applied sciences

MDPI

Article

Hybrid Post-Quantum Key Establishment for Predictive Handover in 6G Satellite IoT Networks

Seungbin Lee [1] **and Jiyoon Kim** [1,2,*]

[1] Department of Computer Science and Engineering, Gyeongsang National University, Jinju 52828, Republic of Korea; dltmdqls1526@gnu.ac.kr
[2] Department of AI Convergence Engineering, Gyeongsang National University, Jinju 52828, Republic of Korea
* Correspondence: jykim92@gnu.ac.kr

Abstract

Sixth-generation (6G) satellite Internet of Things (IoT) networks can provide connectivity to remote sensors and gateways through low Earth orbit (LEO) satellites. Their rapid motion causes frequent handovers that require authentication. We propose a hybrid post-quantum key establishment procedure for predictive handover in 6G satellite IoT networks. The home network authorizes a target LEO satellite without possessing the handover authentication key, which the user equipment (UE) and the serving satellite derive from the existing access-stratum security context. A preliminary handover-preparation request enables target authorization and secure context provisioning. While the UE remains attached to the serving satellite, relayed X-Wing key establishment and mutual key-confirmation messages prepare a hybrid secret bound to the target satellite. After handover to the target satellite, a short authenticated activation exchange derives fresh session keys. ProVerif verifies four symbolic probe-secrecy queries and eight event-correspondence queries covering authentication, authorization, preparation, and message freshness. A separate post-completion access key disclosure model also satisfies four probe-secrecy queries. These symbolic results do not establish the computational security of X-Wing. Compared with on-demand handover using the same X-Wing configuration, our predictive scheme reduces median request-stage local processing from 356.93 to 21.37 µs (94.01%) and request-stage traffic from 3,770 to 470 bytes (87.53%). These results show that predictive preparation shifts hybrid key establishment and its associated signaling to the period before handover, reducing the processing and communication burden during the request stage despite the increased overall cost.

Keywords: 6G; satellite; Internet of Things; hybrid post-quantum cryptography

Check for updates

Academic Editor: Christos Bouras

Received: 13 November 2026
Revised: 25 November 2026
Accepted: 27 November 2026
Published: 28 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Sixth-generation (6G) satellite Internet of Things (IoT) networks can connect remote sensors and gateways beyond terrestrial coverage. In these non-terrestrial networks (NTNs), a fixed gateway may need to change its serving low Earth orbit (LEO) satellite as the constellation moves, even when the sensors are stationary. Satellite visibility windows can be estimated from orbital ephemerides and gateway location, providing an opportunity to prepare authentication before the transition [1,2]. Propagation delay, variable link quality, and intermittent service nonetheless constrain that preparation [3–6].

Authentication and fresh key establishment can extend handover processing. Existing protocols commonly use elliptic-curve Diffie–Hellman (ECDH) to agree on session keys.

This approach produces fresh link keys but assumes the continued hardness of the elliptic-curve discrete logarithm problem. Shor's algorithm challenges that assumption in the presence of a sufficiently capable quantum computer [7]. Post-quantum cryptography (PQC) addresses this challenge, although larger public keys and ciphertexts can increase signaling costs. The National Institute of Standards and Technology (NIST) standardized the module-lattice-based key-encapsulation mechanism (ML-KEM). ML-KEM provides a practical basis for a post-quantum transition [8–10].

A key encapsulation mechanism (KEM) establishes a shared secret between two parties. Hybrid key establishment combines conventional and post-quantum components during the post-quantum transition. Ko et al. [11] developed 5G-AKA-HPQC, a hybrid post-quantum cryptography (HPQC) extension of 5G primary authentication using the X-Wing KEM. X-Wing combines X25519 and ML-KEM-768 and derives a single shared secret [12,13]. The X-Wing public key and ciphertext are larger than those of a standalone ECDH exchange. We execute the hybrid exchange via the current serving connection and reserve the post-attachment window for authenticated key activation. The principal metric in this study is the work remaining after the handover request. Our contribution concerns the security protocol and the redistribution of hybrid key establishment operations across handover stages; it neither proposes a satellite mobility prediction algorithm nor evaluates radio-level handover performance.

Kim et al. [14] proposed an authentication handover protocol via ephemeral ECDH key agreement. Recent studies focus on handover optimization or on space information networks. Predictive preparation and post-quantum satellite authentication have been studied previously. The contribution of this study lies in the specific partitioning of these tasks: target-bound X-Wing preparation under the existing access security context, followed by activation after attachment. We evaluate the reduction in request-stage work and separately account for the preparation overhead.

The proposed satellite IoT handover involves user equipment (UE), a serving LEO satellite (L_1), a target LEO satellite (L_2), and a home network (HN). The HN acts as the authentication controller and approves the target preparation context. The UE and serving LEO independently derive the handover authentication key from the existing access-stratum (AS) key K_{gNB}. The serving LEO delivers this key to the target over a secure control channel. Neither the HN nor the serving LEO obtains the X-Wing shared secret. When a transition in LEO satellite visibility is predicted, a preparation-intent report is generated and forwarded to the HN via L_1. Target authorization and hybrid preparation between the UE and the target satellite are then carried out before handover execution. With the preparation completed, the actual handover requires only a short authenticated activation exchange. The protocol operates at the access-network security layer. The remote sensor gateway can serve as a security component for connectivity within a satellite IoT network, without requiring any changes to the sensing functions at the application layer.

The main contributions are as follows:

1. We design a hybrid post-quantum key establishment scheme for predictive handover in 6G satellite IoT networks, building on the earlier ECDH-based UE–LEO handover structure. The HN authorizes the target, the UE and serving LEO derive *MK* (message authentication code, MAC, key) from the serving K_{gNB}, and only the UE and target LEO establish the hybrid KEM secret.
2. We introduce predictive preparation based on an externally estimated handover time and a conditional preparation time scheduling guideline. KEM messages use the active serving connection, and activation follows target attachment; expiry and target-identity checks are specified but not independently verified by the symbolic queries.

3. We compare security services and analytically derive computation and message propagation costs for Third Generation Partnership Project (3GPP) Xn-interface handover [15,16], the Fast Handover Authentication Protocol (FHAP) [17], and the earlier ECDH-based NTN protocol [14]. A same-platform comparison of on-demand and predictive X-Wing isolates the gain from pre-request preparation and separates reference propagation estimates from measured local processing.
4. We model the UE, serving and target LEOs, and HN as ProVerif processes to verify symbolic confidentiality, injective authentication and freshness correspondences, non-injective authorization correspondences, and the effect of post-completion disclosure of the old Access-Stratum (AS) key on earlier secrets.

Section 2 reviews related work. Sections 3 and 4 present the system model and protocol, respectively. Section 5 examines security properties, Section 6 evaluates performance, and Section 7 concludes the paper.

2. Related Work and Preliminaries

2.1. Standardized and Extensible Authentication

The 3GPP specifies 5G authentication and key agreement (5G-AKA) and Extensible Authentication Protocol AKA-prime (EAP-AKA') for primary authentication. These security contexts support subsequent AS key derivation [15,18]. EAP-based secondary authentication is related to data network access rather than inter-satellite handover [19]. Certificate-based EAP methods using Transport Layer Security (TLS) provide alternative authentication frameworks [20,21]. However, these methods do not define predictive X-Wing handover. RFC 10024 specifies hybrid TLS 1.3 groups that include X25519 and ML-KEM-768 [22]. These groups provide hybrid key establishment. However, these groups do not provide post-quantum peer authentication or an NTN-specific handover schedule. Also, these methods differ from those of the X-Wing KEM used in this study.

2.2. Secure Handover Authentication in NTN

The mobility of LEO satellites differs from that of terrestrial cells because satellite orbits and expected visibility windows can be predicted. Existing NTN studies cover orbit information, predictive coverage, and link measurement for mobility management [2,5,23]. 3GPP standards specify both transparent and regenerative payload architectures. The proposed protocol assumes a regenerative payload capable of supporting cryptographic operations at the target satellite [16,24].

Yang et al. [17] proposed FHAP for high-mobility equipment in 5G satellite-terrestrial integrated networks. The group-based method utilizes the relay node (RN) and access and mobility management function to prepare target access nodes ahead of a railway handover. In the handover process, RN aggregates terminal authentication codes for verification by the target access node. FHAP is a non-hybrid, traditional authentication model, with a different group and relay topology.

Kim et al. [14] proposed an NTN security protocol that includes initial authentication and distinct intra-satellite and inter-satellite handover procedures. In the inter-satellite phase, the UE derives authentication material from its existing access security context and establishes a shared secret with the target satellite through ephemeral ECDH key agreement.

Preemptive post-quantum handover has also been studied. For space information networks, Palaniswamy et al. presented QPCASIN, a continuous-authentication scheme with preemptive handover and Frodo-based primitives [25]. Mamun et al. proposed a security framework for 6G NTN (SkyLock), combining a quantum-resistant authentication protocol, hierarchical key management, and a fast handover security protocol [26]. These studies establish relevant precedents for quantum-resistant handover. Our emphasis is

target-bound X-Wing preparation under an existing access-network key, followed by a short post-handover activation exchange.

2.3. Hybrid Post-Quantum Key Establishment

Federal Information Processing Standard (FIPS) 203 specifies ML-KEM, including ML-KEM-768 [8]. Ko et al. [11] integrate X-Wing into 5G-AKA-HPQC and accordingly measure the initial authentication overhead in a 5G testbed environment. X-Wing combines X25519 and ML-KEM-768 into a general-purpose hybrid KEM and generates a 32-byte shared secret [12]. We apply ML-KEM as the standardized primitive and follow the currently published X-Wing internet draft combination method. The cited specification is an Internet-Draft, rather than a finalized RFC.

RFC 10024 defines X25519MLKEM768, a hybrid key-agreement group for TLS 1.3 [22]. Its combiner and key schedule differ from X-Wing and are used within the TLS 1.3 transcript. X-Wing is a standalone KEM rather than a TLS key-exchange group, so its integration into a handover protocol requires independent peer authentication [12]. Table 1 lists the X-Wing sizes used in the signaling model. The proposed key schedule combines the hybrid secret with the existing handover authentication context.

Table 1. X-Wing parameters.

Item	Size
Public encapsulation key pk_U	1216 bytes
Ciphertext ct_T	1120 bytes
Shared secret Z_H	32 bytes
Decapsulation-key seed	32 bytes
Classical component	X25519
Post-quantum component	ML-KEM-768

The X-Wing public key concatenates the ML-KEM-768 and X25519 public keys, whereas the decapsulation key is a 32-byte seed from which both component secret keys are regenerated. In this proposal, the UE generates an ephemeral X-Wing key pair (sk_U, pk_U), and the target LEO L_2 performs encapsulation under the UE's public key pk_U. Algorithm 1 generates the UE key pair, Algorithm 2 encapsulates at the target, and Algorithm 3 decapsulates at the UE. These definitions follow the X-Wing specification and the description style of Ko et al. [11]. Peer identity must be authenticated separately.

Algorithm 1 X-Wing.KeyGen at the UE

$r \xleftarrow{\$} \{0,1\}^{256}$; $(d, z, x) \leftarrow \text{SHAKE256}(r, 96 \text{ bytes})$
$(sk_M, pk_M) \leftarrow \text{MLKEM768.KeyGenInternal}(d, z)$
$sk_X \leftarrow x$; $pk_X \leftarrow \text{X25519}(sk_X, G)$
$sk_U \leftarrow r$; $pk_U \leftarrow pk_M \,||\, pk_X$
return (sk_U, pk_U)

Algorithm 2 X-Wing.Enc at the target LEO

Input: $pk_U = pk_M \,||\, pk_X$
$e_X \xleftarrow{\$} \{0,1\}^{256}$; $c_X \leftarrow \text{X25519}(e_X, G)$
$ss_X \leftarrow \text{X25519}(e_X, pk_X)$
$(ss_M, c_M) \leftarrow \text{MLKEM768.Enc}(pk_M)$
$Z_H \leftarrow \text{XWC}(ss_M, ss_X, c_X, pk_X)$
$ct_T \leftarrow c_M \,||\, c_X$
return (Z_H, ct_T)

Algorithm 3 X-Wing.Dec at the UE

1: **Input**: $sk_U = r$, $ct_T = c_M \,||\, c_X$
2: Expand r as in Algorithm 1 to recover (sk_M, sk_X, pk_X)
3: $ss_M \leftarrow \text{MLKEM768.Dec}(c_M, sk_M)$
4: $ss_X \leftarrow \text{X25519}(sk_X, c_X)$
5: $Z'_H \leftarrow \text{XWC}(ss_M, ss_X, c_X, pk_X)$
6: **return** Z'_H

The X-Wing combiner is

$$\text{XWC}(a, b, c, d) = \text{SHA3–256}(a \,||\, b \,||\, c \,||\, d \,||\, \ell_{\text{XW}}), \tag{1}$$

where ℓ_{XW} is the six-byte label `5c2e2f2f5e5c` (hexadecimal). The domain separation label is appended, consistently with the implementation. The resulting Z_H has length 32 bytes. Key and ciphertext sizes are listed in Table 1.

For an ephemeral key pair generated by UE, the target satellite encapsulates a random secret, and the UE decapsulates it:

$$(sk_U, pk_U) \leftarrow \mathsf{XWing.KeyGen}(), \tag{2}$$
$$(Z_H, ct_T) \leftarrow \mathsf{XWing.Encaps}(pk_U), \tag{3}$$
$$Z'_H \leftarrow \mathsf{XWing.Decaps}(ct_T, sk_U). \tag{4}$$

Correct execution gives $Z'_H = Z_H$. The X-Wing combiner already incorporates X25519 and ML-KEM-768 material. Accordingly, no second ECDH secret is appended to Z_H in this paper. Explicit authentication remains necessary because KEM encapsulation alone does not identify the sender. We use a hash-based message authentication code (HMAC) and a key derivation function (KDF), specifically the HMAC-based extract-and-expand KDF (HKDF), with domain-separated labels [27–29].

The protocol uses the existing AS context as an authentication anchor rather than adding a public-key signature exchange. This design choice leaves the security of operator control channels as a separate deployment assumption [9].

2.4. Security Requirements

Prior work on NTN handover and hybrid-key authentication motivates mutual agreement, key confidentiality, preparation freshness, and protection of prior sessions after AS-key disclosure [11,14,17]. To evaluate the security of the proposed handover protocol, we define the five requirements summarized in Table 2.

Table 2. Security requirements for predictive satellite IoT handover.

No.	Security requirement	Description
SR1	Mutual authentication	UE and target LEO acceptance events must have distinct matching peer events, with agreement on the session identifier, activation nonces, and established key.
SR2	Hybrid and session-key confidentiality	The hybrid secret Z_H and handover keys K_{HO}, CK_{HO}, and IK_{HO} must remain confidential to a network adversary when active authentication credentials are uncompromised.
SR3	Preparation-to-activation freshness	A UE activation-start event must correspond injectively to a previously confirmed preparation event for the same session. This scoped correspondence does not prove all replay-cache or real-time-expiration behavior.
SR4	Authorization and target binding	Target acceptance must be associated with preceding HN authorization for that target and a target-specific hybrid encapsulation event.
SR5	Post-handover AS-key disclosure secrecy	Disclosure of the previous K_{gNB} after completion must not reveal the established hybrid secret or session keying material under the modeled key-erasure assumptions.

3. System and Threat Models

3.1. Network Architecture

Figure 1 illustrates a fixed remote sensor network in which the IoT gateway acts as the UE. Individual sensors do not execute the handover protocol. As the serving satellite (L_1) approaches the edge of coverage, the gateway prepares for handover to a target satellite (L_2) over its existing connection to L_1. The home network acting as authentication controller authorizes a target-specific context. The gateway and L_1 derive the authentication key MK, which L_1 securely transfers to L_2. Before target attachment, the UE and L_2 perform the X-Wing exchange and key confirmation through the serving connection.

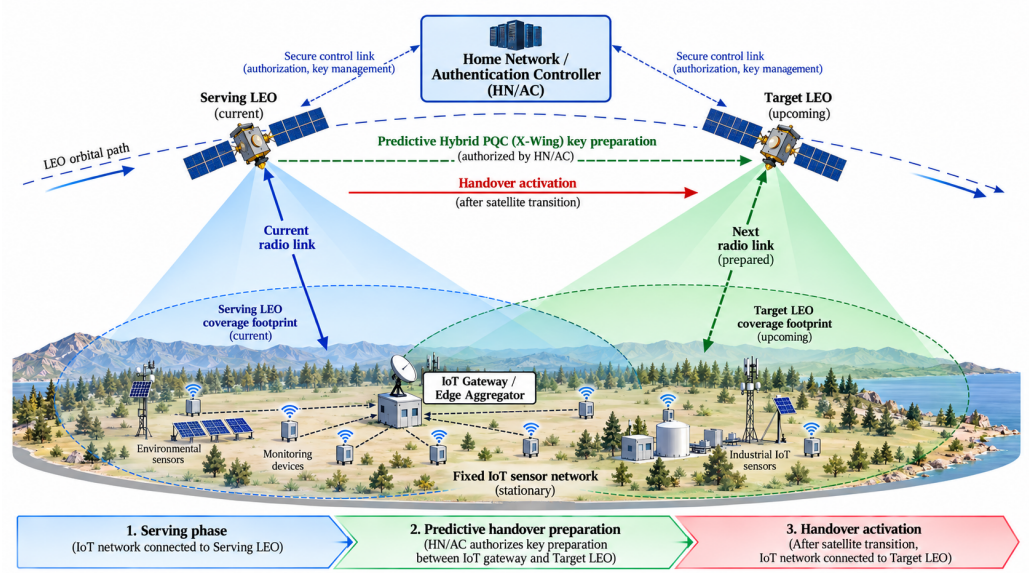

Figure 1. System architecture for predictive handover in 6G satellite IoT networks.

Target selection uses the network mobility procedure; the HN approves the selected target rather than predicting wireless coverage [15,24]. Figure 1 illustrates this architecture. In this context, key management denotes the HN's authorization of key use, rather than the HN's possession of MK. The label "next radio link" (prepared) indicates that the security context for the target radio link has been prepared.

3.2. Security Context and Assumptions

After primary authentication, the UE and serving regenerative LEO share an AS key (K_{gNB}). This key establishment follows the existing 5G security procedure. The handover authentication key (MK) is derived from that AS key and the handover context as a protocol-specific extension to the existing key hierarchy [15].

As specified in the key distribution protocol, the terrestrial HN approved the selected target satellite and did not derive, store, or receive K_{gNB}, MK, or the X-Wing shared secret. The HN maintains authenticated, and integrity-protected control links to both L_1 and L_2. These are the two blue links in Figure 1.

In addition, L_1 and L_2 require a handover context-transfer link that ensures authentication and confidentiality. This link is represented logically within the green preparation phase rather than as a separately labeled physical link. These links can pass through the terrestrial gateway and existing core network functions. MK must be delivered over a quantum-resistant channel. Otherwise, an adversary able to recover MK during the preparation window could impersonate the UE or L_2 and mount an active man-in-the-middle attack on the KEM exchange. Recovering MK after completion does not reveal past session keys, because Z_H remains protected by X-Wing. This is an explicit deployment assumption

rather than a result provided by X-Wing alone. In the single-process implementation, a fresh 32-byte L_1–L_2 AES-256-GCM key is generated with `RAND_bytes` and supplied to both role functions as a stand-in for an already established protected operator association. The encrypted context transfer uses a random 12-byte nonce and a 16-byte authentication tag. Network key establishment, secure provisioning, rollover, and long-term storage of this association key are not implemented or evaluated; operational deployment must provide and refresh that key through a separately protected control-plane mechanism.

In each implementation run, the association key is generated afresh, held in process memory, and cleared after use. Deployment requires authenticated key establishment, secure provisioning, storage, and renewal, with nonce uniqueness under each GCM key.

The UE–L_1 link uses AS integrity protection. The UE–serving-LEO mobility report is conveyed by an existing integrity-protected access-stratum (AS) signaling bearer. We assume that this bearer validates message integrity and enforces a bearer-specific freshness check before forwarding the report to the handover preparation procedure. In the symbolic model, the freshness check is represented by a non-repeating authenticated COUNT value and a corresponding one-use expected receive state. This AS protection is a deployment assumption rather than an additional HPQC message or a newly proposed anti-replay algorithm. Actual Packet Data Convergence Protocol (PDCP) COUNT management, receive windows, packet reordering, and state synchronization are not implemented in ProVerif. The threat model assumes an honest, operator-managed serving L_1 and target L_2; an actively compromised or malicious L_1 is outside its scope. Because L_1 knows MK, it could generate its own X-Wing key pair and nonce, construct a P_1 for an authorized context, and compute a valid τ_1. It could thereby attempt a parallel preparation with L_2 and complete key confirmation using its own decapsulation key, without the legitimate UE's participation. A MAC under MK alone cannot distinguish this attempt from the UE's message. The ProVerif results therefore do not establish security against this insider. Deployment requires a trusted serving node, a unique HN-issued session identifier, and atomic single-use context handling at L_2, including rejection of parallel, consumed, or expired contexts. Single-use enforcement limits duplicate use but cannot by itself prevent a malicious L_1 from winning the first-use race. These stateful controls are deployment requirements and are not fully established by the symbolic queries. Only the UE and L_2 recover the honest-session X-Wing secret Z_H under the stated trust assumptions.

3.3. Adversary Model

In this study, we adopt the Dolev–Yao attack model for the public radio and UE-target relay path [30,31]. The adversary can intercept, inject, replay, suppress, reorder, and modify messages. The adversary can evaluate public cryptographic constructors and apply destructors only when the required inputs and keys are available. Captured communication logs may be stored for later analysis. The HN–LEO control associations and the L_1–L_2 transfer of MK are modeled as authenticated and confidential operator channels. These channels are deployment assumptions rather than proven channel properties. The HN, L_1, and L_2 are honest in the modeled exchange; the HN does not possess K_{gNB}, MK, or Z_H.

In this study, a retrospective disclosure experiment permits exposure of the old K_{gNB} after the modeled session completes and ephemeral KEM material is erased. However, compromise of the authentication key in use or the operator channel is outside the scope of this study. An attacker who possesses the activated K_{gNB} can derive MK and impersonate a peer. We do not claim any security if the active key is compromised. The symbolic X-Wing model idealizes the ML-KEM-768 and X25519 components and their combiner. This model does not verify computational post-quantum security, real-time expiry, jamming,

or resource restrictions [8,12]. Protocol-level checks for target identity, suite and expiry are specified but only the queried subset is formally verified.

3.4. Notation

Table 3 defines the protocol parameters. Handover-specific contexts, nonces, and derived keys are scoped to an authorized session; the existing AS key may be reused across sessions. The proposed context and key derivation labels are distinct from those defined by 3GPP [15].

Table 3. Protocol notation.

Symbol	Definition
UE, L_1, L_2, HN	UE, serving LEO, target LEO, home network
$K_{\text{gNB}}^{(1)}$	Current UE–L_1 access-stratum key
$\mathsf{sid}, \mathsf{uid}, \mathsf{suid}, \mathsf{tid}$	Session, UE, serving, and target identifiers
MR, t_{req}	UE report and its timestamp
t_C	HN context-creation timestamp
$\widehat{t}_{HO}, \mathsf{exp}$	Predicted handover time and context expiry
n_H	Fresh authorization nonce generated by HN
C_{HO}	HN-authorized handover context
PRK_M, MK	Source-key extraction material and handover MAC key
n_U, n_T	UE and target nonces for KEM preparation
pk_U, sk_U	UE ephemeral X-Wing key pair
ct_T, Z_H	Target KEM ciphertext and hybrid secret
P_1, P_2	Encoded UE and target KEM messages
tr	Canonical preparation transcript
PRK, K_C, K_{HO}	Hybrid extraction, confirmation, activation keys
$n_U^{\text{A}}, n_T^{\text{A}}$	Fresh UE and target activation nonces
$CK_{\text{HO}}, IK_{\text{HO}}$	Handover encryption and integrity keys
$\tau_0, \ldots, \tau_7$	Context, preparation, confirmation and activation MACs
T_P	Pre-handover preparation duration

In the implementation, sid and the five nonces n_H, n_U, n_T, n_U^{A}, and n_T^{A} are independently generated 16-byte values. Participant identifiers use an 8-byte experimental encoding, while MK and HMAC tags are 32 bytes. These serialization choices are not standardized 3GPP identifiers [15]. The implementation serializes each packet as a one-byte message-type identifier followed by ordered fields, each prefixed with an unsigned two-byte big-endian length. Specifically, C_{HO} contains sid (16 B), UE/source/target identifiers (8 B each), creation time (8 B), HN nonce (16 B), expiry (8 B), and suite (2 B), with timestamps encoded as unsigned big-endian seconds and a two-byte suite value `0x0101`. The untagged P_1 contains $(\mathsf{sid}, \mathsf{uid}, \mathsf{tid}, n_U, pk_U)$ of lengths $(16, 8, 8, 16, 1216)$ B, and the untagged P_2 contains $(\mathsf{sid}, n_T, ct_T)$ of lengths $(16, 16, 1120)$ B; complete transmitted frames append a full 32-byte HMAC (no truncation). The canonical transcript digest is SHA-256 over the two-byte big-endian length and full bytes of each $C_{\text{HO}}, P_1, P_2, \tau_1, \tau_2$ in that order, where P_1, P_2 exclude their appended HMAC fields. The implementation uses literal ASCII domain labels and retains the full 32-byte SHA-256/HKDF outputs.

The HMAC and extract-and-expand operations in the following equations follow RFC 2104 and RFC 5869, respectively [27,28]; X-Wing key and ciphertext formats follow the cited X-Wing specification [12]. The message identifiers, field lengths, serialization order, and domain labels above are choices of this implementation, not encodings mandated by these standards or 3GPP. In the implementation, the encoding and key schedule are as follows: the serialized context and preparation messages include their one-byte type identifiers, HMAC and HKDF labels are literal ASCII strings without a terminator, and each

transcript component is preceded by its two-byte big-endian length. The included X-Wing reference-vector self-check compares the first 16 bytes of the public key and ciphertext and the complete 32-byte shared secret, including decapsulation. This primitive-level check is not a complete deterministic end-to-end test vector for the handover protocol; such a vector remains future reproducibility work.

4. Proposed Protocol

4.1. Protocol Overview

The protocol distinguishes handover preparation from execution. During preparation, the gateway remains connected to L_1. The protocol assumes that the operator's existing mobility-management procedure supplies a candidate target L_2 and an estimated handover time; neither target prediction nor orbital-state estimation is performed by the proposed key establishment protocol. The gateway reports predicted mobility and participates in HN-authorized UE and L_2 key preparation. After preparation succeeds, the gateway initiates handover execution via a separate execution request. Upon attachment, the UE and L_2 exchange three authenticated messages to activate the stored context. In the on-demand baseline, the same authorization procedure and X-Wing computation are performed upon UE handover request. Both methods share the same authentication suite and roles but differ in scheduling and relay transfer method. This procedure expands on the message-level organization proposed in [14,32].

Existing fifth-generation system (5GS) primary authentication provides K_{gNB} to the UE and L_1 [15]. The HN does not receive this AS key, and the proposed MK, CK_{HO}, and IK_{HO} are handover-specific derivations, not the standard 5G-AKA cipher and integrity keys.

4.2. Predictive Handover Preparation Phase

Step 1 (Prediction and authorization request): While the UE is served by L_1, it sends a protected measurement report (MR) and a timestamp (t_{req}). The source verifies AS integrity and selects a target (L_2) using the operator mobility procedure. The L_1 delivers an authenticated request $(uid, suid, tid, MR, \widehat{t}_{HO}, t_{req})$ to HN. HN verifies the request from L_1 via the protected association, and verifies that the UE, serving node, and target identifier are approved. The report of the UE is authenticated using the existing AS security context. The serving LEO processes the mobility report only after the existing AS layer has validated its integrity and bearer-local freshness. A duplicate report that fails the AS receive-state check is not delivered as a new preparation request.

Step 2 (Authorized context and MK): Upon approval, the HN generates a new session identifier (sid), context generation timestamp (t_C), authorization grant nonce (n_H), expiration time (EXP), and algorithm-suite identifier.

$$C_{\mathrm{HO}} = \mathrm{Encode}(\mathsf{sid}, \mathsf{uid}, \mathsf{suid}, \mathsf{tid}, t_C, n_H, \mathsf{exp}, \mathsf{suite}) \tag{5}$$

The HN sends C_{HO} independently to L_1 and L_2 on authenticated control paths. The serving L_1 verifies the authorization and computes

$$\begin{aligned} PRK_M &= \mathsf{HKDF\text{-}Extract}(0, K_{\mathrm{gNB}}^{(1)}) \\ MK &= \mathsf{HKDF\text{-}Expand}\big(PRK_M, \text{``}NTN - HO - MK\text{''} \,||\, C_{\mathrm{HO}}, 32\big) \end{aligned} \tag{6}$$

L_1 transfers (C_{HO}, MK) to L_2 through the confidential handover-context path. L_2 checks that C_{HO} matches the approval received from the HN and stores MK only until exp. L_1 delivers C_{HO} to UE over the current protected radio link, together with

$$\tau_0 = \mathsf{HMAC}_{MK}(\text{``}HO - CONTEXT\text{''} \,||\, C_{\mathrm{HO}}) \tag{7}$$

The UE computes MK independently based on $K_{gNB}^{(1)}$ and C_{HO}, verifies τ_0, and confirms source, target, algorithm-suite and expiration.

4.3. Pre-Handover Hybrid Key Establishment Phase

Step 3 (UE key generation): With the UE attached to L_1, the UE generates a one-time X-Wing key pair (sk_U, pk_U) and a fresh preparation nonce n_U. The UE forms

$$P_1 = \text{Encode}(\text{sid}, \text{uid}, \text{tid}, n_U, pk_U) \tag{8}$$

$$\tau_1 = \text{HMAC}_{MK}(\text{``}HO - PREP - 1\text{''} \,||\, C_{\text{HO}} \,||\, P_1) \tag{9}$$

The UE sends (P_1, τ_1) through the active serving connection. Prior to computing KEM, L_2 verifies the MAC and the HN-approved context.

Step 4 (Target encapsulation): After checking that sid has not been used and the context has not expired, L_2 performs $(Z_H, ct_T) \leftarrow \text{XWing.Encaps}(pk_U)$ and generates the fresh preparation nonce n_T. L_2 constructs

$$P_2 = \text{Encode}(\text{sid}, n_T, ct_T) \tag{10}$$

$$\tau_2 = \text{HMAC}_{MK}(\text{``}HO - PREP - 2\text{''} \,||\, C_{\text{HO}} \,||\, P_1 \,||\, P_2) \tag{11}$$

The target sends (P_2, τ_2) to the UE through L_1.

Step 5 (Decapsulation and key derivation): The UE verifies τ_2, decapsulates ct_T and restores $Z'_H = \text{XWing.Decaps}(ct_T, sk_U)$. Both participants agree on a canonical transcript

$$\text{tr} = \text{Encode}(C_{\text{HO}}, P_1, P_2, \tau_1, \tau_2). \tag{12}$$

They derive an extraction key and independent role-specific subkeys:

$$PRK = \text{HKDF-Extract}(MK, Z_H), \tag{13}$$

$$K_C = \text{HKDF-Expand}(PRK, \text{``}HO - CONF\text{''} \,||\, \text{H}(\text{tr}), 32) \tag{14}$$

$$K_{HO} = \text{HKDF-Expand}(PRK, \text{``}HO - ACT\text{''} \,||\, \text{H}(\text{tr}), 32) \tag{15}$$

Step 6 (Pre-handover confirmation): The UE sends $\tau_3 = \text{HMAC}_{K_C}(\texttt{"UE-CONF"} \,||\, \text{H}(\text{tr}))$. After verification, L_2 sends $\tau_4 = \text{HMAC}_{K_C}(\texttt{"T-CONF"} \,||\, \text{H}(\text{tr}))$. The UE accepts the preparation only after verifying τ_4. After confirmation, UE and L_2 retain PRK and K_{HO} until activation or expiry and erase ephemeral KEM material. The target maintains a single-use preparation state.

4.4. Fast Inter-Satellite Handover Phase

Step 7 (Attachment and key activation): After attachment to L_2, the UE generates a fresh activation nonce n_U^{A}, distinct from its preparation nonce n_U, and sends $(\text{sid}, n_U^{\text{A}}, \tau_5)$, where

$$\tau_5 = \text{HMAC}_{K_{HO}}(\text{``}UE - ACT\text{''} \,||\, \text{sid} \,||\, \text{tid} \,||\, n_U^{\text{A}}) \tag{16}$$

The target checks that the context is prepared, matches its own identity, has not expired, and has not been activated. The target generates a fresh target activation nonce n_T^A different from n_T, and responds with $(\text{sid}, n_T^A, \tau_6)$, where

$$\tau_6 = \text{HMAC}_{K_{HO}}(\text{``}T - ACT\text{''} \,||\, \text{sid} \,||\, n_U^{\text{A}} \,||\, n_T^{\text{A}}) \tag{17}$$

The UE verifies τ_6 under K_{HO}. Both parties derive the cipher key CK_{HO} and integrity key IK_{HO} using the stored extraction key and fresh activation nonces:

$$CK_{\mathrm{HO}} = \mathsf{HKDF\text{-}Expand}(PRK, \texttt{"HO-CK"} \,||\, \mathsf{sid} \,||\, n_U^{\mathrm{A}} \,||\, n_T^{\mathrm{A}}, 32), \tag{18}$$

$$IK_{\mathrm{HO}} = \mathsf{HKDF\text{-}Expand}(PRK, \texttt{"HO-IK"} \,||\, \mathsf{sid} \,||\, n_U^{\mathrm{A}} \,||\, n_T^{\mathrm{A}}, 32). \tag{19}$$

The UE sends (sid, τ_7) for final confirmation, where $\tau_7 = \mathsf{HMAC}_{IK_{\mathrm{HO}}}(\texttt{"UE-FIN"} \,||\, \mathsf{sid} \,||\, n_U^{\mathrm{A}} \,||\, n_T^{\mathrm{A}})$. The target verifies τ_7 under IK_{HO} before marking the context active. The three-message activation does not repeat X-Wing. The UE enters a pending state after τ_6; loss of τ_7 leaves the target pending, so authenticated completion can be retransmitted without creating a second session.

Step 8 (Completion and cleanup): Once L_2 confirms activation, the mobility-control function updates the serving relationship. The old serving security context is retired according to the underlying AS procedure, and used or expired preparation material is erased. A duplicate activation under the same sid cannot establish an independent live key. If a different satellite is selected, the context for the original target is canceled, and the procedure may proceed with an alternative candidate when a valid target-specific security context is available. An unauthenticated fallback to ECDH-only is disallowed.

4.5. Recovery and Scope

If preparation fails before the predicted transition and the UE remains connected to L_1, the procedure may switch to another eligible candidate. A prepared context may be used only for the target to which it is bound and only while it remains valid; a context for the failed target is not transferred to a different candidate. If a suitable alternative context is not available, additional target-specific preparation is required before activation. An unauthenticated fallback to an ECDH-only exchange is not permitted.

3GPP conditional handover supports one or more candidate cells and cancellation of unused candidates [16]. The optional use of an alternative candidate in our scheme depends on the availability of a valid target-specific context and is not a claim that 3GPP specifies the proposed X-Wing procedure. Switching to an already prepared candidate may reduce the additional delay following a failed preparation; otherwise, establishing the required context can add delay. Alternative-candidate operation and its latency have not been implemented or evaluated in the present study.

4.6. Selection of the Preparation Time

The proposed protocol uses an estimated target-attachment time $\widehat{t}_{HO}$ supplied by an existing operator mobility-management or satellite-visibility prediction mechanism. It does not compute satellite orbits, select a handover-prediction algorithm, or estimate the prediction error. Let T_P denote the time needed to complete HN authorization, X-Wing operations, message transmission, and pre-handover key confirmation. A possible deployment-level preparation-start rule is

$$t_{start} = \widehat{t}_{HO} - Q_{1-\epsilon}(T_P) - \Delta \tag{20}$$

where $Q_{1-\epsilon}(T_P)$ denotes a prospective $(1-\epsilon)$-quantile of the preparation duration, ϵ is a design target for the preparation-overrun probability, and Δ is a guard margin for prediction and scheduling uncertainty. These inputs require deployment-specific calibration. The implementation measures local cryptographic processing, so Equation (20) is a conditional scheduling guideline.

The contribution evaluated here is the placement of HN-authorized hybrid post-quantum key establishment before the request-critical handover stage, given an externally

selected target and sufficient preparation time. Orbit-based prediction accuracy, complete preparation time measurement, and radio-level failure behavior are outside the evaluated scope.

Algorithms 4, 5, 6 and 7 specify the HN, serving LEO, UE, and target LEO procedures, respectively, complementing the message ledger in Table 4.

Table 4. Predictive handover messages.

Phase	Sender → Receiver	Message	Receiver Verification
Prep.	$UE \to L_1$	(MR, t_{req}), AS protection	L_1: report integrity and freshness
Prep.	$L_1 \to HN$	PrepReq (IDs, MR, timestamps)	HN: source authentication and target eligibility
Prep.	$HN \to L_1, L_2$	C_{HO} (individually authenticated)	L_1, L_2: approval authenticity and expiry
Prep.	$L_1 \to L_2$	(C_{HO}, MK) (protected)	L_2: confidential transfer and context match
Prep.	$L_1 \to UE$	(C_{HO}, τ_0)	UE: derive MK and verify τ_0
KEM	$UE \to L_2$ via L_1	(P_1, τ_1); UE key and nonce	L_2: check context, n_U and τ_1
KEM	$L_2 \to UE$ via L_1	(P_2, τ_2); target KEM response	UE: check context and τ_2; decapsulate
Confirm	$UE \to L_2$ via L_1	(sid, τ_3)	L_2: verify τ_3 under K_C
Confirm	$L_2 \to UE$ via L_1	(sid, τ_4)	UE: verify τ_4 under K_C
Execute	$UE \to L_1 \to L_2$	Execution request; target acknowledgement via L_1	L_1, L_2: session, target and protected command
Activate	$UE \to L_2$	$(\text{sid}, n_U^{\text{A}}, \tau_5)$	L_2: context and τ_5 under K_{HO}
Activate	$L_2 \to UE$	$(\text{sid}, n_T^{\text{A}}, \tau_6)$	UE: τ_6; derive $CK_{\text{HO}}, IK_{\text{HO}}$
Activate	$UE \to L_2$	(sid, τ_7)	L_2: final MAC under IK_{HO}
Complete	$L_2 \to$ mobility control	Completion notice	Retire consumed context

Algorithm 4 HN: Target Authorization

Require: Protected PrepReq$(\text{uid}, \text{suid}, \text{tid}, MR, \widehat{t}_{HO}, t_{\text{req}})$
1: **if** source, UE, or target authorization fails **then**
2: **abort**
3: **end if**
4: Generate fresh sid, n_H; set t_C, exp, suite
5: $C_{\text{HO}} \leftarrow \text{Encode}(\text{sid}, \text{uid}, \text{suid}, \text{tid}, t_C, n_H, \text{exp}, \text{suite})$
6: Send C_{HO} to L_1 and L_2 on protected channels
7: Do not derive MK or Z_H

Algorithm 5 Serving LEO L_1: Preparation and Relay

Require: UE report (MR, t_{req}), shared $K_{\text{gNB}}^{(1)}$
1: Verify report integrity; select $L_2, \widehat{t}_{HO}$
2: Trigger preparation at t_{start} by (20)
3: Send protected PrepReq to HN; receive C_{HO}
4: Verify HN authorization and tid, exp, suite
5: Derive MK by (6); send (C_{HO}, MK) securely to L_2
6: Compute τ_0; send (C_{HO}, τ_0) to UE
7: Relay $(P_1, \tau_1), (P_2, \tau_2), \tau_3, \tau_4$ between UE and L_2
8: **if** UE requests handover and prepared context is valid **then**
9: Send authenticated execution command to L_2; coordinate transfer
10: **end if**
11: Never derive Z_H or PRK

Algorithm 6 UE: Hybrid Preparation and Activation

Require: Shared $K_{\text{gNB}}^{(1)}$ and a mobility report MR
1: Send protected (MR, t_{req}) to L_1; receive (C_{HO}, τ_0)
2: Derive MK; verify τ_0 and context fields, including exp
3: Generate n_U and $(sk_U, pk_U) \leftarrow \mathsf{XWing.KeyGen}()$
4: Form $P_1 = (\mathsf{sid}, \mathsf{uid}, \mathsf{tid}, n_U, pk_U)$; compute τ_1
5: Send (P_1, τ_1) to L_2 via L_1
6: Receive (P_2, τ_2); verify τ_2
7: $Z_H \leftarrow \mathsf{XWing.Decaps}(ct_T, sk_U)$; derive PRK, K_C, K_{HO}
8: Compute τ_3 using K_C; send τ_3; verify τ_4
9: **if** attached to approved L_2 with valid context **then**
10: Generate n_U^{A}; compute τ_5 using K_{HO}
11: Send $(\mathsf{sid}, n_U^{\text{A}}, \tau_5)$ to L_2
12: Receive (n_T^{A}, τ_6); verify τ_6
13: Derive $CK_{\text{HO}}, IK_{\text{HO}}$ by (18) and (19)
14: Compute τ_7 using IK_{HO}; send to L_2
15: **end if**
16: Erase ephemeral KEM secrets when no longer required

Algorithm 7 Target LEO L_2: Encapsulation and Acceptance

Require: HN approval C_{HO} and protected (C_{HO}, MK) from L_1
1: **if** approvals differ or context is expired/consumed **then**
2: **abort**
3: **end if**
4: Receive (P_1, τ_1); verify $\mathsf{tid}, \mathsf{sid}, \tau_1$
5: Generate n_T; $(Z_H, ct_T) \leftarrow \mathsf{XWing.Encaps}(pk_U)$
6: Form $P_2 = (\mathsf{sid}, n_T, ct_T)$; compute τ_2 using MK
7: Send (P_2, τ_2) to UE via L_1
8: Derive PRK, K_C, K_{HO}; verify received τ_3
9: Compute τ_4 using K_C; send to UE via L_1
10: Record confirmed, single-use preparation
11: Receive $(\mathsf{sid}, n_U^{\text{A}}, \tau_5)$ after UE attachment
12: **if** prepared, fresh, target-bound, and τ_5 valid **then**
13: Generate n_T^{A}; derive $CK_{\text{HO}}, IK_{\text{HO}}$
14: Compute τ_6 using K_{HO}; send $(\mathsf{sid}, n_T^{\text{A}}, \tau_6)$
15: Receive τ_7 from UE
16: **if** τ_7 is valid **then**
17: Mark active; notify mobility control
18: **end if**
19: **end if**
20: Reject independent re-activation; erase expired state

5. Security Analysis

The five security requirements in Table 2 address authenticated activation, key confidentiality, preparation binding, HN target authorization, and retrospective secrecy after disclosure of a completed session's AS key. The verification scope is narrower than the complete implementation: time-based expiration, durable replay state, and downgrade rejection are specified controls, not fully verified properties.

5.1. ProVerif Model and Source-Level Encoding

We specify the UE, serving LEO, target LEO, and home network as replicated applied-pi-calculus processes in ProVerif. A single access-stratum key K_{gNB} is reused across replicated logical handover instances. To represent the existing AS layer, the UE transmits a symbolic integrity-protected frame carrying the unchanged mobility-report payload, an AS

COUNT, the bearer identity, and the uplink direction. The integrity key is represented by the symbolic constructor `rrcIntKey(kg)`. The serving LEO verifies the frame authentication tag and accepts the report only if its COUNT matches the expected one-use receive state. The process generates a distinct symbolic COUNT for each logical handover instance; therefore, non-repetition and receive-state synchronization are assumed, not independently proved. All radio messages remain on the attacker-controlled public channel `pub`, while HN–LEO authorization and inter-satellite context-transfer links are private trusted channels. The ML-KEM-768 and X25519 operations, X-Wing combiner, HMAC, and HKDF remain symbolic abstractions rather than computational implementations.

Figure 2 shows representative declarations and process composition from the executed model.

```
fun rrcIntKey(bitstring):bitstring.
fun asFrame(bitstring,bitstring,bitstring):bitstring [data].
event ASProtectedSent(bitstring,bitstring).
event ASVerified(bitstring,bitstring).
event UEIssuedMR(bitstring,bitstring).
event SourceAcceptedMR(bitstring,bitstring).
process
  new kg:bitstring;
  !(
    new asCount:bitstring;
    new cSH:channel; new cHS:channel;
    new cHT:channel; new cST:channel;
    (UE(kg,asCount) |
     ServingLEO(kg,asCount,cSH,cHS,cST) |
     HN(cSH,cHS,cHT) | TargetLEO(cHT,cST))
  )
```

Figure 2. Representative declarations and four-role process composition from the executed ProVerif model.

5.2. Security Queries and Event Correspondences

Four private probes `probeZH`, `probeKHO`, `probeCK`, `probeIK` are each encrypted with Z_H, K_{HO}, CK_{HO}, IK_{HO}. The computed ciphertexts are sent via public channels. The inability of the attacker to recover these probes is a symbolic indication of secrecy for the corresponding keys under the abstract encryption model. No direct private session key name is queried. The actual queries use the `attacker` predicate, and the verification result `true` is returned for the non-attacker-known assertions.

We checked four symbolic probe-secrecy properties and eight event correspondence properties. The latter cover UE–target agreement, preparation-before-activation, target authorization, target-specific encapsulation, mobility-report provenance, and AS-frame freshness. The event-correspondence assertions use either non-injective correspondence for provenance/authorization or injective correspondence for one-to-one session and AS-frame matching. For instance, an accepted mobility report must correspond to a previously issued report, and each verified AS frame must be associated with a distinct AS-protected transmission. Representative queries are shown below Figure 3.

```
query attacker(probeZH).
query attacker(probeKHO).
query attacker(probeCK).
query attacker(probeIK).
query qMR1:bitstring, qTS1:bitstring;
  event(SourceAcceptedMR(qMR1,qTS1)) ==>
  event(UEIssuedMR(qMR1,qTS1)).
query qMR2:bitstring, qTS2:bitstring;
  inj-event(SourceAcceptedMR(qMR2,qTS2)) ==>
  inj-event(UEIssuedMR(qMR2,qTS2)).
query qCnt:bitstring, qPayload:bitstring;
  inj-event(ASVerified(qCnt,qPayload)) ==>
  inj-event(ASProtectedSent(qCnt,qPayload)).
```

Figure 3. Representative probe-secrecy, mobility-report, and AS-frame correspondence queries in ProVerif.

The post-completion model conditions release of the old $K_{gNB}^{(1)}$ on the private completion notification from `TargetAcceptedActivation`. The disclosure runs in ProVerif phase 1, and `attacker(probe...) phase 1` queries test retrospective secrecy. The separate event non-reachability diagnostic checks that `ASKeyDisclosed` can actually occur; a false non-reachability verdict is expected. This is a bounded compromise scope: the model excludes live session-secret leakage, operator-channel compromise, and reuse of one AS key across multiple replicated handovers.

5.3. Executed Verification Results

The ProVerif execution of the AS-freshness model completed without errors. All twelve queries returned positive verdicts: four probe-secrecy queries and eight event correspondences. The model includes replicated logical handovers sharing the same K_{gNB}. Under the explicitly modeled one-use AS receive-state and distinct-COUNT assumptions, both mobility-report correspondences and the injective AS-frame correspondence were satisfied. These results establish only the stated symbolic properties. In particular, the model does not implement the full PDCP anti-replay state machine or establish correct COUNT management under arbitrary radio reordering, key reuse across actual serving-node migrations, or active compromise of a serving satellite. The separately executed post-completion AS-key-disclosure model returned positive results for the four phase-1 probe-secrecy queries; its scope and assumptions differ from the main model.

Table 5 shows the security requirements and corresponding ProVerif assertions in the order of the SR sequence. All listed security assertions produced positive verification results under the respective model assumptions.

Table 5. Executed ProVerif queries and verification results.

SR	Verified Property	Query/Test Condition	Result
SR1	UE authentication at target	`qSid1`: injective target acceptance/UE acceptance	true
SR1	Target authentication at UE	`qSid2`: injective UE acceptance/target response	true
SR2	Hybrid-secret probe secrecy	`attacker(probeZH)`	true
SR2	Handover-key probe secrecy	`attacker(probeKHO)`	true
SR2	Encryption-key probe secrecy	`attacker(probeCK)`	true
SR2	Integrity-key probe secrecy	`attacker(probeIK)`	true
SR3	Mobility-report origin correspondence	`qMR1`: `SourceAcceptedMR/UEIssuedMR`	true
SR3	Injective mobility-report acceptance	`qMR2`: injective `SourceAcceptedMR/UEIssuedMR`	true
SR3	Injective AS-frame provenance	`qCnt`: injective `ASVerified/ASProtectedSent`	true
SR3	Preparation before activation	`qSid3`: injective activation start/preparation confirmed	true
SR4	HN target authorization	`qSid4`: target acceptance/HN authorization	true
SR4	Target-specific encapsulation	`qSid5`: target acceptance/encapsulation	true
SR5	Post-completion hybrid-secret probe secrecy	`attacker(probeZH) phase 1`	true
SR5	Post-completion handover-key probe secrecy	`attacker(probeKHO) phase 1`	true
SR5	Post-completion encryption-key probe secrecy	`attacker(probeCK) phase 1`	true
SR5	Post-completion integrity-key probe secrecy	`attacker(probeIK) phase 1`	true

The verified AS-freshness properties depend on the existing AS bearer's integrity checking and non-repeating COUNT with a one-use expected receive state. This formal model does not prove the operational PDCP receive-window algorithm, COUNT rollover, state synchronization among serving satellites, protection of expired contexts, duplicate handover-identifier rejection, algorithm-downgrade resistance, or computational X-Wing security. Post-completion AS-key-disclosure secrecy is evaluated separately under its own phase-1 model and key-erasure assumptions.

6. Performance Evaluation

6.1. Evaluation Scope and Common Accounting Rules

We compare five methods or standards: 3GPP Xn-based handover [15,16], a group-based FHAP [17], the Kim et al. [14] ECDH protocol, and on-demand X-Wing and predictive X-Wing. The 3GPP TS 33.501-based method used standardized key derivation computation; the 3GPP TS 38.300-based method used the Xn signaling phase. The two prior studies use their published message and operation counts. The on-demand and predictive X-Wing variants are implemented and measured on the same test platform. The implementation employs Advanced Encryption Standard in Galois/Counter Mode (AES-GCM) to protect inter-satellite *MK* transmission [33].

For on-demand X-Wing, HN authorization, target provisioning, hybrid exchange, and activation are performed. However, the predictive X-Wing variant completes authorization and hybrid exchange while the UE is attached to L_1. Upon the execution request, the coordination and activation phases are performed. Both methods employ the same cryptographic primitives.

For scheme s, the operation-normalized computation cost is

$$C_s = \sum_{k \in \mathcal{O}} n_{s,k}\, \widetilde{t}_k \tag{21}$$

where $n_{s,k}$ is the source-reported or measured number of operations of type k, and $\widetilde{t}_k$ is the corresponding median operation time measured in the OpenSSL environment. The reference operation counts are taken from 3GPP TS 33.501 (Annex A.11) and TS 38.300 for the key derivation case [15,16], and from Kim et al. [14] for the ECDH case. Both X-Wing implementations utilize a runtime computation counter. For Kim et al., we use the published 15 HMAC and 2 DH operations. For the FHAP, Yang et al. report $n + 5$ handover messages, $(3n + 2)$ hash computations, $(4n + 1)$ KDF calls, and one target-side restore computation for a group of n terminals [17]. Our example used $n = 10$ and maintains the general n dependent formulas. These are source-scoped, operation-weighted estimates rather than measured whole-protocol computation times.

6.2. Security Comparison

Table 6 compares the five methods for the same service categories. In the 3GPP category, NH denotes the next-hop key and NCC the next-hop chaining count. The evaluated ordinary Xn-based handover inherits the AS key management method from the existing 5G security context, including target-dependent key derivation and, when needed, the next-hop chaining procedure [15]. This baseline does not generate an ephemeral hybrid post-quantum secret. It represents ordinary Xn handover, not a separate conditional-handover configuration. Yang et al. [17] use pre-provisioned group credentials and RN-assisted authentication with traditional symmetric-cryptography computations. Kim et al. [14] use HMAC and an ephemeral ECDH for traditional authenticated key agreement. Both X-Wing variants use the HN-approved X-Wing exchange, target and session binding, nonces, and UE–target key confirmation. The difference lies in when the preparation phase takes place. The hybrid KEM confidentiality does not immediately ensure post-quantum peer authentication; the symbolic analysis result and assumptions are detailed in Section 5.

The trust architecture differs across these schemes: 3GPP handover relies on the established operator and AS security context, FHAP on its provisioned group credentials and relay infrastructure, and Kim et al. on their specified satellite authentication context. The two X-Wing variants instead assume HN authorization, trusted operator nodes and protected control links. Their C_{HO} and transcript bind the selected target and session, while the exchanged confirmation MACs explicitly check possession of the derived keys; the

other schemes should be interpreted according to their own message and key derivation definitions rather than assumed to provide identical confirmation semantics. In signaling, the predictive variant transfers the same hybrid material before the request and incurs additional forwarding and execution messages; it reduces request-stage bytes but increases all-stage bytes relative to the on-demand variant, as quantified below. The comparison assumes the trust and disclosure boundaries stated for each configuration; in our schemes, disclosure of an active *MK* or compromise of the serving node is outside the evaluated protection, whereas post-completion AS-key disclosure is assessed only under the separate symbolic model. For example, Lin et al. [34] embed reversible encryption within frequency-domain image compression. Such application-data protection is complementary to the access-layer handover and key establishment protocol considered here. These comparisons describe the evaluated protocol configurations rather than a universal ranking of deployment costs.

Table 6. Comparison of security features among handover schemes.

Scheme	Key Management	Freshness/Forward Secrecy	Hybrid KEM	Early Preparation
3GPP Xn-based handover [15,16]	5G AS context; target-bound $K_{\text{NG-RAN}}$	NH/NCC or source-key chaining; no ephemeral DH	X	X
Yang et al. [17]	Pre-provisioned group keys; RN-assisted MACs	Group nonces; classical key renewal	X	O
Kim et al. [14]	HMAC-secured UE–satellite handover	Nonces; ephemeral ECDH (classical)	X	X
On-Demand HO (X-Wing)	HN approval; UE–target confirmation	Bound nonces; ephemeral hybrid secret	O	X
Predictive HO (Our Proposed)	**HN approval; UE–target confirmation**	**Bound nonces; ephemeral hybrid secret**	**O**	**O**

6.3. Computational Cost

We use unit-operation benchmark costs for the following analytical comparison. Table 7 gives the representative median unit costs. HMAC uses a 96-byte input, and the ECDH proxy is X25519. Because the curve and input configuration differ from those of Kim et al. [14], X25519 serves as a reference DH operation rather than a reproduction of their implementation. For the FHAP comparison, a separate calibration uses 5000 samples and batches of 100 computations. We measure SHA-256 on 96-byte inputs, HMAC-SHA-256 on 96-byte inputs as a KDF proxy, and 512-bit to 256-bit modular reduction in OpenSSL as a proxy for C mod K_{D_j}. These proxies do not reproduce the original terminal hardware or all Chinese Remainder Theorem (CRT) operand lengths. Only packet authentication and key-confirmation HMAC calls are included in the stand-alone HMAC count; the HMAC operations executed internally by HKDF-Extract and HKDF-Expand are included only in their respective HKDF costs, without double-counting. The operation-weighted comparison below covers HMAC, HKDF and X-Wing, for which the corresponding unit costs are reported; the implementation additionally performs transcript SHA-256 and context-transfer AES-GCM operations, but those are excluded from this core-operation estimate because their corresponding protocol-specific unit times were not independently calibrated. The FHAP SHA-256 proxy in Table 7 is not a calibration of the longer handover transcript hash.

For a specific horizontal Xn handover key derivation path in TS 33.501, the source access node and UE each derive the target key using the activated K_{gNB} and target parameters. This procedure invokes the KDF twice under 3GPP Annex A.11, yielding a key derivation cost of 2.744 µs. The target node receives the derived key; the baseline model performs no additional third-key derivation. If an unused NH/NCC pair exists, the vertical key derivation function is applied implicitly in the standard [15]. This result considers only the key derivation procedure. Kim et al. proposed 15 HMAC and 2 DH computations for inter-satellite handover [14]. Substituting the median times of the HMAC and X25519 operations measured in this study yields $15(1.172) + 2(31.868) = 81.316$ µs (using the Table 7

displayed medians). This estimate excludes unreported key-generation cost. For FHAP, the published operation equation is $C_{\mathrm{FHAP}}(n) = (3n+2)t_H + (4n+1)t_{\mathrm{KDF}} + t_{\mathrm{DEC}}$ [17].

Table 7. Cryptographic operation times.

Operation	Median (µs)
HMAC-SHA-256 (96 B)	1.172
FHAP SHA-256 (96 B)	0.298
FHAP HMAC KDF (96 B)	1.030
FHAP CRT reduction (proxy)	0.137
TS 33.501 [15]	1.372
X25519 DH (proxy)	31.868
HKDF-Extract/Expand	1.071/1.642
X-Wing KeyGen	80.581
X-Wing Encaps/Decaps	118.889/80.531

Applying the SHA-256, HMAC and modular-reduction proxies yields an estimate of $C_{\mathrm{FHAP}}(10) = 51.903$ µs ($C_{\mathrm{FHAP}}(1) = 6.777$ µs). Yang et al. present $16.85n + 7.28$ µs (when $n = 10$, 175.78 µs) on a distinct terminal/access node platform. The modular-reduction proxy is not a bit-exact implementation of FHAP target recovery. In addition, the total time measured on the hardware of Yang et al. is not directly comparable with the normalized values of this study. Because of the different protocol scope and operation coverage, these estimates should not be compared directly with the whole-session runtime measurements of the X-Wing implementation.

Table 8 summarizes the request phase operation counts and corresponding weighted estimates. All stand-alone HMAC calls are integrated into one category; HKDF-internal HMAC work is included within the separately measured HKDF-Extract/Expand unit costs and is not counted a second time.

Table 8. Comparison of request-stage computational costs.

Scheme	Operations (Request Stage)	Core Cost (µs)
TS 33.501 [15]	2 derivations	2.744
Yang et al. ($n = 10$) [17]	32 H, 41 KDF, 1 CRT	51.903
Kim et al. [14]	15 HMAC, 2 DH	81.316
On-demand HO	30 HMAC; 4 HKDF-Extract; 10 HKDF-Expand 2 SHA-256; 1 AES-GCM Enc; 1 AES-GCM Dec 1 X-Wing KeyGen; 1 Encaps; 1 Decaps	335.865
Predictive HO (Our proposed)	**14 HMAC; 4 HKDF-Expand**	**22.976**

Note: The core cost weights HMAC, HKDF, and X-Wing only for the two matched variants; SHA-256 and AES-GCM calls are counted but not weighted. Internal HKDF HMAC calls are not included again in the separate HMAC count. Other schemes have different operation scopes.

Using the rounded unit costs in Table 7, the weighted request-stage core cost is 30(1.172) + 4(1.071) + 10(1.642) + 80.581 + 118.889 + 80.531 = 335.865 µs for on-demand handover and 14(1.172) + 4(1.642) = 22.976 µs for predictive handover, a reduction of 93.16%. Predictive preparation performs 18 HMAC, four HKDF-Extract, six HKDF-Expand, and one each X-Wing KeyGen/Encaps/Decaps; the stand-alone execution request accounts for additional HMAC work rather than removing the original exchanges, contributing 315.233 µs of weighted core work. The final completion message contributes two HMAC calls (2.344 µs) in either scheme. Thus, the corresponding all-stage core-onlyestimates are 338.209 and 340.553 µs; these are not full processing-time measurements and omit the

separately counted SHA-256 and AES-GCM work. Predictive preparation moves work out of the handover-critical interval rather than eliminating it.

6.4. Propagation Delay Analysis

We estimate communication latency by multiplying each link message count by its reference one-way propagation latency. This result is an accumulated link-latency estimate rather than an end-to-end latency measurement. The reference environment is a 600-km regenerative LEO NTN. According to 3GPP TR 38.821 (Section 7.1.1), the maximum service-link round-trip propagation is 12.89 ms, corresponding to one-way latency to the UE and LEO of 6.445 ms. According to Section 8.4.1.1, it reports approximately 6 ms one-way for a feeder link at 10-degree elevation. In Section 8.4.1.2, it implies approximately 10 ms one-way as a typical LEO-to-LEO inter-satellite (ISL) value [23]. Therefore,

$$T_{\text{prop}}^{(s)} = 6.445 N_{\text{UE-LEO}}^{(s)} + 10 N_{\text{ISL}}^{(s)} + 6 N_{\text{feeder}}^{(s)} \tag{22}$$

For comparison, LEO and geostationary Earth orbit (GEO) messages in Kim et al. [14] are not actual GEO propagation latency; rather, the LEO-to-LEO reference value is allocated as a simplifying proxy. FHAP used $n + 2$ UE-RN handover messages and three RN-TAN handover messages, corresponding to a reported group-based signaling size of $(32n + 150)$ bytes [17]. Since the UE-RN link is not an NTN service link but rather an onboard or link-for-railway link, this study assumes $d_{\text{UE}-\text{RN}} = 6.445$ ms and $d_{\text{RN}-\text{TAN}} = 10$ ms by defining a mapping between hypothetical topologies. The resulting proxy is $T_{\text{FHAP}}^{\text{proxy}}(n) = 6.445(n + 2) + 30$ ms and does not estimate FHAP deployment latency.

Table 9 compares the resulting reference-link propagation estimates. These inputs represent one fixed illustrative link-delay configuration; we do not sweep satellite altitude, elevation angle, inter-satellite distance, or link availability. The estimates are thus message-count-weighted reference-hop sums, not predictions of interruption time or geometry-specific end-to-end delay. Table 10 reports the measured request-stage and all-stage local processing times.

Table 9. Propagation delay estimate.

Scheme	UE–LEO	ISL	Feeder	Sum
TS 38.300 Xn [16]	2	2	0	32.890
Yang et al. ($n = 10$) [17]	12	3	0	107.340
Kim et al. [14]	4	2	0	45.780
On-demand HO	10	3	3	112.450
Predictive HO (Our proposed)	**5**	**2**	**0**	**52.225**

Following Yang et al. [17], the representative FHAP group ($n = 10$) has 15 handover messages and 470 bytes of data. Its 107.340-ms result in Table 9 is a cross-topology proxy, not an estimate of the proposed protocol or of actual FHAP deployment latency. The TS 38.300 Xn reference calculation includes Handover Request, Handover Request Acknowledge, and the radio resource control (RRC) messages RRCReconfiguration and RRCReconfigurationComplete; the two Path Switch messages are outside this four-message reference sum.

In Figure 8 in Kim et al. [14], messages (3$\tilde{6}$) include a total of 4 UE-LEO transmissions and 2 LEO-GEO transmissions before the final GEO completion notification. The corresponding proxy value is calculated as 45.780 ms. In the implementation `request-to-activation`, the on-demand method needed 16 link traversals, whereas the predictive method needed 7 link traversals. Table 11 shows the sender-transmitted data

volume (in bytes) for each phase and the contribution to the cumulative reference propagation delay. In particular, the predictive method for preparation phases involves 14 link traversals and a 106.670 ms delay before the UE request; both methods add 6.000 ms to the final completion notification. The cumulative reference propagation-hop sum is 118.450 ms for on-demand handover and 164.895 ms for predictive handover. This total result does not include queuing, terrestrial routing, retransmission, and the radio attachment time required.

Table 10. Handover processing time (μs).

Statistic	On-Demand HO	Predictive HO (Our Proposed)
Median	356.93	21.37
95th percentile	530.16	34.72
Total local median	359.22	362.07
Reduction of HO medians		94.01%

Table 11. Communication traffic and propagation contributions by phase.

Scheme	Sender	Pre (B)	HO (B)	Post (B)	Tx (B)	Rx (B)	Prop. (ms)
On-demand HO	UE/IoT gateway	0	1595	0	1595	1508	32.225
	Serving LEO (L_1)	0	546	0	546	293	38.890
	Target LEO (L_2)	0	1373	56	1429	1850	35.335
	HN/AC	0	256	0	256	175	12.000
	Total	**0**	**3770**	**56**	**3826**	**3826**	**118.450**
Predictive HO (Our Proposed)	UE/IoT gateway	1471	207	0	1678	1508	38.670
	Serving LEO (L_1)	3018	136	0	3154	2984	71.780
	Target LEO (L_2)	1246	127	56	1429	1850	42.445
	HN/AC	256	0	0	256	175	12.000
	Total	**5991**	**470**	**56**	**6517**	**6517**	**164.895**

Note: Pre = preparation; HO = UE request to key activation; Post = completion notification; Tx = transmitted bytes; Rx = received bytes; Prop. = reference propagation-hop contribution.

The two corresponding X-Wing implementation cases show that the request-phase propagation sum changes from 112.450 ms to 52.225 ms; this result presents a 53.56% analytical reduction under the fixed-hop-delay model. This reduction reflects the advance authorization and key preparation process, rather than the overall signaling reduction. It is conditional on the stated reference per-hop delays. Upon unit conversion, if operation-weighted request-stage local cryptography addition is used as an optional accounting metric, the request-phase nominal figures are 112.786 ms and 52.248 ms, respectively. These nominal sums are model-based estimates; the improvement rate is presented only between the matched X-Wing implementations.

6.5. Measured Handover Performance

The four roles (UE, serving L_1, target L_2, and terrestrial HN) are emulated in a single-process implementation with OpenSSL 3.5.5, HMAC-SHA-256, HKDF-SHA-256, X-Wing, and AES-256-GCM protection of the inter-LEO *MK* transfer. Application frames use an experimental type–length–value (TLV) format, with a one-byte message type and two-byte length per encoded field. In predictive preparation, UE–target exchanges are forwarded through L_1 and counted once per modeled link traversal (not per real IP-router hop). The implementation checks successful key establishment and selected malformed-message cases.

Measurements used a Linux x86-64 virtualized environment with an AMD EPYC 9V74 CPU, a four-vCPU-equivalent quota, a 4-GiB memory limit, and single-

CPU affinity. The implementation was built using GCC 14.2.0 with `-O2 -Wall -Wno-deprecated-declarations -std=c11` and OpenSSL 3.5.5. X-Wing uses the OpenSSL EVP interfaces for X25519 and ML-KEM-768. The protocol benchmark used 100 untimed warm-up sessions (50 per variant), and the X-Wing primitive benchmark used 200 warm-up cycles. Each variant comprises 15,000 completed sessions. The request-stage timer starts at the UE execution request and ends at target-key activation; it measures local processing and excludes link propagation, core-network routing, radio attachment, and orbital effects.

For the request-stage measurements, the on-demand and predictive interquartile intervals are 344.198–401.687 and 21.192–21.963 μs, respectively. Percentile bootstrap 95% confidence intervals for the medians are [356.597, 357.277] and [21.362, 21.382] μs (2000 resamples). These intervals assume independent session samples and may underestimate variation from temporal correlation and VM load.

Although predictive handover reduces the median request-stage processing time by 94.01%, its total local processing time is slightly higher than that of on-demand handover (362.07 versus 359.22 μs). This is because the X-Wing key establishment operations are performed during preparation rather than eliminated. The additional preparation and execution procedures may also introduce a small processing overhead. Therefore, the proposed scheme reduces processing time during the handover-critical stage rather than the overall computational workload.

Table 11 presents the transmission of each sender in each phase and the contribution to the propagation latency for each method. Each link traversal in the model is allocated signaling based on an original message record. For each modeled link traversal, the same frame length is counted as transmitted (Tx) at the sender and received (Rx) at the receiver. The overall Tx and Rx totals represent the same traffic from the sender and receiver perspectives and are not added together. Messages relayed by the serving LEO are counted on both its inbound and outbound links. The Prop. entries quantify the contribution of each sender to the cumulative reference propagation-hop sum, not measured per-device delays. Phase-level and all-stage totals follow the message ledger. On-demandhas no preparation messages: all 3770 bytes before key activation are transmitted after the UE request, followed by a 56-byte target-to-HN completion notification. In contrast, predictive transfers 5991 bytes while the gateway remains attached to L_1, leaving 470 bytes for the request-to-activation interval and the same 56-byte completion notification. Accordingly, the predictive all-stage total is higher (6517 versus 3826 bytes, an increase of 2691 bytes or 70.33%), although its request-stage traffic is lower by 3300 bytes (87.53%).

For the UE/IoT gateway, the 1595 bytes transmitted in the on-demand scheme consist of a 109-byte handover request/measurement, a 1309-byte ephemeral X-Wing public-key message, a 53-byte UE confirmation, and two activation messages of 71 and 53 bytes. In the predictive scheme, the same 109-, 1,309-, and 53-byte messages are sent during preparation (1471 bytes), whereas a separate 83-byte execution request and the two activation messages (207 bytes) are sent during handover; accordingly, the total UE transmissions increase by only 83 bytes, to 1678 bytes. The UE receives 1508 bytes in either scheme, yielding 3103 and 3186 bytes of bidirectional gateway-facing signaling, respectively (a 2.7% increase).

For the serving LEO L_1, on-demand forwarding and control transmissions total 546 bytes: an HN request (119 bytes), a protected MK transfer to L_2 (163 bytes), a context message to the UE (128 bytes), a source-to-target handover request (73 bytes), and a source-to-UE command (63 bytes). The predictive scheme preserves the same 410 bytes of preparation-related HN/MK/context signaling and 136 bytes of handover control but additionally forwards the 1309-byte X-Wing public key of the UE to L_2, the 1193-byte target ciphertext to the UE, and the two 53-byte confirmation messages in their respective directions. These four additional forwarding transmissions account for exactly

$1309 + 1193 + 53 + 53 = 2608$ bytes; L_1 therefore transmits 3154 instead of 546 bytes. Each forwarded frame is counted on its original UE and L_1 or L_2 and L_1 leg and on its outgoing L_1 leg, rather than as a single logical message.

The target LEO L_2 transmits the same 1429 bytes in both schemes: the 1193-byte X-Wing ciphertext, a 53-byte key confirmation, a 56-byte handover acknowledgement, a 71-byte activation response, and a 56-byte HN completion notification. The predictive scheme moves the ciphertext and key confirmation (1246 bytes) into preparation, leaving 127 bytes for activation-related signaling during the request stage, whereas the on-demand scheme transmits 1373 bytes during the request stage and 56 bytes afterward. Similarly, the HN/AC transmits two 128-byte approval/context messages in both schemes (256 bytes), during the request interval for the on-demand scheme and during preparation for the predictive scheme.

Therefore, the 2691-byte all-stage traffic increase corresponds exactly to the 2608 forwarding bytes of the serving LEO plus the additional 83-byte execution request of the UE; the HN and the target LEO do not increase their total transmitted bytes. The predictive preparation phase involves 14 link traversals and 106.670 ms of cumulative reference propagation, and its request stage involves seven traversals and 52.225 ms; the on-demand scheme has no preparation traversals and requires 16 request-stage traversals and 112.450 ms. Both schemes add a 6.000-ms completion notification. Consequently, the all-stage cumulative propagation increases from 118.450 to 164.895 ms, whereas the modeled request-stage propagation decreases by 53.56%. This is consistent with the measured reduction in local request-stage processing from 356.93 to 21.37 µs (94.01%). The former is a reference hop-delay sum, and the latter is measured local software processing; neither measures end-to-end radio interruption. Predictive preparation thus shifts work to the period of continued serving-LEO access rather than reducing the total amount of cryptographic signaling. The reported gains assume that the operator provides a suitable target estimate and that preparation completes before attachment; they do not quantify the likelihood of these conditions.

7. Conclusions

We presented a hybrid post-quantum key establishment scheme for predictive handover in 6G satellite IoT networks. The scheme prepares a target-bound hybrid secret before satellite handover and activates the resulting context through an authenticated exchange after attachment. For the evaluated single-target case, moving the hybrid exchange to the preparation stage reduced median request-stage local processing from 356.93 to 21.37 µs (94.01%) and request-stage traffic from 3770 to 470 bytes (87.53%) relative to the same-platform on-demand variant. Under the fixed reference-hop model, the request-stage propagation sum decreased from 112.450 to 52.225 ms, while total traffic increased from 3826 to 6517 bytes and the all-stage propagation sum increased from 118.450 to 164.895 ms. These results show that predictive preparation shifts hybrid key establishment and its associated signaling to the period before handover, reducing the processing and communication burden during the request stage despite the increased overall cost.

Under the stated honest-node, protected-channel, and AS freshness assumptions, ProVerif returned positive results for four probe-secrecy and eight correspondence queries. A separate post-completion AS-key disclosure model returned positive results for four probe-secrecy queries under its own assumptions. These symbolic results do not establish the computational security of X-Wing or security against an actively compromised serving satellite. Future work will use a hardware and communication testbed to assess practical operation under representative network conditions and investigate optimization across multiple handover candidates.

https://doi.org/10.3390/app162311847

Author Contributions: Conceptualization, J.K.; methodology, S.L. and J.K.; validation, S.L. and J.K.; formal analysis, J.K.; investigation, S.L.; data curation, S.L. and J.K.; writing—original draft, S.L. and J.K.; writing—review and editing, S.L. and J.K.; visualization, S.L.; supervision, J.K.; project administration, J.K.; funding acquisition, J.K. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea Government (Ministry of Science and ICT, MSIT) (No. RS-2026-25473237) and the Glocal University 30 Project Fund of Gyeongsang National University in 2026.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data presented in this study are available on request from the corresponding author.

Acknowledgments: During the preparation of this manuscript, the authors used OpenAI (San Francisco, CA, USA) ChatGPT (GPT-6) to assist with English-language editing, including improvements to grammar, clarity, readability, and terminological consistency. The authors critically reviewed, revised, and verified all AI-assisted content and took full responsibility for the content and conclusions of the final manuscript.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Giordani, M.; Zorzi, M. Non-Terrestrial Networks in the 6G Era: Challenges and Opportunities. *IEEE Netw.* **2021**, *35*, 244–251. https://doi.org/10.1109/MNET.011.2000493.
2. Lin, X.; Rommer, S.; Euler, S.; Yavuz, E.A.; Karlsson, R.S. 5G From Space: An Overview of 3GPP Non-Terrestrial Networks. *IEEE Commun. Stand. Mag.* **2021**, *5*, 147–153. https://doi.org/10.1109/MCOMSTD.011.2100038.
3. Darwish, T.; Kurt, G.K.; Yanikomeroglu, H.; Bellemare, M.; Lamontagne, G. LEO Satellites in 5G and Beyond Networks: A Review From a Standardization Perspective. *IEEE Access* **2022**, *10*, 35040–35060. https://doi.org/10.1109/ACCESS.2022.3162243.
4. Saad, M.M.; Tariq, M.A.; Khan, M.T.R.; Kim, D. Non-Terrestrial Networks: An Overview of 3GPP Release 17 & 18. *IEEE Internet Things Mag.* **2024**, *7*, 20–26. https://doi.org/10.1109/IOTM.001.2300154.
5. Juan, E.; Lauridsen, M.; Wigard, J.; Mogensen, P. Handover Solutions for 5G Low-Earth Orbit Satellite Networks. *IEEE Access* **2022**, *10*, 93309–93325. https://doi.org/10.1109/ACCESS.2022.3203189.
6. GSMA. *Post-Quantum Cryptography for Non-Terrestrial Networks*; Industry Report, PQ.07; Global System for Mobile Communications Association: London, UK, 2026.
7. Shor, P.W. Algorithms for Quantum Computation: Discrete Logarithms and Factoring. In Proceedings of the IEEE Symposium on Foundations of Computer Science, Santa Fe, NM, USA, 20–22 November 1994; pp. 124–134. https://doi.org/10.1109/SFCS.1994.365700.
8. NIST. *FIPS 203: Module-Lattice-Based Key-Encapsulation Mechanism Standard*; Technical report; National Institute of Standards and Technology: Gaithersburg, MD, USA, 2024. https://doi.org/10.6028/NIST.FIPS.203.
9. Alagic, G.; Barker, E.; Chen, L.; Moody, D.; Robinson, A.; Silberg, H.; Waller, N. *NIST SP 800-227: Recommendations for Key-Encapsulation Mechanisms*; Technical Report; National Institute of Standards and Technology: Gaithersburg, MD, USA, 2025. https://doi.org/10.6028/NIST.SP.800-227.
10. Bos, J.W.; Ducas, L.; Kiltz, E.; Lepoint, T.; Lyubashevsky, V.; Schanck, J.M.; Schwabe, P.; Seiler, G.; Stehlé, D. CRYSTALS-Kyber: A CCA-Secure Module-Lattice-Based KEM. In Proceedings of the IEEE European Symposium on Security and Privacy (EuroS&P), London, UK, 24–26 April 2018; pp. 353–367. https://doi.org/10.1109/EuroSP.2018.00032.
11. Ko, Y.; Pawana, I.W.A.J.; Kwon, H.; Shin, S.; Kim, J.; Baek, J.; You, I. 5G-AKA-HPQC: Hybrid Postquantum Cryptography Protocol for Quantum-Resilient 5G Primary Authentication with Forward Secrecy. *IEEE Internet Things J.* **2026**, *13*, 26977–27000. https://doi.org/10.1109/JIOT.2026.3678536.
12. Connolly, D.; Schwabe, P.; Westerbaan, B.E. X-Wing: General-Purpose Hybrid Post-Quantum KEM. Internet-Draft draft-connolly-cfrg-xwing-kem-11, IETF Crypto Forum Research Group, 2026. Available online: https://datatracker.ietf.org/doc/draft-connolly-cfrg-xwing-kem/11/ (accessed on 26 November 2026).
13. Langley, A.; Hamburg, M.; Turner, S. *Elliptic Curves for Security*; RFC 7748; Internet Engineering Task Force (IETF): Fremont, CA, USA, 2016. https://doi.org/10.17487/RFC7748.

14. Kim, J.; Lee, S.; Ko, Y.; You, I. Semantic-Aware Security Protocol Design for 6G Non-Terrestrial Networks: A Formal Verification Approach. *Int. J. Semant. Web Inf. Syst.* **2025**, *21*, 1–38. https://doi.org/10.4018/IJSWIS.379723.
15. 3GPP. *TS 33.501: Security Architecture and Procedures for 5G System*; Technical Specification; 3rd Generation Partnership Project (3GPP): Sophia Antipolis, France, 2026.
16. 3GPP. *TS 38.300: NR and NG-RAN Overall Description; Stage-2*; Technical Specification; 3rd Generation Partnership Project (3GPP): Sophia Antipolis, France, 2026.
17. Yang, Y.; Cao, J.; Ma, R.; Cheng, L.; Chen, L.; Niu, B.; Li, H. FHAP: Fast Handover Authentication Protocol for High-Speed Mobile Terminals in 5G Satellite–Terrestrial-Integrated Networks. *IEEE Internet Things J.* **2023**, *10*, 13959–13973. https://doi.org/10.1109/JIOT.2023.3262933.
18. Arkko, J.; Lehtovirta, V.; Torvinen, V.; Eronen, P. *Improved Extensible Authentication Protocol Method for 3GPP Mobile Network Authentication and Key Agreement (EAP-AKA')*; RFC 9048; Internet Engineering Task Force (IETF): Fremont, CA, USA, 2021. https://doi.org/10.17487/RFC9048.
19. 3GPP. *TS 23.502: Procedures for the 5G System (5GS)*; Technical Specification; 3rd Generation Partnership Project (3GPP): Sophia Antipolis, France, 2026.
20. Mattsson, J.P.; Sethi, M. *EAP-TLS 1.3: Using the Extensible Authentication Protocol with TLS 1.3*; RFC 9190; Internet Engineering Task Force (IETF): Fremont, CA, USA, 2022. https://doi.org/10.17487/RFC9190.
21. Tschofenig, H.; Kroeselberg, D.; Pashalidis, A.; Ohba, Y.; Bersani, F. *The Extensible Authentication Protocol–Internet Key Exchange Protocol Version 2 (EAP-IKEv2) Method*; RFC (Experimental) 5106; Internet Engineering Task Force (IETF): Fremont, CA, USA, 2008. https://doi.org/10.17487/RFC5106.
22. Kwiatkowski, K.; Kampanakis, P.; Westerbaan, B.E.; Stebila, D. *Post-Quantum Traditional (PQ/T) Hybrid Key Agreement Mechanisms for TLS 1.3*; RFC 10024; Internet Engineering Task Force (IETF): Fremont, CA, USA, 2026. https://doi.org/10.17487/RFC10024.
23. 3rd Generation Partnership Project (3GPP). *Solutions for NR to Support Non-Terrestrial Networks (NTN)*; Technical Report TR 38.821; Version 16.0.0 (Dec. 2019) 3rd Generation Partnership Project (3GPP): Sophia Antipolis, France, 2019.
24. 3GPP. *TS 23.501: System Architecture for the 5G System (5GS)*; Technical Specification; 3rd Generation Partnership Project (3GPP): Sophia Antipolis, France, 2026.
25. Palaniswamy, B.; Karati, A.; Chen, T.Y.; Das, A.K.; Bhargava, B.K. QPCASIN: A Quantum-Defended Privacy-Aware Preemptive Handover-Enabled Continuous Authentication in Space Information Networks. *IEEE Trans. Inf. Forensics Secur.* **2025**, *20*, 6910–6922. https://doi.org/10.1109/TIFS.2025.3583246.
26. Mamun, Q.; Ho, T.D.; Pan, Z.; Shimamoto, S. SkyLock: A Unified Multi-Layer Authentication and Key Management Framework for 6G Non-Terrestrial Networks. In Proceedings of the IEEE Consumer Communications & Networking Conference (CCNC), Las Vegas, NV, USA, 9–12 January 2026. https://doi.org/10.1109/CCNC65079.2026.11366345.
27. Krawczyk, H.; Bellare, M.; Canetti, R. *HMAC: Keyed-Hashing for Message Authentication*; RFC 2104; Internet Engineering Task Force (IETF): Fremont, CA, USA, 1997. https://doi.org/10.17487/RFC2104.
28. Krawczyk, H.; Eronen, P. *HMAC-Based Extract-and-Expand Key Derivation Function (HKDF)*; RFC 5869; Internet Engineering Task Force (IETF): Fremont, CA, USA, 2010. https://doi.org/10.17487/RFC5869.
29. Barnes, R.; Bhargavan, K.; Lipp, B.; Wood, C. *Hybrid Public Key Encryption*; RFC 9180; Internet Engineering Task Force (IETF): Fremont, CA, USA, 2022. https://doi.org/10.17487/RFC9180.
30. Dolev, D.; Yao, A.C. On the Security of Public Key Protocols. *IEEE Trans. Inf. Theory* **1983**, *29*, 198–208. https://doi.org/10.1109/TIT.1983.1056650.
31. Blanchet, B. Modeling and Verifying Security Protocols with the Applied Pi Calculus and ProVerif. *Found. Trends Priv. Secur.* **2016**, *1*, 1–135. https://doi.org/10.1561/3300000004.
32. Kim, J.; Duguma, D.G.; Astillo, P.V.; Park, H.; Kim, B.; You, I.; Sharma, V. A Formally Verified Security Scheme for Inter-gNB-DU Handover in 5G Vehicle-to-Everything. *IEEE Access* **2021**, *9*, 119100–119117. https://doi.org/10.1109/ACCESS.2021.3107308.
33. Dworkin, M. *Recommendation for Block Cipher Modes of Operation: Galois/Counter Mode (GCM) and GMAC*; NIST Special Publication 800-38D; National Institute of Standards and Technology: Gaithersburg, MD, USA, 2007. https://doi.org/10.6028/NIST.SP.800-38D.
34. Lin, Y.; Chen, D.; Zhang, Y.; Li, Y.; Erkan, U.; Toktas, A.; Gao, S. Beyond the Trade-Off: Rate-Efficient and Quality-Preserving Encryption-Embedded Frequency-Domain Image Compression. *Inf. Process. Manag.* **2027**, *64*, 105277. https://doi.org/10.1016/j.ipm.2026.105277.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.