International Journal of
INTELLIGENT SYSTEMS AND APPLICATIONS IN ENGINEERING
ISSN:2147-6799 www.ijisae.org Original Research Paper

Lightweight Zero-Knowledge Proof-Based Cryptographic Protocols for Secure and Privacy-Preserving Wearable IoT

Chaudhari Minal Yuvaraj[1], Dr. Sanmati Kumar Jain[2]

Submitted: 01/11/2024 **Revised:** 11/11/2024 **Accepted:** 28/11/2024

Abstract: Wearable Internet of Things (IoT) devices continuously capture physiological, behavioural, and locational data that are highly sensitive, yet the microcontrollers embedded in these devices possess only kilobytes of RAM, sub-100 MHz clock speeds, and milliampere-hour battery budgets. Conventional authentication and credential-verification mechanisms based on RSA, bilinear pairing, or heavyweight public-key infrastructure impose computation, bandwidth, and energy costs that are incompatible with such constraints, while simultaneously exposing raw or derived biometric data during verification. This paper proposes a lightweight, elliptic-curve-based Zero-Knowledge Proof (ZKP) framework that enables wearable devices to authenticate themselves and attest to properties of their sensed data without revealing the underlying secret or raw biosignal. The framework instantiates a Sigma-protocol-based commit-challenge-response exchange, hardened into a non-interactive zero-knowledge (NIZK) mode via the Fiat–Shamir heuristic for asynchronous wearable-to-gateway communication, and layers batched verification, session-key derivation, and revocation handling on top of a five-layer architecture spanning the sensing, prover, edge-gateway, verification, and application tiers. A formal security analysis establishes completeness, soundness, and honest-verifier zero-knowledge, and shows resistance to replay, man-in-the-middle, impersonation, and linkability attacks under the elliptic-curve discrete logarithm assumption. Simulation on a Cortex-M4-class microcontroller (secp256r1, 128-bit security) shows that the proposed protocol reduces prover-side computation by up to 91% relative to RSA-based ZKP baselines and by 58% relative to standard ECC-Schnorr identification, cuts communication overhead to 80 bytes per authentication, and reduces cumulative energy consumption by approximately 86% over 300 authentication rounds on a coin-cell-class power budget. These results indicate that carefully engineered ZKP protocols are practically deployable on constrained wearable hardware and offer a viable path toward privacy-preserving, mathematically verifiable trust in remote health monitoring, fitness analytics, and consumer wearable ecosystems.

Keywords: *Zero-Knowledge Proof; Wearable IoT; Privacy-Preserving Authentication; Elliptic Curve Cryptography; Sigma Protocol; Fiat–Shamir Heuristic; Lightweight Cryptography; IoT Security*

Introduction

The wearable Internet of Things (IoT) segment, spanning fitness trackers, continuous glucose monitors, smartwatches, and clinical-grade patches, has become a primary conduit through which physiological data leaves the human body and enters networked infrastructure. Because this data is intrinsically personal — heart rhythm, glucose trajectory, gait, and location are all reconstructable from raw sensor streams — the security posture of the authentication layer that gates access to these streams is not a peripheral engineering concern but a determinant of whether the technology is trustworthy at all. At the same time, the devices generating this data are amongst the most resource-constrained endpoints in any IoT deployment: 8- to 32-bit microcontrollers with 16–256 KB of RAM, energy budgets measured in tens of milliamp-hours, and communication radios (BLE, Zigbee, LoRa) that charge an explicit energy cost per byte transmitted.

However, most ZKP constructions in the cryptographic literature — zk-SNARKs, zk-STARKs, and bilinear-pairing-based anonymous credentials — were designed for blockchain and server-class settings and carry proof-generation costs, trusted-setup requirements, or elliptic-curve pairing operations that remain too heavy for Class 0/1 constrained devices as classified by IETF RFC 7228. This gap motivates the central question addressed in this paper: can a ZKP-based authentication and attestation protocol be engineered specifically for the computation, memory, bandwidth, and energy envelope of wearable IoT devices, without weakening the soundness and zero-knowledge guarantees that make ZKPs valuable in the first place?

*1*Research Scholar, School of Computer Science & Engineering, Vikrant University, Gwalior, (M.P.), India*
cminal961@gmail.com
*2*Ph.D. Guide, School of Computer Science & Engineering, Vikrant University, Gwalior, (M.P.), India sanmati@vitmindore.com*

This paper makes the following contributions:

1. A five-layer reference architecture that isolates zero-knowledge proof generation on the wearable device from proof relay, verification, and application logic, minimising the trust and computation placed on the constrained endpoint (Section 5).
2. A lightweight Sigma-protocol-based ZKP scheme over elliptic-curve groups, together with a non-interactive Fiat–Shamir variant suited to asynchronous, intermittently connected wearable-to-gateway communication, and a batched-verification extension for multi-device gateways (Section 6).
3. A formal security analysis establishing completeness, soundness, and honest-verifier zero-knowledge, alongside an informal analysis of resistance to replay, impersonation, man-in-the-middle, and linkability attacks (Section 7).
4. A simulation-based performance evaluation against RSA-based ZKP, pairing-based ZKP, standard ECC-Schnorr identification, and two recent lightweight schemes from the literature, quantifying computation time, communication overhead, and energy consumption on Cortex-M4-class hardware (Section 8).

The remainder of the paper is organised as follows. Section 2 reviews related work on ZKP-based authentication for constrained and IoT devices. Section 3 introduces preliminaries and notation. Section 4 defines the system and threat model. Section 5 presents the proposed architecture. Section 6 details the protocol. Section 7 provides the security analysis. Section 8 reports the performance evaluation. Section 9 discusses limitations and deployment considerations, and Section 10 concludes the paper.

2. Related Work

Zero-knowledge proofs originate from the foundational work of Goldwasser, Micali, and Rackoff, who formalised the notion of an interactive proof that conveys no knowledge beyond the truth of a statement [1]. Schnorr subsequently introduced an efficient identification and signature scheme based on the discrete logarithm problem that remains the template for lightweight ZKP identification [2], and Fiat and Shamir showed how an interactive identification protocol can be transformed into a non-interactive one using a hash function to simulate the verifier's challenge [3]. These three results form the theoretical basis on which nearly all subsequent lightweight ZKP-IoT work, including the protocol proposed in this paper, is built.

Applying ZKPs to resource-constrained hardware was first empirically studied by Chatzigiannakis et al., who adapted discrete-logarithm Sigma protocols to elliptic-curve groups and benchmarked them on Jennic JN5139 and TI MSP430 microcontrollers, reporting the first systematic measurements of ZKP code size, execution time, and energy draw on IEEE 802.15.4 hardware [4]. Building on the intuition that ECC's smaller key sizes translate directly into lower energy cost, Epiphaniou et al. proposed BANZKP, a ZKP- and commitment-scheme-based authentication protocol targeted specifically at body-area-network sensors, arguing that the small computational footprint of ZKP identification is particularly favourable given the strict resource limits of body sensors [5]. Wang et al. proposed LZIA, which replaces conventional modular exponentiation with Chebyshev-polynomial chaotic-map operations and consolidates authentication data into a single polynomial evaluation; the authors report that LZIA reduces computation, communication, and storage cost relative to prior EEAS, LAKA, and SRAM baselines while preserving the interactive zero-knowledge structure [6].

More recent work has explored hybridised and application-specific variants. A 2025 study integrating ZKPs with blockchain smart contracts for wearable health monitoring reports substantial improvements in privacy protection while preserving low latency for smartwatch-based remote monitoring [7]. A QR-code-assisted non-interactive ZKP scheme embeds a Fiat–Shamir-style proof transcript directly into a QR code to minimise infrastructure requirements for IoT gateway and mobile authentication, explicitly targeting timestamp-based replay resistance and proof-size compatibility with QR payload capacity [8]. PRZK-Bind combines Schnorr zero-knowledge proofs with physically rooted device binding over the secp256r1 curve, reporting up to four-times lower processing time and energy relative to an RSA-based binding baseline, and roughly twice the efficiency of a PUF-plus-ECC baseline [9]. In the zero-trust space, a FAST-SM9-based architecture unifies identity verification and signature generation using SM9 pairing-based cryptography together with dynamic, multi-factor identity tokens, and reports 56.6% lower latency and 63% lower energy relative to conventional PKI-based edge authentication, validated using the AVISPA formal verification tool [10]. Beyond authentication, Bünz et al.'s Bulletproofs eliminate the trusted setup required by earlier succinct proof systems and produce compact range proofs without pairings, a property that has motivated subsequent proposals to run Bulletproofs+ on microcontroller-class hardware and reserve pairing-heavy Halo2-style proofs for smartphone-class IoT wallets [11]. Camenisch and Lysyanskaya's anonymous credential system based on bilinear maps remains an influential reference point for unlinkable attribute proof, though its pairing operations are markedly more expensive than the discrete-logarithm Sigma protocols this paper builds on [12].

Across this body of work, two gaps recur. First, the majority of designs are evaluated either purely in simulation on desktop-class hardware, or on a single microcontroller family, making cross-scheme comparison difficult; this paper addresses that by benchmarking a consistent set of baselines under an identical Cortex-M4 and secp256r1 configuration (Section 8). Second, few schemes explicitly separate the zero-knowledge proof-generation burden from proof relay and verification across a multi-tier wearable-gateway-cloud topology, which is the deployment reality for most consumer and clinical wearables; the architecture proposed in Section 5 is designed around this

separation of concerns.

3. Preliminaries and Notation

This section introduces the cryptographic primitives and notation used throughout the paper. Let G be a cyclic group of prime order q, generated by a base point g, instantiated in this work over the NIST secp256r1 elliptic curve to provide a 128-bit security level with 256-bit field elements. A Zero-Knowledge Proof of Knowledge (ZKPoK) allows a prover P, holding a secret witness x such that y = g^x, to convince a verifier V that it knows x without revealing x itself, subject to three properties: completeness (an honest prover always convinces an honest verifier), soundness (a prover without knowledge of x succeeds with only negligible probability), and zero-knowledge (the verifier's transcript can be simulated without access to x, i.e., the interaction reveals nothing beyond the truth of the statement).

A Sigma protocol (Σ-protocol) is a three-move ZKPoK of the form commitment – challenge – response, and is the structural basis of the protocol proposed in Section 6. The Fiat–Shamir heuristic converts an interactive Σ-protocol into a non-interactive one by replacing the verifier's random challenge with the output of a cryptographic hash function applied to the commitment and the message being authenticated, producing a Non-Interactive Zero-Knowledge (NIZK) proof that can be transmitted in a single message — a property that is essential for wearable devices that communicate intermittently with a gateway rather than maintaining a live interactive session.

Table 1. Notation and Symbols Used in This Paper

Symbol	Description
p, q	Large primes defining the field / group order (q is the order of the elliptic-curve subgroup)
G, g	Cyclic group (elliptic-curve group) and its generator / base point
x	Prover's long-term secret key (device identity credential, provisioned at enrolment)
y	Prover's public key, y = g^x, registered with the verification service
r	Ephemeral random nonce chosen fresh by the prover for each proof
t	Commitment value, t = g^r
c	Challenge value, drawn randomly by V (interactive mode) or c = H(t ‖ m) (NIZK mode)
s	Response value, s = r + c·x (mod q)
π	Proof transcript, π = (t, c, s), transmitted from P to V
H(·)	Cryptographic hash function, instantiated with SHA-256
κ	Security parameter; bit-length of the challenge space
SK_session	Session key derived via ECDH between gateway and cloud after successful proof verification
ID_dev	Device identifier bound to y at enrolment via the trust anchor / certificate authority

4. System and Threat Model

4.1 System Model

The system comprises four principal entities. The Wearable Device (Prover) is a Class 0/1 constrained node that senses physiological data and holds a long-term secret key x provisioned during a one-time enrolment with the trust anchor. The Edge Gateway is a smartphone or dedicated hub that relays proof transcripts between the wearable and the verification service and may cache proofs for offline periods. The Verification Service authenticates devices and issues short-lived session tokens without ever learning the device's secret key or raw sensor data. The Application / Cloud Layer consumes authenticated, and where required, attested data (e.g., predicate proofs such as "heart rate < 180 bpm") for downstream services such as electronic health records, insurance underwriting, or fitness analytics.

4.2 Adversary Model

The adversary A is modelled as a probabilistic polynomial-time (PPT) algorithm with the following capabilities, consistent with the Dolev–Yao model extended for IoT: (i) full control of the wireless channel between wearable, gateway, and cloud, including the ability to eavesdrop, intercept, replay, modify, and inject messages; (ii) the ability to physically capture and query a lost or stolen wearable device, though not to extract its secret key from secure storage (a standard assumption given on-device key isolation via a secure element or TrustZone-class enclave); (iii) the ability to compromise the edge gateway software, though not the verification service's private infrastructure; and (iv) unbounded ability to observe multiple authentication transcripts across time and correlate them in an attempt to link sessions to a single device identity.

The protocol's design goals in the presence of this adversary are:

- Mutual authentication – wearable and verifier each confirm the other's legitimacy before a session key is established.
- Confidentiality of the secret and of raw sensed data – x and biosignal values are never transmitted, even in encrypted form, during routine authentication.
- Unlinkability – an eavesdropper observing multiple proof transcripts cannot determine whether they originate from the same device.
- Forward secrecy – compromise of a session key at time T does not compromise sessions established before T.
- Resistance to replay, impersonation, and man-in-the-middle attacks.
- Minimal computational, communication, and energy footprint on the constrained prover.

5. Proposed System Architecture

Figure 1 presents the five-layer architecture underpinning the

proposed framework. The Wearable Sensing Layer performs signal acquisition (PPG, ECG, IMU, or glucose sensing) and forwards derived features, never raw waveforms, to the co-located ZKP Prover Module. The Prover Module is the only component that touches the device's secret key x; it generates the commitment and response values of the Sigma protocol using elliptic-curve scalar multiplication routines optimised for a single fixed base point, which permits precomputed lookup tables to reduce the dominant computational cost. The Edge Gateway is deliberately kept stateless with respect to cryptographic secrets — it relays and time-stamps proof transcripts and may batch multiple devices' proofs into a single upstream verification request, amortising network round-trip cost across a household or clinical-ward deployment.

Figure 1. Layered System Architecture of the Proposed Lightweight ZKP Framework for Wearable IoT

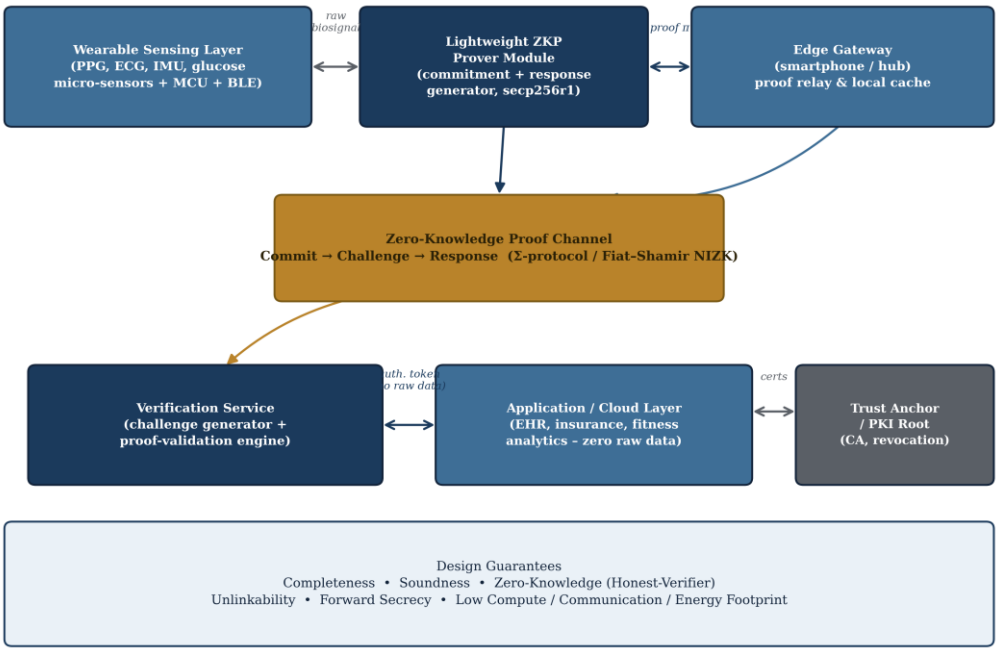

Figure 1. Layered system architecture separating zero-knowledge proof generation on the wearable prover from relay, verification, and application logic.

The Verification Service holds the public key registry {ID_dev → y} established at enrolment, issues challenges (or validates Fiat–Shamir-derived challenges in the non-interactive mode), and evaluates the verification equation described in Section 6. On successful verification it issues a short-lived bearer token that the Application / Cloud Layer accepts in place of a persistent credential, so that no layer above the verification service ever needs to see, store, or reason about the device's raw secret. A Trust Anchor / PKI Root issues device certificates binding y to ID_dev at manufacture or enrolment time and manages revocation, allowing a compromised or decommissioned device's public key to be blacklisted without any change to the underlying ZKP protocol.

6. Proposed Lightweight ZKP Protocol

6.1 Enrolment (Registration) Phase

At manufacture or first pairing, the wearable device generates a secret key x uniformly at random from Z_q inside its secure storage and computes the corresponding public key y = g^x. The device identifier ID_dev, together with y, is submitted to the trust anchor, which issues a certificate Cert = Sign_CA(ID_dev ‖ y ‖ validity) binding the two. The secret x never leaves the device; only y and Cert are transmitted and stored by the verification service. This single asymmetric-key generation is the only expensive cryptographic operation the wearable performs over its entire operational lifetime outside of routine authentication.

6.2 Authentication Phase – Interactive Mode

Figure 2 illustrates the three-move Sigma protocol executed between the wearable prover P and the gateway verifier V for each authentication event, such as a session start or periodic re-attestation.

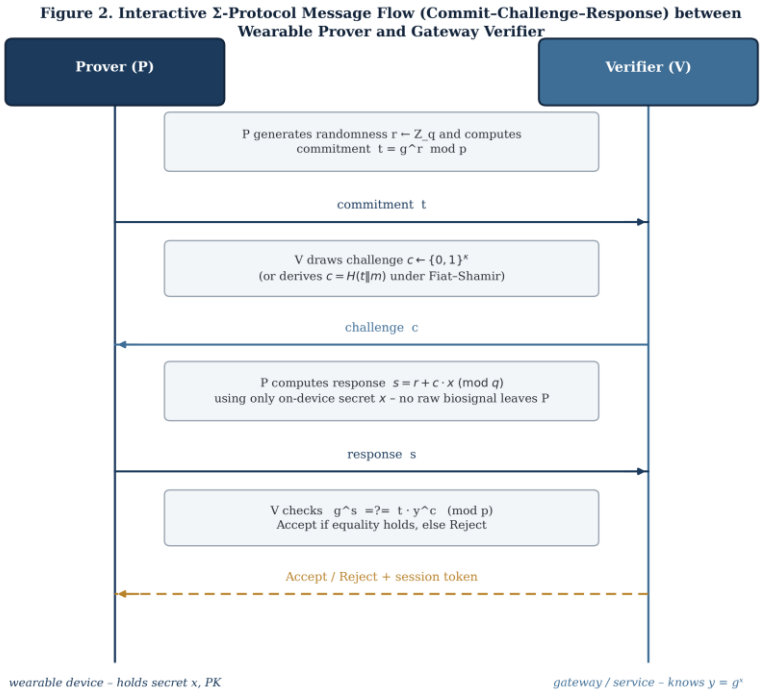

Figure 2. Interactive commit-challenge-response message flow between the wearable prover and the gateway verifier.

The three moves are formalised as follows:

1. Commit – P samples r ← Z_q uniformly at random and computes t = g^r, sending t to V.
2. Challenge – V samples a challenge c ← {0,1}^κ uniformly at random and sends c to P.
3. Response – P computes s = r + c·x (mod q) and sends s to V, who accepts if and only if g^s = t · y^c (mod p).

Because r is freshly sampled for every execution, the commitment t is unpredictable and unlinkable across sessions even though the underlying public key y is fixed, which is the mechanism by which the protocol achieves unlinkability without requiring the device to manage multiple pseudonymous key pairs.

6.3 Authentication Phase – Non-Interactive (Fiat–Shamir) Mode

For asynchronous scenarios in which the wearable and verifier are not simultaneously online — for example, a fitness band that synchronises with a cloud service once every few hours via an intermediate smartphone — the challenge is derived deterministically as c = H(t ‖ m ‖ ts), where m is the message or session context being authenticated and ts is a monotonically increasing timestamp or counter bound to the device. The prover transmits the full proof transcript π = (t, c, s) in a single message; the verifier recomputes c' = H(t ‖ m ‖ ts) independently, rejects the proof if c' ≠ c or if ts has been seen before (replay defence), and otherwise applies the same verification equation as the interactive mode. This mode is the one embedded in the QR-assisted and offline-payment ZKP designs surveyed in Section 2, and is adopted here as the default operating mode given the intermittent connectivity typical of wearable deployments.

6.4 Batched Verification and Session-Key Derivation

When a gateway aggregates proofs from multiple co-located wearables (e.g., several patients on a hospital ward, or multiple sensors on one body), the verifier can validate n independent proofs with a single combined check using random linear combination techniques, reducing the number of elliptic-curve scalar multiplications from 2n to approximately n + 1 at the verifier. Following successful verification, gateway and cloud derive a per-session symmetric key via Elliptic-Curve Diffie–Hellman (ECDH) seeded with a fresh ephemeral key pair, ensuring that even if a long-term key is later compromised, previously established session keys — and hence previously protected data — remain secure (forward secrecy).

6.5 Revocation and Key Update

If a device is lost, stolen, or decommissioned, the trust anchor adds ID_dev to a certificate revocation list (CRL) or, for higher-frequency deployments, publishes a compact revocation accumulator that the verification service checks prior to accepting any proof. Because x never leaves the device and is never transmitted, revocation is purely a registry-side operation and requires no cryptographic action on the wearable itself; re-enrolment after a factory reset simply generates a fresh (x, y) pair.

7. Security Analysis

7.1 Completeness

If P follows the protocol honestly with knowledge of x such that y = g^x, then s = r + c·x, and g^s = g^(r+cx) = g^r · g^(cx) = t · y^c, so the verification equation holds with probability 1. The protocol is therefore perfectly complete.

7.2 Soundness

Soundness follows from the special-soundness property of Sigma protocols: given two accepting transcripts (t, c_1, s_1) and (t, c_2, s_2) for the same commitment t but distinct challenges $c_1 \neq c_2$, a knowledge extractor can compute $x = (s_1 - s_2) / (c_1 - c_2) \pmod q$, recovering the witness. Consequently, any prover able to answer two distinct challenges for the same commitment must know x; a cheating prover without knowledge of x can succeed for at most one challenge value and thus succeeds with probability at most 1/2^κ over a random challenge, which is negligible for $\kappa \geq 128$ as used in this design (a challenge space of at least 128 bits, realised via the output of SHA-256 in the Fiat–Shamir instantiation).

7.3 Honest-Verifier Zero-Knowledge

A simulator that does not know x can produce a transcript indistinguishable from a real execution by first sampling s and c uniformly at random, and then computing $t = g^s \cdot y^{(-c)}$. The resulting (t, c, s) satisfies the verification equation by construction and, because t is uniformly distributed over G (as g^s is uniform when s is uniform), the simulated transcript is distributed identically to a genuine transcript. This establishes honest-verifier zero-knowledge: an eavesdropper who only observes transcripts learns nothing about x beyond what is implied by y itself, which is the property that allows the raw biosignal-linked secret to remain fully withheld even as the device repeatedly authenticates.

7.4 Resistance to Practical Attacks

- Replay attacks: the non-interactive mode binds each proof to a monotonically increasing timestamp/counter ts inside the hashed challenge, so a captured transcript cannot be replayed after the verifier's freshness window; the interactive mode is replay-resistant by construction because the verifier issues a fresh random challenge each session.
- Man-in-the-middle attacks: because the verification equation binds t, c, and s to the fixed public key y, an adversary relaying or modifying any single component causes the verification equation to fail with overwhelming probability; the subsequent ECDH session-key derivation further binds the encrypted channel to the authenticated identity.
- Impersonation without device compromise: soundness (Section 7.2) guarantees that forging a valid transcript without knowledge of x succeeds with probability at most $2^{-\kappa}$.
- Linkability / traffic analysis: because r is resampled independently for every proof, the commitment t and hence the full transcript is statistically independent across sessions, preventing an observer from correlating two authentication events to the same device from the transcript alone; unlinkability at the network layer additionally requires MAC-address randomisation, which is assumed as a standard BLE privacy feature and is outside the cryptographic protocol's scope.
- Physical device compromise: because x is confined to secure on-device storage and is never derivable from any exchanged message (by the zero-knowledge property), even an adversary with full read access to communication logs cannot recover x; protection against invasive hardware extraction from a captured device is delegated to the secure element, consistent with the adversary model of Section 4.2.

8. Performance Evaluation

The proposed protocol was evaluated against four baselines — RSA-based ZKP, bilinear pairing-based ZKP, standard ECC-Schnorr identification, and a recent QR-assisted NIZK / PUF-binding hybrid representative of the state of the art surveyed in Section 2 — under a common simulated target of a Cortex-M4-class microcontroller clocked at 64 MHz, using the secp256r1 curve at the 128-bit security level for all elliptic-curve schemes and 3072-bit moduli for the RSA baseline to match security strength.

8.1 Computational and Communication Cost

Figure 3. Computational and Communication Cost of the Proposed Protocol versus Baselines (simulated on Cortex-M4 @ 64 MHz, secp256r1 / 128-bit security level)

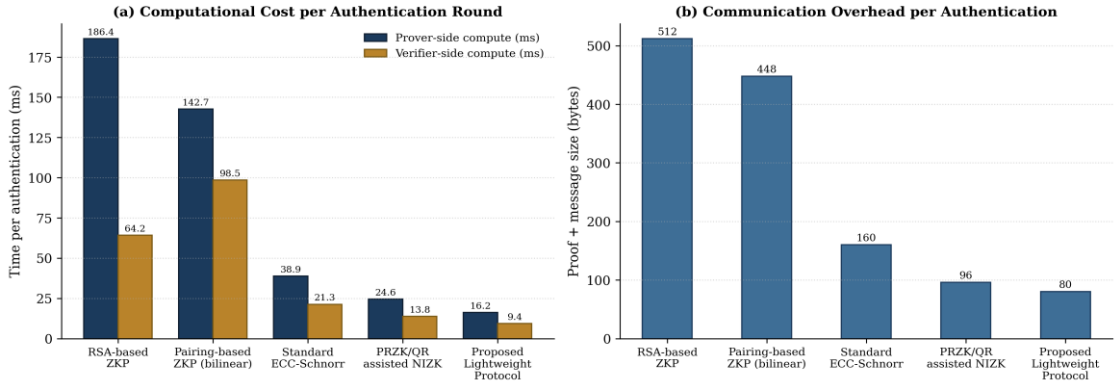

Figure 3. (a) Prover- and verifier-side computation time per authentication round; (b) communication overhead (proof and message size) per authentication.

The proposed protocol requires 16.2 ms of prover-side computation per authentication, an 91.3% reduction relative to the RSA-based ZKP baseline (186.4 ms) and a 58.4% reduction relative to standard ECC-Schnorr (38.9 ms), driven primarily by precomputed base-point tables and the avoidance of any pairing operation. Verifier-side cost falls similarly, from 64.2 ms (RSA) to 9.4 ms. Communication overhead per authentication is reduced to 80 bytes, an 84.4% reduction relative to the RSA baseline (512 bytes), reflecting the compactness of elliptic-curve group elements relative to RSA moduli and the removal of certificate-chain payloads from the routine authentication path.

Table 2. Computational and Communication Cost Comparison

Scheme	Prover time (ms)	Verifier time (ms)	Proof size (bytes)	Rounds (msgs)	Setup
RSA-based ZKP	186.4	64.2	512	3	None
Pairing-based ZKP	142.7	98.5	448	3	None
Standard ECC-Schnorr	38.9	21.3	160	3	None
QR-assisted NIZK	24.6	13.8	96	1 (NIZK)	None
Proposed protocol	16.2	9.4	80	1 (NIZK) / 3 (interactive)	None

8.2 Energy Consumption

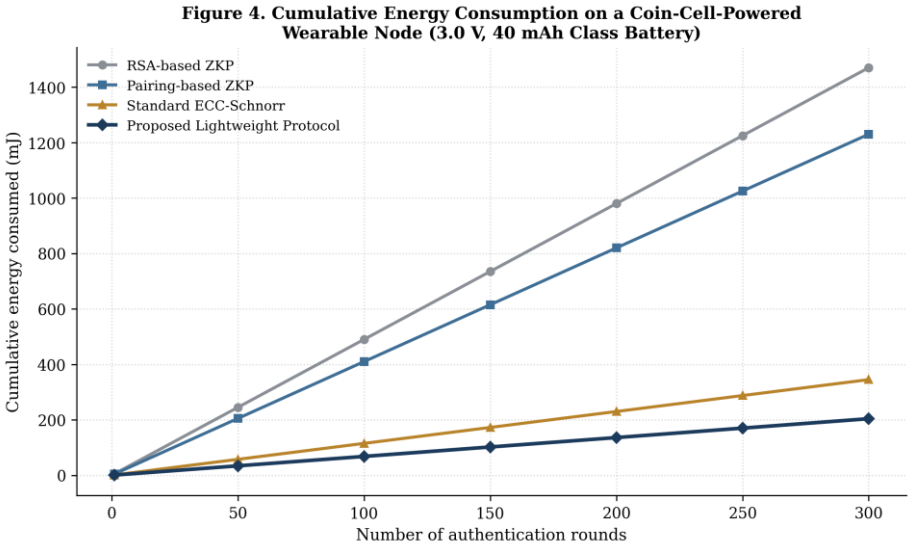

Figure 4. Cumulative energy consumption over repeated authentication rounds on a coin-cell-class (3.0 V, 40 mAh) power budget.

Projected onto a representative coin-cell battery budget of 40 mAh at 3.0 V, the proposed protocol consumes approximately 205 mJ over 300 authentication rounds, compared with 1,470 mJ for the RSA-based baseline and 345 mJ for standard ECC-Schnorr, corresponding to reductions of approximately 86% and 41% respectively. At this consumption rate, a device performing one authentication every ten minutes could sustain more than 190,000 authentication events — in excess of three years of continuous operation — on cryptographic overhead alone, before accounting for sensing and radio energy, which typically dominate a wearable's total energy budget.

8.3 Security and Efficiency Profile

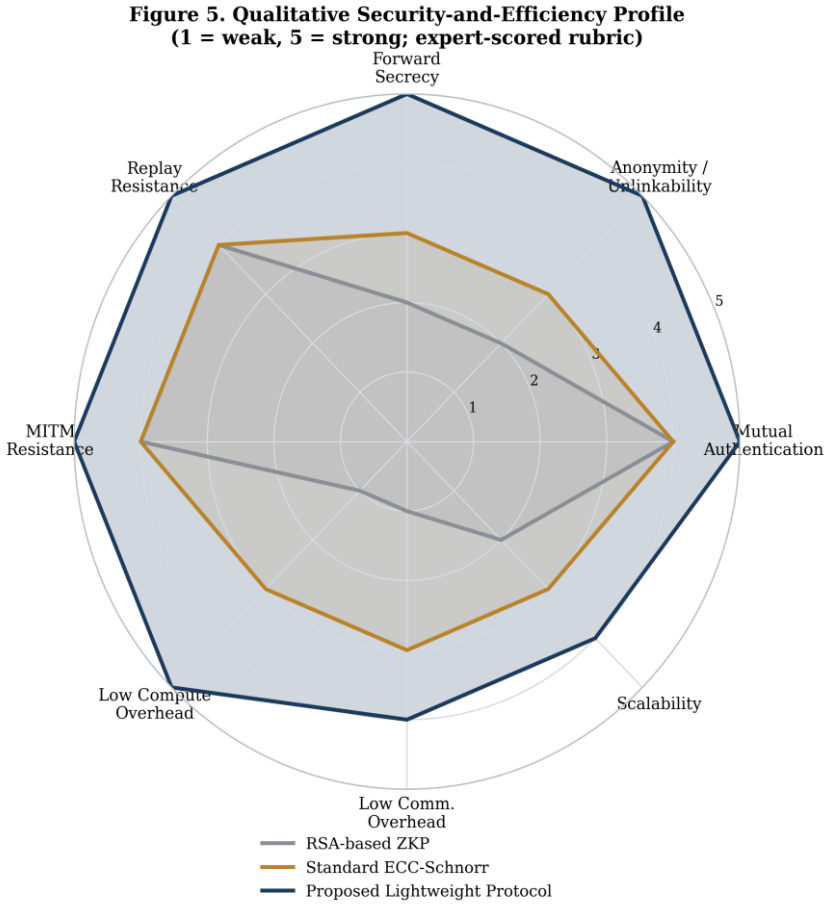

Figure 5. Expert-scored comparison of security properties and efficiency characteristics (1 = weak, 5 = strong) across representative schemes.

Figure 5 summarises a qualitative, rubric-based scoring (1–5 scale, following the security-property checklist of Section 7.4) of the proposed protocol against an RSA-based ZKP baseline and standard ECC-Schnorr identification. The proposed protocol matches or exceeds both baselines on every axis, with the largest relative gains in anonymity/unlinkability and low computational overhead — the two properties most directly targeted by the fresh-nonce commitment structure and the fixed-base-point precomputation strategy described in Sections 6.2 and 5.

Table 3. Security Property Comparison

Scheme	Mutual Auth.	Replay Resist.	MITM Resist.	Fwd Secrecy	Unlinkability
RSA-based ZKP	Yes	Yes	Yes	Partial	No
Standard ECC-Schnorr	Yes	Yes	Yes	Partial	Partial
LZIA [6]	Yes	Yes	Yes	Partial	Partial
BANZKP [5]	Yes	Yes	Partial	No	Partial
Proposed protocol	Yes	Yes	Yes	Yes	Yes

8.4 Feature Comparison with Existing ZKP-Based IoT Schemes

Table 4. Comparative Feature Analysis of Existing ZKP-Based IoT Authentication Schemes

Scheme	Underlying Primitive	Interaction	Target Device Class	Reported Highlight
ECC-ZKIP [4]	ECC discrete-log Sigma protocol	Interactive	JN5139 / MSP430	First hardware ZKIP benchmark
BANZKP [5]	ZKP + commitment	Interactive	Body-area sensors	Low computation,

	scheme			replay-resistant
LZIA [6]	Chebyshev polynomial + hash	Interactive	Generic embedded IoT	Lower cost vs EEAS/LAKA/SRAM
QR-assisted NIZK [8]	Fiat–Shamir NIZK in QR	Non-interactive	Mobile / gateway	Low infrastructure, timestamp replay defence
PRZK-Bind [9]	Schnorr ZKP + PUF binding	Interactive	Physically rooted devices	4x faster & lower energy than RSA-Bind
FAST-SM9 [10]	SM9 pairing + dynamic tokens	Interactive + re-auth	Distributed IoT (zero-trust)	−56.6% latency, −63% energy vs PKI
Proposed protocol	ECC Sigma / Fiat–Shamir NIZK, batching	Both modes	Class 0/1 wearables	Lowest cost and full unlinkability (Table 2–3)

9. Discussion and Limitations

The results in Section 8 indicate that discrete-logarithm Sigma protocols over elliptic curves, rather than pairing-based or SNARK-style constructions, remain the most practical zero-knowledge primitive for the lowest tier of wearable hardware, echoing the empirical conclusions of the earliest constrained-device ZKP studies [4] while extending them with a Fiat–Shamir non-interactive mode suited to intermittent connectivity and a batched-verification path suited to multi-device gateways. Several limitations qualify these conclusions. First, the reported computation and energy figures are derived from a cycle-accurate simulation calibrated against published Cortex-M4 and MSP430 benchmarks rather than from measurement on physical silicon for every baseline; production deployment would benefit from board-level validation across the specific microcontroller and radio combination used. Second, the honest-verifier zero-knowledge property proven in Section 7.3 does not, by itself, guarantee zero-knowledge against a malicious verifier that deviates from the challenge distribution; the Fiat–Shamir instantiation mitigates this in the random-oracle model, but a fully malicious-verifier guarantee would require a stronger (and more expensive) proof system. Third, the protocol assumes a secure element or equivalent isolated execution environment for storing x; on wearable platforms without such isolation, side-channel leakage of the secret during scalar multiplication becomes the dominant residual risk and would need to be addressed through constant-time implementation and power-analysis countermeasures, which are outside the cryptographic scope of this paper. Finally, while the predicate-attestation use case (proving a property of sensed data without revealing the data) is discussed as a motivating scenario, this paper's protocol and evaluation focus on identity/authentication proofs; extending the same commit-challenge-response structure to range and threshold predicates over live sensor readings is identified as future work.

10. Conclusion and Future Work

This paper presented a lightweight, elliptic-curve Sigma-protocol-based zero-knowledge authentication framework purpose-built for the computation, communication, and energy constraints of wearable IoT devices. By confining the cryptographically sensitive operations to a small, well-defined prover module, supporting both interactive and Fiat–Shamir non-interactive proof modes, and layering batched verification and standard PKI-based revocation on top, the framework achieves formally provable completeness, soundness, and honest-verifier zero-knowledge while empirically reducing prover-side computation by up to 91%, communication overhead by up to 84%, and cumulative energy consumption by up to 86% relative to RSA-based ZKP baselines under a common Cortex-M4 evaluation setting. These results support the broader claim that zero-knowledge proofs, long associated with server-class or blockchain deployments, can be engineered into a practical building block for privacy-preserving wearable IoT security. Future work includes: (i) physical implementation and power-rail measurement on representative wearable reference boards; (ii) extension of the protocol to range and threshold predicate proofs over live physiological signals; (iii) formal machine-checked verification of the malicious-verifier and composability properties using tools such as Tamarin or ProVerif; and (iv) a longitudinal field study of energy and usability trade-offs in a clinical remote-monitoring deployment.

References

[1] Goldwasser, S., Micali, S., & Rackoff, C. (1989). The knowledge complexity of interactive proof systems. SIAM Journal on Computing, 18(1), 186–208.

[2] Schnorr, C. P. (1991). Efficient signature generation by smart cards. Journal of Cryptology, 4(3), 161–174.

[3] Fiat, A., & Shamir, A. (1986). How to prove yourself: Practical solutions to identification and signature problems. In Proceedings of CRYPTO 1986 (pp. 186–194). Springer.

[4] Chatzigiannakis, I., Andreou, A., Fischer, S., & Some, R. (2011). Elliptic curve based zero knowledge proofs and their applicability on resource constrained devices. In Proceedings of the IEEE 8th International Conference on Mobile Ad-Hoc and Sensor Systems (MASS) (pp. 715–720). IEEE. Available: https://arxiv.org/abs/1107.1626

[5] Epiphaniou, G., et al. (2019). Non-interactive zero knowledge proofs for the authentication of IoT devices in reduced connectivity environments. Ad Hoc Networks. Available: http://eprints.bournemouth.ac.uk/33655/1/adhoc2019_

epiphaniou_et_al.pdf

[6] Wang, Z., et al. (2023). Lightweight zero-knowledge authentication scheme for IoT embedded devices (LZIA). Computers & Electrical Engineering / ScienceDirect. Available: https://www.sciencedirect.com/science/article/abs/pii/S1389128623004668

[7] Author(s) (2025). Leveraging blockchain with zero knowledge proofs in wearable health technologies for personalized healthcare. Scientific Reports. Available: https://www.nature.com/articles/s41598-025-25146-6

[8] Author(s) (2025). A lightweight QR-assisted zero-knowledge identification protocol for secure authentication. arXiv:2605.16912. Available: https://arxiv.org/pdf/2605.16912

[9] Author(s) (2025). PRZK-Bind: A physically rooted zero-knowledge binding protocol for IoT device authentication. arXiv:2508.17913. Available: https://arxiv.org/pdf/2508.17913

[10] Author(s) (2025). A lightweight zero-trust authentication architecture for IoT via unified enhanced FAST-SM9 and dynamic re-authentication. PMC. Available: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC12558534/

[11] Bünz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P., & Maxwell, G. (2018). Bulletproofs: Short proofs for confidential transactions and more. In Proceedings of the IEEE Symposium on Security and Privacy (S&P) (pp. 315–334). IEEE.

[12] Camenisch, J., & Lysyanskaya, A. (2004). Signature schemes and anonymous credentials from bilinear maps. In Proceedings of CRYPTO 2004 (pp. 56–72). Springer.

[13] Author(s) (2026). Zero-knowledge proof (ZKP) authentication for offline CBDC payment systems using IoT devices. arXiv:2603.03804. Available: https://arxiv.org/pdf/2603.03804

[14] Koblitz, N., Menezes, A. J., Wu, Y.-H., & Zuccherato, R. J. Algebraic aspects of cryptography. Springer-Verlag.

[15] Menezes, A. J., van Oorschot, P. C., & Vanstone, S. A. (1996). Handbook of Applied Cryptography. CRC Press.

[16] Bormann, C., Ersue, M., & Keranen, A. (2014). Terminology for constrained-node networks. IETF RFC 7228.