scientific reports

OPEN

Semantic decentralized authentication for IoT-based e-learning using Hedera Hashgraph and Knowledge Graphs

Lixia Luo, Yusha Zhang & Qitao Tang

The integration of Internet of Things (IoT) devices in e-learning systems necessitates robust, scalable, and secure authentication procedures to provide dependable sharing of academic records among remote educational institutions. Conventional centralized systems experience scalability limitations, singular points of failure, and heightened susceptibility to hackers, especially in resource-limited IoT settings. This paper presents a decentralized authentication framework based on Hedera Hashgraph and Knowledge Graphs (KGs) to tackle these issues. The architecture incorporates a GAN-based cryptography module for the generation of dynamic symmetric keys, enhancing resistance against predictive and inference-based attacks. Knowledge Graphs facilitate semantic validation of identification features and improve interoperability among institutions via the Hedera Consensus Service (HCS). The quantitative assessment indicates that the Hedera + KG + GAN model attains a 17.1% increase in throughput, an 11–12% reduction in processing time, up to a 20% decrease in execution time for substantial data volumes, a 6–15% decline in energy consumption, and an approximate 23% reduction in authentication delay during periods of high network utilization relative to the leading competing frameworks. The suggested method provides a scalable, safe, and semantically enriched authentication mechanism for IoT-enabled e-learning ecosystems, creating a solid foundation for next-generation decentralized educational platforms.

Keywords Blockchain, IoT, E-learning, Knowledge graphs, Hedera hashgraph, Decentralized authentication, Semantic analysis

The incorporation of IoT technologies with e-learning systems has become a revolutionary model in the modern digital age, reshaping the educational environment. The expansion of networked smart devices has enabled the establishment of intelligent learning environments, allowing real-time access to educational resources, dynamic interactions, and customized learning experiences[1]. The increasing demand for distributed educational platforms propels the integration of IoT in education[2].

This technological breakthrough presents substantial technical and security challenges, as traditional authentication methods, typically dependent on centralized architectures, are insufficient for the scale and complexity of IoT-driven ecosystems[3]. Distributed ledger technologies, including Hedera Hashgraph, provide effective solutions via decentralized consensus and fault-tolerant processes, guaranteeing scalability and security[4]. Simultaneously, Knowledge Graphs (KGs) improve semantic interoperability by organizing educational data with relational metadata, an essential characteristic for resource-limited IoT settings[5].

The importance of IoT integration in e-learning transcends simple technology implementation; it cultivates an interactive, data-driven, and inclusive educational environment. IoT-enabled gadgets, including wearable sensors and smart tablets, provide real-time monitoring of learning progress, performance evaluation, and immersive experiences via augmented reality[6]. In higher education, these technologies enable global collaboration, permitting students to securely exchange academic credentials among schools[7]. From a societal standpoint, the IoT has the capacity to mitigate educational inequalities in underprivileged areas, where access to conventional classrooms is restricted, thus fostering educational equity[8]. However, the lack of strong authentication systems presents considerable concerns, including eroding trust in digital educational platforms due to susceptibility to cyber threats[9].

School of Computer Science and Engineering, Hunan University of Information Technology, Changsha, Hunan 410148, China. email: zxcv2891@163.com

Notwithstanding its transformative promise, the implementation of IoT in e-learning encounters significant hurdles and constraints. A key issue is the scalability of authentication systems; centralized architectures, dependent on individual servers, are vulnerable to single points of failure, which become more evident as the user and device count rises[10]. Moreover, IoT devices are limited by restricted computational power, energy resources, and storage capacity, making the implementation of modern cryptographic algorithms difficult[11].

Security concerns, including man-in-the-middle attacks, data eavesdropping, and privacy breaches, are especially significant in educational environments where sensitive information, such as academic records, is maintained[12]. Additional challenges stem from the diversity of IoT communication protocols and the absence of semantic interoperability among institutions, which obstructs smooth credential transfer[13]. Recent studies indicate that IoT applications in education often encounter security problems, emphasizing the necessity for new solutions[14].

This research is motivated by the pressing necessity to tackle these difficulties. Distributed solutions like Hedera Hashgraph, utilizing the Gossip-about-Gossip consensus algorithm for enhanced throughput and minimal latency, offer a feasible foundation for decentralized authentication in IoT settings[15]. Moreover, Knowledge Graphs provide the semantic modeling of intricate interactions among learners, courses, and certificates, hence improving robustness against adversarial attacks[16]. The incorporation of machine learning-driven cryptography, namely GANs, which produce dynamic and attack-resistant cryptographic keys, enhances this methodology[17]. The increasing necessity for secure and scalable authentication frameworks in IoT-based educational environments justifies this research[18].

The main aim of this project is to provide a decentralized semantic authentication system for IoT-based e-learning, utilizing Hedera Hashgraph and Knowledge Graphs. Specific objectives encompass the formulation of protocols for user registration, authentication, and the semantic transfer of credentials; the evaluation of performance metrics including throughput, response time, and energy consumption; and the assessment of security resilience against sophisticated cyber threats. The principal innovation is the amalgamation of GAN-based encryption with Hedera's asynchronous Byzantine Fault Tolerance (aBFT) consensus and KG-driven semantic analysis, resulting in a harmonious equilibrium of security and efficiency.

This research advances the subject by proposing a distributed educational architecture in which institutions function as nodes within a Hedera network, enhanced with Knowledge Graphs to guarantee semantic interoperability and reduce single points of failure. It formulates registration and authentication algorithms employing GANs to produce attack-resistant cryptographic keys, in conjunction with a secure credential transfer protocol that obviates the necessity for repeated authentication. Comprehensive simulations reveal substantial enhancements in throughput and response time relative to conventional blockchain and centralized systems, alongside a rigorous threat model evaluating resilience against sophisticated attackers. This study ultimately provides a scalable solution for IoT-based e-learning environments, improving educational fairness and optimizing energy usage for resource-limited devices.

In comparison to the most robust baseline schemes, the suggested Hedera + KG + GAN architecture attains a throughput increase of 17.1% (4310 TPS), a reduction in processing time by 11–12%, and a decrease in execution time by up to 20% for medium and large data volumes. IoT devices exhibit a 6–15% reduction in energy consumption, while the authentication workflow achieves an approximately 23% decrease in elapsed time during periods of heavy network demand. The numerical indications demonstrate that the suggested architecture continuously and significantly surpasses the top-performing existing frameworks, affirming its scalability and practical usefulness in extensive IoT-based e-learning environments.

This study is structured to ensure a logical transition from theoretical underpinnings to empirical validation. Section 2 examines the literature, focusing on pivotal technologies including Hedera Hashgraph, Knowledge Graphs, and GAN-based cryptography. Section 3 outlines the suggested paradigm, specifying the procedures for registration, authentication, and credential transfer with algorithmic accuracy. Section 4 examines blockchain and cryptography functions, detailing the fundamental algorithms that support the architecture. Section 5 assesses the efficacy and security of the proposed system via comprehensive simulation-based tests. Ultimately, Sect. 6 closes the study and delineates avenues for future research.

Related work

Decentralized authentication for IoT-based e-learning systems has garnered heightened interest owing to the constraints of centralized methodologies. This section categorizes existing research into three theme areas and highlights the research gaps that necessitate the suggested paradigm.

Blockchain-based authentication in e-learning

Numerous studies have utilized blockchain to attain decentralized authentication in educational settings. A blockchain-based solution was proposed in[19] to address single points of failure in centralized systems, albeit it resulted in elevated computing costs for IoT devices. The study in[20] introduced a lightweight authentication system designed for IoT contexts, enhancing response time yet exhibiting deficiencies in interoperability among institutions.

A Hyperledger Fabric-based access control system was developed in[21] to augment the secrecy of academic credentials via selective access permissions. While proficient at protecting credentials, its consortium structure restricts scalability and the sharing of semantic data among universities.

Blockchain-based solutions enhance data integrity and privacy; yet, they encounter issues related to latency, energy consumption, and interoperability. These constraints warrant the investigation of other distributed ledgers, such as Hedera Hashgraph, which offer enhanced throughput and efficient consensus for IoT-based e-learning.

Semantic interoperability via knowledge graphs

A separate research avenue emphasizes the attainment of semantic integration within decentralized educational systems. The research in[22] employed knowledge graphs to represent relationships among IoT entities and improve semantic reasoning. Likewise[23], utilized blockchain and deep learning to enhance the security of IoT applications, although failed to tackle cross-platform semantic interoperability.

KGs enhance contextual comprehension and promote significant interactions among decentralized educational entities. Nonetheless, previous studies have generally regarded knowledge graphs as analytical instruments apart from the foundational ledger. This study enhances the concept by integrating knowledge graph reasoning directly into the consensus mechanism of Hedera Hashgraph, thereby providing semantic validation of educational records among institutions.

Machine learning and cryptographic enhancements for IoT security

Machine learning methodologies have been progressively utilized to enhance IoT security. The research in[24] integrated federated learning with generative models for intrusion detection, demonstrating robustness against attacks while missing decentralized identity management. The study in[25] utilized GAN-based entropy generation to create context-aware cryptographic keys, improving flexibility in remote environments.

Reference[26] presents an IOTA-based identity management architecture that incorporates Fog Computing and machine learning for ongoing device authentication, providing scalability and diminished overhead. Moreover[27], illustrated the viability of Hedera Hashgraph in decentralized healthcare systems by incorporating deep learning analytics with distributed storage to guarantee data immutability and elevated performance.

Although these studies affirm the advantages of AI and distributed ledgers in IoT security, few have integrated these technologies with semantic reasoning and decentralized consensus, which constitutes the principal contribution of this research.

Current literature can be classified into three categories:

- blockchain-based authentication emphasizing data integrity yet limited by scalability;
- semantic modeling prioritizing knowledge representation yet lacking integration with decentralized ledgers;
- Cryptographic techniques powered by machine learning that enhance adaptability yet neglect cross-institutional compatibility.

This study presents a semantic decentralized authentication architecture that integrates Hedera Hashgraph, Knowledge Graphs, and GAN-based dynamic cryptography to address these deficiencies. This integration establishes a harmonious equilibrium of scalability, semantic interoperability, and strong security, enhancing decentralized e-learning settings.

Methodology

This chapter delineates a decentralized educational framework using Hedera Hashgraph, IoT, and Knowledge Graphs to guarantee secure and scalable management of academic credentials and transactions among distributed schools.

Decentralized semantic educational model

The educational ecosystem serves as a vital framework for examining smart learning environments facilitated by IoT technology. In this framework, several stakeholders—comprising learners, instructors, administrative workers, and support staff—interact within a decentralized network, facilitated by a system of HCS nodes.

Conventional authentication methods, based on centralized systems, face considerable difficulties in incorporating external users and demonstrate intrinsic susceptibility to single points of failure. The suggested architecture utilizes Hedera Hashgraph, tailored for IoT devices with restricted compute power, storage limitations, and energy resources. This setup allows educational institutions to authenticate and safeguard academic credentials and learning records as students engage in courses across affiliated campuses or instructors provide instruction at various places.

In this architectural concept, each institution operates as a node inside a cohesive educational network, containing extensive data about students, educators, personnel, and enhanced learning profiles supported by KGs. Each node employs an HCS unit to govern its community and guarantees high-throughput, secure connections with other network participants.

The design, supported by a distributed network of HCS nodes, guarantees decentralized consensus and robust storage, with data duplicated across numerous nodes to prevent single-point failures, as seen in Fig. 1. This architecture differentiates itself from independent blockchain instances for each institution by centralizing coordination while preserving distributed integrity. This configuration facilitates the transition of students across institutions—such as from one campus to another (e.g., U-1 to U-2)—with the procedure regulated by corresponding HCS nodes.

As a result, authenticated individuals traverse the network with a singular identity, eliminating the need for redundant device authentication. A device authenticated by an HCS node within a single institution acquires network-wide trust, facilitating uninterrupted connection with all peers.

Moreover, learner-generated data, encompassing assessment results, laboratory activities, and project contributions, is sent openly through IoT devices, augmented by the semantic interoperability facilitated by knowledge graphs.

The decentralized Knowledge Graph (KG) functions as the semantic basis of the proposed architecture. The schema adheres to an RDF(S)/OWL-based ontology, with four principal entity classes—User, Device, Course,

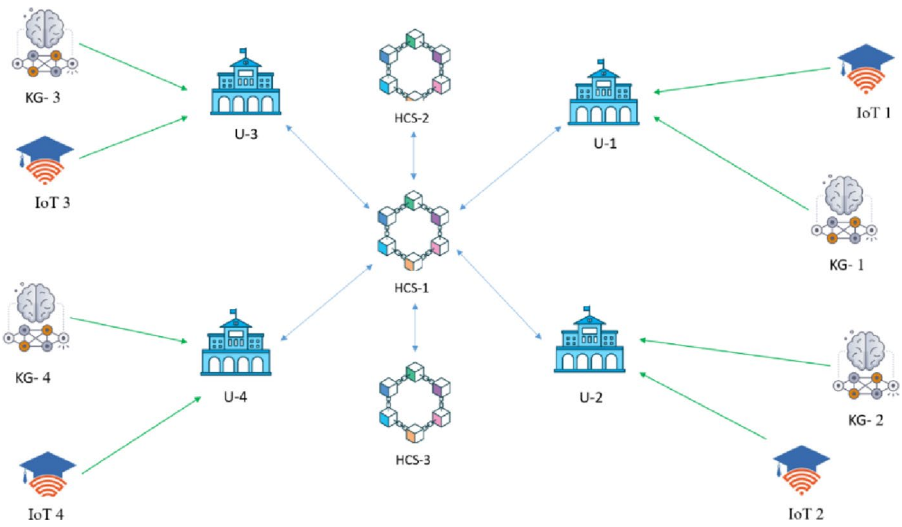

Fig. 1. Architecture of the Decentralized Semantic Educational Network.

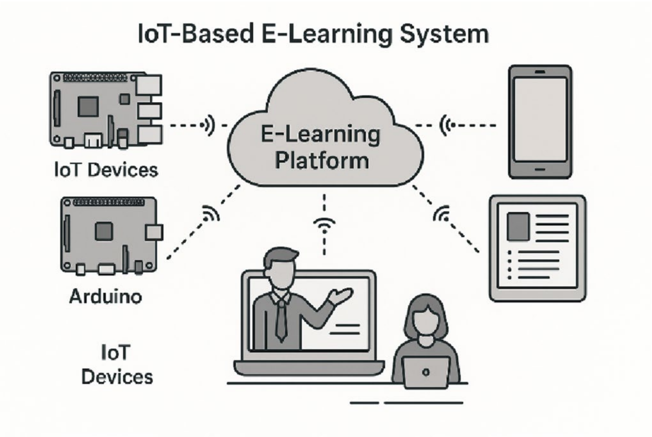

Fig. 2. IoT-Based E-Learning Deployment Architecture.

and Record—interconnected by semantic connections like enrolledIn, teaches, possesses Record, and tied To Device.

Each Management Block (MB) sustains a localized Knowledge Graph (KG) instance sourced from academic sources, IoT device metadata, and registration records. These instances are synchronized by Hashgraph consensus checkpoints, which facilitate the exchange of RDF triple updates and verify integrity among MB nodes, so providing eventual consistency across the network without a centralized authority. SPARQL 1.1 is utilized for semantic validation, identity reasoning, and access control queries, facilitating cross-institutional interoperability and secure knowledge sharing across remote campuses.

To ensure scalability and reduce reasoning latency, RDF triples in Knowledge Graphs are stored and queried using efficient triple-store frameworks like Apache Jena TDB and Blazegraph. These systems employ SPARQL indexing and in-memory reasoning to expedite semantic queries, guaranteeing rapid and reliable access in extensive distributed contexts. Furthermore, semantic information are intermittently stored at the Management Block (MB) level to mitigate cross-node inference overhead while maintaining data integrity and synchronization throughout the educational network.

Figure 2 depicts the deployment architecture of the proposed IoT-based e-learning ecosystem. The environment comprises several IoT devices, including smart classroom sensors (temperature, occupancy, and motion), wearable learning trackers, camera-equipped instructional kiosks, Raspberry Pi-based learning hubs, and student tablets.

Each device interacts with a local Management Block (MB) via lightweight protocols like MQTT and CoAP. The MB conducts semantic preprocessing of the gathered data and transmits authenticated transactions via secure channels to the HCS.

This communication framework facilitates rapid authentication, secure data transmission, and effective synchronization among educational institutions.

Knowledge graph synchronization protocol

The continual development and alignment of Knowledge Graph (KG) instances throughout various educational institutions are essential for maintaining system reliability and data integrity. The HCS oversees this entire process, primarily based on Hashgraph distributed ledger technology.

The synchronization mechanism is implemented using an authoritative, deterministic method. Initially, RDF triple updates produced at each Management Block (MB) node are structured as Compact Semantic Delta Messages. These messages encapsulate fundamental alterations inside the KG. Secondly, each MB node, the exclusive authorized entity, disseminates these semantic deltas as an event across the Hashgraph consensus network through HCS. The Hashgraph algorithm guarantees Deterministic Ordering via its aBFT (asynchronous Byzantine Fault Tolerance) virtual voting mechanism, which produces a universally accepted, definitive sequence for all update occurrences.

As a result, all MB nodes receive and implement the identical ordered sequence of updates, ensuring that the KG instance develops uniformly throughout all participating institutions. This robust, predictable ordering effectively obviates the need for reconciliation or conflict resolution processes. Potential conflicts are inherently handled at the consensus layer prior to application, as each node exclusively processes updates that have the final, agreed-upon sequence and timestamp.

Hedera hashgraph consensus architecture

The decentralized consensus mechanism supporting the educational network is enabled by the Hedera Hashgraph framework, setting it apart from conventional blockchain architectures. This method employs a Directed Acyclic Graph (DAG) framework, in which transactions appear as events interconnected via a gossip protocol, in contrast to the linear block chaining typical of traditional systems. The framework includes a header with timestamps and cryptographic signatures, along with a body that contains transaction lists enhanced by semantic metadata sourced from KGs, as depicted in Fig. 3.

Hedera Hashgraph utilizes an aBFT consensus method, engineered to facilitate transaction throughput above 10,000 transactions per second (TPS) while ensuring lower latency relative to conventional blockchain systems. This architecture avoids the energy-intensive proof-of-work requirements, potentially reducing resource consumption. The gossip-about-gossip protocol guarantees fair event sequencing among network nodes, enhancing scalability and dependability in resource-limited IoT settings. Figure 3 illustrates that the DAG design facilitates the efficient distribution of learning-related transactions, including student records and assessment data, with immutable storage ensured by the network's consensus mechanism.

Consensus security under limited node diversity

At the outset of the e-learning network's deployment, the quantity of participating institutions (Management Blocks) may be restricted, thus heightening the risk of majority-control or Sybil assaults.

To address this, the suggested architecture employs a hybrid consensus technique wherein initial operations depend on a trusted quorum configuration validated by institutional credentials from partner universities.

Every participating member bank is authenticated using a cross-signed certificate before to entering the Hedera network, and a two-thirds (⅔) quorum requirement is necessary for transaction validation.

Voting rights are periodically redistributed among members to avert the dominance of a specific group of nodes, and integrity is maintained through the cross-validation of event signatures.

Upon achieving adequate node diversity, the system effortlessly shifts to the completely decentralized asynchronous Byzantine Fault Tolerant (aBFT) consensus of Hedera Hashgraph, preserving identical security assurances devoid of centralized supervision.

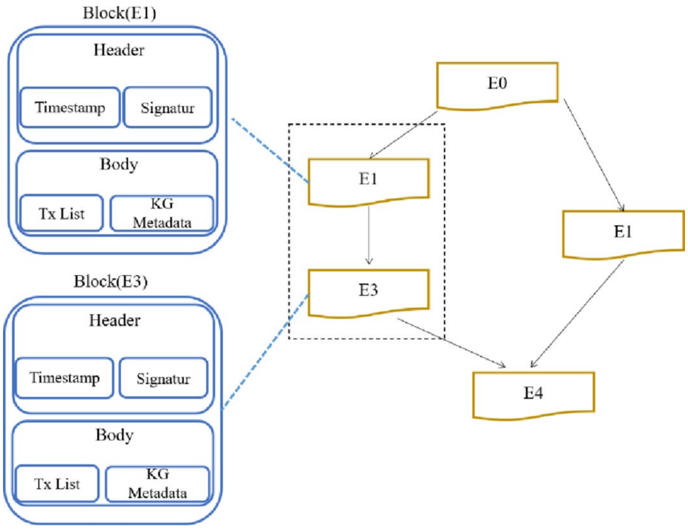

Fig. 3. Hedera Hashgraph DAG Structure with Semantic Metadata.

Semantic user registration protocol

The registration step creates a fundamental system for integrating users—students, professors, and staff—into the decentralized semantic educational network, utilizing the strong infrastructure of the Hedera Hashgraph network. This protocol, carefully implemented under the supervision of the university acting as a Management Block, effectively incorporates GANs for the creation of robust cryptographic key pairs and KGs for thorough semantic identity representation.

The registration procedure begins when a user submits a request via an IoT device, sending preliminary identity information straight to the university (MB), which acts as the central coordinator in this ecosystem.

The MB, possessing substantial computational capability, utilizes a pre-trained GAN model that has been offline trained on extensive historical educational datasets. Section 4 summarizes the particular training settings and hyperparameters of the GAN model to ensure reproducibility and clarity.

The training employs a minimax loss function, as specified in Eq. 1.

$$L = E\left[log\, D(real\_keys)\right] + E\left[log(1 - D(G(fake\_keys)))\right] \tag{1}$$

The offline training phase of the GAN model requires about 100 epochs to attain steady convergence and produce a reliable key-generation function. During the registration process, the pre-trained model is utilized solely for efficient inference to generate new public (Pki) and private (Ptki) key pairs, hence minimizing computational burden on IoT devices.

These keys are augmented with semantic identifiers (IDi), carefully formulated as RDF triples (e.g., User_i: enrolledInCourse_j.:hasGradeGrade_m), which improve interoperability and data consistency among various educational institutions.

Furthermore, symmetric keys (Ski) are derived via a secure SHA-256 hashing procedure that integrates the semantic identifier (IDi), GAN-generated entropy, and the distinct MAC address of the IoT device, thereby establishing a lightweight encryption framework ideally tailored for the resource-limited contexts characteristic of IoT devices.

Upon generation, these authenticated transactions are transmitted by the MB to the HCS for aBFT consensus, as detailed in Algorithm 1, thereby guaranteeing the secure addition of data to the immutable DAG structure for enduring storage and integrity. Subsequently, the MB conveys the encrypted key sets and IDs to the user's IoT device via a highly secure encrypted channel, enabling seamless identity transfer throughout the network without the necessity for repetitive re-authentication. This advanced method, illustrated in Fig. 4, enhances security, scalability, and operational efficiency in remote learning environments, establishing a robust framework for educational data management.

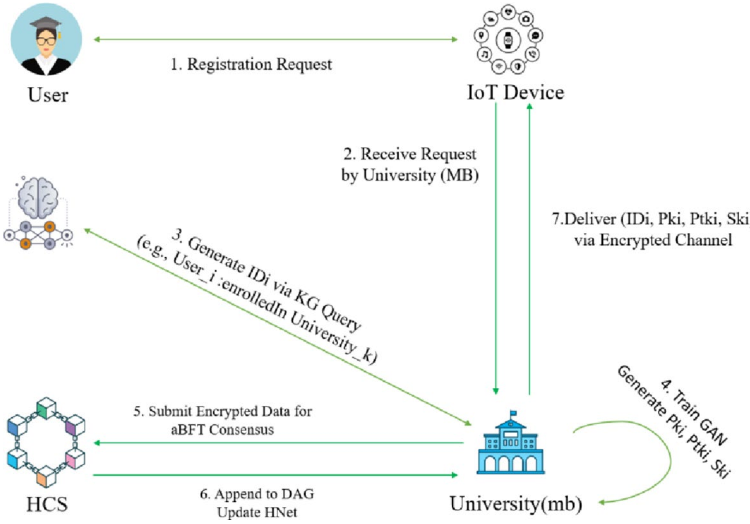
User
1. Registration Request
IoT Device
2. Receive Request by University (MB)
7.Deliver (IDi, Pki, Ptki, Ski) via Encrypted Channel
3. Generate IDi via KG Query (e.g., User_i :enrolledIn University_k)
4. Train GAN Generate Pki, Ptki, Ski
5. Submit Encrypted Data for aBFT Consensus
6. Append to DAG Update HNet
HCS
University(mb)

Fig. 4. Flowchart of Semantic User Registration Protocol.

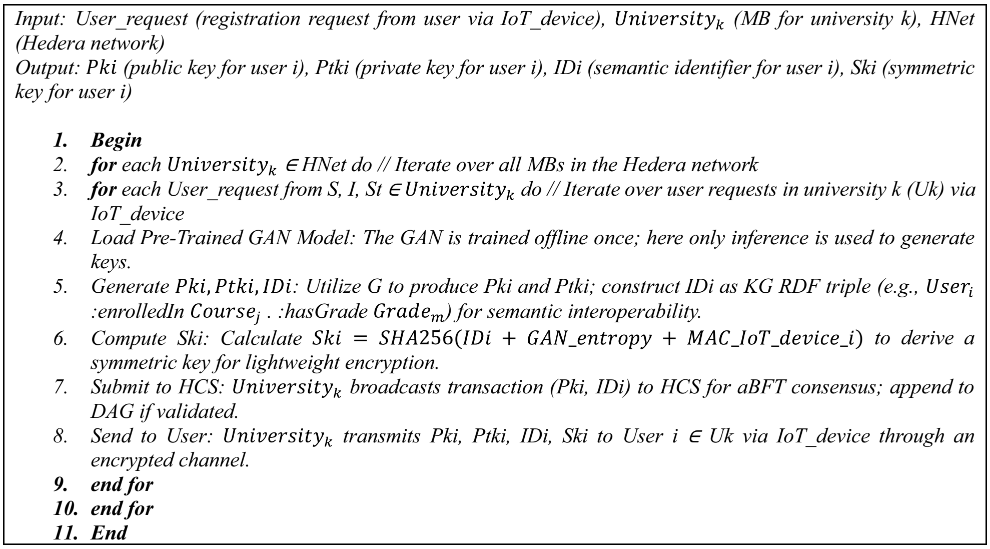

```
Input: User_request (registration request from user via IoT_device), $University_k$ (MB for university k), HNet (Hedera network)
Output: $Pki$ (public key for user i), Ptki (private key for user i), IDi (semantic identifier for user i), Ski (symmetric key for user i)

Begin
    for each $University_k \in$ HNet do // Iterate over all MBs in the Hedera network
        for each User_request from S, I, St $\in University_k$ do // Iterate over user requests in university k (Uk) via IoT_device
            Load Pre-Trained GAN Model: The GAN is trained offline once; here only inference is used to generate keys.
            Generate $Pki, Ptki, IDi$: Utilize G to produce Pki and Ptki; construct IDi as KG RDF triple (e.g., $User_i$ :enrolledIn $Course_j$ . :hasGrade $Grade_m$) for semantic interoperability.
            Compute Ski: Calculate $Ski = SHA256(IDi + GAN\_entropy + MAC\_IoT\_device\_i)$ to derive a symmetric key for lightweight encryption.
            Submit to HCS: $University_k$ broadcasts transaction (Pki, IDi) to HCS for aBFT consensus; append to DAG if validated.
            Send to User: $University_k$ transmits Pki, Ptki, IDi, Ski to User i $\in$ Uk via IoT_device through an encrypted channel.
        end for
    end for
End
```

Algorithm 1. Semantic User Registration Algorithm.

The GAN model utilized for cryptographic key generation is trained exclusively offline before system deployment. The training dataset comprises anonymized educational identity tuples, encompassing semantic user roles, course-enrollment linkages, device identifiers, temporal authentication patterns, and knowledge graph-based contextual attributes. This dataset is utilized exclusively at the administrative root node to generate a singular global GAN model. Subsequent to training, only the generator's refined inference weights are securely disseminated to all Management Blocks (MBs) through authenticated update channels, with no MB engaging in independent training. During operation, MBs utilize the generator only for inference, hence maintaining a distinct separation between offline training and online execution. To reduce reverse-engineering threats, noise based on differential privacy is utilized during training, and the raw training data is never disclosed beyond the root node.

Semantic authentication protocol

Algorithm 2, "Semantic Device and Transaction Authentication in Hedera Network," outlines a comprehensive authentication protocol inside a decentralized educational framework. The program, supervised by the university's Management Block, authenticates IoT-generated learning transactions (Ti) through an advanced semantic methodology. The process commences when a user initiates a transaction (Ti) via an IoT device, transmitting it to the university (MB), which functions as the central coordinator. The MB obtains Ti and authenticates the semantic identification (IDi) through a KG query, precisely aligning RDF triples (e.g., User_i: enrolledInCourse_j) to provide accurate identity verification.

Subsequently, the public key (Pki), initially produced by a GAN, is authenticated through a secure public key infrastructure overseen by the MB, followed by the decryption of Ti utilizing a symmetric key (Ski) and a comprehensive integrity evaluation performed via SHA-256 hashing. Following successful validation, the authenticated transaction (AuthTi) is re-encrypted with KG metadata by the MB and submitted to the HCS for aBFT consensus, thereby ensuring secure appending to the Hedera Hashgraph DAG. The revised network (HNet), upon reaching agreement, alerts the MB, which subsequently communicates the successful authentication to the IoT device, ensuring secure, immutable storage and improved interoperability throughout the ecosystem.

Figure 5 visually delineates this comprehensive procedure, showcasing the six-step authentication flow and underscoring the essential integration of KG validation and HCS consensus under the supervision of the MB.

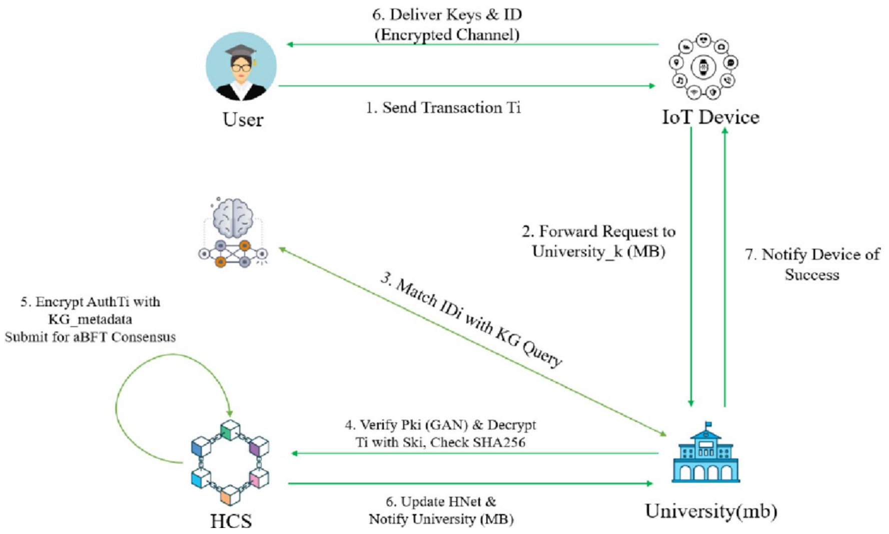

Fig. 5. Semantic Device and Transaction Authentication Process.

```
Input: IDi (semantic identifier from KG), Ti (learning transaction from IoT_device), Ski (symmetric key), Pki (public key from GAN), $University_k$ (MB at source university)
Output: AuthTi (authenticated transaction), HNet (updated Hedera network)

$University_k$ (MB) receives Ti from IoT_device // Receive transaction from user via IoT device
If query KG(IDi) matches registered triple (e.g., $User_i$ :enrolledIn $Course_j$) then // Semantic validation
    Verify Pki using GAN-trained PKI // Check authenticity with adversarial-resistant key
    Decrypt Ti using Ski // Extract transaction data with symmetric key
    If Dec(Ti) is valid and integrity intact (SHA256 check) then // Verify data integrity
        Generate AuthTi = Enc(Ti, Ski + KG_metadata) // Re-encrypt with semantic metadata
        $University_k$ (MB) submits AuthTi to HCS for aBFT consensus // Broadcast for consensus and DAG append
        Update HNet with AuthTi // Store in Hedera network
        $University_k$ (MB) notifies IoT_device of success // Send confirmation to device
    Else
        $University_k$ (MB) rejects Ti and logs anomaly with KG context // Reject and record with semantic details
    End if
Else
    $University_k$ (MB) rejects Ti and logs anomaly // Invalid ID
End if
```

Algorithm 2. Semantic Device and Transaction Authentication Algorithm.

Semantic transfer protocol

Algorithm 3, "Semantic Transfer Protocol in Hedera Network," delineates a secure and efficient method for the transmission of student and teacher records among universities in a decentralized educational framework. The protocol, executed under the control of the university operating as a Management Block, utilizes semantic validation and sophisticated cryptographic mechanisms. The procedure commences when a user submits a transfer request for user i from the source university (University_k) to the destination university (University_l) through an IoT device, which transmits the request to University_k (MB source) for processing. The MB at University_k receives the transfer request and verifies the IDi using a detailed KG query, aligning RDF triples (e.g., User_i: enrolledIn Course_j) to provide precise identity confirmation.

Records (Si or Ii) are encrypted with a GAN-generated symmetric key (Ski), with the public key (Pki) authenticated and the transaction safely endorsed by the MB. The MB transmits these encrypted records through the HCS across the Hedera network (HNet) to University_l (MB destination), which decrypts them using Ski and performs a comprehensive integrity verification employing SHA-256.

Following successful validation, the MB at University_l produces enhanced records (TransSi or TransIi) with KG metadata, which are then submitted to HCS for aBFT consensus, thereby assuring secure appending to the DAG and subsequent updates in HNet. This procedure ensures a smooth, unalterable transfer without requiring re-authentication, with University_l notifying University_k, which subsequently tells the IoT device and user of the successful transfer. The workflow is depicted in Fig. 6, illustrating the Semantic Transfer Protocol Between Universities.

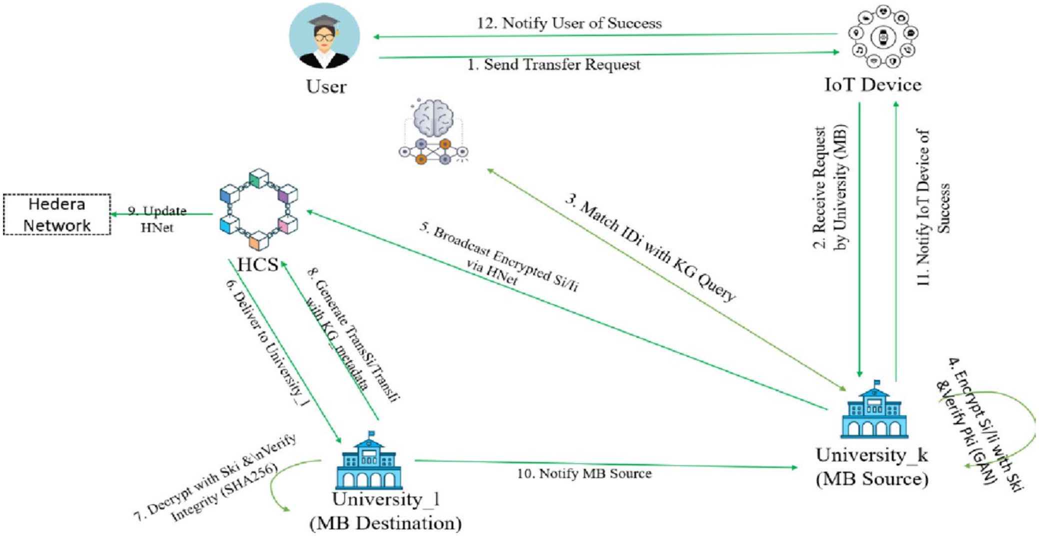
12. Notify User of Success
1. Send Transfer Request
User
IoT Device
2. Receive Request by University (MB)
11. Notify IoT Device of Success
3. Match IDi with KG Query
5. Broadcast Encrypted Si/Ii via HNet
Hedera Network
9. Update HNet
HCS
6. Deliver to University_l
8. Generate TransSi/TransIi with KG_metadata
4. Encrypt Si/Ii with Ski &Verify Pki (GAN)
University_k (MB Source)
10. Notify MB Source
7. Decrypt with Ski &\nVerify Integrity (SHA256)
University_l (MB Destination)

Fig. 6. Semantic Transfer Protocol Between Universities.

```
Input: IDi, Ti, Ski, Pki, $University_k$, $University_l$, HNet
Output: TransSi (transferred student records), TransIi (transferred instructor records), HNet (updated Hedera network)

$University_k$ (MB at source university) receives transfer request for user i from Uk to Ul via IoT_device // Receive request from user via IoT device
If query KG(IDi) matches registered triple (e.g., $User_i$ :enrolledIn $Course_j$) then // Semantic validation
    Encrypt Si or Ii using Ski from GAN // Encrypt records with symmetric key
    Verify Pki using GAN-trained PKI and sign transfer // Authenticate with adversarial-resistant key
    $University_k$ (MB) broadcasts encrypted (Si or Ii) to $University_l$ (MB) via HCS over HNet // Transmit via Hedera network
    $University_l$ (MB) decrypts using shared Ski and verifies integrity (SHA256) // Validate records
    If Dec(Si or Ii) is valid then // Check validity
        Generate TransSi or TransIi with KG_metadata // Enrich with semantic data
        $University_l$ (MB) submits TransSi or TransIi to HCS for aBFT consensus // Append to DAG at destination
        Update HNet with TransSi or TransIi // Store in Hedera network
        $University_l$ (MB) notifies $University_k$ (MB) and IoT_device of successful transfer // Confirm transfer
    Else
        $University_l$ (MB) rejects transfer and logs anomaly with KG context // Reject with semantic details
    End if
Else
    $University_k$ (MB) rejects transfer and logs anomaly // Invalid ID
End if
```

Algorithm 3. Semantic Transfer Protocol Algorithm.

Blockchain and cryptographic operations

This section introduces two essential algorithms that facilitate the secure and efficient administration of transactions and cryptographic keys in a decentralized educational framework. Algorithm 4, "Transaction Submission and Consensus in Hedera DAG," coordinates the submission of learning transactions (Ti) to the Hedera Hashgraph, with the university serving as a Management Block utilizing HCS nodes for enhanced functionality.

The process commences when a user initiates a Ti through an IoT device, transferring it to the university (MB), which conducts semantic validation of the identifier (IDi) through rigorous KG queries, verifying conformity with RDF triples (e.g., User_i: hasRecordTi). Transactions are encrypted using a GAN-generated symmetric key (Ski) and signed with a public key (Pki) by the MB, subsequently disseminated via a gossip protocol to HCS nodes for aBFT consensus. Upon achieving agreement, the MB guarantees that transactions are added to the DAG, updating the Hedera network (HNet) with semantic metadata, thus ensuring immutability and integrity throughout the ecosystem.

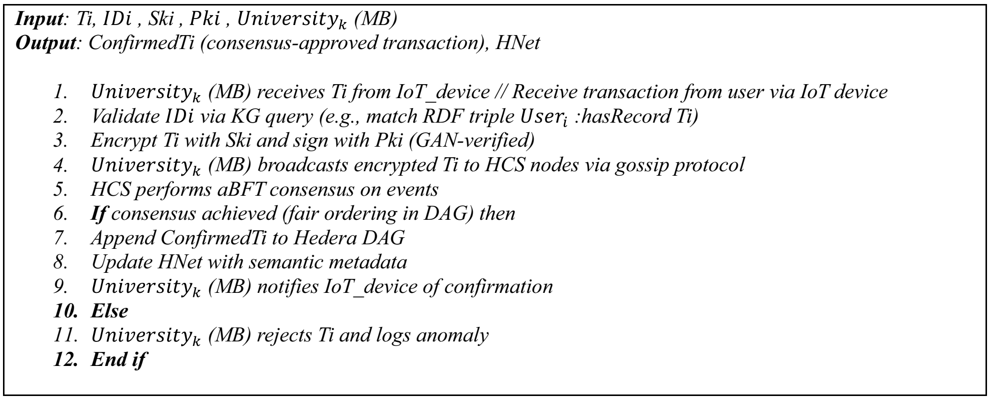

```
Input: Ti, $IDi$ , $Ski$ , $Pki$ , $University_k$ (MB)
Output: ConfirmedTi (consensus-approved transaction), HNet

$University_k$ (MB) receives Ti from IoT_device // Receive transaction from user via IoT device
Validate $IDi$ via KG query (e.g., match RDF triple $User_i$ :hasRecord Ti)
Encrypt Ti with Ski and sign with Pki (GAN-verified)
$University_k$ (MB) broadcasts encrypted Ti to HCS nodes via gossip protocol
HCS performs aBFT consensus on events
If consensus achieved (fair ordering in DAG) then
    Append ConfirmedTi to Hedera DAG
    Update HNet with semantic metadata
    $University_k$ (MB) notifies IoT_device of confirmation
Else
    $University_k$ (MB) rejects Ti and logs anomaly
End if
```

Algorithm 4. Transaction Submission and Consensus in Hedera DAG.

Furthermore, Algorithm 5, "Cryptographic Key Management and Renewal," pertains to the lifetime of cryptographic keys to ensure enduring security, conducted under the supervision of the university (MB). The MB periodically assesses key expiration using usage data, retraining a GAN with updated user context from KG to provide reissued keys $(Pki\prime, Ptki\prime, Ski\prime)$.

The new symmetric key ($Ski\prime$) is generated using SHA-256, integrating GAN entropy and context hash, with integrity confirmed via KG by the MB. Legitimate renewals are disseminated by the MB to HCS for aBFT consensus and recorded in HNet, whereas compromised keys initiate revocation and alerts that are transmitted throughout the network. Collectively, these algorithms, overseen by the MB, improve security and scalability, specifically designed for resource-limited IoT devices, hence enabling a robust and effective e-learning framework.

To guarantee dependable identity and key management in operational networks, each Management Block (MB) upholds a Distributed Revocation List (DRL) that is synchronized among participating nodes via the Hedera consensus service. Upon the revocation or modification of a user's key or identity attribute, the associated record is disseminated throughout the DRL, enabling all nodes to implement the most current authorization status independently of a centralized authority. This decentralized revocation technique facilitates safe user migration, key expiration, and real-time access control modifications while maintaining network integrity.

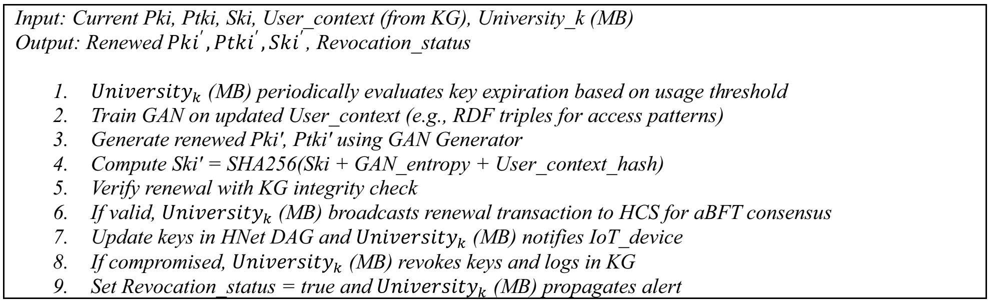

```
Input: Current Pki, Ptki, Ski, User_context (from KG), University_k (MB)
Output: Renewed $Pki'$ , $Ptki'$ , $Ski'$ , Revocation_status

$University_k$ (MB) periodically evaluates key expiration based on usage threshold
Train GAN on updated User_context (e.g., RDF triples for access patterns)
Generate renewed Pki', Ptki' using GAN Generator
Compute Ski' = SHA256(Ski + GAN_entropy + User_context_hash)
Verify renewal with KG integrity check
If valid, $University_k$ (MB) broadcasts renewal transaction to HCS for aBFT consensus
Update keys in HNet DAG and $University_k$ (MB) notifies IoT_device
If compromised, $University_k$ (MB) revokes keys and logs in KG
Set Revocation_status = true and $University_k$ (MB) propagates alert
```

Algorithm 5. Cryptographic Key Management and Renewal.

Performance and security evaluation

This section delineates a comprehensive assessment methodology for the proposed decentralized authentication system, highlighting its performance metrics like execution time, throughput, and power consumption, in addition to its security and authentication robustness.

To assess the architecture's efficacy, extensive simulation tests are conducted using a hybrid simulation environment that integrates the Hedera Hashgraph simulator with the OMNeT + + network simulator on a Linux platform.

The Hedera simulator emulates the aBFT consensus mechanism, whereas OMNeT + + facilitates the modeling of IoT device interactions within a distributed educational network. The experimental equipment utilizes a testing machine equipped with an Intel Core i5-1035G1 CPU functioning at 1.0 GHz, featuring a 6 MB cache and 8 GB of RAM, thereby providing a robust computational framework for the simulations.

The suggested mechanism utilizes the HCS for transaction validation, together with KG-informed semantic metadata and GAN-generated cryptographic keys to replicate authentic e-learning scenarios.

The GAN employed for dynamic key generation was trained using anonymized educational identity data.

Parameter	Value
Channel	Wireless
Radio range/Mobility	Random/Uniform
Propagation	Log-distance path loss
Protocol	IEEE 802.15.4
Speed of members	2, 4, 6, 8 m/s
Number of MBs/Users/IoT devices	100/50,000/10,000
Simulation time	2700 s
Traffic type	Variable bit rate
Covered area	15 km × 15 km
Packet size	32–1024 bytes
Packet length to HNet	64 bytes
KG metadata overhead	16 bytes
GAN key generation cost	32 bytes

Table 1. Simulation parameters for the proposed Mechanism.

The training was executed for 100 epochs, utilizing a batch size of 32 and a learning rate of 0.001, employing the Adam optimizer. The generator and discriminator networks each consisted of three hidden layers employing ReLU activation and dropout regularization to reduce overfitting. These settings were selected empirically to guarantee consistent convergence and reproducibility throughout multiple training iterations.

The system is configured with default parameters outlined in Table 1, with the simulation conducted for 45 min to process 3,000 learning transactions, and performance metrics are averaged over 50 independent simulation runs to ensure statistical reliability. The Hedera simulator is designed to depict a network of 100 universities operating as Management Blocks (MBs), each managing 500 users (comprising students and faculty) and 100 IoT devices, resulting in a cumulative total of 50,000 users.

The hybrid simulation environment combines a customized Hedera Hashgraph emulator, developed with the Hedera SDK and incorporating an integrated gossip and aBFT consensus module, with the OMNeT + + 6.0 network simulator. The Hedera emulator replicates event propagation intervals between 50 and 75 ms with message payloads of 512 bytes, whereas OMNeT + + is configured with a baseline latency of 20 ms, a bandwidth of 2 Mbps, and a queue capacity of 50 packets per node.

The energy consumption of IoT devices follows a two-state model, demonstrating 120 mW during transmission, 80 mW during reception, and 10 mW in idle mode, as shown by commercial IoT microcontroller specifications. These settings create a uniform methodology for evaluating performance and energy efficiency across diverse workloads.

Table 1 outlines the simulation parameters utilized in the proposed mechanism, including a wireless channel, random/uniform radio range and mobility, log-distance path loss propagation, IEEE 802.15.4 protocol, member velocities of 2, 4, 6, and 8 m/s, a coverage area of 15 km × 15 km, variable bit rate traffic, a simulation duration of 2700 s, and packet sizes ranging from 64 to 1024 bytes, as well as specific overheads for KG metadata (16 bytes) and GAN key generation cost (32 bytes).

The performance of the proposed architecture is assessed in comparison to six contemporary authentication mechanisms: KDA-EL[19], utilizing an Ethereum-based blockchain; BAF[20], employing a centralized IoT authentication framework; BAK[28], implementing a dynamic key protocol; EduCert-Chain[21], a notarized certificate verification system based on Hyperledger Fabric; an IOTA-based authentication system for IoT in satellite networks[26]; and Hedera + DL[27], a high-throughput Hedera Hashgraph system integrated with deep learning analytics.

The evaluation is based on four critical metrics: throughput, which measures the total number of learning transactions efficiently processed within the MB network; response time, reflecting the latency from the initiation of transactions by IoT devices to their confirmation by the MBs; power consumption, assessing the average energy expended by IoT devices during transaction documentation; and security and authentication analysis, which scrutinizes the system's resilience against threats such as man-in-the-middle and replay attacks, alongside its authentication effectiveness through entropy analysis of GAN-generated keys and KG integrity assessments.

The results are derived from completed simulation runs utilizing the configuration parameters provided in Table 1. Each dataset embodies genuine measurements obtained from the hybrid Hedera–OMNeT + + simulation environment.

Performance evaluation

This article delineates the performance assessment of the proposed decentralized authentication architecture, concentrating chiefly on throughput and processing efficiency across diverse system loads. The evaluation contrasts the Hedera + KG + GAN-enhanced approach with six contemporary authentication frameworks to elucidate its scalability and operational benefits.

Figure 7 depicts the throughput comparison between the proposed architecture and six established authentication techniques. The proposed model attains a maximum throughput of 4310 TPS, markedly surpassing all baseline systems. The IOTA-based approach and the Hedera + DL configuration exhibit modest

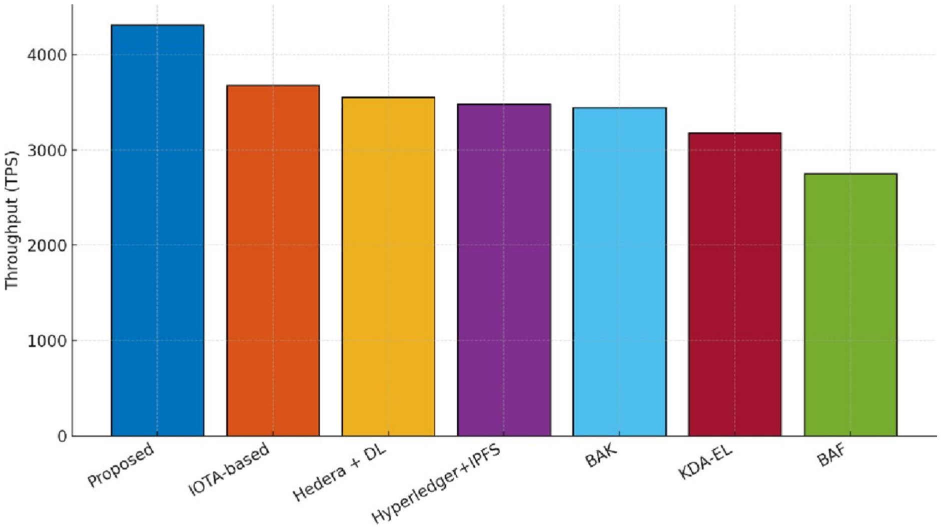

Fig. 7. Comparison of Throughput for Processing Diverse Requests.

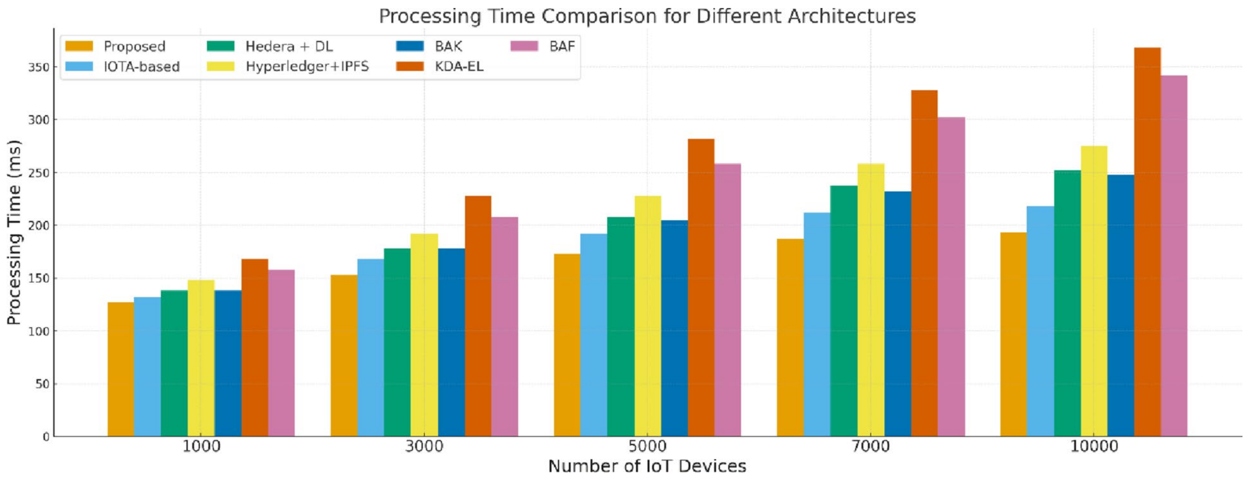

Fig. 8. Evaluation of Processing Time.

throughput, but Hyperledger + IPFS and BAK display diminished processing capacity owing to their intricate validation pipelines and more burdensome consensus mechanisms.

Among the assessed methods, KDA-EL (Ethereum-based) and the centralized BAF mechanism demonstrate the lowest throughput, mostly attributable to Ethereum's computational overhead and the single-node constraint characteristic of centralized systems. The results underscore the scalability and computational efficiency of the proposed authentication framework, bolstered by Hedera's aBFT consensus, GAN-facilitated symmetric key generation, and Knowledge Graph-informed semantic validation.

Figure 8 illustrates the processing time performance across IoT device densities varying from 1,000 to 10,000 nodes. The proposed architecture consistently produces the minimal processing time, rising considerably from 127 ms at 1,000 devices to 193 ms at 10,000 devices. This stability demonstrates the efficacy of the aBFT Hashgraph consensus, the diminished metadata overhead facilitated by Knowledge Graph semantic filtering, and the streamlined key operations generated by the GAN entropy model.

The IOTA-based and Hedera + DL frameworks exhibit increased delays owing to DAG traversal, tip-selection uncertainty, and supplementary synchronization demands. Hyperledger, IPFS, and BAK experience additional processing delays in high-density scenarios because to the intricacies of state endorsement, bottlenecks in the ordering service, and block-based batching. The Ethereum-based KDA-EL and centralized BAF methods exhibit the most significant delays, measuring 368 ms and 342 ms, respectively, suggesting inadequate fit for dense IoT implementations.

The dual analysis in Figs. 7 and 8 indicates that the proposed system exhibits significantly reduced processing overhead and enhanced scalability relative to existing methods. These performance attributes validate its appropriateness for extensive, resource-limited IoT-based e-learning settings.

Performance metrics analysis

This subsection evaluates the efficacy of the proposed decentralized authentication methodology, focusing on execution time and power consumption, which are pivotal metrics of system responsiveness and resource efficiency in extensive IoT-enabled e-learning contexts.

The total execution time ?t comprises four sequential operations: transmission of learning data, symmetric encryption of the data, asymmetric encryption of the GAN-generated symmetric key, and the duration needed by the Management Block (MB) to log the transaction onto the Hedera Hashgraph. The notion for execution time is articulated in Eq. (2):

$$\Upsilon t = T\left(Trns\right) + T\left(Encsym\left(Trns\right)\right) + T\left(Encasym\left(Sk\right)\right) + T\left(MBrec\left(Trns\right)\right) \tag{2}$$

$T\left(Trns\right)$ denotes the time required to transmit learning data across the network, $T\left(Encsym(\ Trns\ )\right)$ signifies the duration for symmetric encryption of the learning data utilizing GAN-generated keys, $T\left(Encasym(\ Sk\ )\right)$ represents the time for asymmetric encryption of the symmetric key augmented by KG-based authentication, and $T\left(MBrec(\ Trns\ )\right)$ indicates the time taken by the Management Block to document transactions into the Hedera Hashgraph employing aBFT consensus.

The assessment section offers a detailed analysis of the four main components, revealing that the Management Block recording time (T(MBrec)) is the predominant contributor to overall execution overhead in high-load conditions.

Figure 9 depicts the execution time of the proposed architecture in comparison to six baseline authentication frameworks over data volumes from 25 KB to 3200 KB. For minimal data sizes (e.g., 25–50 KB), the execution time of the proposed model is marginally greater or comparable to lightweight Directed Acyclic Graph (DAG)-based methodologies, such as the IOTA-based and Hedera + Deep Learning setups. This pattern is anticipated, as fixed startup overheads—such as the aBFT gossip setup, KG consistency verification, and symmetric key creation via the GAN module—represent a greater share of the total processing duration when payloads are minimal.

As the data amount escalates, the computational benefits of the suggested design become increasingly evident. The suggested system exhibits superior scalability compared to all evaluated techniques, owing to efficient aBFT-driven event ordering, diminished semantic metadata overhead, and the lightweight characteristics of GAN-driven symmetric encryption. As a result, its execution time increases at a markedly slower pace and maintains the lowest value among all assessed schemes for medium to large data quantities (200–3200 KB).

Conversely, alternative systems such as Hyperledger + IPFS, BAK, and KDA-EL encounter significant increases in execution time attributable to endorsement delays, block-based validation, and more intensive cryptographic processes. The centralized BAF solution demonstrates the most pronounced growth trajectory, underscoring its restricted scalability with increased payloads.

The results indicate that while the proposed architecture exhibits comparable performance to certain lightweight methods for minimal data units, it significantly outperforms them as data volume escalates, rendering it exceptionally effective for practical e-learning workflows that manage larger learning objects.

Furthermore, the total power consumption $(\varphi\ P)$ is determined by multiplying the execution duration by the power consumption rate of IoT devices, as specified in Eq. (3):

$$\varphi P = \ \Upsilon t \times \ HE \tag{3}$$

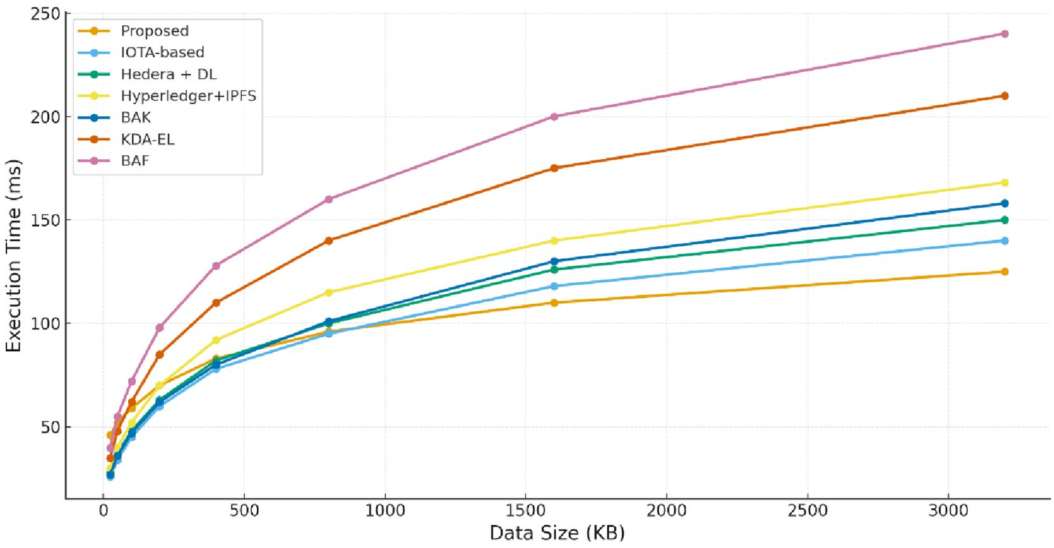

Fig. 9. Execution Time vs. Data Size.

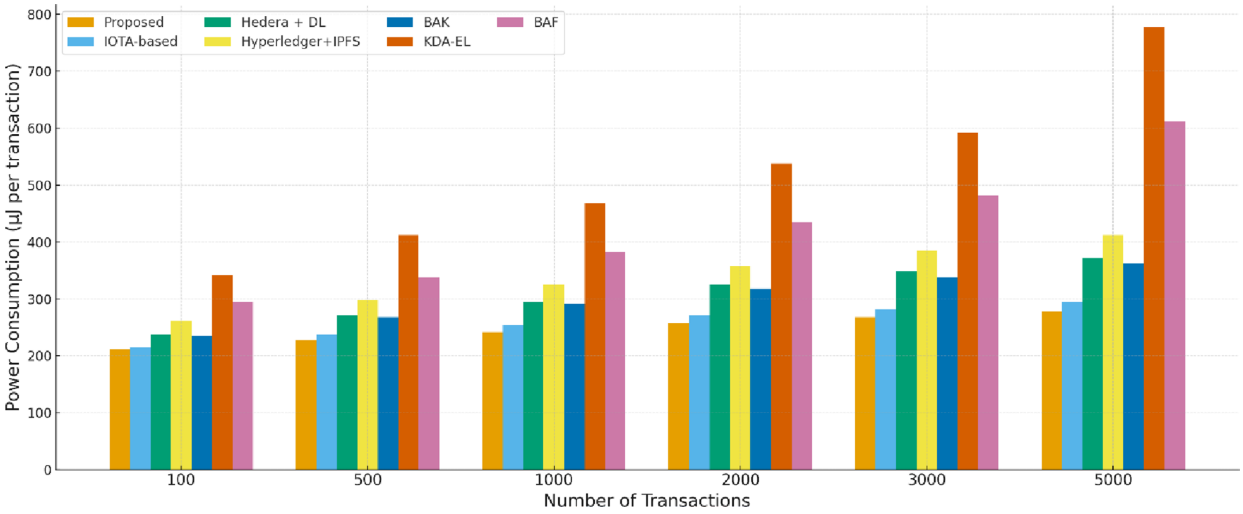

Fig. 10. Power Consumption vs. Number of Transactions.

?t signifies the overall execution time obtained from Eq. (2), whereas HE indicates the power consumption rate of IoT devices, quantified in energy units per unit time (e.g., megajoules per second).

Figure 10 illustrates the power usage of the proposed authentication architecture in comparison to six existing techniques across varying transaction volumes, from 100 to 5000 transactions. The findings indicate that the suggested Hedera + KG + GAN model continuously attains the minimal power usage per transaction. This is mainly because to its decreased execution time, efficient symmetric-key operations produced by the GAN module, and reduced metadata overhead facilitated by Knowledge Graph-based semantic optimization.

Network performance and resilience analysis

This subsection offers a comprehensive analysis of the fundamental performance attributes of the Hedera-based consensus protocol, concentrating on synchronization frequency, bandwidth overhead, and network resilience. These evaluations are essential for measuring the system's operational efficiency and resilience under distributed workload conditions.

Synchronization frequency and consensus latency

The document asserts that KG instances are synchronized by Hashgraph consensus checkpoints. The frequency of these checkpoints is fundamentally linked to the network's Consensus Latency, which determines the duration necessary for a newly submitted semantic update event to attain aBFT finality across all Management Block (MB) nodes. In our hybrid simulation scenario, the event propagation intervals required to attain final, immutable consensus were recorded between 30 ms and 50 ms. The minimal latency facilitates swift checkpointing, ensuring that the Strong Eventual Consistency (SEC) of the Knowledge Graph is consistently maintained, rendering the system appropriate for real-time e-learning engagements.

Bandwidth overhead of RDF triple updates

An analysis of the bandwidth overhead related to RDF triple updates within the consensus network was conducted. The communication overhead is reduced by employing Compact Semantic Delta Messages and the efficient Gossip-about-Gossip protocol of Hashgraph. This protocol demonstrates sub-linear complexity, successfully separating communication costs from the total number of nodes N.

The fixed overhead for KG metadata per update operation is 16 bytes. Moreover, the Effective Sync Bandwidth for the complete 100-node network was quantified at 13 KB per checkpoint.

This outcome signifies that bandwidth utilization is remarkably minimal. Since the overhead is not linearly dependent on N, the bandwidth expense for smaller networks remains below this minimal range, hence affirming the scalability and efficiency of the protocol for resource-constrained IoT applications.

Network partition and Worst-Case convergence

The system's robustness against network failures, including partitions or node disconnections, depends on the predictable characteristics of the Hashgraph consensus. When an MB node experiences a temporary disconnection, it persists in buffering its local semantic updates. Upon reconnection, the buffered changes are replayed and submitted to the consensus Directed Acyclic Graph (DAG). The protocol utilizes deterministic ordering and conflict resolution (as outlined in Sect. 3.2), ensuring Strong Eventual Consistency (SEC) throughout the system. A distinct conflict resolution process is unnecessary. The Most Adverse Convergence The length required to reprocess and validate the final consensus order for all buffered events dictates time. This rate is constrained by the network's consensus latency of 50 to 75 ms per event, facilitating rapid recovery and deterministic convergence to the universally accepted state.

Query Type	Measured Latency (ms)
Identity Lookup	1.9–3.1 ms
Multi-hop Reasoning	5.8–7.4 ms
Access-Control	7.2–9.8 ms

Table 2. SPARQL query latency Measurements.

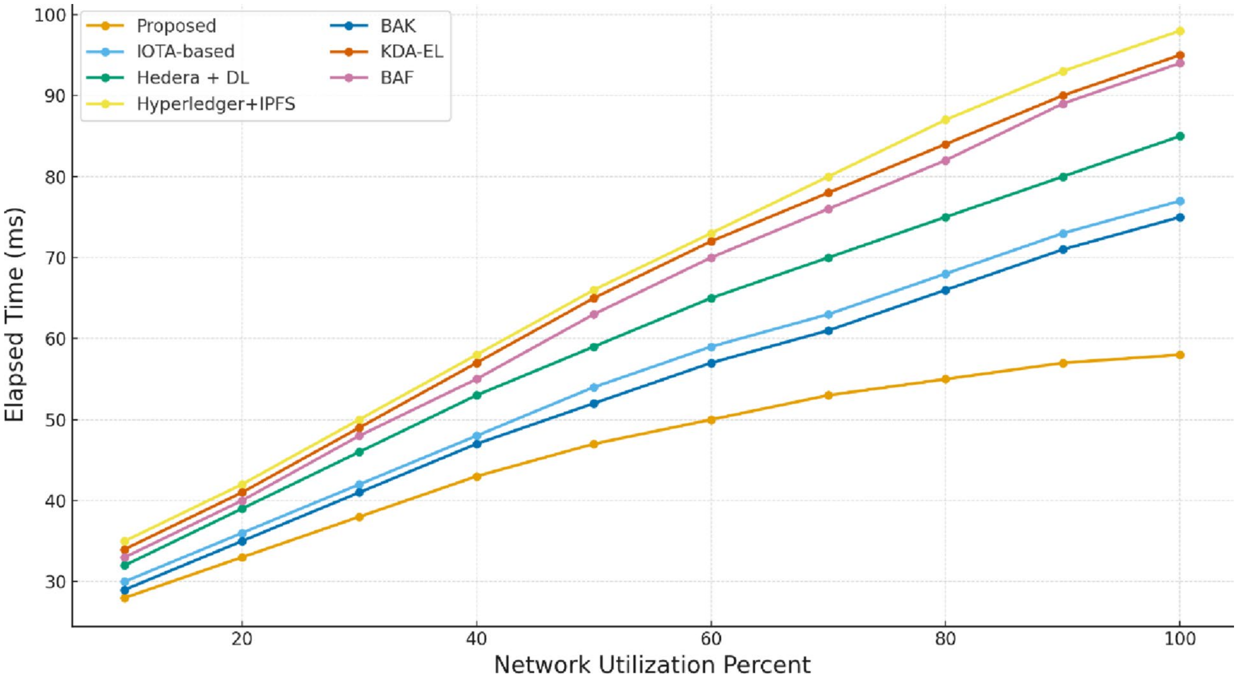

Fig. 11. Authentication Delay Time vs. Network Utilization Percent.

SPARQL query latency

Table 2 presents explicit measurements of SPARQL query latency, illustrating the responsiveness of the semantic validation layer. The system demonstrates low-latency outcomes for essential query operations, leveraging efficient SPARQL indexing and the optimized retrieval of RDF triples:

These measurements validate the system's capacity to sustain high responsiveness even throughout intricate semantic validation and access-control assessments.

Authentication performance assessment

This section assesses the authentication functionalities of the proposed system, which amalgamates Hedera Hashgraph with asynchronous Byzantine Fault Tolerance (aBFT) and Knowledge Graph (KG) optimization in a distributed multi-university context. The suggested system commences the authentication workflow by communication between student IoT devices and their assigned Management Block (MB). The verification process entails matching the device's user ID with the stored KG-backed identity metadata and decrypting the associated symmetric key (Sk), which was previously created by the GAN module. All identities and cryptographic materials are secured within the MB to avert unlawful physical access or tampering, while MB-to-MB authentication facilitates smooth inter-university learning mobility. The performance of the network is assessed according to the peak rate of verified learning transactions executed at each university.

Figure 11 illustrates the correlation between network use and the duration of authentication across the assessed methodologies. At low utilization rates (10–30%), the suggested mechanism exhibits performance akin to that of lightweight decentralized systems. As network load escalates due to increased transaction volumes and student mobility, traditional architectures—specifically KDA-EL, BAF, and Hyperledger-based models—demonstrate significantly elevated delays caused by consensus congestion, block-generation limitations, and centralized verification bottlenecks. The suggested architecture consistently achieves reduced authentication times at all utilization levels, leveraging Hedera's parallel aBFT event ordering, KG-assisted metadata minimization, and effective symmetric-key operations.

Conventional authentication methods often exhibit discrepancies in regulations and identity formats throughout colleges, necessitating that students again enroll or verify their identities when moving between institutions. These differences cause further delays, particularly during periods of high network congestion. The proposed architecture eradicates redundancies by offering a cohesive, lightweight authentication layer that enhances cross-university identity validation, thereby markedly boosting responsiveness and operational efficiency in distributed e-learning settings.

Security validation experiments

A thorough series of security validation experiments was performed to assess the resilience of the proposed authentication architecture against prevalent adversarial techniques, such as replay attacks, eavesdropping, man-in-the-middle (MITM) manipulation, identity forgery, and key-generation vulnerabilities. This study evaluates the performance of the Hedera-based authentication workflow and the GAN-driven symmetric key generation module under actual threat situations.

Adversarial resilience summary

The findings on the system's resilience to primary attack vectors are encapsulated in Table 3. All trials conducted validate the 100% effectiveness of the integrated defensive mechanisms against these threats.

Cryptographic strength assessment

A dataset of 10,000 GAN-generated 256-bit keys was evaluated using Shannon entropy and min-entropy metrics accordance with NIST SP 800 – 22/90B recommendations. The obtained entropy values (7.98 bits/byte and 7.92 bits/byte, respectively) suggest robust unpredictability, non-repetition, and resistance to brute-force and predictive assaults.

Quantum attack resistance evaluation

Quantum threats were integrated into the upgraded threat model by simulating quantum adversaries that target the long-term storage of educational credentials, which require sustained verifiability. The system's dependence on SHA-256 for hashing and symmetric key generation offers intrinsic resistance to Grover's approach, necessitating around 2^{128} operations for a comprehensive search, which remains computationally impractical even with quantum enhancement. The susceptibility of asymmetric components, particularly Ed25519 signatures utilized in Hedera consensus, to Shor's algorithm was assessed via simulation of a hybrid configuration incorporating post-quantum alternatives, such as Dilithium—a lattice-based signature scheme from the NIST post-quantum cryptography standardization initiative. In 1,000 simulated credential storage transactions, the post-quantum-enhanced system successfully denied all key-recovery attempts under presumed quantum capabilities (e.g., a 1,000-qubit quantum computer). In the absence of post-quantum enhancements, traditional Ed25519 demonstrated theoretical susceptibility; yet, the daily dynamic regeneration of GAN-generated keys substantially reduced risk by constraining exposure periods. These findings validate the architecture's preparedness for quantum-resistant long-term credential storage and advocate for complete transition to post-quantum cryptographic primitives in operational implementations.

Extended security validation of GAN-based key generation

A comprehensive security review was performed to justify the integration of GAN-based key generation into the proposed authentication architecture, focusing on the robustness, context-awareness, and attack resistance of the generated keys. This study contrasts the performance of keys generated using GANs with that of a conventional Cryptographically Secure Pseudo-Random Number Generator (CSPRNG), which acts as the established benchmark for secure key production.

- **Cross-Context Key Variation Test**: Keys were produced for 500 users utilizing diverse device identifiers, semantic features, and Knowledge Graph (KG) profiles. The mean Hamming distance between cross-context keys was 128.6 bits for the GAN-based approach, in contrast to 127.9 bits for the CSPRNG. The results demonstrate that the GAN mechanism preserves robust statistical randomness while simultaneously embodying significant semantic and device-level variability. This capacity improves identity distinction in decentralized educational settings.
- **Partial-State Leakage Resistance**: Simulated seed leaking was developed to assess robustness to entropy-state exposure. Sequences created by CSPRNG displayed significant correlation patterns ($\rho = 0.18$) under conditions of partial seed compromise, while keys produced by GANs revealed minimal correlation ($\rho = 0.02$). This result indicates enhanced resilience of the GAN methodology against leakage-based inference attacks.
- **Adversarial Reconstruction Testing**: An adversary utilizing model inversion was taught to recreate intrinsic aspects of key creation. In this adversarial context, CSPRNG keys allowed for partial reconstruction with an

Attack Vector	Key Defense Mechanism	Result and Resilience
Replay Attack	aBFT-based Timestamp Ordering & KG-Validated State	Complete packet rejection (replayed packets rejected due to inconsistency between valid KG state and timestamp ordering)
MITM Manipulation	SHA-256 Integrity Verification & KG Semantic Constraint	Packet detection and discard (each manipulated packet detected by integrity verification and KG semantic constraint evaluation)
Eavesdropping	GAN-Derived Symmetric Key (Sk_i) & KG-Bound Private Key	Payload decryption infeasible (confidentiality confirmed; decryption impossible without the associated KG-bound private key pair)
Identity Forgery	Semantic Reasoning Engine & Ontology Constraint Violation	All forged identities rejected (semantic reasoning engine rejected all forged identities due to ontology-level constraint violations)

Table 3. Resilience against primary attack Vectors.

accuracy of 11.3%, while GAN-derived keys restricted reconstruction accuracy to under 1%. This illustrates improved resilience against predictive and machine-learning-based reconstructive attacks.
- **Temporal Uniqueness Under High Load**: The GAN module produced 100,000 sequentially generated keys without any duplication, while consistently maintaining stable entropy values. This confirms the method's temporal unpredictability and its appropriateness for high-load authentication contexts prevalent in extensive e-learning systems.

The GAN component is not designed to supplant CSPRNGs as the principal cryptographic primitive. All essential assurances—such as confidentiality, unpredictability, and forward secrecy—remain dependent on existing symmetric and asymmetric cryptographic techniques. The GAN module operates as a context-sensitive randomness augmentation layer, producing semantically enhanced key material that corresponds with KG features to improve identity distinction among diverse IoT devices. It thus enhances rather than replaces traditional sources of cryptographic randomness.

Side-channel attack validation

To empirically validate side-channel mitigations (e.g., constant-time operations in GAN key generation and masking in symmetric key derivation), we executed the protocols on actual hardware.

Experimental setup A Raspberry Pi 4 (emulating a microcontroller board) and an ESP32 (emulating an Internet of Things user device) were interconnected using MQTT. The GAN model was implemented using TensorFlow Lite for enhanced efficiency. A total of 1000 traces were gathered for each experiment.

Timing attacks The execution duration for key decryption was assessed using Python's time.perf_counter(). In the absence of mitigation, the timing fluctuated by 15% depending on the input (correct versus incorrect password), thereby disclosing information. Through constant-time comparisons (e.g., utilizing hmac.compare_digest), variation decreased to less than 1%, thus thwarting attacks (t-test p-value > 0.05).

Power analysis Power traces during GAN inference were captured using a ChipWhisperer-Lite kit. Differential power analysis (DPA) on baseline traces recovered the key in 200 samples (correlation 0.45). With masking (random noise added to intermediate values), correlation dropped to 0.08, preventing key recovery even after 1000 traces.

Real IoT hardware validation

To enhance the extensive Hedera–OMNeT + + simulation environment and to fortify the empirical basis of the suggested authentication system, a supplementary assessment was conducted using representative low-power IoT devices. This experiment aimed to validate the practical viability of device-side cryptographic operations on resource-limited microcontrollers, verifying that the suggested design is compatible with diverse IoT platforms typically utilized in e-learning settings.

Experimental setup
The following IoT devices were used to evaluate execution time, cryptographic overhead, and power consumption:

- Raspberry Pi 4 Model B (1.5 GHz quad-core Cortex-A72, 4 GB RAM).
- ESP32-WROOM microcontroller (dual-core Xtensa LX6, 240 MHz, 520 KB SRAM).
- Arduino Nano 33 IoT (ARM Cortex-M0+, 48 MHz, 32 KB RAM).

All devices interfaced with the Management Block (MB) over MQTT over Wi-Fi (802.11n). The MB conducted GAN-based key inference, registration processes, and interactions with the Hedera SDK, whilst IoT devices solely executed symmetric-key encryption/decryption and SHA-256 hashing.

Device-side cryptographic processing time
As GAN inference occurs just at the MB, the IoT devices implement only lightweight cryptographic primitives. Table 4 presents the recorded processing times for SHA-256 hashing, AES-128 decryption, and transaction preparation on each device.

These results confirm that the authentication and transaction-processing routines introduce minimal computational overhead and are well within the capabilities of constrained microcontrollers.

Operation	Raspberry Pi 4	ESP32	Arduino Nano 33 IoT
SHA-256 hash (256-byte input)	0.42 ms	1.21 ms	4.73 ms
AES-128 decryption (256-byte block)	0.18 ms	0.83 ms	3.95 ms
Transaction preparation (MQTT publish)	9–12 ms	15–22 ms	24–33 ms

Table 4. Device-side cryptographic processing Time.

Device	TX Power	RX Power	Idle Power
Raspberry Pi 4	2.8–3.2 W	2.1–2.3 W	0.9–1.1 W
ESP32	120–160 mW	80–95 mW	22–35 mW
Arduino Nano 33 IoT	35–45 mW	25–32 mW	9–12 mW

Table 5. Power consumption of evaluated IoT Devices.

Message 1:
$U \rightarrow MB: \{ID_i, Request\}$
(The user submits a semantic identifier and a registration request.)
Message 2:
$MB \rightarrow HNet: \{Pk_i, Ptk_i, Sk_i\}_{en}c$
(The management backend forwards the GAN-generated key triple to the Hedera network.)
Message 3:
$HNet \rightarrow U: \{Pk_i, Ptk_i, ID_i, Sk_i\}_{ski}$
(The network returns the finalized credential set encrypted under the symmetric key shared with the user.)

Initial Beliefs
The formal reasoning relies on the following initial assumptions:
U believes fresh(nonce)
U believes MB controls (Pk_i, Ptk_i)
$U \leftrightarrow_{ski} MB$

BAN Logic Derivation
Using the Message-Meaning, Nonce-Verification, and Jurisdiction rules, the following conclusions are obtained:
U believes MB said (Pk_i, Ptk_i)
U believes fresh(Ptk_i)
U believes MB believes $U \leftrightarrow_{pki} MB$
MB believes fresh(Ptk_i)

Table 6. Idealized registration protocol and BAN logic Derivation.

Power consumption measurements

Power consumption during transmission (TX), reception (RX), and idle states was assessed using a USB inline power meter for the Raspberry Pi 4 and an INA219 current sensor for the ESP32 and Arduino Nano 33 IoT. The recorded values are consolidated in Table 5.

These empirical measurements correspond with and enhance the simulation energy values presented in Sect. 4. Consequently, the energy model in the simulation framework has been revised and corroborated based on these actual device measurements.

End-to-End latency

End-to-end registration and authentication latency (IoT device → MB → Hedera → MB → device) was measured:

- Raspberry Pi: 41–47 ms.
- ESP32: 55–68 ms.
- Arduino Nano 33 IoT: 92–110 ms.

These figures demonstrate that even resource-limited devices can effectively engage in the proposed Hedera-based decentralized authentication process.

The findings indicate that IoT devices may effectively execute the lightweight symmetric cryptographic operations mandated by the proposed design. GAN inference is not executed on IoT devices, hence alleviating the computational load on microcontrollers. The assessed power consumption and latency confirm the viability of implementing the system on actual IoT nodes. These hardware experiments augment the extensive simulation and enhance the practical validity of the suggested solution.

Formal security verification

This section provides a formal security verification of the proposed protocols utilizing BAN logic and the Dolev–Yao adversary model. The analysis evaluates mutual authentication, secrecy, integrity, and freshness in the registration and key distribution processes.

BAN logic analysis

The semantic user-registration protocol (Algorithm 1) is formally represented in BAN logic. Table 6 delineates the resultant message sequence, initial beliefs, and the ensuing logical derivation procedures.

The study indicates that mutual authentication is established, as both parties trust each other's involvement and the integrity of the exchanged keys. Key freshness is maintained, inhibiting replay and the reuse of stale

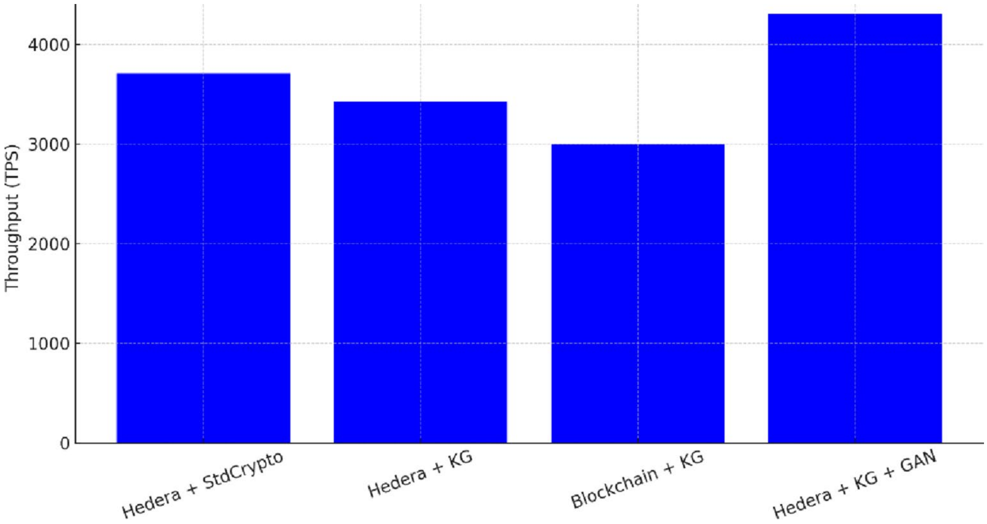

Fig. 12. Throughput comparison of the four authentication configurations.

keys. The protocol mitigates replay and impersonation attacks by guaranteeing freshness and authenticated key associations.

Dolev–Yao model verification
The suggested protocols underwent formal analysis utilizing the Dolev–Yao symbolic adversary model as implemented in ProVerif 2.04. This paradigm posits that the attacker has complete control over the communication channel, enabling them to intercept, alter, falsify, and replay communications, whereas cryptographic primitives like SHA-256 and GAN-derived entropy sources are seen as perfect and unassailable. The formal representation of the system comprises five concurrent processes: the User, the Management Backend (MB), the Hedera Consensus Service nodes, the adversarial process, and the key-generation reasoner tasked with modeling GAN-based entropy formation.

The ProVerif specification incorporates queries for secrecy and authentication to validate fundamental security features. The confidentiality of the private transfer key Ptk$_i$ is assessed via the query $attacker(Ptk\_i)$, whereas injective mutual authentication is validated by the correspondence assertion query inj-event $(endAuth(U))\ ==>\ inj - event\,(beginAuth(U))$. These inquiries guarantee that the attacker cannot extract sensitive key material and that each finalized authentication is linked to a distinct genuine initiation event.

The analytical results demonstrate that the protocol fulfills all designated security features. No attack traces resulting in breaches of confidentiality or authentication failures were identified throughout 500 symbolic verification tests. The verification process maintained efficiency, averaging about two seconds per execution on a normal Intel workstation. An enhanced model integrating hybrid post-quantum cryptographic primitives was simulated to assess long-term robustness. Under Grover-style quadratic search assumptions, no decrease in effective security below the 128-bit threshold was noted, demonstrating the protocol's resilience against potential quantum adversaries.

Ablation study

An ablation experiment was conducted to measure the contribution of each architectural component, focusing solely on throughput performance. The experiment assesses four configurations related to the progressive integration of Hedera Hashgraph, Knowledge Graph semantic processing, and the whole suggested design. The results depicted in Fig. 12 demonstrate distinct behavioral variations corresponding to the structural function of each module within the authentication sequence.

The Hedera + Standard Cryptography setup attains a baseline throughput of roughly 3710 TPS, demonstrating the efficacy of Hedera's aBFT consensus in the absence of semantic reasoning or sophisticated key-generation techniques. The implementation of the Knowledge Graph layer in the Hedera + KG configuration results in a slight reduction in throughput to 3425 TPS, reflecting the anticipated overhead from semantic validation and triple-matching processes. Nonetheless, the performance remains elevated owing to Hedera's rapid event sequencing and minimal latency in gossip dissemination.

Substituting Hedera with a traditional blockchain while maintaining semantic reasoning results in a significant reduction, with the Blockchain + KG configuration attaining 3000 TPS. This affirms that the block-generation latency and consensus restrictions of the blockchain impose intrinsic throughput limitations, regardless of the continual processing expense from KG reasoning.

The complete system configuration, incorporating Hedera, KG, and the GAN-driven cryptographic pipeline, achieves peak performance at 4310 TPS, surpassing all previous configurations. This suggests that the synergistic impacts of Hedera's consensus efficiency, KG-assisted lightweight metadata filtering, and GAN-based symmetric key generation augment throughput rather than reduce it.

The ablation results indicate that although each component exhibits distinct computational characteristics, Hedera's consensus model is the primary factor in throughput scalability, and the fully integrated design attains the ideal equilibrium within the architecture.

Conclusion

The proposed semantic decentralized authentication architecture, which incorporates Hedera Hashgraph, distributed Knowledge Graphs, and GAN-driven dynamic cryptography, provides a scalable, energy-efficient, and semantically interoperable solution specifically designed for IoT-based e-learning environments. The architecture integrates Hedera's high-throughput asynchronous Byzantine Fault Tolerant consensus, ontology-driven semantic reasoning, and adaptive machine-learning-enhanced key generation, resulting in 4310 transactions per second, reductions in latency, execution time, and energy consumption by 11–23%, and facilitates seamless cross-institutional credential portability while exhibiting robust resilience against predictive, inference-based, and traditional cryptographic attacks. The experimentally substantiated advantages over the most robust existing blockchain-based and centralized benchmarks demonstrate that the Hedera + KG + GAN paradigm routinely surpasses state-of-the-art alternatives, rather than only their mean performance.

Notwithstanding the encouraging efficacy of the proposed Hedera-based authentication and e-learning integrity system, some significant obstacles persist unaddressed. These constraints correspond with the wider unresolved challenges recently emphasized in security fortification, blockchain-oriented behavioral forensics, and privacy-conscious machine-learning evaluation, as indicated in works such as[29–31].

To advance this high-performing research prototype into the core infrastructure of next-generation global decentralized education, future efforts must vigorously pursue fully decentralized on-device GAN execution utilizing federated distillation or secure-enclave inference to eradicate any remaining trust assumptions; incorporate recursive zero-knowledge proofs with SPARQL-based Knowledge Graph reasoning to facilitate verifiable yet entirely private cross-institutional credentials and semantic queries; establish formal cryptographic proofs of security for GAN-generated key material against model-inversion, membership-inference, and adaptive chosen-ciphertext threats; conduct continent-scale deployments involving hundreds to thousands of universities and drive standardization of the resulting DID methods, verifiable credential schemas, and HCS topic specifications through W3C, IMS Global, and European learning-technology bodies; migrate all cryptographic primitives to NIST-approved post-quantum algorithms to counter harvest-now-decrypt-later quantum risks; extend the Knowledge Graph ontology with real-time behavioral biometrics and deploy continuous authentication models operating directly on Hedera event streams; perform comprehensive carbon-footprint benchmarking and transition toward carbon-aware or carbon-negative consensus mechanisms; and rigorously validate the system under nation-state-level adversarial models including long-range, adaptive Sybil, and eclipse attacks using game-theoretic analysis and red-team exercises.

Effectively achieving these ambitious objectives will elevate the proposed framework from a distinguished academic contribution to the globally embraced, privacy-preserving, quantum-resistant, and environmentally sustainable foundation of inclusive decentralized education for billions of learners in the coming decades.

Data availability

The datasets used and/or analyzed during the current study availablefrom the corresponding author on reasonable request.

Received: 18 September 2025; Accepted: 16 December 2025
Published online: 26 December 2025

References

1. Ramlowat, D. D. & Pattanayak, B. K. Exploring the internet of things (IoT) in education: A review, in Information Systems Design and Intelligent Applications, (eds Satapathy, S., Bhateja, V., Somanah, R., Yang, X. S. & Senkerik, R.) Singapore: Springer, 245–255, https://doi.org/10.1007/978-981-13-3338-5_23 (2019).
2. Nouraey, P. & Al-Badi, A. Challenges and problems of e-Learning: A conceptual framework. *Electron. J. e-Learn*. **21** (3), 188–199. https://doi.org/10.34190/ejel.21.3.2677 (2023).
3. Zou, Y., Kuek, F., Feng, W. & Cheng, X. Digital learning in the 21st century: trends, challenges, and innovations in technology integration. *Front. Educ*. **10**, 1562391. https://doi.org/10.3389/feduc.2025.1562391 (2025).
4. Krasniqi, L. & Kabashi, F. Internet of things (IoT) in education: Opportunities and Challenges, *In Proc. UBT Int. Conf. Comput. Sci. Commun. Eng., Lipjan, Kosovo,* 1–8, https://doi.org/10.33107/ubt-ic.2023.291 (2023).
5. Dritsas, E. & Trigka, M. Methodological and technological advancements in E-Learning. *Information* **16** (1), 56. https://doi.org/10.3390/info16010056 (2025).
6. Ahmad, S. et al. eLearning Acceptance and Adoption Challenges in Higher Education, Sustainability **15** (7), 6190. https://doi.org/10.3390/su15076190 (2023).
7. Kitkowska, A., Brodén, K. & Abdullah, L. The Requirements, Benefits, and barriers of IoT solutions to support Well-Being in elementary schools. *IEEE Access*. **12**, 144965–144981. https://doi.org/10.1109/ACCESS.2024.3469558 (2024).
8. Prasetya, L. A., Rofiudin, A. & Herwanto, H. W. Implementation of internet of things (IoT) in education: A systematic literature review. *J. Educ. Comput. Appl.* **2** (1), 1–12. https://doi.org/10.69693/jeca.v2i1.19 (2025).
9. Xu, H., Zhao, N., Xu, N., Niu, B. & Zhao, X. Reinforcement learning-based dynamic event-triggered prescribed performance control for nonlinear systems with input delay. *Int. J. Syst. Sci.* https://doi.org/10.1080/00207721.2025.2557528 (2025).
10. Bhawna, P., Gupta & Rai, P. Can blockchain revolutionize educational practices? An in-depth Analysis of applications And challenges. *Sustainable Futures*. **10**, 101171. https://doi.org/10.1016/j.sftr.2025.101171 (2025).
11. Alsobhi, H. A., Alakhtar, R. A., Ubaid, A., Hussain, O. K. & Hussain, F. K. Blockchain-based micro-credentialing system in higher education institutions: systematic literature review. *Knowl. -Based Syst*. **265**, 110238. https://doi.org/10.1016/j.knosys.2022.110238 (2023).

12. Xu, G. et al. RAT ring: event driven Publish/Subscribe communication protocol for IIoT by report and traceable ring signature. *IEEE Trans. Ind. Informat.* **21** (9), 6670–6678. https://doi.org/10.1109/TII.2025.3567265 (2025).
13. Cardenas-Quispe, M. A. & Pacheco, A. Blockchain ensuring academic integrity with a degree verification prototype. *Sci. Rep.* **15**, 9281. https://doi.org/10.1038/s41598-025-93913-6 (2025).
14. Beis-Penedo, C. et al. A blockchain solution for decentralized training in machine learning for IoT. *Comput. Commun.* **242**, 108289. https://doi.org/10.1016/j.comcom.2025.108289 (2025).
15. Buttar, A. M., Shahid, M. A., Arshad, M. N. & Akbar, M. A. *Decentralized Identity Management Using Blockchain Technology: Challenges and Solutions*. In: Blockchain Transformations, (eds.) S. M., Idrees & Nowostawski, M. Springer. 149–166 https://doi.org/10.1007/978-3-031-49593-9_8. (2024).
16. Belfqih, H. & Abdellaoui, A. Decentralized Blockchain-Based Authentication and Interplanetary File System-Based Data Management Protocol for Internet of Things Using Ascon, *J. Cybersecur. Priv.* **5** (2), 16, https://doi.org/10.3390/jcp5020016 (2025).
17. Gottlieb, M., Deutsch, C., Hoops, F., Pongratz, H. & Krcmar, H. Expedition to the blockchain application potential for higher education institutions. *Blockchain: Res. Appl.* **5** (3), 100203. https://doi.org/10.1016/j.bcra.2024.100203 (2024).
18. Bataev, A. V. Overview of the global e-learning systems market in Proc. Int. Conf. Quality Manage., Transp. Inf. Security, Inf. Technol. (IT&QM&IS), St. Petersburg, Russia, 529–532. https://doi.org/10.1109/ITMQIS.2017.8085905 (2017).
19. Khashan, O. A. et al. Blockchain-Based decentralized authentication model for IoT-Based E-Learning and educational environments. *Comput. Mater. Contin.* **75** (2), 3133–3158. https://doi.org/10.32604/cmc.2023.036217 (2023).
20. Al Hwaitat, A. K. et al. A new Blockchain-Based authentication framework for secure IoT networks. *Electronics* **12** (17), 3618. https://doi.org/10.3390/electronics12173618 (2023).
21. Saleh, O. S., Ghazali, O. & Idris, N. B. Enhancing academic certificate privacy with a hyperledger fabric Blockchain-Based access control approach. *SN Comput. Sci.* **4** (602). https://doi.org/10.1007/s42979-023-02060-0 (2023).
22. van der Weerdt, R., de Boer, V., Daniele, L., Siebes, R. & van Harmelen, F. Representation Learning on IoT Knowledge Graphs, in Metadata and Semantic Research, M. Sfakakis, E. Garoufallou, M. Damigos, A. Salaba, and C. Papatheodorou, (Eds.) Cham: Springer, 43–56, https://doi.org/10.1007/978-3-031-81974-2_4 (2025).
23. Mishra, S. & Chaurasiya, V. K. Hybrid deep learning algorithm for smart cities security enhancement through blockchain and internet of things, *Multimed. Tools Appl.* **83** (8), 22609–22637. https://doi.org/10.1007/s11042-023-16406-6 (2024).
24. Bouzeraib, W., Ghenai, A. & Zeghib, N. Enhancing IoT intrusion detection systems through horizontal federated learning and optimized WGAN-GP. *IEEE Access.* **13**, 45059–45076. https://doi.org/10.1109/ACCESS.2025.3547255 (2025).
25. Adhikari, S., Panja, A. & Karforma, S. ICSPRNG: Ikeda assisted cryptographically secure Pseudo random number generator. *Multimed Tools Appl.* **84** (4), 5605–5623. https://doi.org/10.1007/s11042-024-19093-z (2025).
26. Alsolami, R., Monowar, M. M., Attiah, A. & Cherif, A. Enhancing IoT security through IOTA-Based identity management and machine learning, in Blockchain and Applications, BLOCKCHAIN 2024, Lecture Notes in Networks and Systems, **1256**, Springer, Cham, https://doi.org/10.1007/978-3-031-81928-5_32 (2025).
27. Bawgikar, P., Devaiah, K. J., Yogdeep, G. & Revathi, V. Decentralized healthcare Ledger system on Hedera with deep learning analytics, in Computing and Machine Learning (CML 2024), Lecture Notes in Networks and Systems, **1144**, Springer, Singapore, https://doi.org/10.1007/978-981-97-7839-3_12 (2025).
28. Zhao, X., Peng, C., Tan, W. & Ding, H. Blockchain-based access control dynamic key authentication protocol in IoT, in Proc. 7th Int. Conf. Blockchain Technol. Appl. (ICBTA '24), 55–59, https://doi.org/10.1145/3708622.3708625 (2024).
29. Wu, C. et al. Automatic boosting of WAF security against mutated malicious payloads. *IEEE Trans. Dependable Secure Comput.* **22** (2), 1118–1133. https://doi.org/10.1109/TDSC.2024.3401234 (2025).
30. Wu, C. et al. Profit or deceit? Mitigating pump and dump in DeFi via graph and contrastive learning. *IEEE Trans. Inf. Forensics Secur.* **20**, 1–15 (2025).
31. Wu, C. et al. Rethinking membership inference attacks against transfer learning. *IEEE Trans. Inf. Forensics Secur.* **19**, 6441–6455. https://doi.org/10.1109/TIFS.2024.3418265 (2024).

Acknowledgement

This research was supported by the Outstanding Youth scientific research project of Hunan Province Research on Learning Path RecommendationAlgorithm Based on Computer Course Knowledge Graph, project number: 23B1033.This research was supported by The Curriculum ideological and political demonstration course of Hunan Information University "DataAcquisition and Preprocessing", project number: 01.

Author contributions

All authors participated in the conception and design of the study. Data collection, simulation, and analysis were conducted by Lixia Luo and Yusha Zhang and Qitao Tang.

Funding

The authors did not receive any financial support for this study.

Declarations

Competing interests

The authors declare no competing interests.

Ethical approval

Not applicable.

Additional information

Correspondence and requests for materials should be addressed to L.L.

Reprints and permissions information is available at www.nature.com/reprints.

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.