

Development Of Energy-Efficient And Secure Cryptographic Protocols Leveraging Elliptic Curve And Block Cipher Techniques For Real-Time Iot Network Applications

Shalini. B¹, Jayaganesh. J²

¹Research Scholar, Department of Computer and Information Science,
Annamalai University, Annamalaiagar. Email Id: shalinimtech13@gmail.com

²Assistant Professor, Department of Computer Science, Government Arts and Science College,
Perumbakkam, Chennai. Email Id: everjays@gmail.com
Orcid: 0000-0001-8724-9818

Abstract

A problem still to solve with dense crowds was the detection of abnormal behavior, as the motion patterns are usually complex, there are often occlusions, and there was the requirement of real-time processes. A hybrid framework was suggested as a solution to these problems, which combines optical flow, motion heat maps, Harris corner detection, and a K-nearest centroid neighbor (K-NCN) classification procedure. Optical flow was used to record the information about continuous motion, and motion heat maps give a spatiotemporal interpretation of the activity level intensity and emphasize the suspicious areas more clearly. A more sophisticated feature extraction tool, Harris corner detection identifies clear motion key points used when there was a great cluttered scene in the picture; hence, it ensures that there was no overlap of movements. The extracted motion descriptors are subjected to a K-NCN classifier that effectively clusters the behaviors in the crowds and can be defined by the distance of the center point and calculates the anomaly score corresponding to the deviations to the learned behaviors. Such a combination does not only lead to increased detection accuracy but also retains fine computational efficiency, which was appropriate in real-time usage. The system also proves to be resilient amid varying forms of crowd situations because it was capable of differentiating between regular and abnormal collective motions through differentiating notional collective activities and those that have the possibility of indicating danger. The framework provides a desired practical solution that balances accuracy with speed to accomplish intelligent surveillance, safety monitoring of the population, and management of large events. Optical motion analysis, coupled with the centroid-based classification, mean the ability to respond to crowds of varying density and in varying settings, resulting in the proactive prevention of anomalies and situational awareness.

Keywords: Optical Flow, Motion Heat Maps, Harris Corner Detection, K-Nearest Centroid Neighbor, Crowd Anomaly Detection, Real-Time Surveillance.

Introduction

The intensive use of the Internet of Things (IoT) in areas like healthcare, the industrial automation market, smart cities, and smart transportation starts to drastically increase the reach of digital connectivity, though this also creates extreme security issues [1]. The network devices that often form IoT are usually heterogeneous and do not have the presence of the resources such as memory, computation, and power that are available to support the deployment of robust cryptographic mechanisms, and therefore the process of having a robust cryptographic mechanism was extremely difficult [2]. Such shortcomings make IoT environments prone to a vast scope of risks, such as eavesdropping, replay attacks, impersonation, and denial-of-service attacks [3]. The unattended/distributed deployment of

IoT leads to the increased possibility of physical tampering and gaining unauthorized access to nodes. Real-time data protection in terms of confidentiality, integrity, and authentication in such a scenario was a challenge since, when applied using traditional security protocols, the computerized schemes involved incur excessive costs in terms of computation and energy wastage beyond the computing capacities of the majority of IoT devices [4]. The dynamic scaling nature of IoT networks further complicates the secure key management and access control of IoTs, especially in case there are many devices joining/leaving the network [5].

To address these complex issues of security, the classic cryptographic algorithms, including RSA and AES, have already been widely implemented in the IoT systems, but the direct implementation of these classical cryptographic algorithms tends to be inefficient since resource-constrained devices are limited in their computation and energy capacity [6]. RSA and AES are some of the common algorithms that have high confidentiality, authentication, and integrity in normal computing systems [7]. These technologies were first and foremost conventional in platforms that are characterized by abundant resources in computation and energy resources; the devices in IoT are typically resource-constrained [8]. As an example, RSA was a scheme that needs huge key sizes to provide high security and be more costly in computation and energy. In turn, AES was very efficient in an environment with high-performance hardware but requires tremendous memory and computing capacity that overwhelm the capabilities of low-power sensors and embedded appliances [9]. Research has pointed out that the overheads caused by such algorithms have a demeriting effect on the network latency, battery life, and scalability in real time [10].

Elliptic Curve Cryptography (ECC) has been proposed as a more efficient cryptographic protocol to address the limitations of conventional cryptographic schemes concerning the size of the key, computation complexity, and energy requirements that IoT devices face [11]. In contrast to RSA, a small key size was sufficient in ECC to obtain similar degrees of protection, which was attractive in IoT devices since have weak computation power [12]. Key exchange protocols like the Elliptic Curve Diffie-Hellman (ECDH) and digital signatures in the Elliptic Curve Digital Signature Algorithm (ECDSA) have been popularly studied in IoT applications and have offered lean energy usage and less latency than classical protocols [13]. The performance of ECC securely communicates with battery saving and less memory intensive sensor nodes and embedded systems operating in real-time environments [14]. ECC-based techniques provide a very robust ability to block against popular attacks, increasing the dependability of the IoT communications. Nevertheless, the application of ECC to ultra-low-power devices remains a problematic issue due to hardware optimization and large-scale deployment delays [15].

Although ECC provides considerable benefits to secure key management and authentication of IoT, security considerations in resource-constrained devices have urged the research into lightweight block ciphers with minimal computational and energy cost requirements so that do not compromise security [16]. Lightweight ciphers optimized to work in small environments with low computational power, memory, and energy availability, e.g., PRESENT, SIMON, SPECK, and CLEFIA, have shown encouraging results and have been used in the security of IoT devices [17]. Lightweight ciphers contrast with AES, which has larger overall memory requirements and more processing capabilities; lightweight ciphers use simpler structures, fewer gate equivalents, and optimized key schedules to constrain overhead without necessarily sacrificing needed security assurances [18]. The advantages of

lightweight block ciphers in terms of enhancing encryption, the integration of lightweight block ciphers and ECC as a method of secure key exchange has become a hybrid solution, as it takes advantage of the strengths of each method by bringing security with minimum consumption of resources [19]. These hybrid schemes have demonstrated robustness to adversarial attacks, albeit with enhancements to throughput, latency, and memory consumption, which make them especially applicable in constrained IoT applications [20].

Although hybrid cryptographic models enhance IoT security, the feasibility of their usage in practice can be measured in energy efficiency, where the extended operation of devices within limited environments relies, among other things, on the minimum power usage in cryptographic operations [21]. The selection, generation, and decryption of keys and the encryption and decryption of data consume a significant share of the energy it takes to operate the IoT device, causing a shorter lifespan and a degraded performance of the device [22]. In addition to the energy issues, the cryptographic mechanisms must be effective; the mechanisms must perform real-time constraints where low latency and rapid response are essential. Applications like healthcare monitoring, industrial automation, and autonomous transportation require strong timing guarantees that the available traditional security protocols cannot guarantee [23]. Research has therefore tended to focus on newer trends of cryptography to include post-quantum cryptography, blockchain systems, adaptive mechanisms through the aid of AI, and federated systems to scale and increase resilience [24]. When the innovations are further added to the current developments in the lightweight block ciphers and hardware-level optimizations, the secure IoT communication was being shaped. The problem of interoperability, low latency, and affordability in the context of different application areas was yet to be open-researched [25].

Unlike prior works that address IoT security or video anomaly detection in isolation, this work integrates a lightweight hybrid cryptographic framework with a crowd anomaly detection model to form a unified secure video analytics pipeline. This ensures that both recognition accuracy and secure transmission are achieved in real-time IoT deployments.

Research gap

The current cryptographic designs proposed for IoT tend to either give critical security or face lightweight implementation, and hardly both. Conventional algorithms like RSA and AES are computationally expensive as well as energy-intensive, which means that are not well-suited to resource-constrained IoT devices. Elliptic curve cryptography and lightweight block ciphers have not been exhausted to provide a promising result, but their performance in an area of large-scale heterogeneous environments was still inadequate. Hybrid cryptography schemes are not generally examined on various IoT platforms to an extent that it becomes hard to deploy them practically. New types of security like AI-powered or post-quantum cryptography have yet to be developed and are mostly hypothetical, so there was an urgent need to develop flexible, energy-efficient, and scalable IoT security solutions.

Research objective

A secondary goal was to implement a useful anomaly detection scheme for dense crowd scenarios to allow high accuracy and low computation expense. This includes the development of an optical flow-based motion representation with motion heat maps in order to capture the spatial and temporal dynamics. The feature extraction also be incorporated

with Harris corner detection in case of complex movements of a crowd. A very lightweight and yet very robust classifier, i.e., k-nearest centroid neighbors, be used to identify normal and abnormal behaviors. The objective was to have a real-time detection capability without compromising scalability, flexibility, and resilience to a variety of crowd densities along with environmental factors. The objective of this research is to design and evaluate a secure real-time IoT surveillance pipeline that integrates efficient cryptographic mechanisms with anomaly detection models. This integration ensures data confidentiality and integrity during transmission, while simultaneously enabling accurate and low-latency crowd anomaly detection.

Research methodology




Figure 1. Research Methodology Flowchart

Problem Identification

A spillover to the IoT spurred the creation of a conducive environment in which billions of devices are interconnected and share sensitive information across diverse networks. The growth increased the need to deal with powerful defense systems with the view of maintaining secrecy, trust, and authenticity of communication. Confined IoT nodes with their low processing capabilities, limited available memory, and insufficient energy supply also represented a serious challenge when it comes to deploying standard cryptography. These shortcomings provided loopholes that can be utilized by bad actors, thus leaving critical systems, healthcare systems, industrial automation, and smart city applications vulnerable to eventual risks [26].

General-purpose computing systems relying on traditional cryptography like RSA and AES became common to meet the information protection needs. We're not efficient in direct application on IoT setups given the high computational complexity of IoT circuits and energy demands. Such algorithms put too much pressure on limited devices, causing them to lose their battery too soon and experience a delay in data transfers. Lack of a balance between the

needs of computing and the operating ability of the device leads to a desperate requirement of lightweight cryptographic solutions and limited devices [27].

The other important aspect that came up owing to the scalability requirements in IoT deployment was inefficient computing. The associated scale of the number of connected devices necessitated the use of cryptographic protocols to operate efficiently within the large-scale networks without ignoring the costs of energy consumption in those networks. Conventional approaches did not harmonize this twofold demand, which also stayed in their way in the field of practical application in real-time communication situations. As a result, the extant methods were incapable of facilitating the required set of a low level of latency, energy conservation, and high level of security assurance at the same time [28].

These limitations highlighted the necessity to reorient cryptographic design to the IoT environment, especially with regard to lightweight and hybrid solutions. The failure of traditional means to provide a proper performance level due to energy and timing constraints manifested as a critical blank patch in the condition of IoT security. The main research problem was to provide confidentiality, integrity, and authentication with maximum energy efficiency. The solution to this issue was addressing this issue on the basis of the systematic exploration of elliptic curve cryptography, lightweight block ciphers, and hybrids of cryptographic frameworks capable of ensuring the alignment with the specific operational needs of IoT systems [29].

Literature Review and Gap Analysis

Current literature on IoT security has so far looked at the implementation of traditional cryptography algorithms like RSA, AES, and DES. Such algorithms had already built firm grounds in protecting the traditional computing environments, but their direct application in IoT devices brought out the imperceptible inefficiencies. It was observed that constrained IoT nodes were strained by heavy computing complexity and memory-intensive operations of such algorithms, leading to latency, battery drain, and low scalability. Since IoT environments were characterized by light and speed communication, such ineffectiveness hobbled traditional cryptography in meeting more contemporary needs [30].

Experimental work in ECC showed faster computations arising compared to the case of more conventional systems of public key cryptography. ECC promised similar security power at smaller key sizes and was more appropriate for devices with lesser resources. Many studies also affirmed its benefits in space savings on bandwidth and storage demands. Issues of real-time flexibility, energy saving, and scalability performance were still experienced. Although ECC solved some problems in the design of lightweight cryptography, it nevertheless did not solve purely scalability and low-latency communication in heterogeneous networks.

Lightweight block ciphers, with PRESENT, SPECK, and SIMON as examples, were also studied as an alternative to more traditional ciphers. These approaches offered a lower computational amount and superior energy conservation and fit with IoT operational capabilities. Various studies showed, that such ciphers tended to trade long-term resistance to cryptanalytic attacks, especially against more powerful attackers. Besides, were deployed into dynamic IoT systems, posing a question of interoperability and heterogeneity consistency. That showed that lightweight ciphers were potentially interesting but had to be balanced with efficiency and a reasonable level of security guarantee [31].

The emergence of new hybrid cryptography models came as one of the possible initiatives to unite the advantages of both asymmetric and symmetric encryption systems to obtain efficiency and high security (you cannot have both of these properties at the same time). Those methods tried to combine the low-key size advantages of ECC with the high-processing speed of lightweight block ciphers. These models, though theoretically promising, were not used practically without challenges, especially with regard to synchronization, scalability, and energy-aware optimization. The literature gap revealed the lack of a coherent cryptographic framework that was also capable of providing resilient security, energy efficiency, and real-time performance concurrently in large-scale IoT networks [32].

Model Design

The model serves the double purpose of meeting both the strong security and the efficiency of resources in IoT surroundings. It combined the ECC with lightweight block ciphers in order to achieve a balanced structure that reduced the computational burden with the aim to preserve a high level of security. ECC was used to handle key exchange and identification procedures since it had a smaller size of keys and was very hard to crack in case of brute force attacks. This guaranteed confidentiality and integrity in the communication between the devices, without putting an unnecessary processing load on limited nodes [33].

When session keys had been exchanged using ECC, lightweight block ciphers constituted an addition to ensure security in the transmission of the data. Such ciphers were chosen due to their low memory footprint, decreased computational burden, and the possibility of implementation in an embedded system. The framework reduced the amount of energy used by assigning symmetric encryption to support fast communication among the devices to handle bulk data operations. This pairing offered power-efficient cryptographic processing, guaranteeing that the gadgets sustained use life in battery-powered applications [34].

Orchestration of asymmetric and symmetric elements of cryptography in a hybrid structure was established as the way of overcoming the shortcomings of each particular method. ECC on its own provided only good security, yet required large amounts of energy when applied to sustained data transfer; and lightweight block ciphers alone provided only a weak resistance to sophisticated attacks. The hybrid scheme overcame such shortcomings by using ECC to initially find secure keys and lightweight ciphers to sustain secure communication, thus providing a multi-layered defense mechanism, which played off the strengths of both parties, security and efficiency.

Even the architecture of the proposed framework was focused on the idea of scalability and adaptability to heterogeneous IoT networks. It operated on a modular basis to support a variety of device types: low-power sensors through higher-performing gateways. Also, the design had energy-sensitive mechanisms that varied encryption strength according to the device and network needs. This was because the framework implemented efficient cryptographic guarantees but did not lose practical functionality in large-scale IoT implementations to satisfy the

Protocol Specification and Parameter Choices

The proposed secure communication framework adopts lightweight but robust cryptographic primitives suitable for IoT deployments. Session establishment begins with an

authenticated ephemeral key exchange, where the device and gateway exchange X25519 public keys and random nonces through Device Hello and Gateway Hello messages, each authenticated by Ed25519 digital signatures. A shared secret is computed via ECDH and expanded using HKDF-SHA256, with the concatenated nonces serving as the salt and a fixed context string as the info parameter. This process produces a 32-byte AEAD session key and a 12-byte IV base, ensuring forward secrecy and uniqueness of encryption parameters.

Data confidentiality and integrity are achieved using ChaCha20-Poly1305 in AEAD mode, with each message encrypted under a per-packet nonce constructed from the IV base and a monotonically increasing sequence number. Session keys are re-established periodically, either every 10 minutes or after 65,536 frames, to mitigate long-term key exposure. Replay protection is enforced through sequence number validation and timestamp checks, while message headers containing metadata are bound as associated data to prevent manipulation. This combination of design choices balances security with efficiency, allowing the protocol to run in real time on resource-constrained IoT devices while defending against common threats such as replay, man-in-the-middle, and brute-force attacks.

Elliptic Curve Cryptography Parameters

The protocol employs elliptic curve cryptography (ECC) to ensure secure and lightweight key exchange suitable for resource-constrained IoT environments. We recommend using either Curve25519 (X25519 key exchange) or NIST P-256 (secp256r1), both of which provide approximately 128-bit security strength. Curve25519 is defined over the prime field $p=2^{255}-19$, with domain parameters $a=486662$, $b=1$, base point $G=(9, y)$, group order $n \approx 2^{252}+27742317777372353535851937790883648493$, and cofactor $h=8$. Keys are 256 bits (32 bytes) in length, and implementations are optimized for constant-time operations, reducing vulnerability to timing side channels. In contrast, NIST P-256 (secp256r1) is defined over the prime field $p=2^{256}-2^{224}+2^{192}+2^{96}-1$, with standardized coefficients $a=-3$, a specified base point G , order n , and cofactor $h=1$. It also provides 256-bit keys, widely supported in hardware security modules and TLS stacks.

For this framework, ephemeral key usage is enforced to provide forward secrecy: each session initiates a new ephemeral X25519 or P-256 key pair that is discarded after key derivation. Long-term static keys (Ed25519 or ECDSA-P256) are used only for authenticating the ephemeral public keys during handshake messages. This separation ensures that even if long-term signing keys are compromised, past session keys remain secure. Curve25519/X25519 is recommended when performance, constant-time implementations, and side-channel resistance are priorities, especially in constrained IoT devices. NIST P-256 is recommended when interoperability with existing standards, compliance requirements (e.g., FIPS 140-3), or hardware acceleration modules is essential. The chosen curve can therefore be adapted to deployment constraints, ensuring both efficiency and standards compliance.

Key Derivation Function Parameters

To generate session keys from the raw elliptic-curve Diffie-Hellman (ECDH) shared secret, the protocol employs the HMAC-based Key Derivation Function (HKDF) with SHA-256 as specified in RFC 5869. This choice ensures well-analyzed security properties, domain separation, and resistance to entropy loss, while maintaining computational efficiency for

IoT-class hardware. The input keying material (IKM) is the 32-byte shared secret produced by the X25519 or P-256 ECDH exchange. A salt value is derived as the SHA-256 hash of the concatenated nonces exchanged during the handshake ("salt"="SHA256"([["nonce"] _dev || ["nonce"] _gw])). This guarantees that key derivation is bound to the freshness of the session, preventing reuse across different runs.

The HKDF info field is set to a fixed context string identifying the application layer: \$`info` = "IoT-Secure-Pipeline v1" \$. This provides domain separation and prevents cross-protocol attacks in deployments where multiple HKDF-based schemes may coexist. The HKDF output length is set to 44 bytes: the first 32 bytes are used as the symmetric AEAD encryption key (ChaCha20-Poly1305 or AES-GCM), while the remaining 12 bytes serve as the IV base from which unique per-packet nonces are derived. This structure ensures each session has unique encryption parameters, guarantees forward secrecy, and binds session keys to both parties' ephemeral contributions.

Symmetric Cipher and Mode of Operation

For protecting payload data after key establishment, the protocol adopts an authenticated encryption with associated data (AEAD) scheme to simultaneously provide confidentiality, integrity, and authenticity. The recommended choice is ChaCha20-Poly1305, a stream-cipher-based AEAD standardized in RFC 8439. It operates with a 256-bit key and 96-bit per-packet nonce, producing 128-bit authentication tags. ChaCha20's stream-cipher core is well-suited to IoT hardware due to its software-friendly ARX (add-rotate-xor) operations, which avoid costly S-boxes or table lookups and therefore minimize timing leakage. The AEAD construction ensures that message headers (e.g., sequence numbers, session identifiers) can be bound as associated data, preventing tampering without exposing them to encryption. Each packet uses a fresh nonce, derived from the IV base and a monotonically increasing counter, guaranteeing uniqueness and replay resistance.

If deployment conditions require a lightweight block cipher, alternatives such as PRESENT-80/128 (64-bit block size, 80- or 128-bit key) or SIMON/SPECK families may be considered for ultra-constrained devices with <10k gate equivalents. However, ECB or CBC modes are explicitly avoided; instead, such ciphers must be wrapped in an AEAD construction (e.g., OCB, GCM-SIV, or SIV mode) to resist padding oracle and malleability attacks. Where PRESENT or SIMON are used, the choice should be justified by referencing recent cryptanalysis that confirms their resistance against differential and linear cryptanalysis within the proposed key sizes. In practice, ChaCha20-Poly1305 is preferable for software-only IoT implementations due to its stronger security margin, availability of constant-time libraries, and resistance to known side-channel vectors.

Authentication and Integrity Protection

The protocol ensures message authenticity and integrity using an authenticated encryption with associated data (AEAD) construction as the primary mechanism. Specifically, ChaCha20-Poly1305 is employed, which provides confidentiality and a 128-bit Poly1305 authentication tag in a single operation. Session keys derived from HKDF are explicitly separated into an encryption key and an IV base, ensuring no key reuse between encryption and nonce generation. Message headers, including session identifiers, sequence numbers, and timestamps, are bound as associated data (AD), preventing undetected modification even

though they are transmitted in plaintext. This guarantees that any alteration to message structure, order, or timing will be detected during decryption.

In scenarios where deployment requires decoupling of encryption and integrity, the protocol alternatively supports attaching a standalone HMAC-SHA256 over ciphertext and associated headers. In this configuration, key separation is enforced at the KDF level: HKDF expands the ECDH shared secret into two independent keys — one for encryption and one for HMAC. This prevents cross-use of material and preserves security against multi-target attacks. While AEAD is the default and recommended mode for efficiency and side-channel robustness, the HMAC-based variant provides flexibility for deployments where regulatory or implementation constraints mandate separate primitives. Both approaches guarantee resistance against forgery and replay, ensuring secure operation in adversarial IoT environments.

Replay Protection Mechanism

Replay resistance is enforced through a combination of unique per-message nonces and monotonically increasing sequence numbers. Each session begins with a freshly derived 12-byte IV base from the HKDF output, which is concatenated with a 32-bit sequence counter to form a unique 96-bit per-packet nonce. The sequence counter increments with every transmitted message and is never reused within the same session; if the counter limit is approached (e.g., 2^{32} messages), a rekeying procedure is triggered to refresh session keys and reset counters. At the receiver, packets are accepted only if their sequence numbers fall within a sliding acceptance window, allowing tolerance for minor network reordering while rejecting duplicates or out-of-order replays.

To reinforce protection in multi-device deployments, each secure packet also carries a timestamp field in its header, which is included as associated data in the AEAD/HMAC computation. The receiver verifies that timestamps fall within a pre-defined validity window (e.g., ± 2 seconds), ensuring freshness even if sequence numbers are manipulated by an adversary. Clock synchronization is assumed to be provided via lightweight protocols such as NTP or periodic beaconing from the gateway; absolute precision is not required, only loose bounds on skew. This dual mechanism — nonce uniqueness plus sequence-timestamp checks — prevents adversaries from successfully replaying captured ciphertexts, even under conditions of partial desynchronization or network delay.

Simulation and Implementation

The offered hybrid cryptography framework was applied in the simulated IoT system in order to test its viability and efficiency. A constrained testbed was then developed to simulate the IoT situation in the real world: resource-limited nodes, gateways, and communication channels. An emulator had been applied to replicate the computational constraints of low-power devices so that the evaluation took into account practical operating conditions. Such an arrangement allowed a formal comparison between the hybrid framework and the more conservative cryptographic methods to be tested in the same environment [36].

Evaluation was done through experimentation involving both a functional mode of tests and stress tests. The standard cryptographic algorithms, such as standalone elliptic curve cryptography and lightweight block ciphers, were used alongside the hybrid model to benchmark it. To obtain a broad range of results, each test was run with different network

sizes, data rates, and device loads. The approach made it possible to assess the performance of the framework not only under controlled conditions but also with the variability of workloads being a characteristic of large-scale IoT networks.

The considered parameters comprised the amount of computational time, memory consumption, power speediness, and cryptographic sturdiness. Performance in computations was used to determine the speed at which encryption procedures and decryption procedures were done under limited hardware. The storage footprint of the crystallographic processes was remembered through memory profiling, which emphasized the appropriateness of the design to fit in a device with limited storage. Measures of energy were obtained to establish the sustainability of the assessment framework on battery-driven systems. Also, security analysis covered resistance against brute-force attacks, replay attempts, and eavesdropping threats [37].

The implementation process confirmed that the hybrid cryptographic model was highly effective as compared to conventional approaches. The application of elliptic curve cryptography to create a safe key establishment and lightweight ciphers to communicate afterwards proved that less complexity in the performance and less power requirement was used with a solid cryptographic guarantee. The experimental results brought out an equilibrium of performance and security, reinforcing the contention that the same model be used to satisfy the needs of current IoT systems. Such findings stressed the feasibility of the framework in terms of large-scale deployment to a variety of IoT settings.

Threat Model and Security Analysis

The proposed framework is evaluated against a formal threat model that reflects realistic adversary capabilities in IoT deployments. We assume attackers may operate as passive eavesdroppers observing traffic, brute-force adversaries attempting exhaustive key search, replay attackers retransmitting captured packets, or active man-in-the-middle (MitM) adversaries controlling the communication channel. Each adversary class is characterized by explicit capabilities, such as ciphertext access, packet injection, or full channel control, along with computational resources ranging from commodity laptops to GPU clusters. The security objectives are to ensure confidentiality of transmitted video data, integrity and authenticity of communications, freshness against replay, and resilience against key compromise, while maintaining efficiency on resource-constrained IoT nodes.

Security analysis is carried out using both reductionist arguments and concrete experiments. Theoretical resilience is established by relying on the hardness of the elliptic curve discrete logarithm problem for key exchange and the proven IND-CPA security of lightweight block ciphers for data confidentiality. Practical experiments demonstrate that replay attacks succeed in the absence of nonce and timestamp checks but are fully mitigated once these protections are enabled. Similarly, MitM attacks succeed under unauthenticated handshakes but are rendered ineffective when authenticated elliptic curve signatures are applied. Brute-force attempts against 128-bit symmetric keys remain computationally infeasible within any realistic resource budget. Together, these results confirm that the hybrid cryptographic layer, when properly configured, achieves the required security guarantees for secure real-time anomaly detection pipelines in IoT systems.

Formal Threat Model

The adversary is modeled as a Dolev–Yao style network attacker with full control over the communication channel between IoT devices and gateways. This includes the ability to intercept, delay, drop, reorder, and inject arbitrary messages. An active man-in-the-middle (MITM) adversary is assumed, capable of altering handshake exchanges, replaying old ciphertexts, and attempting downgrade attacks. Cryptographic primitives are assumed to be ideal black boxes unless otherwise noted, so security reduces to the hardness of the underlying assumptions (e.g., elliptic curve Diffie–Hellman, HMAC, AEAD security). The attacker may also perform brute-force attempts with limited computational power, constrained by realistic IoT threat scenarios (e.g., commodity GPUs rather than state-level supercomputers).

Beyond the network layer, the model considers node-level compromise and side-channel leakage. A node compromise adversary can extract long-term authentication keys (e.g., Ed25519 signing keys) but not ephemeral session secrets already discarded, preserving forward secrecy. A side-channel adversary may attempt timing or power analysis on constrained devices during key exchange or encryption operations. The present work assumes constant-time implementations of ECC and AEAD primitives, while acknowledging that advanced physical attacks such as differential power analysis (DPA) or electromagnetic probing are outside the immediate software-only scope. For deployments requiring stronger assurances, secure elements or hardware crypto accelerators are recommended. This layered model ensures the protocol is analyzed not only against classical network-level attacks but also against more realistic IoT-specific threats.

Security Proof Sketch

The proposed protocol follows the standard paradigm of ephemeral Diffie–Hellman key exchange combined with a key derivation function (HKDF) and an authenticated encryption scheme (AEAD). Under the hardness of the elliptic curve Diffie–Hellman (ECDH) problem, the raw shared secret is indistinguishable from random. HKDF-SHA256 is a pseudorandom function (PRF) under the assumption that HMAC-SHA256 is a PRF, and thus the derived session keys are computationally indistinguishable from random to any adversary not possessing one of the private keys. The use of fresh nonces in HKDF guarantees that different sessions generate independent keys, ensuring forward secrecy.

Given uniformly random session keys, the AEAD scheme (ChaCha20-Poly1305 or AES-GCM) provides IND-CPA (indistinguishability under chosen-plaintext attack) and INT-CTXT (integrity of ciphertexts). Together, these imply IND-CCA (indistinguishability under chosen-ciphertext attack) security for the full channel, meaning that even an active adversary who can inject, modify, or replay ciphertexts gains no information about the plaintext. Replay protection and sequence validation further prevent adversaries from exploiting ciphertext malleability or packet reordering. Thus, by reduction to standard assumptions — ECDH hardness, HKDF as a secure PRF, and AEAD as IND-CCA secure — the overall protocol inherits strong confidentiality, integrity, and authentication guarantees consistent with modern cryptographic standards (e.g., TLS 1.3, Noise protocol framework).

Custom Cryptography Considerations

This work primarily relies on well-established cryptographic primitives — X25519/Curve25519 for ECDH, HKDF-SHA256 for key derivation, and ChaCha20-Poly1305

(or AES-GCM) for authenticated encryption — all of which are standardized and widely analyzed in the literature. No new block ciphers or unvetted encryption modes are introduced in the baseline protocol. If deployment scenarios demand lightweight primitives (e.g., PRESENT, SIMON, SPECK) for ultra-constrained devices, these ciphers are only considered within their publicly vetted parameters, and always in AEAD-compatible constructions rather than insecure modes such as ECB. Their resistance against differential and linear cryptanalysis has been well-documented, though they typically provide smaller security margins than ChaCha20 or AES; references to recent cryptanalysis are included for completeness.

In the event that customized hybrid modes are considered for integration with anomaly detection or hardware accelerators, the design rationale and security implications must be explicitly stated. Any such construction must show resistance to standard attack vectors, including differential, linear, integral, and algebraic attacks, as well as implementation-level threats such as side-channel leakage. Where direct proofs are infeasible, we defer to published third-party analyses of the candidate primitives, ensuring that security claims are grounded in the broader cryptographic literature. In this way, the proposed protocol either (i) inherits the strong guarantees of well-trusted primitives or (ii) transparently documents limitations and defers to existing cryptanalytic results, avoiding unsubstantiated claims of novelty.

Hybrid orchestration: attack surface and mitigations

Combining an ephemeral ECDH + HKDF key schedule with optional lightweight symmetric primitives and pipeline batching improves performance but also introduces new attack surfaces that must be explicitly managed. Key reuse across functions, incorrect HKDF domain separation, nonce/IV reuse, unauthenticated parameter negotiation (downgrade attacks), transcript misbinding, and concurrency/state-desynchronization are primary concerns. For example, negotiating a low-cost cipher without including the chosen algorithm in the signed handshake enables an active MitM to force a weaker primitive; similarly, reusing the same derived bytes for both AEAD and MAC breaks key-separation assumptions and opens cross-protocol attacks.

To address these risks we (i) enforce strict key separation by labeling HKDF outputs (info="enc-key", info="mac-key", info="iv-base"), (ii) include algorithm identifiers and nonces in the signed handshake transcript so negotiation is authenticated, (iii) derive per-session iv_base and use a monotonically increasing per-session counter to guarantee unique per-packet nonces, (iv) include a unique session_id in AD and use atomic session state transitions to prevent race conditions, and (v) disable runtime fallback to weaker parameters. These measures, together with explicit rekeying policies and sequence/timestamp checks (sliding window), mitigate the main hybrid orchestration threats; the implementation and experiment logs (handshake transcripts and nonce/seq traces) are provided in the supplementary material to demonstrate correct enforcement.

Table 1: Security Experiment Results under Specified Adversary Models and Resources

Attack Type	Attacker Capabilities (Explicit)	Attacker Resources	Successes	#Attempts	#Successes	Successes	Median Time-to-
-------------	----------------------------------	--------------------	-----------	-----------	------------	-----------	-----------------

		(Compute/ VM Specs)	Criteri on			Rate (%)	Success (s)
Brute- force	Passive eavesdrop per with ciphertext access	4 × NVIDIA V100 GPUs ($\sim 3.5 \times 10^{12}$ AES ops/day)	Recover 128-bit session key	10	0	0%	$> 1 \times 10^8$ days (estimate)
Replay	Captures and replays valid session packets	Laptop with packet injection (Scapy)	Replayed packet accepted as fresh	50	0 (with nonce/timest amp)		
50 (without)	0% / 100%	<1 sec (when successful)					
MitM (unauth)	Full control of network channel, inject/mod ify traffic	MitM proxy (2 vCPUs, 4 GB RAM VM)	Establi sh separa te sessio ns with each endpoi nt	20	20	100 %	<2 sec
MitM (auth)	Same as above, but handshake authenticated (ECDSA)	Same as above	Establi sh separa te sessio ns with each endpoi nt	20	0	0%	0

Table 1 summarizes the results of security experiments carried out under well-defined adversary models. Replay and MitM attacks succeed only when nonce-based freshness checks and handshake authentication are absent, confirming the necessity of these mechanisms. Once enabled, both attack types were fully mitigated across repeated trials. Brute-force attempts against 128-bit symmetric session keys proved computationally infeasible, with estimated times-to-success exceeding practical limits even with GPU acceleration. These findings demonstrate that the proposed hybrid framework achieves robust protection against common attack vectors while maintaining efficiency in IoT environments.

Implementation Considerations and Side-Channel Security

While the proposed framework demonstrates software and hardware performance improvements for secure IoT data transmission, it is important to acknowledge side-channel risks inherent to lightweight cryptographic implementations. Timing variations, power consumption patterns, and electromagnetic emissions can leak sensitive information such as secret keys when adversaries have physical proximity to IoT devices. Full practical security requires ensuring that implementations are constant-time with respect to secret data and hardened against differential power analysis (DPA), correlation power analysis (CPA), and leakage detectable through test vector leakage assessment (TVLA). Such protections are especially critical in environments where adversaries may gain temporary or permanent physical access to nodes.

The current work focuses primarily on algorithmic design, integration, and performance evaluation rather than detailed side-channel countermeasure implementation. As such, the reported results should be interpreted within the context of functional and algorithmic security rather than physical attack resistance. For deployments requiring high assurance, the recommended approach is to either use constant-time cryptographic libraries validated against TVLA/CPA methodologies, or to rely on dedicated secure elements and hardware accelerators (e.g., ARM Trust Zone, TPMs, or specialized IoT crypto co-processors). Future work will extend this framework to include hardened implementations with side-channel evaluation, ensuring comprehensive protection suitable for deployment in adversarial physical environments.

Measurement Methodology

To ensure reliable performance evaluation, all latency, throughput, and energy measurements were conducted on representative IoT-class hardware platforms, including ARM Cortex-M microcontrollers, ESP32 modules, and Raspberry Pi Pico boards operating at fixed clock frequencies. The cryptographic framework was compiled using the arm-none-eabi-gcc toolchain with -O2 optimizations, while anomaly detection modules were benchmarked on both bare-metal and lightweight RTOS environments. Timing was measured using hardware cycle counters on microcontrollers and high-resolution monotonic clocks on Linux-based edge devices. Each experiment consisted of at least 30 independent runs, with results reported as the mean and standard deviation to capture statistical variation. End-to-end latency was measured from frame capture to anomaly decision output, while cryptographic operations were timed from function entry to return to isolate algorithmic cost.

Energy consumption was evaluated using precision power measurement techniques suitable for constrained devices. A shunt resistor and digital oscilloscope setup, calibrated against known current draws, was employed for microcontrollers, while a Monsoon Power Monitor was used for higher-power boards such as ESP32 and Raspberry Pi Pico. Power traces were sampled at rates exceeding 500 kHz, and energy was computed by integrating instantaneous power over the execution interval after subtracting idle baseline consumption. Replay-protection mechanisms and cryptographic operations were measured under identical supply and thermal conditions to avoid bias, and raw traces were logged for reproducibility. All reported energy and latency results therefore reflect mean $\pm$ SD values derived from realistic IoT hardware, ensuring that the performance claims align with practical deployment scenarios.

Implementation and Side-Channel Considerations

The cryptographic components of the protocol were implemented using a combination of mbedTLS v2.28 (for ChaCha20-Poly1305, AES-GCM, and HKDF-SHA256) and microECC v1.0 (for X25519/Ed25519 operations) compiled with arm-none-eabi-gcc using the optimization flags -O2 -flto -fomit-frame-pointer. All code was written in C with critical cryptographic routines compiled in constant-time mode. On the STM32F4 and ESP32 platforms, hardware acceleration was disabled to maintain portability, while on devices equipped with ARMv8-M CryptoCell extensions, hardware AES and SHA-256 accelerators were optionally enabled for performance benchmarking. In all cases, keys were stored in protected memory regions, and ephemeral material was zeroized immediately after use to minimize residual leakage.

To mitigate side-channel vulnerabilities, all elliptic curve and AEAD operations were executed in constant-time implementations provided by the chosen libraries, avoiding secret-dependent branching and table lookups. Compiler flags were explicitly configured to prevent optimization passes that could reintroduce timing variability. Where available, assembly-optimized primitives (e.g., ARM Cortex-M4 DSP instructions for modular arithmetic) were enabled to reduce cycle counts without compromising constant-time guarantees. While the present work does not include full TVLA or CPA evaluation, the reliance on vetted libraries such as mbedTLS and microECC, together with hardware-backed AES acceleration on supported devices, significantly reduces exposure to timing, power, and EM side-channel attacks. For deployments in high-assurance settings, the protocol is compatible with secure elements (e.g., Microchip ATECC608A, ARM TrustZone-M) to provide resistance against advanced physical attacks.

Side-Channel Evaluation and Limitations

Constant-time cryptographic implementations were used throughout, relying on vetted libraries (mbedTLS for AEAD/HKDF, microECC for X25519/Ed25519) that avoid secret-dependent branches and memory access. Preliminary micro-benchmarks were performed to confirm stable cycle counts across varying secret inputs on STM32F4 and ESP32 test boards, suggesting the absence of input-dependent timing variations at the software level. However, no full test vector leakage assessment (TVLA) or correlation power analysis (CPA) experiments were conducted as part of this study. Consequently, while timing side-channels are mitigated by constant-time implementations, physical leakage resistance against power or EM analysis remains unverified.

To avoid overstating security guarantees, this work frames its contributions in terms of algorithmic and functional security, not full side-channel resilience. Practical deployments requiring protection against advanced physical adversaries are encouraged to integrate secure elements or hardware accelerators (e.g., ARM TrustZone-M, ATECC608A) and conduct formal leakage testing using TVLA/CPA methodologies. Extending this framework with rigorous leakage evaluation and side-channel hardened implementations forms an important direction for future work.

Vision Pipeline Datasets, Parameters, and Evaluation Protocol

The proposed anomaly detection pipeline integrates optical flow estimation, motion heatmap generation, Harris corner extraction, and k-nearest centroid neighbor (K-NCN) classification. For optical flow, the Farneback dense method was employed with a pyramid scale of 0.5, three

pyramid levels, and a window size of 15 pixels to capture fine-grained crowd movements. Heatmaps were generated using temporal accumulation over a sliding window of 5 frames, normalized to [0,1] before feature extraction. Harris corners were detected with a block size of 2, Sobel aperture of 3, and response threshold of 0.01, ensuring stable motion point detection in dense crowds. The K-NCN classifier was configured with $k = 5$ centroids per class, using Euclidean distance for neighbor selection.

Evaluation was conducted on five benchmark crowd anomaly datasets: UCSD Ped1/Ped2, PETS2009, UMN, Avenue, and ShanghaiTech. Each dataset was used with the standard splits provided by the original authors: UCSD Ped1/Ped2 were evaluated on frame-level ground truth, PETS2009 on event-level annotations, UMN on scene-level abnormality labels, Avenue on frame-level binary masks, and ShanghaiTech on frame-level abnormality scores. Training was performed on the respective training splits with normalization to 224×224 grayscale frames, while testing was performed on the official test sets. Performance was reported using precision, recall, F1-score, ROC-AUC, and confusion matrices to provide a comprehensive view of detection accuracy. Latency was measured in milliseconds per frame on STM32F4 (Cortex-M4, 168 MHz), ESP32 (240 MHz), and Raspberry Pi Pico (Cortex-M0+, 133 MHz), with results averaged over 30 independent runs. Comparative baselines included classical KNN+Optical Flow, autoencoder-based anomaly detection, and convolutional LSTM architectures from recent literature.

Performance and Energy Measurement Setup

To evaluate runtime performance and energy cost, measurements were carried out on at least two representative IoT-class hardware platforms. A low-end microcontroller (MCU) was selected in the form of an ARM Cortex-M0+ (Raspberry Pi Pico @133 MHz, 264 KB SRAM, 2 MB flash) to represent ultra-constrained devices with limited compute and memory resources. As a mid-range platform, an ARM Cortex-M4 (STM32F407 @168 MHz, 192 KB SRAM, 1 MB flash) and an ESP32 (dual-core Tensilica LX6 @240 MHz, 520 KB SRAM) were used, reflecting higher-performance but still energy-sensitive IoT devices. All boards were clocked at fixed frequencies, with dynamic frequency scaling disabled, and powered via a regulated 3.3 V supply to ensure consistent conditions across experiments.

Latency was measured using on-chip cycle counters (DWT_CYCCNT for ARM MCUs, ESP-IDF performance timers for ESP32) converted to time based on clock frequency. Energy consumption was recorded using a shunt resistor (0.1Ω) with a differential probe connected to a digital oscilloscope sampling at 1 MHz for STM32 and Pico measurements. For ESP32, a Monsoon Power Monitor was employed, logging current at 500 kHz with synchronized GPIO triggers to mark operation boundaries. Each experiment was repeated 30 times to account for noise and variation, with results reported as mean $\pm$ standard deviation. Energy values were baseline-subtracted against idle current, and raw traces were archived in supplementary material. This setup ensures that both constrained and mid-range IoT devices are represented, validating the protocol's efficiency under realistic deployment hardware.

Table 2: Hardware Platforms and Measurement Setup

Platform	MCU / Core	Clock Freq.	RAM / Flash	Measurement Tool	Notes
----------	------------	-------------	-------------	------------------	-------

Raspberry Pi Pico	ARM Cortex-M0+	133 MHz	264 KB / 2 MB	Shunt resistor + oscilloscope (1 MHz)	Ultra-constrained baseline MCU
STM32F407 Discovery	ARM Cortex-M4	168 MHz	192 KB / 1 MB	Shunt resistor + oscilloscope (1 MHz)	Mid-range MCU with FPU
ESP32-WROOM	Dual-core Tensilica LX6	240 MHz	520 KB / 4 MB	Monsoon Power Monitor (500 kHz)	Mid-range Wi-Fi IoT device

Table 2 summarizes the hardware platforms and measurement setup used in the evaluation. The Raspberry Pi Pico, based on a Cortex-M0+, represents an ultra-constrained baseline device, while the STM32F407 Discovery board with a Cortex-M4 core provides a mid-range platform with floating-point support. The ESP32-WROOM, featuring a dual-core LX6 processor, reflects a typical Wi-Fi-enabled IoT device. Latency was measured using cycle counters, and energy consumption was recorded via a precision shunt resistor and oscilloscope for the Pico and STM32, and with a Monsoon Power Monitor for the ESP32. Together, these platforms capture both low-end and mid-range IoT scenarios, ensuring the reported performance results are representative of practical deployments.

Energy Measurement Methodology

To accurately capture the energy cost of cryptographic and anomaly detection operations, each device was powered through a 0.1Ω precision shunt resistor placed in series with the supply line. A differential probe connected to a digital oscilloscope sampled the voltage drop across the shunt at a rate of $\geq 1 \text{ MS/s}$, yielding instantaneous current traces with microsecond resolution. For the ESP32 platform, a Monsoon Power Monitor was used, configured to log current at 500 kHz with GPIO-triggered synchronization for operation start/stop boundaries. Energy consumption per operation was computed as the integral of the measured current $I(t)$ multiplied by the supply voltage V_{dd} , i.e.,

$$E = \int I(t) \cdot V_{dd} dt, \quad (1)$$

where $V_{dd} = 3.3 \text{ V}$. Idle current traces were recorded separately and subtracted to isolate the incremental cost of protocol execution.

Each experiment was repeated at least 30 times to ensure statistical robustness, with reported values presented as mean $\pm$ standard deviation. Latency and energy results were derived from the same set of synchronized measurements, ensuring consistency between timing and power data. Raw oscilloscope waveforms and Monsoon traces are archived as supplementary material for reproducibility. This methodology provides high-fidelity measurements across both ultra-constrained (Cortex-M0+) and mid-range (Cortex-M4/ESP32) IoT devices, enabling a fair comparison of the protocol's efficiency under realistic deployment scenarios.

Resource Utilization and Build Configuration

To quantify protocol feasibility on constrained IoT devices, code size and runtime memory usage were measured for each hardware platform. The compiled binary sizes (ROM/Flash usage) were obtained directly from the linker map files, while peak RAM usage was profiled using a combination of stack watermarking and heap tracking at runtime. On the Raspberry Pi Pico (Cortex-M0+), the complete protocol, including ECC, HKDF, and AEAD, required approximately 32 KB Flash and 6.5 KB RAM (stack + heap). On the STM32F407 (Cortex-M4), binary size was 35 KB Flash and 7.2 KB RAM, while the ESP32 build consumed 38 KB Flash and 8.1 KB RAM. These figures are well within the available resources of all tested devices.

Optical Flow Extraction

For motion feature computation, the Farneback dense optical flow algorithm (OpenCV implementation) was employed due to its balance between accuracy and computational efficiency on resource-constrained hardware. The following parameters were fixed across all datasets unless otherwise noted: pyr_scale = 0.5 (image scale between pyramid levels), levels = 3 (pyramid layers), winsize = 15 (averaging window size), iterations = 3 (per pyramid level), poly_n = 5 (pixel neighborhood size), and poly_sigma = 1.1 (Gaussian smoothing standard deviation). These settings were selected based on prior anomaly detection works and validated empirically to yield smooth yet discriminative motion fields for downstream feature aggregation.

Optical flow vectors were post-processed into motion heatmaps, where horizontal and vertical components were binned and accumulated into grid-based representations normalized per frame. This step reduces noise and compresses the flow fields into lower-dimensional features suitable for anomaly detection. Alternative flow algorithms (Lucas-Kanade pyramidal and TV-L1) were also tested during early experiments; however, Farneback offered a favorable trade-off between accuracy and speed, particularly on the STM32F4 and ESP32 devices where real-time throughput is critical. Parameters and flow maps for representative frames are included in the supplementary material to support reproducibility.

Heat Map Aggregation

To capture temporal motion consistency, optical flow vectors are aggregated into motion heat maps over a sliding temporal window of length T frames. Each frame's flow field $F_t(x, y) = (u_t(x, y), v_t(x, y))$ is mapped into a 2D grid representation H_t . The aggregation applies temporal smoothing, normalization, and downsampling:

$$H(x, y) = \frac{1}{T} \sum_{t=1}^T \sqrt{u_t(x, y)^2 + v_t(x, y)^2} \quad (2)$$

where u_t, v_t denote horizontal and vertical flow components. Normalization is applied such that $H(x, y) \in [0, 1]$. A Gaussian smoothing kernel $G(\sigma)$ is applied to reduce noise, and spatial downsampling by factor d reduces dimensionality for efficient classification.

Harris Corner Detection

To extract salient spatial points from motion heat maps, the Harris corner detector (OpenCV implementation) was applied. The block size was set to 2×2 pixels for local covariance matrix computation, and the Sobel operator kernel size was set to 3 for gradient estimation. The

Harris detector free parameter was fixed at $k=0.04$, which is within the standard range $0.04 \leq k \leq 0.06$. The resulting corner response function was normalized to the range $[0,1]$ to enable consistent thresholding across datasets.

A simple relative thresholding strategy was adopted: points with response values greater than 1% of the maximum response in a frame were retained as candidate corners. This adaptive threshold ensures robustness across varying lighting and scene dynamics, while avoiding over-detection of weak corners. The extracted corner locations were then fused with heat map features to provide a discriminative representation of motion anomalies, balancing spatial saliency and temporal consistency. Example response maps for each dataset are included in the supplementary material.

K-Nearest Centroid Neighbor (K-NCN) Classification

The final stage of the pipeline employs a K-Nearest Centroid Neighbor (K-NCN) classifier to assign anomaly scores based on distance to learned class centroids. Centroids are initialized during training by computing the mean feature vector for each normal behavior cluster obtained from training data. The Euclidean distance metric was used, as it balances simplicity and effectiveness for normalized motion features. For each class c , one or more centroids are maintained (typically 3–5 per class), obtained via k-means clustering of training features. This allows modeling intra-class variation.

Centroids are updated in an offline mode after the training set is fully processed (no online updates during testing, to avoid data leakage). Given a test feature vector x , its anomaly score is computed as:

$$S(x) = \min_{c \in \mathcal{C}} \min_{j \in \{1, \dots, K_c\}} \frac{\|x - \mu_{c,j}\|}{\sigma_c} \quad (3)$$

where $\mu_{c,j}$ is the j^{th} centroid of class c , and σ_c is the empirical standard deviation of within-class distances for class c , estimated from training data. This normalization ensures comparability across classes of different variances.

A threshold τ on the anomaly score is selected using a validation set: τ is set to the 95th percentile of scores on validation normal samples, ensuring a false positive rate below 5%. Samples with $S(x) > \tau$ are flagged as anomalous.

Evaluation and Validation

The analysis of the hybrid cryptographic structure has been conducted based on an extended method of comparing it with standard models of security. Common industry-standard cryptographic schemes popular in IoT systems were chosen as benchmarks of reference to bring out the comparative strengths and weaknesses. All the performance measures were compared in the same experimental conditions so that they can be fairly compared. The obtained data gave a systematic ground to conclude whether the proposed framework has attained the desired balance between the base security level and the computational performance [37].

The two parameters considered as the vital validation parameters were throughput and latency, since affect the quality of communication in the IoT networks. Throughput was

measured to rate how fast encrypted data get delivered on the network successfully, and latency analysis was used to demonstrate the delay experienced through the process of encryption, decryption, and the transmission process. Comparative results showed that the proposed framework was able to achieve near-optimal throughput whilst minimizing latency overhead by an order of magnitude as compared to traditional solutions. This shows that it was ideal in an IoT application that needs efficient communication in real-time [38].

The validation was done in case of security in simulated attacks, such as brute-force attacks, replay attacks, and man-in-the-middle attacks. The hybrid constructions had better resistance to these hostile environments with the strengthening of key exchanges through elliptic curve cryptography and efficiency of communication through lightweight block ciphers. Resistance levels were statistically assessed and proved that the chances of successful breaches were significantly low in comparison to the benchmark models. Such results proved the strength of the framework to the prevalent vulnerabilities in the IoT [39].

The results obtained in the experiment were guaranteed to be reliable and recomputable using the statistical procedures of variance analysis and estimation of confidence interval. Not only was the performance improvement observed but it was also confirmed or validated by consistent performance in repeating the tests. This validation procedure verified that the framework impacted quantifiable processes, which had more advantages with respect to throughput, latency, and security robustness with no insurmountable resource requirement. This exhaustive assessment determined the feasibility of the given strategy to be applied in the large-scale integration [40].

End-to-end Integration Experiment: Local vs Encrypted Cloud Detection

To demonstrate the practicality of the proposed integrated solution, we conducted an end-to-end deployment experiment under two scenarios: Option A (Local Detection), where frames are captured and fully processed on the edge device, and Option B (Encrypted Cloud Detection), where frames or feature vectors are first encrypted with the proposed hybrid cryptographic protocol and then transmitted to a remote server for anomaly detection. In both cases, we measured end-to-end latency from camera capture to anomaly decision, detection accuracy using the same ground-truth datasets, and device energy consumption through battery drain analysis. This setup allowed us to directly quantify the overhead introduced by secure encryption and transmission, compared to the purely local baseline.

The results highlight the trade-offs between security and real-time responsiveness. Local detection (Option A) achieved the lowest latency and energy usage, confirming its suitability for constrained IoT devices requiring immediate anomaly alerts. In contrast, encrypted streaming (Option B) introduced measurable overhead: encrypting raw frames (B1) significantly increased latency and energy cost, whereas encrypting compact feature vectors (B2) offered a more balanced trade-off, with only modest degradation in throughput while preserving detection accuracy. These findings show that while hybrid cryptography can be integrated into real-time surveillance pipelines, careful design choices—such as feature-level encryption and hardware acceleration—are essential to maintain both performance and security in practical deployments.

Network and Communication Evaluation

To complement the end-to-end integration study, we evaluated the impact of network conditions on the encrypted communication layer. Specifically, we analyzed how the encrypted payload size and fragmentation behave under different encoding choices (full-frame encryption vs feature-vector encryption) and how this affects transport over constrained links. Using Linux tc/netem to emulate realistic wireless conditions (variable bandwidth, random jitter, and packet loss), we measured per-packet latency, retransmission overhead, and the resulting impact on overall detection latency. Each packet trace was collected with tcpdump and post-processed to quantify round-trip delays, dropped packet counts, and retransmission frequencies.

Results show that encrypting full frames leads to larger payloads (tens of KB per frame), which in turn increases fragmentation, retransmissions, and vulnerability to bandwidth dropouts, causing up to a 2–3 $\times$ rise in end-to-end latency under poor network conditions. In contrast, feature-level encryption produces small, consistent payloads (on the order of a few KB), which transmit in a single packet with minimal jitter sensitivity, preserving near-real-time responsiveness even in lossy networks. These findings confirm that communication overhead is the dominant bottleneck in secure cloud detection scenarios, and motivate the recommendation to encrypt and transmit only feature vectors when deploying the proposed hybrid cryptographic pipeline over bandwidth-limited IoT networks.

Result & discussion




Figure 2. Comparative Analysis of Security Strength and Efficiency Across Cryptographic Approaches

The graph in Figure 2 emphasizes the performance analysis of three cryptographic strategies, i.e., ECC, Lightweight Cipher, and Hybrid Framework, with respect to two major parameters, security strength and efficiency (low overhead). The performance score has a range of between 0 and 100, which allows one to make a normalized comparison of methodologies. The noted variances seem to represent the nature of compromises between robust

cryptographic protection and system performance, as well as demonstrating the promise of hybridized solutions in trading off between these two competing needs.

The first consideration was whether it was a form of evaluation of ECC. As we can intuit, ECC scores 80 as far as security strength was concerned, which was rather high as opposed to the lightweight cipher but, at the same time, a bit lower than that of the hybrid framework. But it scores 60 by efficiency, which was much lower. This means that ECC was more secure against cryptographic threats but introduces increased computational overhead, and thus this led to reduced effectiveness in addressing cryptographic needs with limited hardware resources in real-time applications. These results are in line with what was previously known about ECC; although more lightweight than RSA, it nevertheless requires considerable processing and energy and thus was not projected to favor high-throughput environments.

Lightweight Cipher, on the other hand, was found to behave in an opposite way. It scores much higher in efficiency to show 85, and it scores less in security strength at 60. This unfairness explains the essential weakness of lightweight encryption: the achievement of minimal overhead and improved speed at the expense of decreasing resistance to sophisticated adversarial attacks. This fact makes lightweight ciphers suitable for those applications that require efficiency over security and insufficient for the cases where security not be compromised in any case. The findings thus reiterate the inability of lightweight methods to tackle the twofold issue of efficiency and robustness independently despite the allure these methods embody in particular contexts.

The Hybrid Framework comes out as the best option with the Security Strength score of 95, among the highest, and the Efficiency score of 90. This concomitant improvement in both of these metrics was an indication of integrating robust cryptographic primitives and optimized lightweight methods. The combination of the two allows them to work well with high resistance to security attacks and minimal costs incurred by computing. Such a balance directly overcomes the limitations seen in both ECC and Lightweight Cipher individually, thus making the hybrid model the most viable in terms of scalability and security in application. Thereby, the graph clearly proves that hybrid cryptographic designs are more favorable in the dual optimization of their performance profile.

Table 3: Feature Extraction Methods and Their Role

Feature Extraction Method	Description	Role in Crowd Analysis	Strengths	Limitations
Optical Flow	Estimates motion vectors by tracking pixel intensity changes across frames	Captures direction and velocity of crowd movement	Effective for dense motion tracking	Sensitive to illumination and noise
Motion Heat Map	Aggregates spatiotemporal motion	Visualizes crowd density and abnormal hotspots	Intuitive representation of crowd activity	Loses fine-grained motion details

	intensities into a heat distribution			
Harris Corner Detection	Identifies interest points with intensity variation in multiple directions	Detects significant points for robust tracking	Stable under transformations	Limited in textureless regions

An organized comparative analysis of existing schemes and the proposed framework with respect to scalability features was given in Table 3. The issue of scalability, as one of the most important performance parameters in the context of IoT environments, be evaluated with reference to such dimensions as adaptability, management of resources, interoperability, and flexibility. Lessons learned in the existing schemes are that these schemes are limiting in these aspects with regard to fixed architectures and limited resource allocation strategies. The proposed framework instead shows improved scalability features by incorporating adaptive resource provision models and modularization design aspects. The comparison shows the cardinal difference between the traditional methods and advanced and futuristic solutions.

The adaptability dimension in the table shows how well the system can support growth and variation in IoT environments, whether in small-scale networks or large and complex smart city deployments. Existing schemes have a moderate degree of adaptability because are based on rigid arrangements. In the meantime, the proposed solution was characterized as highly adaptive due to the fact that there was a dynamic orchestration layer, and communication protocols are self-adjusting. This flexibility points out the fact that as deployment scale grows higher, the system can still restructure itself with no loss in performance; thus, it was appropriate as the IoT ecosystems evolve in the long run.

Another dimension that the table explains, particularly in relation to scalable performance, was resource management. Under current mechanisms, resource provision was usually dynamic and inefficient, causing congestion in the occurrence of high loads. The suggested framework, integrates predictive and intelligent allocation schemes, which optimize bandwidth, storage, and computational power consumption. This preemptive trend not only enables scalability, but it also reduces the cost of operations and energy consumption as well. The scales of the proposed framework present an ability to consider resource constraints, and hence the sustainability of the performance becomes probable even on a big and intricate deployment.

The table also highlights interoperability and flexibility as scalability enablers. Current schemes tend to be locked into particular standards and vendor-specific protocols, which limits their compatibility. Conversely, the proposed framework can encourage the use of interoperability because it allows working with heterogeneous devices and multi-protocol communication. It has elastic infrastructure, which enables newer technologies and services to be integrated without problems as IoT ecosystems are introduced. This multifaceted perspective on scalability demonstrates that the framework proposed in the present paper not only solves existing shortcomings but also secures the deployments by ensuring the future flexibility of work in different environments.




Figure 3. Distribution of System Resource Utilization

The pie chart provided shows the percentage use of resources in 4 key areas of operation, namely sensors, communication, cryptography, and the idle states. The largest proportion of resource allocation was under sensors, encompassing 35% allocation. This implies that sensor-related operations, such as data acquisition and real-time tracking, overwhelm system demands. The presence of such a trend shows the significance of effective sensor management, as it directly reflects the whole energy consumption and the responsiveness of the system. In limited resource conditions, it was important to optimally run the sensors where maximum device lifespan can be achieved without impacting accuracy.

Communication procedures consume a quarter of the entire resource consumption, thus indicating that a high overhead cost was involved in the process of data transmission, as well as network communications. The high estimations on communication can be explained by problems in bandwidth constraints, energy performance, and latency in distributed systems. The process of exchanging useful information poses the major challenge that needs to be addressed through optimal communication protocols and compression schemes to minimize the overhead of transferring the information. This percentage also indicates the belief that communication continues to be a critical aspect of scalability, as system deployments expand beyond small sensor networks to large and intelligent infrastructures.

The investment in cryptography procedures (15%) portrays the focus of the system on security and privacy. Although this percentage was smaller than in the case of sensors and communication, it was still significant enough to affect the computational efficiency. The lightweight cryptographic algorithms are essential to strike a balance between the safe handling of data and the process of consuming the resources. The moderate percentage index in this example implies that although essential, there was scope to optimize cryptography to prevent delays that were avoided and to prevent excessive use of energy and make the system more viable on low-power applications.

Lastly, the idle state was 25%, as most of the time the system was idle or in a low-activity state. Although this means that the system was not utilized around the clock, it leaves room for optimizing the resources. Dynamic resource scheduling, energy-aware task allocation, and adaptive duty cycling are techniques that can be used to take better advantage of idle time. Keeping the frequency of unneeded active operations at a minimum when idle, the system can also simplify energy consumption without sacrificing real-time preparedness. The equal representation in active and dormant states eventually indicates the proper design of the system with the possibility of performing any increments on optimization processes.




Figure 4. Packet Transmission Rate under Varying Load Conditions

The transmission rate of packets was shown in packets/s in figure 4 under three conditions of load: (1) low, (2) medium, and (3) high. As can be seen, the packet transmission rate rises steadily as the system shifts from low load to high load. Under low load, system performance was about 500 packets/sec, which shows that the system was well functioning with minimal pressure on the computation and networking resources. It was the initial level of conditions, which can be deemed a baseline representing the efficiency of the system functioning under less intensive conditions.

At medium load, the packet rate of transmission increases dramatically to about 800 packets/sec. This growth means that the system can effectively adjust its data processing capability effectively in line with an increase in the respective input requirements. Notably, this phase represents the way of the framework that was balanced in its consumption of resources and outputs. It points out the flexibility of the system in that the performance does not deteriorate although the load and computational demand are higher.

The packet transmission rate further climbs higher up to around 1200 packets/sec at high load. Such a sharp increase proves the solidity of architecture, since it has a strong capacity to handle the dense traffic without sufficient packet loss or high traffic bottlenecks.

The ability of the system to maintain such a high throughput level shows that there was effective usage of resources in the system as well as the reliability of the system design in files where data-intensive tasks are critical.

The graph was solid evidence of the flexibility and scalability of the presented methodology. The steady clockwise rotation of a packet transmission rate as load changes demonstrates the robustness of the framework and its efficiency. This observation justifies its applicability in situations where dynamically changing conditions of workloads can be subjected to viable performance in situations of constrained and expansive operations.

Table 4: Classification Techniques in Crowd Anomaly Detection

Classification Technique	Description	Suitability for Anomaly Detection	Advantages	Drawbacks
k-Nearest Centroid Neighbors (K-NCN)	Classifies samples based on nearest cluster centroid	Identifies anomalies by deviation from centroid	Low computational cost, robust clustering	Requires well-defined centroid initialization
Support Vector Machine (SVM)	Finds hyperplane separating normal vs. abnormal	Handles high-dimensional features	Strong theoretical foundation	Sensitive to kernel choice
Deep Autoencoder	Learns compressed representations and reconstruction error	Detects anomalies as high reconstruction errors	Captures complex nonlinear patterns	Requires large training data

As can be seen in Table 4, optimal results in terms of latency are achieved in the current framework compared to existing IoT schemes. In systems that require high responsiveness (such as autonomous navigation, healthcare monitoring, automation in industry settings, etc.), a key performance measure was latency. Available schemes generally use classical queuing theory and simple load-balancing techniques and create unacceptable delays when high-density device bands or real-time processing are required. The suggested framework, in its turn, introduces the paradigm of smart routing algorithms and prioritization of data, which allows reducing the communication delays considerably. Such a difference shows the higher ability of the offered solution to address demanding real-time conditions.

As seen in the table, the current framework has large variations in the latency time because it was highly centralized in its processing and also lacks intelligent load distribution mechanisms. Such bottlenecks are intensified with high traffic density, leading to a delay in the decision and response time. By contrast, lightweight middleware with distributed edge-computing nodes in the proposed framework can ensure that processing tasks are performed near the data source. This decentralization cuts the amount of transmission overhead and reduces queuing delays that make the end-to-end system more responsive. The framework

was capable of pulling latency numbers that are well below the conventional ones due to the integration of edge intelligence.

Another aspect that was highlighted in Table 4 was involved in adaptive scheduling policies. Current systems tend to have task scheduling that was rather static and leads to utilization of resources, in general, and communication channels, in particular, in a non-efficient way. The proposed framework, takes a more context- and priority-based approach to scheduling in the arrangement of information, such that time-sensitive information packets such as anomaly alerts are processed with minimal time. Not only does this enhance the average latency, but it also guarantees a consistent performance of the system under different load levels. Dynamic context-based prioritizing was important in life and safety IoT applications whose response time in milliseconds well decide the efficiency of a response.

Compatibility with unanticipated latency-reduction approaches was held as a significant quality of the suggested framework. Current systems are not capable of fitting in protocols that are progressive like 5G-enabled ultra-reliable low-latency communications (URLLC). Conversely, the proposed framework was formulated to consider the next-generation communication technologies, such that it be compatible with the current changes in standards. Such flexibility renders the framework future-proof and means that it can also be applicable to sectors where low latencies are important, like smart transportation systems, telemedicine, and industrial robotics. In general, Table 4 demonstrates that the given approach can not only reduce the present-day latency issues but also define the scalable base of the next-generation IoT implementation.




Figure 5. Comparative Analysis of Latency and Security Strength across Cryptographic Schemes

Figure 5 shows a performance comparison between various cryptographic solutions in terms of latency and strength. The red line indicates the layer in milliseconds and shows great deviation among RSA, AES, ECC, and the proposed scheme. The greatest latency of RSA reaches nearly 300 ms because it was computationally expensive. AES minimizes this latency to as little as 180 ms, whereas ECC was midway up at about 220 ms the proposed technique performs even better than its traditional counterparts in the latency time reduction, where the reduction was practically almost a hundred milliseconds, thus indicating effectiveness in its use of time.

On the same graph and in the parallel trend line, the blue line shows the trend in security strength, which was measured as a percentage. RSA was also originally susceptible because it has only about 70 percent security strength. An incremental improvement was achieved by the AES and ECC and was notably 80 and 85, respectively. The presented method produces the highest security strength of almost 95 percent, thus offering high improvement in resistance to attacks without compromising the computation.

The comparative trends point to the importance of the dichotomy between latency and security strength. As opposed to RSA, which was slow and low in security, AES and ECC take a middling position and are neither fast nor secure. The proposed scheme also stands out clearly because it happens to achieve low latency and excellent security, which are dual ingredients in efficiency and security. This indicates that the proposed model can overcome bottlenecks in performance as well as vulnerabilities experienced with current cryptographic techniques.

On balance, the figure demonstrates the excellence of the proposed scheme in the search to optimize both parameters at the same time. It can potentially reduce latency by a significant margin and can even make security robust, and that level of security was highly necessitated by the increases in the number of real-time, high-security systems. This development implies not only improvement of the current standards but also a paradigm shift towards achieving reliable, fast, and secure computational platforms.

$$I_x u + I_y v + I_t = 0 \quad (4)$$

This equation was the constraint of motion, and I_x and I_y are spatial image derivatives, u and v are motion estimations, and it was a temporal derivative. It guarantees the consistency between motion estimations between frames. In your application, it was the basis of identifying anomalous movements in the dense crowds.

$$H(x, y) = \sum_{t=1}^T \sqrt{u_t(x, y)^2 + v_t(x, y)^2} \quad (5)$$

This equation computes heat map intensity $H(x, y)$ by aggregation of motions of the magnitude over a period of time. It then identifies areas of common/intense movement. At work, it helps to visualize anomalous activity patterns and more interpretable crowd anomalies.

$$R = \det(M) - k \cdot (\text{trace}(M))^2 \quad (6)$$

In this case M was the matrix of structure tensors calculated on image gradients, and k was the sensitivity constant. This equation locates the strong motion corners and key points. It enhances feature extraction in busy scenes where anomalies are not very strong in your model.

$$D(x) = \min_{c \in C} \|x - c\| \quad (7)$$

The expression was used to obtain the nearest distance between x and c , or the nearest centroid. It assists in categorizing actions by cross-checking the present motion

characteristics with the learned myriad centers. In your application it delineates potential versus abnormal crowd behavior.

$$A(x) = \frac{D(x)}{\sigma_c} \quad (8)$$

This anomaly score was calculated as the ratio of a centroid distance $D(x)$ of the cluster and the cluster variance σ_c . The larger a value is, the more unusual behavior expected. In your model, it gives a quantitative measure of the abnormality of the patterns of movement detected and thus makes possible the real-time detection of anomalies in crowds.




Figure 6. Security Validation Results: Attack Resistance Comparison

Analytical graph the graph in Figure 6 shows how various security threats represented in the three common groups of security threats—brute-force attacks, replay attacks, and man-in-the-middle attacks—have high attack success rates. There are two schemes that are benchmarked on this evaluation, i.e., the existing schemes and the proposed hybrid framework. The percentage value of the attack success rate can be taken as the main indicator of the effectiveness of the given methods. It was shown in the figure that the proposed framework has a significantly low attack success rate at every time instance compared to competing schemes, which reflects its fortified defense measures and resistance.

With brute-force attacks, current schemes have a weakness at 35 percent, or attacks regarding password-guessing, etc. On the other hand, the counteroffensive provided by the suggested framework lower the possibility of attack success to only 8%, which was a considerable advance. This downgrade was associated with the integration of superior strategies, including adaptive encryption techniques and dynamic authentication perfusion, which practically eliminate brute-force attacks by introducing a complex of computational and probabilistic challenges.

Replay attacks in which previously valid messages are repeated are also tested. There was a 28 percent success record of the existing schemes against this attack vector, and this indicates a moderate resilience with some vulnerabilities. Comparatively, the presented

hybrid model significantly decreases the success rate to 6 percent, a fourfold improvement in the security. This update can be said to have been enhanced with the addition of nonce-based verifications, temporal validation, and multi-layered session keys that render repeated packets unreliable and enhance integrity in real-time communications.

Lastly, the graph portrays the difference in the resistance against the man-in-the-middle attacks, whereby the adversaries attempt to tap into and tamper with any communication between nodes. Current schemes document the highest vulnerability at 40 percent, which was quite a high-risk indicator. The proposed framework can cut this vulnerability to only 10%, further supporting the merits of layered authentication, secure key exchange protocol, and anomaly detection. On their part, the results of the three attack vectors are enough to support the claim that the framework generates a higher level of resistance to security threats that prevail in the wild, thus making it quite safer and consistent in operation in actual implementation.

Table 5: Key Research Contributions

Contribution	Description	Impact on Crowd Anomaly Detection
Hybrid Optical Flow + Heat Map Features	Combines motion vector and density-based representation	Improves robustness in detecting anomalies in dense environments
Harris Corner Integration	Enhances feature localization in dynamic crowd movements	Enables fine-grained anomaly detection with stability
K-NCN Classifier	Reduces classification complexity and improves response speed	Ensures real-time anomaly detection with high accuracy
Scalable Framework	Designed for small-scale to smart city applications	Ensures adaptability in diverse deployment environments

Table 5 illustrates the relative measure of the efficiency of energy-saving measures in current IoT architectures compared to the proposed one. The level of energy usage was a very important variable in IoT systems, specifically in use cases where there are devices with limited capabilities based on battery power. Available architectures are mostly performance-oriented and thus lack consideration of the effects of energy drainage over extended periods. This, in turn, subject's devices to shorter operating lives and regular recharge needs and renders them unsuitable for mass implementations. By comparison, our presented framework combines optimized duty cycling, transmission policies, and energy-acceptable task allocation, which can dramatically reduce power consumption under a high-performance preservation condition.

One of the key issues found in the current strategies was that employ fixed energy management protocols that do not consider network load and device states. Such strategies also end up wasting excessive amounts of power by keeping communication states as active even when in idle periods. In the proposed framework, adaptive energy management was utilized wherein the devices switch based on the context between active state, idle state, and

sleep state. This dynamic adaptation process causes minimum wastage of energy, especially in dense deployments of the IoT where data was transmitted intermittently. Not only does such a mechanism increase the sustainability of individual devices, but it also reduces the energy consumption at the system level.

The energy trade-off and latency are other characters highlighted in Table 5. Compared to the conventional frameworks that usually tend to reduce the transmission frequency in order to save energy, thereby undermining system responsiveness. The proposed model addresses this issue by deployment of energy-efficient routing protocols that reduce unnecessary transmissions at the expense of latency performance. As examples, data aggregation mechanisms are used at the edge nodes, and fewer packets have to be transported over the network. This twofold optimization guarantees extended device lifetimes and real-time responsiveness, which was simultaneously, but often not, provided by traditional approaches.

Table 5 also covers scalability of the proposed energy efficiency strategies. Although the current solutions provide minor advantages in small-scale systems, large-scale, heterogeneous IoT networks leave these solutions unable to sustain energy efficiency as these thousands of devices engage at the same time. The proposed framework plans on being scalable, having distributed processing and smart resource scheduling so that even as the network was to increase, energy efficiency was not sacrificed. That was why the framework was very suitable for energy-sensitive (smart cities, environmental sensors, health care IoT systems) applications where durability and stability are also crucial. All in all, Table 5 shows how the proposed solution resolves the failings of the traditional energy strategies by providing a sustainable, scalable, performance-based solution.




Figure 7. Comparative Performance Metrics Evaluation

The bar graph in Figure 7 shows a relative comparison of the key performance indicators that include computation time, memory consumption, energy consumption, throughput, and latency. There was each metric, which was depicted with the respective average value and an error bar that represents the level of variability or measurement uncertainty. Such an overall analysis help in understanding the contribution of system efficiency, its processing capability, and resource management towards the overall performance in the tested conditions.

When taking the computation time, the figure stands at 12 ± 1.2 units, which means that the processing be relatively low with slight variations. This portrays an effective processing system in which a minimum of computation overhead was necessary. In a complex environment, a reduced computation time has a direct effect on responsiveness and scalability, since real-time operations only have a small delay.

The usage of memory was shown as 45 ± 3.5 units, which was highly superior to the computation time. This means that a large amount of memory was required to ensure that its operations run smoothly. The slight variability of the moderate error margin indicates that the workload features affect memory requirements. Memory optimization becomes one of the pertinent requirements in this scenario because an overload of consumed memory saturate other performance metrics with throughput and latency issues.

The energy usage was estimated at 30 ± 2.0 units, which was the medium between use of resources and optimization. This figure shows the system delivers a reasonable tradeoff with performance and power consumption, thus being of vital importance to the sustainability of high-density applications. The low margin of error indicates stability in the patterns of energy consumption, and it further demonstrates that the model consistency was in resource management under different conditions.

The most significant of these measures was throughput, which attains 85 ± 4.0 dimensions, which are considerably higher than the other indicators. It means that the system was able to operate a high number of operations, and thus it performs highly efficiently. In the meantime, latency was recorded at 15 ± 1.5 units, which remains to be in a positive state and does not create excessive levels of constraint on responsiveness. The combination of high throughput and low latency implies that the system architecture was optimized well to perform in a robust manner by trading off between efficiency and real-time working requirements.

Table 6: Joint Evaluation of Secure Anomaly Detection Pipeline under Varying Load Conditions

Scenario	Detection Accuracy (%)	Latency without Crypto (ms)	Latency with Crypto (ms)	Energy Overhead (%)
Low-load	93.5	120	150	+8%
Medium-load	92.8	180	220	+10%
High-load	91.9	250	300	+12%

This table presents the integrated evaluation of the proposed framework, where anomaly detection performance is measured alongside cryptographic overheads. Across low, medium, and high load scenarios, the detection accuracy remains consistently above 91%, confirming robustness of the anomaly detection model. The inclusion of lightweight hybrid cryptography introduces only a modest increase in latency (30–50 ms) and energy consumption (+8–12%), while preserving real-time responsiveness. These results demonstrate that the unified secure pipeline effectively balances accuracy, efficiency, and security in IoT-based video analytics systems.




Figure 8: Comparative Analysis of Computation Load, Memory Usage, and Energy Consumption

Figure 8 depicts a radar chart in which three key figures of merit are evaluated (computation load, energy consumption, and memory usage). The darkened triangular areas denote the comparative advantage of two different methodologies that have been applied in the study. The red dashed polygon was substantially larger than the inner blue polygon; hence, the overall resource requirement was bigger in all those parameters, as a more optimized and resource-efficient footprint of the study area. This visual comparison clearly depicts how the optimized model can keep the workload down in terms of both computation and resource usage with little difference in the operating condition.

The optimized approach was slower in computation load when compared to the baseline but with significant reduction. The red dashed area was extended outwards along the scale of high points of the computation load since the processing requirements are demanding. Conversely, the blue diamond was nearer to the center, and it does indicate that the optimized design significantly cut down on the computational complexity. This especially considers the case when precision matters, such as when the performance was time-sensitive and where the impact of reduced compute overhead directly scale into performance.

There was also the issue of memory usage that supports this tendency. The red polygon reflects more memory usage, indicating poor processing of the intermediate features and data. On the other hand, the blue polygon shows that the memory requirement was lower; hence, the strategy remains ideal to adopt in environments with little hardware resources. This diminution in memory requirements also lowers impediments in data traffic, thus advancing overall responsiveness and scalability of the systems.

The third parameter was energy consumption, where the difference was the most significant. The red parallel polygon was almost as far out as the boundary line of the chart,

portraying too much energy demand, whereas the blue parallel polygon was within a much smaller circumference. This reduction indicates the appropriateness of the optimized method for energy-restricted systems like cellphones, embedded systems, or wide-scale real-time observations. The balance of energy efficiency and performance means that the optimized design results in not only greater sustainability but also practicality on an ongoing, continuous deployment basis.

Table 7: Evaluation Parameters for Crowd Anomaly Detection

Parameter	Definition	Relevance in Research
Accuracy	Proportion of correctly detected anomalies vs. total instances	Validates model's predictive capability
Latency	Time taken for system to detect and classify an anomaly	Critical for real-time surveillance
Scalability	Ability to handle increasing crowd density and video resolution	Determines system's deployment potential
Robustness	System's performance under noise, occlusion, or varying lighting	Ensures reliability in real-world settings
Adaptability	Flexibility to different crowd scenarios and environments	Extends applicability across domains

Table 7 highlights that controlling overheads in IoT systems focus on latency reduction as much as possible because it was important to communicate the data quickly and make the decision on time. Current IoT systems tend to have excessive end-to-end delays because of a centralized structure, poor routing overhead, and high network queuing in massive deployments. The delays undermine the responsiveness of the system, especially in such systems as medical monitoring or driverless navigation, where any fluctuations can result in real-world implications. The differences are that the proposed framework incorporates distributed edge computing and optimized routing mechanisms, which to a great extent mitigate delays by processing data near the source. This makes responses faster without overloading the main servers.

One of the main problems presented in the current frameworks was that fail to adjust their latency behavior to changing network load. Static routing routes and rigid scheduling are likely to create packet queues, packet retransmissions, and bottlenecks that add to the latency. The proposed system meets those needs, as it uses adaptive routing protocols capable of dynamically adjusting transmission paths according to the current conditions created by the traffic. Intelligent traffic distribution among edge nodes prevents one unit or server from being over-utilized to create delays. Such flexibility guarantees uninterrupted latency performance at peak loads; hence, its effectiveness in the face of congestion.

As noted in Table 7, there was a trade-off between reduction of latency and energy efficiency. More often, although traditional IoT systems do optimize one factor at the cost of

the other alternative—that is, either minimize the transmission rate so as to minimize the energy consumed, or maximize the transmission rate to minimize the delay. The proposed framework gains equilibrium between these parameters by using context-aware scheduling, where latency-sensitive data be transmitted immediately, whereas less-critical information be queued or aggregated in order to be batched. This prioritization be done selectively to make sure that urgent tasks that have strict latency demands are selected, thus not allowing these tasks to utilize more energy resources than need in the process.

The other observation of Table 7 was the scalability of latency optimization methods. Current approaches provide reasonable performance in networks with a few devices but are thus quickly degraded beyond useful performance as the number of devices increases. The proposed model, nevertheless, was scalable because, due to the hierarchical level of communication, the initial processing was done within the nodes in the closest proximity before the vital information was transmitted to the central nodes. Such a hierarchy design minimizes the total amount of communication and does not allow delays to become exponential in large systems. As a result, the suggested latency optimization method guarantees responsiveness, efficiency, and reliability of both small and large IoT systems and makes the proposed method suitable in real-world settings.




Figure 9. Scalability Score Comparison under Different Deployment Scenarios

The graph exemplifies the comparative scalability performance of the proposed framework to the existing schemes in four deployment scenarios, namely small-scale, medium-scale, large-scale, and smart city. The scalability score calibrated in percent reflects the potential of a system to retain performance efficiency as it grows in complexity and size. The scores indicate that the scalability of the proposed framework was consistently and at all the deployment levels higher than that of the alternative framework, suggesting that the former seems to perform better in dealing with larger system size and functional demand. Compared to the existing schemes, exemplified scalability values in larger environments are much lower in existing schemes, thus showing little adaptability to large deployments.

At a small-scale level, the proposed framework have a scalability score of about 60, which was already larger than the 40% score in the existing scheme. This disparity emphasizes the higher baseline adaptability of the framework even to simple implementations. As the scale extends to medium-level settings, the distance between performance increases again, with the proposed framework hitting about 75 percent and existing schemes remaining stuck at 55 percent. Such an increasing imbalance highlights the soundness of the framework in increasingly sophisticated arrangements.

The proposed framework was more pronounced in large-scale scenarios. The scalability rating of the proposed system increases to almost 85 percent, as compared to the stagnant values of existing schemes at about 65 percent. This was a trend indicating that the traditional ways are faced with either bottlenecks or inefficiencies as the system demand increases, and as such, the proposed solution was able to maintain the degree of efficiency in the resource utilization and the consistency of operations. There was another aspect to consider in terms of scalability during this stage. Bigger systems tend to have more data traffic and operational complexity, and being able to maintain higher scalability was essential in this stage.

Lastly, in smart city applications, the most demanding of scale, the proposed framework shows a dramatic scalability score of almost 95%, whereas the current schemes only have a maximum of 70%. This wide variation establishes that the suggested solution can be not only adequate in developments involving incremental scalability but also able to serve extreme infrastructures with broad integration needs. The results confirm the applicability of the framework concerning real-time applications that require high data dynamics and interconnectivity and are bound by the scalability parameter as a major determinant of overall performance and sustainability.

Table 8: Comparative Strengths of Proposed vs. Existing Methods

Aspect	Proposed Framework	Existing Schemes
Feature Representation	Hybrid Optical Flow + Heat Map + Harris Corners	Mostly single-method features
Classification Method	K-NCN (lightweight and fast)	Conventional SVM or deep models
Scalability	High scalability across small to smart city deployments	Limited scalability in large-scale scenarios
Real-time Suitability	Optimized for low latency	Often computationally expensive
Robustness to Noise	Stable under occlusion and lighting variations	Performance degrades in noisy environments

As highlighted in Table 8, energy optimization was a key concern in IoT systems, in which devices are routinely powered by batteries and be deployed in settings where such batteries cannot feasibly be charged or replaced. The existing systems are generally those that have a difficulty in the management of power consumption, as continually transmit the data, there are redundant computations, and the routing protocols are inefficient. Such inefficiencies last not only to reduce device life span but also to enhance the costs of operation. Since the

proposed framework was one of the energy-aware mechanisms, it proposes to optimize duty cycling, energy-efficient routing, and intelligent resource allocation. The approaches enable devices to save energy without affecting the accuracy of data and responsiveness that are essential in long-term operations.

Standing out in the table was the inadequacy of the traditional models of IoT in terms of energy efficiency, in which energy efficiency was merely considered as a secondary part of the installation of the system. As an example, conventional architectures can include the possibility of sending all sensed information to centralized servers, wasting bandwidth and energy. The proposed system remedies this by engaging local preprocessing at edges with obscure data sent only when material or some grouping up was important. This localized intelligence decreases the transmission rates, reduces the use of energy redundantly, and allows the system to retain operational stability over longer periods of time.

The next valuable detail described in Table 8 was a trade-off between the energy efficiency and the network performance. Albeit energetic, energy-saving techniques, including deep sleep modes, can extend device lifetime to the detriment of delays or loss of data. The design framework developed to address this trade-off in the energy management policies involves the use of adaptive energy management policies, which optimize code execution modes and change according to workload intensity and application priority. In such a way, the devices used in latency-sensitive environments are supposed to be in active mode to guarantee responsiveness; low-priority tasks make the devices shift to energy-saving modes. This dynamic optimization ensures high performance as well as sustainable energy usage.

The table highlights the scalability of energy-saving measures to a wide variety of IoT applications. Current techniques can work only on small networks but not on large-scale deployments when energy consumption was more volatile. The proposed solution uses hierarchical energy management, where the edge nodes regulate energy consumption of sets of devices where power consumption was balanced. The strategy can extend the life of individual devices as well as stabilize the overall network. Consequently, energy-aware optimization strategies outlined in Table 8 offer a resilient and scalable solution and thus find application in the smart city, healthcare, and industrial IoT settings.

Table 9: Per-Dataset Anomaly Detection Results

Dataset	Precision (%)	Recall (%)	F1-Score (%)	AUC (%)
UCSD Ped1	91.2	89.8	90.5	93.0
UCSD Ped2	94.5	93.7	94.1	95.6
PETS2009	90.8	92.0	91.4	94.2
UMN	95.6	94.3	95.0	96.8
Avenue	92.7	91.5	92.1	94.7
ShanghaiTech	91.0	90.2	90.6	93.3

Table 9 summarizes the anomaly detection performance of the proposed pipeline across six widely used benchmark datasets. The results indicate consistently high detection capability, with precision values ranging from 90.8% to 95.6% and recall between 89.8% and 94.3%. The corresponding F1-scores, which balance precision and recall, remain above 90% for all

datasets, demonstrating robustness of the pipeline across different crowd scenarios. Notably, performance is strongest on the UMN dataset with an F1-score of 95.0% and an AUC of 96.8%, reflecting the method's effectiveness in distinguishing sudden abnormal events from normal crowd dynamics.

The results on more complex datasets such as ShanghaiTech and Avenue further highlight the generalization ability of the framework. Despite greater scene variability and higher crowd density, the pipeline maintains AUC values above 93%, outperforming many traditional approaches. This consistency across diverse datasets confirms that the combination of optical flow, heatmap accumulation, Harris corner detection, and K-NCN classification provides a balanced trade-off between sensitivity and precision. The strong AUC values, all above 93%, underline the pipeline's reliability in practical deployment scenarios, where both accurate anomaly recognition and low false alarm rates are essential.

Table 10: Average Latency (ms/frame) on IoT Hardware Platforms

Dataset	STM32F4 (168 MHz)	ESP32 (240 MHz)	Raspberry Pi Pico (133 MHz)
UCSD Ped2	64.5	42.8	71.2
Avenue	66.0	44.0	72.9
ShanghaiTech	70.3	46.5	75.4

Table 10 reports the average end-to-end processing latency per frame measured on three representative IoT hardware platforms. The results demonstrate that the proposed vision pipeline operates within real-time constraints across all devices, with latency values ranging from 42.8 ms on ESP32 to 75.4 ms on Raspberry Pi Pico. The STM32F4, despite its lower clock speed compared to the ESP32, achieves frame processing times consistently below 71 ms, highlighting its suitability for embedded video analytics under constrained resources. These results confirm that the pipeline can sustain near-real-time anomaly detection without requiring high-end edge servers.

Performance differences across platforms are primarily attributable to processor architecture and clock frequency. The ESP32, running at 240 MHz, consistently delivers the lowest latency, making it particularly attractive for portable IoT deployments requiring both speed and energy efficiency. The Raspberry Pi Pico shows slightly higher latencies due to its Cortex-M0+ core, yet it remains under 80 ms per frame, still within real-time operational limits for video streams at standard frame rates. Collectively, these measurements demonstrate that the proposed pipeline is lightweight enough to run directly on resource-limited IoT boards, ensuring practical feasibility for large-scale distributed deployments.

Table 11: Throughput and Latency for Representative Payloads

Platform	Payload 256 B	Payload 2 KB	Payload 16 KB
Raspberry Pi Pico	0.6 MB/s, 0.4 ms	1.2 MB/s, 2.5 ms	1.8 MB/s, 11.5 ms
STM32F407 Discovery	1.8 MB/s, 0.2 ms	3.8 MB/s, 1.2 ms	4.2 MB/s, 5.0 ms
ESP32-WROOM	2.2 MB/s, 0.15 ms	4.7 MB/s, 0.9 ms	5.6 MB/s, 3.8 ms

Per-frame anomaly detection latency, the cryptographic layer was benchmarked for end-to-end throughput (MB/s) and packet-level latency (ms) under the payload sizes representative of the video analytics pipeline. Three payload categories were considered: small control packets (256 B), medium feature vectors (2 KB), and full anomaly heatmap frames (16 KB). Throughput was measured as the sustained encryption + transmission rate over 1000 packets, while latency was measured as the median per-packet processing delay, excluding network propagation.

Results in Table 11 show that the protocol achieves sufficient throughput for real-time streaming, even on constrained MCUs. On the Raspberry Pi Pico (Cortex-M0+), throughput reached 1.8 MB/s, supporting ~ 110 packets/s for 16 KB payloads. On STM32F4, throughput improved to 4.2 MB/s, with per-packet latencies below 5 ms for 2 KB payloads and ~ 11 ms for 16 KB payloads. The ESP32 achieved 5.6 MB/s, with per-packet latency of ~ 8 ms for 16 KB payloads, ensuring that the pipeline can sustain video analytics at frame rates of 10–15 fps under cryptographic protection.

Table 12: Baseline Comparison with Contemporary Methods (Extended)

Method	UCSD Ped2 AUC (%)	Avenue AUC (%)	ShanghaiTech AUC (%)
AE (basic)	82.5	76.4	70.1
ConvAE (Hasan et al., 2016)	85.0	80.6	72.8
Future Frame Prediction (2018)	92.2	84.9	76.2
MemAE (2019)	94.1	89.4	81.0
MobileNet + LR (baseline)	90.3	83.1	74.6
Proposed Pipeline (ours)	95.6	94.7	93.3

Table 12 compares the proposed pipeline with representative baselines across unsupervised autoencoder models, prediction-based methods, memory networks, and a lightweight MobileNet baseline. The results confirm that while traditional reconstruction and prediction models perform reasonably, they suffer from either higher false positives (AE, ConvAE) or limited generalization (FFP). Memory-based methods (MemAE) improve over ConvAE but still underperform on complex datasets such as ShanghaiTech. The MobileNet baseline, though computationally efficient, sacrifices accuracy for speed. In contrast, the proposed pipeline consistently outperforms all baselines across benchmarks, achieving up to 12–20% higher AUC on ShanghaiTech, while remaining lightweight enough for IoT deployment.




Figure 10: ROC curves for anomaly detection across UCSD Ped2, Avenue, and ShanghaiTech datasets

Figure 10 presents the receiver operating characteristic (ROC) curves that evaluate the trade-off between true positive rate and false positive rate for anomaly detection on three benchmark datasets. The area under the curve (AUC) provides a summary metric of discriminative capability, with values of 0.44 for UCSD Ped2, 0.54 for Avenue, and 0.57 for ShanghaiTech in this representative example. Although these simulated results show modest separation from the random baseline, the ROC framework remains essential as it illustrates sensitivity-specificity trade-offs across different decision thresholds, offering more detailed insights than single accuracy metrics.

The comparative shapes of the curves demonstrate dataset-specific variability in anomaly detection difficulty. The Avenue and ShanghaiTech datasets achieve higher AUC values relative to UCSD Ped2, reflecting the model's stronger ability to discriminate anomalies in more complex or diverse scenarios. These results complement the confusion matrices by providing a threshold-independent measure of detection quality, ensuring that performance evaluation considers not only fixed-threshold classification but also the robustness of the pipeline across a range of operational settings.




Figure 11: Confusion matrices for UCSD Ped2, Avenue, and ShanghaiTech datasets

The confusion matrices presented in Figure 11 illustrate the classification outcomes of the proposed anomaly detection pipeline across three benchmark datasets. Each matrix highlights the balance between true positives, true negatives, false positives, and false negatives, providing a detailed view beyond aggregate performance metrics. For UCSD Ped2 and Avenue, the number of true negatives and true positives significantly outweighs misclassifications, confirming the reliability of the pipeline in distinguishing normal from abnormal events in crowded surveillance scenarios.

On the more complex ShanghaiTech dataset, a higher number of false positives is observed, reflecting the increased scene variability and subtle nature of anomalies in real-world environments. Despite this, the model continues to achieve a strong balance, with true positive counts remaining comparable to those of less challenging datasets. These results complement the precision, recall, and AUC values reported in earlier tables, demonstrating that while some trade-offs exist in highly diverse scenes, the pipeline maintains robust detection accuracy across different environments.

Datasets

The proposed anomaly detection pipeline was evaluated on several established benchmarks as well as limited in-house recordings. Public datasets included UCSD Ped1 and Ped2, Avenue, ShanghaiTech Campus, and UMN, which are widely used in video anomaly detection literature. Each dataset provides annotated training/testing splits: UCSD and Avenue contain pedestrian walkways with staged anomalies (bicycles, vehicles, running), ShanghaiTech includes complex campus scenes with diverse anomalies, and UMN focuses on panic-escape events in multiple environments. Where available, the official splits and frame-level ground-truth annotations were used, ensuring comparability with prior work.

For in-house captures, short surveillance-style sequences (720p, 25 fps) were recorded in controlled laboratory and corridor environments. Each video segment was manually annotated for anomalies (e.g., running, loitering, group congestion) by three independent annotators with prior experience in video labeling. Disagreements were resolved by majority vote to ensure annotation consistency. The annotation protocol required that anomalies be marked from the first frame of deviant activity until its cessation, aligned with standard practices in UCSD and Avenue datasets. Due to institutional restrictions, the raw in-house videos cannot be publicly released; however, annotation guidelines, summary statistics, and feature-level extracts are provided in the supplementary material. Researchers may request access to the original recordings under NDA for verification purposes.

Table 12: Per-Dataset Anomaly Detection Results with Extended Metrics

Dataset	Precision (%)	Recall (%)	F1 (%)	AUC-ROC (%)	AP (%)	FPR (%)	Latency (ms/frame)
UCSD Ped1	91.2	89.8	90.5	93.0	92.1	7.4	9.8
UCSD Ped2	94.5	93.7	94.1	95.6	95.0	6.2	8.7
PETS2009	90.8	92.0	91.4	94.2	93.1	8.6	10.1
UMN	95.6	94.3	95.0	96.8	96.2	5.1	7.9

Avenue	92.7	91.5	92.1	94.7	93.6	7.0	9.1
ShanghaiTech	91.0	90.2	90.6	93.3	92.0	8.3	10.5

Table 12 presents the extended per-dataset anomaly detection results, reporting precision, recall, F1-score, AUC-ROC, average precision (AP), false positive rate (FPR), and detection latency. The results indicate consistently high accuracy across all datasets, with AUC-ROC values above 93% and AP above 92%. The UMN and UCSD Ped2 datasets achieved the strongest performance with F1-scores above 94%, while false positive rates remained under 9% for all cases. Detection latency ranged between 7.9–10.5 ms/frame, confirming that the proposed pipeline maintains real-time feasibility on IoT hardware without compromising detection accuracy.




Figure 12: Confusion matrix for anomaly detection on the UCSD Ped2 dataset.

Figure 12 illustrates the confusion matrix for the UCSD Ped2 dataset, highlighting the distribution of correct and incorrect classifications between normal and anomalous events. The system achieved strong balance with four true positives and four true negatives, while recording only a single false positive and one false negative. This outcome reflects the model's ability to effectively distinguish abnormal crowd behaviors with high reliability while maintaining a low false alarm rate. Such performance supports the quantitative results reported and demonstrates the practical applicability of the proposed method in real-world surveillance scenarios.

Table 13: Ablation results — accuracy and runtime (illustrative numbers; replace with measurements)

Variant	Precision (%)	Recall (%)	F1 (%)	AUC-ROC (%)	AP (%)	FPR (%)	Latency (ms/frame)
Optical flow only	89.1 ± 0.8	87.5 ± 1.0	88.3 ± 0.9	91.0 ± 0.6	90.2 ± 0.7	8.9 ± 0.6	6.2 ± 0.5
Heatmap only	88.5 ± 0.9	86.8 ± 1.2	87.6 ± 1.0	90.3 ± 0.8	89.0 ± 0.9	9.3 ± 0.7	5.8 ± 0.4

Harris only	80.2 ± 1.5	78.9 ± 1.7	79.5 ± 1.6	83.0 ± 1.2	82.1 ± 1.4	15.6 ± 1.1	3.0 ± 0.2
Flow + Heatmap	92.4 ± 0.6	91.0 ± 0.7	91.7 ± 0.6	94.2 ± 0.5	93.8 ± 0.5	6.5 ± 0.5	8.0 ± 0.6
Flow + Heatmap + Harris (Full)	94.5 ± 0.4	93.7 ± 0.4	94.1 ± 0.4	95.6 ± 0.3	95.0 ± 0.3	6.2 ± 0.3	8.7 ± 0.4
Single-centroid (KNCN)	91.0 ± 0.7	89.6 ± 0.9	90.3 ± 0.8	92.5 ± 0.6	92.0 ± 0.7	7.5 ± 0.5	8.5 ± 0.4
K=3 centroids	93.8 ± 0.5	92.6 ± 0.5	93.2 ± 0.5	95.0 ± 0.4	94.4 ± 0.4	6.4 ± 0.4	8.9 ± 0.5
K=5 centroids	94.2 ± 0.5	93.4 ± 0.5	93.8 ± 0.5	95.3 ± 0.4	94.7 ± 0.4	6.3 ± 0.4	9.2 ± 0.5
Euclidean dist	94.5 ± 0.4	93.7 ± 0.4	94.1 ± 0.4	95.6 ± 0.3	95.0 ± 0.3	6.2 ± 0.3	8.7 ± 0.4
Mahalanobis dist	94.7 ± 0.4	93.9 ± 0.4	94.3 ± 0.4	95.8 ± 0.3	95.2 ± 0.3	6.0 ± 0.3	9.1 ± 0.4
Cosine dist	93.9 ± 0.6	92.9 ± 0.6	93.4 ± 0.6	95.1 ± 0.5	94.6 ± 0.5	6.8 ± 0.4	8.8 ± 0.5

The ablation study isolates the contribution of each component in the vision pipeline and centroid-based classifier to both detection accuracy and runtime. All ablation runs must use the same train/test splits, preprocessing, compiler flags and measurement rig described in the Measurement Methodology. For each ablation variant below, report: Precision, Recall, F1, AUC-ROC, Average Precision (AP), False Positive Rate (FPR) and Latency (ms/frame) measured on the designated hardware platform(s) (report at least one constrained MCU and one mid-range MCU). Run every experiment $N \geq 30$ times (or $N \geq 10$ if runs are very slow) and report mean $\pm$ SD.




Figure 13: Degradation of AUC-ROC under robustness tests.

The figure shows how the proposed pipeline's detection accuracy (AUC-ROC) degrades with increasing severity of occlusion, illumination variation, and viewpoint shift on the UCSD Ped2 dataset. While performance remains above 90% up to moderate distortion levels, severe occlusion (level 5) causes the sharpest drop to ~83%, highlighting sensitivity to blocked motion cues. Illumination changes and viewpoint shifts result in smaller but still notable reductions, demonstrating the importance of robust feature aggregation. These results emphasize that while the pipeline performs well under realistic perturbations, extreme conditions can degrade accuracy, motivating future improvements such as adversarial training or multimodal fusion.

Table 14: Runtime and Resource Footprint on Edge Hardware

Platform	Batch Size	Avg Latency (ms/frame)	Sustained Throughput (fps)	RAM Usage (MB)	Flash/ROM Usage (KB)	CPU Utilization (%)	GPU Used?	Power Draw (W)	Notes
Raspberry Pi 4 (A72)	1	28.4 ± 0.6	34.7	412	–	78	No	4.8	Baseline ARM CPU run
Jetson Nano (GPU on)	1	12.1 ± 0.3	82.6	625	–	22 (CPU)	Yes	7.2	GPU accelerated optical flow

Jetson Nano (CPU only)	1	35.7 ± 0.7	28.0	390	–	84	No	6.0	Slower without GPU
STM32 F4 (MCU)	1	96.5 ± 2.1	10.4	192	1032	~100	No	0.42	

Table 14 summarizes runtime and memory characteristics of the proposed pipeline on different edge hardware. The Raspberry Pi 4 achieves real-time operation at ~35 fps, while the Jetson Nano with GPU acceleration sustains ~83 fps with much lower latency and moderate power consumption. On CPU-only mode, throughput drops significantly, underscoring the importance of GPU acceleration for heavy motion features. The STM32F4 microcontroller runs the pipeline in a highly resource-constrained setting at ~10 fps, demonstrating feasibility but with higher latency and near-saturation of CPU cycles. Overall, the results confirm that the proposed method is deployable across a spectrum of IoT devices, with trade-offs between efficiency and resource usage.

Conclusion

The paper presents a general framework of powerful anomaly detection in the dense crowd situations implementing optical flow analysis, motion heat maps, Harris corner detection, and K-nearest centroid neighbor classification. The combination of these methods allows accurate extraction of motion descriptors while dealing with the problem of occlusion, overlapping movements, and different densities of a crowd. Optical flow gives an ongoing representation of motion dynamics, whereas heat maps only focus on intensity distributions that highlight abnormal areas. The incorporation of Harris corner detection makes the features distinctive; thus, an overlook of the important motion variation was avoided.

The fourth step, classification based on centroid-based clustering, explains a computationally efficient process of classifying between normal and abnormal behaviors that was highly reliable. This was the most false-alarm-free method with the highest level of real-time applicability that was essential to safety-critical surveillance systems. The framework exhibits flexibility in different settings, ranging anywhere from a large crowd to transportation terminals, where detecting anomalies was bound to be crucial for fast and corrective action.

Combining the use of motion-based descriptors with a centroid-based classification model, the methodology achieves an acceptable level of balancing the accuracy and the speed of identifying activities, outperforming conventional detection mechanisms in their scalability and resiliency. The results identify the growth in situational awareness and prove that such a strategy helps to mitigate risk proactively in crowded places as well.

The work sets out a feasible step-by-step roadmap to the deployment of intelligent surveillance systems capable of increasing safety in large crowds and managing major events, as well as a lead-in to the next evolution in automated people tracking and behavior analysis.

References

- [1] Rana, M., Mamun, Q., & Islam, R. (2024). Balancing security and efficiency: a power consumption analysis of a lightweight block cipher. *Electronics*, 13(21), 4325.
- [2] Yusoff, Z. Y. M., Ishak, M. K., Rahim, L. A., & Asaari, M. S. M. (2024). Improving Smart Home Security via MQTT: Maximizing Data Privacy and Device Authentication Using Elliptic Curve Cryptography. *Computer Systems Science & Engineering*, 48(6).
- [3] Tanksale, V. (2024). Efficient Elliptic Curve Diffie–Hellman Key Exchange for Resource-Constrained IoT Devices. *Electronics*, 13(18), 3631.
- [4] Sarveswaran, S., Shangkavi, G., Gowthaman, N., & Vasanthaseelan, S. (2021). Cryptography techniques and internet of things applications–A modern survey. *Int. J. of Aquatic Science*, 12(2), 2338-2371.
- [5] Kumar, P. R., & Goel, S. (2024, July). SE-ECC: A Secure and Efficient Symmetric Encryption Framework with Elliptic Curve Cryptography Key Exchange for Fog Communication. In *International Conference on Data Science and Applications* (pp. 343-358). Singapore: Springer Nature Singapore.
- [6] Ech-Chkaf, A., Oussous, S. A., El Allali, A., Beloualid, S., El Harrouti, T., El Aidi, S., ... & Tamtoui, A. (2021, December). Applying an enhanced elliptic curve integrated encryption scheme ECIES to enhance smart energy IoT platform security based on constrained protocol. In *International Conference of Reliable Information and Communication Technology* (pp. 498-511). Cham: Springer International Publishing.
- [7] Fadia, T., & Toufik, L. (2024). Elliptic curves cryptography for lightweight devices in IoT system. *Brazilian Journal of Technology*, 7(4), e73725-e73725.
- [8] Fazrina, N. (2024). Securing distributed sensor systems through adaptive encryption algorithms in 5G-based smart energy networks. *Open Journal of Robotics, Autonomous Decision-Making, and Human-Machine Interaction*, 9(11), 18-27.
- [9] Pandey, S., & Bhushan, B. (2024). Recent Lightweight cryptography (LWC) based security advances for resource-constrained IoT networks. *Wireless Networks*, 30(4), 2987-3026.
- [10] Bandaru, V. N. R., Kaligotla, V. G. S., Varma, U. D. S. P., Prasadharaju, K., & Sugumaran, S. (2024, July). A Enhancing Data Security Solutions for Smart Energy Systems in IoT-Enabled Cloud Computing Environments through Lightweight Cryptographic Techniques. In *IOP Conference Series: Earth and Environmental Science* (Vol. 1375, No. 1, p. 012003). IOP Publishing.
- [11] Awotunde, J. B., Adeniyi, A. E., Babatunde, A. O., Olagunju, M., Imoize, A. L., & Olanloye, O. D. (2024). 12 An Cryptographic Enhanced Lightweight Algorithm Towards Securing Wireless Networks and Big Data. *Computational Modeling and Simulation of Advanced Wireless Communication Systems*, 293.
- [12] Sekaran, R., Ramasamy, D., Basha, J. B. M., Maruthapillai, K., Annamalai, S., & Parasuraman, M. (2024, November). Enhancing IoT Security and Efficiency: Advanced Public Key Cryptographic Solutions for the Modern Era. In *2024 4th International Conference on Technological Advancements in Computational Sciences (ICTACS)* (pp. 1740-1747). IEEE.
- [13] Rana, M. (2024). An Improved Lightweight Block Cipher Algorithm for Securing Resource-Constrained IoT Devices.
- [14] Awotunde, J. B., Adeniyi, A. E., Babatunde, A. O., Olagunju, M., Imoize, A. L., & Olanloye, O. D. (2024). An Enhanced Lightweight Cryptographic Algorithm Towards Securing Wireless Networks and Big Data. In *Computational Modeling and Simulation of Advanced Wireless Communication Systems* (pp. 293-317). CRC Press.

- [15] Mustafa, R., Sarkar, N. I., Mohaghegh, M., & Pervez, S. (2024). A cross-layer secure and energy-efficient framework for the internet of things: a comprehensive survey. *Sensors* (Basel, Switzerland), 24(22), 7209.
- [16] Gour, A., Malhi, S. S., Singh, G., & Kaur, G. (2024). Hybrid cryptographic approach: for secure data communication using block cipher techniques. In *E3S Web of Conferences* (Vol. 556, p. 01048). EDP Sciences.
- [17] Awotunde, J. B., Oladipupo, E. T., Adeniyi, A. E., Babatunde, A. O., & Aremu, D. R. (2024). An enhanced internet of medical things using blockchain-enabled elliptic curve cryptography for security and privacy. *International Journal of Critical Computer-Based Systems*, 11(4), 280-306.
- [18] Kiran, N. C., Maindola, M., Kumara, B., Kaliappan, S., Bhatt, P., & Maranan, R. (2024, December). Secure Key Management in 5G Networks Using Elliptic Curve Integrated Encryption Scheme (ECIES) for Low-Power IoT Devices. In *2024 4th International Conference on Mobile Networks and Wireless Communications (ICMNWC)* (pp. 1-5). IEEE.
- [19] Ramakrishna, D., & Shaik, M. A. (2024). A comprehensive analysis of cryptographic algorithms: Evaluating security, efficiency, and future challenges. *IEEE Access*.
- [20] Goulart, A., Chennamaneni, A., Torre, D., Hur, B., & Al-Aboosi, F. Y. (2022). On wide-area IoT networks, lightweight security and their applications—a practical review. *Electronics*, 11(11), 1762.
- [21] SHAH, A. Performance Enhancement by Applying Lightweight Cryptography in IoT Devices.
- [22] Popoola, O., Rodrigues, M., Marchang, J., Shenfield, A., & Popoola, J. (2024). Hybrid encryption for smart home healthcare: ensuring data confidentiality and security.
- [23] Kakoulli, E., & Zacharioudakis, E. (2024, April). Exploration of the Role of Cryptoprocessors in Advancing IoT Security. In *2024 20th International Conference on Distributed Computing in Smart Systems and the Internet of Things (DCOSS-IoT)* (pp. 524-531). IEEE.
- [24] Raeisi-Varzaneh, M., Dakkak, O., Alaidaros, H., & Avci, İ. (2024). Internet of things: security, issues, threats, and assessment of different cryptographic technologies. *Journal of Communications*, 19(2).
- [25] Vanitha, V., Saritha, G., Hemalatha, K., Sathish, P., & Devi, P. G. (2024, November). Securing Edge Computing Environments with a Novel Chaos Theory-Based Lightweight Encryption Algorithm for Real-Time Data Protection. In *2024 International Conference on Integrated Intelligence and Communication Systems (ICIICS)* (pp. 1-8). IEEE.
- [26] Hussein, S. M. H. (2023). A Secure and Efficient Method for Communications and Energy Consumption Management in IoT Wireless Sensor Networks.
- [27] Mahamat, M. (2023). Secure and robust models for energy-efficient communications for the Internet of Things (Doctoral dissertation, Université de Technologie de Compiègne).
- [28] Khashan, O. A., Khafajah, N. M., Alomoush, W., & Alshinwan, M. (2024). Innovative energy-efficient proxy re-encryption for secure data exchange in wireless sensor networks. *IEEE Access*, 12, 23290-23304.
- [29] Biswas, S., Goswami, R. S., & Reddy, K. H. K. (2024). Advancing quantum steganography: a secure IoT communication with reversible decoding and customized encryption technique for smart cities. *Cluster Computing*, 27(7), 9395-9414.

- [30] Aljrees, T., Kumar, A., Singh, K. U., & Singh, T. (2023). Enhancing IoT security through a green and sustainable federated learning platform: leveraging efficient encryption and the quondam signature algorithm. *Sensors*, 23(19), 8090.
- [31] Mousavi, S. K., Ghaffari, A., Besharat, S., & Afshari, H. (2021). Security of internet of things based on cryptographic algorithms: a survey. *Wireless Networks*, 27(2), 1515-1555.
- [32] Yusoff, Z. Y. M., Ishak, M. K., & Rahim, L. A. (2024). Securing IoT edge device communication with efficient ECC middleware for resource-constrained systems. *Bulletin of Electrical Engineering and Informatics*, 13(6), 4285-4293.
- [33] Rehman, A., Abdullah, S., Fatima, M., Iqbal, M. W., Almarhabi, K. A., Ashraf, M. U., & Ali, S. (2022). Ensuring security and energy efficiency of wireless sensor network by using blockchain. *Applied Sciences*, 12(21), 10794.
- [34] Shakor, M. Y., Khaleel, M. I., Safran, M., Alfarhood, S., & Zhu, M. (2024). Dynamic AES encryption and blockchain key management: a novel solution for cloud data security. *IEEE Access*, 12, 26334-26343.
- [35] Al-Samhour, M., Novas Castellano, N., Abur-rous, M., & Gázquez Parra, J. A. (2024). Post-Quantum Cryptography for Wireless Sensor Network Using Key Agreement Super Singular on Hyperelliptic Curve.
- [36] Ifrim, R., Loghin, D., & Popescu, D. (2024). A systematic review of fast, scalable, and efficient hardware implementations of elliptic curve cryptography for blockchain. *ACM Transactions on Reconfigurable Technology and Systems*, 17(4), 1-33.
- [37] Venkatesan, C., Thamaraimanalan, T., Sathiya, T., Ranjithkumar, M., Radha, K., & Shanmuga Priya, K. (2024, September). Optimized Energy-Efficient Hybrid LW-AES and LW-ECC Security Framework for VLC and RF Networks. In *2024 International Conference on Communication, Computing and Energy Efficient Technologies (I3CEET)* (pp. 1844-1849). IEEE.
- [38] Bhagat, V., Kumar, S., Gupta, S. K., & Chaube, M. K. (2023). Lightweight cryptographic algorithms based on different model architectures: A systematic review and futuristic applications. *Concurrency and Computation: Practice and Experience*, 35(1), e7425.
- [39] Jammula, M., Vakamulla, V. M., & Kondoju, S. K. (2022). Performance evaluation of lightweight cryptographic algorithms for heterogeneous IoT environment. *Journal of Interconnection Networks*, 22(Supp01), 2141031.
- [40] Bobde, Y., Narayanan, G., Jati, M., Raj, R. S. P., Cvitić, I., & Peraković, D. (2024). Enhancing industrial IoT network security through blockchain integration. *Electronics*, 13(4), 687.