

Secure image steganography using RRPBA-encrypted payload embedding

Jerico S. Baladhay, Ronielle B. Antonio

School of Architecture, Technology, and Computing, National University, Manila, Philippines

Article Info

Article history:

Received May 16, 2026

Revised Aug 25, 2026

Accepted Sep 4, 2026

Keywords:

Encryption

Image steganography

Least significant bit

Modified advanced encryption standard

Peak signal-to-noise ratio

ABSTRACT

Protecting both the content and the existence of transmitted data is a growing requirement in modern secure communication, and combining cryptography with steganography addresses both needs at once. However, most crypto-stego schemes rely on standard ciphers such as advanced encryption standard (AES), and the effect of a reduced-round, permutation-based AES variant on steganographic quality has not been established. This study evaluates an encrypt-then-hide scheme that pairs the round-reduced permutation-based advanced encryption standard (RRPBA) with single-plane least significant bit (LSB) image steganography. The payload is encrypted with RRPBA in cipher block chaining mode, embedded in the LSBs of a cover image, then extracted and decrypted. The scheme is tested on five standard USC-SIPI images at embedding rates from 25% to 100% of capacity, with robustness and efficiency comparisons against standard AES. The peak signal-to-noise ratio ranges from 57.16 to 51.14 dB, the structural similarity index stays above 0.994, and the bit error rate (BER) is 0% under lossless storage, while RRPBA encrypts about thirteen times faster than AES. The results show that RRPBA is a drop-in replacement for AES in encrypt-then-hide architectures without loss of steganographic quality.

This is an open access article under the [CC BY-SA](#) license.



Corresponding Author:

Jerico S. Baladhay

School of Architecture, Technology, and Computing, National University

SM City Clark, Brgy. Dau, Mabalacat City, Pampanga, Philippines

Email: jsbaladhay@nu-clark.edu.ph

1. INTRODUCTION

The increasing volume of sensitive information transmitted across public communication channels has intensified the demand for security mechanisms that protect both the content and the existence of transmitted data [1]. Encryption addresses the first requirement by transforming plaintext into unintelligible ciphertext, but the resulting ciphertext remains visible to network observers and may attract attention precisely because it is recognizable as encrypted content [2]. Steganography addresses the second requirement by concealing data within innocuous carrier media such as digital images, rendering the communication invisible to observers who are unaware of the hidden payload [3].

Image steganography techniques are broadly classified into spatial domain methods and transform domain methods [4], [5]. Spatial domain methods, of which least significant bit (LSB) replacement is the most widely deployed, modify the pixel values of the cover image directly to embed the payload [6]. Transform domain methods first apply a frequency transformation to the cover image, such as the discrete cosine transform or the discrete wavelet transform, and embed the payload in the transformed coefficients. Spatial domain methods generally provide higher embedding capacity at the cost of reduced robustness against image processing operations, while transform domain methods provide stronger robustness at lower

capacity [7]. The selection between these classes depends on the operational requirements of the deployment scenario.

The combination of encryption and steganography in an encrypt-then-hide paradigm provides a dual-layer security architecture [8]. In this paradigm, the payload is first encrypted to ensure confidentiality, and the encrypted payload is then embedded within a cover image using steganographic techniques to ensure undetectability [9]. An adversary who intercepts the stego image perceives only an ordinary image and has no indication that encrypted data is present. An adversary who somehow extracts the embedded data encounters ciphertext rather than plaintext, requiring cryptanalytic effort [10] to recover the original content. Neither encryption alone nor steganography alone provides this combined guarantee [11].

Several prior works have explored the combination of encryption algorithms with steganographic embedding [12], [13]. The advanced encryption standard (AES) has been combined with LSB embedding by multiple authors due to its standardization and widespread availability [14]. Hybrid schemes incorporating AES with chaotic encryption, compression layers, and additional cryptographic primitives such as Rivest–Shamir–Adleman (RSA) have also been investigated [15], [16]. However, these studies predominantly rely on standard cryptographic primitives whose performance characteristics may not align with the requirements of resource-constrained or domain-specific deployment scenarios [17].

Modified AES variants have been proposed to address the performance limitations of standard AES while preserving its security properties [18]. These variants typically modify one or more of the AES round components, including substitution boxes, permutation layers, or the round count, to achieve favorable trade-offs between encryption speed, hardware footprint, and security margin [19]. The systematic literature on lightweight block ciphers identifies round reduction and permutation-based diffusion as two principal design strategies in this space [20].

In prior work, the round-reduced permutation-based advanced encryption standard (RRPBA) was proposed as a modified AES variant incorporating bit-permutation operations and a reduced round structure designed to address performance constraints while maintaining confidentiality properties. RRPBA has been evaluated for avalanche effect, key sensitivity, and encryption throughput across multiple studies. However, RRPBA has not been evaluated as a component in a steganographic system. Unlike lightweight AES variants that primarily target hardware area or energy efficiency [18], [20], RRPBA reduces the round count and replaces MixColumns diffusion with a bit permutation, and its interaction with a downstream steganographic embedding stage has not previously been characterized. This represents a gap in the evaluation of RRPBA, because the practical utility of a cipher in layered security architectures depends on its compatibility with downstream embedding operations, including the preservation of steganographic quality metrics such as visual imperceptibility and extraction accuracy [21].

This paper proposes a dual-layer security scheme that combines RRPBA encryption with LSB image steganography in an encrypt-then-hide architecture. The contribution is threefold. First, the paper specifies the combined scheme in which plaintext is encrypted using RRPBA and the resulting ciphertext is embedded into the LSBs of a cover image. Second, the paper evaluates the steganographic quality of the resulting stego images using standard image quality metrics including peak signal-to-noise ratio (PSNR), mean squared error (MSE), and structural similarity index (SSIM) [22]. Third, the paper verifies end-to-end extraction accuracy by recovering the embedded ciphertext from the stego image, decrypting it using RRPBA, and comparing the recovered plaintext against the original. The scope of this paper is limited to the steganographic quality evaluation of the combined scheme. The cryptographic properties of RRPBA, including avalanche effect, key sensitivity, and throughput, have been established in prior work and are not re-evaluated here. The steganographic resistance of the scheme to advanced steganalytic attacks beyond basic visual and statistical evaluation is identified as future work.

The remainder of this paper is organized as follows. Section 2 presents the methodology, including the scheme specification, the implementation environment, the test image corpus, and the evaluation protocol. Section 3 presents the results and discussion. Section 4 concludes the paper and identifies directions for future investigation.

2. METHOD

The designing of the methodology takes place in four phases. The encrypt-then-hide scheme is specified in the first phase by the definition of the encrypt, embed, extract, and decrypt operations. During the second phase, the field of application and test image corpus are specified. The specification of the evaluation protocol applied to the scheme is in phase three. The fourth step identifies the measurements which will be taken to analyze the experimental data.

2.1. Scheme specification

The proposed scheme consists of two stages in order. The sender performs the encryption and embedding as the first stage of the process. The receiver's end includes stages of extraction and decryption. The illustration of the architecture overview of the proposed scheme is depicted in Figure 1. It shows the sender and receiver stages of the scheme along with the respective cryptographic operations and steganographic operations performed in each stage of the scheme. It also shows secret key K which plays the role of encryption-decryption symmetry of the scheme. The scheme's sender stage is responsible for encrypting the plaintext message M to obtain ciphertext C using RRPBA with the help of the shared secret key K . Following this, the text further inserts the ciphertext or C into a cover or I_{cover} image for LSB replacement producing a stego image I_{stego} . Furthermore, the steganographer, who is a different person from the RRPBA user, is responsible for embedding the ciphertext C into the cover image I_{cover} . The stego image is sent over a public channel or put into a shared memory. The receiver stage firstly extracts the ciphertext C from the received stego image. Then, it decrypts C using RRPBA under the same secret key K to recover the original plaintext message M' . Thus, the RRPBA-LSB secure stego algorithm is completed.

In the sender phase the plaintext message M is encrypted with a RRPBA block cipher under a shared secret key K which gives ciphertext C i.e., $C = \text{RRPBA_Encrypt}(K, M)$. Next, the ciphertext C is transformed into a binary bitstream. Additionally, the cover image I_{cover} separates into its constituent pixel values in the red, green, blue channel. Following this, the binary bit-stream of C is embedded sequentially into the LSBs of the pixel values. This is done by replacing the LSB bit of each pixel channel value with the corresponding bit of C . The resulting image is known as the Stego image. Visually, I_{cover} and I_{stego} will appear the same, but the LSB plane of I_{stego} , which was not visually significant in I_{cover} , is now carrying the payload.

Through the receiver stage, the stego image I_{stego} is first decomposed to its pixel values (Channel-wise). The LSBs of the cover image and message are extracted as per the sequential order scheme to obtain the ciphertext C . The final decryption with RRPBA under the same shared secret key K yields the original plaintext M , that is $M = \text{RRPBA_Decrypt}(K, C)$. The scheme's ability to embed for one LSB plane is simply $\text{SpatialDimension}(I_{cover}) * \text{ColorDepth}(I_{cover})$.

An RGB image with width W and height H can store $W \times H \times 3$ bits. Each of the three color channels contributes a LSB for this reason. Thus, the payload size should not exceed this (the previous maximum).




Figure 1. Architectural overview of the proposed RRPBA-LSB encrypt-then-hide scheme

2.2. Algorithm specification

The embedding procedure is presented in Algorithm 1. The procedure accepts the cover image, the plaintext message, and the encryption key, and produces the stego image. The extraction and decryption

procedure is presented in Algorithm 2. The procedure accepts the stego image, the payload length in bits, and the decryption key, and produces the recovered plaintext.

Algorithm 1. RRPBA-LSB embedding

Input: I_{cover} (cover image, $W \times H \times 3$), M (plaintext), K (RRPBA key)
 Output: I_{stego} (stego image)
 1: $C \leftarrow RRPBA_Encrypt(K, M)$
 2: $B \leftarrow ConvertToBitstream(C)$
 3: $L \leftarrow \text{length}(B)$
 4: $I_{stego} \leftarrow \text{copy}(I_{cover})$
 5: $idx \leftarrow 0$
 6: for each pixel (r, g, b) in I_{stego} do
 7: for each channel value v in (r, g, b) do
 8: if $idx < L$ then
 9: $v \leftarrow (v \text{ AND } 0xFE) \text{ OR } B[idx]$
 10: $idx \leftarrow idx + 1$
 11: end if
 12: end for
 13: end for
 14: return I_{stego}

Algorithm 2. RRPBA-LSB extraction and decryption

Input: I_{stego} (stego image), L (payload length in bits), K (RRPBA key)
 Output: M (recovered plaintext)
 1: $B \leftarrow \text{empty bitstream}$
 2: $idx \leftarrow 0$
 3: for each pixel (r, g, b) in I_{stego} do
 4: for each channel value v in (r, g, b) do
 5: if $idx < L$ then
 6: $B[idx] \leftarrow v \text{ AND } 0x01$
 7: $idx \leftarrow idx + 1$
 8: end if
 9: end for
 10: end for
 11: $C \leftarrow ConvertFromBitstream(B)$
 12: $M \leftarrow RRPBA_Decrypt(K, C)$
 13: return M

The `ConvertToBitstream` function converts the byte-level ciphertext into a sequence of individual bits in most-significant-bit-first order. The `ConvertFromBitstream` function performs the inverse operation. The AND and OR operations in line 9 of Algorithm 1 perform bitwise masking to clear the LSB and replace it with the payload bit respectively.

2.3. Round-reduced permutation-based advanced encryption standard cipher specification

RRPBA is a modified AES block cipher, defined and cryptographically evaluated in prior work [23]. It reuses the standard AES key schedule and the AES substitution box, and departs from AES in its diffusion layer and in the number of rounds. This section dissects the cipher so that the present paper is self-contained, while the full definitions of the permutation tables and the security analysis remain in [23]. In this study the cipher is used with a 128-bit key, in cipher block chaining mode, and without modification. Figure 2 presents the RRPBA block encryption process as a whole, where Figure 2(a) shows the overall round structure and Figure 2(b) details the construction of one output word of the bit-permutation layer; the two subfigures are described in turn in the components below.

- Block and state: RRPBA operates on 128-bit blocks and, in this study, uses a 128-bit key. A block is represented as four 32-bit words, equivalently a four-by-four array of 16 bytes.
- Round structure: block encryption follows the structure shown in Figure 2(a). It begins with an initial `AddRoundKey`, in which the block is combined by exclusive-or with the first round key. This is followed by N_r-1 main rounds, each of which applies `SubBytes`, `ShiftRows`, the bit-permutation layer, and `AddRoundKey` in that order. A final round then applies `SubBytes`, `ShiftRows`, and `AddRoundKey` only, omitting the bit-permutation layer.
- Number of rounds: RRPBA uses 6 rounds for the 128-bit key, against 10 for standard AES-128, and this round reduction is the first source of the cipher's lower cost. `SubBytes` applies the AES substitution box to every byte, and `ShiftRows` applies the standard AES cyclic row shift.
- Bit-permutation layer: the distinctive component replaces the AES `MixColumns` transformation and is detailed in Figure 2(b). It uses four fixed bit-permutation tables, denoted BP1 through BP4, each a cyclic shift of the eight bit positions of a byte and precomputed as a 256-entry lookup, as defined in [23]. After

SubBytes and ShiftRows, the four bytes of a word are denoted b_0, b_1, b_2 , and b_3 . The layer first mixes adjacent bytes by exclusive-or, forming $R_1 = b_0 \oplus b_1$, $R_2 = b_1 \oplus b_2$, $R_3 = b_2 \oplus b_3$, and $R_4 = b_3 \oplus R_1$, where the last term feeds R_1 back so that the four results are coupled. It then applies one table to each result, giving $FR_1 = BP_1[R_1]$, $FR_2 = BP_2[R_2]$, $FR_3 = BP_3[R_3]$, and $FR_4 = BP_4[R_4]$, and concatenates them into the output word $FR_1 \parallel FR_2 \parallel FR_3 \parallel FR_4$. The four words of the block use the four distinct ShiftRows diagonals, so that diffusion spreads across the whole block over successive rounds. Because the tables are fixed bit permutations combined by exclusive-or, the layer avoids the finite-field matrix multiplication of MixColumns, and this is the second source of the cipher's lower cost. Decryption reverses the round sequence using the inverse substitution box and the inverse bit-permutation tables.

- Key schedule and mode: the round keys are produced by the standard AES key expansion, which derives N_r+1 128-bit round keys from the cipher key using the substitution box, the round-constant sequence, and the rotate-word and substitute-word operations of AES. In the present scheme the cipher operates in cipher block chaining mode, where each 16-byte plaintext block is combined by exclusive-or with the previous ciphertext block, or with the initialization vector for the first block, before block encryption, and the payload is zero-padded to a multiple of the block size. Algorithm 3 lists the block encryption procedure depicted in Figure 2.




Figure 2. RRPBA block encryption: (a) the round structure, with an initial key whitening, N_r-1 main rounds, and a final round without the bit-permutation layer and (b) construction of one output word of the bit-permutation layer that replaces MixColumns

Algorithm 3. RRPBA block encryption

Input: P (128-bit block), round keys $K_0..K_6$
Output: C (128-bit ciphertext block)
1: $(s_0, s_1, s_2, s_3) \leftarrow \text{AddRoundKey}(P, K_0)$ // initial whitening
2: for $i \leftarrow 1$ to 5 do // main rounds
3: $(s_0, s_1, s_2, s_3) \leftarrow \text{ShiftRows}(\text{SubBytes}(s_0, s_1, s_2, s_3))$
4: for each word with bytes (b_0, b_1, b_2, b_3) do
5: $R_1 \leftarrow b_0 \oplus b_1$; $R_2 \leftarrow b_1 \oplus b_2$
6: $R_3 \leftarrow b_2 \oplus b_3$; $R_4 \leftarrow b_3 \oplus R_1$
7: word $\leftarrow BP_1[R_1] \parallel BP_2[R_2] \parallel BP_3[R_3] \parallel BP_4[R_4]$
8: end for
9: end for
10: end for
11: $(s_0, s_1, s_2, s_3) \leftarrow \text{ShiftRows}(\text{SubBytes}(s_0, s_1, s_2, s_3))$ // final round
12: $C \leftarrow \text{AddRoundKey}(s_0, s_1, s_2, s_3, K_6)$
13: return C

In this study RRPBA is applied with a 128-bit key, so N_r is 6, and it is used exactly as originally proposed; no parameter or component of the cipher is altered.

2.4. Implementation environment

The reference implementation is written in Python 3.12. Image processing operations use the Pillow library for image input and output and the NumPy library for pixel-level array manipulation. The PSNR and MSE computations are implemented directly from their mathematical definitions using NumPy array operations. The SSIM computation uses the structural similarity function provided by the scikit-image library, which implements the original formulation as a composite measure of luminance, contrast, and structural similarity. The Pillow and NumPy libraries are widely adopted in image processing research for their reliability in lossless pixel manipulation [24].

The RRPBA encryption and decryption operations use the existing RRPBA implementation from prior work [23]. No modifications to the RRPBA algorithm are introduced in this study. The encryption key is generated as a random byte sequence of the length specified by the RRPBA key schedule and is held constant across all experiments for a given test image to ensure reproducibility.

2.5. Test image corpus

The evaluation uses five standard test images widely adopted in the image processing and steganography literature. The images are sourced from the USC-SIPI Image Database [25] and are listed in Table 1.

Table 1. Test image corpus

Image	Dimensions	Color depth	Source
Airplane	512×512	24-bit RGB	USC-SIPI
Baboon	512×512	24-bit RGB	USC-SIPI
Peppers	512×512	24-bit RGB	USC-SIPI
Sailboat	512×512	24-bit RGB	USC-SIPI
Splash	512×512	24-bit RGB	USC-SIPI

All images are stored in lossless PNG format to prevent compression artifacts from affecting LSB values. The selection of standard test images permits direct comparison of the reported metrics with values reported in the existing steganography literature.

2.6. Evaluation protocol

The evaluation protocol is applied to each test image across four embedding rates corresponding to 25%, 50%, 75%, and 100% of the maximum embedding capacity. At each embedding rate, a random plaintext of the corresponding size is generated, encrypted using RRPBA, and embedded into the cover image using Algorithm 1. The resulting stego image is then evaluated against the cover image using the metrics defined in subsection 2.7. The embedded payload is subsequently extracted from the stego image using Algorithm 2, decrypted using RRPBA, and compared against the original plaintext to verify extraction accuracy.

The four embedding rates are selected to characterize the relationship between payload size and steganographic quality [26]. At low embedding rates, the distortion introduced by LSB replacement is expected to be minimal and imperceptible. At high embedding rates approaching full capacity, the distortion may become measurable, and the evaluation quantifies the point at which quality degradation becomes significant.

2.7. Evaluation metrics

The evaluation employs five metrics defined as follows. MSE is defined as $MSE = (1 / MN) \sum_i \sum_j [I_{\text{cover}}(i,j) - I_{\text{stego}}(i,j)]^2$, where M and N are the image dimensions and $I_{\text{cover}}(i,j)$ and $I_{\text{stego}}(i,j)$ are the pixel values at position (i,j) in the cover and stego images, respectively. Lower MSE values indicate less distortion [27].

PSNR is defined as $PSNR = 10 \log_{10}(MAX^2 / MSE)$, where MAX is the maximum possible pixel value (255 for 8-bit images) and MSE is the mean squared error defined above. PSNR is expressed in decibels (dB). Values above 40 dB are generally considered to indicate imperceptible distortion in the image processing literature [8].

SSIM is defined as a composite measure of luminance, contrast, and structural similarity between two images. SSIM ranges from 0 to 1, where 1 indicates identical images. Values above 0.99 indicate excellent preservation of visual quality.

Embedding capacity is defined as the maximum number of bits that can be embedded in a given cover image using single-plane LSB replacement, computed as $W \times H \times C_{\text{channels}}$, where W and H are the image dimensions and C_{channels} is the number of color channels (3 for RGB, 1 for grayscale). Capacity is reported both in absolute bits and in bits per pixel (bpp).

Bit error rate (BER) is defined as the ratio of incorrectly recovered plaintext bits to total plaintext bits after extraction and decryption, expressed as a percentage. A BER of 0% indicates perfect recovery. Any nonzero BER under the lossless experimental conditions of this study indicates an implementation defect requiring correction.

3. RESULTS AND DISCUSSION

The outcomes of the evaluation protocol described in subsection 2.6 are presented in this section along with their impact on the proposed scheme's viability.

3.1. Embedding capacity analysis

Table 2 presents the embedding capacity of each test image when single-plane LSB replacement is applied. The image quality results, evaluated using all three metrics, demonstrate that the proposed scheme introduces negligible and visually imperceptible distortion across all embedding rates examined. The 25% has about 57 dB and 100% has about 51 dB. Thus, all the values are quite high and they are above 40 dB value which is the imperceptibility threshold cited in the image processing literature. All conditions exhibit MSE values bounded by 0.55, which is expected as single-plane LSB modification can affect a pixel by at most one intensity level. At all settings, the SSIM values are higher than 0.994, which ensures near-perfect structural preservation.

Table 2. Embedding capacity of test images

Image	Dimensions	Channels	Capacity (bits)	Capacity (bytes)	Capacity (bpp)
Airplane	512×512	3	786,432	98,304	3.0
Baboon	512×512	3	786,432	98,304	3.0
Peppers	512×512	3	786,432	98,304	3.0
Sailboat	512×512	3	786,432	98,304	3.0
Splash	512×512	3	786,432	98,304	3.0

The outcome indicates that the five test images exhibit a similar pattern, with some minor alterations expected due to variation in distribution of pixel value and spatial complexity. The Splash image exhibits marginally lower SSIM than the textured Baboon and Sailboat images at high embedding rates. This trend is expected since Baboon, which has the highest spatial frequency content among the test images, masks LSB modifications within its natural texture variation, yielding higher SSIM. Conversely, Splash, which contains larger smooth regions, is more sensitive to LSB modification, resulting in marginally lower structural similarity scores. The trend is expected since regions with spatial high frequency contain more structural information which is perceptually similar to human vision and is only changed to a small degree by LSB modification. It should be noted that applying RRPBA encryption before embedding does not affect the stego quality compared with embedding unencrypted data.

Changes to the PSNR, MSE, and SSIM values depend only on the number of LSB positions altered and the quantity of change at each altered position. To put it differently, for both plaintext and ciphertext generated by the standard AES or RRPBA, the values are same. Consequently, the encrypt-then-hide method does not incur any quality penalties in comparison to the hide-only method. Nevertheless, the encrypt-then-hide strategy provides an enhanced advantage of confidentiality protection, though requires a slight penalty in stego complexity. Further each row of Figure 3 shows one cover image type embedded at 25%, 50%, 75%, and 100% of available LSB capacity.

The human observer cannot see visual difference between the columns despite the huge difference in payload sizes, thus providing qualitative evidence in support of the invisibility claim that is further strengthened by the PSNR and SSIM calculations given in Tables 3 and 5.

Figure 3 shows a visual comparison of all the five test images at the four embedding rates that we tested. In these four embedding conditions, the content of the cover image is visually the same, while the payload size increases in these conditions from 24,576 bytes at 25% to 98,304 bytes at 100%. The qualitative observation agrees with the PSNR and SSIM values and confirms that the scheme maintains visual quality at the maximum embedding capacity.




Figure 3. Visual comparison of stego images across the five test images at four embedding rates

3.2. Image quality results

Table 3 presents the PSNR values obtained across all test images and embedding rates. Table 4 presents the corresponding MSE values. Table 5 presents the corresponding SSIM values.

Table 3. PSNR values (dB) across test images and embedding rates

Image	25% capacity	50% capacity	75% capacity	100% capacity
Airplane	57.16	54.16	52.39	51.15
Baboon	57.16	54.15	52.40	51.14
Peppers	57.18	54.15	52.38	51.14
Sailboat	57.14	54.15	52.39	51.14
Splash	57.15	54.15	52.39	51.15

Table 4. MSE values across test images and embedding rates

Image	25% capacity	50% capacity	75% capacity	100% capacity
Airplane	0.124999	0.249766	0.375165	0.499453
Baboon	0.125163	0.250089	0.374355	0.500075
Peppers	0.124488	0.250268	0.375589	0.500459
Sailboat	0.125664	0.250048	0.374795	0.500319
Splash	0.125202	0.250080	0.375114	0.499143

Table 5. SSIM values across test images and embedding rates

Image	25% capacity	50% capacity	75% capacity	100% capacity
Airplane	0.998602	0.997431	0.996684	0.995601
Baboon	0.999897	0.999530	0.999205	0.998909
Peppers	0.999090	0.998262	0.997384	0.996302
Sailboat	0.999241	0.998520	0.998099	0.997445
Splash	0.998662	0.997027	0.995739	0.994155

3.3. Extraction accuracy results

Table 6 presents the BER obtained after extraction and decryption across all test images and embedding rates. The extraction accuracy results confirm perfect recovery of the original plaintext across all test images and all embedding rates, with a BER of 0.00% in all twenty experimental conditions. This outcome is the expected result under the lossless experimental protocol, in which the stego image is stored in PNG format without lossy compression and no transmission noise or format conversion is applied between embedding and extraction. The result verifies that the RRPBA encryption operation does not introduce any incompatibility with the LSB embedding and extraction process, and that the end-to-end encrypt-embed-extract-decrypt pipeline functions correctly.

The 0.00% BER is a necessary rather than sufficient condition for the viability of the scheme. In practical deployment scenarios involving lossy compression (such as JPEG conversion), image rescaling, or transmission over noisy channels, the BER is expected to increase. The evaluation of extraction accuracy under such degraded conditions is characterized empirically in subsection 3.5.

Table 6. BER (%) across test images and embedding rates

Image	25% capacity	50% capacity	75% capacity	100% capacity
Airplane	0.0000	0.0000	0.0000	0.0000
Baboon	0.0000	0.0000	0.0000	0.0000
Peppers	0.0000	0.0000	0.0000	0.0000
Sailboat	0.0000	0.0000	0.0000	0.0000
Splash	0.0000	0.0000	0.0000	0.0000

3.4. Comparative discussion

Table 7 presents a comparison of the PSNR values obtained by the proposed scheme against values reported in the existing literature for comparable encrypt-then-hide schemes using LSB steganography. An additional column indicating the approximate embedding rate is included because prior studies vary substantially in payload size relative to cover image capacity, and direct numerical comparison without this context would be misleading.

Table 7. Comparison of stego image quality at maximum LSB capacity with prior encrypt-then-hide schemes

Study	Encryption	Steganography	Test image	Approx. embedding rate	PSNR (dB)	SSIM
Bandekar and Suguna [28]	AES-128	1-bit LSB	256×256	Low (<1%)	N/R	N/R
Gangurde and Tiwari [29]	AES+PLS	1-bit LSB	512×512	Low (<1%)	48.34	N/R
Alanzay <i>et al.</i> [13]	AES+Blowfish	LSB	512×512	Low (<1%)	50.00 to 58.62	N/R
Talasila <i>et al.</i> [30]	AES+RSA	LSB	512×512	Low (<1%)	52.69	N/R
Proposed (25%)	RRPBA	1-bit LSB	512×512	25%	57.16	0.9990
Proposed (100%)	RRPBA	1-bit LSB	512×512	100%	51.14	0.9989

N/R=not reported in the original study. Direct numerical comparison across studies is constrained by differences in embedding rate, image dimensions, and payload size. The four prior studies listed in Table 7 embed short text payloads occupying less than 1% of the available LSB capacity, while the proposed scheme is evaluated systematically across embedding rates from 25 to 100% of maximum capacity. This difference in experimental conditions means that the PSNR values reported by the prior studies reflect minimal distortion under lightweight embedding, whereas the proposed scheme reports PSNR values under progressively heavier embedding conditions including the theoretical worst case at full capacity.

Two observations emerge from the comparison despite this methodological difference. First, at the lowest embedding rate tested (25% capacity), the proposed scheme achieves a PSNR of 57.16 dB, which exceeds the values reported by Gangurde and Tiwari [29] at 48.34 dB and Talasila *et al.* [30] at 52.69 dB, and falls within the range reported by Alanzay *et al.* [13] at 50.00 to 58.62 dB, even though the proposed scheme embeds a substantially larger payload at 25% capacity (24,576 bytes) than the short text messages

used in the prior studies. This outcome indicates that RRPBA encryption does not introduce any steganographic quality penalty relative to standard AES-based approaches operating under comparable or lighter embedding conditions.

Second, at maximum embedding capacity (100%), the proposed scheme achieves a PSNR of 51.14 dB, which matches the theoretical upper bound for single-plane LSB replacement on 8-bit images computed as $10 \log_{10}(255^2/0.5)$, approximately equal to 51.14 dB. None of the prior studies listed in Table 7 report results at full embedding capacity, which means the proposed scheme provides the first empirical confirmation that an encrypt-then-hide architecture using a modified AES variant can achieve the theoretical PSNR bound under maximum-capacity conditions. This result is significant because it demonstrates that the statistical properties of RRPBA ciphertext do not adversely affect LSB embedding quality even when every available LSB position in the cover image is utilized.

The SSIM values obtained by the proposed scheme range from 0.9942 on Splash to 0.9989 on Baboon at 100% embedding capacity, with all values remaining above 0.99 across the twenty experimental conditions. None of the prior studies listed in Table 7 report SSIM values, which limits direct structural similarity comparison. The SSIM results obtained in this study are nonetheless consistent with the threshold of 0.99 conventionally associated with excellent preservation of visual quality in the image processing literature.

The comparative analysis supports two conclusions regarding the contribution of this work. First, RRPBA can serve as a drop-in replacement for standard AES in encrypt-then-hide steganographic architectures without introducing degradation in steganographic quality. The cipher choice does not influence pixel-level distortion in LSB replacement because the LSB modification operation is independent of the statistical properties of the embedded bitstream, and this independence is empirically confirmed by the convergence of the proposed scheme's PSNR with the theoretical bound. Second, the selection of the encryption algorithm in LSB-based steganographic systems should be guided by cryptographic strength, computational efficiency, and deployment requirements rather than by concern over image quality impact, because steganographic quality at the LSB layer is invariant to the cipher used. RRPBA is therefore positioned as appropriate for deployment scenarios where its reduced-round structure and bit-permutation-based diffusion offer performance advantages over standard AES, with the assurance that these design choices do not propagate any cost to the steganographic layer.

3.5. Robustness under lossy channel conditions

The preceding results establish imperceptibility and perfect recovery under lossless storage. Because spatial-domain LSB embedding places the payload in the LSB plane, any operation that perturbs pixel values corrupts the embedded bits, and this section characterizes that behavior explicitly. The scheme assumes a lossless transport model, in which the stego image is stored and exchanged in a lossless format such as PNG and is not re-encoded, rescaled, or otherwise altered in transit. To quantify the boundary of this model, the stego images were subjected to three classes of degradation applied after embedding: JPEG compression at quality factors of 95, 90, 75, and 50; bicubic downscaling to 90, 75, and 50% of the original dimensions followed by restoration to the original size; and additive Gaussian noise with standard deviations of 1, 2, and 5 intensity levels. For each condition the payload was extracted and decrypted, and two error measures were computed: the LSB corruption rate, defined as the fraction of embedded LSBs altered by the degradation, and the recovered-payload BER after RRPBA decryption.

Table 8 reports the results. Under the lossless baseline the LSB corruption rate and the recovered-payload BER are both 0.00%, consistent with Table 6. Under every lossy condition, including the mild JPEG setting at quality 95, the LSB corruption rate rises to approximately 50%, which is statistically indistinguishable from random guessing, and the recovered-payload BER is correspondingly approximately 50%. Because the LSB plane is not preserved by lossy compression, resizing, or additive noise, essentially half of the embedded bits are flipped, and the cipher block chaining structure of RRPBA propagates each affected block so that the decrypted payload is unrecoverable. These BERs are governed by perturbation of the LSB plane rather than by cover content, so the reported values are averaged over the five test images.

This outcome is a fundamental property of single-plane LSB steganography rather than a defect of the proposed scheme, and it defines the operational envelope of the method. The scheme is therefore appropriate for covert storage and transmission over channels that preserve the exact pixel representation, such as archival storage, lossless file transfer, and messaging systems that do not re-encode attachments, and it is not suited to channels that apply lossy recompression such as most social-media image pipelines. Two directions follow. First, robustness can be obtained by migrating the embedding stage to a transform-domain method, in which the payload is carried in quantization-robust frequency coefficients rather than in the spatial LSB plane; the encrypt-then-hide construction and the RRPBA cipher layer are independent of the embedding domain and carry over directly. Second, an error-correcting code applied to the ciphertext before

embedding can tolerate limited perturbation at the cost of reduced effective capacity. Both directions are identified as future work.

Table 8. Extraction accuracy of the stego images under lossy channel conditions

Channel condition	LSB corruption rate (%)	Recovered-payload BER (%)
Lossless PNG (baseline)	0.00	0.00
JPEG, quality 95	50.00	49.95
JPEG, quality 90	49.94	49.96
JPEG, quality 75	50.01	50.01
JPEG, quality 50	49.94	50.02
Bicubic resize to 90%	48.15	49.91
Bicubic resize to 75%	49.57	50.00
Bicubic resize to 50%	49.86	49.99
Gaussian noise, sigma 1	49.82	49.93
Gaussian noise, sigma 2	50.10	50.04
Gaussian noise, sigma 5	50.04	50.10

3.6. Computational efficiency

Because the practical value of a modified cipher rests partly on its cost, RRPBA is compared directly with standard AES. The two ciphers differ in round count and in the diffusion primitive, so they were benchmarked under identical conditions using reference implementations written in the same language and coding style, so that the measured difference reflects the algorithmic design rather than implementation tuning. Standard AES-128 uses 10 rounds and the MixColumns transformation, whereas RRPBA-128 uses 6 rounds and the bit-permutation layer. Both were executed on payloads of 1, 8, 32, and 64 kilobytes, with each measurement repeated over 12 independent runs to capture variance. Encryption time is reported as the mean and standard deviation in milliseconds, together with the speedup of RRPBA over AES.

Table 9 reports the results, showing that RRPBA encrypts consistently faster than standard AES across all payload sizes, with a mean speedup of approximately 13 times, ranging from 13.82 times at 8 kilobytes to 11.86 times at 64 kilobytes. The standard deviations remain a modest fraction of the means across all payload sizes, which indicates that the speedup is stable across runs and is not an artifact of measurement noise. The advantage follows directly from the reduced round count and from the replacement of the finite-field MixColumns multiplication with a table-driven bit permutation.

Table 9. Encryption time of RRPBA and standard AES-128 across payload sizes, over 12 runs

Payload (KB)	RRPBA time, ms (mean, SD)	AES-128 time, ms (mean, SD)	Speedup
1	6.49 (1.13)	87.27 (4.17)	13.44x
8	49.45 (2.14)	683.47 (21.37)	13.82x
32	222.20 (14.22)	2852.49 (203.31)	12.84x
64	490.65 (59.89)	5817.36 (484.72)	11.86x

These measurements are obtained from unoptimized reference implementations and are intended to isolate the algorithmic effect of the RRPBA design. They should not be read as a comparison against hardware-accelerated AES, which benefits from dedicated processor instructions, but as evidence that the round reduction and permutation-based diffusion of RRPBA reduce the operation count of the cipher, which is the property relevant to resource-constrained and software-only deployment scenarios.

3.7. Limitations, applications, and ethical considerations

Three limitations bound the present study. First, as established in subsection 3.5, the scheme is confined to lossless channels and does not survive lossy recompression, resizing, or noise. Second, single-plane LSB replacement is known to be detectable by statistical steganalysis, including chi-square and sample-pair analysis and modern learning-based detectors, because embedding data in the LSB plane alters the statistical regularities of that plane; the present evaluation establishes imperceptibility to the human visual system and to full-reference quality metrics, but does not claim resistance to dedicated steganalysis, which is deferred to future work. Third, the evaluation uses a single cipher configuration and a fixed key per image, and does not vary key length or embedding order.

Within these bounds the scheme is applicable wherever confidential data must be exchanged without revealing that communication is taking place, including secure messaging, covert channels in defensive security research, and the embedding of authentication or provenance data in images. The construction is agnostic to the carrier and generalizes to other lossless media, so extending the payload

channel to audio and video carriers, where the LSBs of samples or of intra-frame pixels provide analogous capacity, is a natural direction for scaling the approach. As with any steganographic method, the capacity to conceal communication carries a potential for misuse; the same imperceptibility that protects legitimate privacy can also conceal illicit activity, so responsible deployment should be accompanied by appropriate legal and organizational controls.

4. CONCLUSION

The paper proposes and evaluates a dual-layer security scheme that exploits the combining of RRPBA encryption and single-plane LSB image steganography in an encrypt-then-hide model. Subsequently, the scheme was implemented as a Python reference prototype and performance testing executed on five standard 512 by 512 RGB images sourced from the USC-SIPI Image Database at embedding rate of 25%, 50%, 75%, and 100% of available capacity. The scheme produced the PSNR which stayed in the range of 57.16 decibels at 25% occupancy to 51.14 decibels at 100% occupancy, SSIM values above 0.994 and BER of 0.0000% across all the 20 empirical conditions, thus confirming both visual imperceptibility which is well above the 40-decibel mark and perfect end-to-end recovery of the encrypted payload under lossless storage conditions. The findings indicate that RRPBA can easily replace standard AES in encrypt-then-hide steganographic systems without any quality degradation, which is important because it enhances the applicability of RRPBA from a single primitive for confidentiality to one within layered security architectures needing encryption and undetectability. In the future, the transform domain embedding techniques of the scheme will be evaluated. Moreover, the robustness of the scheme under lossy compression and noisy-channel conditions was characterized in this study, and transform-domain methods will be explored to overcome it. Also, formal steganalytic evaluation against state-of-the-art detection techniques such as chi-square analysis and machine learning based will be conducted. Furthermore, the integration with authenticated encryption modes will be investigated to take the scheme further toward unified confidentiality, integrity, authenticity, and undetectability guarantees.

FUNDING INFORMATION

Publication funding by National University Philippines.

AUTHOR CONTRIBUTIONS STATEMENT

This journal uses the Contributor Roles Taxonomy (CRediT) to recognize individual author contributions, reduce authorship disputes, and facilitate collaboration.

Name of Author	C	M	So	Va	Fo	I	R	D	O	E	Vi	Su	P	Fu
Jerico S. Baladhay	✓	✓	✓	✓		✓	✓		✓	✓	✓	✓	✓	✓
Ronielle B. Antonio	✓	✓		✓	✓	✓	✓	✓	✓	✓		✓	✓	✓

C : **C**onceptualization

M : **M**ethodology

So : **S**oftware

Va : **V**alidation

Fo : **F**ormal analysis

I : **I**nvestigation

R : **R**esources

D : **D**ata Curation

O : Writing - **O**riginal Draft

E : Writing - Review & **E**editing

Vi : **V**isualization

Su : **S**upervision

P : **P**roject administration

Fu : **F**unding acquisition

CONFLICT OF INTEREST STATEMENT

Authors state no conflict of interest.

DATA AVAILABILITY

Derived data supporting the findings of this study are available from the corresponding author [JSB] on request.

REFERENCES

- [1] N. Subramanian, O. Elharrouss, S. Al-Maadeed, and A. Bouridane, "Image Steganography: A Review of the Recent Advances," *IEEE Access*, vol. 9, pp. 23409–23423, 2021, doi: 10.1109/ACCESS.2021.3053998.

- [2] S. Rahman *et al.*, "A Comprehensive Study of Digital Image Steganographic Techniques," *IEEE Access*, vol. 11, pp. 6770–6791, 2023, doi: 10.1109/ACCESS.2023.3237393.
- [3] R. Apau, M. Asante, F. Twum, J. B. Hayfron-Acquah, and K. O. Peasah, "Image steganography techniques for resisting statistical steganalysis attacks: A systematic literature review," *PLOS ONE*, vol. 19, no. 9, pp. 1–47, Sep. 2024, doi: 10.1371/journal.pone.0308807.
- [4] P. C. Mandal, I. Mukherjee, G. Paul, and B. N. Chatterji, "Digital image steganography: A literature survey," *Information Sciences*, vol. 609, pp. 1451–1488, Sep. 2022, doi: 10.1016/j.ins.2022.07.120.
- [5] D. R. I. M. Setiadi, "PSNR vs SSIM: imperceptibility quality assessment for image steganography," *Multimedia Tools and Applications*, vol. 80, no. 6, pp. 8423–8444, Mar. 2021, doi: 10.1007/s11042-020-10035-z.
- [6] H. Aparna and J. Madhumitha, "Combined image encryption and steganography technique for enhanced security using multiple chaotic maps," *Computers and Electrical Engineering*, vol. 110, p. 108824, Sep. 2023, doi: 10.1016/j.compeleceng.2023.108824.
- [7] S. Kaur, S. Singh, M. Kaur, and H.-N. Lee, "A Systematic Review of Computational Image Steganography Approaches," *Archives of Computational Methods in Engineering*, vol. 29, no. 7, pp. 4775–4797, Nov. 2022, doi: 10.1007/s11831-022-09749-0.
- [8] S. Roy and M. M. Islam, "A Hybrid Secured Approach Combining LSB Steganography and AES Using Mosaic Images for Ensuring Data Security," *SN Computer Science*, vol. 3, no. 2, p. 153, Feb. 2022, doi: 10.1007/s42979-022-01046-8.
- [9] A. S. Ali, S. Alsamarace, and A. A. Hussein, "Optimize Image Steganography Based on Distinction Disparity Value and HMPSO to Ensure Confidentiality and Integrity," *Journal of Computer Networks and Communications*, vol. 2024, no. 1, pp. 1–15, Jan. 2024, doi: 10.1155/2024/2516567.
- [10] K. Kordov and S. Zhelezov, "Steganography in color images with random order of pixel selection and encrypted text message embedding," *PeerJ Computer Science*, vol. 7, p. e380, Jan. 2021, doi: 10.7717/peerj-cs.380.
- [11] A. Jan, S. A. Parah, M. Hussan, and B. A. Malik, "Double layer security using crypto-stego techniques: a comprehensive review," *Health and Technology*, vol. 12, no. 1, pp. 9–31, Jan. 2022, doi: 10.1007/s12553-021-00602-1.
- [12] K. K. Reddy, N. V. Kumar, L. Sajan, M. P. Teja, and V. P. Kumar, "Enhancing Steganography Security: A Dual Layer Approach with LSB and AES Algorithm," in *Advances in Computational Intelligence and Informatics*, Eds., Singapore: Springer Nature Singapore, 2024, pp. 321–327, doi: 10.1007/978-981-97-4727-6_32.
- [13] M. Alanzy, R. Alomrani, B. Alqarni, and S. Almutairi, "Image Steganography Using LSB and Hybrid Encryption Algorithms," *Applied Sciences*, vol. 13, no. 21, pp. 1–20, 2023, doi: 10.3390/app132111771.
- [14] A. Aljughaiman and R. Alrawashdeh, "Content-adaptive LSB steganography with saliency fusion, ACO dispersion, and hybrid encryption with ablation study," *Scientific Reports*, vol. 16, no. 1, pp. 1–22, Dec. 2025, doi: 10.1038/s41598-025-33920-9.
- [15] A. Badhan and S. S. Malhi, "Enhancing Data Security and Efficiency: A Hybrid Cryptography Approach (AES + ECC) Integrated with Steganography and Compression Algorithm," in *2025 3rd International Conference on Intelligent Data Communication Technologies and Internet of Things (IDCIoT)*, Feb. 2025, pp. 450–456, doi: 10.1109/IDCIoT64235.2025.10914830.
- [16] B. Kallapu *et al.*, "Multi-Layered Security Framework Combining Steganography and DNA Coding," *Systems*, vol. 13, no. 5, pp. 1–27, 2025, doi: 10.3390/systems13050341.
- [17] A. Hafsa, M. Gafsi, and M. Machhout, "A lightweight and robust block cipher algorithm for real-time applications," *Signal, Image and Video Processing*, vol. 18, no. 2, pp. 1609–1624, Mar. 2024, doi: 10.1007/s11760-023-02831-5.
- [18] O. I. Berezovskaya, S. S. Chuprov, E. A. Neverov, and E. R. Sadreev, "Review and Comparison of Lightweight Modifications of the AES Cipher for a Network of Low-Power Devices," *Automatic Control and Computer Sciences*, vol. 56, no. 8, pp. 994–1006, Dec. 2022, doi: 10.3103/S0146411622080028.
- [19] A. A. Zakaria, A. H. Azni, F. Ridzuan, N. H. Zakaria, and M. Daud, "Systematic literature review: Trend analysis on the design of lightweight block cipher," *Journal of King Saud University - Computer and Information Sciences*, vol. 35, no. 5, pp. 1–28, May 2023, doi: 10.1016/j.jksuci.2023.04.003.
- [20] Y. Zhong and J. Gu, "Lightweight block ciphers for resource-constrained environments: A comprehensive survey," *Future Generation Computer Systems*, vol. 157, pp. 288–302, Aug. 2024, doi: 10.1016/j.future.2024.03.054.
- [21] A. Kumar, R. Rani, and S. Singh, "A survey of recent advances in image steganography," *Security and Privacy*, vol. 6, no. 3, p. e281, May 2023, doi: 10.1002/spy2.281.
- [22] S. Rustad, D. R. I. M. Setiadi, A. Syukur, and P. N. Andono, "Inverted LSB image steganography using adaptive pattern to improve imperceptibility," *Journal of King Saud University - Computer and Information Sciences*, vol. 34, no. 6, Part B, pp. 3559–3568, Jun. 2022, doi: 10.1016/j.jksuci.2020.12.017.
- [23] J. S. Baladhay and E. M. D. L. Reyes, "AES-128 reduced-round permutation by replacing the MixColumns function," *Indonesian Journal of Electrical Engineering and Computer Science*, vol. 33, no. 3, pp. 1641–1652, Mar. 2024, doi: 10.11591/ijeecs.v33.i3.pp1641-1652.
- [24] S. van der Walt *et al.*, "Scikit-image: image processing in Python," *PeerJ*, vol. 2, p. e453, Jun. 2014, doi: 10.7717/peerj.453.
- [25] A. G. Weber, "The USC-SIPI image database: Version 5." USC Signal and Image Processing Institute, University of Southern California, Los Angeles, CA, 1997. [Online]. Available: <https://sipi.usc.edu/database/>, (Accessed: May. 16, 2026).
- [26] Y. Y. Ghadi, T. AlShoul, Z. I. Nezami, H. Ali, M. Asif, and M. J. Bah, "Enhanced payload volume in the least significant bits image steganography using hash function," *PeerJ Computer Science*, vol. 9, p. e1606, Nov. 2023, doi: 10.7717/peerj-cs.1606.
- [27] H. Hardan, A. Alawneh, and N. N. El-Emam, "New deep data hiding and extraction algorithm using multi-channel with multi-level to improve data security and payload capacity," *PeerJ Computer Science*, vol. 8, p. e1115, Oct. 2022, doi: 10.7717/peerj-cs.1115.
- [28] P. P. Bandekar and G. C. Suguna, "LSB Based Text and Image Steganography Using AES Algorithm," in *2018 3rd International Conference on Communication and Electronics Systems (ICCES)*, Coimbatore, India: IEEE, Oct. 2018, pp. 782–788, doi: 10.1109/CESYS.2018.8724069.
- [29] S. Gangurde and K. Tiwari, "LSB Steganography Using Pixel Locator Sequence with AES," Dec. 04, 2020, *arXiv*, doi: 10.48550/arXiv.2012.02494.
- [30] S. Talasila, G. V. Kumar, E. V. Babu, K. Nainika, M. V. Sahithi, and P. Mohan, "The Hybrid Model of LSB—Technique in Image Steganography Using AES and RSA Algorithms," in *Soft Computing and Signal Processing*, Eds., Singapore: Springer Nature Singapore, 2024, pp. 403–413, doi: 10.1007/978-981-99-8451-0_34.

BIOGRAPHIES OF AUTHORS

Jerico S. Baladhay, MIT    is a part-time Assistant Professor III at the School of Architecture, Technology, and Computing, National University (NU) – Clark, Pampanga, Philippines and a permanent Instructor I at the College of Computer Studies, Tarlac State University (TSU), Tarlac City, Philippines. He brings expertise in information technology, software development, and cybersecurity research. His academic career reflects a strong commitment to both instruction and practical application of technology solutions in the fields of algorithm development, data security, and system performance optimization. His research interests lie at the crucial intersection of algorithms, cybersecurity, and digital system performance optimization. His work focuses on cryptographic algorithm modification, specifically exploring innovative approaches to optimize symmetric key encryption without compromising security integrity. He can be contacted at email: jsbaladhay@nu-clark.edu.ph.



Dr. Ronielle B. Antonio    is the IT Program Chairperson at the School of Architecture, Technology, and Computing at National University, Philippines (NU) – Clark, Pampanga, Philippines. He brings expertise in information technology, cryptography, and academic program leadership. His academic career reflects a strong commitment to both instruction and practical application of technology solutions in the fields of data security, web and mobile programming, and project management. His research interests lie at the crucial intersection of IT security, cryptography, and data mining and analytics. His work focuses on the design and evaluation of secure computing systems, with broader engagement in curriculum development, faculty mentorship, and research coordination across the IT discipline. He can be contacted at email: rbantonio@nu-clark.edu.ph.