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.

A Finite-State ARX Image Encryption Scheme Based on Dual Space-Filling Curves

Yifan Ge
Southwest University
Bo Yang
yungbo@swu.edu.cn
Southwest University

Research Article

Keywords:

Posted Date: November 24th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11677435/v1

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

Additional Declarations: No competing interests reported.

A Finite-State ARX Image Encryption Scheme Based on Dual Space-Filling Curves

Yifan Ge[1] and Bo Yang[1*]

[1]College of Electronic and Information Engineering, Southwest University, No. 2 Tiansheng Road, Beibei, Chongqing, 400715, China.

*Corresponding author(s). E-mail(s): yungbo@swu.edu.cn;

Abstract

Finite-state image-encryption generators inevitably produce periodic trajectories, but empirical mixing statistics alone do not certify the control sequence required by an image. We present FSARX, which links exact addition–rotation–exclusive-or (ARX) recurrence analysis to length-dependent control-material admission in a reversible RGB encryption pipeline. Inverse-fiber analysis establishes a one-step transition image of $\mathbf{2^{69}}$ states. Algebraic reduction and exhaustive bit-vector checks identify exactly four fixed points, including three nonzero degeneracies admitted by a zero-preimage-only check. A combined state-return and finite-word period test excludes these degeneracies, guarantees distinct complete states over the accepted prefix, and certifies an emitted-block period greater than half its requested length. The admitted controls drive Feistel block permutation, dual space-filling-curve pixel permutation, joint RGB diffusion, and global reordering. All 17 tested initializations passed at **25 296 921** words, with finite-word period equal to the block length. Experiments on 24 Kodak images and six additional RGB images confirmed exact recovery and full-length finite-word periods. Complete admission increased measured encryption time by 21–27% relative to the baseline implementation. Common-environment CPU experiments quantified the runtime of the complete admission procedure within the encryption pipeline. Multi-position perturbations also revealed image-dependent low-response cases. The contribution is an explicit link between finite-state structure, application-length validation, and reversible image processing, with guarantees confined to the accepted control windows.

Keywords: RGB image encryption, ARX operations, space-filling curves, bidirectional diffusion, finite-state sequence generator

1 Introduction

Image encryption for visual communication must suppress spatial and channel redundancy while permitting exact decryption [1–3]. Recent methods combine nonlinear control, permutation, diffusion, and channel coupling [4–7]. Safdar et al. combine finite-ring constructions with substitution boxes, pixel permutation, and DNA encoding for color images [8]. Song et al. use a three-dimensional Latin cube for simultaneous permutation and diffusion of image batches [9]. These studies illustrate algebraic and multidimensional approaches to image reordering and value transformation. Their reported results depend on the image representation, perturbation protocol, and

implementation. The digital control source introduces a distinct question. Finite-precision generators have ultimately periodic trajectories, so real-valued chaotic arguments alone do not characterize their implemented recurrence [10–12]. Output histograms and perturbation statistics likewise do not identify all fixed states or certify the period of a requested control block. Moreover, a long state trajectory does not imply an equally long output period when the observation map is non-injective. For image encryption, the operational question is therefore what can be guaranteed about the specific finite block allocated to permutation and diffusion. FSARX addresses this question using native unsigned-word ARX arithmetic and an explicit connection between recurrence analysis and control-material admission. Its contributions have three parts. First, exact inverse-fiber analysis determines the transition image, while algebraic fixed-point reduction identifies nonzero period-one degeneracies missed by the baseline rejection rule. Second, a length-dependent admission procedure combines anchor-state return testing with exact finite-word period testing. The resulting guarantees distinguish complete-state repetition, finite output-block periodicity, and eventual cycle periods. Third, the admitted block is allocated to an explicitly reversible RGB pipeline, and experiments evaluate both the guarantees and their processing cost. Original-resolution tests on 30 images, component ablations, and matched-input CPU comparisons provide complementary evidence under stated protocols. The pipeline combines Feistel block permutation, dual space-filling-curve scans, joint RGB diffusion, and global reordering; the central contribution is the connection between exact source analysis and application-length validation. Section 2 derives and tests the control-source properties. Section 3 specifies encryption and inversion, and Section 4 reports image experiments, admission cost, and comparisons with recent public implementations.

2 Finite-State ARX Control Source

2.1 Transition and initialization

Let $M = 2^{32}$, let R_r denote 32-bit left rotation by r bits, and let $\phi = \texttt{0x9E3779B9}$; ϕ is odd and hence invertible modulo M. All word additions and multiplications below are reduced modulo M. The three 32-bit key words K and auxiliary secret words N determine

$$\begin{aligned} x_0 &= K_x \oplus n_0, \\ y_0 &= K_y \oplus R_7(n_1), \\ z_0 &= K_z \oplus R_{13}(n_2), \\ D &= 256 + [(n_0 \oplus n_1 \oplus n_2) \bmod 256]. \end{aligned} \tag{1}$$

For $S_i = (x_i, y_i, z_i)$, the sequential update is

$$\begin{aligned} v_i &= (y_i \oplus z_i) \bmod 32, \\ x_{i+1} &= R_{v_i}(x_i) \oplus [y_i + z_i]_M, \\ y_{i+1} &= [\phi(v_i + (x_{i+1} \oplus z_i))]_M, \\ u_i &= (x_{i+1} \oplus y_{i+1}) \bmod 32, \\ z_{i+1} &= [R_{u_i}(z_i) + (x_{i+1} \oplus y_{i+1})]_M, \\ o_i &= R_{13}(x_{i+1}) \oplus \mathrm{ROTR}_{32}(y_{i+1}, 17) \oplus z_{i+1}. \end{aligned} \tag{2}$$

The D-step initial segment is discarded. State-dependent rotations, carry propagation, and odd multiplication make the update nonlinear over GF(2) while remaining exactly reproducible in unsigned arithmetic. The tuple (K, N) acts through the mixed 96-bit initial state and the eight-bit discard selector. Distinct tuples can therefore be equivalent; their nominal 192 input bits do not establish a 192-bit security level. Reusing identical parameters also reproduces the control material. The experiments below hold these parameters fixed to isolate component effects.

2.2 State image, fixed points, and period bounds

The state space has 2^{96} elements. For a target $S' = (x', y', z')$, the reverse update uniquely recovers

$$\begin{aligned} u &= (x' \oplus y') \bmod 32, \\ z &= \mathrm{ROTR}_{32}([z' - (x' \oplus y')]_M, u), \\ v &= [\phi^{-1} y' - (x' \oplus z)]_M. \end{aligned} \tag{3}$$

The target is in the one-step image $I = \mathrm{Im}(\mathsf{T})$ exactly when $v < 32$. In that case, the 27 high bits of the predecessor's y word are free: for $0 \le k < 2^{27}$,

$$\begin{aligned} y_k &= 32k + [v \oplus (z \bmod 32)], \\ x_k &= \mathrm{ROTR}_{32}(x' \oplus [y_k + z]_M, v). \end{aligned} \tag{4}$$

Each (x_k, y_k, z) maps to S', and these are all its predecessors. Thus every image state has exactly 2^{27} predecessors, giving

$$|I| = 2^{69}, \qquad 1 \le p \le 2^{69}. \tag{5}$$

Because $S_t \in I$ for $t \ge 1$, the transient-plus-cycle bound is

$$\mu + p \le \begin{cases} 2^{69}, & S_0 \in I, \\ 2^{69} + 1, & S_0 \notin I. \end{cases} \tag{6}$$

Here μ is the transient length and p the state-cycle period. All cyclic states lie in I. The image admits exact coordinates (x', z, v) with $64 + 5 = 69$ bits: the reverse formulas uniquely recover z, v and the forward formulas recover y', z'. This coordinate model preserves cycles and transients within I and distinguishes the image size from the nominal 96-bit state representation. The complete preimage of zero is $P_0 = \{(32k, 32k, 0) : 0 \le k < 2^{27}\}$. This follows by reversing the z, y, and x updates in order: $z' = 0$ forces $z = 0$, $y' = 0$ forces $v = 0$, and $x' = 0$ then forces $x = y$ with $y \bmod 32 = 0$. Zero is fixed, but checking only P_0 misses the three nonzero fixed points in Table 1. For a cycle of length p, put $A_i = [y_i + z_i]_M$, $v_i = (y_i \oplus z_i) \bmod 32$, $V = [\sum_i v_i]_{32}$, and $\sigma_i = [\sum_{j=i+1}^{p-1} v_j]_{32}$. Composing the x-updates yields

$$x_0 \oplus R_V(x_0) = \bigoplus_{i=0}^{p-1} R_{\sigma_i}(A_i). \tag{7}$$

If $g = \gcd(V, 32)$, each bit orbit of length $32/g$ on the right must have zero XOR. These parity conditions are necessary and sufficient for closure of the composed x-update for prescribed controls and sums; the y, z updates and control consistency must also hold. Indeed, on each orbit of $r \mapsto r + V \pmod{32}$, one starting bit determines the others. Closure requires zero XOR of the right-hand-side bits along that orbit. Passing all g tests leaves 2^g solutions to the x equation; $V = 0$ requires a zero right-hand side. Full cycle existence additionally requires the other updates and consistent controls. At a fixed point the updates reduce to $x = R_v(x) \oplus [y + z]_M$, $y = [\phi(v + (x \oplus z))]_M$, and $x \oplus y = [z - R_u(z)]_M$. When $u = 0$, $x = y$. In the $u = v = 0$ subbranch, $z = [-x]_M$ and, for nonzero x with $t = \nu_2(x)$, the remaining condition implies $x \equiv -\phi 2^{t+1} \pmod{M}$. The two sides have different 2-adic valuations for $t < 31$, and the right side vanishes modulo M for $t = 31$, so only zero remains. The $v = 0, u = 16$ subbranch reduces to 2048 half-word candidates; exact enumeration excludes each one. The three listed states satisfy the fixed-point equations by substitution. Exhaustive 32-bit Z3 bit-vector queries for $u \neq 0$ and for fixed points outside the four listed states both returned `unsat`; witness replay and independent reduced encodings agreed. Hence these are exactly the four fixed points. For example, $N = (0, 0, 0)$ and $K = (\texttt{0x879D8102}, \texttt{0x879D8102}, \texttt{0x3502030B})$ initialize the third witness and produce a constant output. The baseline zero-preimage check therefore permits a period-one trajectory.

2.3 Length-dependent admission

For requested output length $L \ge 1$, set $H = D + L$ and $\mathcal{R} = P_0 \cup \mathcal{F}$, where $\mathcal{F}$ is the four-state fixed-point set. The revised procedure tests $S_0, \ldots, S_{2H}$ for visits to $\mathcal{R}$; it collects $O[j] = h(S_{D+j+1})$ for $0 \le j < L$ and checks that $S_{H+r} \neq S_H$ for $1 \le r \le H$. It releases no material upon failure. The additional H transitions are validation lookahead.

Proposition 1 *For a deterministic finite-state map, an anchor return $S_{H+r} = S_H$ with $1 \le r \le H$ occurs exactly when $\mu \le H$ and $p \le H$. Consequently, every accepted run has pairwise-distinct states $S_0, \ldots, S_H$.*

Proof When $\mu > H$, S_H is transient and cannot recur. Otherwise it lies on the cycle and first returns after p steps. If two states among $S_0, \ldots, S_H$ were equal, the trajectory would have entered a cycle of period at most H, forcing the rejected anchor return. □

Because h need not be injective, distinct states alone do not bound the output period. For the collected finite word O, define

$$\begin{aligned} q(O) = \min\{q \in \{1, \ldots, L\} : \\ O[j + q] = O[j] \quad (0 \le j < L - q)\}. \end{aligned} \tag{8}$$

The prefix-function π gives $q(O) = L - \pi[L - 1]$ by the longest prefix-suffix equality. Admission additionally requires $q(O) > \lfloor L/2 \rfloor$. This certifies the complete emitted block, including periods that do not divide L. It takes $2(D + L)$ transitions and $O(L)$ exact-word comparisons. When

Table 1 Nonzero fixed points of FSARX

v	$x = y$	z	Constant output
5	0x1860CE3A	0xFC183B3F	0x82C27403
5	0x1A98CE3A	0x2EE83B3F	0x50327520
9	0x879D8102	0x3502030B	0x45A31036

Table 2 Application-length admission results

Quantity	Result
Initializations tested and accepted	17
Discard D	256–511
Rejected-state visits	0
Anchor returns	0
Minimum finite period, each block	$q(O) = L$
Total transitions	860 108 388

the collected block starts on a state cycle, the same condition bounds its cycle-output period from below; otherwise it remains a finite-word guarantee. On a state cycle the fundamental output period p_O divides p, potentially strictly because h is not injective. If collection starts on the cycle, either $p_O \geq L$ or $q(O) \leq p_O < L$; acceptance therefore implies $p_O > \lfloor L/2 \rfloor$. A transient block retains the finite-word guarantee. At $L = 1$ the period threshold is vacuous, but fixed-state rejection still applies. Exhaustive functional graphs of reduced-word analogues ($w = 2, 4, 8$) have maximum cycle periods 5, 9, and 622, respectively (Fig. 1). These are exact for the reduced-width maps, not estimates of the 32-bit period. At the largest evaluated image size (4096^2 pixels, $a = 16$, six Feistel rounds), $L = 25\,296\,921$. Table 2 summarizes the full admission run. The reference start and 16 extended starts all passed: no rejected-state visit or anchor return occurred, and $q(O) = L$ for each output block. Exact unsigned arithmetic, full 96-bit state comparisons, and exact output-word comparisons were used. The checks instantiate both proven finite-window guarantees at the application length.

2.4 Generator diffusion diagnostics

Period and fixed-point results characterize recurrence, while perturbation experiments characterize how control words change over finite observation windows. After the N-dependent discard, 1024 sampled joint-secret initializations (seed 20260723) were each perturbed at all 96 state-bit positions. The paired trajectories were advanced for 64 steps, giving 98 304 one-bit trials. With $d_t = H(S_t, S'_t)$ and $d_0 = 1$, the mean distance reached 18.82 bits after one transition, 45.06 after three, and 47.98 after 64, compared with 48 bits for independent 96-bit states. The mean 32-bit output-word distance rose from 12.00 bits at step one to 15.99 at step 64; the independent-word reference is 16 bits. For M_s pairs of initial Hamming distance h, define the sampled one-step damage map and its local amplification coefficient by

$$\mathcal{D}(h) = \frac{1}{M_s} \sum_{\ell=1}^{M_s} H\big(\mathsf{T}(S_\ell), \mathsf{T}(S'_\ell)\big), \qquad \kappa = \mathcal{D}(1), \qquad H(S_\ell, S'_\ell) = h. \tag{9}$$

Using 98 304 pairs at each of $h \in \{1, 2, 4, 8, 16, 24, 32, 48\}$ gave $\mathcal{D}(1) = 18.8178$, $\mathcal{D}(8) = 43.1529$, and $\mathcal{D}(48) = 48.0042$ bits (Fig. 2). The plotted finite-time indicators are $t^{-1} \ln(\bar{d}_t/d_0)$ and the mean of $t^{-1} \ln(d_t/d_0)$ over trials; their decline reflects distance saturation and the $1/t$ factor. Fixed-rotation, XOR-only, and $\phi = 1$ controls reduced the one-step mean distance to 16.64, 15.95, and 8.85 bits, compared with 18.82 for the reference. These discrete diagnostics quantify early mixing over the sampled states.

For every initial bit i and observed bit j, the bit-diffusion probability was estimated across the same 1024 state pairs. Table 3 summarizes the mean probabilities $\bar{P}^S, \bar{P}^O$ and their mean absolute deviations (MAD) from 0.5 over all state- and output-bit pairs, respectively. Output-bit influence approached the unbiased-bit reference earlier. At $t = 8$, both means were within 0.001 of 0.5 and the MAD values were near the 0.0125 independent-sampling reference.

In a separate 1024-start experiment (seed 20260927; 4096 transitions after each discard), both controls visited all 32 values in every trajectory. Pooled selector entropies were 4.999996 and

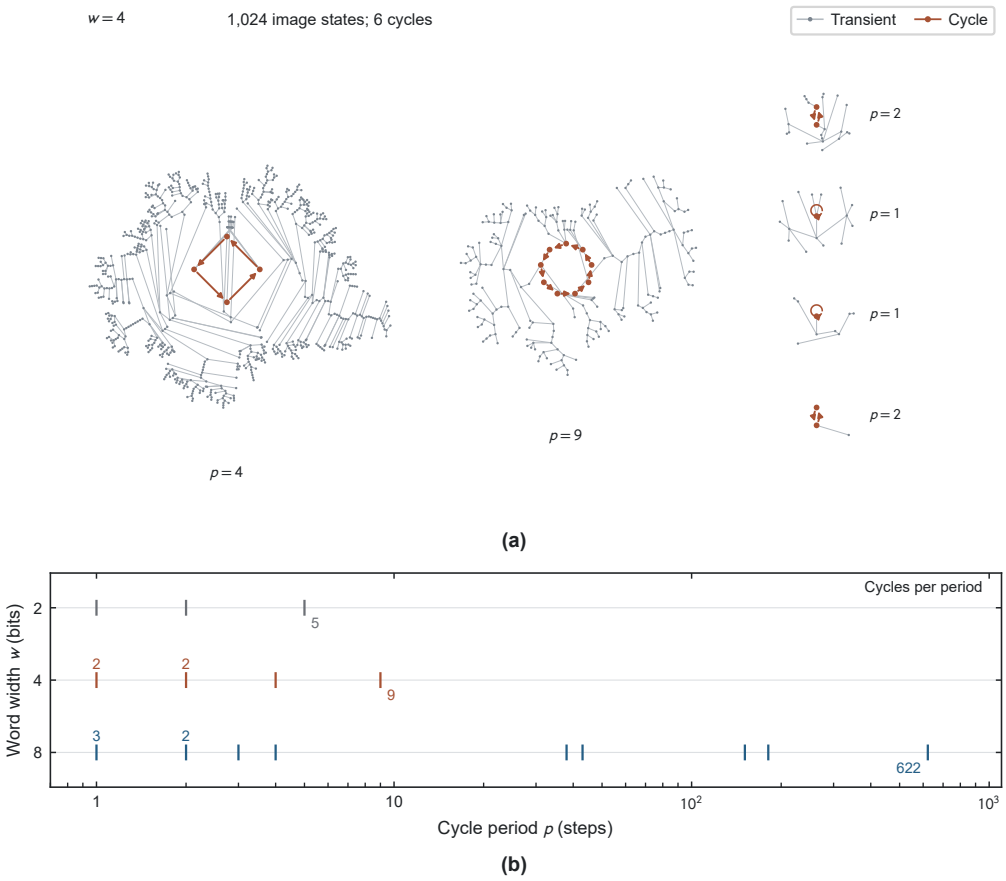

Fig. 1 Exact reduced-width functional graphs and cycle spectra. (a) The $w = 4$ image graph has six cycles with transient trees. (b) Cycle periods for $w = 2, 4, 8$; terminal labels give maxima of 5, 9, and 622

Table 3 State and output bit diffusion

t	$\bar{P}^S$	MAD_S	$\bar{P}^O$	MAD_O
1	0.1960	0.3272	0.3751	0.1700
3	0.4694	0.0368	0.4978	0.0130
8	0.5002	0.0125	0.4994	0.0126
64	0.4998	0.0126	0.4997	0.0124

4.999997 bits, with same-step mutual information 1.45×10^{-4} bits, documenting broad sampled branch coverage. Independent integer checks covered 10 000 states, 50 000 predecessor witnesses, and all 2048 candidates in the exceptional fixed-point branch. Admission checks covered all 288 maps on one to four nodes, all 8190 binary words of lengths one to twelve, the four fixed points, and 24 sampled nonzero-fixed-point predecessors. For accepted parameters, 10 000 control words matched the baseline and a 64×64 RGB integration case preserved ciphertext and exact recovery. The algebra supplies the general guarantee; these checks validate its implementation.

3 Reversible RGB Encryption

3.1 Control allocation and permutation

An $H_0 \times W_0$ 8-bit RGB image is reflection-padded to $H' \times W'$, with both dimensions divisible by block side a. Set $T = H'W'$, $B = T/a^2$, $q_0 = 4R + 2B$, and $m = \lceil (3T + 1)/6 \rceil$, where R is the Feistel round count. The generator supplies $L = q_0 + m + T$ words. Write the collected sequences as $X[j] = x_{D+j+1}$, $Y[j] = y_{D+j+1}$, $Z[j] = z_{D+j+1}$, and $O[j] = o_{D+j}$. The first $4R$ output words key the block Feistel rounds, the

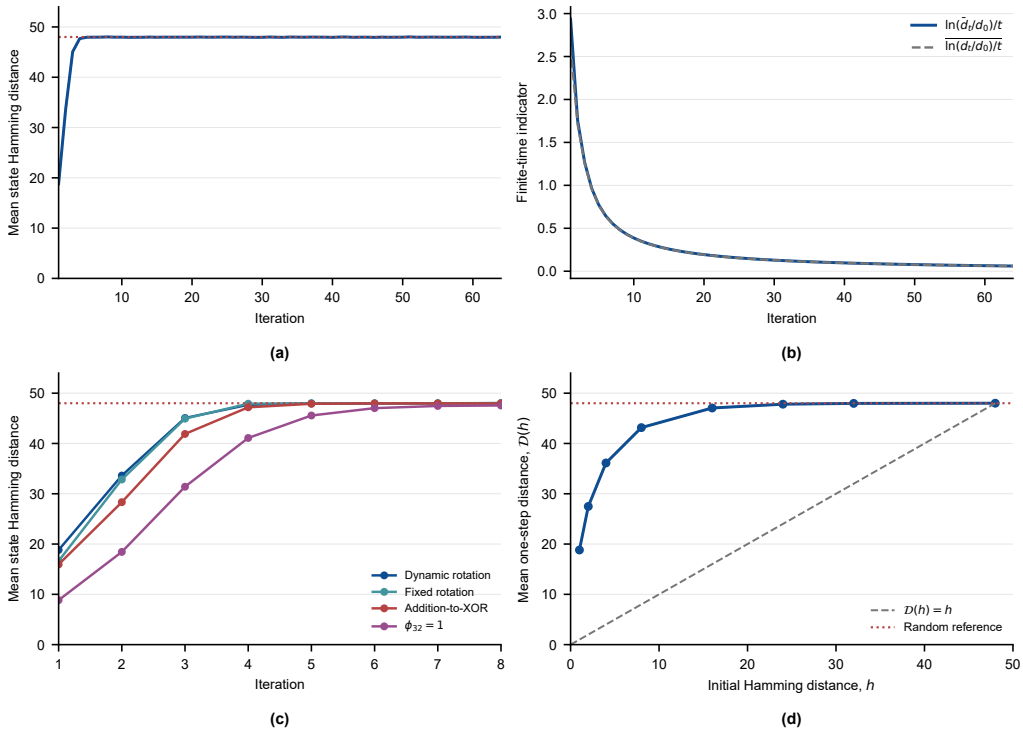

Fig. 2 Finite-state perturbation experiments. (a) Mean state Hamming distance over 98 304 one-bit trials; the dotted line is the independent-state reference. (b) Finite-time diffusion indicator. (c) Early-time component ablations. (d) One-step damage map for eight initial distances; bands are descriptive 95% intervals for sampled means

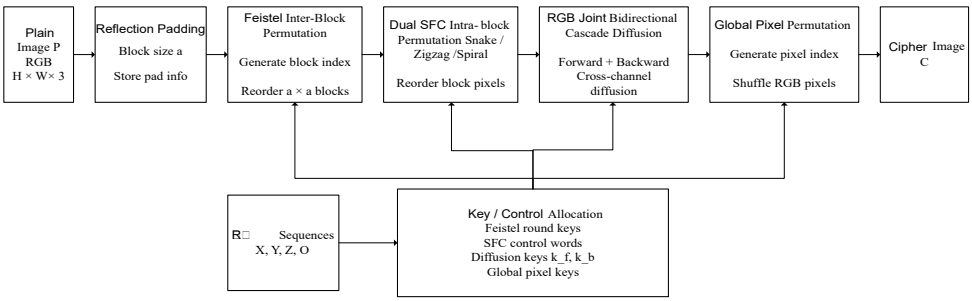

Fig. 3 FSARX encryption pipeline. Control segments from the finite-state generator drive block permutation, intra-block scan-path permutation, joint RGB diffusion, and the final global pixel permutation

next $2B$ select intra-block scan paths, the following m state triples supply diffusion bytes, and the last T state pairs supply global permutation keys. These intervals do not overlap. Figure 3 shows the processing order. For round j, four output words $s_{j,k} = O[4j+k]$ $(k = 0, \ldots, 3)$ define a Feistel function. An index is represented by two h-bit halves, with $h = \lceil \max(1, \lceil \log_2 B \rceil)/2 \rceil$. The round is

$$F_j(r) = \Big[R_{s_{j,2} \bmod 32}([(r \oplus s_{j,0}) + s_{j,1}]_M) \oplus s_{j,3} \oplus [(j+1)\phi]_M\Big]_M, \qquad (10)$$
$$(\ell', r') = (r, \ell \oplus (F_j(r) \bmod 2^h)).$$

Cycle walking restricts the resulting permutation to the B valid block indices [13]. Within block j, $\omega_j^r = O[4R + 2j]$ and $\omega_j^w = O[4R + 2j + 1]$ select read and write paths among snake,

diagonal-zigzag, spiral, and Morton order [14, 15]. Their base-path indices are $r_j = \omega_j^r \bmod 4$ and $w_j = (r_j + 1 + (\omega_j^w \bmod 3)) \bmod 4$, so they differ. Each base path is transformed by one of eight square symmetries, selected by $(\omega_j^r \gg 8) \bmod 8$ or $(\omega_j^w \gg 8) \bmod 8$, respectively. The mapped pixels satisfy $B''_j[P_w(t)] = B'_j[P_r(t)]$ for $0 \le t < a^2$. Both paths are bijections of the block coordinates. The Feistel network permutes its enclosing domain. Cycle walking follows each valid index to the next valid index on the same orbit, so the restricted map is bijective and terminates within 2^{2h} applications. The block assignment is $B'_{\pi(i)} = B_i$. Each SFC coordinate list separately visits every block position once; their composition is therefore invertible. Different base-path families and direction controls vary the orderings without affecting this reversibility argument.

3.2 Joint RGB diffusion and inversion

Serialize the permuted image in interleaved RGB raster order as $P = (p_0, \ldots, p_{n-1})$, $n = 3T$. The next state triples supply two $(n+1)$-byte streams K^F, K^B in big-endian byte order. For each state index $s = q_0 + i$, the forward stream takes $X_s[31:24], Y_s[31:24], Z_s[31:24], X_s[23:16], Y_s[23:16], Z_s[23:16]$; the backward stream takes $X_s[15:8], Y_s[15:8], Z_s[15:8], X_s[7:0], Y_s[7:0], Z_s[7:0]$. Each stream is truncated after $n+1$ bytes. The extra byte initializes its feedback chain, and $m = \lceil (n+1)/6 \rceil$ supplies enough material. With $\mathrm{pos}_j = j \bmod 256$, the forward and backward passes are

$$
\begin{aligned}
c_{-1} &= k_n^F, \quad r_j = (k_j^B \oplus j) \bmod 8, \\
\tau_j &= \mathrm{ROTL}_8(c_{j-1} \oplus k_j^F \oplus \mathrm{pos}_j, r_j), \\
c_j &= (p_j + k_j^F + \tau_j) \bmod 256, \\
f_n &= k_n^B, \quad \rho_j = (k_j^F \oplus j) \bmod 8, \\
\eta_j &= \mathrm{ROTL}_8(f_{j+1} \oplus k_j^B \oplus \mathrm{pos}_j, \rho_j), \\
f_j &= (c_j + k_j^B + \eta_j) \bmod 256.
\end{aligned}
\tag{11}
$$

The c_j pass runs forward and the f_j pass backward. Reshape $F = (f_j)$ into pixels. For pixel i, concatenate the allocated collected words into $g_i = (X[q_0 + m + i] \ll 32) \mid Y[q_0 + m + i]$ and stably sort (g_i, i); the secondary index fixes equal-key order. The resulting permutation π_g gives $I_c(i) = I_d(\pi_g(i))$. Decryption reverses the global permutation, subtracts the two diffusion contributions in reverse pass order, applies the inverse intra-block path assignment and inverse Feistel rounds, and removes reflection padding. In particular, $c_j = (f_j - k_j^B - \eta_j) \bmod 256$ and $p_j = (c_j - k_j^F - \tau_j) \bmod 256$. Each step is bijective under identical controls, so exact plaintext recovery follows. Compute c_j from $n-1$ down to zero using known f_{j+1}, then compute p_j from zero upwards using c_{j-1}. Inverse SFC mapping reads P_w and writes P_r; the original dimensions specify the final crop. Sorting by (g_i, i) makes the global inverse deterministic even for equal 64-bit keys.

4 Experiments and Analysis

4.1 Configuration and image statistics

Figures 4 and 5 show the representative visual and spatial-statistical results used in the following evaluation. The evaluation comprised 24 original RGB PNGs from the Kodak Lossless True Color Image Suite (768×512 or 512×768 pixels) and six additional images: airplane U-2, airplane, couple, house, peppers, and splash. All images retained their original resolutions. The reference settings were $a = 16$, $R = 6$, $K = (\texttt{0x12345678}, \texttt{0x9ABCDEF0}, \texttt{0x0F1E2D3C})$, and $N = (\texttt{0x11111111}, \texttt{0x22222222}, \texttt{0x33333333})$. Each image underwent fresh material generation, complete admission, encryption, and decryption; recovery was checked over every RGB component. Earlier six-image ablations and parameter sweeps are separate experiments whose timings exclude admission. Statistics used ciphertext arrays directly. Global Shannon entropy used 256-bin channel histograms. Expanded adjacent-pixel correlations used all valid horizontal, vertical, and diagonal pairs per channel. Local Shannon entropy (LSE) used up to 30 seeded, non-overlapping, grid-aligned 64×64 windows per channel, or all available windows when fewer existed; its finite-sample uniform reference is approximately 7.955 bits. Figures 4 and 5 illustrate the visual and spatial-statistical changes. On the first million bits of the 1024×1024 RGB ciphertext, the 17 reported NIST SP 800-22 statistics exceeded the threshold at $\alpha = 0.01$ [16] (Table 4). All 30 images were recovered exactly,

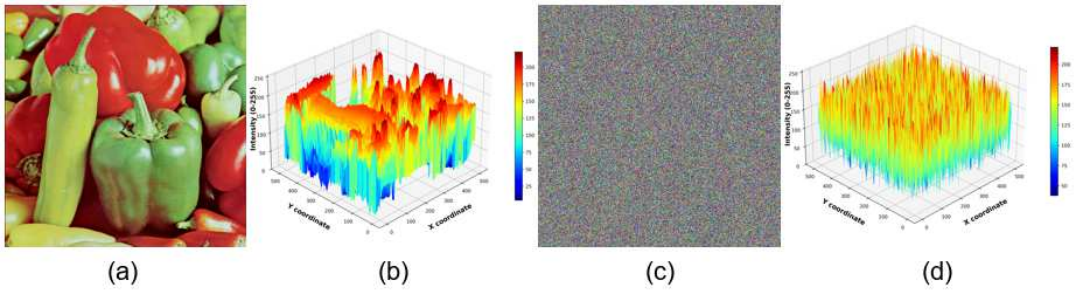

Fig. 4 The peppers test image: (a) plaintext, (b) plaintext intensity surface, (c) ciphertext, and (d) ciphertext intensity surface

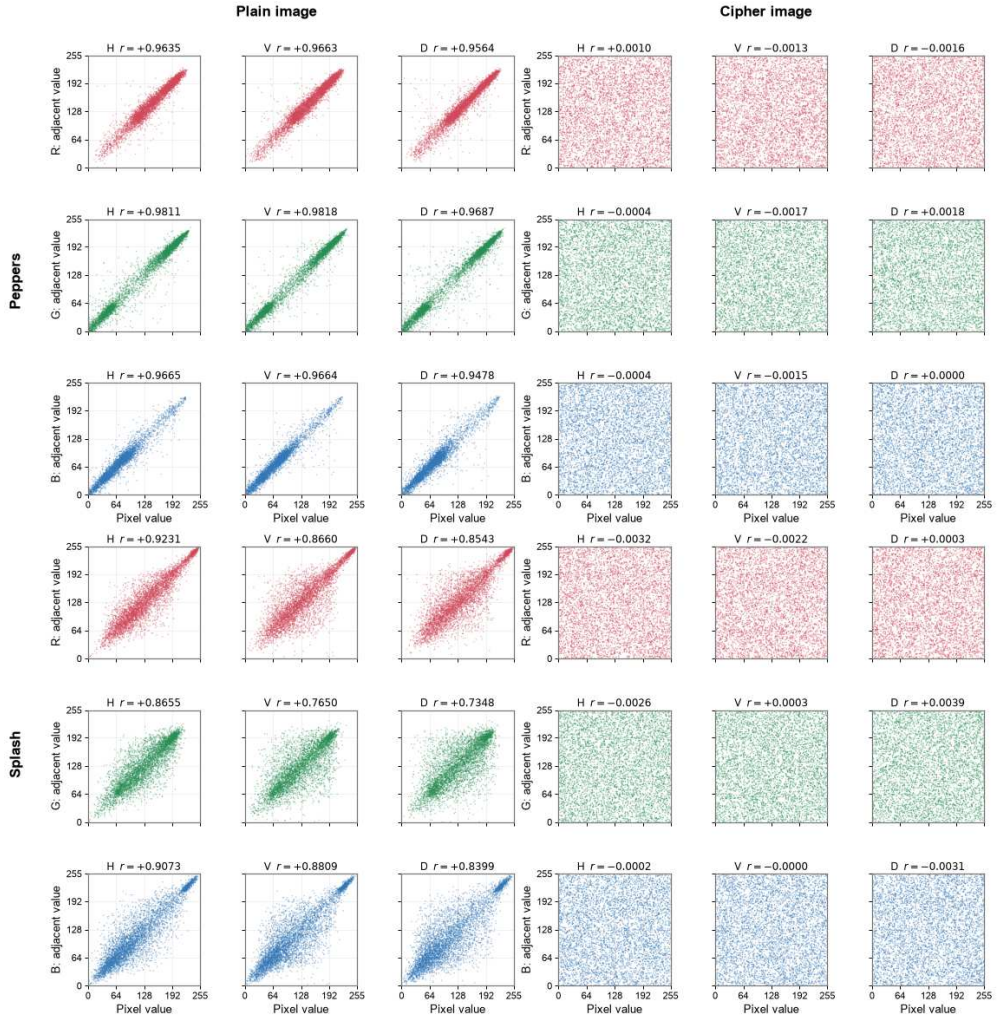

Fig. 5 Adjacent-pixel intensity-pair distributions for peppers and splash before and after encryption. The plotted pairs illustrate the channel-wise correlation measurements reported in the text

and every emitted material block had $q(O) = L$. Table 5 reports means ± sample SD across images after averaging channels and local windows within each image. The maximum absolute neighbor correlation over all images, channels, and directions was 0.011891. These results characterize marginal

Table 4 NIST SP 800-22 test results

Test	p-value
Frequency (monobit)	0.610051
Block frequency	0.243120
Runs	0.077776
Longest run of ones	0.797373
Binary matrix rank	0.441070
Discrete Fourier transform	0.659591
Non-overlapping template	0.641220
Overlapping template	0.657461
Maurer's universal	0.547185
Linear complexity	0.142012
Serial 1	0.220774
Serial 2	0.675732
Approximate entropy	0.712673
Cumulative sums, forward	0.789427
Cumulative sums, backward	0.399445
Random excursions, $x = -4$	0.778849
Random excursions variant, $x = -1$	0.971519

uniformity and low linear neighbor dependence under the fixed reference parameters.

4.2 Plaintext sensitivity and component ablation

The number of pixels change rate (NPCR) and unified average changing intensity (UACI) were computed jointly over all padded RGB ciphertext components. Their independent eight-bit references are 99.6094% and 33.4635%, respectively [17]. The earlier center-pixel R-channel experiment on six images yielded $99.7614 \pm 0.0136\%$ and $33.5236 \pm 0.0372\%$; the full-system row in Table 7 retains this comparison point. For two ciphertexts C_1, C_2 with $n = 3T$ components, the reported metrics were

$$\begin{aligned} \mathrm{NPCR} &= \frac{100}{n} \sum_{k=1}^{n} \mathbf{1}[C_1(k) \neq C_2(k)], \\ \mathrm{UACI} &= \frac{100}{255n} \sum_{k=1}^{n} |C_1(k) - C_2(k)|. \end{aligned} \tag{12}$$

The expanded experiment flipped one seeded random bit separately in R, G, and B at the center, top-edge midpoint, and one seeded random pixel: nine trials per image and 270 in total (seed 20261108 with image-index offsets). Each pair used identical parameters. Altered-image encryptions reused previously admitted material for the same shape and parameters only for these diagnostics, never for timing. Table 6 reports means and sample SD across images after averaging nine trials within each image. Aggregate means conceal perturbation dependence: flipping the center G-component most significant bit in airplane U-2 gave NPCR 54.4233% and UACI 18.2738%; the top-edge G-component counterpart in kodim24 gave 94.1323% and 31.6271%. Fresh, uncached complete-admission runs reproduced both cases. These edits reveal incomplete image-diffusion propagation, independently of the control source's finite-period guarantees. Five single-component ablations changed dynamic rotations to fixed values (A1), bypassed inter-block permutation (A2), used a single SFC read path and raster write path (A3), retained only forward diffusion (A4), or removed the final global permutation (A5). Other controls were held fixed. Ablation statistics are means ± sample SD over six images, except the maximum correlation. Runtime excludes admission and uses 512×512 peppers, three warm-ups, and ten measured runs. Table 7 shows that A4 reduced mean NPCR and UACI to 60.9182% and 20.4529%, isolating the backward pass's role in opposite-direction propagation. A5 leaves NPCR and UACI unchanged because a position permutation preserves the component difference count. A1 yielded higher NPCR but moved farther from its random reference, showing why NPCR alone is insufficient to rank variants.

4.3 Cost, parameter sensitivity, and corruption

For fixed a and R, sequence generation, pixel scans, and diffusion require $O(T)$ work; comparison-based sorting of the global pixel keys costs $O(T \log T)$ time and $O(T)$ storage. Enumerating the cycle-walked block permutation visits at most the $2^{2h} = O(B)$ enclosing indices per round in aggregate, giving $O(BR)$ work. The full algorithm therefore takes $O(T \log T)$ time at fixed a, R. Since $L = 4R + 2T/a^2 + \lceil (3T+1)/6 \rceil + T$, complete admission adds $O(T)$ work and auxiliary storage at fixed a, R. Complete timing used Python 3.12.14 and NumPy 2.3.5. Seeded RGB arrays had side lengths 256–4096 (seed 20261108 plus side length). Each size/implementation used a fresh serial process, three warm-up pairs, and ten measured encryption/decryption pairs, with implementation order alternating by size. Timing

Table 5 Original-resolution ciphertext statistics

Dataset	Exact recovery	Entropy (bits)	LSE (bits)
Kodak24	24/24	7.999529 ± 0.000023	7.954520 ± 0.000296
Original six	6/6	7.998679 ± 0.001220	7.954634 ± 0.000349
All 30	30/30	7.999359 ± 0.000614	7.954543 ± 0.000304

Table 6 Multi-position plaintext sensitivity

Dataset	Trials	NPCR (%)	UACI (%)
Kodak24	216	99.7472 ± 0.1275	33.5119 ± 0.0452
Original six	54	98.9745 ± 2.0680	33.2537 ± 0.7025
All 30	270	99.5927 ± 0.9214	33.4602 ± 0.3127

included padding/cropping, fresh material generation, complete admission, and all image transformations. Tables 8 and 9 report runtime means ± sample SD over ten measured runs. Throughput in Table 8 is $3WH/(2^{20}t_{\text{enc}})$ MiB/s, using mean guarded encryption time t_{enc}. An interrupted 4096^2 baseline batch was archived and excluded before a complete restart. Complete guarded encryption took 21.15–26.76% longer than baseline (Table 8). All five sizes produced identical ciphertext SHA-256 values across implementations, and all 130 warm-up/measured decryptions recovered the inputs exactly. Because material-collection code differs, the increase compares complete implementations rather than isolating one validation switch. At 4096^2, material generation and admission accounted for 50.3013 ± 0.3176 s of 79.6262 ± 0.3304 s total encryption: collection/inline checks took 30.2684 s, lookahead 16.2892 s, and finite-word period testing 3.7437 s (ten-run means). Peak process working set was 1240.1 MiB, including interpreter, inputs, temporaries, and verification data. These measurements establish the substantial preprocessing cost for large still images. A one-factor-at-a-time parameter sweep varied the block side $a \in \{8, 16, 32, 64\}$ at $R = 6$ and the Feistel rounds $R \in \{2, 4, 6, 8, 10\}$ at $a = 16$. Across the block-size sweep, mean NPCR and UACI ranged from 99.7118–99.7745% and 33.4745–33.5345%, respectively. Across the round sweep, the corresponding ranges were 99.7272–99.9084% and 33.5108–33.5697%. Local entropy and adjacent-pixel correlation remained near their reference levels, and no tested setting improved all reported metrics monotonically. The reference $a = 16, R = 6$ configuration is therefore a common comparison point, not an inferred optimum. Ciphertext-corruption tests characterize recovery behavior separately from confidentiality. Masking square ciphertext regions of side ratios 0.25 and 0.50 produced degraded but still visually distinguishable reconstructions, with greater loss at the larger ratio. Salt-and-pepper corruption at probabilities 0.001, 0.01, and 0.1 yielded mean decrypted-image PSNRs of 36.492, 26.225, and 16.374 dB across the two recorded images, with mean pixel-consistency rates of 99.703%, 97.013%, and 72.908%. These measurements describe the tested error conditions and do not substitute for chosen-plaintext, related-key, or authenticated-encryption analysis.

4.4 Common-environment implementation comparison

We evaluated the public implementation associated with Huang et al. [18] and a third-party serial Python reproduction of Gao et al. [19]. All implementations used the CPU environment specified above, without multiprocessing or a GPU. Huang's interface was adapted to in-memory RGB arrays, with plotting and printing disabled and numerical compatibility restored. Core loops and parameters were retained. Statistical comparisons used the same central 256×256 RGB crop from each of 30 images, without interpolation. Every method recovered all 30 crops exactly. Each image received nine identical position/channel/bit perturbations across methods, with effective parameters held fixed within a pair. A separate experiment re-derived Gao's image-dependent parameters for the altered image; these results were kept distinct. This protocol differs from the original-resolution evaluation. In Table 10, entropy is averaged over channels and images. NPCR and UACI are means ± sample SD across 30 image means, each averaging nine edits; the minimum is over all 270 edits. Timing used identical seeded 256^2 and 512^2 RGB arrays, three warm-up pairs and ten measured

Table 7 Component ablation results

Variant	NPCR (%)	UACI (%)	LSE (bits)	max $\lvert r \rvert$	Runtime (s)
A1	99.9678 ± 0.0112	33.6066 ± 0.0387	7.954686 ± 0.004083	0.007378	0.9880 ± 0.0151
A2	99.8132 ± 0.0103	33.5252 ± 0.0599	7.954427 ± 0.004150	0.007723	1.0176 ± 0.0119
A3	99.7579 ± 0.0229	33.5327 ± 0.0542	7.954448 ± 0.003970	0.008363	1.0230 ± 0.0362
A4	60.9182 ± 3.1105	20.4529 ± 1.0418	7.954639 ± 0.003842	0.006000	0.8749 ± 0.0116
A5	99.7614 ± 0.0136	33.5236 ± 0.0372	7.954603 ± 0.004347	0.007405	0.9899 ± 0.0079
Full FSARX	99.7614 ± 0.0136	33.5236 ± 0.0372	7.954606 ± 0.003851	0.011891	1.0167 ± 0.0135

Table 8 Complete runtime and admission overhead

Size	Baseline encryption (s)	Guarded encryption (s)	Guarded decryption (s)	Increase (%)	MiB/s
256^2	0.2128 ± 0.0028	0.2691 ± 0.0013	0.2748 ± 0.0036	26.46	0.6969
512^2	0.8576 ± 0.0042	1.0870 ± 0.0081	1.1131 ± 0.0148	26.76	0.6899
1024^2	3.4981 ± 0.0263	4.3877 ± 0.0453	4.4534 ± 0.0334	25.43	0.6837
2048^2	14.9457 ± 0.7175	18.5855 ± 0.1632	18.6606 ± 0.1694	24.35	0.6457
4096^2	65.7267 ± 0.5240	79.6262 ± 0.3304	80.1174 ± 0.2728	21.15	0.6028

Table 9 Common-environment runtime comparison

Implementation	256^2 Encryption (s)	256^2 Decryption (s)	512^2 Encryption (s)	512^2 Decryption (s)
FSARX, complete admission	0.2701 ± 0.0031	0.2734 ± 0.0029	1.0850 ± 0.0072	1.1023 ± 0.0146
Huang et al., public code	1.4391 ± 0.0571	1.4493 ± 0.0398	5.4888 ± 0.0395	5.5599 ± 0.0531
Gao et al., third-party serial	0.0902 ± 0.0016	0.0631 ± 0.0008	0.4115 ± 0.0043	0.2800 ± 0.0019

Table 10 Fixed-parameter comparison on common crops

Implementation	Entropy (bit)	NPCR (%)	UACI (%)	Minimum NPCR (%)
FSARX, complete admission	7.99717	99.821 ± 0.039	33.535 ± 0.024	99.620
Huang et al., public code	7.99720	86.891 ± 6.483	32.029 ± 1.063	50.000
Gao et al., third-party serial	7.99722	35.271 ± 2.603	11.474 ± 1.437	0.084

pairs per fresh process. Complete material generation and FSARX admission were included; disk I/O, startup/imports and result verification were excluded. No material was cached. Table 9 reports implementation-specific cost, rather than an intrinsic speed ranking of algorithms. The measurements place FSARX between the two selected implementations in runtime. Table 10 characterizes the fixed-parameter protocol. With its default image-dependent parameter re-derivation, the Gao reproduction yielded NPCR $99.6081 \pm 0.0040\%$ and UACI $33.4641 \pm 0.0249\%$ over 30 image means. This separate protocol combines parameter changes with plaintext propagation and is not pooled with the fixed-parameter results.

5 Conclusion

FSARX couples a finite-state ARX generator to reversible RGB permutation and diffusion. Its one-step image has 2^{69} states and exactly four fixed points, including three nonzero degeneracies excluded by admission. Every accepted run has distinct full states over the requested prefix and finite-word output period $q(O) > \lfloor L/2 \rfloor$; 17 application-length runs attained $q(O) = L$. Thirty-image experiments confirmed exact recovery and $q(O) = L$ for every emitted block. Complete guarded encryption took 1.0870 s at 512^2 and 79.6262 s at 4096^2 pixels, with a measured 21–27% increase relative to the baseline implementation. Under a separate matched-input CPU protocol, guarded FSARX encrypted 512^2 RGB arrays in 1.0850 ± 0.0072 s. This was 5.06 times faster than the selected Huang public implementation but took 2.64 times as long as the Gao third-party serial reproduction. These measurements characterize specific implementations and protocols. Multi-position perturbations exposed low-response edits despite near-reference aggregate differential means; addressing this diffusion dependence remains necessary. Future research will examine equivalent-key classes, pursue independent cryptanalysis, and develop authenticated-encryption extensions, building on the established guarantees for accepted finite windows.

Statements and Declarations

Funding. This work was supported in part by the National Natural Science Foundation of China under Grant 62471406.

Competing interests. The authors declare no competing financial or non-financial interests.

Data availability. No data was used for the research described in the article.

References

[1] Fridrich, J.: Symmetric ciphers based on two-dimensional chaotic maps. International Journal of Bifurcation and Chaos **8**(06), 1259–1284 (1998)

[2] Shannon, C.E.: Communication theory of secrecy systems. The Bell system Technical Journal **28**(4), 656–715 (1949)

[3] SaberiKamarposhti, M., Ghorbani, A., Yadollahi, M.: A comprehensive survey on image encryption: Taxonomy, challenges, and future directions. Chaos, Solitons & Fractals **178**, 114361 (2024)

[4] Liu, P., Teng, L., Iu, H.H.-C., Wang, M., Li, Q., Fu, X.: High sensitivity image encryption algorithm based on cascaded chaotic system. Journal of Information Security and Applications **93**, 104153 (2025)

[5] Shahna, K.U., Mohamed, A.: Novel hyper chaotic color image encryption based on pixel and bit level scrambling with diffusion. Signal Processing: Image Communication **99**, 116495 (2021)

[6] Feng, W., Yang, J., Zhao, X., Qin, Z., Zhang, J., Zhu, Z., Wen, H., Qian, K.: A novel multi-channel image encryption algorithm leveraging pixel reorganization and hyperchaotic maps. Mathematics **12**(24), 3917 (2024)

[7] Wen, H., Kang, S., Wu, Z., Lin, Y., Huang, Y.: Dynamic RNA coding color image cipher based on chain feedback structure. Mathematics **11**(14), 3133 (2023)

[8] Safdar, M.U., Shah, T., Ali, A.: A color image encryption scheme hybridizing two finite rings and DNA encoding. Signal, Image and Video Processing **19**(13), 1130 (2025) https://doi.org/10.1007/s11760-025-04739-8

[9] Song, W., Fu, C., Lin, Z., Zhang, Y., Chen, J., Sham, C.-W.: Batch medical image encryption using 3D Latin cube-based simultaneous permutation and diffusion. Signal, Image and Video Processing **18**(3), 2499–2508 (2024) https://doi.org/10.1007/s11760-023-02925-0

[10] Elmanfaloty, R.A., Abou-Bakr, E.: Random property enhancement of a 1d chaotic prng with finite precision implementation. Chaos, Solitons & Fractals **118**, 134–144 (2019) https://doi.org/10.1016/j.chaos.2018.11.019

[11] Wang, S.H., Liu, W.R., Lu, H.P., Kuang, J.Y., Hu, G.: Periodicity of chaotic trajectories in realizations of finite computer precisions and its implication in chaos communications. International Journal of Modern Physics B **18**(17–19), 2617–2622 (2004) https://doi.org/10.1142/S0217979204025798

[12] Li, C., Feng, B., Li, S., Kurths, J., Chen, G.: Dynamic analysis of digital chaotic maps via state-mapping networks. IEEE Transactions on Circuits and Systems I: Regular Papers **66**(6), 2322–2335 (2019) https://doi.org/10.1109/TCSI.2018.2888688

[13] Black, J., Rogaway, P.: Ciphers with arbitrary finite domains. In: Topics in Cryptology–CT-RSA 2002. Lecture Notes in Computer Science, vol. 2271, pp. 114–130. Springer, Berlin, Heidelberg (2002)

[14] Morton, G.M.: A Computer Oriented Geodetic Data Base and a New Technique in File Sequencing. International Business Machines Company, Ottawa, Canada (1966)

[15] Sagan, H.: Space-Filling Curves. Universitext. Springer, New York (1994)

[16] National Institute of Standards and Technology: A Statistical Test Suite for Random and Pseudorandom Number Generators for Cryptographic Applications. National Institute of Standards & Technology, Gaithersburg, MD (2010)

[17] Wu, Y., Noonan, J.P., Agaian, S.: NPCR and UACI randomness tests for image encryption. Journal of Selected Areas in Telecommunications, 31–38 (2011)

[18] Huang, Y., Zhang, Q., Zhao, Y.: Color image encryption algorithm based on hybrid chaos and layered strategies. Journal of Information Security and Applications **89**, 103921 (2025) https://doi.org/10.1016/j.jisa.2024.103921

[19] Gao, S., Zhang, Z., Iu, H.H.-C., Ding, S., Mou, J., Erkan, U., Toktas, A., Li, Q., Wang, C., Cao, Y.: A parallel color image encryption algorithm based on a 2-d logistic-rulkov neuron map. IEEE Internet of Things Journal **12**(11), 18115–18124 (2025) https://doi.org/10.1109/jiot.2025.3540097