CALM: UE-Side Mobility Management for Integrated TN-NTN Networks

Joonyoung Jeong
Department of Mathematics
Korea Military Academy
Seoul, Republic of Korea
13-10113@kma.ac.kr

***Abstract*—Integrated terrestrial and non-terrestrial networks (TN-NTN) are a leading way to keep coverage continuous in 5G-Advanced and beyond. While traffic-aware radio resource management (RRM) has been studied extensively, what its decisions cost the user equipment (UE) has been ignored, even though UE measurement, reporting, and signaling are now among the main drivers of UE energy in NR. We propose CALM, a fast-timescale UE-side mobility control layer on top of a network-side RRM. To our knowledge, CALM is the first scheme to treat the UE-side cost of mobility as an explicit control variable rather than a free byproduct: a UE should spend energy on a radio operation only when the result can change a decision. Each UE forms short-term risk indicators from its own radio observations and applies closed-form thresholds, relaxing measurement while the link is stable, preparing a conditional handover when failure is likely, and turning to the NTN tier only when terrestrial recovery is unlikely. On a 3GPP-standard TN-NTN simulator with executed handovers, CALM matches the reliability of always-prepare baselines at substantially lower UE signaling.**

***Index Terms*—Non-terrestrial networks; 5G-Advanced; 6G; radio resource management; mobility management; conditional handover; measurement relaxation; UE energy efficiency; signaling overhead.**

I. Introduction

Terrestrial networks (TNs) deliver high capacity and low latency in dense deployments. However, they cannot easily keep coverage continuous in, e.g., coverage holes, rural regions, and high-mobility settings. Therefore, non-terrestrial networks (NTNs), built from low Earth orbit (LEO) satellites and aerial platforms, have been studied as complementary access layers for 5G-Advanced and beyond [1]–[4]. In 3GPP NR, mobility is managed through measurement configuration and reporting, radio resource control (RRC) reconfiguration, radio-link failure recovery, and conditional handover (CHO) [5], [6]. Integrating TN and NTN improves coverage, but it raises a new control question: when should a UE stay on TN, when should it prepare a TN handover, and when should it probe or fall back to an NTN candidate?

A major line of work on integrated TN-NTN operation is traffic-aware RRM. A representative framework jointly optimizes the bandwidth split, UE association, transmit power, and TN base-station activation [7]. Its key insight is that NTN both extends coverage and saves TN energy, since selected base stations can reduce power or shut down under low load. More broadly, such RRM targets network-side objectives such as sum-log throughput, fairness, TN energy, and load balance. However, executing an association, handover, or fallback decision is not free at the UE. A UE may have to measure serving and neighbour cells, evaluate reporting criteria, transmit reports, scan beams or candidate cells, probe an NTN candidate, prepare CHO, and exchange RRC signaling. These operations consume UE energy and signaling, yet existing TN-NTN RRM formulations largely abstract them away.

This abstraction matters, because UE power saving is now a first-order design goal in NR. 3GPP studies target power reduction in RRM measurements and measurement relaxation across the IDLE, INACTIVE, and CONNECTED states [8], and NR measurement processing, reporting-criteria evaluation, and report transmission are standardized RRC-configured operations [5], [6]. Therefore, suppressing unnecessary measurements and reports is a mobility control decision in itself, one that trades UE energy and signaling against mobility robustness. Existing TN-NTN RRM, including traffic-aware frameworks [7], models infrastructure-side energy and allocation, but it leaves the UE-side cost of realizing those decisions largely unmodeled.

A second related line of work is learning-assisted mobility robustness. Recent radio-link failure (RLF) prediction frameworks such as GenTrap [9] estimate future link degradation from radio key performance indicators and contextual observations. However, they deliver only a predictor and leave open the protocol-level question of *when* to suppress measurements, prepare CHO, probe NTN, or fall back. In parallel, CHO improves mobility robustness by decoupling handover preparation from execution, at the cost of extra signaling, which fast CHO (FCHO) curbs by reusing prepared configurations [10], [11]. These lines complement TN-NTN RRM, but individually they do not close the loop from predicted RLF or handover-failure (HOF) risk to protocol-level actions. Prediction alone does not say whether to suppress or enhance measurements, prepare CHO, probe NTN, execute fallback, or return to TN, and CHO and FCHO alone do not treat NTN as a candidate tier.

These gaps sharpen when traffic is latency-sensitive and NTN resources are limited. Satellite access extends coverage, but unnecessary NTN residence raises service delay and consumes scarce satellite beam, spectrum, and payload power. A practical policy should not offload to NTN merely because it is

979-8-3195-0167-7/26/$31.00 ©2026 IEEE

available. It should prefer TN whenever connectivity is stable and the quality of service (QoS) is sufficient, and prepare NTN probing or fallback only when the estimated dwell time, RLF/HOF risk, and fallback necessity together indicate that terrestrial recovery is unlikely. We call this principle TN-first. It is not a 3GPP-mandated rule, but a QoS-, latency-, and resource-aware policy motivated by the asymmetric delay and resource profiles of TN and NTN access.

This paper asks whether network-side TN-NTN RRM can be extended with a fast-timescale, signaling-aware mobility control layer that accounts for UE-side energy and signaling while preserving a TN-first use of NTN. Network-side RRM [7] optimizes infrastructure objectives and treats UE measurement and signaling as free. Conditional-handover and link-failure-prediction work [9], [10], [12] improves mobility robustness without pricing the UE actions it triggers. CALM instead makes the UE-side cost of mobility an explicit control variable. To our knowledge, it is the first scheme to close the loop from per-UE risk estimates to cost-aware measurement, conditional-handover, and TN-first NTN decisions on top of a network-side RRM substrate.

To answer this question, we propose CALM (Cost-Aware Lightweight Mobility-control), a signaling-efficient, mobility-aware framework for integrated TN-NTN systems. The main contributions are as follows:

- We formulate a generalized TN-NTN control problem that, for the first time, augments traffic-aware network-side RRM with an explicit UE-side mobility control cost, thereby turning UE energy and signaling from a free byproduct into an explicit objective term.
- We propose CALM, a two-timescale framework in which a slow network-side RRM layer updates association, bandwidth split, transmit power, and BS activation, while a fast layer selects per-UE protocol actions from short-term RLF/HOF risk indicators, TN dwell time, and fallback necessity formed from radio and control-plane observations.
- We give the fast layer a closed-form, preparation-reuse threshold structure (Proposition 1), in which CHO is prepared when terrestrial recovery remains likely, measurement is relaxed when the link is estimated stable, and NTN is probed only when fallback necessity is high, so that the controller is $O(K)$ per slot and needs no rollout.
- We evaluate CALM on a 3GPP-standard TN-NTN simulator with 38.901/ 38.811 channels, executed handovers, and T310/A3 failure semantics. CALM sustains baseline reliability at a fraction of the signaling and measurement cost; the saving comes chiefly from relaxing measurement when the link is stable and is robust across UE mobility, cell density, and a sparse rural-macro deployment.

Section II describes the system model and problem formulation, Section III presents the CALM framework and its threshold policy, Section IV reports the simulation results, and Section V concludes.

II. System Model

Our study focuses on the downlink of an integrated terrestrial and non-terrestrial network.

Sections II-A to II-D restate the network, channel, rate, and energy models of [7], and (25) adapts its objective.

A. Network Model

Following [7], we consider an integrated TN-NTN system with M terrestrial macro base stations (MBSs) and N satellite MBSs (or beams), totalling $L = M+N$ access nodes. Let $\mathcal{T} = \{1, \dots, M\}$ and $\mathcal{S} = \{M+1, \dots, L\}$ denote the terrestrial and satellite sets, so $\mathcal{B} = \mathcal{T} \cup \mathcal{S}$. The UE set is $\mathcal{U} = \{1, \dots, K\}$.

The total bandwidth W is split orthogonally and dynamically between the two tiers; with $\epsilon \in [0, 1]$ allocated to the satellite tier, the bandwidth available at MBS j is

$$W_j = \begin{cases} \epsilon W, & j \in \mathcal{S}, \\ (1-\epsilon)W, & j \in \mathcal{T}. \end{cases} \tag{1}$$

The UE-MBS association variable is

$$x_{ij} = \begin{cases} 1, & \text{if UE } i \text{ is associated with MBS } j, \\ 0, & \text{otherwise.} \end{cases} \tag{2}$$

Under single connectivity,

$$\sum_{j \in \mathcal{B}} x_{ij} = 1, \quad \forall i \in \mathcal{U}, \tag{3}$$

and the load of MBS j is

$$k_j = \sum_{i \in \mathcal{U}} x_{ij}. \tag{4}$$

B. Channel Model

Terrestrial links follow the 3GPP TR 38.901 channel model [13] (distance-dependent path loss, line-of-sight (LoS)/non-LoS (NLoS) conditions, log-normal shadow fading, and outdoor-to-indoor (O2I) penetration loss). Satellite links follow the TR 38.811 NTN channel model [1] (free-space loss over the slant range, clutter, scintillation, and building-entry loss).

Let β_{ij} denote the large-scale channel gain between UE i and MBS j. For $j \in \mathcal{T}$, the large-scale gain is

$$\beta_{ij} = \Big\lfloor G_{\mathrm{TX}} + G_{\mathrm{UE}} + PL^{\mathrm{b}}_{ij} + SF_{ij} + PL^{\mathrm{tw}}_{ij} + PL^{\mathrm{in}}_{ij} + \mathcal{N}(0, \sigma_p^2) \Big\rfloor. \tag{5}$$

where all terms are in dB before conversion to linear scale. G_{TX} and G_{UE} are the MBS and UE antenna gains; PL^{b}_{ij} is the basic outdoor path loss; $SF_{ij} \sim \mathcal{N}(0, \sigma^2_{SF})$ is shadow fading; PL^{tw}_{ij} and PL^{in}_{ij} are the external-wall and indoor penetration losses; and $\mathcal{N}(0, \sigma_p^2)$ is the random component of the penetration loss. The path loss and shadowing parameters depend on the LoS/NLoS condition.

For $j \in \mathcal{S}$, the large-scale gain is

$$\beta_{ij} = \Big\lfloor G_{\mathrm{TX}} + G_{\mathrm{UE}} + PL^{\mathrm{b}}_{ij} + SF_{ij} + CL_{ij} + PL^{\mathrm{s}}_{ij} + PL^{\mathrm{e}}_{ij} \Big\rfloor. \tag{6}$$

where CL_{ij}, PL^{s}_{ij}, and PL^{e}_{ij} are the clutter, scintillation, and building-entry losses, all depending on the LoS condition and on the satellite elevation angle θ_i, which is obtained from the satellite and UE positions.

C. Signal-to-Interference-plus-Noise Ratio (SINR) and Rate

Because the tiers are allocated orthogonal bandwidth, there is no inter-tier interference. The large-scale SINR for UE i on MBS j is

$$\gamma_{ij} = \frac{\beta_{ij} p_j}{\sum_{j' \in \mathcal{I}_j} \beta_{ij'} p_{j'} + \sigma^2}, \tag{7}$$

where p_j is the transmit power per resource element at MBS j, $\mathcal{I}_j$ is the set of MBSs interfering with MBS j within the same tier or reuse group, and σ^2 is the noise power per resource element.

Under equal sharing of W_j among the k_j UEs it serves, the mean throughput of UE i from MBS j is

$$R_{ij} = \frac{W_j}{k_j} \log_2(1 + \gamma_{ij}), \qquad k_j \geq 1, \tag{8}$$

and the UE's served throughput is

$$R_i = \sum_{j \in \mathcal{B}} x_{ij} R_{ij}. \tag{9}$$

D. Energy Consumption Model

Following a standard network energy model [7], the consumption of TN MBS j is the sum of baseline, dynamic, and static active-mode components:

$$Q_j(p_j) = P_0 + p_j + \psi_j \|p_j\|_0, \quad j \in \mathcal{T}, \tag{10}$$

where P_0 is the baseline, p_j the dynamic transmit-power component, ψ_j the static active-mode component, and $\|p_j\|_0$ is one if MBS j is active and zero otherwise, so the total TN energy is $E_{\mathrm{TN}} = \sum_{j \in \mathcal{T}} Q_j(p_j)$. The satellite tier serves as an NTN access layer for service continuity and fallback. CALM focuses on UE-side control cost and TN energy, and discourages satellite resource use through the NTN-use penalty defined below.

E. CALM Mobility-Control Model

The preceding models define the slow RRM layer, and CALM adds a fast-timescale mobility control layer on top, so CALM runs two control loops at different rates. The slow layer is the network-side RRM of [7] built on Sections II-A to II-D: once per slow slot h it re-solves the operating state $\Omega(h)$ below. On the fast layer it adds, CALM selects per-UE protocol actions at each of the H_{f} fast slots t within slot h, while holding $\Omega(h)$ fixed. It is noteworthy that the split is by rate, not by function. $\Omega(h)$ couples all UEs and needs network-wide state, whereas the fast decisions are per-UE and must track the radio link monitoring (RLM) and A3 timers. In our evaluation we use a fast slot of $T_{\mathrm{f}} = 250$ ms and $H_{\mathrm{f}} = 10$, so a slow slot spans 2.5 s, one T_{310} window. Fig. 1 summarizes this structure. The slow RRM state is

$$\Omega(h) = \big(\mathbf{X}(h), \epsilon(h), \mathbf{p}(h)\big), \tag{11}$$

where $\mathbf{X}(h) \in \{0,1\}^{K \times L}$ collects the association variables (2), $\epsilon(h)$ is the bandwidth split, and $\mathbf{p}(h) \in \mathbb{R}^L$ is the transmit-power vector across all MBSs at slot h.

At each fast slot (h,t), UE i constructs an observation vector

$$\begin{aligned}\mathbf{s}_i(h,t) = [&\mathbf{x}^{\mathrm{r}}_i(h,t), \mathbf{x}^{\mathrm{c}}_i(h,t), \mathbf{x}^{\mathrm{b}}_i(h,t), \mathbf{x}^{\mathrm{m}}_i(h,t),\\ &\mathbf{x}^{\mathrm{n}}_i(h,t), \mathbf{x}^{\mathrm{q}}_i(h,t)]\end{aligned} \tag{12}$$

where the components denote radio, control-plane, beam, mobility, NTN-visibility, and RRM-state observations. They may include reference-signal received power and quality (RSRP/RSRQ) and SINR levels and trends, measurement-event history, UE speed and heading, and NTN visibility. Of these, this paper uses only the serving and best-neighbour SINR, the serving-SINR trend, the T310/N310 counters, and TN-cell RSRP, the inputs of the instantiation below.

A risk-estimation map $f_\theta(\cdot)$, computed from the observation, returns short-horizon mobility-risk indicators:

$$\hat{\mathbf{y}}_i(h,t) = f_\theta(\mathbf{s}_i(h,t)), \tag{13}$$

where

$$\hat{\mathbf{y}}_i(h,t) = [\hat{p}^{\mathrm{RLF}}_i(h,t), \hat{p}^{\mathrm{HOF}}_i(h,t), \hat{T}^{\mathrm{TN}}_i(h,t), \hat{F}^{\mathrm{fb}}_i(h,t)]. \tag{14}$$

where $\hat{p}^{\mathrm{RLF}}_i$ and $\hat{p}^{\mathrm{HOF}}_i$ are the estimated RLF and handover-failure (HOF) risks, $\hat{T}^{\mathrm{TN}}_i$ is the estimated TN dwell time, and $\hat{F}^{\mathrm{fb}}_i \in [0,1]$ is the estimated fallback-necessity score.

We instantiate f_θ as an *explicit, deterministic* map of the standardized RLM and A3 state, not a learned model, so that no part of the controller rests on an opaque predictor. Write γ^{s}_i and γ^{n}_i for the serving and best-neighbour SINR, $\Delta\gamma^{\mathrm{s}}_i$ for the per-slot serving-SINR trend, Q_{out} for the out-of-sync threshold, O_{A3} for the A3 offset, and $(c^{310}_i, \tau^{310}_i)$ for the TS 38.331 [6] out-of-sync counter and the residual T310 timer, with thresholds N_{310} and T_{310}. Let $\sigma(x) = 1/(1+e^{-x})$ denote the logistic function and $\zeta > 0$ a small constant. Then

$$\begin{aligned}\hat{p}^{\mathrm{RLF}}_i = {}& \tfrac{1}{2}\,\mathrm{clip}\big(c^{310}_i/N_{310}, 0, 1\big)\\ &+ \tfrac{1}{2}\,\mathbf{1}[\tau^{310}_i > 0]\,\mathrm{clip}\big(1 - \tau^{310}_i/T_{310}, 0, 1\big),\end{aligned} \tag{15}$$

$$\hat{p}^{\mathrm{HOF}}_i = \sigma\big((\gamma^{\mathrm{n}}_i - \gamma^{\mathrm{s}}_i - O_{\mathrm{A3}})/2\big), \tag{16}$$

$$\hat{T}^{\mathrm{TN}}_i = \mathrm{clip}\big((\gamma^{\mathrm{s}}_i - Q_{\mathrm{out}})/\max(-\Delta\gamma^{\mathrm{s}}_i, \zeta),\, 0,\, T_{\max}\big), \tag{17}$$

and $\hat{F}^{\mathrm{fb}}_i = \mathbf{1}[\text{no TN cell exceeds the minimum RSRP}]$. Thus $\hat{p}^{\mathrm{RLF}}_i$ rises monotonically as the RLM procedure advances toward declaration, $\hat{p}^{\mathrm{HOF}}_i$ reads the A3 margin past the offset, and $\hat{T}^{\mathrm{TN}}_i$ extrapolates the serving SINR to Q_{out} at its current slope. Each indicator is a fixed, monotone function of standardized measurements. Therefore, the estimation-noise study of Section IV is a sensitivity analysis around a named estimator rather than a defence of an unspecified box, and the bounded-error robustness in Proposition 1 acquires a concrete error model.

UE i's mobility control action $a_i(h,t)$ is drawn from

$$\begin{aligned}\mathcal{A} = \{&\mathrm{KeepTN}, \mathrm{SuppressMeas}, \mathrm{EnhanceMeas}, \mathrm{PrepareCHO},\\ &\mathrm{ProbeNTN}, \mathrm{FallbackNTN}, \mathrm{ReleasePrep}, \mathrm{ReturnTN}\}.\end{aligned} \tag{18}$$

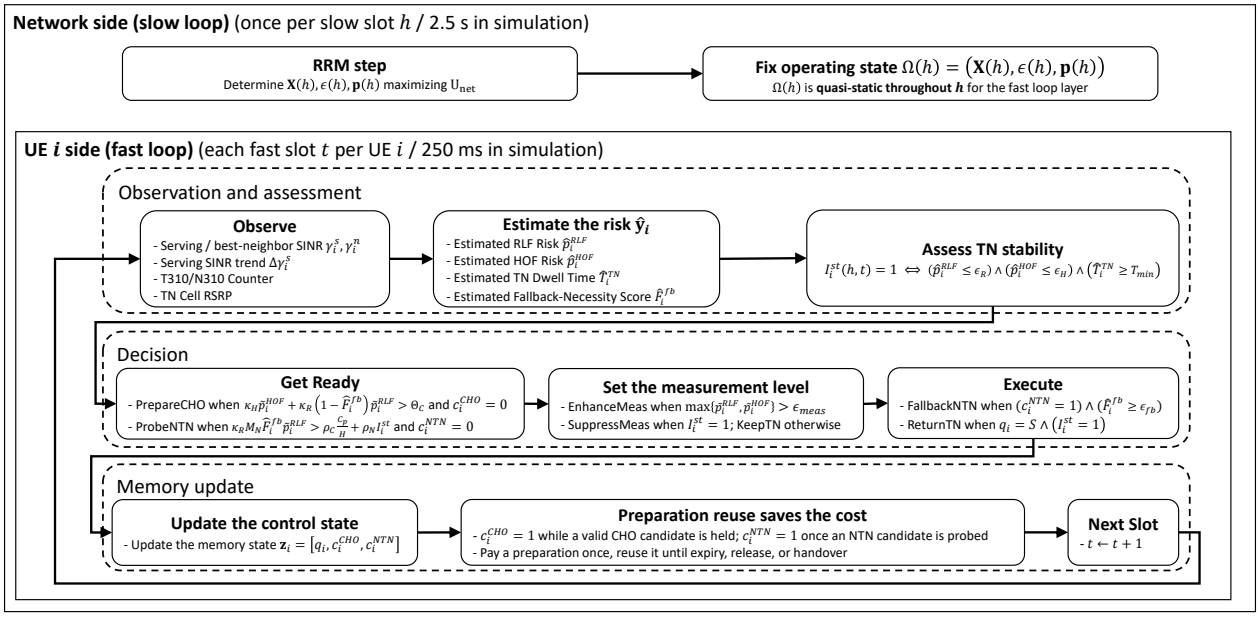


Fig. 1. Two-timescale structure of CALM. The network side re-solves the operating state $\Omega(h)$ (11) once per slow slot and holds it fixed. In each fast slot, every UE observes the radio and RLM/A3 state (12), estimates the risk indicators (14)-(17), assesses TN stability (21), applies the preparation, measurement, and execution rules (31)-(36) with the dwell-anticipated risks (28), and carries the preparation memory (19) into the next slot.

These actions abstract UE-assisted, network-configured mobility control decisions: when to trigger measurement, CHO preparation, NTN probing, fallback, or return-to-TN. The kinds are not exclusive within a slot: CALM decides the CHO and NTN preparations, the measurement level, and an execution component independently, in that order (Algorithm 1), so $a_i(h,t)$ collects the concurrent actions selected from $\mathcal{A}$.

To carry preparation memory across fast slots, the UE mobility control state is

$$\mathbf{z}_i(h,t) = [q_i(h,t), c_i^{\mathrm{CHO}}(h,t), c_i^{\mathrm{NTN}}(h,t)]. \tag{19}$$

where $q_i(h,t) \in \{\mathrm{T}, \mathrm{S}\}$ is the current serving tier, derived from the network-side association [7]; $c_i^{\mathrm{CHO}}(h,t) \in \{0,1\}$ indicates whether a valid network-configured CHO candidate is available, abstracting the standard CHO procedure [5], [12]; and $c_i^{\mathrm{NTN}}(h,t) \in \{0,1\}$ indicates whether an NTN candidate has already been probed or prepared. Unlike c_i^{CHO}, c_i^{NTN} is not a new 3GPP protocol state but an algorithmic memory variable used to avoid repeated NTN probing cost.

This state enables preparation reuse: once CHO or NTN preparation is performed, the corresponding flag remains active until the candidate expires, is released, or is consumed by a handover/fallback action. The abstraction follows CHO/FCHO mechanisms, in which prepared target-cell configurations can be retained or reused to reduce repeated preparation and reconfiguration signaling [10], [11].

F. UE-Side Control Cost and NTN-Use Penalty

We define UE-side control-cost terms that encourage energy- and signaling-efficient mobility control, and an NTN-use penalty that implements the TN-first principle by discouraging NTN probing or residence when TN is estimated to be stable.

Let $N_{\mathrm{m}}, N_{\mathrm{r}}, N_{\mathrm{s}}, N_{\mathrm{p}}, N_{\mathrm{h}}, N_{\mathrm{R}}$ denote the counts of measurement, reporting, scanning, NTN probing, CHO preparation, and RRC events accumulated over all UEs and fast slots under policy π. The UE-side control cost is

$$\begin{aligned} C_{\mathrm{C}}(\pi) =& e_{\mathrm{m}} N_{\mathrm{m}} + e_{\mathrm{r}} N_{\mathrm{r}} + e_{\mathrm{s}} N_{\mathrm{s}} + e_{\mathrm{p}} N_{\mathrm{p}} \\ &+ e_{\mathrm{h}} N_{\mathrm{h}} + e_{\mathrm{R}} N_{\mathrm{R}}. \end{aligned} \tag{20}$$

where $e_{\mathrm{m}}, e_{\mathrm{r}}, e_{\mathrm{s}}, e_{\mathrm{p}}, e_{\mathrm{h}}, e_{\mathrm{R}} \ge 0$ are per-event energy coefficients from the 3GPP TR 38.840 UE power model [8], each taken as the event's *energy* (power $\times$ duration) with one intra-frequency measurement occasion as the reference (e_{m}=1). On the TR 38.840 relative-power scale a measurement occasion (a multi-slot reference-signal window at 50-60 units) and a single-slot report or RRC transmission (250-700 units) carry comparable per-event energies. Therefore, we set e_{r}=e_{R}=2, e_{h}=2.5, e_{s}=1.5, and e_{p}=3 (the inter-frequency/NTN scan and probe also carry a measurement gap and RF retuning), all in units of e_{m}. C_{C} thus measures UE control energy in units of one measurement occasion: at the 250 ms slot, CALM's 39% saving over A3-CHO (Table II) avoids the energy of ≈ 270 measurement occasions per UE per minute. Section IV-D shows the saving is robust to this weighting.

The TN-stability indicator is

$$\begin{aligned} I_i^{\mathrm{st}}(h,t) = \mathbf{1}\big[&\hat{p}_i^{\mathrm{RLF}}(h,t) \le \varepsilon_{\mathrm{R}} \wedge \hat{p}_i^{\mathrm{HOF}}(h,t) \le \varepsilon_{\mathrm{H}} \\ &\wedge \hat{T}_i^{\mathrm{TN}}(h,t) \ge T_{\min}\big]. \end{aligned} \tag{21}$$

i.e., $I_i^{\mathrm{st}} = 1$ when both estimated failure risks are below their thresholds and the estimated TN dwell is sufficiently long.

Let $I_i^{\mathrm{NTN}}(h,t) = 1$ if UE i probes, prepares, or is served by NTN in slot (h,t). The unnecessary NTN-use penalty is

$$C_{\mathrm{N}}(\pi) = \sum_{h,t} \sum_{i \in \mathcal{U}} I_i^{\mathrm{st}}(h,t) I_i^{\mathrm{NTN}}(h,t). \tag{22}$$

It penalizes NTN use when TN is estimated stable, supporting the TN-first policy.

Finally, the interruption cost is

$$C_{\mathrm{I}}(\pi) = \sum_{h,t} \sum_{i \in \mathcal{U}} \tau_i^{\mathrm{int}}(h,t), \tag{23}$$

where $\tau_i^{\mathrm{int}}(h,t)$ is the interruption time from RLF, HOF, fallback, or recovery events; this captures the service disruption that grows if the controller suppresses preparation too aggressively.

G. Generalized Problem Formulation

Let $\bar{R}_i(h)$ denote UE i's average rate over the fast slots in slow slot h:

$$\bar{R}_i(h) = \frac{1}{H_{\mathrm{f}}} \sum_{t=0}^{H_{\mathrm{f}}-1} R_i(h,t). \tag{24}$$

The per-slot network utility is

$$U_{\mathrm{net}}(\Omega(h)) = \sum_{i \in \mathcal{U}} \log(\bar{R}_i(h) + \delta) - \lambda_{\mathrm{T}} E_{\mathrm{TN}}(h), \tag{25}$$

where $\delta > 0$ avoids the log singularity and λ_{T} weights TN energy. The generalized resource- and mobility control problem is

$$\max_{\{\Omega(h)\},\pi} \quad \sum_h U_{\mathrm{net}}(\Omega(h)) - \rho_{\mathrm{C}} C_{\mathrm{C}}(\pi) - \rho_{\mathrm{I}} C_{\mathrm{I}}(\pi) - \rho_{\mathrm{N}} C_{\mathrm{N}}(\pi) \tag{26}$$

$$\begin{aligned} \text{s.t.} \quad & \sum_{j \in \mathcal{B}} x_{ij}(h) = 1, \quad \forall i \in \mathcal{U},\ \forall h, \\ & 0 \le \epsilon(h) \le 1, \quad \forall h, \\ & 0 \le p_j(h) \le p_j^{\max}, \quad \forall j \in \mathcal{B},\ \forall h, \\ & \bar{R}_i(h) \ge R_{\min}, \quad \forall i \in \mathcal{U},\ \forall h, \\ & D_i \le D_{\max}, \quad \forall i \in \mathcal{U}. \end{aligned}$$

Here $\rho_{\mathrm{C}}, \rho_{\mathrm{I}}, \rho_{\mathrm{N}}$ weight UE-side control cost, interruption cost, and unnecessary NTN use. The problem is mixed-integer, non-convex, and sequential because it jointly involves network-side RRM variables and discrete UE-side mobility control actions over time.

III. Designed Solution: CALM

(26) is non-convex, so CALM solves it approximately by separating it into two coupled timescales, following the two-timescale design principle: the fast mobility control layer treats the slow RRM state as quasi-static, while the slow layer sees the fast layer at its short-horizon equilibrium. We do not claim joint global optimality; each layer is solved by a tractable, monotone-improving surrogate.

At the start of slow slot h the network updates the operating state (11); within the slot, CALM holds $\Omega(h)$ fixed and applies the per-UE fast step.

A. Slow Network-side RRM Step

At each slow slot h, CALM computes a network operating point by maximizing the utility defined in Section II:

$$\max_{\Omega(h)} \quad U_{\mathrm{net}}(\Omega(h)). \tag{27}$$

This subproblem updates the association $\mathbf{X}(h)$ (which MBS or beam serves each UE), the bandwidth split $\epsilon(h)$, and the transmit powers $\mathbf{p}(h)$. We implement it by a low-complexity block-coordinate procedure [14]: a TN-first feasible association, then a bandwidth-split and transmit-power update. TN MBS activation is implicit in the power vector ($p_j(h) = 0$ deactivates MBS j via the $\|p_j(h)\|_0$ term). Since only non-decreasing updates of $U_{\mathrm{net}}(\Omega(h))$ are accepted, the slow step performs monotone block-coordinate ascent on a bounded objective and reaches a coordinatewise stationary point under a finite iteration budget [14].

The slow step sets the network operating point on top of which the fast layer acts; it does not decide the individual UE-side actions (e.g., measurement, CHO preparation, or NTN probing), which are decided by the fast step.

B. Fast UE-Side Mobility-Control Step

Within slow slot h, $\Omega(h)$ is fixed. At each fast slot t and for each UE i, CALM builds the observation (12), obtains the risk estimate (14), and selects an action from $\mathcal{A}$.

The fast step is sequential, since a preparation taken now incurs cost immediately yet may be reused in later slots, which CALM captures through the memory state $\mathbf{z}_i(h,t)$. As (26) is intractable to optimize over π directly, CALM replaces it by a per-UE *one-step lookahead* (certainty-equivalent) surrogate [15] that greedily minimizes UE i's expected marginal contribution to (26) over the next H slots. The weights $\rho_{\mathrm{C}}, \rho_{\mathrm{N}}$ carry over directly, whereas ρ_{I} enters through the interruption a preparation averts, namely $\kappa_{\mathrm{H}} = \rho_{\mathrm{I}} \bar{\tau}^{\mathrm{HOF}}$ and $\kappa_{\mathrm{R}} = \rho_{\mathrm{I}} \bar{\tau}^{\mathrm{RLF}}$ with $\bar{\tau}^{\mathrm{HOF}}, \bar{\tau}^{\mathrm{RLF}}$ the expected interruption times of an unmitigated HOF and RLF. The preparation values below are the resulting per-UE marginal surrogates, a heuristic approximation rather than a claim of joint optimality.

1) Dwell-Anticipated Risk: The fast controller acts on the current-slot risk estimates $\hat{p}_i^{\mathrm{RLF}}, \hat{p}_i^{\mathrm{HOF}}$ directly. Because these are read from the radio and RRC state (in particular the standardized RLF-monitoring counters), they already reflect how far a degradation has progressed toward failure, so the controller needs no separate model of the failure process; we set $\tilde{p}_i^{\mathrm{HOF}}(h,t) = \hat{p}_i^{\mathrm{HOF}}(h,t)$. The risk estimator f_θ supplies these indicators from $\mathbf{s}_i(h,t)$.

The only addition over the reactive rule is a dwell-driven anticipation of RLF: even when the instantaneous RLF risk is moderate, a UE with short estimated dwell may soon need fallback or recovery, so it should prepare earlier. Therefore, we define an effective RLF risk as

$$\tilde{p}_i^{\mathrm{RLF}}(h,t) = \max\left\{ \hat{p}_i^{\mathrm{RLF}}(h,t), \sigma_\eta\Big(T_{\min} - \hat{T}_i^{\mathrm{TN}}(h,t) \Big) \right\}, \tag{28}$$

where

$$\sigma_\eta(x) = \frac{1}{1+\exp(-\eta x)} \tag{29}$$

is a logistic urgency with slope $\eta > 0$, which rises as the estimated dwell drops below $T_{\min}$, making the controller more proactive when the UE is about to leave stable TN coverage. The $\max$ in (28) acts as a conservative floor, taking the more pessimistic of the data-driven risk $\hat{p}_i^{\mathrm{RLF}}$ and the dwell-driven urgency $\sigma_\eta(T_{\min} - \hat{T}_i^{\mathrm{TN}})$, thereby preparing the controller earlier exactly when an instantaneous reading would understate the true risk.

2) Preparation Reuse and Preparation Values: The memory state $\mathbf{z}_i(h,t) = [q_i(h,t), c_i^{\mathrm{CHO}}(h,t), c_i^{\mathrm{NTN}}(h,t)]$ lets CALM avoid paying the same preparation cost repeatedly: when $c_i^{\mathrm{CHO}} = 1$ a CHO candidate is already valid, and when $c_i^{\mathrm{NTN}} = 1$ an NTN candidate has already been probed. Therefore, CALM triggers a preparation only on the rising edge of a risk episode, unless the prior candidate has expired, been released, or been consumed by handover or fallback.

Let $c_{\mathrm{h}} = e_{\mathrm{h}} + e_{\mathrm{R}}$ be the CHO preparation cost (including RRC signaling) and $M_{\mathrm{H}} \in (0,1]$ its mitigation effectiveness. CHO covers two TN-recoverable risks: HOF, and the portion of RLF still recoverable on TN. With $\hat{F}_i^{\mathrm{fb}}$ denoting fallback necessity, $1 - \hat{F}_i^{\mathrm{fb}}$ measures the plausibility of terrestrial recovery.

The one-edge surrogate value of preparing a CHO candidate is

$$\Delta J_i^{\mathrm{CHO}}(h,t) = \rho_{\mathrm{C}}\frac{c_{\mathrm{h}}}{H} - M_{\mathrm{H}}\Big(\kappa_{\mathrm{H}}\tilde{p}_i^{\mathrm{HOF}}(h,t) + \kappa_{\mathrm{R}}\big(1-\hat{F}_i^{\mathrm{fb}}(h,t)\big)\tilde{p}_i^{\mathrm{RLF}}(h,t)\Big), \tag{30}$$

where κ_{H} and κ_{R} are the interruption-risk weights associated with HOF and RLF, respectively. CALM prepares a CHO candidate when

$$\Delta J_i^{\mathrm{CHO}}(h,t) < 0 \quad \text{and} \quad c_i^{\mathrm{CHO}}(h,t) = 0. \tag{31}$$

The second condition realizes preparation reuse: the cost is paid only when a new valid candidate must be prepared.

Similarly, let $c_{\mathrm{p}} = e_{\mathrm{p}} + e_{\mathrm{s}}$ be the NTN probing/scanning cost and $M_{\mathrm{N}} \in (0,1]$ its mitigation effectiveness. NTN probing is useful only when fallback is likely, so CALM gates it by the fallback-necessity score:

$$\kappa_{\mathrm{R}} M_{\mathrm{N}} \hat{F}_i^{\mathrm{fb}}(h,t)\tilde{p}_i^{\mathrm{RLF}}(h,t) > \rho_{\mathrm{C}}\frac{c_{\mathrm{p}}}{H} + \rho_{\mathrm{N}} I_i^{\mathrm{st}}(h,t), \tag{32}$$

with the additional condition

$$c_i^{\mathrm{NTN}}(h,t) = 0. \tag{33}$$

The $\rho_{\mathrm{N}} I_i^{\mathrm{st}}$ term penalizes NTN probing when TN is estimated stable; this is what implements the TN-first principle.

3) Measurement Level, Fallback, and Return-to-TN: After the preparation decisions are evaluated, CALM sets the measurement level by a simple threshold rule:

$$a_i^{\mathrm{meas}}(h,t) = \begin{cases} \text{EnhanceMeas}, & \max\{\tilde{p}_i^{\mathrm{RLF}}, \tilde{p}_i^{\mathrm{HOF}}\} > \varepsilon_{\mathrm{meas}}, \\ \text{SuppressMeas}, & I_i^{\mathrm{st}} = 1, \\ \text{KeepTN}, & \text{otherwise}. \end{cases} \tag{34}$$

i.e., measurement is enhanced under high risk, suppressed when TN is estimated stable, and otherwise kept at default.

If an NTN candidate is prepared and fallback necessity is high, CALM executes fallback:

$$a_i(h,t) = \text{FallbackNTN} \quad \text{if} \quad c_i^{\mathrm{NTN}} = 1 \text{ and } \hat{F}_i^{\mathrm{fb}} \geq \varepsilon_{\mathrm{fb}}. \tag{35}$$

Conversely, if a UE on the NTN tier finds TN stable again, CALM returns it to TN:

$$a_i(h,t) = \text{ReturnTN} \quad \text{if} \quad q_i = \mathrm{S} \text{ and } I_i^{\mathrm{st}} = 1. \tag{36}$$

These rules confine NTN usage to cases where fallback is estimated necessary.

C. Threshold Structure of the Fast Policy

The fast preparation rules admit a threshold characterization. We claim no global optimality for the full problem (26); instead, the result shows that the rules are optimal for the per-UE one-step surrogate that CALM actually solves.

Proposition 1 (Threshold-structured fast control): For fixed risk estimate and slow RRM state $\Omega(h)$, the measurement-level, CHO, and NTN decisions of CALM are threshold rules in the risk estimates and the stability indicator, and the CHO and NTN preparation thresholds are the optimal decisions under the one-step surrogate of Section III-B. A CHO candidate is prepared if

$$\kappa_{\mathrm{H}}\tilde{p}_i^{\mathrm{HOF}}(h,t) + \kappa_{\mathrm{R}}\big(1-\hat{F}_i^{\mathrm{fb}}(h,t)\big)\tilde{p}_i^{\mathrm{RLF}}(h,t) > \Theta_{\mathrm{C}}, \tag{37}$$

where

$$\Theta_{\mathrm{C}} = \frac{\rho_{\mathrm{C}} c_{\mathrm{h}}}{H M_{\mathrm{H}}}. \tag{38}$$

An NTN candidate is probed if

$$\tilde{p}_i^{\mathrm{RLF}}(h,t) > \Theta_{\mathrm{N}}, \tag{39}$$

where

$$\Theta_{\mathrm{N}} = \frac{\rho_{\mathrm{C}} c_{\mathrm{p}}/H + \rho_{\mathrm{N}} I_i^{\mathrm{st}}(h,t)}{\kappa_{\mathrm{R}} M_{\mathrm{N}} \hat{F}_i^{\mathrm{fb}}(h,t)}, \tag{40}$$

for $\hat{F}_i^{\mathrm{fb}}(h,t) > 0$. If $\hat{F}_i^{\mathrm{fb}}(h,t) = 0$, NTN probing is never triggered by (32). The measurement level is set by the same indices through the threshold rule (34).

Proof: The CHO rule follows by rearranging $\Delta J_i^{\mathrm{CHO}}(h,t) < 0$ in (30). Dividing both sides by $M_{\mathrm{H}} > 0$ gives (37)-(38). Since $\Delta J_i^{\mathrm{CHO}}$ is affine in the binary prepare/defer decision, the test $\Delta J_i^{\mathrm{CHO}} < 0$ is exactly the $\arg\min$ of the one-step surrogate over that decision, so the CHO threshold is surrogate-optimal; the same argument on the probe value (32) makes the NTN threshold surrogate-optimal. The NTN probing rule follows by rearranging (32). For $\hat{F}_i^{\mathrm{fb}}(h,t) > 0$, division by $\kappa_{\mathrm{R}} M_{\mathrm{N}} \hat{F}_i^{\mathrm{fb}}(h,t)$ gives (39)-(40). If $\hat{F}_i^{\mathrm{fb}}(h,t) = 0$, the left-hand side of (32) is zero, so the probing condition cannot hold. Finally, the effective risks $\tilde{p}_i^{\mathrm{HOF}} = \hat{p}_i^{\mathrm{HOF}}$ and $\tilde{p}_i^{\mathrm{RLF}}$ are monotone in the per-slot risk estimates, and the dwell-urgency term in (28) makes the triggering risk non-increasing in the estimated dwell $\hat{T}_i^{\mathrm{TN}}$, so the threshold form carries over to the underlying risk estimates and dwell. ■

Proposition 1 thus gives the preparation decisions a surrogate-level optimality guarantee and an interpretable threshold form that explains the levers. Because the thresholds are monotone in the per-slot risk estimates, a bounded risk-estimate error shifts the triggering risk only boundedly, so the operating point is robust to estimation error, as our estimation-noise sweep confirms. Raising ρ_{C} lifts both thresholds, making the controller more conservative and cutting signaling; raising H amortizes preparation cost over more slots, prompting earlier preparation; raising ρ_{N} lifts the NTN probing threshold when TN is stable; and as $\hat{F}_i^{\mathrm{fb}}$ decreases, the NTN probing threshold rises, confining NTN use to genuine coverage loss.

Remark (origin of the saving): Of the three rules, the measurement-level threshold dominates the realized UE-cost saving (Table II, column $C_{\mathrm{C}}^{\mathrm{m}}$). It is gated by the stability indicator I_i^{st} rather than by ρ_{C}, so it relaxes routine measurement whenever the terrestrial link reads as stable, and is thus complementary to the CHO and NTN preparation thresholds, which amortize preparation cost and fire only on the rising edge of risk. Consequently, the reactive variant, which keeps the measurement gate but drops dwell anticipation and preparation reuse, already realizes most of the cost saving, whereas the proactive thresholds add earlier preparation as a robustness margin whose reliability gain stays within the confidence interval (Section IV).

D. Special Cases and Complexity

CALM contains the underlying network-side RRM as a special case, since disabling the UE-side mobility control actions and setting $\rho_{\mathrm{C}} = \rho_{\mathrm{I}} = \rho_{\mathrm{N}} = 0$ reduces (26) to the network-utility maximization $\max_{\{\Omega(h)\}} \sum_h U_{\mathrm{net}}(\Omega(h))$, whereas a reactive controller is recovered by setting $H = 1$, removing the dwell-urgency term in (28), and disabling preparation reuse, namely a reactive baseline that does not amortize preparation cost over future slots. With $K = |\mathcal{U}|$ and $L = |\mathcal{B}|$, a greedy slow RRM update costs $O(I_{\mathrm{R}} K L)$ per slow slot for I_{R} iterations, while the fast step evaluates a constant number of closed-form expressions per UE per slot, namely the dwell-anticipated risk, the CHO and NTN thresholds, the measurement level, and the fallback and return-to-TN rules, so its per-slot complexity is $O(K)$ and needs no explicit rollout over the H slots.

E. Proposed Algorithm

Algorithm 1 summarizes CALM as a two-timescale loop. At each slow slot the network-side RRM step sets $\Omega(h)$, attempting terrestrial association first and using the NTN tier only for fallback or load support, after which every UE applies the closed-form threshold rules of Proposition 1 across the fast slots while the C_{C}, C_{I}, and C_{N} counters accumulate as actions and mobility events occur. Each fast decision reduces to an independent threshold comparison.

Algorithm 1 CALM: the two-timescale control loop.

Input: UEs $\mathcal{U}$, nodes $\mathcal{B}$, channel model, traffic/mobility traces, risk map f_θ, horizon H
Output: per-UE actions $a_i(h,t)$ and the cost/failure counters
for each slow RRM slot h **do**
  refresh UE positions, demands, channel gains, loads
  solve the slow RRM step for $\Omega(h) = (\mathbf{X}(h), \epsilon(h), \mathbf{p}(h))$ // network side; the fast layer holds it fixed
  for each fast slot t in h and each UE i **do**
    observe $\mathbf{s}_i$, then $\hat{\mathbf{y}}_i \leftarrow f_\theta(\mathbf{s}_i)$ // current-slot indicators, no rollout
    set $\tilde{p}_i^{\mathrm{HOF}} = \hat{p}_i^{\mathrm{HOF}}$ and $\tilde{p}_i^{\mathrm{RLF}}$ by (28) // current-slot risk + dwell urgency
    if the CHO preparation test (31) fires **then**
      prepare a CHO candidate // pay once on the rising edge, then reuse
    end if
    if the probe test (32) fires and no NTN candidate is held **then**
      probe NTN // only when fallback is truly needed (TN-first)
    end if
    set the measurement level by (34) // suppress while the link reads stable; the main saving
    apply fallback / return-to-TN by (35)-(36), then update $\mathbf{z}_i$ and the counters
  end for
  log $U_{\mathrm{net}}, C_{\mathrm{C}}, C_{\mathrm{I}}, C_{\mathrm{N}}$, failures, delay, NTN residence
end for

IV. Simulation Results

A. Simulation Setup

We evaluate CALM on a traffic-aware network-side system model [7], namely a two-tier network with $M = 7$ terrestrial MBSs and a LEO satellite tier, S-band at 2 GHz, $W = 40$ MHz with orthogonal split (no inter-tier interference). TN links follow 3GPP TR 38.901 [13] urban-macro (UMa) with explicit LoS/NLoS branches under Table 7.4.2-1 LoS probabilities, LoS-conditional shadow fading (σ_{LoS}=4, σ_{NLoS}=6 dB), and high-loss/low-loss O2I penetration for the 80% indoor UEs; the satellite link follows TR 38.811 [1] (elevation-dependent free-space loss, clutter, scintillation, and entry losses). Unlike the three static snapshots of [7], satellite elevation is swept *continuously* over a LEO pass at altitude 600 km. Crucially, handovers are *executed* on the fast timescale. A UE switches to a stronger neighbour once the A3 condition (offset 3 dB, hysteresis 1 dB) has held for the time-to-trigger, or immediately when its serving link approaches out-of-sync. Therefore, the serving cell tracks the UE rather than freezing at the slow RRM association. Ground-truth RLF events follow the 3GPP T310/N310/N311 radio-link-monitoring procedure of TS 38.331 [6], with out-of-sync/in-sync SINR thresholds set to Q_{out}=-8 dB and Q_{in}=-6 dB and counters N_{310}=6, T_{310}=10 slots, N_{311}=1, and HOF is a too-late handover (A3 trigger while already out-of-sync). The two-timescale loop uses 200 slow slots of 10 fast slots each. With a 250 ms fast slot, a UE at 25 m/s moves about 6 m per slot, so T_{310}=2.5 s, and the resulting failure rates fall in the realistic 1-3$\%$ range. There are $K = 60$ mobile UEs. The short-term risk indicators (14) are computed from the per-slot radio and control-plane observations: the serving SINR and its trend, the neighbour-cell A3 margins, and the T310/N310 counters. We perturb the risk indicators with increasing logit-domain Gaussian noise, from an oracle estimate to a heavily corrupted one: the failure probability stays within its 50-seed confidence interval (a shift below 0.02 percentage points) while the UE control cost rises by at most 12%, so reliability is essentially invariant to estimation quality and the cost degrades only gracefully. The

TABLE I
SIMULATION PARAMETERS (UMA DEFAULTS).

Parameter	Value	Parameter	Value
Carrier	S-band, 2 GHz	Bandwidth W	40 MHz
TN deployment	7 MBSs, UMa	Satellite	LEO, 600 km
UEs K	60 (80% indoor)	UE speed	25 m/s
$\sigma_{\mathrm{LoS}}/\sigma_{\mathrm{NLoS}}$	4/6 dB	Fast slot T_{f}	250 ms
Fast slots/slow H_{f}	10	Slow slots	200
A3 offset/hyst.	3/1 dB	Time-to-trigger	1 slot
$Q_{\mathrm{out}}/Q_{\mathrm{in}}$	$-8/-6$ dB	N_{310}, T_{310}, N_{311}	6, 10 slots, 1
Reuse window H	3	TN/NTN delay	5/30 ms
Deadline $D_{\max}$	20 ms	Cost weight ρ_{C}	1

proactive variant prepares over a reuse window $H = 3$. NTN trades coverage for latency: end-to-end TN and NTN service delays are 5 and 30 ms against a deadline $D_{\max} = 20$ ms. Each operating point is averaged over 50 seeds with 95% confidence intervals. Table I summarizes the key parameters. The simulator and the CALM implementation will be released publicly.

CALM (the proactive threshold controller of Proposition 1) is compared against RRM-only (the underlying network-side RRM [7], with no UE-side mobility control; the special case of Section III-D), Always-measure (continuous measurement with always-on CHO and NTN probing), CHO-all (CHO preparation for every UE every slot; an idealized upper bound on always-prepared TN handover), A3-CHO, prediction-based CHO (PCHO), Predictor-only, and CALM (reactive) (the $H = 1$ action-scoring controller of Section III-D). The A3-CHO baseline is the realistic counterpart to CHO-all, triggering CHO only on an A3 measurement event and reusing the configuration in the FCHO style [10], [11], without NTN. A3-CHO + 38.840 relax adds the standardized RRM measurement relaxation of TS 38.304 §5.2.4.9 [16], motivated by the TR 38.840 UE power model [8]: it suppresses measurement on the RSRP-based low-mobility *and* not-cell-edge criteria, the conservative 3GPP-recommended configuration, which isolates CALM's risk/dwell-driven gate from the standard relaxation rule; its s-SearchThresholdP is swept and the cheapest point at equal-or-better reliability is reported. The Predictor-only baseline composes a GenTrap-class predictor [9] with a fixed threshold-reactive rule, which isolates CALM's cost-aware, reuse-based, TN-first control layer from the predictor itself. Prediction-based CHO (PCHO) [17] is a reliability-oriented proactive baseline that measures continuously and prepares a candidate once the extrapolated A3 margin is predicted to trigger within a swept horizon, isolating CALM's cost-awareness from proactive preparation alone. Reported metrics are the UE control-cost proxy C_{C} with its measurement/probing/handover breakdown, the failure probability P_{fail} (a radio link or handover failure), mean service delay, and the $D_{\max}$ violation ratio.

B. Cost-Reliability Trade-off

With handovers executing normally, reliability no longer separates the schemes. Every scheme that prepares at all, namely A3-CHO, CHO-all, Predictor-only, and both CALM variants, reaches 1.2-1.25% P_{fail}, a practically negligible spread of at most 0.04 percentage points. Only RRM-only, which never prepares, fails higher, at 1.48% (Table II). What separates the schemes is the UE control cost they pay to reach that reliability (Fig. 2, UMa group). At $\rho_{\mathrm{C}} = 1$, the reactive variant reaches the common reliability at cost 1.76, against 2.87 for A3-CHO, a 39% signaling reduction, and against 3.40, 7.10, and 14.4 for Predictor-only, CHO-all, and Always-measure. These cost gaps are large and statistically significant (paired t-test, $p < 10^{-3}$), whereas the reliability spread is statistically negligible: a two one-sided test (TOST) confirms the reactive variant and A3-CHO are equivalent in P_{fail} within a 0.1 percentage-point margin ($p < 10^{-15}$). Therefore, reliability saturates, and the UE control cost is what separates the schemes. The proactive variant attains a slightly lower mean failure (1.22%) at a higher cost (2.24), a reliability edge that stays within the 95% confidence interval. Sweeping ρ_{C} traces a continuum between the two CALM points, each of which reaches the common reliability below every baseline. This efficiency is not entirely free, but its price is small: occasional NTN use raises the mean delay from 5.0 to 5.3-5.4 ms and the $D_{\max}$ violation ratio to 1.0-1.4%, both well within the 20 ms deadline. The result holds under stress (Fig. 2, rural-macro (RMa) group). In a sparse RMa deployment with large cells (inter-site distance 7 km), $K = 80$ UEs, $\sigma_{\mathrm{NLoS}} = 8$ dB, and vehicular UEs at 35 m/s, failures rise to 2.3-3.2%, and CALM dominates the realistic baselines on both axes. The reactive variant undercuts A3-CHO on cost (2.08 against 2.90) at a lower failure (2.38% against 2.49%), while the proactive variant matches the best reliability (2.33%) at less than a fifth of the Always-measure cost. The saving is not tied to one deployment: sweeping the cell area (inter-site distance) over 5-9 km and the UE density over 60-110 users at equal-or-better reliability, the reactive variant stays cheaper than A3-CHO throughout, by 15-38%.

TABLE II
OPERATING POINTS UNDER THE REALISTIC 3GPP-STANDARD SCENARIO WITH EXECUTED HANDOVERS (38.901 UMA + T310 RLF + A3 HOF; MEAN OVER 50 SEEDS; CALM AT $\rho_{\mathrm{C}} = 1$). $C_{\mathrm{C}}^{\mathrm{m}}$ IS THE MEASUREMENT COMPONENT OF THE CONTROL COST.

Scheme	C_{C} (proxy)	$C_{\mathrm{C}}^{\mathrm{m}}$ (meas.)	P_{fail} prob.	delay (ms)	$D_{\max}$ viol.
RRM-only	0.52	0.52	0.0148	5.00	0.000
Always-measure	14.40	5.20	0.0121	6.32	0.053
CHO-all (upper bound)	7.10	2.60	0.0123	5.00	0.000
A3-CHO	2.87	2.60	0.0125	5.00	0.000
38.840 relax (both)	2.64	2.36	0.0125	5.00	0.000
Predictor-only	3.40	2.88	0.0124	5.10	0.004
PCHO	3.36	2.60	0.0124	5.00	0.000
CALM (reactive)	1.76	1.15	0.0123	5.26	0.010
CALM (proactive)	2.24	1.15	0.0122	5.36	0.014

C. Where the Savings Come From

The measurement component $C_{\mathrm{C}}^{\mathrm{m}}$ of the control cost in Table II shows that the saving is almost entirely in measure-

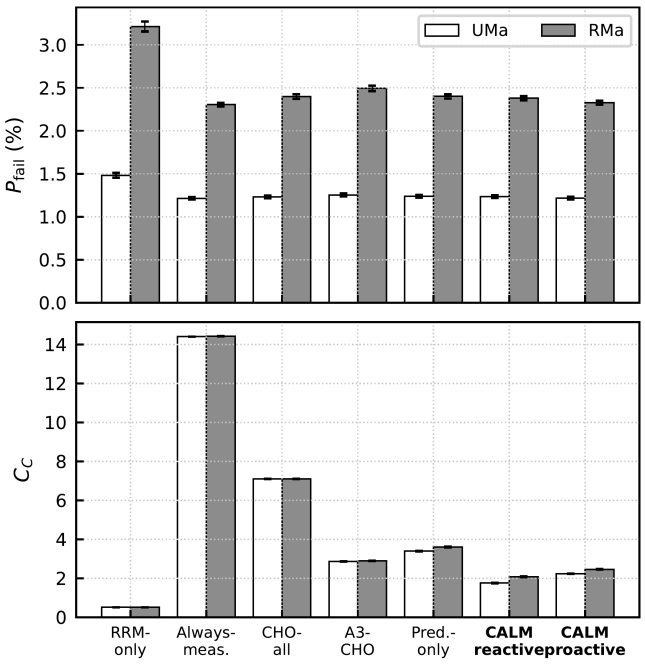


Fig. 2. Per-scheme P_{fail} (top) and UE control cost C_{C} (bottom), benign UMa and stressed RMa grouped per scheme (mean $\pm$ 95% CI, 50 seeds). CALM variants in bold.

ment. CALM measures at cost 1.15, against 2.60 for A3-CHO and 2.88 for Predictor-only, roughly a 55% cut, because the stability indicator I_i^{st} (21) drops measurement to a minimal level whenever the link is estimated to be stable, which is most of the time. A3-CHO already keeps its handover signaling small but measures continuously, and Always-measure marks the upper bound, spending 5.2 on measurement and 4.5 on NTN probing.

This saving is not what the 3GPP standard relaxation already buys. Augmenting A3-CHO with the standardized TS 38.304 relaxation on its conservative low-mobility *and* not-cell-edge criteria [16] lowers the measurement cost only to 2.36 (control cost 2.64), because under realistic UE dynamics the low-mobility test rarely holds long enough to suppress measurement; CALM (reactive) still undercuts this standardized baseline by 33% in UMa (1.76 against 2.64) and by 21% in RMa (2.08 against 2.64) at equal-or-better reliability. Relaxing more aggressively on the not-cell-edge criterion alone can fall below CALM on cost, but only by suppressing measurement away from the cell edge, which leaves P_{fail} at the no-relaxation level, well above what CALM attains. CALM is the only scheme that improves on A3-CHO on *both* axes, because its gate I_i^{st} keys on a risk and dwell estimate, not on serving RSRP alone. Fig. 3 makes this explicit: the standardized relaxation is cheaper than CALM but its P_{fail} plateaus across the entire s-SearchThresholdP sweep and never approaches CALM's, so it does not dominate CALM; CALM is the only scheme that drives P_{fail} below this floor while still undercutting every continuously-measuring baseline on cost.

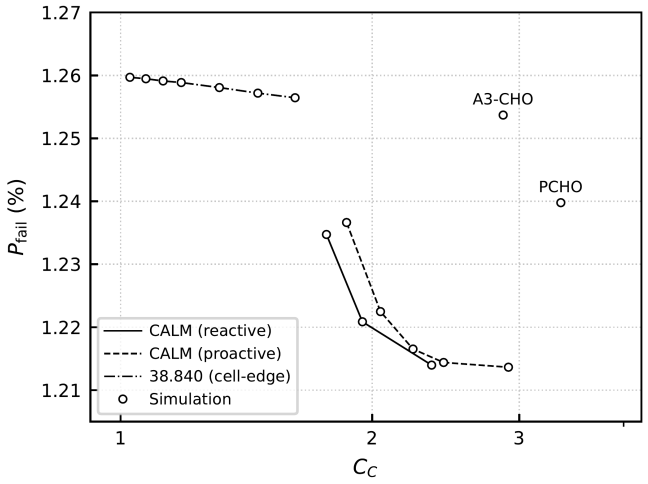


Fig. 3. Cost-reliability frontier (UMa, 50 seeds; CALM swept over ρ_{C}). "38.840 (cell-edge)" is the aggressive cell-edge-only TS 38.304 relaxation swept over s-SearchThresholdP, distinct from the both-criteria row of Table II.

The same gap separates CALM from a reliability-oriented proactive scheme: sweeping PCHO's prediction horizon from 2 to 8 slots lowers the UMa failure only from 1.256% to 1.235% while the cost climbs from 3.05 to 3.77, since longer lead time prepares more candidates and measurement is never relaxed. CALM dominates this whole front on both axes: proactive reaches 1.22% at 2.24, below the cheapest PCHO point, and reactive matches PCHO's best reliability at 1.76, under half its cost (same ordering under RMa). Thus, proactive preparation alone is not enough; the decisive difference is that CALM relaxes measurement to 1.15 against PCHO's 2.60.

The cost advantage is not an artifact of a single operating point. Sweeping the UE speed from 10 to 50 m/s, both CALM variants keep the lowest control cost ($\sim$2.2, against $\sim$2.9 for A3-CHO and $\sim$3.4 for Predictor-only) over the entire range. Even at 80 m/s and a 100 m inter-site distance, where the SINR can fall faster than the prepare-to-complete pipeline, the proactive-minus-reactive P_{fail} gap stays within its 95% confidence interval (at most 0.06 percentage points) while the proactive cost keeps growing. Therefore, we present proactive preparation as a robustness knob, not a reliability headline; the headline is the reactive variant's 39% signaling reduction at matched reliability.

D. Cost-Weight Sensitivity

Because the UE control cost (20) is a weighted sum of event counts, we test whether the saving is an artifact of the chosen weights. Fig. 4 reruns the realistic RMa scenario under three per-event energy settings: the TR 38.840-anchored default, a measurement-dominant setting, and a signaling-heavy setting. Under the physical TR 38.840 weights, CALM (reactive) reaches the common reliability at cost 2.08 against 2.90 for A3-CHO, a 28% reduction, and it stays cheaper at equal-or-better reliability under the measurement-dominant and signaling-heavy weightings as well. The saving naturally

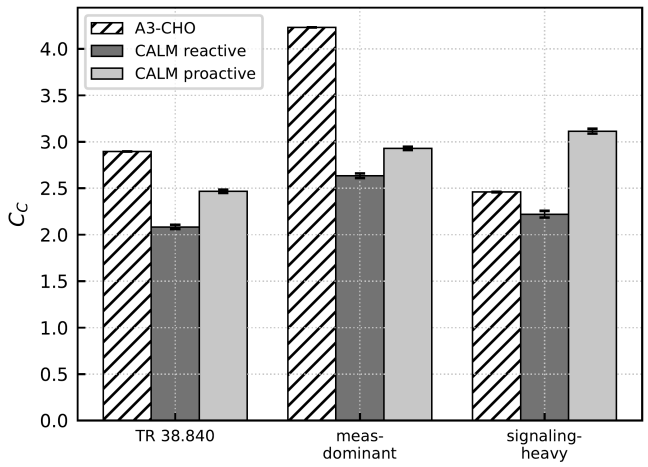


Fig. 4. UE control cost under three per-event energy weightings (realistic RMa, mean $\pm$ 95% CI).

requires measurement to carry non-trivial energy and would vanish only if measurement were made nearly free, a regime the NTN and inter-frequency measurement gap and RF retuning rule out, which is precisely why 3GPP introduced RRM measurement relaxation [8]. The proactive variant, which prepares earlier, is the low-failure operating point and forfeits its cost advantage under the signaling-heavy weighting, where its additional preparation is penalized.

V. Conclusion

We extended network-side TN-NTN RRM with a model of UE-side mobility control cost, and we proposed CALM, a two-timescale framework. Its fast layer reduces to closed-form, preparation-reuse thresholds (Proposition 1) that relax UE signaling while the terrestrial link reads as stable. Disabling the UE-side costs recovers the underlying network-side RRM, and setting $H = 1$ recovers a reactive controller. On a 3GPP-standard TN-NTN simulator with executed handovers, reliability saturates across every preparing scheme in a realistic 1-3$\%$ band, so the decisive difference becomes the UE control cost. CALM matches the realistic event-triggered CHO baseline at about 39$\%$ lower signaling, consistently across UE speed, cell density, and a sparse rural-macro deployment. We scope the claim to this efficiency result on purpose, since the room for reliability or NTN-offload gains is small under realistic conditions. Future work targets environments where the trade-off is sharper: aerial UEs, dense urban blockage, and large-scale signaling-bottleneck regimes.

References

[1] 3GPP, "Study on new radio (nr) to support non-terrestrial networks," 3rd Generation Partnership Project (3GPP), Tech. Rep. TR 38.811, V15.4.0, 2020.

[2] ——, "Solutions for NR to support non-terrestrial networks (NTN)," 3rd Generation Partnership Project (3GPP), Tech. Rep. TR 38.821, V16.1.0, 2021.

[3] ——, "Study on using satellite access in 5G," 3rd Generation Partnership Project (3GPP), Tech. Rep. TR 22.822, V16.0.0, 2018.

[4] O. Kodheli, E. Lagunas, N. Maturo, S. K. Sharma, B. Shankar, J. F. M. Montoya, J. C. M. Duncan, D. Spano, S. Chatzinotas, S. Kisseleff *et al.*, "Satellite communications in the new space era: A survey and future challenges," *IEEE Communications Surveys & Tutorials*, vol. 23, no. 1, pp. 70–109, 2021.

[5] 3GPP, "NR; NR and NG-RAN overall description; stage-2," 3rd Generation Partnership Project (3GPP), Tech. Rep. TS 38.300, V17.5.0, 2023.

[6] ——, "NR; radio resource control (RRC); protocol specification," 3rd Generation Partnership Project (3GPP), Tech. Rep. TS 38.331, V17.6.0, 2023.

[7] H. Alam, A. De Domenico, D. López-Pérez, and F. Kaltenberger, "Optimizing integrated terrestrial and non-terrestrial networks performance with traffic-aware resource management," *arXiv preprint arXiv:2410.06700*, 2024.

[8] 3GPP, "Study on user equipment (UE) power saving in NR," 3rd Generation Partnership Project (3GPP), Tech. Rep. TR 38.840, V16.0.0, 2019.

[9] K. Hasan, K. Papry, T. Trappenberg, and I. Haque, "A generalized GNN-transformer-based radio link failure prediction framework in 5G RAN," *IEEE Transactions on Machine Learning in Communications and Networking*, vol. 3, pp. 710–724, 2025.

[10] S. B. Iqbal, A. Awada, U. Karabulut, I. Viering, P. Schulz, and G. P. Fettweis, "On the modeling and analysis of fast conditional handover for 5G-Advanced," in *2022 IEEE 33rd Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*. IEEE, 2022, pp. 595–601.

[11] S. B. Iqbal, S. Nadaf, A. Awada, U. Karabulut, P. Schulz, and G. P. Fettweis, "On the analysis and optimization of fast conditional handover with hand blockage for mobility," *IEEE Access*, vol. 11, pp. 30 040–30 056, 2023.

[12] J. Stanczak, U. Karabulut, and A. Awada, "Conditional handover in 5G - principles, future use cases and FR2 performance," in *2022 International Wireless Communications and Mobile Computing (IWCMC)*. IEEE, 2022, pp. 660–665.

[13] 3GPP, "Study on channel model for frequencies from 0.5 to 100 GHz," 3rd Generation Partnership Project (3GPP), Tech. Rep. TR 38.901, V17.0.0, 2022.

[14] P. Tseng, "Convergence of a block coordinate descent method for nondifferentiable minimization," *Journal of Optimization Theory and Applications*, vol. 109, no. 3, pp. 475–494, 2001.

[15] D. Bertsekas, *Reinforcement Learning and Optimal Control*. Athena Scientific, 2019.

[16] 3GPP, "NR; user equipment (UE) procedures in idle mode and RRC inactive state," 3rd Generation Partnership Project (3GPP), Tech. Rep. TS 38.304, V18.2.0, 2024.

[17] C. Lee, H.-J. Cho, S. Song, and J.-M. Chung, "Prediction-based conditional handover for 5G mm-Wave networks: A deep-learning approach," *IEEE Vehicular Technology Magazine*, vol. 15, no. 1, pp. 54–62, 2020.