Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

CRC-Verified Adaptive Erasure Decoding for Reliable RS-Protected 16-QAM Image Transmission

Jenifer Darling Rosita Prince Devaraj

jeniferprincedevaraj.research@gmail.com

University of Wales Trinity Saint David - Swansea Campuses https://orcid.org/0009-0006-7265-462X

Stalin Jacob Wilson Packia Dhas

University of Wales Trinity Saint David https://orcid.org/0009-0003-1013-4475

Research Article

Keywords: Reed–Solomon codes, adaptive erasure decoding, CRC verification, reliability-guided decoding, 16-QAM, image transmission

Posted Date: September 15th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11037052/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Abstract

This paper presents a reliability-guided adaptive erasure-decoding method for improving the recovery of Reed–Solomon-protected images transmitted using 16-ary quadrature amplitude modulation. The proposed receiver combines conventional RS(255,223) decoding with block cyclic redundancy check (CRC) verification, nested erasure-set searches, rank-position substitution and a broader reliability-guided search. Reliability is estimated from the distance margin between the nearest and second-nearest constellation points, allowing additional processing to focus on byte positions most likely to contain errors. The method preserves the original RS code structure and applies computationally intensive recovery stages only to codewords that fail direct decoding or CRC verification.

The receiver was evaluated over an additive white Gaussian noise channel using 100 independently seeded trials at each of three measured signal-to-noise ratios. At 15.021 dB, direct RS decoding failed to recover any complete image, whereas the proposed decoder achieved a 92% bit-perfect image-recovery rate. At 15.340 dB, recovery increased from 71% with direct decoding to 99% with adaptive decoding. At 15.667 dB, the corresponding recovery rates were 99% and 100%. Stage-level analysis showed that the nested erasure-set search provided the largest adaptive contribution, while rank-position substitution and broader search recovered additional difficult codewords. Successfully reconstructed images had zero post-decoding bit-error rate, a structural similarity index of 1.000 and infinite peak signal-to-noise ratio. These results demonstrate that reliability-guided erasure processing can substantially improve complete-image recovery near the conventional RS correction limit.

Keywords— Reed–Solomon codes, adaptive erasure decoding, CRC verification, reliability-guided decoding, 16-QAM, image transmission.

1. Introduction

Digital image transmission requires reliable recovery of a large and structured collection of data. Corruption affecting only a small number of image blocks may prevent the complete image from being reconstructed correctly, even when most of the transmitted data have been recovered. This requirement becomes more demanding when higher order modulation is used because each transmitted symbol represents several bits and is more sensitive to displacement across decision boundaries under noisy channel conditions.

Reed–Solomon codes are nonbinary polynomial codes designed for symbol-level error correction [1]. Conventional hard-decision RS decoding, however, discards information concerning the reliability of the detected symbols. When corruption exceeds the conventional correction capability, one or more unresolved codewords may prevent bit-perfect reconstruction of the complete image.

This research develops a CRC-verified adaptive erasure decoder for an RS-protected 16-QAM image transmission system. Reliability is determined using the margin between the nearest and second-nearest squared constellation distances. The resulting measurements are combined at byte level and used to rank potentially erroneous positions within each received RS codeword. Direct decoding is attempted first, after which unresolved blocks are processed progressively through nested erasure-set search, rank-position substitution and broader search for difficult codewords. Every possible recovery must satisfy its block CRC, while a complete-image CRC verifies the final reconstructed output.

The principal contributions of this work are a practical byte reliability measure derived from received 16-QAM samples, a progressive recovery strategy that restricts additional processing to CRC rejected blocks, and the integration of block and complete-image CRC verification into the decoding process. The developed receiver is evaluated against direct RS decoding using independently seeded noise trials at three measured SNR conditions. Performance is assessed

using pre-decoding BER, bit-perfect image-recovery rate, confidence intervals and the contribution of each adaptive recovery stage.

2. Related Work and Research Gap

Reed–Solomon codes were introduced as a class of nonbinary block codes constructed using polynomials over finite fields [1]. Their symbol-based structure makes them effective against corruption affecting several bits within the same symbol. In conventional bounded distance decoding, an RS code containing $n - k$ parity symbols can correct up to $(n - k)/2$ erroneous symbols when the error locations are unknown. When the locations of uncertain symbols are supplied as erasures, the available parity can be used more efficiently. This error and erasure capability provides the foundation for reliability-guided RS decoding.

Generalized minimum distance decoding established that reliability ordered erasures could be used to generate a sequence of decoding attempts from a received word [2]. Channel measurement information was also used to construct alternative error patterns for block code decoding [3]. Later work provided faster generalized minimum distance processing for Reed–Solomon and algebraic geometry codes [11]. Collectively, these studies demonstrated that positions associated with uncertain channel observations can be treated differently from positions supported by stronger decisions.

Algebraic list decoding extended the recovery radius of Reed–Solomon codes beyond the conventional unique decoding limit [5]. Algebraic soft-decision decoding subsequently incorporated probabilistic reliability information by assigning different multiplicities to possible received symbols [4]. Although these methods can provide substantial coding gains, polynomial interpolation, factorisation and multiplicity assignment may introduce considerable computational complexity. Their generality can also be unnecessary when only a small number of codewords within a much larger image frame require additional recovery.

Iterative approaches have also been developed to improve practical RS decoding. One method adapts the binary parity check representation of an RS code according to received bit reliability and applies iterative soft input soft output processing [6]. Another combines belief propagation with algebraic soft decision list decoding to improve performance while reducing complexity relative to a direct algebraic implementation [7]. Algebraic Chase decoding generates test patterns from unreliable positions and applies algebraic interpolation-based decoding to them [8]. Reduced density parity check representations have additionally been used to support practical iterative soft decoding with lower implementation complexity [9].

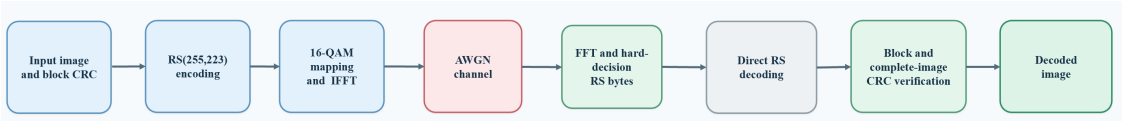
These established approaches demonstrate the benefit of reliability information, but they generally assess decoder performance at codeword or bit level. Complete image reconstruction imposes a stricter requirement because every required image block must be recovered correctly before the output can be considered bit-perfect. Moreover, an algebraically valid RS codeword is not necessarily the originally transmitted message when decoding is attempted beyond the conventional correction radius. An additional integrity test is therefore important when several possible recovery results are examined. CRC methods provide a mechanism for detecting inconsistencies between recovered data and transmitted integrity information [17].

The identified gap is the need for a selective reliability-assisted RS recovery architecture designed specifically for bit-perfect reconstruction of a complete image. Such a receiver must restrict additional processing to unresolved codewords and independently verify each recovered result before including it in the reconstructed output. The present work addresses this requirement by integrating constellation derived byte reliability, progressive RS error-and-erasure decoding, block CRC verification and complete image integrity verification.

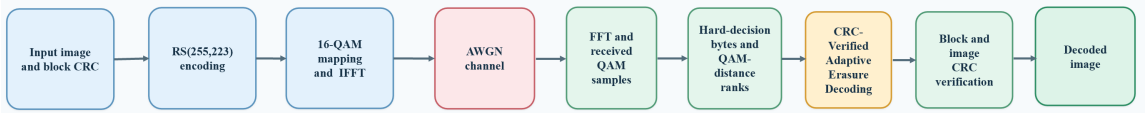
3. System Model and Transmission Architecture

3.1 System Overview

The communication system was developed to transmit a digital colour image through a noisy 16-QAM channel and to compare conventional direct Reed–Solomon decoding with the proposed CRC-verified adaptive erasure decoding method. As illustrated in Figure 1, both receiver configurations use the same image source, framing structure, RS encoder, modulation process and channel model. Using a common transmission path ensures that the observed differences in image recovery are attributable to the decoding method.



(a) Reference receiver



(b) Developed receiver

Figure 1. RS-protected 16-QAM image-transmission system: (a) reference receiver using direct RS decoding and (b) developed receiver using CRC-verified adaptive erasure decoding.

Figure 1 presents the common transmission path and the two receiver configurations. The reference receiver performs direct RS decoding, whereas the developed receiver retains channel reliability information and applies additional processing only to unresolved blocks. The individual transmitter and receiver stages are described in the following subsections.

3.2 Image Framing and RS Encoding

The original colour image contains 135,482 bytes. Before transmission, the image-data are divided into 619 blocks, with each block containing 219 image bytes. A four-byte CRC is appended to every image block, producing a 223-byte message for Reed–

Solomon encoding. Both the block-level and complete-image integrity checks used CRC-32 as implemented by the Python zlib.crc32() function. CRC values were stored in four-byte big-endian format. The block CRC provides an independent integrity check for determining whether a decoded block is consistent with the transmitted data. Any final image block that contains fewer than 219 image bytes is completed using padding before its CRC is calculated.

Each 223-byte message is encoded using an RS(255,223) code implemented over GF(2^8) using the Python reedsolo library with its default primitive polynomial, 0x11D. The encoder adds 32 parity bytes and produces a 255-byte codeword. This configuration can correct up to 16 erroneous bytes when their positions are unknown. Supplying potentially unreliable positions as erasures allows the parity symbols to support error-and-erasure decoding beyond the conventional error-only limit.

The complete protected frame contains one header codeword, 619 image-data codewords and four filler codewords, giving a total of 624 RS codewords. The header carries the information required to interpret and reconstruct the received frame, including the image length and integrity information. Filler codewords complete the transmitted frame structure and are excluded from the reconstructed image. The resulting protected frame contains 159,120 encoded bytes. Because each byte is represented by two 16-QAM symbols, the complete frame produces 318,240 transmitted QAM symbols. The nominal RS code rate is 223/255, or 0.8745, while the effective image-data rate, after accounting for the four-byte block CRC, is 219/255, or 0.8588.

3.3 16-QAM Modulation and Channel Model

Following RS encoding, each encoded byte is separated into two groups of four bits. Each four-bit group is mapped to one complex 16-QAM symbol. The implemented constellation uses in-phase and quadrature amplitude levels of -3, -1, +1 and +3. Therefore, the set of possible constellation symbols is expressed as

$$S = \{a + jb \mid a, b \in \{-3, -1, +1, +3\}\}$$

where a and b represent the in-phase and quadrature components, respectively. The two most significant bits of each four-bit group determine the in-phase component, while the two least significant bits determine the quadrature component. Gray mapping was applied independently to both components using 00 → −3, 01 → −1, 11 → +1 and 10 → +3. The constellation was used without unit-energy normalisation. For equally likely symbols, the average constellation-symbol energy is $E_s = 10$ in the adopted amplitude units. Since each 16-QAM symbol represents four encoded bits, every RS codeword of 255 bytes produces 510 QAM symbols.

The mapped QAM symbols are arranged into blocks and processed using a 16-point inverse fast Fourier transform to generate the time-domain multicarrier signal. At the receiver, the corresponding 16 point fast Fourier transform is applied to recover the complex QAM samples. The transmitter and receiver use discrete Fourier transform processing to implement frequency-division multiplexing through modulation and demodulation [10]. In the present system, the 16-QAM symbols are transmitted using orthogonal subcarriers. Hard decisions are then obtained by assigning each received sample to its nearest ideal constellation point.

Transmission is modelled using an additive white Gaussian noise channel. The received time domain signal can be represented as

$$r[n] = x[n] + w[n]$$

where $r[n]$ is the received sample, $x[n]$ is the transmitted sample and $w[n]$ represents the added Gaussian noise. The GNU Radio Channel Model noise-voltage parameter was used to specify the standard deviation, σ_w, of the generated Gaussian noise. Thus, the corresponding noise variance is σ_w^2. The configured values are simulation amplitude parameters rather than SNR values; the resulting SNR was measured separately as described in Section 5.2. Three values of the noise standard deviation, σ_w, were investigated: 0.130, 0.135 and 0.140. These settings produced measured SNR values of 15.667 dB, 15.340 dB and 15.021 dB, respectively. Increasing the noise amplitude caused greater displacement of the received QAM samples from their ideal constellation points and increased the probability of incorrect symbol and byte decisions. The received complex samples are retained by the developed receiver so that the margin between the nearest and second-nearest squared constellation distances can subsequently be used to measure decision reliability.

3.4 Reference Receiver

The reference receiver represents the conventional hard-decision decoding path shown in Figure 1(a). After the fast Fourier transform, each received complex sample is assigned to its nearest ideal 16-QAM constellation point. The resulting four-bit decisions are combined in pairs to reconstruct the received bytes. These bytes are then arranged into 255-byte codewords according to the transmitted frame structure.

Each received codeword is processed directly by the RS(255,223) decoder without using channel reliability information. Under conventional error-only decoding, a codeword containing no more than 16 erroneous bytes can be corrected. Codewords exceeding this limit may remain unresolved.

Following RS decoding, the four-byte CRC contained in each recovered message is compared with a newly calculated CRC for the corresponding 219 image bytes. A block

is accepted only when these values agree. After all accepted image blocks have been arranged in their original order, filler data are removed and the image is reconstructed using the length information contained in the decoded header. The complete-image CRC is then used to determine whether the reconstructed image is bit-perfect. This receiver provides the reference performance against which the developed CRC-verified adaptive erasure decoder is evaluated.

3.5 Developed Receiver

The developed receiver follows the adaptive decoding path presented in Figure 1(b). It uses the same received 16-QAM samples, frame structure and RS(255,223) codewords as the reference receiver. Unlike the reference receiver, it retains the complex samples obtained after the fast Fourier transform. These samples provide the reliability information used to rank potentially erroneous byte positions. Direct RS decoding is attempted first and only blocks that fail CRC verification are passed to the adaptive decoder described in Section 4.

The recovery process uses only the received QAM samples, received hard-decision bytes, RS code structure and CRC information carried within the transmitted frame. The original transmitted image is used only for experimental performance assessment and is not available to the decoder during recovery.

4. Proposed CRC-Verified Adaptive Erasure Decoder

The proposed decoder applies a progressive recovery procedure to codewords that cannot be accepted through direct RS decoding. As illustrated in Figure 2, processing begins with direct decoding and continues through nested erasure-set search, rank-position substitution and a broader search for difficult codewords.

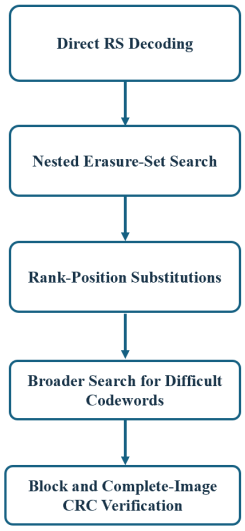


Figure 2. Progressive processing stages of the CRC-verified adaptive erasure decoder.

Each possible recovery is checked using its block CRC before acceptance. After all required blocks have been processed, the reconstructed image is verified using the complete-image CRC.

4.1 Reliability Measurement and Byte Ranking

The developed decoder obtains reliability information directly from the received complex QAM samples. Using channel-derived reliability information to guide erasure-position selection and error-and-erasure decoding is consistent with established Chase and GMD decoding principles [3], [11]. Before reliability calculation, the received samples are equalised using an estimated complex channel gain. For decoder operation, the complex channel gain was estimated using the received samples and their corresponding hard-decision constellation points:

$$\hat{g} = \frac{[\Sigma_k\, y_k(\tilde{s}_k)\, *]}{[\Sigma_k\, |\tilde{s}_k|^2]}$$

where $\tilde{s}_k$ is the hard-decision constellation point associated with the received sample y_k. This estimation uses only receiver-generated decisions and does not require knowledge of the transmitted symbols.

The equalised sample is expressed as

$$z_k = \frac{y_k}{\hat{g}},$$

where y_k is the received QAM sample, $\hat{g}$ is the estimated complex channel gain and z_k is the equalised sample.

For each equalised sample, the squared Euclidean distances to all ideal 16-QAM constellation points are calculated. The nearest and second-nearest squared distances are denoted by $d_{1,k}$ and $d_{2,k}$, respectively. The decision margin is then calculated as

$$\Delta_k = d_{2,k} - d_{1,k}.$$

A large margin indicates that the selected constellation point is clearly separated from the next possible decision. In contrast, a small margin indicates that the received sample lies close to a decision boundary and is therefore less reliable. The unreliability value assigned to the kth QAM symbol is defined as

$$u_k = \frac{1}{\Delta_k + \varepsilon},$$

where $\varepsilon = 10^{-12}$ is a small positive constant used to prevent division by zero. Consequently, a smaller decision margin produces a larger unreliability value.

Each encoded byte is represented by two QAM symbols. The unreliability of byte position $i$is therefore defined as the larger of the two associated symbol unreliability values:

$$U_i = \max(u_{2i}, u_{2i+1}).$$

Using the maximum value ensures that a byte is treated as potentially unreliable when either of its two four-bit components has an uncertain constellation decision. The 255-byte positions within each received RS codeword are arranged in descending order of U_i. This ranked list places the least reliable byte positions first and is subsequently used to select erasure positions during adaptive decoding.

4.2 Direct RS Decoding and CRC Verification

Each received 255-byte codeword is initially processed using conventional RS (255,223) decoding without supplying erasure positions. This first stage recovers codewords that remain within the conventional error correction capability without requiring any additional reliability-guided search. If the RS decoder does not produce a decoded message, the codeword is marked as unresolved and passed to the nested erasure-set search described in Section 4.3.

For each successfully decoded image-data codeword, the resulting 223-byte message is separated into 219 image bytes and the associated four-byte CRC. Let $\widehat{\mathbf{m}}_i$ denote the decoded message for codeword i, expressed as

$$\widehat{\mathbf{m}}_i = [\widehat{\mathbf{p}}_i \parallel \hat{\mathbf{c}}_i],$$

where $\widehat{\mathbf{p}}_i$ contains the 219 recovered image bytes, $\hat{\mathbf{c}}_i$ is the recovered four-byte CRC and $\parallel$ denotes concatenation. A new CRC is calculated from $\widehat{\mathbf{p}}_i$ and compared with $\hat{\mathbf{c}}_i$. The acceptance decision is defined as

$$A_i = \begin{cases} 1, & \text{CRC } (\widehat{\mathbf{p}}_i) = \hat{\mathbf{c}}_i, \\ 0, & \text{otherwise.} \end{cases}$$

A block with $A_i = 1$ is accepted immediately and is not subjected to further processing. A block with $A_i = 0$, or one for which direct RS decoding fails, remains unresolved and proceeds to the adaptive recovery stages. This verification prevents an algebraically valid but incorrect decoded message from being incorporated into the reconstructed image. The transmitted image itself is not used in this decision.

4.3 Nested erasure-set search

When direct RS decoding and CRC verification do not recover a codeword, the ranked byte positions obtained in Section 4.1 are used to construct a sequence of nested erasure sets. Let

$$\boldsymbol{\pi}_i = \left(\pi_{i,1}, \pi_{i,2}, \ldots, \pi_{i,255}\right)$$

denote the byte positions of codeword i, arranged from the highest to the lowest unreliability. For an erasure count s, the trial erasure set is defined as

$$\mathcal{E}_i(s) = \left\{\pi_{i,1}, \pi_{i,2}, \ldots, \pi_{i,s}\right\}, 1 \leq s \leq 32.$$

The sets are nested because each new decoding attempt retains all previously selected positions and adds the next position from the reliability ranking. Therefore,

$$\mathcal{E}_i(1) \subset \mathcal{E}_i(2) \subset \cdots \subset \mathcal{E}_i(32).$$

For each value of s, the RS decoder treats the positions in $\mathcal{E}_i(s)$ as erasures. The maximum of 32 selected positions corresponds to the 32 parity bytes in the RS(255,223) code. If e erroneous bytes with unknown locations remain outside the erasure set, successful error-and-erasure decoding requires

$$2e + s \leq 32.$$

Every message produced by an erasure decoding attempt is subjected to the block CRC test defined in Section 4.2. The search stops immediately when a CRC-valid message is obtained, and the recovered block is accepted. If none of the 32 nested erasure sets produces a valid block, the codeword proceeds to the rank-position substitution stage described in Section 4.4.

4.4 Rank-position substitution

The nested search assumes that the first s positions in the reliability ranking form the most suitable erasure set. However, positions close to the selection boundary may have similar unreliability values, and the correct error location may lie immediately outside the selected set. The rank-position substitution stage therefore generates alternative erasure sets by replacing one boundary position with one nearby unselected position.

For each erasure count s, where $2 \leq s \leq 32$, the initial set is $\mathcal{E}_i(s)$, as defined in Section 4.3. The boundary set $\mathcal{B}_i(s)$ contains up to the final eight positions within $\mathcal{E}_i(s)$, while the outside set $\mathcal{O}_i(s)$ contains up to the next ten positions in the reliability ranking that were not included in the original set. A substituted erasure set is constructed as

$$\mathcal{E}_i'(s; b, o) = (\mathcal{E}_i(s) \setminus \{b\}) \cup \{o\},$$

where $b \in \mathcal{B}_i(s)$ is a selected boundary position and $o \in \mathcal{O}_i(s)$ is a nearby unselected position. This substitution preserves the original erasure count while examining alternatives around the reliability boundary.

Each possible single substitution between the two sets is tested using RS error-and-erasure decoding. Repeated erasure sets are excluded, and every decoded result is subjected to

block CRC verification. The search stops when a CRC-valid message is obtained. If no single substitution produces a valid block for any erasure count, the codeword proceeds to the broader search described in Section 4.5.

4.5 Broader Search and Final Image Verification

Codewords that remain unresolved after rank-position substitution are passed to a broader reliability-guided search. For each erasure count s, where $4 \leq s \leq 32$, the initial erasure set is formed from the first s positions in the reliability ranking. The search considers up to the final 16 positions within this set and up to the next 48 unselected positions in the ranking. Between one and six selected positions are replaced by the same number of unselected positions, thereby preserving the erasure count.

Let $R \subseteq \mathcal{E}_i(s)$ denote the positions removed from the initial erasure set and let A denote the positions added from the unselected pool. A broader trial erasure set is defined as

$$\mathcal{E}_i^{(q)}(s; R, A) = (\mathcal{E}_i(s) \setminus R) \cup A,$$

where

$$| R | = | A | = q, 1 \leq q \leq 6.$$

For each codeword, the Python pseudorandom-number generator was initialised using an integer seed obtained by adding 1,048,576 times the codeword index, the integer part of 1,000,000 times the noise-amplitude setting, and the trial seed. A maximum of 100,000 trial erasure sets is examined for each unresolved codeword.

Repeated erasure sets are excluded. Each trial erasure set is processed using RS error-and-erasure decoding, and every resulting message is subjected to the block CRC test defined in Section 4.2. The search terminates immediately when a CRC-valid message is obtained.

If the broader search does not produce a valid block, a final two-position boundary substitution is attempted. For each erasure count from 4 to 32, two positions are selected from the final four positions of the initial erasure set and replaced by two positions selected from the next six unselected positions in the reliability ranking. Each resulting decoded message is accepted only when its block CRC is valid. A codeword that remains unresolved after this stage causes the corresponding image recovery trial to be classified as unsuccessful.

After all required image blocks have been recovered, the accepted blocks are arranged in their original order. Filler data are removed, and the image is reconstructed using the length information contained in the decoded header. A CRC is then calculated over the complete reconstructed image and compared with the complete-image CRC carried in the transmitted frame. The reconstructed image is accepted as CRC-valid when these CRC values agree. For experimental performance assessment only, bit-perfect recovery is confirmed through byte-for-byte comparison with the original image, which is not available to the decoder during any recovery stage.

5. Experimental Setup

5.1 Implementation and Test Configuration

The experimental system was implemented using GNU Radio 3.8 and Python. A fixed digital colour image containing 135,482 bytes was transmitted in every trial using the framing, RS coding and 16-QAM modulation process described in Section 3. The

reference receiver and the developed receiver were supplied with data obtained from the same transmitted frame and noise realization. Consequently, differences in recovery performance were attributable to the decoding strategy rather than to different channel conditions.

The reference receiver used direct RS decoding followed by CRC verification. The developed receiver used the complete progressive procedure described in Section 4. The reliability ranking, nested erasure-set search, rank-position substitution and broader search parameters were held constant throughout all trials.

5.2 Noise Conditions and Trial Procedure

Three additive noise amplitude settings were investigated: 0.130, 0.135 and 0.140. These settings produced measured SNR values of 15.667 dB, 15.340 dB and 15.021 dB, respectively. For each trial, the SNR was measured using the transmitted QAM symbols and the corresponding received post-FFT complex samples. First, the least-squares complex channel gain was estimated as

$$\hat{g} = \frac{\sum_k s_k^* y_k}{\sum_k |s_k|^2}$$

where s_k is the transmitted QAM symbol, y_k is the corresponding received post-FFT sample, and $(\cdot)^*$ denotes complex conjugation. The fitted signal was defined as $\hat{g}s_k$, and the residual error was calculated as $e_k = y_k - \hat{g}s_k$. The measured SNR was then calculated as.

$$SNR_{dB} = 10log_{10}\left(\frac{mean(|\hat{g}s_k|^2)}{mean(|e_k|^2)}\right)$$

The SNR reported for each noise-amplitude setting is the mean of the values measured across 100 trials using noise seeds 1–100. The same set of seeds was used for all three noise settings, giving a total of 300 trials. The transmitted symbols were used only for experimental SNR measurement and were not supplied to the adaptive decoder.

For every trial, a new noise sequence was generated and applied to the transmitted signal. The same received QAM samples and hard-decision bytes were then used to evaluate both receiver configurations. The developed decoder was not provided with the applied noise amplitude or the original transmitted image during recovery. The original encoded data and image were retained separately only for calculating experimental performance measures.

5.3 Performance Measures

Channel quality before RS decoding was assessed using the pre-decoding bit-error rate. This value was calculated by comparing the received hard-decision encoded bits with the corresponding transmitted encoded bits. It was used only as an experimental channel-performance measure and was not supplied to either decoder.

The principal performance measure was the bit-perfect image-recovery rate, defined as the percentage of trials in which the complete reconstructed image passed the transmitted complete-image CRC. Recovery rates were calculated separately for direct RS decoding and the CRC-verified adaptive erasure decoder. Ninety-five per cent confidence intervals for the observed recovery proportions were calculated using the Wilson score method, with $z = 1.96$ and $n = 100$ trials for each noise condition.

For the developed receiver, the number of blocks recovered by each processing stage was recorded. The reported stages were direct RS decoding, nested erasure-set search, rank-position substitution and broader search. Failed or unresolved blocks were also recorded. These measurements were used to determine the contribution of each stage to complete image recovery and the additional processing required beyond conventional decoding.

6. Results

The experimental results are presented in terms of pre-decoding channel BER, bit-perfect image-recovery rate and the contribution of each adaptive decoding stage. Results were obtained from 100 independently seeded trials at each of the three measured SNR conditions. The conventional reference and developed receivers were evaluated using identical received data in every trial.

6.1 Pre-Decoding Channel Performance

The pre-decoding BER was measured before RS decoding to quantify the level of corruption introduced by the noisy channel. Figure 3 presents the mean pre-decoding BER obtained at each measured SNR. The transmitted encoded bits were used only to calculate this experimental performance measure and were not available to either decoder during recovery.

The mean pre-decoding BER decreased consistently as the measured SNR increased. The most demanding condition, at 15.021 dB, produced a mean BER of approximately 0.0043. This decreased to approximately 0.0033 at 15.340 dB and 0.0025 at 15.667 dB. The monotonic reduction confirms that the three selected conditions represent progressively improving channel quality.

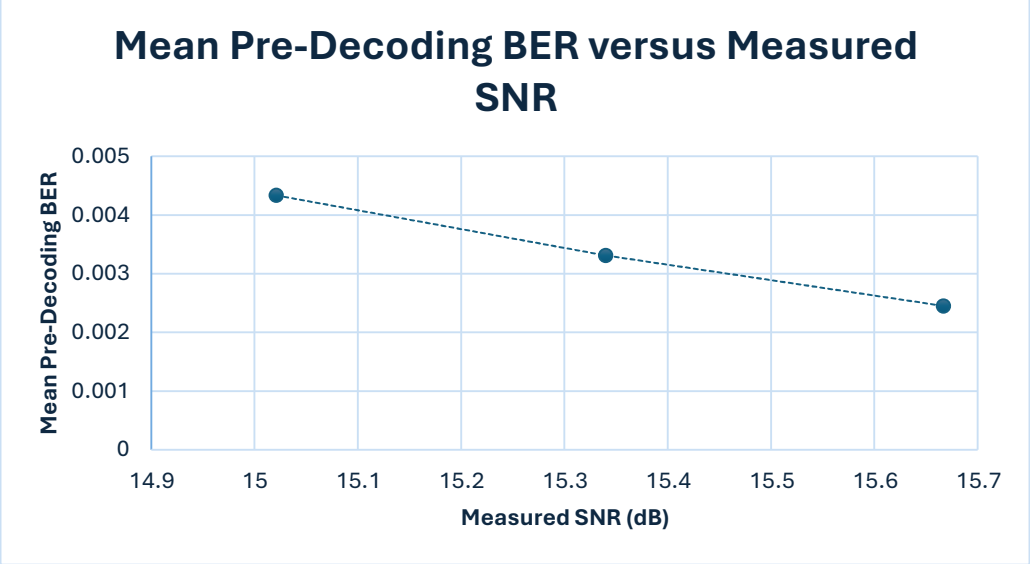


Figure 3. Mean pre-decoding BER as a function of measured SNR for the received 16-QAM signal.

Although the measured BER values were relatively low, errors were distributed across the complete protected image frame. Consequently, a small number of RS codewords could contain more than the 16 erroneous bytes correctable by conventional RS(255,223) error-only decoding. The pre-decoding BER therefore provides the channel-level context for the complete-image recovery results presented in Section 6.2.

6.2 Bit-Perfect Image-Recovery Performance

A recovery trial was classified as successful only when every required image block was recovered and the CRC of the complete reconstructed image agreed with the transmitted complete-image CRC. Figure 4 compares the bit-perfect image-recovery rates achieved by direct RS decoding and the proposed CRC-verified adaptive erasure decoder.

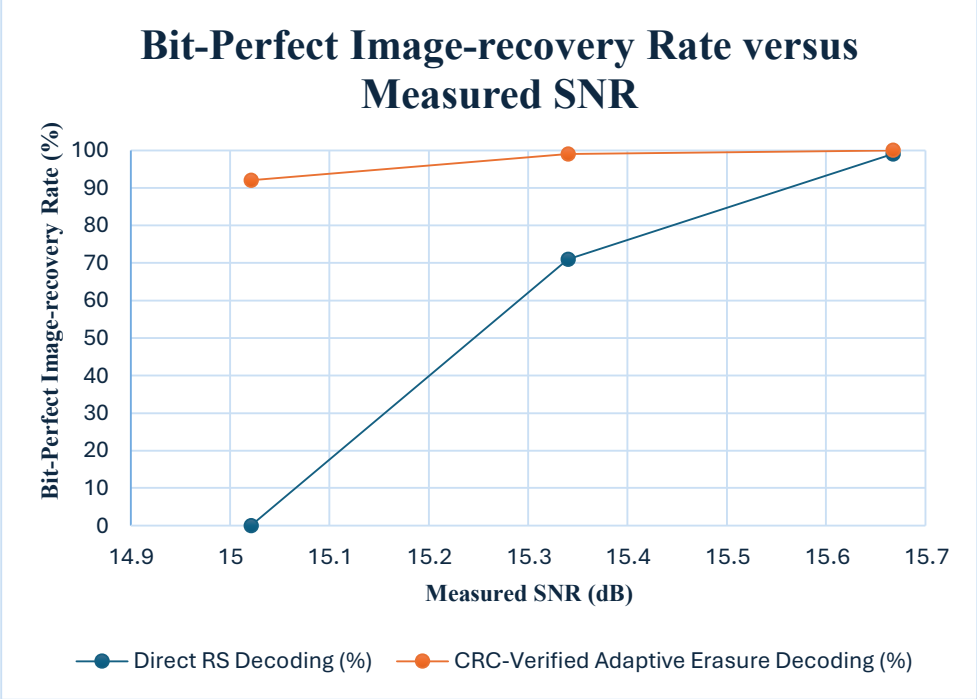


Figure 4. Bit-perfect image-recovery rates obtained using direct RS decoding and CRC-verified adaptive erasure decoding at different measured SNR values.

The connecting lines are included only as visual guides between the measured operating points and do not represent fitted trends. The greatest improvement occurred at the lowest measured SNR of 15.021 dB. Direct RS decoding did not recover any complete image, whereas the proposed decoder achieved a bit-perfect recovery rate of 92%. This represents an improvement of 92 percentage points and demonstrates that the adaptive stages recovered codewords that exceeded the conventional error-only correction capability.

At 15.340 dB, the recovery rate increased from 71% with direct RS decoding to 99% with the proposed decoder, corresponding to an improvement of 28 percentage points. At the highest measured SNR of 15.667 dB, the two receivers achieved recovery rates of 99% and 100%, respectively. The smaller difference at this condition indicates that conventional decoding was already sufficient for nearly all trials when the channel quality was higher.

Table 1. Bit-perfect image-recovery rates with 95% confidence intervals.

Measured SNR (dB)	Direct RS Decoding, % (95% CI)	CRC-Verified Adaptive Erasure Decoding, % (95% CI)
15.021	0 (0.00–3.70)	92 (85.00–95.89)
15.340	71 (61.46–78.99)	99 (94.55–99.82)
15.667	99 (94.55–99.82)	100 (96.30–100.00)

Table 1 presents the recovery proportions together with their 95% confidence intervals. The intervals show a clear separation between the two decoding methods at 15.021 dB and 15.340 dB. At 15.667 dB, the intervals overlap because both receivers approached complete recovery. These results demonstrate that the proposed method provides its largest benefit close to and beyond the conventional RS decoding threshold, while introducing little observable difference when direct decoding is already highly successful.

6.3 Contributions of the Adaptive Recovery Stages

To determine how the overall improvement was achieved, each image-data codeword was assigned to the first decoding stage that produced a CRC-valid block. The recorded categories were direct RS decoding, nested erasure-set search, rank-position substitution, broader search and unresolved failure. The mean number of codewords in each category was calculated over the 100 trials performed at each measured SNR.

Table 2. Mean image-data codeword outcomes over 100 trials at each measured SNR.

Measured SNR (dB)	Directly Decoded CRC-Valid Blocks	Nested Erasure-Set Recovery	Rank-Position Substitution Recovery	Broader Search Recovery	Unrecovered Blocks
15.667	618.99	0.01	0.00	0.00	0.00

15.340	618.67	0.25	0.05	0.02	0.01
15.021	614.50	3.25	0.74	0.43	0.08

The five outcome columns sum to 619 image-data codewords per trial; the header and filler codewords are excluded. At 15.021 dB, direct RS decoding recovered a mean of 614.50 of the 619 image-data codewords per trial. The nested erasure-set, rank-position substitution and broader search stages recovered an additional mean of 3.25, 0.74 and 0.43 codewords, respectively, leaving only 0.08 codewords unresolved. At 15.340 dB, the adaptive stages recovered a combined mean of 0.32 additional codewords per trial, while at 15.667 dB only 0.01 codeword per trial required adaptive recovery.

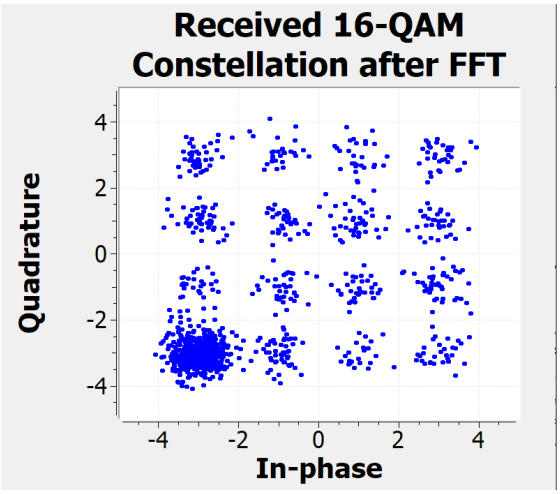
Table 2 shows that direct RS decoding recovered the majority of image-data codewords under all three channel conditions. At the highest measured SNR, almost every codeword was recovered directly, and the adaptive stages were required only occasionally. As the measured SNR decreased, a greater number of codewords exceeded the conventional error-only correction capability and therefore required additional processing.

The nested erasure-set search provided the largest contribution among the adaptive stages because the least reliable byte positions frequently corresponded to actual error locations. Rank-position substitution recovered additional codewords when an erroneous byte lay close to, but outside, the original erasure-set boundary. The broader search was used least frequently; however, its contribution was important because the failure of even one required codeword prevented bit-perfect reconstruction of the complete image.

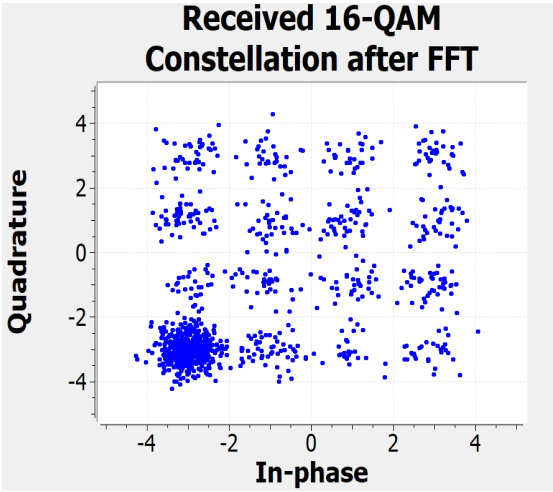
Unresolved codewords occurred primarily at the most demanding channel condition. This is consistent with the 92% complete-image-recovery rate obtained at 15.021 dB. The progressive distribution of recovery effort demonstrates that computationally intensive searches were restricted to a small number of difficult codewords, while most blocks were accepted after conventional direct decoding.

6.4 Received Constellations

The received 16-QAM constellations were examined to illustrate the effect of channel noise on symbol decisions. Figure 5 compares the received symbol distributions at the lowest and highest investigated noise amplitudes. These conditions correspond to measured SNR values of 15.667 dB and 15.021 dB, respectively.



(a) Noise amplitude = 0.130; measured SNR 15.667 dB



(b) Noise amplitude = 0.140; measured SNR 15.021 dB

Figure 5. Received 16-QAM symbol distributions at noise amplitudes of (a) 0.130 and (b) 0.140, corresponding to measured SNR values of 15.667 and 15.021 dB, respectively.

At a noise amplitude of 0.130, the received samples remained more tightly concentrated around the ideal 16-QAM constellation points. At a noise amplitude of 0.140, the distributions became more dispersed, and a greater number of samples approached or crossed the decision boundaries between neighbouring constellation points. This increased dispersion is consistent with the higher pre-decoding BER observed at 15.021 dB.

The constellation plots also illustrate the information used by the proposed reliability measure. Samples positioned close to a decision boundary have a smaller distance margin between their nearest and second-nearest constellation points and are therefore

assigned greater unreliability. Ranking the corresponding byte positions allows the adaptive decoder to direct additional processing towards decisions most likely to be erroneous.

6.5 Representative Image-Recovery Outcomes

Figure 6 presents representative image outcomes at the most demanding measured SNR of 15.021 dB. Panel (a) shows the original transmitted image used only for experimental comparison. Panel (b) shows the bit-perfect image recovered by the CRC-verified adaptive erasure decoder for seed 1. Panel (c) shows an unsuccessful adaptive recovery trial for seed 10, in which one or more required codewords remained unresolved.



(a) Original transmitted image (b) Bit-perfect recovery (c) CRC-invalid corrupted output

Figure 6. Representative image-transmission outcomes at a noise amplitude of 0.140 and a measured SNR of 15.021 dB: (a) original transmitted image, (b) bit-perfect image recovered using CRC-verified adaptive erasure decoding for seed 1, and (c) CRC-invalid corrupted output for seed 10.

For visualisation purposes, unresolved image blocks were replaced with zero-valued bytes. The resulting output failed the complete-image CRC and was therefore classified as an unsuccessful recovery. The successfully recovered image shown in Figure 6 (b) was identical to the original at byte level and satisfied the complete-image CRC. Consequently, its post-decoding BER was zero, its structural similarity index was 1.000 and its peak signal-to-noise ratio was infinite. These values confirm that the method achieved exact data recovery rather than merely producing a visually improved approximation.

The comparison demonstrates the strict nature of complete-image recovery. In the successful trial, every required codeword was recovered and the complete-image CRC was valid. In the failed trial, the inability to recover one or more required codewords prevented verified reconstruction of the complete image. These examples complement the aggregate 92% recovery rate by showing both successful and unsuccessful outcomes under different noise realizations at the same measured SNR condition.

7. Discussion

The results demonstrate that reliability-guided erasure processing can substantially improve complete-image recovery when conventional Reed–Solomon decoding operates near its correction limit. At 15.021 dB, direct RS decoding failed to recover any complete image, whereas the proposed decoder achieved a 92% bit-perfect recovery rate. The improvement decreased as the SNR increased: recovery rose from 71% to 99% at 15.340 dB, while both receivers approached complete recovery at 15.667 dB. The proposed method therefore provides its greatest benefit when only a small number of codewords exceed the conventional error-only correction capability.

The observed improvement agrees with established reliability-assisted decoding principles. Generalized minimum-distance decoding uses reliability-ordered erasures [2], [11], while Chase decoding constructs alternative test patterns using reliability information from the received positions [3]. Algebraic list, soft-decision, iterative and Chase-type methods have also demonstrated that reliability information can extend RS decoding beyond conventional hard-decision processing [4]–[9]. Progressive algebraic soft-decision approaches prioritise promising decoding attempts to reduce unnecessary processing [12], [13], while related Chase-type methods employ structured or lower-complexity test-pattern evaluation [14], [15]. The present method follows this principle by ranking byte positions according to their constellation-distance margins and applying additional processing only to unresolved codewords.

Direct numerical comparison with previous studies is difficult because most report coding gain, BER, codeword-error rate or frame-error rate using different RS codes, modulation schemes, channel models and stopping conditions [4]–[9], [12]–[15]. In this study, success required the correct recovery of all 619 image-data codewords and validation of the complete-image CRC. The reported 92% recovery rate at a pre-decoding BER of approximately 0.0043 is therefore specific to the implemented RS(255,223), 16-QAM, AWGN and image-framing configuration.

The stage-contribution results show that most codewords were recovered through direct decoding. Even at 15.021 dB, only a mean of 4.42 of the 619 codewords required successful adaptive recovery. The nested erasure-set search made the largest contribution, recovering a mean of 3.25 codewords per trial. Rank-position substitution and the broader search recovered fewer codewords but remained important because failure to recover even one required block prevented bit-perfect image reconstruction. This selective escalation is consistent with progressive and Chase-type decoders that allocate additional computation to unresolved or less reliable decoding events [12]–[15]. Previous research evaluates RS list-decoder complexity through interpolation cost and its dependence on the number of instantaneous errors [16]. Therefore, runtime, the number of decoding attempts, and memory usage should be measured before claiming that the proposed method has a quantitative complexity advantage.

CRC verification was essential because an algebraically valid decoded RS codeword is not necessarily the originally transmitted codeword when decoding is attempted beyond the conventional correction radius. Each recovered block was accepted only after block-CRC validation, and the reconstructed image was subsequently checked using the complete-image CRC. Successful adaptive recovery produced a post-decoding BER of zero, an SSIM of 1.000 and an infinite PSNR, confirming exact byte-level recovery rather than merely an improvement in visual quality. Nevertheless, a finite-length CRC provides strong error detection rather than an absolute guarantee because its probability of undetected error is small but nonzero [17].

The constellation plots support the reliability measure used by the decoder. At 15.021 dB, the received samples were more dispersed and more frequently approached the boundaries between neighbouring constellation points. Such samples produced smaller nearest-to-second-nearest distance margins and were assigned greater unreliability. The associated byte positions were consequently prioritised during the adaptive erasure searches.

The main limitations are the use of a single image, one RS configuration, an AWGN channel and three closely spaced SNR conditions. In addition, the broader search was limited to 100,000 trial erasure sets per unresolved codeword, and execution time, memory consumption, and hardware complexity were not measured. Future work should evaluate the method using multiple images and general binary data, wider SNR ranges, different RS code rates, higher-order modulation schemes, and fading and interference channels. The number of decoding attempts and the execution time should also be measured to support direct complexity–performance comparisons with progressive soft-decision, Chase-type, and list-decoding approaches [12]–[16].

8. Conclusion

This study developed a CRC-verified adaptive erasure decoder for reliable RS-protected 16-QAM image transmission over an AWGN channel. The proposed receiver obtains byte-level reliability information from the distance margin between the nearest and second-nearest constellation points. Direct RS decoding is attempted first, after which unresolved codewords are processed progressively using nested erasure-set search, rank-position substitution and a broader reliability-guided search. Block-level CRC verification is applied to every recovered block, while a complete-image CRC verifies the final reconstructed output.

The proposed decoder provided its greatest improvement near the conventional RS decoding threshold. At a measured SNR of 15.021 dB, direct RS decoding did not recover any complete image, whereas the adaptive decoder achieved a 92% bit-perfect image-recovery rate. At 15.340

dB, the recovery rate increased from 71% to 99%. At 15.667 dB, direct and adaptive decoding achieved 99% and 100% recovery, respectively. The nested erasure-set search produced the largest contribution among the adaptive stages, while rank-position substitution and the broader search recovered additional difficult codewords. Successfully reconstructed images achieved zero post-decoding BER, an SSIM of 1.000 and infinite PSNR, confirming exact byte-level recovery.

The results demonstrate that the failure of only a small number of RS codewords can prevent complete-image reconstruction and that reliability-guided processing can recover many of these codewords without changing the original RS code structure. Because additional processing is restricted to unresolved codewords, most received blocks continue to be handled through conventional direct decoding. Future work should evaluate the method using multiple images, wider SNR ranges, different RS code rates, higher-order modulation schemes, and fading and interference channels. Runtime, memory use and number of decoding attempts should also be included to quantify the complexity and implementation cost of each adaptive recovery stage.

References

[1] I. S. Reed and G. Solomon, "Polynomial Codes Over Certain Finite Fields," *Journal of the Society for Industrial and Applied Mathematics*, vol. 8, no. 2, pp. 300–304, 1960, doi: 10.1137/0108018.

[2] G. D. Forney Jr., "Generalized Minimum Distance Decoding," *IEEE Trans. Inf. Theory*, vol. 12, no. 2, pp. 125–131, 1966, doi: 10.1109/TIT.1966.1053873.

[3] D. Chase, "A Class of Algorithms for Decoding Block Codes with Channel Measurement Information," *IEEE Trans. Inf. Theory*, vol. 18, no. 1, pp. 170–182, 1972, doi: 10.1109/TIT.1972.1054746.

[4] R. Koetter and A. Vardy, "Algebraic Soft-Decision Decoding of Reed-Solomon Codes," *IEEE Trans. Inf. Theory*, vol. 49, no. 11, pp. 2809–2825, 2003, doi: 10.1109/TIT.2003.819332.

[5] V. Guruswami and M. Sudan, “Improved Decoding of Reed-Solomon and Algebraic-Geometry Codes,” *IEEE Trans. Inf. Theory*, vol. 45, no. 6, pp. 1757–1767, 1999, doi: 10.1109/18.782097.

[6] J. Jiang and K. R. Narayanan, “Iterative Soft-Input Soft-Output Decoding of Reed-Solomon Codes by Adapting the Parity-Check Matrix,” *IEEE Trans. Inf. Theory*, vol. 52, no. 8, pp. 3746–3756, 2006, doi: 10.1109/TIT.2006.878176.

[7] M. El-Khamy and R. J. McEliece, “Iterative Algebraic Soft-Decision List Decoding of Reed-Solomon Codes,” *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 3, pp. 481–490, 2006, doi: 10.1109/JSAC.2005.862399.

[8] J. Bellorado and A. Kavcic, “Low-Complexity Soft-Decoding Algorithms for Reed-Solomon Codes—Part I: An Algebraic Soft-In Hard-Out Chase Decoder,” *IEEE Trans. Inf. Theory*, vol. 56, no. 3, pp. 945–959, 2010, doi: 10.1109/TIT.2009.2039073.

[9] J. Bellorado, A. Kavcic, M. Marrow, and L. Ping, “Low-Complexity Soft-Decoding Algorithms for Reed-Solomon Codes—Part II: Soft-Input Soft-Output Iterative Decoding,” *IEEE Trans. Inf. Theory*, vol. 56, no. 3, pp. 960–967, 2010, doi: 10.1109/TIT.2009.2039091.

[10] S. B. Weinstein and P. M. Ebert, “Data Transmission by Frequency-Division Multiplexing Using the Discrete Fourier Transform,” *IEEE Transactions on Communication Technology*, vol. 19, no. 5, pp. 628–634, 1971, doi: 10.1109/TCOM.1971.1090705.

[11] R. Koetter, “Fast Generalized Minimum-Distance Decoding of Algebraic-Geometry and Reed-Solomon Codes,” *IEEE Trans. Inf. Theory*, vol. 42, no. 3, pp. 721–736

, 1996, doi: 10.1109/18.490540.

[12] L. Chen, S. Tang, and X. Ma, “Progressive Algebraic Soft-Decision Decoding of Reed-Solomon Codes,” *IEEE Transactions on Communications*, vol. 61, no. 2, pp. 433–442, 2013, doi: 10.1109/TCOMM.2012.100912.110752.

[13] J. Xing, L. Chen, and M. Bossert, “Progressive Algebraic Soft-Decision Decoding of Reed-Solomon Codes Using Module Minimization,” *IEEE Transactions on Communications*, vol. 67, no. 11, pp. 7379–7391, 2019, doi: 10.1109/TCOMM.2019.2927207.

[14] S. Tang, S. Cai, and X. Ma, “A New Chase-Type Soft-Decision Decoding Algorithm for Reed-Solomon Codes,” *Alexandria Engineering Journal*, vol. 61, no. 12, pp. 13067–13077, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1110016822004604

[15] J. Xing, L. Chen, and M. Bossert, “Low-Complexity Chase Decoding of Reed-Solomon Codes Using Module,” *IEEE Transactions on Communications*, vol. 68, no. 10, pp. 6012–6022, 2020, doi: 10.1109/TCOMM.2020.3011991.

[16] Y. Cassuto, J. Bruck, and R. J. McEliece, “On the Average Complexity of Reed-Solomon List Decoders,” *IEEE Trans. Inf. Theory*, vol. 59, no. 4, pp. 2336–2351, 2013, doi: 10.1109/TIT.2012.2235522.

[17] P. Koopman and T. Chakravarty, “Cyclic Redundancy Code (CRC) Polynomial Selection for Embedded Networks,” in *Proceedings of the 2004 International Conference on Dependable Systems and Networks*, 2004, pp. 145–154, doi: 10.1109/DSN.2004.1311885.