A Lightweight Post-Quantum Authentication Framework for 5G Base Station Bootstrapping

Abstract*—The absence of authenticated bootstrapping between User Equipments (UEs) and Base Stations (BSs) in 5G leaves System Information Block (SIB) broadcasts unprotected, enabling fake BS attacks, man-in-the-middle interception, and spoofed emergency alerts. Prior efforts such as Public Key Infrastructure (PKI)-based certificate chains, token-based schemes, and identity-based signatures either impose overhead exceeding 5G's strict packet-size constraints or lack post-quantum (PQ) security. Direct NIST-PQC integration is also infeasible; current standards, such as ML-DSA requires 34 fragmented SIB1 packets and up to 5,282 ms end-to-end delay, and FN-DSA still requires 13 fragments and up to 1,920 ms. We propose EMULSION, a symmetric chained publicly verifiable authentication framework for 5G/6G BS broadcast authentication. EMULSION is the first framework to exploit native 5G architectural features: fixed SIB transmission windows, millisecond-level time synchronization, and eSIM/USIM credential management to achieve genuine PQ security at symmetric-key efficiency. It uses a timed stream loss-tolerant authentication with HMAC chain anchored by a compact PQ signature (MAYO) applied once per epoch, fitting authentication within a single packet with no fragmentation and eliminating certificate transmission entirely. Unlike prior schemes, EMULSION extends to the full SIB family (SIB1-SIB21) at no additional per-broadcast cost, demonstrated over the air on SIB1. Evaluated on a real over-the-air 5G testbed, EMULSION achieves 33x lower end-to-end delay and 31x less communication overhead than ML-DSA, and 12x lower delay and 5.4x less overhead than FN-DSA. We formally prove the security of EMULSION and open-source its implementation for public testing and adaptation.

I. Introduction

Fifth-generation (5G) cellular networks have become the backbone of modern communication infrastructure, with over 3.5 billion subscribers worldwide and projections exceeding 6 billion by 2030 [1]. Before a User Equipment (UE) can place a call, send a message, or initiate any service, it must first acquire System Information Blocks (SIBs), i.e., broadcast messages that carry essential parameters for cell access, resource allocation, and mobility. SIB1 in particular is the bootstrapping message: it is broadcast every 160 ms on the Downlink Shared Channel (DL-SCH) and must be consumed before any authenticated channel (e.g., RRC, NAS) exists [2]. This unauthenticated bootstrap creates a critical window of vulnerability. A Fake Base Station (FBS) can forge or replay SIB1 to mount a family of downstream attacks, including denial-of-service, location tracking, bidding-down, and man-in-the-middle; all rooted in the absence of broadcast-layer authentication [3], [4], [5], [6].

Beyond these threats, the long-term security of 5G authentication is fundamentally at risk from quantum-capable adversaries. The classical public-key primitives supporting virtually all deployed cellular security (e.g., Elliptic Curve Cryptography (ECC)) are rendered insecure by Shor's algorithm on a sufficiently large quantum computer [7]. Unlike TLS 1.3 or PQ-WireGuard, where large Post-Quantum Cryptography (PQC) keys and certificates can be accommodated in a flexible handshake, SIB messages are constrained broadcast frames with fixed, sub-kilobyte size limits, no interactive exchange, and strict timing windows, making a direct PQC integration far from trivial [8]. Migration guidelines from NIST, ETSI, IEEE, and NSA/CISA jointly emphasize that the PQ transition must begin immediately [9], [10], [11], yet no existing 5G bootstrapping mechanism satisfies genuine PQ security under the strict protocol constraints. Securing the 5G SIB broadcast channel against both classical FBS attacks and future quantum-capable adversaries is thus a pressing open problem.

Prior Research. Although prior works have explored SIB authentication and PQ migration for cellular networks, each exhibits at least one of the following limitations: **(A)** Existing approaches for SIB1 authentication rely on traditional signature schemes that do not provide PQ security [12], [13]. **(B)** Direct substitution of classical signatures with NIST-PQC schemes (ML-DSA, FN-DSA) results in signature and public key sizes that vastly exceed the 372-byte size of a single SIB1 transport block, requiring extensive fragmentation across 12–34 packets and incurring multi-second end-to-end delays [8]. **(C)** PQ-AKA proposals [14], [15] target the mutual authentication phase after initial cell access and are orthogonal to the broadcast authentication problem; they do not protect the unauthenticated SIB window. **(D)** Hybrid PQ-KEM approaches [16] aim to reduce key exchange overhead but are not designed for the one-to-many broadcast setting of SIBs, where the base station has no per-UE state.

Thus, none of these approaches simultaneously achieve PQ security, single-packet compactness, and operation without any pre-existing authenticated channel—the three properties required for practical 5G broadcast authentication.

Problem. With the above limitations in mind, in this paper we aim to answer the following research question: *Is it possible to design a post-quantum secure authentication scheme for*

Network and Distributed System Security (NDSS) Symposium 2027
22–26 March 2027, Seoul, Republic of Korea
ISBN 978-1-970672-09-1
https://dx.doi.org/10.14722/ndss.2027.[23|24]xxxx
www.ndss-symposium.org

5G SIB broadcasts that doesn't incur heavy fragmentation, demands no pre-existing authenticated channel, and operates within the real-time constraints of the broadcast periodicity?

A promising Direction. Unlike generic broadcast authentication scenarios, the 5G base station architecture possesses several unique structural features that, when carefully exploited, make efficient PQ broadcast authentication feasible. First, SIBs are transmitted on fixed, periodic windows (e.g., SIB1 every 160 ms), defining natural authentication epochs. Second, modern UEs are equipped with GNSS receivers that provide wall-clock time independently of the cellular base station, and 5G gNBs are themselves GNSS-synchronized for network timing [17]; this gives both endpoints a shared, sub-microsecond time reference that no FBS adversary can manipulate through forged broadcast messages alone. Third, the eSIM/USIM ecosystem already supports secure remote provisioning of long-term cryptographic material into tamper-resistant hardware [18], [19]. These three features can circumvent the fundamental barriers that have stalled prior approaches.

Challenges. Realizing this direction requires overcoming three concrete technical challenges. *(C1) PQ signature overhead.* Even the most compact PQ signature candidates (e.g., MAYO-2 at 180 B) produce signatures that, when combined with the full certificate chain required for bootstrapping trust, far exceed the 372-byte SIB1 packet size. Naively attaching a PQ signature to each SIB broadcast is therefore infeasible. *(C2) Bootstrap without an authenticated channel.* Classical TESLA requires an authenticated delivery of the initial key commitment K_0. In the SIB authentication setting, no such channel exists: the UE has not yet completed RRC setup or NAS registration. Any solution must be self-bootstrapping from the first broadcast packet, relying only on material that can be pre-provisioned before the UE encounters the cell. *(C3) Real-time and loss-tolerance constraints.* The authentication mechanism must complete within the 160 ms SIB1 periodicity to avoid delaying cell access, must tolerate packet loss inherent in wireless broadcast (especially at cell edges), and must not impose significant computational burden on resource-constrained UEs—all while providing PQ security guarantees.

Our Approach. In this paper, we introduce EMULSION, a post-quantum broadcast authentication scheme for 5G that exploits the architectural features identified above. EMULSION employs a TESLA-style one-way key chain built on HMAC-SHA256 to authenticate SIBs within each 160 ms broadcast epoch using only symmetric primitives. The asymmetric cost is concentrated in a single epoch-opening packet that carries a compact PQ signature (MAYO-2 [20]), anchoring the HMAC chain's root of trust. The long-lived MAYO public key PK_{AMF} is pre-provisioned into the UE's eSIM/USIM, eliminating all certificate chains from the over-the-air path. This design makes the AMF (Access and Mobility Management Function) the sole entity performing asymmetric operations as the core key generator; the BS requires no asymmetric cryptography at runtime.

Findings. We implement EMULSION and a comprehensive set of baselines—NIST-PQC standards (ML-DSA, FN-DSA, SLH-DSA) and classical EC-Schnorr—on a full over-the-air 5G testbed (srsRAN, Open5GS, USRP B210). Our evaluation yields the following key results: (1) EMULSION achieves an end-to-end authentication delay of **160.07 ms** versus 1,920 ms for FN-DSA and 5,282 ms for ML-DSA; (2) it requires only **324 bytes** of cryptographic overhead, fitting within a single SIB1 transport block with zero fragmentation; (3) at 10% packet loss, EMULSION's expected retransmission overhead is $\sim$**41 bytes** versus $\sim$1,971 bytes for FN-DSA ($\sim 48\times$ reduction); and (4) demonstrated on SIB1, EMULSION extends to the complete SIB family (SIB1 through SIB21) at no additional per-broadcast cost.

A. Our Contributions

Our key contributions in this paper are as follows:

Key Observations and Main Idea. Unlike generic broadcast authentication scenarios, the 5G base station authentication protocol possesses unique architectural features that, when exploited, can enable efficient authentication mechanisms:

(i) *Fixed SIB transmission intervals.* SIBs are broadcast on fixed periodic windows, e.g., SIB1 every 160 ms with repeated transmissions as short as 20 ms [2]. In 6G, even tighter intervals are anticipated to support sub-millisecond latency targets [21]. These fixed windows define natural, predictable authentication epochs that a time-aware scheme can exploit without any protocol modification. Importantly, the same periodic broadcast structure applies to all SIB types: SIB2 through SIB9 are also transmitted on DL-SCH within transport blocks of similar size, either periodically broadcast or delivered on-demand via RRC signaling [2]. An authentication mechanism that operates at the transport block level, therefore generalizes across the entire SIB family without per-type adaptation.

(ii) *Widely available time synchronization.* TESLA-style protocols require loose time synchronization between sender and receiver. Modern UEs are equipped with GNSS receivers that provide wall-clock time with sub-microsecond accuracy, independent of the cellular base station's broadcast channel. This independence is critical: the System Frame Number (SFN) carried in the MIB on PBCH is unauthenticated and thus susceptible to manipulation by an FBS adversary, whereas GNSS-derived time originates from an external satellite constellation outside the adversary's control. Since GNSS receivers are standard in smartphones and increasingly prevalent in IoT UEs, and since 5G gNBs are themselves GNSS-synchronized for network timing [17], both endpoints share a common, fine-grained time reference that is sufficient to enforce TESLA's safe-packet test without relying on any unauthenticated over-the-air timing signal.

(iii) *eSIM/USIM credential management.* The eSIM infrastructure already supports secure remote provisioning of long-term cryptographic material into tamper-resistant USIMs [18], [19], the same storage used for Authentication and Key Agreement (AKA) provisioning. This provides a natural anchor for long-lived root key material, removing

the need to transmit any certificate or asymmetric public key during SIB authentication.

By exploiting these three features, we propose EMULSION, which harnesses a TESLA-style variant–*Timed Efficient Stream Loss-tolerant Authentication* [22]–relying solely on efficient symmetric primitives (i.e., HMAC) within each fixed SIB transmission window, made viable by 5G's existing precision time synchronization. The HMAC chain is anchored by a compact, PQ digital signature (e.g., MAYO [20], advanced to Round 3 of NIST's additional signature process) applied once at epoch boundaries rather than per-SIB, providing a root of trust without per-broadcast asymmetric overhead. The eSIM/USIM infrastructure securely provisions this root key, eliminating certificate chains entirely. Table I summarizes the performance of EMULSION against representative baselines; Section §VI provides evaluation details. The desirable properties of EMULSION are as follows:

TABLE I: Comparison of authentication schemes for 5G SIB1 bootstrapping.

Scheme	E2E (ms)	Crypto. (B)	Total OTA (B)	PQ Secure	Pkt. Loss Resilient	No Frag. Needed
EC-Schnorr [23]	4.15	256	372	✗	✓	✓
FN-DSA [24]	1920.67	3792	4836	✓	✗	✗
ML-DSA [25]	5282.47	9884	12648	✓	✗	✗
EMULSION ($\|\mathcal{C}\|$=2)	**160.07**	**324**	**744**	✓	✓	✓

Note: E2E delay in milliseconds; overhead in bytes. EMULSION with PK_{MAYO} pre-provisioned in eSIM. PQ counterparts are with 2-level certificates. See Section VI for details.

- **Symmetric-Optimal Computation and Post-Quantum Efficiency.** Per-SIB authentication in EMULSION is driven by HMAC, delivering PQ security and near-optimal computational efficiency simultaneously. The PQ anchor signature (MAYO) is computed only at epoch boundaries, amortizing its cost across all SIBs within the window. EMULSION is $\sim \mathbf{215}\times$ faster than FN-DSA at the BS and $\sim \mathbf{5}\times$ faster at the UE, with a total end-to-end delay of $\mathbf{160.07}$ ms versus 5,282.47 ms for ML-DSA. This translates directly to lower network delay and energy savings for resource-constrained UEs.
- **High Robustness and Security.** (*a*) *Packet loss resilience* compact per-SIB tags yield significantly higher packet loss tolerance than fragmented PQC schemes, with expected retransmission overhead of only $\sim$ 41 bytes at 10% loss versus $\sim$ 1971 bytes for FN-DSA. Beyond tag compactness, TESLA's one-way key chain ensures a UE can authenticate any buffered packet upon receiving its disclosed key, with no retransmission required regardless of intermediate losses. (*b*) *Error-Correction Compactness*: The redundancy size in forward error correction grows linear with the message size. Since EMULSION has smaller tag sizes, it achieves the same level of error correction capability with a lesser redundancy compared to asymmetric alternatives (e.g., $\sim$14 versus $\sim$480 redundancy bytes at 10% loss for EMULSION versus FN-DSA under Reed-Solomon coding). (*c*) *Full SIB coverage*: demonstrated on SIB1, EMULSION extends to the complete SIB family at no additional fragmentation or per-broadcast cost, closing attack surfaces on SIB2–SIB21 that prior schemes leave open.
- **Comprehensive Analysis and Open-Source Evaluation.** We implemented EMULSION and also deployed a full set of baselines–NIST-PQC standards (ML-DSA, FN-DSA, SLH-DSA) and classical schemes (EC-Schnorr)–on an over-the-air 5G testbed (srsRAN, Open5GS), covering cryptographic overhead, communication size, end-to-end delay, time synchronization accuracy, and packet loss behavior. The source code and testbed implementation are publicly released at: https://github.com/Prometheused/EMULSION

II. Background

λ denotes the security parameter (chain key length, $\lambda = 256$). t denotes the truncated MAC tag length ($t = 128$). q_v denotes the number of verification attempts available to an adversary within a disclosure window.

Notations: The symbol $||$ denotes concatenation. $x \xleftarrow{\$} S$ denotes uniform random sampling from set S. λ denotes the security parameter. F and F' denote pseudorandom functions (PRFs) used in TESLA key-chain derivation and MAC-key derivation, respectively. H denotes a collision-resistant hash function (instantiated with SHA-256). $\lfloor \cdot \rfloor$ denotes the floor function. A *time interval* of index i spans the half-open window $[T_0 + i \cdot T_{\text{int}},\ T_0 + (i+1) \cdot T_{\text{int}})$, where T_0 is the epoch start time and T_{int} is the fixed interval duration. Δt denotes the maximum clock-synchronization error between the sender (BS) and the receiver (UE).

A. 5G Cellular Network Architecture

1) Network Entities: A 5G cellular network comprises three principal entities [26]:

- ***5G Core Network (5GC).*** The 5GC orchestrates service delivery, session management, policy enforcement, and subscriber authentication. Most relevant to this work is the Access and Mobility Management Function (AMF), which terminates Non-Access Stratum (NAS) signaling from UEs and manages registration, connection, and mobility procedures.
- ***User Equipment (UE).*** A UE is any subscriber device at the network edge (e.g., smartphone, IoT sensor, or vehicle), provisioned with a Universal Subscriber Identity Module (USIM) storing a permanent identifier and the cryptographic credentials required for mutual authentication with the core network.
- ***Radio Access Network (RAN).*** The RAN consists of base stations (gNBs, denoted BS) and associated UEs, governing radio resource allocation, wireless transmission, and initial access. Critically, the system information messages a BS periodically broadcasts are transmitted in the clear, neither encrypted nor signed, rendering them susceptible to forgery, replay, and manipulation.

2) Protocol Stack and Initial Access: Fig. 1 depicts the 5G initial access sequence and protocol layers. At the BS, the topmost control-plane layer is Radio Resource Control (RRC); at the UE and AMF, the NAS layer is stacked above RRC. A master public key PK_{ID_0} is securely embedded

in the USIM and is publicly verifiable; private keys for the AMF and individual BSs are derived from the master secret key sk_{ID_0} and distributed through authenticated out-of-band channels. During initial access, the BS broadcasts System Information (SI), an RRC-layer message announcing cell-specific configuration parameters, comprising the Master Information Block (MIB) and one or more System Information Blocks (SIBs). The MIB is carried on the Physical Broadcast Channel (PBCH) and supplies the scheduling and decoding parameters needed to locate SIB1, which is transmitted on the Downlink Shared Channel (DL-SCH) with a configurable periodicity of 160 ms and a maximum size of 372 bytes per 3GPP specifications [27].

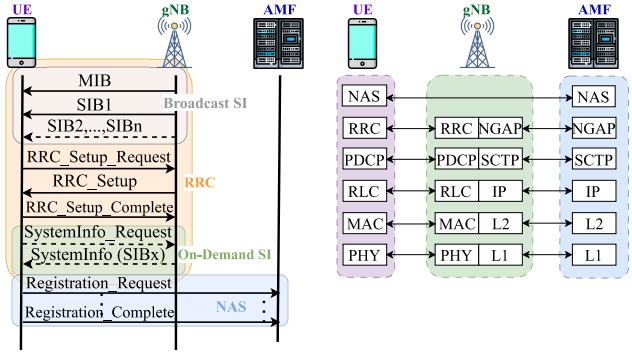

Fig. 1: 5G network connection setup and protocol stack.

The UE decodes the MIB from the PBCH to obtain the System Frame Number (SFN), which provides the timing reference needed to locate SIB1 on the DL-SCH. SIB1 carries the core cell configuration including PLMN identity, cell barring status, and scheduling for additional SIBs; mandatory fields include Cell Selection Info (signal quality metrics) and Cell Access Related Info (PLMN identifiers and cell access status), with optional fields such as IMS-Emergency Support included when applicable [27]. Upon decoding SIB1, the UE initiates the RACH procedure and RRC connection setup, after which NAS registration with the AMF occurs. This ordering is security-critical: *SIB1 is consumed before any authenticated channel exists*, since RRC security activation and NAS authentication both occur strictly after the UE has already acted on the received system information. Any SIB authentication mechanism must therefore be self-contained within the broadcast itself, without relying on prior security context.

Additional SIB messages. Beyond SIB1, the 3GPP specification defines SIB2 through SIB21, each transmitted over DL-SCH in periodic scheduling windows and serving a distinct purpose. For instance, SIB3 carries intra-frequency neighbor cell lists and reselection parameters, SIB4 provides the inter-frequency equivalents, SIB9 delivers GPS and UTC timing information, and SIB15 conveys disaster-roaming configurations. While the authentication framework developed in this paper targets SIB1, the underlying approach generalizes naturally to other SIB types.

B. Building Blocks

TESLA Broadcast Authentication: TESLA (Timed Efficient Stream Loss-tolerant Authentication) [22] achieves sender authentication over broadcast channels using only symmetric primitives, at the cost of loose time synchronization. Its core mechanism combines a *one-way key chain* with *timed key disclosure*: the sender commits to the chain at setup and reveals each key only after a publicly known delay, so that receivers can verify authenticity without shared secrets and without any shared secret or asymmetric operations on every packet.

The chain is built using a one-way function $F\colon \{0,1\}^\lambda \to \{0,1\}^\lambda$ satisfying *(i) one-wayness*: given $K_i = F(K_{i+1})$, recovering K_{i+1} is infeasible; and *(ii) target collision resistance (TCR)*: finding $K' \neq K$ with $F(K') = F(K)$ for a given K is infeasible [22]. Notably, F need not be a keyed PRF. The sender samples $K_N \xleftarrow{\$} \{0,1\}^\lambda$ and derives $K_i \leftarrow F(K_{i+1})$ for $i = N-1,\ldots,0$, yielding the *public anchor* $K_0 = F^N(K_N)$. One-wayness ensures a receiver holding K_i cannot recover any K_j, $j > i$, while forward iteration trivially yields any K_j, $j < i$. For MAC key derivation, a domain-separated function $F'\colon \{0,1\}^\lambda \to \{0,1\}^\lambda$ produces $K'_i = F'(K_i)$ per interval, preventing structural links between chain and MAC keys. We instantiate both F and the per-packet MAC with *HMAC-SHA-256* [28], satisfying one-wayness and TCR under the PRF assumption and aligning with deployed 5G cryptographic stacks [29].

TESLA's security guarantee is that no adversary can forge a MAC tag on any packet that the receiver accepts as authentic, provided the security condition holds (the receiver verifies it received the packet before the corresponding key was disclosed) and F is one-way [22]. TESLA tolerates arbitrary packet loss: any subsequent disclosed key K_j allows the receiver to verify all earlier buffered packets from intervals $\leq j$ by forward-iterating F, regardless of intermediate drops. The original TESLA paper [22] introduces several variants.

Digital Signature Scheme: A digital signature scheme enables a signer to produce an unforgeable authentication tag on a message that any party holding the public key can verify.

Definition II.1. *A Digital Signature Scheme* $\mathtt{SGN}$ *is a triple of PPT algorithms* $(\mathtt{KeyGen}, \mathtt{Sign}, \mathtt{Verify})$*:*

- $(sk, PK) \leftarrow \mathtt{SGN.KeyGen}(1^\lambda)$*: Given security parameter* λ*, outputs a secret signing key* sk *and a public key* PK*.*
- $\sigma \leftarrow \mathtt{SGN.Sign}(sk, m)$*: Given* sk *and* $m \in \{0,1\}^*$*, outputs a signature* σ*.*
- $b \leftarrow \mathtt{SGN.Verify}(PK, m, \sigma)$*: Returns* 1 *(accept) if* σ *is a valid signature on* m *under* PK*, and* 0 *(reject) otherwise.*

Classical signature schemes based on integer factorization or discrete logarithms are broken by Shor's algorithm [30], motivating the adoption of PQ alternatives. NIST concluded its primary PQ standardization process, selecting three signature standards: *ML-DSA* (CRYSTALS-Dilithium) [25], a lattice-based scheme based on Module-LWE and Module-SIS; *SLH-DSA* (SPHINCS+) [31], a stateless hash-based scheme; and *FN-DSA* (Falcon) [32], a compact lattice-based scheme over NTRU rings. To encourage diversity in hard problems, NIST also launched an additional signature competition [33] for schemes not based on structured lattices, with first-round

(a) Fake Base Stations (b) MiTM Attacker (c) Quantum Adversary

Fig. 2: Outline of Our Threat Models.

selections spanning multivariate, code-based, and symmetric-based candidates.

MAYO as an instantiation of SGN: Among the third-round candidates of NIST's additional signature competition [33], *MAYO* [20] is a leading multivariate scheme based on the Oil-and-Vinegar (OV) framework, achieving EUF-CMA security under the hardness of the Multivariate Quadratic (MQ) problem. Its *whipping* technique yields compact signatures (186 B at NIST Level I) and a manageable public key (4,912 B), well-suited for bandwidth-constrained 5G broadcasts. MAYO also supports batch verification of k pairs $\{(m_j, \sigma_j)\}_{j=1}^{k}$ under the same PK at a cost well below k individual verifications.

III. Approach Overview

A. Threat Model

We model the adversary $\mathcal{A}$ as a Quantum Polynomial-Time (QPT) entity with full control over the wireless broadcast channel. Concretely, $\mathcal{A}$ can eavesdrop on all downlink transmissions from any gNB, and may inject, modify, replay, or selectively drop packets at will. $\mathcal{A}$ may also impersonate legitimate BSs to deceive UEs before any cryptographic protections are established. Under this model, $\mathcal{A}$ mounts three classes of attacks, illustrated in Fig. 2:

- **Fake Base Station (FBS) Attacks.** The adversary deploys a rogue base station that spoofs the identity of a legitimate gNB, luring victim UEs into establishing a connection. Once attached, the adversary can launch cascading attacks that exploit vulnerabilities in subsequent protocol stages, including tracking, denial of service, and downgrade attacks [5]. In III-C, we further list the attacks that fall under FBS.
- **Man-in-the-Middle (MitM) Attacks.** A MitM adversary positions itself between a UE and a legitimate gNB, intercepting and potentially altering unprotected signaling traffic. This attack is feasible whenever broadcast messages lack cryptographic authentication, such as digital signatures [34].
- **Quantum-Capable Adversaries.** Our threat model further accounts for adversaries that may eventually gain access to quantum computing resources capable of breaking conventional signature schemes [35].

B. Security Model

Our security analysis relies on the following definitions that correspond to the security properties of our building blocks.

Definition III.1 (EUF-CMA Security of SGN)**.** *A digital signature scheme* $\mathtt{SGN} = (\mathtt{KeyGen}, \mathtt{Sign}, \mathtt{Verify})$ *is existentially unforgeable under adaptive chosen-message attack (EUF-CMA) if for all PPT adversaries* $\mathcal{A}$*, given* PK *and adaptive access to a signing oracle* $\mathcal{O}_S$*, the probability of outputting a valid signature* (m^*, σ^*) *on a fresh* m^* *not queried to* $\mathcal{O}_S$ *is negligible:* $\mathsf{Adv}^{\text{EUF-CMA}}_{SGN,\mathcal{A}}(\lambda) \leq \mathsf{negl}(\lambda)$.

We instantiate SGN with MAYO [20], whose EUF-CMA security reduces to the hardness of the MQ problem, believed to be quantum-resistant. The security proof is generic in SGN and holds for any EUF-CMA-secure instantiation.

Definition III.2 (Security of F and F')**.** *Let* $F : \{0,1\}^\lambda \to \{0,1\}^\lambda$ *denote the one-way chain function and* $F' : \{0,1\}^\lambda \to \{0,1\}^\lambda$ *the MAC-key derivation function, both instantiated with HMAC-SHA-256 [22]. We require these properties:*

- *(i) Pseudorandomness.* F *and* F' *are pseudorandom functions (PRFs): no efficient distinguisher can tell* $F(K,\cdot)$ *or* $F'(K,\cdot)$ *apart from a random function with non-negligible advantage, where* $K \xleftarrow{\$} \{0,1\}^\lambda$*. HMAC-SHA-256 satisfies this under standard assumptions. With* $\lambda = 256$ *bits chain keys, quantum key recovery via Grover's algorithm requires* 2^{128} *sequential oracle invocations, meeting NIST Level I.*
- *(ii) One-Wayness. Given any chain element* K_i*, where* $K_0 = F^\ell(K_\ell)$ *and* $K_i = F^{\ell-i}(K_\ell)$*, no efficient adversary can recover* K_j *for any* $j > i$*:* $\Pr[\mathcal{A}(K_i) = K_{i+1}] \leq \mathsf{negl}(\lambda)$.
- *(iii) Target Collision Resistance (TCR). Given* K*, no efficient adversary can find* $K' \neq K$ *such that* $F(K') = F(K)$*:* $\Pr[\mathcal{A}(K) = K' : F(K') = F(K)] \leq \mathsf{negl}(\lambda)$.

Definition III.3 (EMULSION Security)**.** *Consider experiment* $\mathsf{Exp}^{\text{EUF-CMA}}_{EMULSION,\mathcal{A}}$*: the challenger runs* $EMULSION.\mathtt{Setup}$ *to produce* (sk_{AMF}, PK_{AMF})*, chain* $C = \{K_0, \ldots, K_\ell\}$*, and* $cert_{K_0} \leftarrow SGN.\mathtt{Sign}(sk_{AMF}, \mathsf{info})$ *where* $\mathsf{info} = K_0 \| ID_{BS} \| T_0 \| T_{\mathsf{int}} \| d \| \ell$*, giving* PK_{AMF} *to* $\mathcal{A}$ *(modeling eSIM pre-provisioning).* $\mathcal{A}$ *may adaptively query a broadcast oracle* $\mathcal{O}_B(i,m)$ *receiving honest broadcast packets* $\Pi_i \leftarrow EMULSION.\mathtt{Broadcast}(C, cert_{K_0}, \mathsf{info}, i, m)$*, modeling unrestricted observation of SIB broadcasts.* $\mathcal{A}$ *wins by outputting* (m^*, i^*, Π^*) *such that* $EMULSION.\mathtt{Authenticate}(PK_{AMF}, \Pi^*, t_R, \Delta t) = 1$*,* (m^*, i^*) *was never queried to* $\mathcal{O}_B$*, and* K_{i^*-d} *has not yet been publicly disclosed.* EMULSION *is secure if* $\mathsf{Adv}^{\text{EUF-CMA}}_{EMULSION,\mathcal{A}}(\lambda) \leq \mathsf{negl}(\lambda)$ *for all QPT* $\mathcal{A}$*.*

C. FBS Attack Characterization

An FBS is not itself an attack but a *capability* that enables a set of downstream exploits. Every FBS-driven compromise follows the same structure: the rogue cell broadcasts unverifiable SIBs, the *UE* camps on it, and the adversary exploits pre-security-context signaling. Mubasshir *et al.* [5] enumerate 21 distinct multi-step attacks, all conditioned on this unauthenticated broadcast, spanning denial of service, protocol downgrade, identity exposure, device fingerprinting, battery exhaustion, and public-warning misinformation (Table II). Notably, many attacks depend on SIB2–SIB8, not SIB1 alone: reselection priorities steer victims onto rogue cells, and public-warning payloads reside in SIB6–SIB8. Schemes that protect only SIB1 leave these vectors intact.

Detection-based defenses [5] achieve 96% FBS detection accuracy but are inherently *post hoc*: the *UE* must attach before a trace can be classified. EMULSION solves this pre-emptively: an unverifiable transport block is rejected before cell selection, so the *UE* never camps on the rogue cell. Parameter cloning confers no advantage, since it reproduces content but not a valid authentication tag. The rightmost

TABLE II: FBS-enabled multi-step attacks [5].

Category	Attacks	Dep.	EMULSION
Persistent DoS	Reject, kickoff, counter desync	SIB1	✓
Downgrade	RRC redir., cap. hijack	SIB2–5	✓
Identity exposure	IMSI/SUCI catch, auth relay	SIB1	✓
Fingerprinting	MNmap, cap. hijack	SIB1	✓
Energy drain	Lullaby, energy depletion	SIB1	✓
Misinformation	Panic attack	SIB6–8	✓
Network abuse	X2/Xn flood, HO hijack	SIB1	✓
Injection	Overshadow, paging hijack	PCCH	✓†

Note: ✓ = prevented (forged SIB rejected before cell selection); ✓† = detected but not prevented. Coverage for SIB2–SIB21 follows from the per-SIB-type chain construction (IV-C); measurements in section VI are for SIB1.

column of Table II shows the resulting coverage. Injection-delivered variants (overshadowing, paging hijacking) are reduced to detected denial of service rather than prevention; while paging-message authentication remains orthogonal [27].

D. Research Challenges and Our Solutions

We arrived at EMULSION by following a chain of design constraints, where each resolution exposes the next obstacle. We articulate these challenges and the design decisions they force, which together motivate the EMULSION framework. The technical realization of each is detailed in Section IV.

- **C1: Asymmetric PQ signatures cannot fit the broadcast constraint.** SIB broadcasts are fixed, sub-kilobyte, non-interactive frames transmitted on tight periodic windows (e.g., SIB1 is 372 B every 20–160 ms). Since the 5G security stack must adopt NIST-standardized primitives for long-term deployment, any certificate-based scheme that attaches a per-message NIST-PQC signature or certificate chain exceeds the frame budget, forcing fragmentation across many packets (13–34 packets for NIST-PQC standards) and end-to-end latencies up to 5,282 ms, which is incompatible with 5G bootstrapping. As our performance evaluation confirms (Section VI), this overhead is inherent to per-message asymmetric transmission and cannot be tuned away.
  - **Solution.** We remove the asymmetric cost from the per-message path. Each broadcast is authenticated with a symmetric primitive (HMAC), while a single NIST-selected round three PQ signature (MAYO [20]) is computed once per epoch and amortized across all broadcasts within that epoch. This reduces per-SIB overhead to a single compact tag that fits within one packet with no fragmentation, while retaining NIST-grade PQ security at the root of trust.
- **C2: Symmetric authentication normally presumes a shared secret unavailable at bootstrapping.** Adopting a symmetric approach to satisfy the constraint above requires the sender and receiver to share a key. However, a UE authenticating a BS before any security context is established shares no secret with it, and the broadcast is inherently one-to-many. Conventional symmetric authentication therefore cannot offer public verifiability, and symmetric efficiency and public verifiability appear mutually exclusive.
  - **Solution.** We employ delayed-disclosure one-way key chaining (a TESLA-style variant [22]), which achieves sender authentication over broadcast channels using only symmetric primitives and no shared secret, by committing to a one-way chain and revealing keys after a bounded delay. This yields the central property that most symmetric schemes lack: public verifiability at symmetric-key efficiency. EMULSION thus attaining the efficiency of a symmetric construction without restriction to pre-shared trust, a gain we substantiate through the remaining challenges.
- **C3: Delayed key disclosure requires a trusted timing and anchor reference at initial access.** The TESLA security condition requires the receiver to bound the sender's transmission time and to trust the chain anchor. At initial access, the UE holds neither an authenticated clock nor a pre-shared anchor, and the timing signals it observes (e.g., SFN on the PBCH) are supplied by the still-unverified BS. A design that assumes these away is vulnerable to replay, where an adversary rebroadcasts a previously valid epoch, its disclosed keys, and its SIB messages to a fresh UE.
  - **Solution.** We exploit two trust anchors independent of the unverified base station. For timing, the UE's GNSS receiver provides wall-clock time with sub-microsecond accuracy from a satellite constellation outside the FBS adversary's control, supplying the loose synchronization the TESLA safe-packet test requires [36]. For the chain root, the eSIM/USIM infrastructure [18], [19], which uses the same tamper-resistant storage as AKA credential provisioning, provisions the root public key offline, eliminating over-the-air certificate transmission. Because both references originate outside the radio channel, epoch freshness is bound to state that the attacker cannot forge, closing the replay surface (Section IV-B3).

E. Scope

EMULSION targets SIB broadcast authentication during the initial bootstrapping phase. It is orthogonal to and does not replace 5G-AKA, which provides UE–5GC mutual authentication and session key establishment. Our model does not address side-channel attacks, physical key extraction, relay attacks, jamming, overshadowing, or passive eavesdropping, each of which requires independent, orthogonal defenses.

Our scope excludes authentication and key agreement procedures (e.g., 5G-AKA [37], [14]) and their post-quantum variants [38], [39]. These protocols operate at the NAS layer and provide mutual UE–5GC authentication together with session key establishment, i.e., functions that are orthogonal to broadcast bootstrapping authentication. EMULSION targets the pre-NAS broadcast plane: it ensures that the UE receives a genuine SIB1 from a verified BS before any AKA exchange is initiated, thereby preventing FBS driven bootstrapping from propagating into the key agreement phase. Concretely, an operator can enforce this separation through a policy gate that permits AKA initiation only upon receipt of a fresh, EMULSION-verified SIB1; this linkage is complementary and preserves the cryptographic structure of AKA.

Overall, we do not address vulnerabilities that lie outside the broadcast authentication plane. In particular, the follow-

ing are out of scope: (i) *physical-layer attacks* such as, radio-frequency jamming, signal overshadowing, and downlink eavesdropping, which require orthogonal defenses such as anti-jamming techniques or physical-layer encryption; (ii) *hardware attacks* such as physical key extraction from BS hardware; (iii) *UE-to-BS privacy*, such as IMSI/SUPI catching, which is mitigated by NAS-layer SUCI concealment mechanisms [27]; and (iv) *denial-of-service attacks* at the MAC or RLC sublayers (e.g., resource exhaustion via crafted RACH preambles), which target availability rather than authenticity, (v) GNSS spoofing is explicitly excluded from our threat model, consistent with prior work [40], [36]; co-located GNSS spoofing substantially raises the adversary's cost and complexity beyond FBS deployment alone.

While these attacks lie outside our scope, their effect against EMULSION is bounded to denial rather than deception, since none can cause a UE to accept a forged SIB: forgery requires a valid per-packet tag, and thus a chain key, or a valid anchor signature, and thus sk_{AMF}—neither of which is available to the adversary. Under overshadowing, a forged SIB lacks a valid tag and fails verification, so the UE aborts before RACH and the attack yields a detected denial of service rather than a fraudulent attachment. Jamming and physical-layer interference prevent delivery entirely, causing bootstrapping to fail as with any broadcast scheme. Relay of unmodified genuine SIBs and downgrade via omission during incremental deployment are discussed in Appendix D. Across all these cases, EMULSION reduces the adversary's capability from substituting system information to merely denying it.

IV. The Proposed EMULSION Framework

A. Design Rationale and Comparative Justification

Infeasibility of Direct NIST-PQC Integration for 5G Authentication. Direct application of NIST-PQC signatures to SIB broadcasts is fundamentally precluded by 5G's packet-size constraints. Even the most compact standard, FN-DSA (Falcon-512), requires ≈6 SIB1 fragments and ≈800 ms end-to-end delay in optimized configuration; ML-DSA and SLH-DSA are substantially worse [8].

Leveraging 5G Architectural Features. EMULSION exploits three structural properties of 5G that prior schemes have not exploited. *(i) Fixed SIB transmission windows:* SIB1 is broadcast every 20–160 ms [2], defining natural authentication epochs that a time-aware scheme can exploit without protocol modification. *(ii) Time synchronization:* Modern UEs maintain wall-clock time via GNSS with sub-microsecond accuracy, independent of any base-station broadcast. Since 5G gNBs are themselves GNSS-synchronized [17], both endpoints share a timing reference outside the FBS adversary's control, making the TESLA security condition ($\Delta t \ll T_{\text{int}}$) trivially satisfiable without relying on the unauthenticated SFN. *(iii) eSIM credential provisioning:* The eSIM infrastructure stores long-term cryptographic material in tamper-resistant USIMs [18], [19] at subscription time, allowing the AMF's MAYO public key (4,912 B, stored once offline) to serve as the root of trust with no over-the-air transmission, eliminating certificate chains entirely. Together, EMULSION combines a TESLA-based HMAC chain anchored per epoch (exploiting (i) and (ii)) with a MAYO-certified root key provisioned via eSIM (exploiting (iii)), amortizing the single asymmetric PQ cost across all SIBs in the epoch and reducing per-broadcast overhead to a single HMAC computation in microseconds with no fragmentation.

Symmetric Broadcast Authentication via One-Way Key Chaining. TESLA imposes minimal runtime overhead: one F' evaluation and one HMAC per interval at the BS, and one HMAC verification per packet at the UE, both completing in microseconds with no fragmentation and no additional SIB transmissions. This contrasts sharply with asymmetric alternatives: ML-DSA requires ≈34 SIB fragments and up to 5,282 ms end-to-end delay; FN-DSA, even in optimized hybrid configuration, requires ≈6 fragments and ≈800 ms (see §VI). We adopt the packet-loss-tolerant variant of TESLA with d-interval delayed key disclosure, which is the natural choice given SIB1's fixed broadcast periodicity [2]. HMAC-SHA-256 instantiates both F and the per-packet MAC, satisfying the one-wayness and TCR properties required by TESLA [22], and the TESLA security condition is natively satisfiable since time sync is already native at the RAN level, requiring no additional timing infrastructure.

PQ Root Certification and Anchor Selection. MAYO-2 [20] offers the most favorable size-security trade-off among NIST PQC candidates and standards for our deployment model: its 186 B signature fits within a single SIB1 packet (372 B limit) with headroom for chain parameters, while its 4,912 B public key is stored once in the UE eSIM and never transmitted over the air. Among NIST additional signature candidates [33], MAYO-2 is the only multivariate scheme achieving NIST Level I with a sub-200 B signature; ML-DSA produces 2,420 B and FN-DSA 666 B at the same level, both requiring fragmentation at the root. MAYO's Oil-and-Vinegar structure provides EUF-CMA security under MQ hardness [20], and its batch verification property allows the UE to verify multiple per-epoch chain anchor certificates simultaneously, directly benefiting our per-SIB-type chaining design.

B. EMULSION Framework Initialization

*1) **Entities and Trust Model**:* EMULSION involves three entities: the AMF (which hosts the Key Management Function, CKG, as a logical sub-function), the BS (gNB), and the UE. The AMF is the root of trust: it holds the MAYO secret key sk_{AMF} and is responsible for generating and certifying all cryptographic material distributed to BSs. The UE trusts only the AMF's MAYO public key PK_{AMF}, provisioned into its eSIM once, offline, via the GSMA Remote SIM Provisioning (RSP) protocol [41]. No direct trust relationship between the UE and any BS is assumed; all BS-level trust is derived from the AMF's certification.

*2) **Root Key Generation**:* At system initialization, the CKG runs $(sk_{AMF}, PK_{AMF}) \leftarrow \texttt{MAYO.KeyGen}(1^{\lambda})$. The secret key sk_{AMF} is stored securely within the AMF and never transmitted. The public key PK_{AMF} is provisioned to all UE eSIMs via GSMA RSP and remains valid for the lifetime of

the deployment (or until a scheduled key rotation). All BS-level authentication traces back to this single root key.

Compromise Resilience and Key Lifecycle. EMULSION localizes the effect of compromise. A BS holds only the per-epoch chain material provisioned to it over N2, not any long-term secret, so a compromised BS can forge SIBs only until the epoch expires, after which a fresh certified anchor supersedes it. This contrasts with PKI- and IBS-based schemes, where a leaked BS key remains usable until explicit revocation. Only the AMF root key requires lifecycle management: rotation is performed offline over the GSMA RSP channel already used for provisioning, with UEs holding the current and next public key across a transition window, and root revocation is likewise an RSP update, avoiding per-BS revocation machinery. The remaining open case is cross-operator roaming: since trust is anchored in the home operator's pre-provisioned root, a visited network is authenticated only if its root is provisioned in advance or cross-certified by the home operator; absent this, roaming falls back to unauthenticated SIBs as in current deployments.

*3) **Time Synchronization**:* EMULSION uses TESLA, whose security condition requires the UE to bound the sender's current time with a maximum error Δt. Rather than relying on the SFN carried in the unauthenticated MIB, EMULSION derives its timing reference from the UE's GNSS receiver, which provides wall-clock time with sub-microsecond accuracy from a satellite constellation outside the FBS adversary's control. Since 5G gNBs are themselves GNSS-synchronized for network timing [17], both endpoints share a common time reference that is independent of any over-the-air broadcast, making the TESLA security condition trivially satisfiable for any reasonable T_{int}. We empirically validate this in §VI.

C. *EMULSION* Framework Main Operations

EMULSION is a symmetric chained, publicly verifiable authentication operating in three phases: Setup, Broadcast, and Verify. It is formalized in Algorithms 1–3 with the full protocol flow shown in Fig. 3. The algorithms are stated generically for a message m; in practice, m is any SIB type (SIB1 through SIB21), and the (s) superscript denoting SIB type is omitted for clarity. All chain variables, intervals, and certificates are implicitly indexed by the active SIB type.

Algorithm 1 EMULSION.Setup

$(\mathcal{C},\ cert_{K_0}) \leftarrow \texttt{EMULSION.Setup}(1^\lambda,\ \ell,\ d,\ T_0,\ T_{\text{int}})$: Run by the AMF once per chain epoch per SIB type.

$(sk_{AMF}, PK_{AMF}) \leftarrow \texttt{MAYO.KeyGen}(1^\lambda)$ ▷ Root keypair; PK_{AMF} provisioned to UE eSIM via GSMA RSP
$K_\ell \xleftarrow{\$} \{0,1\}^\lambda$ ▷ Random terminal key; kept secret at BS
for $i = \ell - 1$ **downto** 0 **do**
    $K_i \leftarrow F(K_{i+1})$ ▷ One-way chain; $K_0 = F^\ell(K_\ell)$ is the public anchor
$\mathsf{info} \leftarrow K_0 \| ID_{BS} \| T_0 \| T_{\text{int}} \| d \| \ell$
$cert_{K_0} \leftarrow \texttt{MAYO.Sign}(sk_{AMF},\ \mathsf{info})$
$\mathcal{C} \leftarrow \{K_0, K_1, \ldots, K_\ell\}$
Provision $(\mathcal{C},\ cert_{K_0},\ \mathsf{info})$ to BS over N2 interface

In **Setup** (Alg. 1), the AMF generates the MAYO root key pair and provisions PK_{AMF} to UE eSIMs via GSMA RSP [19]. It constructs a one-way key chain of length ℓ by sampling a random terminal key $K_\ell \xleftarrow{\$} \{0,1\}^\lambda$ and iterating $K_i \leftarrow F(K_{i+1})$ down to the public anchor $K_0 = F^\ell(K_\ell)$. The anchor and TESLA parameters are bound as $\mathsf{info} = K_0 \| ID_{BS} \| T_0 \| T_{\text{int}} \| d \| \ell$ and certified as $cert_{K_0} \leftarrow \texttt{MAYO.Sign}(sk_{AMF}, \mathsf{info})$. The chain $\mathcal{C}$ and certificate are provisioned to the BS over N2; the BS then operates autonomously for the entire epoch.

Per-SIB-type independent chains. EMULSION instantiates one independent chain per SIB type: SIB1 messages are authenticated by chain $\mathcal{C}^{(1)}$, SIB2 messages by $\mathcal{C}^{(2)}$, and so on through SIB21. Each chain is parameterized with its own terminal key, transmission interval T_{int} matching the 3GPP-defined periodicity of that SIB type [2], and disclosure depth d. The AMF runs EMULSION.Setup independently for each type, producing a dedicated certificate $cert_{K_0}$ for each chain anchor, and provisions all 21 chains and their certificates to the BS in bulk over N2 at epoch setup (precomputed offline). The BS holds the full set of chains and operates autonomously for the entire epoch, selecting the appropriate chain for each SIB type at broadcast time. The per-SIB broadcast and verification procedures follow Algorithms 2–3 identically for each type. Coverage extends from SIB1 to all 21 SIB types by adding one MAYO signing operation per type at epoch setup (all offline at the AMF), incurring zero additional over-the-air overhead per broadcast, since each chain's per-SIB cost remains a single HMAC tag regardless of SIB type. MIB coverage is also achieved at zero additional cost by including the 3-byte MIB content in the per-packet HMAC input. Details can be found in Appendix C.

Algorithm 2 EMULSION.Broadcast

$\Pi_i \leftarrow \texttt{EMULSION.Broadcast}(\mathcal{C},\ cert_{K_0},\ \mathsf{info},\ i,\ m)$: Run by the BS at each broadcast interval i.

$K'_i \leftarrow F'(K_i)$ ▷ Derive MAC key via F'
$\tau_i \leftarrow \texttt{HMAC}(K'_i,\ m \| i \| ID_{BS})$
$K_{\text{disc}} \leftarrow K_{i-d}$ **if** $i \geq d$, **else** $\perp$ ▷ Disclose key from d intervals ago
if $i = 1$ **then**
    $\Pi_i \leftarrow m \| i \| \tau_i \| K_{\text{disc}} \| \mathsf{info} \| cert_{K_0}$ ▷ Epoch-opening packet: include anchor certificate
else
    $\Pi_i \leftarrow m \| i \| \tau_i \| K_{\text{disc}}$ ▷ Steady-state packets: HMAC-only
Broadcast Π_i over DL-SCH

In **Broadcast** (Alg. 2), at interval i the BS derives MAC key $K'_i \leftarrow F'(K_i)$ and computes $\tau_i \leftarrow \texttt{HMAC}(K'_i, m \| i \| ID_{BS})$. It discloses K_{i-d}, enabling the UE to retroactively verify earlier buffered packets. The epoch-opening packet ($i = 1$) carries info and $cert_{K_0}$; all subsequent packets carry only m, i, τ_i, and K_{disc}, yielding a payload of ≈131 B that fits within a single SIB packet without fragmentation.

In **Verify** (Alg. 3), the UE performs the following steps. (i) On the epoch-opening packet ($i = 1$), it verifies $cert_{K_0}$ via MAYO.Verify under PK_{AMF} from the eSIM, and stores K_0 and the TESLA parameters as trusted state for this chain. (ii) It checks the TESLA security condition $i_{\max} < i + d$, discarding the packet if violated. (iii) It buffers (m, i, τ_i)

Algorithm 3 EMULSION.Authenticate

$\{1, \perp\} \leftarrow$ EMULSION.Verify$(PK_{AMF}, \Pi_i, t_R, \Delta t)$: Run by the UE upon receiving Π_i.

Parse $\Pi_i \rightarrow (m, i, \tau_i, K_{\text{disc}}, [\text{info}, cert_{K_0}])$ ▷ Bracketed fields present only when $i = 1$
if $i = 1$ **then**
  Parse info $\rightarrow (K_0, ID_{BS}, T_0, T_{\text{int}}, d, \ell)$
  if MAYO.Verify$(PK_{AMF}, \text{info}, cert_{K_0}) \neq 1$ **then**
    return $\perp$ ▷ Anchor not certified by AMF
  Store K_0 and TESLA parameters as trusted state
$i_{\max} \leftarrow \lfloor (t_R + \Delta t - T_0) / T_{\text{int}} \rfloor$
if $i_{\max} \geq i + d$ **then**
  return $\perp$ ▷ Security condition violated: K_{i-d} disclosed
Buffer (m, i, τ_i)
if $K_{\text{disc}} \neq \perp$ **then**
  Let $j \leftarrow i - d$
  if $F^j(K_{\text{disc}}) \neq K_0$ **then**
    return $\perp$ ▷ Disclosed key inconsistent with trusted anchor
  $K'_j \leftarrow F'(K_{\text{disc}})$
  Retrieve buffered (m_j, j, τ_j)
  if HMAC$(K'_j, m_j \| j \| ID_{BS}) = \tau_j$ **then**
    return 1 and ACCEPT(m_j)
return $\perp$

pending key disclosure. (iv) Upon receiving K_{disc}, it verifies chain consistency: $F^j(K_{\text{disc}}) \stackrel{?}{=} K_0$, $j = i - d$. (v) It derives $K'_j \leftarrow F'(K_{\text{disc}})$ and verifies the HMAC tag of the buffered packet. The UE accepts m_j only after all steps pass. The maximum authentication delay is $d \cdot T_{\text{int}}$, configurable per SIB type to match its 3GPP-defined transmission periodicity [2].

Packet Loss Tolerance and Optional FEC. EMULSION adopts the packet-loss-tolerant TESLA variant with d-interval delayed key disclosure [22]: a UE that misses interval i can still verify its buffered m_i upon receiving K_{i-d} in any later packet, with no retransmission required. Each broadcast packet is self-contained, so a single received packet suffices to authenticate all buffered messages with disclosed keys. For deployments requiring additional robustness against burst errors (e.g., dense urban or high-interference environments), EMULSION supports an optional FEC layer applied to the HMAC-only payload before broadcast. Steady-state packets carry ≈131 B, leaving ≈241 B of headroom within the 372 B SIB1 limit [2] to accommodate FEC redundancy without fragmentation. The BS applies FEC.Enc before broadcast and the UE applies FEC.Dec before HMAC verification; all other protocol steps are unchanged. FEC and HMAC serve orthogonal roles (e.g., channel reliability and cryptographic authenticity, respectively) and any standard code fitting within the available headroom may be used, including Reed-Solomon (MDS-optimal), 5G NR polar codes (native UE hardware support [17]), or Raptor codes (rateless, adaptive to variable loss rates). The FEC layer introduces no new cryptographic assumptions and does not affect the PQ security of the HMAC chain or the MAYO root anchor.

V. Security Analysis

We prove EMULSION achieves Definition III.3 under the security properties of Definitions III.1 and III.2. Any winning adversary $\mathcal{A}$ must succeed at one of three tasks: forge the SGN

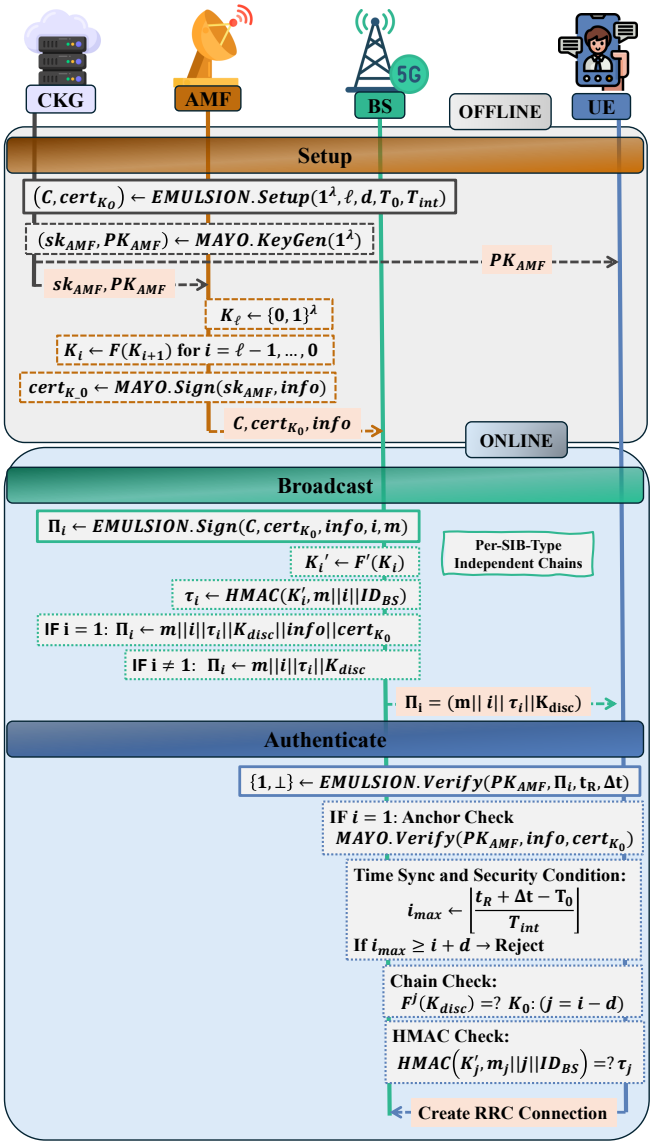

Fig. 3: EMULSION protocol flow.

anchor certificate $cert_{K_0}$ produced by EMULSION.Setup, forge a per-packet HMAC tag on a SIB broadcast by EMULSION.Broadcast without the chain key, or substitute a fraudulent chain key that passes the consistency check in EMULSION.Authenticate. The following lemmas reduce each task to one of the underlying hard problems with full security proofs presented in Appendix.

Lemma 1 (Anchor Certificate Unforgeability)**.** *If $\mathcal{A}$ forges a fresh $cert^*_{K_0}$ accepted by SGN.Verify, passing the certificate check in EMULSION.Authenticate, there exists a QPT $\mathcal{B}_1$ breaking EUF-CMA of SGN:* $\Pr[\mathcal{A} \text{ forges } cert_{K_0}] \leq \mathsf{Adv}^{\text{EUF-CMA}}_{SGN,\mathcal{B}_1}(\lambda)$.

Lemma 2 (Per-Packet MAC Unforgeability)**.** *If $\mathcal{A}$ forges a valid HMAC tag on a fresh (m^*, i^*) not broadcast by EMULSION.Broadcast, without access to K_{i^*-d}, there exists a QPT $\mathcal{B}_2$ breaking PRF security of F':* $\Pr[\mathcal{A} \text{ forges MAC}] \leq \ell \cdot \mathsf{Adv}^{\text{PRF}}_{F',\mathcal{B}_2}(\lambda)$, *where ℓ accounts for*

guessing the target interval i^* uniformly among ℓ epochs.

Lemma 3 (Chain Key Recovery Infeasibility)**.** *If $\mathcal{A}$ outputs $K^*_{\mathsf{disc}} \neq K_{i^*-d}$ with $F^j(K^*_{\mathsf{disc}}) = K_0$ where $j = i^* - d$, passing the chain check in* `EMULSION.Authenticate`*, there exists a QPT $\mathcal{B}_3$ breaking one-wayness or TCR of F:* $\Pr[\mathcal{A}$ forges chain key$] \leq \mathsf{Adv}^{\mathsf{OW\text{-}TCR}}_{F,\mathcal{B}_3}(\lambda)$.

Theorem 1 (`EMULSION` EUF-CMA Security)**.** *For any QPT adversary $\mathcal{A}$ against* `EMULSION` *over a chain of length ℓ, there exist a QPT $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$ such that:* $\mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathit{EMULSION},\mathcal{A}}(\lambda) \leq \mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathit{SGN},\mathcal{B}_1}(\lambda) + \ell \cdot \mathsf{Adv}^{\mathsf{PRF}}_{F',\mathcal{B}_2}(\lambda) + \mathsf{Adv}^{\mathsf{OW\text{-}TCR}}_{F,\mathcal{B}_3}(\lambda) + \mathsf{negl}(\lambda)$.

VI. Performance Evaluation

This section presents a comprehensive evaluation of `EMULSION` against NIST-PQC standards and conventional authentication schemes for 5G initial bootstrapping.

A. Configuration and Experimental Setup

Hardware: We assessed the efficiency of `EMULSION` protocol on a system equipped with a standard desktop equipped with a 12^{th} Gen Intel Core $i7 - 12700H$@$3.50\ GHz$, $16\ GiB$ RAM, a $512\ GiB$ SSD, and Ubuntu $22.04.4\ LTS$. Real network packets were investigated using the Network Signal Guru Android app installed on a OnePlus Nord 5G smartphone [42].

Over-the-Air Testbed. We deploy an SDR-based testbed for over-the-air evaluation using the open-source srsRAN and Open5GS stacks. srsUE and srsgNB run on two USRP B210 devices connected to the same host via USB 3.0, with a Leo Bodnar GPSDO providing a stable 10 MHz reference clock. The srsgNB connects to an Open5GS core and communicates with srsUE over the air. Since commercial UE basebands are closed-source, we adopt a best-effort methodology using this widely used srsRAN and Open5GS platforms, consistent with prior work on 4G/5G bootstrapping security [12], [13].

Libraries: We employed the OpenSSL library[1] for cryptographic primitives such as hash functions and elliptic curve operations (e.g., point multiplication, modular arithmetic), the Open Quantum-Safe library[2] for NIST-PQC schemes used in the baseline comparisons and for the MAYO-2 signing and verification operations within `EMULSION`, and the blst[3] library for the BLS signature.

Parameter Selection: We configured the post-quantum security to NIST Level I [43], which provides quantum resistance approximately equivalent to 128-bit classical security. For the conventionally secure baselines, all elliptic curve operations were performed over *secp256k1*, defined on a 256-bit prime field. In `EMULSION`, both the one-way chain PRF F and the key derivation PRF F' are instantiated as HMAC-SHA256 with distinct key generation to ensure cryptographic independence. MAC tags are truncated to 16 bytes, giving $\approx 2^{128}$ forgery resistance; Grover offers no speedup here, as forging a tag requires online verification attempts against the receiver rather than offline search.

[1]OpenSSL Library: https://openssl-library.org/

[2]Open Quantum-Safe Library: https://openquantumsafe.org/

[3]BLST Library: https://github.com/Chia-Network/bls-signatures

Evaluation Metrics and Rationale: Quantitative metrics include computational costs (signing, verification, and per-packet MAC operations), 5G processing delay (the time network entities spend handling cryptographic material in packets and transmitting them, excluding the cryptographic computations themselves), cryptographic overhead (signature, key, and certificate sizes), total over-the-air (OTA) communication overhead, and end-to-end (E2E) delay. Qualitative evaluation considers system architecture, PQ security guarantees, and resilience to packet loss.

srsRAN Configuration: For the first SIB1 message, the srsRAN gNB utilizes 79 bytes out of the allowed 372 bytes. All subsequent evaluations report the computational and communication overhead for a 79-byte SIB1 message, with 293 bytes available per packet for authentication material, as reflected in the tables. Even considering slightly larger SIB1 configurations observed from real networks, our results remain consistent. Moreover, while we present the evaluation of `EMULSION` on SIB1, the scheme generalizes to all SIB types, as the one-way chain and MAC-based authentication operate independently of the specific SIB content.

Baseline Selection: For PQ baselines, we consider the NIST-standardized *ML-DSA* [25] (lattice-based) and *FN-DSA* [32] (lattice-based), both in homogeneous 2-level certificate chains (e.g., *FN-DSA*-*FN-DSA*, *ML-DSA*-*ML-DSA*) and in a hybrid configuration where the root certificate uses MAYO and the BS uses *FN-DSA* (*FN-DSA*-MAYO) where we allow an optimization in favor of *FN-DSA*-MAYO by provisioning the MAYO public key ($PK_{\texttt{MAYO}}$) in the eSIM, eliminating the need to transmit it over the air, and also a variant augmented with Reed-Solomon erasure coding (FEC) for packet-loss resilience. Hash-based *SLH-DSA* [31], with a 7856-byte signature, is excluded from detailed comparison due to its prohibitively large size and slow execution time ($\sim$11 ms signing, $\sim$0.84 ms verification). For conventionally secure baselines, we include *EC-Schnorr* [23] and *BLS* [44] with aggregation, both deployed in a 2-level certificate hierarchy. While these schemes offer favorable performance, they lack PQ security guarantees and serve as reference points for understanding the computational overhead that `EMULSION` introduces relative to classical alternatives.

B. Experimental Results

TABLE III presents a comprehensive quantitative and qualitative comparison of candidate schemes for SIB1 authentication in the full 5G hierarchical bootstrapping context, covering signing and verification time, 5G processing and transmission delay, cryptographic and total OTA overhead (bytes), and E2E delay. Results are averaged over 10,000 iterations for standalone measurements and 10 iterations for over-the-air testbed runs (due to the requirement for manual intervention).

1) Quantitative Comparison: This section analyzes the computational and communication overhead of `EMULSION` alongside PQ and conventionally secure alternatives.

Computational Costs: A distinguishing feature of `EMULSION` is BS offloading: the MAYO-2 signature ($cert_{K_0}$) is computed

TABLE III: Comparison of candidate signature schemes for authenticating SIB1.

Scheme	System Architecture and Features	Sign Delay	Ver Delay (ms)	5G Delay (ms)	Crypto. Overhead (B)	Total OTA (B)	E2E Delay (ms)	Auth. Success @10% Loss
FN-DSA [24]	2-Level Certificate	0.28 ms	0.15	1920.24	3792	4836	1920.67	25.4%
FN-DSA-MAYO	$PK_{\texttt{MAYO}}$ in eSIM	0.28 ms	0.07	800.24	1749	2232	800.59	53.1%
FN-DSA-MAYO	$PK_{\texttt{MAYO}}$ in eSIM, FEC(n,k)	0.3 ms	0.07	960.24	2051	2604	960.59	85%
ML-DSA [25]	1-Level Certificate	0.12 ms	0.09	3201.32	6152	7440	3201.53	10.9%
ML-DSA [25]	2-Level Certificate	0.12 ms	0.12	5282.23	9884	12648	5282.47	2.8%
EMULSION	$PK_{\texttt{MAYO}}$ in eSIM, $\|\mathcal{C}\| = 2, d = 1$	1.3 μs	0.03	160.04	324	744	160.07	90.0%

Note: E2E delay presents the total time for signature generation, 5G delay, and full verification. Auth. Success reports epoch establishment probability at 10% packet loss. The impact of varying chain length $|\mathcal{C}|$ across different authentication ratios is discussed in detail in the experimental results section.

by the AMF offline during epoch setup, so the BS performs only two HMAC-SHA-256 operations per SIB1 packet (one for key derivation and one for MAC computation) at a per-packet signing cost of $\sim$0.6 μs, roughly three orders of magnitude faster than any signature-based alternative (FN-DSA: 0.28 ms, EC-Schnorr: 0.30 ms, BLS: 0.42 ms). Since a gNB may serve hundreds of cells simultaneously, this negligible per-packet cost has significant practical impact on BS infrastructure load. On the UE side, steady-state verification (all packets after the epoch-opening P_1) requires one HMAC chain check and one MAC verification at $\sim$0.03 ms, with the one-time MAYO-2 certificate verification adding another $\sim$0.03 ms amortized over the epoch. Full PQ certificate chains impose heavier verification: ML-DSA-ML-DSA requires 0.12 ms and FN-DSA-FN-DSA requires 0.15 ms for two-level verification, while conventionally secure schemes incur higher costs due to pairing (BLS: 3.46 ms) or multi-level EC point multiplications (EC-Schnorr: 3.80 ms). The net result is that EMULSION's computational footprint is dominated entirely by the 5G transmission schedule rather than cryptographic processing: the total cryptographic component of EMULSION's E2E delay ($\sim$0.03 ms) is negligible compared to the 160 ms SIB1 broadcast periodicity, as confirmed in Table III.

Communication Overhead: EMULSION uses two packet types. The epoch-opening packet (P_1) carries the SIB1 message (79 B), interval index (4 B), MAC tag (16 B), chain anchor K_0 (32 B), BS identity and metadata ($\sim$34 B), and MAYO-2 certificate (186 B), totaling $\sim$351 B, within the 372-byte limit. Subsequent packets ($P_2, P_3, \ldots$) carry only the message, interval index, MAC tag, and a disclosed key, totaling 131 B with just 52 B of authentication overhead. For a chain of length $|\mathcal{C}|$, total OTA is $(351 - 79) + (|\mathcal{C}| - 1) \times 52$ B per authentication payload; considering full 372-byte SIB1 blocks, $|\mathcal{C}|$=2 requires 744 B over 2 packets, $|\mathcal{C}|$=3 requires 1,116 B over 3 packets, and $|\mathcal{C}|$=4 requires 1,488 B over 4 packets, with every packet fitting within a single SIB1 transport block and zero fragmentation required. By comparison, full PQ certificate chains impose dramatically higher overhead. The ML-DSA-ML-DSA requires 9,884 B of cryptographic material across 34 packets (12,648 B total OTA), and FN-DSA-MAYO with $PK_{\texttt{MAYO}}$ pre-provisioned requires 1,749 B across 6 packets (2,232 B total OTA), rising to 2,051 B across 7 packets (2,604 B total OTA) when Reed-Solomon FEC is added for loss resilience. EMULSION at $|\mathcal{C}|$=2 achieves 324 B cryptographic overhead (5.4× lower than FN-DSA-MAYO and 31× lower than ML-DSA-ML-DSA) and even at $|\mathcal{C}|$=4 remains at 428 B, still 4.1× lower than FN-DSA-MAYO.

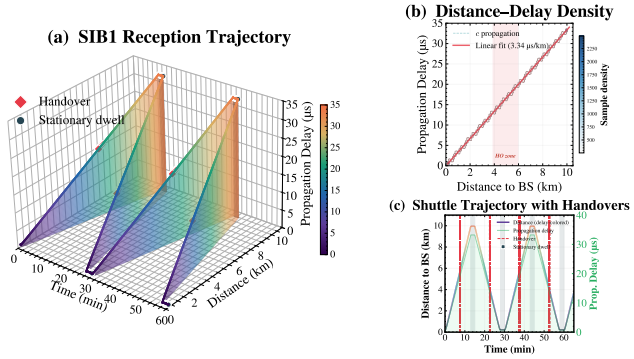

Fig. 4: SIB1 reception characterization on POWDER testbed.

• *5G Delay and E2E Latency:* The dominant cost for all schemes is the 5G delay, driven by the SIB1 broadcast periodicity of 20–160 ms; we evaluate at 160 ms (the common default) as the conservative upper bound. For multi-packet schemes, 5G delay is (packets − 1) × 160 ms plus per-packet processing. ML-DSA-ML-DSA requires 34 packets and incurs 5,282.23 ms ($\sim$5.3 s), entirely impractical for bootstrapping before RACH. FN-DSA-FN-DSA requires 13 packets and 1,920.24 ms; FN-DSA-MAYO with eSIM pre-provisioning reduces this to 6 packets and 800.24 ms, rising to 7 packets and 960.24 ms with FEC. EMULSION at $|\mathcal{C}|$=2 achieves 160.07 ms E2E delay, 5× faster than FN-DSA-MAYO (800.59 ms), 12× faster than FN-DSA-FN-DSA (1,920.67 ms), and 33× faster than ML-DSA-ML-DSA (5,282.47 ms). Even at $|\mathcal{C}|$=4, EMULSION's 480.11 ms remains 1.7× faster than the best PQ certificate chain. Table III reports delay from the epoch-opening packet. A UE arriving mid-epoch waits on average $|C| \cdot T_{int}/2$ more, i.e. 320 ms at $|C| = 2$, still below FN-DSA-MAYO's 800ms. Larger $|C|$ lowers amortized overhead but raises this wait. The wait applies only at first contact: a UE returning to a cell with cached anchor and parameters verifies from the next steady-state packet.

To assess whether distance, mobility, and handover introduce measurable overhead, we conduct a real-network experiment on the POWDER testbed [45]. Fig. 4 shows SIB1 reception traces from a campus shuttle over a $\sim$60-minute route with multiple cell attachments and handovers (See Appendix B for experimental setup). Panel (a) visualizes receptions across time, distance, and propagation delay with visible handover events at cell boundaries. Panel (b) confirms propagation delay

scales linearly at 3.34 μs/km, reaching at most $\sim$35 μs at 10 km (four orders of magnitude below EMULSION's 160 ms E2E delay). Panel (c) traces the UE's distance over the full route, marking handovers and stationary dwells. Across all conditions (with varying distance (0–10 km), speeds, and repeated handovers) the propagation component remains in the tens of microseconds, confirming that neither distance, mobility, nor handover introduces any observable impact on authentication delay for EMULSION or any baseline scheme.

UE-Side Overhead: EMULSION requires the UE to buffer at most d pending messages awaiting key disclosure ($d \times 372$ B; e.g., 372 B for d=1 and 1,116 B for d=3). In contrast, PQ certificate-chain schemes require buffering and reassembling fragments across 6–34 consecutive SIB1 packets before any verification can begin, demanding up to 12,648 B for ML-DSA-ML-DSA and introducing additional sequencing logic into the protocol stack. A single lost fragment forces the UE to discard the entire buffer and restart collection from the next broadcast cycle. In EMULSION, a missed packet affects at most d buffered messages; the next received packet discloses a new key and authentication resumes immediately without any state resynchronization.

2) Qualitative Comparison: This section examines the structural and security properties that differentiate EMULSION from alternative approaches.

Fragmentation and Protocol Compatibility: Every packet of EMULSION, including the epoch-opening packet carrying the MAYO-2 certificate, fits within a single 372-byte SIB1 transport block, eliminating fragmentation entirely and preserving full compatibility with the existing 5G protocol stack without RRC modifications, new message types, or application-layer FEC. PQ certificate-chain schemes are fundamentally incompatible with this constraint: ML-DSA-ML-DSA requires $\sim$10 KB of cryptographic material across 34 packets ($\sim$27$\times$ the transport block size), and even FN-DSA-FN-DSA requires 13 packets. Fragmentation introduces three practical challenges: *(i)* the UE must buffer and order fragments across multiple broadcast cycles, complicating the protocol stack; *(ii)* loss of any single fragment invalidates the entire authentication, since partial PQ signatures cannot be verified; and *(iii)* applying FEC to mitigate loss adds further packets and delay, exacerbating rather than resolving the overhead.

Post-Quantum Security: EMULSION achieves full PQ authentication for SIB broadcasts. The chain anchor K_0 is certified by MAYO-2, whose security rests on MQ hardness, while the one-way chain and per-packet HMAC derive security from the PRF assumption via HMAC-SHA-256, providing 128-bit PQ security at NIST Level I. The entire authentication path from AMF root of trust through per-packet MAC verification is therefore PQ-secure. In contrast, BLS (q-SDH), EC-Schnorr (ECDLP), and Schnorr-HIBS [40] all rely on hardness assumptions broken by Shor's algorithm, offering no long-term security against cryptographically relevant quantum computers.

Asymmetric Cost Amortization: EMULSION concentrates the asymmetric cost in the epoch-opening packet and amortizes it across the chain. The MAYO-2 signature is computed once per epoch by the AMF offline and verified once by the UE; all subsequent packets incur only symmetric operations. As $|\mathcal{C}|$ grows, amortized per-packet overhead converges to 52 B. At $|\mathcal{C}|$=100 ($\sim$16 s at 160 ms periodicity), the epoch-opening premium adds only $\sim$2 B per packet on average. This differs fundamentally from approaches such as EMSS [22], which require a fresh digital signature every k-th packet.

Packet-Loss Resilience: EMULSION provides inherent loss resilience without FEC. A missed steady-state packet P_i affects at most d messages but causes no cascading failure: the next received packet discloses K_i and authentication resumes immediately. If the epoch-opening packet P_1 is missed, the UE waits for the next epoch, which begins with a fresh certified anchor. Certificate-chain schemes have no inherent loss resilience: a single lost fragment out of 6–34 packets causes complete authentication failure, requiring the UE to wait for the next full broadcast cycle. Reed-Solomon FEC improves resilience but at the cost of additional packets, delay, and decoding complexity. Fig. 5 illustrates these differences. Panel (a) shows EMULSION (all $|\mathcal{C}|$) maintaining epoch establishment at 90% even at 10% packet loss, since only P_1 is critical, while ML-DSA-ML-DSA drops to $\sim$17% at 5% loss and FN-DSA-MAYO falls below 50% at 10% loss. Panel (b) shows steady-state per-message authentication: EMULSION with $|\mathcal{C}|$=2, d=1 authenticates $\sim$80% of messages at 20% loss, whereas certificate-chain baselines authenticate nothing until the full chain is successfully reassembled. Moreover, because an EMULSION retry costs only one additional epoch (160 ms), two consecutive attempts raise the success probability to $1-p^2 = 99\%$ ($p = 10\%$ loss) for a total wait of just 320 ms. The equivalent two-attempt probability for FN-DSA-MAYO (6 packets) is only $1-(1-(1-p)^6)^2 = 78\%$, at a cost of $\sim$1,600 ms.

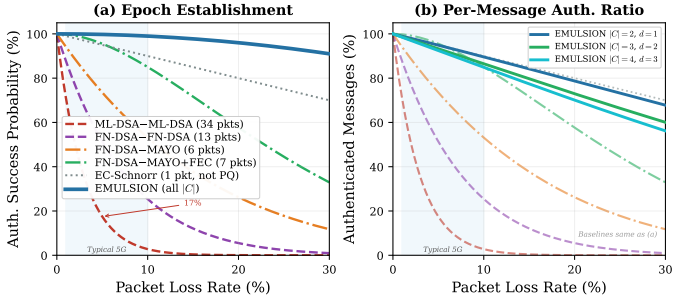

Fig. 5: Packet loss effect on authentication schemes.

Fig. 6 quantifies the delay-resilience tradeoff of Reed-Solomon FEC for FN-DSA-MAYO. Panel (a) shows that FEC variants RS(7,6) through RS(9,6) reduce expected E2E delay under loss compared to the unprotected 6-packet baseline, but all remain above 800 ms even at low loss rates, well above EMULSION's 160 ms at $|\mathcal{C}|$=2. Panel (b) shows that at 10% loss, FN-DSA-MAYO without FEC achieves only $\sim$53% authentication success; while RS(9,6) recovers high success, its E2E delay substantially exceeds all EMULSION configurations. FEC thus improves FN-DSA-MAYO's loss tolerance but cannot close the gap with EMULSION, which requires no error correction overhead.

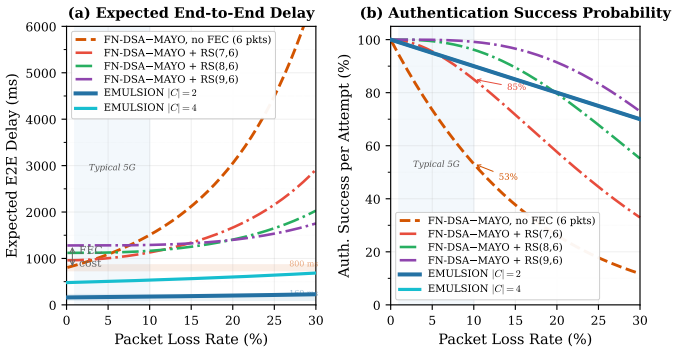

Fig. 6: Tradeoff between end-to-end delay and authentication success based on Reed-Solomon Forward Error Correction (FEC). EMULSION doesn't require FEC and outperforms **FN-DSA-MAYO** on all FEC settings.

Overall, EMULSION achieves PQ SIB authentication at 160.07 ms E2E delay in its minimal configuration ($|\mathcal{C}|$=2), which is 5–33× faster than PQ certificate-chain alternatives. Its per-packet BS cost of $\sim$0.6 μs (two HMAC operations) is roughly ***three*** *orders of magnitude* cheaper than any signature-based scheme, practical for gNBs serving hundreds of cells. Every packet fits withina single SIB transport block, eliminating fragmentation and preserving full 5G protocol stack compatibility. Since EMULSION authenticates over the one-way chain and MAC tags rather than SIB content, it extends naturally to all SIB types without modification, unlike certificate-chain schemes whose fragmentation overhead scales with the number of protected broadcast channels. Collectively, PQ security, zero fragmentation, HMAC-dominated per-packet cost, inherent loss resilience, and full SIB generalizability make EMULSION a practical and deployable solution for 5G bootstrapping authentication in the PQ era.

VII. Related Work

We survey prior work on 5G BS authentication and PQ security for cellular broadcast channels.

PKI-Based BS Authentication. Early proposals adapted PKI frameworks to authenticate SIB messages. Lee et al. [46] and Zheng [47] established certificate-based foundations for mobile network authentication. Hussain et al. [13] provided the first systematic attack analysis of 5G bootstrapping and proposed attaching signatures and certificate chains to SIB1/SIB2. Ross et al. [12] proposed a "broadcast-but-verify" model using a separate signingSIB message to decouple overhead from SIB1. Gao et al. [48] explored delegated signing to reduce per-BS cost, and Wuthier et al. [49] combined multi-factor authentication with blockchain-based certificate delivery. 3GPP has explored PKI-based SIB protection in TR 33.809, though SIB1 remains unprotected in the current RRC specification. All PKI-based schemes share a fundamental limitation: certificate chains for AMF and BS keys routinely exceed the 372-byte SIB1 limit, require fragmentation across multiple packets, and compound verification cost on resource-constrained UEs, while remaining vulnerable to quantum-capable adversaries.

Token- and Symmetric-Based Schemes. BARON [50] employs symmetric tokens for pre-authentication defense via a Closed Trusted Entity to protect connection initialization and handover in 5G, avoiding asymmetric certificate overhead but leaving SIB contents entirely unprotected: broadcast parameters can be tampered without invalidating any token. Al-Mekhlafi [51] extends symmetric security to 5G IoT contexts but similarly does not address SIB broadcast authentication. Song et al. [36] apply TESLA to SIB1 with a Schnorr-like IBS anchor verified once per cell entry, removing per-message signatures but leaving the root of trust on ECDLP and covering SIB1 alone. Substituting a PQ IBS does not help, as none fits a single SIB transport block.

Identity-Based and Certificate-Free Schemes. To eliminate certificate chains, Singla et al. [40] proposed SCHNORR-HIBS, a hierarchical IBS scheme deriving BS and AMF keys from a master key pre-installed in the USIM, achieving compact overhead within a single SIB1 packet. Ramadan et al. [52] explored server-aided IBS to offload UE verification cost. Yu et al. [53], Dong et al. [54], and Sun and Peng [55] further refined two-level HIBSs for LTE/5G. Sengupta and Lakshminarayanan [56] extended this to online-offline threshold IBS for 5G IoT. Darzi et al. [8] presented BORG, a threshold IBS scheme with fail-stop properties that distributes trust across multiple BSs and provides post-mortem forgery detection via a PQ-secured audit log. While IBS schemes achieve the best efficiency among conventional approaches, all rely on ECDLP hardness, are broken by quantum-capable adversaries, and are predominantly evaluated for SIB1 only.

Hybrid Post-Quantum Solutions. A growing body of work combines classical and PQ primitives for 5G/6G security. Vuppala et al. [57] and Ko et al. [58] proposed hybrid schemes pairing classical key exchange with lattice-based KEMs for primary authentication. Scalise et al. [59] analyzed PQ KEMs for 5G/6G core network security, and Attema and de Kock [60] examined PQC deployment challenges in the 5G core. These efforts target unicast session establishment and are structurally incompatible with one-to-many SIB broadcast authentication: KEMs establish shared secrets between two parties, and applying hybrid signatures to SIBs would combine the overhead of two schemes, further exacerbating fragmentation.

VIII. Conclusion and Future Work

We presented EMULSION, a symmetric chained publicly verifiable authentication framework for 5G/6G BS broadcast authentication achieving genuine post-quantum security at symmetric-key efficiency. By exploiting fixed SIB transmission windows, BS-UE time synchronization, and eSIM/USIM credential management, EMULSION harnesses a TESLA-style HMAC chain anchored by a compact MAYO signature applied once per epoch, eliminating certificate chains, avoiding fragmentation, and extending to the full SIB family (MIB and SIB1 through SIB21) at no additional per-broadcast overhead. Evaluated on a real over-the-air 5G testbed, it achieves lower end-to-end delay and communication overhead than ML-DSA with stronger, more durable security guarantees. Future work will extend EMULSION to multi-operator roaming and overshadowing mitigations in the post-quantum setting as 5G transitions toward 6G.

REFERENCES

[1] Ericsson, "Ericsson mobility report – mobile subscriptions outlook," 2023. [Online]. Available: https://www.ericsson.com/en/reports-and-papers/mobility-report/dataforecasts/mobile-subscriptions-outlook

[2] 3GPP, "NR; Radio Resource Control (RRC) Protocol Specification," 3rd Generation Partnership Project, Tech. Rep. TS 38.331 V18.1.0, 2024. [Online]. Available: https://www.etsi.org/deliver/etsi_ts/138300_138399/138331/18.01.00_60/ts_138331v180100p.pdf

[3] S. Hussain, O. Chowdhury, S. Mehnaz, and E. Bertino, "Lteinspector: A systematic approach for adversarial testing of 4g lte," in *25th Annual Network and Distributed System Security Symposium, NDSS 2018*, ser. 25th Annual Network and Distributed System Security Symposium, NDSS 2018. The Internet Society, 2018, publisher Copyright: © 2018 25th Annual Network and Distributed System Security Symposium, NDSS 2018. All Rights Reserved.; 25th Annual Network and Distributed System Security Symposium, NDSS 2018 ; Conference date: 18-02-2018 Through 21-02-2018.

[4] C. Park, S. Bae, B. Oh, J. Lee, E. Lee, I. Yun, and Y. Kim, "{DoLTEst}: In-depth downlink negative testing framework for {LTE} devices," in *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 1325–1342.

[5] K. S. Mubasshir, I. Karim, and E. Bertino, "Gotta detect'em all: Fake base station and multi-step attack detection in cellular networks," in *Proceedings of the 34th USENIX Security Symposium*, 2025.

[6] H. Yang, S. Bae, M. Son, H. Kim, S. M. Kim, and Y. Kim, "Hiding in plain signal: Physical signal overshadowing attack on LTE," in *28th USENIX Security Symposium (USENIX Security 19)*. Santa Clara, CA: USENIX Association, Aug. 2019, pp. 55–72. [Online]. Available: https://www.usenix.org/conference/usenixsecurity19/presentation/yang-hojoon

[7] C. J. Mitchell, "The impact of quantum computing on real-world security: A 5g case study," *Computers & Security*, vol. 93, p. 101825, 2020.

[8] S. Darzi, M. M. Rahman, I. Karim, R. Behnia, A. A. Yavuz, and E. Bertino, "Future-proofing authentication against insecure bootstrapping for 5g networks: Feasibility, resiliency, and accountability," *arXiv preprint arXiv:2510.23457*, 2025.

[9] NIST, "Post-Quantum Cryptography Standards: FIPS 203 (ML-KEM), FIPS 204 (ML-DSA), FIPS 205 (SLH-DSA)," National Institute of Standards and Technology, Tech. Rep. FIPS 203/204/205, 2024. [Online]. Available: https://csrc.nist.gov/projects/post-quantum-cryptography

[10] ETSI, "Cyber Security (CYBER); Quantum-Safe Cryptography (QSC); Efficient Quantum-Safe Hybrid Key Exchanges with Hidden Access Policies," European Telecommunications Standards Institute, Tech. Rep. TS 104 015 V1.1.1, 2024. [Online]. Available: https://www.etsi.org/deliver/etsi_ts/104000_104099/104015/01.01.01_60/ts_104015v010101p.pdf

[11] NSA, CISA, and NIST, "Quantum-Readiness: Migration to Post-Quantum Cryptography," Cybersecurity and Infrastructure Security Agency, Tech. Rep., 2025. [Online]. Available: https://www.cisa.gov/resources-tools/resources/quantum-readiness-migration-post-quantum-cryptography

[12] A. J. Ross, B. Reaves, Y. Nasser, G. Cukierman, and R. P. Jover, "Fixing insecure cellular system information broadcasts for good," in *Proceedings of the 27th International Symposium on Research in Attacks, Intrusions and Defenses*, 2024, pp. 693–708.

[13] S. R. Hussain, M. Echeverria, A. Singla, O. Chowdhury, and E. Bertino, "Insecure connection bootstrapping in cellular networks: the root of all evil," in *Proceedings of the 12th conference on security and privacy in wireless and mobile networks*, 2019, pp. 1–11.

[14] M. T. Damir, T. Meskanen, S. Ramezanian, and V. Niemi, "A beyond-5g authentication and key agreement protocol," in *International Conference on Network and System Security*. Springer, 2022, pp. 249–264.

[15] A. Braeken, A. K. Yadav, and J. Munilla, "A practical transition to post-quantum security in 5G-AKA," *IEEE Transactions on Information Forensics and Security*, vol. 20, pp. 13 071–13 084, 2025.

[16] Y. Ko, I. W. A. J. Pawana, and I. You, "5G-AKA-HPQC: Hybrid post-quantum cryptography protocol for quantum-resilient 5G primary authentication with forward secrecy," *arXiv preprint arXiv:2502.02851*, 2025.

[17] 3GPP, "NR; Base Station (BS) Radio Transmission and Reception," 3rd Generation Partnership Project, Tech. Rep. TS 38.104 V17.9.0, 2023. [Online]. Available: https://www.etsi.org/deliver/etsi_ts/138100_138199/138104/17.09.00_60/ts_138104v170900p.pdf

[18] GSMA, "Embedded SIM Remote Provisioning Architecture," GSM Association, Tech. Rep. SGP.02 V4.2, 2020. [Online]. Available: https://www.gsma.com/solutions-and-impact/technologies/esim/wp-content/uploads/2020/07/SGP.02-v4.2.pdf

[19] ——, "RSP Technical Specification," GSM Association, Tech. Rep. SGP.22 V3.1, 2023. [Online]. Available: https://www.gsma.com/solutions-and-impact/technologies/esim/wp-content/uploads/2023/05/SGP.22-v3.1.pdf

[20] W. Beullens, "Mayo: Practical post-quantum signatures from oil-and-vinegar maps," in *International Conference on Selected Areas in Cryptography*, pp. 355–376.

[21] ITU-R, "Framework and Overall Objectives of the Future Development of IMT for 2030 and Beyond," International Telecommunication Union, Tech. Rep. Recommendation ITU-R M.2160-0, 2023. [Online]. Available: https://www.itu.int/rec/R-REC-M.2160/en

[22] A. Perrig, R. Canetti, J. D. Tygar, and D. Song, "Efficient authentication and signing of multicast streams over lossy channels," in *Proceeding 2000 IEEE symposium on security and privacy. S&P 2000*. IEEE, 2000, pp. 56–73.

[23] C.-P. Schnorr, "Efficient signature generation by smart cards," *Journal of cryptology*, vol. 4, pp. 161–174, 1991.

[24] P.-A. Fouque, J. Hoffstein, P. Kirchner, V. Lyubashevsky, T. Pornin, T. Prest, T. Ricosset, G. Seiler, W. Whyte, Z. Zhang *et al.*, "Falcon: Fast-fourier lattice-based compact signatures over ntru," *Submission to the NIST's post-quantum cryptography standardization process*, vol. 36, no. 5, pp. 1–75, 2018.

[25] T. Dang, J. Lichtinger, Y.-K. Liu, C. Miller, D. Moody, R. Peralta, R. Perlner, A. Robinson *et al.*, "Module-lattice-based digital signature standard," *National Institute of Standards and Technology (NIST), Thinh Dang, Jacob*, 2024.

[26] H. Fourati, R. Maaloul, L. Chaari, and M. Jmaiel, "Comprehensive survey on self-organizing cellular network approaches applied to 5g networks," *Computer Networks*, vol. 199, p. 108435, 2021.

[27] *3GPP RRC Specification*, 2024, https://www.etsi.org/deliver/etsi_ts/138300_138399/138331/18.01.00_60/ts_138331v180100p.pdf.

[28] M. Bellare, R. Canetti, and H. Krawczyk, "Keying hash functions for message authentication," in *Annual international cryptology conference*. Springer, 1996, pp. 1–15.

[29] *3GPP Specification on Security architecture and procedures for 5G System*, 2024, https://www.etsi.org/deliver/etsi_ts/133500_133599/133501/18.06.00_60/ts_133501v180600p.pdf.

[30] P. W. Shor, "Algorithms for quantum computation: discrete logarithms and factoring," in *35th annual symp. on found. of CS*. Ieee, 1994.

[31] D. Cooper *et al.*, "Stateless hash-based digital signature standard," 2024.

[32] D. Soni, K. Basu, M. Nabeel, N. Aaraj, M. Manzano, and R. Karri, "Falcon," *Hardware Architectures for Post-Quantum Digital Signature Schemes*, pp. 31–41, 2021.

[33] NIST, "Additional Digital Signature Schemes – Round 3 Submissions," National Institute of Standards and Technology, Post-Quantum Cryptography Standardization, 2026. [Online]. Available: https://csrc.nist.gov/Projects/pqc-dig-sig/round-3-additional-signatures

[34] D. Rupprecht, K. Kohls, T. Holz, and C. Pöpper, "Breaking lte on layer two," in *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2019, pp. 1121–1136.

[35] S. Darzi, K. Ahmadi, S. Aghapour, A. A. Yavuz, and M. M. Kermani, "Envisioning the future of cyber security in post-quantum era: A survey on pq standardization, applications, challenges and opportunities," *arXiv preprint arXiv:2310.12037*, 2023.

[36] S. Song, M. K. Reiter, and T. Kwon, "Tesla-for-5g: Broadcast authentication for 5g networks using tesla," 2026. [Online]. Available: https://arxiv.org/abs/2606.26528

[37] G. Rossi Figlarz and F. Passuelo Hessel, "Enhancing the 5g-aka protocol with post-quantum digital signature method," in *International Conference on Advanced Information Networking and Applications*. Springer, 2024, pp. 99–110.

[38] A. Braeken, A. K. Yadav, and J. Munilla, "A practical transition to post-quantum security in 5g-aka," *IEEE Transactions on Information Forensics and Security*, vol. 20, pp. 13 071–13 084, 2025.

[39] M. T. Damir and V. Niemi, "On post-quantum identification in 5g," in *Proceedings of the 15th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, ser. WiSec '22. New York, NY, USA:

Association for Computing Machinery, 2022, p. 292–294. [Online]. Available: https://doi.org/10.1145/3507657.3529657

[40] A. Singla, R. Behnia, S. R. Hussain, A. Yavuz, and E. Bertino, "Look before you leap: Secure connection bootstrapping for 5g networks to defend against fake base-stations," in *Proceedings of the 2021 ACM Asia Conference on Computer and Communications Security*, 2021, pp. 501–515.

[41] *SGP.22: GSMA Remote Sim Provisioning*, 2024, https://www.gsma.com/solutions-and-impact/technologies/esim/gsma_resources/sgp-22-v2-7/.

[42] *Network Signal Guru User Manual*, https://m.qtrun.com/docs/NSG_Manual_Aug_2017.pdf.

[43] G. Alagic, D. Apon, D. Cooper, Q. Dang, T. Dang, J. Kelsey, J. Lichtinger, Y.-K. Liu, C. Miller *et al.*, "Status report on the third round of the nist post-quantum cryptography standardization process," 2022.

[44] D. Boneh, B. Lynn, and H. Shacham, "Short signatures from the weil pairing," in *International conference on the theory and application of cryptology and information security*. Springer, 2001, pp. 514–532.

[45] J. Breen, A. Buffmire, J. Duerig, K. Dutt, E. Eide, A. Ghosh, M. Hibler, D. Johnson, S. K. Kasera, E. Lewis, D. Maas, C. Martin, A. Orange, N. Patwari, D. Reading, R. Ricci, D. Schurig, L. B. Stoller, A. Todd, J. Van der Merwe, N. Viswanathan, K. Webb, and G. Wong, "Powder: Platform for open wireless data-driven experimental research," *Computer Networks*, vol. 197, p. 108281, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1389128621003017

[46] C.-C. Lee, I.-E. Liao, and M.-S. Hwang, "An extended certificate-based authentication and security protocol for mobile networks," *Information Technology and Control*, vol. 38, no. 1, 2009.

[47] Y. Zheng, "An authentication and security protocol for mobile computing," in *Mobile Communications: Technology, tools, applications, authentication and security IFIP World Conference on Mobile Communications 2–6 September 1996, Canberra, Australia*. Springer, 1996, pp. 249–257.

[48] H. Gao, Y. Zhang, T. Wan, J. Zhang, and H. Duan, "On evaluating delegated digital signing of broadcasting messages in 5g," in *2021 IEEE global communications conference (GLOBECOM)*. IEEE, 2021, pp. 1–7.

[49] S. Wuthier, J. Kim, J. Kim, and S.-Y. Chang, "Fake base station detection and blacklisting," in *2024 33rd International Conference on Computer Communications and Networks (ICCCN)*. IEEE, 2024, pp. 1–9.

[50] A. Lotto, V. Singh, B. Ramasubramanian, A. Brighente, M. Conti, and R. Poovendran, "Baron: Base-station authentication through core network for mobility management in 5g networks," in *Proceedings of the 16th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, 2023, pp. 133–144.

[51] Z. G. Al-Mekhlafi, M. A. Al-Shareeda, B. A. Mohammed, A. A. Alsadhan, A. Khalil, A. M. Alayba, A. M. S. Saleh, H. A. Alreshidi, and K. Almekhlafi, "Post-quantum lattice-based forward-secure authentication scheme using fog computing in 5g-assisted vehicular networks," 2024.

[52] M. Ramadan, Y. Liao, F. Li, and S. Zhou, "Identity-based signature with server-aided verification scheme for 5g mobile systems," *IEEE Access*, vol. 8, pp. 51 810–51 820, 2020.

[53] C. Yu, S. Chen, Q. Xing, and Z. Wei, "Protecting unauthenticated messages in lte/5g mobile networks: A two-level hierarchical identity-based signature (hibs) solution," *Computer Networks*, vol. 254, p. 110814, 2024.

[54] Y. Dong, R. Behnia, A. A. Yavuz, and S. R. Hussain, "Securing 5g bootstrapping: A two-layer ibs authentication protocol," *arXiv preprint arXiv:2502.04915*, 2025.

[55] Z. Sun and C. Peng, "5g-hcls: An authentication protocol to protect bootstrapping messages in 5g network," in *2025 IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE, 2025, pp. 1–6.

[56] B. Sengupta and A. Lakshminarayanan, "Fast verification of online/offline threshold signatures for 5g iot," in *2024 IEEE International Conference on Advanced Networks and Telecommunications Systems (ANTS)*. IEEE, 2024, pp. 1–6.

[57] R. C. Vuppala, D. Kumar, D. Je, N. Sharma, A. Nigam, and D. Kim, "Post-quantum secure hybrid methods for ue primary authentication in 6g with forward secrecy," in *GLOBECOM 2023-2023 IEEE Global Communications Conference*. IEEE, 2023, pp. 2590–2595.

[58] Y. Ko, I. Pawana, and I. You, "5g-aka-hpqc: Hybrid post-quantum cryptography protocol for quantum-resilient 5g primary authentication with forward secrecy," *arXiv preprint arXiv:2502.02851*, 2025.

[59] P. Scalise, R. Garcia, M. Boeding, M. Hempel, and H. Sharif, "An applied analysis of securing 5g/6g core networks with post-quantum key encapsulation methods," *Electronics*, vol. 13, no. 21, p. 4258, 2024.

[60] T. Attema, B. de Kock, S. M. Jayaprakash, D. Schoinianakis, T. Sijpesteijn, and R. van de Vlasakker, "Post-quantum cryptography in the 5g core," *arXiv preprint arXiv:2512.20243*, 2025.

APPENDIX A
SECURITY PROOFS

We present the full proofs of Lemmas 1–3 and Theorem 1. Each lemma constructs an explicit reduction adversary and analyzes its simulation and success probability.

We use hybrid games $G_0 \to G_1 \to G_2 \to G_3$, where G_0 is the real experiment and $\Pr[G_3 = 1] \le \mathsf{negl}(\lambda)$. In G_1, we abort if $\mathcal{A}$ forges a fresh SGN certificate passing the check in EMULSION.Authenticate; by Lemma 1: $|\Pr[G_0 = 1] - \Pr[G_1 = 1]| \le \mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{SGN},\mathcal{B}_1}$. In G_2, we additionally abort if $\mathcal{A}$ forges a valid HMAC tag on a fresh (m^*, i^*) not produced by EMULSION.Broadcast; since the TESLA security condition (Definition III.2) prevents access to K_{i^*-d} before its disclosure window expires, by Lemma 2: $|\Pr[G_1 = 1] - \Pr[G_2 = 1]| \le \ell \cdot \mathsf{Adv}^{\mathsf{PRF}}_{F',\mathcal{B}_2}$. In G_3, we abort if $\mathcal{A}$ submits a fraudulent K^*_{disc} passing the chain check in EMULSION.Authenticate; by Lemma 3: $|\Pr[G_2 = 1] - \Pr[G_3 = 1]| \le \mathsf{Adv}^{\mathsf{OW\text{-}TCR}}_{F,\mathcal{B}_3}$. After G_3, $\mathcal{A}$ has no remaining winning strategy: certificate forgeries, MAC forgeries, and chain key forgeries are all excluded, covering every verification path through EMULSION.Authenticate (Algorithm 3), so $\Pr[G_3 = 1] \le \mathsf{negl}(\lambda)$. Combining via the triangle inequality yields the bound in Theorem 1.

Lemma 1 (Anchor Certificate Unforgeability)**.** *If $\mathcal{A}$ forges a fresh $cert^*_{K_0}$ accepted by SGN.Verify, passing the certificate check in EMULSION.Authenticate, there exists a QPT $\mathcal{B}_1$ breaking EUF-CMA of SGN:* $\Pr[\mathcal{A}$ *forges* $cert_{K_0}] \le \mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{SGN,\mathcal{B}_1}\lambda)$.

Proof. The reduction exploits the fact that $\mathcal{B}_1$ can simulate the entire EMULSION experiment without knowing sk_{AMF}, since the chain $\mathcal{C}$ is sampled independently of the signing key. We construct a reduction $\mathcal{B}_1$ that uses $\mathcal{A}$ as a blackbox to break EUF-CMA of SGN (Definition II.1).

Setup. $\mathcal{B}_1$ receives a challenge public key PK from its EUF-CMA challenger and sets $PK_{AMF} := PK$. It samples the full key chain honestly by picking $K_\ell \xleftarrow{\$} \{0,1\}^\lambda$ and computing $K_i \leftarrow F(K_{i+1})$ for $i = \ell - 1, \ldots, 0$, yielding anchor $K_0 = F^\ell(K_\ell)$. It computes $\mathsf{info} = K_0 \| ID_{BS} \| T_0 \| T_{\mathsf{int}} \| d \| \ell$ and obtains $cert_{K_0} \leftarrow \mathcal{O}^{\mathrm{SGN}}_S(\mathsf{info})$ via a single query to its own SGN signing oracle. $\mathcal{B}_1$ runs $\mathcal{A}$ on input PK_{AMF}.

Broadcast oracle. For each query $\mathcal{O}_B(i, m)$ from $\mathcal{A}$, $\mathcal{B}_1$ computes $\Pi_i \leftarrow \mathtt{EMULSION.Broadcast}(C, cert_{K_0}, \mathsf{info}, i, m)$ directly using the known chain. No further SGN signing oracle queries are needed.

Extraction. When $\mathcal{A}$ outputs (m^*, i^*, Π^*) containing a fresh $cert^*_{K_0}$ satisfying $\mathtt{SGN.Verify}(PK_{AMF}, \mathsf{info}^*, cert^*_{K_0}) =$

1 for some info^* never submitted to $\mathcal{O}_B$, $\mathcal{B}_1$ outputs $(\mathsf{info}^*, cert^*_{K_0})$ as its SGN forgery.

Analysis. The simulation is perfect: $\mathcal{B}_1$ knows the full chain and answers all broadcast oracle queries honestly. The certificate requires exactly one SGN signing oracle query, permitted in the EUF-CMA game. Whenever $\mathcal{A}$ forges a fresh certificate, $\mathcal{B}_1$ wins its EUF-CMA game: $\Pr[\mathcal{A}$ forges $cert_{K_0}] = \Pr[\mathcal{B}_1$ wins EUF-CMA$] \leq \mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathtt{SGN},\mathcal{B}_1}(\lambda)$. □

Lemma 2 (Per-Packet MAC Unforgeability)**.** *If $\mathcal{A}$ forges a valid HMAC tag on a fresh (m^*, i^*) not broadcast by EMULSION.Broadcast, without access to K_{i^*-d}, there exists a QPT $\mathcal{B}_2$ breaking PRF security of F': $\Pr[\mathcal{A}$ forges MAC$] \leq \ell \cdot \mathsf{Adv}^{\mathsf{PRF}}_{F',\mathcal{B}_2}(\lambda)$, where ℓ accounts for guessing the target interval i^* uniformly among the ℓ chain epochs.*

Proof. We construct $\mathcal{B}_2$ that uses $\mathcal{A}$ to break PRF security of F'.

Setup. $\mathcal{B}_2$ has access to a PRF oracle $\mathcal{O}$ that is either $F'(K^*, \cdot)$ for a hidden random key K^*, or a truly random function $R(\cdot)$. It guesses a target interval $i^* \xleftarrow{\$} \{1, \ldots, \ell\}$ uniformly. It samples $K_\ell \xleftarrow{\$} \{0,1\}^\lambda$, computes $K_i \leftarrow F(K_{i+1})$ for all i, and obtains $cert_{K_0}$ via a single SGN signing oracle query as in Lemma 1. It runs $\mathcal{A}$ on input PK_{AMF}.

Broadcast oracle. For queries $\mathcal{O}_B(i, m)$ with $i \neq i^*$, $\mathcal{B}_2$ computes $\tau_i = \mathtt{HMAC}(F'(K_i), m \| i \| ID_{BS})$ directly. For $i = i^*$, it uses its PRF oracle: $K'_{i^*} \leftarrow \mathcal{O}(K_{i^*})$, then $\tau_{i^*} = \mathtt{HMAC}(K'_{i^*}, m \| i^* \| ID_{BS})$. All other packet fields are computed honestly.

TESLA security condition. By Definition III.2, $\mathcal{A}$ does not have access to K_{i^*-d} before the disclosure window expires. Therefore $\mathcal{A}$ cannot compute $K'_{i^*-d} = F'(K_{i^*-d})$ and must forge τ^* without knowing the MAC key. This is the critical point where the TESLA security condition is invoked in the reduction.

Extraction. When $\mathcal{A}$ outputs a valid forgery τ^* on a fresh (m^*, i^*) at the guessed interval, $\mathcal{B}_2$ distinguishes $\mathcal{O}$ from random: a valid MAC forgery has negligible probability against a truly random function. $\mathcal{B}_2$ outputs "PRF" when $\mathcal{A}$ succeeds and "random" otherwise.

Analysis. The simulation at all $i \neq i^*$ is perfect. At i^*, the simulation is indistinguishable from the real game when $\mathcal{O} = F'(K^*, \cdot)$. The guessing step succeeds with probability $1/\ell$: conditioned on $\mathcal{A}$ forging at any interval, it forges at i^* with probability at least $1/\ell$: $\Pr[\mathcal{A}$ forges MAC$] \leq \ell \cdot \mathsf{Adv}^{\mathsf{PRF}}_{F',\mathcal{B}_2}(\lambda)$. □

Lemma 3 (Chain Key Recovery Infeasibility)**.** *If $\mathcal{A}$ outputs $K^*_{\mathsf{disc}} \neq K_{i^*-d}$ with $F^j(K^*_{\mathsf{disc}}) = K_0$ where $j = i^* - d$, passing the chain check in EMULSION.Authenticate, there exists a QPT $\mathcal{B}_3$ breaking one-wayness or TCR of F: $\Pr[\mathcal{A}$ forges chain key$] \leq \mathsf{Adv}^{\mathsf{OW\text{-}TCR}}_{F,\mathcal{B}_3}(\lambda)$.*

Proof. We construct $\mathcal{B}_3$ that uses $\mathcal{A}$ to break one-wayness or TCR of F.

Two cases. A fraudulent $K^*_{\mathsf{disc}} \neq K_{i^*-d}$ passing $F^j(K^*_{\mathsf{disc}}) = K_0$ implies one of two things. *Case 1 (Preimage):* $\mathcal{A}$ inverts F along the chain, breaking one-wayness (Definition III.2(ii)). *Case 2 (Collision):* $\mathcal{A}$ finds $K^*_{\mathsf{disc}} \neq K_{i^*-d}$ with $F(K^*_{\mathsf{disc}}) = F(K_{i^*-d})$, breaking TCR (Definition III.2(iii)).

Setup. $\mathcal{B}_3$ receives anchor K_0 as its OW/TCR challenge. It samples $K_\ell \xleftarrow{\$} \{0,1\}^\lambda$, computes the full chain $K_i \leftarrow F(K_{i+1})$, and verifies $F^\ell(K_\ell) = K_0$, resampling if necessary. It obtains $cert_{K_0}$ via a single SGN signing oracle query and runs $\mathcal{A}$ on input PK_{AMF}, answering all broadcast oracle queries honestly.

Extraction. When $\mathcal{A}$ outputs $K^*_{\mathsf{disc}} \neq K_{i^*-d}$ with $F^j(K^*_{\mathsf{disc}}) = K_0$, $\mathcal{B}_3$ traces the chain from K^*_{disc} by applying F repeatedly and comparing against the honest chain. The first position k where $F(K^*_k) = F(K_k)$ but $K^*_k \neq K_k$ is a TCR collision (Case 2). If no such position exists, $\mathcal{B}_3$ has found a preimage for K_0 under F^j (Case 1). In either case, $\mathcal{B}_3$ outputs the relevant witness.

Analysis. The simulation is perfect since $\mathcal{B}_3$ knows the full chain. Every fraudulent K^*_{disc} passing the chain check in EMULSION.Authenticate yields a break of one-wayness or TCR: $\Pr[\mathcal{A}$ forges chain key$] \leq \mathsf{Adv}^{\mathsf{OW\text{-}TCR}}_{F,\mathcal{B}_3}(\lambda)$. □

Theorem 1 (EMULSION EUF-CMA Security)**.** *For any QPT adversary $\mathcal{A}$ against EMULSION over a chain of length ℓ, there exist QPT $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$ such that: $\mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathit{EMULSION},\mathcal{A}}(\lambda) \leq \mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathit{SGN},\mathcal{B}_1}(\lambda) + \ell \cdot \mathsf{Adv}^{\mathsf{PRF}}_{F',\mathcal{B}_2}(\lambda) + \mathsf{Adv}^{\mathsf{OW\text{-}TCR}}_{F,\mathcal{B}_3}(\lambda) + \mathsf{negl}(\lambda)$.*

Proof. We define hybrid games G_0, G_1, G_2, G_3, where G_0 is the real $\mathsf{Exp}^{\mathsf{EUF\text{-}CMA}}_{\mathtt{EMULSION},\mathcal{A}}$ experiment. We bound the gap between consecutive games and show $\Pr[G_3 = 1] \leq \mathsf{negl}(\lambda)$.

G_0: Real experiment. $\mathcal{A}$ interacts with the honest EMULSION challenger. $\Pr[G_0 = 1] = \mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathtt{EMULSION},\mathcal{A}}(\lambda)$.

$G_0 \to G_1$: Abort on certificate forgery. G_1 is identical to G_0 except the challenger aborts if Π^* contains a fresh $cert^*_{K_0}$ passing SGN.Verify. Any execution where G_0 outputs 1 but G_1 outputs 0 yields a valid SGN forgery. By Lemma 1: $|\Pr[G_0 = 1] - Pr[G_1 = 1]| \leq \mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathtt{SGN},\mathcal{B}_1}(\lambda)$.

$G_1 \to G_2$: Abort on MAC forgery. G_2 additionally aborts if $\mathcal{A}$'s forgery contains a valid HMAC tag on a fresh (m^*, i^*) not produced by EMULSION.Broadcast. In G_2, $\mathcal{A}$ must use the honest $cert_{K_0}$ (certificate forgeries are excluded by G_1). The TESLA security condition (Definition III.2) ensures $\mathcal{A}$ cannot access K_{i^*-d} before its disclosure window expires, so any valid MAC forgery breaks PRF security of F'. By Lemma 2: $|\Pr[G_1 = 1] - \Pr[G_2 = 1]| \leq \ell \cdot \mathsf{Adv}^{\mathsf{PRF}}_{F',\mathcal{B}_2}(\lambda)$.

$G_2 \to G_3$: Abort on chain forgery. G_3 additionally aborts if $\mathcal{A}$ submits a fraudulent K^*_{disc} passing the chain check $F^j(K^*_{\mathsf{disc}}) \stackrel{?}{=} K_0$ in EMULSION.Authenticate. Any such submission breaks one-wayness or TCR of F. By Lemma 3: $|\Pr[G_2 = 1] - \Pr[G_3 = 1]| \leq \mathsf{Adv}^{\mathsf{OW\text{-}TCR}}_{F,\mathcal{B}_3}(\lambda)$.

G_3: No winning strategy. Certificate forgeries, MAC forgeries, and chain key forgeries are all excluded in G_3, covering every verification path through EMULSION.Authenticate (Algorithm 3). Therefore $\Pr[G_3 = 1] \leq \mathsf{negl}(\lambda)$.

Composition. Applying the triangle inequality: $\mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathtt{EMULSION},\mathcal{A}} = \Pr[G_0 = 1] \le \Pr[G_1 = 1] + \mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathtt{SGN},\mathcal{B}_1} \le \Pr[G_2 = 1] + \ell \cdot \mathsf{Adv}^{\mathsf{PRF}}_{F',\mathcal{B}_2} + \mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathtt{SGN},\mathcal{B}_1} \le \mathsf{negl}(\lambda) + \mathsf{Adv}^{\mathsf{OW\text{-}TCR}}_{F,\mathcal{B}_3} + \ell \cdot \mathsf{Adv}^{\mathsf{PRF}}_{F',\mathcal{B}_2} + \mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathtt{SGN},\mathcal{B}_1}$, which yields the stated bound. □

APPENDIX B
POWDER SETUP

We collected the mobility dataset using the POWDER wireless testbed by deploying an experiment consisting of a core network and multiple base stations installed on rooftop locations, located at the University of Utah [45]. A separate experiment instance was configured with a campus shuttle bus carrying an srsUE device. The shuttle follows a fixed route at regular intervals, enabling controlled and repeatable mobility scenarios across the coverage areas of different base stations. As the shuttle moves away from one base station and approaches another with a stronger signal quality, a handover procedure is triggered, allowing the UE to transition to the better signal provider. During the experiments, all network communications and signaling exchanges were collected as packet capture (PCAP) traces. These traces were subsequently parsed using TShark to generate a structured dataset containing detailed mobility information, including the transmission and reception timestamps of SIB1 messages between the base station and UE, identification of handover events, and related signaling procedures. Additionally, the average speed of the campus shuttle was logged at approximately 29 mph, while the UE-to-base-station distance was estimated using the signal propagation time derived from the collected traces.

APPENDIX C
MIB COVERAGE VIA DEFERRED VERIFICATION.

Although the MIB is carried on the PBCH rather than the DL-SCH, its payload is only 3 bytes (24 bits) [27], small enough to be included directly in the per-packet HMAC input without affecting packet size. Concretely, the BS computes $\tau_i \leftarrow \mathsf{HMAC}(K'_i,\ m \| m_{\mathsf{MIB}} \| i \| ID_{BS})$, where m_{MIB} denotes the current MIB content. The UE necessarily decodes the MIB *before* receiving SIB1 (it is required to locate SIB1 on the DL-SCH), so MIB protection is inherently retroactive: the UE buffers the decoded MIB and, upon successful HMAC verification of the corresponding SIB1 packet, checks that the HMAC verifies under the MIB it decoded from the PBCH. Since m_{MIB} is an input to the HMAC, successful verification confirms that the BS committed to the same MIB content the UE received; a verification failure indicates a spoofed MIB and causes the UE to abort the connection attempt, preventing any further protocol progression (RACH, RRC, NAS) on a fraudulent cell. This deferred-verification model mirrors TESLA's own design, where packets are buffered pending key disclosure, and adds zero communication overhead, since the 3-byte MIB content is absorbed into the existing HMAC computation at negligible cost.

APPENDIX D
RELAY AND DOWNGRADE ATTACKS

Relay Attacks. A relay forwarding genuine, unmodified SIBs may pass our verification, since Algorithm 3 binds the tag to ID_{BS} but not to physical-layer observables (PCI, ARFCN, RF characteristics). However, because the replayed content is authentic, the adversary cannot manipulate cell-selection parameters, reselection priorities, or public-warning payloads without invalidating the HMAC tag; only attacks that require mere UE presence on the rogue cell survive. Binding authentication to a physical-layer anchor is complementary and left to future work.

Downgrade Attacks. During transition or partial rollout, an FBS can broadcast a legacy unauthenticated SIB1, indistinguishable from a non-upgraded cell. We can propose two alternatives here. (i) A UE in strict mode rejects unauthenticated SIBs, preventing downgrade at the cost of denying service on non-upgraded cells; (ii) in permissive mode it prefers authenticated cells but falls back to unauthenticated ones, preserving coverage but allowing bypass. EMULSION supports both via a per-PLMN policy flag provisioned through the eSIM channel: strict enforcement applies only to PLMNs the operator has enrolled, so an FBS cannot downgrade enrolled networks by omitting authentication material, while unenrolled networks remain accessible.