Post-Quantum Self-Updatable Pseudonyms for Conditional Privacy-Preserving Authentication in VANETs From Lattice-Based Credentials

Rui Xiao and Xiaoying Jia

***Abstract*—Authentication in vehicular networks must hide the sender and still allow an authority to identify it when needed. Schemes in which a vehicle derives its own pseudonyms avoid authority contact after registration, but known constructions rest on assumptions that quantum computers break. We present L-SUPA, a lattice-based scheme of this kind. The credential contains a binary key, and the pseudonym of an epoch is a vector of rounded inner products of that key with public pseudo-random vectors; because the rounding is an exact integer relation, the statement that a pseudonym is the rounding of the signed key fits, together with possession of a credential, inside the presentation proof of a practical lattice anonymous-credential system. Messages are signed with a fresh FN-DSA-512 key bound to the pseudonym, and a tracing authority separate from the issuer identifies a sender by scanning stored keys. The relation raises the presentation proof from 79.6 to 81.4 KiB and its verification time by about 3%; on one desktop core a presentation verifies in 0.15 s, the time of about 4500 message-signature verifications, and a message signature verifies in 0.033 ms. A presentation adds 0.32% to the message bytes of a one-hour epoch at 10 Hz for 722-byte messages, the most favorable rate, and messages of one vehicle are linkable within an epoch. The analysis states, goal by goal, whether we have a reduction, a conditional reduction, or a sketch; under the generic hybrid the presentation-layer M-LWE and LWR terms leave margins of about 80 and 176 bits for 2^{30} vehicles over ten years (no claim of 128-bit anonymity).**

***Index Terms*—lattice-based cryptography, post-quantum, VANET, conditional privacy, pseudonym, anonymous credentials**

I. Introduction

Vehicles in a vehicular ad hoc network (VANET) broadcast signed safety messages several times per second. Each message must be authenticated, yet a long-lived identifier would let any listener follow a vehicle. The standard answer is a pool of short-lived pseudonym certificates [1]–[5] with an authority that maps a pseudonym back to a vehicle when a message must be attributed (*conditional privacy*). The pools must be refilled from an authority, and their elliptic-curve signatures fall to a quantum computer [6]; vehicles stay in service for a decade or more, so a post-quantum replacement is needed within their lifetime.

An attractive alternative is to let the vehicle *derive* its pseudonym. In the scheme of Akil et al. [7] a vehicle registers once, receives an anonymous credential and computes a fresh pseudonym per epoch with a proof that it belongs to a valid credential; no authority is contacted again. That construction rests on strong RSA. Lattice CPPA schemes exist [8]–[10], but none lets the vehicle update its pseudonym across epochs without contacting an authority, and none relates signature size to the length of a vehicular message (the size comparison in Section II is carried by the 666 B FN-DSA signature, a property of the standard). Practical lattice credentials have meanwhile become available: Argo et al. [11] report a 6.97 kB signature and a 79.58 KiB presentation. To the best of our knowledge, based on the papers we read in full, none of the surveyed schemes combines a credential-derived pseudonym with a tracing authority separate from the issuer; the closest, Akil et al., rests on strong RSA and keeps tracing with the issuer.

Approach: A vehicle's credential signs a binary key $\mathbf{k}$ of 768 bits; in epoch ε its pseudonym is the vector of 24 rounded inner products of $\mathbf{k}$ with public pseudo-random vectors $H_{\mathsf{pid}}(\varepsilon, j)$. Since the inner product is linear in $\mathbf{k}$ and rounding only splits off low bits, "this pseudonym is the rounding of the signed key" is an *integer relation* that fits inside the credential presentation proof as 24 linear equations and 504 witness bits. A vehicle presents once per epoch with this proof and a fresh FN-DSA-512 key, and signs messages with that key. A separate tracing authority, which receives $\mathbf{k}$ at registration but never usk, recomputes pseudonyms to identify a sender. Section IV states the trust assumptions.

Contributions:

1) *Pseudonyms from LWR rounding, proved inside an anonymous-credential presentation* (Section V-D). The epoch pseudonym PID_ε is defined by adding the integer relation $\langle \boldsymbol{\theta}_{\varepsilon,j}, \mathbf{k} \rangle = 2^{19}\mathsf{cy}_j + 2^{11}\mathrm{PID}_\varepsilon[j] + \mathsf{rm}_j$ to the presentation proof of [11], which proves knowledge of the binary witness carrying $\mathbf{k}$. The rounding-to-exact-relation rewriting is due to Esgin et al. [12]; we add the epoch vectors $\boldsymbol{\theta}_{\varepsilon,j}$, the witness bits shared between the credential relation and the 24 equations, and a tracing authority separate from the issuer. Witness dimension rises from 211 to 219, the proof from 79.58 to 81.44 KiB, and verification by 3.0–3.3% (Section VII-C).
2) *The L-SUPA protocol* (Section V): registration, epoch presentation, per-message signatures with a cache identifier, and conditional tracing. The TRA draws and stores $\mathbf{k}$ but never sees usk, deletes the commitment randomness, and serves only logged and authorized requests

Manuscript received XX Month 2026; revised XX Month 2026; accepted XX Month 2026. Date of publication XX Month 2026; date of current version XX Month 2026.

The authors are with South-Central Minzu University, Wuhan 430074, China (e-mail: xiaorui20010725@163.com; corresponding author: Rui Xiao).

(a trust assumption; the prototype implements no TRA service). The single-thread tracing scan is measured for 2^{20} keys and extrapolated to 2^{30}.

3) *A security analysis whose status is stated goal by goal* (Section VI): for each of five goals we say whether we have a reduction, a black-box argument, a sketch, or a gap. Unlinkability follows from a hybrid of M-LWE and binary-secret LWR with explicit loss factors (in a 2^{30}-vehicle, ten-year scenario the M-LWE loss leaves a margin of 79.6 bits and the LWR loss 176.4 bits under Core-SVP, or 197.6 bits with the estimator's built-in model). The supplementary material checks each change to the original soundness argument.
4) *A prototype and an accounting of what it does and does not show* (Section VII). We extended the reference implementation of [11], fixed an error in its challenge-norm function, measured the proof layers, FN-DSA-512 and the tracing scan on one desktop core, and related message size to a single-frame budget (666 B of the message are the FN-DSA signature, a property of the standard). Scripts, raw logs and the modified prototype are available at https://doi.org/10.5281/zenodo.24433866.

Scope and limitations: The analysis covers the changes to the argument of Argo et al. without replaying its extractor; the anonymity loss is bounded with explicit, not tight, loss factors; the proofs are in the classical random oracle model and assume an honest-but-curious issuer. Revocation and non-frameability are out of scope. The evaluation reports one prototype: the presentation stays large (84 kB), all costs are for presentations disclosing no attribute, and timings come from a desktop CPU.

Sizes are given in bytes, kB (10^3) or KiB (2^{10}); the supplementary material lists the thirteen points where our proof differs from that of Argo et al.

II. Related Work

A. Classical Pseudonym Management

Deployed designs give each vehicle a pool of short-lived pseudonym certificates from an authority [2]: the U.S. Security Credential Management System issues them in batches (e.g. 20 per week) through several authorities that are meant not to link pseudonyms to a vehicle [2], and ETSI/IEEE define the formats and message signing [4], [5]. IFAL pre-issues a lifetime supply and releases each epoch's keys with a 128-bit activation code, so revocation lists are unnecessary and at most two certificates are valid per vehicle; it remains an ECDSA construction with non-colluding enrolment and authorization authorities [3]. All depend on an authority for every pool and on elliptic-curve signatures, which quantum computers break.

B. Self-Derived Pseudonyms

Akil et al. register a vehicle once, hand it an anonymous credential, and let it derive a fresh pseudonym per epoch with a proof of validity [7]; pseudonyms have the form $H(\varepsilon)^{a_0}$, an epoch-dependent hash raised to a credential-bound secret. The construction rests on strong RSA and its security discussion is informal. L-SUPA follows this architecture with a post-quantum credential and gives one pseudonym per credential and epoch.

C. Lattice-Based Schemes

Several lattice CPPA schemes exist. Zhou et al. build an identity-based ring signature with fog-node batch verification and report 20.7 kB signatures at 74 bits under their own parameters [8]. In LRCPA a vehicle chooses its pseudonym $H_1(ID \,\|\, s)$ and registers it with the authority and a signature on its identity, making the trace publicly verifiable [9]; every update is a new registration, each message carries the sender's public matrix, the maximum message length is $165{,}587$ bytes ($n = 100$, $m = 666$, $k = 80$, $q = 101$, 123 bits), only non-repudiation is proved formally (ISIS), and anonymity and unlinkability are informal, at the level of our own sketch-grade analyses (Section VI). LETA targets platoons [10]: ring-LWE aggregate signatures and a hash tree locate a failing member, the TA assigns and re-issues pseudonyms, and unforgeability is proved in the QROM. None of the four lets the vehicle update its pseudonym across epochs without contacting an authority, and none fits the per-message cost into a single-frame budget (LRCPA bounds the message length, not a frame).

For anonymous credentials, LNP22-style lattice proofs [12], [13] lead to the practical credentials of Argo et al. [11]. Esgin et al. rewrite a rounding into an exact relation that is proved in zero knowledge, with lattice VRFs as an application [12]; we use the same rewriting and go beyond it in three respects. The witness bits of the key $\mathbf{k}$ are shared between the credential's signature relation and the 24 integer linear equations that define the epoch pseudonym, so that the pseudonym is tied to the signed credential; the public vectors $\boldsymbol{\theta}_{\varepsilon,j}$ are epoch-dependent, so that each epoch yields a new pseudonym from the same signed key; and a tracing authority, separate from the issuer, recomputes the pseudonyms to identify a sender. We have not studied the VRF construction of that work in detail; our increment does not rely on any property of the VRF. None of these works considers epoch pseudonyms, tracing, or vehicular constraints.

D. Positioning

Table I compares these works on four properties. We do not compare running times or sizes: parameters, platforms and security levels differ. L-SUPA is the only entry with both a post-quantum basis and pseudonyms updated without authority contact; to our knowledge that combination, with a tracing authority separate from the issuer, has not been constructed before. On cost, a presentation is large (Section VII).

Entries reflect our reading of the cited papers, not all of whose final versions we read. The "Sybil resistance" column asks whether one identity can obtain several pseudonyms in one epoch under the scheme's own assumptions; "not discussed" means the paper does not state this.

III. Preliminaries

A. Lattices, M-SIS and M-LWE

For a power of two n and a modulus q let $R = \mathbb{Z}[X]/(X^n + 1)$ and $R_q = R/qR$; vectors are bold and $\|\cdot\|$ is the Euclidean norm. We use the module problems over R_q of rank d as in [11], [13].

Definition 1 (Module short integer solution, M-SIS)**.** *Given $A \xleftarrow{\$} R_q^{d\times m}$ and a bound β, find $\mathbf{z} \in R_q^m \setminus \{0\}$ with $A\mathbf{z} = 0$ and $\|\mathbf{z}\| \le \beta$.*

TABLE I
QUALITATIVE COMPARISON (EA/AA: ENROLMENT/AUTHORIZATION AUTHORITY; CA: CERTIFICATE AUTHORITY; TA: TRUSTED AUTHORITY).

Scheme	PQ basis	Update w/o authority	Tracing	Sybil resistance
SCMS [2]	–	batches	several authorities	by batch size
IFAL [3]	ECDSA	pre-issued	EA, AA	2 certs
Akil et al. [7]	RSA	yes	CA = issuer	per credential
Zhou et al. [8]	Ring-SIS	–	TA	PUF identity
LRCPA [9]	ISIS	re-reg.	CA, public	not discussed
LETA [10]	ring-LWE	–	TA, leader	not discussed
L-SUPA	M-SIS, M-LWE, LWR	yes	TRA $\neq$ issuer	per epoch and credential

Definition 2 (Module learning with errors, M-LWE, knapsack form)**.** *Given* $A \xleftarrow{\$} R_q^{d\times m}$*, distinguish* $A\mathbf{r}$ *for a short* $\mathbf{r}$ *from a uniform vector of* R_q^d*.*

The inhomogeneous variant M-ISIS asks for a short $\mathbf{z}$ with $A\mathbf{z} = \mathbf{u}$; the knapsack form used below is equivalent up to a dimension change. Concrete security is the *Core-SVP* cost ($2^{0.292\beta}$ classical sieving in block size β) [14]; we report classical bits without model-specific polynomial factors. NIST's lattice standards use the same families [15].

B. Argo Credentials and Lattice Zero-Knowledge Proofs

L-SUPA takes its credential layer from Argo et al. [11], who build lattice anonymous credentials from three components; lemma and table numbers refer to the full version [16].

a) Signature (layer sg*):* A stateful signature on a user public key and m_{sg} messages of n_{sg} bits. Each signature carries a binary tag of weight five out of $n_{\mathsf{sg}} = 256$ positions, so the tag space is $\binom{256}{5} \approx 2^{33}$; a signer key is limited to $Q_{\mathsf{sg}} = 2^{32}$ signatures (capacity of [11], not re-derived). Signature size 6971 B (formula value); parameters $(n, d, m, q)_{\mathsf{sg}} = (256, 4, 10, 425801)$.

b) Issuance proof (layer is*):* The user commits to $(\mathrm{usk}, \mathsf{att})$ and proves in zero knowledge that the commitment is well formed; the issuer signs it without learning the message. Zero knowledge is that of [16, Lemma 7.3], against *honest* verifiers only; the soundness error of the issuance proof is given there as well.

c) Presentation proof (layer sh*):* The user proves knowledge of a signature on $(\mathrm{usk}, \mathsf{att})$ without revealing either: a system of linear and quadratic equations over $R_{q_{\mathsf{sh}}}$ with a binary witness, proved in the framework of Lyubashevsky, Nguyen and Plançon (LNP22) [13] (building blocks [12]), made non-interactive with the Fiat–Shamir transform [17], [18].

d) Model of the proof: The proofs hold in the classical random oracle model [18] (extraction rewinds; Argo et al. claim neither QROM nor UC security [11]); we inherit this (Section VI-D). Soundness is bounded by Lemma 7.5 of [16], restated as Eq. (5).

C. Rounded Linear Pseudonyms (LWR)

Learning with rounding [19] replaces LWE's explicit noise by deterministic rounding: for $q_{\mathsf{L}} > p_{\mathsf{L}}$, $\lfloor(\langle \mathbf{a}, \mathbf{s}\rangle \bmod q_{\mathsf{L}}) \cdot p_{\mathsf{L}}/q_{\mathsf{L}}\rfloor$ is indistinguishable from uniform. Reductions from LWE exist [19]; binary-secret LWE keeps hardness at a mildly larger dimension [20]. No reduction covers rounding *and* a binary secret at our parameters ($q_{\mathsf{L}} = 2^{19}$, $p_{\mathsf{L}} = 2^8$, dimension 768, 24 outputs per epoch), so we state it as assumption A3 (Section VI) and measure it against known attacks.

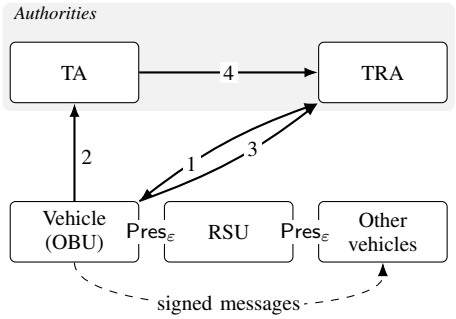

Fig. 1. System model: registration (solid), broadcast (dashed).

D. FN-DSA-512

Messages are signed with Falcon [21], selected by NIST as FN-DSA (no published standard text found [22]), at the level-1 parameter set (degree 512): public keys 897 B, signatures 666 B fixed-length (655 B average compressed). We use the Falcon specification's parameters and sizes and measure the PQClean implementation; the FN-DSA text may differ in details (domain separation, pre-hashing).

IV. SYSTEM MODEL, THREAT MODEL AND DESIGN GOALS

This section fixes the entities, the adversary and the goals; the symbols are defined where they first appear.

A. Entities and Network Model

Four kinds of entities take part (Fig. 1). Registration: (1) the vehicle asks the TRA for its PID key, which the TRA draws at random; (2) it sends its commitment and issuance proof to the TA; (3) it sends the opening to the TRA; (4) the TA forwards $(\mathrm{RID}, \mathrm{upk}, \mathrm{cmt})$ to the TRA (RID: registration identifier, upk: public user key, cmt: commitment). The TRA's confirmation and the return of the credential (steps 8 and 9 of Algorithm 1) are not drawn. Each epoch the vehicle broadcasts its presentation $\mathsf{Pres}_\varepsilon$, which roadside units may cache, then its signed messages. The **TA** issues credentials and holds sk_{TA}; it sees a commitment to the vehicle's keys, never $\mathbf{k}$. The **TRA**, operated by another organization, draws each $\mathbf{k}$, receives the commitment opening at registration, stores $(\mathrm{RID}, \mathbf{k})$ and answers authorized trace requests; it only traces, revocation being out of scope (Section VIII). **Vehicles** with an on-board unit (OBU) hold one credential, $\mathbf{k}$ and usk. **Roadside units (RSU)** relay and cache presentations; they hold no secrets and are trusted only for availability. Registration uses a confidential vehicle–TRA channel and an authenticated TA–TRA channel; the former is a model primitive (e.g., out-of-band enrollment) whose establishment we do not specify. Afterwards vehicles and verifiers use only the broadcast channel, with no authority contact in an epoch. Epochs have fixed length T_{ep} (1 h) and clocks are loosely synchronized.

B. Trust and Adversary Model

TM1 *Setup and vehicles.* Public parameters and TA keys are generated honestly. The adversary corrupts any set of vehicles, learns their keys and credentials, and lets them deviate and collude arbitrarily. Honest vehicles keep their secrets (no side channels).

TM2 *TA: honest-but-curious.* It follows the protocol and issues only after TRA confirmation, but may keep and analyze everything it sees, including registration transcripts, and may be corrupted after the fact (the adversary may hold sk_{TA} in the unlinkability game).

TM3 *TRA: procedure-compliant.* **R1**: draw every $\mathbf{k}$ uniformly, never issue a stored key twice, confirm only when the opened key matches the one issued for that RID, and delete $\mathbf{r}_0, \mathbf{r}_1$ after

verifying the opening. **R2**: serve trace requests only under external authorization and append them to a hash-chained, append-only log.

TM4 *TRA compromise.* A TRA database compromise is possible; Section VI-D states the consequences, mitigations are future work.

TM5 *Channels.* Vehicle–TRA: confidential and TRA-authenticated; TA–TRA: authenticated; broadcast: controlled by the adversary, who may eavesdrop, delay, drop, replay and inject.

TM6 *Collusion.* TA and TRA may collude only as described in Section VI-D, outside the guarantees of Goals 3 and 5(a).

The adversary is computationally bounded and classical for the proofs; M-SIS, M-LWE and M-ISIS are believed hard for quantum adversaries; A3 (binary-secret LWR) has no reduction or published analysis (Section III-C), so its post-quantum status is an assumption. The analysis is not in the quantum random oracle model (Section VI-D).

C. Design Goals

G1 *Authentication.* A receiver accepts a message only if it was sent by a registered vehicle in the claimed epoch under the pseudonym it presents. Replay of a genuine message is limited only by the timestamp check of the receiver, which we do not analyze.

G2 *Same-epoch uniqueness per credential.* One credential yields one pseudonym per epoch.

G3 *Cross-epoch unlinkability.* Given sk_{TA} and all transcripts seen by the TA (the confidential vehicle–TRA channel excluded), an observer cannot tell whether two pseudonyms of different epochs with the same $\mathcal{A}_{\mathrm{pub}}$ belong to one vehicle; different attribute sets disclose $\mathcal{A}_{\mathrm{pub}}$, a quasi-identifier.

G4 *Conditional traceability.* With the TRA's database and an authorization, the identity behind any accepted pseudonym can be found.

G5 *Separation of issuance and tracing.* The TA never learns $\mathbf{k}$; the TRA never learns usk.

G6 *Efficiency.* No authority interaction after registration; a verifier pays one presentation verification per sender and epoch and one signature verification per message.

Section VI states which of G1–G5 have an argument and which only a sketch; G6 is evaluated in Section VII. Revocation, non-frameability and location privacy are excluded.

V. The L-SUPA Scheme

The subscripts sg, sh, is denote the signature, presentation ("show") and issuance layers, and a symbol without subscript is layer-independent. The credential layer is Argo et al. [11] (Section III); new here are the pseudonym function, its proof inside a presentation, and the protocol.

A. Overview and Notation

A vehicle registers once. In each epoch ε of length T_{ep} it derives PID_ε locally from a secret $\mathbf{k}$ bound into its credential, broadcasts one *presentation* $\mathsf{Pres}_\varepsilon$ proving in zero knowledge that PID_ε and a fresh epoch key pk_ε belong to a TA-issued credential, and signs every message with the epoch key. Verifiers check the presentation once per sender and epoch, cache it, and pay one signature verification per message; no authority is contacted after registration.

B. Setup

The TA runs the key generation of the signature scheme of Argo et al., obtaining $(\mathsf{pk}_{TA}, \mathsf{sk}_{TA})$, and publishes the public parameters: common reference strings, the three parameter layers of Section VII-B, and the random oracles H_{pid}, H_{sid}, H_{ctx}. A credential has ten 256-bit message slots; the first three carry the PID key and the rest the public attributes $\mathcal{A}_{\mathrm{pub}}$. The TRA has no encryption key: tracing recomputes pseudonyms from stored keys (Section V-G).

Algorithm 1 Registration (once per vehicle)

1: **Vehicle** $\rightarrow$ **TRA** (confidential channel): RID.
2: **TRA:** draw $\mathbf{k} \xleftarrow{\$} \{0,1\}^{n_{\mathsf{k}}}$ (redraw if $\mathbf{k}$ is in $\mathsf{DB}_{\mathsf{TRA}}$); keep $(\mathrm{RID}, \mathbf{k})$ as pending.
3: **TRA** $\rightarrow$ **Vehicle** (confidential channel): $\mathbf{k}$.
4: **Vehicle:** sample binary usk, $\mathrm{upk} \leftarrow D_s$ usk; $\mathsf{att} \leftarrow (\mathbf{k} \| \mathcal{A}_{\mathrm{pub}})$; compute $(\mathbf{r}_0, \mathbf{r}_1)$ and cmt as in the issuance protocol of [11]; compute π^{is}.
5: **Vehicle** $\rightarrow$ **TA:** $(\mathrm{RID}, \mathrm{upk}, \mathrm{cmt}, \pi^{\mathsf{is}})$.
6: **Vehicle** $\rightarrow$ **TRA** (confidential channel): $(\mathbf{r}_0, \mathbf{r}_1, \mathsf{att})$.
7: **TA** $\rightarrow$ **TRA** (authenticated channel): $(\mathrm{RID}, \mathrm{upk}, \mathrm{cmt})$.
8: **TRA:** check $\mathrm{cmt} - \mathrm{upk} = \mathbf{r}_0 + A'\mathbf{r}_1 + D \cdot \mathsf{att}$ with $\mathbf{r}_0, \mathbf{r}_1, \mathsf{att}$ binary; **reject** if the key part of att differs from the pending $\mathbf{k}$ of this RID; store $(\mathrm{RID}, \mathbf{k})$; **delete** $\mathbf{r}_0, \mathbf{r}_1$; confirm to the TA.
9: **TA:** verify π^{is} and uniqueness of RID; after the TRA's confirmation, issue the credential (a signature with efficient protocols (SEP) on $(\mathrm{usk}, \mathsf{att})$ through the commit-transferable issuance of [11]).

C. Registration

Registration is interactive and runs once (Algorithm 1). The TRA draws $\mathbf{k}$ uniformly, never reissues a stored key and returns it over the confidential channel; the vehicle cannot choose $\mathbf{k}$, which makes the collision bound statistical. The vehicle commits to $(\mathrm{usk}, \mathsf{att})$, $\mathsf{att} = (\mathbf{k} \| \mathcal{A}_{\mathrm{pub}})$, and proves well-formedness with π^{is}; the TA sees only the commitment. The TRA verifies the opening it receives, checks that the key part is the key it issued for this RID, stores $(\mathrm{RID}, \mathbf{k})$, deletes $\mathbf{r}_0, \mathbf{r}_1$ (procedure R1) and confirms. It never sees usk.

The TRA alone draws and stores $\mathbf{k}$, ruling out duplicate keys and keys chosen after the public vectors are known; it is online during issuance (a single point of failure for new registrations only, since registered traffic verifies against pk_{TA} alone).

D. The Pseudonym Function and Its Proof

a) Definition: For epoch ε and $j \in [N_{\mathsf{pid}}]$ let $\boldsymbol{\theta}_{\varepsilon,j} = H_{\mathsf{pid}}(\varepsilon, j) \in \mathbb{Z}_{q_{\mathsf{L}}}^{n_{\mathsf{k}}}$ and

$$\mathrm{PID}_\varepsilon[j] = \left\lfloor \left(\langle \boldsymbol{\theta}_{\varepsilon,j}, \mathbf{k} \rangle \bmod q_{\mathsf{L}} \right) \cdot \frac{p_{\mathsf{L}}}{q_{\mathsf{L}}} \right\rfloor, \qquad q_{\mathsf{L}} = 2^{19},\ p_{\mathsf{L}} = 2^{8}. \quad (1)$$

The pseudonym is the vector $\mathrm{PID}_\varepsilon \in \{0, \ldots, 255\}^{N_{\mathsf{pid}}}$, 24 bytes for $N_{\mathsf{pid}} = 24$. Rounding is the only non-linear step. Like $H(\varepsilon)^{a_0}$ in self-derived pseudonyms (Section II-B), it is a deterministic function of a credential-bound secret and the epoch, so a vehicle cannot choose an epoch's pseudonym independently of its signed key, yet epochs look unrelated to anyone without $\mathbf{k}$ (assumption A3).

b) Integer relation: Because $q_{\mathsf{L}}/p_{\mathsf{L}} = 2^{11}$ is an integer, Eq. (1) is equivalent to the existence of integers $\mathsf{cy}_j \in [0, 1023]$ and $\mathsf{rm}_j \in [0, 2047]$ with

$$\langle \boldsymbol{\theta}_{\varepsilon,j}, \mathbf{k} \rangle \;=\; q_{\mathsf{L}} \cdot \mathsf{cy}_j \;+\; 2^{11} \cdot \mathrm{PID}_\varepsilon[j] \;+\; \mathsf{rm}_j, \quad (2)$$

over $\mathbb{Z}$: the 19 low bits of the inner product split into 8 output bits and 11 remainder bits, and the carry cy_j collects the rest. The carry and remainder bits form the extra binary witness of the presentation proof; $\mathrm{PID}_\varepsilon[j]$ is public. The same rewriting of a rounding into an exact relation is used to prove rounded outputs in [12]. Since $\mathbf{k}$ is binary, $\langle \boldsymbol{\theta}_{\varepsilon,j}, \mathbf{k} \rangle \le n_{\mathsf{k}}(q_{\mathsf{L}} - 1) < 2^{28.6}$, so $\mathsf{cy}_j \le 767$ and ten bits suffice.

c) Embedding in the presentation proof: The presentation proof of [11] is a Fiat–Shamir argument over rings of degree $n_{\mathsf{sh}} = 64$ in the LNP22 framework [13] that already proves knowledge of a binary witness $\mathbf{w}$ containing usk and all credential message slots, in particular $\mathbf{k}$. We extend it in three steps.

1) *Extra witness.* The $21N_{\mathsf{pid}} = 504$ bits $\mathsf{cy}_{j,\iota}, \mathsf{rm}_{j,\iota}$ are appended to $\mathbf{w}$ as $E = 8$ additional ring elements ($8 \cdot 64 = 512$ coefficients). The witness dimension grows from $m_{1,\mathsf{sh}} = 211$ to

219. The remaining 8 coefficients are set to zero. The existing binary check (every coefficient x satisfies $\langle x, x-1\rangle = 0$) covers all 512 coefficients without further equations.

2) *Linear equations.* Each Eq. (2) is a linear relation in $\mathbf{w}$ with public coefficients: the coefficients of $\boldsymbol{\theta}_{\varepsilon,j}$ on the slots of $\mathbf{k}$, $-q_{\mathsf{L}}2^{\iota}$ on carry bit ι, -2^{ι} on remainder bit ι, and the public constant $2^{11}\mathrm{PID}_{\varepsilon}[j]$. The verifier recomputes these coefficients from ε and $\mathrm{PID}_{\varepsilon}$.
3) *Soundness amplification.* The N_{pid} equations are folded into the random linear combination that the proof already uses to merge its constraints: the vector $\gamma_{i,\cdot}$ of the i-th amplification row grows from 262 to $262+N_{\mathsf{pid}} = 286$ entries. The quadratic part of the proof, the commitment dimensions and the number of repetitions $l_{\mathsf{sh}} = 7$ are unchanged.

Over $\mathbb{Z}_{q_{\mathsf{sh}}}$ ($q_{\mathsf{sh}} \approx 2^{57.7}$) the left-minus-right side of Eq. (2) is below $2^{29.1}$ in absolute value, so equality mod q_{sh} implies equality over $\mathbb{Z}$ once the witness is binary; the binary check is what makes the lifting sound (supplementary material).

d) Context binding: A 32-byte string $\mathrm{ctx}_{\varepsilon} = H_{\mathsf{ctx}}(pp\|\varepsilon\|\mathrm{PID}_{\varepsilon}\|\mathsf{pk}_{\varepsilon}\|\mathcal{A}_{\mathrm{pub}})$ is absorbed into every Fiat–Shamir challenge hash. Without it a presentation proof would not be tied to $\mathsf{pk}_{\varepsilon}$, and an attacker could replay a victim's proof under its own epoch key. (Our prototype binds ε, $\mathrm{PID}_{\varepsilon}$ and the epoch key $\mathsf{pk}_{\varepsilon}$ into $\mathrm{ctx}_{\varepsilon}$; $\mathcal{A}_{\mathrm{pub}}$ and the full public parameters are not bound; Section VII-A states the remaining limitation.)

E. Epoch Presentation

At the start of epoch ε the vehicle computes $\boldsymbol{\theta}_{\varepsilon,j}$ and $\mathrm{PID}_{\varepsilon}$, generates a fresh FN-DSA-512 key pair $(\mathsf{sk}_{\varepsilon}, \mathsf{pk}_{\varepsilon})$ that has no algebraic relation to $\mathbf{k}$ or usk, computes $\mathrm{ctx}_{\varepsilon}$, and produces $\pi_{\varepsilon}^{\mathsf{sh}}$, a proof of knowledge of a TA-signed credential on $(\mathrm{usk}, \mathbf{k}\|\mathcal{A}_{\mathrm{pub}})$ together with the N_{pid} relations (2). It broadcasts $\mathsf{Pres}_{\varepsilon} = (\varepsilon, \mathrm{PID}_{\varepsilon}, \mathsf{pk}_{\varepsilon}, \mathcal{A}_{\mathrm{pub}}, \pi_{\varepsilon}^{\mathsf{sh}})$ and, if the access layer allows it, hands the same object to roadside units for caching. A vehicle holding several credentials would obtain several pseudonyms; Section VI states exactly what is and is not enforced.

F. Message Authentication

Let $\mathsf{sid}_{\varepsilon} = \mathsf{trunc}_{128}\big(H_{\mathsf{sid}}(\varepsilon\|\mathrm{PID}_{\varepsilon}\|\mathsf{pk}_{\varepsilon})\big)$. A message with payload m and timestamp ts is sent as

$$(m,\ ts,\ \varepsilon,\ \mathsf{sid}_{\varepsilon},\ \sigma), \qquad \sigma = \mathsf{Sig}_{\mathsf{sk}_{\varepsilon}}(m\|ts\|\varepsilon\|\mathsf{sid}_{\varepsilon}). \tag{3}$$

A receiver looks $\mathsf{sid}_{\varepsilon}$ up in its cache; on a miss it asks a roadside unit for the presentation or waits for the next broadcast, and unknown identifiers are rate-limited. On first sight of $\mathsf{Pres}_{\varepsilon}$ it verifies $\pi_{\varepsilon}^{\mathsf{sh}}$ against pk_{TA}, recomputes $\mathsf{sid}_{\varepsilon}$, and caches $(\mathrm{PID}_{\varepsilon}, \mathsf{pk}_{\varepsilon})$; every later message costs one FN-DSA verification. This resembles the certificate-digest mechanism of the ETSI CAM profile (digest by default, full certificate once per second and on request [4]), with the presentation sent once per epoch instead. Receivers check the timestamp window; replay inside the window is not prevented, and there is no batch verification.

G. Conditional Tracing

Given ε, $\mathrm{PID}_{\varepsilon}$ and an external authorization, the TRA computes $\mathrm{PID}_{\varepsilon}[1]$ for every stored key, keeps the $\approx 2^{-8}$ that agree, checks the survivors on the remaining components and returns the RID of the unique match; each call is appended to a hash-chained log (procedure R2) and measured in Section VII-C. Holding $\mathbf{k}$, the TRA can recompute any vehicle's pseudonym in any epoch (Section VI-D).

H. Correctness and Pseudonym Collisions

Proposition 1 (Correctness)**.** *If the vehicle, TA and TRA follow the protocol, then registration succeeds, an honestly generated* $\mathsf{Pres}_{\varepsilon}$ *is accepted by every verifier, every message of form* (3) *verifies, and* Trace *returns the* RID *of the sender, except with the collision probability stated in the remark below.*

Proof: Eq. (2) holds by construction with $\mathsf{cy}_j = \lfloor\langle\boldsymbol{\theta}_{\varepsilon,j}, \mathbf{k}\rangle/q_{\mathsf{L}}\rfloor \le 767$ and $\mathsf{rm}_j = (\langle\boldsymbol{\theta}_{\varepsilon,j}, \mathbf{k}\rangle \bmod q_{\mathsf{L}}) \bmod 2^{11}$. Acceptance of the proof is the completeness of the presentation proof of [16] (Lemma 7.4 there), unchanged in structure because the added relations are satisfied by the honest witness and the amplification vector is longer but uniform; completeness of the extended statement was checked empirically ($20/20$ honest proofs per configuration). Signature verification is correctness of FN-DSA. The TRA's scan reproduces the sender's $\mathrm{PID}_{\varepsilon}$ from the stored $\mathbf{k}$ by Eq. (1). ∎

Remark 1 (Pseudonym collisions)**.** *Let* $N_{\mathsf{reg}} \le N_{\mathsf{veh}}$ *vehicles be registered and* N_{epc} *epochs be considered. The keys are drawn by the TRA, so the adversary cannot choose them: they are uniform, independent of one another and of the random oracle* H_{pid} *(redrawing on a duplicate changes their distribution by at most* $\binom{N_{\mathsf{reg}}}{2}2^{-768}$*, the probability of any collision). Two binary keys* $\mathbf{k}_1, \mathbf{k}_2$ *that are non-zero and distinct are linearly independent over* $\mathbb{F}_2$*, which fails with probability below* 2^{-766}*. For such keys and a uniform* $\boldsymbol{\theta} \in \mathbb{Z}_{q_{\mathsf{L}}}^{n_{\mathsf{k}}}$*, the map* $\boldsymbol{\theta} \mapsto (\langle\boldsymbol{\theta}, \mathbf{k}_1\rangle, \langle\boldsymbol{\theta}, \mathbf{k}_2\rangle) \bmod q_{\mathsf{L}}$ *is onto* $\mathbb{Z}_{q_{\mathsf{L}}}^2$*, because it is onto modulo* 2 *and* q_{L} *is a power of* 2*. The two inner products are then uniform and independent, and their top* 8 *bits agree with probability* 2^{-8}*. The* N_{pid} *vectors* $\boldsymbol{\theta}_{\varepsilon,j}$ *are independent, so two vehicles share a pseudonym in a given epoch with probability at most* $2^{-8N_{\mathsf{pid}}} + 2^{-766}$*. A union bound over pairs of vehicles and epochs gives*

$$\begin{aligned}\delta_{\mathrm{PID}} &= \binom{N_{\mathsf{reg}}}{2} N_{\mathsf{epc}}\big(2^{-8N_{\mathsf{pid}}} + 2^{-766}\big) \\ &\approx N_{\mathsf{reg}}^2\, N_{\mathsf{epc}}\, 2^{-8N_{\mathrm{pid}}-1}.\end{aligned} \tag{4}$$

For $N_{\mathsf{reg}} = 2^{30}$ *and* $N_{\mathsf{pid}} = 24$ *the bound above is* $2^{60-193} = 2^{-133}$ *per epoch, and over ten years of continuous operation* (87,600 *epochs)* $2^{-116.6}$*. The bound is statistical: it uses neither assumption A3 nor a limit on random-oracle queries. It rests on the keys not being adversarial. If a vehicle could choose its key after seeing the vectors* $\boldsymbol{\theta}_{\varepsilon,j}$*, finding a key that collides with a given pseudonym would be a subset-sum-type problem to which this argument gives no bound, which is why the TRA draws the keys.*

VI. Security Analysis

We analyze L-SUPA in the classical random oracle model. Each result carries one of three labels: *Reduction* (written out, every step identified), *conditional reduction* (written out, but conditional on a lemma or proposition whose extractor or simulator is not replayed), and *Sketch* (outlined, at least one step missing). Levels are classical Core-SVP estimates in the style of [11].

A. Assumptions

A1 *M-SIS* for the signature, presentation and issuance layers (Section VII-B).

A2 *M-LWE*, in the knapsack form used by [11], for the hiding of cmt and for the zero-knowledge property of π^{is} and $\pi_{\varepsilon}^{\mathsf{sh}}$.

A3 *Binary-secret LWR.* For $\mathbf{k} \xleftarrow{\$} \{0,1\}^{n_{\mathsf{k}}}$ and S independent vectors $\boldsymbol{\theta}_i \xleftarrow{\$} \mathbb{Z}_{q_{\mathsf{L}}}^{n_{\mathsf{k}}}$, the pairs $\big(\boldsymbol{\theta}_i, \lfloor(\langle\boldsymbol{\theta}_i, \mathbf{k}\rangle \bmod q_{\mathsf{L}})\, p_{\mathsf{L}}/q_{\mathsf{L}}\rfloor\big)_{i\le S}$ are computationally indistinguishable from $(\boldsymbol{\theta}_i, u_i)_{i\le S}$ with $u_i \xleftarrow{\$} \mathbb{Z}_{p_{\mathsf{L}}}$. We write $\mathsf{Adv}^{\mathrm{LWR}}(T, S)$ for the best distinguishing advantage in time T. This is a *direct* assumption: we claim no reduction from LWE. Related hardness results for LWR [19] and for LWE with binary secrets [20] motivate it but do not cover this parameter set.

A4 *Random oracle model* for $H_{\mathsf{pid}}, H_{\mathsf{sid}}, H_{\mathsf{ctx}}$ and the Fiat–Shamir hashes [18].

A5 *EUF-CMA security* of FN-DSA-512, taken as a black box [21].

A6 Second-preimage resistance of H_{sid} truncated to 128 bits, for targets chosen by honest vehicles (in the ROM, $q_{\mathsf{H}}2^{-128}$ for q_{H} queries), and collision resistance of H_{ctx}. Two presentations of the adversary with equal sid are harmless, because they only affect its own cache entries.

A7 *M-ISIS*: given D_s and upk it is hard to find a short usk with $D_s\mathrm{usk} = \mathrm{upk}$, as in [11].

The trust assumptions TM1–TM6 of Section IV (in particular that the TA is honest-but-curious) are used throughout.

Strength of A3: We ran the lattice estimator of [23] (v0.1.0, archived) on the 768-dimensional instance, modeling the rounding error as uniform noise on $[-1024, 1023]$ (a heuristic; a Gaussian of equal deviation gives the same result). Under Core-SVP [14] the weakest attack is dual-hybrid ($\beta = 707$): $2^{206.4}$ classical and $2^{187.4}$ quantum; usvp ($\beta = 750$) gives $2^{219.0}$. With the estimator's built-in cost model the weakest is again dual-hybrid, at $2^{227.6}$, and the others are no smaller (bdd/bdd-hybrid $2^{235.1}$, usvp $2^{239.8}$, dual $2^{252.8}$, bdd-mitm-hybrid $2^{268.6}$). Coded-BKW ($2^{262.0}$) and meet-in-the-middle (2^{624}) do not compete; Arora–Ge [24] does not apply (vanishing polynomial of degree 2048 in 768 variables, estimated infinite). The estimate is saturated in the sample count ($m = 2^{19}, 2^{20}, 2^{21}$ identical), so the ten-year maximum $S \le N_{\mathsf{pid}} \cdot 87{,}600 \approx 2^{21.0}$ is covered.

B. Goals and Results

Table II summarizes the results. Let $\mathcal{O}_{\mathsf{reg}}$, $\mathcal{O}_{\mathsf{cor}}$, $\mathcal{O}_{\mathsf{pres}}$, $\mathcal{O}_{\mathsf{msg}}$ and $\mathcal{O}_{\mathsf{H}}$ be the oracles for honest registration, corruption, honest presentation, honest signed messages, and the random oracles. N_{veh}, N_{epc}, N_{reg} and N_{pres} denote the numbers of vehicles, epochs per vehicle, registrations and presentations.

Lemma 1 (PID soundness; conditional)**.** *Let a prover make the verifier accept $\pi_\varepsilon^{\mathsf{sh}}$ on a statement containing ctx_ε and PID_ε with non-negligible probability. Then there is an extractor that, in expected polynomial time and except with probability δ_{ks}, outputs a witness with (i) a valid TA signature on $(\mathrm{usk}^*, \mathsf{att}^*)$, (ii) usk^*, $\mathsf{att}^* = (\mathbf{k}^* \| \mathcal{A}^*_{\mathrm{pub}})$ and all carry and remainder bits binary, (iii) Eq. (2) for every j, hence $\mathrm{PID}_\varepsilon[j] = \lfloor (\langle \boldsymbol{\theta}_{\varepsilon,j}, \mathbf{k}^* \rangle \bmod q_{\mathsf{L}}) p_{\mathsf{L}} / q_{\mathsf{L}} \rfloor$ for all j, and, in the variant that discloses attributes, (iv) the disclosed slots of $\mathcal{A}^*_{\mathrm{pub}}$ equal the public $\mathcal{A}_{\mathrm{pub}}$ of the statement (not implemented or measured, supplementary material); or it outputs an M-SIS solution. Here*

$$\delta_{\mathsf{ks}} = \tfrac{2}{|\mathcal{C}_{\mathsf{ch}}|} + q_{\min}^{-n_{\mathsf{sh}}/\kappa} + q_{\min}^{-l_{\mathsf{sh}}} + 2^{-128} + \delta_{\textit{M-SIS}} \approx 2^{-127.1}, \quad (5)$$

the soundness error of Lemma 7.5 of [16] evaluated for our parameters, with $q_{\mathsf{sh}} = q_{1,\mathsf{sh}} q_{\mathsf{sg}}$, $q_{\min} = \min(q_{1,\mathsf{sh}}, q_{\mathsf{sg}}) = q_{\mathsf{sg}} = 425801$, $\kappa = 4$ the number of splitting factors of the modulus, and $\mathcal{C}_{\mathsf{ch}}$ the challenge space. The fixed 2^{-128} term is verbatim from the soundness error of Lemma 7.5 of [16].

Proof sketch: The added relations are linear in the witness. The original extractor [11], [13] outputs a witness for every constant-term equation of the folded system; the binary check ($B^2 < q_{\mathsf{sh}}$, ratio 7×10^{-8}; [16, Lemma 7.5]) makes the coefficients binary, and the lifting of Section V-D turns Eq. (2) modulo q_{sh} into an integer equation. Uniqueness of the decomposition follows from the ranges of cy_j and rm_j, with $\mathrm{PID}_\varepsilon[j] \in [0, 255]$ by the one-byte encoding; the bits of $\mathbf{k}^*$ sit in the same embedding as in signature verification, so $\mathbf{k}^*$ is the signed key. The amplification error does not depend on the number of equations. The supplementary material lists the thirteen differing points; the extractor was not replayed. ∎

Proposition 2 (Extended zero knowledge; conditional)**.** *The presentation argument for the extended statement is zero-knowledge, extending [16, Lemma 7.6], with error $2^{-129.4}$ (rejection sampling) plus a singularity probability below 2^{-10^4} (Section VI-C); the simulator is not re-derived, so the proposition is conditional.*

Goal 1: authentication: In the game $\mathsf{Exp}^{\mathsf{auth}}$ the adversary controls corrupted vehicles, registers them with the honest TA and TRA, and wins by outputting an accepted presentation and message such that (W1) $\mathrm{PID}_{\varepsilon^*}$ belongs to no registered vehicle, or (W2) it equals the pseudonym of an honest vehicle h but was not output by $\mathcal{O}_{\mathsf{pres}}(h, \varepsilon^*)$, or (W3) the presentation is genuine but the message was never signed by $\mathcal{O}_{\mathsf{msg}}(h, \cdot)$.

Theorem 1 (Authentication)**.** *Here $\mathsf{Adv}^{\mathsf{SEP}}$ is the existential-unforgeability advantage against the signature with efficient protocols of [11] (Theorems 5.1 and 5.2 of [16]). Under A1–A7, Lemma 1, and the unforgeability of the signature scheme of [11],*

$$\mathsf{Adv}^{\mathsf{auth}} \le \mathsf{Adv}^{\mathsf{SEP}} + N_{\mathsf{veh}}\mathsf{Adv}^{\textit{M-ISIS}} + N_{\mathsf{pres}}\mathsf{Adv}^{\mathsf{ZK}}_{\mathsf{sh}} + \delta_{\mathsf{ks}} + \delta_{\mathrm{PID}} + N_{\mathsf{pres}}\mathsf{Adv}^{\mathsf{EUF}}_{\textit{FN-DSA}} + q_{\mathsf{H}}2^{-128},$$

In the W2 branch the reduction follows the impersonation-forgery proof of [16] (Section 7): the challenger simulates the presentation oracle and applies the knowledge extractor of Lemma 1 only to the adversary's own forgery, so the two never compose and no separate simulation-soundness lemma is needed; the branch is conditional on Lemma 1 (see Table II and Section VI-D).

Proof: W1: reduce to SEP unforgeability [11]: embed the SEP challenge key as pk_{TA}, answer registrations by querying the SEP signing oracle on $(\mathrm{upk}, \mathsf{att})$ (sampling $\mathrm{usk}, \mathbf{k}$, $\mathrm{upk} = D_s\,\mathrm{usk}$, $\mathsf{att} = (\mathbf{k} \| \mathcal{A}_{\mathrm{pub}})$), answer $\mathcal{O}_{\mathsf{pres}}$ with the zero-knowledge simulator, and run the extractor of Lemma 1 on the forgery (failing with probability δ_{ks}); a never-signed message and its signature form a SEP forgery ($\mathsf{Adv}^{\mathsf{SEP}}$), and a signed message contradicts W1. W2 (Section 7 of [16]): simulate $\mathcal{O}_{\mathsf{pres}}$ and apply the extractor only to the forgery, so the two never compose. If $\mathrm{usk}^* = \mathrm{usk}_h$, guess h (loss N_{veh}) and embed an M-ISIS instance in upk_h; the extracted usk_h solves $D_s\,\mathrm{usk}_h = \mathrm{upk}_h$. If $\mathrm{usk}^* = \mathrm{usk}_w$ for another registered w, then $\mathbf{k}_h \ne \mathbf{k}_w$ yield the same N_{pid}-byte pseudonym, with probability at most δ_{PID} (statistical, since the TRA draws the keys). Otherwise usk^* is unregistered, so its certified signature was never issued—a SEP forgery bounded by $\mathsf{Adv}^{\mathsf{SEP}}$. In all cases $\mathbf{k}^*$ is the key the TRA drew (it confirms only the issued key). The context ctx_ε prevents reusing a victim's proof under a new pk_ε (the prototype binds ε, PID_ε and pk_ε; Section VII-A). W3 is a black-box reduction to A5: guess the vehicle–epoch pair (loss N_{pres}), replace its fresh epoch key by the FN-DSA challenge key pk^* (sound because pk_ε has no algebraic relation to $\mathbf{k}$ or usk and is bound only through ctx_ε), answer the guessed pair's signature queries by the FN-DSA oracle and generate the other pairs' keys; a winning W3 output is a forgery under pk^* unless two presentations collide in sid_ε (A6), giving the last two terms of Theorem 1. The reductions W1 and W2 inherit their only unproven premise from Lemma 1, whose extractor is not replayed (Section VI-D). ∎

Goal 2: same-epoch uniqueness per credential: The adversary registers c corrupted vehicles and wins if it outputs more than c distinct accepted pseudonyms for one epoch. By Lemma 1 every accepted pseudonym is the deterministic image of a signed $\mathbf{k}^*$, so one credential gives one pseudonym per epoch (*Sketch*). The goal does not limit the number of credentials per identity: several identities give several pseudonyms.

Goal 3: cross-epoch unlinkability: In $\mathsf{Exp}^{\mathsf{unl}}$ the adversary may hold sk_{TA} and sees all registrations; it chooses two honest, uncorrupted vehicles i_0, i_1 with equal $\mathcal{A}_{\mathrm{pub}}$ and two epochs $\varepsilon_1^* \ne \varepsilon_2^*$, receives the presentations of both vehicles for ε_1^* and that of i_{cb} for ε_2^* (cb a random bit), and must guess cb; it may not query $\mathcal{O}_{\mathsf{pres}}$ or $\mathcal{O}_{\mathsf{msg}}$ on (i_0, ε_2^*) or (i_1, ε_2^*). Within one epoch the pseudonym is fixed, so the game is cross-epoch only.

Theorem 2 (Cross-epoch unlinkability)**.** *Under A2, A3, A4 and TM2 (honest-but-curious TA),*

$$\mathsf{Adv}^{\mathsf{unl}} \le N_{\mathsf{pres}}\mathsf{Adv}^{\mathsf{ZK}}_{\mathsf{sh}} + N_{\mathsf{reg}}\big(\mathsf{Adv}^{\mathsf{ZK}}_{\mathsf{is}} + \mathsf{Adv}^{\textit{M-LWE}}\big) + N_{\mathsf{veh}}\mathsf{Adv}^{\mathrm{LWR}}(T, S) + \mathsf{negl}. \quad (6)$$

Proof: We bound the adjacent-game gaps of hybrids $\mathsf{Gm}_0, \ldots, \mathsf{Gm}_{3+N_{\mathsf{veh}}}$, where Gm_0 is the real game and, successively: Gm_1 replaces every $\mathcal{O}_{\mathsf{pres}}$ answer with a proof from the simulator of Proposition 2; Gm_2 replaces every registration's issuance proof with the honest-verifier simulator of [16, Lemma 7.3]; Gm_3 replaces the commitment random part $\mathbf{r}_0 + A'\mathbf{r}_1$ with a uniform vector; and Gm_{3+t} ($t = 1, \ldots, N_{\mathsf{veh}}$) replaces vehicle t's pseudonyms with uniform values in all epochs.

Gm_0 *to* Gm_1. Replace the presentations one at a time: a distinguisher with advantage ε embeds the challenge proof of Proposition 2 as the next presentation and simulates the rest, so $|\mathsf{Adv}(\mathsf{Gm}_0) - \mathsf{Adv}(\mathsf{Gm}_1)| \le N_{\mathsf{pres}}\mathsf{Adv}^{\mathsf{ZK}}_{\mathsf{sh}}$.

Gm_1 *to* Gm_2. The same argument over the registrations with [16, Lemma 7.3] (honest-verifier holds by TM2): $|\mathsf{Adv}(\mathsf{Gm}_1) - \mathsf{Adv}(\mathsf{Gm}_2)| \le N_{\mathsf{reg}}\mathsf{Adv}^{\mathsf{ZK}}_{\mathsf{is}}$.

Gm_2 *to* Gm_3. Reduce to M-LWE in knapsack form (A2): on input $(A', u_1, \ldots, u_{N_{\mathsf{reg}}})$, for each registration sample $\mathrm{usk}_i, \mathsf{att}_i = (\mathbf{k}_i \| \mathcal{A}_{\mathrm{pub}})$ yourself, set $\mathrm{upk}_i = D_s\,\mathrm{usk}_i$ and $\mathrm{cmt}_i = u_i + \mathrm{upk}_i + D \cdot \mathsf{att}_i$. The TRA's opening check (Algorithm 1, step 8) cannot be executed—the challenger no longer knows a binary opening—but by TM5 its outcome is outside the adversary's view (the confidential vehicle–TRA channel is excluded; the TA-side transcript is unchanged), so the challenger proceeds as if the TRA had confirmed. The M-LWE secret is the commitment randomness, independent of $\mathbf{k}$, which the reduction knows and uses to compute every PID_ε. Genuine $u_i = \mathbf{r}_0 + A'\mathbf{r}_1$ give Gm_2, uniform u_i give Gm_3, so $|\mathsf{Adv}(\mathsf{Gm}_2) - \mathsf{Adv}(\mathsf{Gm}_3)| \le N_{\mathsf{reg}}\mathsf{Adv}^{\text{M-LWE}}$.

Gm_{3+t-1} *to* Gm_{3+t}. Reduce to A3 with $S \le N_{\mathsf{pid}}N_{\mathsf{epc}}$ samples: on input $(\boldsymbol{\theta}_i, y_i)_{i \le S}$, program H_{pid} so the public vectors of all epochs are the $\boldsymbol{\theta}_i$, set vehicle t's pseudonym components to the y_i, and simulate the others honestly (their $\mathbf{k}$ is known; in Gm_3 only the commitment randomness changed). The $\boldsymbol{\theta}_{\varepsilon,j}$ are public and common to all vehicles, so each step faces one secret on fixed public vectors. Genuine y_i give Gm_{3+t-1}, uniform y_i give Gm_{3+t}, so $|\mathsf{Adv}(\mathsf{Gm}_{3+t-1}) - \mathsf{Adv}(\mathsf{Gm}_{3+t})| \le \mathsf{Adv}^{\mathrm{LWR}}(T, S)$, and A2 and A3 are used one after another with no joint assumption on a shared secret.

Final game. In $\mathsf{Gm}_{3+N_{\mathsf{veh}}}$ all pseudonyms are uniform and all proofs are simulated, so the view is independent of cb: the only cb-dependent objects, the challenge presentations for ε^*_1 and ε^*_2, have simulated proofs (Gm_1) and uniform pseudonyms (Gm_{3+t}), and the adversary cannot query $\mathcal{O}_{\mathsf{pres}}$ or $\mathcal{O}_{\mathsf{msg}}$ on (i_0, ε^*_2) or (i_1, ε^*_2). Hence $\mathsf{Adv}(\mathsf{Gm}_{3+N_{\mathsf{veh}}}) = 0$, and summing over all adjacent pairs gives Eq. (6). ∎

Goal 4: conditional traceability: Correctness: from any accepted Pres^* the TRA's Trace returns the RID of a vehicle whose credential the adversary holds. Uniqueness is the collision bound δ_{PID} of Section V-H, which covers all registered vehicles because the TRA draws the keys. *Sketch*: extract $(\mathrm{usk}^*, \mathsf{att}^*)$ (Lemma 1); it was signed for a vehicle v. By the binding of cmt the key $\mathbf{k}_v$ stored by the TRA equals $\mathbf{k}^*$: a second opening would give a nonzero ternary relation $\Delta\mathbf{r}_0 + A'\Delta\mathbf{r}_1 + D_s\Delta\mathrm{usk} + D\Delta\mathsf{att} = 0$ over $R^4_{q_{\mathsf{sg}}}$ ($n = 256$, $q_{\mathsf{sg}} = 425801$); of the 26 coefficient vectors, 22 ($\Delta\mathbf{r}_1, \Delta\mathrm{usk}, \Delta\mathsf{att}$) are free ternary choices covering 5632 coefficients and the 1024 coefficients of $\Delta\mathbf{r}_0$ are then forced, so the expected number of relations is $3^{5632}(3/425801)^{1024} < 2^{-8598}$ and cmt is statistically binding with no computational assumption. Eq. (2) then makes the scan hit v. This step has no counterpart in [11], which has no tracer; the link between the extracted message and the unique binary opening of cmt is not written out.

Goal 5: separation of issuance and tracing: (a) *The TA does not learn* $\mathbf{k}$: its registration view is independent of $\mathbf{k}$ by Games Gm_1–Gm_3 (*Sketch*, with TM2). (b) *The TRA does not learn* usk: its view adds nothing about usk beyond $\mathrm{upk} = D_s\mathrm{usk}$, and impersonation by a TRA colluding with the TA still requires usk as a witness, hence reduces to inverting upk under M-ISIS (*Sketch*; we only have a first-order indicator of that hardness, *unverified*).

TABLE II
GOALS, ASSUMPTIONS AND STATUS.

Goal	Assumptions	Status
1a Authentication, original relation	A1, A2, A4, A7	Reduction (forgery branch W1, Theorem 1); Reduction (impersonation branch W2, Argo Section 7 template, conditional on Lemma 1)
1b Authentication, PID relation and ctx	A1–A4, A6, A7	Sketch (supplementary material)
1c Message authentication	A5, A6	Reduction (black box, loss N_{pres})
2 Same-epoch uniqueness per credential	A1–A3	Sketch
3 Cross-epoch unlinkability (equal $\mathcal{A}_{\mathrm{pub}}$)	A2, A3, A4, TM2	Reduction, conditional on Proposition 2; losses explicit
4 Conditional traceability	A1, A2, TM3	Sketch
5a The TA does not learn $\mathbf{k}$	A2, A4, TM2	Sketch
5b The TRA does not learn usk	A1, A2, A4, A7	Sketch (M-ISIS unverified)
Revocation, non-frameability	–	out of scope

TABLE III
MARGINS (BITS) OF THE PRESENTATION-PROOF AND LWR TERMS IN EQ. (6).

Scenario	N_{pres}	S	M-LWE$_{\mathsf{sh}}$	LWR$_{206.4}$	LWR$_{227.6}$
2^{30} veh., 10 y	46.4	21.0	79.6	176.4	197.6
2^{30} veh., 10 y, 2 h/day	42.8	17.4	83.2	176.4	197.6

C. Explicit Loss Parameters for Unlinkability

We have no tight reduction for the M-LWE steps of Theorem 2, so the generic hybrid multiplies the advantage by the number of presentations and we report both loss factors as model parameters. Table III gives the margin $b - \log_2(\text{factor})$ for a b-bit assumption (arithmetic only; "y" = years of continuous operation). The assumed hardness is 126 bits for presentation-layer M-LWE (our rerun of the parameter script, Section VII-B; the table of [11] lists 127 for the original configuration, one bit more) and 206.4 / 227.6 bits for LWR (Core-SVP [14], classical, and the estimator's built-in cost model). The other terms of the zero-knowledge bound of [16, Lemma 7.6] are the rejection-sampling loss $2^{-129.4}$ and a singularity probability below 2^{-10^4}; including them, the zero-knowledge term leaves 79.5 bits at 126 bits of M-LWE hardness (the supplementary material). The registration terms of Eq. (6) involve the issuance layer (127 bits) and $N_{\mathsf{reg}} \le N_{\mathsf{veh}} < N_{\mathsf{pres}}$, leaving at least as much margin as the presentation term, provided the commitment-hiding instance is not easier (not checked separately). *Under A2 and A3, the anonymity advantage is bounded by Eq.* (6)*; the LWR margin is at least* 176.4 *bits under Core-SVP (*197.6 *built-in), and the M-LWE margin is as listed in Table III.* We make no claim of "128-bit anonymity".

D. Scope, Trust and Open Gaps

1) *Extractor not replayed.* Lemma 1 adapts the soundness argument of [11], based on [13]. We checked how each change enters its premises and error terms (the supplementary material) but did not re-derive the constants; zero knowledge of the extended statement is argued only informally. The W2 branch of Theorem 1 follows the impersonation-forgery proof of [16] (Section 7): the extractor acts on the adversary's own forgery and the simulator on the oracle, so they never compose; the branch is conditional on Lemma 1.
2) *Knowledge-soundness error.* Eq. (5) is $2^{-127.1}$, about 0.9 bit weaker than the 128-bit target, mainly because of the fixed

2^{-128} term and $2/|\mathcal{C}_{\mathsf{ch}}| = 2^{-128.8}$; both belong to the original proof. The fixed term appears verbatim in Lemma 7.5 of [16], and Argo notes that its knowledge-soundness proofs follow the blueprint of Theorem B.7 of [13]; we re-evaluated all five terms numerically (the supplementary material) but did not re-derive that blueprint, and the treatment of the query-dependent loss of the Fiat–Shamir transform is not established here.

3) *Rewinding, QROM and UC.* The extractor needs two rewinds, so the analysis holds in the classical random oracle model only; the authors of [11] note that straight-line extraction, needed for UC, is not provided.
4) *TA honest-but-curious.* The registration proof is zero-knowledge against honest verifiers only [16, Lemma 7.3]; Goals 3 and 5(a) depend on this. Zero-knowledge against a malicious TA is not claimed.
5) *TRA compromise.* Whoever obtains $\mathsf{DB}_{\mathsf{TRA}}$ can link all recorded traffic of a vehicle, retroactively, to its RID; it cannot impersonate a vehicle (no usk), and a colluding TA can de-anonymize all traffic. R1 (fresh $\mathbf{k}$, confirm only the issued key, delete the randomness) and R2 (audit log, external authorization) limit what a compromise yields; threshold tracing is future work.
6) *Same-epoch linkability.* Within an epoch all messages of a vehicle share one pseudonym; T_{ep} trades overhead against exposure. $\mathcal{A}_{\mathrm{pub}}$ is a quasi-identifier whose range determines the anonymity set. The protocol discloses $\mathcal{A}_{\mathrm{pub}}$ in clear, but the prototype and all measured numbers correspond to a proof that discloses nothing ($|I| = 0$ in the parameter tables of [11]); the public-constant equations for disclosed slots (item (iv) of Lemma 1) are neither implemented nor measured. Disclosing attributes can only shrink the proof: disclosed witness elements become public constants and their equations become public checks, so the reported 81.4 KiB and 0.15 s upper-bound the disclosure variant.
7) *Stateful signer.* With a tag space of 2^{33} and $Q_{\mathsf{sg}} = 2^{32}$ signatures per key, ten years of 2^{30} vehicles re-keying yearly need about $2^{33.3}$ issuances, so issuing keys must rotate and verifiers must accept several pk_{TA}.
8) *Verification denial of service.* An invalid presentation costs as much to reject as a valid one costs to accept (Section VII-C); there is no cheap pre-filter beyond rate limiting.
9) *Not covered:* revocation, non-frameability (a malicious TRA draws the keys and can misattribute messages, for instance by issuing one key to two vehicles), and location privacy beyond pseudonym unlinkability.

VII. Implementation and Evaluation

A. Implementation and Methodology

We extend the C implementation of [11] on FLINT with (i) the pseudonym expansion H_{pid} (SHAKE256, 19-bit outputs), (ii) the $N_{\mathsf{pid}} = 24$ relations of Eq. (2) in prover and verifier, with $E = 8$ extra witness elements and a longer amplification vector ($262 \to 286$), and (iii) a 32-byte context absorbed into the Fiat–Shamir hash. Each iteration runs a full proof and verification; six negative tests (flipped $\mathrm{PID}_\varepsilon[1]$, flipped ctx, another epoch, a different claimed pseudonym under an honest witness, flipped carry and remainder bits) were all rejected ($18/18$), and all honest runs accepted ($20/20$); this is not a soundness proof, as the tampering classes are ours. We do not re-implement the four lattice schemes of Section II; the only baseline is the hypothetical certificate of Section VII-C.

a) Changes to the original code: The challenge-norm function of both layers reduced only the first six coefficients modulo $X^N + 1$ and used a logical OR where an addition was intended, so it always returned 0 or 1 and the rejection loop never fired. We fixed both and added a verifier-side check $\|c\|_1 \le \eta$ (we did not check whether the original prescribes it). The acceptance rate becomes 54.4% (presentation) and 56.5% (issuance), and a prover skipping the check is rejected in 45 of 100 runs, as predicted. The timing effect is below the drift; three archived rounds of 60 iterations show the fixed binary verifying 0.1–3.1% slower.

b) Limitations of the prototype: The prototype binds ε, PID_ε and the 897-byte FN-DSA epoch key pk_ε into ctx; the epoch key is generated per epoch in about 8–11 ms, outside the timed proof. $\mathcal{A}_{\mathrm{pub}}$ and the full public parameters are not bound. Proof timings are unchanged within machine drift: prove minimum 194–215 ms and verify mean 146–160 ms across warm- and cold-boot runs, versus 189–194/141–146 ms for the credential-key binding of the earlier prototype; the binding itself replaces a ≈ 160 KiB hash input with 897 B. Its proof hides every credential slot, so all numbers are for a presentation that discloses nothing. Proofs are not serialized: all sizes are *formula values* (entropy bounds for Gaussian coefficients, bit widths otherwise) and a real encoding will be larger. We built no TRA service, network stack or OBU port.

c) Platform: Timings: Intel Core i5-12600KF, Windows 11 with WSL2 (Ubuntu 24.04, gcc 13.3), one pinned core, no frequency locking; desktop-class, not automotive hardware. Proof and tracing binaries: gcc -O3 -mavx2; FN-DSA: PQClean Makefile defaults. We report minima for proving (rejection sampling; the minimum over 20 runs excludes the about 1.8 rejected attempts per proof on average) and means for verification. Configurations run in two interleaved rounds, drift up to about 8%; FN-DSA uses three rounds and medians, tracing scans three rounds of five passes and means.

B. Parameters and Proof Sizes

Security levels are classical Core-SVP bits including the reduction losses of [11] (from the parameter scripts of that work), and sizes are formula values. The three layers are: signature $(n, d, m, q) = (256, 4, 10, 425801)$; presentation $(64, 23, m_1{=}219, m_2{=}74, \approx 2^{57.7})$; issuance $(64, 20, 104, 58, \approx 2^{37.7})$; the last two use $l = 7$ repetitions. M-SIS levels are $207/179$, 131, 131 (the two signature entries are the two forgery types) and M-LWE 152, 126, 127; proof sizes are 6971 B, 81.44 KiB and 35.99 KiB. The presentation M-LWE level (126 bits) is slightly below 128. Our rerun of the parameter script reproduces our presentation and issuance M-LWE levels (126/127, BKZ blocks 378/379) but not the original table's 127/129 (382/386); we keep our lower values, which keep the margins of Section VI-C conservative. The challenge norms are $\rho = 8$ and $\eta = 93$. Adding the pseudonym relation changes the presentation proof from 79.58 to 81.44 KiB (+2.3%), only through the witness dimension $m_{1,\mathsf{sh}}$ ($211 \to 219$); its knowledge-soundness error is about $2^{-127.1}$ (Eq. (5)), or $2^{-128.3}$ without the fixed 2^{-128} term.

C. Computation Cost

On one core, the registration (issuance proof) takes 127.7–131.1 ms to prove (minimum of 20 runs) and 104.1–109.6 ms to verify (mean of 20 runs, ranges over two rounds); the epoch presentation takes 194.0–210.2 ms and 147.1–156.8 ms. The pseudonym relation raises presentation verification by 3.0–3.3% over the unmodified proof in the same session (measured before the challenge-norm fix; the archived logs reproduce the post-fix timings below but contain no build without the pseudonym relation, so this comparison is not reproduced there). The witness is 3.8% wider, and we have no run that separates the width from the equations, so the 3.3% upper bound covers the added equations and their witness bits together, not the equations alone. Presentation proving takes about 0.2 s and verification about 0.15 s.

a) Per-message signature: FN-DSA-512 (PQClean Falcon, 722-byte messages, medians of three rounds): signing 0.14 ms with AVX2 and 2.65 ms portable; verification 0.033 ms both; per-epoch key generation 4.2 or 8.4–8.9 ms. Every signature verified and every altered message rejected. Portable signing matters for OBUs without SIMD floats; verification is unaffected. Means have a long tail (single samples up to 0.84 ms with AVX2), hence medians. We measured a Falcon-specification implementation [21], which may differ from the final FN-DSA standard.

b) What a receiver pays: One presentation verification per sender and epoch costs about 0.15 s; every later message costs 0.033 ms. For 100 senders at 10 Hz, the benchmark rate of [3], this is about 3.3% of one core plus a presentation verification per new sender ($\approx 15\%$ more if one appears every second; arithmetic, no OBU measurement). An invalid presentation costs as much as a valid one, so receivers need a rate limit on unknown identifiers (Section VI-D).

c) Epoch-boundary burst and a certificate baseline: Epochs are synchronized, so all vehicles in range change pseudonym at once: with 100 senders, 100 presentation verifications (≈ 15 s of one core) and ≈ 8.4 MB to receive, and until verified, messages with an unknown sid_ε cannot be authenticated (Section V-F). Baseline: a hypothetical post-quantum certificate (FN-DSA key 897 B plus issuer signature 666 B, at least 1563 B; no standard defines one) verifying in 0.033 ms, so 100 senders cost ≈ 3.3 ms. Attaching the full certificate once per second as the ETSI profile does [4] gives ≈ 5.6 MB per hour and sender against 84.3 kB for one presentation (≈ 54 certificates), so L-SUPA transmits less for encounters longer than about a minute, while the certificate baseline is about 4500 times cheaper in verification work per new sender; whether that trade is acceptable depends on density and hardware, neither measured.

d) Tracing: The TRA scan computes $\mathrm{PID}_\varepsilon[1]$ for all keys and checks the rest for survivors. On 2^{20} keys (96 B each) one thread needs 21.3–21.8 ms with AVX2 gather (20.5 ns/key), 32–33 ms with table lookup (31.2 ns/key) and 0.51 s branchless (branchy loop 3.1 s; eight-key gather 45 ms). All variants agree with the branchless loop on every key, and a 1000-key sample with a big-integer recomputation. About 2^{-8} survive the first component (4051 of 2^{20}; one fixed-seed data set, shared by the three rounds) and take about 2.4 ms. A 2^{30}-key database (96 GiB) does not fit our 7.8 GiB machine; linear extrapolation gives 22–33.5 s per scan plus 2.5 s for survivors, and a 2^{24}-key run deviates from linearity by up to 8–23%, so the extrapolation is not tighter than that. Scanning is arithmetic-limited: reading 96 MiB takes 3.9 ns/key against 20.5 ns of computation (AVX2 read microbenchmark; WSL2 bandwidth effects not quantified). Reading from disk, authorization and logging were not measured.

D. Communication Cost and Frame Budget

Object sizes (formula values except the 4 B timestamp and epoch fields and the 16 B sid_ε): upk 2432 B, cmt 2432 B, π^{is} 35.99 KiB, the stored credential 6971 B, $\pi_\varepsilon^{\mathsf{sh}}$ 83,390 B, PID_ε 24 B, pk_ε 897 B, a presentation 84,311 B, and a message $666 + 16 + 4 + 4 = 690$ B plus payload. Registration moves about 47 kB once (four transfers; the 7.0 kB credential is not counted). The presentation is 84.3 kB, 106 frames of 800 payload bytes, a frame size that is our assumption: the CAM standard [25] sets none that we found, so a payload F gives $84{,}311/F$ frames (169 at 500 B).

a) Frame budget of a message: A safety message is $690 + \ell_m$ bytes: signature (666), sid_ε (16), timestamp and epoch (4 each), and payload ℓ_m. The conventional $\ell_m = 32$ gives 722 B, a convention of ours that no specification prescribes. Cooperative awareness messages carry such traffic; their standard fixes the interval but we found no byte limit in it [25], nor in the security profile [4]. Secondary reports, not read in full, give payloads of about 200–800 B (averages 300–406 B); we use them only as a range. For $\ell_m > 110$ the message exceeds 800 B, so a single frame needs short payloads or larger frames; we state the property as a conditional and do not claim it for real traffic. An ML-DSA signature of 2420 B (ML-DSA-44 [15]) would not fit an 800-byte budget.

E. Per-Epoch Amortization

The CAM interval lies between 10 and 1 Hz [25]; we use 10 Hz, most favorable for the ratio. At $T_{\mathsf{ep}} = 1$ h a vehicle sends 36,000 messages of 722 B (26.0 MB) plus one 84.3 kB presentation, 0.32% of the message bytes (for $\ell_m = 32/300/406/800$: $0.32/0.24/0.21/0.16\%$); at $T_{\mathsf{ep}} = 300$ s, $3.9/2.8/2.6/1.9\%$. Shorter epochs limit observation but raise the ratio; at 1 Hz all ratios are ten times larger (e.g., 3.2%).

The overhead is small relative to the message bytes, yet 84 kB per vehicle per epoch has to be broadcast or fetched from a roadside unit, which in a dense channel is a load that we did not simulate.

VIII. Discussion

A. Revocation

L-SUPA has no revocation or non-frameability: the TRA identifies a sender but nothing removes a credential, and short-lived credentials are limited by the signatures-per-key bound of [11] (Section VI-D).

B. Deployment Considerations

a) Epoch length: A one-hour epoch keeps the share of presentation traffic at 0.32% for our conventional message size, while a five-minute epoch would cost 3.9% (Section VII-E). The price is that a passive observer can link all messages of a vehicle within an hour. T_{ep} is therefore a parameter that trades privacy against overhead.

b) Distribution of presentations: A receiver needs the presentation before verifying a message; presentations are large, so they should repeat at a low rate or be served by RSUs. The design in Section V-F is a sketch we did not simulate; packing 84.3 kB into frames on a shared channel is the first deployment question, and the epoch-boundary burst and certificate comparison are in Section VII-C.

c) Hardware and key rotation: Presentation proving takes about 0.2 s and verifying about 0.15 s on a desktop core, with no automotive-processor or HSM measurements; the TA must rotate issuing keys and verifiers accept several pk_{TA} (Section VI-D).

C. Limitations

We collect the limitations that matter most for judging the claims of this paper.

1) *Security argument and authorities.* Authentication (Theorem 1) and unlinkability (Theorem 2) follow from reductions conditional on Lemma 1 and Proposition 2, whose extractor and simulator are not replayed; Goals 1b, 2, 4 and 5 rest on sketches, the analysis is classical ROM only, and the knowledge-soundness error is $2^{-127.1}$. The TA is assumed honest-but-curious, and a TRA compromise permits retroactive linking; details are in Section VI-D, the thirteen change points and open items in the supplementary material.
2) *Anonymity bound.* The bound of Eq. (6) multiplies the M-LWE advantage by the number of presentations; we found no attack achieving this loss and no tight reduction, so we make no claim of 128-bit anonymity. The binary-secret LWR assumption has no dedicated cryptanalysis we could find; our estimate ($2^{206.4}$ under Core-SVP, $2^{227.6}$ with the estimator's built-in cost model) uses a lattice estimator with a heuristic noise model, saturated in the sample count up to the ten-year maximum.
3) *Measurements.* One desktop machine under WSL2, one core, no frequency locking; no vehicle hardware; sizes are formula values; the 2^{30}-key tracing time is extrapolated; FN-DSA was measured on a Falcon-specification implementation, not a final standard; the 722-byte message and 800-byte frame are our conventions, unverified against the standards.
4) *Prototype.* The context string binds the epoch, the pseudonym and the epoch public key; the disclosed attribute and the full public parameters are not bound; the verifier-side challenge-norm check is our addition.

D. Future Work

Tracing without a single point of compromise (threshold or partitioned storage of $\mathbf{k}$, rotation with re-issuance); *Proofs* (a replayed extractor for the extended statement, a simulator for the added relations, a tight reduction for the anonymity loss); *Higher security levels* (the parameters target 128 bits only); *Evaluation* (vehicle-hardware measurements, serialization, network simulation of presentation distribution, a comparison with other lattice schemes under identical parameters).

IX. Conclusion

We described L-SUPA, a conditional-privacy authentication scheme for VANETs in which a vehicle derives its epoch pseudonym from a key signed into its credential, proves the derivation inside the presentation proof of a lattice anonymous credential system, and signs messages with a post-quantum signature. The pseudonym is a rounded linear function of a binary key; proving it costs 24 extra linear equations, 8 extra witness ring elements and about 3% more verification time. A tracing authority separate from the issuer recomputes pseudonyms from stored keys and identifies a sender in one scan.

Our analysis is explicit about its status: message authentication follows from a written-out black-box reduction to FN-DSA (Theorem 1), authentication of L-SUPA and cross-epoch unlinkability follow from reductions conditional on the PID-soundness lemma and the zero-knowledge proposition (Theorems 1 and 2), and the pseudonym relation, same-epoch uniqueness per credential, traceability and the issuer–tracer separation rest on sketches, with an anonymity margin of about 80 bits stated explicitly. The per-epoch presentation remains large (84 kB) and the tracing authority is a high-value target; future work on threshold tracing, a replayed proof, and vehicle-grade measurements addresses these. Concretely: a pseudonym collision bound of 2^{-133} per epoch for 2^{30} vehicles, message verification of 0.033 ms, and a single-thread tracing scan of about 22 s (AVX2; 33.5 s portable), extrapolated to 2^{30} stored keys; all timings come from one desktop core.

Appendix
Change-Point Check of the PID Equation

For every point in which the presentation proof of L-SUPA differs from the non-interactive presentation proof of [16], this appendix gives the premise or error term of Lemmas 7.4 (completeness), 7.5 (knowledge soundness) or 7.6 (zero knowledge) of that paper that the difference touches, and what we checked. Three status labels are used: *checked* (against the paper, the code, or arithmetic that can be rerun), *idea* (argued, not written out), and *gap*. We did *not* replay the extractor or the simulator of [11], whose proofs follow [13] (Theorem B.7 and Lemmas B.8–B.11 there). The arithmetic is scripted (22 self-checks). Table IV gives the result.

a) What remains open: (1) *The extractor is not replayed.* Lemma 7.5 of [16] refers to their Section 7.1, which follows the blueprint of [13]; we checked where our changes meet the premises and error terms, not the constants of the proof. (2) *Total error.* The sum is $2^{-127.1}$, about one bit short of the 128-bit target, because of the fixed 2^{-128} term and the challenge-space term; both belong to the original lemma: the fixed 2^{-128} term appears verbatim in Lemma 7.5 of [16], whose proof follows the blueprint of [13] (Theorem B.7); we have not replayed that proof. (3) *Fiat–Shamir query loss.* The statement of Lemma 7.5 contains no number of hash queries, whereas non-interactive knowledge soundness usually loses a factor depending on it; how [11] treats this is not checked. (4) *Zero knowledge* is argued informally (C12). (5) *Zero-knowledge terms.* Besides the M-LWE term, Lemma 7.6 of [16] contains the rejection-sampling loss $\sum_i \epsilon_i/M_i$ and the singularity term $2\delta_{q_{\min}}(m_2, m_2') + \epsilon_{\text{M-LWE}}/(1 - \delta_{q_{\min}}(m_2, m_2 - m_2'))$. With $\epsilon_i = 2^{-130}$ and $M_i = 2$ (the parameter script), the first is $2^{-129.42}$. With the bound $\delta_{q_{\min}}(a, b) \le b\kappa\, q_{\min}^{-(a-b+1)\hat{n}/\kappa}$ of the same proof (p. 60 there), both singularity probabilities are below 2^{-10^4}. So ϵ is $2^{-125.87}$ for 126 bits of M-LWE hardness and $2^{-126.75}$ for 127, and $N_{\mathsf{pres}}\epsilon$ for $N_{\mathsf{pres}} = 2^{46.4}$ leaves 79.5 and 80.4 bits. The rejection-sampling loss alone would leave 83.0 bits. This is arithmetic on the lemma as printed; the lemma itself is not re-derived. (6) *Disclosed attributes.* A presentation that discloses $\mathcal{A}_{\text{pub}}$ needs equations that fix the disclosed slots to public constants; they are not part of the prototype and not covered by C1–C13. (7) *Full* ctx: the prototype now binds ε, PID_ε and the 897-byte epoch key pk_ε; binding $\mathcal{A}_{\text{pub}}$ and the fixed 24-byte wire format of PID_ε are implementation requirements.

References

[1] M. Raya and J.-P. Hubaux, "Securing vehicular ad hoc networks," *Journal of Computer Security*, vol. 15, no. 1, pp. 39–68, 2007.

[2] B. Brecht, D. Therriault, A. Weimerskirch, W. Whyte, V. Kumar, T. Hehn, and R. Goudy, "A security credential management system for V2X communications," *IEEE Transactions on Intelligent Transportation Systems*, vol. 19, no. 12, pp. 3850–3871, Dec. 2018.

[3] E. R. Verheul, C. Hicks, and F. D. Garcia, "IFAL: Issue first activate later certificates for V2X," in *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*, 2019, pp. 279–293.

[4] ETSI, "Intelligent transport systems (ITS); security; security header and certificate formats; release 2," European Telecommunications Standards Institute, Sophia Antipolis, France, Technical Specification TS 103 097 V2.1.1, Oct. 2021. [Online]. Available: https://www.etsi.org/deliver/etsi_ts/103000_103099/103097/02.01.01_60/ts_103097v020101p.pdf

[5] IEEE, "IEEE standard for wireless access in vehicular environments—security services for application and management messages," IEEE, IEEE Standard IEEE Std 1609.2-2025, Oct. 2026, published 2026-10-05; supersedes IEEE Std 1609.2-2022.

[6] P. W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer," *SIAM Journal on Computing*, vol. 26, no. 5, pp. 1484–1509, Oct. 1997.

[7] M. Akil, L. Martucci, and J.-H. Hoepman, "Non-interactive privacy-preserving sybil-free authentication scheme in VANETs," in *Proceedings of the Inaugural International Symposium on Vehicle Security & Privacy (VehicleSec)*. San Diego, CA, USA: Internet Society, 2023.

[8] X. Zhou, Y. Xin, H. Dang, K. Jin, L. Shi, and M. Zheng, "A secure and efficient lattice-based conditional privacy-preserving authentication protocol for VANETs," *Computers, Materials & Continua*, 2026, published online 26 October 2026; volume and pages not yet assigned.

[9] G. Liu, H. Li, J. Le, N. Wang, Y. Liu, and T. Xiang, "LRCPA: Lattice-based robust and conditional privacy-preserving authentication for VANETs," *IEEE Transactions on Vehicular Technology*, vol. 74, no. 3, pp. 4698–4712, 2025.

[10] Q. Li, Y. Xia, X. Liu, X. Chen, Z. Liu, Z. Guo, Z. Liu, and L. Fang, "LETA: A lattice-based efficient and traceable privacy-preserving batch authentication scheme for vehicle platoon in VANETs," *IEEE Journal on Selected Areas in Communications*, vol. 44, pp. 2186–2200, 2026.

[11] S. Argo, T. Güneysu, C. Jeudy, G. Land, A. Roux-Langlois, and O. Sanders, "Practical post-quantum signatures for privacy," in *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2024, pp. 1523–1537.

[12] M. F. Esgin, R. Steinfeld, D. Liu, and S. Ruj, "Efficient hybrid exact/relaxed lattice proofs and applications to rounding and VRFs," in *Advances in Cryptology – CRYPTO 2023*, ser. Lecture Notes in Computer Science. Springer Nature Switzerland, 2023, pp. 484–517.

[13] V. Lyubashevsky, N. K. Nguyen, and M. Plançon, "Lattice-based zero-knowledge proofs and applications: Shorter, simpler, and more general," in *Advances in Cryptology – CRYPTO 2022*, ser. Lecture Notes in Computer Science. Springer Nature Switzerland, 2022, pp. 71–101.

[14] E. Alkim, L. Ducas, T. Pöppelmann, and P. Schwabe, "Post-quantum key exchange—a new hope," in *25th USENIX Security Symposium (USENIX Security 16)*. Austin, TX, USA: USENIX Association, Aug. 2016, pp. 327–343. [Online]. Available: https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/alkim

[15] National Institute of Standards and Technology, "Module-lattice-based digital signature standard," NIST, Federal Information Processing Standards Publication FIPS 204, Aug. 2024.

[16] S. Argo, T. Güneysu, C. Jeudy, G. Land, A. Roux-Langlois, and O. Sanders, "Practical post-quantum signatures for privacy," IACR Cryptol. ePrint Arch., Paper 2024/131, 2024, full version of the CCS 2024 paper; last revised 2025-08-27. [Online]. Available: https://eprint.iacr.org/2024/131

TABLE IV
DIFFERENCES FROM THE PROOF OF ARGO ET AL. AND THEIR STATUS.

	Change	Enters in	Check and result	Status
C1	Witness grows by $E = 8$ binary ring elements, $m_{1,\mathsf{sh}}$: $211 \to 219$	premises on q of Lemma 7.5; M-SIS dimension	norm bound $B^2/q = 7.1 \times 10^{-8}$; second premise $1.12 \to 1.16 \times 10^{-3}$; third premise 0.9569, independent of m_1 (a 4.3% margin that is a property of the original parameters); M-SIS $132 \to 131$, M-LWE 126 unchanged (our rerun of the original configuration also gives 126; the paper lists 127)	checked
C2	Each amplification row $\gamma_{i,\cdot}$ grows from 262 to 286 entries	error term $q_{\min}^{-l}$	$l_{\mathsf{sh}} = 7$, $q_{\min} = 425801$: $q_{\min}^{-l_{\mathsf{sh}}} = 2^{-130.9}$. The bound comes from a random linear combination vanishing, at most $1/q_{\min}$ per row, and does not depend on the number of columns	numbers checked; independence from the number of columns is the same argument as in the original and was not replayed
C3	New equations are linear in the witness, not part of the quadratic matrix	completeness (Lemma 7.4), garbage terms e_0, e_1	code difference only in the combination of γ terms and the constant term; the quadratic part is untouched; honest proofs accepted in $20/20$ runs per configuration	structure checked; completeness is empirical
C4	Binary check covers the new witness	extraction, binary-ness	the sum of the binary check runs over all indices from the secret key to m_1 in prover and verifier; $B^2/q \ll 1$, hence the check lifts to $\mathbb{Z}$	checked
C5	New witness does not appear in the signature verification equation	extracted witness satisfies the signature relation	coefficients of the new witness in that equation are zero; the bits of $\mathbf{k}$ use the same embedding as the message slots, so the extracted $\mathbf{k}^*$ is the signed one	checked (code)
C6	Equation holds mod q_{sh}, lifted to $\mathbb{Z}$	Lemma 1, item (iii)	after C4, both sides are in $[0, 2^{29}]$; difference $< 2^{29.1} \ll q_{\mathsf{sh}}/2 = 2^{56.7}$	checked
C7	Decomposition $x = 2^{19}\mathsf{cy} + 2^{11}\mathrm{PID} + \mathsf{rm}$ is unique	item (iii)	$\mathsf{cy} \le 767$, $\mathsf{rm} \le 2047$, $2^{11} \cdot 255 + 2047 = 2^{19} - 1$; $\mathrm{PID} \in [0, 255]$ is guaranteed by the wire format (24 bytes)	checked
C8	ctx absorbed into the Fiat–Shamir hash	random-oracle input contains the statement	same position in prover and verifier; tampering with one bit of ctx or another epoch is rejected. The prototype absorbs ε, PID_ε and the 897-byte epoch key pk_ε	mechanism checked for ε, PID_ε and the epoch key; binding $\mathcal{A}_{\mathrm{pub}}$ and the full parameters is an implementation gap
C9	Challenge-norm function repaired and enforced by the verifier	$2/\lvert\mathcal{C}_{\mathsf{ch}}\rvert$ and the norm bound of the extractor	acceptance 54.4% at $\eta = 93$; $2/\lvert\mathcal{C}_{\mathsf{ch}}\rvert = 2^{-128.8}$ with the parameter script's challenge-space size (without the half for the norm bound: $2^{-129.8}$)	checked; whether the original paper prescribes the verifier-side check is not checked
C10	Uniform sampler rewritten (8 bytes per sample, mask, rejection)	uniformity of γ and CRS	mask $2^{58} - 1$ with rejection is uniform on $[0, q_{\mathsf{sh}})$ (acceptance 0.81)	checked
C11	Parameters regenerated with the extra witness	$\sigma_1, \sigma_3, B_{256}$	header constants equal the script output with the extra 504 squared-norm units	checked
C12	Zero knowledge (Lemma 7.6)	$m_2 > m_2' = \hat{d} + 256/\hat{n} + l + 1$ with $\hat{d} = d_{\mathsf{sh}}$, $\hat{n} = n_{\mathsf{sh}}$; all terms of ϵ	$74 > 35$; the M-LWE instance depends on $m_2, \hat{d}, l$ only, so 126 bits unchanged; the new terms are public, so the simulator can program h_i as before	idea (premises and error terms checked)
C13	Total error	δ of Lemma 7.5	Eq. (5): $2^{-128.8}$, $2^{-299.2}$, $2^{-130.9}$, 2^{-128}, 2^{-131}, total $2^{-127.13}$ (original configuration: $2^{-127.18}$); without the fixed term $2^{-128.26}$ ($2^{-128.37}$ in the original parameter table). The pseudonym relation changes only the M-SIS term	checked (arithmetic)

[17] A. Fiat and A. Shamir, "How to prove yourself: Practical solutions to identification and signature problems," in *Advances in Cryptology — CRYPTO '86*, ser. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 1987, pp. 186–194, presented at CRYPTO '86; proceedings published 1987.

[18] M. Bellare and P. Rogaway, "Random oracles are practical: A paradigm for designing efficient protocols," in *Proceedings of the 1st ACM Conference on Computer and Communications Security*. ACM, 1993, pp. 62–73.

[19] A. Banerjee, C. Peikert, and A. Rosen, "Pseudorandom functions and lattices," in *Advances in Cryptology – EUROCRYPT 2012*, ser. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2012, pp. 719–737.

[20] Z. Brakerski, A. Langlois, C. Peikert, O. Regev, and D. Stehlé, "Classical hardness of learning with errors," in *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing*. ACM, 2013, pp. 575–584.

[21] P.-A. Fouque, J. Hoffstein, P. Kirchner, V. Lyubashevsky, T. Pornin, T. Prest, T. Ricosset, G. Seiler, W. Whyte, and Z. Zhang, "Falcon: Fast-fourier lattice-based compact signatures over NTRU," Specification v1.2, submission to the NIST Post-Quantum Cryptography standardization project, 2020, cover page dated 01/10/2020. [Online]. Available: https://falcon-sign.info/falcon.pdf

[22] National Institute of Standards and Technology, "Post-quantum cryptography," NIST Computer Security Resource Center, https://csrc.nist.gov/projects/post-quantum-cryptography, 2026, page last updated 2026-10-27; accessed 2026-11-16.

[23] M. R. Albrecht, R. Player, and S. Scott, "On the concrete hardness of learning with errors," *Journal of Mathematical Cryptology*, vol. 9, no. 3, pp. 169–203, Oct. 2015.

[24] S. Arora and R. Ge, "New algorithms for learning in presence of errors," in *Automata, Languages and Programming (ICALP 2011)*, ser. Lecture Notes in Computer Science, vol. 6755. Springer, 2011, pp. 403–415.

[25] ETSI, "Intelligent transport systems (ITS); vehicular communications;

basic set of applications; part 2: Specification of cooperative awareness basic service," European Telecommunications Standards Institute, Sophia Antipolis, France, European Standard EN 302 637-2 V1.4.1, Apr. 2019. [Online]. Available: https://www.etsi.org/deliver/etsi_en/302600_302699/30263702/01.04.01_60/en_30263702v010401p.pdf