Link Adaptation for NB-IoT over GEO NTN: A Joint Throughput and Resource Optimization

Oussama Lachgar
Kinéis
Univ. de Toulouse, CNRS, Toulouse INP, IRIT
Toulouse, France
olachgar@kineis.com, oussama.lachgar@irit.fr

Vincent Deslandes
Kinéis
Toulouse, France
vdeslandes@kineis.com

Rahim Kacimi
Univ. de Toulouse, CNRS, Toulouse INP, IRIT
Toulouse, France
rahim.kacimi@irit.fr

Antonin Le Floch
Kinéis
Toulouse, France
alefloch@kineis.com

André-Luc Beylot
Univ. de Toulouse, CNRS, Toulouse INP, IRIT
Toulouse, France
andre-luc.beylot@irit.fr

***Abstract*—Link adaptation in NB-IoT has been extensively studied for terrestrial networks, where the block-error-rate (BLER) target acts as a constraint under which throughput, resources, energy, or latency are optimized. Extending NB-IoT to non-terrestrial networks (NTN) reshapes this problem in ways the terrestrial literature does not address. In particular, the long GEO round-trip time of 540 ms turns every packet segmentation into a latency penalty; moreover it leaves the scheduler with outdated CSI (channel-state information). While satellite radio resources are scarce: the goal is no longer to maximize throughput alone, but to spend as few resources as possible while meeting the reliability target under delayed CSI.**

To the best of our knowledge, link adaptation for NB-IoT over GEO NTN has not been studied. This paper formulates single-tone NB-IoT uplink link adaptation over a GEO satellite as a joint throughput–resource optimization subject to a BLER target, where the transmission format is selected from the channel-state information (CSI) available at the base station. We propose two low-complexity heuristics that solve the optimization problem. We compare them to a conservative baseline and packet-level optimal bounds we derived using dynamic programming. Our evaluation is driven by a land-mobile-satellite channel model calibrated on real uplink SNR measurements from controlled GEO NB-IoT system. The results show that our heuristics succeed to improve the link performance by adapting to the studied channel. The best heuristic (H2) approaches to the reliability target (0.1) at a fraction of the resource cost of the conservative baseline, cutting resource use by 66% and raising throughput by 157%. Furthermore, it reaches 62% of the clairvoyant optimal throughput, however at about twice the resource cost, quantifying the headroom that the GEO round-trip delay leaves on the table.

***Index Terms*—NTN, NB-IoT, GEO Satellite, Performance, Link Adaptation, Optimisation.**

I. Introduction

Starting from Release 17, 3rd Generation Partnership Project (3GPP) has extended NB-IoT to non-terrestrial networks (NB-NTN), enabling global coverage of remote assets through LEO and GEO satellites [1]. GEO operation is particularly suited for delay-tolerant applications thanks to its fixed earth-relative geometry and single-satellite coverage, but it imposes a round-trip time (RTT) of roughly 540 ms. Such a delay does not make retransmissions impossible, but it makes them expensive: any link adjustment that proceeds by sacrificing packets and reacting to their outcome incurs the full round-trip penalty each time, which is impractical given the cost of retransmission in NTN.

In this regime, the choice of Narrowband Physical Uplink Shared Channel (NPUSCH) transmission format, namely the modulation and coding scheme (MCS), the number of resource units (RU), and the number of repetitions, are the primary lever for trading reliability against the radio resources consumed. In terrestrial networks, this link adaptation (LA) problem for NB-IoT is well studied. Because repetitions are central to NB-IoT coverage, LA must be performed in at least two dimensions: schemes that jointly select the MCS and the repetition number from ACK/NACK feedback outperform repetition-only or straightforward methods [2], and adapting the RU count jointly with the MCS has been exploited for rate–latency optimization [3]. These schemes, however, were designed for the terrestrial channel, where the RTT is negligible and feedback is timely. They rely on trying formats and reacting to failures, which is affordable terrestrially but costly over a GEO link, where each retransmission costs a full round trip.

The non-terrestrial settings reshapes the problem in two ways that the terrestrial LA literature does not address. First, the cost of the format selection: segmenting a packet (which a lower MCS makes more likely) incurs an additional round-trip penalty under the large GEO RTT, while raising the repetition count is costly in itself, all the more so as satellite radio resources are scarce. The objective is therefore no longer to maximize throughput or minimize resources in isolation, as is common terrestrially, but to deliver the data using as few resource units and repetitions as possible. Second, the long GEO round-trip delay means that the channel-state information available to the scheduler may already be outdated when the transmission takes place, especially in a highly dynamic channel, which directly undermines the CSI-driven format

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

selection. On the non-terrestrial side, the adaptations of the NB-IoT air interface required for NTN have been surveyed at a general level [1], and has been characterized at the link-budget and system-performance level [4], [5]; however, the selection of the transmission format (the link-adaptation problem) has not been considered as far. Resource allocation for satellite NB-IoT has been studied for LEO constellations [6], where the effort concentrates on the LEO-specific impairments (mainly the differential Doppler and the limited, time-varying coverage) for the multi-user allocation problem. Efficient scheduling in GEO based LTE systems was treated with focus on Buffer status prediction under long propagation time constraint [7]. A subgroup-based LA approach was proposed to cope with channel diversity in multi-cast 5G-satellite networks [8]. We instead focus on the single-user link optimization in the NB-IoT GEO case, where the challenge shifts to selecting the per-link transmission format, which we cast as a joint throughput–resource optimization. On the other hand, adaptive coding and modulation (ACM) for satellite links is mature and explicitly contends with outdated CSI, through SNR-threshold MODCOD selection [9] or through CSI prediction [10]. In broadband standards such as DVB-S2, however, format selection reduces to a one-dimensional MODCOD choice that a simple look-up table resolves, whereas NB-IoT demands the joint selection of MCS, resource units, and repetitions, a multi-dimensional decision. To the best of our knowledge, link adaptation for NB-IoT over GEO NTN has not been studied, nor framed as a joint throughput–resource optimization under outdated CSI. This paper addresses that gap under a channel model calibrated on real uplink SNR measurements. Our contributions are as follows:

- We formulate and study, for the first time, the NB-IoT uplink link-adaptation problem in the GEO NTN context, where the format $(I_{\mathrm{MCS}}, N_{\mathrm{RU}}, N_{\mathrm{Rep}})$ is chosen from the CSI available at the base station to jointly maximize delivered throughput and minimize the consumed radio resources $\sum N_{\mathrm{RU}} N_{\mathrm{Rep}}$ under a 0.1 block-error-rate.
- We propose a conservative baseline and two lightweight selection heuristics that the base station can run online, and we derive packet-level optimal bounds by dynamic programming to delimit the achievable throughput–resource trade-off.
- Using an LMS channel model calibrated with real SNR measurements from a GEO NB-IoT system in a rural pedestrian scenario, we benchmark the proposed heuristics against an optimal policy and a conservative baseline. We quantify both the gap to the optimum and the gain over the conservative approach.
- We quantify the cost of outdated CSI by comparing the realistic optimum, which decides on the delayed channel sample, against a clairvoyant optimum with perfect CSI, isolating the penalty imposed by delayed CSI.

The remainder of the paper is organized as follows. Section II presents the system model and problem formulation, Section III proposed link adaptation policies, Section IV simulations and results and Section V concludes.

II. System Model and Problem Formulation

We consider uplink data transmission from a single NB-IoT user equipment (UE) to a base station whose radio access network is carried over transparent GEO satellite link. Throughout this work we restrict the study to the *3.75 kHz single-tone* configuration of the NB-IoT uplink data channel (NPUSCH, format 1). This choice is dictated by the link budget: over the long GEO slant range, single-tone transmission concentrates the entire UE power on one $3.75\,\mathrm{kHz}$ subcarrier, maximizing the received power spectral density and thus the link budget. Our channel measurements are based on 3.75 kHz single tone, which is the only configuration that closes the link in our GEO system. Hence, we adopt it as the operating regime of the study; the multi-tone case is left for future work.

A. Link-Adaptation Parameters

For each transmission, the base station configures three link-adaptation parameters, carried to the UE in the uplink grant (DCI format N0). We refer to their joint setting, namely the modulation and coding scheme, the number of resource units, and the number of repetitions, as the *transmission format*, denoted $\phi = (I_{\mathrm{MCS}}, I_{\mathrm{RU}}, N_{\mathrm{Rep}})$. Choosing ϕ for each transmission is the link-adaptation problem studied in this paper.

1) Modulation and coding scheme (MCS): In single-tone operation the MCS index is restricted to $I_{\mathrm{MCS}} \in \mathcal{M} = \{0, 1, \ldots, 10\}$. The index selects both the modulation ($\pi/2$-BPSK or $\pi/4$-QPSK) and the effective coding rate [11], [12]. A lower MCS is more robust (lower required SNR) at the cost of fewer bits per resource unit.

2) Resource unit (RU): The resource unit is the basic time–frequency allocation unit of the NPUSCH. For the 3.75 kHz single-tone configuration, one RU spans a single subcarrier over 16 slots, i.e., a duration $t_{\mathrm{RU}} = 32\,\mathrm{ms}$. The number of allocated resource units N_{RU} is selected through the index I_{RU} from the standardized set $\mathcal{N}_{\mathrm{RU}} = \{1, 2, 3, 4, 5, 6, 8, 10\}$.

3) Transport block size (TBS): The transport block size is not chosen freely: the pair $(I_{\mathrm{MCS}}, I_{\mathrm{RU}})$ maps to a TBS through the standardized lookup of [12],

$$\mathrm{TBS} = f(I_{\mathrm{MCS}}, I_{\mathrm{RU}})\,, \tag{1}$$

reproduced in Table I for the indices accessible in single-tone operation. The largest reachable value is $\mathrm{TBS}_{\max} = 1000$ bits.

4) Repetitions: The transport block can be repeated to improve reliability through soft combining at the receiver. The repetition number N_{Rep} is selected from $\mathcal{N}_{\mathrm{Rep}} = \{1, 2, 4, \ldots, 128\}$ via the index I_{Rep} [12]. More repetitions lower the required SNR but consume proportionally more radio resources.

The decision vector is thus the triple $(I_{\mathrm{MCS}}, I_{\mathrm{RU}}, N_{\mathrm{Rep}}) \in \mathcal{M} \times \mathcal{N}_{\mathrm{RU}} \times \mathcal{N}_{\mathrm{Rep}}$.

TABLE I
TRANSPORT BLOCK SIZE $\mathrm{TBS} = f(I_{\mathrm{MCS}}, I_{\mathrm{RU}})$, SINGLE-TONE OPERATION (3GPP TS 36.213, TABLE 16.5.1.2-2).

I_{MCS}	I_{RU} 0	1	2	3	4	5	6	7
0	16	32	56	88	120	152	208	256
1	24	56	88	144	176	208	256	344
2	32	72	144	176	208	256	328	424
3	40	104	176	208	256	328	440	568
4	56	120	208	256	328	408	552	680
5	72	144	224	328	424	504	680	872
6	88	176	256	392	504	600	808	1000
7	104	224	328	472	584	712	1000	–
8	120	256	392	536	680	808	–	–
9	136	296	456	616	776	936	–	–
10	144	328	504	680	872	1000	–	–

B. Channel and Reliability Model

1) Channel: We characterize the service link from uplink SNR measurements associated with NPUSCH transmissions in a real GEO NB-IoT deployment using the 3.75 kHz single-tone configuration studied here. In the considered transparent GEO architecture, the satellite acts as a relay while the eNB remains on the ground. The measurements were gathered with a body-worn terminal carried by a runner in an open-sky environment. The dominant short-term variability is attributed to body shadowing and terminal orientation: depending on the wearer's orientation relative to the satellite, the body intermittently obstructs the antenna, so the link alternates between high-SNR and deeply shadowed conditions. Statistically, the link operates around a mean of 0.3 dB with a standard deviation of 4.6 dB, and is interrupted by deep, abrupt fades down to -10 dB, with roughly 18% of the samples below -5 dB; the swing between unshadowed and shadowed conditions reaches about 11 dB. This variability makes link adaptation challenging: a format chosen for the mean condition fails during a fade, while one provisioned for the worst case wastes resources during the long unshadowed intervals.

a) From measurements to a generative channel model: The measured records are realistic but fragmented due to periodic traffic and device multiplicity. We therefore use them to calibrate a standard land-mobile-satellite (LMS) channel model able to generate long, continuous SNR sequences. We adopt the three-state Markov model of Fontán et al. [13], also used in 3GPP TR 38.811 [14]. The three states, identified from the measured SNR distribution in Fig. 1, correspond to high-SNR, moderate-shadowing, and deep-shadowing conditions. In our body-worn scenario, the high-SNR state represents an unshadowed terminal orientation rather than a change in satellite visibility. The states are centered at (+4.8), (-0.4), and (-6.7),dB and are occupied 37%, 43%, and 20% of the time, respectively. Their transition probabilities, estimated from continuous channel segments, are reported in Table II. The transition matrix is calibrated at the scheduler observation timescale, so the generated SNR sequence is sampled once per effective GEO round trip, $T_{\mathrm{CSI}} \simeq 0.54$ s. Hence, each scheduling decision relies on CSI aged by approximately one RTT. The calibrated model reproduces the measured SNR distribution and lag-one autocorrelation of (0.77), while providing the long traces required for the evaluation (Fig. 2).

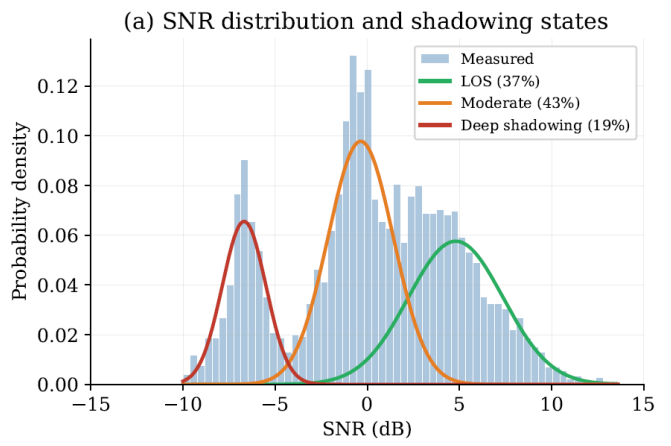


Fig. 1. Measured uplink SNR distribution decomposed into the three LMS states. The high-SNR LOS, moderate, and deep-shadowing states correspond to different body orientations of the worn terminal.

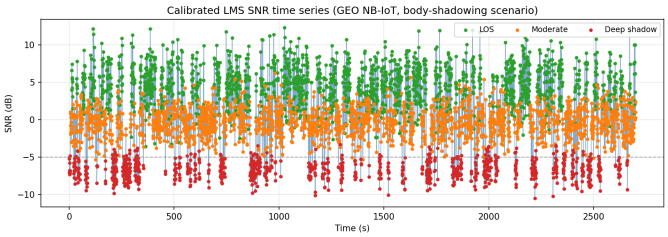


Fig. 2. SNR time series generated by the calibrated LMS model (GEO NB-IoT, 3.75 kHz single-tone, body-worn terminal), with samples colored by shadowing state. The model is calibrated on real measurements and reproduces their targeted statistics (mean ≈ 0.3 dB, 18% of samples below -5 dB)

2) CSI and reliability model: The scheduler selects the transmission format from the most recent SNR measurement, denoted $\hat{\gamma}$. In the considered transparent GEO architecture, this CSI is aged by approximately one round-trip time: the selected format is therefore based on $\hat{\gamma}$, while the transmission experiences the channel at its actual later SNR, γ. The resulting mismatch increases with the channel time variation.

Reliability is modeled through an SNR threshold. From physical-layer link-level simulations based on [15], the BLER-versus-SNR curve of each MCS index defines the target SNR $\gamma^{\star}(I_{\mathrm{MCS}})$ required to meet the BLER target ϵ.

Without repetitions, a transmission is reliable if

$$\gamma \geq \gamma^{\star}(I_{\mathrm{MCS}}). \tag{2}$$

With repetitions, soft combining of N_{Rep} transmissions provides up to a 3 dB gain per doubling of N_{Rep} [16, Sec. 3.2.8.2]. The reliability condition becomes

$$\gamma + 10\log_{10}(N_{\mathrm{Rep}}) \geq \gamma^{\star}(I_{\mathrm{MCS}}). \tag{3}$$

In practice, the scheduler evaluates (3) using the aged estimate $\hat{\gamma}$, whereas the success or failure of the transmission is determined by the actual SNR γ.

TABLE II
CALIBRATED STATE-TRANSITION PROBABILITIES

From \ To	Deep	Moderate	LOS
Deep	0.783	0.204	0.013
Moderate	0.083	0.777	0.140
LOS	0.011	0.159	0.830

C. Problem Formulation

Let B denote the buffer size (in bits) the UE must deliver. Transport blocks have a finite size, bounded by $\mathrm{TBS}_{\max} = 1000$ bits and, for a given format, equal to the corresponding entry of the TBS table (Table I). Whenever the transport block selected for a transmission carries fewer bits than the buffer, the remainder is sent in subsequent transmissions, so the buffer is delivered as a *sequence* of transport blocks. Segmentation is therefore not an external parameter but a consequence of the format decisions relative to the buffer size: a format with a small TBS fragments the buffer into more segments, whereas a larger TBS uses fewer. This matters because every segment incurs one round trip over the GEO link, so the number of segments directly affects latency.

We model the transmission of the buffer as a sequence of K segments. Segment i uses a format $(I_{\mathrm{MCS},i}, I_{\mathrm{RU},i}, N_{\mathrm{Rep},i})$ and carries

$$b_i = \min\Big(f(I_{\mathrm{MCS},i}, I_{\mathrm{RU},i}) \,,\; B - \textstyle\sum_{j<i} b_j \Big) \text{ bits}, \quad (4)$$

where $f(\cdot)$ is the transport-block size of Table I. The time cost of segment i is its airtime plus one round trip,

$$\tau_i = N_{\mathrm{RU},i}\, N_{\mathrm{Rep},i}\, T_{\mathrm{RU}} \;+\; \mathrm{RTT}, \quad (5)$$

with T_{RU} the duration of one resource unit, corresponding to 32 ms. Link adaptation pursues two competing goals: minimizing the radio-resource cost $\sum_i N_{\mathrm{RU},i} N_{\mathrm{Rep},i}$ and maximizing the delivered throughput. Over a GEO link these two goals are unified by a single quantity, the total transmission time: the airtime term $N_{\mathrm{RU},i} N_{\mathrm{Rep},i} T_{\mathrm{RU}}$ *is* the resource cost (scaled by T_{RU}), while the per-segment RTT penalizes fragmentation. Minimizing the total time therefore jointly optimizes throughput and resource use. For a buffer delivered in full, maximizing throughput $B/\sum_i \tau_i$ is equivalent to minimizing $\sum_i \tau_i$. We thus formulate link adaptation as

$$\min_{K,\,\{\phi_i\}_{i=1}^{K}} \quad \sum_{i=1}^{K} \tau_i \quad (6a)$$

$$\text{s.t.} \quad \textstyle\sum_{i=1}^{K} b_i = B, \quad (6b)$$

$$\hat{\gamma}_i + 10\log_{10} N_{\mathrm{Rep},i} \geq \gamma^{\star}(I_{\mathrm{MCS},i}), \;\; \forall i \quad (6c)$$

where $\phi_i = (I_{\mathrm{MCS},i}, I_{\mathrm{RU},i}, N_{\mathrm{Rep},i})$ is the format of segment i and $\tau_i = N_{\mathrm{RU},i} N_{\mathrm{Rep},i} T_{\mathrm{RU}} + \mathrm{RTT}$ its time cost (5). The coverage constraint (6b) requires the segments to carry the whole buffer. The reliability constraint (6c) enforces the SNR requirement (3) on each segment at its decision SNR $\hat{\gamma}_i$, i.e. the SNR available at the base station when the segment is scheduled. Since this decision SNR may differ from the actual γ_i, a format feasible for (6) is not guaranteed reliable at transmission time. A clairvoyant optimal solution that decides on the true SNR is used in Section IV as a bound to isolate the cost of this staleness.

In Problem (6), the format selected for each segment determines both its transmission cost and the number of packet bits served, so segmentation and format selection must be optimized jointly. A greedy strategy, such as always choosing the largest feasible TBS, is not necessarily optimal because the reduction in the number of RTTs may be outweighed by the extra resource or repetition cost required by a larger transport block. Since both the format set and the residual buffer size are discrete, the problem admits an exact dynamic-programming solution over the remaining packet size. This provides the global optimum under the considered model and is used as an optimal benchmark for the proposed heuristics.

III. PROPOSED LINK-ADAPTATION ALGORITHMS

We propose two low-complexity *online* heuristics that solve problem (6) greedily, one segment at a time. We compare them against a conservative baseline and a packet-level optimal bound. A real scheduler cannot see the channel of future segments. It must choose the format of the current segment from the information at hand. The two heuristics therefore work *per segment*: from the residual buffer and the current SNR estimate, they choose the transmission format $\phi = (I_{\mathrm{MCS}}, I_{\mathrm{RU}}, N_{\mathrm{Rep}})$, transmit, and move on. They do not plan the segmentation in advance.

Each decision separates the format ϕ into two parts with different roles. The index pair $(I_{\mathrm{MCS}}, I_{\mathrm{RU}})$ sets, through the TBS table (1), how many bits the transport block carries and at what resource cost $N_{\mathrm{RU}} \cdot N_{\mathrm{Rep}}$. The repetition number N_{Rep} lowers the required SNR without changing the payload, and thus controls the reliability constraint (6c). Each policy therefore builds the format in two stages. The *TBS-selection* stage fixes $(I_{\mathrm{MCS}}, I_{\mathrm{RU}})$ from the residual buffer. The *repetition-selection* stage then sets N_{Rep} from the channel estimate. The policies differ in how each stage is performed.

A. TBS Selection

At each segment, let r be the residual buffer (initially B). The TBS-selection stage searches the single-tone TBS table (Table I) for an index pair $(I_{\mathrm{MCS}}, I_{\mathrm{RU}})$ whose transport block covers $\min(r, \mathrm{TBS}_{\max})$ bits. The segment carries that many bits, and the residual buffer shrinks by the same amount. A buffer larger than $\mathrm{TBS}_{\max}$ is thus served over successive segments. The two heuristics differ only in the order in which the table is scanned. This order sets which resource is minimized first.

1) Minimum-resource rule (H1): The first rule scans the resource-unit columns in increasing order. For each column, it scans the MCS indices from lowest to highest. It returns the first pair that covers the segment target $\min(r, \mathrm{TBS}_{\max})$ (Algorithm 1). Since the outer scan is over N_{RU}, the rule *first minimizes the resource-unit count*, then picks the lowest

Algorithm 1 Minimum-resource TBS selection (H1)

Require: residual buffer r; TBS table $f(\cdot,\cdot)$; RU set $\mathcal{N}_{\mathrm{RU}}$; MCS set $\mathcal{M}$
Ensure: index pair $(I_{\mathrm{MCS}}, I_{\mathrm{RU}})$
$t \leftarrow \min(r, \mathrm{TBS}_{\max})$ ▷ segment target
for $I_{\mathrm{RU}} \in \mathcal{N}_{\mathrm{RU}}$ (increasing N_{RU}) **do**
  for $I_{\mathrm{MCS}} \in \mathcal{M}$ (increasing) **do**
    if $f(I_{\mathrm{MCS}}, I_{\mathrm{RU}})$ is valid **and** $f(I_{\mathrm{MCS}}, I_{\mathrm{RU}}) \geq t$ **then**
      return $(I_{\mathrm{MCS}}, I_{\mathrm{RU}})$ ▷ first feasible: min. N_{RU}

Algorithm 2 Lowest-MCS TBS selection (H2)

Require: residual buffer r; TBS table $f(\cdot,\cdot)$; RU set $\mathcal{N}_{\mathrm{RU}}$; MCS set $\mathcal{M}$
Ensure: index pair $(I_{\mathrm{MCS}}, I_{\mathrm{RU}})$
$t \leftarrow \min(r, \mathrm{TBS}_{\max})$ ▷ segment target
for $I_{\mathrm{MCS}} \in \mathcal{M}$ (increasing) **do**
  for $I_{\mathrm{RU}} \in \mathcal{N}_{\mathrm{RU}}$ (increasing N_{RU}) **do**
    if $f(I_{\mathrm{MCS}}, I_{\mathrm{RU}})$ is valid **and** $f(I_{\mathrm{MCS}}, I_{\mathrm{RU}}) \geq t$ **then**
      return $(I_{\mathrm{MCS}}, I_{\mathrm{RU}})$ ▷ first feasible: min. MCS

sufficient MCS. This gives the most compact pair able to carry the segment. The price is a higher coding rate, hence a higher required SNR, which leaves more of the reliability burden to the repetition stage.

2) Lowest-MCS rule (H2): The second rule scans the MCS indices in increasing order. For each index, it scans the resource-unit columns. It returns the first pair that covers the segment target $\min(r, \mathrm{TBS}_{\max})$ (Algorithm 2). Since the outer scan is over I_{MCS}, the rule *first minimizes the MCS*, then picks the smallest sufficient N_{RU}. This selects the lowest coding rate the segment allows. The block is then more robust and needs a lower SNR before any repetition. The price is a larger N_{RU}.

B. Repetition Selection

Given the index pair $(I_{\mathrm{MCS}}, I_{\mathrm{RU}})$ and the SNR estimate $\hat{\gamma}$ available to the scheduler, the repetition-selection stage sets N_{Rep} so that the reliability constraint (6c) is met. Each MCS index has a reference SNR $\gamma^{\star}(I_{\mathrm{MCS}})$ at which a single transmission meets the BLER target. These values are tabulated from link-level curves (Table III). The stage first computes the margin

$$m = \hat{\gamma} - \gamma^{\star}(I_{\mathrm{MCS}}), \tag{7}$$

that is, how much SNR the link has above (or below) what the chosen MCS needs without repetition. A negative margin means a single transmission would fail and must be backed by repetitions. The repetition number is selected using a staircase rule, where each 3 dB SNR deficit is compensated by doubling N_{Rep}). This rule is intentionally slightly conservative, as it provides a safety margin against small CSI variations. Furthermore, it is open-loop: the repetition number is set directly from the current channel estimate, without any feedback iteration.

TABLE III
REFERENCE SNR $\gamma^{\star}(I_{\mathrm{MCS}})$ (DB) AT WHICH A SINGLE TRANSMISSION ($N_{\mathrm{Rep}} = 1$) MEETS THE BLER TARGET $\epsilon = 10^{-1}$, EXTRACTED FROM THE LINK-LEVEL NPUSCH BLER CURVES (3.75 KHZ SINGLE-TONE, AWGN).

I_{MCS}	$\gamma^{\star}$ (dB)	I_{MCS}	$\gamma^{\star}$ (dB)
0	−1.9	6	3.1
1	−1.3	7	4.1
2	−0.7	8	5.1
3	−0.1	9	6.3
4	1.1	10	7.0
5	2.1		

$$N_{\mathrm{Rep}} = \begin{cases} 1, & m > 0, \\ 2, & -3 < m \leq 0, \\ 4, & -6 < m \leq -3, \\ 8, & -9 < m \leq -6, \\ 16, & -12 < m \leq -9, \\ 32, & -15 < m \leq -12, \\ 64, & m \leq -15. \end{cases} \tag{8}$$

C. Policies

The two stages above combine into the three policies. The *conservative* is our baseline, it does not adapt to the channel at all. It fixes the MCS to the most robust value ($I_{\mathrm{MCS}} = 0$), picks the smallest N_{RU} that covers the buffer at that MCS, and fixes the repetition number to $N_{\mathrm{Rep}} = 4$. These choices follow a worst-case design, the policy serves the link down to about -7.9 dB. This covers about 97% of the measured channel. The baseline is therefore robust but channel-agnostic: it over-provisions repetitions in good conditions and may under-provision them in the deepest fades. The two proposed heuristics, in contrast, keep the adaptive repetition stage (8) and differ only in the TBS rule. *H1* (minimum-resource) starts from the most compact pair; since this pair uses a higher MCS, the margin m is often tight and more repetitions are spent. *H2* (lowest-MCS) starts from the lowest-MCS pair; the required SNR is then low, the margin is usually favorable, and few repetitions are needed. The two heuristics thus place the resource–reliability trade-off differently: H1 adds repetitions on top of a compact pair, while H2 spends on the resource-unit footprint of a robust pair.

All three policies have $O(|\mathcal{N}_{\mathrm{RU}}| \cdot |\mathcal{M}|)$ worst-case complexity for the table scan plus an $O(1)$ repetition lookup, i.e., negligible cost per scheduling decision.

D. Packet-Level Optimal Bound

The heuristics decide each segment greedily. They do not look ahead at the channel of future segments. To measure how much this online operation loses, we compute the exact optimum of problem (6) over the *whole* packet. The format chosen on one segment changes how many bits remain. It

therefore changes how many further segments, each costing one RTT, are needed. A greedy choice is thus generally suboptimal, and the packet must be planned jointly.

Problem (6) has optimal substructure. The minimum-time plan for a residual buffer r at segment s is the best choice of format for segment s, plus the minimum-time plan for the remaining bits at segment $s+1$. Let $V(r,s)$ be that minimum total time. Then

$$V(r,s) = \min_{(I_{\mathrm{MCS}}, I_{\mathrm{RU}}, N_{\mathrm{Rep}})} \Big[\tau_s + V\big(r - b_s,\ s+1\big) \Big], \qquad (9)$$

with τ_s and b_s given by (5) and (4), and $V(0,\cdot) = 0$. The minimization keeps only the formats that satisfy the reliability constraint (6c) at the decision SNR $\hat{\gamma}_s$ of segment s. We solve it by dynamic programming over the discrete residual buffer. This returns the globally optimal segmentation and the per-segment formats. As a bound, the planner is allowed to see the full sequence of decision SNRs of the packet. The online heuristics do not have this advantage, so the planner upper-bounds the performance reachable under the available CSI. Running it with the delayed CSI gives the realistic optimum. Running a clairvoyant variant with the true SNR isolates the cost of CSI staleness (Section IV).

IV. Simulation and Results

A. Simulation Scenario and Methodology

We evaluate the policies of Section III in a simulator driven by the calibrated LMS channel model of Section II-B. We compare five schemes: the conservative baseline, the two online heuristics (H1, minimum-resource; H2, lowest-MCS), and the packet-level optimal planner in two variants, one deciding on aged CSI and one clairvoyant. The online heuristics and the aged-CSI optimum all decide from the SNR observed on the previous transmission, reproducing the one-RTT delayed CSI of Section II-B. The clairvoyant optimum instead decides on the true SNR and is unattainable in practice; it serves only to isolate the cost of CSI staleness.

To make large-scale evaluation tractable, we first characterize the physical layer offline and then use the resulting error model in packet-level simulations. The NPUSCH chain is simulated over the 3.75 kHz single-tone AWGN channel to obtain a BLER-versus-SNR curve for each MCS index, as shown in Fig. 3. The SNR threshold $\gamma^\star(I_{\mathrm{MCS}})$ is defined as the SNR at which the corresponding curve reaches the target BLER $\epsilon = 10^{-1}$, and the resulting values are reported in Table III. During packet-level evaluation, each transmission is resolved by a Bernoulli trial whose error probability is taken from the offline BLER curve associated with the selected MCS index:

$$p_{\mathrm{e}} = \mathrm{BLER}_{I_{\mathrm{MCS}}}\left(\gamma_{\mathrm{eff}}\right). \qquad (10)$$

The effective SNR is computed as

$$\gamma_{\mathrm{eff}} = \gamma_{\mathrm{act}} + 10\log_{10}\left(N_{\mathrm{Rep}}\right), \qquad (11)$$

where γ_{act} is the actual SNR experienced by the transmission and $10\log_{10}(N_{\mathrm{Rep}})$ models the ideal combining gain from repetitions.

TABLE IV
Simulation parameters.

Parameter	Value
Subcarrier spacing	3.75 kHz (single-tone)
RU duration	32 ms (16 slots)
Channel	LMS model (body-worn tracker), Fig. 2
Link abstraction	per-MCS BLER curves
Satellite	GEO transparent
Modulation	QPSK
Repetition set	$\{1, 2, 4, 8, 16, 32, 64\}$
Target BLER	10^{-1}
HARQ retransmissions	Disabled
Number of packets	3000
Buffer size B	$\mathcal{U}[100, 3000]$ bits

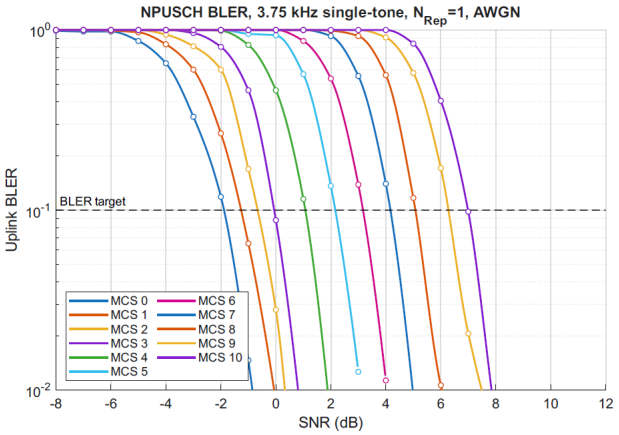


Fig. 3. Link-level NPUSCH BLER versus SNR for each MCS (3.75 kHz single-tone, AWGN, $N_{\mathrm{Rep}} = 1$). The crossings with the target $\epsilon = 10^{-1}$ define the reliability thresholds $\gamma^\star(I_{\mathrm{MCS}})$ of Table III.

The thresholds are extracted for the shortest TBS corresponding to $N_{\mathrm{RU}} = 1$, which provides conservative thresholds because finite-blocklength losses are stronger for short blocks [17]. In our simulations, the maximum threshold shift across N_{RU} is about 1.5 dB. However, this does not affect our benchmark results since applying the $N_{\mathrm{RU}} = 1$ thresholds to larger allocations biases all policies toward over-provisioning repetitions rather than violating the reliability constraint. HARQ retransmissions are disabled, so each transmission is evaluated in single-shot mode: a NACK is counted as a failure. Traffic packets are uniformly drawn between 100 and 3000 bits and segmented when they exceed $\mathrm{TBS}_{\max} = 1000$ bits. All policies use the same traffic and SNR sequences with a fixed seed for paired comparison; the main parameters are listed in Table IV.

We report three key performance indicators.

1) Block error rate (BLER): The BLER is the fraction of transmitted transport blocks that are not correctly decoded over all transmissions of a run,

$$\mathrm{BLER} = 1 - \frac{1}{N_{\mathrm{tx}}} \sum_{i=1}^{N_{\mathrm{tx}}} \mathrm{ACK}_i, \qquad (12)$$

where N_{tx} is the total number of transmitted blocks. It measures how well each policy meets the reliability target ϵ: a policy is reliable when its BLER stays at or below ϵ.

2) Resource occupation: The resource occupation is the airtime, in resource units, consumed to serve the traffic. For a transmission it is $N_{\mathrm{RU}} \cdot N_{\mathrm{Rep}}$; per packet it is summed over the segments and transmissions, and we report the per-packet average,

$$R = \frac{1}{N_{\mathrm{pkt}}} \sum_{\text{packets}} \sum_{\text{tx in packet}} N_{\mathrm{RU}} \cdot N_{\mathrm{Rep}}. \tag{13}$$

It is the quantity the objective (6a) minimizes and the primary efficiency metric: a lower R leaves more of the shared carrier for other devices.

3) Throughput: The throughput is the rate of successfully delivered useful bits, computed as the total delivered payload divided by the total transmission time (including the RTT of each transmission),

$$G = \frac{\sum_{\text{packets}} D_{\text{delivered}}}{\sum_{\text{packets}} T_{\text{tx}}}, \tag{14}$$

where $D_{\text{delivered}}$ is the payload acknowledged and T_{tx} the corresponding airtime plus RTT. Our metric G captures the throughput actually achieved: it rewards policies that deliver data with few repetitions and few failed (hence wasted) transmissions.

Together, the three KPIs span the trade-off of problem (6): BLER on the reliability side, resource occupation on the cost side, and throughput as the net delivered performance that balances the two.

B. Results

We report the three primary KPIs (block error rate, throughput, and resource occupation) for the five schemes. Fig. 4 summarizes the three KPIs. The simulator generate 3000 packets that are uniformly drawn between 100 and 3000 bits, yielding hundreds of block errors per policy. Hence, tight Wilson confidence interval on the BLER (half-width at most ±0.010 at the 95% level) and the relative standard error on resource occupation and throughput of about 1–2 %

1) Reliability and resource trade-off: The five schemes span the whole reliability–resource plane (Fig. 4). The conservative baseline, with its fixed $N_{\mathrm{Rep}} = 4$, drives the BLER far below the target (0.011) but at the highest resource cost ($R = 249$) and the lowest throughput (134 bit/s), since the static repetition count wastes airtime during the long unshadowed intervals. At the opposite end, the aged-CSI optimum minimizes resources ($R = 43$) but, as discussed below, fails the reliability target. The two heuristics operate between these extremes and adapt the format to the observed SNR.

2) Heuristics: Both heuristics adapt the repetition count to the observed SNR, and both land close to the target under aged CSI. H2 (lowest-MCS) reaches a BLER of 0.151 while using only $R = 84$ and achieving 345 bit/s. Relative to the conservative baseline, it cuts resource occupation by 66% and raises throughput by 157%, by replacing the static repetition count with a channel-adaptive one. H1 (minimum-resource) is more compact in MCS–RU terms but, favoring higher-order formats, it requires more repetitions and ends up both less reliable (0.197) and less efficient (261 bit/s at $R = 115$). H2 thus offers the better heuristic trade-off: it is more reliable and more efficient than H1, and it approaches the resource economy of the optimal bounds. Neither heuristic holds the BLER strictly at $\epsilon = 10^{-1}$ under aged CSI, however; H2 stays slightly above the target (0.151 vs. 0.10), reflecting the difficulty of meeting the constraint when the CSI is one RTT old. This residual gap is quantified next.

3) Cost of outdated CSI: The penalty imposed by the GEO round-trip delay is isolated by comparing the two throughput-optimal bounds, which run the *same* packet-level optimization and differ only in the SNR on which they decide. At an *identical* resource occupation ($R \approx 43$), the clairvoyant optimum (perfect CSI) attains a BLER of 0.053 and a throughput of 558 bit/s, whereas the realistic optimum (aged CSI), deciding on the previous SNR sample, degrades to a BLER of 0.281 and 378 bit/s: a 5.3× increase in error rate and a 32% throughput loss caused solely by the staleness of the channel information. The aged-CSI optimum, although the most frugal scheme, thus violates the target by about 2.8× (0.28 vs. 0.10): aggressively minimizing resources on stale CSI is self-defeating. By contrast, H2 stays far closer to the target under the same outdated CSI, reaching 62% of the clairvoyant throughput at 1.9× its resource use. The residual gap to the (unattainable) perfect-CSI bound measures the headroom left by CSI staleness.

4) Discussion: Together, the three KPIs show the central trade-off of link adaptation in this regime. Fixing the format without looking at the channel (conservative) keeps the link reliable but uses far too many resources. Minimizing resources on old CSI (aged-CSI optimum) is cheap but misses the reliability target by a wide margin. The two heuristics sit in between: by adapting the repetition count to the observed SNR, they stay close to the target at a moderate resource cost, and H2 comes near the best resource–throughput trade-off. The gap between the aged-CSI and perfect-CSI optima sets a limit on what any real scheme could gain by reducing CSI staleness, for example through channel prediction. This defines the design space for NB-IoT link adaptation over GEO NTN.

V. Conclusion

We studied uplink link adaptation for single-tone NB-IoT over a GEO satellite. Link adaptation is well studied for terrestrial NB-IoT, but the long GEO round-trip delay changes the problem; and adaptation has been studied in satellite contexts, but for protocols much simpler than NB-IoT. We cast the problem as a joint throughput–resource optimization under a reliability target, where the base station selects the transmission format from outdated channel-state information. We proposed two low-complexity online heuristics, H1 (minimum-resource) and H2 (lowest-MCS), compared them against a conservative baseline, and derived a packet-level optimal bound by dynamic programming. The evaluation

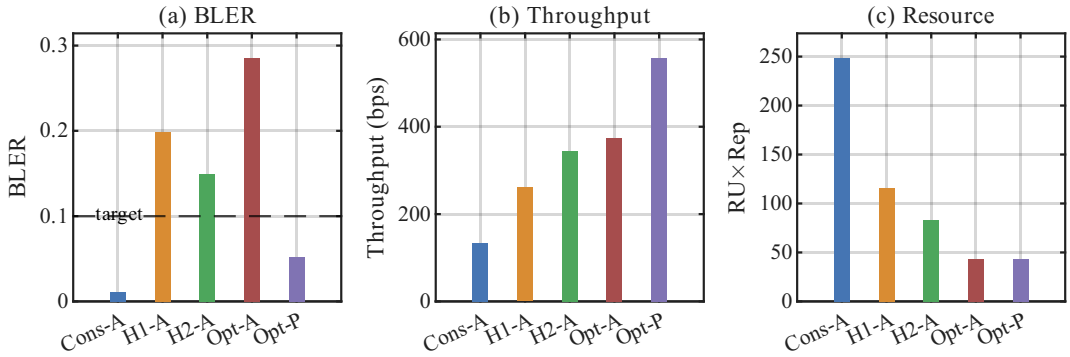


Fig. 4. Primary KPIs of the five schemes over mixed traffic ($B \in [100, 3000]$ bits): (a) block error rate, with the dashed line marking the target $\epsilon = 10^{-1}$; (b) throughput including the GEO round-trip time; (c) resource occupation $R = \sum N_{\mathrm{RU}} N_{\mathrm{Rep}}$. All schemes except the clairvoyant optimum decide on aged CSI (A); the clairvoyant optimum uses the perfect CSI (P)

was driven by a land-mobile-satellite channel model calibrated on real GEO NB-IoT uplink measurements.

The repetition rule is open-loop: it sets the repetition number directly from the current channel estimate, without feedback iteration. This replaces the ACK/NACK loops of terrestrial schemes, which would need many round trips to converge and are too slow at the GEO RTT. Among the online schemes, H2 gave the best trade-off, cutting resource use by 66% and raising throughput by 157% over the conservative baseline while staying close to the reliability target.

The main finding concerns the cost of outdated CSI. At equal resource use, deciding on a channel sample that is one round trip old raised the block error rate by $5.3\times$ and lowered throughput by 32% relative to a clairvoyant optimum. Even the best heuristic could not strictly hold the reliability target under this delayed CSI. The staleness penalty, rather than the choice of policy, is therefore the dominant limit in this regime, and recovering it is the key lever for future designs.

These results are specific to the channel studied here: a body-worn tracker in an open-sky area, whose dynamics are summarized by a lag-one autocorrelation of 0.77. In this regime the channel changes slowly between transmissions, so the previous SNR sample, although one RTT old, still carries useful information. For more dynamic channels, where the autocorrelation is lower and the CSI ages faster, the staleness penalty would be larger and more robust schemes would be needed. Studying such schemes is the subject of our next work.

References

[1] O. Liberg, M. Sundberg, Y.-P. E. Wang, J. Bergman, J. Sachs, and G. Wikström, "Narrowband Internet of Things for non-terrestrial networks," *IEEE Communications Standards Magazine*, vol. 4, no. 4, pp. 49–55, 2020.

[2] C. Yu, L. Yu, Y. Wu, Y. He, and Q. Lu, "Uplink scheduling and link adaptation for narrowband Internet of Things systems," *IEEE Access*, vol. 5, pp. 1724–1734, 2017.

[3] O. Elgarhy, L. Reggiani, H. Malik, M. M. Alam, and M. A. Imran, "Rate-latency optimization for NB-IoT with adaptive resource unit configuration in uplink transmission," *IEEE Systems Journal*, vol. 15, no. 1, pp. 265–276, 2021.

[4] R. Barbau, V. Deslandes, G. Jakllari, J. Tronc, J.-F. Chouteau, and A.-L. Beylot, "NB-IoT over GEO satellite: Performance analysis," in *2020 10th Advanced Satellite Multimedia Systems Conference and the 16th Signal Processing for Space Communications Workshop (ASMS/SPSC)*, pp. 1–8, 2020.

[5] A. Petrosino, G. Sciddurlo, S. Martiradonna, D. Striccoli, G. Piro, and G. Boggia, "WIP: An open-source tool for evaluating system-level performance of NB-IoT non-terrestrial networks," in *2021 IEEE 22nd International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, pp. 236–239, 2021.

[6] O. Kodheli, N. Maturo, S. Chatzinotas, S. Andrenacci, and F. Zimmer, "NB-IoT via LEO satellites: An efficient resource allocation strategy for uplink data transmission," *IEEE Internet of Things Journal*, vol. 9, no. 7, pp. 5094–5107, 2022.

[7] C. Fei, B. Zhao, C. Wu, W. Yu, B. Han, and J. Su, "Towards efficient radio resource scheduling in lte-based satellite mobile communication system," in *2015 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pp. 167–172, IEEE, 2015.

[8] G. Araniti, M. Condoluci, A. Orsino, A. Iera, and A. Molinaro, "Effective resource allocation in 5g-satellite networks," in *2015 IEEE International Conference on Communications (ICC)*, pp. 844–849, IEEE, 2015.

[9] S. Cioni, R. De Gaudenzi, and R. Rinaldo, "Channel estimation and physical layer adaptation techniques for satellite networks exploiting adaptive coding and modulation," *International Journal of Satellite Communications and Networking*, vol. 26, no. 2, pp. 157–188, 2008.

[10] O. A. Coding and M. for Satellite Network with ML-based CSI Prediction, "Optimizing adaptive coding and modulation for satellite network with ML-based CSI prediction," in *2019 IEEE Wireless Communications and Networking Conference (WCNC)*, pp. 1–6, 2019.

[11] 3GPP, "Evolved universal terrestrial radio access (E-UTRA); physical channels and modulation," Technical Specification TS 36.211, 3rd Generation Partnership Project (3GPP), 2020.

[12] 3GPP, "Evolved universal terrestrial radio access (E-UTRA); physical layer procedures," Technical Specification TS 36.213, 3rd Generation Partnership Project (3GPP), 2020.

[13] F. P. Fontan, M. Vázquez-Castro, C. E. Cabado, J. P. Garcia, and E. Kubista, "Statistical modeling of the lms channel," *IEEE Transactions on vehicular technology*, vol. 50, no. 6, pp. 1549–1567, 2001.

[14] 3GPP, "Study on New Radio (NR) to support non-terrestrial networks," Technical Report TR 38.811, 3rd Generation Partnership Project (3GPP), Sept. 2020. Release 15. [Online]. Available: https://www.3gpp.org/dynareport/38811.htm.

[15] MathWorks, "NB-IoT NPUSCH Block Error Rate Simulation." [Online]. Available: https://fr.mathworks.com/help/lte/ug/nb-iot-npusch-block-error-rate-simulation.html. Accessed: Jun. 19, 2026.

[16] O. Liberg, M. Sundberg, Y.-P. E. Wang, J. Bergman, J. Sachs, and G. Wikström, *Cellular Internet of Things: From Massive Deployments to Critical 5G Applications*. Academic Press (Elsevier), 2nd ed., 2020.

[17] Y. Polyanskiy, H. V. Poor, and S. Verdú, "Channel coding rate in the finite blocklength regime," *IEEE Transactions on Information Theory*, vol. 56, no. 5, pp. 2307–2359, 2010.