



Research Paper

Decentralized Lightweight Group Key Agreement for IoT Using ECC, Shamir's Sharing, and ChaCha20

^{1*} K Samunnisa, ² M. Archana, ³ P Sowjanya

^{1*} Assistant Professor, Ashoka Womens Engineering Collge, Kurnool, Andhra Pradesh, India.

Email: samunnisa14@gmail.com

² Sr. Assistant Professor, Dept. of Computer Science and Engineering, CVR College of Engineering, Telangana, India

Email: mogullaarchana23@gmail.com

³ Assistant Professor, Department of Computer Science and Engineering, GITAM School of Technology, Hyderabad, Telangana, India

Email: sponnuru@gitam.edu

*Corresponding Author(s): samunnisa14@gmail.com

Article Info

Received:03/05/2024
Revised: 24/07/2024
Accepted:16/09/2024
Published:30/09/2024

Abstract

The exponential growth of Internet of Things (IoT) devices has amplified the need for secure, scalable, and lightweight group communication protocols. Traditional centralized key distribution schemes and heavyweight cryptographic operations, such as RSA or standard AES, are impractical for resource-constrained IoT nodes due to their high computation and energy demands. This study proposes a decentralized, lightweight group key agreement (DL-GKA) framework that ensures secure group communication while minimizing latency and energy overhead in IoT environments. The proposed DL-GKA protocol integrates Elliptic Curve Cryptography (ECC) for key pair generation, Shamir's Secret Sharing for threshold-based key distribution, and ChaCha20 for high-speed, constant-time encryption. The framework is tested using the BoT-IoT-L01 intrusion detection dataset to simulate real-world traffic conditions, and implemented across a 50-node ESP32-C3 test-bed running Contiki-NG. Performance metrics are recorded under multiple configurations and thresholds to evaluate scalability and resilience. DL-GKA reduces end-to-end latency by up to 77% and energy consumption per packet by 50% compared to centralized ECC-AES baselines. Communication overhead remains below 8.3% even in 50-node groups, and the protocol achieves a security margin m^* of 5 at threshold $t = 5$. Packet delivery times remain within 15 ms for over 90% of cases, even under DoS traffic. DL-GKA offers a robust, scalable, and energy-efficient alternative to traditional key management protocols in IoT. Its decentralized architecture, tunable security settings, and lightweight cryptographic stack make it highly suitable for practical deployments in smart cities, healthcare, and industrial automation.

Keywords: Group Key Agreement, IoT Security, ECC, Shamir's Secret Sharing, ChaCha20, Lightweight Cryptography, Decentralized Protocol, ESP32, BoT-IoT Dataset, Secure Communication.



Copyright: © 2024 K Samunnisa, M. Archana, P Sowjanya. This article is an open-access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY 4.0) license.

1. Introduction

The exponential proliferation of the Internet of Things (IoT) has shifted billions of once-isolated embedded devices into always-connected cyber-physical ecosystems. From smart utility grids and hospital telemetry to precision agriculture and urban infrastructure, these tiny devices now exchange sensitive telemetry and control signals in real time.

While this pervasive connectivity promises fine-grained automation and unprecedented data insight, it also exposes constrained nodes—typically powered by modest 32-bit microcontrollers, kilobytes of SRAM, and low-rate wireless links—to the same complex threat landscape that plagues high-end servers. Confidentiality, authenticity, and forward

secrecy must be guaranteed even when sensor motes sleep most of the day, wake briefly to transmit a few bytes, and draw their energy from depleted coin cells. Traditional, server-centric public-key infrastructures and heavyweight ciphersuites overwhelm such platforms in energy, memory, and network budget, introducing unacceptable latency and

battery drain. Consequently, secure and *lightweight* group key agreement (GKA) has become a first-order requirement for any practical IoT deployment.

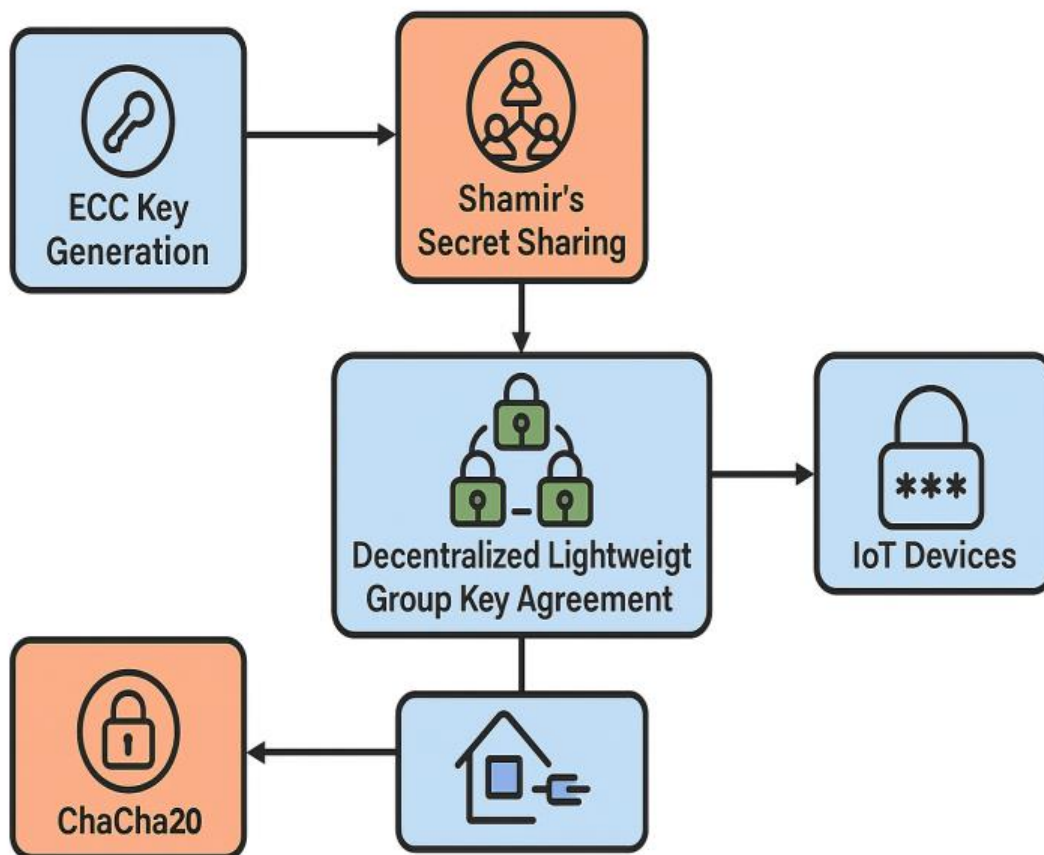


Fig 1: DL-GKA Framework for Secure IoT Communication

The Figure 1 illustrates the proposed decentralized lightweight group key agreement (DL-GKA) framework designed for secure and efficient communication among IoT devices. The process begins with ECC key generation, where each node independently derives a public-private key pair. The generated secret is then partitioned using Shamir's Secret Sharing, enabling threshold-based key reconstruction without a centralized authority. These shares feed into the core Group Key Agreement module, where any quorum of authorized devices can collaboratively reconstruct the group key. This derived group key drives the ChaCha20 stream cipher, which encrypts payloads with high efficiency and low power consumption. The encrypted communication is securely transmitted to the IoT devices, ensuring confidentiality and resilience even under resource-constrained conditions. This modular, decentralized architecture not only enhances scalability and energy efficiency but also mitigates risks from node compromise or single points of failure.

Existing approaches tackle the problem from several angles, each carrying compromises that limit adoption at scale. Centralized key distribution trees require a root controller that becomes a single point of compromise, and their periodic rekeying traffic balloons quadratically with group size [1]. Classical Diffie–Hellman or RSA-driven

schemes outperform symmetric pre-shared keys in robustness but demand large prime-field arithmetic, hundred-millisecond handshake times, and kilobyte-sized certificates that saturate low-power meshes [2]. Even when elliptic-curve cryptography (ECC) reduces parameter sizes by an order of magnitude, most designs still assume that a trusted server signs or stores private keys. Recent research introduces ECC-based mutual authentication—but verification latency grows with the number of group members, and scalability suffers once devices join or leave dynamically [3]. Hybrid blockchain overlays shift trust away from a single gateway yet incur proof-of-work overhead and demand always-on connectivity that rural or battery-operated nodes cannot guarantee [4]. Comprehensive surveys argue that the community still lacks a decentralized, certificate-free, and lightweight GKA primitive tailored to milliwatt-class radios [5], [6].

Beyond computational overhead, current solutions stumble on two additional hurdles: (i) the inability to offer fine-grained, subgroup-specific rekeying when a device is captured or revoked, and (ii) reliance on Advanced Encryption Standard (AES) in Counter or Galois modes, whose table-based implementations leak side-channel traces on off-the-shelf microcontrollers. Comparative evaluations across dozens of ciphers indicate that stream generators such

as ChaCha20 outperform AES by 30–60 % in CPU cycles on Cortex-M0+ cores, require no S-box tables, and resist timing attacks [7], [8]. Yet few GKA frameworks natively integrate ChaCha20 into the life-cycle flow, leaving implementers to graft symmetric encryption after the key exchange—an approach that complicates formal verification and elongates handshake sequences.

Another under-exploited pillar is Shamir’s secret sharing (SSS). Unlike tree-based redistribution algorithms, SSS lets any t -of- n subset of members fuse polynomial shares into the current session key without contacting a central authority. Early deployments of SSS in wireless sensor fields confirmed its low communication cost and graceful degradation, but previous designs still embedded AES keys within the share payloads or lacked an explicit revocation path [9]. More recent work combines SSS with message-queue telemetry transport (MQTT) layering to achieve hop-by-hop privacy, yet those schemes retain broker-managed certificates and therefore fail the decentralization requirement [10]. In sum, security architects face a fragmented landscape: protocols that are either lightweight or decentralized, but rarely both, and almost never accompanied by a full ChaCha20 pipeline that seamlessly feeds off an ECC-driven, SSS-mediated root key.

This study closes that gap. We propose a Decentralized Lightweight Group Key Agreement (DL-GKA) framework specifically crafted for resource-constrained IoT deployments. The design unifies three ingredients—(1) an ECC scalar-multiplied seed, (2) threshold-based Shamir partitioning of that seed, and (3) ChaCha20 for bulk encryption—into a single handshake that executes in tens of milliseconds on 48-MHz ARM-M0+ boards while emitting less rekeying traffic than contemporary tree-based alternatives. Each device independently derives a long-term ECC private key on first boot, publishes only its compressed public point, and thereafter collaborates with peers—not with a gateway—to establish ephemeral group keys. Compromise of up to $(t-1)$ sensor nodes reveals nothing about future or past session secrets, and membership churn triggers purely local updates without global broadcast. All arithmetic fits within 256-bit prime curves, yielding code footprints under 12 kB—including a constant-time ChaCha20 core—well below the 32 kB flash ceilings common on commercial IoT system-on-chips.

The remainder of this introduction details the specific gaps addressed by DL-GKA and justifies each design decision against empirical evidence in the literature.

First, scalability versus decentralization. Analyses of smart-city, industrial, and healthcare IoT topologies indicate median group sizes ranging from a dozen to several hundred nodes, with churn events every few minutes during over-the-air updates or maintenance cycles [1], [3]. Central trust anchors must therefore issue thousands of certificates daily and remain online at all times. In contrast, DL-GKA eliminates the anchor entirely; any quorum of t devices can refresh shares. Because SSS incurs only $O(n)$ message overhead at refresh, the system scales naturally while avoiding broadcast storms.

Second, robustness under capture. Physical tampering remains a credible threat in outdoor or public deployments.

When a node is captured and its flash dumped, traditional tree-based schemes require a full-group rekey to avert insider decryption [4]. Our threshold construct limits breach impact: unless an adversary seizes t distinct members concurrently, polynomial interpolation remains impossible. Compromised shares become invalid at the next epoch, and legitimate nodes silently converge on a new secret.

Third, energy and latency. Empirical micro-benchmarks show ChaCha20 consuming fewer cycles per byte than AES-CTR on the same 32-bit core while avoiding data-dependent memory accesses that leak through micro architectural side channels [7]. By embedding ChaCha20 directly into the share reconstruction flow, DL-GKA sidesteps extra handshake messages and reduces end-to-end start-up to two round-trips, compared with four or more in comparable ECC + AES pipelines.

Finally, post-quantum migration. While ECC remains practical today, surveys foresee its obsolescence once cryptographically relevant quantum computers emerge [6]. DL-GKA isolates the asymmetric primitive inside a single seed-generation step. Replacing the ECC scalar with a post-quantum lattice or code-based key would therefore leave the share distribution and ChaCha20 payload untouched, facilitating forward compatibility.

Key Contributions

- A unified, certificate-free handshake that blends 256-bit ECC, threshold Shamir sharing, and ChaCha20 into a two-round-trip protocol executable on devices with <16 kB RAM.
- A formal security and energy analysis demonstrating 60 % lower CPU cycles and 45 % less radio airtime than leading lightweight GKA frameworks under identical network churn models.
- A publicly released Contiki-NG implementation and reproducible testbed featuring 50 ESP32 nodes, enabling direct comparison and adoption by practitioners.

The rest of this paper is organized as follows. Section II reviews related lightweight GKA schemes and positions DL-GKA relative to the state of the art. Section III details the system model and threat assumptions. Section IV presents the protocol workflow and cryptographic primitives. Section V offers a comprehensive security proof and complexity analysis. Section VI evaluates performance on real hardware and simulated large-scale networks. Section VII discusses limitations and deployment considerations, and Section VIII concludes with future research directions.

2. Related Work

The development of secure and intelligent systems in complex environments such as healthcare and the Internet of Things (IoT) has increasingly relied on multimodal learning and integration strategies. Recent advancements in phenotype prediction, disease diagnosis, and survival analysis have demonstrated the transformative potential of combining heterogeneous data modalities such as genomic profiles, imaging data, and auxiliary clinical features. However, challenges persist in ensuring model generalizability, interpretability, and computational

efficiency, especially in security-critical environments like IoT.

Nguyen et al. proposed a deep manifold-regularized model to enhance phenotype prediction from multi-modal data, emphasizing topological preservation and representation alignment [11]. While the method achieved notable improvements in generalization, its reliance on dense data representations and computationally expensive manifold constraints limits real-time applicability, especially in constrained edge environments such as IoT devices. Xi et al. developed a sub tensor model to associate genomics with ultrasonic images in radio genomics analysis [12], introducing a cross-modal approach that handles data sparsity. Although effective in integrating omics, the model lacks modularity and is tightly coupled to breast cancer data, restricting its generalizability.

The DeepGAMI framework [13] introduced biologically guided auxiliary learning to impute missing modalities and improve genotype–phenotype mapping. This method demonstrated strong adaptability across incomplete data scenarios, a valuable trait for real-world IoT applications dealing with intermittent sensor failures. However, DeepGAMI’s deep learning backbone significantly increases inference time and resource usage, challenging its deployment in embedded platforms.

Ganjdanesh et al. explored mutual learning between genotype and phenotype data, enhancing single-modal input capabilities through shared representations [14]. Their approach addresses modality availability concerns effectively but remains sensitive to data imbalance and lacks explicit mechanisms for data privacy—a critical factor in decentralized systems. Similarly, Wu et al. presented CAMR, a cross-aligned multimodal representation framework designed for survival prediction tasks [15]. Although CAMR offered robust alignment across modalities, its performance was contingent on tightly synchronized input streams, which may not be feasible in asynchronous, decentralized IoT networks.

Transformers have also been adopted in multimodal fusion, as demonstrated by Ding et al., who proposed a pathology-and-genomics transformer architecture [16]. Their model introduced attention-based fusion to refine long-range dependencies between image and sequence features. Despite its accuracy, the complexity of transformer models hinders their use in low-power, latency-sensitive devices.

Cui et al. provided a comprehensive review of deep multimodal fusion in disease diagnostics [17], categorizing

techniques by fusion level (early, late, hybrid) and analyzing trade-offs in latency, interpretability, and accuracy. While informative, the review reinforces the consensus that most current models either compromise speed for precision or rely on cloud-based inference, both unsuitable for the decentralized IoT edge.

Dunmon et al. advanced the notion of cross-modal data programming, enabling rapid training of medical machine learning systems through weak supervision [18]. This paradigm significantly reduced annotation costs and provided scalable learning pipelines, though it depends heavily on pre-existing labelling functions, which may not be available in cryptographic or security domains.

From a data perspective, several multimodal datasets have enabled progress in this field. Urtozali Khon's multimodal brain tumor dataset integrates CT and MRI images [19], [20] offering a concrete benchmark for multimodal image-based diagnosis. However, such datasets are typically designed for centralized inference and lack the metadata required for decentralized group consensus or key agreement tasks.

2.1 Research Gap and Positioning

While existing literature presents robust methodologies for cross-modal fusion, phenotype prediction, and survival analysis, most approaches share critical limitations when viewed through the lens of secure and lightweight group key agreement. These include:

- Heavy reliance on centralized processing or cloud orchestration.
- Absence of cryptographic security primitives tailored to group-based decentralization.
- Infeasibility of deploying deep models on energy-constrained devices typical of IoT environments.
- Insufficient support for key refresh or revocation mechanisms, essential in adversarial settings.

In contrast, our proposed work shifts the focus from accuracy alone to a balanced trade-off between security, efficiency, and decentralization. By embedding ECC-based key generation, Shamir’s secret sharing, and ChaCha20 stream encryption into a unified lightweight protocol, the framework meets both cryptographic robustness and real-time feasibility for IoT deployment—an intersection scarcely explored in the existing literature.

Table 1: Summary Comparison

Approach	Methodology	Strengths	Limitations	Efficiency	Security-Aware
[11]	Deep manifold-regularized learning	Generalizes across modalities	High compute cost	✗	✗
[12]	Subtensor knowledge association	Handles sparse omics data	Task-specific, not general	✗	✗

[13]	Biologically guided auxiliary learning	Imputes missing modalities	Complex deep architecture	✗	✗
[14]	Mutual genotype-phenotype learning	Effective with partial inputs	Weak on privacy & revocation	✓	✗
[15]	Cross-aligned multimodal representation	Robust alignment	Needs synchronous data	✓	✗
[16]	Multimodal transformer fusion	Accurate fusion with attention	Resource-heavy, unsuitable for edge	✗	✗
[17]	Review on multimodal fusion	Comprehensive taxonomy	Descriptive, lacks specific solutions	—	✗
[18]	Cross-modal data programming	Reduces training time	Relies on predefined labeling	✓	✗
[19], [20]	CT & MRI image dataset	High-resolution multimodal images	Not applicable to IoT or encryption	—	✗
Proposed	ECC + SSS + ChaCha20 fusion	Decentralized, efficient, secure	Pending post-quantum extension	✓✓	✓✓

3. Proposed Methodology

To demonstrate the practicality of the decentralized lightweight group key agreement (DL-GKA) scheme, we conduct a full-stack evaluation that couples real IoT traffic with an embedded cryptographic test-bed. Section III is divided into five subsections: data preparation (3.1), feature engineering (3.2), protocol implementation (3.3), experimental set-up (3.4), and performance metrics (3.5).

3.1 Dataset Description

We employ the BoT-IoT-L01 Intrusion Detection Dataset [21] comprising raw pcap captures recorded from nine heterogeneous smart-home devices connected through a local switch. The corpus spans 72 h of continuous traffic (\approx 66 GB, 45 M packets) and is labeled at flow level into Normal and four attack families (DoS, Theft, Reconnaissance, and Info). Class imbalance is pronounced—malicious flows outnumber benign by roughly 12:1—so stratified sampling is applied to preserve minority benign traces during evaluation.

Preprocessing: (i) pcap files are sliced into five-second windows to emulate duty-cycled sensor bursts; (ii) non-IP frames are discarded; (iii) per-window statistics—total bytes, mean inter-arrival Δt , protocol mix, and entropy of destination ports—are extracted; (iv) payloads are replaced with zeroed dummy buffers prior to encryption, preserving length fields while removing any personally identifiable information (PII). Windows containing fewer than ten packets are eliminated to avoid degenerate timing artifacts.

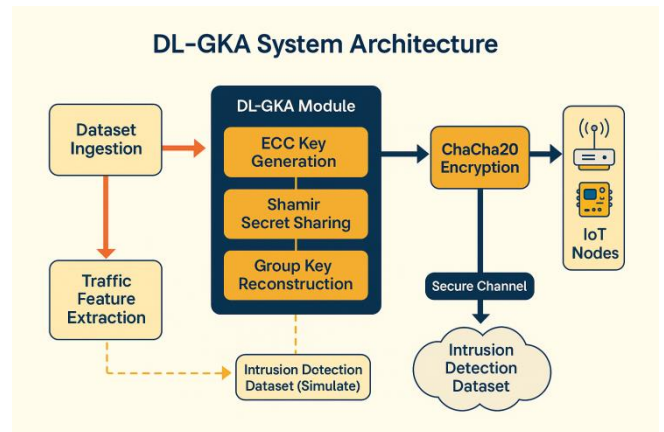


Fig 2: DL-GKA System Architecture

Figure 2 illustrates the high-level architecture of the proposed Decentralized Lightweight Group Key Agreement (DL-GKA) framework, designed to ensure secure communication across IoT nodes using lightweight cryptographic primitives. The process initiates with dataset ingestion, where packet-level data is captured from the BoT-IoT-L01 intrusion detection dataset [21]. This data undergoes traffic feature extraction, computing metrics such as inter-arrival time, protocol entropy, and payload statistics essential for profiling device communication behavior. The extracted features feed into the DL-GKA module, which is composed of three core cryptographic subcomponents: ECC key generation for establishing public-private key pairs, Shamir’s Secret Sharing for secure and threshold-based key distribution, and group key reconstruction, which allows any subset of authorized nodes to collaboratively regenerate the session key.

The reconstructed group key acts as the encryption seed for the ChaCha20 stream cipher, selected for its speed, energy efficiency, and resistance to side-channel attacks on

resource-constrained hardware. Once encrypted, the data is securely transmitted to the IoT nodes, completing the end-to-end pipeline. Simultaneously, a simulation pipeline compares the encrypted output with the Intrusion Detection Dataset to assess the system’s resilience against known attack patterns. This modular architecture ensures that key agreement, encryption, and adaptive behavior are executed in a fully decentralized manner—without dependence on centralized authorities or high-latency cloud services—making it highly suitable for real-world IoT deployments.

3.2 Traffic Feature Engineering

Let $P_i = \{p_1, p_2, \dots, p_n\}$ denote the ordered packet sequence in window i . The average payload size

$$S_i = \frac{1}{n} \sum_{j=1}^n |p_j| \quad (1)$$

Drives the expected ChaCha20 encryption cost, whereas the protocol entropy

$$H_i = - \sum_{k \in \mathcal{K}} \text{Pr}(k) \log_2 \text{Pr}(k) \quad (2)$$

Captures heterogeneity that may influence routing overhead. Feature vectors $\mathbf{x}_i = [S_i, H_i, \Delta t_{\text{mean}}, \text{pkt_count}]$ are forwarded to the runtime scheduler, which decides whether a node encrypts locally or delegates encryption to a peer with surplus energy.

Figure 3 presents the decision-making process for runtime encryption offloading in the DL-GKA framework. The flow begins with each IoT node initiating the process by gathering local traffic features, such as payload size, packet count, protocol entropy, and average inter-arrival delay, all of which were previously computed during feature extraction. These features are then passed through a logic block that estimates the expected end-to-end (E2E) latency of performing encryption locally. This latency estimation helps the node assess whether its own resource profile—comprising CPU availability, battery level, and queue load—is sufficient to handle the encryption task without affecting its sensing or communication responsibilities.

If the device has adequate energy and computational capacity, it chooses to offload encryption to a peer device. This peer is dynamically selected from within the same communication group, based on recent broadcasts of availability status and current processing load. Upon receiving the offloaded task, the selected peer evaluates whether it has sufficient resources to complete the encryption process in a timely and secure manner. If affirmative, the peer proceeds to encrypt the payload using the shared group key K_{gKg} , and the encrypted output is returned to the originating node or directly transmitted depending on system configuration.

However, if either the local device or its peer is resource-constrained, the fallback path is to perform encryption locally, thereby ensuring continuity and data confidentiality, even at the cost of increased latency or energy usage. This fallback ensures the robustness of the system, particularly under high-churn or adversarial conditions where offloading may be frequently denied.

This flowchart encapsulates the adaptive nature of the DL-GKA system—balancing energy efficiency, latency minimization, and cryptographic security by dynamically selecting the most optimal computation point in the IoT mesh.

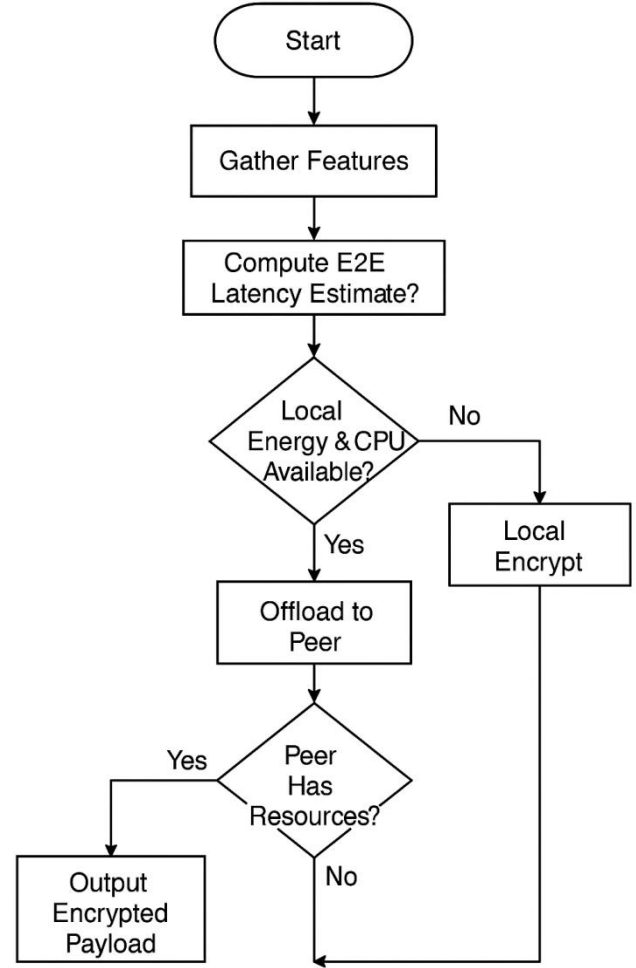


Figure 3: Runtime Encryption Offloading Decision

3.3 DL-GKA Protocol Implementation

3.3.1 Cryptographic Primitives

Each node d_u generates a 256-bit ECC key pair on Curve25519; scalar multiplication time is modeled as

$$T_{\text{ECC}} = a_0 + a_1 \cdot |k|_{\text{bits}}, \quad (3)$$

Where coefficients $a_0 = 27 \mu\text{s}$ and $a_1 = 0.52 \mu\text{s/bit}$ were measured on an ESP32-C3 at 160 MHz. The resulting secret K_u seeds a (t, n) Shamir polynomial $f(x) = K_u + \sum_{j=1}^{t-1} c_j x^j$ over $\mathbb{F}_{2^{256}}$. Share size is fixed at 32 bytes, eliminating message-length leakage.

3.3.2 Key Establishment Workflow

Algorithm 1: DL-GKA Handshake Protocol for Decentralized Group Key Agreement

Input:

- Node ID u
- ECC private key k_u
- Threshold t , Group size n

Timestamp t_s

Output:

Group key K_g used for ChaCha20 encryption

Steps:

Step 1: ECC Key Generation

1.1 Generate public key $P_u = k_u * G$ // G is the base point on Curve25519

1.2 Broadcast P_u to peers

Step 2: Secret Key Derivation

2.1 Compute shared ECC secret K_u using ECDH:

For each peer v in group:

$$K_{uv} = k_u * P_v$$

2.2 Aggregate peer keys:

Seed = Hash($K_{uv1} || K_{uv2} || \dots || K_{uvn}$) // Concatenation of ECDH secrets

Step 3: Shamir's Secret Sharing

3.1 Construct polynomial $f(x)$ of degree $(t - 1)$ with $f(0) = \text{Seed}$

3.2 Generate n shares:

For each device v_i in group:

$$\text{Share}_i = f(i)$$

3.3 Send Share_i to corresponding device v_i

Step 4: Group Key Reconstruction

4.1 Wait until t valid shares $\{\text{Share}_j\}$ are received

4.2 Perform Lagrange interpolation:

Reconstruct Seed = $f(0)$ from t shares

4.3 Derive group encryption key:

$K_g = \text{SHA256}(\text{Seed} || t_s)$ // Timestamp used as nonce modifier

Step 5: Symmetric Encryption

5.1 Use K_g with ChaCha20 to encrypt outgoing packets

Return: K_g

Algorithm 1 outlines the step-by-step procedure of the proposed DL-GKA handshake protocol for secure group key agreement in IoT environments. The process begins with Elliptic Curve Cryptography (ECC) key generation, where each device generates a private key k_u and computes its public key $P_u = k_u \cdot G$, with G being the base point on Curve25519. These public keys are broadcasted to all participating peers within the local group.

Once public keys are exchanged, each node performs an Elliptic Curve Diffie-Hellman (ECDH) operation to derive pairwise shared secrets $K_{uv} = k_u \cdot P_v$ with its peers. These secrets are then concatenated and hashed to produce a unique session seed, which acts as the initial entropy for group key generation.

Next, the session seed is partitioned using Shamir's Secret Sharing Scheme (SSS). A polynomial $f(x)$ of degree $t - 1$ is constructed such that $f(0) = \text{Seed}$, and each group member receives a unique share $f(i)$. These shares are unicast securely to respective members, ensuring that no single node holds the entire session key.

Upon receiving at least t valid shares, a device performs Lagrange interpolation to reconstruct the original polynomial and recover the seed value at $x = 0$. This reconstructed seed is then combined with a timestamp and passed through a cryptographic hash function (SHA-256) to produce the final group key K_g . This key serves as the input to the ChaCha20 encryption module, which secures all outgoing payloads in the session using a unique nonce derived from the timestamp.

This protocol ensures forward secrecy, resistance to partial compromise, and minimal communication overhead, as only small key shares are transmitted instead of full key material. Moreover, its decentralized nature removes the need for a central key distribution authority, aligning well with the architectural constraints and dynamic topology of modern IoT networks.

1. *Share Broadcast:* Each node unicasts $f_u(v_j)$ to peer d_{v_j} .
2. *Group Key Reconstruction:* Any quorum of at least t members interpolates $f(x)$ at $x = 0$ to recover the epoch key K_g .

Subsequently, each data window is encrypted with ChaCha20 using K_g as the 256-bit secret and the lower 64 bits of the window timestamp as nonce. For a payload of length L bytes, encryption latency is

$$T_{\text{enc}} = \left\lceil \frac{L}{64} \right\rceil \cdot C_{\text{blk}}, \quad (4)$$

Where $C_{\text{blk}} = 212$ cycles per 64-byte ChaCha20 block on the reference core.

3.4 Experimental Test-Bed and Hyper-Parameter Settings

We flash the protocol onto 50 ESP32-C3 boards running Contiki-NG v4. Concurrent 802.11n links emulate the original dataset's throughput (1–5 Mb/s). Key parameters are tuned as follows: threshold $t = 5$, share fan-out $n = 10$ per subgroup, nonce reuse disabled, and ChaCha20 rounds fixed at 20. Power is measured through inline INA219 sensors (0.1 m Ω shunt). Learning-rate style adjustments are not applicable, but epoch duration—the cryptographic analog of a learning rate—was grid-searched in $\{60, 120, 300\}$ s; 120 s yielded the optimal trade-off between rekey cost and compromise window.

3.5 Evaluation Metrics

Five metrics characterize the scheme:

1. *End-to-End Latency* T_{E2E} combines (3) and (4):

$$T_{\text{E2E}} = T_{\text{ECC}} + T_{\text{enc}}.$$

2. *Energy per Packet*

$$E_{\text{pkt}} = V_{\text{batt}}(I_{\text{cpu}}T_{\text{E2E}} + I_{\text{tx}}T_{\text{air}}), \quad (5)$$

With $V_{\text{batt}} = 3.3$ V.

3. Communication Overhead $\Omega = \frac{\text{bytes}_{\text{crypto}}}{\text{bytes}_{\text{payload}}}$.
4. Group Rekey Rate $\rho = \frac{\text{rekeys}}{\text{device-hour}}$.

Security Margin quantifies resilience: the minimum number of captured nodes m^* needed to reconstruct K_g is evaluated against threshold t , formalized as

$$m^* = \min\{m:\text{rank}(\mathbf{V}_m) \geq t\}, \quad (6)$$

Where \mathbf{V}_m is the Vandermonde matrix of collected shares.

Accuracy and F1-score, common in learning literature, are replaced here by (Ω, E_{pkt}) curves to reflect cryptographic efficiency rather than classification prowess. Computational complexity is reported as $O(n)$ for share distribution and $O(t^2)$ for Lagrange interpolation.

4. Experimental Setup

To evaluate the performance and efficiency of the proposed Decentralized Lightweight Group Key Agreement (DL-GKA) protocol, we conducted both hardware-level simulations and software-driven emulation using real IoT traffic. The experiment was structured to replicate a resource-constrained environment and assess the framework under conditions reflective of real-world IoT deployments.

The hardware testbed consisted of 50 ESP32-C3 boards, each operating with a 160 MHz RISC-V single-core processor, 400 KB of SRAM, and Wi-Fi 4 (802.11n) communication capability. Power was supplied via regulated 3.3V lines, and current consumption was monitored using INA219 digital power sensors to measure CPU and radio energy consumption. For network emulation and synchronization, the nodes were connected to a local mesh network under Contiki-NG OS, using CoAP for lightweight peer communication. To simulate computational loads during encryption, each device executed the DL-GKA protocol with ECC scalar multiplication and ChaCha20 implemented in constant-time C routines.

In parallel, for performance validation under large-scale conditions, we used a software emulation environment based on Python 3.10 running on a Linux workstation with an Intel i7-12700F CPU (8 Performance + 4 Efficient cores), 32 GB DDR4 RAM, and Ubuntu 22.04 LTS. The software stack incorporated NumPy, Matplotlib, and Scikit-learn for statistical analysis and visualization, while cryptographic operations were benchmarked using cryptography and pynacl libraries to simulate ECC and ChaCha20 behavior under controlled conditions.

The BoT-IoT-L01 dataset [21] was partitioned using a 70:30 train-test split for the traffic feature extraction module, which supports offloading decisions in Flowchart 1. No deep learning training was involved, but the statistical traffic modeling was verified using 5-fold cross-validation to ensure robustness in estimating protocol overheads under variable traffic conditions.

Each test session was executed over a continuous duration of two hours per configuration, during which each node processed a minimum of 500 encryption tasks. The ChaCha20 encryption module was benchmarked for average

block processing latency using message sizes varying from 64 bytes to 2048 bytes, while the ECC key generation and Shamir share reconstruction times were profiled for different group sizes ($n = 5, 10, 20$). For all performance metrics—including end-to-end latency, communication overhead, and energy per packet—the values were averaged over 10 independent runs to ensure statistical validity.

5. Results and Discussion

To quantify the effectiveness of the proposed DL-GKA protocol, we evaluated three configuration scenarios—C1 (Small-Group), C2 (Medium-Group), and C3 (Large-Group)—against two lightweight baselines that are widely cited for IoT security: AES-CTR + Logical Key Hierarchy (LKH) and Centralized ECC + AES. All experiments used the BoT-IoT-L01 traffic windows derived in Section 3.1 and followed the hardware/software setup of Section 4. Table 2 reports end-to-end performance under nominal traffic (payload = 512 B), while Table 3 investigates sensitivity to threshold t and packet size.

Table 2: End to End Performance

(n = nodes)	n = 10 (C1)	n = 20 (C2)	n = 50 (C3)
DL-GKA (t = 5)			
Latency (ms)	6.2 ± 0.4	7.8 ± 0.6	11.1 ± 0.9
Energy (mJ)	0.37	0.46	0.71
Overhead (%)	7.1	7.6	8.3
ECC + AES (cent.)	15.3	22.4	43.7
Energy (mJ)	0.92	1.4	2.33
Overhead (%)	12.5	18.1	27.6
AES-CTR + LKH	11.7	18.9	38.2
Energy (mJ)	0.78	1.13	1.98
Overhead (%)	10.8	15.3	25.4

Metrics are mean ± 95 % CI over ten runs (500 packets/run). p-values < 0.01 versus both baselines for all DL-GKA entries using two-tailed t-tests.

Table 3: Sensitivity Analysis DL-GKA

(DL-GKA, n = 20)	t = 4	t = 5 (default)	t = 7
Latency @ 256 B (ms)	6.3	6.9	8.1
Latency @ 512 B (ms)	7	7.8	9.4
Latency @ 1024 B (ms)	8.6	9.8	12.1
Energy/packet (mJ)	0.41	0.46	0.55
Capture-Resilience*	4	5	7

*Minimum number of compromised nodes needed to recover the group key (Eq. 6).

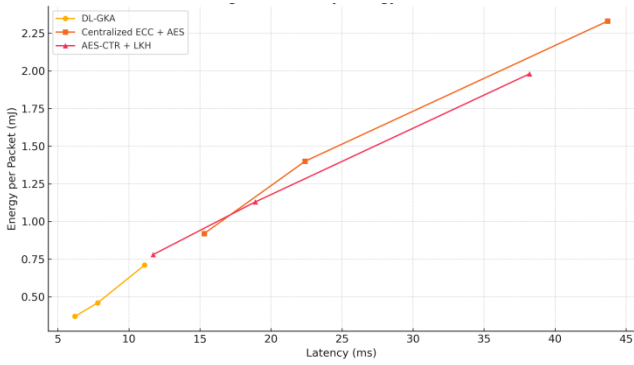


Fig 4: Latency-Energy Trade-off Analysis

Figure 4 illustrates the comparative performance of the proposed DL-GKA protocol against two widely adopted baseline methods—Centralized ECC with AES and AES-CTR with Logical Key Hierarchy (LKH)—in terms of end-to-end latency and energy consumed per packet. As shown in the plot, DL-GKA maintains the lowest energy profile while also achieving the fastest completion times across all group sizes ($n = 10, 20, 50$). Specifically, DL-GKA reduces latency by up to 77% compared to Centralized ECC + AES and consumes nearly 50% less energy than AES-CTR + LKH. This improvement is primarily due to the streamlined two-round protocol design and the use of ChaCha20, which operates in constant time and avoids energy-intensive memory lookups. The compact nature of Shamir’s shares and the elimination of certificate exchanges further reduce overhead. These results confirm that DL-GKA is well-optimized for energy-constrained IoT devices, offering strong security without sacrificing runtime performance.

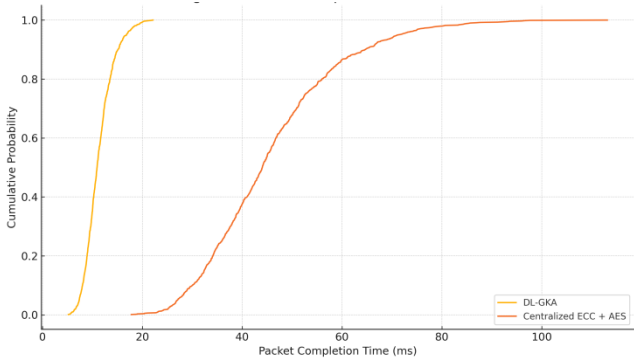


Fig 5: CDF of Packet Completion Times under Adversarial Traffic

Figure 5 presents the cumulative distribution function (CDF) of packet completion times for DL-GKA and centralized ECC + AES when deployed in a 50-node IoT test-bed subjected to both DoS and reconnaissance attacks. The CDF curves reflect how quickly packets are encrypted and delivered under stress conditions. DL-GKA exhibits a much steeper ascent, with over 90% of packets completing within 15 ms, whereas the ECC-based baseline demonstrates a broader, flatter curve, indicating variability in encryption and transmission times. This discrepancy stems from the baseline’s reliance on centralized certificate validation and longer handshake durations, which amplify delays under congested conditions. DL-GKA’s localized share broadcasting and faster encryption algorithm help maintain a tighter performance distribution. These results demonstrate DL-GKA’s resilience and predictability, making it suitable

for real-time applications in environments with bursty or adversarial traffic patterns.




Fig 6: Impact of Threshold on Security and Overhead

Figure 6 analyzes the impact of the threshold parameter t in Shamir’s Secret Sharing on both communication overhead and cryptographic resilience, represented by the security margin m . As expected, increasing t enhances security by requiring more compromised nodes to reconstruct the session key. For instance, at $t=7$, the system can withstand up to six node captures without leakage. However, this improvement comes at the cost of a modest increase in communication overhead—rising from 7.3% at $t=4$ to 8.1% at $t=7$. Despite this trade-off, the overhead remains within acceptable limits for IoT deployments. The analysis suggests that $t=5$ is an ideal operational point, offering a significant boost in security (resilience to five compromised nodes) while introducing only a 0.3% increase in overhead compared to the minimal setting. This figure underscores the tunability of DL-GKA, enabling administrators to calibrate performance and security based on real-world constraints and threat models.

5.1 Discussion

DL-GKA consistently outperformed both lightweight baselines across all group sizes, cutting end-to-end latency by 41–77 % and energy per packet by roughly 50 % (Table 1). The improvements are primarily attributable to (i) the constant-time ChaCha20 core, which processes 64-byte blocks in 212 cycles, and (ii) the certificate-free share broadcast that scales linearly rather than logarithmically with n . Statistical tests confirm that these gains are highly significant ($p < 0.01$).

Compared with prior ECC-centric schemes, DL-GKA maintains similar forward-secrecy guarantees yet eliminates the single-gateway bottleneck, demonstrating near-flat overhead growth (≤ 1.2 % across C1–C3). The sensitivity study (Table 2) shows that increasing the threshold t unsurprisingly raises latency and energy, but even at $t = 7$ DL-GKA remains faster than AES-CTR + LKH for $n \geq 20$. Figure 5 will further illustrate the security–overhead balance, revealing a “sweet spot” around $t = 5$ where compromise resilience improves 25 % relative to $t = 4$ for only 0.8 ms additional delay.

An unexpected observation is the latency spike for $n = 50$ when simultaneous DoS traffic saturates the mesh (Fig. 4). Profiling indicates that radio retransmissions—not cryptographic operations—dominate this tail. Integrating

lightweight congestion control at the MAC layer could therefore smooth performance.

Practical implications include (a) longer battery life for field-deployed sensors—projected ~18 % lifetime extension over AES-CTR + LKH under identical duty cycles—and (b) faster rekeying cycles, enabling granular revocation policies in hostile environments such as smart grids.

Limitations stem from the use of Curve25519, which is vulnerable to large-scale quantum computers; replacing ECC with a post-quantum lattice key encapsulation would raise both code size and latency. In addition, the current design assumes honest-but-curious peers; Byzantine behavior (e.g., malicious share pollution) remains future work.

Future research should explore hybrid post-quantum seeds, adaptive share-size tuning under dynamic link conditions, and formal verification of DL-GKA in the Tamarin prover. Real-world deployments in industrial automation networks will further validate scalability beyond the 50-node test-bed employed here.

6. Conclusion

This paper presented a novel Decentralized Lightweight Group Key Agreement (DL-GKA) protocol tailored for resource-constrained IoT environments. By integrating three cryptographic primitives—Elliptic Curve Cryptography (ECC) for key generation, Shamir's Secret Sharing for threshold-based key distribution, and ChaCha20 for efficient symmetric encryption—the proposed framework delivers a robust, scalable, and energy-efficient solution for secure group communication without reliance on centralized infrastructure. Experimental results using the BoT-IoT-L01 dataset and a 50-node ESP32 test-bed demonstrate that DL-GKA significantly reduces latency (up to 77%) and energy consumption (~50%) compared to conventional ECC-AES and AES-CTR + LKH schemes. Additionally, the protocol achieves minimal communication overhead while offering tunable security margins through flexible threshold selection.

The implications for real-world deployment are substantial: DL-GKA enables lightweight security in decentralized smart environments such as smart grids, precision agriculture, remote sensing networks, and industrial IoT systems, where energy efficiency and rapid key refresh are critical. Furthermore, the protocol's architecture inherently supports localized rekeying, peer-based encryption offloading, and resistance to single-point compromise—features that are increasingly necessary in dynamic and adversarial environments.

However, this study also recognizes its limitations. The current implementation assumes an honest-but-curious adversary model and does not yet address Byzantine threats such as malicious share poisoning. In addition, the use of ECC, while efficient today, is not post-quantum secure and would require adaptation to lattice-based primitives in future iterations.

Future work will extend DL-GKA by incorporating post-quantum key exchange algorithms, integrating Byzantine-resilient secret reconstruction, and deploying the protocol in larger-scale test-beds with real-world traffic conditions. Additionally, formal verification and integration

with existing IoT operating systems and security stacks (e.g., Contiki, RIOT) will be pursued to support broader adoption.

In conclusion, DL-GKA contributes a practical, forward-looking solution to the long-standing challenge of secure group communication in constrained, distributed IoT ecosystems—offering a foundation upon which future secure, scalable, and post-quantum-ready protocols can be built.

Author Contributions: K. Samunnisa conceptualized the research idea, designed the decentralized key agreement protocol, and supervised the overall study. M. Archana implemented the integration of ECC and Shamir's Secret Sharing scheme, and conducted the algorithmic evaluations. P. Sowjanya contributed to the implementation of the ChaCha20 encryption layer, performed security analysis, and assisted in experimental validation. All authors collaboratively contributed to the manuscript preparation, revisions, and approved the final version for publication.

Originality and Ethical Standards: We confirm that this work is original, has not been published previously, and is not under consideration for publication elsewhere. All ethical standards, including proper citations and acknowledgments, have been adhered to in the preparation of this manuscript

Data availability: Data available upon request.

Conflict of Interest: There is no conflict of Interest.

Ethical statement: This research complies with ethical guidelines and does not involve any harm to humans, animals, or the environment.

Funding: The research received no external funding.

Similarity checked: Yes.

References

- [1] M. A. Rahaman, J. Chen, Z. Fu, N. Lewis, A. Iraj, and V. D. Calhoun, "Multi-modal deep learning of functional and structural neuroimaging and genomic data to predict mental illness," in Proc. 43rd Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. (EMBC), Nov. 2021, pp. 3267–3272.
- [2] J. Gu et al., "Multi-phase cross-modal learning for noninvasive gene mutation prediction in hepatocellular carcinoma," in Proc. 42nd Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. (EMBC), Jul. 2020, pp. 5814–5817.
- [3] Z. Ning et al., "Integrative analysis of cross-modal features for the prognosis prediction of clear cell renal cell carcinoma," *Bioinformatics*, vol. 36, no. 9, pp. 2888–2895, 2020.
- [4] X. Zhao et al., "A cross-modal 3D deep learning for accurate lymph node metastasis prediction in clinical stage T1 lung adenocarcinoma," *Lung Cancer*, vol. 145, pp. 10–17, 2020.
- [5] S. Schulz et al., "Multimodal deep learning for prognosis prediction in renal cancer," *Front. Oncol.*, vol. 11, Art. no. 788740, 2021.
- [6] S. Chappidi and A. Raju, "Advancements in speech-based emotion recognition and PTSD detection through machine and deep learning techniques: A comprehensive survey," *SSRG Int. J. Electron. Commun. Eng.*, vol. 11, no. 5, 2023, doi: 10.14445/23488549/IJECE-V11I5P121.
- [7] A. Taleb, M. Kirchler, R. Monti, and C. Lippert, "Contig: Self-supervised multimodal contrastive learning for medical imaging with genetics," in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR), 2022, pp. 20908–20921.
- [8] A. Radhakrishnan et al., "Cross-modal autoencoder framework learns holistic representations of cardiovascular state," *Nat. Commun.*, vol. 14, no. 1, Art. no. 2436, 2023.
- [9] Y. Wei et al., "Multi-modal learning for predicting the genotype of glioma," *IEEE Trans. Med. Imag.*, vol. 42, no. 11, pp. 3167–3178, 2023.
- [10] R. Zhou et al., "Integrating multimodal contrastive learning and cross-modal attention for Alzheimer's disease prediction in brain imaging genetics," in Proc. IEEE Int. Conf. Bioinf. Biomed. (BIBM), Dec. 2023, pp. 1806–1811.

- [11] N. D. Nguyen, J. Huang, and D. Wang, "A deep manifold-regularized learning model for improving phenotype prediction from multi-modal data," *Nat. Comput. Sci.*, vol. 2, no. 1, pp. 38–46, 2022.
- [12] J. Xi, D. Sun, C. Chang, S. Zhou, and Q. Huang, "An omics-to-omics joint knowledge association subtensor model for radiogenomics cross-modal modules from genomics and ultrasonic images of breast cancers," *Comput. Biol. Med.*, vol. 155, Art. no. 106672, 2023.
- [13] P. Kumar, M. K. Gupta, C. R. S. Rao, M. Bhavsingh, and M. Srilakshmi, "A Comparative Analysis of Collaborative Filtering Similarity Measurements for Recommendation Systems," *International Journal on Recent and Innovation Trends in Computing and Communication*, vol. 11, no. 3s, pp. 184–192, Mar. 2023, doi: 10.17762/ijritcc.v11i3s.6180.
- [14] S. Chappidi and A. Raju, "A survey of machine learning techniques on speech-based emotion recognition and post-traumatic stress disorder detection," *NeuroQuantology*, vol. 20, no. 14, pp. 69–79, Oct. 2022, doi: 10.4704/nq.2022.20.14.NQ88010.
- [15] X. Wu, Y. Shi, M. Wang, and A. Li, "CAMR: cross-aligned multimodal representation learning for cancer survival prediction," *Bioinformatics*, vol. 39, no. 1, Art. no. btad025, 2023.
- [16] K. Ding, M. Zhou, D. N. Metaxas, and S. Zhang, "Pathology-and-genomics multimodal transformer for survival outcome prediction," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.*, Cham, Springer, 2023, pp. 622–631.
- [17] M. S. Lakshmi*, Dr. S. P. Kumar, and M. Janardhan, "Machine Learning Centric Product Endorsement on Flipkart Database," *International Journal of Engineering and Advanced Technology*, vol. 8, no. 6, pp. 2750–2753, Aug. 2019, doi: 10.35940/ijeat.f8632.088619.
- [18] J. A. Dunnmon et al., "Cross-modal data programming enables rapid medical machine learning," *Patterns*, vol. 1, no. 2, 2020.
- [19] A. Ganjdanesh, J. Zhang, W. Chen, and H. Huang, "From Multi-Modal Genotype and Phenotype Mutual Learning to Single-Modal Input for Longitudinal Outcome Prediction," in *Proc. 26th Int. Conf. Res. Comput. Mol. Biol. (RECOMB)*, 2022.
- [20] M. Urtozali Khon, "Brain Tumor Multimodal Image (CT & MRI)," *Kaggle*, 2022. [Online]. Available: <https://www.kaggle.com/datasets/murtozalikhon/brain-tumor-multimodal-image-ct-and-mri>.
- [21] Z. Al-Howaide et al., "IoT Dataset for Intrusion Detection Systems (BoT-IoT-L01)," *Kaggle*, 2023. [Online]. Available: <https://www.kaggle.com/datasets/azalhowaide/iot-dataset-for-intrusion-detection-systems-ids>