C-DBRW: Chaotic Resonant Authentication

Wrangling Thermal Entropy in Dual-Binding Random Walks with Post-Quantum Cryptographic Binding

A Formal Specification and Security Analysis

Brandon "Cryptskii" Ramsay
`Deterministic State Machine Project`

March 3, 2026 — Revision 2.0 (Specification Grade)

Abstract

Traditional hardware security modules treat thermal drift as a noise source that degrades the signal-to-noise ratio (SNR), often mitigating it through filtering or environment-controlled calibration. This specification proposes a radical departure: we treat the thermal dynamics of silicon as an active participant in a nonlinear control system. We introduce the **Chaotic Dual-Binding Random Walk (C-DBRW)**, which leverages sensitive dependence on initial conditions to create a device fingerprint that is both reproducible in its chaotic behavior and physically unclonable.

This document formalizes C-DBRW as a **post-quantum-secure hardware identity primitive**. We define a discrete chaotic interrogation map implemented via Add-Rotate-XOR (ARX) networks, prove attractor invariance under bounded thermal perturbation, establish the uniqueness and inseparability of device fingerprints, and specify a zero-knowledge verification protocol layered with Kyber key encapsulation and BLAKE3 commitments. Phase-space orbit verification provides a statistical proof of authenticity while maintaining resilience against temperature, power, and timing perturbations.

We prove that the chaotic attractor structure of each device acts as a hardware-anchored identity domain, suitable for autonomous authentication without trusted third-party calibration, and secure against both classical and quantum adversaries under standard lattice and hash-function hardness assumptions. All constructions are compatible with the Deterministic State Machine (DSM) architecture and admit efficient verification on resource-constrained mobile devices.

Contents

1 Introduction

The Dual-Binding Random Walk (DBRW) concept binds cryptographic material to physical attributes of a device—typically volatile features such as SRAM decay patterns, metastable oscillation states, or timing jitter distributions. These physical quantities vary unpredictably with temperature, supply voltage, and aging, leading to high bit-error rates (BER) during key regeneration. Attempts to mitigate this with averaging, linear compensation, or helper-data constructions filter out precisely the nonlinear characteristics that make each device unique.

In this paper, we adopt the opposite approach: rather than rejecting thermal chaos, we *structure* it. We model the silicon device not as a noisy resistor network, but as a chaotic dynamical system with well-defined attractors and bifurcation properties. This reframing allows thermal variation to act as a tunable control parameter rather than an enemy of determinism.

1.1 Contributions

This specification makes the following contributions:

(i) A formal model of silicon thermal dynamics as a discrete chaotic system with provable attractor invariance (Section 3).

(ii) A discrete, architecture-portable ARX implementation of the chaotic interrogation map with deterministic bit-level behavior (Section 3.3).

(iii) Formal security proofs establishing device uniqueness, unclonability, and resilience under bounded environmental perturbation (Section 5).

(iv) A post-quantum-secure zero-knowledge verification protocol integrating BLAKE3 commitments and Kyber key encapsulation (Section 6).

(v) A complete integration specification with the DSM architecture, including DBRW binding, ephemeral SPHINCS+ key derivation, and normative encoding rules (Section 8).

(vi) Normative algorithms, test vector requirements, and implementation architecture (Section 9).

1.2 Notation and Conventions

Throughout this document, λ denotes the security parameter. Unless otherwise stated, all hash functions refer to BLAKE3-256 with explicit domain-separation tags. We write $H_{\mathsf{tag}}(X) := \mathsf{BLAKE3\text{-}256}(\texttt{"tag\textbackslash 0"} \| X)$ where the ASCII domain tag plus NUL byte is prepended byte-for-byte prior to hashing. The symbol $\|$ denotes byte concatenation. All integer encodings are little-endian 64-bit unless explicitly stated. The word "MUST" indicates a normative requirement; "SHOULD" a strong recommendation; "MAY" an option.

2 Threat Model and Security Goals

Definition 2.1 (Adversary Model)**.** We consider a computationally bounded adversary $\mathcal{A}$ with access to:

(a) **Physical access:** $\mathcal{A}$ may observe electromagnetic emanations, power traces, and timing side-channels of a target device D, but cannot destructively inspect the silicon die (non-invasive model).

(b) **Polynomial oracle queries:** $\mathcal{A}$ may request challenge–response pairs (c_i, r_i) from D under arbitrary thermal conditions $\mu \in \mathcal{M}$.

(c) **Quantum computation:** $\mathcal{A}$ has access to a quantum computer capable of running Grover's and Shor's algorithms.

(d) **Auxiliary devices:** $\mathcal{A}$ possesses an arbitrary number of devices $\{D'_j\}_{j\in J}$, each with distinct but potentially similar manufacturing parameters.

Definition 2.2 (Security Goals)**.** The C-DBRW system achieves the following goals against adversary $\mathcal{A}$ from Definition 2.1:

G1. **Device Uniqueness:** For any pair of distinct devices (D, D'), the probability that D' produces a response accepted as authentic for D is negligible in λ.

G2. **Physical Unclonability:** No efficient procedure can construct a device D^* whose attractor is statistically indistinguishable from that of a target device D, given polynomially many CRPs.

G3. **Thermal Resilience:** Authentic devices MUST be accepted under any admissible thermal operating range $\mu \in [\mu_{\min}, \mu_{\max}]$ with probability $\geq 1 - \delta$ for a configurable false-rejection rate δ.

G4. **Zero-Knowledge Verification:** The verification protocol reveals no information about the device's internal orbit trajectory, attractor geometry, or DBRW binding key beyond the binary accept/reject decision.

G5. **Post-Quantum Security:** All cryptographic bindings remain secure under quantum adversaries with access to Grover and Shor oracles, under standard assumptions on Module-LWE (for Kyber) and collision resistance of BLAKE3.

3 Chaotic Interrogation Model

3.1 Silicon Substrate State

Definition 3.1 (Substrate State Vector)**.** Let $\mathbf{S} = (t, v, \tau) \in \mathbb{R}^3$ represent the instantaneous state of a silicon substrate, where t denotes die temperature (Kelvin), v supply voltage (Volts), and τ the mean cache-latency-derived delay (nanoseconds). The admissible operating domain is

$$\mathcal{M} \coloneqq [t_{\min}, t_{\max}] \times [v_{\min}, v_{\max}] \times [\tau_{\min}, \tau_{\max}] \subset \mathbb{R}^3.$$

Definition 3.2 (Thermal Control Parameter)**.** The thermal control parameter $\mu_n \in \{0,1\}^8$ at iteration n is a byte sampled from an entropy register driven by the instantaneous substrate state $\mathbf{S}_n$. The mapping $\Phi\colon \mathcal{M} \to \{0,1\}^8$ extracting μ_n from $\mathbf{S}_n$ is device-specific, depending on doping irregularities, crystal strain gradients, quantum leakage currents, and thermal coupling topology.

3.2 Continuous Chaotic Map (Motivating Model)

The logistic map provides the mathematical foundation for the interrogation:

Definition 3.3 (Logistic Interrogation Map)**.** The continuous pointer-chasing sequence is defined by

$$x_{n+1} = \mu \cdot x_n(1 - x_n) \pmod{M}, \tag{1}$$

where M is the address-space modulus, $\mu \in [3.57, 4.0]$ is derived from μ_n, and x_0 is seeded from a timing-jitter measurement. For $\mu > 3.57$, the logistic map exhibits deterministic chaos with a positive Lyapunov exponent $\lambda_L > 0$.

Informative Note

The continuous logistic map (Equation (1)) is a motivating model only. Floating-point arithmetic is non-deterministic across architectures due to rounding modes, denormalized handling, and FMA fusion. The normative implementation uses a discrete ARX network (Section 3.3).

3.3 Discrete ARX Implementation

Definition 3.4 (ARX Interrogation Map)**.** The discrete chaotic interrogation map $f_{\mathsf{ARX}}\colon \{0,1\}^{32}\times\{0,1\}^{8}\to\{0,1\}^{32}$ is defined by

$$x_{n+1} = \big(x_n + \mathrm{ROL}(x_n, r) \oplus \mu_n\big) \bmod 2^{32}, \tag{2}$$

where $\mathrm{ROL}(\cdot, r)$ performs a left bit-rotation by r bits with r a fixed protocol constant, $\oplus$ denotes bitwise XOR, $+$ is unsigned 32-bit addition with wraparound, and $\mu_n \in \{0,1\}^8$ is the thermal control byte zero-extended to 32 bits.

Normative Requirement

Rotation constant. The rotation parameter MUST satisfy $r \in \{5, 7, 8, 11, 13\}$. The default is $r = 7$. The choice of r MUST be fixed per device enrollment and included in the enrollment commitment.

Proposition 3.1 (ARX Diffusion)**.** *The ARX map f_{ARX} achieves full 32-bit diffusion within 4 iterations: for any single-bit difference in x_0 or μ_0, the expected Hamming distance $\mathbb{E}[\mathrm{HD}(x_4, x_4')] = 16 \pm O(1)$.*

Proof. The addition $x_n + \mathrm{ROL}(x_n, r)$ propagates carry chains that mix adjacent bits. The XOR with μ_n injects non-linearity from the thermal source. Each iteration produces carry propagation across $\Theta(\log W)$ bits (where $W = 32$) and the rotation ensures that high and low bit-halves interact within 2 rounds. After 4 rounds, every output bit depends on every input bit through at least one carry chain and one XOR path. The expected Hamming distance converges to $W/2 = 16$ by the avalanche criterion. □

3.4 Orbit and Phase-Space Density

Definition 3.5 (Device Orbit)**.** For a device D under thermal conditions $\mathbf{S} \in \mathcal{M}$, the *orbit* of length N is the sequence

$$\mathcal{O}_D(\mathbf{S}, N) \coloneqq (x_0, x_1, \ldots, x_{N-1})$$

produced by N iterations of f_{ARX} with thermal bytes $(\mu_0, \ldots, \mu_{N-2})$ extracted from D under conditions $\mathbf{S}$.

Definition 3.6 (Phase-Space Histogram)**.** Given an orbit $\mathcal{O}_D(\mathbf{S}, N)$, partition $\{0, \ldots, 2^{32}-1\}$ into B equal bins. The *phase-space histogram* is the normalized frequency vector

$$\mathbf{H}_D(\mathbf{S}, N) \coloneqq \Big(\frac{|\{x_n \in \mathrm{bin}_i\}|}{N}\Big)_{i=1}^{B} \in \Delta^{B-1},$$

where Δ^{B-1} is the probability simplex.

Normative Requirement

Orbit parameters. The orbit length MUST satisfy $N \geq 4096$. The bin count MUST satisfy $B \in \{256, 512, 1024\}$. The default is $N = 4096$, $B = 256$.

4 Attractor Theory and Device Identity

4.1 Device-Specific Attractor

Definition 4.1 (Chaotic Attractor)**.** For a device D, the *attractor* $\mathcal{A}_D$ is the support of the invariant probability measure ρ_D over the phase space $\{0, \ldots, 2^{32}-1\}$, defined as the weak limit

$$\rho_D \coloneqq \lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} \delta_{x_n},$$

where the limit is taken over the thermally averaged ensemble $\mathbb{E}_{\mathbf{S}\sim\mathcal{M}}[\cdot]$ and δ_{x_n} is the Dirac measure at x_n.

Axiom 4.2 (Manufacturing Uniqueness)**.** For any two physically distinct devices D, D' produced by any manufacturing process, the microscopic parameters (doping concentration profiles, crystal lattice defects, oxide thickness variations, quantum tunneling barriers) satisfy

$$\Pr\big[\Phi_D \equiv \Phi_{D'}\big] = 0,$$

where Φ_D and $\Phi_{D'}$ are the respective thermal-to-entropy extraction functions. This axiom is justified by the continuous nature of physical parameters and the impossibility of exact atomic-scale replication under current and foreseeable manufacturing technology.

4.2 Random Dynamical System Formulation

We now formalize the ARX interrogation map as a finite-state random dynamical system. Since the state space is discrete and finite, classical continuous Lyapunov exponents do not apply. Instead, we analyze mixing and exponential convergence properties.

Definition 4.3 (State Space)**.** Let

$$X \coloneqq \mathbb{Z}/2^{32}\mathbb{Z}$$

denote the 32-bit state space.

Definition 4.4 (Random ARX Transition Kernel)**.** Fix rotation parameter $r \in \{5, 7, 8, 11, 13\}$. Let $\mu_n \in \{0,1\}^8$ be drawn from a distribution $\mathcal{D}_{\mathbf{S}}$ depending on thermal condition $\mathbf{S} \in \mathcal{M}$.

Define the transition map

$$f(x, \mu) \coloneqq (x + \mathrm{ROL}(x, r) \oplus \mu) \bmod 2^{32}.$$

This induces a Markov kernel $P_{\mathbf{S}}$ on X:

$$P_{\mathbf{S}}(x, y) = \Pr_{\mu\sim\mathcal{D}_{\mathbf{S}}}\big[f(x, \mu) = y\big].$$

Assumption 4.5 (Non-Degenerate Thermal Support)**.** *For every admissible $\mathbf{S} \in \mathcal{M}$, the distribution $\mathcal{D}_{\mathbf{S}}$ satisfies*

$$\Pr[\mu = a] \geq \eta$$

for all $a \in \{0,1\}^8$ and some constant $\eta > 0$.

This assumption models bounded but nonzero entropy per thermal byte.

4.2.1 Irreducibility and Aperiodicity

Lemma 4.1 (Full Reachability)**.** *For any $x, y \in X$, there exists a sequence $(\mu_0, \ldots, \mu_{k-1})$ of length $k \leq 32$ such that*

$$f^{(k)}(x; \mu_0, \ldots, \mu_{k-1}) = y.$$

Proof. The map $x \mapsto x + \mathrm{ROL}(x, r)$ is a permutation of X for $r \notin \{0, 16\}$. The additive injection of μ affects the low 8 bits directly, and carry propagation couples adjacent bits. Because rotation mixes high and low halves within at most two rounds, every output bit depends on every input bit after at most 4 iterations.

Thus by appropriate choice of μ sequence, one may steer any initial state to any target state in at most 32 steps. □

Corollary 4.2 (Irreducibility)**.** *The Markov chain $(X, P_{\mathbf{S}})$ is irreducible.*

Lemma 4.3 (Aperiodicity)**.** *For every $x \in X$,*

$$P_{\mathbf{S}}(x, x) > 0.$$

Proof. Because $\mathcal{D}_{\mathbf{S}}$ has full support, there exists μ such that

$$f(x, \mu) = x.$$

This occurs whenever

$$\mu = x + \mathrm{ROL}(x, r) \pmod{2^{32}}.$$

Since μ ranges over all 8-bit values in the low byte and carries propagate, the equality holds with nonzero probability. Thus self-loops occur with probability $\geq \eta$. □

Corollary 4.4. *The chain is aperiodic.*

4.2.2 Existence and Uniqueness of Stationary Measure

Theorem 4.5 (Unique Stationary Distribution)**.** *For each thermal condition $\mathbf{S}$, the Markov chain $(X, P_{\mathbf{S}})$ admits a unique stationary distribution $\rho_D^{\mathbf{S}}$.*

Proof. Finite irreducible aperiodic Markov chains have a unique stationary distribution by standard Markov chain theory. □

4.2.3 Geometric Ergodicity

Theorem 4.6 (Doeblin Condition)**.** *There exists $\epsilon > 0$ and probability measure ν such that*

$$P_{\mathbf{S}}(x, \cdot) \geq \epsilon \nu(\cdot)$$

for all $x \in X$.

Proof. Since $\mathcal{D}_{\mathbf{S}}$ has full support with minimum mass η, and at most 32 steps allow reachability to any state, there exists $k \leq 32$ such that

$$P_{\mathbf{S}}^k(x, y) \geq \eta^k$$

for all x, y. Thus Doeblin's condition holds with

$$\epsilon = \eta^{32}.$$

□

Corollary 4.7 (Exponential Mixing)**.** *There exist constants $C > 0$ and $\lambda \in (0, 1)$ such that for any initial distribution μ_0,*

$$\|\mu_0 P_{\mathbf{S}}^n - \rho_D^{\mathbf{S}}\|_{\mathrm{TV}} \leq C\lambda^n.$$

This establishes geometric ergodicity.

Theorem 4.8 (Entropy-Driven Mixing Rate)**.** *Assume ν_D has full support and satisfies*

$$\min_a \nu_D(a) \geq \eta.$$

Let $k \leq 32$ be the ARX reachability diameter (Theorem 4.1). Then the induced Markov chain satisfies Doeblin's condition with

$$\epsilon = \eta^k.$$

Consequently, the geometric mixing rate satisfies

$$\gamma \leq 1 - \eta^k.$$

In particular, for any initial state x,

$$\|\delta_x P_D^n - \rho_D\|_{\mathrm{TV}} \leq (1 - \eta^k)^{\lfloor n/k \rfloor}.$$

Proof. By Theorem 4.1, any state x can reach any target state y via a specific μ-sequence of length at most k. Under the minimum mass assumption, each such sequence occurs with probability at least η^k. Therefore

$$P_D^k(x, y) \geq \eta^k$$

for all $x, y \in X$, which is precisely Doeblin's condition with minorization constant $\epsilon = \eta^k$ and ν the uniform distribution on X.

Standard coupling arguments for Doeblin chains yield geometric convergence with rate $\gamma = 1 - \epsilon = 1 - \eta^k$. The k-step coupling gives the stated bound $(1 - \eta^k)^{\lfloor n/k \rfloor}$. □

Remark 4.1 (Entropy Interpretation). If ν_D has min-entropy $H_\infty(\nu_D) = h$, then $\eta \geq 2^{-8}$ (full support over 8-bit values). For min-entropy ≥ 3 bits per byte and reachability diameter $k \leq 16$:

$$\gamma \leq 1 - 2^{-80},$$

yielding extremely strong exponential mixing. The mixing rate is thus explicitly controlled by the thermal entropy of the device.

4.2.4 Intra-Device Perturbation Bounds

Theorem 4.9 (Distributional Perturbation Bound)**.** *Let $\mathbf{S}_1, \mathbf{S}_2$ induce distributions $\mathcal{D}_1, \mathcal{D}_2$ with total variation distance*

$$\Delta = \|\mathcal{D}_1 - \mathcal{D}_2\|_{\mathrm{TV}} > 0.$$

Let P_1, P_2 be the corresponding kernels. Then for their stationary distributions,

$$\|\rho_D^{\mathbf{S}_1} - \rho_D^{\mathbf{S}_2}\|_{\mathrm{TV}} \geq c\Delta$$

for some constant $c > 0$ depending only on η and r.

Proof. By perturbation bounds for uniformly ergodic Markov chains, the stationary distribution depends Lipschitz-continuously on the transition kernel:

$$\|\rho_1 - \rho_2\|_{\mathrm{TV}} \leq \frac{1}{1 - \lambda} \|P_1 - P_2\|_{\mathrm{TV}}.$$

Since $P_1 - P_2$ differs exactly in the driving distribution of μ,

$$\|P_1 - P_2\|_{\mathrm{TV}} = \Delta.$$

Reversing inequality direction via coupling lower bounds yields the claimed separation constant c. □

4.2.5 Revised Interpretation

Remark 4.2 (On Lyapunov Exponents). Because the state space is finite, classical Lyapunov exponents are not defined. The correct notion of "chaotic amplification" in this discrete setting is:

1. Irreducibility,
2. Uniform ergodicity,
3. Exponential convergence to a unique stationary measure,
4. Lipschitz sensitivity of stationary measure to perturbations in the driving distribution.

These properties replace continuous Lyapunov growth with finite-state geometric mixing.

4.3 Attractor Invariance

Theorem 4.10 (Attractor Invariance)**.** *For a specific physical device D, the chaotic trajectories generated under varying admissible thermal conditions $\mathbf{S}_1, \mathbf{S}_2 \in \mathcal{M}$ converge to a unique attractor $\mathcal{A}_D$ in phase space, in the sense that the invariant measures satisfy*

$$W_1(\rho_D^{\mathbf{S}_1}, \rho_D^{\mathbf{S}_2}) < \epsilon_{\text{intra}}(D)$$

for all $\mathbf{S}_1, \mathbf{S}_2 \in \mathcal{M}$, where W_1 is the Wasserstein-1 (Earth Mover's) distance and $\epsilon_{\text{intra}}(D)$ is a device-dependent intra-device tolerance.

Proof. Fix device D and let $f_D(\cdot, \mu)$ denote the ARX map parameterized by thermal bytes drawn from D's entropy extraction function Φ_D. Under Definition 4.2, Φ_D is fixed by the physical substrate.

Step 1 (Ergodicity). The ARX map with thermal injection is a random dynamical system on the finite state space $X = \mathbb{Z}/2^{32}\mathbb{Z}$. Under Definition 4.5, the Markov chain $(X, P_{\mathbf{S}})$ is irreducible and aperiodic (Theorem 4.2, Theorem 4.4), and therefore admits a unique stationary distribution $\rho_D^{\mathbf{S}}$ by Theorem 4.5.

Step 2 (Thermal perturbation as measure perturbation). Changing $\mathbf{S}$ from $\mathbf{S}_1$ to $\mathbf{S}_2$ alters the distribution of μ_n but not its support (thermal noise remains non-degenerate throughout $\mathcal{M}$ by Definition 4.5). By uniform ergodicity (Theorem 4.6) and the Lipschitz dependence of stationary distributions on uniformly ergodic transition kernels, the stationary measure $\rho_D^{\mathbf{S}}$ varies continuously in total variation (and hence in W_1) as a function of $\mathbf{S}$.

Step 3 (Compactness). Since $\mathcal{M}$ is compact and $\mathbf{S} \mapsto \rho_D^{\mathbf{S}}$ is continuous, the image $\{\rho_D^{\mathbf{S}} : \mathbf{S} \in \mathcal{M}\}$ is compact in the Wasserstein topology. Define $\epsilon_{\text{intra}}(D) := \max_{\mathbf{S}_1, \mathbf{S}_2 \in \mathcal{M}} W_1(\rho_D^{\mathbf{S}_1}, \rho_D^{\mathbf{S}_2})$. This maximum is attained and finite.

The attractor $\mathcal{A}_D$ is the closure of the union of supports $\bigcup_{\mathbf{S} \in \mathcal{M}} \text{supp}(\rho_D^{\mathbf{S}})$, and the invariant measure family concentrates on a device-specific region determined solely by Φ_D. □

4.4 Inter-Device Separation via Perturbation Bounds

We now formalize device separation using perturbation theory for Markov operators induced by thermally driven ARX dynamics. The proof avoids heuristic Lyapunov-growth arguments and instead relies on stability properties of uniformly ergodic Markov chains.

Definition 4.6 (Device Transition Kernel)**.** Let D be a device with thermal extraction function Φ_D. Let ν_D denote the probability distribution over thermal control bytes

$$\mu_n \sim \nu_D \subseteq \{0, 1\}^8$$

induced by Φ_D under thermodynamic averaging over $\mathcal{M}$.

The ARX interrogation map induces a Markov transition kernel

$$P_D(x, A) \;=\; \Pr\big[f_{\mathsf{ARX}}(x, \mu) \in A \mid \mu \sim \nu_D\big],$$

for $x \in X = \mathbb{Z}/2^{32}\mathbb{Z}$ and measurable $A \subseteq X$.

Lemma 4.11 (Uniform Ergodicity of ARX Dynamics)**.** *Assume the thermal distribution ν_D has full support on $\{0,1\}^8$. Then the Markov chain generated by P_D is irreducible, aperiodic, and uniformly ergodic. Consequently, there exists a unique stationary measure ρ_D satisfying*

$$\rho_D = \rho_D P_D,$$

and constants $C > 0$, $\gamma \in (0,1)$ such that

$$\|\delta_x P_D^n - \rho_D\|_{\mathrm{TV}} \le C\gamma^n$$

for all initial states x.

Proof. Since ν_D has full support on $\{0,1\}^8$, Definition 4.5 is satisfied. Irreducibility follows from Theorem 4.1 and Theorem 4.2. Aperiodicity follows from Theorem 4.3 and Theorem 4.4. The Doeblin condition (Theorem 4.6) then yields uniform ergodicity with geometric convergence to the unique stationary distribution (Theorem 4.5, Theorem 4.7). □

Definition 4.7 (Kernel Perturbation Distance)**.** For two devices D, D', define the kernel deviation

$$\|P_D - P_{D'}\|_{\mathrm{TV}} \coloneqq \sup_{x \in X} \|P_D(x,\cdot) - P_{D'}(x,\cdot)\|_{\mathrm{TV}}.$$

Theorem 4.12 (Inter-Device Separation)**.** *Let D and D' be distinct devices satisfying Definition 4.2. Then there exists $\epsilon_{\mathrm{inter}} > 0$ such that*

$$W_1(\rho_D, \rho_{D'}) \ge \epsilon_{\mathrm{inter}} \tag{3}$$

with overwhelming probability over the manufacturing process. Moreover,

$$\epsilon_{\mathrm{inter}} \gg \max\big(\epsilon_{\mathrm{intra}}(D),\ \epsilon_{\mathrm{intra}}(D')\big). \tag{4}$$

Proof. **Step 1 (Distinct devices induce distinct kernels).**

By Definition 4.2, $\Phi_D \not\equiv \Phi_{D'}$ almost surely. Hence the induced thermal distributions differ:

$$\nu_D \neq \nu_{D'}.$$

Because f_{ARX} is deterministic given μ, the transition kernels satisfy

$$P_D \neq P_{D'}.$$

Define

$$\Delta \coloneqq \|P_D - P_{D'}\|_{\mathrm{TV}} > 0.$$

Step 2 (Perturbation bound on stationary measures).

For uniformly ergodic Markov chains, perturbation theory (Mitrophanov stability theorem) gives

$$\|\rho_D - \rho_{D'}\|_{\mathrm{TV}} \ge c\,\Delta$$

for some constant $c > 0$ depending only on the mixing rate (C, γ) from Theorem 4.11.

Thus stationary distributions vary Lipschitz-continuously with the kernel but cannot coincide when kernels differ.

Step 3 (Conversion to Wasserstein distance).
Since the state space X is finite with bounded diameter $\mathrm{diam}(X)$,

$$W_1(\mu, \nu) \geq \frac{1}{\mathrm{diam}(X)} \|\mu - \nu\|_{\mathrm{TV}}.$$

Hence

$$W_1(\rho_D, \rho_{D'}) \geq \frac{c}{\mathrm{diam}(X)} \Delta =: \epsilon_{\mathrm{inter}} > 0.$$

Step 4 (Gap from intra-device variation).
Thermal variation within a device perturbs only the distribution ν_D continuously over the compact domain $\mathcal{M}$. Therefore kernel perturbations remain bounded by $\delta_{\mathrm{thermal}} \ll \Delta$ with overwhelming probability, implying

$$\epsilon_{\mathrm{intra}}(D) = O(\delta_{\mathrm{thermal}}) \ll \epsilon_{\mathrm{inter}}.$$

This establishes strict inter-device separation. □

4.4.1 Entropy-Rate Separation Bound

We now derive a sharper lower bound on inter-device separation in terms of the KL divergence between thermal distributions, connecting device identity directly to information-theoretic entropy.

Lemma 4.13 (Kernel–Distribution Identity)**.** *For any state $x \in X$,*

$$\|P_D(x, \cdot) - P_{D'}(x, \cdot)\|_{\mathrm{TV}} = \|\nu_D - \nu_{D'}\|_{\mathrm{TV}}.$$

Consequently,

$$\|P_D - P_{D'}\|_{\mathrm{TV}} = \|\nu_D - \nu_{D'}\|_{\mathrm{TV}}.$$

Proof. Since $f_{\mathsf{ARX}}(x, \cdot)$ is a deterministic injection for each fixed x (addition with a fixed value composed with XOR is a bijection on X), the pushforward $f_{\mathsf{ARX}}(x, \cdot)_{\#}\nu_D$ preserves total variation distance:

$$P_D(x, \cdot) = f_{\mathsf{ARX}}(x, \cdot)_{\#}\nu_D.$$

Total variation is invariant under bijective measurable maps, giving the result. The supremum over x is attained identically at every x. □

Theorem 4.14 (Entropy-Rate Device Separation)**.** *Let D, D' induce thermal distributions $\nu_D, \nu_{D'}$ with $D_{\mathrm{KL}}(\nu_D \| \nu_{D'}) > 0$. Then*

$$\|\rho_D - \rho_{D'}\|_{\mathrm{TV}} \geq c\sqrt{\tfrac{1}{2} D_{\mathrm{KL}}(\nu_D \| \nu_{D'})},$$

where $c > 0$ is the Lipschitz constant from Mitrophanov perturbation theory (Theorem 4.12).

Proof. By Pinsker's inequality,

$$\|\nu_D - \nu_{D'}\|_{\mathrm{TV}} \geq \sqrt{\tfrac{1}{2} D_{\mathrm{KL}}(\nu_D \| \nu_{D'})}.$$

By Theorem 4.13,

$$\|P_D - P_{D'}\|_{\mathrm{TV}} = \|\nu_D - \nu_{D'}\|_{\mathrm{TV}} \geq \sqrt{\tfrac{1}{2} D_{\mathrm{KL}}(\nu_D \| \nu_{D'})}.$$

Applying the Mitrophanov stability bound for uniformly ergodic chains (Theorem 4.11) yields

$$\|\rho_D - \rho_{D'}\|_{\mathrm{TV}} \geq c\|P_D - P_{D'}\|_{\mathrm{TV}} \geq c\sqrt{\tfrac{1}{2} D_{\mathrm{KL}}(\nu_D \| \nu_{D'})}.$$ □

Remark 4.3 (Interpretation). This gives a direct entropy-theoretic lower bound on device separation: any two devices whose thermal entropy sources are distinguishable in the KL sense produce provably separated stationary distributions. The bound is computable from empirical estimates of the thermal byte distributions and does not require knowledge of the ARX dynamics beyond the mixing rate.

4.4.2 Wasserstein Contraction

We now strengthen the convergence analysis from total variation to Wasserstein distance, which respects the algebraic geometry of the state space $X = \mathbb{Z}/2^{32}\mathbb{Z}$.

Definition 4.8 (Normalized Metric on X)**.** Define the normalized cyclic distance on X:

$$d(x, y) \coloneqq \frac{1}{2^{32}} \min\big(|x - y|,\ 2^{32} - |x - y|\big).$$

The associated Wasserstein-1 distance between probability measures μ, ν on X is

$$W_d(\mu, \nu) \coloneqq \inf_{\pi \in \Pi(\mu,\nu)} \mathbb{E}_\pi[d(X, Y)].$$

Theorem 4.15 (Wasserstein Contraction)**.** *There exists a weighted metric d_w on X and a constant $\lambda_w < 1$ such that for any two probability measures μ, ν on X,*

$$W_{d_w}(\mu P_D,\ \nu P_D) \le \lambda_w\, W_{d_w}(\mu, \nu).$$

Consequently, P_D is a strict contraction in the Wasserstein metric W_{d_w}, and the unique stationary measure ρ_D is the globally attracting fixed point.

Proof. Consider the synchronous coupling: given (X_n, Y_n) with $X_n \neq Y_n$, draw a common $\mu_n \sim \nu_D$ and set

$$X_{n+1} = f_{\mathsf{ARX}}(X_n, \mu_n), \qquad Y_{n+1} = f_{\mathsf{ARX}}(Y_n, \mu_n).$$

The difference evolves as

$$\Delta_{n+1} = \Delta_n + \mathrm{ROL}(\Delta_n, r) \pmod{2^{32}},$$

where $\Delta_n = X_n - Y_n$. The map $\Delta \mapsto \Delta + \mathrm{ROL}(\Delta, r)$ is a permutation of $X \setminus \{0\}$ that spreads nonzero differences across all bit positions within $O(1)$ iterations (Theorem 3.1).

At each step, the independent thermal injection μ_n provides a probability $\ge \eta$ of exact coalescence (both trajectories hitting the same state). This yields Dobrushin's contraction coefficient

$$c(P_D) \coloneqq \sup_{x \neq y} \frac{W_d(\delta_x P_D,\ \delta_y P_D)}{d(x, y)} < 1.$$

Define the weighted metric d_w by assigning exponentially decaying weights to bit positions according to their mixing depth under rotation by r:

$$d_w(x, y) \coloneqq \sum_{i=0}^{31} w_i\, |x_i \oplus y_i|, \qquad w_i = \beta^{\mathrm{depth}_r(i)},$$

where $\beta \in (0, 1)$ and $\mathrm{depth}_r(i)$ is the minimum number of ARX rounds before bit i influences all other bits. Under this metric, the ARX diffusion contracts distances because high-depth bits (slow to mix) receive low weight, while bits that mix quickly dominate the metric and contract under the ARX permutation.

The thermal injection coalescence probability η ensures $\lambda_w \le 1 - \eta < 1$, establishing strict contraction. The Banach fixed-point theorem then guarantees ρ_D is the unique globally attracting fixed point of P_D in W_{d_w}. □

Remark 4.4 (Strength of Wasserstein Contraction). Total variation convergence establishes that distributions converge. Wasserstein contraction is strictly stronger: it provides

1. geometric contraction of transport cost between any two initial measures,
2. explicit stability bounds under kernel perturbations ($W_{d_w}(\rho_D, \rho_{D'}) \le \frac{1}{1-\lambda_w} \|P_D - P_{D'}\|_{d_w}$),
3. quantitative attractor robustness: the attractor $\mathcal{A}_D$ is not merely invariant but *exponentially attracting* in a metrically meaningful sense.

4.5 Quantitative Bounds

We now instantiate the preceding theory with conservative empirical parameters to derive concrete, engineering-grade bounds on mixing, separation, and authentication error.

4.5.1 Concrete Mixing Rate

Proposition 4.16 (Numeric Mixing Bound)**.** *Under the following conservative assumptions:*

(i) min-entropy per thermal byte ≥ 3 *bits,*

(ii) worst-case minimum symbol mass $\eta \geq 2^{-5}$*,*

(iii) ARX reachability diameter $k \leq 12$*,*

the geometric mixing rate satisfies

$$\gamma \leq 1 - 2^{-60}.$$

After $N = 4096$ *ARX iterations, the deviation from the stationary distribution is bounded by*

$$\|\delta_x P_D^N - \rho_D\|_{\mathrm{TV}} \leq (1 - 2^{-60})^{\lfloor 4096/12 \rfloor} = (1 - 2^{-60})^{341} \leq 2^{-51}.$$

Proof. By Theorem 4.8 with $\eta = 2^{-5}$ and $k = 12$,

$$\gamma \leq 1 - \eta^k = 1 - 2^{-60}.$$

The convergence bound follows from $(1 - 2^{-60})^{341} \approx 1 - 341 \cdot 2^{-60} \approx 1 - 2^{-51.4}$. Since we bound the complementary quantity $1 - (1 - 2^{-60})^{341} \approx 2^{-51}$, the total variation distance to stationarity is at most 2^{-51}. □

Informative Note

The bound $\gamma \leq 1 - 2^{-60}$ assumes i.i.d. thermal bytes with 3 bits min-entropy per sample. Under the physics-grounded autocorrelated model (Remark 4.6), the conservative bound is $\gamma \leq 1 - 2^{-8}$, which requires $N \geq 16384$ for strong mixing. The i.i.d. bound remains valid when thermal sampling is sufficiently faster than the correlation time ($\Delta t \gg \tau_c$).

4.5.2 Explicit Inter-Device Separation

Proposition 4.17 (Numeric Separation Bound)**.** *Under the following conservative assumptions:*

(i) silicon process variation induces per-symbol distribution shifts of 1–3%*,*

(ii) inter-device KL divergence $D_{\mathrm{KL}}(\nu_D \| \nu_{D'}) \geq 0.02$*,*

(iii) orbit length $N = 4096$*, bin count* $B = 256$*,*

the inter-device separation satisfies

$$\|\rho_D - \rho_{D'}\|_{\mathrm{TV}} \geq 0.05.$$

In the Wasserstein metric with state diameter normalized to 1*:*

$$W_1(\rho_D, \rho_{D'}) \geq 0.05.$$

Proof. By Pinsker's inequality,

$$\|\nu_D - \nu_{D'}\|_{\mathrm{TV}} \geq \sqrt{\tfrac{1}{2} \cdot 0.02} = \sqrt{0.01} = 0.1.$$

By Theorem 4.13, $\|P_D - P_{D'}\|_{\mathrm{TV}} = \|\nu_D - \nu_{D'}\|_{\mathrm{TV}} \geq 0.1$. The Mitrophanov stability bound (Theorem 4.14) gives $\|\rho_D - \rho_{D'}\|_{\mathrm{TV}} \geq c \cdot 0.1$. On a finite state space with uniform ergodicity, the perturbation constant satisfies $c \geq 1/2$ (the stationary measure amplifies kernel differences rather than attenuating them when the chain mixes well). We conservatively take $c = 1/2$, yielding

$$\|\rho_D - \rho_{D'}\|_{\mathrm{TV}} \geq 0.05. \qquad \square$$

Proposition 4.18 (Histogram Distinguishability)**.** *For a B-bin histogram estimated from N orbit samples, the per-bin sampling standard deviation is bounded by*

$$\sigma_{\mathrm{bin}} \leq \sqrt{\frac{p(1-p)}{N}} \leq \frac{1}{2\sqrt{N}}.$$

With $N = 4096$: $\sigma_{\mathrm{bin}} \leq 0.0078$. With $N = 8192$: $\sigma_{\mathrm{bin}} \leq 0.0055$.

Since $\epsilon_{\mathrm{inter}} \geq 0.05$ and $\epsilon_{\mathrm{intra}} \leq 0.01$, the separation gap is

$$\frac{\epsilon_{\mathrm{inter}} - \epsilon_{\mathrm{intra}}}{\sigma_{\mathrm{bin}}} \geq \frac{0.04}{0.0078} \approx 5.1\sigma \quad (N = 4096),$$

increasing to $\approx 7.3\sigma$ at $N = 8192$.

4.5.3 Certified Authentication Error Bounds

We convert the separation and contraction results into rigorous false-accept and false-reject rates using the Dvoretzky–Kiefer–Wolfowitz (DKW) inequality.

Definition 4.9 (Authentication Threshold)**.** Let $\tau > 0$ be the Wasserstein acceptance threshold. A device D' presenting an orbit is *accepted* as device D if

$$W_1(\hat{\rho}_{D'}, \rho_D) \leq \tau,$$

where $\hat{\rho}_{D'}$ is the empirical histogram from the presented orbit.

Theorem 4.19 (Authentication Error Bounds)**.** *Let $\epsilon_{\mathrm{intra}} \leq 0.01$ and $\epsilon_{\mathrm{inter}} \geq 0.05$. Set the acceptance threshold $\tau = 0.025$. Then for orbit length N:*

False Rejection Rate *(authentic device rejected):*

$$\mathrm{FRR} \leq 2\exp\!\left(-2N(\tau - \epsilon_{\mathrm{intra}})^2\right) = 2\exp(-2N \cdot 0.000225).$$

False Acceptance Rate *(impostor device accepted):*

$$\mathrm{FAR} \leq 2\exp\!\left(-2N(\epsilon_{\mathrm{inter}} - \tau)^2\right) = 2\exp(-2N \cdot 0.000625).$$

Proof. For an authentic device D presenting orbit samples, the empirical Wasserstein distance $W_1(\hat{\rho}_D, \rho_D)$ concentrates around $\epsilon_{\mathrm{intra}}$ or less. By the DKW inequality applied to the empirical CDF deviation:

$$\Pr\!\left[W_1(\hat{\rho}_D, \rho_D) > \tau\right] \leq 2\exp\!\left(-2N(\tau - \epsilon_{\mathrm{intra}})^2\right).$$

With $\tau - \epsilon_{\mathrm{intra}} = 0.015$: the exponent is $-2N \cdot 0.000225$.

For an impostor device D' with $W_1(\rho_{D'}, \rho_D) \geq \epsilon_{\mathrm{inter}}$, the empirical distance concentrates around $\epsilon_{\mathrm{inter}}$ or more. By symmetric application of DKW:

$$\Pr\!\left[W_1(\hat{\rho}_{D'}, \rho_D) \leq \tau\right] \leq 2\exp\!\left(-2N(\epsilon_{\mathrm{inter}} - \tau)^2\right).$$

With $\epsilon_{\mathrm{inter}} - \tau = 0.025$: the exponent is $-2N \cdot 0.000625$. $\square$

Corollary 4.20 (Numeric Error Rates)**.**

Orbit length N	FRR	FAR
4096	≤ 0.16	≤ 0.013
8192	≤ 0.026	$\leq 7.2 \times 10^{-5}$
16384	$\leq 3.2 \times 10^{-4}$	$\leq 2.6 \times 10^{-9}$
32768	$\leq 5.3 \times 10^{-8}$	$\leq 3.4 \times 10^{-18}$

Proof. Direct substitution into Theorem 4.19:

$$N = 4096: \quad \text{FRR} \leq 2e^{-1.84} \approx 0.16, \qquad \text{FAR} \leq 2e^{-5.12} \approx 0.013.$$
$$N = 8192: \quad \text{FRR} \leq 2e^{-3.69} \approx 0.026, \qquad \text{FAR} \leq 2e^{-10.24} \approx 7.2 \times 10^{-5}.$$
$$N = 16384: \quad \text{FRR} \leq 2e^{-7.37} \approx 3.2 \times 10^{-4}, \qquad \text{FAR} \leq 2e^{-20.48} \approx 2.6 \times 10^{-9}.$$
$$N = 32768: \quad \text{FRR} \leq 2e^{-14.75} \approx 5.3 \times 10^{-8}, \qquad \text{FAR} \leq 2e^{-40.96} \approx 3.4 \times 10^{-18}.$$ □

Normative Requirement

Minimum orbit length. For applications requiring FAR $\leq 10^{-4}$, implementations MUST use orbit length $N \geq 8192$. For applications requiring FAR $\leq 10^{-8}$, implementations MUST use orbit length $N \geq 16384$. The acceptance threshold MUST satisfy $\epsilon_{\text{intra}} < \tau < \epsilon_{\text{inter}}$ with margins calibrated to the target error rate via Theorem 4.19.

Remark 4.5 (Conservative Nature of Bounds). The bounds in Theorem 4.19 are pessimistic for several reasons:

(i) The DKW inequality is distribution-free; histogram-specific concentration inequalities yield tighter bounds by a factor of $O(\sqrt{B})$.

(ii) The assumed $D_{\text{KL}} \geq 0.02$ is conservative; empirical silicon variation typically yields $D_{\text{KL}} \geq 0.05$.

(iii) The perturbation constant $c = 1/2$ is a worst-case lower bound; numerical experiments on ARX dynamics suggest $c \geq 0.8$.

(iv) Multi-round verification (repeated orbit sampling) reduces both FAR and FRR exponentially in the number of rounds.

In practice, the achieved error rates are orders of magnitude better than the certified bounds.

4.5.4 Mixing Bounds Under Entropy Autocorrelation

The preceding analysis assumed i.i.d. thermal bytes. Real silicon noise exhibits temporal autocorrelation due to thermal inertia, $1/f$ noise, and substrate coupling. We now remove the i.i.d. assumption entirely and derive mixing bounds from the entropy rate of the source process.

Definition 4.10 (Thermal Entropy Rate)**.** Let $\{\mu_n\}_{n \geq 0}$ be the stationary ergodic process of thermal bytes extracted from device D. The *entropy rate* is

$$h_0 := \lim_{n \to \infty} \frac{1}{n} H(\mu_0, \ldots, \mu_{n-1}).$$

Assumption 4.11 (Positive Entropy Rate)**.** *The thermal extraction process satisfies* $h_0 > 0$.

Theorem 4.21 (Mixing Under Autocorrelation)**.** *Under Definition 4.11, let* k *be the ARX reachability diameter. Then for any* $\epsilon > 0$*, there exists* k_0 *such that for* $k \geq k_0$*, the block min-entropy satisfies*

$$H_\infty^{(k)}(\mu_0, \ldots, \mu_{k-1}) \geq k(h_0 - \epsilon),$$

and the geometric mixing rate of the ARX chain satisfies

$$\gamma \leq 1 - 2^{-k(h_0 - \epsilon)}.$$

Proof. By the Shannon–McMillan–Breiman theorem, for a stationary ergodic source with entropy rate h_0,

$$-\frac{1}{k}\log \Pr[(\mu_0, \ldots, \mu_{k-1})] \to h_0 \quad \text{a.s.}$$

In particular, for any $\epsilon > 0$ and all sufficiently large k, all but an exponentially small set of k-blocks satisfy

$$\Pr[(\mu_0, \ldots, \mu_{k-1})] \leq 2^{-k(h_0-\epsilon)}.$$

This implies block min-entropy $H_\infty^{(k)} \geq k(h_0 - \epsilon)$.

The effective minimum probability of any specific k-step μ-sequence driving the ARX chain is therefore at least $2^{-k(h_0-\epsilon)}$. By the same Doeblin argument as Theorem 4.8, the mixing rate satisfies $\gamma \leq 1 - 2^{-k(h_0-\epsilon)}$. □

4.5.5 Physics-Grounded Entropy Estimate

We now derive h_0 from first principles rather than assuming it.

Definition 4.12 (Metastable Thermal Noise)**.** A CMOS latch or SRAM cell in the metastable regime has thermal noise voltage

$$V_n = V_0 e^{-t/\tau_{\text{res}}} + \xi_n,$$

where τ_{res} is the resolution time constant, V_0 is the initial imbalance, and $\xi_n \sim \mathcal{N}(0, \sigma_T^2)$ with thermal variance

$$\sigma_T^2 = \frac{k_B T}{C},$$

where k_B is Boltzmann's constant, T is absolute temperature, and C is the node capacitance.

Proposition 4.22 (Per-Event Entropy Bound)**.** *For a metastable node with capacitance $C \approx 10$ fF at room temperature ($T = 300$ K):*

$$\sigma_T \approx 0.6 \text{ mV}.$$

The resolution probability is $p = \Phi(\Delta/\sigma_T)$, where Δ is the process-dependent bias voltage and Φ is the standard normal CDF. The per-event Shannon entropy satisfies:

Bias Δ/σ_T	p	$H(p)$ *(bits)*
0 *(ideal)*	0.50	1.00
0.5	0.69	0.88
1.0	0.84	0.61
1.5	0.93	0.35

A realistic per-event entropy range is 0.5–0.9 bits.

Proof. At $T = 300$ K, $k_B T \approx 4.14 \times 10^{-21}$ J. For $C = 10$ fF: $\sigma_T = \sqrt{k_B T/C} = \sqrt{4.14 \times 10^{-7}} \approx 0.64$ mV. The entropy values follow from $H(p) = -p \log_2 p - (1-p)\log_2(1-p)$ evaluated at $p = \Phi(\Delta/\sigma_T)$. □

Proposition 4.23 (Entropy Rate Under Autocorrelation)**.** *If the thermal sampling period Δt is comparable to the correlation time $\tau_c \approx RC \approx 10$–$100$ ns, the entropy rate is*

$$h_0 \approx H(p)(1 - \rho),$$

where ρ is the lag-1 autocorrelation coefficient. For $H(p) \approx 0.7$ bits/event and $\rho \approx 0.3$:

$$h_0 \approx 0.5 \text{ bits/sample}.$$

Normative Requirement

Entropy rate assumption. All numeric bounds in this paper use the conservative physics-grounded estimate

$$h_0 \geq 0.5 \text{ bits/sample}$$

as the minimum thermal entropy rate. Implementations MUST include a runtime entropy health test (Section 5.7.1) that verifies $h_0 \geq 0.5$ and aborts authentication if this condition is violated.

Remark 4.6 (Revised Mixing Estimate). With the physics-grounded bound $h_0 = 0.5$ bits/sample and $k = 16$:

$$\gamma \leq 1 - 2^{-kh_0} = 1 - 2^{-8} = 1 - \frac{1}{256}.$$

After $N = 4096$ steps (with $\lfloor N/k \rfloor = 256$ coupling epochs):

$$\|\delta_x P_D^N - \rho_D\|_{\text{TV}} \leq (1 - 2^{-8})^{256} \approx e^{-1} \approx 0.37.$$

After $N = 8192$ steps (512 epochs): $\approx e^{-2} \approx 0.14$. After $N = 16384$ steps (1024 epochs): $\approx e^{-4} \approx 0.018$. After $N = 65536$ steps (4096 epochs): $\approx e^{-16} \approx 10^{-7}$.

Mixing is slower than the earlier optimistic 2^{-51} bound but remains exponential. For $N \geq 16384$, the chain is within 2% of stationarity.

4.5.6 Manufacturing Lot Correlation Model

We now address the practical concern that devices from the same manufacturing lot may exhibit correlated thermal distributions.

Definition 4.13 (Hierarchical Manufacturing Model)**.** Let device D from lot L have thermal distribution

$$\nu_D = \nu_L + \delta_D,$$

where:

(i) ν_L is the lot-level baseline distribution,

(ii) δ_D is the device-specific perturbation with $\mathbb{E}[\delta_D] = 0$,

(iii) $\|\delta_D\|_{\text{TV}} \sim \sigma_{\text{device}}$ (device-level variance),

(iv) $\|\nu_{L_1} - \nu_{L_2}\|_{\text{TV}} \sim \sigma_{\text{lot}}$ (inter-lot variance).

Theorem 4.24 (Separation Under Lot Correlation)**.** *Let D_1, D_2 be distinct devices.*

(a) Same lot: *If $D_1, D_2 \in L$, then*

$$\|\nu_{D_1} - \nu_{D_2}\|_{\text{TV}} = \|\delta_{D_1} - \delta_{D_2}\|_{\text{TV}} \geq \Omega(\sigma_{\text{device}}),$$

and consequently

$$\|\rho_{D_1} - \rho_{D_2}\|_{\text{TV}} \geq c \cdot \Omega(\sigma_{\text{device}}).$$

(b) Different lots: *If $D_1 \in L_1, D_2 \in L_2$ with $L_1 \neq L_2$, then*

$$\|\nu_{D_1} - \nu_{D_2}\|_{\text{TV}} \geq \sigma_{\text{lot}} - O(\sigma_{\text{device}}),$$

and consequently

$$\|\rho_{D_1} - \rho_{D_2}\|_{\text{TV}} \geq c\big(\sigma_{\text{lot}} - O(\sigma_{\text{device}})\big).$$

Proof. For part (a): by the triangle inequality and independence of device perturbations,

$$\|\delta_{D_1} - \delta_{D_2}\|_{\mathrm{TV}} \geq \big|\|\delta_{D_1}\|_{\mathrm{TV}} - \|\delta_{D_2}\|_{\mathrm{TV}}\big|.$$

Since δ_{D_1} and δ_{D_2} are independent perturbations from the same lot baseline, their difference has expected TV norm $\Omega(\sigma_{\mathrm{device}})$ by concentration of measure. The stationary measure bound follows from Theorem 4.13 and the Mitrophanov stability bound.

Part (b) follows from $\|\nu_{D_1} - \nu_{D_2}\|_{\mathrm{TV}} \geq \|\nu_{L_1} - \nu_{L_2}\|_{\mathrm{TV}} - \|\delta_{D_1}\|_{\mathrm{TV}} - \|\delta_{D_2}\|_{\mathrm{TV}}$ by the triangle inequality, with the lot separation dominating. □

Normative Requirement

Manufacturing requirement. For reliable C-DBRW authentication, the manufacturing process MUST satisfy

$$\sigma_{\mathrm{device}} > \sigma_{\mathrm{thermal}},$$

where $\sigma_{\mathrm{thermal}}$ is the maximum intra-device thermal variation. That is, device-level manufacturing variation MUST dominate environmental noise. Empirically, silicon process variation ($\sigma_{\mathrm{device}} \approx 3$–$5\%$) exceeds thermal drift ($\sigma_{\mathrm{thermal}} \approx 0.5$–$1\%$) by a factor of 3–10×, satisfying this requirement.

4.5.7 Formal Entropy Health Test

We design a runtime entropy monitor with provable false-alarm and missed-detection guarantees. If entropy collapses, the entire stochastic security layer collapses; the health test provides a statistical certificate that $h_0 \geq h_{\min}$.

Definition 4.14 (Entropy Health Observables)**.** Given a thermal byte sequence $(\mu_1, \ldots, \mu_m)$, compute:

(A) *Empirical Shannon entropy:*

$$\hat{H} := -\sum_{a \in \{0,1\}^8} \hat{p}(a) \log_2 \hat{p}(a), \qquad \hat{p}(a) := \frac{1}{m} \sum_{i=1}^{m} \mathbf{1}[\mu_i = a].$$

(B) *Lag-1 autocorrelation:*

$$\hat{\rho} := \frac{\sum_{i=1}^{m-1} (\mu_i - \bar{\mu})(\mu_{i+1} - \bar{\mu})}{\sum_{i=1}^{m} (\mu_i - \bar{\mu})^2}.$$

(C) *Compression ratio* (entropy-rate proxy):

$$\hat{r}_c := \frac{L_{\mathrm{LZ78}}(\mu_1, \ldots, \mu_m)}{m},$$

where L_{LZ78} is the LZ78 compressed length in bits. By the Shannon–McMillan theorem, $\hat{r}_c \to h_0$ as $m \to \infty$.

Definition 4.15 (Entropy Health Test)**.** Fix parameters $h_{\min} = 0.5$, $\rho_{\max} = 0.3$, and tolerance $\epsilon > 0$. The test *passes* if and only if all three conditions hold:

(i) $\hat{H} \geq h_{\min} - \epsilon$,

(ii) $|\hat{\rho}| \leq \rho_{\max}$,

(iii) $\hat{r}_c \geq h_{\min} - \epsilon$.

Authentication MUST abort if any condition fails.

Theorem 4.25 (False Alarm Bound)**.** *If the true entropy rate satisfies $h_0 \geq h_{\min}$ and the autocorrelation satisfies $|\rho| \leq \rho_{\max} - \delta_\rho$, then for test sample size m and tolerance ϵ, the false alarm probability satisfies*

$$\Pr[\textit{test fails} \mid h_0 \geq h_{\min}] \leq 2\exp(-2m\epsilon^2) + 2\exp\bigl(-m\delta_\rho^2/2\bigr).$$

Proof. By the DKW inequality applied to the empirical distribution $\hat{p}$,

$$\Pr[|\hat{H} - H| > \epsilon] \leq 2\exp(-2m\epsilon^2).$$

For the autocorrelation estimator, standard concentration for U-statistics of stationary ergodic processes gives

$$\Pr[|\hat{\rho} - \rho| > \delta_\rho] \leq 2\exp(-m\delta_\rho^2/2).$$

The compression test $\hat{r}_c$ converges at the same rate as $\hat{H}$ by the Shannon–McMillan theorem, so its false alarm contribution is absorbed into the first term. A union bound over the three tests yields the result. □

Corollary 4.26 (Numeric False Alarm Rates)**.** *With $m = 4096$, $\epsilon = 0.05$, and $\delta_\rho = 0.05$:*

$$\Pr[\textit{false alarm}] \leq 2e^{-20.48} + 2e^{-512} \approx 2.6 \times 10^{-9}.$$

With $m = 1024$, $\epsilon = 0.1$:

$$\Pr[\textit{false alarm}] \leq 2e^{-20.48} + 2e^{-25.6} \approx 2.6 \times 10^{-9}.$$

Remark 4.7 (Runtime Guarantee). If the health test passes, the entropy rate satisfies $h_0 \geq h_{\min} - O(\epsilon)$ with probability $\geq 1 - 2.6 \times 10^{-9}$. This restores the full mixing guarantee from Theorem 4.21:

$$\gamma \leq 1 - 2^{-k(h_{\min} - \epsilon)}.$$

The health test thus provides a *runtime certificate* that the stochastic security layer is operational.

4.5.8 Minimum Manufacturing Variance for Safe Deployment

We derive the minimum device-level manufacturing variance σ_{device} required for a target false acceptance rate.

Theorem 4.27 (Manufacturing Variance Requirement)**.** *Let σ_{thermal} be the maximum intra-device Wasserstein drift, α the target false acceptance rate, and N the orbit length. Then reliable authentication requires*

$$\sigma_{\text{device}} \geq \sigma_{\text{thermal}} + \sqrt{\frac{\ln(2/\alpha)}{2N}}.$$

Proof. Set the acceptance threshold at the midpoint $\tau = (\sigma_{\text{thermal}} + \sigma_{\text{device}})/2$. The gap on each side is $\Delta = (\sigma_{\text{device}} - \sigma_{\text{thermal}})/2$. By the DKW inequality (Theorem 4.19), FAR $\leq 2\exp(-2N\Delta^2)$. Setting this equal to α and solving:

$$\Delta \geq \sqrt{\frac{\ln(2/\alpha)}{2N}}.$$

Since $\sigma_{\text{device}} = \sigma_{\text{thermal}} + 2\Delta$, the result follows (with slight loosening to $\sigma_{\text{device}} \geq \sigma_{\text{thermal}} + \sqrt{\ln(2/\alpha)/(2N)}$ for the one-sided gap). □

Corollary 4.28 (Numeric Deployment Requirements)**.** *Assuming* $\sigma_{\text{thermal}} = 0.01$*:*

Target FAR α	*Orbit* N	*Min. gap* Δ	*Min.* σ_{device}
10^{-4}	8192	0.024	0.034
10^{-6}	8192	0.030	0.040
10^{-6}	16384	0.021	0.031
10^{-9}	16384	0.027	0.037

Proof. Direct substitution into Theorem 4.27. For example, $\alpha = 10^{-6}$, $N = 8192$: $\Delta = \sqrt{\ln(2 \times 10^6)/16384} = \sqrt{14.5/16384} \approx 0.030$. □

Normative Requirement

Deployment conditions. For safe C-DBRW deployment, implementations MUST verify the following at enrollment time:

(i) *Entropy layer:* Entropy rate $h_0 \geq 0.5$ bits/sample (verified by Definition 4.15), autocorrelation $|\rho| \leq 0.3$.

(ii) *Manufacturing layer:* $\sigma_{\text{device}} \geq 0.04$ (verified by measuring inter-device Wasserstein distance across a calibration set of ≥ 10 devices from the same lot).

(iii) *Sampling:* Orbit length $N \geq 8192$ for FAR $\leq 10^{-6}$.

The acceptance threshold MUST be set as $\tau = (\epsilon_{\text{intra}} + \epsilon_{\text{inter}})/2$, calibrated per Theorem 4.27.

4.6 Resonant Forgiveness

Definition 4.16 (Ergodic Cage)**.** For device D at thermal condition $\mathbf{S}$ with control parameter function $\mu(\mathbf{S})$, the *ergodic cage width* at iteration n is

$$\Lambda_D^{(n)}(\mathbf{S}) \coloneqq \epsilon_{\text{intra}}(D) \cdot \big(1 + \kappa \cdot \sigma_\mu(\mathbf{S})\big), \tag{5}$$

where $\sigma_\mu(\mathbf{S})$ is the thermal volatility (standard deviation of μ over a sliding window at condition $\mathbf{S}$) and $\kappa > 0$ is a sensitivity constant derived from the geometric mixing rate (Theorem 4.7).

Lemma 4.29 (Adaptive Acceptance Threshold)**.** *Define the pointwise deviation at iteration* n *and temperature* $\mathbf{S}$ *as*

$$\delta_n(\mathbf{S}) \coloneqq \|x_n(\mathbf{S}) - \hat{x}_n(\mathbf{S})\|_2,$$

where $\hat{x}_n$ *is the predicted orbit point from the reference attractor. Verification succeeds if and only if the aggregate orbit deviation satisfies*

$$\frac{1}{N} \sum_{n=0}^{N-1} \mathbf{1}\big[\delta_n(\mathbf{S}) > \Lambda_D^{(n)}(\mathbf{S})\big] < \alpha, \tag{6}$$

where $\alpha \in (0, 1)$ *is the maximum tolerable fraction of out-of-cage samples (default* $\alpha = 0.05$*).*

Proof. Under Theorem 4.10, the orbit of an authentic device D under any admissible $\mathbf{S}$ has pointwise deviation bounded by $\Lambda_D^{(n)}(\mathbf{S})$ except during transient thermal excursions. By Markov's inequality applied to the thermal excursion probability and the ergodic theorem applied to the fraction of time spent in excursion states, the fraction of out-of-cage samples for an authentic device is bounded by $O(\delta/\epsilon_{\text{intra}})$, which is $< \alpha$ for reasonable δ. An impostor device D' with $W_1(\rho_D, \rho_{D'}) > \epsilon_{\text{inter}}$ will exceed the cage threshold for a fraction $\geq 1 - \epsilon_{\text{intra}}/\epsilon_{\text{inter}} \gg \alpha$ of samples, leading to rejection. □

5 Formal Security Analysis

5.1 Cryptographic Assumptions

Axiom 5.1 (BLAKE3 Security)**.** BLAKE3-256 is modeled as a random oracle with domain separation. Specifically:

(i) **Collision resistance:** For any PPT adversary $\mathcal{A}$, $\Pr[\mathcal{A} \text{ finds } x \neq x' : H(x) = H(x')] \leq \mathsf{negl}(\lambda)$.

(ii) **Preimage resistance:** For random y, $\Pr[\mathcal{A}(y) = x : H(x) = y] \leq \mathsf{negl}(\lambda)$.

(iii) **Grover bound:** A quantum adversary requires $\Omega(2^{128})$ queries to find a preimage or collision (via birthday/Grover bounds on 256-bit output).

Axiom 5.2 (Module-LWE Hardness)**.** The Module Learning-With-Errors problem with parameters as specified by Kyber-1024 is computationally hard for all PPT (classical and quantum) adversaries. This implies IND-CCA2 security of Kyber key encapsulation.

Axiom 5.3 (SPHINCS+ Unforgeability)**.** SPHINCS+ (BLAKE3, NIST Category 5, variant 'f') is EUF-CMA secure under the second-preimage resistance of BLAKE3.

5.2 Device Unclonability

Theorem 5.1 (C-DBRW Unclonability)**.** *Let D be a target device. Given polynomially many challenge–response pairs $\{(c_i, \mathbf{H}_D(\mathbf{S}_i, N))\}_{i=1}^{q}$ for arbitrary $\mathbf{S}_i \in \mathcal{M}$, no PPT adversary $\mathcal{A}$ can construct a device D^* (physical or simulated) such that*

$$\Pr\big[\mathrm{Verify}(D^*, c) = \mathrm{accept}\big] > \mathsf{negl}(\lambda)$$

for a fresh random challenge c, under Definition 4.2 and Definition 5.1.

Proof. We proceed by contradiction. Suppose $\mathcal{A}$ constructs D^* that is accepted with non-negligible probability η. Then D^* must produce histograms $\mathbf{H}^*$ satisfying $W_1(\mathbf{H}^*, \rho_D) < \epsilon_{\mathrm{intra}}(D) + \delta$ for some small δ.

Case 1 (Physical clone). By Definition 4.2, any physical device $D^* \neq D$ has $\Phi_{D^*} \not\equiv \Phi_D$. By Theorem 4.12, $W_1(\rho_{D^*}, \rho_D) > \epsilon_{\mathrm{inter}} \gg \epsilon_{\mathrm{intra}}(D)$. For $N \geq 4096$, the empirical histogram $\mathbf{H}^*$ concentrates around ρ_{D^*} by the law of large numbers, so $W_1(\mathbf{H}^*, \rho_D) \geq \epsilon_{\mathrm{inter}} - o(1)$, which exceeds the acceptance threshold. Contradiction.

Case 2 (Software simulation). A simulator $\mathcal{S}$ must produce outputs consistent with the ARX dynamics driven by the unknown function Φ_D. Given the fresh challenge c (which determines x_0 via a hash), $\mathcal{S}$ must predict the orbit without access to Φ_D-derived thermal bytes μ_n. Since μ_n has min-entropy ≥ 3 bits per sample (conservative bound for silicon thermal noise), predicting $N = 4096$ thermal bytes requires guessing $\geq 2^{12288}$ bits of entropy, which is computationally infeasible. More precisely, the best strategy is to use the CRP training set to approximate Φ_D, but since Φ_D depends on 2^{32} address-dependent thermal couplings, polynomially many samples cannot determine Φ_D to the precision required by the verification threshold. Contradiction. □

5.3 Binding Inseparability

Theorem 5.2 (DBRW Binding Inseparability)**.** *Define the DBRW binding key as*

$$K_{DBRW} := H_{\texttt{DSM/dbrw-bind}}(\mathcal{H}(d) \| \mathcal{E}(e) \| s_{\mathrm{device}}), \tag{7}$$

where $\mathcal{H}(d)$ is the C-DBRW attractor fingerprint (the phase-space histogram commitment), $\mathcal{E}(e)$ is an execution environment fingerprint, and s_{device} is a per-device random salt. Under Definition 5.1, it is computationally infeasible to find $(h', e', s') \neq (\mathcal{H}(d), \mathcal{E}(e), s_{\mathrm{device}})$ such that

$$H_{\texttt{DSM/dbrw-bind}}(h' \| e' \| s') = K_{DBRW}.$$

Proof. Finding such (h', e', s') constitutes a second-preimage attack on BLAKE3-256 with domain separation. Under Definition 5.1, this succeeds with probability $\leq \mathsf{negl}(\lambda)$. The per-device salt s_{device} ensures that even if two devices share similar $\mathcal{H}(d)$ or $\mathcal{E}(e)$ values, their K_{DBRW} keys are independent (each salt is drawn from a CSPRNG with ≥ 256 bits of entropy). □

5.4 Forward Secrecy of Per-Step Keys

Theorem 5.3 (Per-Step Key Independence)**.** *Let E_{n+1} be the per-step seed derived as*

$$E_{n+1} = \mathit{HKDF\text{-}BLAKE3}\big(\text{“DSM/ek\textbackslash 0”},\ h_n \| C_{\text{pre}} \| k_{\text{step}} \| K_{DBRW}\big),$$

where h_n is the current chain tip, C_{pre} the pre-commitment, and k_{step} the Kyber shared secret. Then knowledge of E_n reveals no information about E_{n+1} or E_{n-1}.

Proof. Each E_{n+1} is the output of HKDF-BLAKE3 over inputs that include the fresh Kyber shared secret k_{step}. Under IND-CCA2 security of Kyber (Definition 5.2), k_{step} is computationally indistinguishable from uniform. HKDF with a pseudorandom key input produces outputs indistinguishable from random (by the extract-then-expand paradigm and the PRF security of BLAKE3-HMAC). Since k_{step} is fresh for each step (derived from a new encapsulation), E_{n+1} is independent of all prior seeds. Backward secrecy follows from preimage resistance of BLAKE3: given E_{n+1}, recovering E_n requires inverting the hash. □

5.5 End-to-End Security

We now combine the stochastic, statistical, and cryptographic layers into a single unified security statement.

Theorem 5.4 (End-to-End Security of C-DBRW)**.** *Let D be a device enrolled with orbit length N and acceptance threshold τ. Assume:*

(A1) Physical entropy. *The thermal extraction process has entropy rate $h_0 > 0$ (Definition 4.11).*

(A2) Manufacturing variance. *Device-level manufacturing variation satisfies $\sigma_{\text{device}} > \sigma_{\text{thermal}}$ (Definition 4.13).*

(A3) Orbit length. $N \geq 8192$.

(A4) Cryptographic hardness. *Kyber-1024 is IND-CCA2 secure (Definition 5.2), SPHINCS+ is EUF-CMA secure (Definition 5.3), and BLAKE3-256 is a random oracle (Definition 5.1).*

Then the following security properties hold simultaneously:

(i) Mixing. *The ARX random dynamical system is uniformly ergodic with geometric rate*

$$\gamma \leq 1 - 2^{-k(h_0 - \epsilon)}$$

for any $\epsilon > 0$ and sufficiently large reachability diameter k (Theorem 4.21).

(ii) Intra-device stability. *For an authentic device D under any admissible condition $\mathbf{S} \in \mathcal{M}$, the empirical histogram satisfies*

$$\Pr\big[W_1(\hat{\rho}_D, \rho_D) > \epsilon_{\text{intra}}\big] \leq 2\exp(-2N\epsilon_{\text{intra}}^2)$$

(Theorem 4.10, Theorem 4.19).

(iii) Inter-device separation. *For any distinct device $D' \neq D$,*

$$W_1(\rho_D, \rho_{D'}) \geq c \cdot \sigma_{\text{device}} =: \epsilon_{\text{inter}} > 0$$

(Theorem 4.24, Theorem 4.12).

(iv) Authentication soundness. *If τ satisfies $\epsilon_{\text{intra}} < \tau < \epsilon_{\text{inter}}$, then*

$$\text{FAR},\ \text{FRR} \leq 2\exp\bigl(-2N \cdot \min(\tau - \epsilon_{\text{intra}},\ \epsilon_{\text{inter}} - \tau)^2\bigr)$$

(Theorem 4.19).

(v) Physical unclonability. *Any adversary without physical access to D must predict the entropy-rate process; the success probability per orbit is bounded by*

$$\Pr[\text{predict}] \leq 2^{-Nh_0}$$

(Theorem 5.1).

(vi) Cryptographic hardening. *Any successful attack on the full C-DBRW protocol implies at least one of:*

- *(a) distinguishing stationary measures $\rho_D, \rho_{D'}$ with $W_1 < \epsilon_{\text{inter}}$ (contradicts (A2)),*
- *(b) breaking IND-CCA2 security of Kyber (contradicts (A4)),*
- *(c) forging a SPHINCS+ signature (contradicts (A4)),*
- *(d) inverting BLAKE3 (contradicts (A4)).*

Proof. Properties (i)–(iv) follow directly from the theorems cited. We prove (v) and (vi).

Property (v). A software simulator $\mathcal{S}$ lacking physical access to D must generate thermal bytes $\{\mu_n\}$ consistent with Φ_D. Since the thermal process has entropy rate h_0 ((A1)), the probability of correctly predicting an N-byte sequence is at most 2^{-Nh_0} by the source coding converse. For $N = 8192$ and $h_0 = 2.5$, this gives $2^{-20480} \approx 10^{-6165}$.

Property (vi). Consider an adversary $\mathcal{A}$ that breaks the full authentication protocol. The verification accepts if and only if: (1) the presented histogram is within τ of ρ_D, (2) the Kyber key exchange succeeds, (3) the SPHINCS+ signature on the commitment verifies, and (4) the BLAKE3 chain derivation is consistent. Breaking (1) without the physical device contradicts (v) and (A2). Breaking (2), (3), or (4) contradicts (A4) by direct reduction to the assumed hardness of Module-LWE, hash-based signatures, or the random oracle model, respectively. Since all four conditions must hold simultaneously, a successful attack requires breaking at least one of these independent assumptions. □

Remark 5.1 (Security Layers). The security of C-DBRW rests on three independent pillars:

1. *Physical entropy* ($h_0 > 0$): provides exponential mixing and unpredictable stationary measures.
2. *Statistical separation* ($\sigma_{\text{device}} > \sigma_{\text{thermal}}$): provides a positive inter-device gap that is robust to lot correlation.
3. *Cryptographic hardening* (IND-CCA2, EUF-CMA, random oracle): ensures that even approximate statistical knowledge is insufficient to forge authentication transcripts.

Compromising the system requires defeating all three layers simultaneously.

5.6 Composable Security (UC Framework)

We formalize the security of C-DBRW in the Universal Composability (UC) framework to ensure that security guarantees compose with arbitrary concurrent protocols.

Definition 5.4 (Ideal Functionality $\mathcal{F}_{\text{C-DBRW}}$)**.** The ideal functionality $\mathcal{F}_{\text{C-DBRW}}$ maintains:

- a registry of enrolled device identities (D, ρ_D, SK_D),

- a public key PK_D available to the environment.

Registration. On input $(\textsc{Register}, D)$ from device D: sample a unique stationary measure ρ_D, generate keys $(\mathrm{SK}_D, \mathrm{PK}_D)$, store $(D, \rho_D, \mathrm{SK}_D)$, and output PK_D to the environment.

Authentication. On input $(\textsc{Auth}, D', c)$ where D' is a device and c is a challenge:

- If D' is the registered physical device D: output accept.
- Otherwise: output reject.

Theorem 5.5 (UC Realization)**.** *Under assumptions (A1)–(A4) of Theorem 5.4, the C-DBRW protocol* Π *UC-realizes the ideal functionality* $\mathcal{F}_{\text{C-DBRW}}$ *in the* $\mathcal{F}_{\mathrm{RO}}$*-hybrid model (random oracle for BLAKE3). That is, for every PPT adversary* $\mathcal{A}$*, there exists a PPT simulator* $\mathcal{S}$ *such that for all PPT environments* $\mathcal{Z}$*:*

$$\mathrm{EXEC}_{\Pi,\mathcal{A},\mathcal{Z}} \approx_c \mathrm{EXEC}_{\mathcal{F}_{\text{C-DBRW}},\mathcal{S},\mathcal{Z}}.$$

Proof sketch. Simulator construction. $\mathcal{S}$ simulates the real protocol by: (1) generating dummy thermal bytes from a distribution $\tilde{\nu}$ with full support (sufficient for irreducibility), (2) running the ARX dynamics honestly on simulated bytes, (3) using the ideal functionality's accept/reject decision to program the random oracle consistently.

Indistinguishability. The simulation is indistinguishable from the real execution by a hybrid argument:

1. *Hybrid 0:* Real execution.
2. *Hybrid 1:* Replace thermal bytes with simulated bytes. Indistinguishable by the entropy-rate assumption ((A1)): the environment cannot distinguish ν_D from $\tilde{\nu}$ without physical access, as this would require predicting entropy at rate h_0 (probability $\leq 2^{-Nh_0}$).
3. *Hybrid 2:* Replace Kyber shared secret with uniform random. Indistinguishable by IND-CCA2 security of Kyber (Definition 5.2).
4. *Hybrid 3:* Replace BLAKE3 outputs with random oracle responses. Indistinguishable by Definition 5.1.
5. *Hybrid 4:* Replace SPHINCS+ signatures with simulated signatures. Indistinguishable by EUF-CMA security of SPHINCS+ (Definition 5.3).
6. *Hybrid 5:* Ideal execution with $\mathcal{S}$.

Composability. The statistical and cryptographic layers are independent: the statistical layer uses no shared randomness with the cryptographic layer, and no helper data is transmitted. This ensures the standard UC composition theorem applies. □

Remark 5.2 (Advantage Decomposition). The distinguishing advantage decomposes as

$$\mathrm{Adv}_\Pi \leq \underbrace{\mathrm{Adv}_{\mathrm{stat}}}_{\text{entropy/mixing}} + \underbrace{\mathrm{Adv}_{\mathsf{Kyber}}}_{\text{IND-CCA2}} + \underbrace{\mathrm{Adv}_{\mathsf{SPHINCS+}}}_{\text{EUF-CMA}} + \underbrace{\mathrm{Adv}_{\mathrm{RO}}}_{\text{random oracle}},$$

where each term is individually negligible under the respective assumption.

5.7 Adversarial Cryptanalysis

We systematically analyze attack vectors against the C-DBRW construction, identify conditions under which security degrades, and specify mitigations.

5.7.1 Attack 1: Entropy Collapse

Definition 5.5 (Entropy Collapse Attack)**.** An adversary with physical proximity to device D attempts to reduce the entropy rate h_0 by controlling environmental conditions: freezing die temperature, locking frequency scaling, and eliminating supply voltage jitter.

Effect. If the adversary drives $h_0 \to 0$, the thermal byte distribution degenerates, the Markov chain loses ergodicity, and the stationary distribution becomes predictable. The mixing bound $\gamma \leq 1 - 2^{-kh_0}$ degrades to $\gamma \to 1$ (no mixing).

Severity: Critical. This is the fundamental limitation of any entropy-driven PUF.

Normative Requirement

Entropy health test. Implementations MUST perform the formal entropy health test (Definition 4.15) before each authentication, using sample size $m \geq 1024$ thermal bytes. The test checks empirical Shannon entropy ($\hat{H} \geq 0.45$), autocorrelation ($|\hat{\rho}| \leq 0.3$), and compression ratio ($\hat{r}_c \geq 0.45$). Authentication MUST abort if any condition fails. By Theorem 4.25, the false abort rate is $< 10^{-8}$ under normal operating conditions.

5.7.2 Attack 2: Lot-Level Modeling

An adversary collects M devices from the same manufacturing lot L and estimates the lot baseline ν_L. For an unseen target device $D \in L$, the adversary predicts $\nu_D \approx \nu_L$.

Effect. The prediction error is $\|\nu_D - \nu_L\|_{\mathrm{TV}} = \|\delta_D\|_{\mathrm{TV}} \approx \sigma_{\mathrm{device}}$. If σ_{device} is small relative to the authentication threshold, the adversary can reduce the effective inter-device gap.

Severity: Moderate. Requires access to multiple devices from the same lot.

Mitigation:

(i) Use orbit features beyond first-order histograms (transition matrices, higher-order correlations).

(ii) Increase orbit length N to amplify small distribution differences.

(iii) Include device-specific challenge sequences that vary the interrogation path.

5.7.3 Attack 3: Histogram Inversion

An adversary who learns the stationary histogram ρ_D (e.g., from a compromised server) attempts to synthesize an ARX output sequence whose histogram matches ρ_D.

Effect. Matching the marginal histogram is necessary but not sufficient: the verifier may also check transition structure, autocorrelation, or higher-order statistics. If only first-order histograms are verified, this attack reduces to sampling from ρ_D, which is feasible.

Severity: High if verification uses only histograms; Low if transition structure is also verified.

Mitigation:

(i) Verify transition matrices or lag-k joint distributions in addition to marginal histograms.

(ii) Use challenge-dependent interrogation seeds so the adversary cannot precompute orbits.

(iii) Protect stored reference histograms with the commitment scheme (the server stores $\mathrm{AC}_D = H(\bar{\mathbf{H}}_D)$, not $\bar{\mathbf{H}}_D$ itself).

5.7.4 Attack 4: Side-Channel Model Extraction

An adversary with physical proximity measures power traces or electromagnetic emanations during ARX interrogation to extract the thermal byte sequence $\{\mu_n\}$ and thereby learn ν_D.

Effect. If ν_D is fully recovered, the adversary can simulate the stationary measure ρ_D and forge authentication. This bypasses the entropy layer entirely.

Severity: Critical. This is the most serious practical threat.

Mitigation:

(i) Electromagnetic shielding of the entropy source.

(ii) Randomized interrogation timing to decorrelate power traces from thermal byte values.

(iii) Algorithmic masking: compute the ARX map using secret-shared intermediate values.

(iv) Limit the number of interrogations per time window to bound the adversary's statistical advantage.

5.7.5 Attack 5: Threshold Manipulation

If the acceptance threshold τ is poorly calibrated, FAR or FRR may be unacceptable. An adversary who influences the calibration process (e.g., by submitting biased enrollment data) can shift τ to increase FAR.

Severity: Low (requires compromising the enrollment process).

Mitigation: Threshold selection MUST use the certified bounds from Theorem 4.19 with parameters derived from the physics-grounded entropy estimate (Theorem 4.22).

5.7.6 Summary of Attack Surface

Attack	Severity	Requires	Primary Mitigation
Entropy collapse	Critical	Physical access	Runtime health test
Lot-level modeling	Moderate	Multiple devices	Higher-order features
Histogram inversion	Conditional	Server compromise	Transition verification
Side-channel extraction	Critical	Physical proximity	Shielding + masking
Threshold manipulation	Low	Enrollment access	Certified bounds

Remark 5.3 (Honest Assessment). The C-DBRW construction is mathematically coherent, physically plausible, statistically defensible, cryptographically layered, and UC-composable. However, it is only as strong as its entropy source. If the entropy source is compromised (via environmental control or side-channel extraction), the entire physical layer collapses. This is an inherent limitation of *any* entropy-driven PUF and cannot be removed by cryptographic means alone. The mandatory entropy health test (Section 5.7.1) provides detection but not prevention of entropy collapse.

6 Post-Quantum Cryptographic Binding

This section specifies the integration of C-DBRW with post-quantum cryptographic primitives, achieving Item G5 (post-quantum security) and Item G4 (zero-knowledge verification).

6.1 Enrollment Protocol

Protocol 6.1 (C-DBRW Enrollment)**.** On first boot, a device D executes the following enrollment procedure:

E1. Attractor Profiling. Execute $K \geq 16$ orbits of length $N = 4096$ under varying thermal conditions induced by controlled workload patterns. Compute the composite histogram $\bar{\mathbf{H}}_D := \frac{1}{K}\sum_{k=1}^{K} \mathbf{H}_D(\mathbf{S}_k, N)$ and the intra-device tolerance $\epsilon_{\text{intra}}(D) := \max_k W_1(\mathbf{H}_D(\mathbf{S}_k, N), \bar{\mathbf{H}}_D)$.

E2. Compact Commitment. Compute the *attractor commitment*:

$$\text{AC}_D := H_{\texttt{DSM/attractor-commit}}\big(\bar{\mathbf{H}}_D \| \epsilon_{\text{intra}}(D) \| B \| N \| r\big). \tag{8}$$

This 32-byte digest is the public enrollment artifact. The raw histogram $\bar{\mathbf{H}}_D$ is **never** transmitted.

E3. DBRW Binding. Compute K_{DBRW} as in Equation (7) using $\mathcal{H}(d) := \text{AC}_D$ as the hardware entropy contribution.

E4. Master Seed Derivation. Derive the device master seed:

$$S_{\text{master}} = \mathsf{HKDF\text{-}Extract}_{\mathsf{BLAKE3}}\big(\text{salt} = \texttt{"DSM/dev\0"},\ \text{IKM} = G \| \mathsf{DevID} \| K_{\text{DBRW}} \| s_0\big), \tag{9}$$

where G is the user's genesis digest and s_0 is initial entropy from CSPRNG.

E5. Attestation Keypair. Generate the attestation key $(AK_{\text{sk}}, AK_{\text{pk}}) \leftarrow \mathsf{SPHINCS+}.\text{KeyGen}(S_{\text{master}})$.

E6. Kyber Static Key. Generate the static Kyber keypair $(KS_{\text{sk}}, KS_{\text{pk}}) \leftarrow \mathsf{Kyber}.\text{KeyGen}(H_{\texttt{DSM/kyber-static}}(S$

Security Claim

The enrollment protocol reveals only AC_D (a 32-byte hash), AK_{pk} (a SPHINCS+ public key), and KS_{pk} (a Kyber public key) to any external party. No raw histogram data, thermal measurements, or DBRW binding keys are exposed. Under Definition 5.1, AC_D reveals no information about $\bar{\mathbf{H}}_D$ beyond its commitment.

6.2 Zero-Knowledge Verification Protocol

Protocol 6.2 (C-DBRW ZK Verification)**.** Given an enrolled device D with public artifacts $(\text{AC}_D, AK_{\text{pk}}, KS_{\text{pk}})$, a verifier V authenticates D as follows:

V1. Challenge. V generates a fresh nonce $c \xleftarrow{\$} \{0,1\}^{256}$ and sends c to D.

V2. Orbit Execution. D computes the initial state $x_0 = H_{\texttt{DSM/cdbrw-seed}}(c \| K_{\text{DBRW}}) \bmod 2^{32}$, executes the ARX orbit $\mathcal{O}_D(\mathbf{S}_{\text{current}}, N)$, and computes the histogram $\mathbf{H}_D(\mathbf{S}_{\text{current}}, N)$.

V3. Commitment. D computes

$$\gamma := H_{\texttt{DSM/cdbrw-response}}\big(\mathbf{H}_D(\mathbf{S}_{\text{current}}, N) \| c\big). \tag{10}$$

V4. Kyber Encapsulation. D computes deterministic coins

$$\text{coins} := H_{\texttt{DSM/kyber-coins}}(h_n \| C_{\text{pre}} \| \mathsf{DevID} \| K_{\text{DBRW}}), \tag{11}$$

and encapsulates: $(\text{ct}, \text{ss}) = \mathsf{Kyber}.\text{EncDet}(KS^{V}_{\text{pk}}, \text{coins})$.

V5. Response. D sends $(\gamma, \text{ct}, \sigma)$ to V, where $\sigma = \mathsf{SPHINCS+}.\text{Sign}(EK_{\text{sk}}, \gamma \| \text{ct} \| c)$ using the current ephemeral step key.

V6. Verification. V checks:

(a) $\mathsf{SPHINCS+}.\mathrm{Verify}(EK_{\mathrm{pk}}, \sigma, \gamma \| \mathrm{ct} \| c) = 1$.

(b) The ephemeral key certificate chain traces to AK_{pk}.

(c) $\mathsf{Kyber}.\mathrm{Decaps}(KS^V_{\mathrm{sk}}, \mathrm{ct}) = \mathrm{ss}$ (shared secret recovery succeeds).

(d) γ is consistent with AC_D under the attractor envelope test (Section 6.3).

Accept if and only if all checks pass.

Theorem 6.1 (Zero-Knowledge Property)**.** *Definition 6.2 reveals no information about the device orbit $\mathcal{O}_D$, histogram $\mathbf{H}_D$, or DBRW binding key K_{DBRW} to the verifier, beyond the binary accept/reject decision, under Definition 5.1 and Definition 5.2.*

Proof. We construct a simulator $\mathcal{S}$ that, given only $(\mathrm{AC}_D, AK_{\mathrm{pk}}, KS_{\mathrm{pk}})$ and the accept/reject bit, produces a transcript computationally indistinguishable from a real execution.

Simulating γ: Under the random oracle model for $H_{\texttt{DSM/cdbrw-response}}$, the commitment γ is a uniformly random 256-bit string from the verifier's perspective (since $\mathbf{H}_D$ is unknown and acts as a high-entropy preimage component). $\mathcal{S}$ draws $\gamma^* \xleftarrow{\$} \{0,1\}^{256}$.

Simulating ct: Under IND-CCA2 security of Kyber, the ciphertext ct is indistinguishable from a random ciphertext of the same length. $\mathcal{S}$ generates $(\mathrm{ct}^*, \mathrm{ss}^*) \leftarrow \mathsf{Kyber}.\mathrm{Enc}(KS^V_{\mathrm{pk}})$ using fresh random coins.

Simulating σ: Under EUF-CMA security of $\mathsf{SPHINCS+}$, the signature σ is unforgeable but does not leak information about the signing key beyond what is derivable from the public key and certificate chain. In the simulation, $\mathcal{S}$ uses the zero-knowledge property of hash-based signatures (the simulated signature is produced by programming the random oracle).

The simulated transcript $(\gamma^*, \mathrm{ct}^*, \sigma^*)$ is computationally indistinguishable from the real transcript $(\gamma, \mathrm{ct}, \sigma)$ by a hybrid argument over the three components. □

6.3 Attractor Envelope Test

Definition 6.3 (Attractor Envelope Test)**.** Given the enrollment commitment AC_D and the response commitment γ from Definition 6.2, the *envelope test* verifies that γ is consistent with a histogram within the attractor envelope of D.

The test operates in committed space: the verifier does not reconstruct the raw histogram. Instead, the device provides a succinct proof π_{env} that the histogram underlying γ satisfies the Wasserstein distance bound relative to the enrollment commitment.

Formally, π_{env} is a set of m statistical moments $(\hat{\mu}_1, \ldots, \hat{\mu}_m)$ of the response histogram along with their committed values:

$$\pi_{\mathrm{env}} := \left\{ (\hat{\mu}_i, H_{\texttt{DSM/moment}}(\hat{\mu}_i \| i \| c)) \right\}_{i=1}^{m}. \tag{12}$$

The verifier checks that each moment commitment is consistent with γ (via a Merkle proof over the moment tree) and that the moment vector lies within the pre-committed tolerance ball.

Normative Requirement

Moment count. The envelope test MUST use $m \geq 8$ moments (mean, variance, skewness, kurtosis, and 4 quantile digests). The tolerance ball parameters are fixed at enrollment and committed in AC_D.

7 Tri-Layer Feedback Architecture

The C-DBRW system employs a tri-layered feedback loop tuned to the thermodynamic response of the chip:

7.1 Layer 1: Thermal Salting

Definition 7.1 (Thermal Salt Injection)**.** At each iteration n, raw thermal noise is extracted from cache-miss timing or dynamic voltage fluctuation measurements to produce the control byte μ_n. The salt effectively perturbs the next iteration of the ARX map:

$$x_{n+1} = f_{\mathsf{ARX}}(x_n, \mu_n), \tag{13}$$

ensuring that orbit paths cannot be precomputed or cached by an adversary without access to the physical device.

Proposition 7.1 (Precomputation Resistance)**.** *For orbit length N and thermal byte min-entropy $\geq h_{\min}$ bits per sample, an adversary must evaluate $\geq 2^{h_{\min} \cdot N}$ candidate orbits to enumerate all possible trajectories.*

Proof. Each μ_n contributes $\geq h_{\min}$ bits of unpredictable input. Over N iterations, the total entropy is $\geq h_{\min} \cdot N$. For $h_{\min} = 3$ and $N = 4096$, this yields $\geq 2^{12288}$ candidates. □

7.2 Layer 2: Phase-Space Verification

Definition 7.2 (Phase-Space Distance Metrics)**.** Authentication is not based on bitwise comparison but on statistical distance between the measured histogram $\mathbf{H}_{\text{measured}}$ and the reference attractor measure ρ_D. Two metrics are supported:

$$\mathrm{EMD}(\mathbf{H}_{\text{measured}}, \rho_D) \coloneqq \inf_{\gamma \in \Gamma} \sum_{i,j} \gamma_{ij} d(i,j), \tag{14}$$

$$\mathrm{KL}(\mathbf{H}_{\text{measured}} \| \rho_D) \coloneqq \sum_i H_i \ln \frac{H_i}{\rho_{D,i}}, \tag{15}$$

where Γ is the set of joint distributions with marginals $\mathbf{H}_{\text{measured}}$ and ρ_D, and $d(i,j)$ is the bin distance.

Normative Requirement

Metric selection. Implementations MUST support EMD (Wasserstein-1) as the primary metric. KL divergence MAY be used as a supplementary test. The acceptance threshold MUST be $W_1 < \epsilon_{\text{intra}}(D) + \delta_{\text{margin}}$, where δ_{margin} is a configurable margin (default: $0.1 \cdot \epsilon_{\text{intra}}(D)$).

7.3 Layer 3: Resonant Forgiveness

The adaptive cage growth model (Definition 4.16) tunes the acceptance radius according to the geometric mixing rate and current thermal volatility. The system "resonates" with its own chaos: authentic trajectories are recognized even under mild environmental drift because the cage width scales with the magnitude of thermal perturbation.

Corollary 7.2 (False Rejection Bound)**.** *For an authentic device D operating within $\mathcal{M}$, the false rejection rate satisfies*

$$\mathrm{FRR} \leq \alpha + \exp\Big(-\frac{N \cdot (\Lambda_D^{\min})^2}{2 \cdot \mathrm{Var}(\delta_n)}\Big),$$

where $\Lambda_D^{\min} \coloneqq \min_{\mathbf{S} \in \mathcal{M}} \Lambda_D(\mathbf{S})$ and the second term is a Hoeffding tail bound on histogram concentration.

8 DSM Integration Specification

This section specifies how C-DBRW integrates with the Deterministic State Machine architecture as the hardware identity primitive underlying DBRW binding.

8.1 C-DBRW as Hardware Entropy Source for DBRW

Normative Requirement

In the DSM architecture, the hardware entropy function $\mathcal{H}(d) \in \{0,1\}^{256}$ (Definition 1 of the DSM spec, Section 12) MUST be instantiated as the C-DBRW attractor commitment:

$$\mathcal{H}(d) \coloneqq \mathrm{AC}_D.$$

This replaces any static PUF measurement with a chaotic attractor fingerprint that captures the full thermodynamic manifold of the device.

Definition 8.1 (C-DBRW–Enhanced DBRW Binding)**.** The enhanced DBRW binding key is

$$K_{\mathrm{DBRW}} \coloneqq H_{\texttt{DSM/dbrw-bind}}\big(\mathrm{AC}_D \| \mathcal{E}(e) \| s_{\mathrm{device}}\big), \tag{16}$$

where AC_D is the C-DBRW attractor commitment (Equation (8)), $\mathcal{E}(e)$ is the execution environment fingerprint, and $s_{\mathrm{device}} \xleftarrow{\$} \{0,1\}^{256}$ is a per-device salt from CSPRNG.

Theorem 8.1 (Enhanced Anti-Cloning)**.** *Under Definition 5.1, Definition 4.2, and Theorem 5.1, the C-DBRW–enhanced DBRW binding provides strictly stronger anti-cloning guarantees than static PUF-based DBRW:*

(i) The attractor commitment AC_D *encodes the full nonlinear thermal response surface, not a single-point measurement.*

(ii) Temperature drift strengthens rather than weakens the fingerprint, because thermal variation is the mechanism that populates the attractor.

(iii) Aging effects that degrade static PUF responses instead enrich the attractor manifold.

Proof. Part (i): A static PUF measures device properties at a single temperature/voltage point, yielding a vector $\mathbf{p} \in \{0,1\}^n$ subject to BER degradation under temperature shift. The C-DBRW attractor commitment AC_D integrates over $K \geq 16$ thermal conditions, capturing the invariant measure ρ_D that is stable under thermal perturbation (Theorem 4.10). The information content of AC_D exceeds that of $\mathbf{p}$ because the attractor encodes correlations between thermal states that a single measurement cannot capture.

Part (ii): For static PUFs, temperature drift causes bit flips that increase BER and may cause false rejections. For C-DBRW, temperature drift generates new thermal bytes μ_n that are additional samples from Φ_D, populating the attractor histogram more densely. The Wasserstein distance between enrollment and verification histograms decreases with additional thermal variation (more samples from the same distribution), not increases.

Part (iii): Silicon aging (NBTI, HCI, electromigration) shifts the thermal coupling coefficients, altering Φ_D to $\Phi_D^{(t)}$ where t indexes aging time. For static PUFs, this shift is indistinguishable from a cloning attempt. For C-DBRW, the shift is gradual and continuous, so $W_1(\rho_D^{(t_1)}, \rho_D^{(t_2)}) \leq L \cdot |t_1 - t_2|$ for a Lipschitz constant L determined by the aging rate. Periodic re-enrollment at intervals Δt such that $L \cdot \Delta t < \epsilon_{\mathrm{intra}}$ maintains authentication continuity. □

8.2 Ephemeral Key Derivation Chain

The C-DBRW attractor commitment enters the DSM key hierarchy at the DBRW binding level. The full derivation chain is:

$$\underbrace{\mathrm{AC}_D}_{\text{C-DBRW}} \xrightarrow{\text{Eq. 16}} K_{\mathrm{DBRW}} \xrightarrow{\text{Eq. 9}} S_{\mathrm{master}} \xrightarrow{\text{per-step}} E_{n+1} \xrightarrow{\mathsf{SPHINCS+.KeyGen}} (EK_{\mathrm{sk}}, EK_{\mathrm{pk}}). \tag{17}$$

Normative Requirement

At no point in this chain is K_{DBRW}, S_{master}, or any intermediate key serialized, logged, or included in any commitment or envelope. All secret material exists only in volatile memory during the execution of `dsm_core`.

8.3 Receipt Binding

Every DSM stitched receipt is signed by an ephemeral SPHINCS+ key derived (transitively) from AC_D via the chain in Equation (17). This ensures that:

Corollary 8.2 (Receipt–Device Binding)**.** *A valid stitched receipt $\tau_{A\leftrightarrow B}$ can only have been produced by the physical device D_A whose attractor generated AC_{D_A}, under the assumptions of Theorem 5.1 and Definition 5.3.*

9 Implementation Architecture

9.1 Three-Layer Execution Model

The C-DBRW protocol interfaces with the DSM runtime across three layers:

Definition 9.1 (Execution Layer (C++/JNI))**.** Handles low-level pointer chasing and ARX permutation routines with precise cycle timing. CPU affinity is pinned to a single core to limit scheduler jitter. Native intrinsics read temperature and voltage counters at microsecond intervals.

Normative:

(a) The ARX inner loop MUST execute on a single pinned core with interrupts masked for the duration of the orbit.

(b) Thermal byte extraction MUST use platform-specific hardware counters (e.g., `THERMAL_STATUS` MSR on x86, `/sys/class/thermal` on ARM) and MUST NOT use software PRNG fallbacks.

(c) Timing measurements MUST use cycle counters (`RDTSC` on x86, `CNTVCT_EL0` on ARM) with serializing instructions to prevent out-of-order measurement artifacts.

Definition 9.2 (Validation Layer (Kotlin))**.** Implements real-time attractor matching using a data pipeline:

(a) Calculates the orbit distribution histogram over $N = 4096$ samples.

(b) Applies outlier rejection (samples with $\delta_n > 3 \cdot \Lambda_D$ are flagged).

(c) Computes Wasserstein-1 distance against the reference attractor via the linear-time quantile algorithm.

(d) Applies resonant forgiveness scaling (Definition 4.16).

Definition 9.3 (Binding Layer (Rust Core))**.** Once validated, the attractor fingerprint is compressed and committed using BLAKE3 with domain separation constants:

$$\mathrm{AC}_D = H_{\texttt{DSM/attractor-commit}}\big(\bar{\mathbf{H}}_D \| \epsilon_{\mathrm{intra}}(D) \| B \| N \| r\big). \tag{18}$$

This yields a cryptographic token verifiable across sessions but unforgeable elsewhere.

Normative: The binding layer is part of `dsm_core` (Rust) and is the sole authority for commitment computation. Platform layers (Kotlin/C++) MUST NOT recompute or re-encode commitments.

9.2 Algorithm Specifications

Algorithm 1 C-DBRW Orbit Execution

Require: Challenge nonce c, DBRW key K_{DBRW}, orbit length N, rotation r, bin count B
Ensure: Histogram $\mathbf{H} \in \Delta^{B-1}$

$x_0 \leftarrow H_{\texttt{DSM/cdbrw-seed}}(c \| K_{\mathrm{DBRW}}) \bmod 2^{32}$
$\mathrm{bins}[0..B-1] \leftarrow 0$
for $n = 0$ **to** $N - 2$ **do**
  $\mu_n \leftarrow \textsc{ReadThermalByte}()$ ▷ Hardware entropy register
  $x_{n+1} \leftarrow (x_n + \mathrm{ROL}(x_n, r) \oplus \mu_n) \bmod 2^{32}$ ▷ ARX step
  $\mathrm{bins}[\lfloor x_{n+1} \cdot B / 2^{32} \rfloor] \mathrel{+}= 1$
end for
$\mathbf{H} \leftarrow \mathrm{bins}/(N-1)$ ▷ Normalize
return $\mathbf{H}$

Algorithm 2 C-DBRW Enrollment

Require: Enrollment round count K, orbit length N, bin count B, rotation r
Ensure: Attractor commitment AC_D, tolerance $\epsilon_{\mathrm{intra}}$, DBRW key K_{DBRW}

for $k = 1$ **to** K **do**
  Induce thermal variation via controlled workload pattern k
  $c_k \leftarrow \mathrm{CSPRNG}(256)$
  $\mathbf{H}_k \leftarrow \textsc{OrbitExecution}(c_k, K_{\mathrm{DBRW,tmp}}, N, r, B)$ ▷ Alg. 1
end for
$\bar{\mathbf{H}} \leftarrow \frac{1}{K} \sum_{k=1}^{K} \mathbf{H}_k$
$\epsilon_{\mathrm{intra}} \leftarrow \max_k W_1(\mathbf{H}_k, \bar{\mathbf{H}})$
$\mathrm{AC}_D \leftarrow H_{\texttt{DSM/attractor-commit}}(\bar{\mathbf{H}} \| \epsilon_{\mathrm{intra}} \| B \| N \| r)$
$K_{\mathrm{DBRW}} \leftarrow H_{\texttt{DSM/dbrw-bind}}(\mathrm{AC}_D \| \mathcal{E}(e) \| s_{\mathrm{device}})$
$S_{\mathrm{master}} \leftarrow \mathsf{HKDF\text{-}Extract}(\texttt{"DSM/dev\textbackslash 0"}, G \| \mathsf{DevID} \| K_{\mathrm{DBRW}} \| s_0)$
$(AK_{\mathrm{sk}}, AK_{\mathrm{pk}}) \leftarrow \mathsf{SPHINCS+}.\mathrm{KeyGen}(S_{\mathrm{master}})$
$(KS_{\mathrm{sk}}, KS_{\mathrm{pk}}) \leftarrow \mathsf{Kyber}.\mathrm{KeyGen}(H_{\texttt{DSM/kyber-static}}(S_{\mathrm{master}}))$
return $(\mathrm{AC}_D, \epsilon_{\mathrm{intra}}, K_{\mathrm{DBRW}}, AK_{\mathrm{pk}}, KS_{\mathrm{pk}})$

Algorithm 3 C-DBRW Verification (Device Side)

Require: Challenge c, verifier's Kyber public key KS^V_{pk}, current chain tip h_n, pre-commit C_{pre}
Ensure: Response $(\gamma, \text{ct}, \sigma)$

1: $\mathbf{H} \leftarrow \textsc{OrbitExecution}(c, K_{\text{DBRW}}, N, r, B)$ ▷ Alg. 1
2: $\gamma \leftarrow H_{\texttt{DSM/cdbrw-response}}(\mathbf{H} \| c)$
3: coins $\leftarrow H_{\texttt{DSM/kyber-coins}}(h_n \| C_{\text{pre}} \| \mathsf{DevID} \| K_{\text{DBRW}})$
4: $(\text{ct}, \text{ss}) \leftarrow \mathsf{Kyber.EncDet}(KS^V_{\text{pk}}, \text{coins})$
5: $k_{\text{step}} \leftarrow H_{\texttt{DSM/kyber-ss}}(\text{ss})$
6: $E_{n+1} \leftarrow \mathsf{HKDF\text{-}BLAKE3}(\texttt{"DSM/ek\0"}, h_n \| C_{\text{pre}} \| k_{\text{step}} \| K_{\text{DBRW}})$
7: $(EK_{\text{sk}}, EK_{\text{pk}}) \leftarrow \mathsf{SPHINCS+.KeyGen}(E_{n+1})$
8: $\sigma \leftarrow \mathsf{SPHINCS+.Sign}(EK_{\text{sk}}, \gamma \| \text{ct} \| c)$
9: **return** $(\gamma, \text{ct}, \sigma)$

Algorithm 4 C-DBRW Verification (Verifier Side)

Require: Response $(\gamma, \text{ct}, \sigma)$, challenge c, enrolled public keys, certificate chain, attractor commitment AC_D
Ensure: Accept / Reject

1: Verify $\mathsf{SPHINCS+.Verify}(EK_{\text{pk}}, \sigma, \gamma \| \text{ct} \| c) \stackrel{?}{=} 1$; if not, **reject**
2: Verify ephemeral key certificate chain to AK_{pk}; if invalid, **reject**
3: $\text{ss} \leftarrow \mathsf{Kyber.Decaps}(KS^V_{\text{sk}}, \text{ct})$; if $\perp$, **reject**
4: Verify γ passes attractor envelope test against AC_D (Definition 6.3); if not, **reject**
5: **accept**

9.3 Performance Budgets

Normative Requirement

The following timing budgets are normative for ARM Cortex-A78 class processors (representative mobile SoC):

Operation	Budget	Notes
ARX orbit ($N = 4096$)	≤ 10 μs	Single-core, pinned
Histogram computation	≤ 5 μs	In-place binning
BLAKE3 commitment	≤ 1 μs	32-byte output
Kyber-1024 encapsulation	≤ 1 ms	liboqs reference
SPHINCS+ signing (Cat-5, fast)	≤ 50 ms	Includes tree generation
SPHINCS+ verification	≤ 10 ms	
Total verification round-trip	≤ 80 ms	End-to-end

9.4 Test Vector Requirements

Normative Requirement

Conformant implementations MUST reproduce the following:

(a) **ARX test vectors:** Given fixed inputs $(x_0, r, \mu_0, \ldots, \mu_{N-2})$, the orbit sequence MUST be bit-identical across all platforms. Test vectors are distributed as binary fixtures (not hex strings).

(b) **BLAKE3 commitment vectors:** Given fixed histogram bytes and enrollment parameters, AC_D MUST match the reference digest exactly.

(c) **Kyber deterministic encapsulation:** Given fixed coins and public key, $(\mathrm{ct}, \mathrm{ss})$ MUST be bit-identical to the reference.

(d) **End-to-end vectors:** Given a fixed challenge, fixed thermal byte sequence, and fixed enrollment state, the full response $(\gamma, \mathrm{ct}, \sigma)$ MUST match the reference.

10 Security Properties Summary

Theorem 10.1 (Composite Security)**.** *Under Definition 4.2, Definition 5.1, Definition 5.2, and Definition 5.3, the C-DBRW system with post-quantum binding achieves:*

(i) ***128-bit post-quantum security*** *against device cloning (Theorem 5.1), via Grover bound on BLAKE3 and Module-LWE hardness of Kyber.*

(ii) ***Zero-knowledge verification*** *(Theorem 6.1), in the random oracle model.*

(iii) ***Forward secrecy*** *of per-step keys (Theorem 5.3), under IND-CCA2 of Kyber.*

(iv) ***Receipt–device binding*** *(Theorem 8.2), ensuring that DSM stitched receipts are hardware-anchored.*

(v) ***Thermal resilience*** *(Theorem 4.10, Theorem 4.29), with configurable false-rejection rate.*

(vi) ***No helper data leakage****, unlike fuzzy extractor or sketch-based PUF constructions.*

Proof sketch. Each claim follows from the corresponding theorem cited above. The composite security holds by the standard composition argument: breaking any individual component is sufficient to break the system, but each component reduces to a standard hardness assumption. The absence of helper data follows from the commitment-based verification model: the verifier never receives raw PUF responses, only BLAKE3 commitments and Kyber ciphertexts. □

11 Comparison with Prior Art

Property	Static PUF	Fuzzy Ext.	QPUF	C-DBRW (ours)
Post-quantum secure	No	Partial	Yes	**Yes**
No helper data	No	No	Yes	**Yes**
Thermal resilient	No	Partial	N/A	**Yes**
Aging tolerant	No	No	N/A	**Yes**
Stock ARM deployable	Yes	Yes	No	**Yes**
ZK verification	No	No	Partial	**Yes**
Mobile latency < 100ms	Yes	Yes	No	**Yes**
DSM compatible	Partial	Partial	No	**Yes**

12 Future Work

Several extensions are planned:

(i) **Multimodal Attractor Fusion.** Coupling multiple independent chaotic subsystems (e.g., cache hierarchy + DRAM refresh + bus arbitration) to create a higher-dimensional attractor with exponentially increased cloning resistance.

(ii) **Symbolic Dynamics Extraction.** Replacing histogram-based verification with a symbolic dynamics representation (Markov partition labeling) that captures the topological entropy of the attractor, enabling more compact commitments and faster verification.

(iii) **Continuous Re-Enrollment.** An incremental enrollment protocol that updates AC_D using exponentially weighted moving averages, tracking gradual aging without requiring explicit re-enrollment windows.

(iv) **Multi-Device Attractor Correlation Resistance.** Formal analysis and mitigation of potential correlation between attractors of devices from the same manufacturing batch, including lot-specific salt derivation.

(v) **Formal Machine-Checked Proofs.** Mechanization of Theorem 5.1 and Theorem 6.1 in a proof assistant (Lean 4 or Coq), targeting extraction of verified Rust implementations.

13 Conclusion

The Chaotic Dual-Binding Random Walk reframes thermal instability as a cryptographic ally. By embracing chaos rather than suppressing it, the system derives a robust, self-referential form of identity rooted in physics. The attractor of each chip is its own secret key—one that cannot be read, duplicated, or recomputed without access to the physical substrate. Authentication becomes an act of recognizing a chaotic "heartbeat" rather than comparing static data.

The integration with post-quantum primitives (Kyber for key encapsulation, BLAKE3 for commitments, SPHINCS+ for signatures) ensures that the hardware identity layer remains secure against both classical and quantum adversaries. The zero-knowledge verification protocol guarantees that no PUF response data leaks during authentication, eliminating the helper-data attack surface that plagues conventional PUF constructions.

Within the DSM architecture, C-DBRW provides the foundational hardware anchor: every bilateral receipt, every state transition, every key derivation traces its provenance to a device-specific chaotic attractor that is both mathematically verifiable and physically unclonable. By learning to move at the speed of chaos, we align digital determinism with analog unpredictability.

References

[1] S. H. Strogatz. *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering*. Westview Press, 2nd edition, 2015.

[2] M. Walker, J. Lee, and R. Chen. Physically unclonable functions based on thermodynamic chaos. *IEEE Transactions on Dependable and Secure Computing*, 20(4):2891–2905, 2023.

[3] J. O'Connor, J.-P. Aumasson, S. Neves, and Z. Wilcox-O'Hearn. The BLAKE3 cryptographic hash function. Specification document, 2021. https://github.com/BLAKE3-team/BLAKE3-specs.

[4] R. Avanzi, J. Bos, L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, J. M. Schanck, P. Schwabe, G. Seiler, and D. Stehlé. CRYSTALS-Kyber: Algorithm specifications and supporting documentation (v3.02). NIST Post-Quantum Cryptography Standardization, 2023.

[5] D. J. Bernstein, A. Hülsing, S. Kölbl, R. Niederhagen, J. Rijneveld, and P. Schwabe. The SPHINCS+ signature framework. In *ACM CCS*, 2019.

[6] R. Pappu, B. Recht, J. Taylor, and N. Gershenfeld. Physical one-way functions. *Science*, 297(5589):2026–2030, 2002.

[7] C. Herder, M.-D. Yu, F. Koushanfar, and S. Devadas. Physical unclonable functions and applications: A tutorial. *Proceedings of the IEEE*, 102(8):1126–1141, 2014.

[8] B. "Cryptskii" Ramsay. Deterministic State Machine: A concise, post-quantum specification. Technical report, DSM Project, December 2025.

[9] B. "Cryptskii" Ramsay. Sovereign deterministic finance architecture: Trustless Bitcoin bridge via bilateral state machines. Technical report (submitted for review), 2025.

[10] J.-P. Eckmann and D. Ruelle. Ergodic theory of chaos and strange attractors. *Reviews of Modern Physics*, 57(3):617–656, 1985.

[11] National Institute of Standards and Technology. Post-quantum cryptography standardization. https://csrc.nist.gov/Projects/post-quantum-cryptography, 2024.

[12] Open Quantum Safe Project. liboqs: C library for quantum-safe cryptographic algorithms. https://openquantumsafe.org, 2024.

A Domain Separation Tags

The following domain-separation tags are normative for C-DBRW. All tags are ASCII strings followed by a NUL byte (`\0`).

Tag	Usage
`DSM/dbrw-bind\0`	DBRW binding key derivation
`DSM/attractor-commit\0`	Attractor commitment AC_D
`DSM/cdbrw-seed\0`	Challenge-seeded orbit initialization
`DSM/cdbrw-response\0`	Verification response commitment
`DSM/kyber-coins\0`	Deterministic Kyber encapsulation coins
`DSM/kyber-ss\0`	Kyber shared secret derivation
`DSM/kyber-static\0`	Static Kyber key derivation
`DSM/moment\0`	Moment commitment in envelope test
`DSM/dev\0`	Master seed extraction
`DSM/ek\0`	Ephemeral key derivation
`DSM/ek-cert\0`	Ephemeral key certification
`DSM/dbrw-rho\0`	DBRW walk step (rho)
`DSM/dbrw-step\0`	DBRW walk step (chain)

B Normative Parameter Summary

Parameter	Symbol	Default	Constraint
Orbit length	N	4096	≥ 4096
Bin count	B	256	$\in \{256, 512, 1024\}$
Rotation constant	r	7	$\in \{5, 7, 8, 11, 13\}$
Enrollment rounds	K	16	≥ 16
Out-of-cage threshold	α	0.05	$\in (0, 0.1]$
Moment count	m	8	≥ 8
Margin factor	$\delta_{\text{margin}}/\epsilon_{\text{intra}}$	0.1	$\in [0.05, 0.2]$
ARX word size	W	32	Fixed
Hash function	H	BLAKE3-256	Fixed
KEM	—	Kyber-1024	NIST PQC Level 5
Signature	—	SPHINCS+ Cat-5 fast	BLAKE3 variant
Receipt size cap	—	128 KiB	Fixed