H. (2021). A combined multi-agent system for distributed multi-project scheduling problems. *Appl. Soft Comput.*, 107, 107402. <https://doi.org/10.1016/J.ASOC.2021.107402>.
- [30] Costa, D., Carrido, D., & Silva, D. (2021). Efficient Secure Communication for Distributed Multi-Agent Systems. , 543-552. <https://doi.org/10.5220/0010375605430552>.
- [31] Barbosa, J., Leitão, P., Ferreira, A., Queiroz, J., Galdes, C., & Coelho, J. (2018). Implementation of a Multi-Agent System to Support ZDM Strategies in Multi-Stage Environments. 2018 IEEE 16th International Conference on Industrial Informatics (INDIN), 822-827. <https://doi.org/10.1109/INDIN.2018.8471948>.
- [32] Dimarogonas, D., Frazzoli, E., & Johansson, K. (2012). Distributed Event-Triggered Control for Multi-Agent Systems. *IEEE Transactions on Automatic Control*, 57, 1291-1297. <https://doi.org/10.1109/TAC.2011.2174666>.
- [33] Sahingoz, O., & Sonmez, A. (2006). Fault Tolerance Mechanism of Agent-Based Distributed Event System. , 192-199. https://doi.org/10.1007/11758532_27.
- [34] Cristea, V., Pop, F., Dobre, C., & Costan, A. (2011). Distributed Architectures for Event-Based Systems. , 11-45. https://doi.org/10.1007/978-3-642-19724-6_2.
- [35] Long, Z. (2017). Improvement and Implementation of a High Performance CQRS Architecture. 2017 International Conference on Robots & Intelligent System (ICRIS), 170-173. <https://doi.org/10.1109/ICRIS.2017.49>.
- [36] Navabi, S., & Nayyar, A. (2020). A Dynamic Mechanism for Security Management in Multi-Agent Networked Systems. *IEEE INFOCOM 2020 - IEEE Conference on Computer Communications*, 1628-1637. <https://doi.org/10.1109/INFOCOM41043.2020.9155295>.
- [37] Lu, A., & Yang, G. (2020). Secure State Estimation for Multiagent Systems With Faulty and Malicious Agents. *IEEE Transactions on Automatic Control*, 65, 3471-3485. <https://doi.org/10.1109/TAC.2019.2945032>.
- [38] Zheng, Z., Xie, S., Dai, H., Chen, X., & Wang, H. (2017). An Overview of Blockchain Technology: Architecture, Consensus, and Future Trends. 2017 IEEE International Congress on Big Data (BigData Congress), 557-564. <https://doi.org/10.1109/BIGDATAACONGRESS.2017.85>.
- [39] Yang, Y., Jin, K., Liang, W., Liu, Y., Li, Y., & Hosam, O. (2023). A Review of Blockchain-based Privacy Computing Research. 2023 IEEE 10th International Conference on Cyber Security and Cloud Computing (CSCloud)/2023 IEEE 9th International Conference on Edge Computing and Scalable Cloud (EdgeCom), 241-246. <https://doi.org/10.1109/CSCloud-EdgeCom58631.2023.00049>.
- [40] Li, W., Guo, H., Nejad, M., & Shen, C. (2020). Privacy-Preserving Traffic Management: A Blockchain and Zero-Knowledge Proof Inspired Approach. *IEEE Access*, 8, 181733-181743. <https://doi.org/10.1109/ACCESS.2020.3028189>.
- [41] Sun, X., Yu, F., Zhang, P., Sun, Z., Xie, W., & Peng, X. (2021). A Survey on Zero-Knowledge Proof in Blockchain. *IEEE Network*, 35, 198-205. <https://doi.org/10.1109/MNET.011.2000473>.
- [42] Wang, H., & Liao, J. (2021). Blockchain Privacy Protection Algorithm Based on Pedersen Commitment and Zero-knowledge Proof. *Proceedings of the 2021 4th International Conference on Blockchain Technology and Applications*. <https://doi.org/10.1145/3510487.3510488>.
- [43] Goldwasser, S., Micali, S., & Rackoff, C. (1989). The Knowledge Complexity of Interactive Proof Systems. , 203-225. <https://doi.org/10.1137/0218012>.
- [44] Yang, X., & Li, W. (2020). A zero-knowledge-proof-based digital identity management scheme in blockchain. *Comput. Secur.*, 99, 102050. <https://doi.org/10.1016/J.COSE.2020.102050>.
- [45] Fotiou, N., Siris, V., & Polyzos, G. (2021). Capability-based access control for multi-tenant systems using OAuth 2.0 and Verifiable Credentials. 2021 International Conference on Computer Communications and Networks (ICCCN), 1-9. <https://doi.org/10.1109/ICCCN52240.2021.9522214>.
- [46] Ra, G., Kim, T., & Lee, I. (2021). VAIM: Verifiable Anonymous Identity Management for Human-Centric Security and Privacy in the Internet of Things. *IEEE Access*, 9, 75945-75960. <https://doi.org/10.1109/ACCESS.2021.3080329>.