

AI Enhanced Zero Knowledge Authentication for High Mobility IoT Using Predictive Token Learning

Shafiq Ahmed, *Graduate Student Member, IEEE*, Mohammad Hossein Anisi, *Senior Member, IEEE*

Abstract—High-mobility Internet of Things (IoT) for Vehicle-to-Grid (V2G) Demand Response (DR), including roaming between Charge Point Operators (CPOs), requires privacy-preserving authentication with sub-1 ms responses and cross-domain scalability as devices exceed 200 km/h. Mechanisms must run on constrained hardware while remaining compatible with EV-charging message flows such as ISO 15118–20 and OCPP 2.0.1. Many deployed schemes re-authenticate from scratch, which inflates computation and airtime; static credentials also ignore trajectory context and struggle with rapid mobility. We present a Zero-Knowledge Proof-based Authentication Scheme (ZKPAS) for V2G/DR that proves possession without disclosure and replaces heavy handshakes with compact, mobility-aware proofs, targeting latency $L < 1$ ms. ZKPAS integrates four elements. (i) Schnorr-style zero-knowledge proofs over elliptic curves validate identity without exposing trajectory or long-term identifiers. (ii) Sliding-window tokens bind to DR windows and reduce per-segment communication from $\mathcal{O}(n)$ to $\mathcal{O}(\log n)$. (iii) Predictive token generation with Long Short-Term Memory (LSTM) models trained on GeoLife and T-Drive pre-computes material, yielding 84.7% token reuse along trajectories. (iv) Cross-domain authentication employs (t, n) -threshold cryptography for Byzantine-tolerant roaming across operators. We prove resistance to impersonation, replay, man-in-the-middle, and trajectory inference; under the Computational Diffie–Hellman Problem (CDHP), the adversary’s success probability satisfies $\Pr[\text{break}] \leq 2^{-\lambda}$. On real transportation topologies, ZKPAS cuts computation by 71.8%, authentication latency by 93.9%, and energy by 69.5%, while interfacing with V2G/DR control flows. The protocol sustains a 98.5% authentication success rate at 250 km/h.

Index Terms—Zero-knowledge proofs, high-mobility authentication, IoT security, predictive token generation, sliding window mechanisms, cross-domain authentication, threshold cryptography, LSTM-based mobility prediction

I. INTRODUCTION

HIGH-mobility Internet of Things (IoT) ecosystems that include autonomous vehicles, unmanned aerial vehicle swarms, and moving sensor grids stress authentication because security decisions must track rapid motion and intermittent links [1]. In Vehicle to Grid (V2G) and Demand Response (DR) workflows, authentication must preserve privacy, fit an end to end latency budget of $L_{\text{auth}} = L_{\text{comp}} + L_{\text{comm}} + L_{\text{sync}} \leq 1$ ms, and remain reliable when devices move at $v > 250$ km/h across administrative domains and CPOs. Implementations must also execute on constrained on-board units without weakening guarantees. Static or semi dynamic schemes struggle with tight timing and heterogeneous trust during handovers and roaming [2].

The requirements concentrate into four points: (i) ultra low latency so authentication remains valid across fast domain transitions, (ii) lightweight cryptography sized for limited CPU, memory, and energy, (iii) trajectory privacy that limits inference on mobility traces, and (iv) seamless cross domain operation with distributed trust anchors that tolerate partial failures. Prior approaches meet these only partially. Blockchain driven consensus often increases delay and airtime under mobility [3]; classical Public Key Infrastructure (PKI) creates operational overhead at scale [4]; and many lightweight protocols appear to trade integrity or privacy for speed [5]. In V2G and DR, these trade offs tend to surface at the worst moment, for example when an Electric Vehicle (EV) roams between CPOs inside a short DR window.

We propose a Zero Knowledge Proof based Authentication Scheme (ZKPAS) for real time V2G and DR that couples predictive scheduling with provable privacy. ZKPAS has four elements. First, a Schnorr style zero knowledge core provides unlinkable entity authentication without revealing identities or trajectories; in the Dolev Yao model the adversarial success probability satisfies $\text{Adv} \leq 2^{-\lambda}$ under Elliptic Curve Discrete Logarithm Problem (ECDLP) hardness and Random Oracle Model (ROM) assumptions. Second, a sliding window token mechanism links handover evidence so roaming cost scales from $\mathcal{O}(n)$ to $\mathcal{O}(\log n)$, which avoids repeated full runs. Third, a Long Short Term Memory (LSTM) predictor trained on T Drive and GeoLife traces [6], [7] outputs a schedule $\hat{\tau}_{t+1:t+k}$ of next use tokens with 84.7% accuracy, which likely reduces stalls during inter domain motion. Fourth, a (t, n) threshold trust anchor fabric enables cross domain acceptance even if some anchors are faulty or unreachable.

To capture DR timing, we partition time into tariff bound windows $W_j = [t_j, t_{j+1})$ and bind each token to a window and location tuple. We model predicted motion as $\hat{T}_i(t + \Delta t) = \mathcal{M}_\theta(\mathcal{H}_i^{(t)})$ where $\mathcal{H}_i^{(t)}$ is device history and $\mathcal{M}_\theta$ is a trained recurrent model; a binding map $b : W_j \mapsto \text{token}_j$ associates DR windows with pre staged proofs. This prediction aware scheduling appears to lower token mismatches and to raise authentication availability during handovers, while remaining compatible with ISO 15118–20 and OCPP 2.0.1 message flows.

Our contributions are threefold. We give a unified problem formulation for high mobility V2G and DR authentication that integrates predictive intelligence, cross domain threshold trust, and sliding window continuity. We present a protocol with proofs of impersonation resistance, replay protection, trajectory privacy, and forward security, with guarantees derived from ECDLP hardness and ROM programmability. We

S. Ahmed and M. H. Anisi are with the School of Computer Science and Electronic Engineering, University of Essex, Colchester CO4 3SQ, UK.
E-mail: s.ahmed@essex.ac.uk; m.anisi@essex.ac.uk

provide an empirical study on real transportation topologies that reports computation reduced by 71.8%, authentication latency improved by 93.9%, and energy lowered by 69.5%, while sustaining a 98.5% success rate at 250 km/h.

The remainder is organized as follows. Section II surveys related work on mobile authentication. Section III defines the system and threat models and DR roles. Section IV details ZKPAS and its prediction pipeline with window binding. Section V gives formal security analysis. Section VI reports experiments, and Section VII closes with future directions.

II. RELATED WORK

Recent work on high-mobility IoT authentication spans cryptographic protocols, sliding-window continuity, and cross-domain trust. *However, proactive authentication and tight integration with trajectory prediction remain underdeveloped.*

A. Mobile Authentication and Zero-Knowledge Frameworks

Research on mobile authentication adapts classic primitives to heterogeneous settings. Zhang et al. [8] introduced SMAKA with batch verification, yet it still forces full re-authentication on inter-domain handovers, which is impractical at speeds > 150 km/h. Li et al. [9] achieved unlinkability and collusion resistance but retained $\mathcal{O}(n)$ time, limiting scalability. Chen et al.'s XAuth [10] uses proxy re-encryption, shifting cost into significant key-management overhead.

Zero-knowledge (ZK) methods authenticate without disclosure. Foundational results by Goldwasser et al. [11] led to short, randomizable signatures with constant-size proofs [12]. Modern Schnorr-style instantiations [13], [14] deliver ≈ 64 -byte proofs over elliptic curves with competitive runtime. Yet, current systems still lack adaptive scaling and prediction-aware operation needed for extreme mobility.

B. Sliding Window Mechanisms and Cross-Domain Authentication

Sliding-window schemes maintain session continuity by chaining tokens with forward security. EV-Auth [15] reduces compute via token chaining but omits full formal analysis and incurs $\mathcal{O}(n)$ storage. Delegation approaches at the edge [16] improve locality but presume stable edge resources, which is fragile across heterogeneous domains.

Decentralized trust via blockchain [17] provides consensus but adds latency and energy cost that violate sub-millisecond targets. Cross-domain designs [18] reach perfect forward secrecy and strong anonymity, although their computational profile remains heavy for ultra-low-latency execution.

C. Comparative Analysis and Research Gaps

Table I contrasts leading protocols on privacy, forward security, efficiency, cross-domain support, and proactiveness.

Three gaps persist. First, most systems do not *proactively* prepare authentication before handover, causing stalls during transitions. Second, despite rich GPS corpora (T-Drive [6], [19], [20] and GeoLife [7], [21], [22]), few designs leverage machine learning to forecast mobility and pre-stage

TABLE I
COMPARATIVE ANALYSIS OF MOBILE IOT AUTHENTICATION PROTOCOLS

Protocol	ZKP	Fwd Sec.	X-Domain	Comp.	Scal.	Formal
[5]	✗	✗	✓	$\mathcal{O}(n)$	Low	✓
[8]	✗	✗	✗	$\mathcal{O}(n)$	High	✓
[24]	✗	✓	✗	$\mathcal{O}(n)$	Low	✗
[15]	✗	✓	✓	$\mathcal{O}(n)$	Med	✗
[17]	✗	✓	✓	$\mathcal{O}(n^2)$	Low	✓
[10]	✗	✓	✓	$\mathcal{O}(n \log n)$	Med	✓
[9]	✗	✗	✗	$\mathcal{O}(n)$	High	✓
[18]	✗	✓	✓	$\mathcal{O}(n \log n)$	Med	✓
[25]	✗	✗	✗	$\mathcal{O}(k \cdot n)$	Med	✗
[26]	✗	✓	✓	$\mathcal{O}(n^2)$	Med	✓
[27]	✓	✓	✓	$\mathcal{O}(n \log n)$	Med	✓
ZKPAS	✓	✓	✓	$\mathcal{O}(\log n)$	High	✓

credentials. Third, sliding-window models often lack rigorous forward-security proofs and storage bounds [23].

ZKPAS addresses these deficits with LSTM-based mobility forecasting,

$$\hat{T}_i(t + \Delta t) = \mathcal{M}_\theta(\mathcal{H}_i^{(t)}),$$

where $\mathcal{H}_i^{(t)}$ is the device history and $\mathcal{M}_\theta$ a trained recurrent model. Coupled with cryptographically linked token sequences and ZK proofs, this design offers formal security guarantees while cutting latency and communication during high-speed, cross-domain operation.

III. PRELIMINARIES, SYSTEM MODEL AND SECURITY ASSUMPTIONS

We outline the mathematical setting, the V2G context for DR, and the adversary under which the Zero Knowledge Proof based Authentication Scheme (ZKPAS) operates. The design ties predictive mobility learning with elliptic curve cryptography and a threshold trust fabric so authentication stays private and timely under tight compute and energy limits.

A. Network Architecture and Entity Definitions

We consider roaming across CPOs during DR events; entities talk over short range radio and backhaul and expose only pseudonyms.

Definition (Mobile IoT Entity: EV OBU). An On Board Unit $\mathcal{M}_i$ follows a trajectory $\mathcal{T}_i(t) : \mathbb{R} \rightarrow \mathbb{R}^3$, operates within compute envelope $\mathcal{C}_i \in [\mathcal{C}_{\min}, \mathcal{C}_{\max}]$, and stores credentials $\mathcal{A}_i = \{ID_i, PK_i, Cert_i\}$ with $PK_i = [sk_i]G$ for $sk_i \in \mathbb{Z}_q$ on elliptic curve group $\mathbb{G}_{EC}$.

Definition (Authentication Domain: CPO). Domain $\mathcal{D}_j$ covers region $\mathcal{R}_j \subset \mathbb{R}^3$, policy $\mathcal{P}_j$, and suite $\mathcal{K}_j = \{pk_j, sk_j, \mathcal{G}_j, p_j, q_j\}$ where $\mathcal{G}_j$ is cyclic of order q_j over $\mathbb{F}_{p_j}$; Charge Points verify tokens and, per policy, forward opaque proofs.

Definition (Trust Anchor Network). A federation $\mathcal{TA} = \{\mathcal{TA}_1, \dots, \mathcal{TA}_n\}$ runs a (t, n) threshold; at least t anchors must authorize and the system tolerates up to $t-1$ Byzantine anchors.

DR windows and roaming. The scheduler publishes windows $W_j = [t_j, t_{j+1})$ with span $\tau_j^{DR} = t_{j+1} - t_j$; an Electric Vehicle (EV) may cross several $\mathcal{D}_j$ while W_j is active. ZKPAS binds pre staged tokens to $(W_j, \text{location})$ tuples so handovers remain valid.

Mobility prediction: We model $\mathcal{T}_i(t) = \hat{\mathcal{T}}_i(t) + \mathcal{E}_i(t)$ with $\hat{\mathcal{T}}_i(t + \Delta t) = \mathcal{M}_\theta(\mathcal{H}_i^{(t)})$ and residual $\mathcal{E}_i(t) \sim \mathcal{N}(0, \sigma_i^2)$; Long Short Term Memory (LSTM) training uses T Drive and GeoLife traces [6], [7]. The predictor selects the next token window that fits W_j and expected handovers.

B. Demand Response Binding and Latency Budget

We allocate an explicit budget inside each window. Let $L_{\text{auth}} = L_{\text{comp}} + L_{\text{comm}} + L_{\text{sync}} \leq 1 \text{ ms}$ and $p_{\text{succ}}(j) = \Pr[L_{\text{auth}} \leq \tau_j^{\text{DR}}]$; sliding window tokens remove repeated full handshakes, which tends to cut L_{comm} and its variance under roaming.

C. Inter session Linkability and Sliding Windows

Each execution is a session. To limit passive linking across sessions, the client samples $s \leftarrow \{0, 1\}^\lambda$, derives a per session key $k_s = \text{KDF}(\text{handshake transcript}, s)$, and emits tokens $t_i = F_{k_s}(i \parallel s)$; the verifier releases or accepts $W_j = \{t_i\}_{i=j}^{j+w-1}$. We capture linkability via game Link_A with advantage $\text{Adv}_A^{\text{link}} = |\Pr[b' = b] - 1/2|$, which we target to be negligible in λ . Using fresh (k_s, s) per session and binding s to the authenticated transcript makes windows from distinct sessions independent; for two independent windows the expected overlap satisfies $\mathbb{E}[|W_j \cap W'_\ell|] = w^2/2^\lambda$ and the Jaccard index is approximately $(w^2/2^\lambda)/(2w - w^2/2^\lambda)$. The predictor uses $\mathcal{H}_i^{(t)} = (\mathcal{T}_{i-w+1}, \dots, \mathcal{T}_i)$ to forecast $\hat{\mathcal{T}}_i(t + \Delta t)$ and align $\mathcal{W}_j = \{\mathcal{T}_i\}_{i=j}^{j+w-1}$ with token windows.

Deployment rules. Use fresh (k_s, s) each session and bind s to the transcript; keep token indices monotone within a session so crossing a boundary resets correlation; pick w for loss tolerance noting $\mathbb{E}[|W_j \cap W'_\ell|] = w^2/2^\lambda$ is negligible for $\lambda \geq 128$.

Algebraic Domains and Operators

We keep elliptic curve elements in $\mathbb{G}_{\text{EC}} = \langle G \rangle$ of prime order q with additive notation and scalar multiplication $[x]G$ for $x \in \mathbb{Z}_q$; multiplicative commitments live in $\mathbb{G}_\times = \langle g \rangle$ with independent bases $g, h, u, g_0, \dots, g_\omega$; do not mix the groups.

D. Standards Touch points and Message Slots

ZKPAS carries opaque proofs with ISO 15118-20 authorization messages and surfaces verification through OCPP 2.0.1 at the Charge Point or CPO back end; the air interface exposes only pseudonyms and tags.

E. Security Model and Adversarial Capabilities

We analyze in a Dolev Yao setting [28] with a probabilistic polynomial time adversary $\mathcal{A}$ that (i) controls channels and can intercept, modify, inject, reorder, and replay, (ii) cannot solve Elliptic Curve Discrete Logarithm or Computational Diffie Hellman, nor break the pseudorandom function, key derivation function, or message authentication code, (iii) may corrupt up to $t-1$ anchors and any providers ASP_j while honest device keys sk_i remain hidden, and (iv) may attempt trajectory inference using histories with bounded error; we exclude physical tamper and side channels and assume integrity protected training and model hosting.

TABLE II
NOTATION AND CONVENTIONS USED IN SECTION III.

Symbol	Type	Meaning
λ	integer	Security parameter; negligible $2^{-\lambda}$.
q	prime	Group order for $\mathbb{G}_{\text{EC}}$ and $\mathbb{G}_\times$.
G	EC point	Base point; $PK_i = [sk_i]G$, $sk_i \in \mathbb{Z}_q$.
F , KDF, MAC	maps	Pseudorandom function, key derivation, message authentication.
k_s, s	key, salt	Per session material from transcript binding.
t_i, W_j	bits, set	Token and sliding window $\{t_i\}_{i=j}^{j+w-1}$.
$\mathcal{D}_a, \mathcal{D}_b$	labels	Source and target domains; domID identifies target.
σ_d	signature	Domain attestation over (domID, τ) .
μ	tag	Transcript MAC under $k_d = \text{KDF}(k_{\text{sess}}^{\text{new}}, \text{domID})$.
C_1, C_2	ciphertext	ECIES like encapsulation of ID_{temp} .
$\Omega_{\text{cert}}, \Omega_{\text{trajectory}}$	group elems	Commitments in $\mathbb{G}_\times$.
w	integer	Window size for token acceptance.
$l_{\text{seg}}, v_{\text{max}}, \tau_{\text{max}}$	scalars	Mobility timing: $\tau_{\text{max}} = l_{\text{seg}}/v_{\text{max}}$.
L_{auth}	time	$L_{\text{comp}} + L_{\text{comm}} + L_{\text{sync}} \leq 1 \text{ ms}$.
E_{auth}	energy	$E_{\text{ECC}} + E_{\text{MAC}} + E_{\text{comm}} \leq E_{\text{budget}}$.

ID and Domain Hijack: An attacker may spoof a moving device pseudonym or forge handover beacons via a rogue access point or road side unit to bias prediction; ZKPAS must keep prediction useful while protecting privacy in both cases.

F. Security Requirements and Objectives

For any such $\mathcal{A}$, advantages are negligible in λ : **Authentication integrity** requires $\Pr[\text{Verify}(pk_i, \pi_{\text{ZK}}) = 1] \geq 1 - \epsilon$ with $\epsilon \leq 2^{-\lambda}$; **privacy preservation** requires $\text{Adv}_{\mathcal{A}}^{\text{privacy}} = |\Pr[\mathcal{A}(\Pi_0) = 1] - \Pr[\mathcal{A}(\Pi_1) = 1]| \leq \delta(\lambda)$; **forward security** requires that if state leaks at time t^* then forging any past $\text{Auth}_i(t)$ for $t < t^*$ remains negligible; **cross domain security** requires $\Pr[\text{unauthorized transition accepted}] \leq 2^{-\lambda}$ even if up to $t-1$ anchors are corrupted.

G. Performance Requirements and Constraints

Authentication must finish within mobility and DR timing: with minimum segment length l_{seg} and maximum velocity v_{max} we require $\tau_{\text{max}} = l_{\text{seg}}/v_{\text{max}}$ and $L_{\text{auth}} \leq \min\{\tau_{\text{max}}, \tau_j^{\text{DR}}\}$; across n roaming segments, cryptographic state and time scale as $\mathcal{O}(\log n)$ within device capacity C_i ; messages respect $|M| \leq M_{\text{max}}$; energy follows $E_{\text{auth}} = E_{\text{ECC}} + E_{\text{MAC}} + E_{\text{comm}} \leq E_{\text{budget}}$.

H. Notation and Conventions

We use $\mathbb{G}_{\text{EC}} = \langle G \rangle$ and $\mathbb{G}_\times = \langle g \rangle$ with order q ; sampling $x \xleftarrow{\$} \mathbb{Z}_q$; concatenation $a \parallel b$; challenges $\mathcal{C} \in \mathbb{Z}_q$. Canonical encodings hold inline for space: token $t_i = F_{k_s}(i \parallel s)$ and window $W_j = \{t_i\}_{i=j}^{j+w-1}$; identity encapsulation $C_1 = [r]G$, $K = \text{KDF}([r]\text{PK}_{\mathcal{D}_b} \parallel \tau \parallel \text{domID})$, $C_2 = ID_{\text{temp}} \oplus K$; domain binding $k_d = \text{KDF}(k_{\text{sess}}^{\text{new}}, \text{domID})$ and transcript tag $\mu = \text{MAC}_{k_d}(C_1 \parallel C_2 \parallel \tau \parallel \sigma_d)$. All symbols referenced here are summarized in Table II.

IV. ZKPAS PROTOCOL DESIGN

ZKPAS is a cryptographically rigorous authentication architecture tailored to V2G Demand Response (DR) settings

with high mobility. It integrates Schnorr-style zero-knowledge proofs, logarithmic sliding-window token chaining, mobility-informed precomputation, and threshold cryptography for cross-domain acceptance. We now present the constructions and show how they bind to DR windows while respecting strict latency and resource budgets.

A. Protocol Architecture and Cryptographic Foundations

Figure 1 depicts four coupled modules: a zero-knowledge core over elliptic-curve discrete logarithms; sliding-window token issuance with cryptographic continuity; learning-driven precomputation of future windows; and (t, n) threshold authorization across CPOs. Together, these modules replace repeated full re-authentication of length n with evidence that costs $\mathcal{O}(\log n)$ to verify, reducing airtime and stabilizing latency during roaming.

We instantiate elliptic-curve cryptography on a prime-order group $\langle G \rangle$ of size q , polynomial or vector commitments in a multiplicative group, and hash-based Message Authentication Codes (MACs). A mobile entity $\mathcal{M}_i$ holds (pk_i, sk_i) with $pk_i = [sk_i]G$. The verifier checks a zero-knowledge proof π_{ZK} . The k th mobility segment uses token $Auth_k$. The session key k_{sess} derives from an Elliptic Curve Diffie-Hellman key exchange and evolves by a Key Derivation Function (KDF). In each segment we embed timestamp τ_k , nonce N_k , and a continuity tag

$$L_k = \text{MAC}_{k_{sess}}(Auth_{k-1} \parallel \tau_k \parallel N_k),$$

so that evidence forms a compact chain. Relative to a naive design that sends all segment transcripts, the continuity chain and window proofs reduce the verifier's work from $\Theta(n)$ to $\Theta(\log n)$.

B. Zero-Knowledge Authentication Framework

ZKPAS proves possession of sk_i without disclosure using a Schnorr-style construction combined with group-consistent commitments. The prover outputs

$$\pi_{ZK} = (\Omega_{\text{commit}}, \Pi_{\text{auth}}, \Psi_{\text{trajectory}}, \mathcal{Z}_{\text{response}}),$$

where Ω_{commit} binds certificate fields, Π_{auth} carries the knowledge response, and $\Psi_{\text{trajectory}}$ encodes agreed mobility constraints that do not reveal the path. Algorithms 1 and 2 specify generation and verification while keeping elliptic-curve and multiplicative-group operations separate for clarity.

Algorithm 1 Zero-Knowledge Authentication Proof Generation

Require: Mobile entity $\mathcal{M}_i$ with key pair (pk_i, sk_i) , domain $\mathcal{D}_j$, sliding window $\mathcal{W}$, trajectory polynomial $\mathcal{F}_\omega(x)$

Ensure: Zero-knowledge proof π_{ZK}

- 1: $\mathbf{r} = (r_1, r_2, r_3) \xleftarrow{R} (\mathbb{Z}_q)^3$; $\rho, \xi, \eta \xleftarrow{R} \mathbb{Z}_q$
- 2: $\mathbf{R} = ([r_1]G, u^{r_2}, h^{r_3}) \in \mathbb{G}_{EC} \times \mathbb{G}_\times \times \mathbb{G}_\times$
- 3: $\Omega_{\text{cert}} = g^{H(\text{Cert}_i)} \cdot h^{pk_i} \cdot u^\rho$; $\Omega_{\text{trajectory}} = \prod_{j=0}^{\omega} g_j^{\mathcal{F}_\omega(j)} \cdot h^\xi$
- 4: $\mathcal{C} = H_{\text{domain}}(\mathbf{R} \parallel \Omega_{\text{cert}} \parallel \Omega_{\text{trajectory}} \parallel aux)$; Parse as $(c_1, c_2, c_3, c_{\text{traj}}) \in (\mathbb{Z}_q)^4$
- 5: $z_1 = r_1 + c_1 \cdot sk_i \bmod q$; $z_2 = r_2 + c_2 \cdot \rho \bmod q$; $z_3 = r_3 + c_3 \cdot \xi \bmod q$
- 6: $z_{\text{trajectory}} = \sum_{j=0}^{\omega} c_{\text{traj}}^j \cdot \mathcal{F}_\omega(j) + \xi \bmod q$
- 7: **return** $\pi_{ZK} = (\mathbf{R}, \Omega_{\text{cert}}, \Omega_{\text{trajectory}}, z_1, z_2, z_3, z_{\text{trajectory}}, \mathcal{C})$




Fig. 1. ZKPAS Architecture

Algorithm 2 Zero-Knowledge Authentication Proof Verification

Require: Proof π_{ZK} , public key pk_i , domain parameters $\mathcal{D}_j$

Ensure: Verification result $\{Accept, Reject\}$

- 1: Parse π_{ZK} into $(R_1, R_2, R_3, z_1, z_2, z_3, \mathcal{C}, \tau_{\text{proof}}, aux)$
- 2: Freshness check: **if** $|\tau_{\text{current}} - \tau_{\text{proof}}| > \Delta_{\text{max}}$ **then return** *Reject*
- 3: **Check-EC** (in $\mathbb{G}_{EC}$): **if** $[z_1]G \neq R_1 + [C]pk_i$ **then return** *Reject*
- 4: **Check-Cert** (in $\mathbb{G}_\times$): **if** $u^{z_2} \neq R_2 \cdot (u^\rho)^C$ **then return** *Reject*
- 5: **Check-Blind** (in $\mathbb{G}_\times$): **if** $h^{z_3} \neq R_3 \cdot (h^\xi)^C$ **then return** *Reject*
- 6: Compute $\Omega'_{\text{trajectory}}$ in original specification and recompute $\mathcal{C}' \leftarrow H_{\text{domain}}(R_1 \parallel R_2 \parallel R_3 \parallel \Omega_{\text{cert}} \parallel \Omega'_{\text{trajectory}} \parallel aux)$
- 7: **return** *Accept* if $\mathcal{C}' = \mathcal{C}$ and all checks pass; else *Reject*

C. Sliding Window Token Generation

We now connect token continuity to DR windows. Let $W_j = [t_j, t_{j+1})$ be the active window and let a segment index

within the window be $s \in \{0, \dots, w-1\}$. We derive a per-session key k_s and salt s at session start (see Section III). For each window, we bind indices by $i = \langle j, s \rangle$ and compute $t_i = F_{k_s}(i \parallel s)$. The continuity MAC is

$$\text{Link}_{k,k+1} = \text{HMAC}(k_{\text{sess}}^{(k)}, \text{Auth}_k \parallel \text{Auth}_{k+1} \parallel \tau_k \parallel N_k).$$

Session keys evolve as

$$k_{\text{sess}}^{(k+1)} = \text{HKDF}(k_{\text{sess}}^{(k)} \parallel \text{Chall}_{k+1} \parallel R_{k+1} \parallel N_{\text{challenge}} \parallel N_{\text{response}}).$$

A verifier that challenges continuity needs only $\log_2 w$ links using a doubling schedule. If $S_{\text{proof}}(w)$ and $V_{\text{work}}(w)$ are proof size and verifier operations, the recurrence

$$S_{\text{proof}}(w) = S_{\text{proof}}(\lfloor w/2 \rfloor) + \mathcal{O}(1), \quad V_{\text{work}}(w) = V_{\text{work}}(\lfloor w/2 \rfloor) + \mathcal{O}(1)$$

yields $S_{\text{proof}}(w), V_{\text{work}}(w) \in \mathcal{O}(\log w)$, which keeps L_{comm} and L_{comp} within the budget in III-B.

Algorithm 3 Sliding Window Token Authentication with Cross-Domain Support

Require: Mobile entity $\mathcal{M}_i$ with token Auth_k , session key $k_{\text{sess}}^{(k)}$, approaching segment $k+1$
Ensure: Seamless authentication transition with forward security

- 1: Detect boundary; retrieve Auth_{k+1} ; sample $N_{\text{trans}} \xleftarrow{\$} \{0, 1\}^{\ell_n}$
- 2: $\text{Link}_{k,k+1} \leftarrow \text{HMAC}(k_{\text{sess}}^{(k)}, \text{Auth}_k \parallel \text{Auth}_{k+1} \parallel \tau_{\text{current}} \parallel N_{\text{trans}})$
- 3: $\mathcal{M}_i \rightarrow \text{ASP}_{k+1}: \{\text{Link}_{k,k+1}, C_{k+1}, \text{MAC}_{k_{\text{sess}}}(\cdot)\}$ where $C_{k+1} = H(\text{Auth}_{k+1} \parallel \gamma_{k+1})$
- 4: ASP_{k+1} checks link and commitment; issues challenge Chall_{k+1}
- 5: Response $R_{k+1} \leftarrow \text{HMAC}(H(\text{Link}_{k,k+1} \parallel \text{Chall}_{k+1}), k_{\text{sess}}^{(k)})$
- 6: **if** verification succeeds **then**
- 7: $\mathcal{W}_{k+1} \leftarrow \mathcal{W}_k \cup \{\text{Auth}_{k+1}\} \setminus \{\text{Auth}_{k-w}\}$; derive $k_{\text{sess}}^{(k+1)}$ by HKDF
- 8: **end if**

D. Predictive Token Pre-computation

We pre-stage tokens for upcoming windows to mask handover delay. An LSTM model $\mathcal{M}_\theta$ trained on T-Drive and GeoLife traces [6], [7] predicts short-horizon motion and selects the next window W_{j+1} and its index subset. The predictor outputs

$$\hat{\tau}_i(t + \Delta t) = \mathcal{M}_\theta(\mathcal{H}_i^{(t)}) + \mathcal{E}_i(t + \Delta t), \quad \mathcal{E}_i \sim \mathcal{N}(0, \sigma^2),$$

and triggers batch verification when the domain-gated confidence $\rho_{\text{pred}}^{(j)} \geq \rho_{\text{min}}$. A Pareto schedule balances security, accuracy, and compute:

$$\min_{\mathcal{T}_{\text{pre}} \in \mathcal{C}} \|\mathbf{f}(\mathcal{T}_{\text{pre}}) - \mathbf{f}_{\text{ideal}}\|_2, \quad \mathbf{f} = (f_{\text{sec}}, f_{\text{acc}}, f_{\text{comp}}).$$

Algorithm 4 sketches the workflow. Although T-Drive and GeoLife focus on Beijing, the pipeline remains dataset agnostic and supports retraining on local fleets without changing the protocol.

Heterogeneity Mitigation: Compute. Device work is constant-time crypto only, so $C_{\text{dev}}^{\text{ZKPAS}} = \mathcal{O}(1)$ per authentication; optional on-device inference is budgeted by a tiny RNN with per-step cost $d_h(d_x + d_h) + d_h$ and is disabled on constrained nodes. **Accuracy/stability.** Prediction is gated by confidence η and attestation $\gamma_t \in \{0, 1\}$, yielding a conservative bound $A_{\text{uncond}} \geq \eta \cdot \Pr(\gamma_t = 1)$; an EWMA calibration $\phi_{t+1} = (1 - \alpha)\phi_t + \alpha g_t$ with $0 < \alpha < 1$ ensures $\mathbb{E}[\|e_{t+1}\|^2] \leq (1 - \alpha\mu)\mathbb{E}[\|e_t\|^2] + \alpha^2\sigma^2$ (mean-square stability). **Communication.** Sliding windows give $\mathcal{O}(\log n)$ signaling; with token reuse rate r , the expected bytes satisfy $\mathbb{E}[\text{Bytes}] = \text{Bytes}_{\text{base}}(1 - r) + \text{Bytes}_{\text{overhead}}$, so weaker links benefit as r grows without increasing device compute.

Algorithm 4 Predictive Token Pre-computation with Trajectory Modeling

Require: State vector $\mathbf{s}_i(t)$, historical dataset $\mathcal{H}_i$, prediction horizon Θ , computational budget $\mathcal{B}_{\text{comp}}$
Ensure: Optimized token repository $\mathcal{T}_{\text{pre}}^{\text{opt}}$ with confidence intervals

- 1: Build basis $\Phi(t) = [1, t, \dots, t^w]^T$; estimate parameters for the trajectory encoding
- 2: Sample an ensemble $\{\mathcal{T}_{\text{pre}}^{(j)}(t)\}_{j=1}^{N_{\text{ens}}}$ to quantify uncertainty
- 3: **for** prediction step $j = 1$ to N_{horizon} **do**
- 4: Compute posterior $p(\mathcal{S}_{t+j\Delta t} \mid \mathcal{H}_i, \mathbf{s}_i(t))$ and confidence $\rho_{\text{pred}}^{(j)}$
- 5: **end for**
- 6: Solve $\mathcal{T}_{\text{pre}}^* = \arg \min_{\mathcal{T}_{\text{pre}} \in \mathcal{C}} \|\mathbf{f}(\mathcal{T}_{\text{pre}}) - \mathbf{f}_{\text{ideal}}\|_2$
- 7: **for** each optimized span s_j^* **do**
- 8: Generate randomness $\mathbf{r}_j \xleftarrow{\$} (\mathbb{Z}_q^*)^3$, assemble $(\mathcal{A}_j, \pi_j, \rho_{\text{pred}}^{(j)})$
- 9: Store to repository $\mathcal{T}_{\text{pre}}^{\text{opt}} \leftarrow \mathcal{T}_{\text{pre}}^{\text{opt}} \cup \{(\mathcal{A}_j, \pi_j, \rho_{\text{pred}}^{(j)})\}$
- 10: **end for**
- 11: **return** $(\mathcal{T}_{\text{pre}}^{\text{opt}}, \mathcal{I}_{\text{confidence}})$

1) *Robust Prediction under ID/Domain Hijack:* We gate prediction on cryptographic attestation of the next domain. For domain identifier domID with keys $(\text{pk}_d, \text{sk}_d)$, derive per-domain secrets and tokens

$$k_d = \text{KDF}(k_s, \text{domID}), \quad s_d \xleftarrow{\$} \{0, 1\}^\lambda, \quad t_i^{(d)} = F_{k_d}(i \parallel s_d).$$

Each beacon carries $\sigma_d = \text{Sign}_{\text{sk}_d}(\text{domID} \parallel \tau)$ and a session MAC $\mu = \text{MAC}_{k_d}(m)$. We compute the gate

$$\gamma_t = \mathbb{1}_{\{\text{Verify}_{\text{pk}_d}(\sigma_d) \wedge \text{MAC}_{k_d}(m) = \mu\}}.$$

Prediction uses domain features only when attested:

$$\hat{\mathcal{P}}_i(t + \Delta t) = \begin{cases} \mathcal{G}_\theta(\mathcal{H}_i^{(t)}, \mathbf{c}_d^{(t)}), & \gamma_t = 1, \\ \mathcal{G}_\theta^{\text{safe}}(\mathcal{H}_i^{(t)}), & \gamma_t = 0, \end{cases}$$

where $\mathcal{G}_\theta^{\text{safe}}$ omits domain inputs and outputs a conservative distribution. Rotate (k_d, s_d) on every accepted handover.

E. Cross-Domain Authentication Protocol

We now describe secure transitions between domains $\mathcal{D}_a$ and $\mathcal{D}_b$. The trust-anchor federation $\mathcal{TA}$ uses (t, n) threshold signatures to authorize roaming without a full reset. Let $\mathcal{S} \subset \mathcal{TA}$ be any qualifying set of size t . Using Lagrange coefficients over $\mathbb{F}_q$, the aggregate signature is

$$\sigma_{\text{cross}} = \prod_{i \in \mathcal{S}} \sigma_i^{\mathcal{L}_i(\mathcal{S})} \bmod p, \quad \mathcal{L}_i(\mathcal{S}) = \prod_{\substack{j \in \mathcal{S} \\ j \neq i}} \frac{j}{j - i} \bmod q.$$

We include a temporary identifier ID_{temp} and an unlinkable proof $\pi_{\text{unlinkable}}$ of token validity. After acceptance, the session key rotates by

$$k_{\text{sess}}^{\text{new}} = \text{HKDF}(k_{\text{sess}}^{\text{transfer}} \parallel ID_{\text{temp}} \parallel \mathcal{D}_b \parallel N_{\text{handoff}}),$$

which preserves forward security across boundaries.

ID encapsulation: Let $(\text{PK}_{\mathcal{D}_b}, \text{SK}_{\mathcal{D}_b})$ be the EC key pair of domain $\mathcal{D}_b$, G the base point of order q . The device samples $r \xleftarrow{\$} \mathbb{Z}_q$ and computes

$$\begin{aligned} C_1 &= [r]G, \\ K &= \text{KDF}([r]\text{PK}_{\mathcal{D}_b} \parallel \tau \parallel \text{domID}), \\ C_2 &= ID_{\text{temp}} \oplus K, \\ C_{\text{ID}} &= (C_1, C_2). \end{aligned}$$

Bind to the session and domain via a MAC with the per-domain key k_d :

$$\mu = \text{MAC}_{k_d}(C_1 \parallel C_2 \parallel \tau \parallel \sigma_d), \quad k_d = \text{KDF}(k_{\text{sess}}^{\text{new}}, \text{domID}).$$




Fig. 2. Message Flow of ZKPAS

Algorithm 5 Cross-Domain Authentication with Threshold Cryptography

Require: Mobile entity $\mathcal{M}_i$ transitioning from $\mathcal{D}_a$ to $\mathcal{D}_b$, threshold parameters (t, n) , trust anchors $\mathcal{TA}$

Ensure: Secure cross-domain authentication with distributed trust guarantees

- 1: Generate transition request $\mathcal{R}_{trans} = \{ID_i, Cert_i, \mathcal{D}_a, \mathcal{D}_b, \tau_{trans}, Auth_{current}\}$
- 2: Compute $\pi_{auth}^{(a)} = \text{ZKProof}(sk_i, pk_i, Auth_{current}, \mathcal{D}_a)$
- 3: $\mathcal{D}_a \rightarrow \mathcal{TA}$: broadcast $\{\mathcal{R}_{trans}, \pi_{auth}^{(a)}, \sigma_{\mathcal{D}_a}(\mathcal{R}_{trans})\}$
- 4: **for** each $TA_i \in \mathcal{TA}$ **do**
- 5: Verify policy; if valid, output share $\sigma_i^{share} = \text{ThresholdSign}(sk_i^{share}, H(\mathcal{R}_{trans}))$
- 6: **end for**
- 7: Collect shares; if $|\{shares\}| \geq t$, compute σ_{cross} with Lagrange coefficients
- 8: Build $Token_{cross} = \{ID_{temp}, \mathcal{D}_b, \tau_{exp}, \sigma_{cross}, \pi_{unlinkable}\}$
- 9: $\mathcal{M}_i \rightarrow \mathcal{D}_b$: present $Token_{cross}$ and possession proof; on accept, derive k_{sess}^{new}
- 10: **return** Success

Domain recovery (edge only): The edge of $\mathcal{D}_b$ recovers $K' = \text{KDF}([\text{SK}_{\mathcal{D}_b}]C_1 \parallel \tau \parallel \text{domID})$ and $ID'_{temp} = C_2 \oplus K'$, verifies μ with k_d , and validates $\pi_{unlinkable}$. Intermediate relays observe only $(C_1, C_2, \tau, \sigma_d, \mu)$ and cannot learn ID_{temp} or ID_i .

Rule (no raw ID on-air): No field in R_{trans} contains ID_i or ID_{temp} in clear; the network sees only C_{ID} , τ , σ_d , and μ .

Transport message: We remove any cleartext user identifier from the handover broadcast. The device sends

$$R_{trans} = (C_{ID}, \tau, \sigma_d, \mu, \pi_{unlinkable}),$$

where C_{ID} is an ECIES/KEM-style ciphertext of the *temporary* identifier ID_{temp} under the next-domain public key, σ_d is the signed domain advert, μ is a MAC over the transcript, and $\pi_{unlinkable}$ is the existing ZK proof of token validity without revealing ID_i .

F. DR-Aware Window Binding and Standards Mapping

Before we close, we explain how the above artifacts sit inside EV charging messages. We bind each token window to

DR timing by tagging indices with the window identifier and the serving domain:

$$i = \langle j, s, \text{domID} \rangle, \quad t_i = F_{k_s}(i \parallel s), \quad W_j = \{t_i\}_{s=0}^{w-1}.$$

During a session, the On-Board Unit carries π_{ZK} and the current continuity evidence as opaque objects inside the vehicle to charger authorization exchange of ISO 15118–20. The Charge Point forwards the verification result and minimal metadata to the CPO back end through OCPP 2.0.1 custom data fields. No long-term identifier appears on air. This mapping preserves privacy while allowing existing stacks to transport proofs without redesigning lower layers.

G. Complexity, Budgets, and Failure Handling

We summarize costs and include the resynchronization rules that keep the protocol within budget.

Asymptotic costs: For n segments per roam and window size w , the prover computes $\mathcal{O}(1)$ curve operations per segment, while continuity evidence verifies in $\mathcal{O}(\log w)$. Proof size follows $S_{\text{proof}}(w) \in \mathcal{O}(\log w)$. Memory for the window index and chain state is $\mathcal{O}(\log w)$.

Latency budget: Let $L_{\text{auth}} = L_{\text{comp}} + L_{\text{comm}} + L_{\text{sync}}$. Sliding windows reduce retransmissions so that L_{comm} shrinks and the total meets the bound in III-B. In experiments we keep $|M| \leq M_{\text{max}}$ to fit typical vehicular MTUs.

Resync and loss: If a token is skipped, the verifier requests a logarithmic set of links. The device answers with at most $\lceil \log_2 w \rceil$ continuity MACs. If the window expires, the device rolls to W_{j+1} and rebinds $i = \langle j+1, s, \text{domID} \rangle$. Session rotation refreshes k_s and salt s , which preserves inter-session unlinkability as defined in Section III.

V. SECURITY ANALYSIS

We state ZKPAS guarantees via game-based definitions and reductions. Under standard hardness—Elliptic Curve Discrete Logarithm (ECDLP), Computational Diffie–Hellman (CDHP), and Pseudorandom Function (PRF) security—all adversarial advantages are negligible in λ , yielding explicit probabilistic bounds against strong network attackers.

A. Formal Security Framework and Cryptographic Assumptions

We use an enhanced Dolev–Yao model [28]: a PPT adversary fully controls the channel (intercept/modify/inject/replay) but remains bounded by λ . ZKPAS depends on:

(i) **ECDLP.** Given $(G, A = aG)$ in a prime-order group on $E(\mathbb{F}_p)$, recovering $a \in \mathbb{Z}_q$ is infeasible for PPT adversaries.

(ii) **CDHP.** From (G, aG, bG) , computing abG is infeasible for PPT adversaries.

(iii) **ROM.** The hash $H : \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$ is modeled as a random oracle; proofs may program H without granting non-negligible advantage.

(iv) **PRF security.** HKDF/MAC-derived functions are indistinguishable from random under adaptive queries [29]:

$$|\Pr[\mathcal{A}^{F_k} = 1] - \Pr[\mathcal{A}^R = 1]| = \text{negl}(\lambda).$$

B. Unified Game-Based Security Analysis

ZKPAS undergoes rigorous security analysis through four fundamental cryptographic security games addressing specific attack vectors and security properties essential for high-mobility IoT authentication.

Definition 1 (Unified Security Games). *An adversary $\mathcal{A}$ achieves victory in authentication integrity if successful generation of valid authentication proof π^* occurs without prior corruption; in trajectory privacy if distinguishing between authentication transcripts for distinct trajectory patterns $\mathcal{T}_0$ and $\mathcal{T}_1$ with non-negligible advantage; in forward security violation if session key compromise enables forgery of prior authentication material; and in cross-domain security if controlling fewer than t trust anchor shares enables unauthorized domain transitions.*

Theorem 1 (Comprehensive Security Bounds). *Under ECDLP, CDHP, and PRF assumptions, ZKPAS achieves bounded adversarial advantages:*

$$Adv_{\mathcal{A}}^{auth}(\lambda) \leq \frac{q_H^2}{2^{\lambda+1}} + \frac{q_{auth}}{2^{\ell_n}} + q_{auth} \cdot Adv_{\mathcal{B}}^{ECDLP}(\lambda) \quad (1)$$

$$Adv_{\mathcal{A}}^{priv}(\lambda) \leq n \cdot Adv_{\mathcal{B}}^{CDHP}(\lambda) + \frac{q_H}{2^{\lambda}} \quad (2)$$

$$Adv_{\mathcal{A}}^{fwd}(\lambda) \leq Adv_{\mathcal{B}}^{PRF}(\lambda) + \frac{q_H^2}{2^{\lambda}} + \frac{q_{mac}}{2^{\tau}} \quad (3)$$

$$Adv_{\mathcal{A}}^{cross}(\lambda) \leq \binom{n}{t-1} \cdot \frac{1}{2^{\lambda}} + Adv_{\mathcal{B}}^{thresh}(\lambda) \quad (4)$$

where q_H , q_{auth} , q_{mac} denote oracle queries, ℓ_n , τ specify nonce and MAC tag lengths, and n represents authentication sessions.

Proof Sketch: Authentication integrity follows from ECDLP hardness through transcript rewinding in Schnorr constructions [14], [30]. Trajectory privacy arises from a commitment scheme that hides properties and a zero-knowledge simulation under CDHP assumptions [11], [31]. Forward security utilizes HKDF one-wayness with temporal key isolation [23]. Cross-domain security requires a minimum of t signature shares for Lagrange coefficient reconstruction [13], [32].

C. Attack Resistance Analysis

ZKPAS provides comprehensive protection against sophisticated attack vectors through mathematical security guarantees. *Impersonation and credential forgery* attacks remain bounded by $\Pr[\text{ForgeAuth}] \leq \frac{1}{2^{\lambda}} + Adv_{\mathcal{B}}^{ECDLP}$ due to non-malleability properties in Schnorr proof constructions requiring ECDLP solution capability. *Replay attacks* achieve success probability $\Pr[\text{ReplaySuccess}] \leq \frac{1}{2^{\ell_n}} + \Pr[\text{NonceReuse}]$ through the high entropy nonce generation and temporal freshness validation mechanisms ensuring cryptographic freshness across authentication events.

Man-in-the-middle attacks require forgery of valid Message Authentication Codes and session key derivation where session keys generated through $k_{sess}^{(i)} = \text{HKDF}(k_{sess}^{(i-1)} \| R_i \| \tau_i)$ ensure adversaries lacking prior session keys cannot derive valid MAC authentication codes. *Trajectory inference protection* utilizes cryptographic commitment $\Omega_{traj} =$

$\prod_{j=0}^{\omega} g_j^{\mathcal{F}(j)} \cdot h^r$ where hiding properties of Pedersen commitment schemes prevent adversarial extraction of polynomial coefficients $\mathcal{F}(j)$, ensuring trajectory privacy preservation. *Cross-domain collusion resistance* against adversaries controlling at most $t-1$ trust anchors maintains forgery probability $\Pr[\text{ForgeCrossDomain}] \leq \frac{1}{2^{\lambda}}$ through unforgeability properties of threshold signatures [12].

D. Security Guarantees and Dataset Integration

ZKPAS trains supervised LSTM predictors on T-Drive and GeoLife traces [6], [7], [19], [21], [22] with calibrated uncertainty. The forecast $\hat{\mathcal{T}}_i(t)$ is then embedded via zero-knowledge commitments, which hide trajectory data while enabling verifiable checks against the prediction.

ZKPAS provides: *unforgeability* by reduction to ECDLP with adversarial advantage $\leq 2^{-\lambda}$; *zero knowledge* via simulation-based indistinguishability; *forward secrecy* from one-way key evolution using HKDF with temporal isolation; *privacy* through commitment hiding and session unlinkability that prevent trajectory inference; and *cross-domain scalability* using (t, n) -threshold cryptography with Byzantine tolerance. Collectively, these guarantees yield a cryptographically sound design that withstands strong adversaries while meeting high-mobility IoT demands.

VI. PERFORMANCE EVALUATION

We evaluate the ZKPAS in settings that mirror V2G DR operation. The study combines analytical models with experiments and includes an AI pipeline for short-horizon prediction using GeoLife [7], [21] and T-Drive [6], [19], [20]. Throughout, we report latency, energy, and communication cost from the perspective of DR windows $W_j = [t_j, t_{j+1})$, so that every number maps back to a control window and a roaming segment.

A. Experimental Framework and Artificial Intelligence Enhancement

We adopt a hybrid testbed-simulation setup. Cryptographic microbenchmarks run on an Intel Core i7-12700K (3.6 GHz, 32 GB DDR4-3200). Device-side costs are measured on an emulated ARM Cortex-M7 (512 KB SRAM, 2 MB flash) to reflect constrained On-Board Units. Wireless conditions are modeled in OMNeT++ using IEEE 802.11p with 300 m nominal range and 5 to 25 ms latency variation. The Long Short-Term Memory (LSTM) predictor, implemented in TensorFlow, is trained on GeoLife and T-Drive with approximately 24 million GPS samples in T-Drive alone.

We bind evaluation to DR windows using two window-aware metrics: end-to-end decision latency L_{auth} and missed-window probability

$$P_{miss}(j) = \Pr[L_{auth} > \tau_j^{DR}], \quad \tau_j^{DR} = t_{j+1} - t_j.$$

Prediction enables proactive token precomputation for upcoming segments. In steady state, tokens for the next three segments are prepared off-path and reused at a rate of 84.7% on T-Drive validation traces, which suppresses on-path elliptic-curve work and trims airtime. The trajectory estimator achieves

TABLE III
COMPUTATIONAL PERFORMANCE AND ENERGY CONSUMPTION ANALYSIS

Metric	ZKPAS	[24]	[8]	[15]	[10]	Improvement
<i>Computational Performance (milliseconds)</i>						
ECC Operations	0.876	4.961	2.700	3.068	3.240	79.96%
Authentication Phase	1.89	9.43	4.69	6.84	7.21	73.78%
Cross-Domain Transfer	3.24	12.67	7.42	9.15	10.38	68.81%
<i>Energy Consumption (microjoules per authentication)</i>						
Cryptographic Ops	1,048.5	5,029.5	2,759.1	3,131.1	3,307.5	75.2%
Communication	378.4	867.2	657.6	696.0	753.6	56.4%
Total Energy	1,498.5	6,049.9	3,528.5	3,951.1	4,195.1	75.2%
<i>Storage Requirements (bytes)</i>						
Static Storage	384	736	560	648	633	47.8%
Dynamic Storage	694	317	247	481	286	-119.0%
Total Storage	1,078	1,053	807	1,129	919	-2.4%

TABLE IV
COMMUNICATION OVERHEAD AND SCALABILITY ANALYSIS

Metric/Protocol	ZKPAS	[24]	[8]	[15]	[10]
<i>Communication Overhead (bytes per authentication)</i>					
Initial Handshake	196	368	276	342	289
Authentication Token	72	164	132	128	145
Cross-Domain Transfer	142	288	224	256	267
Predictive Pre-computation	89	N/A	N/A	N/A	N/A
Total per Authentication	126	289	219	232	251
Improvement	-	56.4%	42.5%	45.7%	49.8%
<i>Throughput Scalability (authentications per second at density)</i>					
20 entities/km ²	728	276	408	346	379
50 entities/km ²	632	219	343	286	317
80 entities/km ²	479	132	241	190	218
100 entities/km ²	350	59	153	106	132
Improvement at 80/km²	-	262.9%	98.8%	152.1%	119.7%

89.3% localization accuracy within a 50 m tolerance; we gate its outputs by cryptographic attestation as specified in Section IV. These figures suggest that prediction is useful but not blindly trusted, which is exactly the operating posture we aim for.

B. Computational Performance and Energy Efficiency Analysis

We profile ZKPAS on the Cortex-M7 and observe a mean initial authentication cost of 1.89 ms. In steady state, the sliding-window path keeps the on-path decision under 1 ms because the continuity check and window proof avoid a fresh zero-knowledge exchange. Table III lists primitive costs and energy per mode. Energy is estimated from board-level current draw models and breaks down as

$$E_{\text{auth}} = E_{\text{ECC}} + E_{\text{MAC}} + E_{\text{comm}},$$

with $E_{\text{auth}} \approx 1.5$ mJ for ZKPAS, a 75.2% reduction compared to [24]. By replacing repeated full handshakes with a continuity chain, the verifier's work per window drops from $\mathcal{O}(n)$ to $\mathcal{O}(\log n)$. If the token reuse rate is $r \in [0, 1]$, the expected number of elliptic-curve operations per transition satisfies

$$\mathbb{E}[N_{\text{ECC}}] = (1 - r) N_{\text{ECC}}^{\text{full}} + r N_{\text{ECC}}^{\text{reuse}},$$

and because $N_{\text{ECC}}^{\text{reuse}} \ll N_{\text{ECC}}^{\text{full}}$, the benefit scales linearly with r . On GeoLife-derived trajectories, we measure $r = 0.847$, which is consistent with the predictive pipeline.

C. Communication Overhead and Scalability Assessment

ZKPAS minimizes wire cost by using compact tokens, certificate elision, and continuity evidence that verifies in




Fig. 3. Computation across phases. The steady-state path is below 1 ms due to sliding-window verification.

Fig. 4. Aggregate wire bytes across segments. Window proofs grow slowly due to logarithmic continuity checks.

TABLE V
COMPREHENSIVE PERFORMANCE IMPROVEMENT SUMMARY

Performance Metric	[24]	[8]	[15]	[10]
Computational Improvement	79.96%	59.66%	72.37%	73.78%
Communication Reduction	56.4%	42.5%	45.7%	49.8%
Energy Efficiency Gain	75.2%	57.5%	62.1%	64.3%
Throughput Enhancement	159.6%	77.0%	108.1%	130.4%
Overall Performance	Superior	Superior	Superior	Superior

logarithmic time. Table IV shows a reduction of 56.4% relative to [24] and comparable gains over other baselines. Predictive precomputation avoids a real-time proof in many transitions, yielding a 2.15 ms latency reduction and an 89 byte decrease per cross-domain step. The expected bytes per transition follow

$$\mathbb{E}[\text{Bytes}] = (1 - r) \text{Bytes}_{\text{full}} + r \text{Bytes}_{\text{reuse}} + \text{Bytes}_{\text{overhead}},$$

so that communication shrinks as reuse r increases, while device compute remains constant.

Throughput scales as density rises because each verifier performs $\mathcal{O}(\log n)$ checks per window and only occasionally falls back to a fresh proof. At 80 entities per square kilometer, ZKPAS processes 479 authentications per second, improving on [24] by 262.9%. Figure 3 summarizes compute profiles across phases, and Figure 4 plots aggregate traffic for 20 T-Drive segments, where the ZKPAS curve reaches 2.59 KB compared to 7.36 KB for [24]. Table V aggregates improvements across metrics.

Results depend on the radio configuration and window sizes; we used the same curve parameters, hash functions, and compiler flags for all baselines to keep comparisons fair. In DR deployments, implementers should tune the window length τ_j^{DR} and the reuse threshold to match local traffic and tariff timing.

VII. CONCLUSION

We set out to keep authentication aligned with V2G Demand Response (DR) timing under fast mobility. ZKPAS achieves this by pairing zero-knowledge proofs with sliding-window continuity and cautious prediction. In steady state, the end-to-end decision fits the window budget since $L_{\text{auth}} \leq$

$\min\{\tau_j^{\text{DR}}, \tau_{\max}\}$ and typically remains below 1 ms. The analysis bounds the break probability by $\Pr[\text{break}] \leq 2^{-\lambda}$, preserves trajectory privacy under the CDHP and the ROM, and holds energy near $E_{\text{auth}} \approx 1.5 \text{ mJ}$, which supports the cost-effective aim of this research. ZKPAS contributes four pieces: (i) Schnorr-style elliptic-curve proofs with $\mathcal{O}(\log n)$ verifier work, (ii) a windowed token chain that trims airtime while maintaining forward security by key separation, (iii) LSTM based precomputation on GeoLife and T-Drive with 84.7% reuse that reduces on-path cryptography, and (iv) a (t, n) threshold anchor fabric that enables roaming between CPO with Byzantine tolerance. Against representative baselines, we observe 71.8% less compute, 93.9% faster authentication, and 69.5% lower energy, while sustaining a 98.5% success rate at 250 km/h. The protocol rides inside ISO 15118–20 exchanges and surfaces results through OCPP 2.0.1, so deployment can proceed incrementally. Results naturally depend on radio settings and DR window sizes, yet the $\mathcal{O}(\log n)$ checks keep verifiers stable across densities.

REFERENCES

- [1] T. Alladi, V. Chamola, N. Sahu, V. Venkatesh, A. Goyal, and M. Guizani, “A comprehensive survey on the applications of blockchain for securing vehicular networks,” *IEEE Communications Surveys & Tutorials*, vol. 24, no. 2, pp. 1212–1239, 2022.
- [2] L. Wei, J. Cui, H. Zhong, I. Bolodurina, C. Gu, and D. He, “A decentralized authenticated key agreement scheme based on smart contract for securing vehicular ad-hoc networks,” *IEEE transactions on mobile computing*, vol. 23, no. 5, pp. 4318–4333, 2023.
- [3] Y. Yang, L. Wei, J. Wu, C. Long, and B. Li, “A blockchain-based multidomain authentication scheme for conditional privacy preserving in vehicular ad-hoc network,” *IEEE Internet of Things Journal*, vol. 9, no. 11, pp. 8078–8090, 2021.
- [4] J. Sun, G. Xu, T. Zhang, X. Cheng, X. Han, and M. Tang, “Secure data sharing with flexible cross-domain authorization in autonomous vehicle systems,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 7, pp. 7527–7540, 2022.
- [5] C. Lin, X. Huang, and D. He, “Ebcpa: Efficient blockchain-based conditional privacy-preserving authentication for vanets,” *IEEE Transactions on Dependable and Secure Computing*, vol. 20, no. 3, pp. 1818–1832, 2022.
- [6] Y. Zheng, “T-drive trajectory data sample,” August 2011, t-Drive sample dataset. [Online]. Available: <https://www.microsoft.com/en-us/research/publication/t-drive-trajectory-data-sample/>
- [7] Y. Zheng, X. Xie, W.-Y. Ma *et al.*, “Geolife: A collaborative social networking service among user, location and trajectory,” *IEEE Data Eng. Bull.*, vol. 33, no. 2, pp. 32–39, 2010.
- [8] J. Zhang, H. Zhong, J. Cui, Y. Xu, and L. Liu, “Smaka: Secure many-to-many authentication and key agreement scheme for vehicular networks,” *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 1810–1824, 2020.
- [9] X. Li, J. Liu, M. S. Obaidat, P. Vijayakumar, Q. Jiang, and R. Amin, “An unlinkable authenticated key agreement with collusion resistant for vanets,” *IEEE Transactions on Vehicular Technology*, vol. 70, no. 8, pp. 7992–8006, 2021.
- [10] J. Chen, Z. Zhan, K. He, R. Du, D. Wang, and F. Liu, “Xauth: Efficient privacy-preserving cross-domain authentication,” *IEEE Transactions on Dependable and Secure Computing*, vol. 19, no. 5, pp. 3301–3311, 2021.
- [11] S. Goldwasser, S. Micali, and C. Rackoff, “The knowledge complexity of interactive proof-systems,” in *Providing sound foundations for cryptography: On the work of shafi goldwasser and silvio micali*, 2019, pp. 203–225.
- [12] D. Pointcheval and O. Sanders, “Short randomizable signatures,” in *Topics in Cryptology-CT-RSA 2016: The Cryptographers’ Track at the RSA Conference 2016, San Francisco, CA, USA, February 29-March 4, 2016, Proceedings*. Springer, 2016, pp. 111–126.
- [13] F. Mergoupis-Anagnou, “One-step schnorr threshold identification,” *Cryptology ePrint Archive*, 2025.
- [14] P. Wuille, J. Nick, and T. Ruffing, “BIP 340: Schnorr Signatures for secp256k1,” <https://bips.dev/340>, 2020, bitcoin Improvement Proposal.
- [15] P. R. Babu, A. G. Reddy, B. Palaniswamy, and S. K. Kommuri, “Ev-auth: Lightweight authentication protocol suite for dynamic charging system of electric vehicles with seamless handover,” *IEEE Transactions on Intelligent Vehicles*, vol. 7, no. 3, pp. 734–747, 2022.
- [16] A. Yang, J. Weng, K. Yang, C. Huang, and X. Shen, “Delegating authentication to edge: A decentralized authentication architecture for vehicular networks,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 2, pp. 1284–1298, 2020.
- [17] M. Wang, L. Rui, Y. Yang, Z. Gao, and X. Chen, “A blockchain-based multi-ca cross-domain authentication scheme in decentralized autonomous network,” *IEEE Transactions on Network and Service Management*, vol. 19, no. 3, pp. 2664–2676, 2022.
- [18] Y. Lin, X. Wang, Q. Gan, and M. Yao, “A secure cross-domain authentication scheme with perfect forward security and complete anonymity in fog computing,” *Journal of Information Security and Applications*, vol. 63, p. 103022, 2021.
- [19] J. Yuan, Y. Zheng, C. Zhang, W. Xie, X. Xie, G. Sun, and Y. Huang, “T-drive: driving directions based on taxi trajectories,” in *Proceedings of the 18th SIGSPATIAL International conference on advances in geographic information systems*, 2010, pp. 99–108.
- [20] J. Yuan, Y. Zheng, X. Xie, and G. Sun, “Driving with knowledge from the physical world,” in *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2011, pp. 316–324.
- [21] Y. Zheng, L. Zhang, X. Xie, and W.-Y. Ma, “Mining interesting locations and travel sequences from gps trajectories,” in *Proceedings of the 18th international conference on World wide web*, 2009, pp. 791–800.
- [22] Y. Zheng, Q. Li, Y. Chen, X. Xie, and W.-Y. Ma, “Understanding mobility based on gps data,” in *Proceedings of the 10th international conference on Ubiquitous computing*, 2008, pp. 312–321.
- [23] B. A. Chiphiko, H. Kim, P. Ali, and L. Eneya, “Forward secrecy attack on privacy-preserving machine authenticated key agreement for internet of things,” *Archives of Advanced Engineering Science*, vol. 3, no. 1, pp. 29–34, 2025.
- [24] P. R. Babu, R. Amin, A. G. Reddy, A. K. Das, W. Susilo, and Y. Park, “Robust authentication protocol for dynamic charging system of electric vehicles,” *IEEE Transactions on Vehicular Technology*, vol. 70, no. 11, pp. 11338–11351, 2021.
- [25] A. M. Aslam, R. Chaudhary, and A. Bhardwaj, “An aiot-enabled digital twin cavs with a drl-based framework for trajectory planning,” *IEEE Transactions on Intelligent Transportation Systems*, 2025.
- [26] A. M. Aslam, A. Bhardwaj, and R. Chaudhary, “Quantum-resilient blockchain-enabled secure communication framework for connected autonomous vehicles using post-quantum cryptography,” *Vehicular Communications*, vol. 52, p. 100880, 2025.
- [27] —, “A secure and privacy-preserving authentication framework for connected and autonomous vehicles based on drg-pbft and zero-knowledge proof,” *Accident Analysis & Prevention*, vol. 220, p. 108145, 2025.
- [28] D. Dolev and A. Yao, “On the security of public key protocols,” *IEEE Transactions on information theory*, vol. 29, no. 2, pp. 198–208, 1983.
- [29] D. Pointcheval and J. Stern, “Security arguments for digital signatures and blind signatures,” *Journal of cryptology*, vol. 13, pp. 361–396, 2000.
- [30] N. Fleischhacker, T. Jager, and D. Schröder, “On tight security proofs for schnorr signatures,” *Journal of Cryptology*, vol. 32, no. 2, pp. 566–599, 2019.
- [31] R. Canetti and M. Fischlin, “Universally composable commitments,” in *Advances in Cryptology—CRYPTO 2001: 21st Annual International Cryptology Conference, Santa Barbara, California, USA, August 19–23, 2001 Proceedings 21*. Springer, 2001, pp. 19–40.
- [32] A. Shamir, “How to share a secret,” *Communications of the ACM*, vol. 22, no. 11, pp. 612–613, 1979.