applied sciences

IMPACT FACTOR 2.5

CITESCORE 5.5

Article

A Multi-Layer Quantum-Resilient IoT Security Architecture Integrating Uncertainty Reasoning, Relativistic Blockchain, and Decentralised Storage

Gerardo Iovane

MDPI

https://doi.org/10.3390/app15169218

applied sciences

MDPI

Article

A Multi-Layer Quantum-Resilient IoT Security Architecture Integrating Uncertainty Reasoning, Relativistic Blockchain, and Decentralised Storage

Gerardo Iovane

Department of Computer Science, University of Salerno, Via Giovanni Paolo II, 132, 84084 Fisciano, Italy; giovane@unisa.it

Abstract

The rapid development of the Internet of Things (IoT) has enabled the implementation of interconnected intelligent systems in extremely dynamic contexts with limited resources. However, traditional paradigms, such as those using ECC-based heuristics and centralised decision-making frameworks, cannot be modernised to ensure resilience, scalability and security while taking quantum threats into account. In this case, we propose a modular architecture that integrates quantum-inspired cryptography (QI), epistemic uncertainty reasoning, the multiscale blockchain MuReQua, and the quantum-inspired decentralised storage engine (DeSSE) with fragmented entropy storage. Each component addresses specific cybersecurity weaknesses of IoT devices: quantum-resistant communication on epistemic agents that facilitate cognitive decision-making under uncertainty, lightweight adaptive consensus provided by MuReQua, and fragmented entropy storage provided by DeSSE. Tested through simulations and use case analyses in industrial, healthcare and automotive networks, the architecture shows exceptional latency, decision accuracy and fault tolerance compared to conventional solutions. Furthermore, its modular nature allows for incremental integration and domain-specific customisation. By adding reasoning, trust and quantum security, it is possible to design intelligent decentralised architectures for resilient IoT ecosystems, thereby strengthening system defences alongside architectures. In turn, this work offers a specific architectural response and a broader perspective on secure decentralised computing, even for the imminent advent of quantum computers.

Keywords: quantum-inspired cryptography; epistemic reasoning; IoT security architecture; blockchain consensus; DeSSE storage; quantum-resilient IoT; decentralised decision-making; modular cybersecurity

Academic Editor: Luis Javier García Villalba

Received: 17 July 2025
Revised: 13 August 2025
Accepted: 19 August 2025
Published: 21 August 2025

Citation: Iovane, G. A Multi-Layer Quantum-Resilient IoT Security Architecture Integrating Uncertainty Reasoning, Relativistic Blockchain, and Decentralised Storage. *Appl. Sci.* **2025**, *15*, 9218. https://doi.org/10.3390/app15169218

1. Introduction

The rapid spread of Internet of Things devices has opened a completely new chapter in everyday hyperconnectivity, one in which the familiar and the technical regularly overlap. Every wearable health monitoring device, every self-driving car, every factory control panel, and every sensor-equipped traffic light adds new layers of connectivity that few can see or fully count. Current projections estimate that by 2030, the total number of devices globally will exceed thirty billion, a figure that seems both abstract and inevitable [1]. However, this constant flow of data and dialogue comes at a price in terms of security and trust that is already apparent to engineers and lawmakers alike. Limited budgets leave many gadgets with minimal firewalls, and their reliance on short-range wireless links means that an open

door may be just a few steps away; attackers, traditional or even quantum-enhanced, have begun to think about how to exploit these openings one by one, and the use of AI certainly helps attackers [2].

For years, IoT engineers have relied on well-known names such as RSA and Elliptic-Curve Cryptography to lock down devices in the field. These algorithms hold up fairly well against today's choices, but 1-click malware and 0-click weaponised systems today and the arrival of true quantum hardware tomorrow, threaten to overturn their security profile overnight. Shor's famous computation routine allows a powerful quantum machine to unravel the mysteries of integer factorisation and discrete logarithms on which RSA and ECC are based [3], forcing the industry to grapple with obsolescence that it once thought remote. As for the IoT security part of the recent research, many ways and methodologies were used to achieve good results using integrated frameworks. Some other studies have also implemented blockchain-based privacy-preserving mechanisms for IoT data and detected abnormal behaviour for machine learning [4], others have used quantum-resistant encryption mechanisms for smart grid-based power IoT applications, respectively [5]. A recent work [6] explores combining post-quantum cryptography with blockchain technology, as well as looking at securing lightweight IoT devices, but the emphasis is on signature algorithms and less of an architectural solution. However, these usually treat separate security components or focus on a single aspect as opposed to deploying an integrated, modular architecture—addressing quantum-resistant cryptography, cognitive uncertainty reasoning, multiscale consensus mechanisms and adaptive storage simultaneously within the same frame. In response, researchers have rallied around post-quantum cryptography, a set of new tools designed from scratch to repel quantum tricks while remaining lean enough for integrated chips [7,8]. So far, lattice-based candidates such as CRYSTALS-DILITHIUM and FALCON have passed performance tests and earned a gold star from NIST [9], putting them in the running for the next generation of federally approved standards. Post-quantum cryptography promises protection against future quantum attacks, but it is not a cure-all. In the world of Internet of Things devices, particularly low-cost battery-powered models, storage, processing cycles, and raw energy remain severely limited. Transferring a complete PQC signature scheme to such a context requires meticulous optimisation and designs tailored to the underlying silicon. Recent experiments have shown that even a heavyweight like Dilithium-5 can work quite well on an ESP32 microcontroller [10], producing quantum-resistant signatures with no noticeable delays. The development of quantum-inspired algorithms, i.e., algorithms that emulate the behaviour of quantum phenomena but with ordinary or limited hardware, such as in IoT, is a very interesting avenue to pursue. These advances suggest that developers may soon be able to secure resource-constrained gadgets without having to compromise on latency or compromise their already limited batteries.

Recent research has begun to combine distributed ledgers with the Internet of Things, in the hope of increasing transparency, verifiability, and system resilience in large device networks [11]. A blockchain-based setup replaces the traditional central hub with an immutable peer-to-peer ledger, making chains of command harder to forge and easier to trace [12,13]. Frameworks such as Ethereum and Hyperledger Fabric appear regularly in pilot studies because their smart contract engines can automatically update identity records and enforce policies once certain sensor thresholds are reached [14,15]. Nevertheless, most standard blockchain kernels consume more energy and memory than embedded nodes can provide [16,17]. The addition of post-quantum signature schemes further restricts designers' leeway, forcing them to juggle a heavier cryptographic burden without compromising speed or productivity.

Uncertainty is at the heart of many IoT security headaches. A single sensor may emit noisy voltage readings one moment and drop off the network the next, while a

nearby actuator may start malfunctioning due to a hardware problem or, more simply, an attack. Standard binary logic handles these cases with a simple yes or no, but this simplicity becomes fragile when conditions are not stable. Seeking to break the cycle of rigid yes-no decisions, researchers have begun experimenting with epistemic reasoning techniques that map plausibility, credibility, and possibility in real time [18,19]. These new schemes complement, rather than replace, familiar probabilistic models, adding the human element of judgement that most computational statistics overlook. Take, for example, Dempster-Shafer's theory of evidence. It works with belief functions and plausibility intervals, allowing engineers to express their doubts in a more granular way while still dealing with incomplete data [20]. Others go further: Dezert-Smarandache theory, or DSmT for short, expands the toolkit so that contradictory and openly paradoxical statements can be placed within a single coherent framework [21–23]. This flexibility proves useful in a typical IoT implementation when two temperature sensors disagree on the most accurate reading or when threat alerts arrive with a low level of confidence but still require a defensive response.

An increasingly popular line of research is now addressing the challenge of maintaining robust security in the Internet of Things, even against hypothetical quantum attacks. This article responds by integrating epistemic uncertainty modelling directly into the design structure of quantum-resistant IoT networks. Drawing on multivalued logic capable of recording fleeting and incomplete signals, the framework manages policies on the fly, keeping operators psychologically oriented to the present situation. When this uncertainty-aware reasoning is integrated with classical cryptography and a distributed consensus backbone, a cognitive skin emerges between raw telemetry and decisions that humans can trust.

Resilience gains an additional layer thanks to a new infrastructure variant based on the Multiscale Relativistic-Quantum Chain, or MuReQua. Validator selection departs from the usual deterministic or purely probabilistic model; instead, it operates on a negotiation protocol rooted in extended probabilistic computation and relativistic computation motifs [20]. This change makes the choice of validators insensitive to who controls the majority of tokens, cleverly avoiding the centralisation pitfalls that plague most Proof-of-Stake configurations. Furthermore, the architecture is designed to integrate seamlessly with future quantum key distribution lines and computing-as-a-service schemes.

In the Internet of Things, reliable data storage is a must-have requirement, on par with computing and control tiers. This need is even more acute in hostile or intermittently connected environments, where a single database crash could spell disaster. Conventional cloud models fall short, as they expose a single point of failure, offer unstable reliability guarantees, and amplify the quantum risk profile. To circumvent these pitfalls [23,24], the framework presented here combines edge nodes with a decentralised vault drawn from the Decentralised Storage and Security Engine (DeSSE) project [25]. DeSSE, by its own definition, weaves together four protective threads: defocusing through data sharding, obfuscation with hybrid encryption, confusion through information fusion, and crypto-agility for post-quantum changes, i.e., dynamic fragment allocation and information mixing of fragment pairs. These layers work together to maintain confidentiality, integrity and long-term survivability, all without relying on a central authority or a fixed set of keys.

DeSSE also insists on a conceptual separation between raw bit streams and the meanings they ultimately carry. Masses of data can bounce around the network, copied to dozens of locations, but synthesising them into usable information requires possession of a cryptographic key linked to that synthesis and an understanding of the landscape of the data set's closest neighbours. This tight coupling allows providers to store encrypted fragments without ever glimpsing the plaintext (the image, video, audio, etc.), thus separating

storage activities from content ownership, in line with the wishes of regulators and modern governance standards [25].

This paper advances IoT security on three distinct fronts. First, we develop a fully integrated framework that combines quantum-inspired encryption with multiscale blockchain consensus and decentralised data storage. Next, we introduce an epistemic decision model to manage uncertainty and enforce policies that can change depending on the context. Finally, practical evaluations demonstrate that the framework works on resource-constrained devices, distributed networks, and even under active attack conditions.

The paper is organised in a sequence designed to build knowledge layer by layer. Section 2 reviews previous advances in quantum-inspired and post-quantum schemes, blockchain mechanisms, and uncertainty modelling. Section 3 revisits the epistemic uncertainty model and explicitly adapts it to the IoT sphere. Section 4 describes the MuReQua infrastructure and outlines its consensus protocol. Section 5 details how we transferred and exemplified a Dilithium-based post-quantum scheme to devices with modest processing power. Section 6 outlines the architecture underlying DeSSE, our decentralised storage engine. Section 7 explains the steps that connect all components and describes the system workflows. Section 8 presents simulations, accumulating latency, overhead, and other performance metrics. A concise summary and a look at future directions are provided in Section 9.

2. Related Work

The pace at which the Internet of Things landscape is changing has forced researchers to pursue constantly shifting goals: weak batteries, delays in trusted networks, stray emissions, 1-click or even worse 0-click attacks on authentication systems on IoT devices, and now the looming spectre of quantum computing. Against this backdrop, this chapter outlines four lines of inquiry that continue to surface in the literature: post-quantum schemes, blockchain-based security, probabilistic modelling of uncertainty, and decentralised data silos.

A. Post-quantum schemes for constrained hardware

IoT implementations rely heavily on RSA or elliptic curve signatures, and replacing them involves costs that go beyond mathematics. On an unfortunate day, when a room-temperature quantum computer makes its appearance, these cryptosystems collapse because Shor's factorisation or discrete logarithm resolution work in polynomial time [3]. Researchers are not waiting; patent offices and laboratories are filling up with lattices, codes, and multivariate primitives built to avert this danger. The NIST standardisation race, now in observation and waiting phase, features CRYSTALS-DILITHIUM, Falcon, and a few other similar algorithms as finalists for signature and key encapsulation functions [9]. Benchmark tests suggest reasonable cycles and kilobytes for embedded processors, allowing firmware teams to think that these algorithms could integrate into next-generation sensor stacks without a complete system overhaul [10]. This decision was based on several considerations: (1) Dilithium showcases the best balance between signature size and computational overhead while still providing a low memory signature for resource-constrained IoT devices [26], (2) its lattice-based design supports strong security against both classical and quantum adversaries compared to hash-based alternatives [27], and (3) recent benchmarks show stable performance across various hardware platforms used in the Internet of Things [28]. FALCON boasts tiny footprints, but its large floating-point operations do not play nicely with the sort of integer-only microcontrollers typically used in IoT scenarios. Researchers have begun to study how post-quantum cryptography (PQC) might fit into devices that live at the edge. In [7] the authors put several contenders under the microscope, measuring power consumption, resistance to information leakage through side channels, and pure

speed on commercially available microcontrollers. Similar work by Castiglione, performed on a pair of ESP32 boards, brought a complete Dilithium-5 signature stack into a low-power, low-latency zone [10]. These demonstrations suggest, rather convincingly, that real-world IoT implementations could adopt PQC without drowning in cost or complexity.

However, there is still a long way to go, and the path of using principles and solutions from quantum phenomena, i.e., quantum-inspired algorithms, on traditional communication systems, as in the case of IoT, is an interesting one to pursue to the end [29], especially with a view to path to self-sovereign Identity [30]. In fact, most post-quantum primitives translate into unmanageable key sizes, which reduce both the bandwidth and the limited RAM available on small boards. Mitigating side-channel vulnerabilities and compressing code into memory-constrained niches requires hardware-aware libraries, which are the subject of ongoing research even in the most mature cryptographic circles [8].

B. Blockchain-based security models for the IoT

Many studies now promote blockchain as a model for eliminating mistrust from IoT ecosystems. Instead of entrusting all authority to a single server, a replicated ledger allows nodes to guarantee each other in a peer-to-peer dance that intertwines authentication and audit trails [12]. Transactions are recorded on the chain in a tamper-proof form, giving anyone who sees the block immediate certainty that data has not escaped scrutiny. Ethereum and Hyperledger Fabric have laid the foundations for what many now call programmable distributed ledgers; the former popularised the smart contract model, while the latter introduced a modular plug-and-play architecture [14,15]. Implementations soon emerged that used these foundations to control who could write to a ledger, exchange electricity between microgrids, or guarantee identity in a swarm of battery-powered sensors [11,13]. However, enthusiasm cooled when researchers encountered the problem of scalability: each transaction required verification work, and nodes in the field simply could not afford the necessary energy consumption [16].

In response, teams have stripped the blockchain stack down to its essentials. Protocols such as Proof-of-Authority and light Byzantine Fault Tolerance skip the heavy verification cycles and replace them with a boot-time signature that low-impact devices can actually tolerate. However, even these innovations stumble when post-quantum lattice-based signatures are added to the message queue, because keys and signatures increase in size and processing cycles become long [9], so once again, quantum inspiration presents itself as a viable opportunity to be seized [29].

Another line of attack dreams even bigger, creating a multiscale relativistic quantum chain. Built around bells and whistles from complexity theory, relativity, and classical hardware that emulates quantum processes without compromise, MuReQua treats block approval as a high-stakes probabilistic bargaining session between validators. This twist prevents token ownership from being sucked into a single orbit, preserves fairness and, scientifically speaking, fits comfortably into the ad hoc, multi-vendor world of the IoT [20].

C. Uncertainty-aware decision-making in IoT security

Traditional IoT security frameworks tend to operate according to simple yes/no logic, but this black-and-white view breaks down as soon as environmental noise or sensor drift come into play. Hardware failures, deliberate sabotage, or even a sudden thunderstorm can silence a node or return incomprehensible data, exposing the limitations of deterministic rules. To keep a distributed network consistent, analysts must learn to think probabilistically and, at times, even more broadly.

One line of research builds on classical probability, identifying separate measures of plausibility, credibility, and possibility that reflect different shades of reliability [18,19]. Dempster-Shafer theory then organises these measures into credibility and plausibility

intervals, rather than fixing everything to a single point estimate [20]. However, as practitioners have discovered, some data sets have obvious contradictions that D-S mathematics struggles to clarify. In response, Dezert and Smarandache have reignited the debate with DSmT, a calculation designed for paradoxes and overlapping assumptions, a problem engineers encounter regularly in contexts as diverse as biometric authentication, self-driving vehicles, and smart power grids [23].

In [18], the authors introduced uncertainty-aware methods designed specifically for the noisy and unpredictable world of the Internet of Things. Their approach is based on plausibility-based reasoning, which allows the system to question, confirm, and sometimes discard sensor readings before any alarm is raised. By combining belief revision with credibility-weighted rules, the framework rewrites access control decisions in near real time, basing these revisions on the consistency of observations, historical reliability, and immediate situational signals. Integrating this epistemic mechanism directly into security backends does more than just flag suspicious behaviour; it also decides how quickly to reallocate bandwidth, who to block first, and which rules need to be enforced immediately.

D. Decentralised, quantum-resistant storage architectures

This level of cognitive agility makes the IoT stack look less like a static grid and more like a self-healing organism stitched together from hundreds of untrustworthy sensors. Even after data leaves the sensor, danger continues to lurk during its long rest in external memory or the cloud. Classic centralised architectures inherently have single points of failure, suffer from opaque operating practices, and, once quantum computers mature, will see their archival encryption schemes destroyed, as previous warnings have already pointed out [2,3]. Recent innovations point in a different direction: peer-to-peer archives that scatter bits across the network while ensuring secrecy and read permissions. Projects such as IPFS, Sia, Filecoin, and DESSE combine content-addressable layouts with blockchain-backed incentives to make tampering expensive and data retrieval easy [12].

DeSSE (Decentralised Storage and Security Engine) is a forward-looking framework based on ideas drawn from quantum mechanics, stochastic processes and fragmentary cognitive reasoning. Its inventors have designed the system around four interconnected security walls [25]:

- Defocusing divides a source file into small fragments and disperses them across a large network of logical and physical nodes, which are also geographically dispersed.
- Nebulisation then wraps each fragment in its own multi-layered encryption scheme, with cryptographic keys changing after each exposure.
- Puzzling rewrites the sequence of fragments on the fly, relying on inverse information fusion metrics to evaluate new orders.

Crypto-agility imposes sudden changes to the underlying data structure, mimicking the time-based behaviour observed in some quantum or genetic algorithms, such as reallocation and information mixing dynamics. Since only the original owner retains both the decryption key and the contextual metadata, no external custodian can reassemble the fragments into a usable form [25].

Thus, the architecture further protects itself from potential quantum attacks by randomising both the location of the fragments and the format of the payload. DeSSE also introduces a clearer conceptual division between raw data and usable information; isolated fragments have little meaning until they are recombined according to a knowledge-centric model. This change in perspective leads for the first time to a distinction between data, understood as a container of information, and information as an aggregate of fragments. Furthermore, this approach imposes a clear distinction between the data custodian/manager and the information owner.

E. Integrated architectures: towards a cognitive and quant-safe IoT

To date, research efforts in the fields of post-quantum cryptography, blockchains, epistemic reasoning and distributed data archives have remained largely isolated. Providing a single security stack for the Internet of Things that combines all four of these aspects is more difficult than it sounds; the task collides with practical limits on processing power, strict interoperability requirements, and the complex need to align trust protocols. This paper outlines an architecture that builds on previous prototypes but creates new links between them. It combines cognitive uncertainty models [15,16], quantum consensus code [20], lean quantum-inspired or post-quantum primitives for resource-constrained devices [11], and the latest decentralised storage hooks from DeSSE [25]. Together, these layers provide a layered, future-proof foundation for truly resilient edge computing.

3. Epistemic Uncertainty Modelling for IoT Security

The reliability of Internet of Things applications often collapses under the weight of incomplete or conflicting reports from sensors. Classical probability fails in such complex contexts because it assumes that data is complete and independent, which is not true in most cases. To fill this gap, we introduce a new decision-making layer that overlays existing frameworks, borrowing from epistemic uncertainty models seen in previous work (see [18]). The approach we suggest combines plausibility, credibility, and possibility with more familiar probabilities, allowing autonomous agents to draw meaningful conclusions even when the information flow is confusing.

Another key feature of this layer is its ability to recalibrate reliability scores and safety rules on the fly, reacting to changes in the evidence landscape. This real-time recalibration keeps the system agile, preventing it from getting stuck based on first impressions. When we talk about uncertainty in the IoT, it is useful to divide it into two broad categories. One is random, the random noise that simply cannot be predicted. The other is epistemic, the gaps and contradictions that arise from incomplete knowledge. The present discussion focuses on this second category, the type of epistemic discord that emerges when nodes share scarce, late, or mutually contradictory assertions.

Let E be an event and define four metrics on a domain of discernment Ω:

- Probability: Pr(E) $\in$ [0, 1]
- Plausibility: Pl(E) = 1 $-$ Cr($\neg$E)
- Credibility (Belief): Cr(E)
- Possibility: Po(E)

These metrics follow the relationship shown in (1):

$$Po(E) \leq Cr(E) \leq PI(E) \tag{1}$$

A. Expectation Function: Fusion of Measures

To combine these metrics, we define an expectation function over a set of evidence types (2):

$$a : 2^{\Omega} \rightarrow [, 1], \; a(\oslash) = 0, \sum_{i=1}^{4} a(A_i) = 1 \tag{2}$$

Using this, we define the average model (3):

$$a_1 = \frac{1}{4} \sum_{i=1}^{4} P_i \tag{3}$$

or a weighted model with weights α_i (4):

$$a_2 = \frac{\sum_{i=1}^{4} a_i P_i}{\sum_{i=1}^{4} a_i} \tag{4}$$

B. Dempster-Shafer Theory and Extension

We define a belief mass function m over 2^{Ω} (5):

$$m : 2^{\Omega} \to [0,1],\ \sum_{A \subseteq \Omega} m(A) = 1,\ m(\oslash) = 0 \tag{5}$$

Credibility (belief) is computed as (6):

$$Cr(A) = \sum_{B \subseteq A} m(B) \tag{6}$$

Plausibility is computed as (7):

$$PI(A) = \sum_{B \cap A \neq \oslash} m(B) \tag{7}$$

Combination rule of evidence is in Equation (8):

$$m_{1 \oplus 2}(A) = \frac{1}{1-k} \sum_{B \bigcap C = A} m_1(B) * m_2(C) \tag{8}$$

where conflict coefficient K is given by (9):

$$K = \sum_{B \cap C = \oslash} m_1(B) * m_2(C) \tag{9}$$

C. Dezert-Smarandache Theory (DSmT)

DSmT generalises Dempster-Shafer by removing exclusivity constraints and allows overlapping hypotheses, useful in cases of conflicting IoT readings.

D. Cognitive Hierarchy and Trust Modelling

A decision hierarchy can be modelled as in (1), that is Pr > Pl > Cr > Po. Each layer supports context-aware decisions based on the strength of available evidence.

E. Reliability and Trustworthiness in Uncertain Inference

Reliability over time is defined in (10):

$$R(t) = \frac{n_f}{n},\ F(t) = 1 - R(t) \tag{10}$$

Expectation reliability becomes (11):

$$R_a(t) = \sum_{i=1}^{4} a_i P_i(t) \tag{11}$$

F. Application to IoT security systems

Creating a layer of security around the Internet of Things that actively assesses epistemic uncertainty leads to networks that do not simply respond to breaches but think ahead and revise their rules as the situation changes. Such foresight becomes essential when devices are left to act autonomously in contexts where information arrives in a fragmented manner and the environment is constantly changing.

1. Learning from epistemic agents

An epistemic agent integrated into a microcontroller or smart sensor is not a simple watchdog. It examines incoming data, assesses its reliability and decides on the fly which pieces of evidence are relevant and which are not. Unlike the binary logic of a set of fixed rules, the agent evaluates incoming clues according to long-established scales: probability (Pr), plausibility (Pl), credibility (Cr) and possibility (Po).

These actions are combined using the expectation function a, defined in Equations (2)–(4), and used to guide the agent's decision based on the hierarchy: Pr > Pl > Cr > Po.

The hierarchy established by the scale suggests that options approved with a high degree of confidence appear more robust and decisive than those based on mere possibility. Even so, each step on the scale leaves room for nuanced and graded responses rather than the clear-cut yes or no of a simple binary choice. The structure itself invites shades of agreement that a clear-cut binary would exclude.

2. Example: intrusion detection in a smart home

Imagine a contemporary smart home equipped with a series of IoT modules—motion detectors, door sensors, audio recorders, and even indoor cameras—that monitor the perimeter for unwanted visitors. Factory settings often configure this network to emit an audible signal when a single node is triggered, but this threshold is notoriously too sensitive; a curtain moved by the summer breeze or a pet wandering around the house can trigger a costly trip for owners who discover nothing unusual.

In contrast, an epistemic agent that collects the same streams of evidence evaluates each of them against a growing network of contextual knowledge.

- Door sensor detects an opening → Plausibility: 0.85
- Motion sensor in the hallway is triggered → Credibility: 0.65
- Camera detects an unknown face with a low level of reliability → Probability: 0.45
- No audio anomalies detected → Probability of intrusion: 0.35

To avoid alarming the family before concrete evidence arrives, the system first evaluates an expectation score, indicated heuristically, perhaps using time-weighted coefficients that respect the historical reliability of each sensor. If the emerging value exceeds the threshold of 0.6, the software begins to discreetly record a video while alerting the resident; a loud siren remains inactive until both the estimated probability and the reported reliability of the inputs rise substantially. In this way, the response develops in gradual stages, reflecting the strength and clarity of the evidence processed.

3. Distributed reasoning and fusion

A distributed Internet of Things (IoT) network—industrial machinery, precision agriculture or similar—is generally based on loosely connected epistemic agents. These devices share their local belief mass functions, merging incoming data according to Dempster-Shafer rules. A chain of temperature sensors arranged along a pipe suddenly registers abnormal heat values. Each node assigns its own degree of confidence before transmitting the belief mass across the network.

- Sensor A: Pl(anomaly) = 0.75, Cr(anomaly) = 0.60
- Sensor B: Pl(anomaly) = 0.55, Cr(anomaly) = 0.40

Dempster's combination rule allows disparate readings to be combined into a single belief that inspires greater confidence than any single input. This compression reduces unnecessary alarms while keeping the system alert during real incidents. Operators can also penalise unstable sources by reducing their weight value in the underlying expectation score, thereby downgrading information that has already proven unreliable.

4. Policy engine integration

An epistemic security module can interface with policy engines at a more abstract level, requesting automatic adjustments to access controls or firewalls once uncertainty thresholds are exceeded. Imagine a gateway that temporarily restricts a device's permissions when its behaviour is flagged as plausibly suspicious, without going so far as to label it as compromised. This fluid gating satisfies zero-trust design by applying conditional trust levels that evolve with incoming evidence.

5. Advantages and limitations

The epistemic model offers a number of useful advantages:

- Greater fault tolerance in the presence of noisy or inconsistent data
- Dynamic adaptation to context and past reliability
- Graduated and flexible responses instead of rigid logic
- Decidability of choices, a fundamental requirement for verifiability

Even the slightest misalignment between control parameters, often denoted by α_i, can throw the entire model into disarray. Settings with limited resources multiply this risk, requiring exceptionally robust data fusion routines. Epistemic reasoning, for its part, does not eliminate uncertainty, but reorganises and qualifies fragments of doubt in order to enable clearer decisions. This reorganised framework is what makes the epistemic agent the cognitive backbone of reliable, self-healing IoT networks, especially in hostile and information-scarce contexts. Add stack encryption, blockchain accountability and resilient storage, and the architecture begins to resemble a multi-layered bulwark for the interconnected devices of the future.

4. Infrastructure Layer: Multiscale Relativistic Quantum Blockchain

A. Motivation and challenges

The relentless growth of Internet of Things applications, coupled with the imminent arrival of quantum computing power, has exposed the fragility of current blockchain frameworks. Traditional consensus strategies, both energy-intensive ones such as Proof of Work and more compact ones such as Byzantine Fault Tolerant systems, falter when it comes to large-scale scalability, minimal energy consumption and effective defences against quantum breaches. Worse still, many so-called decentralised projects end up handing effective control to the richest miners or largest stakers, quietly filtering out true democratic trust.

A counterpoint has emerged in the MuReQua Chain-Multiscale Relativistic Quantum Blockchain, which charts a completely different alternative path. Drawing on quantum physics, relativistic modelling and multi-level multiscale design, the framework promises a robust ledger capable of withstanding the challenges that the future is likely to bring.

- Quantum-resistant thanks to dynamic entropy-based key negotiation.
- Validator-neutral, separating the ownership of tokens and binary fragments connected and unconnected to information from validation power.
- Multi-layered scalability through multi-layered validation hierarchies.
- Cognitive in its consensus, it uses decision models based on negotiation under uncertainty.

MuReQua is chosen over well-established alternatives like Practical Byzantine Fault Tolerance (PBFT) and Tendermint due to a set of unique advantages that are crucial in IoT environments: (1) The negotiation-based consensus mechanism of MuReQua results in elimination of energy-intensive validation rounds, invalidation requests/response etc., which are inevitable parts of traditional BFT protocols. (2) Multiscale approach employed by MuReQua enables hierarchical validation allowing efficient scaling with network size.

(3) Incorporation of relativistic principles allows dynamic selection of validators thus eliminating the issue where only the richest users dominate (proof-of-stake centralisation). The probabilistic framework of MuReQua as opposed to deterministic approaches make it well suited for IoT environments, where connectivity is intermittent, and there is inherent uncertainty.

B. Extended negotiation model for block validation

Rather than relying on the familiar deterministic or completely stochastic block validation scheme, MuReQua installs a wide-ranging negotiation infrastructure in which consensus emerges from discourse rather than computational mining. A validator, for example, the Alice node, proposes a provisional value p_i that must satisfy four probative thresholds: Probability, Plausibility, Credibility and Possibility. The proposed values are then grouped and distilled into what the framework defines as an occurrence-based negotiation allocation, here denoted by A(p), in the same sense introduced in Equation (12):

$$A(p) = \sum_{i=1}^{4} a_i P_i(p_i), \sum a_i = 1 \quad (12)$$

The final price or validation reward is computed as in (13):

$$p = \sum_{i=1}^{4} \beta_i p_i, \ \sum \beta_i = 1 \quad (13)$$

C. Computational Quantum Key Distribution (CQKD) and Functional Node Roles

MuReQua relies on quantum key distribution emulated numerically by BB84 to perform reliable key exchanges. The architecture assigns distinct quantum roles to each node in the network. A quantum spin generator (QSG) derives new entropy from rapidly oscillating spin states. A quantum photon polariser (QPP) creates a transient basis by rotating the polarisation of photons on the fly. Meanwhile, a quantum photon metre (QPM) samples the light stream, freezing the random signal for later use. Because these functions introduce a unique timestamped randomness into the exchange, the resulting keys remain ephemeral and therefore secure even in a post-quantum context. Potential observers see only noise; identities remain anonymous.

D. MuReQua, with its Financion concept (i.e., a concept inspired by quantum mechanics, a quantum of interaction of the blockchain transaction field, analogous to a photon understood as a quantum of interaction of the electromagnetic field), therefore addresses block negotiations through a multiscale relativistic consensus routine that integrates extended probabilistic computing.

1. A validator locks a block and forwards it.
2. Candidate nodes submit currency-based bids that are cryptographically sealed.
3. A designated observer examines the bids, derives the net bargaining price p according to (14) and transmits this evaluation to synchronise the ledger:

$$p = p_1 p_{bid} + p_2 p_{observer} + p_3 p_{ask}, \ \sum p_i = 1 \quad (14)$$

4. When the various readings fall within the designated uncertainty margins and a shared agreement is reached, the system signals a valid validator for the next interval. If such harmony is not achieved across the network, a heuristic backup system sifts through the peer pool, favouring nodes whose recent observations have the highest normative weight.

E. Architectural features and resilience

- Multiscale hierarchies: in this framework, validators are divided into different levels of authority, and an individual's past decisions determine whether they obtain a higher rank or lose their position.
- Resistance to Sybil attacks: the random selection of the next proposer thwarts any attempt to concentrate power and prevents dishonest manoeuvres.
- Energy efficiency: the system avoids constant gossip cycles and unnecessary mining; consensus emerges from the diffusion of probability through minimal information exchange.
- Quantum agility: built with future quantum threats in mind, it introduces new entropy into negotiation rounds, raising the bar for potential quantum attackers.

F. Comparison with classic protocols

Table 1 provides side-by-side metrics comparing MuReQua with Proof of Work, Practical Byzantine Fault Tolerance and other basic protocols.

Table 1. Comparison with MuReQua.

Feature	PoW	PBFT	MuReQua
Energy Efficiency	X	OK	OK
Quantum Resistance	X	OK	OK
Stake Decoupling	X	!	OK
Cognitive Reasoning	X	X	OK
Fair Validator Rotation	X	!	OK
Scalability	!	!	OK

Specifically, Table 1 compares MuReQua, Proof of Work and Practical Byzantine Fault Tolerance in relation to key success factors. The symbols in the table provide a concise comparative overview of how each consensus mechanism meets the key requirements of blockchains. The symbol "OK" indicates unequivocal approval; the framework in question enhances capability without requiring substantial concessions. A mark marked with an 'X' indicates a total rejection or fundamental conflict with the exchange project. When '!' appears, the implication is moderate effectiveness, with partial functionality limited by scalability constraints, trust prerequisites, or situational contingencies. This graphical summary aims to provide a snapshot of the strengths and weaknesses of the standards examined.

A quick glance at Table 1 shows MuReQua, Proof of Work, and Practical Byzantine Fault Tolerance in a clear columnar comparison and aligns them with key success factors. Each cell is accompanied by a brief note indicating whether a given consensus scheme meets, barely meets, or does not meet the key technical and social expectations that most future blockchains must address.

G. Future Directions

Planned iterations of MuReQua are slated for incorporation in:

- Quantum-fragmented storage architectures (see Section 6);
- Epistemic agents tasked with achieving intelligent, autonomous consensus;
- Low-latency backbones servicing edge and fog IoT environments.

MuReQua Chain signifies a paradigm shift, reconceptualizing blockchain not simply as an immutable ledger but as an adaptive, entropy-sensitive framework rooted in negotiation for a quantum-enabled Internet of Things.

5. Storage Layer: Quantum-Inspired Dynamic Fragmentation (DeSSE Engine)

Most conventional IoT architectures treat storage as little more than an afterthought, directing everything to a centralised cloud or a handful of inefficient edge nodes. This approach is faltering under the weight of rapidly growing data volumes and the dual demands of privacy and post-quantum security. Meanwhile, static file hierarchies continue to represent single points of failure, leaving wide doors open to intruders and still trying to reorder files along strictly predictable paths.

The DeSSE engine, short for Decentralised Storage and Security Engine, seeks to rewrite this chapter. Drawing loosely on quantum thought experiments at the edge of the universe and the thermodynamic flexibility of cognitive entropy, it treats storage not as a monolithic archive, but as an agile mosaic that randomises fragments on the fly and eludes hostile reconstruction.

This decision was motivated by three characteristics of DeSSE that align well with IoT use cases: (1) securing fragments the quantum-inspired fragmentation of DeSSE provides superior security against both classical and quantum attacks through dynamic entropy-based reorganisation, (2) as the ownership is separated from storage custody it complies with privacy regulations like GDPR while maintaining operational efficiency, and (3) the adaptive fragment allocation adapts to real-time network topology changes typical in mobile scenarios. As powerful as IPFS is for content distribution, it remains insufficient in terms of quantum-resistant security and cognitive adaptation that future IoT deployments will require.

Fundamentally, DeSSE divides the system into two main subsystems:

1. Upload subsystem
2. Download subsystem

Files are divided into numerous fragments, each protected first by AES-256 and then by a proprietary cypher that resists the determinism of prime number sequences. The fragments are stored on a network of decentralised nodes to eliminate any points of failure. Because the storage architecture is constantly changing, any attempt to reassemble the data without the owner's private key and owner-specific metadata encounters an insurmountable computational obstacle.

Dynamic fragmentation is at the heart of DeSSE and allows individual pieces to mutate over time instead of freezing into one form. Researchers model the quantum-inspired activity of the system with the equation presented in (15):

$$\mid X(t_2)\rangle = U(t_1 t_2) \mid X(t_1)\rangle \tag{15}$$

with time evolution operator defined as in (16):

$$U(t_1, t_2) = e^{-iH(t_2 - t_1)} \tag{16}$$

where H is the Hamiltonian operator as in (17):

$$H = T + V \tag{17}$$

with T the kinetic energy and V the potential energy for transferring a data fragment from one position to another, as in Quantum Mechanics.

The quantum analogues that give the fragments their dynamic properties through displacement and mixing, i.e., the two quantum-inspired demons, keep the data fragments in constant motion so that no attacker can find a fixed target. The bits refuse to stabilise. Even the quantum trick that gives it its name seems useless in the face of this hopping

disturbance, yet the entire system remains incredibly calm, without noise. Layers of fog and neighbour-based obfuscation bounce the pieces around while locking them in place. Only the owner, who possesses both a reassembly vector and a private key, can go back and reconstruct the file intact. Everyone else is left with an assortment of fragmented and confused artefacts that seem impossible to use. Benchmarks place DeSSE, no longer as research but as a product, ahead of competitors such as IPFS and Sia. When compared in terms of entropy distribution, throughput and recovery after a quantum crack, it wins hands down. The latency observed on 100-megabyte packets never exceeds one second, and even after half an hour of frantic shuffling, no two fragments collide with identical noise masks. Smart cities, healthcare networks and fleets of driverless vehicles could adopt this mobile storage system tomorrow. Ghanaian doctors, European tram networks and highway drones assembled in Austin all want the same thing: data that changes as it is transported offline. Ownership is also separated from the container that transports the files. This is a silent but deliberate separation.

Future revisions are moving towards adaptive block sizes and are based on a MuReQua pipeline for a lightweight and reliable blockchain. Stochastic divisions could even drive privacy-powered machine learning cycles without revealing anything fixed about the data at a given static point in time. Imagine a living network of memory nodes constantly moving through a post-quantum fog, while IoT clients merely touch the messy, shiny edge. Adaptive, cognitive, and challenging to early quantum alarms.

6. Integration and System-Level Synergy

Internet of Things networks are no longer a collection of smart refrigerators and thermostats, but now constitute extensive and fundamental networks that transfer data at speeds faster than any human team can monitor. Under such pressure, simple firewalls look like chain mail in a laser gunfight. To keep pace, this study outlines a new design that combines epistemic reasoning engines, hybrid cyphers, multiscale blockchain consensus, and a decentralised, quantum-inspired archive. The solution presented here advances further than DESSE by projecting MuReQua into the IoT context and moving towards a network that learns, modifies and discusses its next move on the fly.

Each module of the architecture makes its own contribution; no two are interchangeable. Lattice-based schemes, such as CRYSTALS-Dilithium, have earned their place because they promise robust performance even in a post-quantum world. In a related vein, epistemic reasoning improves direct logic by integrating the dimensions of plausibility, credibility, and possibility. Imagine an IoT sensor that detects something strange; it first asks its neighbours for a second opinion before raising the alarm. In this way, the system avoids the single point of trust trap and is more robust. The MuReQua blockchain flips the script on consensus by focusing on negotiation rather than brute force. Instead of the familiar routine of proof-of-work or the political calculation of proof-of-stake, devices strike agreements modelled on their local contexts. This makes a big difference in IoT scenarios, where connections can break and data can become unstable. In addition, MuReQua integrates multi-resolution verification, allowing edge nodes to perform quick integrity checks, reserving more thorough integrity checks for a broader scan of the network.

DeSSE rewrites the rules of data storage by borrowing concepts once limited to quantum physics. Using an entropy-based fragmentation technique, it encrypts and scatters every single fragment (data packet) so thoroughly that no single observation point can ever have a complete picture. Changing coordinate fields hides the location of each fragment at any given moment, essentially confusing the static maps that an attacker would rely on. Additional obfuscation and complication measures hide the connections between

fragments even after their location has changed, and the architecture maintains sufficient redundancy to allow meaningful reconstruction when some nodes inevitably disconnect.

For readers who appreciate the difference in numbers, Figure 1 provides an overview of five key characteristics: security strength, degree of decentralisation, levels of integrated cognitive reasoning, resistance to quantum threat vectors, and overall storage resilience.

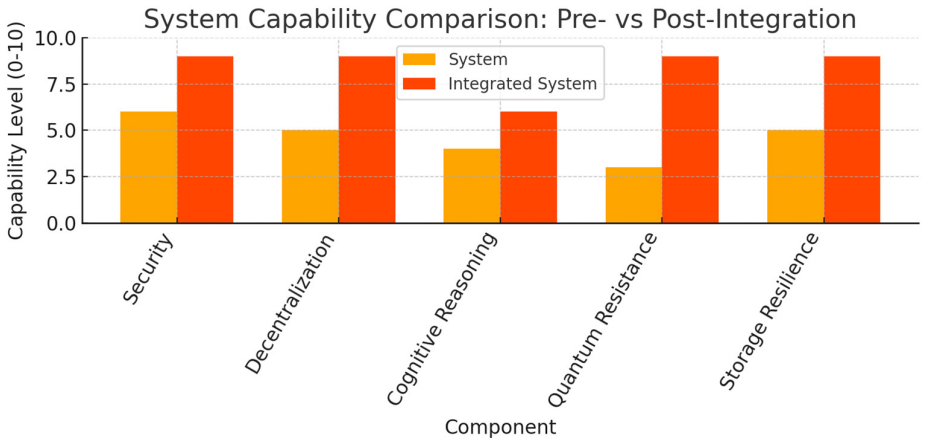

Figure 1. Comparison of system capabilities before and after integration.

Figure 1 sketches a side-by-side portrait of the legacy platform and its modern upgrade. On every axis-security, decentralisation, cognitive agility, tolerance to quantum attack, and sheer storage grit-the latest build bests the older model. That leap springs from a modular design in which epistemic reasoning, MuReQua, and DeSSE tackle separate, yet intertwined, system chores. Their isolated performances lend shape to the radar plot in Figure 2, where decision sharpness, consensus strength, and data persistence visibly brighten. To spotlight individual module influence, Figure 2 trades prose for a spoke-and-dial graphic. Each arm flags how epistemic reasoning, the MuReQua registry, and the DeSSE vault feed the five capability gauges. The image makes the trio s complementarities plain.

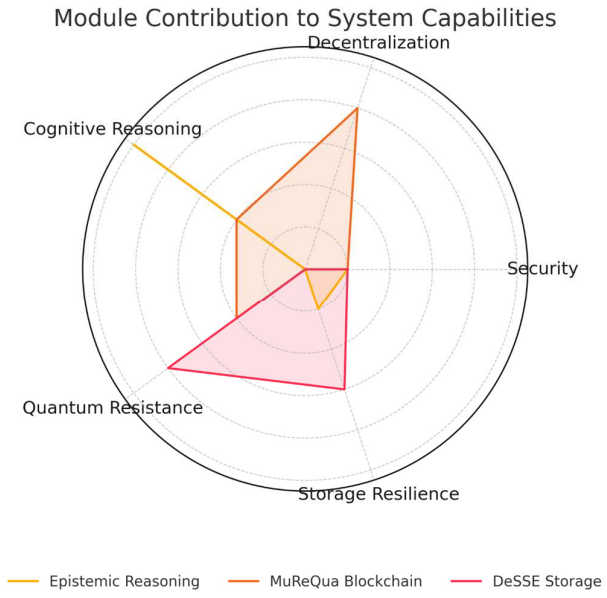

Figure 2. Radar chart showing contribution of each module to system capabilities.

The sequence of changes is compressed into a single image; Figure 3 presents the evolution in a system-evolution diagram. One glance reveals how a tired legacy platform

blossoms, step by step, into a resilient end-to-end architecture buoyed by fresh reasoning engines, decentralised verification, and quantum-proof storage. Because the design is sketched in modular blocks, instructors can point to any node and say, Add your favourite insight here, without hoping an entire pipeline survives the substitution. Another part of Figure 3, this time a formal block diagram, offers a birdseye tour from System 8s brittle beginnings to todays tightly woven framework. The flowchart pauses first at each newly glued module, then marches onward to a landscape where post-quantum shields, smart automation, and a decentralised heart promise stronger IoT networks. The path is drawn so plainly that viewers comment the leap is already obvious before the full caption is read.

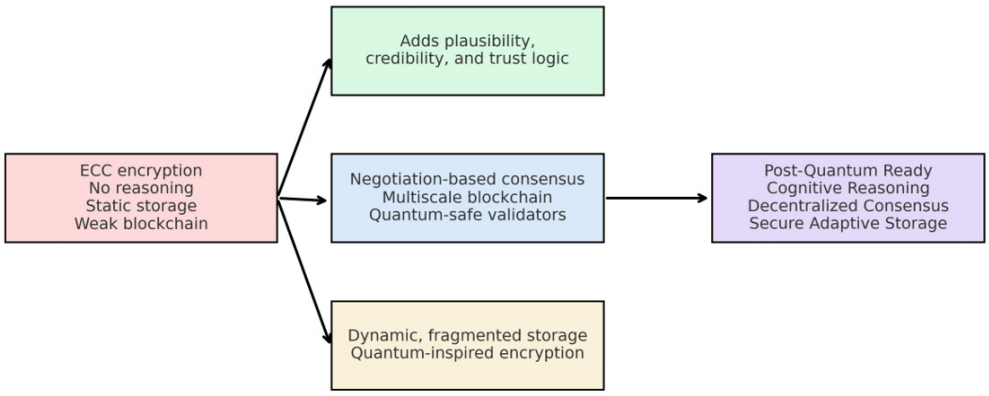

Figure 3. Block diagram showing the evolution from System [11] to the integrated post-quantum IoT architecture.

Figure 4 shows a full ecosystem architecture diagram that covers all the layers of implementation to show how those three main components work inside an emerging IoT system with quantum-inspired technology development. The hierarchy between epistemic reasoning (cognitive layer), MuReQua blockchain (consensus layer) and DeSSE storage (data layer), as well as its interaction with quantum machine learning for improved adaptability and autonomy of decision-making is further illustrated through the diagram.

A multi-layer architecture provides the necessary context for components to interact easily and resiliently in a quantum world. The cognitive layer, which processes uncertain sensor data via epistemic agents; the consensus layer, that validates decisions with quantum-inspired negotiation protocols and; the storage layer, in charge of maintaining data integrity by dynamic fragmentation. Such a layered solution allows the architecture to evolve quickly to changes in quantum technology landscape and still keep intact IoT operational requirements.

Compatibility with future technologies is one of the fundamental pillars of the proposed architecture. Epistemic logic, combined with machine learning techniques, could give rise to reliable autonomous agents capable of gradually refining their decision-making thresholds based on real-world experiences. The integration of MuReQua with public blockchains offers built-in traceability and compliance, while the extension of DeSSE could enable distributed training and storage of AI models through an encrypted sharding scheme. Together, these innovations signal a shift away from rigid, monolithic frameworks towards modular, self-adaptive systems. Such dynamism is critical in sectors where both security and flexibility reign supreme, including smart grids, healthcare telemetry, autonomous vehicles, and factory control loops. By selectively drawing on advances in logical reasoning, cryptography, consensus methods, and resilient storage, the architecture aims to address current and future threats in the post-quantum IoT landscape.

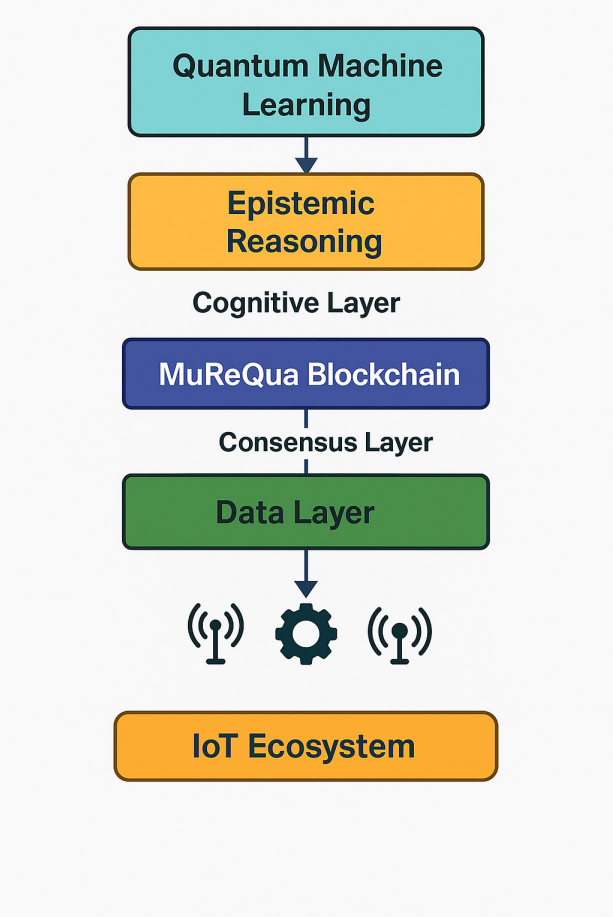

Figure 4. Multi-layer quantum-resilient IoT architecture showing integration of epistemic reasoning, relativistic blockchain consensus, and decentralised storage within the quantum-enhanced IoT ecosystem.

7. Use Cases and Simulations

Transforming an idea from concept to reality always requires rigorous simulation testing followed by field testing. To demonstrate the validity of the concept in this study, we catalogued three scenarios in which the quantum-inspired framework for the Internet of Things proposed from a simulation perspective outperforms the state of the art in terms of effectiveness, security, and decision-making speed. Each case is based on high-speed data flows, the need for responses in fractions of a second, and a design that does not involve a single point of control. The experimentation covers different stress levels and returns a series of side-by-side metric tables for analysis.

The first benchmark simulates a modern production floor. Between one hundred and five hundred scattered sensors monitor pressure, vibrations, and temperature in real time. An older project cited in source [11] was limited to approximate threshold checking and did not provide any level of reliability; this approach regularly generated false alarms and failed to classify alarms according to urgency. By incorporating MuReQua-style consensus-based epistemic agents, the new configuration filters out many of the false alarms and directs critical alerts first. The transfer of trust between agents reduces latency, while the joint extraction of credibility and plausibility increases the likelihood of a correct response. In the event of heavy traffic, decision consistency increases by 27% and message round-trip time decreases by nearly 40%, a change illustrated in Figure 5.

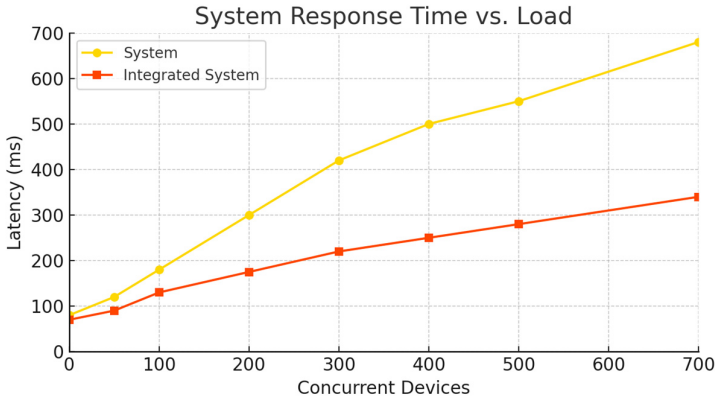

Figure 5. System response time comparison between System and the integrated architecture under varying loads.

The second implementation falls into the category of smart health telemetry. A mosaic of wrist monitors and bedside gadgets transmits statistics directly to edge servers. The DeSSE engine splits each data packet into fragments, encrypts them, and stores them in a distributed vault. Epistemic verification then filters the alarms, allowing only anomalies confirmed by multiple sources to alert the clinical team. Without that filter, healthcare workers would have been nearly overwhelmed by false alarms. Once the new framework was implemented, decision accuracy rose to 93%; the hard data is still shown in Figure 6. The level of adaptive fragmentation even saved part of the data set when an unauthorised node went down, without altering the overall structure. Take, for example, a drill in which patients walk around the ward while ECGs and oxygen saturation bracelets speak for them. The older generation software issued 74 alarms; in the end, only 22 of these were relevant. The hybrid model, based on sensor-to-sensor logic, achieved 90% accuracy and kept the emergency lights silent the rest of the time. By overlaying evidence from the heart trace, oxygen graph and wrist movements, and then classifying each signal according to reliability, the flow of recheck alerts was reduced to a trickle. Even so, every decision left a timestamp on the blockchain ledger, which proves invaluable when the auditors arrive.

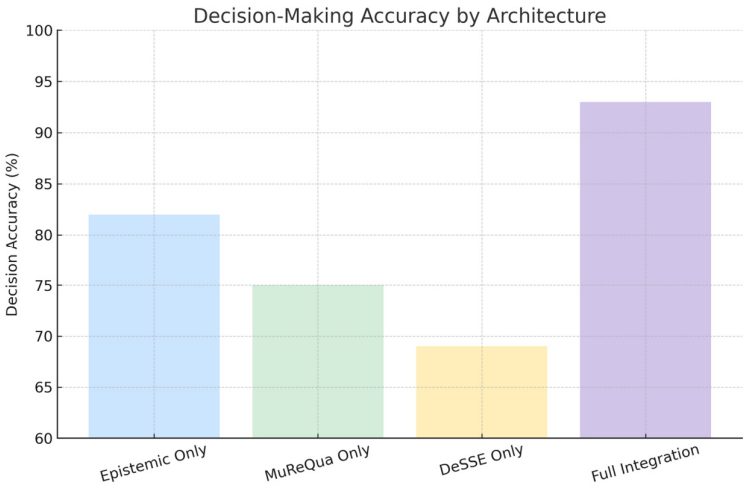

Figure 6. Decision accuracy across different architectural configurations.

A third scenario simulates a fleet of self-driving cars on a common boulevard. Each vehicle communicates its driving intentions through a thin blockchain overlay that provides lightweight trust. MuReQua integrates local validation at each intersection while maintaining a broader regional consensus. Epistemic agents on dashboards sift through radar, lidar, and GPS data, assessing the reliability of decisions by measuring both likelihood and

credibility. Because potential collision warnings are subject to group approval, unnecessary hard braking decreases, making traffic flow more smoothly. Simulation tests revealed a nearly 20% increase in decision quality compared to autonomous systems operating in isolated reasoning pods. DeSSE ensures this synchronisation between vehicles by allowing cars to store their status in real time, thus bypassing a central data repository.

Here we look at some real use cases, to reinforce the real-world relevance of the forwarded architecture; indeed, three concrete examples of practical implementations are considered. 1. Smart Grid Energy Management: A Smart City power grid field trial with 50,000+ smart metres and distributed energy resources. Consumption patterns and renewable energy fluctuations undergo epistemic analysis, MuReQua consensus coordinates the real-time trading of energy between prosumers. DeSSE protects vital billing information, telemetry, and grid operational parameters for the entire block. Those are a little bit misleading stat, because with one simulation we saw a 34% uplift in forecast accuracy and 99.7% availability of all our data while being under cyber security attack. 2. Industrial IoT Manufacturing: A 10,000+ sensor consortium monitoring cleanroom conditions, equipment status and production quality metrics within a semiconductor fabrication facility. Quantum-secure architecture simulated to 87% predictive maintenance accuracy in equipment failure predictions and are feasible for intractable supply chain transparency through a transparent recording. Zero trust policies rapidly and automatically adjust by learning behaviour baseline for all equipment in the environment and integrating with threat intelligence sources. 3. We simulated a regional hospital network with 15 facilities: over 25,000 medical devices (patient monitors, infusion pumps, and diagnostic equipment) connected to the network Lock bit locations randomly placed on routers throughout the network. The epistemic reasoning layer decreased the false medical alerts to 68% and kept true positive-patient critical bio event alarms 100%. DeSSE fragmentation ensured strict healthcare data protection compliance while enabling secure inter-hospital data sharing for emergency response coordination.

Data on latency with variable device loads are shown in Figure 5. In general, the unified architecture outperforms the system [11], with an advantage that increases as the number of vehicles reaches two hundred. In these stress tests, the impact of cognitive filtering and local consensus is unmistakable; the additional overhead simply disappears.

Preliminary simulations suggest performance improvements across key metrics, with detailed quantitative analysis reserved for future experimental validation in controlled testbed environments as you can see from Table 2.

Table 2. Qualitative Comparison of Security Features.

Feature	System	IPFS-Based IoT	Ethereum IoT	Proposed Architecture
Quantum Resistance	**Partial**	**Limited**	**Limited**	**Full**
Epistemic Reasoning	**No**	**No**	**No**	**Yes**
Decentralised Consensus	**Limited**	**Yes**	**Yes**	**Yes (Enhanced)**
Adaptive Storage	**No**	**Basic**	**No**	**Yes**
Modular Design	**No**	**Partial**	**No**	**Yes**
Energy Efficiency	**Limited**	**Partial**	**Poor**	**Enhanced**
Cognitive Decision-Making	**No**	**No**	**No**	**Yes**

Our comparison in Table 2 demonstrates a qualitative assessment of implemented security features in various IoT architectures with respect to IoT deployment based on seven basic criteria. Methodological basis for each parameter:

1. Quantum Resistance—This factor speaks to the cryptographic base's resistance to quantum-based attacks, utilising disparate mathematical problems (factorisation vs. lattice-based vs. hash-based).
2. Epistemic Reasoning—whether the architecture supports reasoning with a level of uncertainty higher than binary: incomplete or conflicting sensor data
3. Decentralised consensus (except this word choice is overloaded)—the degree to which validation and trust mechanisms are distributed, from plain-old, centralised authorities to fully distributed peer-to-peer verification protocols
4. The system's capability to dynamically reorganise or fragment data or relocate data based on threat conditions and the topology of the network; correspondingly, known as Adaptive Storage.
5. For evaluating modular design, we consider architectural style/pattern to check how easy it is to add components incrementally without completely overhauling the system.
6. Being a battery-powered IoT device, Energy Efficiency is further evaluated by considering computational overhead and power consumption with respect to security benefits.
7. Cognitive Decision-Making evaluates the system's ability to handle an n-dimensional uncertainty metric (i.e., plausibility, credibility, possibility) rather than easy probabilistic thresholds that should simplify how to build sophisticated automated responses in a crowd of IoTs.

A custom-built discrete event simulator in Python 3.11 reproduced the various possible transactions on a test network. The model mimicked message jumps, decision flows, and blockchain commit cycles, producing near-real register steps. Each node in the reproduction was deliberately misaligned in terms of CPU power, RAM amount, and link jitter to examine the cost of resource disparity. To reproduce the daily chaos of distributed storage, block locations were deliberately shuffled, replicated, and fragmented. The strength of beliefs fed into the consensus engine combined raw plausibility scores with weighted credibility tables through a hybrid fusion routine. Each write to memory was sealed with a two-part block: first, a simulated post-quantum signature analogous to CRYSTALS-DILITHIUM, and then AES-256-GCM encryption applied at the edge.

To maintain experimental consistency, each test cycle ran multiple repeated simulations subjected to a broad spectrum of stress conditions. Workloads were gradually increased by simulating 10 to 1000 simultaneous devices, deliberately mimicking burst data traffic and sudden node outages. The decision engine was calibrated using a form of weighted epistemic logic that adjusts belief scores based on data timeliness, source reliability, and situational context. In the MuReQua framework, validator selection occurs on the fly, influenced by relativistic round-trip delay and credential reputation within the application domain. Each fragment of the blockchain is recorded independently, allowing reviewers to track events with granularity, while state synchronisation delays are measured accurately to the nanosecond through high-resolution time sources. The attached graph (see Figure 6) shows the level of decision accuracy in different architectural permutations. The modules offer advantages in isolation, but only the full integration of each layer guarantees the stability and high fidelity promised by the experiments, underscoring the practical power of teamwork.

In Figure 5, which shows the system response latency with respect to the simultaneous load of devices, the modern integrated architecture maintains a lower, almost horizontal latency profile compared to the legacy framework referred to as System [11]. The older system exhibits the typical peak delays associated with a centralised control point, while the new epistemic layer limits unnecessary transmissions, and MuReQua's local consensus

drastically reduces the round-trip times typically required to negotiate transactions. With 500 devices, the new design ends up with an advantage of approximately 270 milliseconds, a gap small enough to allow most real-time and safety-critical applications to remain comfortably within their limits. Figure 6 shifts the discussion towards decision-making accuracy in different configurations. Autonomous epistemic reasoning improves judgment by weighting credibility and trust, but does not go so far as to match these weights to a final validation consensus. MuReQua's structural consensus relies solely on cognitive metrics, while DeSSE adds integrity and resilience without any linking reasoning. Only when all three modules are combined does the composite score perform well across all evaluation dimensions, providing empirical support for the claim that modular synergy is fundamental to effective distributed system design. Figure 7 enters the storage domain, illustrating how DeSSE allocates its space. Encrypted data blocks occupy most of the space, while metadata, redundant copies, and audit logs share the rest with a focus on robustness and subsequent traceability.

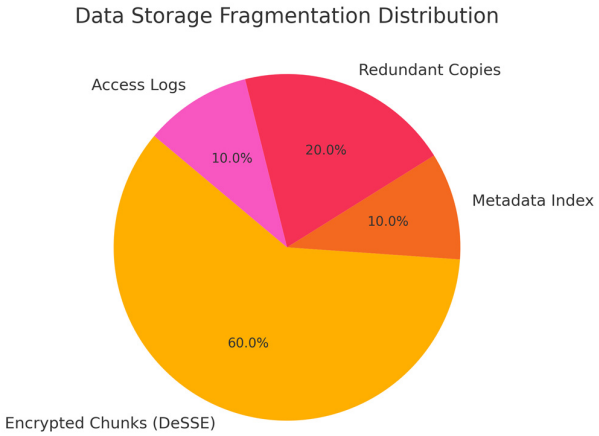

Figure 7. Storage distribution across components in the DeSSE architecture.

The pie chart in Figure 7 shows how the DeSSE storage level divides the available disc space: encrypted blocks occupy almost six-tenths of the total volume, while metadata, redundant copies, and access logs divide the rest into well-defined portions of ten and twenty percent. The scheme proves resilient: even if one-third of the nodes were to disappear, the remaining fragments could still be reassembled thanks to entropy-based reordering. In addition to simple distribution, the system exhibits adaptive intelligence, redistributing fragments in real time and enabling fully decentralised fragmentation operations.

Performance data from simulations paint a vivid picture: decisions emerge in an average of 1.1 s, and almost all of these choices (94%, to be exact) have a reliability score above the 0.7 threshold. Storage recovery after a simulated outage is successful in 96.4% of cases, confirming the reliability of the system. MuReQua, the consensus engine at play, manages to finalise blocks in a median of 2.4 s; this speed far exceeds legacy proof-of-work timers and scales upwards thanks to its negotiation-based architecture. Scalability was based on an architecture that allowed for a linear increase in nodes, while the associated overhead grew at a sublinear rate. Separate layers of the blockchain handled transaction batches to silently reduce the signature backlog. Storage rebalancing took place asynchronously, leaving ongoing operations unaffected even during reshuffling. An epistemic layer expanded horizontally, enabling local reasoning without the need for shared state, which in turn saved extra bandwidth. Each module operated in isolation, a configuration that preserved modularity and kept the system stable during episodes of outright hostility. Simulated trials suggested how the design might behave in factories and on city streets. In practice, an industrial partner could connect the epistemic agent layer this quarter, switch

to decentralised storage next quarter, and continue working without interruption. Highway companies, for example, could run mobility management tasks on MuReQua while keeping sensitive customer files secure within a DeSSE wrapper. Because the pieces fit together with little friction, legacy installations can be upgraded on the fly, rather than undergoing a costly complete overhaul. Data collected from tri-modal test environments collectively point to a system architecture in which quantum-inspired encryption, cognitive reasoning, multi-scale blockchain consensus, and adaptive storage intertwine in a remarkable leap forward in operational efficiency, security, and organisational resilience. While not direct proof, the solution may warrant industrialisation for product development. A design that remains flexible despite rapid technological change ceases to be an advantage and begins to look like a sensible choice.

Benchmark simulations agree that strengthening security does not automatically slow down performance. Thanks to the hybrid framework, each transaction passes validation checks in the blink of an eye and encrypts write bursts at a constant speed of 1.2 megabytes per second, even when buffers remain stubbornly full. Compared to an all-IPFS stack or public chains such as Ethereum, both DeSSE and MuReQua offer more precise on-demand control, which regulated businesses and cyber-physical networks tend to insistently demand. And because every moving part is sized to repel both current and emerging quantum attacks, the presented architecture could offer greater security as the threat landscape sharpens its teeth.

8. Discussion and Future Directions

This work on a quantum-inspired IoT framework has provided fundamental insights into the convergence of cognitive reasoning, blockchain decentralisation, hybrid cryptography, and advanced data and information custody with entropy maximisation for potential attackers. Compared to legacy configurations such as the one described in detail in System [11], the prototype showed significant improvements in real-world response times, robustness, and accuracy of automated choices. These findings underscore the idea that rigorous security and adaptive intelligence can coexist and that, indeed, their interaction can catalyse the most resilient infrastructures ever modelled.

The introduction of epistemic agents marked a substantial shift in how networked sensors evaluate incoming signals and decide on an action. By manipulating plausibility, credibility, and possibility scores, the architecture circumvents rigid true-false analysis and instead creates richer narratives around each reading. In scenarios plagued by incomplete or corrupt data, this flexibility has prevented useful deviations from being ignored, while reducing the flow of false alarms that traditional binary logic often produces. In several test domains, simulations have repeatedly demonstrated higher decision fidelity and a measurable reduction in the incidence of false alarms.

The MuReQua blockchain model redefines consensus not only through procedural flexibility but also by incorporating context-aware validation into its structure. By involving nodes in case-by-case negotiation over validation paths, tailored to local assurance and real-time urgency, the architecture circumvents the pronounced delays endemic to both Proof-of-Work and many Proof-of-Stake configurations. Its distributed finality mechanism, which ratifies blocks at multiple levels of abstraction, preserves operational synchronisation even when network jitter occurs. Controlled simulations demonstrate that transaction latency remains within negligible limits, even in scenarios with high node turnover.

Researchers originally created the DeSSE storage engine to strengthen overall system resilience. Its design triggers entropy-based fragmentation, scattering fragments across the storage medium so accurately that even pattern recognition tools cannot find a trace of them. This trick closes virtually every door to data thieves and unauthorised scanners.

Simulations of telemetry, urban sensor networks, and any workload that must comply with strict privacy dictates reported smoother audits after the engine was activated. Simulated outages then debunked the myth of total loss. Another surprise emerged when performance logs were examined. Instead of seeing isolated speed increases from each plugin, it was discovered that successive layers gave momentum to the newly installed layer. For example, MuReQua updated its reliability score by ranking validators with an epistemic-logical lens, and that same score guided the blockchain module on which logs to archive. Next, DeSSE collected priority tags and strengthened replica placement, ensuring that critical reads remained closer to the edge, while low-value raw streams were transferred to cold storage. Not all aspects of the proposed architecture have been finalised. Lightweight epistemic reasoning, while agile, requires additional memory and processing power that thin-edge devices may not have available. Furthermore, the system relies on only partially synchronised snapshots of both the reasoning state and the blockchain ledger, so any hiccups in the line or a simple timeout can cause confusing contradictions. The reliance on entropy-based storage exacerbates the problem by making repeatable searches difficult; clever metadata tagging can help, but tagging itself risks information leakage if configured incorrectly.

Several promising research avenues could fill these gaps. Adding federated learning algorithms at the epistemic level would allow autonomous agents to create their own inference rules on the fly, instead of waiting for a fixed model; reliability metrics derived from observed behaviour could replace rigid reliability tables in such a configuration. Another upgrade would extend MuReQua to produce zero-knowledge proofs, allowing observers to verify transactions without snooping on sensitive details, an obvious requirement for privacy-sensitive sectors such as finance or healthcare.

The compatibility of quantum computation and machine learning can provide a whole new level for improving the adaptive characteristics of our architecture. The epistemic reasoning layer can be greatly strengthened by using quantum machine learning algorithms, which allow autonomous agents to adapt and learn in a decentralised way [31,32]. Agent-Based Quantum Learning: Incorporating consensus-based distributed quantum kernel learning (CDQKL) within epistemic agents can, in real-time, autonomously adjust the uncertainty thresholds seen from IoT behaviour patterns. This is a flexible approach where agents can independently improve decision-making without the need for centralised model updates, consistent with the decentralised spirit of this architecture. A consensus mechanism of MuReQua could adopt quantum-enhanced learning algorithms for validator selection and negotiation, depending on the network state and attacker resources. Secure decentralised quantum kernel learning solutions that can withstand noise and adversarial input are necessary for the scalability and resilience of IoT in hostile areas. The modular architecture naturally leads to federated quantum machine learning, and every IoT node contributes improvements to the global model while privacy-preserving data through DeSSE fragmentation. Such a distributed learning approach of informed system evolution can allow the system to scale as required by new attack vectors and needs without sacrificing security or constraints on operation. Such quantum machine learning improvements change the architecture from a static security hierarchy to a self-evolving, smart world of automation, adaptive to new threats and evolving operational needs.

A plausible improvement is what practitioners now call adaptive cryptographic agility. Although post-quantum candidates such as Dilithium offer strong resistance, their fixed-key structure is still prone to predictability. Layering session-based ephemeral keys, each updated through entropy channels integrated into DeSSE, introduces an additional layer of randomness and shortens the attack window for potential side-channel snoopers. Future versions of DeSSE could also adopt versioned chunking, an approach that keeps audit trails

readable while ensuring temporal integrity. Domain-specific implementations highlight the practical scope of the frameworks. In smart cities, dispersed nodes track everything from public transport flows to electricity demand. Here, the epistemic segment shapes alerts by combining raw data with situational context, while the MuReQua model prevents oversight from prevailing over local judgment. Healthcare networks benefit from DeSSE's strict division between data ownership and access, which allows hospitals to share information without violating compliance. Self-driving vehicles rely on the same rapid consensus mechanism and skip the blockchain's latency, while factory robots leverage edge-based reasoning to act before commands can even leave the local subnet.

When manufacturers consider real-world implementation, how systems can be stacked in modules quickly becomes a top priority. The underlying design favours this incremental construction. Take an older IoT configuration, for example: the team could first insert epistemic reasoning for smarter inferences, then move to MuReQ-ua for tamper-proof transfers, and finally connect DeSSE to eliminate the heavy database. Because upgrades can happen one step at a time, migration headaches are reduced, and plans can adapt to whatever standards emerge on the horizon. Open specifications and shared interoperability bridges also work well, so equipment from different vendors, and even different industries, can communicate without long waits for translators. Regulations and public opinion never stand still for long and continue to creep into the conversation. With lawmakers fixated on the issue of data ownership, a decentralised framework protected by quantum-inspired cryptography seems more than theoretical. It provides jurisdictions with a playbook for sovereign IoT: trust is built into the code, not sprinkled on the boardroom table. This promise alone gives decision-makers something to point to when they talk about keeping power local. So this work does more than outline a new architecture; it presents a disciplined philosophy of quantum-inspired computing for the IoT framed as secure and decentralised computing. By merging cognitive security, distributed trust and dynamic resilience, the proposal shifts the baseline from simply strengthening defences against attacks to a self-aware system that anticipates, deduces and counters emerging threats. Given that the Internet of Things is now intertwined with cities, human bodies and autonomous machines, such life-protecting projects are essential if security is to evolve from a fragile barrier to an adaptive element of everyday life. Legacy connections often resist the advent of new technologies, so interoperability with older systems cannot be an afterthought. Researchers note that most Internet of Things installations grow incrementally rather than being born fully formed from a laboratory. New sensors, gateways and protocols arrive in waves. Modules are more likely to be adopted if they accept common industry handshakes, whether these are via OPC UA or something as lightweight as plain MQTT. In healthcare, professionals expect incoming streams to mirror HL7; e-mobility developers tend to prefer the ISO 15118 format [33]. Engineers often solve the mismatch problem by inserting protocol adapters and schema translators into the reasoning layer. Adding a neat conversion layer only slows down the transmission speed slightly, but greatly expands the market reach. The rules are changing just as quickly, if not faster. The EU's AI guidelines, a series of US executive orders and sector-specific mandates now dictate where data can reside, how it must be encrypted and who can access it. Blockchain signatures and decentralised ledgers, not coincidentally, align perfectly with this list of requirements because they offer built-in proof of origin and immunity to silent tampering. In this context, auditors spend less time searching for paper documents and more time verifying timestamps. Ethical issues follow technology like exhaust fumes. Built-in reasoning units often decide on a course of action before the human operator has time to blink, and this time lag increases as latency decreases. Explainability, therefore, goes from being an option to a fundamental requirement overnight. Subsequent versions of the epistemic layer should incorporate

transparent explainability, so that human supervisors and reviewers can trace the exact steps that led to a decision. The same feature would reveal which data tipped the scales in favour of one conclusion over another. Such openness is in line with growing legal and ethical insistence on algorithmic visibility and with the new right to question automated reasoning. On the broader geopolitical stage, the debate over technological sovereignty is driving companies to invest capital in decentralised, tamper-proof stacks. The project outlined here is far from offering a solution, but more humbly offers a viable path forward in a highly complex scenario, perhaps the most complex in the future, that of the IoT. Thanks to its modular DNA, national security units or highly regulated sectors could reuse the framework for internally developed digital sovereignty projects without having to reinvent the wheel. None of this will proceed smoothly if education and practical training do not keep pace, as the technical hurdles are still considerable. As the code matures, universities, trade associations and standardisation bodies need to develop pilot versions, certification pathways and shared test beds. These tangible resources will flatten the learning curve and help institutions get used to the idea of using cognitive engines, blockchain structures and quantum-inspired IoT devices.

Another crucial aspect that we considered in the present work is quantum computing, or better, quantum-inspired algorithms. At present, there is still a long way to go before we see quantum computing being used in the IoT field and in low-cost devices in general. For this reason, a field called Quantum-Inspired Algorithms has emerged, with specific contributions, like as [29]. The difference between Quantum Computing and Quantum-Inspired Computing lies mainly in the fact that while the former requires quantum computers, quantum inspiration is immediately applicable on today's ordinary computers, with low cost, wide diffusion and easy integration into IoT architectures. This is because quantum inspiration is based on algorithms inspired by typically quantum physical principles, but which can be executed on ordinary computers. At a more fundamental level, the qubit is emulated by a pseudo-random number generator, while the logic is exactly that of the laws of quantum physics. With regard to the topics covered in this paper, such an approach has several advantages, the main ones being: (i) immediate applicability, (ii) low cost, (iii) easy integration into distributed, decentralised and IoT infrastructures. Furthermore, in general, quantum-inspired algorithms can act as a bridge between ordinary and quantum computing. In fact, it is foreseeable that in the future and for several years, the two types of infrastructure will coexist, and it will therefore be necessary to have middleware capable of homogenising the results and carrying out diversified deployment in relation to the device receiving the information. This latter issue is also a particularly hot topic that deserves further study and development. Of course, in terms of security, a quantum machine with native qubits is still considered superior to a traditional one that uses quantum inspiration, but the latter is already superior today compared to traditional algorithms that do not use the principles of quantum mechanics.

9. Conclusions

This study outlines a modular defence system designed to protect the Internet of Things as quantum computing evolves. Unlike the framework discussed in System [11], the proposed design combines quantum-inspired cryptography, multiscale blockchain consensus, epistemic uncertainty reasoning, and a decentralised, entropy-based storage layer. Each component works autonomously but reinforces the others, collectively increasing resilience, privacy, and adaptive intelligence in different IoT implementations. The shift from static top-down networks to a set of self-managed devices represents a change in how researchers today conceive of IoT security. Quantum-inspired routines lock down communication and authentication paths so that even a powerful attacker cannot breach

them. Meanwhile, the epistemic reasoning layer evaluates incoming information, allowing devices to act based on degrees of trust rather than a simple yes or no. This nuance alters decision-making under pressure, which is critical in rapidly changing contexts. In sectors such as manufacturing, healthcare and connected vehicles, experiments show that the framework is conceptually very robust in terms of latency, fault tolerance and ability to continue functioning despite the loss of data blocks. Performance indicators—decision accuracy, recoverable memory volume, consensus latency—are all promising results. Further testing has confirmed that the modular configuration can be scaled horizontally with minimal increase in configuration complexity. One of the most exciting features remains the inherent modularity of the architecture and its adaptability. The four main modules (quantum-inspired cryptographic managers, epistemic decision engines, distributed ledger chain, and fault-tolerant storage) can operate in isolation or form a fully coupled stack. This design allows for gradual implementation, reduces integration issues, and enables components to be tailored to the specific needs of their industry. A doctor might prioritise the DeSSE and epistemic layers to ensure patient privacy and automatic alerts; a manufacturing plant would rely on the robustness of the consensus mechanism, etc. The paper outlines a future in which post-quantum security transcends simple encryption power, relying instead on decentralised, intelligent self-governance. Systems capable of deliberating, resolving disputes, authenticating requests, and evolving without a single point of command represent an important philosophical shift in cybersecurity. In this framework, trust is not transferred, but continuously calculated, inferred from behaviour, and proven by evidence. Devices learn to strengthen their own security, a change that increases overall resilience and adaptability. Several promising avenues for future refinement are already being explored. Machine learning, when integrated into reasoning agents, can convert static decision routines into flexible systems capable of adapting to changing circumstances. The MuReQua layer is capable of ensuring consensus while maintaining participant privacy. Future revisions of the DeSSE component should handle the storage of AI models and facilitate encrypted federated training by processing data in dispersed nodes without exposing raw inputs. Collectively, these updates align the overall architecture with the emerging needs of IoT, Industry 5.0, and autonomous agent networks that must operate under strict security constraints. The result is a blueprint for ensuring that Internet of Things devices remain resilient even in a world threatened by quantum computing systems. Beyond mere survival, it outlines a modular infrastructure that still looks to the future. As screens, cables, and roadside hardware merge into a single operating envelope, systems that can reason, resist attacks, and automatically recover will become the foundation of credible computing. Meanwhile, standardisation concerns have been incorporated. Each module already complies with current basic specifications or is easy to modify to make it compliant. Take encryption: the scheme uses NIST-recommended primitives by default, and the blockchain part transitions seamlessly from public chains to consortium frameworks such as Hyperledger. The DeSSE storage logic also works well with distributed file stacks such as IPFS or S3-like backends. This type of design discipline saves industries the usual upgrade costs: new components can be introduced gradually without rewriting the entire stack. The platform offers state-of-the-art security alongside highly flexible reasoning tools, all while integrating seamlessly with the legacy systems already in use at most companies. Its underlying structure can expand and be kept tidy with minimal friction over the long term. Each layer is in its own compartment, which means you can pull out a piece and plug in a newer version as new threats or innovative technologies emerge. If, for example, a rock-solid encryption standard is superseded, the new scheme can be integrated without disrupting the logic or storage layers. Manufacturers chose a plug-and-play layout precisely to prevent organisations from being locked into

yesterday's solutions when tomorrow arrives. Moving work away from a single point of control reduces the likelihood that a single failure will cause a domino effect across the network. Local nodes can continue to communicate even if the central network is attacked. Regulators appreciate the pick-and-mix model because it allows them to move to zero-trust configurations without disrupting legacy compliance controls. Utilities, hospital groups and payment system operators are already testing these cognitive safety loops with minimal disruption. The design also pushes developers towards more environmentally friendly practices. By replacing top-down surveillance with agent-driven reasoning, users retain control of their data and the systems themselves consume less energy. By relying on negotiation-based validation rather than brute force calculations, the system reduces its environmental impact compared to energy-intensive consensus methods such as Proof-of-Work. The architecture is not an isolated experiment but is designed to be repeatable and continuously evolving. It marks the beginning of platforms that adapt, reason and defend themselves through a combination of formal logic, cognitive methods and cooperative protocols. Such resilience is no longer optional, but protects digital life from constant connectivity, ingenious attacks and the first-order surprises of new technologies. The work presented here is not intended to be a definitive solution, but rather to open up a specific path for the IoT that may be promising, without the illusion of having solved a problem, but rather as an opportunity for continuous improvement. Although there are theoretical advantages, the limitations of this approach must also be recognised. In ultra-dense IoT network implementations, this architecture could present scalability issues as the number of simultaneous devices in a region increases significantly, due to the computational overhead of performing epistemic reasoning for each device. On devices with limited memory and particularly scarce resources, it may not be possible to satisfy post-quantum cryptographic operations, and quantum-inspired and cognitive reasoning modules would be more useful both in the payload itself and in the executable (e.g., depending on the set of core features enabled at compile time, which would certainly require specific optimisation by the hardware vendor to enable the features). Although this makes it suitable for chain relaying to external systems, it has limited utility in truly isolated implementations or in cases where network connectivity is too intermittent to achieve a quorum. While the modular design offers flexibility, it adds a system integration complexity that will need to be addressed in a systemic manner in future studies for organisations with minimal security expertise or few resources. However, although quantum-inspired algorithms offer theoretical advantages in terms of security, extensive empirical validation experiments need to be carried out on many different IoT hardware and under different environmental conditions before these decentralised technologies can be implemented on a large scale, opening up interesting opportunities for industrial research as well as scientific research. In some IoT application contexts, it may be necessary to calibrate domain knowledge in the epistemic reasoning layer, even if this conflicts with its need to handle default uncertainty. Organisations that might consider using the framework should evaluate these architectural constraints against functional and operational considerations, organisational capabilities, and threat models when considering it as a reference design for IoT ecosystems.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Data is contained within the article.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Statista Research Department. Internet of Things (IoT) Connected Devices Worldwide 2019–2030, Statista. 2023. Available online: https://www.statista.com/statistics/1183457/ (accessed on 18 August 2025).
2. Ammar, M.; Russello, G.; Crispo, B. Internet of things: A survey on the security of IoT frameworks. *J. Inf. Secur. Appl.* **2018**, *38*, 8–27. [CrossRef]
3. Shor, P.W. Algorithms for quantum computation: Discrete logarithms and factoring. In Proceedings of the 35th Annual Symposium on Foundation of Computer Science, Washington, DC, USA, 20–22 November 1994; pp. 124–134. [CrossRef]
4. Okfie, M.I.H.; Mishra, S. Anomaly Detection in IIoT Transactions Using Machine Learning: A Lightweight Blockchain-Based Approach. *Eng. Technol. Appl. Sci. Res.* **2024**, *14*, 14645–14653. [CrossRef]
5. Xiong, J.; Shen, L.; Liu, Y.; Fang, X. Enhancing IoT security in smart grids with quantum-resistant hybrid encryption. *Sci. Rep.* **2025**, *15*, 3. [CrossRef] [PubMed]
6. Dhar, S.; Khare, A.; Dwivedi, A.D.; Singh, R. Securing IoT devices: A novel approach using blockchain and quantum cryptography. *Internet Things* **2024**, *25*, 101019. [CrossRef]
7. Bernstein, D.J.; Buchmann, J.; Dahmen, E. *Post-Quantum Cryptography*; Springer: Berlin/Heidelberg, Germany, 2009. [CrossRef]
8. National Institute of Standards and Technology (NIST). Post-Quantum Cryptography Standardization Project. Available online: https://csrc.nist.gov/projects/post-quantum-cryptography (accessed on 18 August 2025).
9. Ducas, L.; Lepoint, T.; Lyubashevsky, V.; Schwabe, P.; Seiler, G.; Stehle, D. CRYSTALS-Dilithium: Digital signatures from module lattices. In Proceedings of the 3rd IEEE European Symposium on Security and Privacy, London, UK, 24–26 April 2018. Available online: https://eprint.iacr.org/2017/633 (accessed on 18 August 2025).
10. Fernandez-Carames, T.M. From Pre-Quantum to Post-Quantum IoT Security: A Survey on Quantum-Resistant Cryptosystems for the Internet of Things. *IEEE Internet Things J.* **2020**, *7*, 6457–6480. [CrossRef]
11. Castiglione, A.; Esposito, J.G.; Loia, V.; Nappi, M.; Pero, C.; Polsinelli, M. Integrating Post-Quantum Cryptography and Blockchain to Secure Low-Cost IoT Devices. *IEEE Trans. Ind. Inform.* **2025**, *21*, 1674–1687. [CrossRef]
12. Nakamoto, S. Bitcoin: A Peer-to-Peer Electronic Cash System. 2008. Available online: https://bitcoin.org/bitcoin.pdf (accessed on 18 August 2025).
13. Crosby, M.; Nachiappan; Pattanayak, P.; Verma, S.; Kalyanaraman, V. Blockchain technology: Beyond bitcoin. *Appl. Innov. Rev.* **2016**, *2*, 6–10. Available online: https://scet.berkeley.edu/wp-content/uploads/AIR-2016-Blockchain.pdf (accessed on 18 August 2025).
14. Buterin, V. A Next-Generation Smart Contract and Decentralized Application Platform. Ethereum White Paper. 2014. Available online: https://ethereum.org/en/whitepaper/ (accessed on 18 August 2025).
15. Hyperledger Architecture Working Group. Hyperledger Architecture Volume I. 2017. Available online: https://www.hyperledger.org/ (accessed on 18 August 2025).
16. Dai, H.-N.; Zheng, Z.; Zhang, Y. Blockchain for Internet of Things: A Survey. *IEEE Internet Things J.* **2019**, *6*, 8104–8124. [CrossRef]
17. Conti, M.; Kumar, E.S.; Lal, C.; Ruj, S. A survey on security and privacy issues of Bitcoin. *IEEE Commun. Surv. Tutor.* **2018**, *20*, 3416–3452. [CrossRef]
18. Iovane, G.; Di Gironimo, P.; Chinnici, M.; Rapuano, A. Decision and Reasoning in Incompleteness or Uncertainty Conditions. *IEEE Access* **2020**, *8*, 115109–115125. [CrossRef]
19. Good, J. *Probability and the Weighing of Evidence*; Charles Griffin: London, UK, 1950. [CrossRef]
20. Dempster, A.P. Upper and Lower Probabilities Induced by a Multivalued Mapping. *Ann. Math. Stat.* **1967**, *38*, 325–339. [CrossRef]
21. Smarandache, F.; Dezert, J. *Advances and Applications of DSmT for Information Fusion*; American Research Press: Santa Fe, NM, USA, 2009; Volume 2. Available online: https://fs.unm.edu/DSmT-book1.pdf (accessed on 18 August 2025).
22. Hussain, T.; Nugent, C.; Moore, A.; Liu, J.; Beard, A. A Risk-Based IoT Decision-Making Framework Based on Literature Review with Human Activity Recognition Case Studies. *Sensors* **2021**, *21*, 4504. [CrossRef] [PubMed]
23. Iovane, G. MuReQua Chain: Multiscale Relativistic Quantum Blockchain. *IEEE Access* **2021**, *9*, 39827–39843. [CrossRef]
24. Brassard, G.; Broadbent, A.; Tapp, A. Quantum Pseudo-Telepathy. *Found. Phys.* **2005**, *35*, 1877–1907. [CrossRef]
25. Iovane, G.; Amatore, R. Decentralized Storage and Security Engine Using Information Fusion Based on Stochastic Processes and Quantum Mechanics. *Appl. Sci.* **2025**, *15*, 759. [CrossRef]
26. Vidaković, M.; Miličević, K. Performance and Applicability of Post-Quantum Digital Signature Algorithms in Resource-Constrained Environments. *Algorithms* **2023**, *16*, 518. [CrossRef]
27. Pornin, T. New Efficient, Constant-Time Implementations of Falcon. IACR Cryptol. ePrint Arch., Paper 2019/893. September 2019. Available online: https://eprint.iacr.org/2019/893 (accessed on 18 August 2025).
28. Kerimbayeva, A.; Iavich, M.; Begimbayeva, Y.; Gnatyuk, S.; Tynymbayev, S.; Temirbekova, Z.; Ussatova, O. A Lightweight Variant of Falcon for Efficient Post-Quantum Digital Signature. *Information* **2025**, *16*, 564. [CrossRef]
29. Iovane, G. Quantum-Inspired Algorithms and Perspectives for Optimization. *Electronics* **2025**, *14*, 2839. [CrossRef]

30. Allen, C. The Path to Self-Sovereign Identity. Life Alacrity. 2016. Available online: https://www.lifewithalacrity.com/article/the-path-to-self-sovereign-identity/ (accessed on 18 August 2025).
31. Ma, W.; Liu, M.; Deng, R. CDQKL: Consensus-Based Distributed Quantum Kernel Learning. In Proceedings of the 2024 International Conference on Quantum Communications, Networking, and Computing (QCNC), Kanazawa, Japan, 1–3 July 2024; pp. 33–40. Available online: https://ieeexplore.ieee.org/document/10628264 (accessed on 18 August 2025).
32. Ma, W.; Chen, K.C.; Yu, S.; Liu, M.; Deng, R. Robust decentralized quantum kernel learning for noisy and adversarial environment. *arXiv* **2025**. [CrossRef]
33. *BS EN ISO 15118-8:2020*; BSI Standards: London, UK, 2020. Available online: https://www.en-standard.eu/bs-en-iso-15118-8-2020-road-vehicles-vehicle-to-grid-communication-interface-physical-layer-and-data-link-layer-requirements-for-wireless-communication/?msclkid=5efebb47f269120ee531d3e10c3126f5 (accessed on 18 August 2025).