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.

From hyperchaotic expansion to stream-cipher failure: observability bounds, finite-precision dynamics and cryptanalysis of a coupled torus map

Annas Malik

annas.waseem@ucp.edu.pk

University of Central Punjab

Research Article

Keywords:

Posted Date: November 26th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11569236/v1

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

Additional Declarations: No competing interests reported.

From hyperchaotic expansion to stream-cipher failure: observability bounds, finite-precision dynamics and cryptanalysis of a coupled torus map

Annas W. Malik [1*]

[1*]Faculty of Information Technology and Computer Science, University of Central Punjab, Lahore, Pakistan.

Corresponding author(s). E-mail(s): annas.waseem@ucp.edu.pk;

Abstract

Positive Lyapunov exponents are often cited as evidence that a chaotic map yields secure keystreams. We test this inference on a four-dimensional sine-coupled expanding torus map (SCUT-4D) whose hyperchaos is proved: for $\mathbf{0 \le \varepsilon < 1/\pi}$ every exponent exceeds $\mathbf{\ln(3 - 2\pi\varepsilon) > 0}$, and the map is topologically conjugate to an uncoupled linear endomorphism with entropy $\mathbf{\ln 945}$. For its integer version, which outputs only the top bits of each coordinate, we prove how fast the hidden state can be inferred. Without coupling, the uncertainty about the state falls by $\mathbf{\log_2 945 \approx 9.88}$ bits per observed step, up to a bounded error. For every $\mathbf{0 \le \varepsilon < 1/\pi}$, interval arithmetic confines the state to a box whose width shrinks by the factor $\mathbf{3 - 2\pi\varepsilon}$ per step, the same constant that bounds the exponents. The expansion that proves hyperchaos therefore also sets how fast the state is revealed. An attack built on this recovered the unique state in all 150 trials, including 128-bit states from 80–128 known bytes, in at most 0.17 s; a generic guess-and-determine attack gains only $\mathbf{2^4}$–$\mathbf{2^6}$ over exhaustive search of the hidden bits. Exact functional graphs and cycle sampling up to 64-bit states show many small basins and cycles several times shorter than for keyed random maps, by an amount that depends on how the coupling table is rounded. Backward search finds the $\mathbf{\approx 500}$ equivalent keys that random-mapping theory predicts for a 1000-step warm-up, and session states collide about $\mathbf{2^{9.7}}$ times more often than uniform. Proven hyperchaos guarantees that nearby states separate; it does not make the output hard to predict.

Keywords: Hyperchaos, Lyapunov exponents, Expanding torus maps, Observability, Finite-precision dynamics, Chaos-based cryptography

1 Introduction

Chaos-based encryption has been studied for almost three decades [1–3]. Sensitive dependence on initial conditions, mixing and broadband spectra resemble the confusion and diffusion that Shannon required of ciphers [4], and the literature has grown into a steady stream of image, speech and video schemes built on new chaotic and hyperchaotic systems; recent examples in this journal include memristive conservative hyperchaos with scrambling [5], quantum bit-level image cryptosystems driven by two-dimensional hyperchaotic maps [6], DNA-coded encryption from modified Sprott systems [7], robust-chaos speech encryption [8] and image secure-communication systems that combine compressed sensing with encryption [9]. Many reuse compact maps such as the 2D logistic-adjusted-sine map [10]. The usual evaluation measures information entropy, adjacent-pixel correlation, NPCR/UACI, key sensitivity and statistical test suites. Cryptanalysts have repeatedly warned that these statistics describe how ciphertext looks, not how hard it is to predict [11–13], many published schemes have been broken by known- or chosen-plaintext attacks [14–19], and Everett [20] recently gave an efficient interval-refinement attack on stream ciphers built from a single chaotic map of the interval, together with a complexity-theoretic formulation of the problem.

A second line of work concerns digital degradation. A chaotic map realised with b-bit arithmetic is a map on a finite set, so every orbit is eventually periodic. Round-off-induced periodicity [21], the degradation of piecewise-linear maps [22], the periodicity of floating-point logistic generators [23] and the description of digital chaos as a state-mapping network [24, 25] show that finite-state dynamics can differ qualitatively from the real dynamics; remedies such as higher-dimensional digital constructions [26] and perturbation-based anti-degradation schemes [27] remain active topics in this journal. Analytical tools such as Jacobi stability, invariant sets and hidden-attractor analysis [28] are increasingly used to characterise real dynamics, but they are seldom carried through to the digital implementation that a cipher actually runs.

The two lines rarely meet at the level of the dynamical quantity used to justify security. The Lyapunov spectrum is the most widely quoted certificate of "hyperchaos", yet it measures the asymptotic growth of infinitesimal perturbations of the real system [29]. Four gaps follow. (i) Exponents are almost always estimated numerically, and the coupling terms that are said to create hyperchaos are rarely shown to matter. (ii) The finite-state map that produces the keystream is seldom specified exactly enough (arithmetic, rounding, tables, update order) to be reproduced. (iii) How much an observer of truncated output learns about the hidden state per step is not quantified in dynamical terms. (iv) Security is inferred from statistics rather than tested against an explicit adversary.

This paper addresses these gaps for one system that is simple enough to analyse. We do not propose a cipher. We construct a four-dimensional torus map with sine coupling (working name SCUT-4D) whose hyperchaos can be proved, define its integer version bit for bit, and analyse the stream cipher that a typical design would build from it. The contributions are:

- **Proven real dynamics.** A uniform singular-value bound gives four exponents of at least $\ln(3 - 2\pi\varepsilon) > 0$ for $0 \le \varepsilon < 1/\pi$; topological conjugacy to $\mathbf{D} = \mathrm{diag}(3, 5, 7, 9)$

gives $h_{\mathrm{top}} = \ln 945$, 384 fixed points and an upper bound on the exponent sum (Sect. 3). Conjugacy does not fix the exponents; numerically they agree with $\ln d_i$ to within the estimation error (Sect. 6.1).

- **How fast the state is revealed.** For the integer map with truncated output we prove upper and lower bounds on the remaining uncertainty: without coupling, hidden bits disappear at $\log_2 945$ bits per step (Proposition 6); for all $0 \le \varepsilon < 1/\pi$, interval arithmetic shrinks the uncertainty by the factor $3 - 2\pi\varepsilon$ per step (Proposition 7). For this system the claim that expansion reveals the state is therefore proved rather than assumed, and the interval-refinement idea of [20] acquires explicit rates for a coupled four-dimensional map.
- **Attacks with measured cost.** An interval attack recovers 64- to 128-bit states in polynomial time. A generic guess-and-determine attack serves as a baseline, and its cost is split into what truncation exposes for any generator and what the dynamics add. Equivalent keys, collisions between sessions and a formal distinguishing game are quantified (Sects. 6.4–6.5).
- **Finite-state dynamics compared with random controls.** Exact functional graphs, basin statistics and cycle sampling up to 64-bit states are compared with keyed random functions and random permutations, and with other ways of rounding the coupling (Sect. 6.2).

All parameters used here are small research instances chosen so that they can be analysed, and the attacks target only the SCUT-4D interfaces of Sect. 4; we make no claim about any other published scheme, including those cited above. Nonce reuse and malleability are generic properties of unauthenticated XOR streams and are reported only as controls.

2 Related work and research gap

Requirements for chaos-based cryptosystems were systematised by Alvarez and Li [11]: an exactly specified implementation, a key space described in bits, and resistance to known- and chosen-plaintext attacks. Preishuber et al. [12] showed empirically that many chaos-based image and video ciphers are slower than AES while offering no security advantage, and that common metrics do not separate secure from insecure designs; Özkaynak [13] reviewed the same issues for this journal, and Dhall and Yadav [19] recently reviewed cryptanalyses of substitution–permutation image ciphers. Specific cryptanalyses, for example of Fridrich-type permutation–diffusion ciphers [14], logistic-map image ciphers [15], pixel-bit scrambling [16] and plaintext-related schemes [17, 18], show that the weakness that breaks a scheme is usually structural rather than statistical. The closest prior work is Everett's interval-refinement attack [20], which tracks the set of seeds consistent with the symbolic output of a unimodal interval map on a dyadic grid and casts chaos-based generators in modern cryptographic language.

On the dynamical side, degradation of digital chaotic maps has been characterised through cycle statistics [22, 23] and functional graphs [24, 25]; random-mapping theory [30] provides the natural null model. Expanding maps are among the best understood chaotic systems: expanding maps of tori are topologically conjugate to linear ones [31], they possess a unique absolutely continuous invariant measure [32], and their entropy

Table 1 Positioning against prior analyses of chaos-based ciphers and digital chaos

Work	Object	Output / attack model	Finite-precision model	Guarantee or result
Preishuber et al. 2018 [12]	Chaos-based image and video ciphers (empirical survey)	Statistical metrics, known attacks	As published per scheme	Empirical; metrics do not separate secure from insecure designs
Li et al. 2019 [24]	Digitised chaotic maps	None (dynamics only)	Fixed-point state-mapping networks	Network structure of finite-precision maps
Everett 2024 [20]	PRGs from one chaotic map of the interval	Symbolic output, known plaintext	Dyadic fixed point, interval analysis	Efficient interval-refinement attack; complexity-theoretic framing
Zhou & Yu 2024 [18]	A plaintext-related chaotic image cipher	Divide-and-conquer cryptanalysis	As specified by the attacked scheme	Attack on one published scheme
Dhall & Yadav 2024 [19]	SPN-based image ciphers	Systematic review of cryptanalyses	–	Review
Shen et al. 2026 [27]	Anti-degradation for digital chaos	Statistical and dynamical evaluation	64-bit phase accumulator, exact period 2^{64}	Period guarantee; no adversarial model
This work	Coupled 4-D expanding torus map with proven exponents	Truncated or raw output, known prefix; games	Frozen integer tables, $b \le 32$; float variants	Proven observability rates (Props. 6, 7); polynomial-time state recovery

obeys Ruelle's inequality and Pesin's formula [33, 34]. These results are standard [35]. We use them, together with the exact integer version of the map, to connect a provable expansion rate with how quickly the state can be inferred from output and how much it costs to recover it. Table 1 positions the paper against the closest analyses, and Table 2 contrasts the evaluation dimensions of representative recent schemes with this work.

3 The SCUT-4D map and its real dynamics

3.1 Definition

Let $\mathbb{T}^4 = (\mathbb{R}/\mathbb{Z})^4$ with coordinates $x = (x_1, \dots, x_4)$, multipliers $d = (3, 5, 7, 9)$ and coupling $\varepsilon \ge 0$. Indices are cyclic ($x_5 \equiv x_1$). The map $F_\varepsilon : \mathbb{T}^4 \to \mathbb{T}^4$ is

$$x_i(n+1) = \big[d_i x_i(n) + \varepsilon \sin\big(2\pi x_{i+1}(n)\big)\big] \bmod 1, \qquad i = 1, \dots, 4. \tag{1}$$

Each coordinate is an integer-multiplier circle map driven by its cyclic successor, a template common in coupled-map designs. Away from the identification points the

Table 2 Evaluation dimensions of representative chaos-based encryption studies and of this work

Study	Dynamical evidence	Finite-precision model	Adversarial tests reported	Authentication
Alexan et al. 2024 [36]	Fractional-order hyperchaotic Chen system	Not specified at bit level	Entropy, NPCR/UACI, NIST suite, throughput	Not reported
Zhang et al. 2026 [7]	Equilibria, bifurcation, LE spectra	Discretised system (Sect. 2.5 of [7])	Differential, noise/occlusion, key sensitivity	Not reported
Huang et al. 2026 [8]	Robust chaos, complexity measures	Not specified at bit level	Key sensitivity, "attack resistance", statistics	Not reported
Shen et al. 2026 [27]	Max. LE, phase-space reconstruction	64-bit phase accumulator, exact period 2^{64}	Dynamical, security and randomness evaluation	Not reported
This work	Proven bounds, conjugacy, QR spectra, observability theorems	Frozen integer tables, $b = 3$–32, hashes	State and key recovery, games, collisions	AEAD baseline

Entries summarise each paper's abstract and section headings as published; "not reported" means not identified there, not that the property is absent. No attack on these schemes is claimed.

Jacobian is

$$DF_\varepsilon(x) = \mathbf{D} + \mathbf{E}(x), \qquad \mathbf{D} = \operatorname{diag}(3,5,7,9), \qquad E_{i,i+1}(x) = 2\pi\varepsilon\cos(2\pi x_{i+1}), \quad (2)$$

with all other entries of $\mathbf{E}$ zero. The determinant is $\det DF_\varepsilon(x) = 945 - (2\pi\varepsilon)^4 \prod_j \cos(2\pi x_j)$, because the only non-diagonal permutation contributing is the 4-cycle, whose sign is -1.

3.2 Uniform expansion

Proposition 1 *For $0 \le \varepsilon < 1/\pi$ and every $x \in \mathbb{T}^4$, $v \in \mathbb{R}^4$,*

$$(3 - 2\pi\varepsilon)\,\|v\| \;\le\; \|DF_\varepsilon(x)v\| \;\le\; (9 + 2\pi\varepsilon)\,\|v\|.$$

Hence every Lyapunov exponent of every orbit, where it exists, lies in $[\ln(3-2\pi\varepsilon), \ln(9+2\pi\varepsilon)]$, and all four are positive.

Proof Write

$$\mathbf{E}(x) = 2\pi\varepsilon\,\mathbf{C}(x)\mathbf{P}, \qquad \mathbf{C}(x) = \operatorname{diag}(\cos 2\pi x_2, \cos 2\pi x_3, \cos 2\pi x_4, \cos 2\pi x_1),$$

where $\mathbf{P}$ is the cyclic permutation matrix ($P_{i,i+1} = 1$). Since $\mathbf{P}$ is orthogonal and $\|\mathbf{C}\|_2 \le 1$, $\|\mathbf{E}(x)\|_2 \le 2\pi\varepsilon$. Weyl's inequality for singular values gives $\sigma_{\min}(\mathbf{D}+\mathbf{E}) \ge \sigma_{\min}(\mathbf{D}) - \|\mathbf{E}\|_2 \ge 3-2\pi\varepsilon$ and $\sigma_{\max}(\mathbf{D}+\mathbf{E}) \le 9+2\pi\varepsilon$. Iterating, $\|DF_\varepsilon^n(x)v\| \ge (3-2\pi\varepsilon)^n\|v\|$, so every exponent is at least $\ln(3 - 2\pi\varepsilon)$, which is positive for $\varepsilon < 1/\pi$. □

At $\varepsilon = 0.1$ the bound is $\ln 2.3717 = 0.8636$ nats/step. The bound is conservative: the global minimum of $\sigma_{\min}(DF_{0.1})$ over the torus, found by grid search and Nelder–Mead refinement, is 2.9545 (Fig. 2c). The same constant $3 - 2\pi\varepsilon$ reappears as the contraction rate of an observer's uncertainty in Proposition 7.

3.3 Topological structure

Proposition 2 *For $0 \le \varepsilon < 1/\pi$, F_ε is a C^∞ expanding endomorphism of $\mathbb{T}^4$ topologically conjugate to the linear map $L_{\mathbf{D}} : x \mapsto \mathbf{D}x \bmod 1$. Consequently F_ε has degree 945, topological entropy $h_{\mathrm{top}}(F_\varepsilon) = \ln 945 \approx 6.8512$ nats/step, exactly $|\det(\mathbf{D} - \mathbf{I})| = 384$ fixed points, dense periodic points and topological mixing. F_ε has a unique absolutely continuous invariant measure μ_ε, and for μ_ε-almost every x the exponents satisfy $\lambda_1 + \lambda_2 + \lambda_3 + \lambda_4 = h_{\mu_\varepsilon} \le \ln 945$.*

Proof The family $F_{t\varepsilon}$, $t \in [0, 1]$, is a homotopy through expanding maps (Proposition 1) from $L_{\mathbf{D}}$ to F_ε, so F_ε induces $\mathbf{D}$ on $\pi_1(\mathbb{T}^4) = \mathbb{Z}^4$. Shub's theorem [31] states that an expanding map of a torus is topologically conjugate to the expanding linear endomorphism given by its action on π_1; see also [35]. Degree, number of fixed points, density of periodic points, mixing and topological entropy are conjugacy invariants; for $L_{\mathbf{D}}$ the degree is $\det \mathbf{D} = 945$, the fixed points number $|\det(\mathbf{D} - \mathbf{I})| = 2 \cdot 4 \cdot 6 \cdot 8 = 384$, and $h_{\mathrm{top}}(L_{\mathbf{D}}) = \sum_i \ln d_i = \ln 945$ [37]. Existence and uniqueness of the absolutely continuous invariant measure follow from Krzyżewski and Szlenk [32]. Pesin's formula for endomorphisms [34] gives $h_\mu = \sum_i \lambda_i^+$, and the variational principle bounds h_μ by h_{top}. □

Remark 1 The conjugacy in Proposition 2 is in general only a homeomorphism, and Lyapunov exponents are not topological invariants. Proposition 2 therefore does not imply $\lambda_i = \ln d_i$; it fixes only the entropy and the periodic-orbit structure. The agreement of the spectrum with $\ln d_i$ reported in Sect. 6.1 is a numerical finding.

Proposition 2 has a simple consequence. Up to a continuous change of coordinates, the "hyperchaotic coupled" map is four independent multiplications. The coupling neither creates new topological dynamics nor raises the entropy above that of the uncoupled system. Any security argument that rests on the coupling must therefore come from somewhere other than topological dynamics.

3.4 A smooth-gated control variant

Gating weights between chaotic terms are often described as fuzzy memberships. To test whether such a gate changes the finite-state behaviour, we use a smooth weight $\mu(x) = [1 + \cos(2\pi x)]/2$ (no inference rules are involved, so we call the variant smooth-gated) and

$$\begin{aligned} x_i(n+1) &= \big[d_i x_i + \varepsilon g_i(x)\big] \bmod 1, \\ g_i(x) &= \mu(x_i)\sin(2\pi x_{i+1}) + \big[1 - \mu(x_i)\big]\sin(2\pi x_{i+2}). \end{aligned} \tag{3}$$

Proposition 3 *For the map (3), $\|Dg(x)\|_2 \le B = 2\pi\sqrt{6} \approx 15.39$ for all x, so the map is uniformly expanding with every exponent at least $\ln(3 - \varepsilon B)$ whenever $\varepsilon < 2/B \approx 0.1299$, and it is conjugate to $L_{\mathbf{D}}$ in that range.*

Proof Row i of Dg has entries $\partial g_i/\partial x_i = \mu'(x_i)[\sin 2\pi x_{i+1} - \sin 2\pi x_{i+2}]$ with $|\mu'| \le \pi$, so $|\partial g_i/\partial x_i| \le 2\pi$; $|\partial g_i/\partial x_{i+1}| \le 2\pi\mu(x_i)$; $|\partial g_i/\partial x_{i+2}| \le 2\pi[1 - \mu(x_i)]$. Every row sum is at

most 4π and every column sum at most 6π, so $\|Dg\|_2 \le (\|Dg\|_1\|Dg\|_\infty)^{1/2} \le 2\pi\sqrt{6}$; Weyl's inequality gives $\sigma_{\min} \ge 3 - \varepsilon B$, and the homotopy argument of Proposition 2 applies. □

The Jacobian of (3) was verified symbolically (SymPy; maximum absolute difference 1.8×10^{-15} at 50 random points), and the largest sampled value of $\|Dg\|_2$ over 10^6 random points was $8.885 < B$. The variant serves only as a negative control.

4 Exact digital realisation, observability and stream interfaces

4.1 The integer map

Security properties can only be measured for an exactly specified arithmetic model. With $q = 2^b$, the state is $u \in \mathbb{Z}_q^4$ and

$$u_i(n+1) = \big[d_i u_i(n) + R_q\big(u_{i+1}(n)\big)\big] \bmod q, \qquad R_q(t) = \operatorname{rint}\big(q\varepsilon \sin(2\pi t/q)\big), \quad (4)$$

where rint rounds half to even. For $b \le 24$ the table R_q is computed once in IEEE-754 double precision (NumPy 2.5.4) and frozen, with its SHA-256 digest published (e.g. $b = 8$, $\varepsilon = 0.1$: `8cfa15011220...`); for $b = 32$ the same float64 expression is evaluated on demand, because a table would need 16 GiB. Independent C and Python implementations were cross-checked byte by byte; the test vector $u(0) = (17, 89, 121, 203)$, $b = 8$, $\varepsilon = 0.1$ gives $u(1) = (72, 193, 54, 45)$, $u(2) = (190, 222, 145, 174)$ and first repetition after 5 986 transient steps on a cycle of length 484.

4.2 Structural results

Proposition 4 *Let F_q be the map (4). (a) If $R_q(q-t) = -R_q(t)$ for all t, then $F_q(-u) = -F_q(u)$. (b) If $R_q(0) = R_q(q/2) = 0$, the sixteen states $u \in \{0, q/2\}^4$ are fixed points of F_q for every $b \ge 1$ and every ε. (c) If $q\varepsilon \le 1/2$, then $R_q \equiv 0$, F_q is a bijection of $\mathbb{Z}_q^4$, and for $b \ge 3$ every orbit is purely periodic with period dividing 2^{b-2}.*

Proof (a) is immediate from (4). (b) d_i is odd, so $d_i(q/2) \equiv q/2 \pmod q$, and the coupling term vanishes at 0 and $q/2$. (c) $|q\varepsilon \sin(\cdot)| \le 1/2$, and rint maps $[-1/2, 1/2]$ to 0 under round-half-to-even. Then $u_i \mapsto d_i u_i$ is a bijection because d_i is a unit modulo 2^b. The unit group $(\mathbb{Z}/2^b\mathbb{Z})^\times$ has exponent 2^{b-2} for $b \ge 3$, so $d_i^{2^{b-2}} \equiv 1$ and $F_q^{2^{b-2}} = \mathrm{id}$. □

All reference tables satisfy the hypotheses of (a) and (b). Part (c) is a precision threshold: for $\varepsilon = 0.025$ every precision $b \le 4$, and for $\varepsilon = 0.05$ every $b \le 3$, yields the uncoupled permutation, whose periods are at most 2^{b-2} although the real map has exponents $\ln 3$ to $\ln 9$.

Lemma 5 *For $v \in \mathbb{Z}_q^4$ define $h_i(s) = d_i^{-1}\big(v_i - R_q(s)\big) \bmod q$, where d_i^{-1} is the inverse of d_i modulo q. The preimages of v under F_q are in bijection with the fixed points t of $g_v = h_1 \circ h_2 \circ h_3 \circ h_4$ on $\mathbb{Z}_q$, via $t \mapsto \big(t, h_2(h_3(h_4(t))), h_3(h_4(t)), h_4(t)\big)$. All preimages can be listed with $4q$ table look-ups.*

Proof $F_q(u) = v$ is equivalent to $u_i = h_i(u_{i+1})$ for $i = 1, \dots, 4$ (cyclically). Guessing $u_1 = t$ fixes $u_4 = h_4(t)$, $u_3 = h_3(u_4)$ and $u_2 = h_2(u_3)$, and consistency requires $u_1 = h_1(u_2)$, i.e. $g_v(t) = t$. □

Because every state has exactly one image, the preimage sets of distinct states are disjoint and the inverse graph rooted at any state is a tree; backward search therefore needs no memoisation, and its cost is $4q$ look-ups per visited node.

4.3 Observability of truncated output

The truncation extractor (Sect. 4.4) reveals the top h bits of each coordinate, $Z_t = \lfloor u(t)/2^k \rfloor$ componentwise with $k = b - h$, and hides $4k$ bits. We ask how fast observations remove the uncertainty about the hidden state $U = u(0)$, taken uniform on $\mathbb{Z}_q^4$; Z_0^{T-1} denotes $(Z_0, \dots, Z_{T-1})$ and $|I|$ the number of integers in an interval.

Proposition 6 (Observability without coupling) *Let $\varepsilon = 0$ and $h \geq 4$. Then*

$$\sum_{i=1}^{4} \max\Big\{0,\ k - \log_2 \frac{d_i^T - 1}{d_i - 1}\Big\} \ \leq\ H\big(U \mid Z_0^{T-1}\big) \ \leq\ \sum_{i=1}^{4} \log_2\big(\lfloor 2^k d_i^{-(T-1)} \rfloor + 1\big).$$

In particular $H(U \mid Z_0^{T-1}) = 4k - (T-1)\log_2 945 + \delta_T$ with $-1.30 \leq \delta_T \leq 4$ while every coordinate retains hidden bits: hidden bits disappear at $\log_2 945 = h_{\mathrm{top}}(F_\varepsilon)/\ln 2$ bits per step up to a bounded error.

Proof For $\varepsilon = 0$ the coordinates evolve independently, $u_i \mapsto d_i u_i \bmod q$, so $H(U \mid Z_0^{T-1}) = \sum_i H(U_i \mid Z_{i,0}^{T-1})$. Fix i, write $d = d_i$ and $x = U_i/q$. Then $u_i(t)/q = d^t x \bmod 1$ exactly, and $Z_{i,t} = \lfloor 2^h (d^t x \bmod 1) \rfloor$ is the symbol of $d^t x$ in the partition $\mathcal{P}$ of $[0,1)$ into 2^h intervals of length 2^{-h}: the observations are the itinerary of x under $f(x) = dx \bmod 1$, sampled on the grid $q^{-1}\mathbb{Z}_q$. *Lower bound.* Points with different itineraries are separated by a point of $B = \{m 2^{-h} d^{-j} : 0 \leq j < T,\ m \in \mathbb{Z}\} \cap [0,1)$, which has at most $2^h \sum_{j<T} d^j = 2^h (d^T - 1)/(d-1)$ elements, so $Z_{i,0}^{T-1}$ takes at most that many values. As Z is a function of U_i, $H(U_i \mid Z) = H(U_i) - H(Z) \geq b - h - \log_2 \frac{d^T - 1}{d-1}$, and $H \geq 0$. *Upper bound.* Let $A_T(x)$ be the set of points with the same first T symbols as x; $A_1(x)$ is an interval of length 2^{-h}. On $A_1(x)$ the map f is affine with slope d, and its image is an arc of length $d 2^{-h}$; $A_{T-1}(fx)$ lies in an element of $\mathcal{P}$, and $(d+1)2^{-h} \leq 10/16 < 1$, so the two arcs meet in a single interval. Hence $A_T(x)$ is an interval of length at most $|A_{T-1}(fx)|/d$, and by induction at most $2^{-h} d^{-(T-1)}$. It contains at most $\lfloor 2^k d^{-(T-1)} \rfloor + 1$ grid points, over which the posterior of U_i is uniform. The constants follow from $\frac{d^T - 1}{d-1} \leq d^{T-1} \frac{d}{d-1}$, $\log_2 \prod_i \frac{d_i}{d_i - 1} < 1.30$ and $\log_2(\lfloor y \rfloor + 1) \leq \log_2 y + 1$ for $y \geq 1$. □

With coupling the coordinates interact, and an exact entropy computation is no longer available. The following constructive bound holds for the whole proven range and is achieved by an efficient algorithm.

Proposition 7 (Interval observability with coupling) *Let $0 \le \varepsilon < 1/\pi$, $h \ge 4$ and $\rho = (3-2\pi\varepsilon)^{-1}$. For observations Z_0^{T-1} of (4), initialise integer intervals $I_i(t) = [2^k Z_{i,t},\, 2^k Z_{i,t} + 2^k - 1]$, write $R_q[I] = [\min_I R_q, \max_I R_q]$, and repeat, for all i and $t < T-1$, the backward update*

$$I_i(t) \leftarrow \operatorname{hull}\big\{u \in I_i(t) :\ (d_i u + r) \bmod q \in I_i(t+1) \text{ for some } r \in R_q[I_{i+1}(t)]\big\}$$

until no interval changes. The true trajectory never leaves the intervals, the procedure terminates, and at termination

$$\max_i |I_i(0)| \ \le\ \rho^{T-1} 2^k + \frac{3}{2 - 2\pi\varepsilon}.$$

Consequently $H(U \mid Z_0^{T-1}) \le 4 \log_2\big(\rho^{T-1} 2^k + 3/(2-2\pi\varepsilon)\big)$ throughout the proven range.

Proof Soundness and termination. If $u_{i+1}(t) \in I_{i+1}(t)$, then $R_q(u_{i+1}(t))$ lies between the minimum and maximum of R_q over $I_{i+1}(t)$, so the true $u_i(t)$ survives every update; intervals only shrink, so the iteration terminates. *Widths.* Let $w_i(t) = |I_i(t)|$ at termination. Since $|\sin a - \sin b| \le |a - b|$ and rint moves a value by at most $1/2$, the range of R_q over $I_{i+1}(t)$ has length $\Delta \le 2\pi\varepsilon(w_{i+1}(t) - 1) + 1$ (for the float64 tables up to a term below 10^{-8} that does not change the integer bounds). The admissible residues $d_i u \bmod q$ form an arc of length $w_i(t+1) - 1 + \Delta$; as u ranges over $I_i(t)$, $d_i u$ ranges over an integer interval of length below $9 \cdot 2^k$, and $(9 + 1 + 2\pi\varepsilon)2^k < 2^{k+4} \le q$ for $h \ge 4$, so it meets at most one translate of the arc. The admissible u therefore form one interval with at most $(w_i(t+1) - 1 + \Delta)/d_i + 1$ elements, and at termination

$$w_i(t) \le \big(w_i(t+1) + 2\pi\varepsilon\, w_{i+1}(t)\big)/d_i + 1.$$

In vector form $w(t) \le \mathbf{D}^{-1} w(t+1) + \mathbf{A} w(t) + \mathbf{1}$ with $\mathbf{A} = 2\pi\varepsilon \mathbf{D}^{-1}\mathbf{P} \ge 0$ and $\|\mathbf{A}\|_\infty = 2\pi\varepsilon/3 < 1$, so $w(t) \le (\mathbf{I} - \mathbf{A})^{-1}(\mathbf{D}^{-1} w(t+1) + \mathbf{1})$ and $\|w(t)\|_\infty \le (\|w(t+1)\|_\infty/3 + 1)/(1 - 2\pi\varepsilon/3) = \rho\|w(t+1)\|_\infty + 3\rho$. With $\|w(T-1)\|_\infty \le 2^k$, induction gives $\|w(0)\|_\infty \le \rho^{T-1} 2^k + 3\rho/(1-\rho)$, and $3\rho/(1-\rho) = 3/(2 - 2\pi\varepsilon)$. The posterior support lies in $\prod_i I_i(0)$. □

Corollary 8 (Data and work for state recovery) *For $0 \le \varepsilon < 1/\pi$, $h \ge 4$ and $T \ge T^* = 1 + \lceil k/\log_2(3 - 2\pi\varepsilon) \rceil$ observed steps, the box $\prod_i I_i(0)$ contains at most $\lfloor 1 + 3/(2 - 2\pi\varepsilon) \rfloor^4$ states (81 at $\varepsilon = 0.1$), among which the true state is found by forward verification. Each sweep of the updates costs $8(T-1)$ interval operations; adding the analogous forward update can only shrink the intervals further.*

Remark 2 Proposition 6 is information-theoretic and two-sided: no observer, however powerful, can do better, and the rate is exactly the topological entropy of the real map. Proposition 7 is one-sided and constructive: an efficient observer achieves a per-coordinate rate of at least $\log_2(3 - 2\pi\varepsilon)$, the expansion floor of Proposition 1. Topological entropy concerns the real map; the relevant digital quantity is the block entropy of the output process, $H(Z_0^{T-1}) = 4b - H(U \mid Z_0^{T-1})$, which by Proposition 6 grows as $4h + (T-1)\log_2 945$ until the state is determined and then stops growing. The interval-refinement idea itself is that of [20] for one-dimensional maps; the propositions add explicit rates and show that coupling of strength below $1/\pi$ does not prevent it.

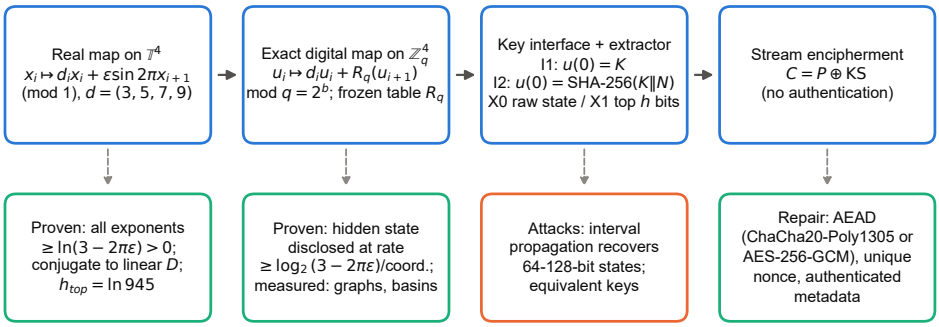


Fig. 1 Study design. Top row: the real SCUT-4D map, its exact b-bit realisation, the key interfaces I1/I2 with extractors X0/X1, and XOR encipherment. Bottom row: what is established at each stage: analytical results (Sects. 3, 4.3), finite-state measurements (Sect. 6.2), attacks (Sects. 6.4, 6.5) and the authenticated-encryption baseline

4.4 Key interfaces, extractors and encipherment

We study two key interfaces and two output extractors that cover common practice. In interface **I1** the $4b$-bit key is the initial state, $u(0) = K$, as in designs that use initial conditions as keys. In interface **I2**, $u(0)$ consists of the first $4b$ bits of SHA-256$(K\|N)$ for a 128-bit key K and a 96-bit nonce N. In both interfaces $W = 1000$ warm-up iterations are discarded. Extractor **X0** outputs all state bits (big-endian per coordinate, $4\lceil b/8 \rceil$ bytes per step); extractor **X1** outputs the top h bits of each coordinate and hides $k = b - h$ bits per coordinate. Images are serialised row-major (channel-interleaved for colour) as raw pixel bytes without container headers and enciphered as $C = P \oplus KS$. All demonstrated parameters are research instances: with $b = 16$ the session state has 64 bits whatever the length of the master key in I2, and $b = 32$ is used to show that 128-bit states fare no better. Figure 1 summarises the chain from the real map to the cipher.

4.5 Adversary model and security games

The adversary knows the algorithm, ε, the tables, W and the extractor (Kerckhoffs' principle) [38]. Plaintext knowledge consists of the first m pixel bytes of an image, for example a uniform border row, a known logo or a chosen all-zero image; container headers are not part of the encrypted payload in our serialisation. We use two games. *State recovery* G_{SR}: given one ciphertext and its first m plaintext bytes, the adversary wins if it outputs the session state at the first output block (equivalently, all remaining keystream). *Known-prefix distinguishing* G_{KP}: the adversary chooses two messages M_0, M_1 of equal length that share their first m bytes; the challenger draws $\beta \in \{0, 1\}$ and a fresh nonce N, and returns $(N, M_\beta \oplus KS(K, N))$; the adversary outputs β' and has advantage $|2\Pr[\beta' = \beta] - 1|$. G_{KP} needs no encryption-oracle queries beyond the challenge, so any non-negligible advantage contradicts IND-CPA security [38]. We report data (m), time, memory, success rates with exact Clopper–Pearson intervals and, for I1, the set of keys consistent with the recovered state. Integrity requires

authenticated encryption [39]; flipping a ciphertext bit of any XOR stream flips the corresponding plaintext bit, which we report only as a control.

5 Methods

5.1 Lyapunov spectra and numerical precision

Spectra of (1) and (3) were computed with the QR (Benettin) method [40, 41] for $\varepsilon \in \{0, 0.025, 0.05, 0.1, 0.15, 0.2, 0.25, 0.3\}$, using 100 uniformly random initial points per ε (seed 20261030), 1 000 transient iterations during which both the state and the QR frame are evolved, and 100 000 accumulation steps; finite-time exponents over 1 000-step windows were stored. Twenty further seeds were run for 10^6 steps at $\varepsilon = 0.025, 0.1, 0.3$, with the deviation from $\ln d_i$ recorded at $n = 10^4, \ldots, 10^6$ to separate $O(1/n)$ terms from asymptotic shifts. As a precision control, a reference orbit was computed for 2000 steps with mpmath at $p = 2048$, 4096, 6810 and 8192 bits from the same dyadic initial point. With Lipschitz constant $L = 9 + 2\pi\varepsilon$ and elementary operations accurate to a few units in the last place, the forward error of the p-bit orbit after n steps is at most about $2^{-p+2}L^n L/(L-1)$, i.e. 2^{-273} for $p = 6810$, $n = 2000$; floating-point orbits of an expanding map are pseudo-orbits shadowed by true orbits [35].

5.2 Functional graphs, basins and cycle statistics

For $4b \le 24$ the functional graph of F_q was enumerated exhaustively: image fraction, cyclic states, cycles, longest cycle and tail, exact mean rho length over uniformly random starting states, in-degree distribution, number of components (one per cycle), basin sizes (by pointer doubling) and the basin share of the sixteen fixed points. For $b = 7, 8$ the image fraction was computed with a bitmap. Rho lengths $\mu + \lambda$ (tail plus cycle) were sampled with Brent's algorithm [42] from random states: 1000 for $b \le 10$, 500 for $b = 11$–12, 200 for $b = 13$–16 (the first 30 at $b = 15, 16$ from the original run, 170 added), with a censoring limit of 2^{36} steps. Three null models act on the same $4b$-bit space, with fin the SplitMix64 finaliser: **H**, $x \mapsto \mathrm{fin}(x + \kappa) \bmod 2^{4b}$, a random-function surrogate only for $4b < 64$ (for $4b = 64$ it is a permutation and was not used); **X**, $x \mapsto [\mathrm{fin}(x + \kappa_1) \oplus \mathrm{fin}(x + \kappa_2)] \bmod 2^{4b}$, a random-function surrogate for every b, run with five independent key pairs; and **P**, a keyed xorshift–multiply permutation computed modulo 2^{4b}, run with three keys. Exhaustive graphs confirm that X behaves as a random function (image fraction $1 - 1/e$) and P as a permutation (image fraction 1). Their expected rho lengths are $\sqrt{\pi N/2}$ and $(N+1)/2$ [30]. Growth slopes of $\log_2$ median rho length in b carry bootstrap 95% intervals from resampling starts within each b.

5.3 Observability

For $k = 4, 5$, $b \in \{8, 9, 12, 16\}$ and $\varepsilon \in \{0, 0.1\}$, the survivor set $S_T(U)$ of states consistent with Z_0^{T-1} was computed exactly by enumerating the 2^{4k} states that share Z_0 and filtering them step by step, for 400 ($k = 4$) or 150 ($k = 5$) uniformly random states U. Because the posterior is uniform on $S_T(U)$, $H(U \mid Z_0^{T-1}) = \mathbb{E} \log_2 |S_T(U)|$,

the probability that a single guess is correct is $\mathbb{E}\,|S_T(U)|^{-1}$, and the guessing entropy is $\mathbb{E}(|S_T(U)|+1)/2$.

5.4 Attacks

X0 state recovery. XOR of the known prefix with the ciphertext gives $u(W+1)$ directly; the public update predicts the rest. **Interval attack (X1).** Proposition 7 is implemented with both backward and forward updates; the range of R_q over an interval is attained at its endpoints or at $q/4$, $3q/4$, because rint is monotone and the sine is monotone on quarter periods. If the box at $t=0$ holds at most 2^{14} states they are verified by forward simulation; otherwise the widest interval is split and both halves are propagated again (branch-and-bound with a budget of 2^{16} boxes). Every step over-approximates, so the true state is never discarded. Random states were drawn for $b \in \{16, 20, 24, 32\}$, $h \in \{4, 8\}$ and T up to 40; the minimum T was determined in a separate sweep. **Guess-and-determine baseline (X1).** Given $T=8$ observed blocks, the low parts (l_1, l_2) of u_1, u_2 are enumerated and filtered by the top bits of $u_1(1)$; survivors are extended by l_3 and filtered by $u_2(1)$, $u_1(2)$; then by l_4 and checked against all observations. Work is the number of candidate states evaluated; it uses only the sparsity of the coupling, not its size. Trials: 50 at $b = 12, 16$ with $h=8$ on image data, 10–20 per further (b, h) on simulated output, including $k=10$ ($b=16$, 20) and $k=11$ ($b=16$; 3 trials of about 2 h each). For an ideal random-function transition with the same output rate, generic search over the 2^{4k} states consistent with Z_0 is the reference; the exposure $2^{4b} \to 2^{4k}$ is generic to truncation. **Equivalent-key enumeration (I1).** From the recovered state $v = u(W+1)$, depth-first search applies Lemma 5 level by level to list $F_q^{-(W+1)}(v)$, all keys that generate the observed keystream; memory is the DFS stack ($W+1$ states) plus the output list, and the reported cost is visited nodes times $4q$ look-ups. **Distinguishing game.** G_{KP} was run under I2 with fresh nonces at $b=16$ for X0 ($m=8$, 200 trials) and X1 ($m=32$, 60 trials, guess-and-determine), with 4096-byte messages sharing a random known prefix. Test images were drawn from scikit-image 0.26 (camera, CC0; astronaut, public domain; coffee, CC0; chelsea, CC0) plus resized and synthetic images; SHA-256 digests are archived.

5.5 Sessions, integrity, statistics and performance

Multi-session collisions under I2 were estimated in two ways. The collision probability of session states after the warm-up is $\sum_v p(v)^2 = \mathbb{E}_U |F_q^{-(W+1)}(F_q^{W+1}(U))|/N$, estimated without bias by the mean equivalence-class size of uniformly random keys; at $b=8$ it was also measured by simulating $Q = 2^{14}$–2^{17} sessions. Malleability (10 000 random bit flips; targeted 64×64 block forgery) and nonce reuse were tested against ChaCha20-Poly1305 [43] and AES-256-GCM [44] (Python `cryptography` 50.0.2) as generic controls. Nine tests of NIST SP 800-22 Rev. 1a [45] were implemented and validated against its worked examples and against `os.urandom` and ChaCha20; each generator was tested on 100 sequences of 10^6 bits from non-overlapping stream segments, and a test is passed only if the pass proportion is at least $0.99 - 3(0.99 \cdot 0.01/100)^{1/2} = 0.960$ *and* the p-values are uniform (χ^2 P-value $\geq 10^{-4}$). NIST has

decided to revise the suite and to reject its use for assessing cryptographic generators [46]; we report it only for comparability. Throughput was measured on 64 MiB buffers (best of five), single-threaded.

5.6 Arithmetic variants

To separate arithmetic from dynamical effects, the coupling table was rebuilt with three other rules: truncation toward zero, float32 evaluation followed by rint, and floor. Odd symmetry and $R_q(0) = R_q(q/2) = 0$ were checked for every table; functional graphs ($b \le 6$) and rho lengths ($b = 8, 10, 12$, 200 starts) were measured as above. Equation (1) was also iterated directly in float32 and float64 with the platform `libm`, with Brent cycle detection censored at 10^8 steps (20 starts each).

5.7 Software and hardware

Experiments ran on an Intel Core i5-1145G7 @ 2.60 GHz with 7.7 GB RAM under Windows-11-10.0.26300-SP0, Python 3.12.10, NumPy 2.5.4, SciPy 1.18.1, SymPy 1.14, mpmath 1.3 and GCC 13.2.0 (MinGW-W64), -O2 -ffp-contract=off (`-O2 -ffp-contract=off`). All random choices use recorded seeds; code, frozen tables with SHA-256 manifest, raw per-trial results, pinned package versions and figure scripts are provided as Online Resource 1.

6 Results

6.1 The coupling does not measurably change the real spectrum

At $\varepsilon = 0.1$ the 100-seed survey gives $\lambda = (2.197225, 1.945910, 1.609438, 1.098612)$ nats/step against $(\ln 9, \ln 7, \ln 5, \ln 3) = (2.197225, 1.945910, 1.609438, 1.098612)$; over the whole grid the largest deviation of any exponent from $\ln d_i$ is 7.0×10^{-6} and the exponent sum lies in $[6.851181, 6.851185]$, below $\ln 945 = 6.851185$ as Proposition 2 requires (Fig. 2). The smallest finite-time exponent in any 1000-step window was 1.0980, well above the rigorous floor $\ln(3 - 2\pi\varepsilon)$. The longer runs (10^6 steps, 20 seeds) place the deviations below 2.5×10^{-10}, 2.0×10^{-8} and 5.1×10^{-6} nats/step at $\varepsilon = 0.025$, 0.1 and 0.3 (largest standard error 8.9×10^{-7}), and the deviation of the sum from $\ln 945$ is -5.9×10^{-6} at $\varepsilon = 0.3$, close to the second-order estimate $-(2\pi\varepsilon)^8/(32 \cdot 945^2) = -5.6 \times 10^{-6}$ obtained by expanding $\ln|\det DF_\varepsilon|$ under Lebesgue measure. The first version of this analysis started the QR frame at the identity after the state transient; the frame realignment then added an $O(1/n)$ bias of up to 1.2×10^{-4} that was largest at small ε. All values reported here evolve the frame through the transient. In the precision study the 2048- and 4096-bit orbits first deviate from the 8192-bit orbit by more than 10^{-12} after 638 and 1282 steps, against 627 and 1254 predicted by $p/\log_2 L$; the 6810-bit orbit coincides with it for all 2000 steps. The 2000-step exponent estimates of the four orbits nevertheless differ by at most 1.2×10^{-5} nats/step, and float64 estimates from the same initial points agree with the reference to within 1.1×10^{-5}: individual orbits are lost after a few hundred steps, but their averages are not. The smooth-gated variant at $\varepsilon = 0.1$ gives $(2.1969, 1.9454, 1.6079, 1.0946)$.

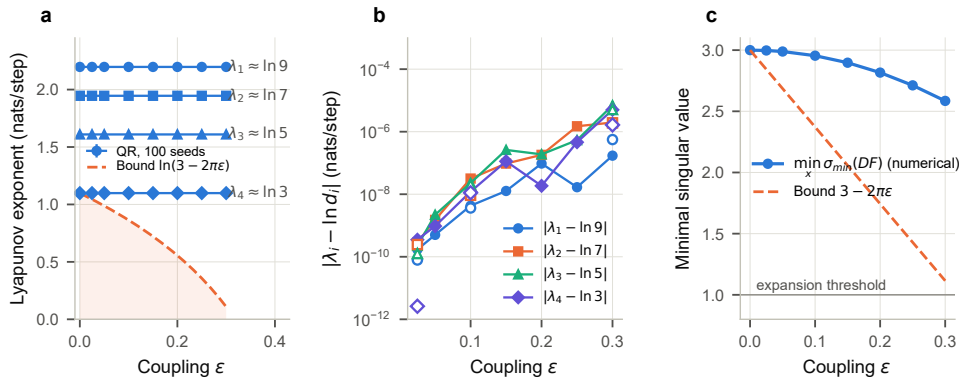


Fig. 2 **a** QR estimates of the four Lyapunov exponents of the real SCUT-4D map (markers: mean over 100 seeds; bars: min–max, smaller than the markers) and the rigorous lower bound $\ln(3 - 2\pi\varepsilon)$ of Proposition 1 (dashed; shaded region excluded). **b** Absolute deviation of each exponent from $\ln d_i$ (100 seeds, 10^5 steps; open markers: 20 seeds, 10^6 steps) on a logarithmic scale. **c** Numerical global minimum of the smallest singular value of DF_ε compared with the bound $3 - 2\pi\varepsilon$; values above 1 imply uniform expansion

A plausible explanation of the near-invariance is perturbative: $\mathbf{E}(x)$ has zero diagonal, so the first-order change of each exponent vanishes, and the only closed coupling loop is the full 4-cycle, which enters at order ε^4 with the factor $\prod_j \cos 2\pi x_j$ of mean zero under Lebesgue measure. We do not prove this; the data only bound the shifts. The rigorous bound of Proposition 1 is far from tight (Fig. 2c), but it is a proof, whereas the spectrum is an estimate.

Bifurcation structure. Figure 3 shows bifurcation diagrams of the real and the integer map. Throughout the proven range $\varepsilon < 1/\pi$ the real map has no bifurcation: no periodic window appears, the smallest exponent stays above the rigorous bound (minimum 1.099 nats/step), and a 400-bin histogram of the asymptotic x_1 is uniform within sampling noise (mean largest bin deviation 0.45, as expected for about 50 points per bin). Beyond the proven range, up to $\varepsilon = 1.5$, the histogram stays uniform within the same noise (mean largest deviation 0.46) and all exponents remain positive (smallest 0.980): expansion persists numerically there but is no longer proved. The three-dimensional diagrams make the contrast with the integer map visible. The real attractor fills the torus for every coupling (panel a), whereas the integer map takes values on a lattice whose occupied points change abruptly with ε: at $b = 4$ the reached cycles visit between 4 and 256 of the 256 pairs (u_1, u_2) (median 204; panel b). The finite-state analogue of a bifurcation diagram is the set of cycle lengths reached (panel d): it jumps between unrelated values from one tabulated coupling to the next, with no continuation in ε; at $\varepsilon = 0.1$ the reached cycles have lengths 140, 171, 286, 1324, 1676 for $b = 6$ and 484, 572, 2706, 4620, 4942 for $b = 8$, while below the precision threshold of Proposition 4(c) all lengths divide 2^{b-2}.

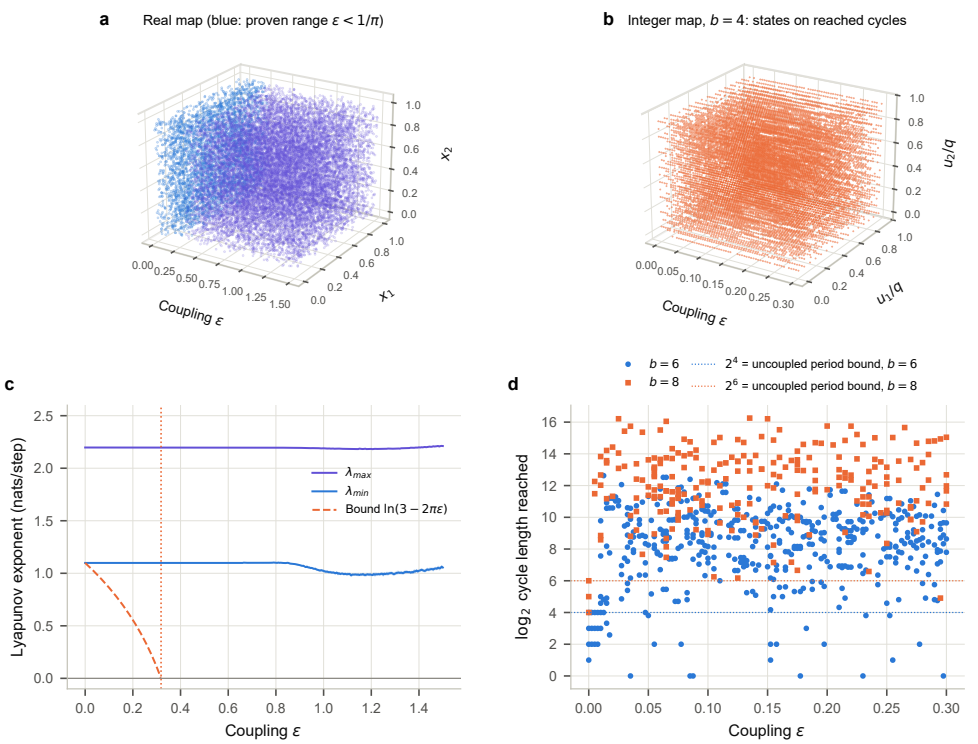


Fig. 3 Bifurcation structure. **a** Three-dimensional bifurcation diagram of the real map: (ε, x_1, x_2) after 1000 transient steps for 300 random initial points per $\varepsilon \in [0, 1.5]$ (blue: proven range $\varepsilon < 1/\pi$). **b** The same diagram for the integer map with $b = 4$: pairs $(u_1/q, u_2/q)$ visited on the cycles reached from 100 random states. **c** Largest and smallest Lyapunov exponents of the real map (20 seeds, 3000 steps; band: ± 1 SD of λ_{min}) with the rigorous bound $\ln(3 - 2\pi\varepsilon)$. **d** Lengths of the cycles reached by the integer map from 200 random states for $b = 6$ and $b = 8$; dotted: the period bound 2^{b-2} of the uncoupled map

6.2 Finite precision: many small basins and short cycles

Exhaustive enumeration confirms Proposition 4 (Table 3). With $\varepsilon = 0$ the map is a permutation whose longest cycle is exactly 2^{b-2}; with $\varepsilon = 0.025$ the graphs for $b = 3, 4$ coincide with the uncoupled case, as the threshold $q\varepsilon \le 1/2$ predicts. With $\varepsilon = 0.1$ the image fraction falls from 0.863 ($b = 3$) to 0.659 ($b = 8$), staying above the random-function value $1 - 1/e \approx 0.632$. The graphs split into many small components (Table 4): at $b = 6$ SCUT-4D has 104 components and a largest basin of 5.4% of the states, against 8 components and 71.7% for the random-function control (theory: about $\frac{1}{2} \ln N \approx 8$ components and a largest share of 0.758 [30]). The in-degree distribution is narrower than the Poisson(1) law of a random function (variance 0.61 against 1), i.e. the map is closer to injective, yet its cycles are shorter. The sixteen fixed points of Proposition 4(b) do not explain this: their basins hold 2.0% of the states at $b = 4$ and a fraction 9.5×10^{-7} at $b = 6$.

Table 3 Exact functional-graph statistics (all 2^{4b} states enumerated); control rows are means over independent keys

Map	b	Image fraction	Cyclic states	Cycles	Fixed points	Longest cycle	Longest tail	Mean rho
SCUT, $\varepsilon = 0$	4	1.0000	65 536	20 544	128	4	0	3.5
SCUT, $\varepsilon = 0$	6	1.0000	16 777 216	1 347 648	128	16	0	13.9
SCUT, $\varepsilon = 0.1$	4	0.8025	1 984	72	16	100	117	116.1
SCUT, $\varepsilon = 0.1$	6	0.7294	31 760	104	16	1 676	2 095	1751.1
Smooth-gated, $\varepsilon = 0.1$	4	0.5763	402	100	32	96	252	107.8
Smooth-gated, $\varepsilon = 0.1$	6	0.5723	2 650	138	32	549	4 383	1716.5
Random function (5 keys)	4	0.6327	345	6.4	1.8	255	417	361.7
Random function (5 keys)	6	0.6321	3 856	8.4	0.6	2 185	7 234	4190.8
Random permutation (3 keys)	4	1.0000	65 536	14.7	2.3	36 660	0	25893.9
Random permutation (3 keys)	6	1.0000	16 777 216	18.0	1.3	11 087 619	0	8333329.6

Random function: map X of Sect. 5; random permutation: map P. Expected values for a uniform random mapping on $N = 2^{24}$ states: image fraction ≈ 0.632, cyclic states and mean rho $\approx \sqrt{\pi N/2} \approx 5133$.

Table 4 Basin and in-degree structure of the exact functional graphs ($\varepsilon = 0.1$; random function X with key 0)

Map	b	Comp.	Largest basin	Basin entropy (bit)	Fixed-point basin share	In-degree 0	In-degree 1	In-degree variance
SCUT-4D	4	72	0.092	4.38	2.0×10^{-2}	0.198	0.613	0.41
Random function	4	7	0.826	0.71	3.1×10^{-5}	0.368	0.369	1.00
SCUT-4D	5	72	0.064	4.13	3.1×10^{-5}	0.250	0.521	0.54
Random function	5	4	0.574	1.45	0	0.368	0.368	1.00
SCUT-4D	6	104	0.054	4.86	9.5×10^{-7}	0.271	0.493	0.61
Random function	6	8	0.717	0.94	1.2×10^{-7}	0.368	0.368	1.00
Random mapping (theory)	–	$\frac{1}{2}\ln N$	0.758	–	–	0.368	0.368	1.00

Cycle sampling extends the comparison to 64-bit states (Table 5, Fig. 4). Fitting $\log_2$ of the median rho length for $b \geq 6$ gives a slope of 1.93 (95% CI 1.92–1.94) bits per bit of coordinate precision for SCUT-4D and 2.02 (2.01–2.03) for the five-key random-function control (theory 2.00). At $b = 16$ the median rho length of SCUT-4D is 9.55×10^8 steps (200 starts, 0 censored), against 5.23×10^9 for the control and $\sqrt{\pi N/2} = 5.38 \times 10^9$ in theory. Orbits concentrate on few cycles: at $b = 8$ every

Table 5 Rho lengths at large precision: SCUT-4D ($\varepsilon = 0.1$) against the random-function control X (five keys) and theory

b	n (cens.)	SCUT median	95% CI	X median	X key medians	n_X (cens.)	$\sqrt{\pi N/2}$	Ratio
12	500 (0)	7.4e6	6.9e6–7.7e6	2.7e7	1.5e7–6.2e7	500 (0)	2.1e7	2.8
13	200 (0)	2e7	1.9e7–2.1e7	6.6e7	5e7–1.3e8	200 (0)	8.4e7	4.1
14	200 (0)	1e8	9.9e7–1.1e8	2.8e8	2e8–5.5e8	200 (0)	3.4e8	3.3
15	200 (0)	2.8e8	2.5e8–3.1e8	1.4e9	1.2e9–1.6e9	60 (0)	1.3e9	4.7
16	200 (0)	9.5e8	8.8e8–9.9e8	5.2e9	4e9–6e9	30 (0)	5.4e9	5.6

Ratio: random-mapping expectation divided by the SCUT-4D median. CI: bootstrap over starts. Numbers in e-notation, e.g. 1.0e9 = 1.0×10^9; censored runs (limit 2^{36} steps) are excluded from medians.

sampled state ended on one cycle of length 484. The smooth gate shortens cycles even more and is not considered again. Table 6 varies only the rounding of the coupling table. The ratio of the random-mapping expectation to the measured median rho length at $b = 8, 10, 12$ is 6.2, 3.3, 2.9 for the reference table, 2.6, 3.4, 1.4 for truncation, 6.2, 3.3, 2.9 for a float32 table and 0.9, 1.1, 0.8 for floor rounding. The float32 table coincided entry by entry with the reference table for every tested b, so its row only confirms reproducibility. Both odd-symmetric tables give cycles shorter than random at all three precisions, whereas floor rounding, which breaks $R_q(q - t) = -R_q(t)$, gives ratios close to one. The shortfall therefore comes from the integer arithmetic and goes with the odd symmetry of Proposition 4(a); since the two symmetric rules still differ by up to a factor of two at equal b, symmetry alone does not determine its size. Direct floating-point iteration of (1) behaves differently again: none of 20 float32 orbits and none of 20 float64 orbits closed within 10^8 steps (20 and 20 censored). The observability results of Sect. 4.3 and the attacks of Sect. 6.4 do not depend on cycle lengths, and the raw-output attack applies to any deterministic public update, including floating-point ones.

6.3 Observability

The exact survivor counts agree with Proposition 6 in every configuration (Fig. 5a): at $\varepsilon = 0$, $k = 4$ and $T = 2$ the conditional entropy is 6.19 bits (bounds 5.09–7.17; rate prediction $4k - \log_2 945 = 6.12$), and the values for $b = 8$, 12 and 16 coincide within their standard errors, as the proposition implies: uncertainty depends on the hidden bits k, not on the key length $4b$. With $\varepsilon = 0.1$ the entropy lies inside the same bounds but 0.4–0.6 bit lower at $T = 2$ (5.65 bits for $k = 4$), and the state becomes unique earlier: at $T = 3$, 53% of states are determined against 15% without coupling. Coupling therefore reveals the state slightly faster, presumably because it passes information from the fast coordinates to the slow one. For $k = 5$ the probability that one guess is correct rises from 0.001 at $T = 2$ to 0.383 at $T = 3$ and 0.963 at $T = 4$. At $b = 16$,

Table 6 Effect of the rounding rule of the coupling table ($\varepsilon = 0.1$)

Table rule	Odd symmetric	Fixed points ($b = 6$)	Image fraction ($b = 6$)	Cycles ($b = 6$)	Mean rho ($b = 6$)	Ratio $b = 8$	Ratio $b = 10$	Ratio $b = 12$
rint (reference)	yes	16	0.729	104	1751	6.2	3.3	2.9
truncate toward zero	yes	32	0.723	184	1901	2.6	3.4	1.4
float32 sine, then rint	yes	16	0.729	104	1751	6.2	3.3	2.9
floor	no	20	0.703	27	6780	0.9	1.1	0.8

Ratio: random-mapping expectation $\sqrt{\pi N/2}$ divided by the median rho length (200 starts).

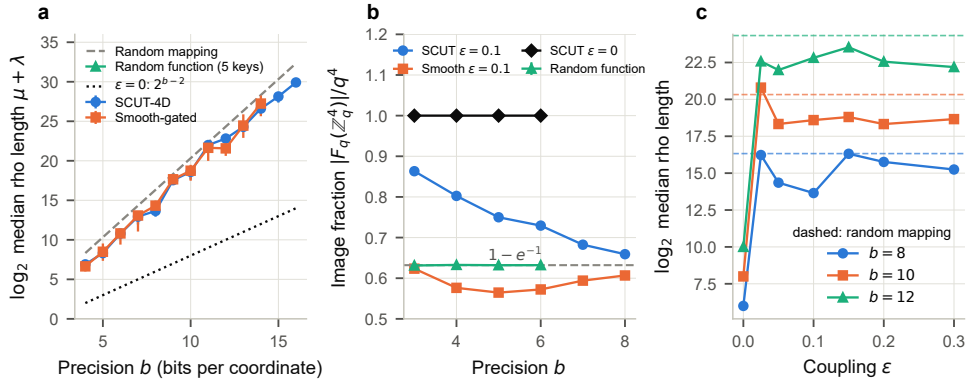


Fig. 4 **a** Median rho length (tail + cycle) versus precision b (bars: interquartile range) for SCUT-4D and the smooth-gated variant, compared with the random-mapping expectation $\sqrt{\pi N/2}$, $N = 2^{4b}$, the random-function control X (line: pooled median; band: range of the five key medians) and the period bound 2^{b-2} of the uncoupled map. **b** Exact image fraction versus b (control: mean and range over five keys). **c** Median rho length versus coupling ε for $b = 8, 10, 12$ (dashed: random-mapping expectation)

$k = 8$, the median survivor counts of the guess-and-determine runs, 3 784 434, 2 942, 14, 1, 1 for $T = 2, \ldots, 6$, follow the same rate (Fig. 5b).

6.4 State recovery

Raw output (X0). Four known bytes at $b = 8$ and eight at $b = 16$ determine $u(W + 1)$; every remaining byte was predicted exactly in 100/100 (95% CI [0.964, 1.000]) and 100/100 (CI [0.964, 1.000]) image encryptions.

Truncated output (X1): interval attack. The interval attack recovered the true state in 150 of 150 trials, and the state was unique in 150 (Table 7). No trial needed more than 1 box; propagation alone reduced the 128-bit state at $b = 32$, $h = 8$ to a single candidate after 1 736 interval operations and 82 ms (median, Python). The

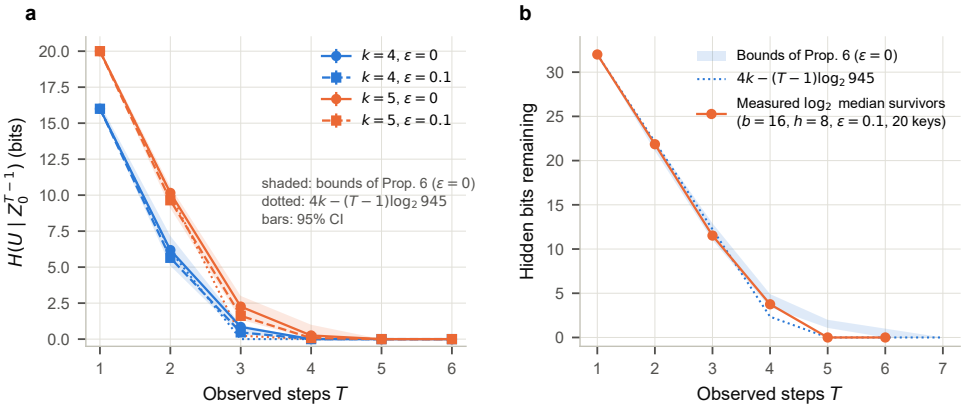


Fig. 5 Observability of the truncated output. **a** Conditional entropy $H(U \mid Z_0^{T-1})$ of the hidden state versus the number of observed steps for $k = 4, 5$ ($b = 16$), without coupling (circles) and with $\varepsilon = 0.1$ (squares); shaded: bounds of Proposition 6; dotted: $4k - (T-1)\log_2 945$; bars: 95% CI. **b** Hidden bits remaining at $b = 16$, $k = 8$: bounds of Proposition 6 and $\log_2$ median survivor counts at $\varepsilon = 0.1$

minimum number of observed steps at which every trial returned the unique true state was $T = 7$ for $b = 16$, $h = 8$, $\varepsilon = 0.1$; $T = 5$ for $b = 16$, $h = 8$, $\varepsilon = 0.3$; $T = 10$ for $b = 24$, $h = 8$, $\varepsilon = 0.1$; $T = 12$ for $b = 32$, $h = 8$, $\varepsilon = 0.1$, against $1 + 4k/\log_2 945 = 4.2, 4.2, 7.5, 10.7$, the number of steps after which the output carries enough information to fix the state, and the sufficient condition of Corollary 8 ($T \geq 8, 52, 14, 21$). At $b = 16$, $h = 8$ the attack needs 32 known bytes and about 3 ms, where the guess-and-determine baseline needs $2^{26.5}$ candidate evaluations. This works because of the structure used in Proposition 7: the coupling is monotone on quarter periods and varies slowly, so interval arithmetic can follow the expansion backwards, and the cost no longer grows with k.

Guess-and-determine baseline and the generic exposure. Truncation alone confines the state to 2^{4k} candidates at $t = 0$; this exposure is generic and would hold for any generator with the same output rate. The guess-and-determine attack, which exploits only the sparsity of the coupling, improves on generic search over those candidates by 4.5–6.4 bits (Table 8). A least-squares fit over all 295 trials, now including $k = 10$ and 11, gives $\log_2 W = -3.7 + 3.73k$ (bootstrap 95% CI of the slope 3.69–3.77; leaving out any one value of k changes the predicted median at that k by at most 0.8 bits). The first filter keeps a fraction 0.282, 0.277, 0.279, 0.264, 0.302, 0.292 of the (l_1, l_2) guesses for $k = 4, 5, 6, 7, 8, 9$, close to $1/d_1 = 1/3$, because a hidden low part reaches the observed high bits only through multiplication by d_i. Comparing the two attacks separates three things: what truncation exposes for any generator ($2^{4b} \to 2^{4k}$), the modest gain from the sparsity of the coupling, and the polynomial cost that the dynamics allow.

Distinguishing game. In G_{KP} with fresh nonces the adversary won 200/200 challenges for X0 (8 known bytes; advantage 1.00, 95% CI 0.963–1) and 60/60 for X1 (32 known bytes; advantage 1.00, CI 0.881–1), with one challenge, no further oracle queries and median times of 2.6 ms and 1.33 s.

Table 7 State recovery for the truncation extractor X1: generic exposure, guess-and-determine (G&D, $T = 8$) and interval attack

b	h	k	$\log_2$ generic	G&D $\log_2 W$	G&D success	T	Interval success	Ops	ms
16	8	8	32	26.5	50/50	8	20/20	224	3
16	6	10	40	34.2	12/12	–	–	–	–
20	10	10	40	34.0	10/10	–	–	–	–
16	5	11	44	37.6	3/3	–	–	–	–
16	8	8	32	26.5	50/50	12	20/20	352	5
20	8	12	48	–	–	16	20/20	600	8
24	8	16	64	–	–	16	20/20	600	8
24	8	16	64	–	–	24	20/20	920	12
16	4	12	48	–	–	24	20/20	920	21
24	4	20	80	–	–	40	10/10	1 872	46
32	8	24	96	–	–	32	10/10	1 736	82
32	4	28	112	–	–	40	10/10	2 496	158

State size $4b$ bits. Generic: states consistent with the first observation, 2^{4k}; the step from 2^{4b} to 2^{4k} is the exposure common to every truncated generator. Interval attack: median interval operations and median time in Python; known bytes $= 4Th/8$.

Table 8 Guess-and-determine baseline for X1 ($T = 8$ observed steps)

b	h	k	Recovered	Unique	Median $\log_2 W$	Range	$4k$	Speed-up (bit)	Median s
12	8	4	50/50	50/50	11.5	10.1–11.8	16	4.5	0.00
16	12	4	20/20	20/20	11.3	10.3–11.7	16	4.7	0.00
20	16	4	20/20	20/20	11.3	10.8–11.8	16	4.7	0.00
24	20	4	20/20	20/20	11.5	10.7–11.7	16	4.5	0.00
16	11	5	20/20	20/20	14.8	11.6–15.1	20	5.2	0.00
16	10	6	20/20	20/20	18.7	16.9–18.8	24	5.3	0.00
16	9	7	20/20	20/20	22.5	19.4–22.7	28	5.5	0.04
16	8	8	50/50	50/50	26.5	23.0–26.6	32	5.5	0.77
20	12	8	20/20	20/20	26.4	25.1–26.6	32	5.6	0.70
24	16	8	20/20	20/20	26.5	25.2–26.6	32	5.5	0.70
16	7	9	10/10	10/10	30.1	27.9–30.5	36	5.9	9.41
16	6	10	12/12	12/12	34.2	32.2–34.5	40	5.8	336.30
20	10	10	10/10	10/10	34.0	31.3–34.5	40	6.0	165.62
16	5	11	3/3	3/3	37.6	36.1–38.5	44	6.4	7252.22

Speed-up: $4k$ minus median $\log_2 W$, i.e. the gain over exhaustive search of the states consistent with the first observation. Rows with $h = 8$, $b \leq 16$ use image data; others use simulated output. Time: single-threaded C.

6.5 Keys and sessions

Equivalent keys (I1). Backward search from the recovered state lists every key that generates the observed keystream; in all trials the generating key was among them (100/100 at $b = 8$, 15/15 at $b = 16$), which checks the completeness of the enumeration, and any listed key decrypts. The class sizes have median 827 (IQR 514–1497, maximum 4181) at $b = 8$ and 603 (426–800, maximum 2326) at $b = 16$. Because keys are sampled uniformly, these samples are size-biased; their harmonic

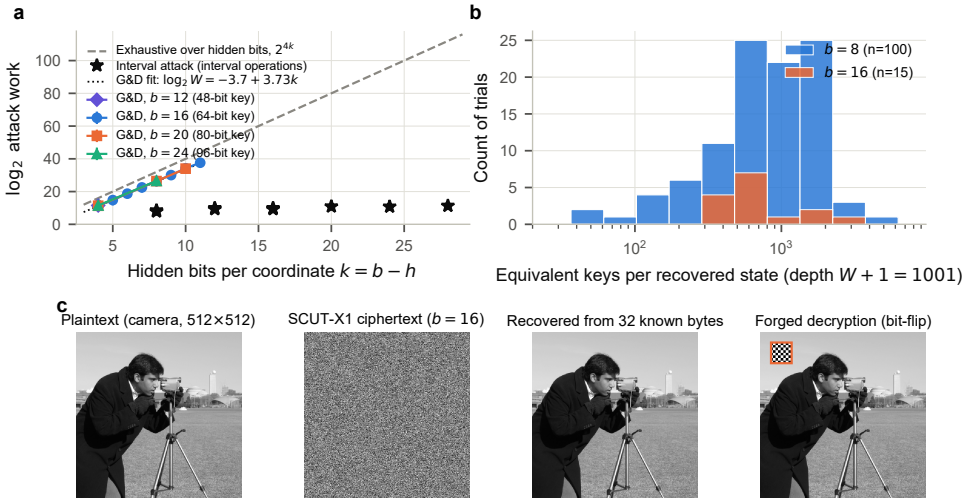


Fig. 6 **a** Cost of state recovery for X1 versus hidden bits per coordinate k: guess-and-determine work (markers: median; bars: range over trials) with its least-squares fit, interval-attack operations (stars) and exhaustive search over the states consistent with the first observation. **b** Equivalent keys found by backward search with Lemma 5 for $b = 8$ and $b = 16$. **c** Plaintext, SCUT-X1 ciphertext ($b = 16$), image recovered from 32 known bytes, and decryption of a ciphertext forged by XOR-ing a chosen 64×64 pattern into a known region (box)

means estimate the mean class size over states in the image of F_q^{W+1}, 461 and 554, close to the random-mapping value $1/(1 - \tau_{W+1}) = 503$, where $\tau_0 = 0$, $\tau_{j+1} = e^{\tau_j - 1}$ [30]. The equivalent keys are therefore what any non-injective map produces after a 1000-step warm-up, not a defect specific to SCUT-4D. The median search at $b = 16$ visited 4.65×10^5 nodes ($4q$ look-ups each) in 67 s.

Multi-session collisions (I2). Hashing a 128-bit key and a nonce into a $4b$-bit state cannot provide more than $4b$ bits of session-state entropy, and the warm-up reduces it further. The probability that two session states coincide after the warm-up is the mean class size divided by N: at $b = 16$ the collision entropy is 54.3 bits instead of 64, so a collision of session states, which reveals $P_1 \oplus P_2$ exactly as nonce reuse does, is expected with probability 1/2 after about $1.8 \times 10^8 \approx 2^{27.4}$ sessions instead of $2^{32.2}$. The simulation at $b = 8$ confirms the estimate: with $Q = 2^{17}$ sessions, 1774 colliding pairs were observed after warm-up against 2017 predicted and 2 expected for uniform states.

Generic stream-cipher controls. With fresh nonces, a state recovered in one session predicted another session's keystream only at chance level (byte agreement 0.0039, $n = 50$); re-using (K, N) revealed P_2 from P_1 in 50/50 trials; the unauthenticated cipher accepted all 10 000 random bit flips and the targeted forgery, whereas both AEADs rejected all (95% upper bound 3.7×10^{-4}). These are standard properties of unauthenticated XOR streams and are not attributed to the dynamics.

Table 9 Descriptive ciphertext statistics (reported for comparability only; not security evidence)

Image	Cipher	Entropy (bit)	r_H	r_V	r_D	NPCR % (1-bit P change)	NPCR % (1-bit K change)	UACI % (K change)
camera	SCUT-X1	7.9994	−0.0023	−0.0017	−0.0031	0.0004	99.60	33.41
camera	ChaCha20	7.9994	+0.0014	+0.0015	−0.0002	0.0004	99.59	33.41
astronaut (RGB)	SCUT-X1	7.9998	+0.0013	+0.0007	+0.0005	0.0001	99.61	33.46
astronaut (RGB)	ChaCha20	7.9998	−0.0031	−0.0003	−0.0011	0.0001	99.61	33.44
zeros	SCUT-X1	7.9993	−0.0030	+0.0035	+0.0019	0.0004	99.61	33.51
zeros	ChaCha20	7.9992	+0.0003	+0.0026	+0.0006	0.0004	99.61	33.37
checkerboard	SCUT-X1	7.9993	−0.0032	+0.0015	−0.0025	0.0004	99.60	33.44
checkerboard	ChaCha20	7.9993	−0.0014	−0.0014	+0.0013	0.0004	99.62	33.50

Table 10 Pass proportions of nine SP 800-22 tests (100 sequences $\times 10^6$ bits; acceptance threshold 0.960)

Test	os.urandom	ChaCha20	X0, $b = 16$	X1, $b = 16$	X1, $b = 12$	X0, $b = 8$
Frequency	1.00	1.00	0.98	0.98	0.99	0.36
Block frequency	0.99	0.99	0.97	0.55	0.63	0.77
Runs	0.98	0.99	1.00	0.99	0.97	0.12
Longest run	1.00	0.99	0.99	1.00	1.00	0.14
Serial 1	0.98	0.98	1.00	0.99	1.00	0.03
Serial 2	0.99	0.98	1.00	0.98	0.99	0.03
Approximate entropy	0.99	1.00	0.98	0.96	1.00	0.04
Cumulative sums	0.99	1.00	0.99	0.99	1.00	0.36
DFT	1.00	0.98	0.98	0.99	0.98	0.10
Prop. passed	9/9	9/9	9/9	7/9	8/9	0/9
Both criteria	9/9	9/9	8/9	7/9	8/9	0/9
Min. uniformity P	7.7×10^{-3}	3.3×10^{-2}	1.9×10^{-12}	3.8×10^{-130}	1.2×10^{-122}	2.6×10^{-175}

X0/X1 columns are SCUT-4D streams. Prop. passed: tests meeting the pass-proportion criterion. Both criteria: tests that also have uniform p-values ($P \geq 10^{-4}$). Proportions are rounded; 0.96 in the approximate-entropy row of X1, $b = 16$, is 96/100, just below 0.9602.

6.6 What the customary statistics report

SCUT-X1 and ChaCha20 are indistinguishable by the usual descriptive metrics (Table 9; camera image: entropy 7.9994 versus 7.9994 bits, $|r| < 0.004$), although SCUT-X1 falls to 32 known bytes; NPCR for a one-bit plaintext change is 0.0004% for both, as for any stream cipher. The NIST tests give mixed verdicts (Table 10). The $b = 8$ raw stream, which lies on a 484-step cycle, fails all nine tests. The $b = 16$ raw stream X0 meets the pass-proportion criterion in 9/9 tests and both criteria in 8/9; only its block-frequency p-values are non-uniform ($P = 1.9 \times 10^{-12}$), although eight known bytes break it. The truncated stream X1 fails the block-frequency test at $b = 16$ and 12 (pass proportions 0.55 and 0.63) and at $b = 16$ misses the approximate-entropy proportion by one sequence. A diagnostic over ten further keys traced the failure to the expansion: byte frequencies are uniform (mean $\chi^2 = 257.2$, 255 d.f.), but the Hamming weight of the top byte of coordinate i is correlated with its value one step later (coefficients 0.040, 0.024, 0.017, 0.013 for $d_i = 3, 5, 7, 9$; SD across keys ≤ 0.0015), close to $0.12/d_i$, because a top byte is roughly d_i times its predecessor plus a carry.

6.7 Performance

On the stated CPU, ChaCha20-Poly1305 encrypted and authenticated at 1579 MB/s and AES-256-GCM at 2243 MB/s (OpenSSL, likely using vector and AES instructions), whereas the SCUT keystream alone ran at 321 MB/s (X1, $b = 16$) and 402 MB/s (X0, $b = 8$) in a plain C loop. The comparison is between implementations, not algorithms, and supports only the statement that the authenticated baseline is not slower here. The table-based realisation also has a memory cost that grows as $4 \cdot 2^b$ bytes: 256 KiB at $b = 16$, 64 MiB at $b = 24$ and 16 GiB at $b = 32$.

7 Discussion

What the Lyapunov exponents do and do not tell us. Proposition 1 is stronger evidence of hyperchaos than the numerical spectra usually reported, and it holds for every orbit. Propositions 6 and 7 show what the same expansion means for someone who sees only the top bits of each coordinate: without coupling the hidden state is revealed at the rate of the topological entropy, and throughout the proven coupling range an efficient algorithm reveals it at least at the rate set by $3-2\pi\varepsilon$. The exponents guarantee that nearby states move apart. A keystream needs something different: future output must be hard to predict from past output. For a uniformly expanding map with truncated output and slowly varying coupling, the first property works against the second, and the expansion constant measures by how much.

What is generic and what is specific. Several observed weaknesses are generic: truncation exposes $4h$ bits per step for any generator; equivalent keys arise for any non-injective map iterated before output, in the numbers that random-mapping theory predicts; a $4b$-bit session state limits collision resistance whatever the master-key length; nonce reuse and malleability affect every unauthenticated XOR stream. The specific weaknesses come from the dynamics: an expanding map refines what the observer sees, so the sets of states consistent with the output are products of intervals that interval arithmetic can follow backwards at polynomial cost (Propositions 6, 7); and with odd-symmetric rounding the coupling splits the functional graph into many small basins and short cycles, a property of the integer version that changes with the rounding rule (Table 6). This is the interval-refinement principle of [20], extended here to a coupled four-dimensional map with explicit rates.

Implications for design and reporting. Our evidence supports four practices for chaos-based security papers: (i) specify the integer map, tables and rounding exactly, with test vectors; (ii) report the finite-state functional graph or cycle statistics against random-function and random-permutation controls; (iii) attack the actual output interface with known or chosen plaintext, including interval and constraint-propagation attacks that exploit expansion and smooth coupling, and report data, time and memory; (iv) use a standard AEAD for confidentiality and integrity, and if a chaotic stage is retained, place it inside the AEAD envelope with domain-separated keys and compare it with the same AEAD without the stage.

Limitations. (1) The analysis concerns one map family with integer multipliers and sine coupling, chosen because it can be analysed; dissipative attractors and maps without uniform expansion need their own analysis, and we do not claim that all

chaotic ciphers are insecure. (2) Proposition 7 needs $h \geq 4$ and $\varepsilon < 1/\pi$; outside this range we provide no bound. (3) SAT/SMT and algebraic attacks were not attempted; for X1 they are unnecessary at the tested sizes. (4) The near-equality $\lambda_i \approx \ln d_i$ is numerical; its perturbative explanation is not proved. (5) Cycle statistics at $b = 15, 16$ rest on 200 starts per precision and ten or twenty per control key; direct floating-point iteration was only searched for cycles up to 10^8 steps. (6) Only nine SP 800-22 tests were implemented, from the specification rather than with the reference software, and NIST no longer endorses the suite for cryptographic generators [46]. (7) Throughput compares a plain C loop with optimised AEAD libraries.

8 Conclusions

A four-dimensional sine-coupled expanding torus map was used to test the inference "provably hyperchaotic, hence suitable as a keystream generator". The real map has four provably positive Lyapunov exponents and entropy $\ln 945$; it is conjugate to four uncoupled multiplications, and the coupling barely changes its numerical spectrum. When the integer version outputs only the top bits of each coordinate, the hidden state is revealed at the rate of the topological entropy without coupling and, for every coupling in the proven range, at least at the rate set by the expansion bound. An interval-arithmetic attack built on this recovered 64- to 128-bit states in a fraction of a second. The integer map has many small basins and cycles several times shorter than those of random controls, by an amount that depends on the rounding rule; its equivalent keys and session-state collisions follow random-mapping theory. Positive Lyapunov exponents show that nearby states separate, not that the output is secret: when only part of the state is visible, the same expansion sets how fast the rest is revealed. Authenticated encryption with unique nonces remains the right baseline for image confidentiality and integrity.

Supplementary information. Online Resource 1 (ZIP) contains the C and Python source code, frozen coupling tables with SHA-256 manifest, test vectors, raw per-trial results, pinned package versions and figure scripts.

Declarations

Funding. The author declares that no funds, grants, or other support were received during the preparation of this manuscript.

Competing Interests. The author has no relevant financial or non-financial interests to disclose.

Author Contributions. Annas W. Malik is the sole author and carried out the conceptualisation, methodology, formal analysis, software, investigation, visualisation and writing, and approved the final manuscript.

Data Availability. The datasets generated and analysed during the current study (raw per-trial results, frozen coupling tables with SHA-256 manifest, test vectors) and all source code are provided as Online Resource 1 and are archived

in the Zenodo repository (https://doi.org/10.5281/zenodo.23814229; version 2: https://doi.org/10.5281/zenodo.23815286). The test images are public scikit-image data (version 0.26).

Code Availability. Source code (C and Python) is included in Online Resource 1 and archived with the data under the same DOI.

Ethics approval and dual-use statement. Not applicable (no human or animal subjects). The attacks are demonstrated only against the purpose-built SCUT-4D interfaces defined in this paper and do not target any deployed system or other published scheme.

References

[1] Fridrich, J.: Symmetric ciphers based on two-dimensional chaotic maps. Int. J. Bifurc. Chaos 8(6), 1259–1284 (1998). https://doi.org/10.1142/S021812749800098X

[2] Kocarev, L.: Chaos-based cryptography: a brief overview. IEEE Circuits Syst. Mag. 1(3), 6–21 (2001). https://doi.org/10.1109/7384.963463

[3] Jakimoski, G., Kocarev, L.: Chaos and cryptography: block encryption ciphers based on chaotic maps. IEEE Trans. Circuits Syst. I Fundam. Theory Appl. 48(2), 163–169 (2001). https://doi.org/10.1109/81.904880

[4] Shannon, C.E.: Communication theory of secrecy systems. Bell Syst. Tech. J. 28(4), 656–715 (1949). https://doi.org/10.1002/j.1538-7305.1949.tb00928.x

[5] Kong, S., Sun, J.: Adaptive image encryption scheme via 6D memristive conservative hyperchaos and ISO scrambling. Nonlinear Dyn. 114, 1106 (2026). https://doi.org/10.1007/s11071-026-12954-2

[6] Chen, H., Ma, S., Huang, Q., Liao, B., Chen, C., Huang, L., Chen, L.: A plaintext-aware quantum bit-level image cryptosystem based on 2D-GHSMM hyperchaotic map and involutive Bagua transform. Nonlinear Dyn. 114, 1174 (2026). https://doi.org/10.1007/s11071-026-13053-y

[7] Zhang, X., Hou, Y., Wang, G., Ye, X.: An improved Sprott-E chaotic system with coexisting attractors for regional DNA image encryption. Nonlinear Dyn. 114, 1105 (2026). https://doi.org/10.1007/s11071-026-13012-7

[8] Huang, Y., Qi, Y., Yang, F., Li, Z., Fan, M.: Application of recursive nonlinear coupled system driven by transcendental number-based robust chaos in speech encryption. Nonlinear Dyn. 114, 1231 (2026). https://doi.org/10.1007/s11071-026-13118-y

[9] Liu, W., Li, C., Sun, K., Wang, H.: Image secure communication system based on sparse threshold estimation and batch compressed sensing for IoT. Nonlinear Dyn. 114, 773 (2026). https://doi.org/10.1007/s11071-026-12644-z

[10] Hua, Z., Zhou, Y.: Image encryption using 2D Logistic-adjusted-Sine map. Inf. Sci. 339, 237–253 (2016). https://doi.org/10.1016/j.ins.2016.01.017

[11] Alvarez, G., Li, S.: Some basic cryptographic requirements for chaos-based cryptosystems. Int. J. Bifurc. Chaos 16(8), 2129–2151 (2006). https://doi.org/10.1142/S0218127406015970

[12] Preishuber, M., Hütter, T., Katzenbeisser, S., Uhl, A.: Depreciating motivation and empirical security analysis of chaos-based image and video encryption. IEEE Trans. Inf. Forensics Secur. 13(9), 2137–2150 (2018). https://doi.org/10.1109/TIFS.2018.2812080

[13] Özkaynak, F.: Brief review on application of nonlinear dynamics in image encryption. Nonlinear Dyn. 92, 305–313 (2018). https://doi.org/10.1007/s11071-018-4056-x

[14] Solak, E., Çokal, C., Yildiz, O.T., Biyikoğlu, T.: Cryptanalysis of Fridrich's chaotic image encryption. Int. J. Bifurc. Chaos 20(5), 1405–1413 (2010). https://doi.org/10.1142/S0218127410026563

[15] Li, C., Xie, T., Liu, Q., Cheng, G.: Cryptanalyzing image encryption using chaotic logistic map. Nonlinear Dyn. 78, 1545–1551 (2014). https://doi.org/10.1007/s11071-014-1533-8

[16] Li, C., Lin, D., Lü, J.: Cryptanalyzing an image-scrambling encryption algorithm of pixel bits. IEEE MultiMedia 24(3), 64–71 (2017). https://doi.org/10.1109/MMUL.2017.3051512

[17] Li, C., Lin, D., Lü, J., Hao, F.: Cryptanalyzing an image encryption algorithm based on autoblocking and electrocardiography. IEEE MultiMedia 25(4), 46–56 (2018). https://doi.org/10.1109/MMUL.2018.2873472

[18] Zhou, R., Yu, S.: On the divide-and-conquer attack of a plaintext related image chaotic encryption scheme. Nonlinear Dyn. 112, 11571–11594 (2024). https://doi.org/10.1007/s11071-024-09646-0

[19] Dhall, S., Yadav, K.: Cryptanalysis of substitution-permutation network based image encryption schemes: a systematic review. Nonlinear Dyn. 112(17), 14719–14744 (2024). https://doi.org/10.1007/s11071-024-09816-0

[20] Everett, S.: On the use of dynamical systems in cryptography. Chaos Solitons Fractals 183, 114952 (2024). https://doi.org/10.1016/j.chaos.2024.114952

[21] Grebogi, C., Ott, E., Yorke, J.A.: Roundoff-induced periodicity and the correlation dimension of chaotic attractors. Phys. Rev. A 38(7), 3688–3692 (1988). https://doi.org/10.1103/PhysRevA.38.3688

[22] Li, S., Chen, G., Mou, X.: On the dynamical degradation of digital piecewise linear chaotic maps. Int. J. Bifurc. Chaos 15(10), 3119–3151 (2005). https://doi.org/10.1142/S0218127405014052

[23] Persohn, K.J., Povinelli, R.J.: Analyzing logistic map pseudorandom number generators for periodicity induced by finite precision floating-point representation. Chaos Solitons Fractals 45(3), 238–245 (2012). https://doi.org/10.1016/j.chaos.2011.12.006

[24] Li, C., Feng, B., Li, S., Kurths, J., Chen, G.: Dynamic analysis of digital chaotic maps via state-mapping networks. IEEE Trans. Circuits Syst. I Regul. Pap. 66(6), 2322–2335 (2019). https://doi.org/10.1109/TCSI.2018.2888688

[25] Li, C., Tan, K., Feng, B., Lü, J.: The graph structure of the generalized discrete Arnold's cat map. IEEE Trans. Comput. 71(2), 364–377 (2022). https://doi.org/10.1109/TC.2021.3051387

[26] Wang, Q., Yu, S., Li, C., Lü, J., Fang, X., Guyeux, C., Bahi, J.M.: Theoretical design and FPGA-based implementation of higher-dimensional digital chaotic systems. IEEE Trans. Circuits Syst. I Regul. Pap. 63(3), 401–412 (2016). https://doi.org/10.1109/TCSI.2016.2515398

[27] Shen, Q., Zhang, Y., Xue, R., Li, X., Zhang, Y., Liu, W.: Counteracting digital chaotic degradation by dual delayed feedback and irrational rotation perturbation. Nonlinear Dyn. 114, 1175 (2026). https://doi.org/10.1007/s11071-026-13060-z

[28] Lu, X., Yang, Q.: Jacobi stability, invariant sets, hidden chaos in a radio-physical oscillator system. Nonlinear Dyn. 114, 1195 (2026). https://doi.org/10.1007/s11071-026-12973-z

[29] Eckmann, J.-P., Ruelle, D.: Ergodic theory of chaos and strange attractors. Rev. Mod. Phys. 57(3), 617–656 (1985). https://doi.org/10.1103/RevModPhys.57.617

[30] Flajolet, P., Odlyzko, A.M.: Random mapping statistics. In: Quisquater, J.-J., Vandewalle, J. (eds.) Advances in Cryptology – EUROCRYPT '89. Lecture Notes in Computer Science, vol. 434, pp. 329–354. Springer, Berlin (1990). https://doi.org/10.1007/3-540-46885-4_34

[31] Shub, M.: Endomorphisms of compact differentiable manifolds. Am. J. Math. 91(1), 175–199 (1969). https://doi.org/10.2307/2373276

[32] Krzyżewski, K., Szlenk, W.: On invariant measures for expanding differentiable mappings. Stud. Math. 33(1), 83–92 (1969). https://doi.org/10.4064/sm-33-1-83-92

[33] Ruelle, D.: An inequality for the entropy of differentiable maps. Bol. Soc. Bras. Mat. 9, 83–87 (1978). https://doi.org/10.1007/BF02584795

[34] Liu, P.-D.: Pesin's entropy formula for endomorphisms. Nagoya Math. J. 150, 197–209 (1998). https://doi.org/10.1017/S0027763000025113

[35] Katok, A., Hasselblatt, B.: Introduction to the Modern Theory of Dynamical Systems. Cambridge University Press, Cambridge (1995). https://doi.org/10.1017/CBO9780511809187

[36] Alexan, W., El-Damak, D., Gabr, M.: Image encryption based on Fourier-DNA coding for hyperchaotic Chen system, Chen-based binary quantization S-box, and variable-base modulo operation. IEEE Access 12, 21092–21113 (2024). https://doi.org/10.1109/ACCESS.2024.3363018

[37] Bowen, R.: Entropy for group endomorphisms and homogeneous spaces. Trans. Am. Math. Soc. 153, 401–414 (1971). https://doi.org/10.1090/S0002-9947-1971-0274707-X

[38] Katz, J., Lindell, Y.: Introduction to Modern Cryptography, 3rd edn. Chapman and Hall/CRC, Boca Raton (2020). https://doi.org/10.1201/9781351133036

[39] Bellare, M., Namprempre, C.: Authenticated encryption: relations among notions and analysis of the generic composition paradigm. J. Cryptol. 21(4), 469–491 (2008). https://doi.org/10.1007/s00145-008-9026-x

[40] Benettin, G., Galgani, L., Giorgilli, A., Strelcyn, J.-M.: Lyapunov characteristic exponents for smooth dynamical systems and for Hamiltonian systems; a method for computing all of them. Part 1: Theory. Meccanica 15(1), 9–20 (1980). https://doi.org/10.1007/BF02128236

[41] Wolf, A., Swift, J.B., Swinney, H.L., Vastano, J.A.: Determining Lyapunov exponents from a time series. Physica D 16(3), 285–317 (1985). https://doi.org/10.1016/0167-2789(85)90011-9

[42] Brent, R.P.: An improved Monte Carlo factorization algorithm. BIT 20(2), 176–184 (1980). https://doi.org/10.1007/BF01933190

[43] Nir, Y., Langley, A.: ChaCha20 and Poly1305 for IETF Protocols. RFC 8439. RFC Editor (2018). https://doi.org/10.17487/RFC8439

[44] Dworkin, M.: Recommendation for Block Cipher Modes of Operation: Galois/Counter Mode (GCM) and GMAC. NIST Special Publication 800-38D.

National Institute of Standards and Technology, Gaithersburg (2007). https://doi.org/10.6028/NIST.SP.800-38D

[45] Bassham, L., Rukhin, A., Soto, J., Nechvatal, J., Smid, M., Barker, E., Leigh, S., Levenson, M., Vangel, M., Banks, D., Heckert, A., Dray, J., Vo, S.: A Statistical Test Suite for Random and Pseudorandom Number Generators for Cryptographic Applications. NIST Special Publication 800-22 Rev. 1a. National Institute of Standards and Technology, Gaithersburg (2010). https://doi.org/10.6028/NIST.SP.800-22r1a

[46] National Institute of Standards and Technology: Decision to revise NIST SP 800-22 Rev. 1a. Computer Security Resource Center (2022). https://csrc.nist.gov/news/2022/decision-to-revise-nist-sp-800-22-rev-1a. Accessed 30 October 2026