Performance Analysis of Satellite HARQ under Partial Feedback Conditions

Estefanía Recayte*, Carla Amatetti†, Amira Alloum‡

* Institute of Communications and Navigation, German Aerospace Center (DLR), Weßling, Germany
† Department of Electrical, Electronic, and Information Engineering, University of Bologna, Bologna, Italy
‡ Qualcomm, France

***Abstract*—Non-terrestrial networks (NTNs) have emerged as a key enabler for achieving ubiquitous, seamless, and continuous connectivity, especially following the inclusion of satellite communication in the 5G ecosystem by the 3GPP starting with Release 17. This integration poses significant challenges, particularly in adapting terrestrial mechanisms such as the hybrid automatic repeat request (HARQ) protocol to the satellite environment, where significant round-trip delays prevent its straightforward application. In this work, we investigate the performance of HARQ in satellite communication scenarios under three feedback configurations: fully enabled, fully disabled, and partially disabled. We propose a novel approach for the partially disabled case, employing optimal erasure coding to reduce feedback overhead while improving transmission efficiency compared to the fully disabled configuration. Through analytical evaluation and simulations, we highlight the key trade-offs and performance metrics of the system across the various configurations studied.**

I. Introduction

Non-terrestrial networks (NTNs) have recently been recognized as a key enabler for achieving seamless, continuous, and ubiquitous communications. This recognition is a result of the inclusion of satellite connectivity in the 5G ecosystem by the third generation partnership project (3GPP), beginning with Release 17, where satellites were officially integrated as a standard component of next-generation wireless networks. This integration not only introduces new challenges but also highlights the need for novel and alternative solutions when adapting terrestrial technologies to the satellite environment [1]. In terrestrial networks, data transmission reliability is typically ensured through the use of the hybrid automatic repeat request (HARQ) protocol. However, the use of HARQ in satellite networks is not straight forward mainly due to the non negligible round trip time (RTT) that satellite communications experience due to the long distance between the user equipments (UEs) and the satellite [2].

To improve the feasibility of HARQ implementation in satellite networks, Release 18 of the 3GPP standard introduced enhancements targeting the challenges posed by the high round-trip time in satellite communications compared to terrestrial systems. This increased latency causes conventional HARQ mechanisms to suffer from significant throughput degradation due to stalling. To address this, two main changes were adopted. First, the number of parallel HARQ processes supported was increased to 32 [3], which may be sufficient to mitigate stalling in certain NTN scenarios, such as in some communication links with low-Earth orbit (LEO) satellites. However, this extension may still be insufficient to accommodate the longer delays typical of geostationary satellite links. The second decision involves disabling feedback (FB) [4], [5], allowing transmissions to continue without the HARQ process waiting for an acknowledgment (ACK) or negative acknowledgment (NACK) signals. The disabling feedback is configurable on a per UE and HARQ process basis. This offers greater flexibility in managing feedback signaling and reducing unnecessary overhead in scenarios where reliable feedback is difficult to maintain.

These adaptations aim to better align HARQ operation with the unique constraints of satellite links. However, they may not be sufficient to ensure the efficient operation of the protocol. HARQ for satellite communications has been the subject of recent studies, e.g., in [6]–[8]. In [6] it is shown that HARQ, although effective in terrestrial 5G NR, suffers from significant performance limitations over satellite links. This observation motivates the need for adapting HARQ mechanisms to the specific characteristics of non-terrestrial networks. The work in [6] shows how that carefully designed HARQ scheduling can substantially improve reliability and efficiency in NTN. On the other hand, the idea of incorporating additional coding within HARQ has already been analyzed in the terrestrial domain, for example through fountain-code-based proposed in [9] and analyzed in terms of throughput in [10].

While the study of outer coding for HARQ in NTN remains largely unexplored, this work investigates the performance of the HARQ protocol combined with maximum-distance separable (MDS) codes in a satellite communication scenario to assess these limitations and explore potential improvements. Different from previous literature, we analyze and compare the system under three different feedback configurations: fully enabled, fully disabled, and partially disabled. In particular, we propose the integration of optimal erasure codes, i.e. MDS, into a partially disabled HARQ configuration, which achieves a balance between the two extremes: it reduces the feedback overhead compared to the fully enabled case, while at the same time enhancing transmission efficiency relative to the fully disabled case.

II. Background

We begin by introducing the key concepts of the hybrid automatic repeat request protocol and the maximum separable code.

A. Hybrid automatic repeat request

HARQ is a retransmission protocol designed to enhance data reliability by combining forward error correction (FEC) with retransmission mechanisms. At the transmitter, redundant bits are appended to the data, enabling the receiver to exploit this redundancy for error detection and correction. If residual errors cannot be corrected, the receiver buffers the corrupted data and requests a retransmission. Reliability is further improved as HARQ combines information from multiple transmissions, typically using one of the following techniques.

The first, *chase combining*, involves retransmitting the same information, allowing the receiver to perform soft combining, such as maximum ratio combining, to enhance decoding performance. This technique requires less memory resources because only the combined information needs to be stored.

The second technique, *incremental redundancy*, transmits additional coded bits at each retransmission, effectively lowering the coding rate and improving the error correction capacity. However, this method increases the complexity of the decoding process and demands more memory resources from the UE.

The HARQ protocol assigns a transport block (TB) to each process. The processes are transmitted, and once a TB is successfully decoded, an acknoledgment (ACK) message is sent, allowing the corresponding process to be reused for a new transmission. When a negative acknoledgment (NACK) is received, one of the previously mentioned techniques is applied. HARQ over NTN networks suffers from stalling due to stop-and-wait behavior, where all processes may be paused while the satellite waits for user feedback to determine whether a retransmission is needed or a new process can be scheduled. This issue is primarily caused by the long round-trip time inherent to satellite links.

B. Maximum distance separable codes

MDS codes are optimal in the sense that they achieve the Singleton bound [11], meaning that they achieve the maximum possible data recovery for a given code length and number of input blocks. In practice, an MDS(N, n) code encodes n input blocks into N outputs, from which any combination of n are sufficient to reconstruct the original data. Such codes can be particularly effective in supporting the HARQ process by reducing or eliminating the need for continuous feedback, while still ensuring successful data recovery as long as a sufficient number of encoded blocks are successfully received. MDS-HARQ feedback is minimized since the receiver only needs to send a single acknowledgment once any n encoded blocks have been successfully received and the original data can be fully reconstructed, eliminating the need for feedback after every individual transmission. We investigate this technique as an alternative to existing chase combining and incremental redundancy methods.

III. System Model

We consider an NTN scenario involving a LEO satellite and a UE. The LEO satellite is assumed to have a regenerative payload, meaning it is equipped with onboard processing capabilities. The UE employ time division multiplexing, which prevents it from transmitting and receiving simultaneously, whereas the LEO operates in full-duplex frequency division duplex mode allowing it to transmit and receive at the same time. Our analysis focuses on the downlink channel, where the satellite is required to reliably transmit data to the device. The transmission is carried out using TBs of equal size and duration $t_{\text{TB}} = 1$ ms, with each transport block representing a separate HARQ transmission process. As agreed in Release 17 [3], $N_p = 32$ supported processes are assumed to mitigate and stop-and-wait due to the round trip time (RTT) and to maintain a continuous transmission flow. Table 5 in [6] presents satellite scenarios, including orbit parameters and elevation angles, where stop-and-wait behavior does not occur because the number of HARQ processes is sufficient to maintain continuous transmission, i.e. fewer than $N_p = 32$. The propagation time of a TB to reach the receiver depends on the distance between the UE and the LEO, and is denoted by t_{prop}.

Each transport block experiences independent Rayleigh fading. The channel amplitude is h, assumed to be constant over each block. This results in a varying signal-to-noise ratio denoted by Γ, computed as $\Gamma = \frac{h^2 P}{\sigma^2}$ where $h^2 P$ is the received power and σ^2 is the Gaussian noise power. The random variable Γ follows an exponential distribution, with probability density function

$$f_\Gamma(x) = \begin{cases} \frac{1}{\bar{\Gamma}} \exp\left\{ -\frac{x}{\bar{\Gamma}}, \right\} & x > 0 \\ 0 & x \leq 0 \end{cases}$$

where $\bar{\Gamma}$ is the average SNR.

TBs are protected by a channel code with rate R and each block is successfully decoded when the condition $\log_2(1 + \bar{\Gamma}) \geq \mathsf{R}$ is satisfied. From which, we can also define $\gamma_{\mathsf{th}} = 2^{\mathsf{R}} - 1$ as the decoding threshold. In this way, we define $\mathsf{P}_{\mathsf{dec}}(\Gamma, \gamma_{\mathsf{th}})$ as the probability that a TB is successfully decoded at the receiver evaluated as

$$\mathsf{P}_{\mathsf{dec}}(\Gamma, \gamma_{\mathsf{th}}) = \Pr\{\Gamma \geq \gamma_{\mathsf{th}}\} = \int_{\gamma_{\mathsf{th}}}^{\infty} f_\Gamma(x)\, dx = e^{-\frac{\gamma_{\mathsf{th}}}{\bar{\Gamma}}}.$$

For simplicity, we refer to $\mathsf{P}_{\mathsf{dec}}(\Gamma, \gamma_{\mathsf{th}})$ as $\mathsf{P}_{\mathsf{dec}}$ in the rest of the paper.

It is assumed that the satellite has continuously has data to transmit during its visibility windows, denoted by W_{tx}, within which transmission occur. A total of TB_{new} new processes are transmitted over the transmission windows W_{tx}, of which $\mathsf{TB}_{\text{succ}}$ are successfully received. When feedback is enabled, the UE transmits an ACK or NACK message to the satellite, indicating the outcome of the last TB reception. The feedback duration is indicated with t_{FB}.

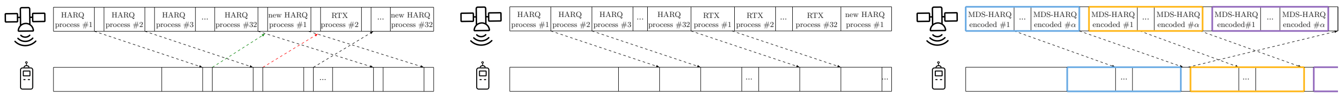

(a) HARQ feedback enabled (b) HARQ feedback disabled (c) MDS-HARQ

Fig. 1. HARQ schemes: (a) feedback enabled, (b) feedback disabled, and (c) MDS-HARQ.

IV. HARQ

In this section, we present the three different HARQ approaches considered in this work. For an easy comparison between schemes, HARQ is assumed to be non adaptive meaning that there is no changes in the modulation, coding and bandwidth. A truncated HARQ scheme is also assumed where the maximum number of retransmission is set to $r_{\max}$. We compare the performance under the following configurations: (*A*) feedback is enabled for all processes, (*B*) feedback is disabled for all processes, and (*C*) feedback is partially enabled, with the proposed MDS-HARQ scheme applied.

A. HARQ feedback enabled

In the feedback enabled configuration illustrated in Fig. 1a, the satellite transmits N_p processes, a process is transmitted by means of a transport block. The UE attempts to decode each upon reception. When the UE successfully decodes the i-th porcess, it sends an ACK, prompting the scheduler to allocate a new TB for the i-th HARQ process. If decoding fails, due to severe fading affecting the received block, the UE sends a NACK and stores the received TB. The satellite then retransmits an identical replica of the failed TB. The UE applies combination between the stored and the retransmitted block using maximal-ratio combining (MRC) and attempts decoding again. If decoding fails again, the information is preserved and incrementally updated with each subsequent retransmission. This process continues iteratively until successful decoding is achieved or the maximum number $r_{\max}$ of retransmissions is reached.

The SNR after MRC of i retransmissions is given by

$$\Gamma_{\mathsf{MRC}} = \frac{P}{\sigma^2} \sum_{l=1}^{i} |h_l|^2$$

where h_l denotes the channel amplitude of the l-th retransmission. Transmission is successful when $\Gamma_{\mathsf{MRC}} \geq \gamma_{th}$.

B. HARQ feedback disabled

In the feedback disabled configuration, as illustrated in Fig. 1b, the UE attempts to decode each received TB upon reception. Unlike the feedback enabled case, no acknowledgment (ACK/NACK) is sent to the satellite. Consequently, the satellite blindly transmits a predetermined number $r_{\max}$ of retransmissions for each TB. If the initial decoding attempt fails, the UE employs MRC to combine the received retransmissions, thereby enhancing the probability of successful decoding. However, in cases where the UE successfully decodes a TB before all $r_{\max}$ retransmissions are received, the remaining transmissions result in unnecessary use of spectral and energy resources.

C. MDS-HARQ - Partial Feedback

A possible implementation of MDS-HARQ is illustrated in Fig. 1c. The proposed MDS-HARQ partial feedback scheme constructs N encoded blocks from n original processes using a MDS(N, n) code, providing the satellite with $N-n$ redundant blocks for potential retransmissions. To fairly compare this scheme with respect to the others, we assume that the number of encoded blocks are $N = n \cdot r_{\max}$. This coding enables reliable reconstruction of the original n processes at the UE from any n successfully received out of the N encoded TB.

The procedure of the scheme can be described as follows: the satellite encodes a group of $n \leq N_p$ processes into N encoded blocks, and transmits $\alpha \geq n$ of them to the device, as illustrated in the figure, where colors refer to different encoding processes. Feedback is received after $\alpha + N_p$ processes[1], indicating whether the group of n processes has been successfully reconstructed. If decoding fails, additional blocks from the remaining $N - \alpha$ redundancy are transmitted until success or until all blocks have been sent. An enhanced feedback mechanism may additionally be employed, specifying the number of failed transmissions and enabling the LEO to transmit the corresponding redundancy proactively.

The main advantage of the scheme is two-fold: first, it reduces the amount of feedback required by a factor of n, second, it allows the transmission of predetermined redundancy in advance, improving reliability and efficiency. Note that unlike the conventional HARQ schemes, the MDS-HARQ proposed does not rely on combining techniques at the receiver.

In this context, we investigate the following feasible MDS-HARQ scheme implementations: (i) a feedback-based scheme, in which the satellite transmits α TBs without redundancy and receives a single feedback message for the encoded group. Based on this feedback, either a retransmission is performed or a new encoded processes are transmitted. (ii) A feedback-free scheme, in which the satellite transmits α TBs with a priori redundancy based on the average number of TBs required for successful decoding, eliminating the need for any feedback exchange. (iii) A feedback-based smart scheme similar to (i) assuming that the NACK contains information on the number of unsuccessful TBs. Upon reception of the NACK, the corresponding number of blocks is transmitted.

[1]Processing time of the blocks is omitted for simplicity.

V. MDS-HARQ Analysis and performance metrics

We provide next a performance evaluation of the proposed MSD-HARQ scheme. We derive the average number of transmissions needed for successfully receive n processes.

Let T denote the random variable corresponding to the number of MDS-HARQ encoded TBs received upon the successful reconstruction of the n original data blocks in the proposed approach. The average number of transmissions required to successfully decode the n is denoted by $\bar{T}_{\text{MDS-HARQ}}$ and is given by

$$\bar{T}_{\text{MDS-HARQ}} = \mathbb{E}[T] = \sum_{\tau=0}^{\infty} \tau \Pr\{T=\tau\}. \tag{6}$$

The expression $\Pr\{T=\tau\}$ represents the probability that successful decoding occurs exactly when τ encoded transport blocks are received. According to the (N,n)-MDS proposed scheme, decoding is not possible if fewer than n blocks are received Therefore

$$\Pr\{T=\tau\} = 0 \text{ for } \tau < n.$$

When exactly n encoded TBs are received, successful reception requires all them to be correctly decoded, that is

$$\Pr\{T=\tau\} = (\mathsf{P}_{\mathsf{dec}})^n \text{ for } \tau = n.$$

For the case $\tau > n$, we denote the probability of successful decoding by P_τ. Let J be the random variable representing the number of transport blocks among the first n that are not successfully decoded. Then, $N-J$ successfully TBs are successfully received, and J follows a binomial distribution with parameters n and $\mathsf{P}_{\mathsf{dec}}$

$$\Pr\{J=j\} = \binom{n}{n-j} (\mathsf{P}_{\mathsf{dec}})^{n-j} (1-\mathsf{P}_{\mathsf{dec}})^j .$$

When $j > 0$, additional encoded TBs are needed. Let us denote with $A = T-n$ denote the number of additional blocks transmitted beyond the initial n. Among these $T-n$ additional blocks, exactly J must be successfully decoded, with the $j-$th one being the final block that completes the decoding. The probability of this event is

$$\Pr\{A=\tau-n|J=j\} = \binom{\tau-n-1}{j-1} (\mathsf{P}_{\mathsf{dec}})^j (1-\mathsf{P}_{\mathsf{dec}})^{\tau-n-j} .$$

Thus, the total probability of successful decoding at $\tau > n$ is

$$P_\tau = \sum_{j=1}^{\tau-n} \Pr\{A=\tau-n|J=j\} \Pr\{J=j\}$$

this becomes

$$P_\tau = \sum_{j=1}^{\tau-n} \binom{n}{n-j} \binom{\tau-n-1}{j-1} (\mathsf{P}_{\mathsf{dec}})^n (1-\mathsf{P}_{\mathsf{dec}})^{\tau-n}.$$

Summing up we have

$$\Pr\{T=\tau\} = \begin{cases} 0 & \text{if } \tau \le n, \\ (\mathsf{P}_{\mathsf{dec}})^n & \text{if } \tau = n, \\ P_\tau & \text{if } \tau > n. \end{cases}$$

and if we plug in these results into (6) we obtain the average of transmission needed for the MDS scheme proposed in (5) at the top of the page.

Performance metrics

Throughput: The throughput, η, is defined as the ratio of successfully received processes (or TB) to the total duration of the transmission window, written as

$$\eta = \frac{\mathsf{TB}_{\text{succ}}}{\mathsf{W}_{\text{tx}}}.$$

It should be noted that both retransmissions and feedback decrease throughput by consuming time resources without delivering new processes.

Block error rate: The block error rate, BLER, is defined as the ratio of failed processes to the total transmitted processes, written as

$$\mathsf{BLER} = \frac{\mathsf{TB}_{\text{fail}}}{\mathsf{TB}_{\text{new}}}$$

Average Delay: The average delay, D, is the expected total time it takes for a process to be successfully received.[2]

In the MDS-HARQ scheme, when redundancy is transmitted in advance and no feedback is in place, the delay $\mathsf{D}_{\text{MDS-HARQ}}$ per process can be evaluated as

$$\mathsf{D}_{\text{MDS-HARQ}} = \bar{T}_{\text{MDS-HARQ}} \cdot t_{\text{TB}}.$$

In the HARQ feedback-enabled, the delay per process, denoted as $\mathsf{D}_{\text{HARQ-FB}}$, can be evaluated as

$$\mathsf{D}_{\text{HARQ-FB}} = \mathsf{t}_{\text{prop}} + t_{\text{FB}} + (\bar{T}_{\text{FB}} - 1) \cdot N_p \cdot (t_{\text{TB}} + t_{\text{FB}})$$

where $\bar{T}_{\text{FB}} - 1$ represents the average number of retransmissions required for successful decoding. If no retransmissions are needed, for example at high SNR, the average delay reduces to $\mathsf{t}_{\text{prop}} + t_{\text{TB}}$. Otherwise, the delay increases according to the number of retransmissions, each following the transmission of all the N_p processes and their corresponding feedback. The evaluation of $\bar{T}_{\text{FB}}$ is not straightforward, as in the MDS-HARQ case, and will therefore be performed via simulation.

Similarly, in the HARQ scheme without feedback, the delay per process, $\mathsf{D}_{\text{HARQ-woFB}}$, can be expressed as

$$\mathsf{D}_{\text{HARQ-woFB}} = \mathsf{t}_{\text{prop}} + t_{\text{FB}} + (\bar{T}_{\text{FB}} - 1) \cdot N_p \cdot t_{\text{TB}}$$

where $\bar{T}_{\text{woFB}}$ denotes the average number of transmissions combined before decoding. Note that in this scheme, no feedback time is included.

[2]The time required for physical-layer operations, including MRC and decoding, is neglected.

$$\bar{T}_{\text{MDS-HARQ}} = \mathbb{E}[T] = n \cdot \left(e^{-\frac{\gamma_{\text{th}}}{\Gamma}}\right)^n + \sum_{\tau=n+1}^{\infty} \tau \sum_{j=1}^{\tau-n} \cdot \binom{n}{n-j} \binom{\tau-n-1}{j-1} \left(e^{-\frac{\gamma_{\text{th}}}{\Gamma}}\right)^n \left(1-e^{-\frac{\gamma_{\text{th}}}{\Gamma}}\right)^{\tau-n}. \tag{5}$$

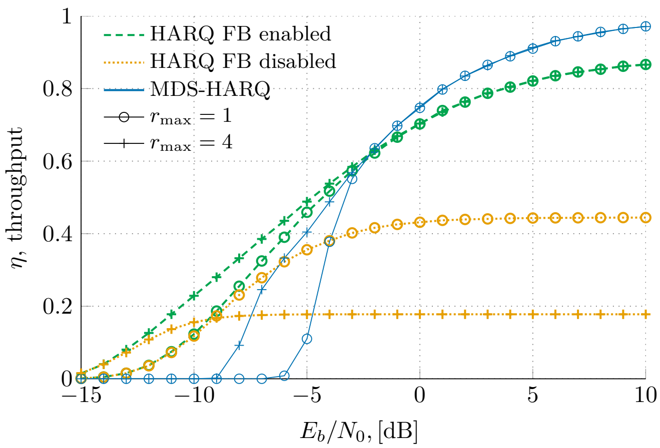

Fig. 2. Throughput as a function of E_b/N_0 for HARQ with feedback enabled, with feedback disabled and for MDS-HARQ for $r_{\max} = 1, 4$ retransmissions.

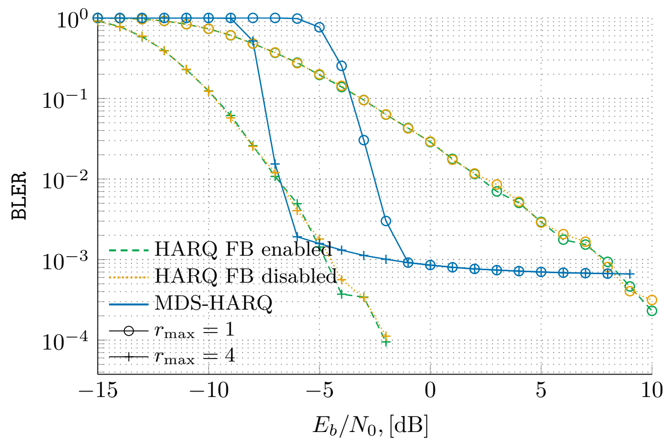

Fig. 3. Block error rate as a function of E_b/N_0 for HARQ with feedback enabled, with feedback disabled and for MDS-HARQ for $r_{\max} = 1, 4$ retransmissions.

Numerical Results

In our first set of results, we compare the throughput, block error rate, and delay for three schemes: HARQ with feedback enabled, HARQ with feedback disabled, and MDS-HARQ. In this setting, we assume that the MDS-HARQ encodes $n = 32$ processes into $N = n + r_{\max}$ encoded transport blocks, $\alpha = 32$ encoded blocks are transmitted, and feedback is sent after the reception of the last block.

For all the results, we consider unitary power $P = 1$ and a code rate of $\mathsf{R} = 1/3$. We set the feedback time to $t_{\text{FB}} = 0.1250$ ms, which is one-eighth of the TB duration $t_{\text{TB}} = 1$ ms, the $t_{\text{prop}} = 6$ ms and the visibility window to $\mathsf{W}_{\text{tx}} = 50 \times 10^4$ TBs, which defines the simulation duration.

In Fig. 2, the throughput η is shown as a function of E_b/N_0 for HARQ with feedback enabled, HARQ with feedback disabled, and MDS-HARQ, for $r_{\max} = 1$ and $r_{\max} = 4$ repetitions. Here, $\eta = 0$ indicates that none of the transmitted processes are successfully received, while $\eta = 1$ indicates the best performance, where the maximum transmittable processes in the visibility window W_{tx} are correctly received.

From the plot, we can identify two distinct performance regimes. At very low SNR, i.e., E_b/N_0 from -15 dB to -3 dB, traditional HARQ with feedback outperforms MDS-HARQ. This is because at low SNR, performing MRC over retransmissions is highly effective, allowing the receiver to accumulate enough signal energy to successfully decode the transmitted blocks. In the higher SNR regime, from -3 dB to 10 dB, MDS-HARQ shows superior throughput, as it reduces the multiple feedback rounds and ensures successful reception as soon as a sufficient number of encoded blocks are received, improving transmission efficiency. The performance of the HARQ-FB disabled is degraded because that time resources are wasted in blind retransmission, where decoding has already succeed and new processes cannot be allocated. Finally, we note that only at low SNR values ($E_b/N_0 < -3$ dB) does the difference between $r_{\max} = 1$ and $r_{\max} = 4$ significantly affect the performance of HARQ with feedback enabled and MDS-HARQ. Indeed, at higher SNRs, a single retransmission is sufficient to achieve high throughput.

In Fig. 3, the block error rate is shown as a function of E_b/N_0 for HARQ with feedback enabled, HARQ with feedback disabled, and MDS-HARQ, for $r_{\max} = 1, 4$ retransmissions. As expected, the BLER improves with increasing SNR. The plot shows that for a fixed $r_{\max}$, the performance of HARQ with feedback enabled and feedback disabled coincides. This is because both schemes achieve decoding once the accumulated energy from retransmissions is sufficient, which on average occurs after the same number of retransmissions, regardless of whether feedback is present or not. In contrast, MDS-HARQ with four retransmissions can achieve lower BLER at low SNR by allowing successful recovery once a sufficient number of encoded blocks are received.

In Fig. 4, the delay D is shown as a function of E_b/N_0 for HARQ with feedback enabled, HARQ with feedback disabled, and MDS-HARQ, for $r_{\max} = 1, 4$ retransmissions. In this case, we observe that HARQ, both with and without feedback, consistently outperforms the MDS-HARQ scheme. Furthermore, for SNR values above -5 dB, the performance is nearly identical across all schemes, since a single retransmission

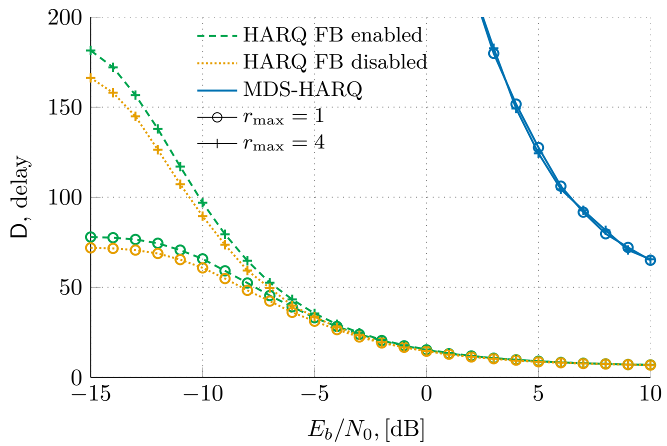

Fig. 4. Delay as a function of E_b/N_0 for HARQ with feedback enabled, with feedback disabled and for MDS-HARQ for $r_{\max} = 1, 4$ repetitions.

is sufficient and decoding occurs at the same time in both feedback-enabled and feedback-disabled cases.

MDS-HARQ is penalized in terms of delay because the receiver must wait to decode the entire group of n processes; in the case of retransmissions, an additional n TBs must be received before decoding can occur. To reduce this delay, we next propose alternative implementations of the MDS-HARQ scheme and evaluate their performance.

In Fig. 5, the throughput, block error rate, and delay are shown for three different implementations of the MDS-HARQ scheme. The MDS average scheme without feedback transmits, at each SNR, α the number of transport blocks transmitted is the average number of encoded blocks from (5), without considering any feedback. The MDS-HARQ with smart feedback scheme, instead of sending a simple NACK, specifies the exact number of additional encoded transport blocks needed for successful decoding, so that the next retransmission sends precisely this number of blocks. Finally, the MDS-HARQ with feedback scheme follows the standard approach described previously, where each NACK triggers the transmission of one additional encoded TB.

In this case, the results show that the smart feedback configuration outperforms the previous implementation in terms of BLER and delay, while achieving comparable throughput. Furthermore, when compared to traditional HARQ schemes with feedback enabled or disabled, the proposed smart FB implementation demonstrates superior performance in terms of BLER. Specifically, for SNR values above –3 dB with four retransmissions and above –1dB with a single retransmission, the MDS-HARQ with smart feedback achieves a negligible BLER, significantly outperforming conventional HARQ approaches, as illustrated in Fig. 3. This result validates the effectiveness of the smart feedback strategy, confirming its potential to enhance reliability while minimizing unnecessary retransmissions. The MDS-HARQ scheme, however, continues to suffer from higher delays due to the need to decode the entire group of blocks collectively.

Finally, we observe that implementing a transmission the average number of transmitted blocks without feedback is not convenient, as the variability of the channel can lead to either under- or over-transmission. In some cases, fewer blocks than necessary may be sent, resulting in decoding failures, while in others, more blocks than required are transmitted, increasing delay and reducing transmission efficiency. This highlights the advantage of using feedback-based, such as MDS-HARQ with smart feedback, which adjust the number of transmitted blocks according to actual channel conditions.

In our final results, we show the performance of the MDS-HARQ scheme with smart feedback for different numbers of encoded processes, namely $n = 8, 15, 24,$ and 32. As observed from the plot, at high SNR values, throughput is essentially independent of n. However, in the intermediate SNR range (approximately -10 to -3 dB), a slightly better performance is observed for $n = 8$ when only one retransmission is allowed. This improvement arises because the probability of successfully receiving all required blocks is higher when fewer blocks are transmitted, reducing the likelihood of retransmissions.

For the BLER performance analysis, different values of n exhibit similar behavior, with a slight gain observed for $n = 32$ when one retransmission is allowed at SNR values above -7 dB, and for four retransmissions at SNR values above -3 dB.

Our final results show that lower values of n lead to an improvement in delay. This is primarily because the probability of successful decoding increases when fewer blocks are involved.

In our final remarks, we highlight the potential of alternative HARQ process configurations and their benefits when some delay can be tolerated. Another advantage of MDS-HARQ is its scalability for broadcasting to multiple users: retransmissions do not need to be dedicated to specific users, as any additional encoded packet can be useful to all recipients. This makes MDS-HARQ a more efficient scheme for multi-user satellite communications.

As future work, the evaluation MDS-HARQ in a multi-user scenario is planned, exploring potential improvements in its implementation. The insights gained from this study have been valuable in understanding the proposed MDS-HARQ scheme, providing a solid foundation for these further investigations.

VI. Conclusions

This work investigated the performance of different HARQ protocols and proposed the MDS-HARQ in satellite communication scenarios. While conventional HARQ performs reliably at very low SNRs, MDS-HARQ significantly improves BLER and throughput by allowing the receiver to reconstruct data once enough encoded blocks are received, reducing dependence on multiple feedback rounds. The smart feedback mechanism further boosts performance by precisely indicating the additional blocks needed for decoding, lowering block error rates and counteracting delays. These results highlight MDS-HARQ's potential to enhance reliability in next-generation satellite communications.

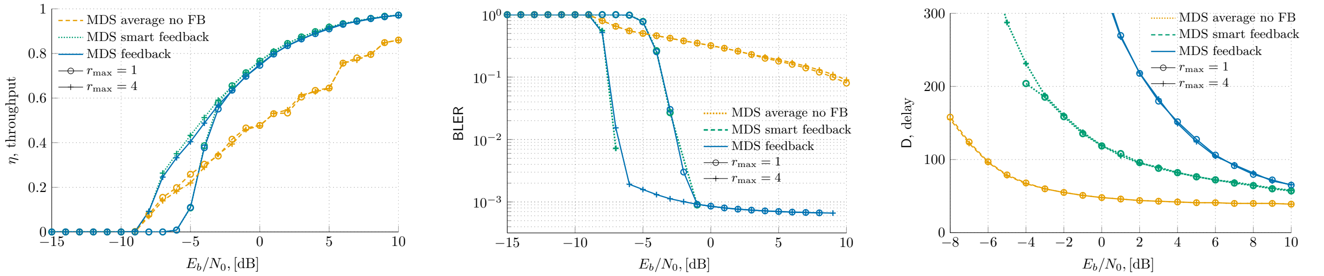

(a) MDS-HARQ throughput (b) MDS-HARQ block error rate (c) MDS-HARQ delay

Fig. 5. MDS-HARQ implementation comparision: (a) MDS-HARQ throughput, (b) feedback MDS-HARQ block error rate, and (c) MDS-HARQ delay.

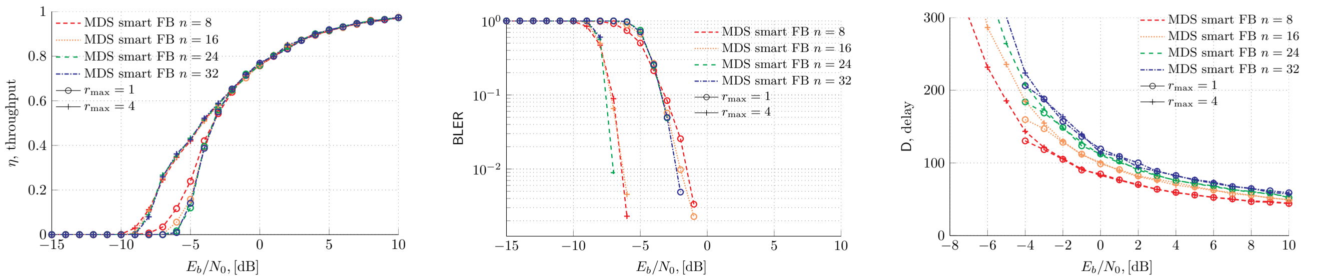

(a) Throughput (b) Block error rate (c) Delay

Fig. 6. MDS-HARQ with smart feedback: (a) throughput, (b) block error rate, and (c) delay.

ACKNOWLEDGEMENTS

This work has been funded by the 6G-NTN project, which received funding from the SNS JU under the European Union's Horizon Europe research and innovation program under Grant Agreement No 101096479. The views expressed are those of the authors and do not necessarily represent the project. The Commission is not liable for any use that may be made of any of the information contained therein.

REFERENCES

[1] A. Guidotti, A. Vanelli-Coralli, M. Conti, S. Andrenacci, S. Chatzinotas, N. Maturo, B. Evans, A. Awoseyila, A. Ugolini, T. Foggi, L. Gaudio, N. Alagha, and S. Cioni, "Architectures and key technical challenges for 5g systems incorporating satellites," *IEEE Transactions on Vehicular Technology*, vol. 68, no. 3, pp. 2624–2639, 2019.

[2] 3GPP, "Solutions for NR to support NTN," 3rd Generation Partnership Project (3GPP), Technical Specification TS 38.821, Jan. 2020, available: https://www.3gpp.org/DynaReport/38821.htm.

[3] S. Cioni, X. Lin, B. Chamaillard, M. E. Jaafari, G. Charbit, and L. Raschkowski, "Physical layer enhancements in 5G-NR for direct access via satellite systems," *International Journal of Satellite Communications and Networking*, vol. 41, no. 3, pp. 262–275, 2023.

[4] "Solutions for NR to support non-terrestrial networks (NTN)," 3GPP, Technical Report TR 38.821, Dec. 2019, https://www.3gpp.org/DynaReport/38821.htm.

[5] G. A. Medina-Acosta *et al.*, "3GPP Release-17 physical layer enhancements for LTE-M and NB-IoT," *IEEE Commun. Standards Mag.*, vol. 6, no. 4, pp. 80–86, Dec. 2022.

[6] R. Tuninato, G. Maiolini Capez, N. Mazzali, and R. Garello, "5G new radio for non-terrestrial networks: Analysis and comparison of HARQ and RLC ARQ performance over satellite links," *IEEE Access*, vol. 13, pp. 75 400–75 415, 2025.

[7] G. Prasad, V. R. Chandrika, L. Lampe, and G. Vos, "Enhanced hybrid automatic repeat request scheduling for non-terrestrial IoT networks," *IEEE Open Journal of the Communications Society*, vol. 4, pp. 2802–2814, 2023.

[8] A. Vanelli-Coralli, A. Guidotti, T. Foggi, G. Colavolpe, and G. Montorsi, "5G and beyond 5G non-terrestrial networks: trends and research challenges," in *2020 IEEE 3rd 5G World Forum (5GWF)*, 2020, pp. 163–169.

[9] H. Chen, R. G. Maunder, Y. Ma, R. Tafazolli, and L. Hanzo, "Hybrid-ARQ-aided short fountain codes designed for block-fading channels," *IEEE Transactions on Vehicular Technology*, vol. 64, no. 12, pp. 5701–5712, 2015.

[10] D. Sejdinovic, V. Ponnampalam, R. J. Piechocki, and A. Doufexi, "The throughput analysis of different IR-HARQ schemes based on fountain codes," in *2008 IEEE Wireless Communications and Networking Conference*, 2008, pp. 267–272.

[11] R. Singleton, "Maximum distance q-nary codes," *IEEE Transactions on Information Theory*, vol. 10, no. 2, pp. 116–118, 1964.