

LPG: Raise Your Location Privacy Game in Direct-to-Cell LEO Satellite Networks

Quan Shi^{1*}, Liying Wang^{2*}, Prosanta Gope^{3†}, Qi Liang⁴, Haowen Wang⁴, Qirui Liu², Chenren Xu^{2,5},
Shanguang Wang⁴, Qing Li^{4†}, Biplab Sikdar¹

¹National University of Singapore ²Peking University

³University of Sheffield ⁴Beijing University of Posts and Telecommunications

⁵Key Laboratory of High Confidence Software Technologies, Ministry of Education (PKU)

Abstract

Multi-tenant direct-to-cell (D2C) Low Earth Orbit (LEO) satellite networks pose significant risks to users' location privacy by linking Mobile Network Operator (MNO)- managed identities with Satellite Network Operator (SNO)- visible locations. Existing privacy solutions are ill-suited to the resource-constrained hardware and orbital dynamics of these satellite environments. We present LPG (Location Privacy Game), the first protocol-layer solution offering user-configurable location privacy for D2C LEO. LPG achieves this via identity-location decoupling: SNOs provide connectivity without visibility of user identity, while MNOs manage service and billing without access to precise location information. LPG enables offline secure authentication and key agreement without revealing user identity to satellites, supports user-configurable location disclosure at chosen geographic granularity for essential service needs, and ensures fair billing between MNOs and SNOs through privacy-preserving settlement. Our implementation on a real-world in-orbit LEO satellite and commercial mobile phones demonstrates that LPG is practical and viable in resource-constrained, highly-dynamic LEO environments.

1 Introduction

Direct-to-cell (D2C) Low Earth Orbit (LEO) satellite networks [58, 65, 68] are transforming global connectivity by enabling standard consumer devices (e.g., phones, IoT terminals) to seamlessly access satellite connectivity, especially in remote areas. A key enabler for this expansion is the multi-tenant model [4, 31, 50] (Figure 1), where Mobile Network Operators (MNOs) lease satellite infrastructure from Satellite Network Operators (SNOs), addressing key deployment barriers: orbital slot scarcity [20], spectrum licensing constraints [49], and prohibitive constellation costs [12]. This partnership, supported by regulators [21, 74] and industry [9, 11, 14, 16, 33, 34, 66, 71], enables MNOs to extend




Figure 1: In direct-to-cell multi-tenant LEO satellite architecture, SNOs and MNOs constitute a *dual-entity* setting. Our LPG provides *user configurable location disclosure*.

coverage without satellite deployment while granting SNOs access to licensed spectrum and subscribers.

However, this multi-tenant D2C integration creates severe location privacy risks. By migrating cellular functions like Radio Access Networks (RAN) and Core Networks (CN) onto SNO-managed satellites [38, 41, 50, 68, 72, 77–79], the architecture inherently couples MNO-managed identities with SNO-visible locations (§2.1). This coupling grants SNOs powerful capabilities to monitor user movements via physical triangulation, signal characteristic, or protocol metadata [29, 32, 35, 38–42], while enabling MNOs to reconstruct trajectories from settlement reports [45] (§2.2). The transnational nature of D2C multi-tenancy amplifies these risks, leading to severe consequences, such as massive fines for location data misuse [24, 51, 75], and threats to the individuals in conflict zones [23, 35]. The growing geopolitical importance of the satellite internet has made user geolocation an urgent concern for citizens and governments [27]. While prior works have highlighted physical-layer localization risks [32, 35, 36], the severe location privacy concerns at the cellular protocol layer remain largely unaddressed. This underscores the critical need for protocol-layer D2C location privacy.

Achieving this robust privacy, however, is in fundamental tension with the operational demands of satellite D2C service. The extreme orbital dynamics, limited onboard resources and the transnational nature [38, 41, 79] require a solution that is lightweight, resilient, and offers flexible, non-binary privacy (§2.2). Crucially, such a solution must be implemented

*Both authors contributed equally to this research.

†Corresponding authors: p.gope@sheffield.ac.uk, qingli@bupt.edu.cn

Table 1: Comparison of the state-of-the-art terrestrial and satellite 4G/5G network architectures.

Scheme	OE	Operator	Broker	Location Privacy			Functionality		Viability		P1: Identity-Location Decoupling P2: Flexible Location Disclosure P3: Session Unlinkability F1: Fair Billing and Fraud Prevention F2: ID-based Services with MNO Control V1: Resource-Constrained Operation V2: Resilience to Orbital Dynamics ●: Support, ○: Does Not Support OE: Operating Environment
				P1	P2	P3	F1	F2	V1	V2	
MVNO [2, 80]	Terrestrial	MNO	MVNO	○	○	●	○	●	○	○	
PGPP [61]	Terrestrial	MNO	MVNO	●	○	●	○	○	●	○	
LOCA [45]	Terrestrial	MNO	MVNO	●	○	●	●	●	○	○	
Starlink [69]	D2C LEO	SNO	MNO	○	○	○	●	●	●	○	
SpaceCore [38]	D2C LEO	SNO	MNO	○	○	○	●	●	●	○	
MOSAIC [41]	D2C LEO	SNO	MNO	○	○	○	○	○	●	●	
Ours (LPG)	D2C LEO	SNO	MNO	●	●	●	●	●	●	●	

without compromising basic network functions such as connectivity, identity-based services, and billing fairness [41, 61].

Therefore, we systematically identify that an effective D2C LEO location privacy framework must be holistically co-designed to achieve three objectives: (1) *Location Privacy*, preventing entities (MNOs, SNOs) tracking users while allowing flexible disclosure; (2) *Service Functionality*, maintaining essential cellular operations like fair billing, QoS, etc.; and (3) *System Viability*, ensuring practical deployment on resource-constrained satellites within dynamic orbital environments. We have established the following set of requirements:

Category 1: Robust Protocol-Layer Location Privacy.

◦ *P1: Identity-Location Decoupling.* The framework must ensure: (1) SNOs, while providing connectivity, cannot access or infer the user’s stable identity. (2) MNOs, while managing user identity-based services, cannot access the user’s precise real-time location or trajectory.

◦ *P2: Flexible Location Disclosure.* Compared to binary privacy models [45, 60, 61, 80], the D2C system needs to empower users to prove their presence within geographic boundaries at selectable levels of granularity. In cross-jurisdictional multi-tenant architecture, this is crucial for balancing user privacy with operational necessities such as regulatory compliance [31, 49, 76] and emergency service [22].

◦ *P3: Session Unlinkability.* To prevent entities from reconstructing user trajectories (even for anonymous users), the framework must ensure that the same user’s distinct service sessions and cryptographic primitives are unlinkable. This prevents the correlation of protocol-layer metadata by entities across multiple sessions to infer user trajectory.

Category 2: Uncompromised Service Functionality.

◦ *F1: Fair Billing and Fraud Prevention.* The system must incorporate mechanisms for fair and accountable billing. This includes preventing SNOs from falsifying service claims (i.e. overbill) to MNOs, preventing MNOs from falsifying settlement (i.e. repudiation), and preventing UEs from gaining unauthorised access through token forgery or replay.

◦ *F2: Identity-based Services with MNO Control.* The framework must allow MNOs to manage identity-based policies and services (e.g., QoS tiers, billing, and lawful intercept) without requiring MNOs to access precise UE location data.

Category 3: D2C LEO System Viability.

◦ *V1: Resource-Constrained Operation.* Given the severe





Figure 2: BUPT-3: Our in-orbit LEO satellite. Figure 3: An overview of cellular service procedure in LPG

limitations in computational power [53, 54, 64, 78, 79], energy [43, 73], and thermal dissipation [37, 73, 79] for operating in space, all privacy-preserving mechanisms must be lightweight and efficient for the operator SNO.

◦ *V2: Resilience to Orbital Dynamics.* Due to the extreme velocity of LEO satellites ($\sim 7.9\text{km/s}$), the solution must maintain its functionality and security guarantees despite the dynamic satellite topology, intermittent satellite-ground connectivity, and potentially immature Inter-Satellite Links (ISLs) of LEO constellations [38, 41, 72, 78]. This requires that the SNO and UE could establish an offline attachment without contacting the remote MNO.

While various multi-tenant architectures have been proposed for cellular privacy, a systematic comparison (Table 1) shows that no existing approach holistically satisfies all the critical requirements for D2C LEO networks.

Existing terrestrial cellular architectures, multi-tenant MVNO-MNO partnerships [2] and privacy solutions (e.g., PGPP [61], LOCA [45], PGUS [80]), though achieve identity-location decoupling between entities (MNO, MVNO), they typically assume continuous infrastructure connectivity and sufficient computational resources (e.g., for ZK-SNARK proof [45]). These assumptions are fundamentally incompatible with the intermittent connectivity and resource-constrained onboard Commercial Off-The-Shelf (COTS) computing environment of D2C LEO satellites [37, 78, 79]. Furthermore, many adopt all-or-nothing binary privacy models [45, 61, 80] for location, which cannot flexibly balance operational requirements and privacy.

Simultaneously, state-of-the-art D2C LEO satellite architectures (e.g., Starlink [66, 69], SpaceCore [38], MOSAIC [41]) primarily focus on improving connectivity and efficiency. Although these works use anonymous tokens for pay-as-you-go services, they generally lack guarantees against fine-grained

trajectory reconstruction. The ability of entities to link tokens, metadata, movement patterns, and settlement reports to reconstruct trajectories at the protocol layer remains a significant concern. Furthermore, these solutions do not prevent SNO overbilling, which raises fraudulent billing risks.

Contribution. We present **LPG** (Location Privacy Game). To the best of our knowledge, this is the *first framework to preserve location privacy* at the protocol layer in multi-tenant D2C LEO networks. Our contributions are:

(1) We introduce *a novel protocol-layer approach for Identity-Location decoupling* in D2C satellite networks, which is the first framework to enable SNOs to provide connectivity without accessing user identities while allowing MNOs to manage subscriber services without visibility into precise locations.

(2) Based on identity-location decoupling, our proposed LPG framework transcends the binary privacy model, uniquely supporting *User-Configurable Location Disclosure*, which balances privacy with essential cross-border regulatory and emergency service requirements.

(3) We develop specialised Dual Token mechanism (Dual-Token), Offline Authentication and Key Agreement (OAKA) protocol, Session Management Tables (SMTs) and Scalable Collaborative Aggregate Claiming (SCAC) optimised for resource-constrained and highly dynamic LEO environment, while ensuring session unlinkability and verifiable fair billing (§4). An overview is shown in Figure 3.

(4) We implement LPG on a real in-orbit LEO satellite equipped with COTS hardware (Figure 2) and commercial mobile phones, demonstrating its practicality under actual orbital conditions. Our evaluations show that LPG achieves flexible location privacy with minimal overhead, providing a viable path toward location privacy-preserving multi-tenant satellite operations (§5).

2 Background

2.1 Multi-tenant D2C Architecture

D2C satellite networks extend cellular connectivity to standard consumer devices via orbiting satellites that implement 4G/5G functions. The *multi-tenant* model (MNOs leasing SNO infrastructure) addresses deployment barriers like orbital slot scarcity [20], spectrum licensing [49], and constellation costs [12]. While the division of functions between SNOs and MNOs is critical, our core privacy goal is to prevent any single entity from linking a user’s identity with their precise location. However, as D2C architectures evolve to improve connectivity by migrating more cellular functions (traditionally handled within the MNO’s domain) onto the SNO satellite platform, they inherently couple user identity and location, enabling real-time tracking. Figure 4 illustrates this evolution:

◦ *Transparent Pipe* (Figure 4A): Early D2C systems (e.g. Iridium [58], Globalstar [65], ViaSat [71]) use satellites as RF relays. This model offers strong physical separation, but suffers from severe limitations, including limited coverage, no 4G/5G availability, and backhaul saturation [38].

◦ *Onboard RAN-DU or Full RAN* (Figure 4 B1/B2): Modern architectures like Starlink [68] (Radio Access Network-Distributed Unit, RAN-DU, B1) and 3GPP Rel-19 proposed systems (Full RAN, B2) [50] deploy RAN functions on satellites. This improves service but gives SNOs protocol-layer visibility such as identifiers and cell transition events, which can be exploited to infer user regional locations with increasing precision [32, 35, 42] (§2.2).

◦ *Full Onboard RAN and Core Functions* (Figure 4C): The most advanced designs [38, 41, 72, 77–79] further integrate Core Network (CN) functions (e.g., session management, aspects of authentication) directly onboard SNO-managed satellites. This architecture enables satellite self-service, but introduces the most significant privacy challenges: SNOs now operate a single administrative domain that processes both real-time, RAN-derived location data and identity-linked CN functions. This increases the SNO’s capability to track users, raising location privacy concerns (§2.2).

This functional evolution eliminates the natural separation between location and identity. The most advanced architectures (Figure 4C) therefore require a protocol-layer solution to re-enforce this separation.

2.2 Location Privacy Concerns

The migration of cellular functions onto D2C LEO satellites (§2.1) exposes users to significant location privacy risks, which is amplified in dual-entity multi-tenancy (Figure 5). These concerns include:

SNO-Enabled Tracking through Protocol-Layer. SNOs operate on-board RAN and CN functions and provide connectivity, thus inherently observing granular protocol-layer information by the following ways:

◦ *Beam Localization.* SNOs inherently observe RAN-level information, such as beam identifiers and the timing of cell transition events. Passive analysis of these beam transition patterns can narrow down a user’s location to a region of a few kilometers (e.g., RECORD [32] achieves below 4km radius). With multiple satellites or direct signal triangulation capabilities [35], this precision can be further enhanced.

◦ *Signal Characteristic Analysis.* SNOs can observe various uplink signal characteristics from the UE, including Timing Advance, Doppler Shift and Received Signal Strength [42]. SNOs can correlate these with known satellite trajectories and antenna patterns to further refine a UE’s location, potentially on a continuous basis.

◦ *Geospatial IP Address.* The most advanced architectures such as SpaceCore [38] and MOSAIC [41] (and other systems [29, 39, 40]) use 128-bit geospatial IP addresses for




Figure 4: Comparison of the state-of-the-art cellular network function splits in D2C.




Figure 5: Concerns of coupling.

routing, which consists of the UE identifier in its cell location. This creates a binding between a network identifier (even if a Temporary Mobile Subscriber Identity, TMSI) and a defined geographic service region. SNOs directly observe these IP addresses for every active session, which reveals a clear regional trajectory for the SNO over time.

◦ *Session Linkability*. If a user’s different sessions can be linked by the SNO through other persistent protocol-layer metadata, SNOs can construct more comprehensive and long-term user trajectory. (e.g., a user frequently traveling between their workplace and home. An entity might identify the user by linking this travel and time pattern with address data.)

MNO-Enabled Tracking via Settlement. In a typical multi-tenant billing [3] environment, SNOs must report per-session usage data to MNOs to prevent billing fraud. If these reports include service location data, MNOs can link the location to their authenticated subscribers and reconstruct detailed regional trajectories for their identified users. This directly violates GDPR data minimisation principles [13, 48] and exposes users to significant location privacy risks [23, 24, 45, 51, 61, 75].

Amplified Risks in Multi-Tenant D2C. D2C Multi-tenancy usually operates across national borders, serving users in diverse international jurisdictions. It creates fundamental challenges for binary (all-or-nothing) privacy models, which force a choice between complete location disclosure or complete location hiding. However, some regions mandate data minimisation principles (e.g., GDPR [13], CCPA [17]) that prohibit unnecessary location misuse, while others require precise location sharing for emergency services (e.g., E911 [22]), or coarse-grained geofencing for spectrum licensing [31, 49] and digital content (e.g., streaming video [76]). This causes high compliance costs and liability risks for the entities.

3 System and Threat Model

3.1 System Model

We consider a *multi-tenant* direct-to-cell LEO satellite network architecture, consistent with emerging industry trends and standardisation efforts [4, 9, 11, 14, 16, 31, 33, 34, 41, 47, 66, 68, 71, 74]. LPG involves three types of participants:

UE: Standard mobile phones or IoT devices, equipped with

a *tamper-resistant* SIM card, seeking cellular connectivity via the D2C satellite network. UEs are subscribers of an MNO. The SIM card acts as a trusted component for storing data and enforcing cryptographic protocols locally.

SNO: An entity operating a LEO satellite constellation and its associated ground infrastructure, providing orbital connectivity to UEs within their satellite coverage.

MNO: An entity responsible for managing subscriber services, including user identity, subscriptions, billing, and value-added services. In D2C satellite networks, MNOs extend their terrestrial coverage by cooperating with SNOs for orbital connectivity, enabling service for subscribers beyond traditional ground infrastructure coverage.

In this *multi-tenant* D2C LEO architecture, service access and authentication face a critical challenge where traditional solutions [45, 61, 80] conflict with the high-latency and intermittent connectivity (Figure 6). Conventional approaches like terrestrial LOCA [45] migrated to space (Figure 6b) rely on *online attachment*, requiring multiple high-latency round-trips between UE, SNO satellite, and remote MNO core network for real-time authentication [5, 6]. This interactive signaling is slow and risks signaling storms in LEO networks with moving beams [38, 41]. Instead, LPG employs *Offline Attachment* (Figure 6c), aligning with 3GPP NTN’s evolution toward on-orbit gNB/CN functions [5, 50].

LPG achieves this through four components. During Dual-Token registration, UEs and SNOs register MNO-authorized policy-embedded tokens (§4.1.4.2). After establishing a standard Radio Resource Control (RRC) connection with On-board NTN gNB [5], the UE encapsulates its unlinkable token in a Non-Access Stratum (NAS) message to operate *Offline Authentication and Key Agreement (OAKA)* (§4.3), allowing SNO satellites to verify tokens locally and statelessly. Once OAKA establishes secure sessions, *Session Management Tables (SMTs)* manage anonymous sessions, supporting user-configurable *Zero-Knowledge Location Proofs (ZKLP)* [19] for flexible, legal-compliant location disclosure (§4.4). For billing, *Scalable Collaborative Aggregate Claiming (SCAC)* enables private, verifiable SNO-MNO settlement optimized for resource-constrained satellites (§4.5).

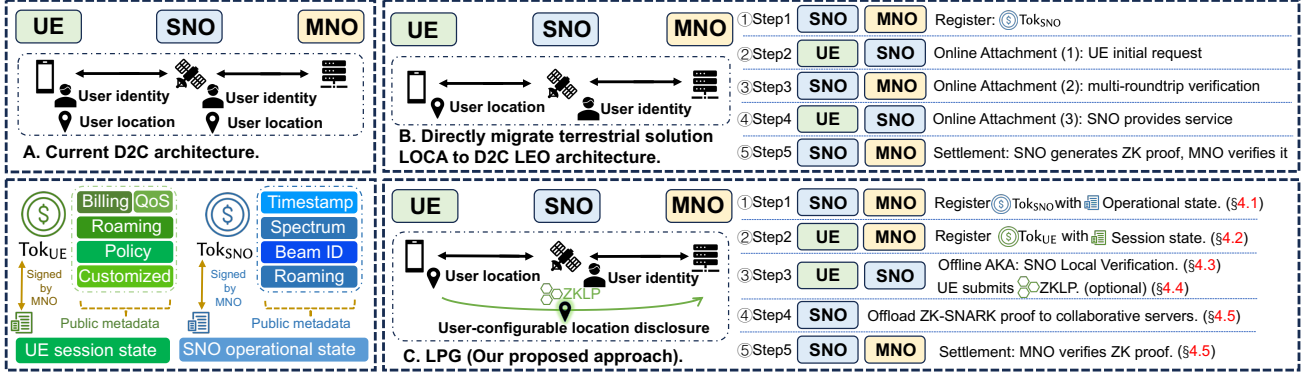


Figure 6: Comparison of architectures: current D2C, LOCA[45]-migrated, and our proposed LPG.

Integration: LPG operates at the NAS layer. It replaces the standard, high-latency AKA procedure with the efficient, offline OAKA protocol after a successful RRC attachment. It assumes sufficient underlying RRC security for the initial link setup [5, 6].

3.2 Threat Model

LPG considers the following threat models:

User Location Privacy Threats. We adopt a common threat model in the security community (e.g., [45, 52, 59, 80]), where SNOs and MNOs are considered as *semi-honest* (honest-but-curious) and *non-colluding*. These entities correctly follow protocol specifications but will attempt to infer maximum user location and identity information from legitimately received data. This assumption is particularly realistic in D2C LEOs’ cross-border partnerships, where MNOs are typically bound by national privacy regulations to prevent foreign SNOs from accessing their UE’s identity-location relations. Given these regulatory enforcement, LPG provides protocol-layer guarantees to prevent each entity from inferring identity-location relations through legitimate protocol interactions. We assume SNOs and MNOs as follows:

- (1) *SNO’s Capabilities:* SNO always has visibility in user location during operation (§2.2). However, SNO attempts to: (i) *Infer UE’s MNO-managed Identity:* SNO attempts to determine the stable identity (e.g., SUPI) of UEs; (ii) *Reconstruct UE Trajectories and Link Sessions:* SNO attempts to link multiple anonymous sessions, potentially initiated with different anonymous tokens to the same UE, thereby reconstructing its movement patterns or service trajectory over time. This could be achieved through the analysis of protocol-derived metadata (e.g., timing of attachments and detachments) and the characteristics of anonymous tokens.
- (2) *MNO’s Capabilities:* MNO has access to persistent user identifiers, subscription info, and historical service records. MNOs attempt to correlate identity with precise location or trajectories mainly via settlement (§2.2).

Specifically, in cross-border D2C scenarios, strict regulatory frameworks (e.g., GDPR [48]) and national security policies legally prohibit the unauthorized sharing of user data between foreign SNOs and domestic MNOs, enforcing the *non-collusion* assumption.

Billing and Settlement Integrity Threats. A malicious SNO might attempt to over-report service usage (i.e. *over-billing*) to MNOs for financial gain. A malicious MNO might under-report or repudiate legitimate usage to SNOs to reduce settlement costs. A malicious UE phone might gain unauthorized service or evade billing by replaying or forging tokens. For MNO-issued, UE-held (e)SIM card, we assume that SIM is a tamper-resistant secure element [81], which is resistant to extraction, duplication, and manipulation.

Communication Channel Threats. We consider the standard Dolev-Yao adversaries [18] controlling the communication channels between UEs and SNOs in attachment (§4.3). The adversary can intercept, modify, replay, and inject messages. Keys are assumed to be secure.

Outsourced Computation Threats. For the settlement (§4.5), some cryptographic computations (ZK-SNARK proof generation) are outsourced to a distributed set of servers. Here, we consider a semi-honest security model for these servers: a threshold number may be corrupted and collude [26, 44, 55], but they generally follow the protocol.

Remark: Specific threats are out of scope: (1) *Denial-of-Service (DoS) Attacks:* LPG mitigates specific DoS vectors, such as signalling storms from CN [41], but comprehensive DoS protection is not our goal. (2) *Global External Adversaries:* Defending against a global adversary capable of correlating ephemeral identifiers across unrelated datasets to track users is beyond LPG.

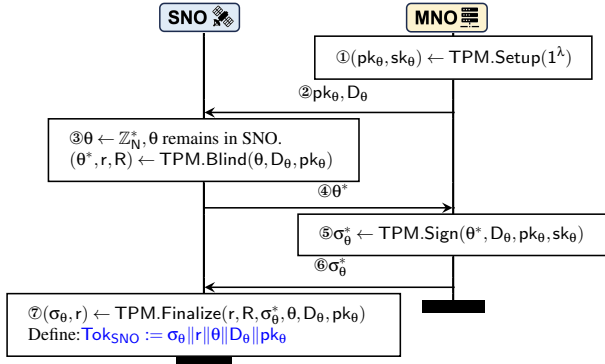


Figure 7: SNO token Tok_{SNO} generation in DualToken.

4 LPG Framework Design

4.1 SNO Tokens Generation

For pay-as-you-go D2C LEO services [41], SNOs need verifiable, time-bound MNO authorisation without continuous MNO oversight. LPG introduces SNO Tokens (Tok_{SNO}), part of the **DualToken**. Unlike prior D2C needing tight infrastructure [50, 68], Tok_{SNO} enables SNO self-service without continuous MNO involvement by using TPM (RSA Anonymous Tokens with Public Metadata [8]). It allows an MNO to issue authorisations to an SNO while binding them to specific operational parameters (D_θ). This registration occurs before service, and the resulting Tok_{SNO} is passed to the satellite upon entering its ground gateway's coverage area.

Protocol Description. Figure 7 shows (Tok_{SNO}) generation via the TPM [8] during SNO-MNO registration:

①② **MNO Setup.** MNO generates RSA modulus N and key pair of TPM, sharing public key pk_θ with UE.

③④ **SNO Blinding and Request.** The SNO generates an ephemeral nonce $\theta \in \mathbb{Z}_N^*$, blinds it with metadata D_θ to compute a blinded message σ_θ^* , and sends σ_θ^* to the MNO, locally storing θ and blinding parameters (r_s, R_s).

⑤⑥ **MNO Verification and Signing.** Upon receiving σ_θ^* , the MNO validates the parameters in metadata D_θ against the SNO's pre-established profile. If valid, the MNO computes and returns the TPM signature.

⑦ **SNO Finalize.** SNO unblinds the signature, producing the final token $\text{Tok}_{\text{SNO}} = (d_\theta, r_s, \theta, D_\theta, pk_\theta)$.

Metadata (D_θ) It contains critical operational parameters, including the issuer MNO ID (linked to pk_θ), coverage region (geographic boundaries), spectrum licences (e.g., frequency bands, power levels), validity period (temporal bounds and timestamps), and other necessary operational or regulatory parameters. This allows the UE to offline validate the SNO token and record satellite operation details, which can then be submitted to MNOs upon network connection for verification related to billing and regulatory auditing.

Security Properties. SNO Token Generation directly invokes the TPM to ensure (1) *Unforgeability*. Only tokens issued by MNO are valid; (2) *Metadata Binding*. The Tok_{SNO} is cryptographically tied to specific operational parameters in D_θ , and the forged D_θ will fail verification.

4.2 UE Tokens Generation

UE token Tok_{UE} is another component of **DualToken**, designed to decouple identity and location. It allows UEs to prove service entitlements to SNOs anonymously while preventing session linkability. Our novel SIM-enforced unlinkable Tok_{UE} extends TPM [8] to a three-party context (SIM, Phone, MNO), embedding MNO-authorized policies, preventing reuse, and ensuring unlinkability. For initial service access where direct MNO connection is unavailable, MNOs can preload a batch of signed tokens and policies into the (e)SIM during manufacturing or via offline prepaid provisioning, enabling immediate offline attachment.

Protocol Description. Figure 8 shows Tok_{UE} generation via our SIM-enforced scheme during UE-MNO registration:

①② **MNO Setup.** The MNO generates an RSA modulus $N = p \cdot q$. MNO generates the public key $pk_\eta = N$ and the secret key $sk_\eta = \phi(N)$. The MNO also defines the hash function H_{MD} for mapping metadata D_η to exponents e_{md} .

③ **SIM generates nonce.** The MNO-issued trusted SIM generates and internally stores a fresh, secret random nonce $\eta \in \mathbb{Z}_N^*$, marking it as "pending". The original nonce η never leaves the SIM.

④⑤⑥ **UE Blinding and request.** The SIM chooses a small base g and computes $\mathcal{E} = g^\eta \bmod N$, which acts as the core message component derived from the secret η . The phone receives \mathcal{E} and determines the required public metadata D_η based on the subscription plan. The phone generates a fresh random string r_{bit} , and $R \in \mathbb{Z}_N^*$, and computes the blinded message components: $M' = \mathcal{E} || r_{\text{bit}}$ and $\eta^* = H_{\text{M}}(M') \cdot R^{e_{\text{md}}} \bmod N$, where $e_{\text{md}} = H_{\text{MD}}(D)$, while H_{M} is a standard message encoding function [8] that securely maps messages to group elements. Finally, the phone sends the request (η^*, D_η) to the MNO.

⑦⑧ **MNO Verification and Signing.** The MNO verifies that the plaintext D_η is consistent with the known UE's subscription plan and policies. If invalid, reject. If valid, the MNO computes $e_{\text{md}} = H_{\text{MD}}(D)$ and $d_{\text{md}} = (e_{\text{md}})^{-1} \bmod \phi(N)$. The MNO generates a signature $\sigma^* = (\eta^*)^{d_{\text{md}}} \bmod N$ and responses it to the phone. The phone forwards σ^* along with the blinding factor R to the SIM.

⑨⑩ **SIM finalisation and One-Time Enforcement** Upon receiving σ^* and R , the SIM checks if the internal flag for η is still "pending". If not (meaning η was already used or compromised), it aborts. If pending, it computes the final unblinded signature $\sigma_\eta = \sigma^* \cdot R^{-1} \bmod N$, and irreversibly marks the nonce η as "used" to prevent reuse. The SIM then outputs σ_η to the phone, which assembles the final UE token

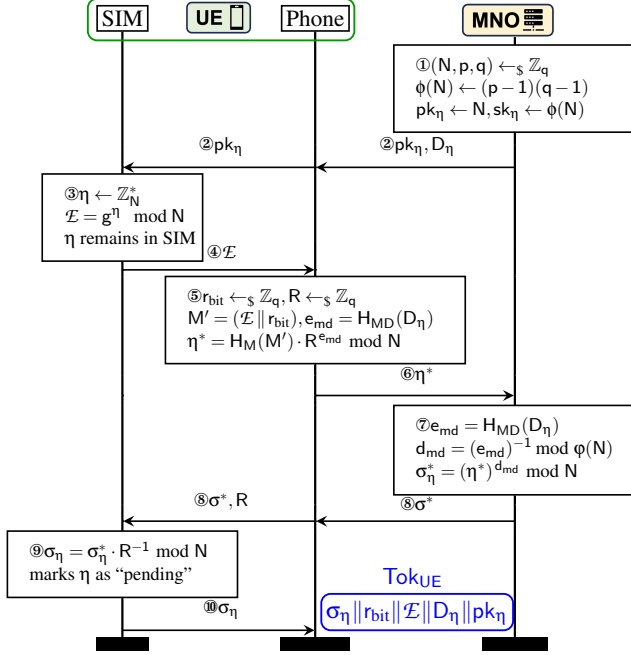


Figure 8: UE token Tok_{UE} generation in DualToken.

$\text{Tok}_{\text{UE}} = (\sigma_\eta, r_{\text{bit}}, \mathcal{E}, D_\eta, \text{pk}_\eta)$.

Metadata (D_η) It embeds UE-specific session state and MNO-defined policies, such as subscription level (e.g., service tier), QoS profile ID (Quality of Service parameters), allowed roaming regions (geographical permissions), and billing rules (e.g., Protocol Data Unit, charging), as well as other policy elements like lawful intercept flags. This enables the SNO to provide customised cellular service tailored to the UE’s entitlements without needing to query remote MNOs.

Ephemeral ID Encryption. UEs pre-compute an MNO-encrypted ephemeral ID $\text{Enc}_{\text{pk}_{\text{MNO}}}(\text{EphID})$ using a periodically broadcast MNO public key pk_{MNO} . This EphID is crucial for session management (§4.4) and is updated upon MNO registration or the broadcast of a new key.

Tok_{UE} is used in OAKA (§4.3), where its unique nonce η identifies a session. The embedded session states allow a UE to initiate a new session with a different SNO by simply submitting a new token, enabling seamless attachment to any available satellite.

Security Properties. The security properties of the Tok_{UE} are summarised in Theorems 1, 2. The formal proofs are given in the full version of this paper [63].

Theorem 1 (Unforgeability of Tok_{UE}) Under the Strong One-More RSA assumption, Tok_{UE} is unforgeable. A PPT adversary making at most ℓ signing oracle queries cannot produce $\ell + 1$ valid tokens with non-negligible probability.

Theorem 2 (Unlinkability of Tok_{UE}) Assuming the random oracle model and the RSA assumption, Tok_{UE} is unlinkable. A PPT adversary cannot link a Tok_{UE} back to the specific UE

with negligible probability.

4.3 Secure Offline Attachment Procedure

LPG’s Offline Authentication and Key Agreement (OAKA) protocol enables secure UE-SNO sessions without real-time MNO involvement, prevents token reuse, and facilitates downlink connectivity without revealing fine-grained location. OAKA operates at the NAS layer, providing mutual authentication, key exchange, and Dolev-Yao resistance using standard NAS messages for compatibility [5, 41].

Protocol Description. In OAKA (Figure 9), UEs/SNOs exchange DualTokens (Tok_{UE} and Tok_{SNO}) and utilize the Elliptic Curve Integrated Encryption Scheme (ECIES) [28] for secure key exchange and message encryption. The protocol initiates only after the UE successfully transition to RRC_CONNECTED state by completing the standard RRC setup procedures (RRCSetupRequest/RRCSetupComplete) with on board NTN gNB. Then, instead of sending standard 5G Registration Requests, UEs encapsulate OAKA requests containing Tok_{UE} within ULInformationTransfer NAS messages after monitoring SNO advertisements on broadcast channels (RRC System Information):

①② **SNO Broadcasts Keys and Tokens.** (Executed periodically by SNO) The SNO generates a fresh ephemeral key pair ($\text{spk}_s, \text{ssk}_s$) for encryption, which is updated periodically. Then, the protocol employs a standard digital signature scheme (e.g., ECDSA), denoted by Sig.Sig for signing and Sig.Verify for verification. The SNO retrieves its certificate Cert_{SNO} ($\text{ID}_{\text{SNO}}, \text{pk}_C$) and sk_C , and signs the ephemeral key, which generates a signature σ_s . Finally, SNO broadcasts $(\text{spk}_s, \sigma_s, \text{Cert}_{\text{SNO}})$ using the customizable fields in the System Information Block (e.g. SIB19) transmitted on the Downlink Shared Channel (DL-SCH), which supports periodic broadcasts without modifying the standard RRC protocol semantics.

③④ **UE Authentication and Token Submission.** (Executed by UE upon receiving the broadcast message) Upon receiving $(\text{Cert}'_{\text{SNO}}, \text{spk}'_s, \sigma'_s)$, the UE (Phone) verifies $\text{Cert}'_{\text{SNO}}$ chain to root CA. Then, the phone verifies the ephemeral key signature: $\text{valid_sig} \leftarrow \text{Sig.Verify}(\text{pk}'_C, \text{ID}'_{\text{SNO}} || \text{spk}'_s, \sigma'_s)$. If valid_sig verification fails, UE aborts. Otherwise, the phone generates ephemeral keys $(\text{spk}_u, \text{ssk}_u) \leftarrow \text{ECIES.KGen}(1^\lambda)$. The phone asks $\text{Tok}_{\text{UE}} = (\sigma_\eta, r_{\text{bit}}, \mathcal{E}, D_\eta, \text{pk}_\eta)$ from SIM. The SIM provides it only if nonce η is “pending”. If not available or already used, abort. Then the phone constructs the request message. This message includes its ephemeral public key spk_u , its UE token Tok_{UE} , and its current SNO-specific serving cell ID (same as the Earth-Fixed Geographic Cells from SpaceCore [38]), we define it as initial_cell_id . The entire payload $\text{msg}_{\text{req}} = \text{spk}_u || \text{Tok}_{\text{UE}} || \text{initial_cell_id}$ is then encrypted: $\text{req} \leftarrow \text{ECIES.Enc}(\text{spk}'_s, \text{msg}_{\text{req}})$. Finally, the phone transmits request req to SNO via NAS ULInformationTransfer [7].

⑤⑥ **SNO Token Verification and Response.** (Executed

by SNO upon receiving req) Upon receiving the encrypted request req from UE, the SNO attempts to use ephemeral private key ssk_s to decrypt it: $msg_{req} \leftarrow ECIES.Dec(ssk_s, req)$. If decryption fails, the SNO aborts. If successful, the decrypted message is parsed to its components: the UE's ephemeral public key spk'_u and the anonymous token Tok_{UE}' , and the $initial_cell_id$. Then, SNO parse Tok_{UE}' into $(\sigma_\eta, r_{bit}, \mathcal{E}, D_\eta, pk_\eta)$. SNO verifies this UE token as follows: The SNO computes $e'_{md} = H_{MD}(D'_\eta)$ and checks whether $\{(\sigma'_\eta)^{e'_{md}} \bmod N'\}$ is equal to $H_M(\mathcal{E}' \parallel r_{bit}')$. If equal, the SNO knows it is an MNO-issued token.

After verification, SNO extracts the policies D'_η and retrieves its operational token $Tok_{SNO} = (\sigma_\theta, r_\theta, \theta, D_\theta, pk_\theta)$. Crucially, the SNO now uses the decrypted $initial_cell_id$ to determine the appropriate downlink channel to send its response (resp) to the UE. The SNO records this $initial_cell_id$ in its session management table (§4.4) associated with the anonymous session. SNO sets response $msg_{resp} = Tok_{SNO}$ and uses ephemeral public key spk'_u to encrypt it into resp, then transmits resp to UE via DLInformationTransfer. Upon successful verification of Tok_{UE} , the SNO's on-board authentication function immediately authorises network access without MNO involvement. The SNO then sends an RRCReconfiguration message to the UE, establishing the necessary data radio bearers (DRBs) and access parameters.

⑦⑧ UE Validates SNO Token and SIM Enforcement.

Upon receiving and successfully applying the RRCReconfiguration message, the UE responds with an RRCReconfigurationComplete and completes the secure session setup. With this authenticated data path now active, the UE uses its UE ephemeral private key ssk_u to recover the SNO token Tok_{SNO} . If decryption fails, it aborts. Otherwise, the UE phone verifies the token by running TPM.Verify algorithm. If verification fails, it aborts. Otherwise, the phone validates SNO operational parameters in D'_θ . Then the phone performs ECDH key exchange using its private key ssk_u and the received SNO public key spk'_s to derive the shared secret S'_{shared} (it must equal S_{shared} due to ECIES containing the key agreement mechanism ECDH inside). Next, a Key Derivation Function (KDF) is used to derive the session key from the shared secret: $K_{sess} \leftarrow KDF(sp'_s, \text{"LPG Session Key"})$. K_{sess} is used to encrypt and protect subsequent communication content, which indicates that both sides are authenticated correctly and the secure channel is ready. All subsequent cell location updates due to UE mobility within the SNO's network are encrypted with K_{sess} , ensuring downlink paging.

The phone signals a successful authentication notification linked to \mathcal{E} on its SIM card. Upon receiving the notification from the phone, the SIM card atomically checks the status of the nonce η linked to \mathcal{E} . If "pending", SIM irrevocably marks η as "used" to prevent token replay attack. If not "pending", SIM signals an error and the phone aborts.

To prevent UE-SNO collusion and incorrect billing, the SIM securely logs the SNO's Tok_{SNO} identifier (Batch_ID

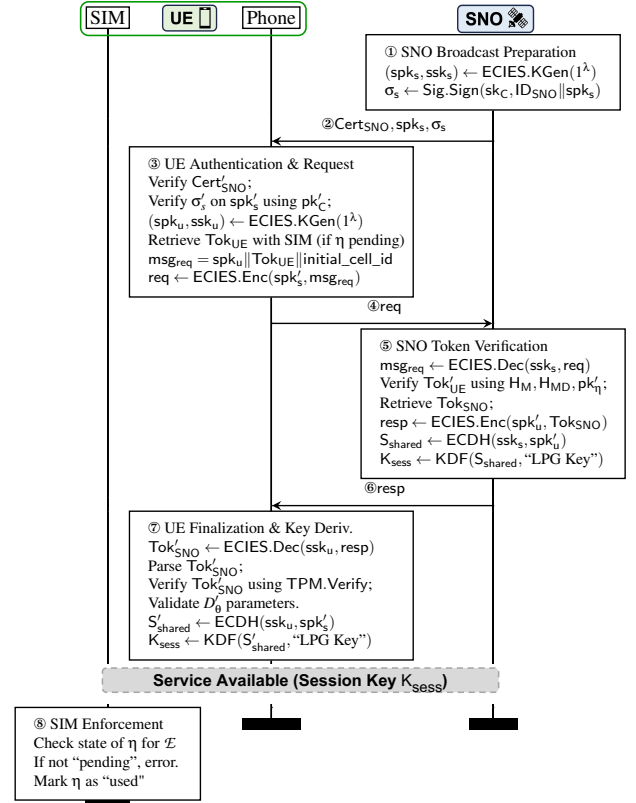


Figure 9: OAKA protocol between UE and SNO

from §4.4.2) and a counter after each OAKA. This SIM-verified record is later reported to the MNO, enabling fraud detection by cross-checking against SNO's SCAC claims (§4.5).

To mitigate SNO trajectory inference through auxiliary information (e.g., timing analysis, moving patterns), LPG adopts a maturing solution from LOCA [45] named *Periodic Attachment*: UEs periodically initiate new OAKA sessions using fresh Tok_{UE} at regular intervals (reattachment period P), preventing SNOs from linking consecutive sessions. To defeat timing-based correlation attacks, UEs randomise their attachment timing within a configurable window W_d around each scheduled reattachment. A higher $\frac{W_d}{P}$ ratio increases the difficulty of identity inference. Furthermore, during each periodic reattachment, the SNO strictly validates the UE's current serving $Cell_id$ (from the Earth-Fixed Geographic Cell [38]) against the allowed region policies defined in the token metadata D_η , ensuring continuous enforcement of geographic restrictions. We formally prove in the full version of this paper [63] that LPG prevents SNO from linking OAKA events to specific UEs, ensuring session unlinkability ($P3$).

Security Properties. The security properties of the OAKA protocol are summarised in Theorems 3, 4, and 5. The formal proofs are given in the full version of this paper [63].

Theorem 3 (SIM-enforced one-time use for Tok_{UE}): As-

suming the SIM card is honest, the Tok_{UE} prevents the double spending of a specific token with overwhelming probability. This security property prevents UE from replaying the same Tok_{UE} for a new OAKA session.

Theorem 4 (Mutual Authentication and Secure Key Exchange): The UE and SNO mutually authenticate each other via a valid Tok_{UE} and a valid Tok_{SNO} . A shared session key K_{sess} is established, known only to the UE and SNO.

Theorem 5 (Resilience against Dolev-Yao Adversaries): An active channel adversary cannot impersonate either party, derive K_{sess} , or cause either party to accept a session.

4.4 Session Management

After OAKA establishes secure sessions (§4.3), LPG’s Session Management Tables (SMTs) provide carrier-grade services, downlink connectivity, and flexible location privacy. This is achieved through three key aspects: *i*) user-configurable ZKLP-based location disclosure for flexible privacy, *ii*) SNO-side cell location recording for operational needs, and *iii*) session recording for settlement.

4.4.1 User ZKLP Configuration

Traditional D2C systems expose fine-grained user location to MNOs (§2.2), while binary privacy models [45, 61, 80] are ill-suited for flexible privacy needs. To address this, LPG integrates ZKLP as follows:

ZKLP Mechanism. Let the UE’s geographic position be $\mathbf{P}_u = (\text{lat}_u, \text{lon}_u)$. ZKLP allows the UE to generate a zero-knowledge proof π_{zk} attesting that \mathbf{P}_u lies within a specific hexagonal region $\text{Region}(\text{hexID}, \text{res})$ without revealing \mathbf{P}_u . The Region is identified by hexID from the Uber H3 Discrete Global Grid System (DGGS) [70] at a user-chosen resolution level $\text{res} \in [0, 15]$. \mathcal{H} can be defined by its centre \mathbf{c}_i and vertices (e.g., $\mathbf{v}_0, \dots, \mathbf{v}_5$ or $\mathbf{v}_6, \dots, \mathbf{v}_{11}$, Figure 10).

UE Disclosure. During an active OAKA session (§4.3), the UE decides *if* and *at what granularity* (res) it wishes to disclose its location range. Specifically, following the ZKLP specification [19], $\text{res} = 0$ corresponds to the coarsest granularity (millions of km^2), while $\text{res} = 15$ offers the finest sub-meter precision ($< 1 \text{ m}^2$). A higher res implies a smaller hexagonal cell and more precise disclosure, while a lower res provides stronger privacy within a larger area (e.g., $\text{res} = 11$ is $\sim 21.44\text{m}$; $\text{res} = 1$ is $\sim 360\text{km}$). This allows balancing privacy preferences with service requirements.

ZKLP Proof Generation. If a UE decide to disclose, it generates a proof $\pi_{ip} \leftarrow \text{ZKLP.Prove}(\text{pp}, (\text{lat}, \text{lon}), \text{res})$ using public parameters pp . To securely link this disclosure to its MNO-managed profile without revealing its identity to the SNO, the UE includes its MNO-encrypted ephemeral identifier $\text{Enc}_{\text{pk}_{\text{MNO}}}(\text{EphID})$ (from §4.2). The UE then sends the ZKLP proof π_{ip} , encrypted identifier $\text{Enc}_{\text{pk}_{\text{MNO}}}(\text{EphID})$,




Figure 10: SMTs mechanism and ZKLP disclosure

claimed hexagonal region hexID , and granularity res to the serving SNO over the secure channel established with K_{sess} .

SNO Verification and Record. The SNO then verifies the proof: $\text{ZKLP.Verify}(\text{pp}, \pi_{ip}, \text{hexID}, \text{res})$. If valid, the SNO records the ZKLP data in its session tables. Crucially, the SNO cannot decrypt $\text{Enc}_{\text{pk}_{\text{MNO}}}(\text{EphID})$, preserving user identity privacy. This data is later forwarded to the MNO (§4.4.2), enabling the MNO to verify regional presence.

4.4.2 SNO Session Recording and Management

To manage sessions and prepare for settlement, the SNO uses Session Management Tables (SMTs) (Figure 10), which support billing and disclosure while preserving UE anonymity. SMTs operate on two concepts:

- *Session & SID*: A session is a single UE service period initiated via OAKA (§4.3) using a specific Tok_{UE} (§4.2). Each session is uniquely identified by a Session Identifier (SID), linked to the unique nonce η (and \mathcal{E}) and its Tok_{UE} .

- *Session Group & Batch_ID*: A session group comprises multiple independent sessions authorised under the same Tok_{SNO} (§4.1) and processed as a single settlement unit. Each session group is uniquely identified by a Batch_ID, linked to the nonce θ within its Tok_{SNO} .

SNO-recorded Session Table (e.g., Table 2): The SNO locally maintains a log for all active and recent anonymous sessions. Each entry maps a single session, which contains:

- SID: The unique session identifier binding to the nonce η (and \mathcal{E}) within its authorizing Tok_{UE} .

- Enc_EphID: Stores $\text{Enc}_{\text{pk}_{\text{MNO}}}(\text{EphID})$ if the UE disclosed location via ZKLP (§4.4.1), otherwise None.

- Cell_id: A identifier for Earth-Fixed Geographic Cell from SpaceCore [38], which is the last known cell ID provided by the UE via updates `initial_cell_id` in OAKA (§4.3). This is used by the SNO for paging and downlink connectivity.

- Batch_ID: A session group identifier. It links all the sessions to its authorizing Tok_{SNO} and session group.

- H(Secrets): A SNO-computed per-session cryptographic commitment $H(\text{secret}_i)$ for settlement, where $\text{secret}_i = \text{PRF}(k_\theta, \mathcal{E} \parallel \theta)$. H is a hash function, PRF is a secure pseudo-random function (e.g., HMAC-SHA256), and k_θ is an SNO-secret key tied to θ . H is a hash function. Both PRF and H

Table 2: Example: SNO-recorded Session Table

SID	Enc_EphID	Cell_ID	Batch_ID	H(Secrets)	Usage	Duration	QoS	Region Proof (ZKLP)
S1	Enc _{pk_{MNO}} (ID Alice)	(row1, col1)	B123	H(secret ₁)	0.5GB	20min	99%Uptime, etc.	$\pi_{B123,S1}$: Prove HexID=R45, Res=15
S2	None	(row2, col2)	B123	H(secret ₂)	1.2GB	35min	90%Uptime, etc.	None
S3	Enc _{pk_{MNO}} (ID Bob)	(row3, col3)	B123	H(secret ₃)	0.8GB	30min	97%Uptime, etc.	$\pi_{B123,S3}$: Prove HexID=C09, Res=3

Table 3: Example: Aggregated Usage Table

Entry_ID	H(Secrets)	Usage	QoS	Batch_ID
1	H(Secrets ₁)	0.5 GB	99%	B123
2	H(Secrets ₂)	1.2 GB	90%	B123
3	H(Secrets ₃)	1.8 GB	95%	B456

Table 4: Example: Location Disclosure Table

Enc_EphID	Usage	Verified Region (ZKLP)	Individual QoS
Enc _{pk_{MNO}} (ID Alice)	0.5GB	Hex ID:R45 (Res=15)	99%Uptime, etc.
Enc _{pk_{MNO}} (ID Bob)	0.8GB	Hex ID:C09 (Res=3)	97%Uptime, etc.
Enc _{pk_{MNO}} (ID Carol)	1.8GB	Hex ID:C03 (Res=4)	95%Uptime, etc.

are known to SNO and MNO.

Usage: Resource consumption records. (data volume, etc.)

QoS: Quality of Service parameters.

ZKLP_Region: Stores $\{\pi_p, \text{hexID}, \text{res}\}$ if UE provided a ZKLP; otherwise None.

As shown in Figure 10, this internal SNO-recorded Session Table is the foundation for generating two externally shared tables for the MNO:

Aggregated Usage Table (e.g., Table 3): Periodically, the SNO generates this AggTable as public input for the SCAC settlement process (§4.5). This table comprises all sessions served within a given settlement period for a specific MNO, and we denote its total number of entries as n_{sess} . For each Batch_ID, it aggregates usage data. Crucially, instead of exposing individual SID, each entry representing a session includes: (1) Entry_ID, An index for the table; (2) H(secret), the hash from the SNO-recorded Session Table; (3) Usage; (4) QoS; (5) Batch_ID.

Location Disclosure Table (e.g., Table 4): This table, also shared with the MNO, facilitates user-configurable location disclosure. For sessions with ZKLP proofs, it contains: (1) Enc_EphID, the MNO-encrypted ephemeral ID; (2) Usage; (3) QoS; (4) ZKLP_Region. It consists of the hexID and resolution res from the UE-submitted ZKLP, along with the proof π_p . This allows the MNO to decrypt EphID, link the disclosure to the subscriber, and verify their presence without learning any other location information.

Lawful interception. LPG can also enable lawful interception. If a Law Enforcement Agency (LEA) [30] needs to interact with the MNO to resolve the target identity for network identifier purposes. The MNO then identifies all the Tok_{UE} (issued by MNO as discussed in §4.2) and session states, then coordinates with the SNO. The SNO can then provide the LEA with location data for sessions matching the specified tokens without learning the user’s identity. At the same time, the MNO remains unaware of precise locations beyond LEA-authorized disclosures. This preserves the identity-location decoupling principle while meeting regulatory requirements.

In conclusion, LPG’s SMT mechanism enables SNO to manage anonymous sessions and downlink paging, facilitates MNO billing and policy enforcement, and provides UEs with

flexible control over their location privacy.

4.5 Scalable Collaborative Aggregate Claim

Fair and verifiable billing is crucial in a privacy-preserving D2C LEO system. Traditional settlement conflicts with LPG’s privacy goals by relying on per-user records that reveal UE trajectories. Conversely, simple anonymous reports are unverifiable, allowing malicious SNOs to overbill, while heavy-weight solutions like LOCA [45] are computationally infeasible on resource-limited satellites. To address these challenges, LPG introduces Scalable Collaborative Aggregate Claiming (SCAC), adapting and extending ideas from collaborative ZK-SNARK [26, 44, 55] and LOCA [45]. SCAC allows an SNO to prove to an MNO that it legitimately served an aggregated batch of user sessions for a certain total usage, without revealing individual session details. To ensure viability, SCAC offloads the heavy ZK-SNARK proof generation from the satellite to terrestrial servers.

Aggregate Claiming. For each settlement period, the SNO generates an aggregate claim to show that for the specific MNO who issues the tokens, the SNO has legitimately served k (the number of sessions) different sessions of this MNO’s subscribers, consuming \mathcal{Z} total usage, proven by ZK-SNARK proof π_{agg} . (\mathcal{X} represents the system parameter for maximum number of k). The aggregate nature of the claim creates obfuscation: while MNO observes the total usage \mathcal{Z} , however the number of valid session subsets that could produce this sum is sufficiently large, thus it’s difficult for MNO to infer whether an individual session is part of SNO’s claim. This prevents the MNO from reconstructing user trajectories and learning location data except for ZKLP. The ZK-SNARK proof π_{agg} convinces the MNO that: (1) The SNO possesses k distinct private session secrets (secret _{i}). (2) Each of these k secrets corresponds to a unique entry in AggTable (as identified by H(secret _{i}) matching the entry’s hash of secret field). (3) The sum of the usage values from these k corresponding entries in AggTable equals the claimed total \mathcal{Z} . This proof, combined with AggTable, prevents SNO from overbilling.

ZK-SNARK Circuit Design. We leverage Scalable Collaborative ZK-SNARK (SCZK) [44] to generate the proof π_{agg} .

This step needs to translate SNO’s aggregate claiming mentioned above into a zero-knowledge circuit C_{LPG} that encodes the relation \mathcal{R}_{LPG} . The formal detail is as follows:

- *Public Input* (\mathcal{X}_{pub}). The public statement consists of three components: *i*) *AggTable* (e.g., Table 3): It contains k entries. Each entry includes a hash $H(\text{secret}_i)$, individual usage information, and Batch ID, but not user identity; *ii*) k : The actual number of sessions claimed by the SNO in this instance; *iii*) \mathcal{Z} : The total usage amount claimed by the SNO for these k sessions.

- *Private Witness* (\mathcal{W}). The SNO’s private knowledge of k session secrets $\{\text{secret}_i\}_{i=1}^k$, where each secret_i was generated during the corresponding i -th claimed session (in §4.4).

- *Relation* (\mathcal{R}_{LPG}). The relation checks if public statement $\mathcal{X}_{pub} = (\text{AggTable}, k, \mathcal{Z})$ is consistent with private witness $\mathcal{W} = \{\text{secret}_i\}_{i=1}^k$. It holds if and only if:

- (i) *Session Authentication*: Each claimed secret corresponds to a unique entry in *AggTable*. There exists an injective mapping $\sigma: \{1, \dots, \mathcal{K}\} \rightarrow \{1, \dots, |\text{AggTable}|\}$ such that: $H(\text{secret}_i) = \text{AggTable}[\sigma(i)].H(\text{secret})$.

- (ii) *Usage Integrity*: To prevent overbilling, it must ensure \mathcal{Z} correctly reflects the sum of usage from legitimately served sessions. The sum of usage values from matched entries equals the claimed total: $\sum_{i=1}^k \text{AggTable}[\sigma(i)].\text{Usage} = \mathcal{Z}$.

Scalable Collaborative Proof Generation. To address the computational limitations of LEO satellites, SCAC employs the SCZK framework [44]. SCZK enables multiple distributed servers to jointly generate a ZK-SNARK proof without any single server learning the entire witness. It operates in a load-balanced setting among N_{ser} servers by first having the entity holding the witness (i.e. the SNO) encode its private witness \mathcal{W} into N_{ser} secret shares using Packed Secret Sharing (PSS) [25]. These shares are then distributed to the servers, which engage in a secure multi-party computation (MPC) protocol to collectively execute the underlying ZK-SNARK’s proving on these shares.

In our SCAC design, the SNO transforms its private witness $\mathcal{W} = \{\text{secret}_i\}_{i=1}^k$ into N_{ser} shares using Packed Secret Sharing: $\{\mathcal{W}_1, \mathcal{W}_2, \dots, \mathcal{W}_{N_{ser}}\} = \text{PSS.Share}(\mathcal{W}, \mathcal{L})$, where $\mathcal{L} < N_{ser}/2$ is the maximum number of corrupted servers. During satellite-ground communication windows, the SNO distributes these shares along with the public input \mathcal{X}_{pub} (containing *AggTable*, k , and \mathcal{Z}) to N_{ser} terrestrial servers. The servers collaboratively execute $\pi_{agg} = \text{SCZK.Prove}(\text{pp}, \mathcal{X}_{pub}, \mathcal{W}_1, \dots, \mathcal{W}_{N_{ser}})$ without revealing \mathcal{W} to any individual prover. Each prover performs only $O(|C|/N_{ser})$ operations where $|C|$ is the circuit size, achieving scalable proof generation. The resulting proof π_{agg} can be verified by the MNO using standard ZK verification.

Asynchronous Submit. To minimize settlement latency in the intermittent connectivity environment of LEO networks, the settlement procedure can be completed asynchronously, and the ZK-SNARK proof π_{agg} can be generated with delay. After each session, the SNO locally computes $H(\text{secret}_i)$.

Only during satellite-ground communication windows, SNO upload the *AggTable* to its ground infrastructure and servers. It avoids relying on the unstable ISLs.

Security Properties. The security properties of SCAC are formalized in **Theorem 6~9**, the formal proofs are provided in the full version of this paper [63]. The SCAC protocol guarantees *soundness* to prevent SNO overbilling, *completeness* to prevent MNO repudiation, and *zero-knowledge* to protect user location privacy during settlement. The collaborative proof generation is also proven secure against *a threshold of corrupted servers* ($\mathcal{L} < N_{ser}/2$). In conclusion, SCAC enables fair, efficient, and privacy-preserving settlement for the dynamic D2C LEO environment.

Theorem 6 (Soundness and SNO Overbilling Prevention) A malicious SNO can only generate a valid ZK-SNARK proof π_{agg} for an invalid aggregate claim with a negligible probability. Therefore, it prevents SNO overbilling.

Theorem 7 (Completeness and MNO Repudiation Prevention) A legitimate aggregate claim from an honest SNO with a correctly generated ZK-SNARK proof π_{agg} will always be verified as valid, thus the MNO cannot repudiate billing.

Theorem 8 (Zero-Knowledge): The proof π_{agg} reveals no information about the individual sessions that contribute to the aggregate beyond what is explicitly claimed (k, \mathcal{Z}).

Theorem 9 (Collaborative Security): The proof generation is secure against up to $\mathcal{L} < N_{ser}/2$ corrupted servers.

4.6 Location Privacy Analysis

LPG establishes robust identity-location decoupling. We prove this with two core theorems, which are formally defined and proven in Appendix A.1, A.2:

First, **Theorem 10** (Identity Hiding against SNO) shows that LPG provides protocol-layer anonymity for UEs against the SNO. The proof demonstrates that unlinkable Tok_{UE} can prevent the SNO from linking observed sessions to a UE’s stable identity. While this does not prevent location inference from physical-layer signals, it ensures that within the protocol, a UE’s identity remains hiding to the SNO.

Theorem 10 (Identity Hiding against SNO). LPG ensures two properties: (1) *Identity-hiding in Protocol Content*: Assuming that *i*) the DualToken scheme provides (ϵ_1, t_1) -unlinkability, *ii*) ECIES provides (ϵ_2, t_2) -IND-CPA security, and *iii*) SNOs and MNOs are non-colluding. Then all content observed by the SNO during its interaction with a UE cannot make SNO guess the UE identity, that is UE holds $(\epsilon_1 + \epsilon_2)$ -Identity-Hiding against SNO. (2) *Resilience Against Trajectory Inference*: Assuming UE operational strategies and factors satisfy the conditions of Lemma 1, the probability $\text{Pr}_{\text{link_traj}}$ that SNO can efficiently reconstruct a specific UE’s identifying trajectory of N_{attach} anonymous sessions is negligible for sufficiently large N_{attach} or ND.

We formalize the resilience against such trajectory inference using the following lemma [45]:

Lemma 1 (Trajectory Linking Ambiguity [45]) Let an observer attempt to reconstruct a target entity’s trajectory by linking a sequence of N_{attach} events, where each event is associated with a cryptographically unlinkable identifier. If the target entity initiates new events periodically (average period P) with randomized timing within an effective window Wd , and NU other entities generate similar unlinkable events in the observer’s view, then: (1) The average number of candidate events, ND , that the observer must consider for each true event of the target can be lower-bounded by $ND \geq 1 + NU \cdot (Wd/P)$. (2) The probability, $\text{Pr}_{\text{link_traj}}$, that the observer correctly links all N_{attach} events to the target’s true trajectory is approximately $\text{Pr}_{\text{link_traj}} \leq (1/ND)^{(N_{\text{attach}}-1)}$ (for $N_{\text{attach}} > 1$), assuming optimal guessing by the observer based solely on these operational parameters.

Second, **Theorem 11** (Location Obscurity from MNO) shows that LPG prevents the MNO from learning a UE’s precise SNO-observed location. The proof is based on the fact that the MNO’s view is restricted to either user-consented, granular ZKLP disclosures or aggregate settlement claims. The privacy of these claims is guaranteed by the zero-knowledge property of SCAC and disaggregating AggTable to infer individual session details.

Theorem 11 (Location Obscurity against MNO). Assuming that *i*) (k, n_{sess}) satisfy the conditions in Lemma 2, *ii*) SCZK provides (ϵ_3, t_3) -Zero-Knowledge, *iii*) PRF is (ϵ_4, t_4) -secure. Thus, the UE holds $(\epsilon_3 + \epsilon_4)$ -Location-Obscurity against MNO. In LPG, the MNO cannot determine the UE’s precise real-time trajectory, apart from ZKLP regions the UE voluntarily discloses.

We formalize the ambiguity inherent in the aggregate settlement claim using the following lemma [15, 45]:

Lemma 2 (Aggregate Sum Ambiguity [15, 45]) Let S_{table} be the SNO’s internal session table containing n_{sess} actual user sessions, where each session $j \in S_{\text{table}}$ has an associated usage value v_j drawn from a distribution with sufficient entropy. Consider an aggregate claim (k, Z) which asserts that a subset $Y \subset S_{\text{table}}$ of k sessions has $\sum_{y_i \in Y} v_{y_i} = Z$. If the ratio k/n_{sess} falls within a non-trivial theoretical range $[L_{\text{val}}, U_{\text{val}}]$ (where $0 < L_{\text{val}} \leq U_{\text{val}} < 1$), then the expected number of distinct subsets $Y' \subset S_{\text{table}}$ of size k such that $\sum_{y'_i \in Y'} v_{y'_i} = Z$ grows exponentially with n_{sess} . Empirical analysis shows that for $k/n_{\text{sess}} \in [1/150, 1/2]$, the exponential base $\alpha > 1.1$, providing strong obfuscation.

Consequently, LPG breaks the correlation between UE identity and its real-time location and trajectory, while allowing for user-consented disclosure.

5 Evaluation

We prototype LPG with a real in-orbit LEO satellite, a ground-based server, and two commodity phones (Figure 11). We evaluate LPG’s performance and overhead in terms of scalability, latency, CPU usage, power and temperature dynamics.




Figure 11: Prototype of LPG

5.1 LPG Prototype

We prototype LPG with a real in-orbit LEO satellite. The basic information of the satellite platform used in our experiments is summarized in Table 5.

Table 5: Basic Information of BUPT-3 Satellite.

Category	Parameter	Value
Orbit	Orbit Type	Low Earth Orbit
	Altitude (km)	525
	Inclination (°)	97.418
	Eccentricity	0.00125
Satellite Platform	Dimensions (mm)	78×840×601.8
	Mass (kg)	63
	Onboard Power (Wh)	720
	TT&C Rates (kbps)	4.8 (uplink), 9.6 (downlink)
	Data Rates (Mbps)	1 (uplink), 100 (downlink)
Payload	Max Power Limit (W)	250 (total for 2 servers)
	Max Surface Temp. (°C)	45
	Payload Processor	8-core ARM @2.3 GHz
	Memory	DDR4 16 GB @2400 MHz
	Storage	256 GB+
	Operating System	Custom Embedded Linux OS

Before in-orbit deployment, we developed and tested all experimental workloads on terrestrial workstations. Each containerized workload passed a multi-stage validation process (standalone, hardware-in-the-loop, full-system integration) to ensure correctness. Our satellite has been operating in orbit and is undergoing platform stability validation.

SNO. Our satellite BUPT-3 (TY46) has operated stably in orbit since 2025-05-17 04:12 UTC. It hosts two onboard commodity servers operating within the resource envelopes summarized in Table 5. On one server, we implement LPG’s Tok_{SNO} generation (§4.2), OAKA attachment (§4.3), session management (§4.4) and SCAC packet secret sharing (§4.5). We also deploy LOCA [45] onboard as a proof generation baseline.

MNO. We implement the MNO protocol stack on a terrestrial commodity XH320 V2 server equipped with two Intel Xeon E5-2470 v2 processors, 32 GB of DDR4 RAM and 557.9 GB of disk storage. On the server, we implement LPG’s MNO setup, verification and signing logic (§4.1, §4.2), and SCAC proof generation (§4.5).

UE. We use various commodity smartphones, including Xiaomi 14 Ultra with satellite voice calls via Tiantong GEO and messaging via Beidou GEO satellites, and Google Pixel 9 Pro with satellite emergency SOS messaging. We develop a

simple test app that coordinates SIM/eSIM profiles via TelephonyManager APIs [10], implementing LPG’s Tok_{UE} generation (§4.2), SIM-enforced consumption in OAKA (§4.3), and ZKLP location disclosure (§4.4).

5.2 Satellite Experiments and Data Collection

We collect two main categories of data from our in-orbit experiments: *payload-layer data* and *platform-layer telemetries*. Payload data includes performance metrics from our LPG prototype (e.g., latency, resource usage), augmented by diagnostic logs and thermal data. Platform telemetries cover satellite-level metrics like thermal regulation, power subsystem behavior, and attitude control. We also generate custom telemetry, such as command counters and success/failure indicators, reflecting workload status. All telemetry and experimental outputs are stored onboard, periodically packaged, and downlinked to our ground station via TT&C and a dedicated X-band link.

5.3 Results and Analysis

5.3.1 Registration

LPG’s DualToken registration (§4.1, §4.2) enables identity-location decoupling (*PI*) with practical overheads. Figure 12a compares the number of UE token nonces storable on various commercial SIM/eSIM cards based on their RAM capacities [1], which is a constraint noted in SIM-based token systems [41, 61]. LPG achieves comparable storage capacity to them, which is sufficient for practical user needs, aligned with the common prepaid SIM on the market [67], and session unlinkability (*P3*) can be supported through sufficient token freshness without changing existing UE hardware. To maintain token freshness within limited storage, the SIM automatically recycles storage by deleting the oldest used nonces after session termination, allowing for sustainable operation. Figures 12b and 12c show that the generation time for a Tok_{SNO} and a Tok_{UE} is 80 ms and 33 ms, with size of ~ 130 and ~ 130 bytes. Although LPG increases token generation time, the registration overhead remains significantly smaller than typical UE session durations (a few minutes [41, 79]), making the additional cost practically acceptable.

5.3.2 Attachment

LPG’s OAKA protocol (§4.3) introduces an acceptable 140 ms average extra latency for DualToken exchange (Figure 14a). Given that the baseline LEO attachment latencies are 40 \sim 300ms [41], the total time remains well under the 500ms threshold for minimal application impact [46]. The latency is a negligible fraction of typical 7 \sim 8 minute communication windows, making it suitable for periodic reattachments 4.3. OAKA’s communication overhead is also lightweight (366-byte request and 233-byte response), imposing a negligible

impact on data communication links (NAS layer) with at least 1Mbps bandwidth (Table 5). Figure 14b shows that during OAKA operations, CPU power stays below 35W and temperature remains under 55°C. This ensures stable COTS device operation within thermal and power budgets (Table 5), validating LPG’s viability in resource-constrained space (*VI*). The overhead is acceptable considering OAKA provides critical identity-location decoupling (*PI*) and MNO-controlled identity-based service (*F2*). LPG achieves effectively zero handover latency through its token-based design. When a UE needs to switch satellites, it simply initiates a new session with the next available satellite using its existing Tok_{UE}, without requiring any inter-satellite coordination or state transfer, which is resilient to orbital dynamics (*V2*).

5.3.3 Settlement

We evaluate LPG’s SCAC protocol (§4.5), which ensures efficient, privacy-preserving settlement. We also evaluate LOCA baseline [45] across various configurations of \mathcal{K} (maximum claimed sessions per SNO each time) and \mathcal{N} (total session table entries). The results demonstrate LPG’s superior scalability, reduced overhead, and acceptable on-orbit system effects.

Scalability. LPG readily scales to large $\mathcal{K} \times \mathcal{N}$ settings, whereas the LOCA fails when $\mathcal{K} \times \mathcal{N}$ exceeds 8192. LOCA’s scalability is limited by its in-orbit ZK-SNARK proof generation, which exhausts satellite CPU and memory and causes prohibitive latency. In contrast, LPG achieves high scalability by offloading computation-intensive tasks.

SNO overhead. LPG reduces satellite processing latency by orders of magnitude ($\sim 10^4 \times$) compared to LOCA (Figure 15a). LPG’s SCAC only requires the satellite to perform lightweight PSS [25], distributing secret shares to servers (§4.5). At similar scales, LOCA’s latency and resource demands would be prohibitive for in-orbit deployment.

Server overhead. LPG’s ground server latency is orders of magnitude lower than the baseline ($\sim 10^2 \times$), primarily due to distributing proof generation across 32 servers (Figure 15b). Its SCZK-compatible HyperPlonk framework also yields a more efficient circuit. As Figure 15c shows, the constraint count grows near-linearly with \mathcal{K} for fixed \mathcal{N} s (256, 512, 1024), indicating high scalability. The constraint count is dominated by SHA256 hash computations ($\sim 35,000$ per session) and session matching, scaling as $|C_{LPG}| \approx \mathcal{K} \times 35000 + O(\mathcal{K} \times \mathcal{N})$.

Satellite system effects. During settlement, LPG’s impact on the satellite’s CPU usage, power, and temperature remains well within safe operational limits [79] (Figure 15d, 15e, 15f). CPU power stays below 40W (well under the 250W satellite’s operational limit, as Table 5 specifies), and CPU temperature stays under 62°C (which keeps payload surface $< 45^\circ\text{C}$). For smaller $\mathcal{K} \times \mathcal{N}$ values, LPG shows lower power and temperature than LOCA. Conversely, for larger $\mathcal{K} \times \mathcal{N}$ values, LOCA’s metrics appear lower as heavy computation saturates




Figure 12: Registration evaluation.

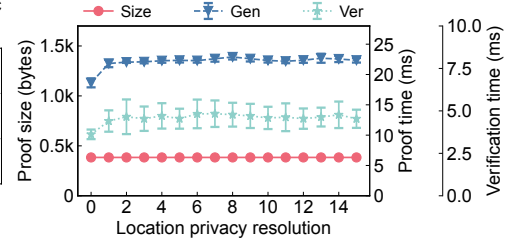
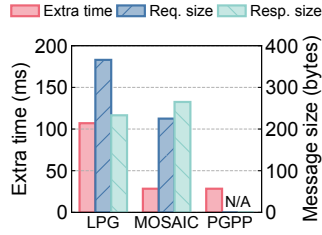



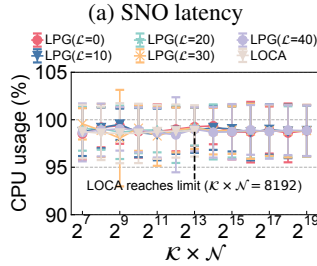
Figure 13: Impact of resolution.



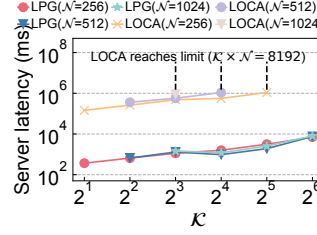
(a) Extra time and message size




(b) CPU power & temperature



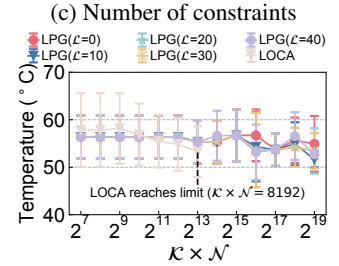
(c) SNO latency



(d) CPU usage



(e) CPU power



(f) CPU temperature

Figure 14: Attachment.

Figure 15: Satellite in-orbit evaluation of settlement.

the resources, triggering the satellite platform's emergency system that artificially caps further increases. Finally, our evaluation across corrupted parties $0 \sim 40$ shows a reasonable corruption threshold L has minimal impact.

Impact of disclosure resolution. We evaluate LPG's performance across different location disclosure resolutions ($res \in \{0, 1, \dots, 15\}$) (Figure 13). The ZKLP proof size remains consistently small, ~ 430 bytes across all resolutions. Proof generation and verification times average about 23 ms and 2.5 ms. This consistency allows UEs to choose any desired location granularity on commercial mobile phones.

6 Discussion

Adaptability in Divergent Constellations. Due to fundamental physics limitations, LPG's cellular 4G/5G service model is inapplicable to MEO/GEO satellites. However, *User-Configurable Location Disclosure* (§4.4) remains highly valuable for MEO/GEO systems supporting IoT connectivity, which could embed ZKLP [19] in sessions. (e.g., oil rigs can prove the "North Sea" without revealing exact coordinates to maintain business confidentiality when using MEO [62].)

Adaptability in Different Operation Models. Although LPG assumes a multi-tenant model, it can also adapt vertically integrated SNO models [56, 57], where SNOs control both satellite infrastructure and customer relationships. Here, LPG's dual-entity framework transforms into internal privacy firewalls, enabling operators to implement internal separation between network operations and customer billing/analytics divisions for regulatory compliance and audit transparency, thus attracting users who prioritise privacy.

Adaptability in D2C LEO evolution. LPG adapts the industry progress and future trends. Although LPG is designed for resource-constrained operation (V1) and immature ISLs (V2), more powerful onboard computing devices and mature ISLs critically enhance LPG's operations. They enable LPG to have the following enhancements: (1) distributed SCAC proof generation (§4.5) across satellite networks, reducing computational bottlenecks; (2) faster token updates to enhance location privacy ($P \downarrow, ND \uparrow, P_{\Gamma_{link_traj}} \downarrow$). The trend towards *onboard CN functions* improves our motivation [5, 50]. Since more identity-sensitive operations within CN migrate to satellites, the protocol-layer privacy becomes essential.

7 Conclusion

This paper proposes LPG, the first framework that provides identity-location decoupling and user-configurable location disclosure in multi-tenant D2C LEO networks. LPG is viable on real-world LEO satellites and commercial mobile phones, providing practical location privacy without compromising essential services or regulatory compliance. We view this as a first step towards location-privacy-preserving global satellite connectivity, and believe our protocol-centric design can complement physical and signal-layer privacy technologies.

8 Acknowledgement

This work is supported by A*STAR, CISCO Systems (USA) Pte. Ltd and National University of Singapore under its CISCUS Accelerated Digital Economy Corporate Laboratory (Award I21001E0002). This work is also supported by Science and Technology Innovation Program of Xiongan New Area (Grant No. 2024XAGG0025), National Key Research and Development Plan, China (Grant No. 2023YFB2903902), Research Initiation Project for Introduced Talents of BUPT (Grant No. 2025KYQD09), the National Natural Science Foundation of China (Grant No. U25B2040, 62272010, 62502012, 62572067), and the China Postdoctoral Science Foundation (Grant No. GZC20251070). The corresponding authors of this work are Qing Li and Prosanta Gope. We sincerely thank all anonymous reviewers for their valuable suggestions. We thank the RHONE emulator for supporting our pre-launch validation experiments, and Tiansuan Constellation for facilitating the in-orbit implementation.

Open Science

We make our research artifacts accessible in <https://zenodo.org/records/17906146>, including the (i) on-board experimental workflows and procedures; (ii) telemetry and telecommand specifications and methods; (iii) end-to-end scripts for phone and satellite experiments; (iv) evaluation datasets (temperature, power, and experiment logs). These artifacts enable the reproduction of our evaluation results and further experimentation on similar in-orbit satellites. To protect the operational security of the live in-orbit satellite, we do not release the full source code of communication protocols and telemetry platform at this moment. Once the academic satellite completes its in-orbit phase and the additional code and documentation pass provider review, we plan to release them in <https://github.com/LPGSatellite/LPG> to facilitate community reuse.

Ethical Considerations

Stakeholders and Potential Impact. Our research involves several stakeholders: subscribers gain stronger privacy protection against tracking, while MNOs and SNOs face reduced visibility and increased operational complexity. Regulators must balance privacy with public safety. We acknowledge that our work protects the privacy of individual users, but may complicate legitimate law enforcement. We carefully consider this tension in our design, ultimately prioritising individuals from unauthorised tracking while maintaining structural affordances for lawful oversight when properly authorised.

Dual-Use Concerns. Enhanced privacy may facilitate the evasion of legitimate monitoring. Following GDPR data minimisation principles, LPG empowers users to control their disclosure levels. However, to accommodate varying jurisdictional requirements, our approach allows MNOs to enforce regional restriction policies for service access. Crucially, the framework remains compatible with lawful intercept mechanisms for targeted monitoring when legally mandated.

Responsible Disclosure. We identified vulnerabilities in existing satellite protocols and disclosed them responsibly to advance scientific understanding and guide secure development. To ensure reproducibility without compromising the operational security of our live academic satellite system (BUPT-3, TY46), we provide all evaluation artifacts (e.g., scripts and datasets) but selectively withhold satellite-specific code. These components will be released after the satellite completes its in-orbit phase.

Data Protection. Our experiments prioritise data minimisation. We conducted evaluations on our academic LEO satellite platform for research purposes, and its specifications align with existing LEO satellites equipped with COTS computing devices. We exclusively used devices controlled by informed research team members, ensuring no uninformed users were exposed to risk. We collect only timing and performance metrics necessary to assess feasibility, and all identifiers are anonymised. This approach validates real-world viability while strictly preventing risk exposure to the public.

References

- [1] sysmoisim-sja2. <https://osmocom.org/projects/cellular-infrastructure/wiki/SysmoISIM-SJA2>.
- [2] Google fi wireless, 2025. <https://fi.google.com/>.
- [3] 3GPP. Tr23.401: General packet radio service (gprs) enhancements for evolved universal terrestrial radio access network (e-utran) access, 2015.
- [4] 3GPP. Tr23.737: Study on architecture aspects for using satellite access in 5g, 2019.
- [5] 3GPP. Tr38.821: Solutions for nr to support non-terrestrial networks (ntn), 2019.

- [6] 3GPP. Tr38.331: Nr; radio resource control (rrc); protocol specification, 2022.
- [7] 3GPP. Tr24.501: Non-access-stratum (nas) protocol for 5g system (5gs), 2025.
- [8] Ghous Amjad, Kevin Yeo, and Moti Yung. Rsa blind signatures with public metadata. *PETS*, 2025.
- [9] Kim Ancin. 5g+leo: Verizon and project kuiper team up to develop connectivity solutions, 2021. <https://www.verizon.com/about/news/5g-leo-verizon-project-kuiper-team>.
- [10] Android. Android telephonymanager api, 2025. <https://tinyurl.com/5aswc3b7>.
- [11] Evelyn Janeidy Arevalo. Salt partners with spacex to provide starlink satellite-to-cellular connection through-out switzerland, 2023. <https://www.tesmanian.com/blogs/tesmanian-blog/swiss>.
- [12] Tran Bao and Attorney Patent. Rocket launch costs (2020-2030):how cheap is space travel becoming?, 2025. <https://patentpc.com/blog/rocket-launch-costs-2020-2030-how-cheap-is-space-travel-becoming-latest-pricing-data>.
- [13] Leda Bargiotti, Inge Gielis, Bram Verdegem, Pieter Breyne, Francesco Pignatelli, Paul Smits, and Ray Boguslawski. Guidelines for public administrations on location privacy: European union location framework. Technical report, Joint Research Centre, 2016.
- [14] Jeff Baumgartner. At&t will tap into oneweb’s satellite network to reach remote areas, 2021. <https://www.lichtreading.com/satellite/at-t-will-tap-into-oneweb-s-satellite-network-to-reach-remote-areas>.
- [15] Camila CS Caiado and Pushpa N Rathie. Polynomial coefficients and distribution of the sum of discrete uniform variables. In *Eighth Annual Conference of the Society of Special Functions and their Applications, Pala, India, Society for Special Functions and their Applications*, 2007.
- [16] Hawthorne Calif. First spacex satellites launch for breakthrough direct to cell service with t-mobile, 2024. <https://www.t-mobile.com/news/un-carrier/first-spacex-satellites-launch-for-breakthrough-direct-to-cell-service-with-t-mobile>.
- [17] CCPA. California proposes ccpa update on location data rules, 2025.
- [18] D. Dolev and A. Yao. On the security of public key protocols. *IEEE TIT*, 1983.
- [19] Jens Ernstberger, Chengru Zhang, Luca Ciprian, Philipp Jovanovic, and Sebastian Steinhorst. Zero-knowledge location privacy via accurate floating-point snarks. In *IEEE S&P*, 2025.
- [20] ESA. Space debris by the numbers, 2025. https://www.esa.int/Space_Safety/Space_Debris/Space_debris_by_the_numbers.
- [21] FCC. Fcc advances supplemental coverage from space framework, 2024. <https://www.fcc.gov/document/fcc-advances-supplemental-coverage-space-framework>.
- [22] FCC. Interim 911 requirements for supplemental coverage from space, 2024. <https://www.fcc.gov/Interim911Requirements-SupplementalCoveragefromSpace>.
- [23] Electronic Frontier Foundation. Satphones, syria, and surveillance, 2012. <https://www.eff.org/deeplinks/2012/02/satphones-syria-and-surveillance>.
- [24] Lorenzo Franceschi-Bicchierai. Us fines telcos \$200m for sharing customer location data without consent, 2024. <https://techcrunch.com/2024/04/30/us-fines-telcos-200m-for-sharing-customer-location-data-without-consent/>.
- [25] Matthew Franklin and Moti Yung. Communication complexity of secure computation. In *ACM STOC*, 1992.
- [26] Sanjam Garg, Aarushi Goel, Abhishek Jain, Guru-Vamsi Policharla, and Sruthi Sekar. zkSaas: zero-knowledge snarks as a service. In *USENIX Security*, 2023.
- [27] Emma Gatti and Mark Linder. Starlink’s rise as a geopolitical disruptor, 2025. <https://www.e-ir.info/2025/02/26/starlinks-rise-as-a-geopolitical-disruptor/>.
- [28] Víctor Gayoso Martínez, Luis Hernández Encinas, and Carmen Sánchez Ávila. A survey of the elliptic curve integrated encryption scheme. 2010.
- [29] Sitian Huang, Dongchao Ma, Yuzhu Jin, and Mingwei Xu. Location-based prefix aggregation in satellite-ground networks. In *IEEE IPCCC*, 2023.
- [30] Francesco Intoci, Julian Sturm, Daniel Fraunholz, Apostolos Pyrgelis, and Colin Barschel. P3li5: practical and confidential lawful interception on the 5g core. In *2023 IEEE Conference on Communications and Network Security (CNS)*, pages 1–9. IEEE, 2023.
- [31] ITU. Presentations-wrs, 2024. <https://www.itu.int/wrs-24/programme/presentations/>.

- [32] Eric Jedermann, Martin Strohmeier, Vincent Lenders, and Jens Schmitt. Record: A reception-only region determination attack on leo satellite users. In *USENIX Security*, 2024.
- [33] Hema Kadia. Esa, telesat, and amarisoft achieve first 5g 3gpp ntn link over leo, 2025. <https://tecknexus.com/5gnews-all/esa-telesat-and-amarisoft-achieve-first-5g-3gpp-ntn-link-over-leo/>.
- [34] KDDI. Ffirst spacex satellites launch for direct to cell service with kddi, 2024. https://newsroom.kddi.com/english/news/detail/kddi_pr-1074.html.
- [35] David Koisser, Richard Mitev, Marco Chilese, and Ahmad-Reza Sadeghi. Don't shoot the messenger: Localization prevention of satellite internet users. In *IEEE S&P*, 2024.
- [36] David Koisser, Richard Mitev, Nikita Yadav, Franziska Vollmer, and Ahmad-Reza Sadeghi. Orbital trust and privacy:sok on pki and location privacy challenges in space networks. In *USENIX Security*, 2024.
- [37] Qing Li, Shanguang Wang, Chenren Xu, Xiao Ma, Mengwei Xu, Ao Zhou, Ruolin Xing, Boyuan Yang, Zuo Zhu, Ying Zhang, et al. Exploring real-time satellite computing: From energy and thermal perspectives. In *IEEE RTSS*, 2024.
- [38] Yuanjie Li, Hewu Li, Wei Liu, Lixin Liu, Yimei Chen, Jianping Wu, Qian Wu, Jun Liu, and Zeqi Lai. A case for stateless mobile core network functions in space. In *ACM SIGCOMM*, 2022.
- [39] Yuanjie Li, Lixin Liu, Hewu Li, Wei Liu, Yimei Chen, Wei Zhao, Jianping Wu, Qian Wu, Jun Liu, and Zeqi Lai. Stable hierarchical routing for operational leo networks. In *MobiCom*, 2024.
- [40] Lixin Liu, Hewu Li, Yuanjie Li, Zeqi Lai, Yangtao Deng, Yimei Chen, Wei Liu, and Qian Wu. Geographic low-earth-orbit networking without qos bottlenecks from infrastructure mobility. In *IEEE/ACM IWQoS*, 2022.
- [41] Lixin Liu, Yuanjie Li, Hewu Li, Jiabo Yang, Wei Liu, Jingyi Lan, Yufeng Wang, Jiarui Li, Jianping Wu, Qian Wu, et al. Democratizing direct-to-cell low earth orbit satellite networks. In *NSDI*, 2024.
- [42] Weisen Liu, Zeqi Lai, Qian Wu, Hewu Li, Yuxuan Weng, Wei Liu, Qi Zhang, Jihao Li, Yuanjie Li, and Jun Liu. Mind the Location Leakage in LEO Direct-to-Cell Satellite Networks . In *IEEE S&P*, 2025.
- [43] Weisen Liu, Zeqi Lai, Qian Wu, Hewu Li, Qi Zhang, Zonglun Li, Yuanjie Li, and Jun Liu. In-orbit processing or not? sunlight-aware task scheduling for energy-efficient space edge computing networks. In *IEEE INFOCOM*, 2024.
- [44] Xuanming Liu, Zhelei Zhou, Yinghao Wang, Yanxin Pang, Jinye He, Bingsheng Zhang, Xiaohu Yang, and Jiaheng Zhang. Scalable collaborative zk-snark and its application to fully distributed proof delegation. In *USENIX Security*, 2025.
- [45] Zhihong Luo, Silvery Fu, Natacha Crooks, Shaddi Hasan, Christian Maciocco, Sylvia Ratnasamy, and Scott Shenker. Loca: A location-oblivious cellular architecture. In *NSDI*, 2023.
- [46] Zhihong Luo, Silvery Fu, Mark Theis, Shaddi Hasan, Sylvia Ratnasamy, and Scott Shenker. Democratizing cellular access with cellbricks. In *ACM SIGCOMM*, 2021.
- [47] Lynk. Our partners: World class mnos, 2025. <https://lynk.world/our-partners/>.
- [48] Data Privacy Manager. 20 biggest gdpr fines so far [2025], 2025. <https://dataprivacymanager.net/5-biggest-gdpr-fines-so-far-2020/>.
- [49] Sue Marek. At&t leases spectrum to ast spacemobile, 2023. <https://tinyurl.com/mr356tha>.
- [50] Gino Masini. How we reached a common vision on the architecture for 5g non-terrestrial networks in 3gpp rel-19, 2024. <https://www.ericsson.com/en/blog/2024/10/ntn-payload-architecture>.
- [51] Kieren McCarthy. Your mobile network broke the law by selling location data and may be fined millions... or maybe not, shrugs fcc, 2020. https://www.theregister.com/2020/02/04/fcc_location_data/.
- [52] Payman Mohassel and Yupeng Zhang. Secureml: A system for scalable privacy-preserving machine learning. In *IEEE S&P*, 2017.
- [53] James Murphy, John E Ward, and Brian Mac Namee. Low-power boards enabling ml-based approaches to fdir in space-based applications. 2021.
- [54] NASA. 2022 smallsat avionics chapter, 2022. <https://www.nasa.gov/wp-content/uploads/2023/05/8.-smallsat-avionics-2022.pdf>.
- [55] Alex Ozdemir and Dan Boneh. Experimenting with collaborative zk-snarks: Zero-knowledge proofs for distributed secrets. In *USENIX Security*, 2022.

- [56] Daniel Pereira. Starlink business model, 2024. <https://businessmodelanalyst.com/starlink-business-model/>.
- [57] Jason Rainbow. The shifting landscape for satellite capacity wholesalers, 2023. <https://spacenews.com/connecting-the-dots-the-shifting-landscape-for-satellite-capacity-wholesalers/>.
- [58] Jason Rainbow. Iridium pivots to standardized direct-to-device satellite services, 2024. <https://spacenews.com/iridium-pivots-to-standardized-direct-to-device-satellite-services/>.
- [59] Ling Ren, Muhammad Haris Mughees, and I Sun. Simple and practical amortized sublinear private information retrieval using dummy subsets. In *ACM CCS*, 2024.
- [60] Paul Schmitt, Jana Iyengar, Christopher Wood, and Barath Raghavan. The decoupling principle: a practical privacy framework. *ACM HotNets*, 2022.
- [61] Paul Schmitt and Barath Raghavan. Pretty good phone privacy. In *USENIX Security*, 2021.
- [62] SES. O3b meo, 2025. <https://www.ses.com/our-coverage/o3b-meo>.
- [63] Quan Shi, Liying Wang, Prosanta Gope, Qi Liang, Haowen Wang, Qirui Liu, Chenren Xu, Shangguang Wang, Qing Li, and Biplab Sikdar. LPG: Raise your location privacy game in direct-to-cell LEO satellite networks. *Cryptology ePrint Archive*, Paper 2025/2255, 2025.
- [64] Anton Shilov. Nvidia's jetson ai board is ready to go to space, 2021. <https://www.tomshardware.com/news/nvidias-jetson-ai-board-is-ready-to-go-to-space>.
- [65] Spacenews. Globalstar soars on apple's \$1.7 billion satellite investment, 2024. <https://spacenews.com/globalstar-soars-on-apples-1-5-billion-satellite-investment/>.
- [66] SpaceX. SpaceX invites world's carriers to collaborate: No more cell phone dead zones, 2022.
- [67] SpeedTalk. Flexible "pay as you go" mobile phone plans. <https://speedtalkmobile.com/pay-as-you-go-phone-plans/>.
- [68] Starlink. High-speed internet around the world, 2025. <https://www.starlink.com>.
- [69] Starlink. Starlink direct to cell, 2025. <https://www.starlink.com/business/direct-to-cell>.
- [70] Uber. H3: Uber's hexagonal hierarchical spatial index, 2018. <https://www.uber.com/en-DE/blog/h3/>.
- [71] Anne Wainscott-Sargent. Sizing up the satellite-to-cell opportunity, 2022. <https://interactive.satellite.com/via/december-2022/sizing-up-the-satellite-to-cell-opportunity>.
- [72] Chao Wang, Xiao Ma, Ruolin Xing, Sisi Li, Ao Zhou, and Shangguang Wang. Delay-and resource-aware satellite upf service optimization. *IEEE TMC*, 2024.
- [73] Liying Wang, Qing Li, Yuhan Zhou, Zhaofeng Luo, Donghao Zhang, Shangguang Wang, Xuanzhe Liu, and Chenren Xu. Emulating space computing networks with rhone. In *USENIX ATC*, 2025.
- [74] Jess Weatherbed. Starlink's direct-to-cell satellite service is the first to receive fcc approval, 2024. <https://www.theverge.com/2024/11/27/24307394/starlink-spacex-tmobile-direct-to-cell-satellite-fcc-approval>.
- [75] Zack Whittaker. Us cell carriers are selling access to your real-time phone location data, 2018. <https://www.zdnet.com/article/us-cell-carriers-selling-access-to-real-time-location-data/>.
- [76] Privacy World. Eu digital services act in full force, 2024. <https://www.privacyworld.blog/2024/02/eu-digital-services-act-in-full-force/>.
- [77] Jiasheng Wu, Shaojie Su, Xiong Wang, Jingjing Zhang, and Yue Gao. Accelerating handover in mobile satellite network. In *IEEE INFOCOM*, 2024.
- [78] Ruolin Xing, Xiao Ma, Ao Zhou, Schahram Dustdar, and Shangguang Wang. From earth to space: A first deployment of 5g core network on satellite. *China Communications*, 2023.
- [79] Ruolin Xing, Mengwei Xu, Ao Zhou, Qing Li, Yiran Zhang, Feng Qian, and Shangguang Wang. Deciphering the enigma of satellite computing with cots devices: Measurement and analysis. In *MobiCom*, 2024.
- [80] Yang Yang, Quan Shi, Prosanta Gope, Behzad Abdolmaleki, and Biplab Sikdar. Pcus: Pretty good user security for thick mvnos with a novel sanitizable blind signature. In *IEEE S&P*, 2025.
- [81] Hexuan Yu, Changlai Du, Yang Xiao, Angelos Keromytis, Chonggang Wang, Robert Gazda, Y Thomas Hou, and Wenjing Lou. Aaka: An anti-tracking cellular authentication scheme leveraging anonymous credentials. In *NDSS*, 2024.

Identity-Hiding: $\text{Game}_{\text{IH}}(\lambda, \mathcal{A})$

- A challenger \mathcal{C} generates system parameters and establishes MNO-SNO agreements in LPG. An adversary \mathcal{A} controls an SNO for λ .
- Challenge Phase:
 - (i) \mathcal{A} outputs two UE identities id_0, id_1 and auxiliary information aux . (e.g. internal MNO identifiers)
 - (ii) \mathcal{C} selects $b \leftarrow \{0, 1\}$ uniformly at random.
 - (iii) \mathcal{C} executes LPG protocols with UE_{id_b} (UE with id_b) for polynomial round of sessions.
 - (iv) \mathcal{A} observes all SNO-visible protocol transcripts.
- \mathcal{A} outputs a guess $b' \in \{0, 1\}$, wins if $b' = b$.

Figure 16: Definition of Identity-Hiding $\text{Game}_{\text{IH}}(\lambda, \mathcal{A})$

A Security Analysis

A.1 Theorem 10: Definition and Proof

Theorem 10 establishes that LPG provides robust identity hiding against a semi-honest SNO via two key aspects: protection of identity within the protocol’s content, and resilience against trajectory inference attacks.

In an LPG session, all information observed by the SNO during its interaction with a UE is computationally indistinguishable from information generated by an ideal functionality that does not use the UE’s MNO-managed persistent identity. This holds under three assumptions that: *i*) the MNO and SNO does not collude; *ii*) the Tok_{UE} scheme is unlinkable; *iii*) the ECIES encryption scheme is IND-CPA secure.

We provide a formal definition of Identity-Hiding in Theorem 10, and give the game-based proof.

Definition 3 (Identity-Hiding). The advantage of a PPT adversary \mathcal{A} in breaking the Identity-Hiding game of a scheme Π is defined as: $\mathcal{A}_{\text{IH}}^{\text{IH}}(\lambda) = |\Pr[\mathcal{A} \text{ wins } \text{Game}_{\text{IH}}(\lambda, \mathcal{A})] - 1/2|$. (Game_{IH} is defined in Figure 16) A scheme Π is Identity-Hiding to \mathcal{A} if for any PPT adversary \mathcal{A} , $\mathcal{A}_{\text{IH}}^{\text{IH}}(\lambda)$ is negligible in λ .

Proof of Theorem 10. (1) *Protocol content.* We prove this through a sequence of hybrid games \mathbf{G}_0 - \mathbf{G}_4 , showing that each adjacent pair is computationally indistinguishable for the SNO adversary \mathcal{A} . Let $\text{Adv}(\mathbf{G}_i)$ be \mathcal{A} ’s advantage in game \mathbf{G}_i .

- Game \mathbf{G}_0 : The real Identity-Hiding Game_{IH} for UE_{id_b} .
- Game \mathbf{G}_1 : Same as \mathbf{G}_0 , but when UE_{id_b} generates or uses Tok_{UE} , it is replaced by tokens perfectly simulated by an unlinkability simulator. Due to the (ϵ_1, t_1) -unlinkability of Tok_{UE} , SNO cannot distinguish which one is linked to id_0 , so $|\text{Adv}(\mathbf{G}_0) - \text{Adv}(\mathbf{G}_1)| \leq \epsilon_1$.
- Game \mathbf{G}_2 : Same as \mathbf{G}_1 , but all ephemeral keys $(\text{spk}_u, \text{ssk}_u)$ used by UE_{id_b} in OAKA and the mutually agreed session key K_{sess} , are replaced with uniformly random independent keys. The SNO observes spk_u encrypted via ECIES [28]. Due to the (ϵ_2, t_2) -IND-CPA security

of ECIES and the security of ECDH, $|\text{Adv}(\mathbf{G}_1) - \text{Adv}(\mathbf{G}_2)| \leq \epsilon_2$.
- Game \mathbf{G}_3 : Same as \mathbf{G}_2 , but if UE_{id_b} submits ZKLP proofs $\pi_{i,p}$, these are replaced by proofs generated by a zero-knowledge simulator for ZKLP. By the zero-knowledge property of ZKLP, $|\text{Adv}(\mathbf{G}_2) - \text{Adv}(\mathbf{G}_3)| \leq \text{negl}(\lambda)$.
- Game \mathbf{G}_4 : Same as \mathbf{G}_3 , but the challenger now internally uses $\text{UE}_{\text{id}_{1-b}}$ instead of UE_{id_b} . In \mathbf{G}_3 , all cryptographic components visible to \mathcal{A} that could depend on b have been replaced by simulated or random values. The view of SNO consists of simulated unlinkable tokens, communications encrypted with random keys, and simulated zero-knowledge proofs. None of these depend on b . Thus, \mathcal{A} ’s view in \mathbf{G}_3 is identically distributed whether $b = 0$ or $b = 1$. Hence, $\text{Adv}(\mathbf{G}_3) = \text{Adv}(\mathbf{G}_4)$.

In \mathbf{G}_4 , since \mathcal{A} ’s view is independent of the challenge bit b , \mathcal{A} ’s probability of guessing b correctly is exactly $1/2$. So, $\text{Adv}(\mathbf{G}_4) = 0$. Therefore, $\mathcal{A}_{\text{IH}}^{\text{IH}}(\lambda) \leq \epsilon_1 + \epsilon_2 + \text{negl}(\lambda)$.

(2) *Trajectory Inference.* A semi-honest SNO adversary might attempt to link many anonymous sessions over long periods to form a trajectory that uniquely identifies a UE [45]. UEs submit fresh and unlinkable Tok_{UE} for distinct OAKA sessions; the Tok_{UE} tokens serve as unlinkable session identifiers. Therefore, the UE further employs operational strategies in LOCA [45] (periodic reattachment with average period P , randomised attachment timing within window Wd) in an environment with NU other users. By Lemma 1, the probability $\Pr_{\text{link_traj}}$ that the SNO can correctly link N_{attach} such anonymous sessions is $\leq (1/ND)^{(N_{\text{attach}}-1)}$, where $ND \geq 1 + NU \cdot (Wd/P)$. For realistic system parameters where $ND > 1$ (e.g., even a small NU or Wd/P ratio makes $ND \geq 2$, such as when $Wd/P > 0.03$, $NU > 40$, then the number of candidate detachments ND is greater than 2), this probability decreases exponentially with N_{attach} . Thus, for any trajectory long enough to be behaviorally identifying (e.g., $N_{\text{attach}} > 6$), the probability of the SNO correctly reconstructing it or inferring identity is negligible (< 0.01). \square

A.2 Theorem 11: Definition and Proof

Theorem 11 asserts that a semi-honest MNO cannot determine a UE’s precise SNO-observed location from protocol interactions. The MNO’s view is restricted to either user-consented ZKLP disclosures or aggregate settlement data.

Definition 4 (Location-Obscurity). The advantage of a PPT adversary \mathcal{A} in breaking the Location-Obscurity game of a scheme Π is defined as: $\mathcal{A}_{\text{LO}}^{\text{LO}}(\lambda) = |\Pr[\mathcal{A} \text{ wins } \text{Game}_{\text{LO}}(\lambda, \mathcal{A})] - 1/2|$. (Game_{LO} is defined in Figure 17) A scheme Π holds Location-Obscurity to \mathcal{A} if for any PPT adversary \mathcal{A} , $\mathcal{A}_{\text{LO}}^{\text{LO}}(\lambda)$ is negligible in λ .

Proof of Theorem 11. We analyze the information channels available to the MNO adversary \mathcal{A} in the Location-Obscurity Game to distinguish between location traces L_0, L_1 for a target UE*.

- Game \mathbf{G}_0 : The real Location-Obscurity Game_{LO} .

Location-Obcurity: $\text{Game}_{\text{LO}}(\lambda, \mathcal{A})$

- A challenger \mathcal{C} generates system parameters. An adversary \mathcal{A} controls an MNO for λ , it observes polynomial LPG executions including DualToken, SMTs and SCAC.
- Challenge Phase:
 - (i) \mathcal{C} selects a target UE* known to \mathcal{A} .
 - (ii) \mathcal{C} selects two distinct location patterns L_0, L_1 (e.g., different geographic regions, mobility patterns)
 - (iii) \mathcal{C} selects $b \in \{0, 1\}$ and runs UE* with pattern L_b .
 - (iv) \mathcal{A} observes aggregate claims, ZKLP disclosures, but no SNO-observed fine-grained location information.
- \mathcal{A} outputs a guess $b' \in \{0, 1\}$, wins if $b' = b$.

Figure 17: Definition of Location-Obcurity $\text{Game}_{\text{LO}}(\lambda, \mathcal{A})$

\mathbf{G}_1 : Same as \mathbf{G}_0 , but when generating SCAC proofs $\pi_{\text{agg},i}$, it is replaced by the proof generated by the zero-knowledge simulator $\mathcal{S}_{\text{SCZK}}$. $\mathcal{S}_{\text{SCZK}}$ receives only public statements including (\mathcal{X}_i, Z_i) and AggTable , and all other protocol elements remain unchanged. The adversary that efficiently distinguishes \mathbf{G}_0 from \mathbf{G}_1 can serve as a zero-knowledge distinguisher against SCZK. By the zero-knowledge property of SCZK, $|\text{Adv}(\mathbf{G}_0) - \text{Adv}(\mathbf{G}_1)| \leq \epsilon_3$.

◦ Game \mathbf{G}_2 : Same as \mathbf{G}_1 , but all hash values $H(\text{secret}_j)$ in aggregate usage tables are replaced with uniformly random values $r_j \leftarrow \{0, 1\}^\ell$. The secret values are computed as $\text{secret}_j = \text{PRF}(\text{key}_\theta, \eta_j \| \text{metadata}_j)$ where: key_θ is derived from SNO token Tok_{SNO} , η_j is UE token nonce in SIM card. The adversary \mathcal{A} that efficiently distinguishes \mathbf{G}_1 from \mathbf{G}_2 can efficiently distinguish $\text{PRF}(\text{key}, \cdot)$ and random values, breaking the pseudorandomness. By the (ϵ_4, t_4) -security of PRF, $|\text{Adv}(\mathbf{G}_1) - \text{Adv}(\mathbf{G}_2)| \leq \epsilon_4$.

◦ Game \mathbf{G}_3 : Same as \mathbf{G}_2 , but challenge bit b is ignored, LPG execution always use location pattern L_0 for UE*. All information visible to \mathcal{A} is now independent of the true challenge. Now in \mathcal{A} 's view, the only information potentially distinguishing L_0 from L_1 comes from the aggregate values (\mathcal{X}_i, Z_i) . By Lemma 1, for session tables with N_{total} entries and claims of size \mathcal{X}_i , when $|\mathcal{X}_i|/N_{\text{total}} \in [L_{\text{val}}, U_{\text{val}}]$, the expected number of distinct subsets summing to Z_i grows as $\alpha^{N_{\text{total}}}$ for some $\alpha > 1$. Therefore, the probability that \mathcal{A} can identify the specific subset corresponding to UE*'s sessions, and thus distinguish location patterns is at most $\alpha^{-\Omega(N_{\text{total}})}$. This is robust even if the user submits ZKLP in SMTs (§4.4). ZKLP disclosures provide only user-chosen information and do not leak SNO-observed fine-grained locations.

In \mathbf{G}_3 , \mathcal{A} 's view is independent of challenge bit b , so $\text{Adv}(\mathbf{G}_3) = 0$. Therefore, $\mathcal{A}_{\text{LO}}^{\text{LO}}(\lambda) \leq \epsilon_3 + \epsilon_4 + \alpha^{-\Omega(N_{\text{total}})}$. For large N_{total} (e.g., $N_{\text{total}} = 4096$ with session groups of duration $T = 20\text{s}$), the exponential term dominates, ensuring that the advantage of \mathcal{A} is negligible. \square

Finally, based on the results in security analysis, we provide a practical example of parameter deployment as follows.

Corollary 1 (Privacy Parameters for Deployment). For a

practical LPG deployment with $n_{\text{sess}} = 4096$ total sessions in the settlement table, session groups of duration $T \geq 20\text{s}$, and claim ratios $k/n_{\text{sess}} \in [1/150, 1/2]$, LPG achieves:

(1) *Identity Hiding (SNO)*: The advantage $\mathcal{A}_{\text{IH}}^{\text{IH}}$ is negligible, assuming the security of the underlying Tok_{UE} and ECIES primitives.

(2) *Trajectory Linking Resistance (SNO)*: The probability $\text{Pr}_{\text{link_traj}} < 0.01$ for trajectories longer than 6 attachments, assuming moderate environmental parameters ($W_d/P > 0.03$, $N_U > 40$) [45].

(3) *Location Obscurity (MNO)*: The advantage $\mathcal{A}_{\text{LO}}^{\text{LO}}$ is negligible (set $\alpha > 1.1$), dominated by the exponential term from aggregate sum ambiguity (Lemma 2).

These concrete bounds demonstrate that LPG provides strong, quantifiable location privacy guarantees while maintaining practical system performance.

A.3 Additional Threat Analysis

Side-Channel Threat. LPG is designed to provide robust protocol-layer location privacy (NAS and above). It prevents network operators from tracking users at the cellular network protocol level. LPG have not modeled side-channel adversaries at the physical or signal layer [32, 35, 42] (§2.2). Several effective physical-layer countermeasures already exist to safeguard D2C user location privacy, such as identifier changing, signal randomization, encryption, phantom routing, and the generation of fake traffic [32, 36, 42]. By integrating these cross-layer defenses with LPG's protocol-layer solutions, we can achieve comprehensive location privacy for D2C users, representing a promising avenue for future research.

Resilience to SIM Compromise. The one-time-use property of Tok_{UE} depends on SIM tamper resistance. If a SIM is cloned or its state is manipulated, a malicious UE may replay Tok_{UE} . To mitigate this risk, MNOs could choose to deploy a MOSAIC-like server-side defense [41]: during settlement, they perform cross-SNO deduplication to flag reuse of the same Tok_{UE} nonce (or its commitment, $H(\text{secret})$) by multiple SNOs or across different Tok_{SNO} batches, indicating a multi-spend. This check enables rapid detection and revocation of compromised credentials, containing the impact of SIM compromise.

Collusion between SNO and MNO. Similar to prior split-architecture privacy solutions [45, 61, 80], LPG cannot prevent location leakage if the SNO and MNO actively collude. In LPG, the MNO holds the mapping between the user's stable identity and their authorized tokens (Tok_{UE}), while the SNO observes the precise location and token usage during the session. If these two entities were to share their internal states, they could trivially re-associate the user's identity with their location trajectory. However, such active collusion is usually prevented by strict legal frameworks (e.g., GDPR [48]) and the nature of cross-border D2C partnerships (§3.2), where operators are legally bound to protect subscriber data.