International Journal of Web Information Systems

# Energy: reducing latency in IoT DLTs for AI-driven real-time solutions

Francisco Moya Perez, Francisco José Quesada Real, Luis Martínez López and Fco Javier Estrella Liebana
*Department of Computer Science, University of Jaén, Jaén, Spain*

Received 25 November 2024
Revised 26 March 2025
Accepted 28 March 2025

## Abstract

**Purpose** – Integrating Internet of Things (IoT) networks with distributed ledger technology (DLT) and artificial intelligence (AI) presents critical challenges, particularly related to latency, scalability, hardware constraints and data security. Efficient data ingestion and validation are essential to enable real-time AI processing. The main contribution of this paper is the proposal of the *Energy* consensus algorithm, designed to minimize both latency and energy consumption in such environments.

**Design/methodology/approach** – Energy is a consensus algorithm tailored for public directed acyclic graph-based DLTs in IoT contexts. It introduces a flexible transaction validation mechanism that reduces or bypasses Proof of Work requirements. The algorithm's performance is experimentally compared with IOTA under varying payload conditions.

**Findings** – Results show that Energy significantly reduces latency and energy consumption, especially for small payloads, which are common in IoT applications. These findings demonstrate Energy's ability to enhance transaction efficiency and support real-time AI model updates based on verified IoT data streams.

**Research limitations/implications** – Future work should investigate the scalability of Energy in larger and more heterogeneous IoT ecosystems, as well as its compatibility with different AI frameworks. Evaluating its performance under diverse network conditions and hardware setups would further strengthen the generalizability of the results.

**Practical implications** – The Energy algorithm enables continuous AI model updates while ensuring data integrity, traceability and low latency. Its adaptability makes it a suitable solution for large-scale IoT deployments requiring secure and efficient data processing.

**Originality/value** – This paper presents a novel consensus algorithm that bridges the requirements of IoT, DLT and AI, with a particular focus on improving latency and energy efficiency. Energy offers a robust approach for optimizing data flow and transaction processing in real-time, AI-driven IoT systems.

**Keywords** Consensus algorithm, Latency reduction, AI-driven IoT systems, DLT, DAG, Distributed Ledger Technology (DLT)

**Paper type** Research paper

## 1. Introduction

Artificial intelligence (AI) is currently one of the most prominent fields of research because its technologies enable the analysis, interpretation and management of vast amounts of daily data. It has led to the creation of powerful knowledge models that allow for real-time decision-making and the generation of various types of content, whether written, audio or visual, without limitations (Harshvardhan *et al.*, 2020). Nowadays, AI models usually

This work was partially supported by the Research Project TED2021–132073B-I00 PHADAS funded by MCIN/AEI/10.13039/501100011033 and NextGenerationEU/PRTR.

*Disclosure statement:* No potential conflict of interest is reported by the authors.

International Journal of Web Information Systems

1744-0084
DOI 10.1108/IJWIS-11-2024-0332

require large volumes of data to create and continue improving them. Despite the massive amounts of data generated daily worldwide, not all is accessible due to data protection issues, requiring explicit consent for their use (Aldoseri *et al.*, 2023). Thus, one of the main challenges in generating AI models is obtaining accessible data initially and continuously over time.

One of the technologies that generates the highest data volumes today is the Internet of Things (IoT). These environments consist of a series of sensors of different types, detecting what is happening in the real world and propagating it throughout the network, which creates a network of devices/information. Numerous IoT-based applications have been developed and deployed in real world, such as innovative environments, smart cities and industrial warehouses, among others (Bail *et al.*, 2021; Shahab *et al.*, 2022; Yang *et al.*, 2020). These environments transmit large amounts of information from diverse sensors, characterized by a high degree of heterogeneity. This involves numerous sensors with varying metrics and formats. Therefore, interpreting data generated within an IoT network is complex and challenging.

Using the millions of heterogeneous data points generated in an IoT network to create AI models could enable extracting knowledge from the flowing information and continuously improve the generated model with diverse and ongoing data streams (Abdallah *et al.*, 2020; Farahani *et al.*, 2021; Makhdoom *et al.*, 2019). However, a significant challenge arises: valuable data for creating AI models must be valid, reliable and correct from collection to processing. Furthermore, the data may be sensitive, protected by privacy regulations or require explicit rights for collection and manipulation (Zantalis *et al.*, 2019). IoT networks suffer from specific security issues at various points, including end devices, sensors, communication channels, network protocols, denial of service (DoS) and software (Wu *et al.*, 2024). Therefore, finding alternatives to guarantee data reliability is necessary to create valid AI models.

In response to these challenges, the research community has proposed using distributed ledger technology (DLT) alongside IoT to mitigate security problems (Shah *et al.*, 2019). DLT refers to the creation of distributed registries characterized by key features such as decentralization, immutability, traceability, transparency and security. These properties enable the development of applications that drive the digital transformation of processes across various domains (Raikwar *et al.*, 2019; Rose *et al.*, 2015), potentially fitting well with IoT environments to reduce security issues. A wealth of research exists on this topic, with multiple applications exploring the best of both worlds, whether in frameworks (Moya *et al.*, 2023) or various fields such as agriculture, health-care or smart environments (Muñoz-Higueras *et al.*, 2024; Ullo and Sinha, 2020; Zhu *et al.*, 2019). The most popular and widely adopted proposal is IOTA (Popov, 2018), a public DLT based on a directed acyclic graph (DAG) designed for IoT and broadly used within the community across various applications.

IoT applications that rely on DLT infrastructures present several challenges. The first relates to the hardware of the sensors, which are typically devices with minimal or no computing power and minimal energy capacity, often using batteries or cells. This minimal computing power directly impacts the Proof of Work (PoW) mechanism included in several DLT consensus algorithms or as a means of preventing spam. Studies have explored ways to reduce this impact (Chatrapathi *et al.*, 2023; Sravya *et al.*, 2023), but mitigation often involves externalizing the PoW or adopting consensus algorithms that do not depend on PoW. The second challenge concerns the sheer volume of data generated in IoT environments and its scalability impact on DLT. The scalability trilemma in DLTs (Monte *et al.*, 2020) highlights the challenge of balancing scalability,

security and decentralization simultaneously. Improving one of these properties often compromises the other two, limiting system efficiency. This dilemma is central to the design of distributed networks, and although solutions like DAGs have attempted to address it, the problem remains a structural challenge.

Moreover, most DLT proposals are blockchain-based, which entails specific scalability challenges, while DAG-based solutions offer greater flexibility and scalability. As noted above, IoT environments generate such massive amounts of data that they could cause latency and scalability issues within DLT (Khan *et al.*, 2021; Monte *et al.*, 2020). Mitigation could involve using DAG instead of blockchain, optimizing the consensus algorithm and reducing or eliminating PoW as an anti-spam measure (Moya *et al.*, 2024; Popov, 2018).

Our contribution lies in the design, implementation and evaluation of *Energy*, a novel consensus algorithm for public DAG-based DLTs. We hypothesize that *Energy*, specifically designed for low-latency and high-scalability data ingestion, is well suited for IoT applications, enabling continuous and heterogeneous data flows necessary for real-time AI model creation. This integration aims to ensure data quality, security and compliance with data protection regulations. The evaluation of *Energy* focuses on key aspects such as transaction speed, scalability, ingestion latency, energy consumption, data protection and its ability to support real-time AI model development.

Based on this hypothesis, the main contributions of our work are as follows:

- *Optimized transaction validation for constrained devices:* to address hardware limitations in IoT environments, we propose a flexible mechanism within *Energy* that allows PoW to be either delegated to the network (in exchange for a fee) or bypassed entirely with a fee. For scalability, fee-based transactions can skip PoW, requiring only the validation of pending transactions in the network.
- *IoT-oriented abstraction layer:* we introduce an abstraction layer tailored for IoT systems that enables seamless and transparent interaction with the DLT infrastructure. This layer simplifies integration and accelerates deployment in real-world applications.
- *Support for real-time AI model generation:* we tackle key challenges in current DLT-based IoT applications by enabling the generation of AI models from verified, trustworthy data streams, while preserving compliance with data protection principles.

To build our proposal, we will implement a DAG-based DLT with the proposed consensus algorithm to test its functionality and verify our hypothesis. Additionally, we will adapt the Phonendo platform (Moya *et al.*, 2023), adding Energy as a Phonendo publisher to provide the abstraction layer for IoT and DLT integration for data ingestion. Once all pieces are joined, we will conduct experiments and evaluations to assess whether consensus can be achieved with devices that do not perform PoW or delegate PoW, allowing us to adjust parameters to fit the scenario. First, we will conduct scalability tests by injecting one million transactions into three different nodes to verify that the consensus algorithm ingests all traffic with low latency, balancing the various situations that may arise in an IoT environment. Second, we will conduct tests related to the hardware and energy cost challenges to assess whether the proposal adds value to hardware capabilities/energy consumption.

The structure of the paper is as follows: Section 2 introduces the foundational concepts and challenges associated with AI, IoT and DLT. Section 3 presents our proposed *Energy* consensus algorithm, structured around four main components designed to optimize IoT data integration with DAG-based DLTs. This section also explores how *Energy* facilitates seamless integration with AI models, enabling continuous training using verified and secure

IoT data. In Section 4, we discuss experimental results and evaluate *Energy*'s performance compared to IOTA, focusing on transaction latency, scalability and energy efficiency. Finally, Section 5 summarizes our findings and discusses potential applications and future directions for *Energy* in real-time AI model development.

## 2. Background

The intersection of AI, IoT and DLT has generated significant interest within the research community, as each of these technologies addresses specific challenges in data management, analysis and security. AI transforms large volumes of data into valuable knowledge, while IoT generates massive amounts of real-time information from heterogeneous devices. However, the security and reliability of data in IoT networks present significant challenges.

In this context, DLTs have been proposed to ensure data integrity and traceability, offering properties such as transparency and decentralization. The following section reviews the fundamental concepts of each technology and previous work exploring their integration, highlighting the challenges and proposed solutions in the literature.

### *2.1 Artificial intelligence and the data challenge*

AI has grown at an impressive rate in recent years, becoming a vital tool for many sectors. The reason behind this boom is its ability to process enormous amounts of data and transform them into something useful, whether for making real-time decisions, generating creative content such as texts or images or predicting behaviors in complex situations. This capability has made AI indispensable in medicine, industry, marketing and technology (Jiang *et al.*, 2017; Zhu *et al.*, 2019).

The foundation of AI is data, which serves as the essential input for AI models to learn, detect patterns and improve over time. AI's ability to function effectively requires a large amount of high-quality information during the initial training phase and subsequent adjustments to adapt to changing environments or new situations. However, gathering this data poses significant challenges. The massive volume of daily information is not always easily accessible or suitable. Privacy concerns, data protection laws and regulations often limit access to sensitive information (Zantalis *et al.*, 2019), obtaining the proper consent needed for use is challenging. It also raises ethical and legal concerns, mainly when data is collected without the explicit consent of the individuals or entities that generated it (Wu *et al.*, 2019). Balancing the need for accessible data with privacy regulations, such as GDPR, remains a significant challenge for AI. Furthermore, not all data is appropriate for training AI models.

Data quality refers to the extent to which data meets specific standards of accuracy, consistency, completeness, relevance and timeliness, making it suitable for use in particular applications. Data quality is crucial in AI, as this data serves as the raw material for training models (Mattioli *et al.*, 2022). Poor-quality data, which may be incomplete, erroneous or inconsistent, can lead to biased or unreliable results, directly affecting the model's ability to learn and make accurate decisions. Reliable, accurate and error-free data is critical, as inaccurate or incomplete information can lead to unreliable outcomes, particularly in high-stakes applications like medical diagnosis or autonomous driving.

The importance of data quality lies not only in its accuracy but also in its ability to represent the context in which it was generated faithfully. Additionally, AI models need up-to-date data as environments change over time, and models require information that reflects those variations. Data quality impacts the precision of predictions and increases the complexity of training models, forcing researchers to spend more time and resources

cleaning and preparing the data before it can be used. Therefore, both the quality and quantity of data play critical roles in the performance and trustworthiness of AI systems.

Since AI requires large amounts of data to function correctly, IoT is an exciting alternative as it can be a constant and valuable source of information. IoT devices generate real-time data from all kinds of sensors, and this data could continuously feed AI models.

### *2.2 Internet of Things*

Technological advancements have significantly improved the stability and capacity of communication networks, enabling a wide range of devices to share information continuously. This progress has driven the development of machine-to-machine communication, where devices autonomously exchange data. Collectively, these interconnected devices form the IoT ecosystem (Ryang and Yun, 2016).

IoT-based solutions have significantly impacted various sectors and have the potential to impact many aspects of human life, including transportation, logistics, smart environments, as well as social and health-care applications (Abbas *et al.*, 2021; Golpîra *et al.*, 2021; Hajjaji *et al.*, 2021; Sravya *et al.*, 2023; Vermesan *et al.*, 2022; Zha *et al.*, 2018). IoT systems can be broadly classified by their mobility and performance. Mobile IoT networks, like those in vehicles, often face signal loss due to movement. Stationary systems, such as those in smart buildings, benefit from stable connections and fixed sensor locations (Zha *et al.*, 2016). Similarly, IoT devices range from low-cost, low-capacity sensors, like temperature or humidity monitors, to high-performance devices, including smartphones and personal computers, which offer greater versatility.

One of the defining characteristics of IoT networks is the massive volume of devices and data they generate. With projections estimating up to 30.9 billion IoT devices by 2025 (Sunny and Scott, 2025), managing this vast network requires decentralized approaches, particularly in large-scale environments like smart cities, where central management is impractical (Wang *et al.*, 2019). However, decentralization introduces challenges, such as unstable and unpredictable connections, as devices may experience varying signal strength or availability.

Another significant challenge is the interoperability between IoT devices. Given the wide range of communication protocols and software these devices use, ensuring seamless communication is often problematic, resulting in unreliable connections and operational instability. This lack of standardization complicates the integration and management of diverse IoT systems (Banafa, 2016).

From a hardware perspective, IoT devices are typically designed with specific scenarios in mind, making them difficult to modify or repurpose for other uses. Many IoT devices are built with fixed hardware configurations optimized for specific tasks, limiting their ability to handle complex computations or adapt to new security requirements. Attempting to upgrade their hardware for additional capabilities could also increase the risk of security vulnerabilities (Alsaadi and Tubaishat, 2015).

In terms of security, IoT systems are vulnerable to various attacks. Devices can be physically compromised, allowing attackers to access sensitive data such as network keys (Alsaadi and Tubaishat, 2015). Communication channels can also be intercepted if not properly encrypted, potentially revealing private data (Mehta *et al.*, 2011). Network protocols may contain vulnerabilities that attackers could exploit to launch man-in-the-middle or blackhole attacks (Zhao *et al.*, 2020). Furthermore, attackers could manipulate sensor data or introduce false information into the network, undermining the system's integrity (Zhang *et al.*, 2014). DoS attacks, which overwhelm system resources and prevent regular operation, poses a severe threat, especially given the limited processing power of

many IoT devices (Namvar *et al.*, 2016). Finally, software vulnerabilities within IoT systems could be exploited, allowing attackers to alter operations or controls (Makhdoom *et al.*, 2018).

Addressing these security challenges with traditional architectures is often highly complex, if possible. Therefore, exploring alternative architectures and paradigms is necessary to tackle these issues effectively from new perspectives. DLT technology was raised as an alternative solution for ensuring data integrity and traceability in IoT environments (Zhu *et al.*, 2019).

### *2.3 Distributed ledger technology*

DLT is founded on principles of cryptography (Ramamurthy, 2020) and distributed database systems (Popov, 2018). Cryptography ensures data integrity and nonrepudiation, while distributed database systems consist of a network of nodes that store data without a central authority. Due to these characteristics, DLTs have gained significant popularity in recent years (Ramos-Cruz *et al.*, 2024).

One of the most relevant aspects of DLTs is their immutability, meaning that once data is recorded, it cannot be altered or deleted. Additionally, they operate decentralized, without a central authority overseeing or controlling the network. Each participant in the network has a copy of the records, promoting complete transparency. To maintain consistency, the network uses consensus mechanisms that allow participants to agree on the network's state. The records within the network are protected against modification, enabling participants to trust the integrity of the data. Each data entry has different timestamps associated with it, depending on its status, from the moment it enters the system until it is validated by the network. Furthermore, transactions are secured through individual encryption and asymmetric cryptography to prevent nonrepudiation (Goldreich, 2001).

Among the various types of DLT structures, the most prominent are blockchain (Nakamoto, 2008) and DAG (Popov, 2018). While blockchain has been the foundational structure for many cryptocurrencies, DAG-based DLTs, such as IOTA's Tangle, offer innovative approaches to overcome some of the inherent challenges of traditional blockchains.

Blockchain consists of blocks linked by cryptographic hashes, ensuring the accuracy and immutability of records. Miners authenticate transactions using consensus mechanisms like PoW, but blockchain faces scalability issues and high energy consumption due to node validation requirements (Khan *et al.*, 2021).

In contrast, DAG-based structures like IOTA's Tangle do not require dedicated miners, as participants collectively maintain the ledger. While offering high scalability and low fees, their security and maturity are still under development. Preventing DDoS attacks is crucial, as they can disrupt networks and erode trust. PoW is commonly used in DLTs to mitigate SPAM and maintain network integrity (Chatrapathi *et al.*, 2023; Sravya *et al.*, 2023).

*2.3.1 Distributed ledger technology consensus algorithms.* Consensus algorithms in DLT enable multiple nodes to agree on the state of a distributed network, addressing critical issues such as consistency, Byzantine faults and activity. Consistency ensures that all nodes maintain a uniform copy of the data, which is crucial in networks with distributed nodes and unreliable conditions (Gilbert and Lynch, 2002; Lamport, 1977). Byzantine faults refer to the ability of some nodes to behave maliciously or defectively, where the algorithm must ensure that the network can still reach consensus despite these failures (Lamport *et al.*, 1982). Activity measures node participation in network tasks, such as transaction validation, ensuring that the network remains operational and unaffected by node inactivity (Liu *et al.*, 2018).

Consensus algorithms are categorized into two types: permissioned and permissionless. The former prioritizes efficiency and control, while the latter, slower but highly decentralized,

is ideal for public networks like IoT. We focus on the latter, as they allow the integration of devices without centralized authorities, ensuring decentralization and security in real-world scenarios. Table 1 details the main permissionless DLT consensus algorithms.

## 3. Using energy for Internet of Things data integration in artificial intelligence-based applications

Our proposal is based on the four main blocks depicted in Figure 1. First, we have a public IoT environment with heterogeneous sensors capturing real-world data. Second, we adapt the Phonendo platform (Moya *et al.*, 2023) to process, verify and transmit the data to the DLT infrastructure. Third, we implement a permissionless DAG-based DLT with Energy, a novel consensus algorithm tailored to specific IoT environments. Fourth, the information made

**Table 1.** Comparison of main permissionless DLT consensus algorithms

| Algorithm | Description | Improvement |
|---|---|---|
| Proof of Work (PoW) (Nakamoto, 2008) | Nodes solve complex transaction validation problems, ensuring high security at the cost of significant energy use and low throughput | High security, but high energy consumption |
| Proof of Stake (PoS) (Uddin *et al.*, 2021) | Blocks are validated by nodes holding more cryptocurrency, reducing energy use compared to PoW | Greater energy efficiency |
| Delegated Proof of Stake (DPoS) (Senart, 2025) | Nodes vote for representatives to validate blocks, enhancing efficiency and scalability with partial centralization | Improved scalability, though with reduced decentralization |
| Proof of Humanity (PoH) (Arjomandi-Nezhad *et al.*, 2021) | Nodes with higher contributions are more likely to validate blocks, aligning economic input with validation | Encourages active participation |
| Practical Byzantine Fault Tolerance (PBFT) | Consensus requires two-thirds agreement, offering high efficiency but facing scalability issues due to message volume | High efficiency and throughput; difficulty scaling |
| Delegated Byzantine Fault Tolerance (DBFT) (Crain *et al.*, 2018) | Similar to PBFT but with selected nodes validating blocks, improving efficiency for private networks | Better efficiency for private networks, some centralization |
| Proof of Weight (PoWeight) (24) | Nodes with more resources have greater influence; Filecoin uses this model based on stored data | Resource-based consensus allocation |
| Proof of Burn (PoB) (Karantias *et al.*, 2020) | Miners burn cryptocurrency to validate blocks, reducing energy needs compared to PoW | Reduced energy consumption |
| Proof of Capacity (PoC) (Dziembowski *et al.*, 2015) | Miners store solutions on hard drives; greater storage increases block validation chances | Incentivizes storage over energy use |
| Proof of Importance (PoI) (Quesada-Real *et al.*, 2025) | NEM's PoI selects nodes based on activity, coin holdings and transactions, encouraging active use | Promotes network activity and participation |

**Source(s):** Authors' own work

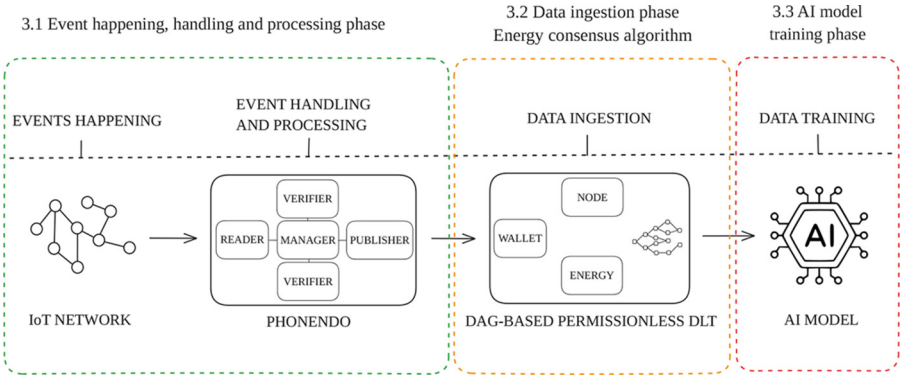


**Figure 1.** Main blocks of proposed architecture
**Source:** Authors' own work

available through the DLT facilitates AI model initialization and subsequent improvement. Below, we examine these components in detail.

### *3.1 Event happening, handling and processing phase*

In an IoT environment with numerous heterogeneous sensors, each type of sensor responds to events relevant to its function, for example, temperature sensors detect thermal changes, motion sensors register physical movement and humidity sensors react to moisture levels. These events are triggered under specific conditions suited to each sensor's capabilities. For the sake of clarity, to handle and process events, we use Phonendo (Moya *et al.*, 2023), a platform comprising various software services that manage the entire data lifecycle, from collecting data from wearable IoT devices to its publication on the DLT. One of the main advantages of using Phonendo is its architecture, which is built around independent modules with distinct responsibilities (Reader, Manager, Storage, Verifier and Publisher), which provides flexibility in enhancing or enriching particular modules. Figure 2 depicts the dataflow since the event occurs until the data is published in a DLT.

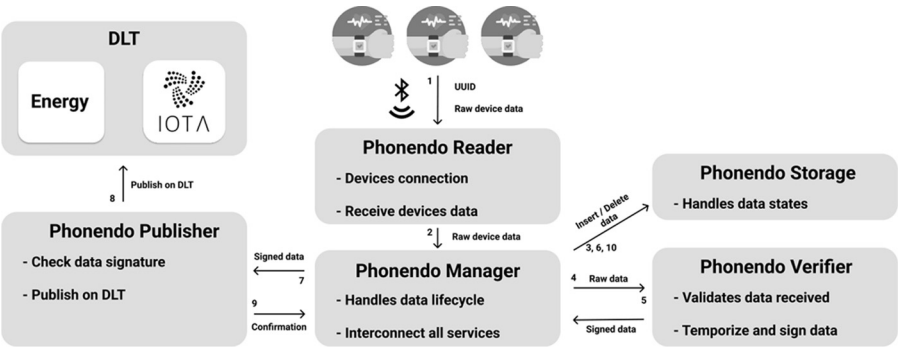


**Figure 2.** Phonendo adapted for multi-publishers
**Source:** Authors' own work

Initially, Phonendo's Reader service receives and securely transmits the event to the Manager service. The Manager service functions as an orchestration layer, overseeing the different states the event undergoes. Subsequently, the Verifier service validates and signs the data, as a way to act as a trusted authority. The Storage service records each state change to safeguard the system's overall state. Finally, the Publisher service publishes the data on the DLT. Given the modular nature of Phonendo's architecture, integration with different DLTs is straightforward, requiring only adaptation of the Publisher. In our case, we have tailored the Phonendo Publisher to publish on both IOTA and Energy. It would allow us to build multipublisher applications. This approach ensures comprehensive data management and verification from when the event occurs until it is published on the DLT.

### *3.2 Data ingestion phase – energy consensus algorithm*

Once the data is sent to the DLT, the next phase begins: data ingestion. As previously discussed, this phase faces several key challenges: (i) hardware with limited processing and energy capacity, often using batteries and (ii) ingestion latency and scalability on DLT ingesting a huge IoT data volume.

Regarding (i), limited computing power directly affects the PoW mechanism integrated into several DLT consensus algorithms or used as an anti-spam measure (Moya *et al.*, 2024). Studies have examined methods to reduce this impact (Chatrapathi *et al.*, 2023; Sravya *et al.*, 2023), but mitigation often involves either externalizing the PoW or adopting consensus algorithms that do not rely on PoW.

Regarding (ii), software applications in real-world environments are often limited by the volume of traffic they can handle, which depends on the software implementation and, more importantly, the scalability of the underlying infrastructure. Additionally, as the volume of incoming data grows, ingestion latency becomes an issue when the system reaches capacity, causing delays and compromising real-time functionality, affecting temporal data consistency. Mitigation could involve using DAG instead of blockchain, optimizing the consensus algorithm and reducing or eliminating PoW as an anti-spam measure.

We propose a DAG-based permissionless DLT solution with a two-tiered algorithm addressing the previously identified challenges. One approach tackles both issues by enabling the addition of transactions without PoW, using a fee system based on the amount of data being transmitted. This efficient fee system reassures the economic viability of the solution, reducing latency and increasing network scalability by streamlining the transaction validation and consensus process. The consensus algorithm distributes the fast-ingestion fee among nodes approving the transaction, incentivizing participation and contributing to a self-sustaining ecosystem.

The second approach allows the addition of transactions via PoW in situations where computational or energy resources are sufficient or when rapid data ingestion into the network is not required. Additionally, both approaches require each new transaction to approve two prior transactions inherent to the DAG structure. Unlike traditional DLTs, no dedicated miners are needed, being network participants who maintain the system.

Another critical point is how Energy generates the PoW associated with each transaction. The purpose of generating PoW in our proposal is to ensure that the computational effort required for validating a transaction is fair and proportional to the data being sent. This is achieved by requiring a specific number of leading zeros in the resulting hash, enforced through a `nonce` – a variable value appended to the transaction data to alter the hash output. The requirement for these leading zeros ensures that generating a valid hash involves substantial computational effort, helping prevent attackers from flooding the network with invalid transactions without incurring significant costs. Thus, the nonce and the number of

leading zeros together serve as control mechanisms, making the PoW process progressively harder based on data size and requiring computational resources that deter spam and double-spending attacks.

The PoW generation process can be mathematically described as follows. Given the `transactionData` string $T$, the PoW algorithm aims to find a `nonce` $n$ such that the SHA-256 hash function [National Institute of Standards and Technology (NIST), 2015] $H$ satisfies equation (1):

$$H(T \parallel n) < 2^{256-b} \tag{1}$$

$\parallel$ represents the concatenation of the transaction data $T$ and the `nonce` $n$, and $b$ is the number of bits that need to be zero in the hash prefix. The function $H$ denotes the SHA-256 hash function, which generates a 256-bit digest from any input.

To dynamically determine the required number of leading zeros based on the size of the transaction data, we calculate $b$ with equation (2):

$$b = 4 \times \left( \max \left( 1, log_2 \left( \frac{|T|}{10} \right) \right) + 1 \right) \tag{2}$$

$|T|$ is the length of the transaction data in characters, and $log_2$ represents the binary logarithm. This formula ensures that the difficulty of the PoW task increases logarithmically with the size of the transaction data, making it harder to find a valid `nonce` as the data size grows.

The expected number of attempts to find a valid `nonce` can be estimated with equation (3):

$$E(b) = 2^b \tag{3}$$

This means that, on average, $2^b$ attempts will be required to find a `nonce` $n$ that satisfies the condition for the PoW. The process of attempting different `nonce` values continues until the hash $H(T \| n)$ begins with the required number of zeros, as shown in Algorithm 1.

**Algorithm 1:** Proof of Work Generation Algorithm

```
Initialize nonce n ← 0;
Compute the target prefix length b;
while nonce < maxIterations do
    Generate the hash H(T ∥ n);
    if the hash starts with b leading zeros then
        return valid nonce n;
    Increment n;
return "Proof of work failed after max iterations";
```

The time complexity of this process is proportional to the number of leading zeros $b$, which directly affects the number of iterations required. As $b$ increases, the computational difficulty of the PoW increases exponentially, ensuring that the network remains secure even as transaction volumes grow (Nakamoto, 2008).

Given the importance of balancing computational complexity and network efficiency, it is crucial to ensure consistent transaction validation. To this end, we introduce a common validation process that standardizes the initial verification steps shared across different transaction ingestion approaches. This process establishes a consistent baseline for transaction validation, enabling specialized logic to be applied for either fast data ingestion or network stability, as detailed in Algorithm 2.

**Algorithm 2:** Common Transaction Validation

**Input:** Transaction data
**Output:** Boolean value indicating transaction validity

```
Procedure CommonValidation()
    Validate transaction sign: ;
    if signature verification using from address public key fails then
        return False;
    Validate 2 pending transactions: ;
    if 2 pending transactions are not validated then
        return False;
    Validate from address: ;
    if from address is not valid then
        return False;
    Validate to address: ;
    if to address is not valid then
        return False;
    return True;
```

Both previous Energy approaches are further detailed in the coming subsections.

*3.2.1 Fast ingestion approach.* The fast ingestion approach addresses the resource-limited hardware and latency challenges by allowing transactions to be added without relying on PoW, instead using a fee structure determined by the volume of data transmitted. This fee mechanism not only ensures the economic sustainability of the system but also minimizes latency and enhances network scalability by optimizing the transaction validation and consensus processes. Additionally, the consensus algorithm allocates the ingestion fees to nodes that approve transactions, encouraging node participation and fostering the development of a self-sustaining, decentralized ecosystem.

The sequence diagram in Figure 3 illustrates the fast data ingestion approach. As per the diagram, the *Wallet* software calculates a fee based on the transaction data

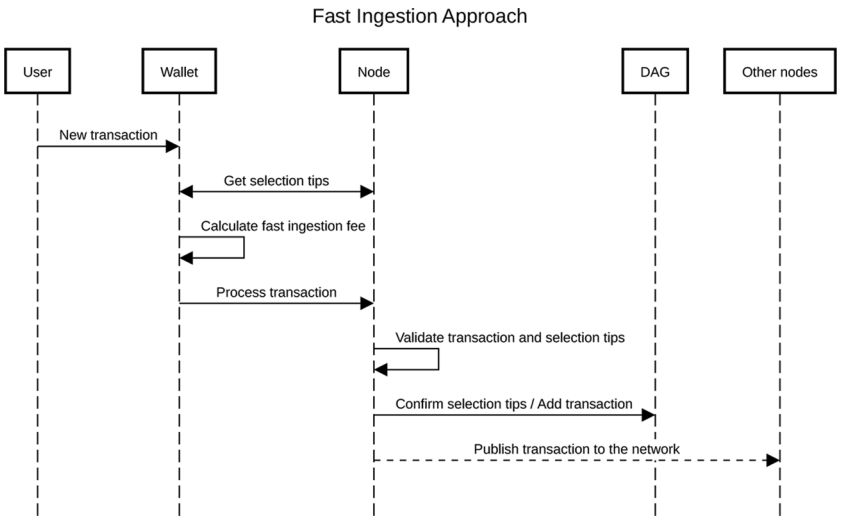


**Figure 3.** Energy – fast ingestion approach
**Source:** Authors' own work

when it receives a new transaction. The transaction is then sent to the *Node* software, where it undergoes verification and is added to the pending confirmation transaction list, also known as *selection tips*. Once the transaction is confirmed, it is sent to the *DAG* and communicated with the rest of the network via PubSub (Eugster *et al.*, 2003). Similar processes are followed by other nodes in the network that receive the transaction.

The response time of the consensus algorithm depends on the completion of specific validation steps unique to the fast data ingestion approach. Following the common validation process, Algorithm 3 outlines the additional steps: verifying that the transaction type is "Fast" and meets specific requirements, and ensuring that sufficient balance is available to cover associated fees and token transfers.

**Algorithm 3:** Transaction validation - Fast Data Ingestion Approach

```
Input: Transaction data
Output: Boolean value indicating transaction validity
Procedure ValidateFastTransaction()
    if CommonValidation() returns False then
        return False;
    Validate Fast Transaction Type: ;
    if transaction type is not "Fast" and does not meet transaction
     requirements then
        return False;
    Validate fee and token availability: ;
    if insufficient balance in from address to cover fee and transfer then
        return False;
    return True;
```

*3.2.2 Network stability approach.* The network stability approach is ideal when network security and robustness precede low-latency demands. It implements a strict PoW mechanism to prevent spam, requiring PoW for each transaction to enhance security. However, performing PoW on resource-constrained devices increases energy consumption and costs, as limited processing power results in longer computational times and higher energy use, raising electricity bills and maintenance expenses. Despite these challenges, this approach is attractive for organizations managing large data volumes, as it eliminates transaction fees, enabling seamless global data transfer across public and private sectors.

The network stability approach sequence diagram is detailed in Figure 4. According to this diagram, whenever the *Wallet* software receives a new transaction, it calculates the PoW required based on the transaction data. Then, the transaction is forwarded to the *Node* software, where it goes through a verification process and is added to the selection tips. Once the transaction is confirmed, it is sent to the *DAG* and shared with the rest of the network through PubSub. Other nodes in the network that also receive the transaction follow similar processes.

Network stability is critical for systems handling millions of requests, including potential spam attacks. To face it, Algorithm 4 builds on the common validation process and outlines additional steps specific to the network stability approach: reviewing the transaction type, associated requirements and ensuring the correctness of the provided PoW; and verifying that sufficient balance exists to cover the token transfer.

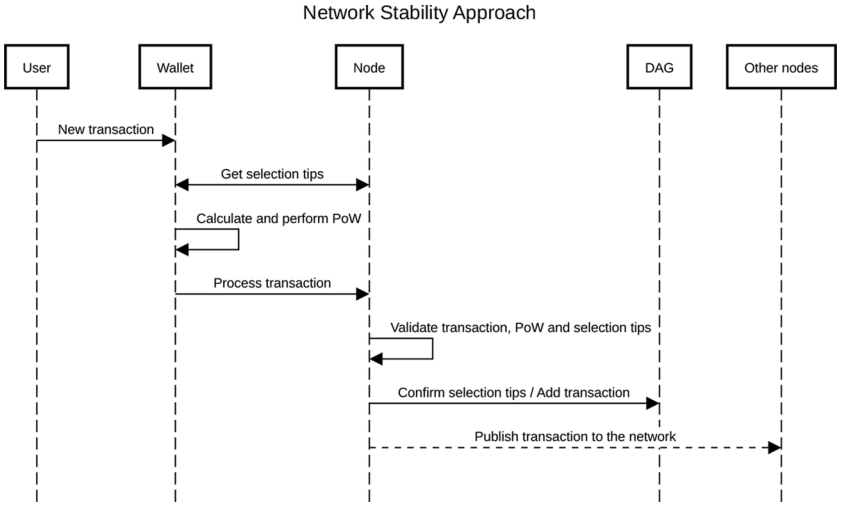


**Figure 4.** Energy – network stability approach
**Source:** Authors' own work

**Algorithm 4:** Transaction validation - Network Stability Approach

**Input:** Transaction data
**Output:** Boolean value indicating transaction validity

```
Procedure ValidateNetworkTransaction()
    if CommonValidation() returns False then
        return False;
    Validate Standard Transaction Type: ;
    if transaction type is not "Standard" and does not meet transaction
     requirements then
        return False;
    Validate PoW: ;
    if provided PoW is not correct or is missing then
        return False;
    Validate token availability: ;
    if insufficient balance in from address to cover transfer then
        return False;
    return True;
```

### *3.3 Artificial intelligence model training phase*

Integrating Energy into an AI model learning process is feasible by leveraging all existing data from an IoT network stored in the DLT at a specific time. The integration of Energy enables AI models to update continuously as new transactions are generated within the IoT network. This process can be optimized by incorporating real-time data and strengthening machine learning with reliable and verified information at every update. Instead of conducting intermittent training on large data batches, Energy facilitates the creation of continuous learning models, where each approved transaction in the network serves as a new source of information to enhance model predictions. This approach eliminates the need to rebuild or retrain models from scratch whenever new data becomes available.

Incremental learning algorithms are ideal for these applications (Xu *et al.*, 2022), adjusting parameters in real-time and effectively managing computational load. These algorithms

enable rapid adaptation to data changes, crucial for applications in dynamic environments such as energy management or industrial maintenance.

Sliding window techniques (Larimer *et al.*, 2017) complement this adaptability, allowing models to update by integrating recent transactions and discarding old data continuously. This is particularly useful in contexts where data integrity is critical, such as in predictive maintenance. Recurrent neural networks (Fu *et al.*, 2016), like long short-term memory or gated recurrent units, optimized for data sequences, can be continuously trained to adjust to temporal patterns in IoT data, providing accurate and adaptive predictions under the data security that Energy guarantees.

For instance, in an industrial maintenance prediction environment, every new transaction reporting the status of a specific sensor can immediately feed the AI model. Thus, model predictions adjust dynamically without disrupting their operation. Additionally, the trust in data provided by Energy's consensus ensures that the model is only fed with valid and accurate information, enhancing its effectiveness.

This system fosters a positive feedback loop, where the AI model adapts in real-time to new conditions or events detected by the IoT network. With each transaction validated in the network, the model continues to evolve without compromising the security or integrity of the data used for continuous improvement.

### *3.4 Implementation*

A consensus algorithm's real-world implementation and evaluation are crucial to assessing its potential. This section outlines the practical deployment of the proposed DAG-based, permissionless DLT consensus algorithm, focusing on the implementation details and performance evaluation under various conditions. Energy code is publicly accessible at Github [1].

*3.4.1 Software infrastructure.* The algorithm was implemented in a fully developed DLT system, written in Go (go.dev. Golang), chosen for its concurrency, cross-platform compatibility and high performance. The architecture is microservices-based, as illustrated in Figure 5, enabling modular scalability and real-time communication via lightweight APIs and asynchronous messaging. Its decentralized design optimizes performance and resilience, dynamically adapting to demand. Users interact through digital wallets, connecting to nodes that verify and broadcast transactions in a publish/subscribe system. The DAG structure supports concurrent transactions, enhancing throughput and reducing latency while ensuring ledger integrity.

*3.4.2 Network integration.* The network uses *P2P PubSub* and *REST* protocols, leveraging the LibP2P framework (libp2p.io. Libp2p) for peer discovery and communication. Two main topics are defined for node synchronization: *wallet-create-topic* for managing wallet creation and *new-transaction-topic* for broadcasting transactions. The communication layer is built on REST APIs, defined via OpenAPI. *The Wallet API* [2] manages wallet creation, transaction handling and node interactions, while the *Node API* [3] oversees node operations, including transaction validation and processing.

## 4. Evaluation

This section evaluates the Energy solution in terms of performance, scalability and energy efficiency. Key metrics – including transaction latency, consensus accuracy and energy consumption – are assessed through parameterized tests and comparative analyses. The goal is to validate the system's ability to support high transaction volumes with minimal latency and energy usage, particularly in IoT and large-scale infrastructure scenarios. The next subsection compares Energy with IOTA, focusing on processing time and energy

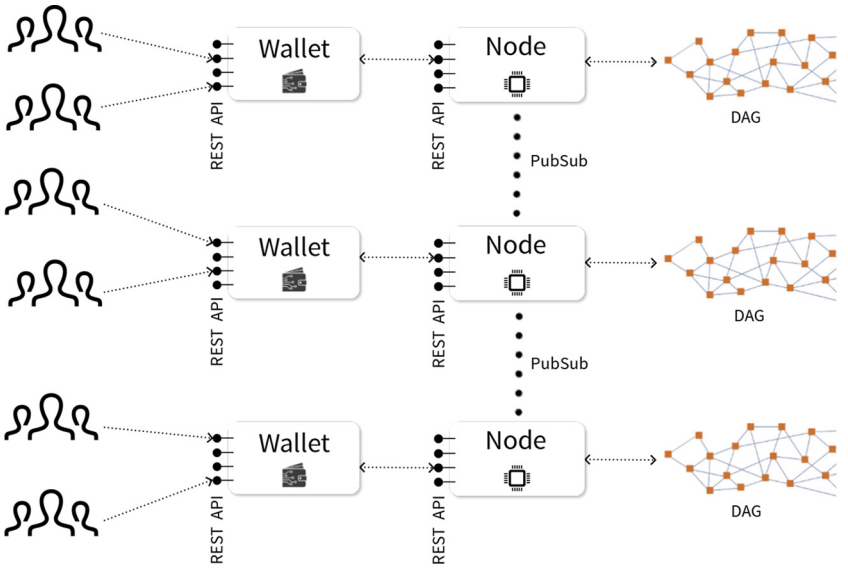


**Figure 5.** Software architecture diagram
**Source:** Authors' own work

consumption under different conditions, demonstrating Energy's potential as a more efficient and sustainable alternative for real-world applications.

*4.1 Evaluating latency and consensus in the proposed solution*

Once the proposed solution has been developed and integrated, it is crucial to validate its proper functionality and the correctness of our hypothesis. To achieve this, we conducted a parameterized test by sending numerous random transactions to the network. The goal was two-fold: first, to verify the system's latency and scalability; and second, to ensure that all transactions were successfully verified, with the consensus algorithm operating correctly.

Initially, a Bash script (gnu.org. Gnu bash) was developed to test the scalability and resilience of the transaction ingestion in the proposed Energy solution. The script, accessible via this link [4], progressively sends requests to the Wallet software through the Energy instance. Figure 6 outlines a scenario generating 1,000,000 transactions with a size of 300 bytes, which three active Energy nodes process. The nodes are preinitialized with a minimal number of transactions to ensure validation during the test. The open-source tool Vegeta Load testing (Sengupta *et al.*, 2020) orchestrates the test according to specific criteria.

Transactions are generated synthetically using a Python script [5]. It produces unique transactions on disk and randomly assigns them to one of the three active nodes, ensuring an unbiased workload distribution. The script also determines the transaction type (fast ingestion or network stability), selects a random valid destination address, the number of tokens and the data to be transferred. Generated transactions can be found in the following link [6].

The test results are shown in Tables 2 and 3, providing insight into transaction processing times. Table 2 outlines the time percentiles (P50, P90, P95 and P99) related to transaction ingestion latency when a transaction is submitted through a node, reflecting how quickly the

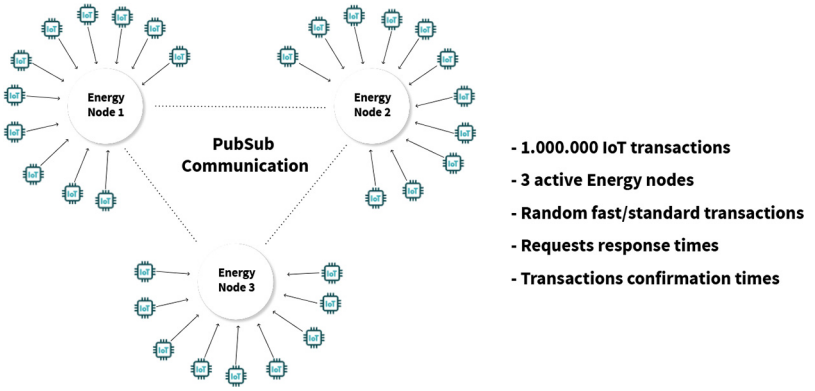


**Figure 6.** Energy load test
**Source:** Authors' own work

**Table 2.** Add a transaction to the network

| Nodes | Transactions | P50 | P90 | P95 | P99 |
|---|---|---|---|---|---|
| Node 1 | 337,504 | 0.34 ms | 2.01 ms | 2.75 ms | 34.22 ms |
| Node 2 | 328,470 | 0.35 ms | 2.04 ms | 2.78 ms | 38.21 ms |
| Node 3 | 334,404 | 0.35 ms | 2.03 ms | 2.77 ms | 37.27 ms |

**Source(s):** Authors' own work

**Table 3.** Transaction confirmation

| Nodes | Transactions | P50 | P90 | P95 | P99 |
|---|---|---|---|---|---|
| Node 1 | 337,504 | 33.34 ms | 139.45 ms | 233.19 ms | 2254.38 ms |
| Node 2 | 328,470 | 33.54 ms | 142.16 | 242.91 ms | 3649.54 ms |
| Node 3 | 334,404 | 33.55 ms | 141.79 ms | 241.96 ms | 4437.69 ms |

**Source(s):** Authors' own work

network processes new transactions. Table 3 shows the time percentiles for transaction confirmation after validation by subsequent transactions in the DAG. Even in the worst-case scenario, confirmation times remain under 4.5 s, while the 95th percentile shows confirmation times of less than 250 ms. Additionally, the DAGs of the three nodes tested can be viewed through the following link [7].

The experimental setup was designed to simulate a realistic scenario where traffic flows are distributed across multiple nodes. By randomly assigning transactions to the three active nodes, we ensured that the workload was balanced and that no single node was overloaded. This approach allowed us to evaluate the system's performance under conditions that mimic real-world deployments.

The scalability of the Energy solution is inherently supported by its DAG-based architecture (Zaeem and Barber, 2020). Unlike traditional blockchain systems, where sequential block

validation can create bottlenecks, the DAG structure allows for the parallel processing of transactions. This parallelism ensures that as the number of nodes increases, the system can handle higher transaction volumes without significant increases in latency. These architectural features make the system suitable for large-scale deployments in real-world scenarios, such as IoT networks or public infrastructure systems. By reducing transaction latency in DAG-based DLTs, Energy significantly enhances network scalability. This improvement not only demonstrates the algorithm's own scalability but also contributes to a substantial increase in the scalability of the underlying DLT.

*4.2 IOTA and energy processing time and energy consumption comparative*

In this subsection, we performed a time and energy cost analysis, comparing our Energy solution to IOTA. IOTA was selected as the benchmark for comparison because our baseline case should perform similarly. Due to its implementation is more refined, we can obtain more precise references for the required energy consumption than if we relied on our implementation. Unlike other DLTs, IOTA provides a lightweight architecture optimized for IoT, enabling low-power devices to participate in the network through techniques such as Masked Authenticated Messaging [8] and Streams [9]. Its extensive documentation and community support allowed us to use robust implementations, ensuring that our comparison with Energy was fair and reliable.

The efficiency of the Energy consensus algorithm compared to IOTA can be mathematically explained in terms of latency and energy consumption, considering both the payload size and the PoW difficulty. Each system uses a unique hashing algorithm and approach to difficulty adjustment, directly affecting their time and resource efficiency performance, particularly when managing small and large payloads.

IOTA uses the Curl-P-81 algorithm [10], a specific trinary hash function designed by IOTA for its transaction trinary structure. The minimum weight magnitude (MWM) serves as a baseline difficulty level, requiring that each hash meets a minimum of 11 leading zeros on mainnet. Additionally, IOTA imposes a maximum transaction payload size of 32 KB. Any payload exceeding 1.6 KB is split into multiple fragments of 1.6 KB, each fragment undergoing independent PoW with the baseline MWM difficulty.

The PoW latency in IOTA, therefore, increases linearly with payload size. For a payload of size $x$ (in KB), the total PoW time can be expressed with equation (4):

$$T_{\mathrm{IOTA}}(x) = T_{\mathrm{hashbase}} \times D(x) \tag{4}$$

where

- $T_{\mathrm{hashbase}}$ represents the base hash time.
- Difficulty is determined as $D(x) = 11 \times \frac{x}{1.6}$,

Consequently, IOTA energy consumption is directly proportional to this PoW time using equation (5):

$$E_{\mathrm{IOTA}}(x) = T_{\mathrm{IOTA}}(x) \times C_{\mathrm{energyfactor}} \tag{5}$$

where $C_{\mathrm{energyfactor}}$ is a constant representing the energy per unit of time.

On the other hand, Energy relies on SHA-256 for its PoW, with difficulty scaling logarithmically with payload size. Unlike IOTA, Energy's difficulty does not have a fixed upper limit or fragmentation requirement. Instead, the difficulty begins at a minimum of 1 zero and increases according to a logarithmic function using equation (6):

$$b = 4 \times \left( max\left(1, log_2\left(\frac{x \text{ x } 1024}{10}\right)\right) + 1 \right) \quad (6)$$

where $b$ denotes the required number of leading zeros, and $x$ is the payload size in KB.

This formula ensures that the difficulty of the PoW task increases logarithmically with the payload size. The constants in the formula were determined through experimental evaluation to balance computational cost, scalability and security:

- The constant 4 is a scaling factor that adjusts the overall difficulty to ensure a reasonable computational effort.
- The divisor 10 normalizes the payload size to avoid excessively high difficulty for small payloads.
- The $max\left(1, log_2(\ldots)\right)$ ensures a minimum difficulty of 1 zero, even for tiny payloads.

The PoW time in Energy thus scales as shown in equation (7):

$$T_{\text{Energy}}(x) = T_{\text{hashbase}} \text{ x } 2^b \quad (7)$$

The energy consumption for Energy is also proportional to the PoW time using equation (8):

$$E_{\text{Energy}}(x) = T_{\text{Energy}}(x) \text{ x } C_{\text{energyfactor}} \quad (8)$$

This logarithmic scaling is mainly designed to enhance efficiency for small payloads, commonly encountered in IoT, while remaining adaptable for larger payloads. Furthermore, the energy consumption of Energy transactions could be reduced using the fast transaction approach provided by the consensus algorithm.

Using different payload sizes, we sent 50 transactions through the IOTA network using a Raspberry Pi 3 B with a 16 GB Class 10 SD card. This setup emulates resource-constrained environments, such as those typical in IoT. The experiment revealed an average latency of 103.25 s per transaction and an energy consumption of 0.1434 W per transaction. These results highlight the resource-intensive nature of IOTA's PoW, particularly in low-power devices. Table 4 presents the processing times in seconds for both IOTA and Energy across varying levels of PoW difficulty, for payloads ranging from 0.3 KB to 3.5 KB, simulating typical data transmission scenarios in IoT environments. The corresponding results are visually depicted in Figure 7.

Similarly, Table 5 provides the energy consumption in watts for both systems under the same PoW levels and payload sizes. These energy results are represented graphically in Figure 8.

**Table 4.** Comparative energy and IOTA processing times for small payloads

| Size (KB) | IOTA proc. time (s) | Energy proc. time (100%) | Energy proc. time (70%) | Energy proc. time (50%) | Energy proc. time (30%) |
|---|---|---|---|---|---|
| 0.3 | 16.73 | 3.93 | 2.75 | 1.97 | 1.18 |
| 1.0 | 25.81 | 7.44 | 5.20 | 3.72 | 2.23 |
| 1.8 | 34.88 | 10.61 | 7.43 | 5.31 | 3.18 |
| 2.5 | 43.96 | 13.12 | 9.18 | 6.56 | 3.94 |
| 3.5 | 53.03 | 15.63 | 10.94 | 7.81 | 4.69 |

**Source(s):** Authors' own work

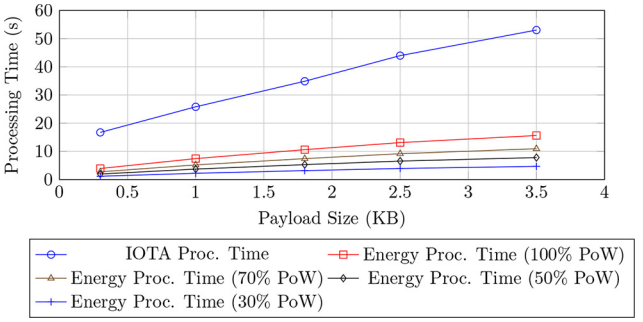


**Figure 7.** Comparative energy and IOTA processing time for small payloads
**Source:** Authors' own work

**Table 5.** Comparative energy and IOTA energy consumption for small payloads

| Size (KB) | IOTA E. consump. (W) | Energy E. consump. (100%) | Energy E. consump. (70%) | Energy E. consump. (50%) | Energy E. consump. (30%) |
|---|---|---|---|---|---|
| 0.3 | 0.0232 | 0.0054 | 0.0038 | 0.0027 | 0.0016 |
| 1.0 | 0.0358 | 0.0103 | 0.0072 | 0.0051 | 0.0031 |
| 1.8 | 0.0484 | 0.0147 | 0.0103 | 0.0074 | 0.0044 |
| 2.5 | 0.0611 | 0.0182 | 0.0127 | 0.0091 | 0.0055 |
| 3.5 | 0.0737 | 0.0217 | 0.0152 | 0.0109 | 0.0065 |

**Source(s):** Authors' own work

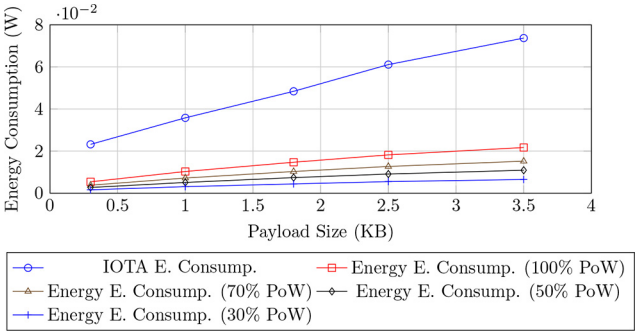


**Figure 8.** Comparative energy and IOTA energy consumption for small payloads
**Source:** Authors' own work

Following, we will continue analyzing large payloads, specifically 150 KB, 200 KB, 500 KB and 800 KB, using both IOTA and Energy systems across different levels of PoW. These larger payloads are typical of more substantial data transfers in IoT environments, where resource constraints are still a factor, but throughput demands increase. Due to IOTA's fixed MWM difficulty in PoW execution, each 1.6 KB segment must undergo independent PoW, leading to higher latency and energy consumption as payload size increases. In contrast,

Energy dynamically adjusts PoW difficulty with payload size and offers a fee-based option to bypass PoW, significantly reducing processing times and energy use, especially as the proportion of fee-based transactions increases. Table 6 displays the processing times for IOTA and Energy across different PoW levels for large payloads. Figure 9 visually compares these results, showing how Energy reduces processing times significantly with increased fee-based transactions.

Table 7 presents the energy consumption in watts for both IOTA and Energy with varying PoW levels for large payloads. Figure 10 illustrates the efficiency gains in Energy's energy use as the proportion of transactions with PoW decreases.

*4.2.1 Case study: real-time monitoring and optimization of a public railway network.* A country's railway network is a critical public infrastructure, connecting cities, towns and

**Table 6.** Comparative energy and IOTA processing times for large payloads

| Size (KB) | IOTA proc. time (s) | Energy proc. time (100%) | Energy proc. time (70%) | Energy proc. time (50%) | Energy proc. time (30%) |
|---|---|---|---|---|---|
| 150 | 484.86 | 173.87 | 121.71 | 86.93 | 52.16 |
| 200 | 645.81 | 206.38 | 144.47 | 103.19 | 61.91 |
| 500 | 1,614.53 | 443.22 | 310.25 | 221.61 | 132.97 |
| 800 | 2,583.25 | 602.75 | 421.93 | 301.37 | 180.82 |

**Source(s):** Authors' own work

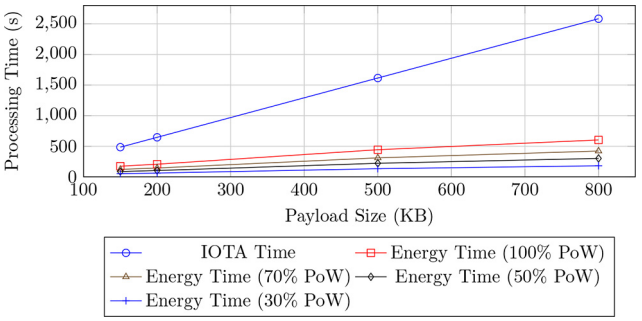


**Figure 9.** Comparative energy and IOTA processing times for large payloads
**Source:** Authors' own work

**Table 7.** Comparative energy and IOTA energy consumption for large payloads

| Size (KB) | IOTA E. consump. (W) | Energy E. consump. (100%) | Energy E. consump. (70%) | Energy E. consump. (50%) | Energy E. consump. (30%) |
|---|---|---|---|---|---|
| 150 | 0.6739 | 0.2413 | 0.1689 | 0.1206 | 0.0724 |
| 200 | 0.8997 | 0.2860 | 0.2002 | 0.1430 | 0.0858 |
| 500 | 2.2741 | 0.6156 | 0.4309 | 0.3078 | 0.1847 |
| 800 | 3.6484 | 0.8378 | 0.5865 | 0.4189 | 0.2513 |

**Source(s):** Authors' own work

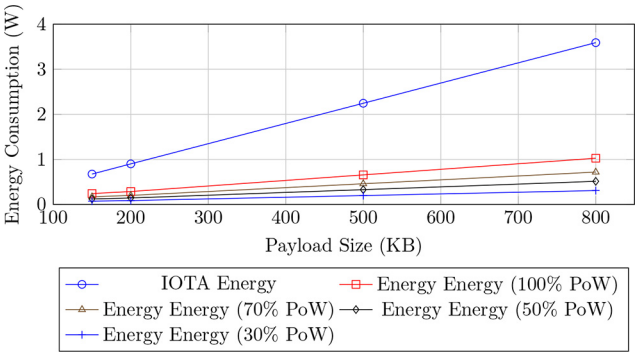


**Figure 10.** Comparative energy and IOTA energy consumption for large payloads
**Source:** Authors' own work

regions while supporting millions of passengers and freight operations daily. This network consists of tracks, signals, traffic lights, electrical cabling and stations, managed by a government ministry. Meanwhile, public and private companies operate the trains that traverse this infrastructure, creating a complex ecosystem of stakeholders.

For this case study, we envision a future where the railway network is equipped with IoT sensors to monitor infrastructure and train movement in real time. The data is recorded in a DAG-based DLT for transparency and traceability, while AI models analyze the information to optimize operations and maintenance. We assume the following parameters:

- Each IoT sensor sends a 300-byte update every minute for 24 h. With 10,000 sensors deployed across the network, this results in 14,400,000 transactions daily.
- Each train sends a 300-byte location update every 10 s during an average 2-h operation period. With 500 trains in operation, this results in 18,000,000 transactions daily.

In total, the system processes approximately 32,400,000 transactions daily. Assuming each 300-byte transaction in IOTA consumes approximately 0.0232 W, the daily energy cost for IOTA would be:

$$32,400,000 \text{ transactions} \times 0.0232 \text{ W} = 750,880 \text{ Wh} \, (750.88 \text{ kWh}).$$

With Energy set to 100% PoW, each 300-byte transaction consumes approximately 0.0054 W, resulting in:

$$32,400,000 \text{ transactions} \times 0.0054 \text{ W} = 174,960 \text{ Wh} \, (174.96 \text{ kWh}).$$

Assuming an average cost of €0.15 per kWh, the daily and annual costs would be as follows:

- *IOTA*: Daily cost of €112.63; annual cost of €41,108.95.
- *Energy (100% PoW)*: Daily cost of €26.24; annual cost of €9,577.60.

This results in an estimated daily savings of €86.39 and an annual savings of €31,531.35, representing a 76.7% cost reduction with Energy compared to IOTA under similar conditions. These savings demonstrate the economic viability of the proposed system, particularly in scenarios requiring large-scale, real-time data processing.

This case study highlights how IoT, DLT and AI can transform a public railway network into an intelligent, transparent and efficient system. By addressing challenges such as maintenance, coordination and transparency, the proposed solution ensures that the infrastructure remains safe, sustainable and accessible to all.

*4.2.2 Case study: future real-time traffic control in Madrid.* Madrid has approximately 56,709 traffic lights [11] and an average daily traffic of 10.5 million vehicles [12]. For this case study, we envision a future where traffic lights are equipped with low-capacity IoT hardware while vehicles are equipped with smart capabilities to share location data. This setup illustrates the impact of deploying a real-time traffic control system using our proposal with a DAG-based DLT for transparency and traceability, and an AI model for predictions and traffic control system adjustments. We assume the following parameters:

- Each traffic light sends a 300-byte update every minute for 24 h. One thousand four hundred forty updates per day per light turn in 81,658,560 transactions daily.
- Each vehicle sends a 300-byte location update every 10 s during an average 2-h period of daily driving. Seven hundred twenty updates per day per vehicle turn in 7,560,000,000 transactions daily.

In total, this system would process approximately 7,641,658,560 transactions daily. Assuming each 300-byte transaction in IOTA consumes approximately 0.0232 W, the daily energy cost for IOTA would be: 7,641,658,560 transactions $x$ 0.0232 $W = 177,295.47\,kWh$.

With Energy set to 100% PoW, each 300-byte transaction consumes approximately 0.0054 W, resulting in: 7,641,658,560 transactions $x$ 0.0054 $W = 41,264.96\,kWh$.

Assuming an average cost of €0.15 per kWh, the daily and annual costs would be as follows:

- *IOTA*: Daily cost of €26,594.32; annual cost of €9,705,946.80.
- *Energy (100% PoW)*: Daily cost of €6,189.74; annual cost of €2,259,267.10.

This results in an estimated daily savings of €20,404.58 and an annual savings of €7,446,679.70, representing a 76.7% cost reduction with Energy compared to IOTA under similar conditions. This case study suggests that Energy performs effectively in scenarios requiring energy and cost efficiency, particularly when using equal PoW conditions. While these results highlight Energy's advantages, its efficiency could potentially be enhanced further through optional fee-based transactions. By allowing specific transactions to bypass PoW, Energy can further reduce latency and energy consumption. However, it is important to balance the introduction of fee-based transactions with potential additional costs to maintain system cost-effectiveness. This flexibility positions Energy as a potentially adaptable, sustainable and economically viable solution for large-scale, real-time urban traffic management within IoT-enabled environments.

## 5. Conclusions

This work presents *Energy*, a novel consensus algorithm optimized for IoT environments through a DAG-based DLT architecture that minimizes latency, enhances scalability and reduces energy consumption. By allowing the bypassing of PoW via fee-based transactions and dynamically adjusting PoW difficulty based on payload size, *Energy* enables efficient transaction processing in resource-constrained settings. This mechanism facilitates seamless integration with AI systems that require continuous access to secure, traceable and real-time data, thereby supporting applications in which AI must interpret and respond rapidly to dynamic conditions.

Preliminary experimental results indicate that *Energy* performs effectively with payloads under 32 KB, achieving notable reductions in latency and energy consumption through adaptive PoW. Compared to IOTA, *Energy* demonstrates substantial improvements in both efficiency and responsiveness. Its fast-ingestion mechanism, enabled by optional transaction fees, further reduces processing time, reinforcing its suitability for demanding IoT scenarios where AI models rely on timely and validated data streams.

Future research will focus on extending the capabilities of *Energy* through the development of AI models trained on both historical and real-time transactions. These models will aim to detect SPAM attacks and enhance the overall stability and performance of the network. Additionally, we plan to explore AI-driven strategies for dynamically regulating PoW requirements according to current network conditions and device capacities, further improving energy efficiency and responsiveness.

To conduct a more detailed validation of the scalability of the proposed solution, forthcoming experiments will be carried out in larger and more complex environments. These will include performance evaluations across networks of varying sizes – specifically with 10, 50 and 100 nodes – and simulations of real-world scenarios such as smart cities and national railway systems. Stress testing under extreme workloads will be used to identify potential bottlenecks and guide system optimization. Furthermore, collaborations with existing IoT projects will enable real-world deployments of Energy and the collection of empirical data on its scalability and performance.

These efforts aim to provide a comprehensive understanding of Energy's behavior in large-scale deployments and to confirm its suitability for supporting real-time, AI-driven applications in diverse, high-demand environments.

**Notes**

1. https://github.com/sinbad2-ujaen/energy
2. https://github.com/sinbad2-ujaen/energy/blob/main/wallet/wallet-openapi.yaml
3. https://github.com/sinbad2-ujaen/energy/blob/main/node/node-openapi.yaml
4. https://github.com/sinbad2-ujaen/energy/blob/main/loadtest/energy.sh
5. https://github.com/sinbad2-ujaen/energy/blob/main/loadtest/generate_transactions.py
6. https://drive.google.com/file/d/1tj-_aW2gsD57YEGXd_trfiF7fVCuqg-v/view?usp=drive_link
7. https://drive.google.com/drive/folders/1d1R4xneS281hJvCTCpxnzPNMj1cf_vjT?usp=drive_link
8. https://github.com/iotaledger/MAM
9. https://github.com/iotaledger/streams
10. https://wiki.iota.org/tips/tips/TIP-0012/
11. https://datos.madrid.es/portal/site/egob/menuitem.c05c1f754a33a9fbe4b2e4b284f1a5a0/?vgnextoid=1e69ee425d7c6410VgnVCM2000000c205a0aRCRD&vgnextchannel=374512b9ace9f310VgnVCM100000171f5a0aRCRD&vgnextfmt=default
12. www.comunidad.madrid/media/transportes/dossier2023.pdf

**References**

Abbas, A., Alroobaea, R., Krichen, M., Rubaiee, S., Vimal, S. and Almansour, F.M. (2021), "Blockchain-assisted secured data management framework for health information analysis based on internet of medical things", *Personal and Ubiquitous Computing*, Vol. 28 No. 1, pp. 1-14.

Abdallah, M., Dobre, O.A., Ho, P., Jabbar, S., Khabbaz, M.J. and Rodrigues, J. (2020), "Blockchain-enabled industrial Internet of Things: advances, applications, and challenges", *IEEE Internet of Things Magazine*, Vol. 3 No. 2, pp. 16-18.

Aldoseri, A., Al-Khalifa, K.N. and Hamouda, A.M. (2023), "Re-thinking data strategy and integration for artificial intelligence: concepts, opportunities, and challenges", *Applied Sciences*, Vol. 13 No. 12, p. 7082.

Alsaadi, E. and Tubaishat, A. (2015), "Internet of Things: features, challenges, and vulnerabilities", *International Journal of Advanced Computer Science and Information Technology*, Vol. 4 No. 1, pp. 1-13.

Arjomandi-Nezhad, A., Fotuhi-Firuzabad, M., Dorri, A. and Dehghanian, P. (2021), "Proof of humanity: a tax-aware society-centric consensus algorithm for blockchains", *Peer-to-Peer Networking and Applications*, Vol. 14 No. 6, pp. 3634-3646.

Bail, R., Kovaleski, J.L., da Silva, V.L., Pagani, R.N. and Chiroli, D.M. (2021), "Internet of Things in disaster management: Technologies and uses", *Environmental Hazards*, Vol. 20 No. 5, pp. 493-513.

Banafa, A. (2016), "IoT standardization and implementation challenges", *IEEE Internet of Things Newsletter*, Vol. 2016, pp. 1-10.

Chatrapathi, R., Ramkumar, M., Jayakumar, D. and Pooja, R. (2023), "Reinforcing IoT security through machine learning based spam detection", *2023 2nd International Conference on Smart Technologies and Systems for Next Generation Computing (ICSTSN)*, pp. 1-6.

Crain, T., Gramoli, V., Larrea, M. and Raynal, M. (2018), "Dbft: Efficient leaderless byzantine consensus and its application to blockchains", *2018 IEEE 17th International Symposium on Network Computing and Applications (NCA)*, pp. 1-8.

Dziembowski, S., Faust, S., Kolmogorov, V. and Pietrzak, K. (2015), "Proofs of space", *IACR Cryptol. ePrint Arch*, Vol. 2013, p. 796.

Eugster, P.T., Felber, P.A., Guerraoui, R. and Kermarrec, A. (2003), "The many faces of publish/subscribe", *ACM Computing Surveys*, Vol. 35 No. 2, pp. 114-131.

Farahani, B., Firouzi, F. and Luecking, M. (2021), "The convergence of IoT and distributed ledger technologies (DLT): opportunities, challenges, and solutions", *Journal of Network and Computer Applications*, Vol. 177, p. 102936.

Fu, R., Zhang, Z. and Li, L. (2016), "Using LSTM and GRU neural network methods for traffic flow prediction", *2016 31st Youth Academic Annual Conference of Chinese Association of Automation (YAC)*, pp. 324-328. *IEEE*.

Gilbert, S. and Lynch, N. (2002), "Brewer's conjecture and the feasibility of consistent, available, partition-tolerant web services", *ACM SIGACT News*, Vol. 33 No. 2, pp. 51-59.

Goldreich, O. (2001), *Foundations of Cryptography: Volume 2, Basic Applications*, Vol. 2. Cambridge University Press.

Golpîra, H., Khan, S.A.R. and Safaeipour, S. (2021), "A review of logistics internet-of-things: current trends and scope for future research", *Journal of Industrial Information Integration*, Vol. 22, p. 100194.

Hajjaji, Y., Boulila, W., Farah, I.R., Romdhani, I. and Hussain, A. (2021), "Big data and IoT-based applications in smart environments: a systematic review", *Computer Science Review*, Vol. 39, p. 100318.

Harshvardhan, G., Mahendra, K.G., Manjusha, P. and Siddharth, S.R. (2020), "A comprehensive survey and analysis of generative models in machine learning", *Computer Science Review*, Vol. 38, p. 100285.

Jiang, F., Jiang, Y., Zhi, H., Dong, Y., Li, H., Ma, S., Wang, Y., Dong, Q., Shen, H. and Wang, Y. (2017), "Artificial intelligence in healthcare: past, present and future", *Stroke and Vascular Neurology*, Vol. 2 No. 4.

Karantias, K., Kiayias, A. and Zindros, D. (2020), "Proof-of-burn", Financial Cryptography and Data Security: 24th International Conference, FC 2020, Kota Kinabalu, Malaysia, February 10-14, 2020 Revised Selected Papers 24, pp. 523-540, *Springer*.

Khan, D., Jung, L.T. and Hashmani, M.A. (2021), "Systematic literature review of challenges in blockchain scalability", *Applied Sciences*, Vol. 11 No. 20.

Lamport, L., Shostak, R. and Pease, M. (1982), "The Byzantine generals problem", *ACM Transactions on Programming Languages and Systems*, Vol. 4 No. 3, pp. 382-401.

Lamport, L. (1977), "Proving the correctness of multiprocess programs", *IEEE Transactions on Software Engineering*, Vol. SE-3 No. 2, pp. 125-143.

Larimer, D., Hoskinson, C. and Larimer, S. (2017), "Bitshares: a peer-to-peer polymorphic digital asset exchange", *BitShares*, available at: BitShares%20A%20Peer-to-Peer%20Polymorphic%20Digital%20Asset%20Exchange.pdf

Liu, M., Yu, F., Teng, Y., Leung, V.C.M. and Song, M. (2018), "Computation offloading and content caching in wireless blockchain networks with mobile edge computing", *IEEE Transactions on Vehicular Technology*, Vol. 67 No. 11, pp. 11008-11021.

Makhdoom, I., Abolhasan, M., Abbas, H. and Ni, W. (2019), "Blockchain's adoption in IoT: the challenges, and a way forward", *Journal of Network and Computer Applications*, Vol. 125, pp. 251-279.

Makhdoom, I., Abolhasan, M., Lipman, J., Liu, R.P. and Ni, W. (2018), "Anatomy of threats to the Internet of Things", *IEEE Communications Surveys and Tutorials*, Vol. 21 No. 2, pp. 1636-1675.

Mattioli, J., Robic, P. and Jesson, E. (2022), "Information quality: the cornerstone for ai-based industry 4.0", Procedia Computer Science, 201:453-460, 2022. The 13th International Conference on Ambient Systems, Networks and Technologies (ANT) / The 5th International Conference on Emerging Data and Industry 4.0 (EDI40).

Mehta, K., Liu, D. and Wright, M. (2011), "Protecting location privacy in sensor networks against a global eavesdropper", *IEEE Transactions on Mobile Computing*, Vol. 11 No. 2, pp. 320-336.

Monte, G.D., Pennino, D. and Pizzonia, M. (2020), "Scaling blockchains without giving up decentralization and security: a solution to the blockchain scalability trilemma", *Proceedings of the 3rd Workshop on Cryptocurrencies and Blockchains for Distributed Systems*.

Moya, F., Quesada, F.J., Martínez, L. and Estrella, F.J. (2024), "Aspmi: an adaptable spam protection mechanism for IoT scenarios", In International Conference on Ubiquitous Computing and Ambient Intelligence, *Springer*, pp. 909-919.

Moya, F., Quesada, F.J., Martínez, L. and Estrella, F.J. (2023), "Phonendo: a platform for publishing wearable data on distributed ledger technologies", *Wireless Networks*, Vol. 30 No. 7.

Muñoz-Higueras, C., Serradilla-Gil, A.M., Moreno-Colmenero, P. and Quesada-Real, F.J. (2024), "Integrating IoT and DLT to enhance patient wait time traceability in radiotherapy oncology", *International Conference on Ubiquitous Computing and Ambient Intelligence*, *Springer*, pp. 932-942.

Nakamoto, S. (2008), *Bitcoin: A Peer-to-Peer Electronic Cash System*, Satoshi Nakamoto.

Namvar, N., Saad, W., Bahadori, N. and Kelley, B. (2016), "Jamming in the Internet of Things: a game-theoretic perspective", In 2016 IEEE Global Communications Conference (GLOBECOM), *IEEE*, pp. 1-6.

National Institute of Standards and Technology (NIST) (2015), "Secure hash standard (SHS)", FIPS PUB 180-4.

Popov, S. (2018), "IOTA whitepaper v1.4.3", *New Yorker*, Vol. 81 No. 8, pp. 1-28.

Quesada-Real, F.J., Moya-Pérez, F., Rodriguez-Garcia, M. and Dutta, B. (2025), "A transparent and ecologically sustainable DLT-based approach for tendering processes", *JUCS - Journal of Universal Computer Science*, Vol. 31 No. 3, pp. 277-297.

Raikwar, M., Gligoroski, D. and Kralevska, K. (2019), "SoK of used cryptography in blockchain", *IEEE Access*, Vol. 7, pp. 148550-148575.

Ramamurthy, B. (2020), *Blockchain in Action*, Manning Publications.

Ramos-Cruz, B., Quesada-Real, F.J., Rodriguez-Garcia, M., Andreu-Pérez, J. and Martínez, L. (2024), "Combining distributed ledger technologies and differentially private sketching techniques for securing health monitoring", *International Conference on Ubiquitous Computing and Ambient Intelligence*, *Springer*, pp. 920-931.

Rose, K., Eldridge, S. and Chapin, L. (2015), "The Internet of Things: an overview", *The Internet Society (ISOC)*, Vol. 80 No. 15, pp. 1-53.

Ryang, H. and Yun, U. (2016), "High utility pattern mining over data streams with sliding window technique", *Expert Systems with Applications*, Vol. 57, pp. 214-231.

Senart, T. (2025), Vegeta load testing.

Sengupta, J., Ruj, S. and Das Bit, S. (January 2020), "A comprehensive survey on attacks, security issues and blockchain solutions for IoT and IIoT", *Journal of Network and Computer Applications*, Vol. 149, p. 102481.

Shah, S.A., Seker, D.Z., Hameed, S. and Draheim, D. (2019), "The rising role of big data analytics and IoT in disaster management: recent advances, taxonomy and prospects", *IEEE Access*, Vol. 7, pp. 54595-54614.,

Shahab, S., Agarwal, P., Mufti, T. and Obaid, A.J. (2022), "SIoT (social Internet of Things): a review", *ICT Analysis and Applications*, pp. 289-297.

Sravya, P., Shree, R., Madhuri, G., Vani, G. and Vidya, J. (2023), "An efficient spam detection technique for IoT devices using machine learning", *International Journal of Innovative Research in Advanced Engineering*, Vol. 10 No. 5.

Sunny, N. and Scott, K. (2025), "Ppcoin: peer-to-peer crypto-currency with proof-of-stake".

Uddin, M.A., Stranieri, A., Gondal, I. and Balasubramanian, V. (2021), "A survey on the adoption of blockchain in IoT: challenges and solutions", *Blockchain: Research and Applications*, Vol. 2 No. 2, p. 100006.

Ullo, S.L. and Sinha, G.R. (2020), "Advances in smart environment monitoring systems using IoT and sensors", *Sensors*, Vol. 20 No. 11, p. 3113.

Vermesan, O., Friess, P., Guillemin, P., Gusmeroli, S., Sundmaeker, H., Bassi, A., Soler, I., Mazura, M., Harrison, M. and Eisenhauer, M. (2022), "Internet of Things strategic research roadmap", *In Internet of Things-Global Technological and Societal Trends from Smart Environments and Spaces to Green ICT*, River Publishers, pp. 9-52.

Wang, X., Zha, X., Ni, W., Liu, R.P., Guo, Y.J., Niu, X. and Zheng, K. (2019), "Survey on blockchain for Internet of Things", *Computer Communications*, Vol. 136, pp. 10-29.

Wu, X., Duan, R. and Ni, J. (2024), "Unveiling security, privacy, and ethical concerns of ChatGPT", *Journal of Information and Intelligence*, Vol. 2 No. 2, pp. 102-115.

Wu, Y., Chen, Y., Wang, L., Ye, Y., Liu, Z., Guo, Y. and Fu, Y. (2019), "Large scale incremental learning", *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 374-382.

Xu, J., Gu, B. and Tian, G. (2022), "Review of agricultural IoT technology", *Artificial Intelligence in Agriculture*, Vol. 6.

Yang, D., Long, C., Xu, H. and Peng, S. (2020), "A review on scalability of blockchain", *Proceedings of the 2020 2nd International Conference on Blockchain Technology*, pp. 1-6.

Zaeem, R. N. and Barber, K.S. (2020), "The effect of the GDPR on privacy policies: recent progress and future promise", *ACM Transactions on Management Information Systems (TMIS)*, Vol. 12 No. 1, pp. 1-20.

Zantalis, F., Koulouras, G., Karabetsos, S. and Kandris, D. (2019), "A review of machine learning and IoT in smart transportation", *Future Internet*, Vol. 11 No. 4, p. 94.

Zha, X., Zheng, K. and Zhang, D. (2016), "Anti-pollution source location privacy preserving scheme in wireless sensor networks", *2016 13th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON), IEEE*, pp. 1-8.

Zha, X., Ni, W., Wang, X., Liu, R.P., Guo, Y.J., Niu, X. and Zheng, K. (2018), "The impact of link duration on the integrity of distributed mobile networks", *IEEE Transactions on Information Forensics and Security*, Vol. 13 No. 9, pp. 2240-2255.

Zhang, K., Liang, X., Lu, R. and Shen, X. (2014), "Sybil attacks and their defenses in the Internet of Things", *IEEE Internet of Things Journal*, Vol. 1 No. 5, pp. 372-383.

Zhao, S., Blaabjerg, F. and Wang, H. (2020), "An overview of artificial intelligence applications for power electronics", *IEEE Transactions on Power Electronics*, Vol. 36 No. 4, pp. 4633-4658.

Zhu, Q., Loke, S.W., Trujillo-Rasua, R., Jiang, F. and Xiang, Y. (2019), "Applications of distributed ledger technologies to the Internet of Things: a survey", *ACM Computing Surveys (CSUR)*, Vol. 52 No. 6, pp. 1-34.

**Further reading**

Go.dev (2022), "Golang", available at: https://go.dev/, (accessed 4 June 2025).

GNU Project (2022), "GNU Bash Reference Manual, Version 5.2", available at: www.gnu.org/software/bash/manual/ (accessed 4 June 2025).

libp2p.io (2025), "LibP2P", available at: https://libp2p.io/, (accessed 4 June 2025).

NEM Project (2025), "Proof of Importance", available at: https://nemproject.github.io/nem-docs/pages/Whitepapers/NEM_techRef.pdf (accessed 4 June 2025).

Özsu, M.T. and Valduriez, P. (1999), *Principles of Distributed Database Systems*, Springer, Vol. 2.

Protocol Labs (2017), "Filecoin: a decentralized storage network", available at: https://research.protocol.ai/publications/filecoin-a-decentralized-storage-network/protocollabs2017a.pdf (accessed 23 May 2025).

Statista.com (2022), "Internet of things (IoT) connected devices installed base worldwide from 2015 to 2025", available at: https://www.statista.com/statistics/471264/iot-number-of-connected-devices-worldwide/ (accessed 4 June 2025).

**Corresponding author**

Francisco Moya Perez can be contacted at: fpmoya@ujaen.es